



NAS RESEARCH / COMPUTATIONAL GENOMICS

The ClinVar Conflict Atlas

Distribution, severity and observed persistence of public germline classification disagreement

This observational computational study measures germline classification disagreement in three preserved ClinVar archives. In the latest archive, 179,613 eligible variant-condition-inheritance units include 224 positive-versus-benign disagreements (0.125%). Disagreement involving uncertain significance and adjacent categories is reported separately. Historical follow-up begins at baseline observation, and loss of eligibility is distinguished from resolution. Source representation differences are retained and assessed for cohort impact. The study characterizes public records; it does not reassess pathogenicity or estimate patient risk.

DATE / VERSION / AUTHOR

2026-10-10 / v1.0.1 / review candidate / NaS

NOT PEER REVIEWED; INDEPENDENT REVIEW PENDING

Unpublished review candidate prepared for public scrutiny. Internally prespecified, not externally registered. Editorial corrections preserve the scientific code, cohorts and estimates. Historical archives are analyzed as retrieved; original-distribution byte identity has not been independently established.

Observed disagreement

The latest archive contains 179,613 eligible condition-inheritance units across 154,037 distinct variants. These are submitted-record denominators, not patient or population denominators.

Definition	Eligible units	Fraction (model-based 95% interval)
Positive versus benign	224 / 179,613	0.125% (0.109% to 0.141%)
Positive opposed by benign or uncertain	5,517 / 179,613	3.072% (2.989% to 3.158%)
Any cross-band disagreement	32,684 / 179,613	18.197% (18.004% to 18.387%)
Adjacent-only category difference	34,003 / 179,613	18.931% (18.738% to 19.122%)

Unit intervals use the prespecified VariationID-cluster bootstrap. Counts are exact for the preserved eligible records; the intervals describe a model beyond that finite archive.

XML snapshot date	Eligible units / variants	T1 units	T1/T2/T3 units
2025-08-31	159,438 / 137,134	214	29,626
2026-03-02	169,976 / 146,266	223	31,879
2026-08-29	179,613 / 154,037	224	32,684

Snapshot-level differences combine changing archive coverage and eligibility with changing classifications. Paired historical correspondence is assessed separately.

Distinct-variant estimates

Definition	Distinct variants	Fraction (model-based 95% interval)
Positive versus benign	224 / 154,037	0.145% (0.128% to 0.166%)
Positive opposed by benign or uncertain	5,207 / 154,037	3.380% (3.291% to 3.472%)
Any cross-band disagreement	31,161 / 154,037	20.230% (20.030% to 20.431%)
Highest severity: adjacent only	31,073 / 154,037	20.172% (19.973% to 20.374%)

Each variant is counted once at its highest-severity eligible unit. The T4 variant estimate excludes variants containing any T1/T2/T3 unit. The distinct-variant denominator comprises variants with at least one eligible multi-organization unit; it is not all ClinVar variants.

These model-based 95% intervals use the prespecified Wilson method. Unit estimates on the preceding page use variant-cluster bootstrap intervals and a different denominator.

Research question

This observational computational study measures disagreement among submitted germline variant classifications in three preserved ClinVar archives. Its principal estimand is the fraction of eligible variant-condition-inheritance units containing a pathogenic or likely pathogenic classification opposed by a benign or likely benign classification. Secondary estimands describe disagreement involving uncertain significance, adjacent categories, and distinct variants with at least one affected unit.

The study characterizes public records. It does not reassess pathogenicity, determine which organization is correct, measure patient outcomes, or estimate clinical risk. Gene, condition, organization and resolution-association analyses are descriptive exploratory screens.

Sources and prespecification

The source set includes matched VCV and RCV XML, variant-summary, submission-summary and conflicting-interpretation tables for September 2025, March 2026 and September 2026. The XML internal snapshot dates are August 31, 2025; March 2, 2026; and August 29, 2026. These dates determine elapsed follow-up and assertion evaluation ages. Archive month, posting date and retrieval date are distinct.

The protocol and four dated addenda were internally frozen with recorded checksums. This was internal prespecification, not external registration. The source-reconciliation addendum followed targeted discrepancy inspection and preceded study-level outcome distributions. The original protocol and earlier editions remain preserved. Parser and plotting repairs are disclosed separately. Cross-sectional cohort and severity definitions remain unchanged. After outcome distributions had been inspected, unresolved declared-versus-present assertion counters prompted an additional conservative source-quality restriction on historical comparability; its timing and affected records are reported explicitly.

