Employee Performance Study

Manager.jmp

A firm either promotes managers in-house or recruits managers from outside firms. It has been claimed that managers hired outside the firm perform better than those promoted inside the firm, at least in terms of internal employee evaluations.

Two candidates for promotion have comparable backgrounds, but one comes from within the firm and the other comes from outside. Based on this sample of 150 managers, which should you promote?

The results of a two-sample comparison of in-house promotions (*Origin* = "Internal") to those brought in from outside the firm (*Origin* = "External") support the conjecture that those recruited outside the firm perform better.

Mngr Rating by Origin

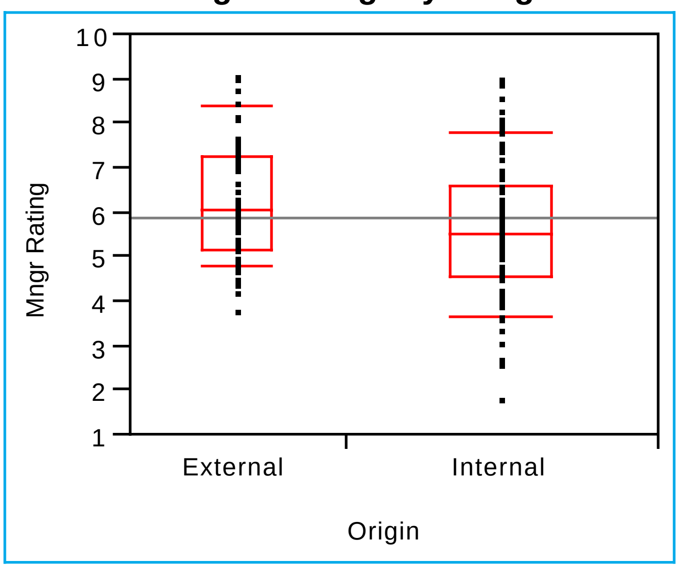

	Difference	t-Test	DF	Prob>|t|
Estimate	0.72	2.984	148	0.0033
Std Error	0.24			
Lower 95%	0.24			
Upper 95%	1.19			

Assuming equal variances

Level	Number	Mean	Std Error (pooled)
External	62	6.32	0.18
Internal	88	5.60	0.15

Surprising though it may seem, regression analysis with a categorical factor gives the same results as the initial *t*-test! To force JMP to do the regression (rather than a two-sample *t*-test), use the *Fit Model* platform with *Origin* as the single predictor.

Response: Mngr Rating

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	5.96	0.12	49.67	<.0001
Origin[External-Internal]	0.36	0.12	2.98	0.0033

The leverage plot is also of interest with a categorical predictor. The vertical scale is that from the prior comparison boxplot, and the groups are located at the so-called "least squares means," which in this case are the original group means (compare to the *t*-test output).

Origin

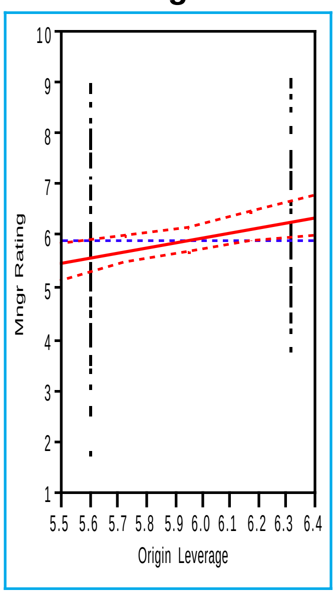

Least Squares Means

Level	Least Sq Mean	Std Error	Mean
External	6.321	0.184	6.321
Internal	5.605	0.154	5.605

Both analyses lead to the same conclusion. The *t*-test indicates that the external managers are doing significantly better:

	Difference	t-Test	DF	Prob>|t|
Estimate	0.72	2.984	148	0.0033

Level	Number	Mean	Std Error (pooled)
External	62	6.32	0.18
Internal	88	5.60	0.15

Equivalently, via regression (from the *Fit Model* output), we obtain the same *t*-statistic for the size of the difference in the averages. In this case, though, the estimated regression coefficient for *Origin* is one-half of the difference 0.72 between the two means.

