

Creating the Time to Event ADaM Dataset: The Nuts and Bolts

Nancy Brucken, PharmaNet/i3, Ann Arbor, MI

Paul Slagle, United BioSource Corp., Ann Arbor, MI

ABSTRACT

The ADaM Basic Data Structure can be used to create far more than just laboratory and vital signs analysis datasets. Often, the biggest challenge is the development of efficacy datasets, and of the commonly-used efficacy datasets, creation of a time-to-event (TTE) dataset presents many interesting problems. These TTE datasets are frequently used in survival analysis, for example, to generate Kaplan-Meier curves for oncology reports. As a continuation of this series, we will provide directions for creating the ADaM-based TTE dataset, as well as help in using that structure to support both validation and statistical review of the results. Prior experience with ADSL and creation of other BDS datasets is expected.

INTRODUCTION

If you have been following our topics over the last couple of years, you will already have an appreciation for the structure of the ADaM datasets and how they can be used, along with metadata, to support the systematic development of tables to support clinical trials. We have discussed how the Subject Level data set (ADSL) provides a basic structure for defining who is, and isn't, in the analysis using the population flags. We have covered the Basic Data Structure (BDS) associated with such data as vital signs and how some data sets, such as adverse events, are allowed to differ from the BDS. That is old news, so we are going to pass on reviewing those topics.

But if all we did were demographics, adverse events, and lab tables then our reporting would be kind of boring, right? Determining if the compounds we are testing are working as expected is the most interesting part of the studies. A few years ago, much of the analysis being performed was simply based on whether there was a significant difference in an analytical result between the populations being tested, such as a reduction in lipid levels from baseline or decrease in the average blood pressure reading between two populations. Compared to some of today's analytics, these were simple tests.

Today many of the analytics incorporate time in the analysis. More common in oncology studies, the question becomes whether the time to a specific event, often death or the lack of progression in the tumor, is longer in one population when compared to another. Other cases might include how long a specific Alzheimer's treatment prevents the onset of the disease compared to placebo. And while it might seem simple just to capture the first date of treatment and compare it to the date of the event, the source of that data can span multiple pages of CRFs. There are additional considerations, such as whether the subject maintained the treatment until the event occurred. All of these are taken into account for the analysis, and in many cases, require the programmer to produce output showing all potential conditions for the study team.

Because of the uniqueness of the time-to-event analysis, we are going to present this subject by itself. As we have done in the past, we will start with the end in mind by presenting the report and the report metadata. From that point, we will work through the construction of the BDS datasets needed for analysis using the Vital Signs data. Once that is complete, we will take on the challenge of a simplified progression-free survival analysis data set. All of our presentations will be based on the draft ADTTE Guidelines published by CDISC in March 2011 using ADaM version 2.1.

STARTING AT THE END

As we have taught in prior workshops, let's start by looking at what we have to produce. First check out the Statistical Analysis Plan (Appendix 1). In Section 4 of the SAP you will find the efficacy parameter described as "Time until the first diastolic blood pressure ≤ 90 mmHg is achieved." Reading on, you will find the definition for the time until the first diastolic blood pressure is reached in Section 6.4. You will notice that not only will we need to determine the date on which each subject achieves the ≤ 90 mmHg reading, we will also need to censor subjects who do not achieve that by Week 24 at Week 24. Also, if a subject discontinues prior to Week 24, we will censor that subject's data at the last date read.

Table 14.2.5.1 is the specific table related to the Kaplan-Meier analysis described in the SAP. Notice from this table, we are going to need some data that is not directly captured in the ADVS analysis data set. Specifically, we are going to be creating a variable called CNSR (censor) in order to identify how many subjects were censored for each group. We are also going to need to identify the number of weeks between when a subject first took the medication and when their diastolic blood pressure reached 90 mmHg.

