Chapter 14
Establishing a Shelf Life and Setting Lot-Release Specifications

William Egan and Timothy Schofield

14.1 Introduction

Potency is a critical quality attribute of a vaccine and those attributes that impact potency are essential to assuring vaccine quality. Over its entire shelf life, a vaccine must exceed a minimal potency value and, when defined, cannot exceed a maximal potency.[1] Thus, over their shelf lives, vaccines must remain either above a minimum value of potency or within a defined range of potency. The upper and lower potency limits are established through clinical studies, with the lower limit based on efficacy and the upper limit based on known or potential safety concerns. As illustrated in Fig. 14.1, the vaccine's potency at release (expressed as a $TCID_{50}$ value in the Figure for a hypothetical live attenuated viral vaccine) and its associated shelf life are linked by the rate at which the vaccine, under defined storage conditions, loses potency. Thus, for the hypothetical viral vaccine in Fig. 14.1, with a lower allowed potency limit of 3,000 $TCID_{50}$s, its effective shelf life is nearly 36 months if released at 5,000 $TCID_{50}$s or *ca*, 19 months if released at 4,000 $TCID_{50}$s. The choice of the shelf life and its associated release value would be the manufacturer's decision (assuming that both release values were shown to be safe). In establishing a release potency value and shelf life specification, the uncertainties in the potency determination at release and rate of potency loss must also be

[1] Although maximal potency values are set for the majority of vaccines, certain vaccines have no defined upper potency limit.

W. Egan (✉)
Novartis Vaccines and Diagnostics, 350 Massachusetts Avenue,
Cambridge, MA 02139, USA
e-mail: William.egan@novartis.com

T. Schofield
MedImmune, One MedImmune Way, Gaithersburg, MD 20878, USA
e-mail: SchofieldT@MedImmune.com

B.K. Nunnally et al. (eds.), *Vaccine Analysis: Strategies, Principles, and Control*,
DOI 10.1007/978-3-662-45024-6_14

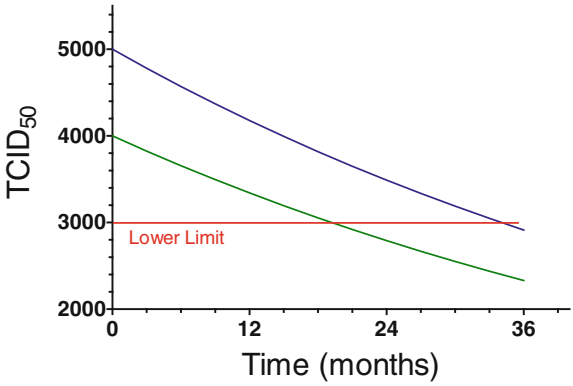

Fig. 14.1 Time dependence of the potency value for a hypothetical vaccine that would be released at either 5,000 or 4,000 $TCID_{50}$ units

established and accounted for; these uncertainties are not illustrated in Fig. 14.1. This chapter will focus on the determination of a vaccine shelf life and release specification, including the effects of statistical uncertainties using potency as an example attribute.

14.2 Establishing the Release Specification and Shelf Life

14.2.1 Establishing the Rate of Loss of Potency

Consider a hypothetical live attenuated viral vaccine and its associated potency values ($TCID_{50}$ values; herein expressed as their natural logarithm) as a function of time for three vaccine Lots as depicted in Fig. 14.2 (the $TCID_{50}$ values depicted in Fig. 14.2 are presented in Appendix). The immediate question is how to use these three data sets to establish a shelf life and associated release specification for the

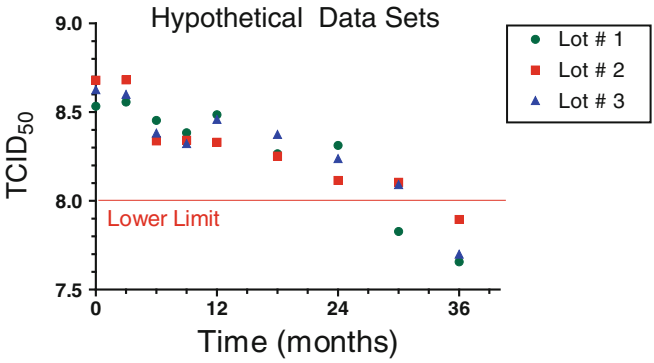

Fig. 14.2 Simulated time dependence of the potency ($TCID_{50}$ values; expressed as their natural logarithm) for three Lots of a hypothetical live attenuated viral vaccine

vaccine. Most simply, it could be concluded that if the vaccine were released at 5,000 $TCID_{50}$s [ln(5000) ≈ 8.52], or more, as permitted, then the vaccine's potency would remain within specification, i.e., above 3,000 $TCID_{50}$s [ln(3000) ≈ 8.0], for 24 months (although two of the Lots remained above this lower limit at 30 months, this was not the case for Lot #1 which fell below the lower limit at the 30-month time point); the shelf life could be set as 24 months, the last time point for which all three Lots remained within specification.

The above-described model for setting a shelf life, which may be termed a "compliance model" as it sets a shelf life by the length of time that the vaccine's potency measurements remain above (are compliant with) the assigned lower limit, while workable and having the advantage of simplicity, suffers several drawbacks.[2] First, there is an overreliance placed on the one or few data points that appear outside of the specification window; moreover, the totality of stability information that is contained within the other data points, namely, how potency is changing with time, is ignored. Second, there is a realization that increasing the number of Lots (or the number of time points for a given lot) that are tested will result in an increased probability that a test result, due simply to the random errors inherent in the assay, will fall outside the specification window and, therefore, result in a shortened shelf life. This discourages the collection of additional data. Third, a sense of the confidence or extent of certainty that may be placed in the shelf life is lacking. Finally, the shelf life of the vaccine would be undefined if the vaccine were to be released at a lower potency value.

