A Scientific and Statistical Analysis of Accelerated Aging for Pharmaceuticals: Assessment of Error Estimation Methods

Kenneth C. Waterman*, Jon T. Swanson, Blake L. Lippold
FreeThink Technologies, Inc. Branford, CT

Methods of predicting the uncertainty in ambient shelf-life (probability of passing a specification limit at a given storage time) from accelerated data using a two-step process (estimation of isoconversion then Arrhenius fitting) were evaluated. An extrema method is shown to provide a reasonable estimation of the isoconversion regression interval in the general case of a constant relative standard deviation in degradant levels, where there is a minimum standard deviation equal to the limit of detection. The rate distribution used for the Arrhenius calculation was determined either from a distribution of isoconversion times or from a distribution of degradant values at the means of the isoconversion times. Least-squares fitting to the Arrhenius equation in both linear and exponential forms were calculated. The non-linear fitting weights higher temperature points more heavily than the linear fitting and therefore carries greater scientific risk when there is non-Arrhenius behavior. Since the linear fitting is also less computationally intensive, it is preferred over non-linear fitting, even with weighting of the latter.

INTRODUCTION

In Part I of this series, we examined the ability of different calculation methods to accurately determine shelf-life at ambient conditions when extrapolating degradation behavior from high temperature (accelerated) conditions[1]. We focused on somewhat complex, but common, kinetics, where drug degradation involved both primary and secondary degradation processes. In that paper, it was shown that in order to achieve accurate projections in the majority of cases, it was necessary to use an isoconversion approach. In this process, the degradation as a function of time at each temperature was used to estimate the time to reach the specification limit. These times, in turn, were used to determine isoconversion rates, which were used with the Arrhenius equation to estimate ambient shelf-life (time to reach the specification limit). While instances that follow simple kinetics could use alternative methods for calculation of the ambient behavior (i.e., bypass the intermediate calculation of the isoconversion time), assumptions of simple kinetics are not justified in a large percentage of cases. Estimations using an isoconversion rate calculation will be as accurate as the alternative methods when degradation kinetics are simple, but significantly superior when they are not.

In the present paper, we examine different methods for calculating probability distributions for ambient behavior based on model projections from accelerated conditions. Unlike the case with accuracy using model data, where there is an objective assessment of how close an answer is to the true value, we can only compare different methods to each other based on scientific and statistical judgment. We hope that a calculated probability of passing with a given model data set will correspond to the actual likelihood that such a system will indeed pass. Even if a statistical model could be made that were perfect in its predictability, in practice, the ambient conditions are subject to variability themselves which complicate the way in which the models are used. These sources of ambient variability include the following: analytical errors associated with the measurements, lot-to-lot variability, and the range of actual storage conditions nominally labeled the same. It should be noted that the ICH guidelines[2] allow for ±2°C and ±5%RH. This means that for an average solid drug product having an activation energy of 29 kcal/mol (121 kJ/mol) and a B value (i.e., relative humidity, RH, sensitivity) of 0.04, the rate at 28°C/60%RH would be 2.7 times slower (shelf-life 2.7 times longer) than the rate at 32°C/70%RH (open), yet both would still be within the guidelines for controlled 30°C/65%RH ICH storage conditions.

In Part I, we discussed the role that RH plays in determining reaction rates for solids. The RH dependence can be complex, since it depends on the packaging of the material and is therefore likely to change with time[3]. As with the first paper, we again make the simplifying assumption that the RH remains constant with the implication that different statistical treatments will be similar, though more complex, when considering RH changes as a function of time. In the present paper, we also make the simplifying assumption that the actual temperatures and times at the accelerated conditions are exact, with all the uncertainties condensed into the degradation measurements.

Determining the distribution of model outcomes involves a propagation of errors beginning with the uncertainty estimates for the individual data points (amount of degradant formed or loss of active). In an ideal world, repeated measurements would be made at each time, and each of those individual points would be used in the calculations. In practice, a sufficient number of repeats are seldom carried out to allow use of this procedure. It is therefore necessary to assume some distribution of error in the degradant or active level about the time points used in an accelerated study.

While early work in accelerated aging mostly involved calculations based on a loss of potency[4-8], current shelf-life determinations for most drug substances and drug products are instead limited by formation of degradation products. Potency variability in many dosage forms is dominated by content uniformity (variability) of the dosage forms themselves (e.g., tablets, capsules). As such, the error bars in the values often represent fixed errors. With purity analyses, the precision of degradant levels is more often limited by sample preparation and measurement than by the samples themselves. In this case an assumption of a constant relative standard deviation (RSD) is more accurate. This assumption is necessarily tempered by a minimal error bounded by a limit of detection (LOD). We therefore make the assumption, based on the most common scientific situation, that the error bars for degradant levels at each time point are represented by a normal distribution centered at the entered value with a standard deviation being the larger of the RSD or LOD. Since this is a more difficult statistical situation than an assumption of a constant error, it is likely that any methods that are appropriate for such cases will also apply to the simpler case of a constant error or having no minimum error.