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	5.96	0.12	49.67	<.0001
Origin[External-Internal]	0.36	0.12	2.98	0.0033

To understand the interpretation of the coefficient estimate attached to the categorical factor, write out the model for each group represented by the categorical factor. In this case, there are only two groups (internal and external). Reverse the signs of the estimate associated with "Internal" (as suggested by the minus sign preceding it).

External   Fit = 5.96 + 0.36 = 6.32 = average for "External"

Internal   Fit = 5.96 – 0.36 = 5.60 = average for "Internal"

Thus, the coefficient of the categorical factor shows how the fitted intercept for each group differs from the overall intercept.

The convention used by JMP for reporting the coefficients of categorical variables requires some explanation. The notation Origin[External-Internal], in spite of its labeling, indicates that we are judging the effect of *Origin* being zero relative to the average of the groups, not to the other group directly. Since the coefficient associated with Origin[External-Internal] is the distance of this category from the average and there are only two groups, the distance between the two groups is twice the value shown.

Other factors appear to influence the rating score as well. For example, the rating is correlated with the salary of the employee. The coded scatterplot (use the *Color marker by Col...* command from the *Rows* menu to use *Origin* as the color and marker variable) suggests that the internal managers are paid less (internal managers are coded with the green +'s). Evidently, external managers were recruited at higher salaries.

Mngr Rating by Salary

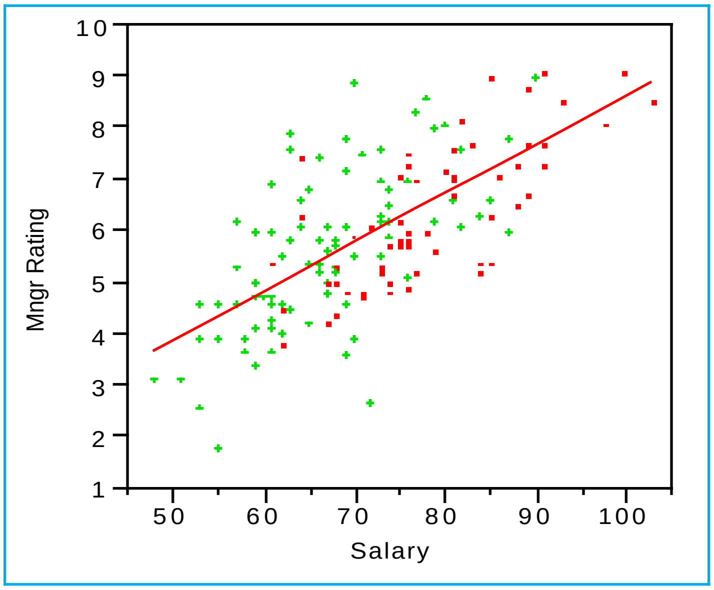

Linear Fit

RSquare	0.467
RSquare Adj	0.464
Root Mean Square Error	1.088
Mean of Response	5.901
Observations (or Sum Wgts)	150

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	-0.898	0.603	-1.49	0.139
Salary	0.095	0.008	11.40	0.000

To confirm the impression from the coded scatterplot, here is the comparison of salaries of the two groups. Indeed, the externally recruited managers have significantly higher salaries.

Salary by Origin

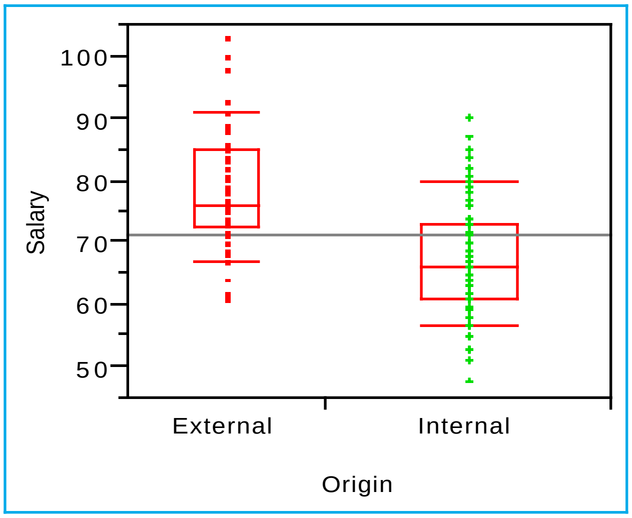

t-Test

	Difference	t-Test	DF	Prob>|t|
Estimate	11.457	7.581	148	<.0001
Std Error	1.511			
Lower 95%	8.470			
Upper 95%	14.444			

Assuming equal variances

Level	Number	Mean	Std Error
No	62	78.35	1.16
Yes	88	66.90	0.97

Thus, it appears that our initial assessment of performance has mixed two factors. The performance evaluations are related to the salary of the manager. Evidently, those more well paid are in higher positions and receive higher evaluations. Since the externally recruited managers occupy higher positions (presumably), we do not have a "fair" comparison.

