

Safety Estimands and Analysis Dataset Design: Advancing Safety Evaluation in Clinical Trials

Kun Ren, AstraZeneca R&D China;
Zeyu Guang, AstraZeneca R&D China

ABSTRACT

With the adoption of the ICH E9(R1) estimand framework (ICH 2019), there is increasing interest in applying estimand principles to safety evaluation in clinical trials. Traditional safety analyses are often descriptive and may not explicitly define the safety question or treatment effect of interest, particularly in the presence of intercurrent events such as treatment discontinuation, dose modification, rescue medication, or death. This paper introduces the concept of safety estimands, outlines the background for their use, and discusses their application in safety analyses. It also highlights key considerations for analysis dataset design within the CDISC ADaM framework (CDISC 2016), with particular focus on time-to-event dataset structures to support estimand-aligned derivations and transparent handling of intercurrent events. Integration of the safety estimand framework with the aligned ADaM design strengthens the scientific and operational quality of safety analyses.

INTRODUCTION

The ICH E9(R1) addendum, through its estimand framework (ICH 2019), has formalized the requirement to precisely define treatment effects in clinical trials, extending beyond efficacy to safety evaluation. However, safety analyses in many studies remain largely descriptive, often relying on simple tabulations or percentages of patients with adverse events. While operationally convenient, these summaries do not always clearly define the clinical question of interest.

For safety endpoints, this uncertainty in interpretation is particularly important because the occurrence and timing of adverse events are often influenced by factors beyond the randomized treatment itself. Premature study treatment discontinuation, initiation of concomitant or rescue medication, treatment modification, and death can all affect what data are observed and how those data should be interpreted. Without a predefined strategy for handling these events, the reported safety results may combine different clinical questions within a single analysis. This can reduce interpretability and make comparisons across studies or across analysis approaches more difficult (Ratitch et al. 2020).

Applying the estimand framework to safety can address this issue by aligning the study objective, the handling of intercurrent events, the target population, and the population-level summary (ICH 2019). While this framework has been widely adopted in efficacy analyses, its use in safety analyses is still evolving. Applying estimand thinking to safety requires careful consideration because safety questions are often broader, may involve multiple endpoint types, and are closely tied to treatment exposure and follow-up duration.

This paper introduces the use of estimand principles in safety analyses and discusses how these principles can be implemented in analysis dataset design under the CDISC ADaM framework. Attention is given to the handling of intercurrent events and to the derivation of variables that support estimand-aligned analysis periods, especially for time-to-event safety analyses. The goal is to improve both the scientific clarity and operational consistency of safety analyses by ensuring that analysis datasets transparently reflect the intended estimand.

BACKGROUND

The estimand framework, formalized in ICH E9(R1), requires that any treatment effect of interest be defined through five attributes: the target population, the endpoint (or variable), the specification of intercurrent events and the strategies for handling them, the population-level summary measure, and the

treatment condition under which the effect is evaluated. Together, these attributes ensure that the clinical question of interest is stated precisely and that the statistical method is aligned with that question. In efficacy analyses, this framework is now increasingly embedded in study objectives, endpoint definitions, and statistical methodology.

In contrast, safety analyses in many trials remain largely descriptive and do not explicitly identify the scientific question being addressed. Several specific limitations can be identified.

First, safety analyses often lack a clearly defined objective. In many trials, safety is summarized through incidence tables or percentages of patients with at least one adverse event, but the exact treatment effect represented by these summaries is not stated. Without an explicit estimand, it is unclear whether the analysis is targeting risk while patients remain on treatment, risk regardless of treatment discontinuation, or a hypothetical risk under alternative post-randomization scenarios. This absence of specificity limits both the interpretability and the reproducibility of the analysis.

Second, commonly used summary measures may not isolate the treatment effect of interest. A simple proportion of patients with an adverse event may reflect not only the pharmacological effect of the treatment, but also disease-related background risk, patient characteristics, and differences in observation time. In chronic or severe diseases, the observed adverse event burden may be heavily influenced by the underlying condition rather than the investigational treatment alone. As a result, crude incidence proportions may conflate treatment-related risk with baseline clinical risk.

