

Supporting Information for: Recovering Parametric Inference for Skewed, Small-Sample, and Heteroscedastic Data: The Coefficient-of-Variation Screen and the Log-Scale Variance Ratio (ρ). William J. Dwyer.

SUPPORTING INFORMATION

Appendix S1. The Coefficient-of-Variation Screen and the Log-Scale Variance Ratio ρ

The Log-Scale Variance Ratio ρ : A Derivation and Its Relation to the Coefficient-of-Variation Ratio and $\ln\text{CVR}$

W. J. Dwyer

Mathematical appendix to “Recovering Parametric Inference for Skewed, Small-Sample, and Heteroscedastic Data: The Coefficient-of-Variation Screen and the Log-Scale Variance Ratio (ρ)”

Summary

This note derives, step by step, the quantity that governs the variance-model choice for two-group comparisons of positive, right-skewed (approximately log-normal) data: the log-scale variance ratio, $\rho = \ln(1 + \text{CV}_2^2)/\ln(1 + \text{CV}_1^2)$. Beginning from the moments of the log-normal distribution, we establish the fundamental identity $\sigma^2 = \ln(1 + \text{CV}^2)$ relating a log-normal variable's log-scale variance to its coefficient of variation, show that the coefficient of variation is therefore a pure dispersion parameter independent of the mean, and use this to separate three quantities that are easily conflated: the **CV ratio** CV_2/CV_1 (the quantity a reader computes from a published table; a screen and diagnostic), the **log coefficient-of-variation ratio** $\ln\text{CVR} = \ln(\text{CV}_2/\text{CV}_1)$ (the established meta-analytic effect size and prior art; the logarithm of the CV ratio), and the **mean-free log-scale variance ratio** ρ (the decision-relevant parameter that the CV ratio feeds, the estimand the variance-ratio F-test targets, mean-free because the group means cancel, which fixes the pooled Student versus the Welch form of the log-scale t-test; its sample analogue s_2^2/s_1^2 is the test statistic). We prove the small-coefficient-of-variation approximation $\rho \approx (\text{CV ratio})^2$, exhibit the exact and asymptotic relationships among the three, and close with a worked numerical example. The three quantities are one family: the CV ratio is the input, $\ln\text{CVR}$ is its logarithm, and ρ is the log-scale variance ratio that actually drives the variance-model decision. All identities are verified numerically.

1. Introduction and the three-way distinction

For a positive, right-skewed outcome that is approximately log-normal, dispersion is naturally expressed by the coefficient of variation, $\text{CV} = \sigma_X / \mu_X$, the ratio of the natural-scale

standard deviation to the natural-scale mean. When two groups are compared, three related quantities built from their coefficients of variation arise, and they play distinct roles. They are easily confused because each is a monotone transformation of the same underlying contrast, yet only one of them is the parameter on which the choice between the pooled and the Welch t-test turns. The purpose of this note is to derive that statistic from first principles and to fix the relationships among the three.

The three quantities are:

- (i) the **coefficient-of-variation ratio**, CV_2/CV_1 — the quantity recoverable directly from two published means and standard deviations; it serves as a screen and a heteroscedasticity diagnostic;
- (ii) the **log coefficient-of-variation ratio**, $\ln CVR = \ln(CV_2/CV_1)$ — the established effect size of the meta-analysis of variation [5,6], i.e., the logarithm of the CV ratio, constructed to be combined across studies; and
- (iii) the **mean-free log-scale variance ratio**, $\rho = \ln(1 + CV_2^2)/\ln(1 + CV_1^2)$ — the ratio of the two groups' variances on the logarithmic scale (mean-free because the means cancel), which is the scale on which log-normal data are analyzed and on which equal variance is the condition for pooling; ρ is the decision-relevant parameter (the estimand the variance-ratio F-test targets), and $\rho \approx (CV \text{ ratio})^2$ for small CV.

The remainder establishes (ii) and (iii) from the log-normal moments, proves the small-CV approximation, and states the exact relations among all three.

2. Preliminaries: moments of the log-normal distribution

Let X be a positive random variable such that its logarithm is normally distributed,

$$Y = \ln X \sim N(\mu, \sigma^2),$$

so that $X = e^Y$ is log-normal with log-scale location μ and log-scale variance σ^2 .

Throughout, μ and σ^2 are the parameters on the logarithmic scale, while $\mu_X = E[X]$ and $\sigma_X^2 = \text{Var}[X]$ denote the mean and variance on the natural (untransformed) scale. We use the normal moment generating function, which for $Y \sim N(\mu, \sigma^2)$ is

$$E[e^{tY}] = \exp(\mu t + \frac{1}{2}\sigma^2 t^2), \text{ for all real } t. \quad (1)$$

Evaluating the generating function at $t = 1$ and $t = 2$ gives the first two natural-scale moments of X :

$$E[X] = E[e^{\wedge}Y] = \exp(\mu + \frac{1}{2}\sigma^2), \quad (2)$$

$$E[X^2] = E[e^{\wedge}2Y] = \exp(2\mu + 2\sigma^2). \quad (3)$$

The natural-scale variance follows by $\text{Var}[X] = E[X^2] - (E[X])^2$:

$$\text{Var}[X] = \exp(2\mu + 2\sigma^2) - \exp(2\mu + \sigma^2) = \exp(2\mu + \sigma^2) \cdot [\exp(\sigma^2) - 1]. \quad (4)$$

Equations (2) and (4) are the only facts about the log-normal distribution that the derivation requires; both are standard [1,2].