As an alternative procedure, and to overcome the above-noted deficits, the time and potency data sets may be fit to a mathematical model (regression analysis) to obtain an estimate of the time dependence of the potency loss. The regression analysis would utilize all of the data and not rely on a single point or select set of data points. From the fitting process, an estimate for the mathematical model as well as a measure of the confidence in that estimate, expressed as a standard deviation or confidence interval on the regression line, is obtained. Non-mechanistically associated mathematical models, for example, fitting the data to a polynomial equation, or mechanistically associated models, such as a particular chemical kinetic model (with its accompanying mathematical form), can be explored.

Changes in vaccine potency arise from various processes. For a protein-based vaccine, the changes may derive from denaturation/aggregation, hydrolysis, deamidation, oxidation, disulfide interchanges, or other transformations to which proteins are subject. For a polysaccharide-based vaccine (including polysaccharide-protein conjugates), the changes in potency may derive, inter alia, from hydrolysis of the saccharide main chain or side chain, or loss of particular appended groupings (such as O-acetyl or pyruvate). For a live attenuated virus vaccine, changes in

[2] The World Health Organization, in its Guidelines on Stability Evaluation of Vaccines, has noted that "In many countries, expiry periods of vaccine products are calculated by testing a predefined number of Lots, at pre-defined intervals, and designating the expiry period as the first time at which a stability measurement falls below an acceptable threshold". The WHO Guideline also notes that "This approach has the advantage of simplicity, but may yield spurious results due to assay variability".

potency may derive from various changes which inhibit the ability of the virus to enter cells or to replicate within those cells. Such changes may be modeled—and have been modeled—as processes following particular kinetic rate equations. Oftentimes, these changes in potency, typically losses in potency, may be modeled as a first-order reaction, i.e., as a single exponential process. However, at times, more complicated kinetic models (for example, biphasic kinetics) are necessary. In this chapter, we will illustrate general concepts, treating changes in potency as simple exponential processes, i.e., as first-order kinetic processes.

14.2.1.1 First-Order Kinetics

The simplest plausible model for the loss of vaccine potency would be based on first-order kinetics, where the potency of the vaccine follows a rate law of the form

$$P(t) = P(0)e^{-kt} \tag{14.1}$$

where,

$P(t)$ is the vaccine potency at time, t
$P(0)$ is the vaccine potency at time zero
k is the rate constant for the loss of potency
e is the base of the natural logarithm

In Eq. 14.1, $P(0)$ is the time that the kinetic measurements are begun. Equation 14.1 may be linearized by taking the natural logarithm of both sides of the equation. Thus,

$$\ln P(t) = \ln P(0) - kt \tag{14.2}$$

Linear regression of a [time, potency] data set to Eq. 14.2 will provide estimates for the slope and y-intercept, representing, respectively, the estimated rate constant for the loss of potency, k, and the natural log of the value of the potency at $t = 0$, $\ln P(0)$. As an example, the data from Fig. 14.2 (see Appendix) may be analyzed for a particular lot, for example, Lot 1. The least-squares analysis provides the best estimates of the parameters of the linear model and the uncertainties associated with those parameters (see Fig. 14.3). The regression line crosses the lower limit for potency at *ca.* 27 months, a value that might reasonably be regarded as a shelf life. The ICH Guidance on stability (Q1E) recommends, however, that the uncertainty in the regression analysis needs also to be considered. Thus, given a decrease in potency over time, the shelf life may be set as the point in time where the one-sided 95 % confidence limit on the regression line intersects the prescribed lower limit for potency—approximately 22 months for Lot 1.[3] (The dashed line in Fig. 14.3 presents

[3] ICH Q1E notes that, "for an attribute known to decrease with time, the lower one-sided 95 % confidence limit should be compared to the acceptance criterion. For an attribute known to increase with time, the upper one-sided 95 % confidence limit should be compared to the acceptance

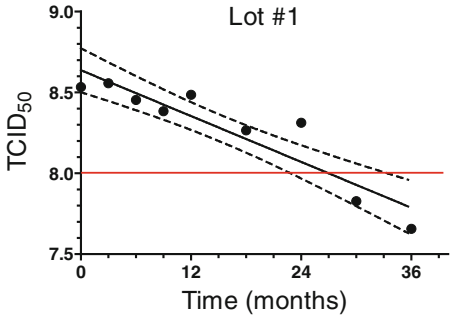

	Lot # 1
Best-fit values	
Slope	-0.02364 ± 0.003705
Y-intercept when X=0.0	8.637 ± 0.07165
X-intercept when Y=8.0	26.97
1/slope	-42.31
95% Confidence Intervals	
Slope	-0.03240 to -0.01487
X-intercept when Y=8.0	21.96 to 35.69
Goodness of Fit	
r^2	0.8532
Sy.x	0.1310
Is slope significantly non-zero?	
F	40.70
DFn, DFd	1.000, 7.000
P value	0.0004
Deviation from zero?	Significant

Fig. 14.3 Linear least-squares fit of time and potency data set for *Lot 1*. The *regression line* is the *solid line* and the 90 % confidence interval about that *regression line* is represented by the *dashed lines*. The least-squares values from the regression analysis are presented in the accompanying Table. In the accompanying Table, "Sy.x" represents the standard deviation of the residuals

Table 14.1 Calculated slope and intercept, with associated standard deviations, for Lots 1, 2, and 3, as well as combined Lots 1, 2, and 3

Lot number	Slope	Standard deviation of the slope	Intercept	Standard deviation of the intercept
1	−0.0236	±0.0037	8.637	±0.0716
2	−0.0194	±0.0027	8.601	±0.0521
3	−0.0207	±0.0034	8.629	±0.0663
1, 2, 3, combined	−0.2125	±0.0018	8.622	±0.0487

the two-sided 90 % confidence limits about the regression line; the lower branch of the two-sided 90 % confidence limit is equivalent to the lower one-sided 95 % confidence limit. A similar analysis can be carried out for Lots 2 and 3 and, as can be noted in Table 14.1, slightly different estimates for the rate in potency loss are observed, with slightly decreased rates of potency loss seen for Lots 2 and 3. Under suitable conditions, as defined in ICH Stability Guidance (ICH Q1E), the data from various vaccine Lots may be combined to obtain a better estimate of the parameters of the kinetic model.[4] This was done for the three Lots (which meet the ICH criteria for

(Footnote 3 continued)
criterion. For an attribute that can either increase or decrease, or whose direction of change is not known, two-sided 95 % confidence limits should be calculated and compared to the upper and lower acceptance criteria". In the examples of the time dependence of potency presented in this chapter, a loss in potency with time is considered.

