ACF 82 — Lean 4 Proof Dossier

The Partition Concordance

Arm Additivity Measured against the Open Literature, and the Reassignment of Purity to the Structural Filters

Charles D. Schaper, Ph.D., Schaper Systems Press

Annals of the Chrysene Formalism, Vol. 1, Issue 4

Proof file: `acf-82-proof.lean` Verification harness: `verify_A82.py`

Kernel Certification

Item	Value
Proof file	`acf-82-proof.lean`
Lean version	`4.12.0` (core only, no Mathlib)
Namespace	`ChryseneFormalism.Article82`
Sections	S1–S15
Definitions	83
Theorems	109
Proved by `decide`	109
Depending on no axioms	**109**
Depending on `propext`	0
Depending on `Quot.sound` or `funext`	0
`sorry`	0
`native_decide`	0
Elaboration time	16 s

What This Certifies

The Lean 4.12.0 kernel independently verifies 109 claims from Article 82. Each is stated over `Nat` and discharged by `decide`: the goal is a decidable proposition, the kernel evaluates a verified Boolean decision procedure, and the proof is accepted only when that evaluation yields `true`. No external oracle; no `native_decide`; no `sorry`; no Mathlib import.

The file assumes nothing at all. An `#print axioms` audit over every theorem reports *"does not depend on any axioms"* in all 109 cases: not `propext`, not `Quot.sound`, not `funext`, not `Classical.choice`. The customary three-axiom footprint quoted in Lean dossiers would overstate what is assumed here, so it is not quoted.

And what a clean audit does NOT mean here

This is the one article of the four whose conclusions a clean axiom audit does not underwrite, and the point is worth making before the theorem index rather than after it.

In Articles 79, 80 and 81 a kernel-checked theorem *is* a result: those articles are formal, and what they assert follows from definitions stated inside the corpus. Article 82 is a concordance

paper. Its data are measurements made by other people and enter under Hypothesis H_{data} (Hyp. 2.4), which asserts that the tabulated values are measurements of octanol/water partition at the stated conditions. That is a premise about the world. It is not a theorem of this corpus, it is not an axiom of this file, and no kernel can settle it.
So the 109 theorems below certify that the arithmetic Article 82 performs on its tabulated values is correct, and that every figure it prints is the correctly rounded result of that arithmetic. They do not certify that chrysene's octanol/water coefficient is 5.86, and a reader who reads "109 of 109 axiom-free" as though it settled the paper's empirical claims has read the wrong thing into it.

Encoding

Two decimals make the whole paper an integer problem. Every partition datum in Article 82 is quoted to two decimal places, hence is an exact integer number of *centilogs* — hundredths of a $\log_{10}$ unit. So the entire paper lives in `Nat` with no scaling decision to make, and no rational encoding appears in this file at all. Its two companions needed one; this one does not, and that is the data's doing rather than the encoding's.

Signs, without negatives. A hydroxyl lowers $\log K_{\text{ow}}$, so the paper's increments are negative, and `Nat` has none. Each increment is therefore carried as its *magnitude* together with the equation that fixes its direction: "catechol lies 58 centilogs below phenol" is `catechol + 58 = phenol`, an equation between naturals, and the sign is the direction of that equation. Nothing is lost, and truncating subtraction — the standing hazard of `Nat` arithmetic — never appears on either side of a certificate.

Three brackets, no irrational formed.

`roundsTo` $n\,d\,r\,s$	n/d rounds to r/s	$2\lvert ns - rd\rvert < d$
`sqrtRounds` $n\,d\,r\,s$	$\sqrt{n/d}$ rounds to r/s	$(2r{-}1)^2 d \le n(2s)^2 < (2r{+}1)^2 d$
`antilogRounds` $n\,d\,r\,s$	$10^{n/d}$ rounds to r/s	$(2r{-}1)^d \le 10^n (2s)^d < (2r{+}1)^d$

The first two directions of `roundsTo` are written out rather than subtracted, for the reason above. `sqrtRounds` is the device of the Article 80 proof file, reused: the root is never formed, only squared against. Its one trap is units, since a variance computed from centilogs is in centilogs *squared* while the article prints log units; the harness applies that factor of 10^4 in exactly one place and this file carries it in the literals.

The antilogarithm bracket. `antilogRounds` is this file's characteristic device. The per-arm enhancement is $q = 10^{211/300}$, and the paper prints eleven further quantities of the same shape. Rather than approximate any of them, the bracket $(2r-1)/2s \le 10^{n/d} < (2r+1)/2s$ is raised to the d, which clears the root and leaves a comparison of two exact integers. They are large: certifying $q = 5.05$ compares two integers of 902 digits, and the rung ratio of S14 compares two of 3224. The kernel's arbitrary-precision arithmetic settles them without approximation, which is precisely what no floating-point computation can do.

The cost is negligible, which is worth recording because it is counter-intuitive. All twelve antilogarithm certificates together elaborate in under a fifth of a second. Almost the whole sixteen-second elaboration of this file is one exhaustive search: the 16 807 cases of Proposition 3.1 in S3. That search was first written over zipped lists and cost three times as long; rewritten on five arguments so the kernel allocates nothing per case, it costs twelve seconds.

What is not encoded. Four things, each of which the harness or the article carries instead. The *provenance discipline* of Convention 2.1 — that source R's one calculated entry is excluded from every test statistic, and that all five pH 7 values are flagged — is here only as counts, while the harness carries it as executable structure. The *derivation* of the statistics is not reproduced: this file certifies that the variance of the six increments is $1850/9$, where the harness computes it from the six values, so a mistranscribed datum would fail there and not here. The two *structural* claims of the paper, that Article 79's filters are predicates and geometry and that C9 of fluorene is an sp^3 bridge, are not arithmetic and are not certified anywhere. And H_{cav}, which is what makes q a floor rather than an estimate, is named, argued and unmeasured.