Bing Pharmaceutical Company BP3304-002		Date: ddMONyyyy Program: sasresults.SAS Page X of Y	
14.2.5.1 Time to Achievement of Diastolic Blood Pressure ≤ 90 mmHg Modified Intent-to-Treat Population			
	BP3304 (N = xx)	Placebo (N = xx)	
	n(%)	n(%)	
N	xx	xx	
No. with DBP ≤ 90 mmHg	xx (xx.x)	xx (xx.x)	
No. of Censored	xx (xx.x)	xx (xx.x)	
Time to DBP ≤ 90 mmHg (Weeks)			
Median	xx.xx	xx.xx	
95% CI of Median	(xx.xx, xx.xx)	(xx.xx, xx.xx)	
2.5-97.5%ile	xx.xx - xx.xx	xx.xx - xx.xx	
Min, Max	xx.x, xx.x*	xx.x, xx.x*	
<small>Reference: Listing 16.2.6.2 Note: Time to diastolic blood pressure ≤ 90 mmHg is calculated using Kaplan-Meier method. 95% CI for median is computed using Brookmeyer and Crowley's method. * = censored value, C.I. = Confidence interval, DBP = diastolic blood pressure.</small>			

But what if the subject never achieves a value of 90 mmHg? What if the subject not only was unable to achieve a 90mmHg diastolic pressure but had pressure readings that continued to escalate, causing the subject to change their dosing regimen or even leave the study? What about subjects who failed to follow the protocol requirements at some point during the study? These subjects will still be counted, but their results will be censored. Not only will we need to account for these subjects in this table, but often there is a table or listing showing all of the reasons subjects were censored.

One last item to consider when developing this table is supporting the statisticians on the study. Due to the complexity of creating this type of analysis, in the statisticians are very interested in knowing which visits are being used to determine the number of weeks until the event. What if the SAP were to include a specification that only the first date is to be used where two subsequent diastolic readings were at or below 90mmHg? Tracking which record to use is very important not only to the statisticians, but also to the regulatory agency reviewers. It's also important for maintaining the sanity of the programming department when debugging.

SO NOW LET'S LOOK AT THE METADATA

Let's start with the Results Metadata. From looking at that table, we can see there are going to be two main sections. The first section shows basic frequencies of the number of subjects in the population, the number who were censored and the number who were not censored. The second section takes the same data and performs the Kaplan-Meier analysis. So if we look at the table below, we can see some of the key fields that will need to be completed in the Results Metadata.

Display Identifier	Result Identifier	PARAMCD	Analysis Variable	Selection Criteria	Documentation	Programming Statements
14.2.5.1	N, Number with DBP < 90 mmHg, Number censored	DBP	CNSR	MITTFL='Y' and PARAMCD = 'DBP90'	SAP Section 6.3: Time until first DBP ≤ 90 mmHG is achieved.	PROC FREQ DATA=ADTTE; TABLES TRTP * CNSR;
14.2.5.1	Time to DBP ≤ 90 mmHg	DBP	AVAL	MITTFL='Y' and PARAMCD = 'DBP90'	SAP Section 6.3: Kaplan-Meier analysis of time to event endpoints.	PROC LIFETEST DATA=ADTTE ALPHA=0.05 METHOD=KM; BY TRTPN; TIME AVAL * CNSR(1);

Figure 2, Analysis Results Metadata

Once we have the Results Metadata, let's take a look at the Variable Metadata. As you can tell, the ADTTE data set has the same fundamental layout of the BDS described in the ADaM specifications. If we just look at the basic variables that will need to be included, we have the variables from ADSL showing the population for each subject as

well as the variables used in the analysis. We assign a value of DBP90 to the PARAMCD, which relates to the Results Metadata we have described above. We calculate the AVAL as the number of weeks between STARTDT, the Time to Event Origin/Date, and the ADT, or the Analysis Date. Note that the analysis date is assigned to either the date on which the 90 mmHg value is achieved, the date of the Week 24 vital signs reading, or the last date of treatment.