[4] Analysis of covariance (ANCOVA) can be employed, where time is considered the covariate, to test the differences in slopes and intercepts of the regression lines among batches. Each of these

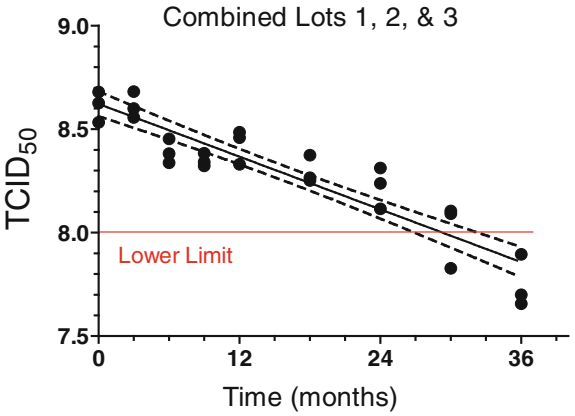

Fig. 14.4 Least-squares regression analysis of *Lots 1*, *2*, and *3* combined. The *dashed lines* represent the two-sided 90 % confidence limits of the *regression line*

combining of the data sets) of this example and the results are displayed in Fig. 14.4 (see also Table 14.1). Through the combination of the three Lots, it can be seen that the confidence limits about the regression parameter estimates have decreased. An additional comment about the results of these data analyses may be made. For any given time point over the vaccine shelf life, the potency estimate that is provided by the regression line may be regarded as a better estimate of the potency at that time than the measured value (for example, see Fig. 14.3) or the average of the measured values if replicate determinations are made (for example, see Fig. 14.4). This consideration is of importance when considering potential "out-of-specification" values at any given time point, and especially at time points approaching the end of shelf life for those vaccines that exhibit a loss of potency over time.

As noted, the regression analysis is based on a mathematical model. Various results from the regression analysis may be used to assess the adequacy of the model; for example, an analysis of the residuals (the residuals should be randomly distributed about the regression line and their standard deviation should be approximately equal to the SD of the potency assay).

The value of the coefficient of determination (the r^2 value) or the correlation coefficient (r) is often taken as a measure of the adequacy of the linear model. However, as is noted in the following Sect. 14.2.2.1, there are concerns associated with the sole use of this evaluation metric.

As shown in Fig. 14.4, the regression line intersects the lower assigned potency limit at nearly 30 months, while the lower one-sided 95 % confidence limit intersects the lower limit at approximately 27 months. These are slightly different values than were obtained from the analysis of Lot 1 (or of Lots 2 and 3).

(Footnote 4 continued)
tests should be conducted using a significance level of 0.25 to compensate for the expected low power of the design due to the relatively limited sample size in a typical formal stability study.

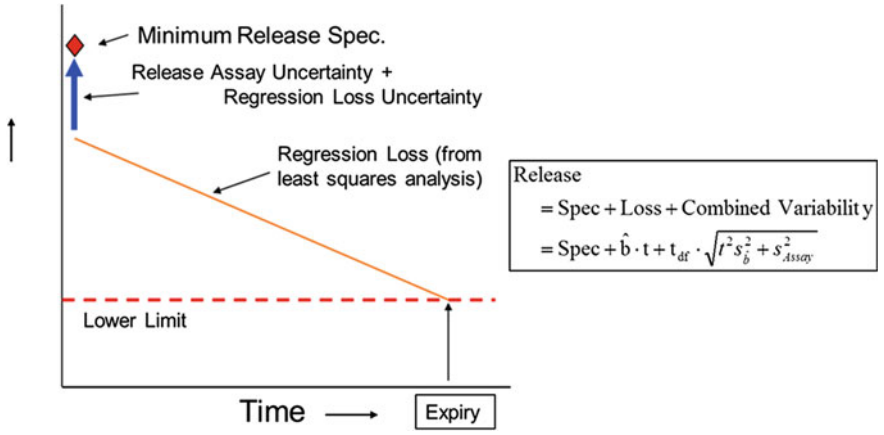

Fig. 14.5 Calculation of a minimum release specification as presented in the WHO Guidelines of stability for vaccines. The release specification is determined by the estimated rate of loss of potency over time as well as the uncertainties in the estimation of that rate of potency loss and the release assay

14.2.2 Setting a Shelf Life and Associated Lot-Release Specification

In the introduction of this chapter, we noted that, although commonly used, several drawbacks attend the use of a "compliance model" to setting a release specification and shelf life. We now present an alternative paradigm, the "Expiry Model", which has also been described in the aforementioned WHO Guideline on Vaccine Stability. The various elements of the Expiry Model are illustrated in Fig. 14.5. Basically, with a target shelf life as a goal, a minimum release value is calculated by accounting for the loss of potency over the period of the proposed shelf life and the combined uncertainties in the rate of potency loss and the determination of the potency at release. Knowing the rate at which potency is lost (from the least-squares regression analysis), a potency value at $t = 0$ is determined (the y-intercept at $t = 0$); uncertainties in (i) the rate of potency loss and (ii) the release potency assay are then added in. The combined uncertainty is given as a statistical multiplier (associated with 95 % confidence[5]) times the square root of the sum of the variances of the individual determinations for the rate of potency loss and the release potency assay. Thus, we have:

$$\text{Release Potency} = \text{Lower Potency Limit} + \hat{b}_1 \times t + t_{\alpha,df} \sqrt{(t \times s_b)^2 + (s_a)^2}$$

[5] Although other confidence limits might be chosen, 95 % is a generally agreed-to default value.

