

Wharton School
Wharton Research Scholars Journal

University of Pennsylvania

Year 2004

Simpson's Paradox – A Survey of Past,
Present and Future Research

Yao Hua Ooi
University of Pennsylvania



Simpson's Paradox – A Survey of Past, Present and Future Research

By –Yao Hua Ooi
Advisor – J. Michael Steele
Wharton Undergraduate Research Scholars
WH-299-301
April 2004

Abstract

Simpson's paradox refers to the reversal of a statistical relationship between two variables in sub-populations when the sub-populations are combined and analyzed as a population. This article is intended to provide a broad survey of the past, present and future research surrounding the issue. Real data from a discrimination litigation case is examined to identify the occurrence of the paradox.

Section 1

Introduction

Simpson's Paradox was first discovered by Karl Pearson in 1899. The paradox derives its name from "The Interpretation of Interaction in Contingency Tables" [1]. Simpson's Paradox refers to the reversal of a statistical relationship between two variables in sub-populations when the sub-populations are aggregated and analyzed as a population. Alternatively, Simpson's Paradox can be thought of as the reversal of a statistical relationship between two variables when additional factors are added in the analysis.

Blyth states Simpson's Reversal Paradox as follows [2].

It is possible to have:

$$P(A | B) > P(A | B')$$

and have at the same time both: (1)

$$P(A | B \text{ and } C) \leq P(A | B' \text{ and } C)$$

$$P(A | B \text{ and } C') < P(A | B' \text{ and } C'),$$

where $P(A | B)$ is the probability of event A conditional on event B and the prime indicates complements

First consider a fictitious example from [2] that illustrates the paradox:

A doctor in California conducted a new treatment on patients in two different cities, California (C) and Seattle (C'). As each patient from California became available, he assigned him to the new treatment with probability 0.91, leaving him to the standard treatment with probability 0.09. In the same way, he assigned each patient in New Jersey to the new treatment with probability 0.01, leaving him to the standard treatment with probability 0.99. (These probabilities were expected to give him about the number of patients he could handle in each city). Observing the survival rate on the patients in both cities, the following results were obtained:

Table 1
Number of Patients

Outcome	Treatment	
	Standard	New
Died	5950 (54%)	9005 (89%)
Survived	5050 (46%)	1095 (11%)
Total	11000 (100%)	10100 (100%)

Based the results in Table 1, it appears that the new treatment is clearly inferior to the standard treatment, since only 11 percent of the patients receiving the new treatment survived, compared to 46 percent under the standard treatment. However, when the data is disaggregated for the two cities, the opposite conclusion is reached.

Table 2
Number of Patients

Outcome	Patient Type C		Patient Type C'	
	Treatment		Treatment	
	Standard	New	Standard	New
Died	950 (95%)	9000 (90%)	5000 (50%)	5 (5%)
Survived	50 (5%)	1000 (10%)	5000 (50%)	95 (95%)
Total	1000 (100%)	10000 (100%)	10000 (100%)	100 (100%)

From the data in Table 2, patients under the new treatment in California (C) had a survival rate of 10 percent compared to 5 percent for those under the standard treatment, while patients under the new treatment in Seattle (C') had a survival rate of 95 percent compared to 50 percent for those under the standard treatment. Thus, it appears that the new treatment is a superior treatment to the standard treatment, a conclusion contrary to the one derived from examining the aggregated data.

The paradox can be explained once we see what has occurred: the C patients are much less likely to recover, and since the new treatment was given mostly to C patients, the new treatment showed a low survival rate since it was tried out mostly on the most seriously ill patients. If event A is survival and B is the new treatment, we have

Simpson's Paradox as stated in (1):

$$P(A | B) = 0.11 < P(A | B') = 0.46$$

$$P(A | B \text{ and } C) = 0.10 > P(A | B' \text{ and } C) = 0.05$$

$$P(A | B \text{ and } C') = 0.95 > P(A | B' \text{ and } C') = 0.50$$

One tends to reason that this is intuitively impossible because:

$$P(A | B) = \text{An equally weighted average of } P(A | BC) \text{ and } P(A | BC')$$

$$P(A | B') = \text{An equally weighted average of } P(A | B'C) \text{ and } P(A | B'C')$$