Third, safety presentations often make limited use of comparator information. Without a structured comparison framework, it is difficult to distinguish the contribution of treatment from the background event profile of the study population. This issue is particularly relevant when pooled summaries are presented across multiple studies, populations, or follow-up durations, where heterogeneity further complicates interpretation.

Fourth, conventional safety analyses frequently do not account adequately for differences in follow-up time. For many safety endpoints, the probability of observing an event depends strongly on both the length of exposure and the duration of surveillance. A patient observed for one month and a patient observed for one year contribute differently to the interpretation of risk. Yet standard summaries often do not address this distinction, particularly when treatment of duration differs across arms or when intercurrent events shorten the observation period.

These limitations suggest that safety analyses would benefit from the same conceptual discipline already applied to efficacy. In practice, the need is not simply to produce more outputs, but to define a more meaningful and clinically relevant safety question. The safety estimand content can often be traced to the protocol sections on objectives, endpoints, and estimands, as well as the safety analysis section of the statistical analysis plan (SAP). Ensuring consistency across these documents and the downstream ADaM implementation is essential for transparent analysis.

Table 1 summarizes the contrast between current common practice and the areas where an estimand-based safety framework may add value.

Aspect	Current Practice	Problems/Challenges	Suggested Need/Focus
Rigor in analysis	High for efficacy, often estimand-based under ICH E9(R1); lower for safety	Safety objectives and analytic approaches are often less explicitly defined	Bring the same level of rigor to safety as to efficacy
How AE data are presented	Often tabulated by frequency or summarized as crude percentages	Results may aggregate heterogeneous trials, populations, or observation periods	Explicitly address trial and patient-level differences
Comparator data	Sometimes limited or not emphasized	Difficult to separate treatment effect from background risk	Use comparator information to clarify interpretation

Aspect	Current Practice	Problems/Challenges	Suggested Need/Focus
Time and duration	Not always standardized across analyses	AE risk depends strongly on exposure and follow-up time	Explicitly account for follow-up duration
Interpretability	Often based on pooled or crude summaries	Causal interpretation may be unclear and risks may be misread	Define a clinically meaningful safety question

Table 1. Current Practice and Key Challenges in Safety Analyses

SAFETY ESTIMAND DEFINITION

The objective of a safety estimand is to define the treatment effect of interest for a safety endpoint in a manner that is clinically relevant, statistically coherent, and operationally implementable. Unlike efficacy, where a single primary estimand typically drives the study design, safety evaluation often requires more than one estimand to address complementary clinical questions about the same endpoint.

In the context of this paper, two complementary estimands are defined for safety analyses, corresponding to two analysis perspectives. The on-treatment estimand targets the risk of an adverse event under a scenario within the period of pharmacological exposure, including active treatment and a prespecified residual risk window after the last dose. The on-study estimand targets the overall safety experience following randomization, regardless of whether treatment is continued — a treatment policy framing that captures the full observed safety profile. These two estimands are not competing alternatives; rather, they address different questions and are intended to be used together.

Attribute	Specification for Safety Estimand
Population	All randomized patients in the target population
Treatment condition	Randomized treatment conditions, preferably aligned with the efficacy estimand.
Endpoint (variables)	Occurrence of a safety event, such as an AE or laboratory abnormality, before a pre-defined timepoint; for repeated safety assessments, the measurement distribution at each timepoint.
Intercurrent events	Premature study treatment discontinuation: addressed using two complementary approaches: treatment policy based on the on-study period, and hypothetical based on the on-treatment period with censoring/exclusion after discontinuation plus a pre-specified residual risk window. Initiation/change of concomitant medication including rescue medication: handled using a treatment policy approach. Death: handled hypothetically by censoring at death and not imputing post-death data.
Population-level summary	Event probability before a pre-defined timepoint, treatment difference, and where relevant HR; for repeated measures, the distribution of measurements at each timepoint.

Table 2. Safety Estimand Specification

The target population for safety is aligned with the population defined for efficacy, unless there is a specific scientific reason to do otherwise. The treatment condition is similarly aligned with the efficacy estimand, reflecting the randomized treatment assignment.

The intercurrent events relevant to safety — premature treatment discontinuation, initiation or change of concomitant medication, and death — are the attributes that most distinguish the two estimands from each other. The strategies selected for each ICE and the corresponding data processing rules are

detailed in the section below.