The first step in shelf-life determination is calculating the distribution of rates at each condition based on a distribution of isoconversion times. We can consider two sources of error in fitting the degradant versus time data. The first is the standard error of the fit, which measures the variance from the best fit line. Determining this value is complicated when the maximum of the RSD and the LOD is used as an estimate of the uncertainty of the degradant at each time, rather than true individual measurements. The second source of error comes from the range of best fit lines than can be produced based on different samples of the population. This is the error of the mean or the confidence interval (CI). If all degradant values are independent, normal and have the same standard deviation (σ) the confidence interval can be expressed in closed form (Equation 1)[9]:

$$CI = \sqrt{\frac{1}{n} + \frac{(d_o - \bar{d})^2}{\sum (d_i - \bar{d})^2}} \qquad (1)$$

where d_0 is one of the measured points, $\overline{d}$ is the average value for the degradant level, and n is the number of points. If a large number of points are sampled, the estimate of the error is zero at the mean value of d, since CI is an estimate of the error of the mean. When estimating the uncertainty in a point not measured, the error at that point will be at least as large as the error introduced from the noise in the data as reflected in the regression interval, RI, shown in Equation 2[9]

$$RI = \sigma\sqrt{1+\frac{1}{n}+\frac{(d_p-\bar{d})^2}{\sum(d_i-\bar{d})^2}} \qquad (2)$$

where d_p represents a new point whose value we are predicting. This value will still be a minimum at the mean, but the minimum for very large numbers of samples is now the estimate of the standard deviation σ. As d_p is moved further from the mean, the values of the CI and RI converge.

In the present paper, we examine three methods for calculating the uncertainty in isoconversion. In the first method, we use a stochastic (Monte-Carlo) procedure[10]. This method estimates the CI; it does not provide an estimate of the RI. In this method, degradant values (assumed to be normally distributed) are sampled at each time point and used to produce a best fit line. Repeated sampling produces a distribution of possible values for the isoconversion time. The resulting isoconversion distribution will not, in general, be normally distributed. One potential approach to working with such non-normal distributions can be derived from the fact that the distribution in degradant (or potency) at any specific time can be calculated and will be normal. At the "true mean" isoconversion time (i.e., the average value when an infinite number of observations are made), a normal distribution of degradant levels can be determined. This distribution can be converted to a distribution of rates by dividing each value of the degradant distribution by the mean isoconversion time.

In the second estimation process, a non-stochastic method is used. Rather than use random points, representative points are used, specifically at plus and minus one standard deviation. This will generate 2^n fits to the data (where n is the number of time points used). This will in turn generate isoconversion times and rates corresponding to one standard deviation from a normal distribution (or the 68% confidence interval for a non-normal distribution). Again, this is an estimate of the CI, rather than the RI. The results obtained by this method are identical to the previous method except that only the values at plus or minus one standard deviation are determined, rather than a complete distribution. As with the stochastic approach, the distribution in isoconversion times will not in general be normal; however, the distribution in degradant (potency) levels at the mean isoconversion value will be.

The third method for estimation is an "extrema method", which is similar to the non-stochastic method described above in that representative points at plus and minus one standard deviation are used to generate 2^n lines. The difference is that either the line to cross the specification limit corresponding to the minimum of the 2^n isoconversion times is used to define a standard deviation for the isoconversion distribution (assumed to be normal about the mean) or the line to cross the mean isoconversion time corresponding to the maximum of 2^n degradant values is used to define a standard deviation for the degradant distribution. While this is not a RI calculation, it does insure that predicted points do not have associated errors less than that of the measured values. Likewise, the standard deviation estimated at the entered degradant/time points is close to the entered values. Figure 1 shows how the extrema method is used to estimate uncertainty in an example with two specification limits.

Figure 1

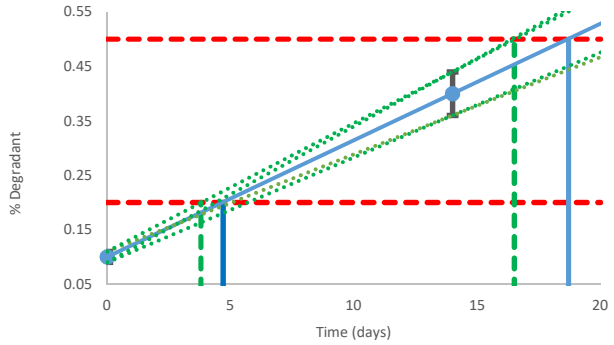

One of the methods of estimating the error in isoconversion involves using a one-sided extrema process illustrated in this figure. In this case, the zero time point data (measured value of 0.10) has an error bar (one standard deviation) of 0.01%, while the 14-day point (measured value of 0.40%) has an error bar of 0.04% (i.e., both have a 10% relative standard deviation). The no-error line to two specification limits (0.20 and 0.50%) is reached at 4.7 and 18.7 days, respectively. Four extrema lines are generated from the error bars. The shortest time for an extrema line to hit the specification limits are at 3.8 and 16.4 days for 0.20 and 0.50% degradant, respectively. The error bars for the isoconversions are estimated (as a normal distribution) as the difference between the no-error isoconversion and the shortest time isoconversions, in this case, 0.9 and 2.3 days, respectively.

