coefconv

Comprehensive Marginal Effects for Stata

Includes coefconv_plot visualization companion

Version 1.0.0
April 2026
Stata 14 or higher
Statistical Software Components (SSC)

Dr Noman Arshed
Senior Lecturer, Department of Business Analytics
Sunway Business School, Sunway University
nouman.arshed@gmail.com

Contents

1 Introduction

Regression coefficients are often reported as raw slopes without further transformation. A raw slope answers 'how much does Y change for a one-unit increase in X?' — but one unit may be arbitrary, comparisons across predictors in different units are impossible, and practical significance depends entirely on context rarely stated.

coefconv solves this by computing 26 re-expressions of every slope coefficient, organized into seven interpretation families. Each family answers a different substantive question — from 'how large is this in standard-deviation units?' (Family 1) to 'what would happen if the average person moved from the 25th to 75th percentile of this variable?' (Family 7). All inputs are drawn automatically from Stata's e() results and restricted to the estimation sample.

The companion program **coefconv_plot** produces three simultaneously open named graphs that assess which coefficients are both statistically significant and practically meaningful: a standardized-slope forest plot with Cohen (1988) benchmarks, a Pratt R^2 decomposition bar chart, and a per-variable discrete-effects ladder chart.

Guiding principle

A coefficient is *meaningful* on two independent dimensions: statistical (the CI does not cross zero) and practical (the effect is large enough to matter in context). coefconv addresses both simultaneously.

Notation

Symbol	Definition
β	Raw OLS slope coefficient
β*	Fully standardized slope = β × σX / σY
X■ , Y■	Sample means, restricted to e(sample)
σX , σY	Sample standard deviations of X and Y
r(Y,X)	Zero-order Pearson correlation
X_pq	q-th percentile of X in estimation sample
IQR	Interquartile range = X_p75 – X_p25
ΔX / ΔY	Change in X or predicted change in Y = β × ΔX

grate	Growth rate (default 0.01 = 1%)

2 Installation

From SSC (recommended)

```
. ssc install coefconv, replace
```

Manual installation

Place coefconv.ado, coefconv_plot.ado, coefconv.sthlp, and coefconv_plot.sthlp in your personal adopath:

OS	Path
Windows	C:\ado\personal\
Mac / Linux	~/ado/personal/

Or add dynamically:

```
. adopath + "path/to/folder"
```

Verify:

```
. which coefconv
. help coefconv
```

Dependency

coefconv_plot calls coefconv internally. Both .ado files must be in the same adopath. No other external packages are required.

3 coefconv — Syntax and Options

Syntax

```
. coefconv [, grate(#) quantiles(numlist) delta(numlist)
saving(filename[, replace]) notable format(fmt)]
```

Run immediately after a regression command before any other estimation command overwrites e().

Options

Option	Description
`grate(#)`	Growth rate for default ΔX = grate × X■ in Family 7. Default: 0.01 (1%).
`quantiles(numlist)`	Integer percentiles to add to default {10 25 50 75 90} for Family 6 quantile displacements.
`delta(numlist)`	Custom ΔX values for every predictor in Family 7. Each value generates one additional row.
`saving(file[,replace])`	Save wide results dataset: one row per predictor, one column per effect type, all with variable labels.
`notable`	Suppress all display. Use when accessing r() scalars programmatically.
`format(fmt)`	Stata numeric format for displayed values. Default: %12.6f.

How coefconv extracts its inputs

Source	What is extracted
e(b)	Coefficient vector; column names give predictor names
e(depvar)	Name of the dependent variable
e(sample)	Estimation sample; all stats restricted to it
e(r2)	R-squared for Pratt decomposition heading
e(cmd)	Estimator name; triggers non-linear warning
correlate	Zero-order correlations r(Y,X) for Pratt
summarize	Mean, SD, min, max for every predictor

_pctile	Quantiles for Families 6 and 7

4 Effect Families Reference

All numerical examples use: regress hourly_wage education experience with β = 1.20 (education), X■ = 14.00, σX = 3.00, Y■ = 18.00, σY = 6.00, r(Y,X) = 0.55.

4.1 Family 1 — Raw and Standardized Slopes

Raw slope (β)

`β = 1.20`

One additional year of education is associated with $1.20/hour higher wages.

Fully standardized (β*)

`β* = 1.20 × 3.00/6.00 = 0.60`

A 1-SD increase in education is associated with 0.60 SDs higher wages. By Cohen (1988) criteria, 0.60 is a large effect.