The population-level summary depends on the endpoint type. For event-based safety endpoints, the primary summary is the probability of experiencing the event before a predefined time point, with treatment-group comparison expressed as a risk difference or hazard ratio. For repeated measurements such as laboratory parameters, the summary may focus on the proportion of patients with clinically notable abnormalities at specific time points. In all cases, the summary measure should be interpretable in relation to the clinical question defined by the estimand.

DATA HANDLING FOR INTERCURRENT EVENTS

To operationalize the safety estimands defined in the above section, each intercurrent event must be addressed in two steps: first, a strategy is selected based on the clinical question of interest; second, that strategy is translated into explicit data processing rules that govern inclusion, exclusion, or censoring of post-event data in the analysis dataset. Table 3 presents the complete mapping of ICE categories, strategies, and data handling rules.

1.1 Premature Study Treatment Discontinuation

Premature treatment discontinuation is the most consequential intercurrent event for safety analyses because it directly determines the length of observed treatment exposure and is the attribute that differentiates the two estimands defined in above section.

For the on-treatment estimand, a hypothetical strategy is applied. This targets the question: what would the adverse event risk have been had the patient continued treatment as planned? Operationally, data observed after treatment of discontinuation are excluded or censored following a prespecified residual risk window (e.g., 30 days after the last dose). The residual window accounts for the possibility that treatment-related events may manifest with a delay. The choice of hypothetical strategy — rather than simply labeling this an "on-treatment" analysis — reflects the ICH E9(R1) taxonomy: because the analysis conditions on a counterfactual scenario (continued treatment), it aligns with the hypothetical strategy definition.

For the on-study estimand, a treatment policy strategy is applied. All data collected after discontinuation are retained regardless of subsequent treatment decisions. This captures the full safety experience following randomization, including any consequences of stopping treatment or switching therapies.

1.2 Initiation or Change of Concomitant Medication

For initiation or change of concomitant medication, including rescue medication, a treatment policy strategy is applied under both estimands. The rationale is that concomitant medications represent part of the real-world treatment experience; their use does not alter the clinical question of interest for safety. Operationally, no censoring or exclusion is introduced at the time of medication change. All post-event data remain in the analysis dataset.

1.3 Death

For non-fatal safety endpoints, death represents a competing event that precludes further observation of the endpoint of interest. A hypothetical strategy is applied under both estimands: observation time is censored at the date of death and no post-death data are imputed. This reflects the clinical question of what the risk of the safety event would be in a setting where death does not preclude its occurrence.

Where mortality itself is part of the safety question (e.g., cardiovascular death as a component of a composite safety endpoint), death would instead be incorporated into the endpoint definition rather than treated as an intercurrent event. Such cases should be addressed on an endpoint-specific basis.

No.	ICE Category	On-Treatment Estimand	On-Study Estimand	Data Processing Rule
1a	Premature study treatment discontinuation	Hypothetical Strategy	-----	Censor/Exclude after discontinuation
1b	Premature study treatment discontinuation	-----	Treatment policy	Include all available data

No.	ICE Category	On-Treatment Estimand	On-Study Estimand	Data Processing Rule
2	Initiation or change of concomitant medication, including rescue medication	Treatment policy strategy	Treatment Policy	Include all available data
3	Death	Hypothetical strategy	Hypothetical	Censor at death; no imputation

Table 3. Intercurrent Event Categories and Data Handling Rules

To Implement the ICE strategies defined above, the corresponding analysis periods must be represented as derivable quantities in the ADaM dataset structure. At the subject level, this requires defining end dates for each analysis period. Table 4 provides an illustrative example.

USUBJID	AP01SDT (Period 01 Start Date)	AP01EDT (Period 01 End Date)	OTR01EDT (On-treatment End Date for Period 01)
ABC001	09OCT2024	08JAN2025	08JAN2025
ABC002	22OCT2024	24APR2025	30DEC2024
ABC003	10OCT2024	03JAN2025	03JAN2025
ABC004	12DEC2024	06MAR2025	06MAR2025

Table 4. Example of Subject-level End Dates in ADSL

For most subjects, the planned period end date (AP01EDT) and the on-treatment end date (OTR01EDT) are identical, indicating that treatment was not prematurely discontinued. For subject ABC002, however, OTR01EDT occurs earlier than AP01EDT because of premature treatment discontinuation — the intercurrent event handled under the hypothetical strategy (ICE 1a). The source for this derivation is the corresponding disposition record in SDTM DS. The corresponding SDTM disposition information for subject ABC002 is shown in Table 5.

