

A graphic comparison of the Fieller and Delta intervals for ratios of parameter estimates.

Joe Hirschberg and Jenny Lye
Economics, University of Melbourne.

September 2017

Introduction

Fieller's 1954 proposal for the use of an inverse test to construct confidence intervals (Cis) for the ratio of normally distributed statistics has been shown to be superior to the application of the Delta method in several applications. In this presentation, we demonstrate how a simple graphic exposition can be generated to illustrate the relationship between the Delta and the Fieller Cis. The advantage of the graphical presentation over the numeric result available in the *fieller* Stata command (Coveney 2004) is that it may indicate how the level of significance can be changed to result in finite upper and lower bounds. And to be used in cases where the Fieller results are not defined.

We define two normally distributed random variables as:

$$[\hat{\gamma}, \hat{\beta}] \sim N\{[\gamma, \beta], \Omega\}, \text{ where } \Omega = \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix}$$

Where the quantity of interest is the ratio defined as $\theta = \frac{\gamma}{\beta}$

Which has been estimated by $\hat{\theta} = \frac{\hat{\gamma}}{\hat{\beta}}$

The objective of this analysis is to determine the confidence interval for $\hat{\theta}$.

The estimation of the confidence interval can be done via the application of the Delta approximation method (i.e. the use of the Stata *nlcom* command) or alternatively they may employ the Fieller method for the approximation of the confidence interval.

Many papers have compared these approximations and found that they coincide for cases when the denominator variable is estimated with a low relative variance (i.e. $t_{\beta} = \frac{\hat{\beta}}{\hat{\sigma}_2}$ is large (> 3)). However, in other cases the Fieller has been shown to provide superior coverage (See Hirschberg and Lye 2010) for comparisons). Note that Fieller intervals are not forced to be symmetric as are the Delta intervals.

The Inverse Test for the ratio (Fieller)¹

One method for determining the value of the ratio is to rewrite the formula $\hat{\theta} = \frac{\hat{\gamma}}{\hat{\beta}}$ as $\hat{\gamma} - \hat{\theta}\hat{\beta} = 0$.

Thus we could use the values of $[\hat{\gamma}, \hat{\beta}]$ and try different values of $\theta = \theta_i$ to search for the value of θ_i where $\hat{\gamma} - \theta_*\hat{\beta} = 0$ is true then conclude that $\hat{\theta} = \theta_*$.

¹ See our paper Lye and Hirschberg (2012) for more examples of similar inverse tests.

The benefit of this approach is that the linear function $\hat{y} - \theta_i \hat{\beta}$ of normally distributed random variables is normally distributed.

While the ratio of normally distributed random variables $\hat{\theta} = \frac{\hat{y}}{\hat{\beta}}$, is

distributed via a Cauchy distribution that does not have finite moments. However, this does not preclude the construction of probability statements concerning $\hat{\theta}$.

We can test our candidate values of θ_i via the hypothesis defined by:

$$H_0 : \hat{\gamma} - \theta_i \hat{\beta} = 0$$

$$H_1 : \hat{\gamma} - \theta_i \hat{\beta} \neq 0$$

Under H_0 we can construct a test statistic defined by the ratio of the linear combination evaluated at θ_i :

$$\hat{T}^2 = \left(\hat{\gamma} - \theta_i \hat{\beta} \right)^2 / \widehat{\text{Var}} \left(\hat{\gamma} - \theta_i \hat{\beta} \right)$$

Would be distributed as the square of a standard normal or a Chi-square with 1 degree of freedom.

We “invert the problem” by determining the value(s) of θ_i where the test statistic is equal to the value of the square of standard normal for the level of α ($z_{\alpha/2}^2$):

$$z_{\alpha/2}^2 = (\hat{\gamma} - \theta_i \hat{\beta})^2 / \widehat{\text{Var}}(\hat{\gamma} - \theta_i \hat{\beta})$$

as the limits of the confidence interval.

We designate these limits as $\tilde{\theta}_0$. Using the estimate of

$$\hat{\Omega} = \begin{bmatrix} \hat{\sigma}_1^2 & \hat{\sigma}_{12} \\ \hat{\sigma}_{12} & \hat{\sigma}_2^2 \end{bmatrix}, \text{ this is equivalent to the quadratic in } \tilde{\theta}_0:$$

$$\left(\hat{\beta}^2 - \hat{\sigma}_2^2 z_{\alpha/2}^2\right) \tilde{\theta}_0^2 + \left(2\hat{\sigma}_{12} z_{\alpha/2}^2 - 2\hat{\gamma}\hat{\beta}\right) \tilde{\theta}_0 + \left(\hat{\gamma}^2 - \hat{\sigma}_1^2 z_{\alpha/2}^2\right) = 0$$

The two roots of this quadratic define the potential solutions for $\tilde{\theta}_0$ and provide the limits of the confidence interval.

The solution to this quadratic are defined by:

$$\{\tilde{\theta}_{0U}, \tilde{\theta}_{0L}\} = \frac{(\hat{\sigma}_{12} z_{\alpha/2}^2 - \hat{\beta}\hat{\gamma}) \pm z_{\alpha/2} \sqrt{\hat{\beta}^2 \hat{\sigma}_1^2 + \hat{\gamma}^2 \hat{\sigma}_2^2 + \hat{\sigma}_{12}^2 z_{\alpha/2}^2 - \hat{\sigma}_1^2 \hat{\sigma}_2^2 z_{\alpha/2}^2 - 2\hat{\beta}\hat{\gamma}\hat{\sigma}_{12}}}{(\hat{\beta}^2 - \hat{\sigma}_2^2 z_{\alpha/2}^2)}$$

Computing the solution to this quadratic equation is widely available (i.e. the *polyval* command).

However, the quadratic may not have real roots. Furthermore, the numbers generated may not provide much intuition as to how the interval may change with differing values of α .

The Line-plot Version of the Fieller

The alternative approach to computing the roots of the implied quadratic equation is the line-plot approach. This approach returns to the formulation:

$$z_{\alpha/2}^2 = \frac{(\hat{\gamma} - \tilde{\theta}_0 \hat{\beta})^2}{\widehat{\text{Var}}(\hat{\gamma} - \tilde{\theta}_0 \hat{\beta})}$$

After taking the square-root of both sides we have:

$$(\hat{\gamma} - \tilde{\theta}_0 \hat{\beta}) \pm z_{\alpha/2} \sqrt{\widehat{\text{Var}}(\hat{\gamma} - \tilde{\theta}_0 \hat{\beta})} = 0$$

Where $(\hat{\gamma} - \theta_i \hat{\beta}) \pm z_{\alpha/2} \sqrt{\widehat{\text{Var}}(\hat{\gamma} - \theta_i \hat{\beta})}$ defines the confidence interval of the linear combination $(\hat{\gamma} - \theta_i \hat{\beta})$ for different θ_i s.

Thus by plotting the line defined by $\hat{\gamma} - \theta_i \hat{\beta}$ over different values of θ_i around the value of the estimated ratio $\hat{\theta} = \frac{\hat{\gamma}}{\hat{\beta}}$ we can find the confidence bounds $\{\tilde{\theta}_{0U}, \tilde{\theta}_{0L}\}$ at the values of θ_i where the confidence interval of the line defined by:

$$(\hat{\gamma} - \theta_i \hat{\beta}) \pm z_{\alpha/2} \sqrt{\widehat{\text{Var}}(\hat{\gamma} - \theta_i \hat{\beta})}$$

cuts the zero-axis line.

