Symmetric Coefficients of Variation for Binomial Estimates October 2026

Connor Doherty, Justin McIllece
Office of Employment and Unemployment Statistics, U.S. Bureau of Labor Statistics, 4600 Silver Hill Road, Washington, DC 20212-0002

Abstract
Dating back decades, the U.S. Bureau of Labor Statistics (BLS) determines the publishability of Current Population Survey (CPS) labor force statistics by the weighted size of the estimation base. Recently, the effectiveness of this publication rule has attenuated as response rates have declined from around 95 percent to less than 70 percent, leading the BLS to seek more robust data quality standards based on coefficients of variation (CVs), as recommended in U.S. Census Bureau Statistical Quality Standards. However, CVs for binomial data, like labor force participation or unemployment rates, are inherently asymmetric, as the CV of a rate p and its complement $1-p$ can be quite different despite having equal variances. In this paper, we propose the geometric mean of CV(p) and CV($1-p$) as a symmetric coefficient of variation for binomial estimates, resulting in publication standards that only depend on design effects and response counts, and demonstrate its relative stability compared to CV(p).

Key Words: Current Population Survey, coefficient of variation, symmetric CV, binomial, geometric mean, design effect

1. Introduction

Coefficients of variation (CVs) are useful in quantifying the reliability of an estimate. Some government agencies, such as the U.S. Census Bureau, utilize CVs as the basis for a portion of their statistical quality standards (U.S. Census Bureau, 2026). The U.S. Bureau of Labor Statistics (BLS) is interested in developing similar standards for determining publishability for their Current Population Survey (CPS) estimates. However, Kish (1995) warns that caution must be taken when working with CVs: They can be large and unstable for estimates with means near 0, and CVs for binomial estimates will differ for p and $1-p$. We propose a symmetric CV that remedies these issues, is robust to changes in level, and is simple to calculate. In this paper, we will use the CPS to motivate the development of the symmetric CV, compare the symmetric CV to other CV variations, and examine symmetric CVs applied to CPS estimates.

1.1 CPS Information

The CPS is a monthly household survey run jointly by the BLS and the U.S. Census Bureau. Each month, tens of thousands of labor force estimates are produced, but some of these are too volatile for official publication. Currently, the BLS suppresses publication of CPS rates based on the weighted size of the estimation base: the denominator of the rate. No level estimates are suppressed. The current publication standard for monthly rates is a minimum weighted estimation base of 75,000 persons. Internal research from the 1990s found that this estimation base requirement corresponded to a minimum of about 35 to 40 respondents at the time and was deemed sufficient reliability for publication. However, CPS weights have increased since the minimum base standard was created, in part due to decreasing response rates seen across many U.S. federal surveys (U.S. Bureau of Labor Statistics, 2025). Response rates historically stood around 95 percent but have fallen over the past two decades to under 70 percent today, seen below. This, among other factors, has caused the criterion to weaken over time, and the same guidelines now correspond to a minimum response count of about 20 individuals, necessitating an updated publishability criterion for CPS.