USUBJID	DSTERM	DSCAT	DSSCAT	DSSTDTC
AB002	TERMINATION DUE TO STUDY DESIGN	DISPOSITION EVENT	DRAGA/PLACEBO	2024-12-20

Table 5. Example of Associated SDTM DS Record

This example illustrates a key principle: the distinction between the on-treatment and on-study estimands, defined conceptually in terms of ICE strategies, becomes operationally visible as a difference in subject-level end dates. These dates then propagate downstream into analysis of period flags, record-level sub setting, and time-to-event censoring logic. The detailed ADaM implementation of these derivations is described in the section below.

IMPLEMENTING SAFETY ESTIMANDS IN ADAM

Building on the safety estimands and the handling of intercurrent events, the next step is to translate these concepts into analysis datasets. In practice, this requires that estimand-specific decisions, particularly those related to observation windows, intercurrent event handling, and censoring rules, be represented explicitly in derived analysis variables rather than embedded in output-specific programming. ADaM therefore serves as the operational bridge between the scientific question defined in the protocol and SAP and the statistical results ultimately presented.

At the subject level, ADSL provides the foundation for estimand-aligned derivation. Variables such as treatment start date, randomization date, treatment end date, study follow-up end date, disposition-derived discontinuation date, and death date should be available to support the derivation of subject-

specific analysis periods. Of particular importance is the explicit derivation of separate end dates for the on-treatment and on-study perspectives. The on-treatment end date reflects the period during which the subject is considered at risk under the hypothetical strategy for premature treatment discontinuation, often including a prespecified residual risk window after the last dose. The on-study end date reflects the full follow-up period relevant to the treatment policy strategy. By making these derivations visible in ADSL, subsequent analysis datasets can rely on shared and traceable period definitions.

At the record level, source-aligned safety ADaMs such as ADAE, ADLB, and ADVS may include estimand-specific analysis flags that identify whether each observation falls within the relevant observation window. Sponsor-standard variables such as ONTRTFL and ONSTDYFL, or equivalent flags, can be derived from the subject-level dates in ADSL. This approach enables a common set of source records to support multiple estimands without requiring repeated or inconsistent period logic in downstream programs.

For time-to-event safety analyses, a dedicated ADTTE dataset provides the clearest implementation because event definition, censoring logic, and follow-up time are all represented directly in analysis-ready records. This is especially useful when the same safety endpoint is analyzed under both on-treatment and on-study perspectives. In such cases, the distinction between estimands should be reflected through separate parameter records or analysis categories within ADTTE, each with its own derivation of analysis origin, event window, and censoring boundary.

TTE DATA DESIGN FOR SAFETY ANALYSES

Time-to-event analysis is particularly useful in safety evaluation because it incorporates not only whether a safety event occurs, but also when it occurs relative to a defined analysis origin. This feature is important in safety settings where follow-up may vary substantially across subjects because of treatment discontinuation, death, withdrawal, or administrative study end. In such settings, a simple incidence proportion may obscure meaningful differences in observation time, whereas a time-to-event framework aligns more naturally with estimand-based questions that depend on how the risk period is defined.

Within an estimand-aligned safety framework, a time-to-event endpoint is defined through several closely linked components: the analysis origin, the qualifying event, the observation window, and the censoring rule. The analysis origin may be treatment start date, randomization date, or another protocol-defined date, depending on the scientific question. The event of interest is typically the first occurrence of a qualifying safety event within the estimand-specific observation window. The censoring rule reflects the end of follow-up in the absence of the event, including administrative end of observation and any subject-specific truncation implied by the estimand. Accordingly, the estimand is implemented not only through the choice of statistical method, but through the derivation of the time-to-event variables themselves.

CORE TTE STRUCTURE FOR SAFETY ENDPOINTS

ADTTE follows the CDISC ADaM Basic Data Structure (CDISC 2021) and is designed to be analysis-ready, with one record per subject per parameter. For each record, all variables required for statistical analysis are available without additional derivation in the analysis program. This structure is particularly suitable for Kaplan-Meier estimation, Cox proportional hazards modeling, and related survival-type analyses used to summarize safety event risk over time.