Figure 2

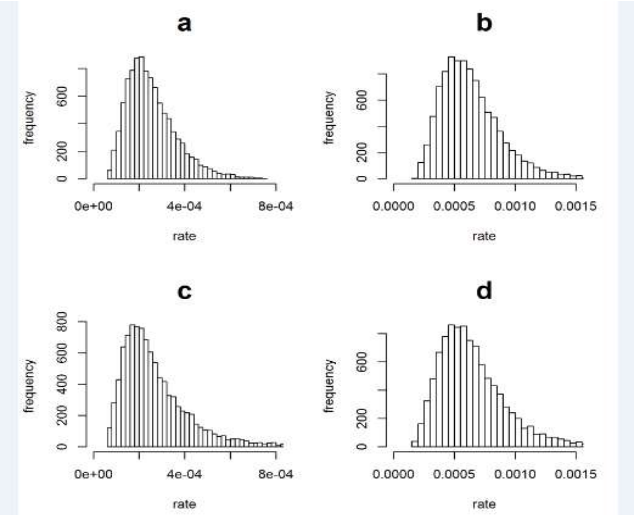

Comparison of rate distributions based on three-temperature (60, 70, 80°C) accelerated aging with a fixed relative standard deviation (RSD = 10%), minimum error (limit of detection, LOD = 0.02%) and an activation energy of 25 kcal/mol (104.67 kJ/mol). A single non-zero time point (at 10 days) was used at each temperature. Degradant values were chosen such that the 60°C point exactly equaled the specification limit in each case. Distributions were calculated for two specification limits to project to 25°C rates: a) using an isoconversion time distribution with a 0.2% specification limit; b) using an isoconversion time distribution with a 0.5% specification limit; c) using a normal degradant distribution about the specification limit of 0.2% at the zero-error isoconversion time; d) using a normal degradant distribution about a specification limit of 0.5% at the zero-error isoconversion time.

With error estimates for the isoconversion rates determined (whether normally distributed or not), the next step in the error propagation is to determine the fit to the Arrhenius equation. Two

methods of fitting are considered; with all error estimates determined using a Monte-Carlo simulation: a linear, least-squares fit to the logarithmic form of the Arrhenius equation; and a nonlinear least-squares fit using the exponential form of the Arrhenius equation 5,8.

Once a distribution of rates is determined for the behavior at room temperature, the probability of passing at a given time or the uncertainty interval around the mean predicted shelf-life can be determined. In the present paper, this was done in two ways: assuming that the distribution is sufficiently normal and computing probabilities from the normal distribution, and determining the cumulative probability distribution (CPD) from the actual distribution and computing probabilities from that distribution.

METHODS

Degradant Distributions. Degradant distributions were calculated in three ways: fixed standard deviation (fixed), relative standard deviation (RSD), and RSD bounded by the limit of detection (LOD). The fixed method used a standard deviation of 0.02% for each degradant value. The RSD method used 10% of the degradant value as the standard deviation. The RSD bounded by the LOD method used the larger of the RSD or 0.02% (LOD) for the standard deviation. For stochastic calculations, the distribution of degradant was represented by a normal distribution of 10,000 points with a mean of the provided degradant value and a standard deviation calculated by one of the listed methods. For non-stochastic calculations, the distribution was represented by values at plus or minus one standard deviation from the provided degradant value.

Stochastic Isoconversion Calculations. A matrix of the degradant values was generated from the degradant distributions calculated as described above. This provided a matrix of n x 10000 values where n was the number of time points. The time values were assumed to have no error. A linear least squares fit was then done with the degradant values as the dependent and time as the independent variables. This provided a set of 10000 coefficients, which were then used to calculate the distribution of isoconversion times at the given specification limit (the specification limit divided by the slope of the fit line).

Non-stochastic Isoconversion Calculations. A matrix of the degradant values was determined by generating the 2^n permutations of plus and minus one standard deviation of the degradant at each time point. A linear least squares fit was done to obtain a set of 2^n coefficients. These coefficients were then used to solve for either the distribution in degradant (intersection with the zero-error isoconversion time) or isoconversion (intersection with the specification limit).

Extrema Isoconversion Calculations. The extrema method repeated the non-stochastic isoconversion calculations to produce the distributions of degradant or isoconversion. Rather than use the full distribution, a standard deviation was calculated. For isoconversion, this is the mean isoconversion time minus the minimum value of the isoconversion distribution. For degradant, this is the maximum value of the degradant distribution minus the specification limit. These values were then converted to normal distributions in isoconversion or degradant (using the mean isoconversion or the specification limit as the mean, respectively).

Linear Shelf-life Calculations. The distribution in rates was calculated either using an isoconversion distribution or a degradant distribution. The rates were obtained from the isoconversion distribution by dividing the specification limit by the isoconversion distribution. The rates were obtained from the degradant distribution by dividing the degradant distribution by the mean isoconversion time. The Arrhenius parameters were then calculated from a linear least-squares fit of the linear form of the equation (Equation 3).

$$lnk = lnA - \frac{E_a}{R}\frac{1}{T} \qquad (3)$$

The rate at 25°C was calculated from this equation and converted to a shelf-life by dividing the specification limit by the rate.

Non-linear Shelf-life Calculations. The distribution in rates was calculated in the same manner as for the linear shelf-life calculations. The Arrhenius parameters were then calculated from a non-linear least-squares fit of the equation (Equation 4).

$$k = Ae^{-\frac{E_a}{R}\left(\frac{1}{T}\right)} \qquad (4)$$

The rate at 25°C was calculated from this equation and converted to a shelf-life by dividing the specification limit by the rate.