Dataset Name	Variable Name	Variable Label	Variable Type	Display Format	CodeList / Controlled Terms	Source / Derivation
ADTTE	PARAMCD	Parameter Code	Text	\$8.		Set to 'DBP90'
ADTTE	AVAL	Analysis Value	Float			ADT - STARTDT + 1
ADTTE	STARTDT	Initial Date of Event Risk	Integer	DATE9.		ADSL.TRSTDT
ADTTE	ADT	Analysis Date	Integer	DATE9.		Set to ADVS.ADT for first DBP record with ADVS.AVAL <= 90 and ADY > 1. If there are no such records, then set to date of last Week 24 DBP record (ADVS.ADT where PARAMCD='DIABP' and AVISIT=7 and DTYPE=' '); if no Week 24 DBP record, then set to ADSL.TRTEDT.
ADTTE	CNSR	Censoring Indicator	Integer	1.0	0,1	Set to 0 if patient has at least 1 post-baseline ADVS record with DBP < 90 mmHg; set to 1 if patient has at least 1 post-baseline ADVS record; else leave blank.
ADTTE	EVNTDESC	Event or Censoring Description	Text			Set to 'DBP <= 90 mmHg' if CNSR=0; set to 'COMPLETED STUDY WITHOUT DBP <= 90 mmHg' if CNSR=1; else leave blank.

Figure 3, Analysis Variable Metadata

Since we have all of these possibilities for the analysis date, we also need to use the CNSR variable to define whether the analysis date represents the subject achieving the goal criterion, which in this case is the 90 mmHg diastolic pressure reading. If the subject reaches the goal, then the CNSR variable is assigned a '0', meaning they were not censored. Otherwise the CNSR variable is assigned a value of '1', indicating the subject was censored.

In addition, if the value for CNSR is set to 1, then the EVNTDESC variable contains the reason for censoring. This will help the team summarize why subjects were censored. But even with these extra variables, the ADTTE data set still looks very much like any other ADaM data sets we have created.

ADDING TRACEABILITY TO METADATA

Traceability can be described as the key value of the ADTTE dataset. Earlier in the paper, we discussed how traceability is important when evaluating where the results came from. In this case, we need to identify the source of the Analysis Date. The key to all of this is the SRCDOM variable, which identifies the analysis data set being used to populate that analysis date variable. SRCVAR is used to identify which variable is used from the dataset identified in SRCDOM. The last item of accounting is the SRCSEQ variable, which identifies which record from the SRCDOM dataset is used to populate the specific record in the ADTTE dataset. Confusing? Let's try to explain it a little more.

Let's start by looking at CNSR. When CNSR is set to '0', the subject reached our criteria, so SRCDOM is set to ADVS, SRCVAR is set to ADT, and SRCSEQ is set to the value in VSSEQ for the ADVS record used to identify the analysis date. This is the simplest condition for populating these variables.

If CNSR is set to '1' then we need to check if the subject completed the

Dataset Name	Variable Name	Variable Label	Variable Type	Display Format	CodeList / Controlled Terms	Source / Derivation
ADTTE	SRCDOM	Source Domain	Text	\$8.	ADVS, ADSL	Set to 'ADVS' if ADT comes from ADVS; set to 'ADSL' if ADT comes from ADSL.
ADTTE	SRCVAR	Source Variable	Text	\$8.	ADT, TRTEDT	Set to 'ADT' if ADT comes from ADVS; set to 'TRTEDT' if ADT comes from ADSL.
ADTTE	SRCSEQ	Source Sequence Number	Integer			Set to ADVS.VSSEQ of the record containing the value of ADVS.ADT selected for ADT; if ADT comes from ADSL.TRTEDT, then leave blank.

Figure 4, Traceability variables in Variable Metadata

study. If the subject made it to Week 24, then the values in these variables represent the record in ADVS containing the Week 24 results for this subject. If, however, the subject discontinues early, then we are recording the date of last treatment from the ADSL dataset, so SRCDOM is set to ADSL, SRCVAR is set to TRTEDT, and SRCSEQ is set to missing (since ADSL contains only one record per subject).

Why is this information so important? Well let's take a look at subject A02 in our sample study. In this case, the subject never achieves the 90 mmHg results and the subject discontinued before visit 7 (Week 24). As a result, their analysis date will be generated from the ADSL dataset. When reviewing the results, if the only thing in the analysis data set is the analysis date, how would you know that value cannot be found in the ADVS dataset?