For safety applications, the core ADTTE variables should be defined in relation to the endpoint definition and the estimand-specific analysis period. Table 6 summarizes the role of key variables commonly used to implement safety time-to-event endpoints (CDISC 2021).

Variable Name	Variable Label	Type	Core	CDISC Notes
STARTDT	Time to Event Origin Date for Subject	Num	Perm	The original date of risk for the time-to-event analysis. This is generally the time at which a subject is first at risk for the event of interest evaluation (as defined in the Protocol or Statistical Analysis Plan). For example, this may be the randomization date or the date of first study

				therapy exposure.
ADT	Analysis Date	Num	Perm	Analysis date of event or censoring associated with AVAL in numeric format.
AVAL	Analysis Value	Num	Req	AVAL is the elapsed time to the event of interest from the origin. For example, if AVAL is measured in days, AVAL would be ADT – STARTDT + 1.
CNSR	Censor	Num	Cond	CNSR is a required variable for a time-to-event analysis dataset. For example, CNSR = 0 for event and CNSR > 0 for censored records.
EVNTDESC	Event or Censoring Description	Char	Perm	Describe the event of interest or an event that warrants censoring. Values for EVNTDESC will be defined in the metadata as sponsor-defined controlled terminology

Table 6. Core ADTTE Variables for Safety Time-to-Event Endpoints

These variables become more interpretable when linked to a concrete derivation example. Consider the endpoint time to first serious adverse event for subject ABC002 who experienced two adverse events: a non-serious event before treatment discontinuation and a serious adverse event after the on-treatment window but before the end of study follow-up. Table 7 summarizes the key dates and source adverse event records relevant to the derivation.

USUBJID	TRTSDT	OTR01 EDT	AESTDTC	AEDE COD	AESER	ONTRTFL	ONST DYFL
ABC002	22OCT2024	30DEC2024	15DEC2024	Nausea	N	Y	Y
ABC002	22OCT2024	30DEC2024	10JAN2025	Pneumonia	Y	N	Y

Table 7. Illustrative Source Data and Analysis Eligibility for Time to First Serious Adverse Event

In this example, the first adverse event occurs within both the on-treatment and on-study observation windows, but it does not qualify for the endpoint because it is not serious. The second adverse event qualifies as a serious adverse event and occurs within study follow-up, but after OTR01EDT. It is therefore eligible under the on-study estimand but not under the on-treatment estimand. The resulting records are shown in Table 8.

USUBJID	PARAM	PARCAT3	STARTDT	ADT	CNSR	EVNTDESC	CNSDT DSC	AVAL
ABC002	Time to First Serious Adverse Event	On-treatment	22OCT2024	30DEC2024	1	First SAE	End of on-treatment period	70
ABC002	Time to First Serious	On-study	22OCT2024	10JAN2025	0	First SAE		81

USUBJID	PARAM	PARCAT3	STARTDT	ADT	CNSR	EVENTDESC	CNSDTDSC	AVAL
	Adverse Event							

Table 8. Illustrative Records for Time to First Serious Adverse Event

This example shows how the core ADTTE variables are derived in an estimand-specific way. Under the on-treatment estimand, no serious adverse event is observed within the relevant observation window, so ADT is the censoring date, CNSR = 1, and CNSDTDSC records the reason for censoring. Under the on-study estimand, the serious adverse event on 10JAN2025 remains eligible, so ADT is the event date and CNSR = 0. The distinction between estimands is therefore reflected directly in the analysis-ready data rather than embedded only in downstream analysis code. In practice, the same logic can be extended to other safety endpoints and, for adverse event analyses, can be implemented within a dedicated ADTTEAE structure.

INTEGRATED ADTTEAE STRUCTURE FOR SAFETY TIME-TO-EVENT ANALYSES

A practical approach for safety time-to-event analyses is to represent distinct safety questions through separate parameters within a common ADTTE framework. Within this structure, each record represents a subject-level time-to-event endpoint evaluated over a defined analysis period. The dataset integrates the endpoint definition, adverse event category, MedDRA hierarchy level, analysis period, and core time-to-event variables in a standardized manner. This design supports both overall and hierarchical adverse event analyses and makes the derivation logic explicit in the analysis dataset.