Frozen source files retain URLs, sizes, hashes and retrieval metadata. Historical results describe the archived bytes as retrieved. Original-distribution byte identity has not been independently established; later archive modification dates limit claims about exactly what users saw at the original release.

Cohort construction

The primary cohort comprises current, directly classified human SimpleAlleles with germline classifications, identified organizations, recognized standard labels, provided assertion criteria, and exact membership in a specific resolvable MedGen condition set. Exclusively somatic origins are excluded; mixed and unspecified origins are retained. Structural and unmapped SimpleAlleles remain eligible when other requirements are met. Haplotype or genotype classifications are not projected onto their component alleles.

A unit combines VariationID, release-specific RCV membership and the exact normalized inheritance set. Complete condition sets remain intact. Known inheritance modes and unspecified inheritance are separate strata. At least two qualifying organizations are required after filtering. Repeated submissions from one organization cannot create between-organization disagreement. An organization with different standardized classification codes within a unit is excluded from primary comparisons and retained in a separate sensitivity definition.

Labels are preserved verbatim alongside prespecified mappings. P and LP are the positive band; B and LB are the benign band; U is uncertainty. Recognized within-band compounds retain their compound identity. VUS-high, VUS-mid and VUS-low map to U while preserving the subtype. Special, cross-band compound and unrecognized labels are not converted into additional votes.

Disagreement and denominators

Each between-organization pair is assigned T1 for positive versus benign, T2 for positive versus uncertain, T3 for uncertain versus benign, or T4 for P versus LP or B versus LB. The unit receives its highest present category in that order. T0 means no derived disagreement; it does not require identical source text. T4 is adjacent-category disagreement and remains separate from cross-band disagreement.

Unit proportions count each eligible condition-inheritance unit once. Variant proportions count each distinct VariationID once, using any affected eligible unit as the event. The variant denominator includes only variants with at least one eligible multi-organization unit. Neither denominator represents all ClinVar variants, tested patients or a population sample. Overlapping strata and fractional allocations are identified explicitly.

Estimation and exploratory analyses

Finite-record counts are exact for the preserved cohort. Unit-proportion intervals use 2,000 percentile bootstrap resamples clustered by VariationID with seed 20260926, retaining repeated units from the same variant together. Distinct-variant proportions use Wilson intervals. These intervals describe model-based uncertainty beyond the observed database; they do not turn record counts into patient prevalence estimates.

Gene screens use eligible single-gene variants and the presence of any T1/T2 unit. Expected counts sum leave-one-gene-out reference rates stratified by organization count, variant type, expert involvement and maximum valid evaluation age. Reference strata below 50 variants pool age, then variant type, then organization count. Genes below 30 eligible variants are not ranked, and incomplete reference coverage remains visible. No multiplicity-adjusted significance or causal interpretation is claimed.

Condition analyses retain complete identifier sets; a variant may contribute to multiple sets. Fractional weights divide disputed variants across their gene or condition memberships. Organization analyses divide each disputed unit across its qualifying organizations and report participation relative to eligible units containing that organization. These quantities describe participation and ascertainment, not responsibility for error. Collection method and reported organization type are overlapping source categories.

Evaluation ages and sensitivities

Evaluation dates remain separate from creation, submission and record-update dates. Missing, future or unsupported date representations do not contribute an age and are never replaced by update dates. Cross-sectional eligibility does not require a valid evaluation date. Oldest and newest evaluation ages, completeness, and recent-versus-old pair patterns are reported; evaluation age is not duration of disagreement.

The twelve definitions comprise primary; all review levels; explicit germline origin; clinical-testing method; exclusion of expert/guideline assertions; allowance of within-organization inconsistency; condition/inheritance pooling; RCV/inheritance pooling; variant-wide grouping; exclusion of compounds and VUS subtypes; and five- and two-year evaluation windows. Age-filter retention preserves the original primary denominator and separates reduced disagreement from loss of the two-organization eligibility requirement.

Historical comparisons

Historical correspondence requires stable VariationID, complete condition identifiers and inheritance, unique matching, and compatible sequence and accession-lineage metadata. Missing variants, unavailable condition identities, lost eligibility and non-comparable records remain separate from resolution. Proven paired-source assertion absence and unresolved declared-versus-present assertion counts have distinct coverage statuses. Comparisons touching either status are non-comparable and cannot generate resolution events. The counter-based restriction was added after outcome inspection as conservative quality control, while preserving observed cross-sectional results.

A cross-band unit resolves only when a later comparable observation is eligible at T0 or T4; those two destination categories remain distinct. A change from T1 to T2 is a severity reduction, not cross-band resolution. Successive snapshots reveal observed persistence and interval-observed changes, not uninterrupted disagreement or transient changes between releases.