The easiest way to adjust statistically for the differences in salary paid to the internal and external managers is to fit separate lines in the plot of *Rating* on *Salary*. (Use the grouping option at the bottom of the options offered by the fitting button in the *Fit Y by X* view.) The top fit in this plot is for the internal managers. The plot suggests that at comparable salaries (as a proxy for position in the company), the internal managers are doing better than those externally recruited. Summaries of the two fitted lines appear on the next page.

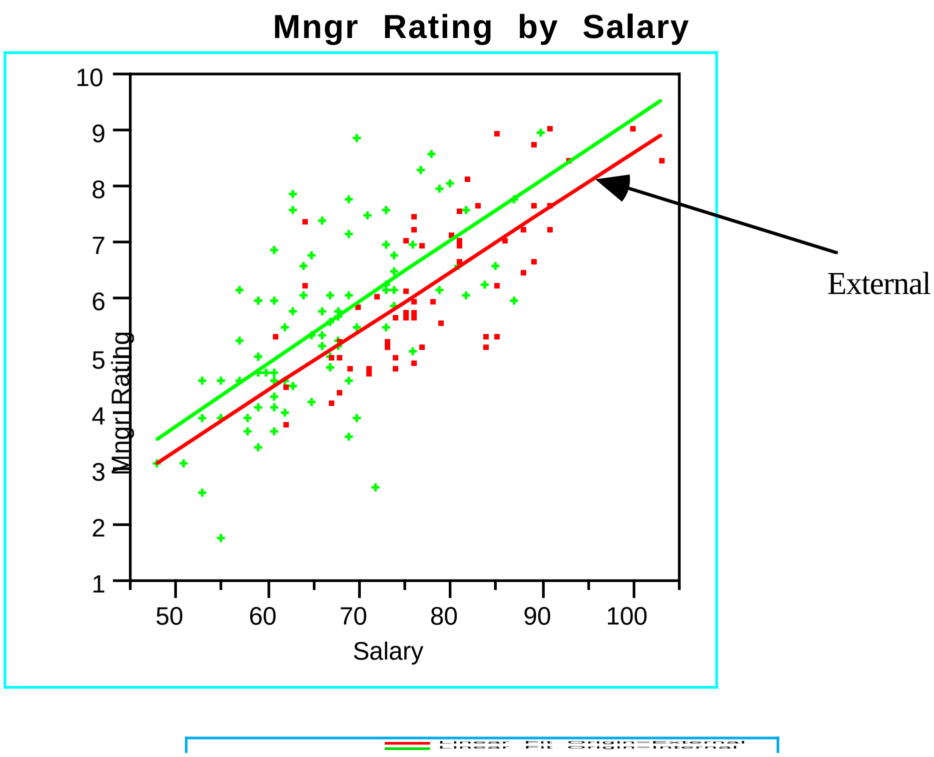

Linear Fit, External Managers

Mngr Rating = -1.9369 + 0.10539 Salary

RSquare	0.542
RSquare Adj	0.535
Root Mean Square Error	0.915
Mean of Response	6.321
Observations (or Sum Wgts)	62

Term	**Estimate**	**Std Error**	**t Ratio**	**Prob>|t|**
Intercept	-1.937	0.986	-1.96	0.0542
Salary	0.105	0.012	8.43	<.0001

Linear Fit , Internal Managers

Mngr Rating = -1.6935 + 0.10909 Salary

RSquare	0.412
RSquare Adj	0.405
Root Mean Square Error	1.171
Mean of Response	5.605
Observations (or Sum Wgts)	88

Term	**Estimate**	**Std Error**	**t Ratio**	**Prob>|t|**
Intercept	-1.694	0.949	-1.78	0.0779
Salary	0.109	0.014	7.76	<.0001

The difference between the two lines is roughly constant (the slopes are about the same). The difference between the intercepts is –1.937 – (–1.694) = –0.243. Is this a significant difference?