Parameterization and Hierarchical Representation

Within ADTTEAE, distinct safety questions may be represented as separate parameters within a common analysis structure. These parameters may correspond to clinically meaningful adverse event categories, such as all adverse events, serious adverse events, adverse events leading to discontinuation of investigational product, or events defined by severity. In practice, the analysis parameter may be represented in PARAM, with broader adverse event groupings captured in PARCAT1 or a comparable sponsor-defined categorization variable. Sponsors may implement these analyses either within a single parameterized dataset or as a family of structurally similar datasets, such as an ADTTEAExx series in Table 9, with each dataset corresponding to a predefined adverse event category. Regardless of implementation, the underlying principle remains the same: each parameter should correspond to a clearly defined safety question.

Dataset	Category (PARCAT1)	Value Levels
ADTTEAEA	All AEs	Time to first AE of any type
ADTTEAEB	SAEs	Time to first serious adverse event
ADTTEAEC	SAEs with outcome death	Time to first fatal SAE
ADTTEAED	AEs leading to discontinuation of IP	Time to first AE leading to treatment discontinuation
ADTTEAEE	AEs possibly related	Time to first treatment-related AE
ADTTEAEF	AEs with any intensity	Time to first AE with intensity
ADTTEAEG	AEs with moderate or severe intensity	Time to first moderate/severe AE
ADTTEAEH	AEs with severe intensity	Time to first severe AE
ADTTEAEI	AEs with missing intensity	Time to first AE with missing intensity

Table 9. Illustrative Parameter Categories for Adverse Events Analyses

The same adverse event category may also be evaluated at different levels of the MedDRA hierarchy, since an effect observed at the overall category level may be driven by one or more specific organ systems or preferred terms. To support this, multiple hierarchy levels may be represented

through PARCAT2 or an equivalent categorization variable in Table 10. At the broadest level, an overall record represents time to first qualifying event within the category of interest. More detailed records may then be created for System Organ Class (SOC) using AESOC or AESOCCD, for High Level Group Term (HLGT) using AEHLGT or AEHLGTCD, for High Level Term (HLT) using AEHLT or AEHLTCD, and for Preferred Term (PT) using AEDECOD or AEPTCD. At each hierarchy level, the endpoint remains the time to first qualifying event within that level, so a subject may contribute multiple records within the same analysis framework.

PARCAT2 Level	Analysis Unit	Source Variables	Record Definition
Overall	Subject within safety category and analysis period	Dataset-level AE subset	One record per subject per parameter and analysis period, representing time to first qualifying event in the category
SOC	Subject within System Organ Class and analysis period	AESOC, AESOCCD	One record per subject per SOC per analysis period, representing time to first qualifying event within the SOC
HLGT	Subject within High Level Group Term and analysis period	AEHLGT, AEHLGTCD	One record per subject per HLGT per analysis period, representing time to first qualifying event within the HLGT
HLT	Subject within High Level Term and analysis period	AEHLT, AEHLTCD	One record per subject per HLT per analysis period, representing time to first qualifying event within the HLT
PT	Subject within Preferred Term and analysis period	AEDECOD, AEPTCD	One record per subject per PT per analysis period, representing time to first qualifying event within the PT

Table 10. Hierarchical Representation of Safety Event Parameters in ADTTEAEx

Analysis Period Representation

The distinction between the on-treatment and on-study estimands becomes operational in ADTTEAE through the explicit representation of the analysis period. This may be implemented using PARCAT3 or another sponsor-standard categorization variable that identifies the analysis period associated with each record.

For the on-treatment period, the analysis origin is typically TRTSDT, and follow-up continues through the subject-specific on-treatment end date, such as OTR01EDT, which incorporates treatment discontinuation and any predefined residual risk window. Only adverse events occurring on or before that date are eligible to be counted. For the on-study period, the analysis origin is generally RANDDT or AP01SDT, and follow-up continues through the end of study observation, such as AP01EDT. Under this treatment policy perspective, premature treatment discontinuation does not truncate follow-up, and post-discontinuation events remain eligible if they occur within the study's observation window.