Among 29,626 baseline prevalent cross-band units, 719 have a first observed eligible T0/T4 state: 452 at 183 days and 267 first observed at 363 days. Another 682 lack eligible comparable endpoint follow-up before an observed resolution. Seven of the units observed at T0/T4 in the intermediate archive return to an eligible cross-band state in the final archive; consequently, 712 baseline units have a resolved final endpoint. First observed resolution and resolved final endpoint are different outcomes. These counts identify changes at the three observed archives and do not establish clinical correctness, permanent resolution or the absence of intervening changes.

Exploratory time-to-first-observed-resolution estimates begin at baseline detection, not true conflict onset. When at least 50 resolutions are observed, the interval-censored Turnbull fit retains support intervals, masses and numerical diagnostics. Right-continuous survival bounds are conditional on fitted masses; they are not sampling confidence intervals or bounds over all possible nonunique likelihood solutions. The independent-censoring assumption may be strained by loss of eligibility or database changes.

Expert involvement and record updates in the previous 365 days are measured at interval start. Unknown update dates remain unknown. Associations report observed comparable-pair denominators, losses, clustered risk-difference intervals and odds ratios when estimable without a zero-cell correction. These are observational associations. Changes accompanying resolution do not establish what caused it.

Validation framework

Validation separates extraction fidelity, implementation agreement and scientific interpretation. Release-specific XSD checks, identifier and membership checks, and paired-source reconciliation precede cohort construction. A separate literal SQL implementation reconstructs the primary units and categories without calling the production grouping or tier functions. Exact numerical agreement is required; shared normalized inputs mean this is computational agreement, not independent clinical validation.

Matched tabular files, official aggregate fields and source-provided VCF representations provide additional checks with their own scopes. All nonzero discrepancies are retained, reviewed for cohort impact and disclosed. XML remains the prespecified primary source. Official aggregate conflict and this study's derived condition-inheritance disagreement are different quantities; an unsuccessful reconstruction of an official rule cannot be called validation of that rule.

The case sample selects the twenty lowest seeded hashes in each of T1, T2, T3, T4, T0 and excluded strata, with shortfalls disclosed. Independent raw VCV and RCV field extraction is followed by explicit inspection of labels, organizations, membership, conditions, inheritance, dates, review status and aggregate representation. Examples are the first two validated cases per tier in the prespecified order. Source verification is performed within the same analysis workflow and does not constitute separate human, expert or clinical review.

The GRCh38 coordinate check compares AlleleID and the source-provided VCF position and allele fields by VariationID. XML HGVS start/stop strings are not substituted for left-normalized VCF fields. This tests agreement between supplied representations; it does not independently normalize alleles against a reference sequence. Coordinate strings do not determine cohort or historical identity, and structural or unmapped variants remain eligible outside the VCF scope.

All 120 prespecified cases were explicitly inspected against independently extracted raw VCV fields and actual RCV membership. The sample contains 20 cases in each of T1, T2, T3, T4, concordant and excluded strata; all sampled source fields and derived grouping rules agreed. This inspection occurred within the same computational workflow. No independent human, expert, clinical or peer review has been completed.

Sensitivity definitions

Definition	Eligible units	T1 units	T1 fraction
all review	211,605	740	0.350%
allow internal	179,613	224	0.125%
clinical testing	160,619	143	0.089%
condition inheritance pooled	197,762	291	0.147%
explicit germline	140,083	162	0.116%
no compound subtype	179,607	224	0.125%
no expert	176,038	222	0.126%
primary	179,613	224	0.125%
rcv inheritance pooled	197,762	291	0.147%
recent 2y	22,047	23	0.104%
recent 5y	99,978	87	0.087%
variant wide	850,502	1,265	0.149%

Each definition has its own eligible denominator. Age-filter attrition on the original primary denominator is reported separately; loss of a second qualifying organization is not resolution. Pooling changes the scientific comparison and must not be interpreted as a correction of condition-specific results.

Cohort attrition

6,959,945 extracted assertions receive one first-applicable disposition. 2,826,024 pass assertion-level filters, before organization consistency handling and the two-organization requirement.

First applicable disposition	Assertions
Nonspecific or unresolved condition	3,126,390
Passes assertion-level filters	2,826,024
Outside the germline classification branch	534,697
Assertion criteria not provided	426,307
Nonstandard or unsupported label	34,487
Exclusively somatic origin	8,530
Not directly classified SimpleAllele	3,510

These dispositions are mutually exclusive and sum to the extracted assertion total. The companion overlapping reason flags answer a different question. Assertion-level passage does not imply final unit eligibility.