3. The fundamental identity: $\sigma^2 = \ln(1 + \text{CV}^2)$

The coefficient of variation of X is, by definition, the natural-scale standard deviation divided by the natural-scale mean, $\text{CV} = \sqrt{\text{Var}[X]} / E[X]$; equivalently $\text{CV}^2 = \text{Var}[X] / (E[X])^2$. Substituting (4) for the numerator and the square of (2) for the denominator,

$$\text{CV}^2 = \text{Var}[X] / (E[X])^2 = \{ \exp(2\mu + \sigma^2) \cdot [\exp(\sigma^2) - 1] \} / \{ \exp(2\mu + \sigma^2) \}. \quad (5)$$

The factor $\exp(2\mu + \sigma^2)$ is common to numerator and denominator and cancels exactly, leaving

$$\text{CV}^2 = \exp(\sigma^2) - 1. \quad (6)$$

Solving (6) for the log-scale variance gives the identity on which everything else rests:

$$\exp(\sigma^2) = 1 + \text{CV}^2, \text{ hence } \sigma^2 = \ln(1 + \text{CV}^2). \quad (7)$$

Equation (7) expresses the log-scale variance σ^2 entirely in terms of the natural-scale coefficient of variation CV. It is exact for any log-normal variable, requires no knowledge of the mean, and is invertible: $\text{CV} = \sqrt{\exp(\sigma^2) - 1}$. A numerical check confirms it — simulating $Y \sim N(\mu, \sigma^2)$ and forming $X = e^{\wedge}Y$ reproduces $\text{CV} = \sqrt{\exp(\sigma^2) - 1}$ and $\sigma^2 = \ln(1 + \text{CV}^2)$ to within Monte-Carlo error across $\sigma = 0.3$ to 1.5 and several values of μ .

4. The coefficient of variation is a pure dispersion parameter

A direct corollary of (6) is that the coefficient of variation, and through (7) the log-scale variance, depends only on σ and not at all on the location μ :

$$\text{CV}^2 = \exp(\sigma^2) - 1 \text{ is independent of } \mu. \quad (8)$$

Two log-normal variables that share the same σ but differ in μ (that is, that differ only by a multiplicative shift in scale, since changing μ multiplies X by a constant $e^{\wedge}\{\Delta\mu\}$) have identical coefficients of variation. The coefficient of variation is therefore a pure **dispersion** (shape) parameter, free of **location**. This is the property that makes it usable from summary statistics: a

difference in coefficient of variation between two groups reflects a genuine difference in relative dispersion, not a difference in their means. (Simulation confirms invariance: with $\sigma = 0.7$ fixed, $\mu = 0, 2, 5$ all give $CV \approx 0.795$.)

5. Why the natural-scale variance ratio confounds dispersion with the mean

On the natural scale, the standard deviation is tied to the mean through the coefficient of variation: from $CV = \sigma_X / \mu_X$ we have $\sigma_X = CV \cdot \mu_X$. For two groups indexed 1 and 2, the ratio of natural-scale variances is therefore

$$\sigma_{\{X,2\}}^2 / \sigma_{\{X,1\}}^2 = (CV_2 \cdot \mu_{\{X,2\}})^2 / (CV_1 \cdot \mu_{\{X,1\}})^2 = (CV_2 / CV_1)^2 \cdot (\mu_{\{X,2\}} / \mu_{\{X,1\}})^2. \quad (9)$$

Equation (9) factors the raw variance ratio into a **dispersion** term, $(CV_2 / CV_1)^2$, and a **location** term, $(\mu_{\{X,2\}} / \mu_{\{X,1\}})^2$. A large natural-scale variance ratio can therefore be produced almost entirely by a difference in means, even when the two groups share the same coefficient of variation. The natural-scale variance ratio is thus not a clean measure of heteroscedasticity. The remedy is to move to the logarithmic scale, where, as the next section shows, the location term disappears.

6. The log-scale variance ratio ρ

Log-normal data are analyzed on the logarithmic scale, where $Y = \ln X$ is normal and the two-sample t-test applies. On that scale the variance of group i is, by the fundamental identity (7),

$$\sigma_i^2 = \ln(1 + CV_i^2), \quad i = 1, 2. \quad (10)$$

The between-group log-scale variance ratio — the quantity an F-test compares to decide whether the two log-scale variances may be treated as equal — is the ratio of (10) for the two groups:

$$\rho \equiv \sigma_2^2 / \sigma_1^2 = \ln(1 + CV_2^2) / \ln(1 + CV_1^2). \quad (11)$$

Comparison of (11) with (9) shows precisely what the logarithm accomplishes. The location terms $\mu_{\{X,i\}}$, which dominate the natural-scale ratio, are absent from ρ ; the log-scale variance ratio depends on the coefficients of variation alone. It is therefore the genuine, mean-free measure of dispersion inequality between the two groups, and $\rho = 1$ exactly when the two groups are homoscedastic on the log scale. Because ρ is a function of the coefficients of variation only, it is recoverable from two published means and standard deviations without the raw data.