This design results in paired records for the same subject and endpoint under different estimands. For subjects who complete treatment as planned, the on-treatment and on-study records may be identical. For subjects who discontinue treatment prematurely, however, the event status, event or censoring date, and derived time-to-event value may differ. The estimand distinction is therefore fully visible in the analysis dataset and does not depend on hidden logic in analysis programs.

Similar estimand-based design principles also extended to other sponsor-defined safety time-to-event datasets, including laboratory, vital signs, and electrocardiogram-based structures, although the qualifying event definitions in those settings depend on measurement-based abnormality criteria rather than clinically reported adverse events.

PRACTICAL CONSIDERATIONS FOR STATISTICAL PROGRAMMERS

Several implementation considerations are important when designing safety ADTTE datasets. First, derivations of analysis period boundaries should be centralized in ADSL and reused consistently across all downstream safety datasets. This reduces the risk of discrepancies between analyses and ensures that all estimand-aligned outputs rely on the same subject-level definitions.

Second, traceability to source records should be preserved wherever possible. When multiple adverse event records qualify for a given parameter, the algorithm used to identify the first qualifying event should be clearly defined and reproducible. This includes explicit handling of ties, partial dates, and competing dates such as death or end of follow-up.

Third, parameter naming and categorization should be sufficiently structured to support both analysis and review. The combination of PARAMCD, PARCAT1, PARCAT2, and PARCAT3 should allow the user to identify the event category, hierarchy level, and estimand-specific analysis period directly from the dataset. This minimizes dependence on external documentation when interpreting records.

Fourth, care should be taken to distinguish between intercurrent events and endpoint components. Death, for example, may function as a censoring event for a nonfatal safety endpoint, but it may also be part of a composite endpoint or a separate safety endpoint in its own right. The ADTTE implementation should reflect this distinction consistently in event definitions and censoring codes.

Finally, scalability should be considered in the dataset design. A robust ADTTExx structure should allow extension to new safety topics, additional MedDRA-based groupings, or alternative estimands without changing the underlying framework. This is particularly important in development programs where safety review questions evolve across studies and over time.

CONCLUSION

The ICH E9(R1) estimand framework provides a useful foundation for improving the clarity and interpretability of safety analyses in clinical trials. For safety endpoints, intercurrent events such as treatment discontinuation, concomitant medication changes and death can substantially affect what is observed and how results should be interpreted. Explicit definition of safety estimands helps ensure that the clinical question of interest is clearly aligned with the corresponding analytical approach.

This paper has examined how estimand principles can be applied to safety evaluation and how those principles can be implemented within the CDISC ADaM framework. By representing estimand-specific analysis periods, event definitions, and censoring rules explicitly in ADaM, especially in ADTTE, the distinction between on-treatment and on-study perspectives becomes transparent and traceable. Such an approach improves consistency in programming, reduces reliance on output-level logic, and supports more reproducible and scientifically meaningful safety analyses.

REFERENCES

- International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use. 2019. ICH E9(R1) Addendum on Estimands and Sensitivity Analysis in Clinical Trials to the Guideline on Statistical Principles for Clinical Trials. Geneva: ICH.
- Clinical Data Interchange Standards Consortium. 2016. *ADaM Basic Data Structure (BDS) for Time-to-Event (TTE) Analyses*. Version 1.0. Austin, TX: CDISC.
- Ratitch, B., Bell, J., Mallinckrodt, C., et al. 2020. "Choosing Estimands in Clinical Trials: Putting the ICH E9(R1) into Practice." *Therapeutic Innovation & Regulatory Science*, 54:324–341.
- Hedman K, Lisovskaja V and Nyström P. A safety estimand for late phase clinical trials where the analysis period varies over the subjects. *Clin Trials*. 2024; 21(4):483-490.

ACKNOWLEDGMENTS

The authors would like to sincerely thank Haoyong Yu, senior statistical programmer II at AstraZeneca for his invaluable contributions to this project and constructive review. Special thanks to our supervisor Yiwen Wang for her supports on this manuscript.

RECOMMENDED READING

- *CDISC ADAM IG v1.3*
- *CDISC ADAM OCCDS v1.1*

CONTACT INFORMATION

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

Kun Ren
AstraZeneca R&D China
kun.ren1@astrazeneca.com

Zeyu Guang
AstraZeneca R&D China
zeyu.guang@astrazeneca.com

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