Such an assumption is incorrect because the averages have different weightings given by the following equation:

$$P(A | B) = P(C | B) \cdot P(A | BC) + P(C' | B) \cdot P(A | BC') \tag{2}$$

$$P(A | B') = P(C | B') \cdot P(A | B'C) + P(C' | B') \cdot P(A | B'C')$$

In the above example,

$$P(C | B) = 100/101$$

$$P(C' | B) = 1/101$$

$$P(C | B') = 1/11$$

$$P(C' | B') = 10/11$$

Thus the surprising fact that an average of 0.10 and 0.95 is so much smaller than an average of 0.05 and 0.50 is easily explained by showing according to (2):

$$0.11 = 100/101 \cdot 0.10 + 1/101 \cdot 0.95$$

$$0.46 = 1/11 \cdot 0.05 + 10/11 \cdot 0.50$$

Here the paradox could not happen if B and C were independent, i.e., if the proportion receiving the new treatment were the same for C and C' patients. However, here the paradox results due to the dependence or the interaction of B and C.

A second example is given by Gardner [3] and is shown in the following table:

Table 3

	Box 1	Box 2	Box 3	Box 4	Box 1 + 3	Box 2 + 4
No. of black chips	5	3	6	9	11	12
No. of white chips	6	4	3	5	9	9

To maximize the likelihood of drawing a black chip, box 1 is preferred over box 2 (since $5/11 > 3/7$) and box 3 is preferred over box 4 (since $6/9 > 9/14$). Yet, when the contents of box 1 and 3 are combined and the contents of box 2 and box 4 are combined, the second combination is preferred over the first one (since $12/21 > 11/20$). The paradox can be stated in terms of positive numbers. It is possible to have:

$$\frac{a+b}{c+d} > \frac{e+f}{g+h}$$

and at the same time have both:

(3)

$$\frac{a}{c} \leq \frac{e}{g}$$

and:

$$\frac{b}{d} \leq \frac{f}{h}$$

The equivalence of (1) and (3) can be seen by setting:

$$a = P(C | B) \cdot P(A | BC)$$

$$b = P(C' | B) \cdot P(A | BC')$$

$$c = P(C | B)$$

$$d = P(C' | B)$$

$$e = P(C | B') \cdot P(A | B'C)$$

$$f = P(C' | B') \cdot P(A | B'C')$$

$$g = P(C | B')$$

$$h = P(C' | B')$$

Section 2 provides four real-life examples where Simpson's Paradox occurs. The examples, which encompass a broad range of fields, provide a motivation for further understanding the consequences of bias in sampling associated with the paradox. Section 3 discusses other literature related to Simpson's Paradox, including a geometric interpretation and alternative data presentation methods. In Section 4, analysis on a real-life data set from a gender discrimination case is presented in light of the paradox. Section 5 discusses Simpson's Paradox in a broader light as a topic in sampling and suggests several avenues for further research.

Section 2

Real-life examples of Simpson's Paradox

There have been numerous papers written to describe examples of Simpson's paradox, both using fictitious data [2,4,5] as well as real data [6,7,8,9,10,11]. In this section, four examples of Simpson's paradox in real-life are provided to illustrate the point that the phenomena occurs across sampling studies in different fields of study.

Simpson's Paradox in Jury Selection

The first example relates to a survey on jury composition in New Zealand by the New Zealand Department of Justice in September 1993 [6]. The issue of particular interest in the survey was the question of representation of the Maori tribe, the indigenous people of New Zealand. The original impetus for the research was a suspicion that Maori was under-represented on juries. The Department of Justice surveyed the composition of juries and the pool of the jurors the juries were selected from during a period in September and October 1993. Looking the results nationally, the report noted that "9.5 percent of people living within the jury districts were Maori. This compares with 10.1 percent of Maori in the pool of potential jurors. It is tempting to conclude, therefore, that Maori were adequately represented in the jury pool". Since the draft report gave the Maori proportion for each court district, Westbrook further analyzed the data using 1991 Census data and found the following results:

Table 4**Percentage Maori Ethnic Group**

District	Eligible Population (aged 20 – 64)	Jury Pool	Shortfall
Whangarei	17.0	16.8	0.2
Auckland	9.2	9.0	0.2
Hamilton	13.5	11.5	2.0
Rotorua	27.0	23.4	3.6
Gisborne	32.2	29.5	2.7
Napier	15.5	12.4	3.1
New Plymouth	8.9	4.1	4.8
Palmerston North	8.9	4.3	4.6
Wellington	8.7	7.5	1.2
Nelson	3.9	1.7	2.2
Christchurch	4.5	3.3	1.2
Dunedin	3.3	2.4	0.9
Invercargill	8.4	4.8	3.6
All Districts	9.5	10.1	-0.6

Source: Unpublished 1991 Census tables, and Trial By Peers

Looking at the disaggregated data, the Maori are under-represented in every district. However, when the data is aggregated at the national level, Maori is over-represented. The paradox arises because the jury pool size in each district is not proportional to the population of the cities. Certain districts have a much higher juror pool to population ratio than others. It turns out that those districts with a high proportion of Maori tend to have relatively large jury pools, and this pulls the proportion of Maori in the jury pool for the combined districts above the proportion of Maori in the population.

Example of Simpson’s Paradox when a Covariate is Ignored

The second example involves a study on smoking habits and survival rates [7]. Between 1972 and 1974 a one-in-six survey of the electoral roll was carried out in Whickham, a mixed urban and rural district near Newcastle upon Tyne, United Kingdom (Tunbridge et al., 1977). A follow-up study was conducted twenty years later

(Vanderpump et al. 1995). Appleton examines the study data, restricting the analysis to 1314 women smokers in the district for simplicity.

The following study results were obtained:

Table 5

Relationship Between Smoking Habits and 20-Year Survival in 1314 Women

	<i>Smoker</i>		
	Yes	No	Total
Dead	139	230	369
Alive	443	502	945
	582	732	1314
Survival Rate	76.1%	68.6%	

The results imply a significant protective effect of smoking because the survival rate for smokers is higher (76 percent) than the survival rate of non-smokers (69 percent).

The explanation for this baffling result is provided when an extra variable, age, which is strongly related to survival, is added to the analysis. Table 6 shows the effect of adding the variable to the analysis:

Table 6

Numbers of Women Smokers and Nonsmokers in Different Age Groups, Showing their 20-year Survival Status

Age	18 – 24		25 – 34		35 – 44		45 – 54		55 – 64		65 – 74		75+	
Smoker?	+	-	+	-	+	-	+	-	+	-	+	-	+	-
Dead	2	1	3	5	14	7	27	12	51	40	29	101	13	64
Alive	53	61	121	152	95	114	103	66	64	81	7	28	0	0
Survival Rate (%)	96	98	98	97	87	94	79	85	56	67	19	22	0	0

When examined by age group, the survival rate for non-smokers is higher in almost all cases. The paradox occurs as a higher proportion of non-smokers studied belong to older age groups, in which survival rates for both smokers and non-smokers are significantly lower compared to younger age groups. Thus the important covariate, namely age, was ignored in the aggregated data study, causing a misleading conclusion.

The study also suffers from a selection bias: the small proportion of older women smoking is likely to be due not only due to a low proportion in that cohort being smokers, but also to those who had smoked being less likely to survive to be seen in the original study.

Simpson's Paradox in Graduate Admissions

The third example of Simpson's paradox, involves a study of sex bias in graduate admissions at the University of California, Berkeley in 1973. Bickel, Hammel and O'Connell investigate student applications to graduate study at the University of California, Berkeley, for the fall 1973 quarter [8].

Two assumptions were made before delving into analysis of the data. The first assumption was that in any given discipline male and female applicants do not differ in respect of their intelligence, skill, qualifications, promise, or any other attribute deemed legitimately pertinent to their acceptance as students. This assumption is necessary for a meaningful study on gender bias, since otherwise the difference in acceptance of applicants by sex could be attributed to differences in their qualifications, intelligence, and so on. The second assumption is that the sex ratios of applicants to the various fields of graduate study are not importantly associated with any other factors in admission.