7. The small-coefficient-of-variation approximation: $\rho \approx (\text{CV ratio})^2$

The relationship between ρ and the square of the CV ratio is made precise by expanding the logarithm. For $|u| < 1$ the Maclaurin series is $\ln(1 + u) = u - \frac{1}{2}u^2 + \frac{1}{3}u^3 - \dots$. Taking $u = CV^2$,

$$\ln(1 + CV^2) = CV^2 - \frac{1}{2}CV^4 + \frac{1}{3}CV^6 - \dots = CV^2 \cdot (1 - \frac{1}{2}CV^2 + \dots). \quad (12)$$

Substituting (12) into (11) for both groups and cancelling the leading CV^2 factors,

$$\rho = [CV_2^2(1 - \frac{1}{2}CV_2^2 + \dots)] / [CV_1^2(1 - \frac{1}{2}CV_1^2 + \dots)] = (CV_2/CV_1)^2 \cdot [1 - \frac{1}{2}(CV_2^2 - CV_1^2) + \dots]. \quad (13)$$

Two consequences follow. First, in the small-coefficient-of-variation limit ($CV \rightarrow 0$) the bracketed correction tends to 1, so

$$\rho \rightarrow (CV_2/CV_1)^2 = (\text{CV ratio})^2 \text{ as } CV \rightarrow 0. \quad (14)$$

The log-scale variance ratio is approximately the **square** of the CV ratio when the coefficients of variation are small. Second, because every higher-order term in (12) is subtracted, $\ln(1 + CV^2) < CV^2$ for all $CV > 0$, so ρ grows **sub-quadratically** in the CV ratio as the coefficients of variation increase: the same CV ratio implies a smaller log-variance ratio at larger CV. Equivalently, near-normal (small-CV) data are the least forgiving of a CV-ratio departure, since there ρ tracks the full square; strongly skewed data tolerate a wider CV ratio before the log-variances become as unequal.

8. $\ln\text{CVR}$ and its relation to ρ

The log coefficient-of-variation ratio of the meta-analysis of variation is, by definition, the logarithm of the CV ratio [5,6]:

$$\ln\text{CVR} = \ln(CV_2/CV_1) = \ln CV_2 - \ln CV_1. \quad (15)$$

It is an effect size — a single signed number that summarizes the dispersion contrast between two groups (adjusting for the mean–variance relationship that the coefficient of variation already removes) and that can be combined across studies in a random-effects meta-analysis. Its relation to ρ follows from the small-CV approximation. Taking logarithms of (14), for small CV,

$$\ln \rho \approx 2 \cdot \ln(CV_2/CV_1) = 2 \cdot \ln\text{CVR}, \text{ so } \ln\text{CVR} \approx \frac{1}{2} \cdot \ln \rho. \quad (16)$$

Thus in the small-CV regime the three quantities are monotone transformations of one another: $\rho \approx (\text{CV ratio})^2$, $\ln\text{CVR} = \ln(\text{CV ratio})$, and $\ln\text{CVR} \approx \frac{1}{2} \ln \rho$. The correspondence (16) is asymptotic, not exact; it degrades as the coefficients of variation grow, where the exact expression (11) must be used. The worked example below makes the divergence concrete.

The distinction in role, however, is exact and does not depend on the size of the CV. $\ln\text{CVR}$ is an **effect size**, an outcome to be estimated and pooled. ρ is a **decision-relevant parameter** (the estimand of the variance-ratio test), read before any test is run to select the variance model. They answer different questions even though, for small CV, they are reparametrizations of the same contrast.

9. Why ρ selects the pooled versus the Welch test

On the logarithmic scale the two samples are normal, $Y_{\{i\}} \sim N(\mu_i, \sigma_i^2)$. The pooled, equal-variance (Student) two-sample t-test [7] assumes $\sigma_1^2 = \sigma_2^2$, that is, $\rho = 1$; under that assumption and normality it is exact and is the uniformly most powerful unbiased test of equal log-scale means. When $\rho \neq 1$ the comparison is the Behrens–Fisher problem transported to the log scale: the pooled test is no longer exact, and its true size departs from the nominal level by an amount that depends jointly on ρ and on the allocation of sample size between the groups — liberal when the smaller group is the more variable and conservative when the larger group is, with balanced designs robust to any ρ [8]. The Welch approximate-degrees-of-freedom test [9] restores calibration without assuming equal variances.

The statistic that adjudicates this choice from data is the ratio of sample variances on the log scale. Under the null hypothesis of equal log-scale variances ($\rho = 1$) and normality,

$$F = s_2^2 / s_1^2 \sim F(n_2 - 1, n_1 - 1), \quad (17)$$

so comparing s_2^2/s_1^2 to the quantiles of the F distribution is the formal test of homoscedasticity, and s_2^2/s_1^2 is the sample estimate of the population quantity ρ in (11). In the language of estimation the roles are clear: ρ is the **parameter** — the estimand — and the F-ratio s_2^2/s_1^2 is its **estimator and the corresponding test statistic**. ρ is therefore a decision-relevant **parameter**, not a sample statistic; the present method uses it as a design rule by reading it from published coefficients of variation rather than by testing it. The decision between pooled and Welch turns on ρ , and ρ is exactly what the coefficients of variation deliver through (11). This is the sense in which the CV ratio (the input) feeds ρ (the decision parameter): one reconstructs ρ from the published coefficients of variation and reads off the variance model.