Cumulative Probability Estimates. Estimates of the value of a distribution at 15.86553%, 50% and 84.13447% confidence (corresponding to +1 SD, the median and -1 SD of a normal distribution) were calculated by sorting the distribution values and returning the value at the appropriate points. Where the exact value fell between two points, a weighted average of the two closest values was used.

All calculations were done using the R statistical package[11]. Normal distributions were determined using the function *rnorm*. Linear fits were determined using the function *lm*. Non-linear fits were determined using the function *nls*. An initial guess for the *nls* function was generated using a linear fit of the data (*lm*). Histograms were generated using the function *hist*.

RESULTS AND DISCUSSION

I. Isoconversion Error

While Equations 1 and 2 provide closed forms for calculating the confidence and regression intervals, the situation is more complicated for real stability studies. Degradant values are typically measured at just a few time points, and the measured standard deviation at each point is generally not the same. The methods outlined above were evaluated for handling these more complex situations.

In the simplest case, three time points are used that exactly follow zero-order kinetics (0, 0; 5, 0.2; 10 days, 0.40% degradant) with the error distribution assumed to be the same at each point (standard deviation of 0.02%). We can calculate the isoconversion error bars using the five methods described previously, the confidence interval (CI) and regression interval (RI) from the formulae shown earlier (Eqs. 1-2), and using stochastic, non-stochastic and extrema methods and sampling the degradant distributions. The results of these calculations are shown in Table 1 for both the interpolation region (at 5 days) and extrapolation region (40 days) for comparison. As expected, the CI calculated by the formula, the stochastic and non-stochastic methods produce the same results (within the limits of the sample size of the stochastic method). The CI at the isoconversion time produces an uncertainty in the value that is considerably less than the actual measured uncertainty at that point. This argues for use of the RI as the more appropriate and more conservative indication of the uncertainty. Unfortunately, calculation of the RI is complex when one cannot assume the same standard deviation at all points. Also, the RI will have larger estimates of the error at the other measured points (0 and 10 days) than was entered. One might thus also consider taking the larger of the standard deviation and the confidence interval, which has the property of providing an estimate close to the entered values in the interpolation region of the fit and an estimate close to the CI in the extrapolation region. The extrema method is a variation of this idea. As can be seen in the table, the extrema method provides results that are closer to the RI in the interpolation region than the other methods. In the extrapolation region, the extrema method is more conservative than the other methods and more conservative than the RI. This conservative estimation of extrapolated prediction intervals will be important when the potential for non-linearity in the results is considered.

Table 1

Calculation Method	5-Days (Interpolation)	40-Days (Extrapolation)
Formula Regression Interval	0.0231%	0.1017%
Formula Confidence Interval	0.0115%	0.0997%
Stochastic	0.0116%	0.0989%
Non-Stochastic	0.0115%	0.0997%
Extrema	0.0200%	0.1467%

Using a fixed standard deviation of 0.02% with perfect linear data (0, 0.00; 5, 0.20; 10 days, 0.40% degradant), prediction intervals for the degradant level at interpolated and extrapolated conditions are calculated using different methodologies.

We next examine the same data set, but add the more realistic assumptions of either a fixed RSD, or an RSD with a minimum error equal to the LOD. In these cases, we cannot readily calculate either the CI or RI by equations 1 and 2. The results using the remaining methods are shown in Table 2. As seen in the table, several observations can be made. In all three methods, the impact of the minimum error is significant. This emphasizes the importance of including this term in any calculations. If no minimum error is assumed, the extrapolation error bars are similar for all three methods. When a minimum error is assumed, only the extrema method provides an interpolated error bar which is greater than this minimum error. This means that with the stochastic and non-stochastic calculations, the error bar dips below the LOD, which would seem hard to justify scientifically.

Table 2

Calculation Method	Minimum Error	5-Days (Interpolation)	40-Days (Extrapolation)
Stochastic	0.00%	0.0147%	0.1524%
	0.02%	0.0163%	0.1655%
Non-Stochastic	0.00%	0.0149%	0.1535%
	0.02%	0.0163%	0.1660%
Extrema	0.00%	0.0200%	0.1600%
	0.02%	0.0267%	0.2233%

Using perfect linear data (0, 0.00; 5, 0.20; 10 days, 0.40% degradant), prediction intervals for the degradant level at interpolated and extrapolated conditions are calculated using different methodologies. In this case, the standard deviation at each measured point is either assumed to be a constant percentage (relative standard deviation of 10%), or assumed to be a constant percentage (relative standard deviation of 10%) with a minimum error of either 0.0% or 0.02%, the latter corresponding to an assumed limit of detection for the degradant.