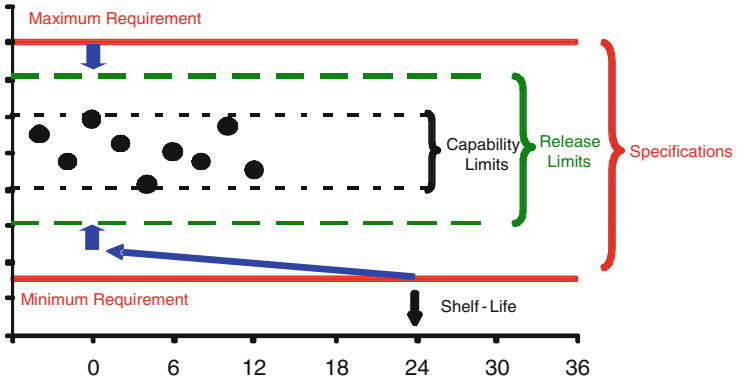

Fig. 14.6 Ideal relationship between the minimum release value and manufacturing window

where

$\hat{b}_1$ is the least-squares slope of the regression line
t is the shelf life
$t_{\alpha,df}$ is the critical t-value for the desired certainty with the degrees of freedom associated with the regression
s_b is the standard deviation of the regression line slope
s_a is the standard deviation of the release assay potency.

Similarly, an upper release limit can be calculated to assure (with 95 % confidence) that the true potency does not exceed an upper bound.

Ideally, the thus-calculated release potency value will embrace the range of the manufacturing capability. If, however, the lower release limit is not less than the lower value of the manufacturing capability, it will be necessary to either decrease the shelf life and recalculate a minimum release limit or, alternatively, reject those Lots where the manufactured value falls below the needed release value. This relationship between manufacturing capability and release value is illustrated in Fig. 14.6.

14.2.2.1 A Statistical Interlude

For purposes of establishing an estimate for the potency loss over time, as shown in Figs. 14.3 and 14.4, the mathematical expression describing the loss of potency, Eq. 14.1, was first transformed to a linear form, Eq. 14.2; the transformed data were then analyzed to provide estimates for k and $P(0)$. While the transformation provides an equation which can be fit using simple linear regression, this logarithmic transformation also generates data that are commonly more suitable for statistical modeling. An assumption of least-squares linear regression is that the variance is uniform across levels of response and that the error values (the differences between

the measured and fit values) are normally distributed. Potency measurements are frequently log-normally distributed. A log transformation of such measurements will therefore have attributes which satisfy the assumptions of the modeling.

In Figs. 14.3 and 14.4, in addition to providing the least-squares regression line, the confidence interval (CI) about that regression line was provided (the dashed lines in the Figures). The confidence interval represents the uncertainty in the estimation of the kinetics model. In this regard, the point in time where the lower confidence bound intersects the minimum potency requirement is the maximum time that we are assured (assured with probability ≥0.95) that the vaccine remains above the minimum potency (Fig. 14.5).

In selecting a set of time points for the determination of the rate of potency loss, time points presented by the ICH Guidance on stability, were used (these are generally 0, 3, 6, 9, 12, 18, 24, 36, 48 months). Although generally regarded by regulators and manufacturers as an acceptable set of time intervals, from the point of view of minimizing the error in the estimated shelf life, they are not always ideal. The equation for the lower confidence bound (at time t) is as follows:

$$\mathrm{LCB}(t) = \hat{y}(t) - t_{\alpha,\mathrm{df}} \cdot s \cdot \sqrt{\frac{1}{n} + \frac{(t-\bar{t})^2}{\sum (t_i - \bar{t})^2}}$$

where

$\hat{y}(t)$ is the estimated response from the linear regression equation,
$t_{\alpha,df}$ is a critical value from the *t*-distribution,
s is the regression means square error (an estimate of assay variability), and
$\bar{t}$ is the average time

The limiting factor in this equation is: $\frac{(t-\bar{t})^2}{\sum (t_i-\bar{t})^2}$. This is reduced by either decreasing the numerator or increasing the denominator. The numerator is minimized by concentrating time points near the expected shelf life, *t*. It will be shown later that the denominator is maximized by concentrating time points at the beginning and end of the shelf life period.

The point of showing that alternative designs may provide better estimates of the shelf life than ICH is not to argue against using ICH intervals. These are de facto regulatory intervals. Rather the point is to illustrate that better estimates can be achieved in stability studies through reduction of uncertainty (variability), which can be managed through stability study design. For example, this may be achieved by addition of extra data points near the expected time of expiry or past the expected time of expiry.

We have noted that the simplest plausible kinetic model for loss of potency would be a first-order kinetic model. However, the loss of potency for any particular vaccine may not follow first-order kinetics and more complicated kinetic models may be necessary. The first question we should address is whether a first-order kinetic model is *adequate* to describe the data and then whether a first-order kinetic

model is able to predict the potency at a later time point. To the first part of this question, statistical tests may be applied to the model, namely, we can ask whether the log-transformed data are linear. A common statistic r^2 (r-square) is not ideal for detecting nonlinearity. This is because r-square is impacted by a number of factors: (1) the steepness of the slope; (2) the variability about the regression (s); and (3) the variability (range) in the time points. The variability about the regression (s) is closest to capturing the aspect of graphical linearity, and may therefore be used to monitor linearity. Alternatively one may use residual plots or fit a model with a quadratic (curvature) term in the model. While a test of significance of the quadratic term would indicate statistically significant curvature, an equivalence approach using a range on the quadratic coefficient (equivalence margin) is better.

14.2.2.2 A Chemical Kinetics Interlude

At times, it may appear that it is not necessary to log-transform the data; that is, that the data are adequately fit by a simple linear model. Chemically, this would correspond to the vaccine losing potency via a zero-order kinetic model, wherein $P(t) = P(0) - kt$. Although zero-order reactions are rare, and generally derive from surface catalysis, the appearance of zero-order like kinetic behavior is not surprising as it mimics the initial portion of an exponential process, i.e., first-order kinetics. The similarity between a zero-order reaction and the initial portion of a first-order reaction is readily seen by expanding the exponential (see Eq. 14.2) in a Taylor series and noting the close correspondence of time values to a zero-order model when

$$e^{-kt} = 1 - \frac{kt}{1!} + \frac{(kt)^2}{2!} - \frac{(kt)^3}{3!} + \cdots$$

values of kt are small; see Table 14.2.