Gene screen coverage

2,975 genes have eligible single-gene variants. 691 meet the 30-variant descriptive volume threshold; 691 of these have complete leave-one-gene-out reference coverage. Expected counts remain non-estimable where pooled reference coverage is incomplete.

Gene	Eligible variants	Observed T1/T2	Expected T1/T2	Observed / expected
BRCA2	7,770	52	284.95	0.18
APC	4,597	12	171.32	0.07
BRCA1	3,884	43	127.26	0.34
NF1	3,109	111	115.56	0.96
TSC2	2,958	29	121.44	0.24

The five examples are selected by the largest eligible denominators, with GeneID as the tie breaker. They are not a ranking of disease risk, classification accuracy or statistical significance. Standardization describes the archive and cannot remove its ascertainment bias.

Concentration of disputed participation

Domain and observed unit	HHI	Top-five share	Top-ten share
Genes: fractional variant participation	0.0060	12.58%	19.07%
Condition sets: fractional variant participation	0.0093	15.47%	23.88%
Organizations: fractional unit participation	0.0587	39.42%	54.42%

The gene and condition fractional weights sum to 5,207 and 5,207 disputed variants, respectively. Organization weights sum to 5,517 disputed units.

Shares use the domain-specific fractional T1/T2 participation denominator. Gene membership can be missing; condition sets may overlap; organization participation uses eligible units rather than distinct variants. HHI sums squared normalized participation shares. Different domains cannot be interpreted as interchangeable samples or as attribution of error.

Evaluation-age completeness

Evaluation-age coverage	Observed count
Eligible units	179,613
Units with at least one unavailable age	6,579
Units with no valid evaluation age	103
Unavailable assertion-unit ages	7,089 / 425,878

Assertion-unit counts represent qualifying assertion occurrences in unit contexts, not unique patients, organizations or SCV accessions. Unavailable ages include absent or unusable evaluation dates. Oldest/newest curves include a unit when at least one valid age is available; they do not require complete dates for every assertion.

Cross-band unit pair indicator	Units / cross-band units
At least one pair with both sides within two years	2,296 / 32,684
At least one pair: over five years versus within two years	10,448 / 32,684
At least one pair with both sides over five years	2,755 / 32,684

Pair indicators can overlap within the same unit. Only between-organization cross-band pairs with valid ages on both sides can satisfy an age indicator. Their absence does not prove an old or recent origin of disagreement.

Age-filter retention

Both rows begin with the same 32,684 primary cross-band units. The filter changes which assertions and organizations can participate in a unit; it does not observe a temporal resolution event.

Filter applied to primary cross-band units	Cross-band retained	Eligible T0/T4 after filtering	Insufficient organizations
Two-year evaluation window	2,296	929	29,459
Five-year evaluation window	16,282	1,372	15,030

The age-window modes recompute eligible denominators across their full candidate cohorts. This retention table instead follows the original primary cross-band units and separates loss of qualifying organizations from remaining eligible at T0/T4. Reduced disagreement after filtering is age-sensitive under that definition, without establishing that old evidence caused it.

Observed historical changes

The baseline prevalent cross-band cohort contains 29,626 uniquely matched units. The archives show 719 first eligible T0/T4 observations and 682 units lost or non-comparable before an observed resolution.

Quantity	Observed value
Observed first-resolution fraction	2.427%
Upper loss-to-follow-up sensitivity fraction	4.729%

Only snapshot-observed resolution is identified. Transient resolution and recurrence between snapshots cannot be excluded. Follow-up loss bounds do not bound all unobserved transient events.

Follow-up begins with the baseline observation of a prevalent cross-band disagreement; it does not identify the true onset of disagreement.

The interval-censored Turnbull fit uses an experimental library implementation. Its envelopes bound survival conditional on the fitted support masses; they are not confidence intervals.

Accompanying assertion removals, additions, label changes and update-date associations describe record changes. They do not determine the cause of disagreement or resolution.

Historical associations: 2025-09 to 2026-03

The interval-start cross-band cohort contains 29,626 eligible units. Exposure is measured at the start archive. Each exposure table retains the full starting cohort; rates and risk differences condition on observed comparable follow-up. A unit is recent when any qualifying interval-start SCV has a valid update age within 365 days; older requires all qualifying SCV update ages to be valid and over 365 days. With no valid recent update, an unavailable date or no dates creates the unknown category, excluded from the recent-versus-older effect comparison. A recent unit can contain another SCV with an unavailable update date.