10. Worked numerical example

Take a control group with coefficient of variation $CV_1 = 1.0$ and a disease group with $CV_2 = 2.27$ (the values arising from a bilirubin comparison in the related work). The log-scale variances follow from (7):

$\sigma_1^2 = \ln(1 + 1.0^2) = \ln 2 = 0.6931$, $\sigma_2^2 = \ln(1 + 2.27^2) = \ln 6.1529 = 1.8169$,
and the log-scale variance ratio from (11) is

$$\rho = 1.8169 / 0.6931 = 2.62.$$

The genuine log-scale heteroscedasticity is thus a 2.62-fold difference in variance. Three comparisons sharpen the interpretation. First, the small-CV approximation (14) gives $(CV_2/CV_1)^2 = 2.27^2 = 5.15$, which overshoots $\rho = 2.62$ by a wide margin precisely because the coefficients of variation here are far from small, so the exact logarithmic form is required. Second, the natural-scale variance ratio, from standard deviations of 3.4 and 0.5, is $(3.4/0.5)^2 = 46.2$ (more than an order of magnitude larger than ρ) because, by (9), most of it is the roughly threefold difference in means rather than dispersion. Third, the effect size is $\ln CVR = \ln(2.27) = 0.82$, whereas $\frac{1}{2} \cdot \ln \rho = \frac{1}{2} \cdot \ln 2.62 = 0.48$; the small-CV identity (16) is visibly inaccurate at these large coefficients of variation, as expected. A 2.62-fold log-variance inequality places this comparison outside the pooled test's tolerance under realistic sample-size imbalance, so the Welch form of the log-scale t-test is the calibrated choice.

11. Summary: the family of quantities

The three quantities derived above form a single family built from the two groups' coefficients of variation. They are monotone transformations of the same underlying contrast in the small-CV limit, yet they occupy distinct roles, summarized in Table 1.

Quantity	Definition	Role	Small-CV relation
CV ratio	CV_2/CV_1	Input; screen and heteroscedasticity diagnostic, read from a published table	—
$\ln CVR$	$\ln(CV_2/CV_1)$	Effect size (prior art); a dispersion contrast to be estimated and pooled across studies	$\ln CVR = \ln(\text{CV ratio}); \approx \frac{1}{2} \ln \rho$
ρ (log-scale variance ratio)	$\ln(1 + CV_2^2) / \ln(1 + CV_1^2)$	Decision-relevant parameter (estimand of the variance-ratio F-test); selects pooled (Student) vs Welch	$\rho \approx (\text{CV ratio})^2$

Table 1. The coefficient-of-variation ratio, the log coefficient-of-variation ratio ($\ln CVR$), and the log-scale variance ratio ρ . The CV ratio is the input a reader computes; $\ln CVR$ is its logarithm and the established meta-analytic effect size; ρ is the ratio of log-scale variances that drives the variance-model decision and is approximately the square of the CV ratio for small CV.

Definitions and identities computed by rho_toolkit.py ($\sigma^2 = \ln(1 + CV^2)$, eq 7; $\rho = \ln(1 + CV_2^2)/\ln(1 + CV_1^2)$, eq 11; $\ln CVR = \ln(CV_2/CV_1)$, eq 15); each identity verified in verify_rho_derivation.py.

The relationships are shown graphically in Figure 1, computed directly from the accompanying software. Panel (a) is the fundamental identity (7), the natural-scale coefficient of variation mapped to the log-scale variance, with the small-CV line $\sigma^2 \approx CV^2$ for comparison. Panel (b) plots ρ against the CV ratio for several baseline coefficients of variation, against the $(CV \text{ ratio})^2$ reference of (14): the small-CV curves hug the square, while larger coefficients of variation bend sub-quadratically below it. Panel (c) overlays the three family members as transforms of the CV ratio — the linear input, the logarithmic effect size $\ln CVR$, and the decision parameter ρ , together with $\frac{1}{2} \ln \rho$ to display the small-CV correspondence (16). Panel (d) is the operational consequence: across ρ , the pooled (Student) log-scale t-test under a four-to-one sample-size imbalance leaves the acceptable size band — liberal for $\rho < 1$, conservative for $\rho > 1$ — while a balanced pooled test and the Welch test hold the nominal level, so it is ρ , read from the coefficients of variation, that selects the variance model.