To determine if the difference in intercepts is significant (assuming equal slopes), we use an *analysis of covariance*, a multiple regression analysis that combines a categorical variable with other covariates. In this case, the coefficient of the categorical variable *Origin* indicates that the in-house managers are performing better if we adjust for salary.

Again, keep the JMP coding convention in mind. The notation Origin[External-Internal] indicates that we are judging the effect of *Origin* being zero relative to the average of both groups, not to the other group directly. Since the coefficient associated with Origin[External-Internal] is the distance from the average, and there are only two groups, the distance from the other group is twice the value shown, –0.514. A picture helps.

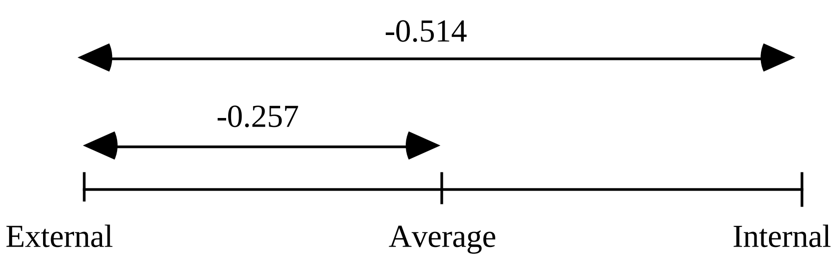

Response: Mngr Rating

RSquare	0.488
RSquare Adj	0.482
Root Mean Square Error	1.070
Mean of Response	5.901
Observations (or Sum Wgts)	150

Term	Estimate	Std Error	t Ratio	Prob>|t|
Intercept	-1.843	0.706	-2.61	0.01
Origin[Ext-Int]	-0.257	0.105	-2.46	0.01
Salary	0.107	0.010	11.14	0.00

Assuming the common slope for *Salary*, the difference in intercepts is significant.

As before with just the single categorical predictor, it is helpful to write out the fitted model for each group. When the model is augmented by adding *Salary* to the equation, the fitted model has a new interpretation. The effect of adding just *Salary* and not its interaction with the categorical variable *Origin* is to fit parallel lines to the two groups. Here is the repeated summary from the *Fit Model* command with *Origin* and *Salary*:

Term	**Estimate**	**Std Error**	**t Ratio**	**Prob>|t|**
Intercept	-1.843	0.706	-2.61	0.01
Origin[Ext-Int]	-0.257	0.105	-2.46	0.01
Salary	0.107	0.010	11.14	0.00

The fits for the two groups are

External Fit = (-1.843 – 0.257) + 0.107 *Salary* = –2.1 + 0.107 *Salary*

Internal Fit = (-1.843 + 0.257) + 0.107 *Salary* = –1.586 + 0.107 *Salary*

The fitted lines are parallel because we have made an important assumption that *forces* them to be parallel. The coefficient of the categorical factor is *one-half of the difference* between the intercepts of the two fits. As in the first example (where the regression contained only the variable *Origin*), the coefficient of the categorical factor represents the difference of the intercept of each group from the overall intercept. From this result, the difference between the fitted intercepts is significant. Only now the model adjusts for the *Salary* effect, and the internal managers are doing significantly better (higher) than the external managers (*t*=-2.46). The adjustment for *Salary* reverses the impression conveyed by the initial *t*-test.

Leverage plots are still useful, though they are a bit harder to understand when other predictors are in the model. Before, the leverage plot for *Origin* showed two discrete categories; the plot showed two columns. Now, with the other factor added, the categories are less well defined (albeit still rather distinct when viewed in color).

Why? As we have seen, *Salary* and *Origin* are related. Now the values along the horizontal axis of the leverage plot are adjusted for the impact of *Salary*.