Based on these two assumptions, expected frequencies of males and female applicants admitted and denied were calculated on the assumption that male and female applicants have equal chances of admission to the university. The computation is shown in the following table:

Table 7. Decisions on applications to Graduate Division for fall 1973, by applicant gender

Applicants	Outcome				Difference	
	Observed		Expected			
	Admit	Deny	Admit	Deny	Admit	Deny
Men	3738	4704	3460.7	4981.3	277.3	-277.3
Women	1494	2827	1771.3	2549.7	-277.3	277.3

Based on this computation, 277 fewer women and 277 more men were admitted than would have been expected under the assumption noted. The chi-square value for this table is 110.8, confirming that the likelihood that the bias against women applicants observed appears due to chance alone is extremely low. Thus, we could conclude that a bias against female applicants existed in the fall 1973 admissions. Given the hypothesis that this bias existed, the researchers further examined to data to determine the responsible parties that made admission decisions. The outcome for an application for admission to graduate study is determined mainly by the faculty of the department to which each applicant applies. However, when a study was done at the departmental level, it was found that there were only 4 departments that had a bias against female applicants, accounting for a deficit of 26 in women admitted, and 6 departments that had a bias against male applicants, accounting for a deficit of 64 men. Analysis on the disaggregated data clearly contradicted with the findings of the aggregated data, leading the researchers to further investigate their assumptions and identify faults with the method of examination.

In performing their further analysis, they considered an alternative to aggregating the data across the 85 departments and computing a statistic. Instead, they computed a statistic for each department first and aggregated those using methods described by Fisher

in [12]. In using this approach, the evidence for campus-wide bias in favor of men is extremely weak. In fact, there is evidence of bias in favor of women.

The paradoxical results exist because the second assumption in the initial analysis was faulty. The fact is that not all departments are equally easy to enter. The odds of gaining admission to different departments are widely divergent. In addition, the odds of getting into each department are strongly associated with the tendency of men and women to apply to different departments in different degree. The departments that were hardest to get into had a higher proportion of women applicants while the departments that were the easiest to get into had a lower proportion of women applicants. Furthermore, this phenomenon is even more pronounced in departments with large number of applicants. Therefore, the apparent bias against women applicants overall is due to the fact that a much larger proportion of women apply to departments that have lower admission rates, rather than actual bias in admission policies among departments.

The authors then proceeded to reanalyze Table 7, using all the data leading to us, by estimating the number of women expected to be admitted to a department by multiplying the estimated probability of admission of any applicant (regardless of sex) to that department by the number of women applying to it. For example, if the chances of getting into a department were 50 percent for all applicants, and if 300 women applied to it, the expected number of women admitted would be 150, assuming there is no sex bias in admissions. By doing this computation for each department separately, and summing the results, the actual number of women admitted as a whole was 60 higher than the estimated number of women admitted, showing a slight bias toward admitting women applicants. The researchers proceeded to consider data for the entire campus from the

years 1969 to 1973 and found little evidence of bias of any kind, except for 1973, where there seemed to be a slight bias toward women.

Thus, in this example examination of aggregate data on graduate admissions to the University of California, Berkeley, for fall 1973 shows a clear but misleading pattern of bias against female applicants. If the data are properly pooled, taking into account the autonomy of the departmental decision-making, and correcting for the tendency of women applicants to apply to graduate departments that have a lower admission rate for applicants of either sex, there is a slight bias toward women applicants.

This example, as well as a number of others [9], show how researchers need to be wary of Simpson's Paradox when performing studies on issues of discrimination and bias, since the level of aggregation in data analysis can bias the conclusions derived from such studies.

Section 3

Recent Research on Simpson's Paradox

Besides articles describing real examples of Simpson's Paradox, there is also a body of literature devoted to discussing better ways to interpret Simpson's Paradox, as well as methods that can be used to avoid the misrepresentation of data associated with the paradox.