TAKING A LOOK AT ONCOLOGY DATA

So far we have taken a rather simple approach to creating the ADTTE dataset using the vital signs data. But time to event data is frequently used in oncology trials to evaluate results of overall survival or progression free survival as assessed by radiologic evidence. The good news is that the basic structure of the ADTTE is the same as what we have presented so far.

On the other hand, maintaining the traceability for all the analysis data can be much more challenging. Depending on the design of the CRF, you may have many more sources for pulling values. For a simple design, you might only need to collect the data of death from an end of study record. If a subject withdrew from the study early, you might have to scan more than a single CRF to determine the last known survival date.

For this study, we're going to amend the SAP (like that never happens in the real world). Assume that in another clinical trial, this compound was found to have an effect on metastatic liver cancer. We now have a wonder drug on our hands and the sponsor, now giddy with excitement, has hastily amended the protocol. Our analysis plan has now been amended with an additional secondary efficacy parameter:

To assess the effect of BP3304 in combination with existing cancer treatments on progression-free survival using assessments based on RECIST v1.0

Based on the protocol, and in an effort to keep this simple, the study will measure efficacy using investigator-assessed radiologic image assessment. We will also assume that for this study, the imaging assessment will be based solely on a single tumor and not on multiple tumors or other measurements. For those unfamiliar with oncology trials, let's review some of the terms now being used:

- 1) RECIST v1.0 criteria is a standard used in oncology studies for assessing tumors. The investigator will use the images to measure the lesions being treated. Based on that measurement and the investigators assessment, the lesion, and therefore the disease, will be graded as:
 - a. CR – Complete Response. The disappearance of all target lesions.
 - b. PR – Partial Response. In this case the size of the lesion has diminished from the previous assessment by a noticeable amount.
 - c. PD – Progressive Disease. In this case, the lesion has now increased in size from the previous assessment.
 - d. SD – Stable Disease. In this case, no noticeable change has occurred in the lesion size.
- 2) Progression-Free Survival. The National Cancer institute defines PFS as “The length of time during and after a treatment in which a patient is living with a disease that does not get worse.” To simplify this term for us, it is the time until an oncologist assesses a subject with a RECIST grade of PD.

For this workshop, we're going to walk through an example using an oncology study that estimates progression-free survival (PFS) as a secondary efficacy parameter. We will further refine the definition within the context of the time-to-event analysis by saying that PFS is the amount of time until the following events occur:

- the disease progresses (in other words, PD)
- the patient dies (worse than PD)
- the patient withdraws from the study due to lack of efficacy

Censoring occurs when any of the following situations happens:

- the patient completes the study without the disease having progressed
- the patient begins taking a non-study anti-cancer medication
- the patient withdraws from the study for any reason other than lack of efficacy

Starting with the end in mind, let's take a look at table 14.2.5.3, which has been designed to summarize time to Progression-Free Survival. Now that we have the design, along with the definition of Progression-Free Survival, we are now ready to define the metadata. We are going to skip past the report metadata, as that will be similar to what

you have already seen above. And while we are going to create a dataset called ADPFS, you can surmise that it will look amazingly close to the ADTTE dataset described above. However, as you start going thru the above specifications, you will begin to get an idea of the effort required to collect all of the data you will be working with. As a primary data source, you will want to use data from the Results (RS) domain.

Specifically, we will want to use the RSTESTCD of TRGRESP which contains the response of the target lesion. Looking at subject A02, we find that they have an assessment of PD on Nov 2, 2009. However, you can see that subject A01 never achieved a response of PD. For that subject we will need to assess

when their last visit occurred and use that date with the censoring flag set. What if they were lost to follow-up? What if they did not survive?

In order to support traceability and maintain any sense of sanity, we'll be creating an intermediate dataset, called ADEVENT.

Within this dataset, we will create a PARAMCD of MILESTNE, or milestone. For each subject, we will identify the milestone dates from the different data sources. Using the example above, you can see that we have identified subject A01 as having a milestone of SCREENING on June 23, 2009. We have also shown that the last recorded image date had a value of PR on November 17, 2009 and there was a follow-up visit on January 12, 2010. Based on the above PFS criteria, we know that if a subject is not assessed as PD, we should use the last known date alive. Thus, we set ANL01FL to Y on the record that we will use. We also have the sources for the dates and assessments indicated in the SRCDOM, SRCVAR, and SRCSEQ variables. For subject A02, we have the date of the screening visit, and we have determined the subject did show disease progression, so ANL01FL is assigned a value of "Y" for that date.