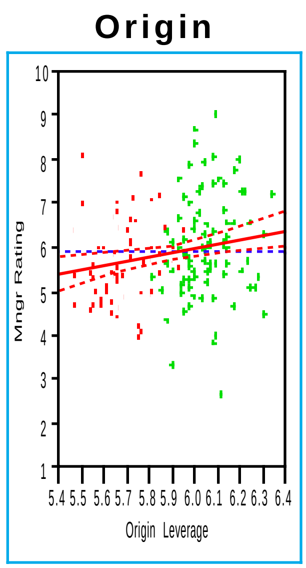

Note the inversion in the ordering of the means. Initially, the mean for in-house managers was smaller (5.60 versus 6.32). After adjusting for the salary effect as given by the least squares means, the ordering is reversed (6.11 versus 5.60)

Origin	Least Sq Mean	Std Error	Mean
External	5.60	0.15	6.32
Internal	6.11	0.12	5.60

The leverage plot for *Origin* provides further confirmation that the difference in intercepts is significant, though small relative to the variation.

The residual plot and the normal probability plot of the residuals suggest that the assumptions underlying this fitted model are reasonable.

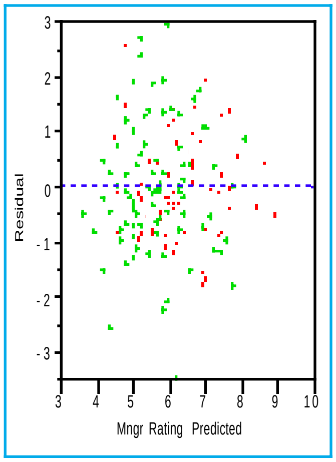

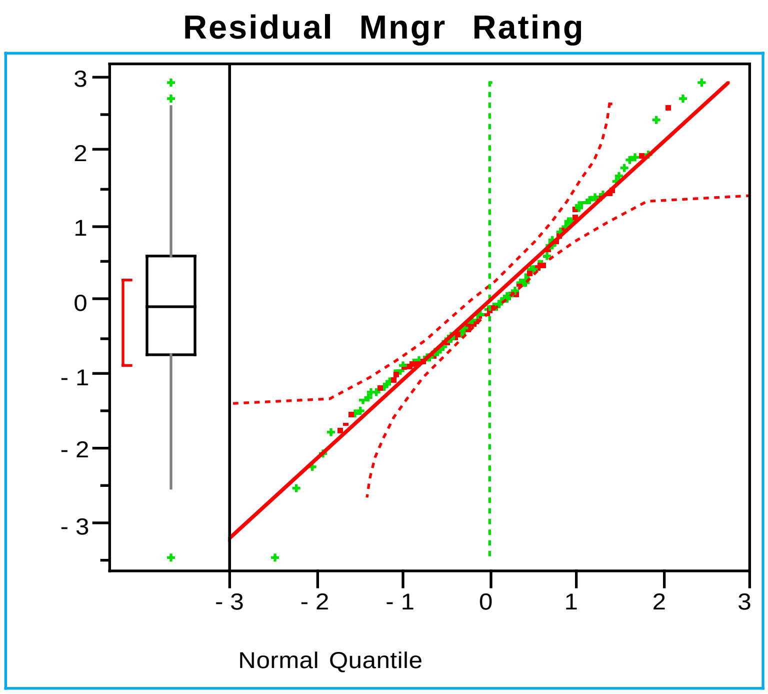

But what of the important assumption regarding parallel fits? Though clearly similar, we should check this assumption. They are not always going to be so similar, and we need a method to check this crucial step of the analysis. Doing so requires that we add an *interaction* to the analysis (also known as a *cross-product* term). An interaction allows the slopes to differ in the multiple regression fit.

An interaction is added to the fitted model by the following sequence of steps in the *Fit Model* dialog, assuming that the dialog from the prior fit is still available (with *Mngr Rating* as the response and *Origin* and *Salary* as predictors). The steps are

(1) add *Salary* to the set of predictors. *Salary* shows up as a predictor twice;

(2) select the *Salary* term just added to the collection of predictors (click on it);

(3) select the categorical factor *Origin* from the list of variables in the upper left corner of the dialog box;

(4) with both selected, the "Cross" button becomes highlighted. Click on this button to add the cross-product or interaction term to the model; and

(5) run the model as usual.

Here is the summary of the fit. The added variable is not significant, implying that the slopes are indeed roughly parallel. Note the high variance inflation factors in the VIF column.