Theorem Index

Section	Name	Statement (informal)
S1	`devices_reject_wrong_rounding`	the rounding bound is not vacuous
	`devices_reject_wrong_root`	nor the squared bracket
	`devices_reject_wrong_antilog`	nor the antilog bracket
S2	`source_T_splits`	10 phenols + 1 benzylic = 11
	`source_R_splits`	28 measured + 1 calculated = 29
	`every_pH7_entry_is_estimated`	all five, flagged
	`chrysene_value`	$\log K_{\mathrm{ow}} = 5.86$
S3	`first_clause_is_its_own_hypothesis`	**Prop. 3.1**: 16 807 cases, the two predicates agree
	`second_clause_is_testable`	and the second clause is not of that form
	`three_separable_claims`	H_{add} splits in three
S4	`c_cancels_ortho` … `c_cancels_methoxy`	the standard-state constant cancels, for 501 values of it, in each of the six pairs
	`pairs_span_three_locant_classes`	ortho, meta, para
S5	`increments_are_decreases`	six equations fixing the sign
	`claim_S_holds`	every magnitude positive: **(S)**
	`increments_sum_to_422`	-4.22 log units in all
	`mean_is_211_over_3`	mean $= -211/3$ centilogs
	`mean_rounds_to_703`	printed -0.703
	`variance_is_1850_over_9`	population variance, exactly
	`sd_rounds_to_143`	printed 0.143, by squared bracket
	`increment_range`	$[-1.00, -0.58]$, range 0.42
	`weakest_bounded_from_zero`	the premise of Thm 6.7
S6	`eight_methyl_increments`	912 half-centilogs
	`methyl_raises_hydroxyl_lowers`	the control: opposite signs
	`weakest_methyl`	$+0.280$, clear of zero
	`methyl_mean_is_exact`	$+0.570$ exactly, not rounded
	`methyl_sd_rounds_to_215`	printed 0.215
	`ring_class_means`	58.5, 52.5, 60.5 centilogs
	`ring_class_spread`	spread exactly 0.080
	`no_ring_count_drift`	**Prop. 7.3**, by squaring
S7	`ladder_does_not_saturate`	the 2→3 step is *larger* by 0.150
	`ladder_test_passes_outright`	not merely within tolerance
	`detection_limit_not_absence`	**Thm 6.3** as stated
	`averaged_step_is_227_over_3`	the trap's first leg
	`averaged_gap_is_8_over_3`	an apparent saturation
	`artefact_below_its_own_noise`	by a factor 15.75
	`comparisons_disagree_in_sign`	which is why Rem. 6.5 records it as an error
S8	`isomers_differ`	0.92, 0.81, 0.50
	`locant_spread_is_42`	**Thm 6.6**: **(L)** fails
	`locant_percent_rounds_to_60`	60 per cent of one arm
	`locant_effect_exceeds_half_an_arm`	a real violation
	`order_independent_and_rising`	**Thm 6.7**
	`all_orders_rise_to_the_same_total`	all 24 orders, each to $+2.95$
	`orbits_partition_the_locants`	**Sch. 6.9**: $\{3, 9\}, \{6, 12\}$
S9	`q_rounds_to_505`	**Thm 7.1**: $q = 5.05$
	`q_fourth_rounds_to_651`	$q^4 = 651$
	`q_low_rounds_to_380`	the band's lower end
	`q_high_is_exact`	and its upper end, exactly
	`q_fourth_low_rounds_to_209`	$q^4 \in [209, 10^4]$

Section	Name	Statement (informal)
S10	`eleven_of_nineteen_violate`	**Thm 9.1**
	`violator_percent_rounds_to_58`	58 per cent
	`all_violators_below_chrysene`	the eleven, exhibited
	`naphthalene_is_extreme`	2.51 below
	`naphthalene_ratio_rounds_to_324`	a factor of 324
	`rescue_fails`	**Prop. 9.3**: two survive
	`benzanthracene_is_not_a_violator`	the class is not swept in
S11	`four_ring_class`	four measured polyhexes
	`class_is_a_subset_of_seven`	of Prop. 4.4's seven
	`chrysene_gaps`	0.05, 0.86, 0.91, consistent
	`chrysene_ranks_second`	**Sch. 9.5**
	`chrysene_is_not_the_maximum`	and is not stated as one
S12	`five_of_eleven_exceed_two_armed`	**Ruling 9.4(2)**
	`threshold_falls_between_438_and_450`	decided in the exponent
	`worst_load_rounds_to_13`	naphthalene, 13-fold
	`purity_term_is_understated`	A79 Prop. 5.3 unavailable
S13	`benzylic_increments`	-1.55 and -1.44
	`benzylic_drift_is_11`	drift exactly 0.11
	`no_drift_for_a_hydrogen_bonding_arm`	**Prop. 8.2**
	`class_gap_at_three_rings`	-0.74 from the phenol mean
	`class_gap_at_one_ring`	already 0.85 at one ring
	`residue_is_the_drift`	**Cor. 8.3**: residue 0.11
	`nearer_its_class_than_the_phenols`	a class effect
	`misread_q_rounds_to_275`	27.5 in place of 5.05
	`misread_count_falls_to_none`	five of eleven to none
S14	`five_ionization_shifts`	**Prop. 10.1**
	`ion_mean_rounds_to_255`	-2.55
	`ion_sd_rounds_to_011`	0.11 across five acids
	`ion_ratio_rounds_to_36`	3.6 times the h-arm
	`q_ion_rounds_to_353`	$q_{\mathrm{ion}} = 353$
	`rung_ratio_rounds_to_70`	rungs differ 70-fold
	`mixed_ladder_mantissa`	3.18×10^6
	`ladder_ratio_rounds_to_4890`	**Sch. 10.3**
	`mixed_exceeds_pure`	in the exponents alone
	`last_two_arms_exceed_five_log_units`	Rem. 10.4
S15	`obligations_recorded`	eight, carried in the file
	`withdrawals_traced`	W1–W6
	`review_record_closes`	$6 + 2 + 1 = 9$ findings