Start exposure / category	Resolved	Persistent	Lost / non-comparable
Expert / Absent	445	28,660	289
Expert / Present	7	218	7
Prior update / Older	91	6,572	33
Prior update / Recent	361	22,306	263
Prior update / Unknown date	0	0	0

Start exposure	Resolved / observed: exposed; comparison	Risk difference (percentage points)	Odds ratio
Expert involvement at interval start	7 / 225; 445 / 29,105	1.582 (-0.533 to 4.038)	2.068
SCV update within preceding 365 days	361 / 22,667; 91 / 6,663	0.227 (-0.100 to 0.544)	1.169

The expert comparison is present versus absent; the update comparison is recent versus older. Parenthesized ranges are model-based 95% intervals from 2,000 VariationID-cluster bootstrap resamples. Odds ratios are descriptive point estimates without a zero-cell correction or reported confidence interval. An SCV update date does not establish a new clinical evaluation. The comparisons are exploratory observational associations; ascertainment, changing eligibility and loss can bias them. They do not establish a causal benefit of expert involvement or recent updates.

Historical associations: 2026-03 to 2026-09

The interval-start cross-band cohort contains 31,879 eligible units. Exposure is measured at the start archive. Each exposure table retains the full starting cohort; rates and risk differences condition on observed comparable follow-up. A unit is recent when any qualifying interval-start SCV has a valid update age within 365 days; older requires all qualifying SCV update ages to be valid and over 365 days. With no valid recent update, an unavailable date or no dates creates the unknown category, excluded from the recent-versus-older effect comparison. A recent unit can contain another SCV with an unavailable update date.

Start exposure / category	Resolved	Persistent	Lost / non-comparable
Expert / Absent	289	30,908	412
Expert / Present	8	253	9
Prior update / Older	112	9,977	159
Prior update / Recent	185	21,184	262
Prior update / Unknown date	0	0	0

Start exposure	Resolved / observed: exposed; comparison	Risk difference (percentage points)	Odds ratio
Expert involvement at interval start	8 / 261; 289 / 31,197	2.139 (0.252 to 4.353)	3.382
SCV update within preceding 365 days	185 / 21,369; 112 / 10,089	-0.244 (-0.483 to -0.005)	0.778

The expert comparison is present versus absent; the update comparison is recent versus older. Parenthesized ranges are model-based 95% intervals from 2,000 VariationID-cluster bootstrap resamples. Odds ratios are descriptive point estimates without a zero-cell correction or reported confidence interval. An SCV update date does not establish a new clinical evaluation. The comparisons are exploratory observational associations; ascertainment, changing eligibility and loss can bias them. They do not establish a causal benefit of expert involvement or recent updates.

Historical associations: 2025-09 to 2026-09

The interval-start cross-band cohort contains 29,626 eligible units. Exposure is measured at the start archive. Each exposure table retains the full starting cohort; rates and risk differences condition on observed comparable follow-up. A unit is recent when any qualifying interval-start SCV has a valid update age within 365 days; older requires all qualifying SCV update ages to be valid and over 365 days. With no valid recent update, an unavailable date or no dates creates the unknown category, excluded from the recent-versus-older effect comparison. A recent unit can contain another SCV with an unavailable update date.

Start exposure / category	Resolved	Persistent	Lost / non-comparable
Expert / Absent	698	28,024	672
Expert / Present	14	208	10
Prior update / Older	150	6,458	88
Prior update / Recent	562	21,774	594
Prior update / Unknown date	0	0	0

Start exposure	Resolved / observed: exposed; comparison	Risk difference (percentage points)	Odds ratio
Expert involvement at interval start	14 / 222; 698 / 28,722	3.876 (0.687 to 7.445)	2.702
SCV update within preceding 365 days	562 / 22,336; 150 / 6,608	0.246 (-0.173 to 0.669)	1.111

The expert comparison is present versus absent; the update comparison is recent versus older. Parenthesized ranges are model-based 95% intervals from 2,000 VariationID-cluster bootstrap resamples. Odds ratios are descriptive point estimates without a zero-cell correction or reported confidence interval. An SCV update date does not establish a new clinical evaluation. The comparisons are exploratory observational associations; ascertainment, changing eligibility and loss can bias them. They do not establish a causal benefit of expert involvement or recent updates.