In cases where the degradation data do not perfectly fit the regression line, these methods yield the same estimated error as when the data are a perfect fit. There can be circumstances when this means that the fitted value at a measured time point will not overlap the standard deviation at that point. This can be interpreted to mean that either the model used to fit the data is not correct or that the estimation of the standard deviation at that point is not correct. While this potential issue can be visually identified in graphs of the data (with error bars) when the measured data extend to the specification limit, the situation can be difficult to recognize in situations where the specification limit is only reached by extrapolation. As an example of this situation, consider the case where a primary degradant degrades to secondary products. With exact data based on the rate constants (k_1 = 0.000113%/d and, k_2 = 0.011250%/d), the primary degradant will curve [1]. A linear fit to the data is by all accounts quite good (R^2 = 0.998), yet provides an estimation of the mean isoconversion time equal to 51.6 days compared to the true value of 62.0 days (k of 0.0097%/d versus the true k of 0.0081%/d). The inability of simple methods to distinguish a true linear behavior from one that only appears linear is one of the reasons that designing experiments to provide degradant levels at the specification limit is important for accuracy. Nonetheless, there will undoubtedly be situations where some or all of the accelerated conditions require extrapolation to estimate isoconversion. The confidence of the predictions should therefore take some account of the potential for the model used for fitting the data (e.g., linear) being in error, especially when there is significant extrapolation. It is worthwhile to examine the above example to see which error bar estimates include the true value within the uncertainty interval. These calculations are shown in Table 3. As seen in the table, only the extrema method provides an uncertainty estimate which is near the true rate at one standard deviation from the mean.

Table 3

Method	Mean Rate (%/d)	Standard Deviation (%/d)
Stochastic	9.67×10^{-3}	0.94×10^{-3}
Non-Stochastic	9.68×10^{-3}	0.95×10^{-3}
Extrema	9.69×10^{-3}	1.46×10^{-3}

Assuming a primary degradant that subsequently undergoes secondary degradation (k_1 = 0.000113%/d and, k_2 = 0.011250%/d at 50°C, E_a = 25.0 kcal/mol; time points at 0, 3, 7, 14 and 28 days), prediction intervals for the degradant formation rate at a specification limit of 0.50% are calculated assuming a relative standard deviation of 10% with a minimum error equal to the the limit of detection (0.02%). The true rate is 8.1 X 10^{-3}%/d, which means that only the extrema method includes the true rate within approximately one standard deviation.

The final issue to be addressed with respect to the precision of the isoconversion estimation is the distribution. While it is assumed that there is a normal distribution of the measured degradant (potency) levels at each time point, this is not necessarily true of the isoconversion times. In order to determine a distribution in rates and shelf-life at ambient conditions, one must propagate a distribution to the Arrhenius calculations. The extrema method provides a single value as an estimate of the uncertainty of the isoconversion time and %degradant at the mean isoconversion time. Two methods were considered for generating distributions from this estimate. While the distribution in isoconversion times will generally not be normal, the distribution will be nearly normal when the values are a sufficient number of standard deviations from zero. One option is to assume the isoconversion distribution will be sufficiently normal and generate a distribution using the mean estimate of the isoconversion and the determined error estimate as the standard deviation. Another option is to take advantage of the normality of the degradant levels by defining a degradant distribution at the mean estimate of the isoconversion time. This distribution in degradant levels can be used to estimate the distribution in isoconversion rates (dividing the degradant distribution by the mean isoconversion time). It should be noted that the relation that we really want is between isoconversion at accelerated conditions and isoconversion at ambient conditions. For convenience this is converted into a rate by using the specification limit divided by the isoconversion time. If we use the degradant distribution at the mean isoconversion time directly to calculate the rate (degradant distribution divided by mean isoconversion time), this approximates, but is not identical to the rates calculated from the isoconversion time distribution itself. The rate distributions from the true isoconversion distribution and approximated from the degradant distribution will diverge as the rate decreases.

Additional complexity arises for these low rate conditions. For example, the assumption of normality of the degradant distribution

results in negative values of degradant levels. Since the distribution now includes zero, there will be a discontinuity in the reciprocal of the degradant distribution; i.e., it will include an infinite isoconversion time. Calculations in this paper have been done assuming rates that are sufficiently high to avoid many of the added statistical and scientific complexity required to handle cases when there is very little change in measured degradant levels with respect to the error bars.

II. Arrhenius Fitting Error

Once we have a set of isoconversion rates and their corresponding distributions at different temperatures ($k_{iso_{T2}}$), we can proceed to estimate isoconversion rates at ambient temperature ($k_{iso_{T1}}$, where the shelf-life is the specification limit divided by $k_{iso_{T1}}$) using the Arrhenius equation, as described either in Equation 5 or 6 (E_a is the activation energy, A is the collision frequency, and R is the gas constant).

$$k_{iso_{T_1}} = Ae^{-E_a/R(1/T_2)} \quad (5)$$

$$k_{iso_{T_1}} = k_{iso_{T_2}} e^{-E_a/R(1/T_1 - 1/T_2)} \quad (6)$$

Equation 5 represents a standard form of the Arrhenius equation with the pre-exponential collision frequency term A. Equation 6 avoids solving for the collision frequency by relating the rate at an elevated temperature to that at each of the accelerated temperatures[5]. While it may be tempting to treat each measured degradant point as an independent rate value in the overall fitting (as was employed in early literature[5]), this approach is intrinsically inaccurate for the general case of non-simple kinetics[1]. As an example of the issue, consider the case where one has only two temperatures with many time points far from isoconversion. Using each independent point to calculate rates, these points would contribute directly to an overall Arrhenius fitting. If the assumption of that rate form were in error (usually assumed to be zero or first order), the overall resulting fit would be erroneous. To make matters worse, because many points were used, the error bars would be small around the wrong projected shelf-life. With the isoconversion approach, the number of points alone would not make the final precision high if those points do not provide a precise estimate of the isoconversion. In this case, the points far from isoconversion will still have a large error bar due to the extrapolation. Isoconversion will in general provide better accuracy, but greater error bars than a point estimate approach, and therefore is more appropriate for making critical stability-related decisions.