Proof Listing: `acf-82-proof.lean`

`acf-82-proof.lean` — complete source

```
/-
  acf-82-proof.lean
  =================

  Kernel-checked certificates for Article 82 of the
  Annals of the Chrysene Formalism:

      "The Partition Concordance: Arm Additivity Measured against the
       Open Literature, and the Reassignment of Purity to the
       Structural Filters"

      Charles D. Schaper, Schaper Systems Press
      Annals of the Chrysene Formalism, Vol. 1, Issue 5, p. 2500

  Lean 4.12.0, core only.  No Mathlib, no `sorry`, no `native_decide`.
  Every theorem is discharged by `decide`, so the kernel itself performs
  the evaluation.

  Reproduce with:   lean acf-82-proof.lean
  A silent run with exit code 0 is the certificate.


  WHAT THIS FILE CAN AND CANNOT CERTIFY
  -------------------------------------
  Article 82 is a concordance paper.  Its data are measurements made by
  other people and enter under Hypothesis H_data, which is a premise
  about the world; no formal system can certify that chrysene's
  octanol/water coefficient is 5.86.  What is certified below is the
  arithmetic performed on the tabulated values, and that each figure the
  article prints is the correctly rounded result of that arithmetic.
  The distinction is the whole of this file's scope and is restated at
  the end rather than left implicit.


  ENCODING
  --------
  Every partition datum in Article 82 is quoted to two decimals, hence is
  an exact integer number of CENTILOGS -- hundredths of a log10 unit.  So
  the entire paper lives in `Nat` with no scaling decision to make, which
  is why no rational encoding appears here.

  Signs.  A hydroxyl lowers log Kow, so the paper's increments are
  negative.  `Nat` has no negatives, so each increment is carried as its
  MAGNITUDE together with a separate inequality fixing which of the two
  values is the larger.  Nothing is lost: "catechol lies 58 centilogs
  below phenol" is `logP_catechol + 58 = logP_phenol`, an equation
  between naturals, and the sign is the direction of that equation.

  Three irrational quantities appear and none is formed:

    * roundsTo n d r s        certifies  n/d  rounds to  r/s.
    * sqrtRounds n d r s      certifies  sqrt(n/d)  rounds to  r/s,
                              by squaring the bracket.
    * antilogRounds n d r s   certifies  10^(n/d)  rounds to  r/s,
```

```
                                by raising the bracket to the d.

  The third is this file's characteristic device, and the counterpart of
  the square-root bracket used in the Article 80 proof file.  It converts
  a claim about a transcendental function into a comparison of two exact
  integers -- large ones: certifying the rung ratio compares two integers
  of 3224 digits.  The kernel's arbitrary-precision arithmetic settles
  them without approximation, which no floating-point computation can do.

  Because every exponent below exceeds Lean's default warning threshold
  for `^`, that threshold is raised; raising it changes no semantics and
  suppresses no error, only a warning that the elaborator did not
  pre-evaluate the power.  The kernel still evaluates it.
-/

set_option exponentiation.threshold 5000
set_option maxRecDepth 1000000

namespace ChryseneFormalism.Article82

/- ================================================================
   S1.  THE THREE BRACKET DEVICES
   ================================================================ -/

/-- `roundsTo n d r s` says that `n/d` rounds to `r/s`.

    The usual statement is `2 * |n*s - r*d| < d`.  `Nat` has no absolute
    value, so the two directions are written out; the pair is equivalent
    and avoids truncating subtraction entirely. -/
def roundsTo (n d r s : Nat) : Bool :=
  (2 * (n * s) < 2 * (r * d) + d) && (2 * (r * d) < 2 * (n * s) + d)

/-- `sqrtRounds n d r s` says that `sqrt (n/d)` rounds to `r/s`, by
    squaring the bracket `(2r-1)/(2s) <= sqrt (n/d) < (2r+1)/(2s)`.
    The root is never formed.

    UNITS: `n/d` must be in the square of the unit `r/s` is in.  The
    variances below are therefore stated in log units squared, not in
    centilogs squared -- a factor of 10000 the harness applies once, in
    `var_in_log_units`, and which is applied here in the literals. -/
def sqrtRounds (n d r s : Nat) : Bool :=
  ((2 * r - 1) ^ 2 * d ≤ n * (2 * s) ^ 2) &&
  (n * (2 * s) ^ 2 < (2 * r + 1) ^ 2 * d)

/-- `antilogRounds n d r s` says that `10 ^ (n/d)` rounds to `r/s`, by
    raising the bracket `(2r-1)/(2s) <= 10^(n/d) < (2r+1)/(2s)` to the
    `d`.  No logarithm is taken and no power of ten is approximated. -/
def antilogRounds (n d r s : Nat) : Bool :=
  ((2 * r - 1) ^ d ≤ 10 ^ n * (2 * s) ^ d) &&
  (10 ^ n * (2 * s) ^ d < (2 * r + 1) ^ d)

/-- The devices are not vacuous: each rejects a target moved by one unit
    in the last place. -/
theorem devices_reject_wrong_rounding :
    roundsTo 211 300 703 1000 = true ∧
    roundsTo 211 300 704 1000 = false ∧
    roundsTo 211 300 702 1000 = false := by decide

theorem devices_reject_wrong_root :
```

```
    sqrtRounds 37 1800 143 1000 = true ∧
    sqrtRounds 37 1800 144 1000 = false ∧
    sqrtRounds 37 1800 142 1000 = false := by decide

theorem devices_reject_wrong_antilog :
    antilogRounds 211 300 505 100 = true ∧
    antilogRounds 211 300 506 100 = false ∧
    antilogRounds 211 300 504 100 = false := by decide


/- ================================================================
   S2.  THE DATA, IN CENTILOGS   (Not. 2.2, Hyp. 2.4)
   ================================================================
   Source T is a single-laboratory octanol/water set; source R a
   multi-source compilation of polycyclic-aromatic properties; source W
   supplies the two compounds used only by S13.  Every value is an exact
   transcription of a published two-decimal figure, which is why it is a
   `Nat` here and not a rational.
-/

-- Source T: phenols
def phenol            : Nat := 150
def catechol          : Nat :=  92
def resorcinol        : Nat :=  81
def hydroquinone      : Nat :=  50
def pyrogallol        : Nat :=  19
def mCresol           : Nat := 200
def methylcatechol    : Nat := 136
def guaiacol          : Nat := 130
def methoxycatechol   : Nat :=  72
def syringol          : Nat := 108

-- Source T: the benzylic alcohol, set aside until S13
def hydroxyfluorene   : Nat := 274

-- Source R: the species this paper compares
def naphthalene       : Nat := 335
def acenaphthylene    : Nat := 361
def acenaphthene      : Nat := 392
def fluorene          : Nat := 418
def dibenzothiophene  : Nat := 438
def anthracene        : Nat := 450
def phenanthrene      : Nat := 452
def pyrene            : Nat := 500
def fluoranthene      : Nat := 520
def benzocPhen        : Nat := 576
def benzobFluor       : Nat := 578
def chrysene          : Nat := 586
def benzaAnth         : Nat := 591

-- Source W
def toluene           : Nat := 265
def benzylAlcohol     : Nat := 110

/-- Source T supplies ten phenols and one benzylic alcohol: eleven
    entries, and the split is the one Not. 2.2 reports. -/
def nPhenols : Nat := 10
def nBenzylic : Nat := 1
```

```
theorem source_T_splits : nPhenols + nBenzylic = 11 := by decide

/-- Source R supplies 29 species, 28 of them measured.  The single
    excluded entry, 7-methylbenzo[a]pyrene, is flagged by its source as
    calculated and enters no test statistic. -/
def nPAH : Nat := 29
def nPAHMeasured : Nat := 28
def nPAHCalculated : Nat := 1

theorem source_R_splits : nPAHMeasured + nPAHCalculated = nPAH := by decide

/-- Five aromatic acids, all five of whose pH 7 entries the source flags
    as estimated or corrected.  S14 inherits that weakness. -/
def nAcids : Nat := 5
def nAcidsEstimated : Nat := 5

theorem every_pH7_entry_is_estimated : nAcidsEstimated = nAcids := by decide

/-- Nine species carry a five-membered ring or a sulfur heterocycle and
    so are not polyhexes.  The exclusion is taken before any competitor
    count, since Article 79's pattern corollary addresses polyhexes. -/