Record changes accompanying resolution

Units with accompanying change	Baseline to March	March to September	Baseline to September
Resolved interval endpoints	452	297	712
Baseline SCV label revision	401	135	513
Baseline SCV absent at endpoint	13	126	139
Explicit endpoint deletion record	13	126	139
Present SCV no longer qualifying in same unit	47	40	72
Newly qualifying SCV	23	29	97
Baseline SCV review-status change	34	35	56
Baseline SCV evaluation-date change	420	170	627
Qualifying organization removed	60	166	211
Qualifying organization added	23	29	97

Each cell counts resolved units with at least one accompanying change, rather than the number of assertions or patients. Categories overlap and must not be summed into causal mechanisms. Denominators are comparable cross-band-to-T0/T4 interval endpoints, not the 719 first-observed-resolution events. A unit can contribute to more than one interval comparison.

An absent baseline SCV is distinguished from a present record no longer qualifying for the same condition-inheritance unit. Explicit deletion requires an endpoint DeletedSCV entry. Newly qualifying SCVs need not be newly created. Classification, review, membership, inheritance and organization changes can alter eligibility. These snapshots do not establish the exact order, timing or cause of the observed changes.

Source representation and validation

XML is the prespecified source for cohort construction. Matched tabular files test representation agreement, but differences are retained rather than overwritten. Legacy labels, versions, duplicate SCV rows and evaluation dates have been reviewed for assertion and variant-level cohort impact. A source discrepancy outside the qualifying assertions does not justify excluding every eligible unit at the same variant.

The normalized category `somatic_or_other` records absence of the modeled `GermlineClassification` element. An empty `germline-label` field and the recorded `no classification` provided review status describe the inspected extraction fields; this review does not adjudicate labels in other classification branches. The category alone does not establish biological sample origin. Germline classification branch and reported origin are separate primary-cohort fields.

The matched March submission table omits 530,566 assertions without a modeled germline classification present in VCV. None is a retained primary assertion. The baseline table omits 6,665 assertions without a modeled germline classification, also outside qualifying primary assertions; otherwise eligible units at some of those variants remain included. These observations establish representation differences between preserved files, not their underlying cause.

The September submission table lacks matching rows for 531,028 XML assertions with no germline label and no classification provided. None contributes to an eligible primary unit. The smaller version, label, duplicate, evaluation-date and pair-label differences also implicate no eligible primary units. Contribution flags differ for 1,045 assertions without a modeled germline classification; their variants contain 344 otherwise eligible primary units, which remain retained. Contribution-to-aggregate is metadata rather than a primary eligibility or severity input. The direct variant, `AlleleID`, aggregate-label/review and pair-identity comparisons show zero differences. These findings characterize the preserved files and do not establish why the representations differ.

The missing-assertion variants contain 3,836 otherwise eligible primary units in March (2,942 distinct variants) and 3,964 in September (3,009 distinct variants). Among those retained units, 665 and 673, respectively, have T1/T2/T3 disagreement. The missing assertions themselves are assertions without a modeled germline classification outside primary eligibility. Sharing a variant with those records does not make the qualifying assertions absent or invalidate the observed unit tiers. These units remain retained; variant-overlap counts describe review scope rather than changed outcomes. The missing-assertion and contribution-flag scopes can overlap and must not be summed as unique affected units.

The matched GRCh38 VCF check evaluates 4,468,152 records: 4,466,321 have exact source-provided VCF fields, 688 are not comparable as directly classified `SimpleAlleles`, and 1,143 retain representation differences. Of those differences, 1,064 have the same chromosome, position and reference allele but XML REF equals ALT while VCF ALT is a dot; 51 use a contig accession in VCF in place of the XML chromosome label with matching inspected position and alleles; and 28 replace non-ACGTN XML alternate bases with N in VCF. The dot marker denotes no alternate allele in the VCF 4.1 specification. All 27 otherwise eligible primary units at discrepant variants occur in the first group: 25 T0 and two T3. Separately, 63 eligible variants covering 64 units are absent from VCF (57 T0, two T2 and five T4). These units remain retained under the XML-based cohort. Exact differences

remain nonzero; source-field agreement does not constitute independent reference-aware normalization or clinical adjudication, and the less specific N representation does not prove exact allele equivalence.


Declared submission totals sometimes exceed the number of current assertion elements physically present in the classified VCV records. An independent raw-file count audit confirms agreement with the normalized extracted counts. In the baseline archive, 31 records each have one fewer present element than declared; the same 28 records are affected in March and September. Three eligible baseline units and one unit in each later archive occur at affected records. None of these units has observed T1/T2/T3 disagreement. Missing assertion identities, organizations and classifications are not inferred from these counters. Their cause remains unresolved. The examined schema and NCBI guides do not establish a formal equality contract between this counter and current assertion-element cardinality. The difference alone is not proof of an omitted assertion.