An Example of the ratio of the means of a common variable.

The means of different size samples of a common variable is a special case where the covariance between the means (σ_{12}) is set to zero and the variances are equal ($\sigma_1^2 = \sigma_2^2$). Here we use the property of a least squares regression with two dummy variables and no intercept.

$$Y_t = \gamma D_{1t} + \beta D_{2t} + \varepsilon_t$$

This specification will estimate the means of Y for two groups when we assume the variance of Y is the same for both groups.

The estimate of the ratio of the means is defined by the ratio of the two estimated parameters $\hat{\theta} = \frac{\hat{\gamma}}{\hat{\beta}}$.

By using a regression we can employ the *margins* command, as well as the *nlcom* command to provide a comparison to the Delta method.

The Stata code below reads the simulated blood pressure data, creates two dummies by the variable sex and computes the two means via a regression.

This data has an indicator for sex and we first create two dummies for each sex. Then we generate a new variable `_xx_` that will only be used here. The regression estimates the two means (note that `_xx_` is not used in the regression but we need it for the *margins* command also we need to drop the constant in this case)²

```
tabulate sex, generate(Sex)
gen _xx_ = 0
label var _xx_ "Ratio"
reg bp Sex1 Sex2 _xx_ , noconstant
```

² See program `Fieller_for_ind_corr_means.do` for details..

We first use the *nlcom* command to determine the Delta method 95% confidence interval of the ratio.

```
nlcom _b[Sex1] / _b[Sex2] , level(95)
```

This results in:

bp	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_nl_1	.5301097	.4622018	1.15	0.251	-.3757891	1.436009

Thus the ratio is .5301097 and the Delta CI is -.375789 to 1.436009 – a symmetric interval.

To generate the line-plot version of the Fieller we can use the *margins* command to estimate the value of the linear combination $\hat{\gamma} - \theta_i \hat{\beta}$ where `_xx_` is the value of θ_i along with the 95% confidence interval.

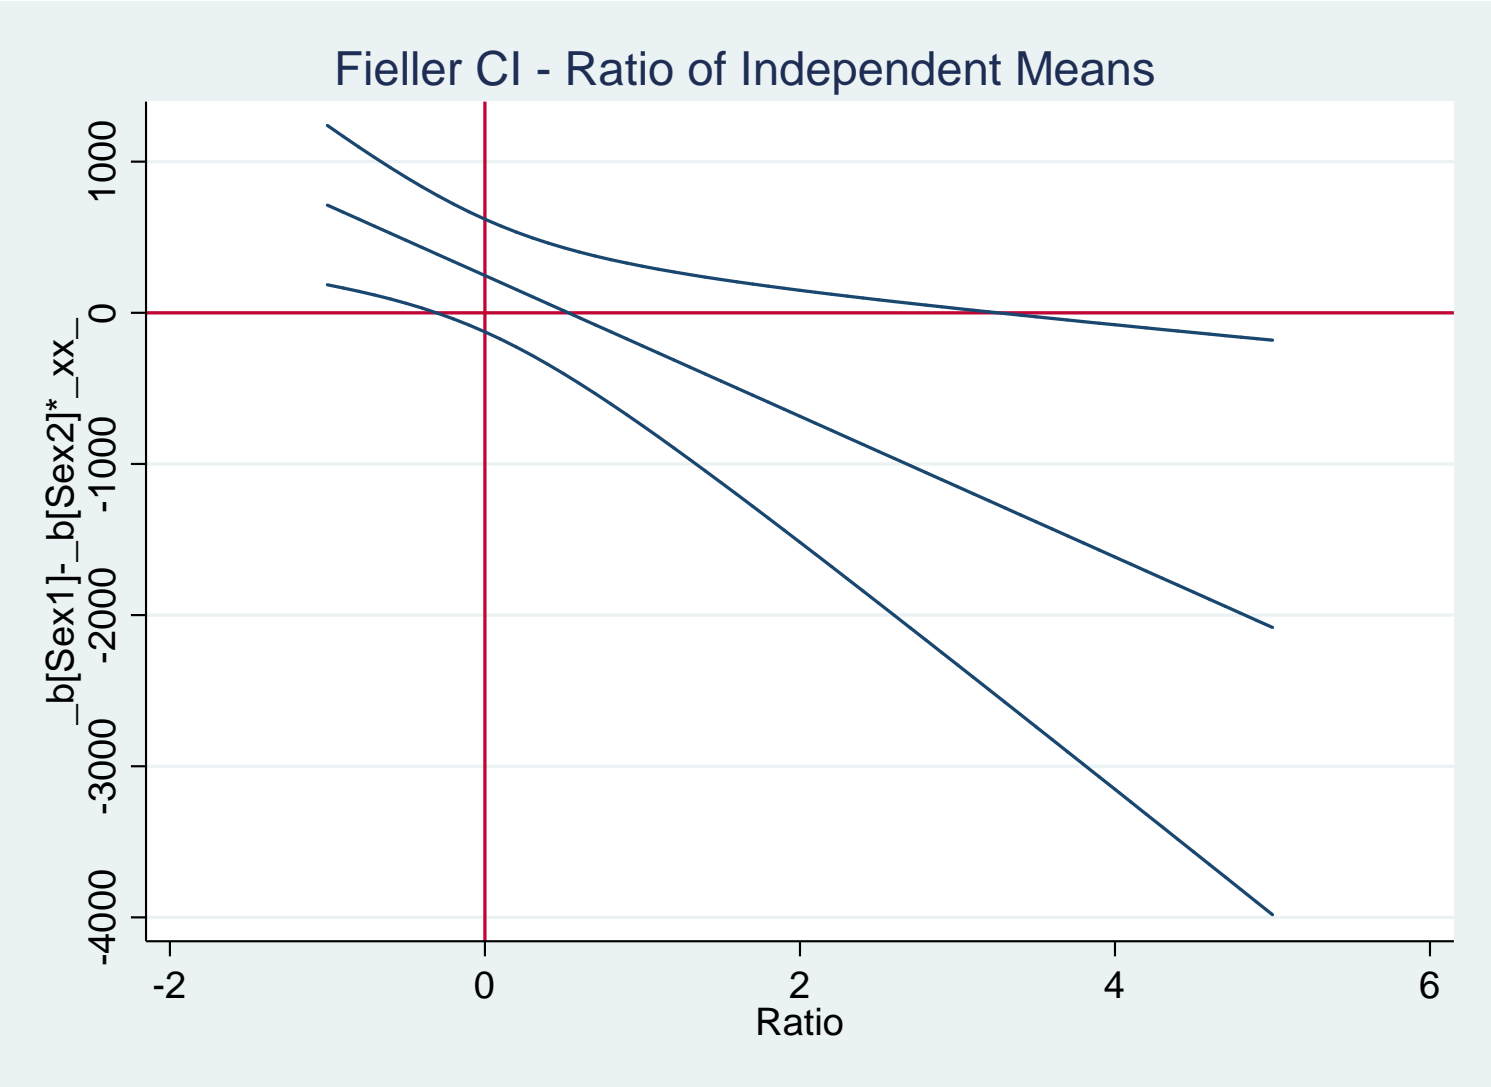
```
margins, expression (_b[Sex1] - _b[Sex2]*_xx_)  
at(_xx_ = (-1(.1)5)) level(95)
```

Note that we have set the limits of the expression to be between -1 and 5 as the limiting values over which we evaluate the possible values of the confidence interval.