Figure 1. The log-scale variance ratio ρ and its family, computed from the accompanying toolkit. (a) The fundamental identity $\sigma^2 = \ln(1 + CV^2)$ (equation 7) with the small-CV approximation $\sigma^2 \approx CV^2$. (b) $\rho = \ln(1 + CV_2^2)/\ln(1 + CV_1^2)$ versus the CV ratio for baseline $CV_1 = 0.1, 0.5, 1.0, 2.0$, against the small-CV limit $(CV \text{ ratio})^2$

(dashed): near-normal data track the square, strongly skewed data fall below it (sub-quadratic). (c) The CV ratio (input), $\ln\text{CVR} = \ln(\text{CV ratio})$ (effect size), and ρ (decision parameter) as functions of the CV ratio at $\text{CV}_1 = 1$, with $\frac{1}{2} \ln \rho$ illustrating the small-CV relation $\ln\text{CVR} \approx \frac{1}{2} \ln \rho$. (d) Type-I error of the pooled (Student) and Welch log-scale t-tests under the null as a function of ρ (log axis): the pooled test mis-calibrates with ρ under a 1:4 imbalance — liberal for $\rho < 1$, conservative for $\rho > 1$ — while the balanced pooled test and Welch hold the 0.04–0.06 band; 4,000 Monte-Carlo replicates per point. Rendered by the deposited `fig_rho_family.py`.

In one sentence: the coefficient-of-variation ratio is the input, $\ln\text{CVR}$ is its logarithm and the prior-art effect size, and the log-scale variance ratio ρ is the quantity that actually drives the pooled-versus-Welch decision.

Software availability

Every identity in this note can be re-derived and re-checked by a reviewer with the accompanying open-source package, which establishes each result twice — symbolically with SymPy (exact algebra) and numerically with NumPy/SciPy (Monte-Carlo). It comprises `rho_toolkit.py`, in which each function names the equation it implements ($\sigma^2 = \ln(1 + \text{CV}^2)$, equation 7; ρ , equation 11; $\ln\text{CVR}$, equation 15); `verify_rho_derivation.py`, which prints a plain-language verification report, computing each numbered quantity from the stated inputs and printing it beside the value derived here so a reader can read the agreement directly, including a size study confirming that the pooled log-scale t-test mis-calibrates with ρ under sample-size imbalance while Welch's test holds its nominal level (Section 9); a figure generator, `fig_rho_family.py`, that produces Figure 1 directly from the toolkit; and a README mapping each check to the manuscript. Running “`python verify_rho_derivation.py`” prints the full report, and “`python fig_rho_family.py`” regenerates the figure.

12. Verification reference table

Each numbered result is reproducible from the deposited code. The table records how each was established (P = symbolic proof; S = simulation; S->P = simulation then proof), the routine that carries it, and the command to reproduce it. Symbolic identities are checked with SymPy; the distributional and size results are checked by Monte-Carlo; no image tools are used for figures.

Result	How	Code routine	Command	Reported value
$\sigma^2 = \ln(1 + \text{CV}^2)$ (eq 7)	P	<code>verify_rho_derivation.py</code>	<code>python verify_rho_derivation.py</code>	$\sigma^2 = \ln(1 + \text{CV}^2)$ (symbolic)
CV is mean-free (eq 8)	P, S	<code>verify_rho_derivation.py</code>	(same)	$d/d\mu \text{CV}^2 = 0$; CV invariant to μ
Raw variance ratio	P	<code>verify_rho_derivation.py</code>	(same)	$(\text{CV}_2/\text{CV}_1)^2$

$= (CV2/CV1)^2$ ($\mu2/\mu1$) ² (eq 9)		on.py		$(\mu2/\mu1)^2$ (symbolic)
$\rho \rightarrow (CV \text{ ratio})^2$ as $CV \rightarrow 0$ (eq 14)	P	verify_rho_derivati on.py	(same)	limit = $(CV2/CV1)^2$
$\ln CVR = (1/2) \ln$ ρ_{approx} (eq 16)	P	verify_rho_derivati on.py	(same)	$\ln((CV \text{ ratio})^2) =$ $2 \ln CVR$
$s2^2/s1^2 \sim F(n2-$ $1, n1-1)$ under $\rho = 1$ (eq 17)	S->P	verify_rho_derivati on.py	(same)	KS $p > 0.01$ vs $F(n2-1, n1-1)$
Welch holds nominal size; pooled departs under imbalance when $\rho \neq 1$	S	verify_rho_derivati on.py	(same)	Welch ~ 0.05 ; pooled departs

Table 2. Verification reference table. P = proof, S = simulation, S->P = simulation then proof. Each row names the routine, the command that runs it, and the value the software reports for a reviewer to read against the derivation.