The two baseline RCV-only SCVs identified in the frozen reconciliation addendum are separately documented evidence of paired-source absence. A declared-versus-present counter discrepancy alone is weaker evidence and receives a distinct unresolved-source-count status. The primary cross-sectional estimand uses observed qualifying records. A conservative post-outcome quality restriction makes historical comparisons touching either type of coverage issue non-comparable, preventing these source issues from creating resolution events. This restriction and its chronology are disclosed separately from internally frozen scientific definitions. Supplementing the two baseline RCV-only assertions leaves primary eligibility unchanged: neither affected variant contributes an eligible primary unit. In the focused source-completeness sensitivity, VariationID 978268 gains one eligible variant-wide T4 unit; other evaluated definitions are unchanged. This supplemental effect is separate from the primary estimates.

The unresolved declared-versus-present counter restriction makes seven pairwise transition records non-comparable: three from baseline to March, one from March to September and three from baseline to September. They represent three condition-inheritance identities at two variants, all observed at T0 or T4 rather than cross-band disagreement. All seven carry the unresolved-source-count status; none is created by the separately documented paired-source assertion-absence status. The observed eligible unit tiers remain in the cross-sectional results. These transition counts describe the conservative quality rule and do not establish the underlying cause of the counter differences.


Observed official aggregate classifications are preserved. The limited aggregation approximation fails for some legacy or compound representations and is undecidable when equally ranked expert assertions have different condition scopes. The observed official conflict flag and a derived between-organization condition-inheritance tier address different scopes. Exact reproduction of ClinVar's official aggregation algorithm has not been established. In September, the limited approximation has 29 non-null disagreements and 5,961 undecidable VCV/RCV aggregate-record comparisons. Observed official labels are retained; these results do not validate replication of ClinVar's aggregation rules.

Disagreement categories




Highest-severity category per eligible variant-RCV-inheritance unit; denominator 179,613. T0 is no derived disagreement, not necessarily identical source labels. P/LP = pathogenic/likely pathogenic; U = uncertain significance; B/LB = benign/likely benign.

Gene-level observed and expected disagreement




Each point is a gene with at least 30 eligible single-gene variants and complete reference-stratum coverage. Expected counts use leave-one-gene-out backgrounds stratified by organization count, variant type, expert involvement and evaluation age, with prespecified pooling. Dashed line: observed equals expected. Descriptive screen; no significance or clinical-risk claim. Both axes use symmetric logarithmic scaling with a linear region for counts from zero to one.

Evaluation ages




Empirical distributions of the oldest and newest nonmissing relevant evaluation dates per eligible unit. These are evaluation ages, not durations of conflict. Units with missing dates are retained in the study and their counts are reported separately; curves omit missing values.

Observed release-to-release states




Baseline cross-band units followed across three snapshots. Missing variants, ineligibility and non-comparable identities remain separate from observed resolution. Changes between snapshots are interval-observed; this plot cannot reveal transient resolution or recurrence between releases. Bars show observed state at each snapshot for the same 29,626 baseline cross-band units, rather than cumulative first resolutions. The intermediate and final eligible T0/T4 totals are 452 and 712; 719 first resolutions are observed across follow-up, with seven units returning to cross-band disagreement by the final snapshot.

Official review-status strata




T1/T2 units divided by all eligible units in each official RCV review-status category; numerators and denominators are shown. Bars are 95% percentile intervals from 2,000 resamples clustered by VariationID, with seed 20260926; they account for repeated units within the same variant. Finite-record counts are exact; uncertainty is model-based. Review-status association is not causal.

Fractional participation




Entities are ranked by disputed fractional participation. A variant is divided across its genes or disputed condition sets; each disputed unit is divided across qualifying organizations. Domains use different units and ascertainment, so curves are descriptive and do not establish responsibility for disagreement.

Complete condition-set distribution



Fifteen complete condition sets with the most eligible variants, selected by denominator rather than conflict count. Each variant is counted once within a condition set, with T1/T2 in any eligible inheritance-specific unit. Variants may contribute to multiple condition sets. Names are source labels, not clinical reinterpretations.

Assertion disposition



Mutually exclusive first-applicable assertion dispositions, before organization deduplication and the requirement for at least two qualifying organizations per unit. Additional exclusion flags may overlap and are reported separately. This panel does not equate assertion eligibility with final unit eligibility.

T1 source examples

The examples below are the first two verified cases in the prespecified sampled order. They illustrate the source grouping and label rule; they do not reassess clinical pathogenicity.