Using this approach, we have created a dataset that can be used by the programming team to indicate which dates are being

correctly used for analysis. This data set can also be used to produce a table enabling the statisticians and programming team to review which dates are being used and determine if additional criteria are needed for selecting visits.

14.2.5.3 Progression-Free Survival Based on Investigator Analysis
Modified Intent-to-Treat Population

	BP3304 (N=xxx)	Placebo (N=xxx)
Time to progression/death (mo)*:		
	95% CI*	95% CI*
25th Percentile	xx.x [xx.x, xx.x]	xx.x [xx.x, xx.x]
Median	xx.x [xx.x, xx.x]	xx.x [xx.x, xx.x]
75th Percentile	xx.x [xx.x, xx.x]	xx.x [xx.x, xx.x]
p-value^b	0.xxxxx	
Number of subjects by censor status:	n	
	(%)	
Not Censored (Progressed/Dead)	xxx (xx.x)	
Progressed	xxx (xx.x)	
Dead	xxx (xx.x)	
Censored	xxx (xx.x)	
Discontinued Study	xxx (xx.x)	
Initiated Nonstudy Anticancer Trt	xxx (xx.x)	
Missed Visit	xxx (xx.x)	

STUDYID	DOMAIN	USUBJID	RSDTC	RSTESTCD	RSTEST	RSORRES	RSTRESC
BP3304	RS	BP3304-A01	6/23/2009	TRGRESP	Target Response	SD	SD
BP3304	RS	BP3304-A01	9/22/2009	TRGRESP	Target Response	PR	PR
BP3304	RS	BP3304-A01	11/17/2009	TRGRESP	Target Response	PR	PR
BP3304	RS	BP3304-A01	12/15/2009	TRGRESP	Target Response		
BP3304	RS	BP3304-A02	6/8/2009	TRGRESP	Target Response	SD	SD
BP3304	RS	BP3304-A02	6/15/2009	TRGRESP	Target Response	PR	PR
BP3304	RS	BP3304-A02	8/10/2009	TRGRESP	Target Response	PR	PR
BP3304	RS	BP3304-A02	11/2/2009	TRGRESP	Target Response	PD	PD

USUBIID	ASEQ	ADT	EVNTDESC	PARAMCD	AVALC	ANL01FL	SRCDOM	SRCVAR	SRCSEQ
BP3304-A01	1	6/23/2009	Initiated Non-Study Anti-Cancer Therapy	MILESTNE	SCREENING		SV	SVSTDTC	1
BP3304-A01	2	11/17/2009	Last Image Assess	MILESTNE	PR		RS	RSORRES	7
BP3304-A01	3	01/12/2010	Follow-UP Visit	MILESTNE	Follow-Up	Y	DS	DSSTDTC	5
BP3304-A02	4	6/8/2009	Initiated Non-Study Anti-Cancer Therapy	MILESTNE	SCREENING		SV	SVSTDTC	10
BP3304-A01	5	11/17/2009	Last Image Assess	MILESTNE	PD	Y	RS	RSORRES	10
ADEVENT Example									

Now that we have all of this information for each subject, we can create the ADPFS dataset. Since we will be using the TRTSDT variable for each subject from ADSL, all we need from the ADEVENT dataset are those records where ANL01FL has been flagged. That provides the information we need to create values for AVAL when PARAMCD is PFSDUR, or PFS duration. And since the source of our calculation comes from ADEVENT, the SRCDOM is set to ADEVENT while the SRCVAR is assigned as ADT to indicate the date variable being used. The value for SRCSEQ will be the ASEQ value in ADEVENT for the record being used. This dataset is now ready to be called by a table program using the value of AVAL to create the KM analysis for Progression-Free Survival.

CONCLUSION

Creating the Time-to-Event data sets can seem quite daunting at first, especially when creating some of the datasets for oncology trials. However, the basic approaches we have described above should take away some of the mystique from the process. As with any of the analysis datasets, begin with the end in mind. From there, don't be afraid to create interim datasets that can be used to provide traceability and sanity checks.