Another difference in this work versus the earlier work[5] is our focus on the distribution of rates (Equations 5, 6) rather than shelf-lives. This is useful in order to handle situations where there are changes in rates over time, such as when RH changes with time (permeable packaging [3]), or when a drug product experiences an excursion during shipping or handling.

Equations 5 and 6 can be solved either in exponential or linearized forms. If three or more temperatures are used, the estimation of the uncertainty will be different between the two fitting methods. If one had perfect data (i.e., where a line could be fit through three temperature points exactly), the point estimate of the shelf-life would be identical using either method. Estimating an uncertainty interval requires sampling points from the distribution in rates. Even when the regression line is a perfect fit, the sampled points will not generate perfect fit lines. A least-squares fit minimizes the square root of the sum of the squares of the difference between the actual and computed value. For the linear fit, the difference that is being minimized is a difference in ln k. For the non-linear fit, the difference is in k. This will affect which best fit line is chosen and ultimately the distribution of values used to estimate the uncertainty.

In a series of calculated experiments described below, we examine two of the approaches for determining the error bars for the isoconversion rates described above: a normal isoconversion distribution (the mean is the zero-error isoconversion value, and the standard deviation is from the extrema method) and a normal degradant distribution (the mean is the specification limit, and the standard deviation is determined from the extrema method). In a plot of degradant versus time, the former case is normalized about the x-axis, and the latter case is normalized about the y-axis. For both cases, differences between fitting the linear and non-linear forms of the Arrhenius equation are calculated.

The first test case uses only a single time point with a measured degradant point (0.02% fixed error bar) and the zero time point (with zero uncertainty) at two temperatures (60 and 70°C) with an activation energy of 25 kcal/mol (104.7 kJ/mol). The degradant values at 10 days are 0.2000% and 0.6015%, at 60°C and 70°C, respectively. This results in a value for ln A of 33.87 and a shelf-life at 25°C of 2.31 years. Only the linearized Arrhenius equation was used in this example, since both linear and non-linear methods will give the same results when only two temperature points are used. Using the two methods described above, the distribution of isoconversion values and corresponding predicted shelf-lives were calculated using a cumulative distribution function (CDF), with results shown in Table 4. The distribution in isoconversion values is nearly normal. The distribution of shelf-life values is not, as can be seen by the difference between the mean and median values. As expected, using the median produces values closer to the zero-error shelf-life. From the distribution in rates (not shown), a probability of passing can be determined at any given time; however, the distribution cannot be assumed to be normal and probabilities need to be determined from the CDF to accurately reflect the probability of passing.

Table 5 shows the case where the zero time value also has an uncertainty of 0.02% (LOD) associated with it (other values are the same as in the previous experiment). In this case, the standard deviation at 60°C remains close to that without the zero-point error, despite the additional uncertainty. The values calculated from both methods remain similar to each other. For these first two examples, the median in either case is close to the true value. While the two methods differ in their results, there is no over-riding reason for selection of one over the other; both provide reasonable estimations of the uncertainty in shelf-life at ambient conditions. Again, it is clear that the shelf-life distribution with either method is not normal and supports use of a CDF to determine the probability that the specification limit will not be exceeded at a given time.

When a third temperature is included, the linear least squares and non-linear least squares fit of the Arrhenius equation will give different uncertainty estimates. In applying these methods, the distribution in both k and ln k are assumed to be sufficiently normal to apply a least-squares fitting procedure. Using three temperatures (60, 70, 80°C) and a constant 10%RSD with a 0.02% LOD the distribution in rates in shown in Figure 2 with the results summarized in Table 6. In all cases, the distributions in shelf-life values are not normal (as illustrated by the significant difference between the mean and median), with the median being very close to the actual value. From these calculations, it remains difficult to distinguish between the two methods for estimating the error in the isoconversion values. With non-linear fitting, the error bars become significantly larger. This occurs because the degradant rates are much higher at higher temperatures. A small change in an absolute rate at low temperature will not affect the least-squares sum very much, but will result in a large percentage error in the prediction of the rate. The high temperature values are effectively fit more tightly resulting in larger intervals at ambient temperature. The logarithmic fitting tends to compress these values effectively giving more weight to the lower temperature values than with the non-linear fitting. It is certainly possible to use some form of weighting for the non-linear least-squares procedure. Ultimately, however, this would bring the results from a much more complex calculation into line with those achieved by the simpler linear fitting process: applying weighting schemes to the non-linear fit would just reproduce the linear fit at higher computational cost. Another issue to consider when applying non-linear fitting is that when there are any deviations from ideal Arrhenius behavior (e.g., change in mechanism with temperature),

the points that would make any such deviation most likely to be inaccurate at room temperature will be weighted most heavily. This consideration also favors using the linear fitting.