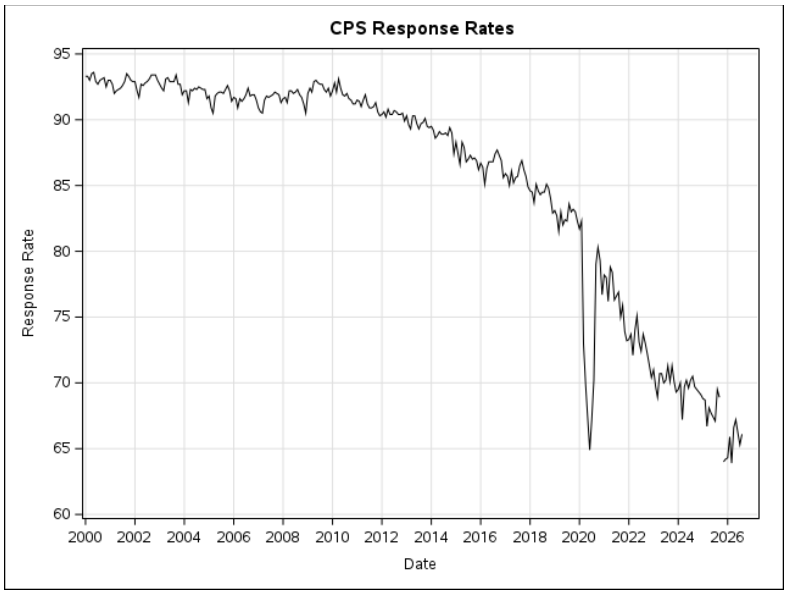


Figure 1: Plot of CPS response rates over time from January 2000 to October 2026. Note that data were not collected in October 2025 due to a lapse in appropriations, creating a break in series.

The CPS relies on a complex sample design: a state-based, multistage stratified design with one primary sampling unit per stratum and clustered sampling of households. This induces a design effect on CPS estimates. The design effect measures how the complex sampling design affects the variance of an estimate, compared to a simple random sample of the same size. CPS design effects are assumed constant for a given series and fall between 0.17 and 2.51, as seen in the histogram below[1]. The design effect is used in the calculation of CPS standard errors, and thus will be used in the calculation of the symmetric CV. For more details on the CPS design and variance estimation see Technical Paper 77 (2019).

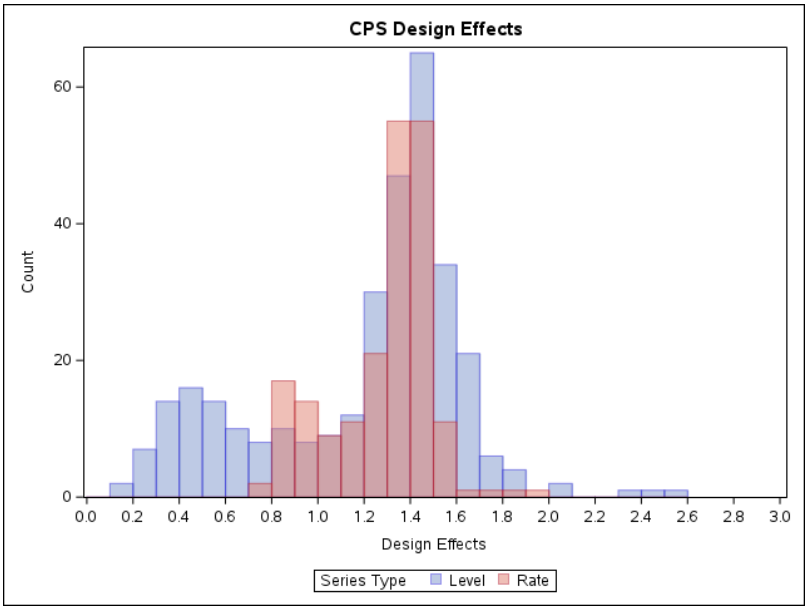


Figure 2: Histogram of design effects for the 521 non-veteran CPS series contained in *The Employment Situation* report. The design effects for level series are in blue and the rate series' design effects are in red.

[1] The design effects calculated here account for the poststratification, or "benchmarking," of weights to external population controls, significantly reducing the variances of the numerator for series highly correlated with the controls, as shown in Table 2-4.3 of Technical Paper 77 (2019). This is why numerous CPS design effects are well below 1.0.

2. Coefficients of Variation

A coefficient of variation is a standardized measure of variance, calculated as the standard deviation divided by the mean:

$$CV = \frac{\sigma}{\mu}$$

CVs are useful as they are a relative measure, not an absolute one. By dividing by the mean, units of measurement are cancelled, allowing for comparisons of variability (Kish, 1995) or for evaluating reliability of data (U.S. Census Bureau, 2026). Since the CPS is a joint survey between BLS and Census, we began with the *U.S. Census Bureau Statistical Quality Standards* (2026) as the guide for our publishability criteria, in which requirement F1-7 states that "serious data quality issues related to sampling error occur when the estimated coefficients of variation (CV) for the majority of the key estimates are larger than 30 percent." We desire a similar criterion that determines publishability of a series based on a single CV measure.