A Geometric Interpretation of Simpson's Paradox

A. Tan provides a geometric interpretation of Simpson's Paradox in [4]. Applying his geometric approach to illustrating Simpson's Paradox in our first example, we get the following figure:




Figure 1 A geometric interpretation of Simpson's Paradox

where:

B_r = survival rate of those receiving new treatment in California

B_1 = survival rate of those receiving new treatment in Seattle

A_r = survival rate of those receiving standard treatment in California

A_1 = survival rate of those receiving standard treatment in Seattle

b_r = proportion of California patients in study who receive new treatment

b_1 = proportion of California patients in study who receive standard treatment

a_r = proportion of Seattle patients in study who receive new treatment

a_1 = proportion of Seattle patients in study who receive standard treatment

A = survival rate for those receiving standard treatment

B = survival rate for those receiving new treatment

From the diagram, we observe that even though $A_r < B_r$ and $A_1 < B_1$, we see that $A > B$ (draw horizontal lines through A and B), due to the significant differences between a_r and b_r , the proportion of patients who receive the new treatment in Seattle and California respectively. From the diagram, we would conclude that the new treatment is superior to the standard treatment if the equal proportions of patients in both states receive the new treatment, i.e. if $a_r = b_r$. Simpson's Paradox occurs when the difference between a_r and b_r is large enough such that $A > B$ when $A_r < B_r$ and $A_1 < B_1$. The geometric interpretation of Simpson's Paradox provides extra information by showing the minimum difference between a_r and b_r needed as a condition for the reversal in association to exist. The difference is found for any given level of b_r by drawing a horizontal line through the point of intersection between the vertical line emerging from b_r on the horizontal axis and the diagonal line connecting B_r and B_1 . By examining the point where this horizontal line intersects the diagonal line connecting A_r and A_1 , we can then determine the value of minimum value of a_r on the horizontal axis that leads to the reversal in association.

Representation of Data

In addition, a number of methods have been suggested to avoid the problem of Simpson's Paradox when aggregating data over sub-populations. One such method is proposed by Westbrooke in relation to the New Zealand jury study given in Section 2 [6]. The original presentation of Table 4 did not incorporate the difference in the ratio of juror pools to population size. To overcome this problem, Westbrooke proposes presenting the expected number of Maori in the jury pool assuming the selection of random. When this approach is used, the following table is obtained:

Table 8
Actual Number of Maori in Jury Pool Compared with the Expected Number

	Actual	Expected
Whangarei	28	28
Auckland	74	76
Hamilton	23	27
Rotorua	79	91
Gisborne	23	25
Napier	15	19
New Plymouth	4	9
Palmerston North	7	14
Wellington	28	33
Nelson	1	2
Christchurch	111	15
Dunedin	4	6
Invercargill	3	5
All Districts	300	350

The expected number of Maori juror for each district is calculated by multiplying the size of the overall juror pool in the district by the proportion of Maori in the local population. The advantage of using this approach is that the number of jurors – actual and expected can be added to get an overall total that means something. The percentages used

before hid the real situation of Maori under-representation. These totals show under-representation of Maori at both the national level, as well as at each local level bar one.

Section 4

Simpson's Paradox an Employment Litigation Case

Description of Data

A real set of data was analyzed to identify the possibly occurrence of Simpson's Paradox. The data was used in the 2003 Penn Law School Moot Case Competition and involved the case Nussbaum v. Arkansas College of Medicine. Nussbaum was a woman pediatrician who sued the Arkansas College of Medicine claiming that as a woman, she had been discriminated against both in terms of salary and in being passed over for promotion. The plaintiff's argument hinged on the fact that female salaries were significantly lower compared to male salaries within the college on average. The average male salary in 1985 was \$194914, while the average female salary in 1985 was \$130877.

The variables in the dataset included information on the department each faculty member belonged too, their gender, whether they had a primarily research or clinical emphasis, board certification, publication rate, experience level, professorship rank and their salaries in 1984 and 1985. There were a total of 261 observations in the dataset.