Although a zero-order model may be seen to adequately fit the available data, we note that such a model would be highly unusual from a chemical kinetics viewpoint, but, more importantly, that extrapolating zero-order kinetics to longer times, for which data are not available, would inappropriately lead to a markedly reduced estimate for the shelf life. This example also serves as a caution in overly relying on extrapolated data, even when the model appears to provide a reasonably good fit to the existing data.

Table 14.2 Values of $A_0 \exp(-kt)$ and $A_0[1-kt]$ as a function of kt

kt	$A_0 \exp(-kt)$	$A_0[1-kt]$
0.05	95.1	95.0
0.10	90.5	90.0
0.15	86.1	85.0
0.20	81.9	80.0
0.25	77.9	75.0

14.3 The Temperature Dependence of Reaction Rates and the Arrhenius Equation

The rates of chemical reactions are, in general, dependent on temperature, increasing with increasing temperature. It is a relatively well-accepted generalization, about which we will soon have more to say, that the rate of a reaction approximately doubles with each 10° increase in temperature. As changes in vaccine potency derive ultimately from chemical transformations, potency changes are also temperature dependent. This temperature dependence is generally described in terms of the Arrhenius[6] equation, shown below:

$$k = Ae^{-E_a/RT} \tag{14.3}$$

where

k is the reaction rate constant
R is the universal gas constant
T is the temperature in degrees Kelvin
E_a is the activation energy for the reaction
A is the pre-exponential factor (the units of A are those of the rate constant, k; for a first-order reaction, the units are time^{-1})

The values of E_a and A are considered constant for a given reaction, a reasonable approximation given the relatively limited temperature range that is generally investigated in kinetic studies of vaccines. Therefore, if E_a and A are known, the reaction rate at any temperature may be calculated. The values of E_a and A may be determined measuring the reaction rate at two different temperatures, generating two equations for the two unknowns. In practice, the reaction rate at a number of temperatures is determined and the [rate, temperature] data set are then analyzed by a least-squares fitting procedure to a linearized form of the Arrhenius equation, gotten by taking the natural log of both sides of Eq. 14.3; thus

$$\ln(k) = \ln(A) - \frac{E_a}{R} \cdot \frac{1}{T} \tag{14.4}$$

From a least-squares fit of $\ln(k)$ versus $1/T$, one obtains a slope, Ea/R, and intercept, $\ln(A)$; see Fig. 14.7. With the values of E_a and A in hand, as noted above, the rate constant at any other temperature may be calculated, either within or outside of the investigated temperature range; it is important to note, however, that the further outside of the investigated temperature range, the greater the uncertainty in the predicted reaction rate, as indicated by the 95 % confidence limits about the regression line shown in Fig. 14.7. It should be noted that the accuracy of the

[6] The relationship between reaction rates and temperature was developed by the Swedish physical chemist, Svante August Arrhenius (1859–1927).

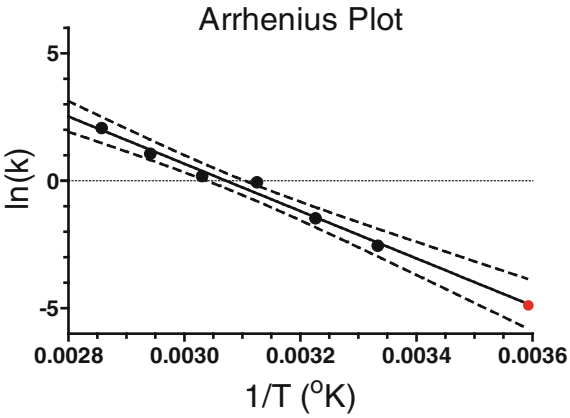

Fig. 14.7 A plot of ln(k) versus 1/T for a hypothetical vaccine displaying Arrhenius-type behavior. From such a plot, values of E_{act} and A may be obtained. An extrapolated rate, at a lower temperature, is shown in *red*

simple least-squares fit presented in Fig. 14.7 is premised upon determining the reaction rates at each temperature with the same degree of accuracy; since reaction rates decrease with decreasing temperature, increasing amounts of time are needed to monitor a fixed level of change (and hence constant error in the rates) as the temperature decreases. Because of time constraints, this is seldom done and more sophisticated statistical approaches should be considered, namely, attaching appropriate weighting factors to the various the data points.

As noted above, the statement is generally made that reaction rates increase approximately twofold for each 10° increase in temperature. Although a useful rough approximation, the actual increase in rate depends on the value of the activation energy and on the temperature range that is investigated. For two specific temperatures, T_1 and T_2, we may write the Arrhenius equation as:

$$\ln(k_1) = \ln(A) - \frac{E_a}{RT_1} \tag{14.5a}$$

and

$$\ln(k_2) = \ln(A) - \frac{E_a}{RT_2} \tag{14.5b}$$

Subtracting Eqs. 14.5a from 14.5b, we arrive at the expression:

$$\frac{k_2}{k_1} = \exp\left\{\frac{E_a(T_2 - T_1)}{RT_2T_1}\right\} \tag{14.6}$$

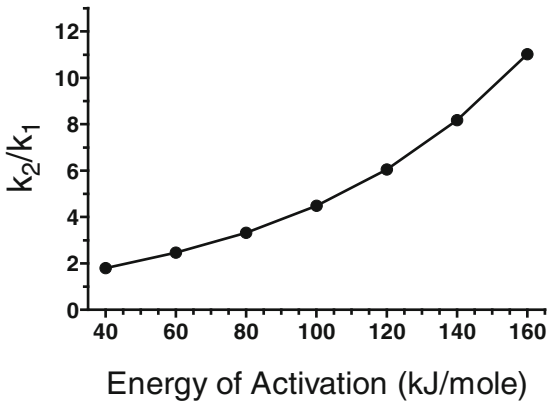

Fig. 14.8 The *n*-fold increase in reaction rate due to a 10° increase in temperature, from 5 to 15 °C, as a function of the energy of activation

The *n*-fold increases in rate (the value of k_2/k_1) as a function of selected values for the activation energy for an increase in temperature from 5 to 15 °C are displayed in Fig. 14.8. We also note that as the average temperature for the 10° interval increases (for example, from 100 to 110 °C vs. 5 to 15 °C), the *n*-fold change in rate will decrease.