By changing the mean that is used, many variations on the standard CV can be calculated. A non-exhaustive list of alternative means are the arithmetic, geometric, and harmonic means. These alternative means may have properties that are desirable for use as a publishability criterion. However, many of them run into the same issues. One such issue is that CVs are, in general, asymmetric for a binomial rate p and its complement $1 - p$ (Kish, 1995). The standard errors for these values are equal, but the means are not necessarily the same. This can be a problem for rates very close to zero or one. For example, in the standard CV, a rate that is close to one will have a CV that approaches the standard error as the rate approaches one. At the same time, the CV for the complement approaches infinity. If CVs are used to determine publishability of estimates, this scenario would indicate that the rate p is reliable and may be published, but the complement is not. This is unintuitive, as the rates can be calculated from one another. Additionally, rates close to zero or one also result in CVs that are unstable and sensitive to changes in level.

2.1 Symmetric CVs

To remedy the above issues, we propose a symmetric CV. To ensure symmetry of CVs for a rate and its complement, the geometric mean is used.

$$\mu = \sqrt{p(1-p)}$$

The binomial standard error for a complex survey design is described in Technical Paper 77 (2019),

$$\sigma = \frac{\sqrt{p(1-p)}}{\sqrt{\frac{n}{D_{eff}}}} = \sqrt{\frac{D_{eff}p(1-p)}{n}}$$

where n is the response count and D_{eff} is the design effect. Using these in the calculation of a CV yields

$$CV_{symm} = \frac{\sigma}{\mu} = \sqrt{\frac{D_{eff}p(1-p)}{np(1-p)}} = \sqrt{\frac{D_{eff}}{n}}$$

We call this the symmetric CV. By construction, the symmetric CV is equal for a rate p and its complement $1 - p$. The resulting CV also is only a function of the design effect and response count of a given estimate, not the survey estimate, making it robust to changes in level. Since the design effect of a series is assumed to be constant (McIllece, 2016), the only variability in this CV over time comes from fluctuating response counts. This same formulation can also be obtained by taking the geometric mean of the standard CVs of p and $1 - p$.

Additionally, the inverse of the radicand, $\frac{n}{D_{eff}}$, is known as the effective sample size in complex survey design. This represents the sample size needed for a simple random sample to obtain equal variance given by the complex sample. Thus, publishability criteria utilizing the symmetric CV are based on a "true" sample size unobscured by sample design, a rather appealing property making for interpretable criteria.

There are two potential issues that we have noted for the proposed symmetric CV. First, highly seasonal series can still see large fluctuations in the CV over time. Since the symmetric CV is only a function of the response count over time, months where the response count is low will see the CV increase, with the change reflecting the magnitude of seasonality in the series. This is intentional, as the data quality does decrease when response count decreases, and the symmetric CV reflects that. However, this can pose an issue when using CVs for publishability criteria. It is desired that a series remains publishable throughout an entire year and not exceed a selected CV threshold for publication during only part of the year. We believe that this issue can be remedied using the annual average of the symmetric CV. For the purposes of the CPS, publishability of series would be determined once annually, at the start of the calendar year. Using the annual average of the CV means the decrease in data quality from seasonality is not ignored, and the decision is simpler due to using a single value, a desirable property when determining publishability for tens of thousands of series annually.

The second issue is that while symmetric CVs are truly symmetric for rate estimates, they are not necessarily symmetric for level estimates. The symmetric CV was designed primarily for rates, but they can still be applied to levels with this caveat in mind. For the purposes of the CPS, this is not a concern, as only rates are suppressed, not levels. This is not expected to change when introducing new publishability criteria.

Note that these two issues are present for other types of CVs as well and are not unique to symmetric CVs.

There is an important distinction between calculating symmetric CVs for levels and rates: the two types of series use different response counts. For levels, one would simply use the response count for that series, but for rates, one would use the response count for the base of the rate. For example, the unemployment rate is calculated as the unemployment level divided by the civilian labor force level. The symmetric CV for the unemployment rate would use the response count for the civilian labor force level, while the unemployment level's symmetric CV would be calculated using the response count for the unemployment level. Another example that may be unintuitive is the symmetric CV for the total employment-population ratio, calculated as the employment level divided by the population level. This series is a rate and would use the total response count for the survey, as this corresponds to the population estimate in the denominator. This distinction means that levels and rates for the same series have different CV magnitudes. Thus, different thresholds for publication are needed for level series and rate series.