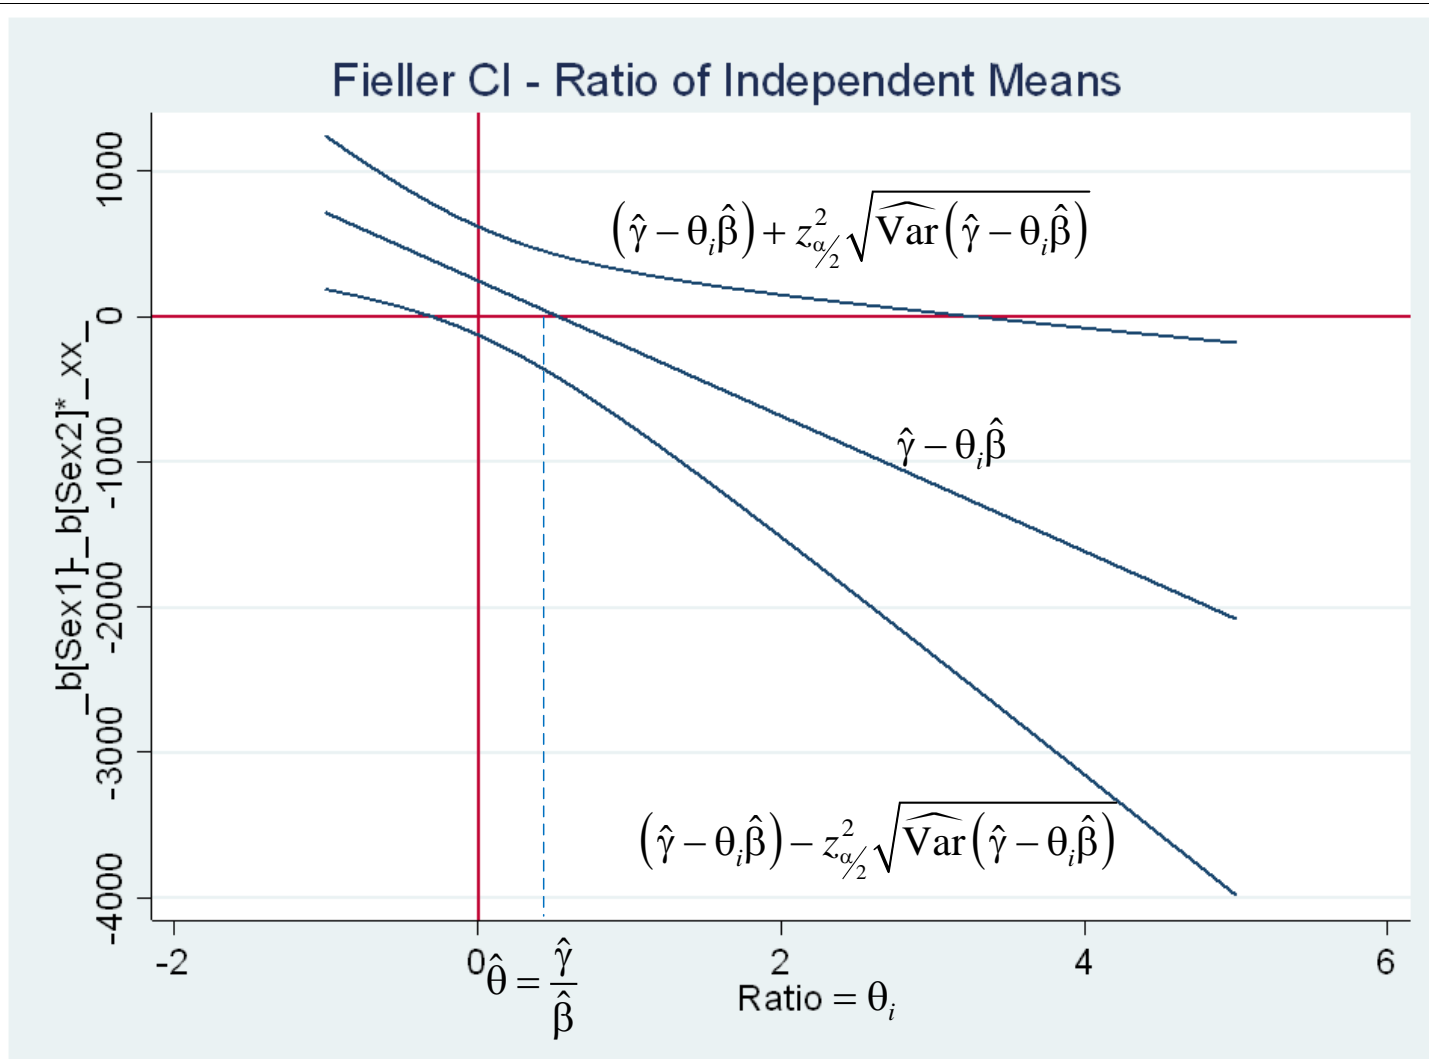
Next we use the *marginsplot* command to plot the linear combination along with the 95% interval.

```
marginsplot, yline(0) xline(0) recast(line)
recastci(rline) name(Ratio_ind, replace)
title(Fieller CI- Ratio of Independent Means)
```