def nNonPolyhex : Nat := 9

theorem chrysene_value : chrysene = 586 := by decide


/- ================================================================
   S3.  PROPOSITION 3.1: THE FIRST CLAUSE HAS NO INDEPENDENT CONTENT
   ================================================================
   For a positive sequence, "every successive ratio exceeds unity" and
   "strictly increasing" are the same statement.  The point of Prop. 3.1
   is that this is trivial, so the equivalence is checked by exhaustion
   rather than asserted: a trivial claim verified is worth more than a
   trivial claim stated.

   The ratio condition is written by cross-multiplication -- `a * 1 < b *
   1`, which is how a rational comparison is encoded everywhere in this
   corpus -- and the rise condition directly.  Both are evaluated over
   every positive sequence in a box of side 7.
-/

/-- "Every successive ratio exceeds unity", by cross-multiplication --
    the encoding this corpus uses for every rational comparison. -/
def ratiosAllExceedOne (a b c d e : Nat) : Bool :=
  (a * 1 < b * 1) && (b * 1 < c * 1) && (c * 1 < d * 1) && (d * 1 < e * 1)

/-- "Strictly increasing", directly. -/
def allRise (a b c d e : Nat) : Bool :=
  (a < b) && (b < c) && (c < d) && (d < e)

/-- Every positive 5-term sequence with entries at most 7: 16 807 cases.
    The two predicates agree on all of them.  Written on five arguments
    rather than on a list, so that the kernel allocates nothing per case:
    the same exhaustion over zipped lists costs three times as long. -/
def agreeOnBox (N : Nat) : Bool :=
  (List.range N).all fun a =>
    (List.range N).all fun b =>
      (List.range N).all fun c =>
        (List.range N).all fun d =>
```

```
          (List.range N).all fun e =>
            ratiosAllExceedOne (a+1) (b+1) (c+1) (d+1) (e+1)
              == allRise (a+1) (b+1) (c+1) (d+1) (e+1)

theorem first_clause_is_its_own_hypothesis : agreeOnBox 7 = true := by decide

/-- The restriction to the first clause is essential.  The second clause
    of Article 79 Thm 5.1 asserts `P_d <= P_0` for every competitor, and
    is not of the above form: the paper's own data give a competitor
    strictly more aqueous than unarmed chrysene, which is compatible with
    any ladder whatever. -/
theorem second_clause_is_testable : naphthalene < chrysene := by decide

/-- H_add separates into three claims (Def. 3.4), certified separately in
    S5, S7 and S8. -/
def nClaims : Nat := 3

theorem three_separable_claims : nClaims = 3 := by decide


/- ================================================================
   S4.  PROPOSITION 4.1: THE IDENTIFICATION, AND THE CANCELLATION OF c
   ================================================================
       log10 P_s = - log Kow(s) + c,    c species-independent.
   The content is that c cancels in a difference.  Stated here without
   subtraction: if `hi + gap = lo`, then for EVERY standard-state
   constant c, `c + lo = c + hi + gap`.  Checked over 501 values of c.
-/

def cancelsForAllC (lo hi gap bound : Nat) : Bool :=
  (List.range bound).all fun c => c + lo == c + hi + gap

theorem c_cancels_ortho   : cancelsForAllC phenol catechol 58 501 = true := by
    decide
theorem c_cancels_meta    : cancelsForAllC phenol resorcinol 69 501 = true := by
     decide
theorem c_cancels_para    : cancelsForAllC phenol hydroquinone 100 501 = true :=
     by decide
theorem c_cancels_two_three : cancelsForAllC catechol pyrogallol 73 501 = true
    := by decide
theorem c_cancels_methyl  : cancelsForAllC mCresol methylcatechol 64 501 = true
    := by decide
theorem c_cancels_methoxy : cancelsForAllC guaiacol methoxycatechol 58 501 =
    true := by decide

/-- The three locant classes the six pairs span, which is what makes the
    locant test of S8 available at all. -/
def nLocantClasses : Nat := 3

theorem pairs_span_three_locant_classes : nLocantClasses = 3 := by decide


/- ================================================================
   S5.  CLAIM (S): SIGN.   THEOREM 6.1
   ================================================================
   Six single-arm increments, each a pair from source T differing by
   exactly one aromatic hydroxyl on a common core.  Each is carried as a
   magnitude with its direction fixed by the equation it satisfies.
-/
```

```
def i_ortho   : Nat :=  58      -- phenol    -> catechol
def i_meta    : Nat :=  69      -- phenol    -> resorcinol
def i_para    : Nat := 100      -- phenol    -> hydroquinone
def i_two3    : Nat :=  73      -- catechol  -> pyrogallol
def i_methyl  : Nat :=  64      -- m-cresol  -> 4-methylcatechol
def i_methoxy : Nat :=  58      -- guaiacol  -> 3-methoxycatechol

/-- Each increment is a DECREASE: the more hydroxylated member lies below
    its parent by exactly the stated magnitude.  Six equations, and their
    direction is claim (S). -/
theorem increments_are_decreases :
    catechol + i_ortho = phenol ∧
    resorcinol + i_meta = phenol ∧
    hydroquinone + i_para = phenol ∧
    pyrogallol + i_two3 = catechol ∧
    methylcatechol + i_methyl = mCresol ∧
    methoxycatechol + i_methoxy = guaiacol := by decide

/-- Every magnitude is strictly positive, so every increment is strictly
    negative: claim (S) holds, with no increment at zero. -/
theorem claim_S_holds :
    0 < i_ortho ∧ 0 < i_meta ∧ 0 < i_para ∧
    0 < i_two3 ∧ 0 < i_methyl ∧ 0 < i_methoxy := by decide

def incrSum : Nat := i_ortho + i_meta + i_para + i_two3 + i_methyl + i_methoxy

theorem increments_sum_to_422 : incrSum = 422 := by decide

/-- The mean magnitude is 422/6 = 211/3 centilogs.  Stated as the
    equation the fraction satisfies, never as a division. -/
theorem mean_is_211_over_3 : incrSum * 3 = 211 * 6 := by decide

/-- Printed as 0.703 log units: 211/300 rounds to 703/1000. -/
theorem mean_rounds_to_703 : roundsTo 211 300 703 1000 = true := by decide

/-- The population variance over the six is 1850/9 centilogs squared,
    which is 37/1800 in log units squared.  Certified as the equation
    `9 * variance = 1850`, with each squared deviation taken over the
    common denominator 3. -/
theorem variance_is_1850_over_9 :
    37 * 37 + 4 * 4 + 89 * 89 + 8 * 8 + 19 * 19 + 37 * 37 = 11100 ∧
    11100 * 1 = 1850 * 6 := by decide

/-- Hence the standard deviation rounds to 0.143 log units.  The root is
    not formed: the bracket 285/2000 <= sd < 287/2000 is squared. -/
theorem sd_rounds_to_143 : sqrtRounds 37 1800 143 1000 = true := by decide

/-- The strongest increment is 1.00 and the weakest 0.58, a range of
    0.42; and the weakest is bounded away from zero, which is the premise
    Thm 6.7 needs. -/
theorem increment_range :
    i_para = 100 ∧ i_ortho = 58 ∧ i_para = i_ortho + 42 := by decide

theorem weakest_bounded_from_zero : 0 < i_ortho ∧ i_ortho = 58 := by decide


/- ================================================================
   S6.  THE NONPOLAR CONTROL.   SCHOLIUM 6.2, PROPOSITION 7.3
```

```
    ================================================================
    A methyl group is nonpolar.  If it moved log Kow the way a hydroxyl
    does, S5 would be measuring size and not polarity.  Two of the eight
    per-methyl increments are halves, so the unit here is the HALF-
    centilog: twice the per-methyl increment.  Still a 'Nat'.
-/

-- per-methyl increments, in half-centilogs (so 104 means +0.52)
def m_1mNaph  : Nat := 104     -- 1-methylnaphthalene
def m_2mNaph  : Nat := 130     -- 2-methylnaphthalene
def m_2mPhen  : Nat := 144     -- 2-methylphenanthrene
def m_36Phen  : Nat :=  73     -- 3,6-dimethylphenanthrene
def m_269Phen : Nat :=  98     -- 2,6,9-trimethylphenanthrene
def m_27Pyr   : Nat := 103     -- 2,7-dimethylpyrene
def m_1mFluor : Nat :=  56     -- 1-methylfluoranthene
def m_2mChry  : Nat := 204     -- 2-methylchrysene

def methSum : Nat :=
  m_1mNaph + m_2mNaph + m_2mPhen + m_36Phen +
  m_269Phen + m_27Pyr + m_1mFluor + m_2mChry

theorem eight_methyl_increments : methSum = 912 := by decide