3. Results

In this section we apply the symmetric CV to CPS series. We analyzed 521 monthly series from January 2003 to April 2025. Three topics were of interest: comparing symmetric CVs with other types of CVs, determining CV threshold values for suppression, and the stability and evolution of symmetric CVs over time.

3.1 Comparison to Other CVs

Below is a graph comparing the symmetric CV for the national unemployment rate to the standard CV, harmonic CV, and arithmetic CV, which use the standard, harmonic, and arithmetic means, respectively. The standard CV and harmonic CV are quite noisy and sensitive to changes in level. They also dropped significantly during the COVID-19 pandemic, indicating an increase in data quality during a time when base response counts were far lower than usual: an unintuitive result. The volatile nature of these CVs makes them undesirable for use in publishability criteria. The

arithmetic CV sees similar volatility, but to a lesser extent. This CV does correctly indicate a decrease in reliability during the COVID-19 pandemic, however. These three CVs are a function of p, either in the mean or in the standard error, explaining part of their volatility.

In contrast, the symmetric CV is extremely stable, with the only fluctuations over time coming from changes in observed response counts. The decrease in CPS response rates, leading to decreased response counts, can be seen in the slight increase of this CV over time, which is particularly notable after 2020. Like the arithmetic CV, the symmetric CV also indicates decreased reliability during the COVID-19 pandemic, due to a decrease in base response counts. We believe this stability and reliance on observed response counts make the symmetric CV appealing to use for publishability criteria.

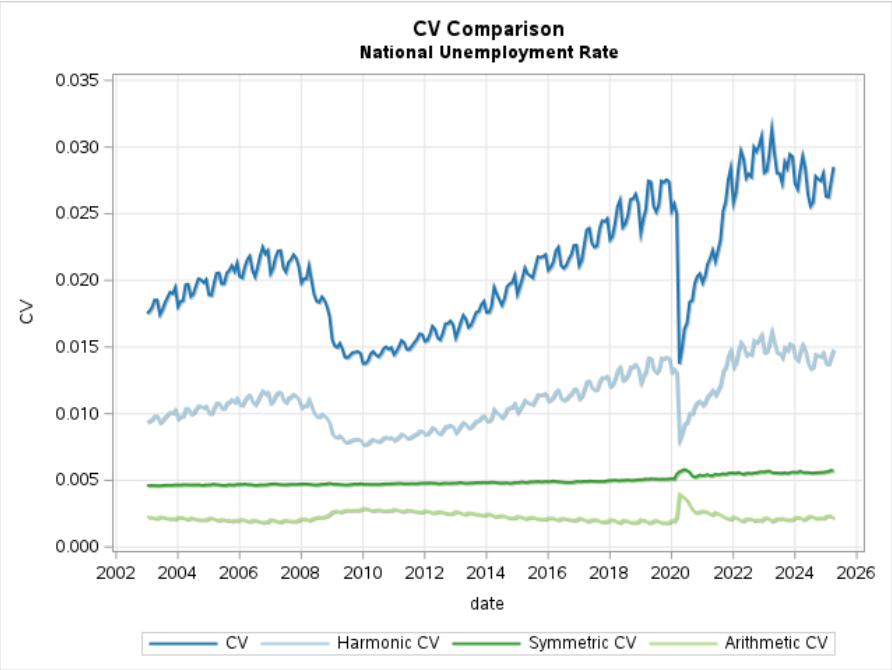


Figure 3: Comparison of four different CVs for the national unemployment rate. The symmetric CV is notably more stable than the other CVs.

3.1 Symmetric CV Thresholds

As stated before, different symmetric CV thresholds are needed for level series and rate series due to the difference in magnitude. Below are our current proposed thresholds for each type of series. Note that research and implementation discussions are still ongoing, and these do not represent the final CPS publishability criteria. While the CPS does not currently suppress levels, an acceptable CV threshold for level series was found should there be interest in level suppression.