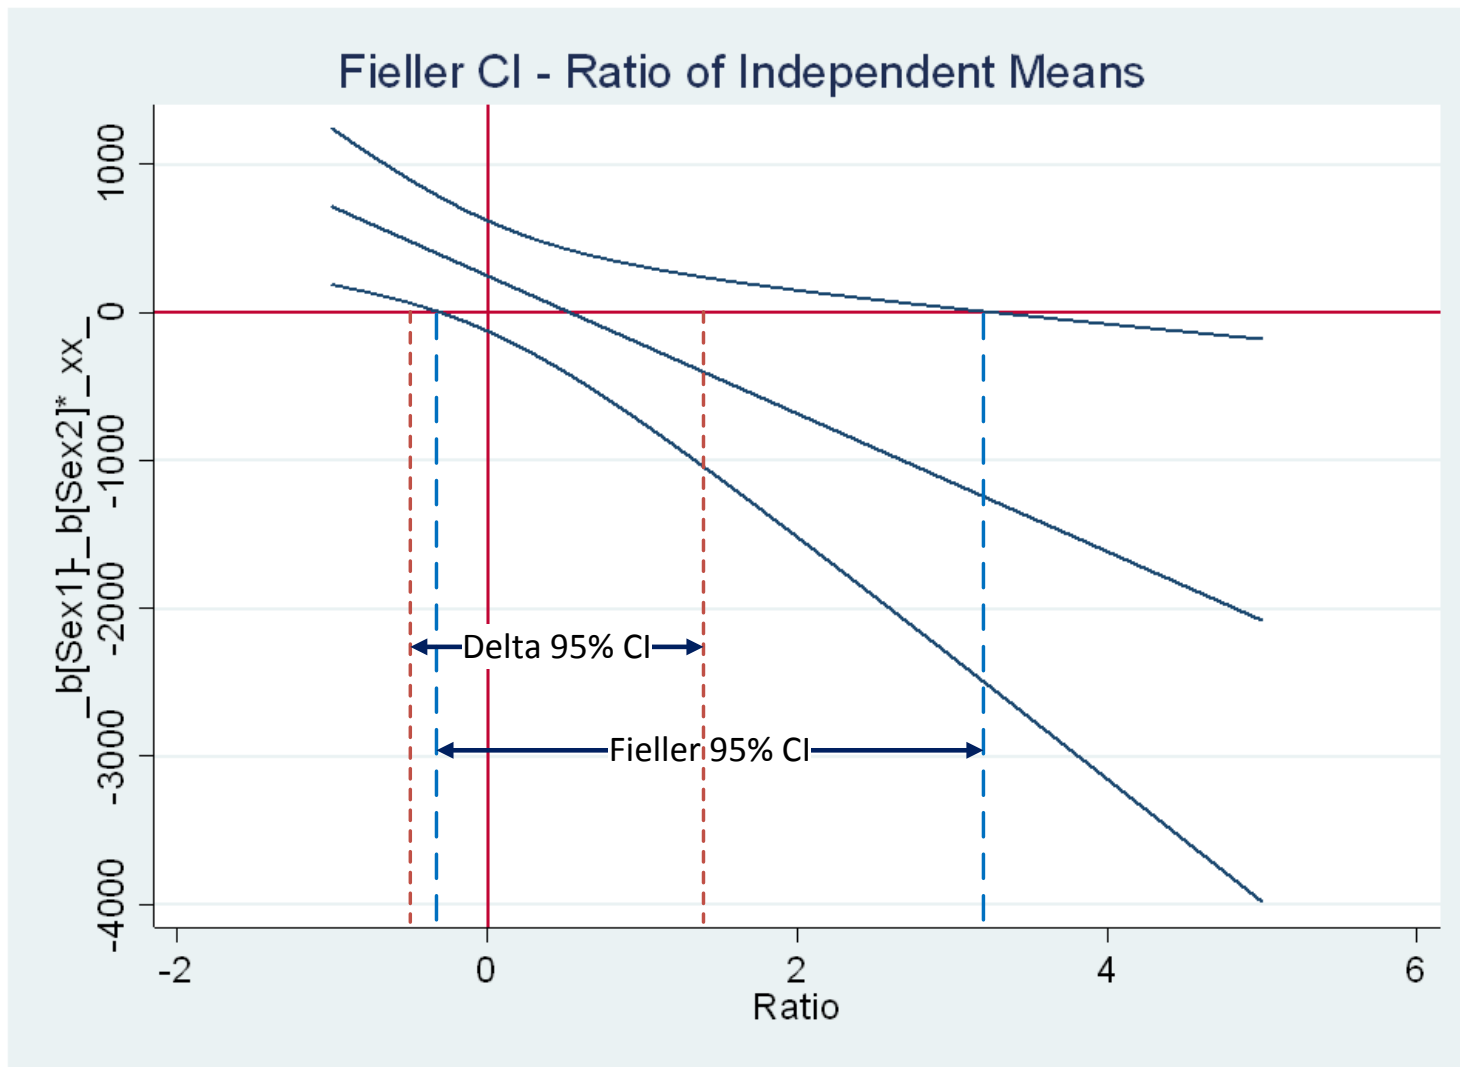
We add the zero-zero axis to evaluate our hypothesis.



This is the Stata produced graph.



The items plotted.



The identification of the 95% Fieller and Delta CIs.

Alternatively, the 99% Delta interval is listed as:

```
nlcom _b[Sex1] / _b[Sex2] , level(99)
```

will always result in a finite symmetric CI:

bp	Coef.	Std. Err.	z	P> z	[99% Conf. Interval]	
_nl_1	.5301097	.4622018	1.15	0.251	-.6604432	1.720663

The results of the *fieller* command however:

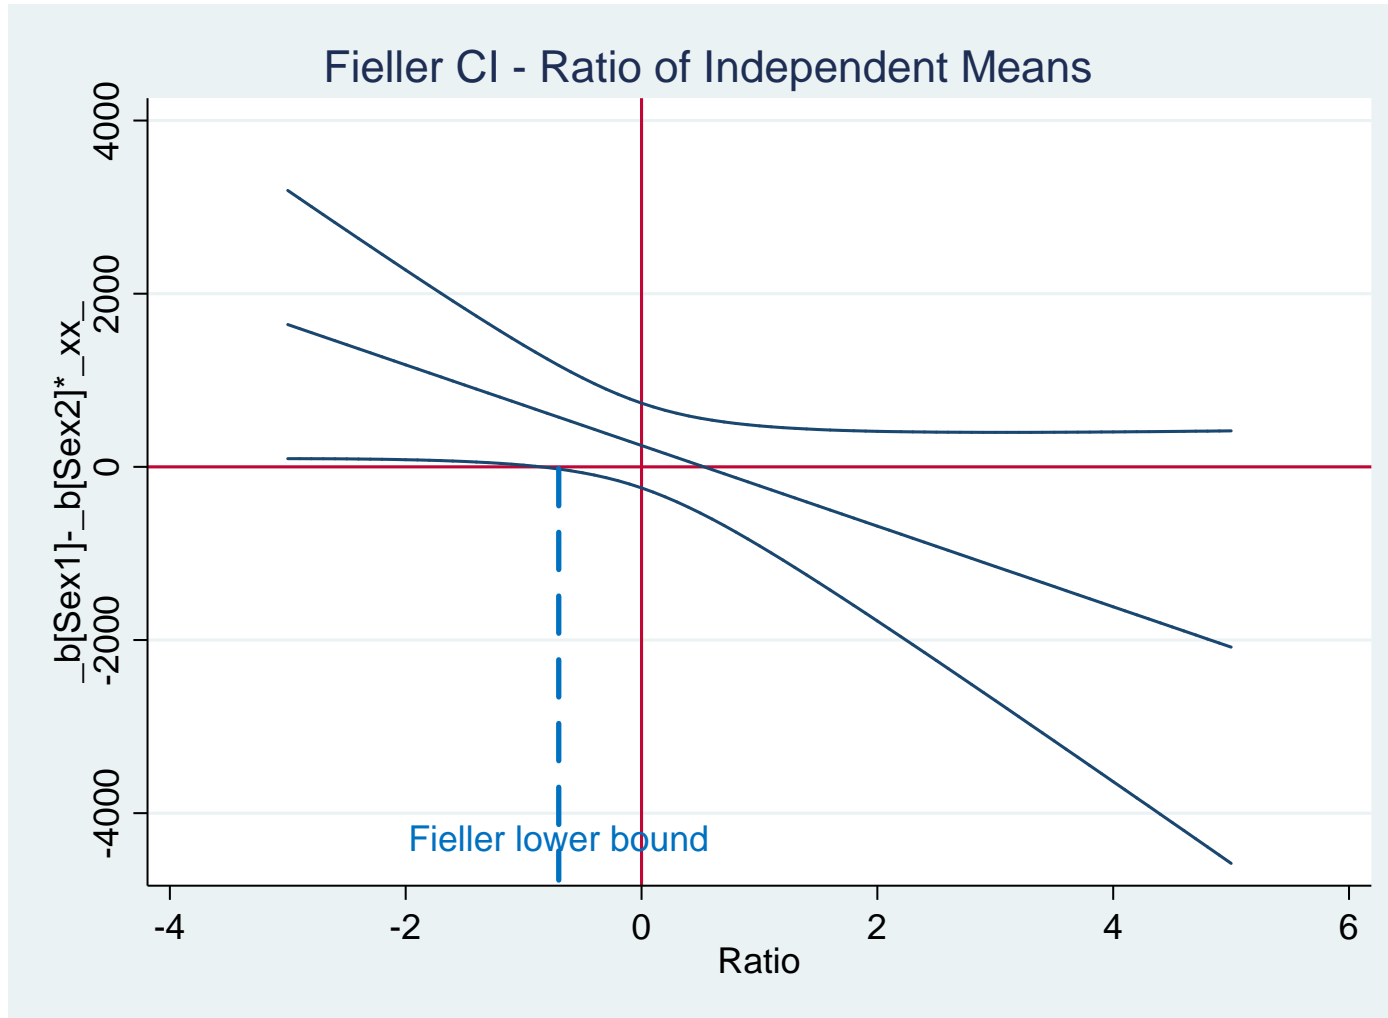
```
fieller bp , by(sex) level(99)
```

results in:

Quotient not statistically significantly different from zero at the 99 level.