Findings

Using simple regression, I first validated that the salaries of 1984 and 1985 are closely correlated (R² value of 99.8%). I then proceeded to perform analysis only on the 1985 salary data. Using an approach similar to the Berkeley admissions case, I performed analysis at the departmental level, assuming that departments were the main decision-makers in deciding salary levels.

Since there was no strong correlation between experience or publication rate and salaries, I performed simple comparisons without explicitly adjusting for experience levels and publication rates. Within each of the six departments, male salaries tended to be higher than female salaries on average. However, the difference between average male and female salary were of much smaller magnitude, compared to the difference in male and female salary when the data was aggregated as a whole. Although there is no obvious case of Simpson's Paradox in this case, since male salaries were higher than female salaries at both the disaggregated, departmental level, as well as the aggregated level, there is a "weak" Simpsonian effect, which I define as a reduction in the strength of association between two variables when the analysis is considered at an appropriate disaggregated level, that ameliorates the difference between male and female salaries when examined at the departmental level. The additional factor of analysis in this case is the fact that a much large proportion of men tend to work in higher-paying departments compared to women, thus making the difference in average salaries even more pronounced at the aggregate level. Even though no conclusive evidence against gender discrimination can be found in this case, the small number of conclusions, coupled with the unclear relationship between experience and salary levels complicate the task of drawing clear conclusions in the context of the litigation suit.

Section 5

Conclusions

Simpson's Paradox is a relevant topic of research since its occurrence has important implications in business surveys, litigation, social science research and the public policy domain. Present research focuses on methods of finding quantitative measures of data used in identifying the occurrence of the paradox, as well as methods of avoiding inconsistent results when aggregating data, specifically in the case of collapsing multi-dimensional contingency tables [13,14].

Haunsperger and Saari points out that the inconsistent behavior associated with reaching different conclusions at different levels of aggregation in data analysis occurs in many, if not most statistical decision processes [15]. They propose a simple theory to determine whether statistical decision processes admit such inconsistencies and to find data restrictions that avoid outcomes such as Simpson's Paradox.

Ultimately, the issue of bias and discrimination in any statistical study is more subtle than one might imagine, and great care must be taken to identify and minimize any prejudicial treatment in data analysis [8]. Appropriate corrective measures should then be taken to provide analysis from which meaningful conclusions can be drawn.

References:

1. Simpson, E.H. (1951). The Interpretation of Interaction in Contingency Tables. *Journal of Royal Statistical Society Ser.B*, 13, No. 2 (1951), 238-41
2. Blyth, Colin R. (1972). On Simpson's paradox and the sure-thing principle. *Journal of the American Statistical Association*, 67, 364-366
3. Gardner, M. (1976). Mathematic Games. *Scientific American* 234, no. 3: 119-24.
4. Tan, A. (1986). A geometric interpretation of Simpson's paradox. *The College Mathematics Journal*, 17, 340-341
5. Sunder, S. (1983). Simpson's reversal paradox and cost allocation. *Journal of Accounting Research*, 21, 222-233, 1983.
6. Westbrooke, Ian (1998). Simpson's paradox: An example in a New Zealand survey of jury composition. *Chance, New Directions for Statistics and Computers*, 11, 40-42
7. Appleton, David R., French, Joyce M. and Vanderpump, Mark P. J. (1996). Ignoring a covariate: An example of Simpson's paradox. *The American Statistician*, 50, 340-341
8. Bickel, P. J., Hammel, E. A., and O'Connell, J. W. (1975). Sex bias in graduate admissions: Data from Berkeley. *Science*, 187, 1975, 398-404.
9. Knapp, Thomas R. (1985). Instances of Simpson's paradox. *The College Mathematics Journal*, 16, 209-211
10. Wagner, Clifford H. (1982). Simpson's paradox in real life. *The American Statistician*, 36, 46-48
11. Meier, Paul and Dinardo, Lawrence C. (1990). Simpson's paradox in employment litigation. *ASA Proceedings of the Social Statistics Section*, 66-69. *American Statistical Association (Alexandria, VA)*

12. Fisher, R.A. (1932). *Statistical Methods for Research Workers*. Oliver and Boyd, London, ed.4
13. Shapiro, S. H.(1982). Collapsing contingency tables -- a geometric approach. *The American Statistician*, 36, 43-46, 1982.
14. Whittemore, A. S. (1978). Collapsibility of multi- dimensional contingency tables. *Journal of the Royal Statistical Society, Ser. B.*, 40, 328-340, 1978.
15. Haunsperger, D. B. and Saari, D. G. (1991). The lack of consistency for statistical decision procedures. *The American Statistician*, 45(3), 252-255, 1991.