We found that a symmetric CV threshold of 0.30 is suitable for level series, which corresponds to Census's statistical quality standards for key estimates, and a threshold of 0.08 is suitable for rate series. These values capture the noisy series with very low response counts, with some of the captured level series having less than 10 respondents each month, such as employed unpaid family workers in agriculture. This selection was performed manually, by taking note of series that are known to have problems with response counts or reliability. In total, four of the 322 examined level series exceed this CV threshold of 0.30 and two of 199 examined rate series exceed the CV threshold of 0.08.

These thresholds can be converted to minimum effective sample sizes by squaring the inverse of the symmetric CV. For levels, the corresponding minimum effective sample size is 11.1. For rates, the effective sample size is 156.25 for the base. Since the design effect varies for each series, the observed sample size that corresponds to these effective sample sizes also varies, but they will generally be larger as CPS design effects tend to be greater than one, with some exceptions as noted in footnote one.

3.2 Symmetric CVs Over Time

In this section we examine symmetric CVs of individual series over time. As seen before, symmetric CVs are significantly less variable than other formulations of CVs. They do still fluctuate over time, due to inconsistent response counts from month to month. This is less prominent for large, highly aggregated series, such as the monthly national unemployment rate, seen below. The annual averages of the symmetric CVs are plotted in green. If implemented, these are likely what the CPS publishability criteria would be based on, to simplify decision-making and avoid potential seasonality issues.

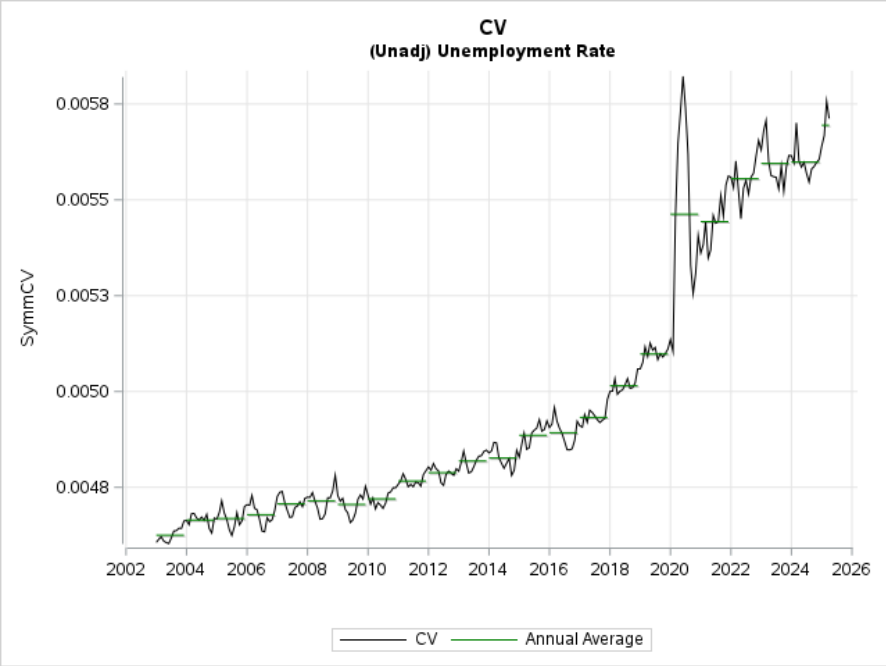


Figure 4: Symmetric CV for the monthly national unemployment rate from January 2003 to April 2025.

The symmetric CV for the unemployment rate of Hispanic or Latino teens is plotted below. This is a much smaller, noisier series than the national unemployment rate, seen in the larger and more varied CV values, but these are still more stable than for other CV types. This series also has clear seasonality, seen by the CV dropping significantly each summer as unemployment response counts increase due to the end of the school year. This is where using the annual average symmetric CV is particularly useful, and these are stable over time, hovering around 0.055 for most of the series.