As noted, the value of having established the Arrhenius parameters is that the reaction rate may then be calculated at any temperature. Thus, one may determine the Arrhenius parameters at high temperatures, where reactions occur rapidly, and then calculate the expected reaction rate at lower temperatures, where reactions occur slowly. Thus, one may determine in several months, what might otherwise require several years. This is illustrated in Table 14.3, where, for a first-order reaction, and a representative value of E_{act} and A, the change in rate is provided as a function of temperature; the change in the reaction half-life[7] is also provided in the Table.

The question may be raised why, given this potentially enormous saving in time, that accelerated kinetics are not used in establishing a product shelf life. There are two major reasons why this is not done. The first is a practical one. Due to experimental errors in determining the reaction rates at various measured temperatures, there is an associated uncertainty in the determined Arrhenius parameters and, accordingly, uncertainties in the extrapolated rates; the greater the extent of the extrapolation, the greater the uncertainty. In general, a sufficiently accurate shelf life can simply not be gotten by any practical use of the Arrhenius equation. There is also a theoretical component, in that Arrhenius behavior may only be approximately followed; that is, the Arrhenius parameters may themselves be temperature dependent. Although the Arrhenius equation may not be used for setting a shelf life, it does have a number of valuable uses. First, it is of considerable values in developing a vaccine formulation, where the primary concern is an increased shelf

[7] The half-life for a reaction is the time taken for one-half of a reactant to be consumed; for a first-order reaction, the half-life, $\tau_{½}$, is equal to ln (2)/k.

Table 14.3 Rate constants and associated half-lives for a reaction with an E_{act} of 90 kJ/mole and pre-exponential factor, A, of 1×10^9 s^{-1}

Temperature (°C)	Rate constant, k	Half-life, $\tau_{½}$ (months)
5	3.2×10^{-2}/month	21.6
15	12.3×10^{-2}/month	5.6
25	43.4×10^{-2}/month	1.6
35	131.4×10^{-2}/month	0.5

life. The trends that are observed at higher temperature will, in nearly all cases, persist at lower temperatures. Thus, a formulation that provides greater stability at 35 °C will provide a greater stability at 5 °C. A comparison between two formulations may be carried out in several months at 35 °C, whereas the comparison might require several years at 5 °C. Second, an accelerated stability study is of considerable use in establishing product comparability following a manufacturing change. If the vaccines are comparable at a higher temperature, for example 35 °C, they are very likely to be comparable at 5 °C. Again, observing a change and showing that the change is comparable in the two vaccines, is more readily accomplished at the higher temperature(s). Finally, from values of the Arrhenius parameters, one may estimate the extent of potency losses that may occur during excursions from given storage conditions.

14.4 Stability Studies and the Product Life Cycle

The goals of stability studies vary during the product life cycle. Prelicensure, at the earlier stages of product development, the goals of stability studies are related to knowing stability over the course of clinical trials (and thus knowing the potency of vaccine that clinical trial subjects will receive and have received) and developing a vaccine formulation that maximizes product stability. At the time of licensure, the goals of stability studies are related to establishing a shelf life and release specifications as well as demonstrating manufacturing consistency. Postlicensure, the goals of stability studies are to demonstrate manufacturing consistency as evidenced by similar potency profiles over time (i.e., the annual stability studies) and to support the comparability of the vaccine following manufacturing process changes.

Prelicensure and licensure. As noted above, at the prelicensure stage of development, the goals of stability studies are to describe the vaccine stability over the course of clinical trials, to develop a formulation that maximizes the stability of all active components of the vaccine, and to demonstrate manufacturing consistency through a consistent stability profile. Finally, a shelf life, under defined storage conditions, must be established at the time of licensure.

As stated, vaccines are formulated to maximize stability. For this purpose, namely, developing a formulation that maximizes vaccine stability, it is useful to carry out the stability studies at temperatures that are higher than those that are intended for storage. At the higher temperatures, changes in vaccine potency are

detected more rapidly and, generally, more precisely. The trends in stability at the higher temperatures will, with rare exception, remain at the lower temperatures; that is, the more stable formulation at the higher temperature will be the more stable formulation at the lower temperature (as illustrated in Fig. 14.9). Indeed, based on the previous discussion of *n*-fold changes in rates with temperature, the ratio of rates would be expected to become slightly more disparate at lower temperatures. Although it would be sufficient to study the various formulations at a single higher temperature, it is, nonetheless, useful to study a range of temperatures, both to demonstrate that Arrhenius behavior is followed over the studied temperature range (for example, 15–45 °C) as well as to actually determine the Arrhenius parameters, E_{act} and A, so as to be better able to extrapolate kinetic behavior to the lower temperatures (and, thus, an estimate for the expected shelf life) and to be able to estimate stability at intermediate temperatures (and thus be able to determine the consequences of unexpected temperature excursions).

During the final phases of clinical development, it is necessary to demonstrate manufacturing consistency, for which stability studies play an important role. In this regard, it is necessary to select the best metric for demonstrating consistency. A relevant metric of consistency might be a comparison of the slopes of the regression lines from the stability studies of the various vaccine Lots (generally three Lots are evaluated for manufacturing consistency). While this seems a reasonable choice of metric, a question then arises as to the basis for comparing the thus-obtained slopes. The slopes might be compared for poolability as outlined in the ICH Guidance (Q1E). Alternatively, the slopes might be evaluated for equivalence within a predefined margin. A fuller discussion of these two approaches is provided in the next section on postlicensure stability studies.