Confidence interval not estimable

The equivalent plot however does provide more information:



From this plot we can observe a finite lower bound with an infinite upper bound

An Example of the ratio of the means from the same number of observations.

This example data contains data on observations that are designated as “before” and “after” for the same individual. We can use a seemingly unrelated regression which relaxes the assumptions ($\sigma_{12} = 0$, and $\sigma_1^2 = \sigma_2^2$). However, it does require that the number of observations be equal.

To use this same example data as above, we first need to reshape the data from a “long” format to a “wide” format using the *reshape* command.

```
reshape wide bp, i(patient) j(when)
```

Then we estimate two equations using a seemingly unrelated regressions (SUR) specification with the *sureg* command:

```
sureg (bp1 _xx_) (bp2 _xx_ ) , corr
```

Again, we include the `_xx_` variable thus we are interested in the constants in equation. Note that although the SUR procedure is a form of GLS the estimates in this case are the same as OLS because the regressors are the same in both equations, however the resulting covariance between the intercepts will be estimated as non-zero and the errors for each equation will be used to estimate separate variances.

The Delta interval can be determined for this ratio can be estimated via the *nlcom* command where we identify the estimates of the constants from each equation.

```
nlcom [bp1]_b[_cons] / [bp2]_b[_cons]
```

Resulting in

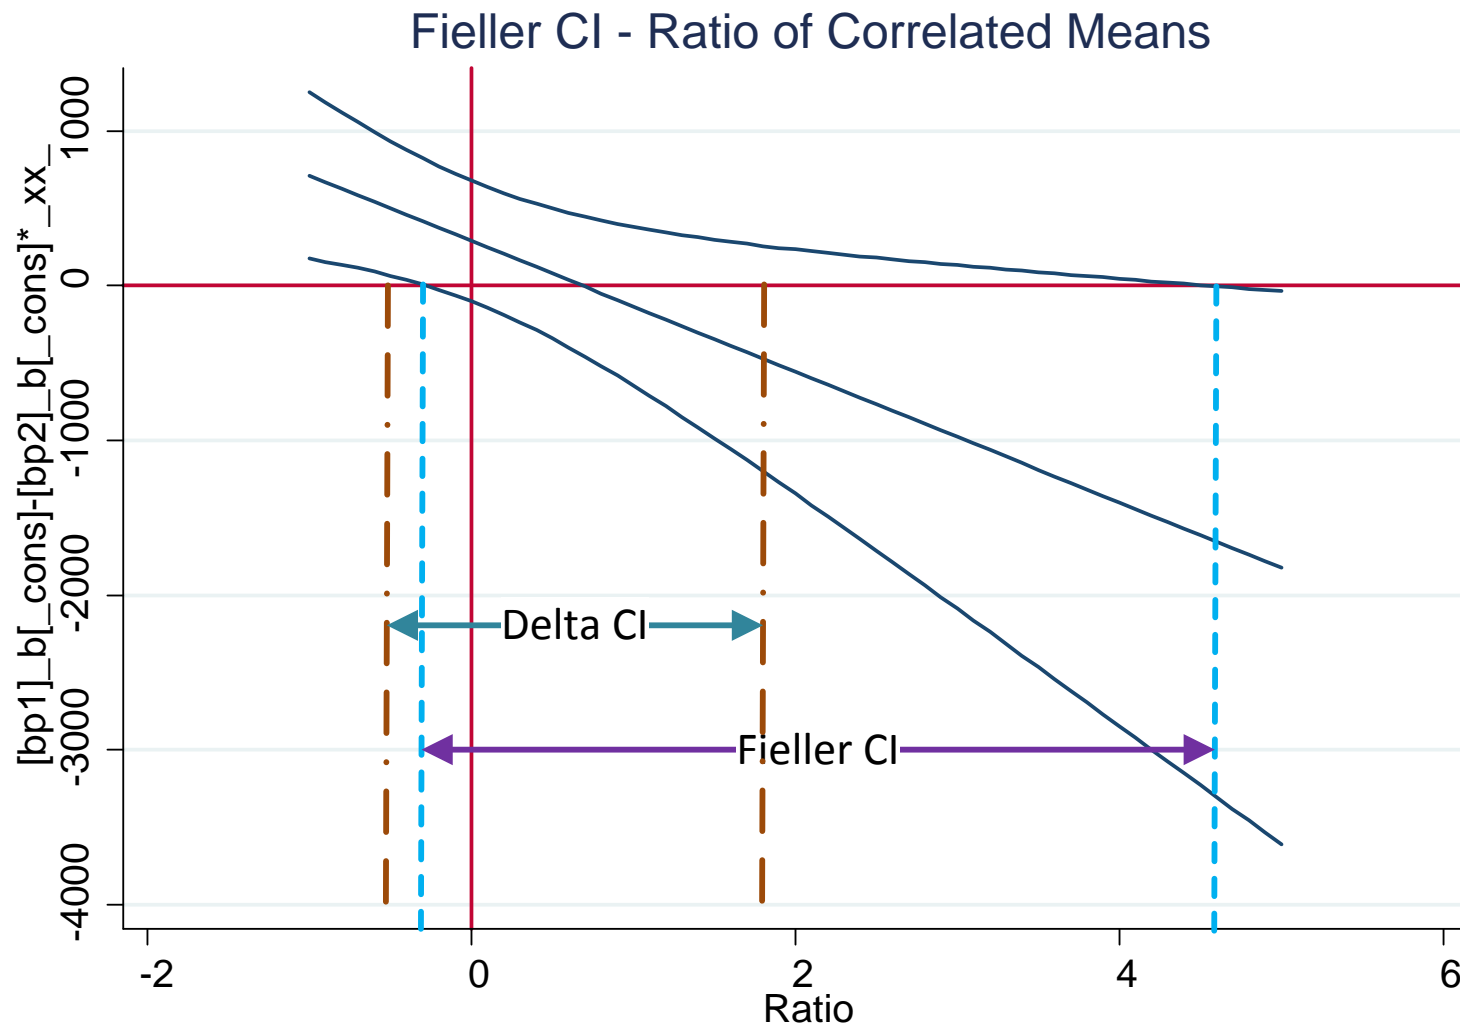
	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_nl_1	.6867762	.5416124	1.27	0.205	-.3747645	1.748317

The ratio is estimated as .6867762 and the CI is approximated as -.374 to 1.748.

As before we now use the *margins* and *marginsplot* commands to estimate and plot the confidence interval for the linear combination. The code for these steps are:

```
margins, expression ([bp1]_b[_cons] -  
[bp2]_b[_cons]*_xx_) at(_xx_=(-1(.1)5)) level(95)
```

```
marginsplot, yline(0) xline(0) recast(line)  
recastci(rline) name(Ratio_corr, replace)  
title(Fieller CI for Ratio of Correlated Means)
```



The comparison of the CIs for the correlated means.

The turning point of a quadratic.

A common case of the ratio of parameter estimates would be where one has fit a quadratic specification and would like to determine the location of the turning point.