REFERENCES

CDISC Analysis Data Model, Version 2.1, available at <http://www.cdisc.org>

The ADaM Basic Data Structure for Time-To-Event Analysis, available at <http://www.cdisc.org>

National Cancer Institute website; PFS is defined at <http://www.cancer.gov/dictionary?cdrid=44782>.

Imaging Response Criteria, RECIST Guidelines can be found at <http://imaging.cancer.gov/clinicaltrials/imaging>

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Nancy Brucken
PharmaNet/i3
5430 Data Court, Suite 200
Ann Arbor, MI 48108
(734) 757-9045

nbrucken@pharmanet-i3.com

Paul Slagle
United BioSource Corp.
2200 Commonwealth Blvd, Suite 100
Ann Arbor, MI 48105
(734) 944-8940 x1604

Paul.Slagle@unitedbiosource.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.

APPENDIX 1 – STATISTICAL ANALYSIS PLAN

STATISTICAL ANALYSIS PLAN
Bigg Pharmaceutical Company
BP3304-002
Page 7

- The change in diastolic blood pressure from baseline to Weeks 4, 8, 12, 16 and 20 of the Treatment Phase, Week 24 of the Follow-up Phase, and end of study. Diastolic blood pressure data will be obtained from the vital signs dataset.
- The change in systolic blood pressure from baseline to Weeks 4, 8, 12, 16 and 20 of the Treatment Phase, Week 24 of the Follow-up Phase, and end of study. Systolic blood pressure data will be obtained from the vital signs dataset.
- Time until the first diastolic blood pressure ≤ 90 mmHg is achieved. Patients who do not achieve diastolic blood pressure ≤ 90 mmHg by Week 24 will be censored at that visit. Patients who discontinue from the study prior to Week 24 for reasons not related to study conduct, or discontinue secondarily to an adverse event, will be censored at the time of discontinuation and treated as randomly censored.
- Proportion of patients achieving diastolic blood pressure ≤ 90 mmHg at Week 20 of Treatment and at Week 24 of Follow-up.
- The change in diastolic blood pressure from baseline to Weeks 4, 8, 12, 16 and 20 of the Treatment Phase, Week 24 of the Follow-up Phase, and end of study in patients over 65 years old.
- The change in diastolic blood pressure from baseline to Weeks 4, 8, 12, 16 and 20 of the Treatment Phase, Week 24 of the Follow-up Phase, and end of study in patients diagnosed with hypertension more than 10 years prior to enrollment in the study.

Summary statistics (including number of patients, mean, SD, minimum and maximum by visit and by treatment group of the raw and change-from-baseline values will be provided for both diastolic and systolic blood pressure. Baseline will be defined as the average of three measurements taken on Day 1 before the first dose of study drug; change will be defined as the post-baseline value minus the baseline value. If there are no measurements taken on Day 1, the value taken at the screening visit will be used as the baseline value.

For change in diastolic and systolic blood pressure from baseline, a last observation carried forward (LOCF) analysis using an analysis of covariance (ANCOVA) model will be performed. The ANCOVA model will have fixed effects for treatment and baseline therapy strata and a covariate for baseline blood pressure.

Time to event endpoints (time to first diastolic blood pressure ≤ 90 mmHg) will be analyzed using Kaplan-Meier (KM) survival methods (SAS PROC LIFETEST) (Collett 2003, Allison 1995). The KM algorithm will be applied to derive the median event time and the 95% confidence interval for the median for each treatment group. A small number of events (if median is non-estimable) would limit the use of the Kaplan-Meier method to provide reliable information. In this case, descriptive statistics (n, mean, median, standard deviation, min, max) will be provided. Differences in the survival distributions between treatment groups will be assessed using 2-sided log-rank and Wilcoxon tests with alpha set at the 0.05 level of significance acknowledging the differential emphasis each method places upon earlier versus later events. The event time will be considered a censored time for patients who either complete the study or discontinue prior to achieving the defined event of interest.

Figure 1, KM text from SAP