X-standardized (β × σX)

`1.20 × 3.00 = 3.60`

A 1-SD increase (3 extra years) raises wages by $3.60/hour. Y stays in original units.

Y-standardized (β / σY)

`1.20 / 6.00 = 0.20`

Each extra year of education raises wages by 0.20 standard deviations. X stays in original units.

4.2 Family 2 — Elasticity and Semi-Elasticity

These express effects as percentage changes and are scale-free.

Point elasticity (ε = β × X■ / Y■)

`ε = 1.20 × 14.00 / 18.00 = 0.933`

A 1% increase in education from its mean (+0.14 years) is associated with a 0.933% increase in wages. Since ε < 1, education is inelastic at the means.

X-semi-elasticity (β × X■)

`1.20 × 14.00 = 16.80 [divide by 100 for 1% ΔX → +$0.168/hr]`

Arises naturally in log-linear models. Divide by 100 to get ΔY per 1% increase in X from its mean.

Y-semi-elasticity ((β/Y■) × 100)

```
(1.20 / 18.00) × 100 = 6.667%
```

Each extra year of education raises wages by 6.67% of the mean wage. Common in linear-log models where Y is a level.

4.3 Family 3 — Basis-Point and Percentage-Point Effects

Rescaled versions of β for variables measured as rates or proportions.

Basis-point effect (β ÷ 10,000)

```
1.20 / 10,000 = 0.000120
```

Effect of a 1-basis-point (0.01 pp) increase in X. Most meaningful when β is large (e.g. β = 500 on an interest rate → 0.05 per bp).

Percentage-point effect (β ÷ 100)

```
1.20 / 100 = 0.0120
```

Effect of a 1-percentage-point increase. Standard for unemployment rate, inflation, tax rates.

Per-1%-of-X effect (β × X■ ÷ 100)

```
1.20 × 14.00 / 100 = 0.168
```

Effect of a 1% increase in X from its mean. A natural 'small change' benchmark when X■ >> 1.

4.4 Family 4 — Relative and Proportional Effects

Proportional ME (β / Y■)

```
1.20 / 18.00 = 0.0667
```

A 1-unit increase in X raises Y by 6.67% of the mean outcome. Useful for comparing effects across outcomes on different scales.

% of mean-Y ME ((β/Y■) × 100)

```
(1.20 / 18.00) × 100 = 6.667%
```

Percentage version of the above. Easier to communicate to non-technical audiences.

4.5 Family 5 — Variance and Importance Measures

These decompose R^2 to quantify each predictor's unique contribution, accounting for correlations among predictors.

Squared std. coefficient (β^{*2})

```
0.60² = 0.36
```

Approximates unique variance share under orthogonality. Overestimates with correlated predictors; prefer Pratt's measure.

Product measure (β × r_XY)

```
1.20 × 0.55 = 0.660
```

Unstandardized Pratt numerator. Sums to R^2. Negative values identify suppressor variables.

Pratt numerator (β* × r_XY)

```
0.60 × 0.55 = 0.330
```

Divide by $\Sigma(\beta^* \times r)$ across predictors to get Pratt%. If $R^2 = 0.42$, then Pratt% = 0.330/0.420 = 78.6%.

Pratt % of R^2 (summary table)

```
78.6% in this example
```

Printed after all per-variable blocks. All percentages sum to 100% of R^2. A negative Pratt% confirms a suppressor variable.

4.6 Family 6 — Quantile Displacement Effects

For a linear model the slope is constant, so Family 6 asks: if X were at a specific quantile rather than the median, how different would the predicted Y be? This maps the slope onto realistic distributional distances.

p50 → p10

`β × (10–14) = –4.80`

Workers at p10 of education earn $4.80/hr less than median workers.

p50 → p25

`β × (12–14) = –2.40`

Workers at p25 earn $2.40/hr less than median workers.

p50 → p75

`β × (16–14) = +2.40`

Workers at p75 earn $2.40/hr more than median workers.

p50 → p90

`β × (18–14) = +4.80`

Workers at p90 earn $4.80/hr more than median workers.

4.7 Family 7 — Discrete Change Effects

Growth-rate ΔX (β × grate × X■)

`1.20 × 0.01 × 14 = 0.168`

A 1% growth in education from its mean raises wages by $0.168/hr. Override with grate(0.05) for a 5% growth scenario.

IQR effect (β × IQR)

`1.20 × (16–12) = 4.80`

Moving from Q25 to Q75 raises wages by $4.80/hr. Most robust to outliers; preferred for policy communication.