For example, the typical specification would be of the form:

$$Y_t = \alpha + \gamma X_t + \beta X_t^2 + \varepsilon_t$$

The turning point occurs when the marginal effect of the X variable is zero:

$$\frac{\partial Y_t}{\partial X_t} = \gamma + 2\beta X_t, \text{ thus } X^* = -\frac{\gamma}{2\beta}$$

Where X^* is the value of X where the sign of the marginal effect reverses – in other words where the turning point occurs.

For example, using the Stata example data set “auto” we can estimate a quadratic between a car’s price and the weight of the car – here we only consider the domestic cars in the data.³

³ See Turning_point.do for the full set of code for this example.

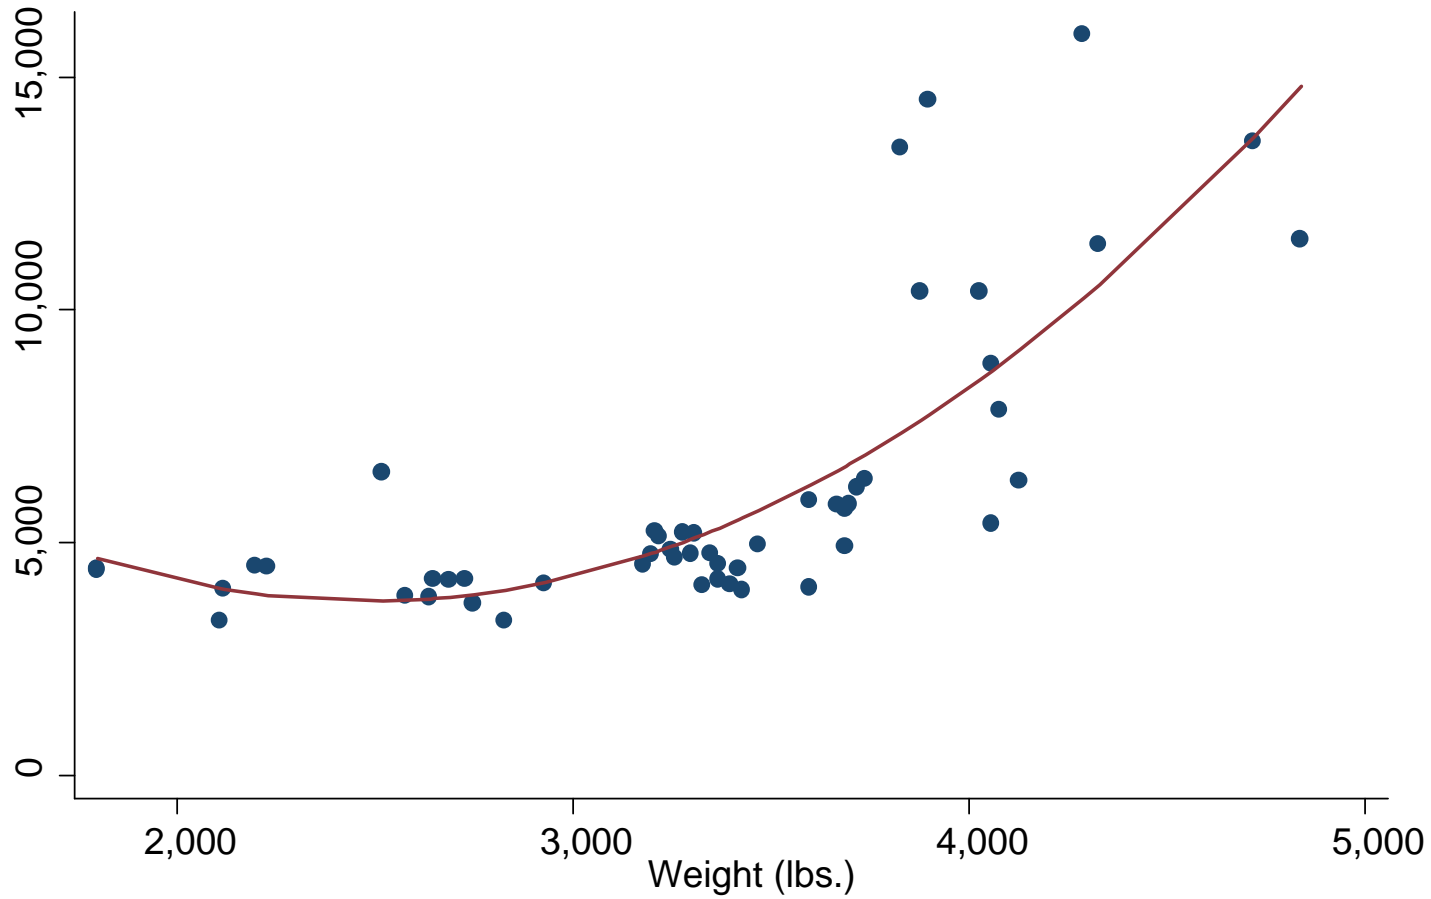
First we estimate a quadratic equation where weight2 is the square of weight:

```
reg price weight weight2
```

Source	SS	df	MS	Number of obs	=	52
Model	300678877	2	150339439	F(2, 49)	=	39.08
Residual	188515923	49	3847263.74	Prob > F	=	0.0000
Total	489194801	51	9592054.92	R-squared	=	0.6146
				Adj R-squared	=	0.5989
				Root MSE	=	1961.4

price	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
weight	-9.841423	2.851178	-3.45	0.001	-15.57108	-4.111766
weight2	.0019851	.0004367	4.55	0.000	.0011075	.0028626
_cons	15933.98	4555.578	3.50	0.001	6779.208	25088.76

The fitted regression line and data are shown below:



Price to weight relationship

To find the turning point we can use the following code:

```
nlcom (_b[weight] / ( -2*_b[weight2] ) ) , level(95)
```

Where the *nlcom* command provides the Delta estimation of the CIs for the turning point as:

price	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_nl_1	2478.865	193.4764	12.81	0.000	2099.658	2858.072