References

1. Aitchison J, Brown JAC. The Log-normal Distribution, with Special Reference to Its Uses in Economics. Cambridge: Cambridge University Press; 1957.
2. Johnson NL, Kotz S, Balakrishnan N. Continuous Univariate Distributions, Volume 1. 2nd ed. New York: Wiley; 1994. (Log-normal moments, Chapter 14.)
3. Limpert E, Stahel WA, Abbt M. Log-normal distributions across the sciences: keys and clues. BioScience. 2001;51(5):341-352. doi:10.1641/0006-3568(2001)051[0341:LNDATS]2.0.CO;2.
4. Koopmans LH, Owen DB, Rosenblatt JI. Confidence intervals for the coefficient of variation for the normal and log normal distributions. Biometrika. 1964;51(1-2):25-32. doi:10.1093/biomet/51.1-2.25.
5. Nakagawa S, Poulin R, Mengersen K, Reinhold K, Engqvist L, Lagisz M, Senior AM. Meta-analysis of variation: ecological and evolutionary applications and beyond. Methods Ecol Evol. 2015;6(2):143-152. doi:10.1111/2041-210X.12309.
6. Senior AM, Viechtbauer W, Nakagawa S. Revisiting and expanding the meta-analysis of variation: the log coefficient of variation ratio, $\ln CVR$. Res Synth Methods. 2020;11(4):553-567. doi:10.1002/jrsm.1423.
7. Student. The probable error of a mean. Biometrika. 1908;6(1):1-25. doi:10.1093/biomet/6.1.1.
8. Box GEP. Some theorems on quadratic forms applied in the study of analysis of variance problems, I: effect of inequality of variance in the one-way classification. Ann Math Stat. 1954;25(2):290-302. doi:10.1214/aoms/1177728786.
9. Welch BL. The generalization of 'Student's' problem when several different population variances are involved. Biometrika. 1947;34(1-2):28-35. doi:10.1093/biomet/34.1-2.28.

Appendix S2. The Polygamma Bridge — the gamma-model companion to the coefficient-of-variation screen

The Polygamma Bridge

A gamma-model companion to the coefficient-of-variation screen

William J. Dwyer · ORCID 0009-0004-0855-7222 · cross-cutting derivation note (companion to M1, M3)

Abstract

The M1 screen reads a single natural-scale statistic, the coefficient of variation $CV = SD/\text{mean}$, and converts it to a log-scale variance through the log-normal identity $\sigma^2 = \ln(1 + CV^2)$. This note shows that the same conversion has an exact gamma-model counterpart written with the **trigamma function**: for a gamma variate the log-scale variance equals $\psi'(1/CV^2)$. The two bridges agree to second order in CV but split at order CV^4 with opposite sign (the log-normal value falls below CV^2 , the gamma value rises above it), so the gap is itself a model-discrimination statistic that upgrades M1's two-way honesty check into a three-way screen (log-normal vs gamma vs neither). The same shape parameter $k = 1/CV^2$ is shown to be the quantity the M3 count-overdispersion screen reads (the negative binomial is a gamma–Poisson mixture with dispersion index $D = 1 + \mu/r$), unifying the continuous-CV and count-dispersion screens under one latent gamma; and the same moment-matching-to-a-gamma move is shown to underlie the Welch–Satterthwaite (M1, M34), Box (1954), and Greenhouse–Geisser (M37) approximations. Skewness-from-CV laws and a skewness bracket are stated with their exact scope. Every numbered result is reproduced by the deposited `verify_polygamma_bridge.py`. The underlying special-function identities are classical; the contribution is the unification and its operationalization as a portable screen, not new mathematics.

1. Setup: the M1 identity recalled

Let $X > 0$ be a positive outcome with natural-scale mean μ , standard deviation σ_n , and coefficient of variation $CV = \sigma_n/\mu$. M1 (the CV screen) is built on the log-normal bridge: if X is log-normal, then $Y = \ln X$ is exactly normal with variance

$$\sigma^2 = \ln(1 + CV^2), \quad (1)$$

so the log-scale spread is recoverable from the natural-scale CV alone. The screen's honesty check compares the directly estimated log variance s_y^2 against the value (1) implied by the CV; close agreement supports the log-normal model. The question of this note is what the **gamma**

family — the other canonical positive, right-skewed model, and the mixing law behind overdispersed counts and frailty survival — says about the same quantity.

2. The gamma bridge and the $O(CV^4)$ split (Result 1, P)

Let $X \sim \text{Gamma}(k, \theta)$ with shape k and scale θ , so $\mu = k\theta$ and $\sigma_n^2 = k\theta^2$. Then

$$CV^2 = 1/k, \quad \text{i.e.} \quad k = 1/CV^2, \quad (2)$$

a defining feature of the gamma family: the CV fixes the shape and nothing else. The log-scale moments of a gamma variate are classical (the log-gamma distribution): with $Y = \ln X$,

$$E[Y] = \psi(k) + \ln \theta, \quad \text{Var}[Y] = \psi'(k), \quad (3)$$

where ψ is the digamma and ψ' the trigamma function. Substituting (2), the gamma counterpart of the M1 identity (1) is the exact statement

$$\text{Var}[\ln X] = \psi'(1/CV^2) \text{ (gamma)} \quad \text{vs} \quad \ln(1 + CV^2) \text{ (log-normal)}. \quad (4)$$

The split theorem

Expand both sides for small CV. With $u = CV^2$, $\ln(1 + u) = u - u^2/2 + u^3/3 - \dots$. The trigamma asymptotic $\psi'(k) = 1/k + 1/(2k^2) + 1/(6k^3) - \dots$ with $k = 1/u$ gives $\psi'(1/u) = u + u^2/2 + u^3/6 + \dots$. Hence

$$\psi'(1/CV^2) - \ln(1 + CV^2) = CV^4 - (1/6)CV^6 + \dots, \quad (5)$$

Both models give $\text{Var}[\ln X] = CV^2$ to second order, but they part company at order CV^4 with **opposite sign**: the log-normal log-variance sits below CV^2 (coefficient $-1/2$), the gamma log-variance above it (coefficient $+1/2$), and their difference is $+CV^4$ to leading order; the gamma log-variance exceeds the log-normal for every $CV > 0$. Numerically the gap tracks CV^4 closely in the screen's working range: at $CV = 0.3$ it is 0.0080 ($CV^4 = 0.0081$); at $CV = 0.5$ it is 0.0607 ($CV^4 = 0.0625$). [verify: Result 1]