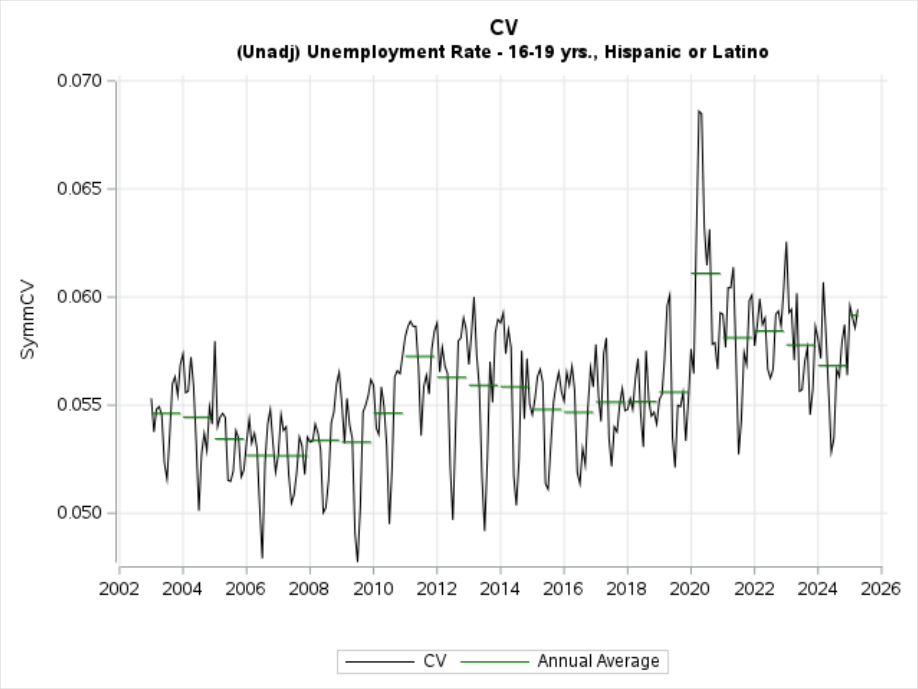


Figure 5: Symmetric CV for the monthly unemployment rate of Hispanic or Latino teens from January 2003 to April 2025.

An example of a level series that would not be fit for publication under the proposed CV thresholds is seen below. Unpaid family workers in agriculture is a series that regularly has fewer than 10 respondents, which can be seen by the high symmetric CV values and high volatility.

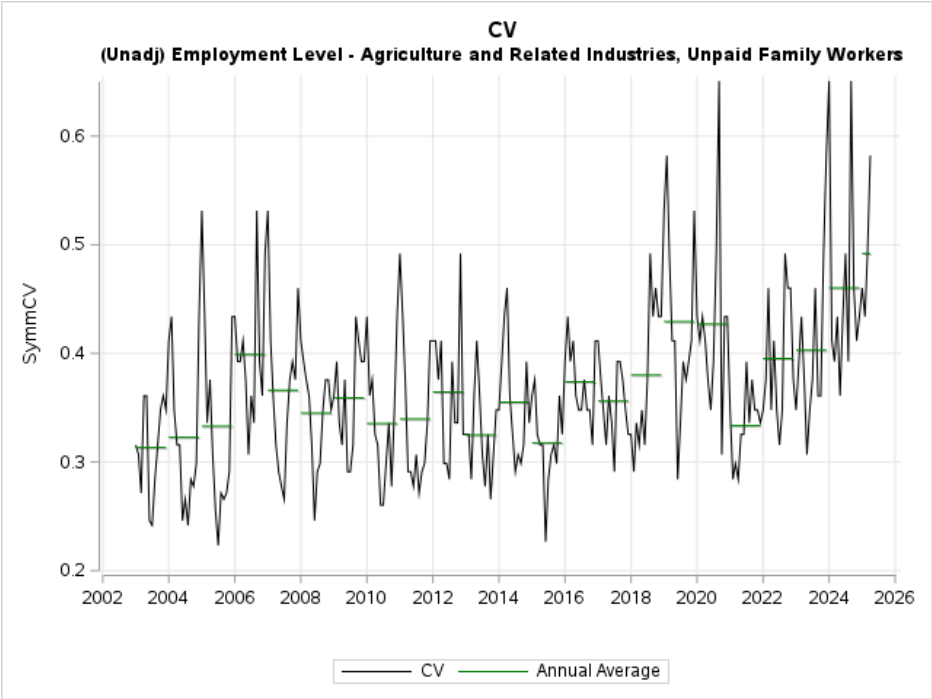


Figure 6: Symmetric CV for the monthly employment level of unpaid family workers in agriculture from January 2003 to April 2025.