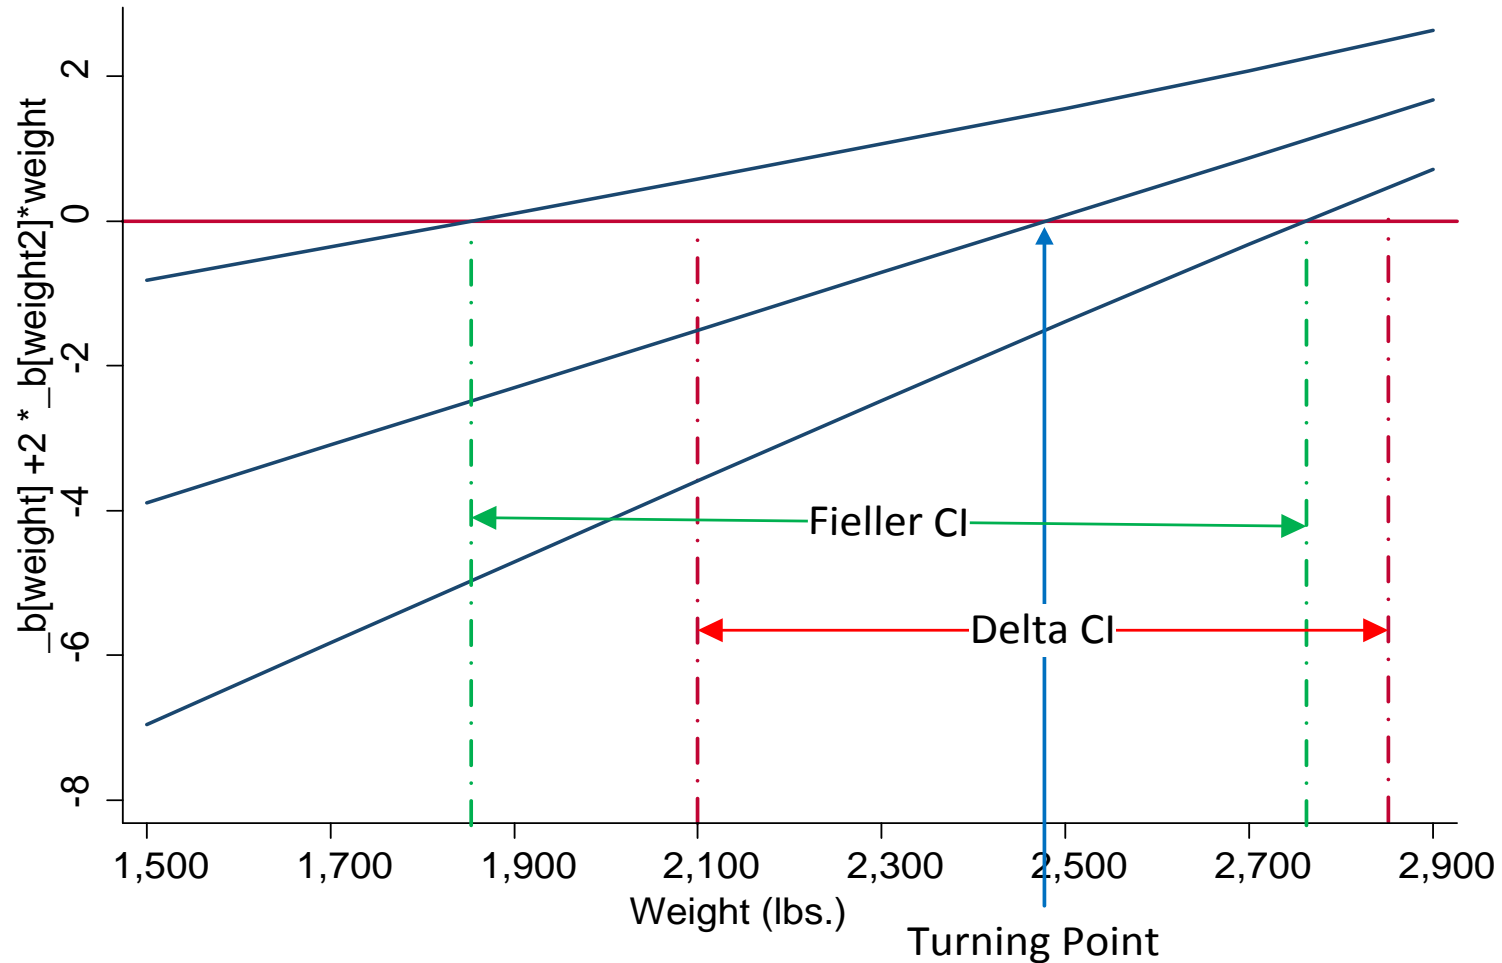
The turning point is estimated at 2,478.865 pounds and the approximate 95% CI is 2,099.658 to 2,858.072.

Alternatively, we can plot the marginal effect of the change in weight as it influences the price. Using the Stata code:

```
margins, expression (_b[weight] + 2
*_b[weight2]*weight) at (weight = (1500(200)3000))
level(95)
```

```
marginsplot, yline(0) recast(line) recastci(rline)
name(Turning_point, replace) title(Fieller CI -
Turning point of quadratic )
```

Fieller CI - Turning point of quadratic



Note that the Fieller CI is asymmetric and has a lower limit that is less than the Delta interval. (see Lye and Hirschberg 2012 for more examples).

An Example of the 50% dose and Willingness to Pay.

The 50% dose problem was the original (1944) application that Fieller considered when he first proposed this method. In that case, he was interested in the dose of a medication/poison that would result in a greater than 50% chance that the subject would live/die.

In the willingness to pay case the variable applied is the price paid (e.g. Hole 2007). How much would someone pay for a product/service? Thus, the data would be a dichotomous response to a continuous stimulus-drug-price.

Typically, this problem is modelled by a probability of occurrence using a logit or probit estimator.

$$P(Y_t = 1) = \alpha + \gamma X_t + \sum_{i=1}^k \beta_i W_{it} + \varepsilon_t$$

Where X is the variable dose/price and the W s are other covariates.

The question of the 50% dose is posed as: What value of X (designated as X^*) makes this relationship

$$1/2 = \alpha + \gamma X^* + \sum_{i=1}^k \beta_i W_{it}$$

true?

Solving for X^* we have the ratio given by:

$$X^* = \frac{1/2 - \left(\alpha + \sum_{i=1}^k \beta_i W_{it} \right)}{\gamma}$$

The ratio can be evaluated at specific values of W for cases of interest or at the means.

As an example, we use data on the survivors of the Titanic disaster where the dose/price is the fare paid for the trip (see Frey et al 2010). We fit the model to the data with covariates for sex and age.⁴

```
probit survive fare age female
```

Which results in:

Probit regression		Number of obs	=	1,033	
Log likelihood = -525.5508		LR chi2(3)	=	345.42	
		Prob > chi2	=	0.0000	
		Pseudo R2	=	0.2473	
survived	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
fare	.0059587	.0010804	5.52	0.000	.0038411 .0080763
age	-.006552	.0031211	-2.10	0.036	-.0126692 -.0004348
female	1.440331	.0910842	15.81	0.000	1.261809 1.618852
_cons	-.8152731	.1074258	-7.59	0.000	-1.025824 -.6047223

⁴ See the program Willingness_to_pay.do for the full code for this example.

Using the Delta method for the 50% fare for the average age and assuming 50% gender split:

```
nlcom ( - (_b[_cons]+ _b[age]*m_age + _b[female] * .5) / _b[fare] ), level(95)
```

Results in the estimated CI on the 50% fare of:

survived	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
_nl_1	48.67789	8.092075	6.02	0.000	32.81771	64.53806

Thus, we find a fare of 48.7 pounds with an approximate 95% interval of (32.82 to 64.54).