The Polygamma Bridge — the gamma companion to the coefficient-of-variation screen




Figure 1. Panel A: the log-scale variance implied by the natural-scale CV is $\ln(1 + CV^2)$ under log-normality (below CV^2) and $\psi'(1/CV^2)$ under the gamma model (above CV^2); the shaded band is the model-discrimination region and its width is $\approx CV^4$. Panel B: skewness is CV times a family factor — 2 for the gamma (lower edge), $3 + CV^2$ for the log-normal (upper edge) — with the inverse-Gaussian ($3 \cdot CV$) lying inside the bracket. Rendered by the deposited `fig_polygamma_bridge.py`.

3. The log-mean offset (Result 2, P)

The same second-order-agreement / fourth-order-split pattern governs the back-transform bias. The offset between the log-mean and the log of the arithmetic mean is

$$E[\ln X] - \ln \mu = \psi(1/CV^2) - \ln(1/CV^2) \text{ (gamma)} \quad \text{vs} \quad -\frac{1}{2} \ln(1 + CV^2) \text{ (log-normal)}. \quad (6)$$

Both equal $-CV^2/2$ to leading order — the familiar half-variance correction that distinguishes the geometric (median) from the arithmetic mean — and split at CV^4 with opposite sign (gamma $-1/12$, log-normal $+1/4$). This is the gamma-model footnote to M1's geometric-vs-arithmetic-mean distinction and to the Cox/Land back-transformed interval. [verify: Result 2]

4. The operational upgrade: a three-way eq-7 screen

M1's honesty check is two-way: compare the data's s_y^2 to $\ln(1 + CV^2)$. Equations (4)–(5) make it three-way at no extra cost — every input (s_y^2 , CV) is already in hand:

$$\text{compare } s_y^2 \text{ against } \ln(1 + CV^2) \text{ [log-normal]} \text{ and } \psi'(1/CV^2) \text{ [gamma]}. \quad (7)$$

Because the two model values straddle CV^2 (log-normal below, gamma above), the **position of the observed s_y^2 between them is informative**: nearer $\ln(1 + CV^2)$ points to multiplicative/log-normal structure (use the log-scale t and its back-transform); nearer $\psi'(1/CV^2)$ points to gamma/additive-shape structure (a gamma GLM with log link is the better-calibrated

rescue); outside the interval on either side flags a heavier or lighter tail than both and routes to the distribution-free backstop. This third line is added to the M1 confidence-interval calculator (the eq-7 panel) so the verdict is shown directly to the analyst.

5. One latent gamma under two screens (Result 4, P/S)

The shape parameter $k = 1/CV^2$ is not unique to the continuous screen. The negative binomial — the rescue M3 routes to for overdispersed counts — is exactly a **gamma–Poisson mixture**: a Poisson whose rate is drawn from a Gamma(shape r) law. Its dispersion index is

$$D = Var/mean = 1 + \mu/r = 1 + \mu \cdot CV_{mix}^2, \quad CV_{mix}^2 = 1/r, \quad (8)$$

so the M3 dispersion screen reads $1/r$ — the squared CV of the **same kind of latent gamma** whose continuous-scale CV the M1 screen reads. Gamma frailty in survival (the shared-frailty models M4/M8/M9 route to) is the identical mixing operation on the hazard. The three screens — continuous dispersion (M1), count overdispersion (M3), and frailty heterogeneity (M4/M8/M9) — are therefore all reading the shape of one latent gamma, and $1/k = CV^2$ is the common coordinate. A simulation confirms the empirical negative-binomial dispersion matches $1 + \mu/r$ across a grid. [verify: Result 4]

6. The match-to-a-gamma theme (Result 5, S)

A second, methodological thread runs through the scale and dependence screens: when an exact sampling distribution is an awkward nonnegative combination of chi-squares, the standard move is to **approximate it by a single scaled chi-square — a gamma — matched on its first two moments**. The Welch–Satterthwaite degrees of freedom (M1's log-scale Welch test and M34's Behrens–Fisher screen), Box's (1954) theorems on quadratic forms (M1 reference 17), and the Greenhouse–Geisser ϵ correction for sphericity (M37) are three instances of the same gamma moment-matching. The trigamma and digamma functions of §2–§3 are precisely what quantify the error of such matches on the log scale, which is why the gamma function is the analytic engine — not merely a normalizing constant — beneath this part of the catalog. [verify: Result 5]

7. Skewness from the CV: laws and a bounded conjecture (Results 3 & 6)

On the natural scale the CV also fixes skewness up to a family factor. Two exact laws (Result 3, P):

$$\gamma_1 = 2 \cdot CV \text{ (gamma, exact)} \quad \gamma_1 = (CV^2 + 3) \cdot CV \text{ (log-normal, exact)}. \quad (9)$$