Response: Mngr Rating

RSquare	0.489
RSquare Adj	0.478
Root Mean Square Error	1.073
Mean of Response	5.901
Observations (or Sum Wgts)	150

Term	Estimate	Std Error	t Ratio	Prob>|t|	VIF
Intercept	-1.815	0.724	-2.51	0.013	0.0
Origin[Ext-Int]	-0.122	0.724	-0.17	0.867	66.1
Salary	0.107	0.010	10.99	0.000	1.4
Salary*Origin[Ext-Int]	-0.002	0.010	-0.19	0.850	64.5

Once again, to interpret the coefficients in this model, write out the fits for the two groups.

"External": Fit = (-1.815-0.122) + (0.107 - 0.002) Salary = –1.937 + 0.105 Salary

"Internal": Fit = (-1.815+0.122) + (0.107 + 0.002) Salary = –1.693 + 0.109 Salary

These are the same two fits that we obtained by fitting regression lines to the two groups separately at the start of the analysis. The addition of the interaction allows the slopes to differ rather than be forced to be equal (parallel fits). The small size of the interaction (t = –0.19) indicates that the slopes are indeed essentially the same (as noted in the graph of the two fits). Since the fits are evidently parallel, we should remove the interaction and focus our interpretation upon the model with just *Salary* and *Origin.*

The presence of the interaction leads to substantial collinearity in this example and suggests that the difference in intercepts is not significant. The collinearity is also apparent in the associated leverage plots (not shown here). The residual plots from this last regression both look fine with a reasonable approximation to normality.

Finally, we need to investigate one more aspect of the fitted model. When analyzing grouped data, we need to confirm that the variability is consistent across the groups. In this example, the two sets of residuals appear have similar variation. Remember, with groups of differing size, compare the heights of the boxes, not the range of the data.

Residual Mngr Rating by Origin

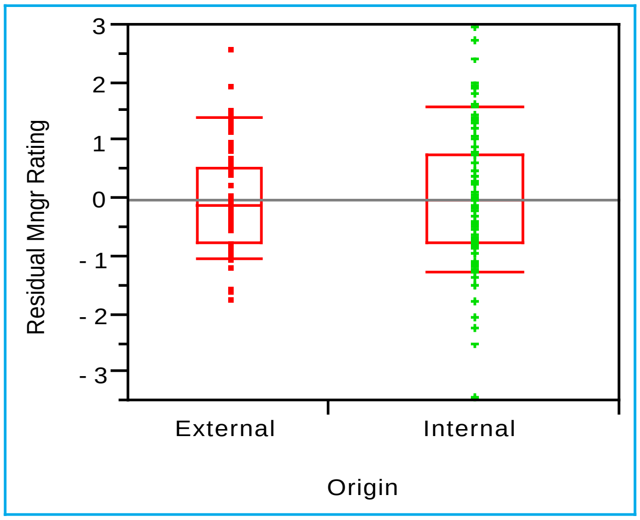

Means and Std Deviations

Level	Number	Mean	Std Dev	Std Err Mean
External	62	-0.000	0.908	0.115
Internal	88	-0.000	1.164	0.124

A firm either promotes managers in-house or recruits managers from outside firms. It has been claimed that managers hired outside of the firm perform better than those promoted inside the firm, at least in terms of internal employee evaluations.

Do the available data on a sample of 150 managers support this claim?

We find that the fits to the two groups are indeed parallel and conclude from the prior fit (with *Salary* and *Origin*) that internal managers are doing better once we adjust for differences in position within the firm. That is, if we adjust for differences in salary, the in-house managers are doing better (t=-2.5, p=0.01) than their colleagues who are recruited from the outside at higher salaries.

On the other hand, why control for only differences in salary? The following regression shows what occurs if years of experience is used as an additional control. The internal and external managers are doing about the same!

Response: Mngr Rating

RSquare	0.560
RSquare Adj	0.551
Root Mean Square Error	0.995
Mean of Response	5.901
Observations (or Sum Wgts)	150

Parameter Estimates

Term	Estimate	Std Error	t Ratio	Prob>|t|	VIF
Intercept	-2.948	0.695	-4.24	<.0001	0.0
Origin[Ext - Int]	-0.014	0.109	-0.13	0.8979	1.8
Salary	0.110	0.009	12.22	<.0001	1.4
Years Exp	0.120	0.025	4.88	<.0001	1.4