At the time of licensure, a shelf life for the vaccine, under defined storage conditions, needs to have been established. This shelf life must be established from data gathered under the intended storage conditions and over the period of the intended shelf life. Some extrapolation of the data in establishing the shelf life may be warranted; however, further data to justify the extrapolation would then be needed. In practice, to reduce the error in estimating the shelf life, having additional

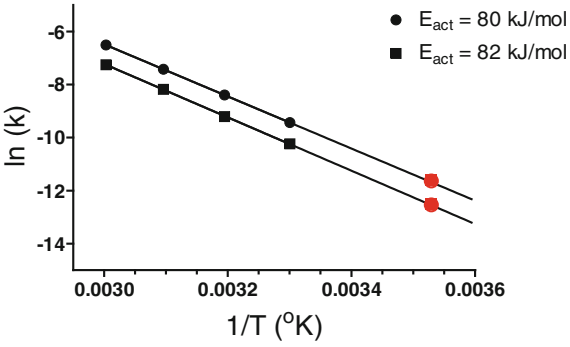

Fig. 14.9 Arrhenius plots of two hypothetical vaccine formulations having slightly different energies of activation for a process leading to a loss of potency

data points near the intended expiry period (additional to those indicated by the ICH Guidance) as well as having data from time points in excess of the intended shelf life are useful.

Postlicensure. Postlicensure, annual stability studies are required. Each year randomly selected vaccine Lots are required to undergo stability testing over the product's shelf life. The goal of these studies is primarily to demonstrate that the stability profile of the postlicensure produced vaccine is comparable to that of vaccine produced at the time of licensure and to provide ongoing support for the shelf life that was established at the time of licensure (the precise goals of postlicensure stability studies have not been delineated in either regulation or guidance).

With this goal in mind—demonstrating comparability to previously manufactured Lots—the question may be raised as to what might constitute the best metric for defining comparability. One commonly employed approach is to simply demonstrate that the vaccine maintains its potency at all measured time points (generally, those stated in the ICH Guidance), falling between the established upper and lower limits. This is akin to the "compliance model" that was previously mentioned. While this may help to confirm the appropriateness of the designated shelf life (assuming that the potency for all time points remain within the upper and lower bounds), several problems attend this approach. The first is that an actual change in stability, albeit one that is consistent with the assigned shelf life, may be overlooked. From a cGMP perspective, this is undesirable. If there is a change in the stability profile, it should be known and further studied to determine its root cause and potential effect on other vaccine attributes.[8] The second problem relates to an observed out-of-specification value, a situation that becomes increasingly more probable as the expiry period is approached. The question may be posed whether the product actually is out of specification because of a change in the product stability profile, or, alternatively, appears to be out of specification because of the variability in the potency assay. The potency at that time point or a subsequent time point may be remeasured. If the value for the subsequent measurement is within specifications, the problem of deciding which measurement is "correct" remains. The results of additional measurements may help to assure certainty in a decision; however, such additional testing may require a significant amount of time, especially if the potency assay is animal based. The point to be made, however, is that there is additional stability information that is contained within the previous potency determinations and these should be brought to bear on this latter problem.

A more meaningful metric of comparability may be to compare the rates of potency loss of prelicensure Lots of vaccine (and that were used in efficacy studies) with those being currently marketed. In such an exercise, the slopes from the respective regression analyses would be compared. This comparison might be made in different manners. One might, for example, employ methods and acceptability criteria to determine if the data were "poolable", i.e., the slopes were parallel, as in

[8] Although potency tests are designed to indicate the expected effects in the intended recipient population, they are, in reality, often imperfect predictors.

the ICH Q1E Guidance. Alternatively, the data might be the subject of an equivalence test. We would like to comment on these two methods. The above-mentioned test for poolability, while often used, suffers from a significant drawback. On the one hand, for regression analyses having scant and variable data, one is not able to conclude that the slopes are different and hence the data are regarded as poolable; on the other hand, for regression analyses with abundant and precise data, minor differences, which may not be practically significant (as opposed to statistically significant) from a consideration of vaccine use, will be detected and the data will be considered noncomparable. These possibilities are illustrated in Fig. 14.10.

Thus, there is seen to be a risk of masking real differences due to data variability or inadequate study design or the risk of highlighting an insignificant difference when the assay is precise or there is an adequate study design. Such an approach appears counterintuitive in that there is a reward for excess variability and a penalty for good precision. An alternative approach would be to evaluate the equivalence of the slopes within predefined limits, as illustrated in Fig. 14.11.

Postlicensure, various manufacturing changes are made, for which it is necessary to demonstrate product comparability. Stability studies form an integral part of that comparability study. The basic principles that have been given above for evaluating annual stability studies will equally apply to comparability studies following manufacturing changes. Of course, some changes are made to enhance product stability and it would then be necessary to demonstrate a superiority.

Returning to the theme of demonstrating comparability, either in the annual stability studies or following manufacturing changes, and accepting use of the slope as a metric for comparability, the question may be posed as to whether that comparability should be demonstrated at the intended storage temperature or at an elevated temperature, or both. In general, stability studies at elevated temperatures form part of the comparability assessment following manufacturing changes; data under the intended storage conditions are also utilized and do form the primary basis for the assigned shelf life. In contrast, the annual stability studies are carried out under the licensed storage conditions, generally 5 ± 3 °C. For a fixed time increment, there will be a greater change in vaccine potency that is observed at a

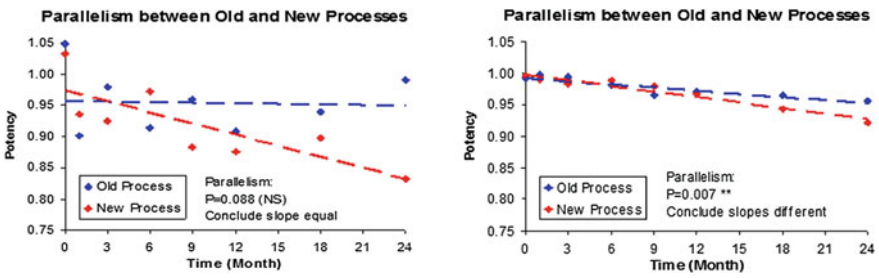

Fig. 14.10 Hypothetical time, potency data sets wherein a significant difference in rates of potency loss may be masked by limited, imprecise data (*left hand panel*) or minor difference in potency loss is considered statistically significant because of precise data (*right hand panel*)