The gamma gives the smaller skewness, the log-normal the larger, at any CV. This motivates a **conjecture** (Result 6, $S \rightarrow P$, stated with its scope): for the multiplicative-error class that the screens target, skewness is bracketed by

$$2 \cdot CV \leq \gamma_1 \leq (CV^2 + 3) \cdot CV. \quad (10)$$

The inverse-Gaussian (skewness $3 \cdot CV$) lies strictly inside the bracket for all tested CV, supporting (10) on the multiplicative family. The bracket is **not universal**, and the note says so plainly: a Weibull with shape > 1 ($CV < 1$) falls below the lower edge (e.g., Weibull shape 2 has $CV = 0.523$ but skewness 0.631, well under $2 \cdot CV = 1.045$) and can even be left-skewed. So (10) is a useful bracket for the gamma-to-log-normal multiplicative range that sharpens the CV/five-number-summary skewness diagnostic (M1; Shi et al. 2023), not a theorem for all positive distributions. [verify: Result 6]

8. Provenance and reproducibility

Results 1–3 are proofs (P); the series identities are produced symbolically. Result 4 is a proof with a simulation check (P/S); Result 5 is an algebraic identity verified numerically (S); Result 6 is a simulation-then-proof conjecture ($S \rightarrow P$) reported with an explicit counterexample fixing its scope. The deposited script `verify_polygamma_bridge.py` regenerates every number above and exits 0 only if all thirteen checks pass. Consistent with the program's posture, the special-function identities (trigamma log-variance, gamma–Poisson, Box 1954) are established; what is new here is the unification — one shape parameter $k = 1/CV^2$ threading the CV, dispersion, and frailty screens — and the three-way operational screen of §4. No new mathematics is claimed.

References

1. Olsson U. Confidence intervals for the mean of a log-normal distribution. *J Stat Educ.* 2005;13(1). doi:10.1080/10691898.2005.11910638.
2. Abramowitz M, Stegun IA. *Handbook of Mathematical Functions*. Washington, DC: National Bureau of Standards; 1964. (Digamma and trigamma functions, §6.3–6.4.)
3. Johnson NL, Kotz S, Balakrishnan N. *Continuous Univariate Distributions*, Vol. 1. 2nd ed. New York: Wiley; 1994. (Gamma and log-gamma moments.)
4. Box GEP. Some theorems on quadratic forms applied in the study of analysis of variance problems, I. *Ann Math Stat.* 1954;25(2):290-302. doi:10.1214/aoms/1177728786.
5. Satterthwaite FE. An approximate distribution of estimates of variance components. *Biometrics Bull.* 1946;2(6):110-114. doi:10.2307/3002019.
6. Greenhouse SW, Geisser S. On methods in the analysis of profile data. *Psychometrika.* 1959;24(2):95-112. doi:10.1007/BF02289823.

7. Hilbe JM. Negative Binomial Regression. 2nd ed. Cambridge: Cambridge University Press; 2011.
8. Shi J, Luo D, Wan X, et al. Detecting the skewness of data from the five-number summary and its application in meta-analysis. Stat Methods Med Res. 2023;32(7):1338-1360.
doi:10.1177/09622802231172043.
9. Dwyer WJ. Comparable-Strata Designer software (M-series). Zenodo. Concept DOI
10.5281/zenodo.20709963.

Appendix S3. The log-domain distributional atlas — summary, with the full development archived in the deposited derivations record

The two-coordinate screen of this paper — the log-normal value $\sigma^2 = \ln(1 + CV^2)$ and its gamma companion $\psi'(1/CV^2)$ — is the first corner of a larger construction. Because the coefficient of variation pins the shape of any one-parameter positive scale family, every log-scale cumulant becomes a function of the coefficient of variation, so each family traces a curve in the plane of log-skewness and log-excess-kurtosis, and a sample's observed log-cumulants triage which family fits. Several families carry a shape-independent log-skewness signature — 0 for the log-normal and log-logistic, -1.14 for every Weibull, $+2$ for every Pareto tail — so one additional log-moment is a near-categorical discriminator; the deposited classifier model_atlas() ranks the families from the observed coefficient of variation and log-cumulants and routes each to its matched rescue. Because the log-scale fingerprint survives where the variance does not, the reading extends continuously into the heavy-tailed regime the coefficient of variation cannot see.

This log-domain distributional atlas is developed in full — the family loci and $CV \rightarrow \text{Var}(\ln X)$ bridges, the generalized-gamma encompassing frame, the classifier, six public-data worked examples, and the figures — in a companion paper by the author, “The Log-Domain Distributional Atlas: Triaging Skewed Positive Outcomes to Their Optimal Family and Rescue from Summary Statistics.” The complete derivation and the verification script `verify_distribution_atlas.py`, which regenerates every closed form, are archived in the deposited derivations record (Zenodo, concept DOI 10.5281/zenodo.20710253). The special-function and log-cumulant identities are classical; the contribution is their unification as one coefficient-of-variation-and-log-cumulant atlas. No new mathematics is claimed here.