Below is another level series that exceeds the proposed CV threshold in recent years. The response counts are plotted alongside the symmetric CV, and the inverse relationship between the two is clear.

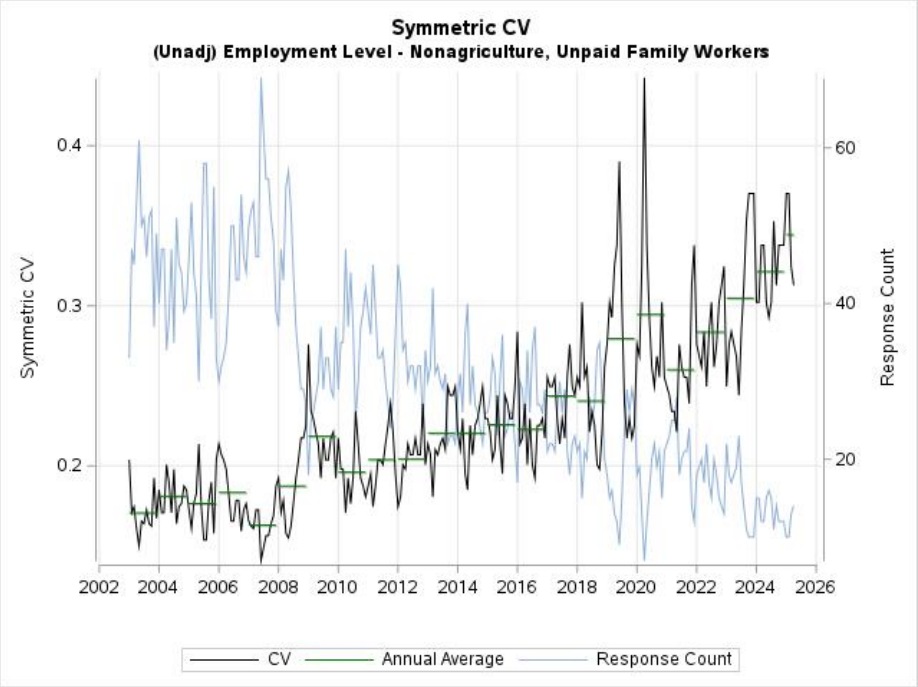


Figure 7: Symmetric CV and response counts for the monthly employment level of unpaid family workers in nonagriculture from January 2003 to April 2025.

4. Conclusion

We have proposed a symmetric CV for use as a reliability measure. This CV is symmetric for binomial rates, robust to changes in level, and simple to calculate. It was found to be more stable than other CV variations, since over time it is only a function of observed response counts, assuming a static design effect. The symmetric CV also has an intuitive calculation, being the inverse of the square root of the effective sample size of each series. The stability and interpretability make the proposed symmetric CV attractive for use in publishability criteria.

Internal implementation discussions will continue, and the proposed CV threshold values for CPS may be revised. Future research may involve seeking a more elegant method for dealing with seasonality or a more automated method for determining threshold values.

References

Kish, L. (1995). *Survey Sampling*. Wiley.

McIllece, J.J. (2016). "Calculating Generalized Variance Functions with a Single-Series Model in the Current Population Survey," https://www.bls.gov/osmr/research-papers/2016/pdf/st160130.pdf.

U.S. Bureau of Labor Statistics (2025). CPS Response Rates. https://www.bls.gov/cps/methods/response_rates.htm.

U.S. Census Bureau (2019). Design and Methodology, Current Population Survey, Technical Paper 77.

U.S. Census Bureau (2026) U.S. Census Bureau Statistical Quality Standards. https://www2.census.gov/about/policies/quality/quality-standards.pdf.