Fig. 14.11 A comparison of the slopes obtained from linear least-squares regression analyses of two vaccine Lots. Having selected a Δ-value that is considered as an allowable variation in the rate of potency loss relative to a reference value (indicated by the *shaded rectangle*), the confidence intervals of the test samples may be compared. The *blue line* depicts a comparable value in that the depicted confidence interval about the point estimate for newly measured slope is within the agreed-to range; the *red line* depicts a noncomparable value in that the confidence extends outside this range

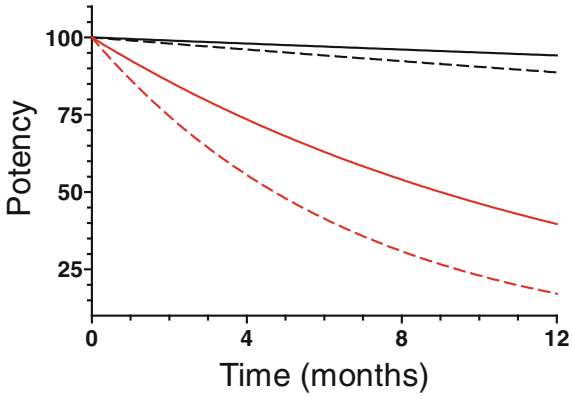

Fig. 14.12 Depiction of a twofold change in the rate of potency loss as might occur at a low temperature (*black lines*) and as that twofold rate change might appear at an elevated temperature, corresponding to a 20° increase in storage temperature (*red lines*)

higher temperature than at a lower temperature; for example, for a vaccine studied at 5 and 25 °C, we might expect an approximately 10 to 20-fold decrease in potency at 25 °C relative to that observed at 5 °C. Given that the error in measuring potency is the same for all samples, e.g., whether there has been a 5 % loss in potency or a 50 % loss in potency, then the ability to detect a change in the rate of potency loss is greater for the sample that has undergone the greater change. An argument might therefore be advanced that annual stability studies might best be carried out at elevated temperatures, provided the goal of such studies is to monitor the consistency of manufacture (as evidenced by the consistency of the stability profile). This is illustrated in Fig. 14.12. Consider a vaccine that had a modest loss in potency (approximately 6 %/year; $k = 0.005$ month^{-1}) and that, due to an unsuspected manufacturing variation, the rate doubled.[9] Given a modest uncertainty in potency

[9] In the above example, activation energies, E_a, of 94.28 and 92.67 kJ/mole (with a pre-exponential factor of 10^{-9}) were chosen to mimic the kinetics at 5 °C and to subsequently calculate the rates at 25 °C.

determinations, it would be difficult to detect a twofold change within 1 year at 5 °C, as is illustrated by the solid and dashed black lines in Fig. 14.12. For this hypothetical vaccine having the Arrhenius parameters presented in Footnote 6, a 20° increase in temperature would lead to an approximately 15-fold increase in rate as depicted by the solid and dashed red lines in Fig. 14.12. The twofold change in rate at 5 °C would be more likely to be detected at the increased temperature within the 12-month period and, possibly, sooner. In summary, a reasonable case may be advanced for conducting annual stability studies at elevated temperatures, either in lieu of studies at the storage temperature or in addition to them.

14.5 Concluding Remarks

The principle goal of all stability studies is to ensure vaccine quality throughout the vaccine's shelf life. This entails having established an acceptable lower limit for potency and knowing the rate at which the vaccine loses potency, and, finally, accounting for the uncertainties in the determination of potency at release and the rate of potency loss. Done properly, such stability studies reduce the risk to the vaccine recipient of receiving a subpotent vaccine while reducing the risk to the manufacturer of rejecting suitably potent vaccines.

Stability studies begin with establishing an appropriate mathematic model for the data; in this chapter, we have exemplified various principles using a first-order kinetic model, the simplest feasible model for loss of vaccine potency; we have noted that more complex models may be needed—for example, biphasic kinetics or second order kinetics (as might be encountered in protein oxidation). Various statistical tests to determine the adequacy of the model are then employed—for example, testing for linearity of the log-transformed data or evaluating that the residuals are normally distributed. Given the adequacy of the kinetic model, the uncertainties associated with the key parameters can be determined; confidence limits on the slope of the regression line can be determined and, together with the known variance of the release assay, can be used to establish a lower release limit and shelf life for the vaccine.

Although the shelf life of the vaccine needs to be established at the intended storage temperature, the use of accelerated degradation studies (at higher temperatures than intended for storage) were seen to be useful in developing a vaccine formulation to optimize stability and in situations where it is necessary to establish the consistency or comparability of the vaccine's stability profile. From a study of the vaccine's degradation rate as a function of temperature, the Arrhenius parameters for the loss of potency can be determined and used to estimate vaccine stability at other temperatures (although it is not used in establishing the shelf life).

The importance of statistical considerations have been emphasized in this chapter. Such considerations are necessary for understanding the risk that is associated with interpreting stability data and establishing a lower release limit. A basic principle in evaluating stability that has been emphasized in this chapter is that a set

of measurements are collected over time to provide estimates of stability parameters (rates of potency loss and the associated uncertainty in those rates) and not to assess compliance of individual data points with specifications. Measuring uncertainty in connection with stability estimates places an emphasis on the design of experiments together with the goals of the study, for the end purpose of delivering high quality vaccines.

14.6 Appendix

See Table 14.4.

Table 14.4 Time and potency data for Lots 1, 2, and 3 as shown in Fig. 14.2, prior to logarithmic transformation

Time (months)	$TCID_{50}$ for Lot 1	$TCID_{50}$ for Lot 2	$TCID_{50}$ for Lot 3
0	5082	5882	5575
3	5201	5893	5431
6	4693	4177	4367
9	4376	4193	4118
12	4844	4148	4718
18	3885	3836	4335
24	4073	3346	3783
30	2507	3309	3268
36	2115	2684	2208