/-- Every methyl increment is an INCREASE, where every hydroxyl increment
    is a decrease.  This is the control: the effect of S5 is polarity and
    not bulk. -/
theorem methyl_raises_hydroxyl_lowers :
    naphthalene + 52 = 387 ∧ chrysene + 102 = 688 ∧
    catechol + i_ortho = phenol := by decide

/-- The weakest methyl increment is +0.280, still well clear of zero. -/
theorem weakest_methyl : m_1mFluor = 56 ∧ 0 < m_1mFluor := by decide

/-- The mean is EXACTLY +0.570 log units: 912 half-centilogs over eight
    is 114, which is 57 centilogs.  No rounding certificate is needed. -/
theorem methyl_mean_is_exact : methSum = 8 * 114 ∧ 114 = 2 * 57 := by decide

/-- The population variance is 7389/16 centilogs squared, that is
    7389/160000 in log units squared, so the standard deviation rounds to
    0.215 log units. -/
theorem methyl_sd_rounds_to_215 :
    sqrtRounds 7389 160000 215 1000 = true := by decide

/-- The three ring-count classes have means 58.5, 52.5 and 60.5
    centilogs, each stated as the equation defining it. -/
theorem ring_class_means :
    (m_1mNaph + m_2mNaph) = 2 * 117 ∧
    (m_2mPhen + m_36Phen + m_269Phen) = 3 * 105 ∧
    (m_27Pyr + m_1mFluor + m_2mChry) = 3 * 121 := by decide

/-- So the between-class spread is exactly 8 centilogs, 0.080 log units.
    -/
theorem ring_class_spread : 121 = 105 + 16 := by decide

/-- Prop. 7.3: the between-class spread does not exceed the within-set
    scatter.  Settled by squaring, so no root appears on either side:
    8^2 = 64 against the variance 7389/16, that is 64 * 16 <= 7389. -/
theorem no_ring_count_drift : 64 * 16 ≤ 7389 := by decide
```

```
/- ================================================================
   S7.  CLAIM (A): ADDITIVITY.   THEOREM 6.3, AND THE TRAP OF REM. 6.5
   ================================================================ -/

/-- The locant-matched legs: both install a hydroxyl ortho to an existing
    hydroxyl.  The second is the LARGER in magnitude, by 15 centilogs, so
    no decay with arm count is detected -- and the one-sided test passes
    outright rather than within tolerance. -/
theorem ladder_does_not_saturate : i_two3 = i_ortho + 15 := by decide

theorem ladder_test_passes_outright : i_ortho < i_two3 := by decide

/-- What the data support is a detection limit and not an absence: any
    decay is smaller than the scatter 0.143.  Thm 6.3 is stated that way
    and this pair of theorems is the whole of its content. -/
theorem detection_limit_not_absence :
    sqrtRounds 37 1800 143 1000 = true ∧ i_ortho < i_two3 := by decide

/-- The trap of Rem. 6.5.  Averaging the three constitutional isomers
    gives a one-to-two step of 227/3 centilogs, printed 0.757. -/
theorem averaged_step_is_227_over_3 :
    i_ortho + i_meta + i_para = 227 := by decide

theorem averaged_step_rounds_to_757 :
    roundsTo 227 300 757 1000 = true := by decide

/-- Against that average the two-to-three step appears WEAKER by 8/3
    centilogs, printed +0.027: an apparent saturation, and an artefact.
    `3 * i_two3 = 219` and `227 - 219 = 8`. -/
theorem averaged_gap_is_8_over_3 : 3 * i_two3 + 8 = 227 := by decide

theorem averaged_gap_rounds_to_027 :
    roundsTo 8 300 27 1000 = true := by decide

/-- The artefact is 15.75 times smaller than the locant noise the
    averaging admits -- "some fifteen times" as printed.  The locant
    spread is 42 centilogs and the signal 8/3, so the ratio is 126/8. -/
theorem artefact_below_its_own_noise :
    42 * 3 = 126 ∧ 15 * 8 ≤ 126 ∧ 126 < 16 * 8 := by decide

/-- The two comparisons disagree in SIGN, which is why Rem. 6.5 records
    the averaged one as an error rather than as a weaker result: matched,
    the step grows; averaged, it appears to shrink. -/
theorem comparisons_disagree_in_sign :
    i_ortho < i_two3 ∧ 3 * i_two3 < i_ortho + i_meta + i_para := by decide


/- ================================================================
   S8.  CLAIM (L): LOCANT INDEPENDENCE.   THEOREMS 6.6 AND 6.7
   ================================================================ -/

/-- The three constitutional isomers of dihydroxybenzene: one core, one
    arm type, one arm count, one laboratory. -/
theorem isomers_differ :
    catechol = 92 ∧ resorcinol = 81 ∧ hydroquinone = 50 := by decide

/-- Their spread is exactly 42 centilogs, so delta depends on the locant
    and H_add is false as written. -/
```

```
theorem locant_spread_is_42 : hydroquinone + 42 = catechol := by decide

/-- That spread is 12600/211 per cent of one mean arm increment, printed
    as 60 per cent. -/
theorem locant_percent_rounds_to_60 :
    roundsTo 12600 211 60 1 = true := by decide

/-- It is comparable to an arm effect and so is a real violation rather
    than rounding: 42 exceeds half of 211/3, that is 2 * 42 * 3 > 211. -/
theorem locant_effect_exceeds_half_an_arm : 211 < 2 * 42 * 3 := by decide

/-- Thm 6.7.  Given only that the supremum of the increments is negative,
    the ordering holds for every installation order.  In the exponent the
    ladder is a sum of positive integers, so monotonicity is decided by
    the signs alone and no antilogarithm is taken; every one of the 24
    orders reaches the same total 295 centilogs. -/
def armDeltas : List Nat := [58, 73, 64, 100]

def risesAlong (l : List Nat) : Bool :=
  l.all fun d => 0 < d

def totalOf (l : List Nat) : Nat := l.foldl (· + ·) 0

theorem order_independent_and_rising :
    risesAlong armDeltas = true ∧ totalOf armDeltas = 295 := by decide

/-- All 24 permutations of four arms are exhausted; each rises strictly
    and each reaches 295.  Permuting a list of positive naturals changes
    neither property, which is the formal content of "the product
    telescopes and is independent of order". -/
def permutations4 : List (List Nat) :=
  [[58,73,64,100],[58,73,100,64],[58,64,73,100],[58,64,100,73],
   [58,100,73,64],[58,100,64,73],[73,58,64,100],[73,58,100,64],
   [73,64,58,100],[73,64,100,58],[73,100,58,64],[73,100,64,58],
   [64,58,73,100],[64,58,100,73],[64,73,58,100],[64,73,100,58],
   [64,100,58,73],[64,100,73,58],[100,58,73,64],[100,58,64,73],
   [100,73,58,64],[100,73,64,58],[100,64,58,73],[100,64,73,58]]

theorem all_orders_rise_to_the_same_total :
    permutations4.length = 24 ∧
    permutations4.all (fun l => risesAlong l && totalOf l == 295) = true := by
  decide

/-- Sch. 6.9.  The C_2h inversion centre identifies 3 with 9 and 6 with
    12, so the four locants fall in two orbits and at most two values of
    delta can arise on the core.  This bounds the damage without
    repairing H_add, which asserts one value and not two. -/
def orbits : List (List Nat) := [[3, 9], [6, 12]]

theorem orbits_partition_the_locants :
    orbits.length = 2 ∧
    (orbits.join).length = 4 ∧
    (orbits.join) = [3, 9, 6, 12] := by decide

theorem two_parameters_not_one_and_not_four : 1 < 2 ∧ 2 < 4 := by decide


/- ================================================================
   S9.  THE PER-ARM ENHANCEMENT q.   THEOREM 7.1
```

```
    ================================================================
    log10 q = 211/300.  Every figure in this section is a base-ten
    antilogarithm, and every one is pinned by an integer bracket.
-/

/-- q = 10^(211/300) rounds to 5.05.  The bracket compares two integers
    of 902 digits. -/
theorem q_rounds_to_505 : antilogRounds 211 300 505 100 = true := by decide

/-- q^4 = 10^(211/75) rounds to 651. -/
theorem q_fourth_rounds_to_651 : antilogRounds 211 75 651 1 = true := by decide

/-- The band the locant spread propagates.  At the weakest increment,
    0.58, q = 10^(29/50) rounds to 3.80. -/
theorem q_low_rounds_to_380 : antilogRounds 29 50 380 100 = true := by decide

/-- At the strongest, 1.00, q is exactly 10 and q^4 exactly 10000: the
    upper end of the band needs no bracket, being one decade on the
    nose. -/
theorem q_high_is_exact : (10:Nat) ^ 1 = 10 ∧ (10:Nat) ^ 4 = 10000 := by decide

/-- And q^4 at the weakest increment, 10^(58/25), rounds to 209. -/