Full-range effect (β × (X_max–X_min))

`1.20 × (20–8) = 14.40`

Maximum possible effect across observed data. Sensitive to outliers; useful as an upper ceiling.

±1 SD (β × σX)

`1.20 × 3.00 = 3.60`

Standard social-science effect-size unit. Spans approx. p16–p84. Numerically identical to the X-standardized slope in Family 1.

±2 SD (β × 2σX)

```
1.20 × 6.00 = 7.20
```

Spans approx. p2–p98. Recommended by Gelman & Hill (2007) for comparing continuous predictors to binary ones.

Custom ΔX (delta() option)

```
. coefconv, delta(2 5) → β×2 = 2.40; β×5 = 6.00
```

Specific externally defined changes. For an intervention of exactly N units, delta(N) gives the direct predicted outcome.

5 Complete Effects Reference Table

All 26 effects computed by coefconv, with formula, unit of interpretation, and effect family.

#	Effect name	Formula	Unit	Fam.
1	Raw slope	β	Y-units per X-unit	F1
2	Fully standardized	β×σX/σY	SDs of Y per SD of X	F1
3	X-standardized	β×σX	Y-units per SD of X	F1
4	Y-standardized	β/σY	SDs of Y per X-unit	F1
5	Point elasticity	β×X■/Y■	% Y per % X	F2
6	X-semi-elasticity	β×X■	Y-units per % X	F2
7	Y-semi-elasticity	(β/Y■)×100	% Y per X-unit	F2
8	Basis-point effect	β/10000	Y-units per bp	F3
9	Pct-point effect	β/100	Y-units per pp	F3
10	Per-1%-of-X	β×X■/100	Y-units per 1% X	F3
11	Proportional ME	β/Y■	Fraction of Y■	F4
12	% of mean-Y ME	(β/Y■)×100	% of Y■ per X-unit	F4
13	Squared std. coef.	β^{*2}	Approx. var. share	F5
14	Product measure	β×r_XY	R^2 decomp.	F5
15	Pratt numerator	β*×r_XY	Pratt raw	F5
16	Pratt % of R^2	β*r / Σ(β*r)	% of R^2	F5
17	p50 → p10	β×(X_p10−X_p50)	Y-units	F6

18	p50 → p25	β×(X_p25−X_p50)	Y-units	F6
19	p50 → p75	β×(X_p75−X_p50)	Y-units	F6
20	p50 → p90	β×(X_p90−X_p50)	Y-units	F6
21	Growth-rate ΔX	β×grate×X■	Y-units	F7
22	IQR effect	β×IQR	Y-units	F7
23	Full-range	β×(X_max−X_min)	Y-units	F7
24	±1 SD	β×σX	Y-units	F7
25	±2 SD	β×2σX	Y-units	F7
26	Custom ΔX	β×ΔX_user	Y-units	F7

F1–F7 = Effect Family. All formulas evaluated at estimation-sample means/quantiles. Y■ = 0 causes Families 2 and 4 to return missing.

6 Interpretation Guide

A coefficient is **meaningful** on two independent dimensions. The table below maps each dimension to the relevant coefconv output and recommended thresholds.

Dimension	Tool	Threshold / Guidance
Statistical significance	Family 1 β^* CI (Graph 1 forest plot)	CI must not cross zero. Significance rules out a zero effect at the chosen confidence level. It does NOT imply practical importance.
Effect size (standardized)	Family 1 β^* (Graph 1 forest plot)	Cohen (1988): Small ≥ 0.10 | Medium ≥ 0.30 | Large ≥ 0.50. Domain-specific thresholds may differ.
Relative importance	Family 5 Pratt% (Graph 2 bar chart)	Pratt% < 5% = small unique contribution even if β is significant. Negative Pratt% = suppressor variable; retain it.
Practical magnitude	Family 7 IQR effect (Graph 3 ladder)	Compare IQR ΔY to Y■ or a domain minimum important difference (e.g. 2 points on a 100-point clinical scale).

Pratt's decomposition and suppressors

Pratt's (1987) method decomposes R^2 into signed additive contributions: $\text{Pratt}_j = (\beta^*_j \times r_j) / \Sigma(\beta^*_k \times r_k)$. All Pratt percentages sum to exactly 100% of R^2. A **negative** Pratt index indicates a **suppressor variable**: a predictor whose bivariate correlation with Y is weaker or opposite in sign relative to its partial relationship, because it absorbs error variance from other predictors. Suppressor variables improve model fit and must not be removed based on a negative Pratt index alone.