Variation 2676534 / RCV003476465

Condition set: Bardet-Biedl syndrome 6 [C1858054]. Inheritance set: unspecified.

Qualifying source assertions: 1006: Likely pathogenic (SCV004194880.2); 508298: Likely benign (SCV005061820.1).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

Variation 38930 / RCV000032181

Condition set: Koolen-de Vries syndrome [C1864871]. Inheritance set: unspecified.

Qualifying source assertions: 500035: Pathogenic (SCV002516587.1); 507782: Benign (SCV005061832.2).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

T2 source examples

The examples below are the first two verified cases in the prespecified sampled order. They illustrate the source grouping and label rule; they do not reassess clinical pathogenicity.

Variation 1408134 / RCV001909813

Condition set: Tay-Sachs disease [C0039373]. Inheritance set: unspecified.

Qualifying source assertions: 500031: Uncertain significance (SCV002178604.2); 507830: Likely pathogenic (SCV007609608.1).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

Variation 1807424 / RCV005645397

Condition set: Hypophosphatasia [C0020630]. Inheritance set: unspecified.

Qualifying source assertions: 508293: Uncertain significance (SCV006587149.1); 508671: Likely pathogenic (SCV006334926.1).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

T3 source examples

The examples below are the first two verified cases in the prespecified sampled order. They illustrate the source grouping and label rule; they do not reassess clinical pathogenicity.

Variation 418249 / RCV001441245

Condition set: Holocarboxylase synthetase deficiency [C0268581]. Inheritance set: unspecified.

Qualifying source assertions: 167595: Uncertain significance (SCV003813234.3); 500031: Likely benign (SCV001644169.5).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

Variation 231987 / RCV000215163

Condition set: Hereditary cancer-predisposing syndrome [C0027672]. Inheritance set: unspecified.

Qualifying source assertions: 505849: Uncertain significance (SCV000904710.5); 61756: Likely benign (SCV000275997.7).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

T4 source examples

The examples below are the first two verified cases in the prespecified sampled order. They illustrate the source grouping and label rule; they do not reassess clinical pathogenicity.

Variation 251050 / RCV000237719

Condition set: Hypercholesterolemia, familial, 1 [C0745103]. Inheritance set: Autosomal dominant inheritance.

Qualifying source assertions: 500536: Pathogenic (SCV000294502.2); 505909: Likely pathogenic (SCV000322878.1).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

Variation 1076952 / RCV005867012

Condition set: X-linked Alport syndrome [C4746986]. Inheritance set: unspecified.

Qualifying source assertions: 500104: Pathogenic (SCV007097327.2); 507830: Likely pathogenic (SCV006581776.1).

The retained organizations satisfy the frozen review and grouping rules. Raw VCV fields and actual raw RCV membership were checked separately. Missing evaluation dates remain missing.

Limits of interpretation

ClinVar is a submitted-record archive with uneven coverage and dependence among organizations, variants and condition sets. Eligibility restricts the denominator to comparable multi-organization units. Disagreement may reflect evolving evidence, inheritance, disease context, penetrance, submitter scope or review processes. Record disagreement is not synonymous with a biological error.

The analysis does not access all underlying evidence or determine the correctness of any classification. Missing dates and unsupported date representations constrain age analyses. Source omissions, aggregate-rule limitations and differences between archived representations require explicit disposition. Descriptive adjustment cannot remove ascertainment bias.

Three snapshots provide sparse follow-up. Archive revision, changing eligibility and losses limit historical inference. Fitted survival bounds do not identify true conflict onset or an exact resolution date. Computational replay establishes reproducibility of calculations under the frozen inputs and code; it does not establish clinical validity or independent scientific endorsement.

No independent human, expert, clinical or peer review has been completed. Computational agreement and replay do not establish correctness of a biological classification.

Reproducibility

The companion package identifies the preserved sources by retrieval URL, byte size and SHA-256, includes portable protocol distribution editions with original-freeze provenance, and provides the analysis code, dependency lock and scientific tables. A fresh source-to-export replay is evaluated separately from figure and paper rendering. Refer to the package validation receipt for its exact stage scope and results.

The atlas specification describes an interactive research interface. No deployed atlas or clinical decision tool is claimed.

References

ClinVar download and identifier guide

ClinVar classification representation

ClinVar review status

Landrum et al. ClinVar: updates to support classifications of both germline and somatic variants. *Nucleic Acids Research* 2025;53(D1):D1313-D1321. doi:10.1093/nar/gkae1090

NCBI data usage policy

VCF version 4.1 specification, section 1.4.1. HTS specifications; printing c101c79, 5 November 2025.