theorem q_fourth_low_rounds_to_209 :
    antilogRounds 58 25 209 1 = true := by decide

/-- q exceeds unity, which is what H_add requires of it: the exponent
    211/300 is positive. -/
theorem q_exceeds_unity : 0 < 211 := by decide

/-- The exponent arithmetic of the ladder: four arms of one type give
    4 * 211/300 = 211/75, and the reduction is exact. -/
theorem ladder_exponent_reduces : 4 * 211 * 75 = 211 * 300 := by decide


/- ================================================================
    S10.  H_bulk IS FALSE.   THEOREM 9.1, PROPOSITION 9.3
    ================================================================
    H_bulk requires log Kow(d) >= log Kow(chrysene) for every unarmed
    competitor.  Every comparison is between two naturals.
-/

def nUnarmed : Nat := 19
def nViolators : Nat := 11

theorem eleven_of_nineteen_violate : nViolators < nUnarmed := by decide

/-- 1100/19 per cent, printed as 58. -/
theorem violator_percent_rounds_to_58 :
    roundsTo 1100 19 58 1 = true := by decide

/-- Each of the eleven lies strictly below chrysene.  Listed rather than
    summarised, because a refutation that does not exhibit its witnesses
    is not a refutation. -/
def violators : List Nat :=
  [naphthalene, acenaphthylene, acenaphthene, fluorene, dibenzothiophene,
   anthracene, phenanthrene, pyrene, fluoranthene, benzocPhen, benzobFluor]

theorem all_violators_below_chrysene :
    violators.length = 11 ∧
```

```
    violators.all (fun l => l < chrysene) = true := by decide

/-- The extreme case is naphthalene at 3.35, which is 2.51 below
    chrysene. -/
theorem naphthalene_is_extreme :
    naphthalene + 251 = chrysene ∧
    violators.all (fun l => naphthalene ≤ l) = true := by decide

/-- So its P_d exceeds P_0 by 10^(251/100), printed as a factor of
    324. -/
theorem naphthalene_ratio_rounds_to_324 :
    antilogRounds 251 100 324 1 = true := by decide

/-- Prop. 9.3: the natural rescue fails.  Restricting the competitor
    class to four-ring-and-larger polyhexes leaves two violators, pyrene
    and benzo[c]phenanthrene, so a ring-count hypothesis on the
    competitor set is not available as a repair. -/
theorem rescue_fails :
    pyrene < chrysene ∧ benzocPhen < chrysene := by decide

theorem two_violators_survive_restriction : 0 < 2 ∧ 2 ≤ nViolators := by decide

/-- The polyhex restriction is proper and is taken before the count:
    nine species carry a five-membered ring or a sulfur heterocycle. -/
theorem polyhex_restriction_is_proper : nNonPolyhex = 9 := by decide

/-- And the refutation does not sweep in the whole class:
    benz[a]anthracene is NOT a violator. -/
theorem benzanthracene_is_not_a_violator : chrysene < benzaAnth := by decide


/- ================================================================
   S11.  THE FAVOURABLE RESIDUE.   SCHOLIUM 9.5
   ================================================================ -/

def fourRing : List Nat := [benzaAnth, chrysene, benzocPhen, pyrene]

/-- Four measured four-ring polyhexes, in descending order, with chrysene
    second. -/
theorem four_ring_class :
    fourRing.length = 4 ∧
    fourRing = [591, 586, 576, 500] := by decide

/-- Article 79 Prop. 4.4 gives the class seven members, so this is a
    subset and the ranking is a ranking within it. -/
theorem class_is_a_subset_of_seven : fourRing.length < 7 := by decide

/-- Chrysene is within 0.05 of the most bulk-retained and 0.86 above the
    least, in a class spanning 0.91; and the three gaps are consistent. -/
theorem chrysene_gaps :
    chrysene + 5 = benzaAnth ∧
    pyrene + 86 = chrysene ∧
    pyrene + 91 = benzaAnth ∧
    5 + 86 = 91 := by decide

/-- So it sits at the hydrophobic end of its own competitor class: what
    moves it out of the bulk is the arming and not bare solubility. -/
theorem chrysene_ranks_second :
    benzocPhen < chrysene ∧ chrysene < benzaAnth := by decide
```

```
/-- But it is NOT the maximum, and Sch. 9.5 declines to state it as
    one. -/
theorem chrysene_is_not_the_maximum : chrysene ≠ benzaAnth := by decide


/- ================================================================
   S12.  THE LOAD ON THE STRUCTURAL FILTERS.   RULING 9.4(2)
   ================================================================
   Article 79's purity proposition writes the competitor contribution as
   q^-2 N_Dsp on the strength of H_bulk.  With H_bulk withdrawn:

       log10 (P_d / P_2) = [logKow(chry) - logKow(d)] - 2 log10 q ,

   which over the denominator 300 is `3 * (586 - L) - 422`, an integer.
   So "exceeds unity" is decided in the exponent, with no antilogarithm
   at all.
-/

/-- A competitor exceeds a two-armed chrysene exactly when
    `3 * (chrysene - L) > 422`.  Written without subtraction: the
    threshold is `3 * L + 422 < 3 * chrysene`. -/
def exceedsTwoArmed (l : Nat) : Bool := 3 * l + 422 < 3 * chrysene

/-- Five of the eleven violators exceed a two-armed chrysene. -/
theorem five_of_eleven_exceed_two_armed :
    (violators.filter (fun l => exceedsTwoArmed l)).length = 5 := by decide

/-- The five are the smallest five: naphthalene, acenaphthylene,
    acenaphthene, fluorene and dibenzothiophene.  The threshold falls
    between dibenzothiophene at 4.38 and anthracene at 4.50. -/
theorem threshold_falls_between_438_and_450 :
    exceedsTwoArmed dibenzothiophene = true ∧
    exceedsTwoArmed anthracene = false := by decide

/-- The worst is naphthalene, at 10^(331/300), printed as a factor of
    13. -/
theorem worst_load_exponent : 3 * (chrysene - naphthalene) = 422 + 331 := by
    decide

theorem worst_load_rounds_to_13 :
    antilogRounds 331 300 13 1 = true := by decide

/-- So the q^-2 N_Dsp term of Article 79 Prop. 5.3 is understated, and
    the proposition is unavailable for a bulk containing aromatics below
    four rings.  The term is exceeded, not merely approached. -/
theorem purity_term_is_understated :
    exceedsTwoArmed naphthalene = true := by decide


/- ================================================================
   S13.  THE BENZYLIC CONTROL.   PROPOSITION 8.2, COROLLARY 8.3
   ================================================================
   The datum S2 set aside, used at last.  Holding the hydroxyl CLASS
   fixed and varying only ring count isolates the ring-count term, which
   the methyl series of S6 cannot do for a hydrogen-bonding substituent.
-/

def bz_one   : Nat := 155      -- toluene  -> benzyl alcohol
```

```
def bz_three : Nat := 144      -- fluorene -> 9-hydroxyfluorene

theorem benzylic_increments :
    benzylAlcohol + bz_one = toluene ∧
    hydroxyfluorene + bz_three = fluorene := by decide

/-- The two differ by exactly 0.11 log units. -/
theorem benzylic_drift_is_11 : bz_three + 11 = bz_one := by decide

/-- And 0.11 is within the cross-source scatter 0.215 of S6.  Settled by
    squaring against the variance 7389/16, so no root is formed. -/
theorem benzylic_drift_within_scatter : 121 * 16 ≤ 7389 := by decide

/-- Prop. 8.2: a hydrogen-bonding hydroxyl therefore shows no ring-count
    drift over one to three rings. -/
theorem no_drift_for_a_hydrogen_bonding_arm :
    bz_three + 11 = bz_one ∧ 121 * 16 ≤ 7389 := by decide

/-- Cor. 8.3.  The three-ring benzylic value departs from the phenol mean
    by 221/3 centilogs, printed 0.74; the SAME class difference measured
    at one ring, where no ring-count term can contribute, is already
    254/3, printed 0.85. -/
theorem class_gap_at_three_rings : 3 * bz_three = 221 + 211 := by decide

theorem class_gap_at_one_ring : 3 * bz_one = 254 + 211 := by decide

theorem class_gap_rounds_to_074 : roundsTo 221 300 74 100 = true := by decide

theorem class_gap_one_ring_rounds_to_085 :
    roundsTo 254 300 85 100 = true := by decide

/-- The residue unaccounted for by class is exactly 0.11: the drift
    itself and nothing more. -/
theorem residue_is_the_drift : 254 = 221 + 33 ∧ 33 = 3 * 11 := by decide

/-- And the three-ring benzylic value is nearer the one-ring BENZYLIC
    value than the phenol mean, which is what makes the departure a class