Elasticity conventions

coefconv reports a **point elasticity** evaluated at sample means. Do not confuse the three related measures:

Measure	Formula	Both proportional?	Y in levels	X in levels
Full elasticity	β×X■/Y■	Yes	—	—
X-semi-elasticity	β×X■	No	Y in levels	X proportional Δ

Y-semi-elasticity	(β/Y■)×100	No	Y proportional Δ	X in units

7 coefconv_plot — Syntax and Graphs

Syntax

```
. coefconv_plot [, grate(#) level(#) scheme(str)
saving(stub[, replace])
noSTD noPRATT noEFFects]
```

Options

Option	Description
`grate(#)`	Growth rate passed to coefconv. Default: 0.01.
`level(#)`	CI confidence level for forest plot. Default: 95.
`scheme(str)`	Stata graph scheme. Example: scheme(s1color).
`saving(stub[,replace])`	Saves stub_std.gph, stub_pratt.gph, and stub_eff_[var].gph per predictor.
`noSTD`	Skip Graph 1 (standardized slopes forest plot).
`noPRATT`	Skip Graph 2 (Pratt importance bar chart).
`noEFFects`	Skip Graph 3 (per-variable discrete effects charts).

Named graphs (all open simultaneously)

Each graph receives a unique name() so all remain open at once in Stata's Graph window. Use the dropdown or type `graph display [name]` to switch.

Graph	Name	Command to bring to front
Standardized slopes forest	ccv_std	graph display ccv_std
Pratt importance chart	ccv_pratt	graph display ccv_pratt
Discrete effects: mpg	ccv_eff_mpg	graph display ccv_eff_mpg
Discrete effects: weight	ccv_eff_weight	graph display ccv_eff_weight

Graph 1 — Standardized Slopes Forest Plot

Name: ccv_std

Displays β^* for each predictor as a diamond marker with a $\pm z \times SE(\beta^*)$ confidence interval. **Navy** diamonds are significant at the chosen level; **gray** are not. Vertical dashed lines mark Cohen's small (0.20), medium (0.50), and large (0.80) benchmarks. A predictor with a large navy diamond past the 0.50 line is both statistically significant and practically important.

Graph 2 — Pratt Importance Bar Chart

Name: ccv_pratt

Horizontal bars showing each predictor's Pratt percentage of R^2, sorted ascending so the largest appears at the top. **Navy** bars are productive predictors; **cranberry** bars are suppressor variables (negative Pratt index). All bars sum to 100% of R^2.

Graph 3 — Discrete Effects Ladder

Name: ccv_eff_[varname]

One graph per predictor. Nine ΔY scenarios on a common axis in original Y units, sorted by $|\Delta Y|$ ascending. The bottom bar (growth-rate scenario) is the minimum realistic benchmark; the IQR bar is most policy-relevant; the full-range bar is the theoretical ceiling. **Navy** = positive ΔY; **cranberry** = negative ΔY.

8 Saved Results (r() scalars)

After running coefconv (with or without notable), the following scalars are accessible via r():

Scalar	Description
r(N)	Number of observations in e(sample)
r(r2)	R-squared from the estimated model
r(ymean)	Mean of the dependent variable
r(ysd)	SD of the dependent variable
r(pratt_tot)	Sum of all Pratt numerators (= R^2)
r(b_varname)	Raw slope β for each predictor
r(bstd_varname)	Standardized slope β* for each predictor
r(elas_varname)	Point elasticity at means for each predictor
r(ysemi_varname)	Y-semi-elasticity for each predictor
r(pratt_n_varname)	Pratt numerator for each predictor
r(pratt_pct_varname)	Pratt % of R^2 for each predictor

Example

```
. sysuse auto, clear
. regress price mpg weight foreign
. coefconv, notable
. display "Elasticity of mpg: " r(elas_mpg)
. display "Pratt% of weight: " r(pratt_pct_weight)
. display "Std. slope foreign: " r(bstd_foreign)
. display "Model R-squared: " r(r2)
```

9 Worked Examples

Example 1 — Basic use

```
. sysuse auto, clear
. regress price mpg weight foreign
. coefconv
```

Runs all 26 effect types for mpg, weight, and foreign. Ends with a Pratt summary table showing each predictor's share of the model's R^2 = 0.500.

Example 2 — Custom growth rate and extra quantiles

```
. coefconv, grate(0.05) quantiles(5 95)
```

Changes the Family 7 growth benchmark to 5% and adds p5 and p95 to the Family 6 quantile displacement block.