The final case to consider is a complex kinetics example of secondary degradation. In this case, we compare the same methodologies (assuming the most likely scenario of a fixed RSD with a minimum LOD) where some of the conditions more accurately reflect the true isoconversion values than other conditions. The results of these calculations are shown in Table 7. The table shows results with two different specification limits. With the larger specification limit, some of the conditions require extrapolations. The result is that the predictions are not as accurate. What is reassuring, however, is that in spite of this, the fitting still gives predictions that include the correct values within one standard deviation.

CONCLUSIONS

A two-step process of first estimating the isoconversion distribution at different conditions, then using the rates derived from this distribution to fit the Arrhenius equation was employed to determine uncertainty intervals for ambient shelf-life. The intermediate calculation of isoconversion is required on scientific grounds, as it produces a more accurate estimate of shelf-life. As a result of this requirement, different methods of producing isoconversion distributions were evaluated including stochastic, non-stochastic and extrema methods. While there is no objective value to indicate which method is correct, each method was compared to the regression interval (RI) and the results of cases where fitting involves complex kinetics. An extrema method for estimating error in the isoconversion times (times to hit the specification limit) was found to give results in reasonable agreement with the regression interval (RI) in situations where the RI could be calculated (constant, fixed standard deviation, SD), yet allow for extension of calculations to include the more realistic situation of a constant relative standard deviation (RSD) with a minimum standard deviation equal to the limit of detection where the RI could not be directly determined. This method was used either to calculate a distribution of isoconversion times (times to hit the specification limit) or a distribution of degradant levels (values around the “true” isoconversion time). Based on the present evaluation, there is no over-riding reason to select either option, with the differentiation to be determined based on cases where the degradant levels are as low as the standard deviation. Linear and non-linear least-squares fitting of the Arrhenius equation were also evaluated. Fitting the non-linear form of the Arrhenius equation inherently places more weighting on high temperature conditions, resulting in larger estimated uncertainty when extrapolating to ambient conditions. Since the non-linear fitting is also more computationally challenging and more likely to be inaccurate when there are deviations from ideal Arrhenius behavior, use of the linear form of the Arrhenius equation is recommended. Distributions of uncertainty intervals show that in many cases, even when a normal distribution of isoconversion values are used, a normal distribution of ambient projected rates or shelf-life estimates will not generally be observed.

1. Waterman KC, Swanson JT, Lippold BL. A scientific and statistical analysis of accelerated aging for pharmaceuticals part 1: accuracy of fitting methods. *J. Pharm. Sci.* 2014; **103**(10):3000-3006. DOI: 10.1002/jps.24075.
2. International Conference on Harmonisation of Technical Requirements for Registration of Pharmaceuticals for Human Use: Harmonised Tripartite Guideline, Stability testing of new drug substances and products Q1A(R2), 6 February 2003.
3. Waterman KC, MacDonald BC. Package selection for moisture protection for solid, oral drug products. *J. Pharm. Sci.* 2010; **99**:4437-4452. DOI: 10.1002/jps.22161.
4. Porterfield RI, Capone, JJ. Application of kinetic models and Arrhenius methods to product stability evaluation. *Med. Device Diag. Ind.* 1984: 45-50.
5. King SP, Fung H, Kung M. Statistical prediction of drug stability based on nonlinear parameter estimation. *J. Pharm. Sci.* 1984; **73**:657-662. DOI: 10.1002/jps.2600730517.
6. Sundberg R. Statistical aspects on fitting the Arrhenius equation. *Chemometrics Int. Lab. Systems* 1998; **41**:249-252.
7. Gil-Alegre ME, Bernabeu JA, Camacho MA, Torres-Suarez AI. Statistical evaluation for stability studies under stress storage conditions. *Il Farmaco* 2001; **56**:877-883.
8. Shimizu Y, Tamura T, Ono M, Kasai O, Nakajima T. Application of nonlinear fitting and selection of the most fitted equation by AIC in stability test of pharmaceutical ingredients. *Drug Dev. Ind. Pharm.* 2002; **28**(8):931-937. DOI: 10.1081/DDC-120006425.
9. Huntsberger DV, Billingsley P. 1977. Elements of Statistical Inference, 4th Edition. London: Allyn & Bacon, Inc. p. 275.
10. Waterman KC, Carella AJ, Gumkowski MJ, Lukulay P, MacDonald BC, Roy MC, Shamblin SL. Improved protocol and data analysis for accelerated shelf-life estimation of solid dosage forms. *Pharm. Research* 2007; **24**(4):780-790. DOI: 10.1007/s11095-006-9201-4.
11. R Core Team (2013). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL http://www.R-project.org/

Table 4

	Isoconversion time distribution					Normalized degradant levels at zero-error isoconversion time				
	CDF[a]			Normal[b]		CDF[a]			Normal[b]	
	15.9%[c]	Median	84.1%[c]	Mean	σ	15.9%[c]	Median	84.1%[c]	Mean	σ
60°C (isoconversion, days)	9.1	10.0	11.1	10.1	1.0	9.1	10.0	11.1	10.0	1.0
70°C (isoconversion, days)	3.2	3.3	3.4	3.3	0.1	3.2	3.3	3.4	3.3	0.1
ln A	30.5	33.8	37.4	34.0	3.5	30.5	33.8	37.4	34.0	3.5
E_a (kcal/mol)	22.7	25.0	27.4	25.1	2.4	22.7	25.0	27.4	25.1	2.4
25°C Shelf-life (yrs)	1.4	2.3	3.9	2.7	1.7	1.4	2.3	3.9	2.7	1.7

Comparison of two shelf-life prediction methods based on two-temperature accelerated aging: a distribution of isoconversion times (intercepts with the specification limit); and a normalized distribution in degradant levels where the extrema lines intersect the time for the regression fit to the specification limit. The degradation is based on a single time point at 10 days for both temperatures (E_a = 25.0 kcal/mol) with a 10% relative standard deviation at those points (0 standard deviation at the origin). The true shelf-life at 25°C is 2.31 years. [a] Cumulative distribution function. [b] Values for mean and standard deviation (σ) calculated assuming a normal distribution. [c] Time to the indicated percent area under the curve based on the CDF.