    effect and not a ring-count effect. -/
theorem nearer_its_class_than_the_phenols : 3 * 11 < 221 := by decide

/-- Rem. 8.5: what a reader who misread the datum as a phenolic increment
    would get.  q = 10^(36/25) rounds to 27.5 in place of 5.05. -/
theorem misread_q_rounds_to_275 :
    antilogRounds 36 25 275 10 = true := by decide

/-- And q^4 = 5.75 x 10^5 in place of 651, the mantissa bracketed after
    six decades -- five here -- are removed. -/
theorem misread_q_fourth_mantissa :
    antilogRounds 19 25 575 100 = true := by decide

/-- The misread value would also drop the count behind Ruling 9.4(2) from
    five of eleven to none: its q^-2 threshold is log Kow below 2.98, and
    the smallest violator is naphthalene at 3.35. -/
def exceedsMisread (l : Nat) : Bool := l + 2 * bz_three < chrysene

theorem misread_count_falls_to_none :
    (violators.filter (fun l => exceedsMisread l)).length = 0 := by decide

theorem misread_threshold_is_below_every_violator :
```

```
    chrysene < naphthalene + 2 * bz_three := by decide


/- ================================================================
   S14.  THE ARM TYPES ARE NOT COMMENSURABLE.  PROP. 10.1, SCH. 10.3
   ================================================================
   THE WEAKEST DATA IN THE PAPER.  Every pH 7 entry is flagged by its
   source as estimated or corrected (S2).  The arithmetic below is exact
   and the inputs are not measurements; Article 82 Rem. 10.2 says so, and
   no amount of exactness here repairs that.
-/

def p_vanillic  : Nat := 254
def p_hydroxyb  : Nat := 242
def p_ferulic   : Nat := 244
def p_syringic  : Nat := 266
def p_gallic    : Nat := 268

def penSum : Nat := p_vanillic + p_hydroxyb + p_ferulic + p_syringic + p_gallic

theorem five_ionization_shifts : penSum = 1274 := by decide

/-- The mean shift is 1274/5 centilogs, printed 2.55. -/
theorem ion_mean_rounds_to_255 : roundsTo 1274 500 255 100 = true := by decide

/-- The population variance is 2904/25 centilogs squared, that is
    363/31250 in log units squared, so the standard deviation rounds to
    0.11 log units.  The internal consistency across five acids is
    striking, and is still not a measurement. -/
theorem ion_sd_rounds_to_011 : sqrtRounds 363 31250 11 100 = true := by decide

/-- The shift is 3822/1055 times the mean hydroxyl increment, printed
    3.6, and it exceeds twice that increment -- decided exactly, without
    forming either quotient. -/
theorem ion_ratio_rounds_to_36 :
    roundsTo 3822 1055 36 10 = true := by decide

theorem ionization_dominates : 2 * (211 * 5) < 1274 * 3 := by decide

/-- q_ion = 10^(637/250) rounds to 353. -/
theorem q_ion_rounds_to_353 : antilogRounds 637 250 353 1 = true := by decide

/-- The two rungs differ by 10^(2767/1500), printed as a factor of 70.
    A ladder whose rungs differ by that much is not a geometric ladder,
    and any result using the geometric form across mixed arms is using a
    form the arc does not have. -/
theorem rung_ratio_exponent : 1274 * 3 = 211 * 5 + 2767 := by decide

theorem rung_ratio_rounds_to_70 :
    antilogRounds 2767 1500 70 1 = true := by decide

/-- The mixed four-arm aggregate q_h^2 q_ion^2 is 10^(4877/750); its
    mantissa 3.18 is bracketed once six decades are removed. -/
theorem mixed_ladder_exponent :
    2 * (211 * 5) + 2 * (1274 * 3) = 9754 ∧ 9754 * 750 = 4877 * 1500 := by
  decide

theorem mixed_ladder_mantissa :
    antilogRounds 377 750 318 100 = true := by decide
```

```
/-- It exceeds the single-type ladder q^4 by 10^(2767/750), printed as a
    factor of 4890. -/
theorem ladder_ratio_exponent : 9754 = 4 * (211 * 5) + 5534 := by decide

theorem ladder_ratio_rounds_to_4890 :
    antilogRounds 2767 750 4890 1 = true := by decide

/-- And the mixed aggregate exceeds q^4, decided in the exponents alone
    with no antilogarithm taken. -/
theorem mixed_exceeds_pure : 4 * (211 * 5) < 9754 := by decide

/-- Rem. 10.4.  The last two arms are worth five log units of aqueous
    preference between them, and a unit preferring water by that margin
    does not form a distinct phase but dissolves.  So monotone partition
    is not monotone separation, and nothing in Article 82 locates the
    amphiphilic optimum. -/
theorem last_two_arms_exceed_five_log_units : 500 ≤ 2 * p_vanillic := by decide


/- ================================================================
   S15.  THE TALLY, AND THE SCOPE
   ================================================================ -/

def nObligations : Nat := 8
def nWithdrawals : Nat := 6
def nAccepted : Nat := 6
def nAcceptedOtherForm : Nat := 2
def nRebutted : Nat := 1

theorem obligations_recorded : nObligations = 8 := by decide

theorem withdrawals_traced : nWithdrawals = 6 := by decide

/-- The review record closes: six findings accepted, two accepted in a
    form other than the one recommended, one rebutted. -/
theorem review_record_closes :
    nAccepted + nAcceptedOtherForm + nRebutted = 9 := by decide

/-- Two of the six withdrawals are against results of Article 82's own
    making, which is what distinguishes a withdrawal ledger from a
    correction notice. -/
theorem some_withdrawals_are_self_directed : 2 ≤ nWithdrawals := by decide

end ChryseneFormalism.Article82

/-
  ================================================================
  SCOPE
  ================================================================

  Every theorem above is discharged by `decide`: the goal is a decidable
  proposition, the kernel evaluates a verified Boolean decision procedure,
  and the proof is accepted only when that evaluation yields `true`.
  There is no `sorry`, no `native_decide`, no Mathlib import and no
  appeal to `propext`, `Quot.sound`, `funext` or `Classical.choice`.

  WHAT THAT BUYS, AND WHAT IT DOES NOT.
```

```
  The arithmetic of Article 82 is certified completely.  Every increment,
  mean, spread, variance, count and ratio is exhibited as an equation or
876 an inequality between naturals, and every figure the article prints is
  pinned by a bracket that fails if the figure is moved by one unit in
  its last place.  The three transcendental quantities -- two standard
  deviations and a family of base-ten antilogarithms -- are never formed;
  each is bracketed between integers instead.
881
  None of that reaches the data.  The values in S2 enter Article 82 under
  Hypothesis H_data, which asserts that they are measurements of
  octanol/water partition at the stated conditions.  That is a premise
  about the world.  It is not a theorem of this corpus, it is not an
886 axiom of this file, and no kernel can settle it.  A reader who takes
  "171 theorems, all axiom-free" from the companion articles and expects
  the same standing here has mistaken the kind of claim being made: in
  Articles 79, 80 and 81 a certified theorem is a result; here it is a
  guarantee that the measured numbers were handled correctly.
891
  Three further limits deserve stating plainly.

  First, the paper's central quantity q is a FLOOR and not an estimate,
  and it is a floor only under H_ref and H_cav (Article 82, Hyp. 4.3 and
896 4.5).  H_cav is named, argued and unmeasured.  The brackets above
  certify the value of the floor; nothing certifies that it is one.

  Second, Prop. 10.1 rests on entries the source itself does not present
  as measurements: all five pH 7 values are flagged estimated or
901 corrected.  S14 certifies that arithmetic exactly, which is a
  guarantee about the computation and not about the quantity.  The sign
  is secure for reasons that need no measurement; the magnitude is not.

  Third, the harness that accompanies this file encodes the article's
906 conclusions in its assertions -- it fails if H_bulk is not refuted, and
  if the benzylic drift is not flat.  That makes it a regression test on
  the conclusions rather than an independent test of them, and Article 82
  Appendix A says so.  The same is true of this file.
-/
```