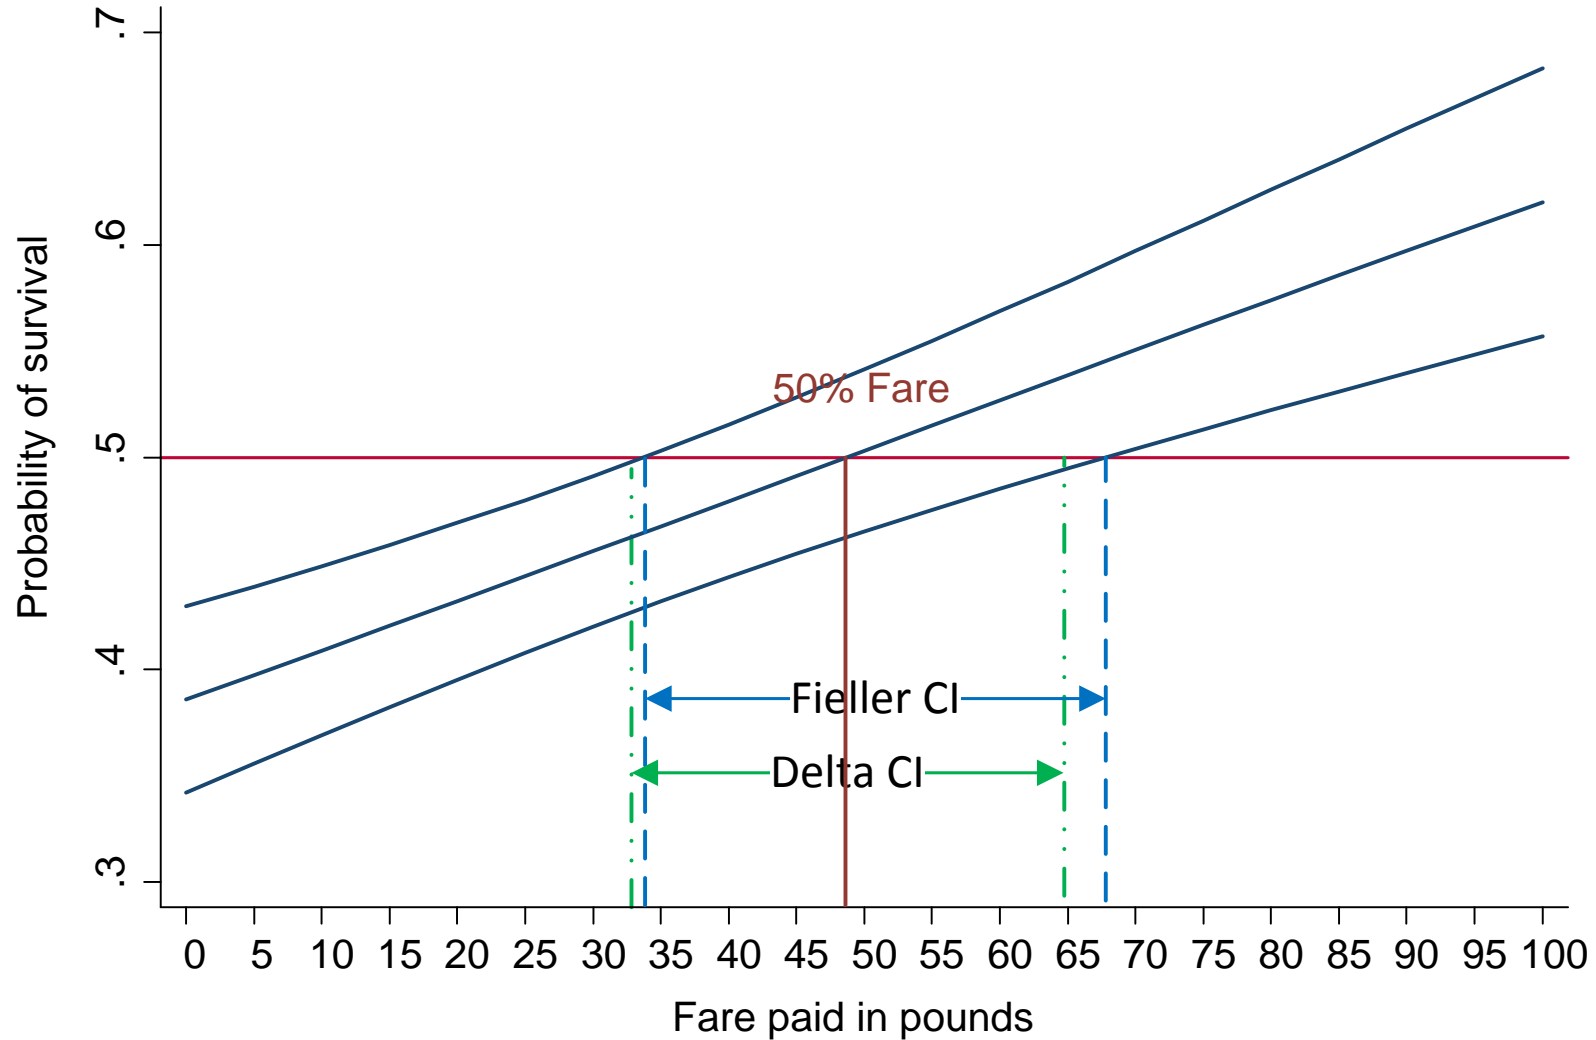
To estimate the Fieller interval we use the *margins* command to predict the probability of survival:

```
margins, expression ( normal( _b[_cons]+_b[fare]*fare  
+ _b[age]*age + _b[female] * female ) ) at(age  
(mean) female = .5 fare=(0(5)100)) level(95)
```

Followed by the *marginsplot* command:

```
marginsplot, yline(.5) recast(line) recastci(rline)
```


Fieller CI - limit of 50% dose



In this case the Delta and Fieller intervals coincide quite closely.

Conclusions

In addition to the examples discussed here we have also added the routine used in our 2017 paper that estimates the indirect least-squares estimate for the just identified IV estimator which can be shown to be the ratio of parameters estimated from two seemingly unrelated regressions.⁵

The types of code for these cases varies depending on the best way to proceed. Using the *margins* and the *marginsplot* command is a very compact method for generating these plots. The limits of the margins may need some adjustment after the first attempt so that the bounds are plotted within the plot window. In addition, the *level* option allows the construction of CIs with different values of α .

⁵ See the program `Indirect_Least_squares.do` for this example.

All programs/data used in this talk are available at:

https://www.online.fbe.unimelb.edu.au/t_drive/stata_examples

do files	Data
Fieller_for_ind_corr.do	bplong_example.dta
Indirect_Least_Squares.do	maketable4.dta
Turning_point.do	auto_domestic.dta
Willingness_to_pay.do	titanic1.dta

References

- Coveney, Joseph, (2004),” FIELLER: Stata module to calculate confidence intervals of quotients of normal variate”, downloaded 25/9/17, from : <https://ideas.repec.org/c/boc/bocode/s448101.html> .
- Fieller, E. C., (1944), “A Fundamental Formula in the Statistics of Biological Assay and Some Applications”, *Quarterly Journal of Pharmacy and Pharmacology*, 17, 117-123.
- _____,(1954), “Some Problems in Interval Estimation,” *Journal of the Royal Statistical Society, Ser. B*, 16, 174–185.
- Frey, B. S., D. A. Savage, and B. Torgler, 2010, “Interaction of natural survival instincts and internalized social norms exploring the *Titanic* and *Lusitania* disasters”, *Proceedings of the National Academy of Science*, 107, 4862-4865.
- Hole, A. R., 2007, “A comparison of approaches to estimating confidence intervals for willingness to pay measures”, *Health Economics*, 16, 827-840.
- Hirschberg, Joseph G. and Jenny N. Lye, 2010, “A Geometric Comparison of the Delta and Fieller Confidence Intervals”, *The American Statistician*, 64, 234-241.
- _____, 2017, “Inverting the Indirect - The Ellipse and the Boomerang: Visualising the Confidence Intervals of the Structural Coefficient from Two-Stage Least Squares”, *Journal of Econometrics*, 199, 173-183.
- Lye, Jenny N. and Joseph G. Hirschberg, 2010, “Two Geometric Representations of Confidence Intervals for Ratios of Linear Combinations of Regression Parameters: An Application to the NAIRU, *Economics Letters*, 108, 73-76.
- _____, 2012, “Inverse Test Confidence Intervals for Turning Points: A Demonstration with Higher Order Polynomials”, Dek Terrell and Daniel Millimet (ed.) 30th Anniversary Edition (*Advances in Econometrics*, Volume 30) Emerald Group Publishing Limited, pp 59-95.