Table 5

	Isoconversion time distribution					Normalized degradant levels at zero-error isoconversion time				
	CDF[a]			Normal[b]		CDF[a]			Normal[b]	
	15.9%[c]	Median	84.1%[c]	Mean	σ	15.9%[c]	Median	84.1%[c]	Mean	σ
60°C (isoconversion, days)	9.08	10.00	11.09	10.10	1.07	9.09	10.00	11.07	10.09	1.05
70°C(isoconversion, days)	3.07	3.32	3.57	3.32	0.25	3.09	3.32	3.59	3.34	0.26
ln A	29.88	33.98	38.48	34.17	4.32	29.68	33.83	38.17	33.93	4.26
E_a (kcal/mol)	22.31	25.07	28.10	25.20	2.91	22.18	24.97	27.90	25.04	2.87
25°C Shelf-life (yrs)	1.33	2.35	4.37	2.92	2.22	1.30	2.29	4.22	2.81	1.99

Comparison of two shelf-life prediction methods based on two-temperature accelerated aging: a distribution of isoconversion times (intercepts with the specification limit); and a normalized distribution in degradant levels where the extrema lines intersect the time for the regression fit to the specification limit. The degradation is based on a single time point at 10 days for both temperatures (E_a = 25.0 kcal/mol) with a 10% relative standard deviation at those points (0.02% standard deviation minimum). The true shelf-life at 25°C is 2.31 years. [a] Cumulative distribution function. [b] Values for mean and standard deviation (σ) calculated assuming a normal distribution. [c] Time to the indicated percent area under the curve based on the CDF.

Table 6

		Isoconversion time distribution				Normalized degradant levels at zero-error isoconversion time			
		15.9%[a]	Median	84.1%[a]	Mean	15.9%[a]	Median	84.1%[a]	Mean
Linear	Shelf-life (years) at 25°C	1.43	2.31	3.86	2.70	1.35	2.34	3.84	2.61
	Rate at 25°C (10^{-4}%/day)	3.83	2.38	1.42	2.62	4.05	2.34	1.43	2.80
Non-linear	Shelf-life (years) at 25°C	0.90	2.33	7.12	5.41	0.81	2.36	7.48	5.90
	Rate at 25°C (10^{-4}%/day)	6.21	2.35	7.70	3.57	6.74	2.32	0.73	3.92

Comparison of shelf-life prediction methods based on three-temperature (50, 60, 70°C) accelerated aging: a distribution of isoconversion times (intercepts with the specification limit); and a normalized distribution in degradant levels where the extrema lines intersect the time for the regression fit to the specification limit. Both linear and non-linear methods for fitting the Arrhenius equation were applied. The degradation is based on a single time point at 10 days for each temperature (E_a = 25.0 kcal/mol) with a 10% relative standard deviation at those points (0.02% standard deviation minimum). The true shelf-life at 25°C is 2.31 years with a rate of 2.37 X 10^{-4}%/day. [a] Time to the indicated percent area under the curve based on the cumulative distribution function.

Table 7

		25°C shelf-life (yrs) calculated using distribution in isoconversion times				25°C shelf-life (yrs) using normalized degradant levels at zero-error isoconversion time				True 25°C shelf-life (yrs)
Degradant	**Specification limit**	**15.9%**	**Median**	**84.1%**	**Mean**	**15.9%**	**Median**	**84.1%**	**Mean**	
Primary	0.2%	1.03	1.37	1.81	1.49	0.98	1.36	1.92	1.45	**1.43**
	0.5%	2.09	3.16	5.38	4.32	1.94	3.23	5.89	4.21	**4.45**
Secondary	0.2%	1.55	2.06	2.76	2.16	1.50	2.07	2.82	2.17	**2.02**
	0.5%	3.48	4.75	7.12	5.40	2.91	4.70	8.70	6.12	**4.01**

Comparison of shelf-life prediction uncertainty methods when the degradant formation is not linear; i.e., when the isoconversion times at each condition are not exactly determined by the regression lines. In this example, a primary degradant subsequently undergoes secondary degradation (k_1 = 0.000113%/d and, k_2 = 0.011250%/d at 50°C, E_a = 25.0 kcal/mol; time points at 0, 3, 7, 14 and 28 days; using 50, 60, 70°C), assuming a relative standard deviation of 10% with a minimum error equal to the the limit of detection (0.02%). Only the two points that bracket the specification limit are used at each condition with distributions using the extrema method (extending the last two points when the specification limit is not hit). In each case, a cumulative distribution function was used to calculate the ambient distributions.