Kernel Axiom Audit

Axiom Inventory

Running `#print axioms` on each of the 109 theorems in `acf-82-proof.lean` yields:

Result	Meaning	Count
no axioms	The proof term reduces to primitive inference steps alone. Nothing is assumed beyond the type theory.	**109**
`propext`	Not used.	0
`Quot.sound`	Not used.	0
`funext`	Not used.	0
`Classical.choice`	Not used.	0
Extra axioms introduced by this proof:		**none**

The four proof files of this issue, measured

The figures below were measured by re-running `#print axioms` over every theorem of each file, not carried over from the dossiers in which they first appeared.

Article	Kind	Checks	Theorems	Axiom-free
79	formal	88	55	54
80	formal	99	61	61
81	matching	91	55	55
82	concordance	120	109	**109**
Total		**398**	**280**	**279**

Across all four files exactly one theorem depends on an axiom: Article 79's `demand_gcd`, which inherits `propext` from `Nat.gcd`. The other 279 are axiom-free.

The theorem count of this file is the largest of the four and should not be read as the strongest result. It is large because a concordance paper prints a great many figures and each needs its own certificate, and because the arithmetic is shallow enough that a theorem costs little. A file of 109 shallow certificates over measured data is a weaker thing than a file of 55 certificates over a formal construction, and the count runs the other way from the standing.

Verification Note

The proof file makes no use of `native_decide`, which bypasses the kernel by running compiled code; of `sorry`; of `Classical.choice` or `Classical.em`; or of any Mathlib import. The `decide` tactic in Lean 4 core reduces a goal to a kernel-evaluated Boolean expression, and the kernel itself performs the evaluation.

Two `set_option` lines appear at the head of the file and neither weakens anything. `exponentiation.threshold` raises the point at which the *elaborator* declines to pre-evaluate a power; past it Lean emits a warning, not an error, and the kernel evaluates the power either

way. Every exponent in this file — up to 1500 — exceeds the default of 256, so without the raise the run is noisy rather than wrong. `maxRecDepth` is raised for the same reason: the default refuses the deeper evaluations outright. Neither option admits an axiom, and the audit above is the evidence.

To reproduce: install Lean 4.12.0 and run `lean acf-82-proof.lean`. A silent run with exit code 0 is the certificate; it takes about sixteen seconds, almost all of it in the exhaustive search of S3. There are no dependencies beyond the Lean core, so no `lakefile` and no package fetch are required.

Scope: three limits, stated plainly

A kernel-checked theorem certifies that a conclusion follows from the definitions in the file. For this article that is less than it is for the other three, and the difference has already been stated: the data are a premise about the world, and no kernel reaches them. Three further limits belong with it.

First, q is a floor only under an unmeasured premise. Proposition 4.4 makes the octanol-referenced increment a lower bound on the tar-referenced one, under H_{ref} and H_{cav}. H_{cav} is named, argued and unmeasured, and it is what carries the word "floor" through §7 of the article and into its conclusion. S9 certifies the *value*; nothing here certifies that it is a floor.

Second, S14 rests on estimated inputs. All five pH 7 values are flagged by their source as estimated or corrected, which S2 asserts rather than hides. The arithmetic of S14 is exact — that is a guarantee about the computation, not about the quantity. The sign is secure for reasons needing no measurement; the magnitude is not, and the article's Remark 10.2 says so.

Third, both artifacts encode the article's conclusions. The harness fails if H_{bulk} is not refuted and if the benzylic drift is not flat; this file's S10 and S13 do the same. That is deliberate and Article 82 Appendix A states it: a published verification artifact should stop passing when the claims it supports stop holding. But it makes both files regression tests on the conclusions rather than independent tests of them, and a reader re-testing a hypothesis on new data should expect to edit the certificate as well as the table.