Example 3 — Custom delta-X scenarios

```
. coefconv, delta(500 2000)
```

Adds two 'Custom ΔX' rows in Family 7 for every predictor, showing $\beta \times 500$ and $\beta \times 2000$ directly in Y units.

Example 4 — Save wide results dataset

```
. coefconv, saving(coefconv_results, replace)
. use coefconv_results, clear
. list varname beta elasticity pratt_pct, noobs
```

Exports a dataset with one row per predictor and 26 effect columns, all carrying variable labels for codebook clarity.

Example 5 — Programmatic scalar access

```
. coefconv, notable
. display r(elas_mpg)
. display r(pratt_pct_weight)
```

Silent run; access any scalar directly. Useful in loops collecting effects across nested models.

Example 6 — Visualization — all three graphs

```
. coefconv_plot
```

Opens ccv_std, ccv_pratt, and one ccv_eff_[var] per predictor simultaneously. Switch between graphs with the Graph window dropdown.

Example 7 — Forest plot only, 90% CI

```
. coefconv_plot, nopratt noeffects level(90)
```

Produces only the standardized slopes forest plot with 90% CIs.

Example 8 — Save graphs for a paper

```
. coefconv_plot, saving(paper_graphs, replace)
```

Saves paper_graphs_std.gph, paper_graphs_pratt.gph, and paper_graphs_eff_[var].gph. Re-open later with graph use paper_graphs_std.

Example 9 — After ivregress

```
. ivregress 2sls price (mpg = gear_ratio) weight foreign
. coefconv, grate(0.03)
```

coefconv reads 2SLS coefficients and VCE from e(). SE(β*) in the forest plot reflects 2SLS standard errors.

Example 10 — After areg (absorbed fixed effects)

```
. sysuse nlsw88, clear
. areg wage age hours tenure, absorb(industry)
. coefconv_plot, noeffects
```

Works with areg. noeffects suppresses per-variable plots, showing only the forest plot and Pratt chart.

Example 11 — Loop: elasticities across models

```
. local models `" "mpg" "mpg weight" "mpg weight foreign" "'
. local col = 0
. foreach mvars of local models {
. local col = `col' + 1
. regress price `mvars'
. coefconv, notable
. display r(elas_mpg)
. }
```

Collects elasticities across three nested models. Extend with matrix storage for a publication-ready sensitivity table.

10 Compatibility and Limitations

Compatible estimators

coefconv works with any estimator that populates e(b), e(V), e(depvar), and e(sample):

Estimator	Notes
regress	Full support; all 26 effects and SEs valid
ivregress	Full support; SE(β*) reflects IV variance-covariance
areg	Full support including absorbed-FE models
xtreg, fe	Full support for within-estimator slopes
xtreg, re	Full support for GLS slopes
xtgls	Full support
logit/probit	Runs with warning; slopes ≠ marginal effects in non-linear models. Consider margins after estimation.
tobit	Runs with warning; slopes are latent-index effects

Known limitations

- Elasticities and proportional effects (Families 2 and 4) are undefined when Y■ = 0 and are returned as missing.
- Standardized slopes require $\sigma X > 0$ and $\sigma Y > 0$. Constant predictors (SD = 0) generate a warning and SD-based effects are set to missing.
- Factor variables (i.region, i.industry) are skipped in Families 2–7 because they cannot be summarized by a single mean and SD. Their raw β appears in Family 1.
- coefconv_plot requires coefconv to be installed in the same adopath.
- Pratt importance requires non-zero R^2. With $R^2 = 0$ all Pratt percentages are undefined.
- The Pratt decomposition can produce percentages above 100% for any single predictor when suppressors are present (the sum still equals 100%).

11 References

Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum Associates.

Gelman, A., & Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press.

Pratt, J. W. (1987). Dividing the indivisible: Using simple symmetry to partition variance explained. In T. Pukkila & S. Puntanen (Eds.), *Proceedings of the Second International Conference in Statistics* (pp. 245–260). University of Tampere.

Thomas, D. R., Hughes, E., & Zumbo, B. D. (1998). On variable importance in linear regression. *Social Indicators Research*, 45(1–3), 253–275.

Stata Corp. (2023). *Stata Base Reference Manual, Release 18*. Stata Press.

Contact: Dr Noman Arshed — nouman.arshed@gmail.com
Sunway Business School, Sunway University
Bug reports, suggestions, and contributions are welcome.