Appendix:

Study of Nussbaum v. Arkansas College of Medicine

Key:

0 = Female

1 = Male




Figure 2 Regression between 1984 and 1985 salaries

— Linear Fit

Linear Fit

$$\text{Sal85} = -419.512 + 1.1024317 \text{ Sal84}$$

Summary of Fit

RSquare	0.998511
RSquare Adj	0.998505




Figure 3 Analysis of Salary Levels (Aggregated)

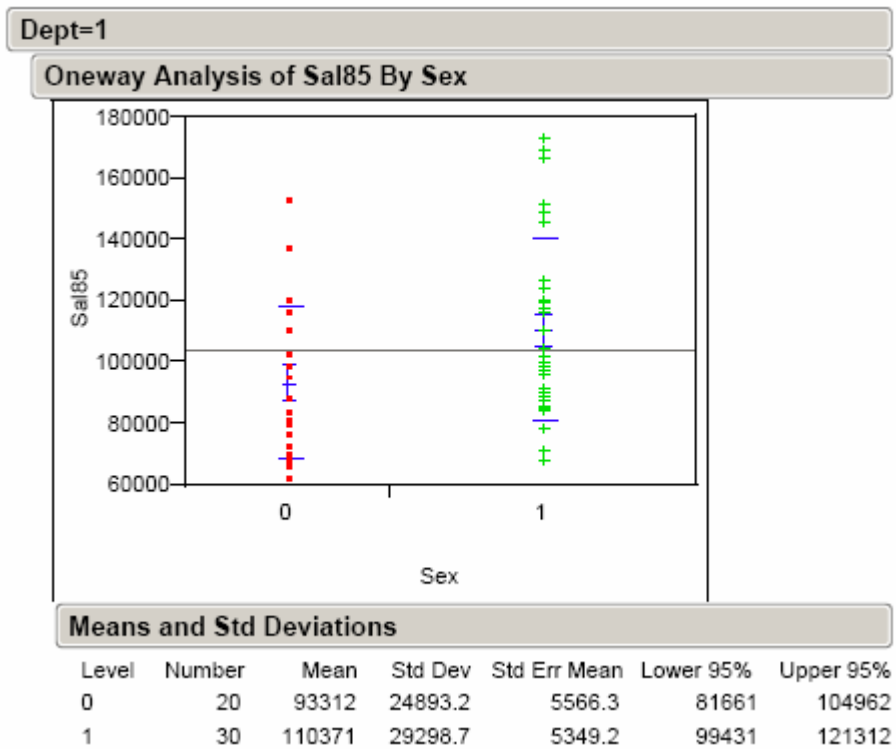



Figure 4 Analysis of Salary Levels for Department 1 (Biochemistry/Molecular Biology)

Dept=2

Oneway Analysis of Sal85 By Sex




Means and Std Deviations

Level	Number	Mean	Std Dev	Std Err Mean	Lower 95%	Upper 95%
0	20	77406	27043.6	6047.1	64749	90062
1	20	105031	31379.9	7016.8	90344	119717

Figure 5 Analysis of Salary Levels for Department 2 (Physiology)

Dept=3

Oneway Analysis of Sal85 By Sex



Means and Std Deviations

Level	Number	Mean	Std Dev	Std Err Mean	Lower 95%	Upper 95%
0	11	103769	30451.9	9182	83311	124227
1	10	150187	42958.0	13585	119457	180918

Figure 6 Analysis of Salary Levels for Department 3 (Genetics)

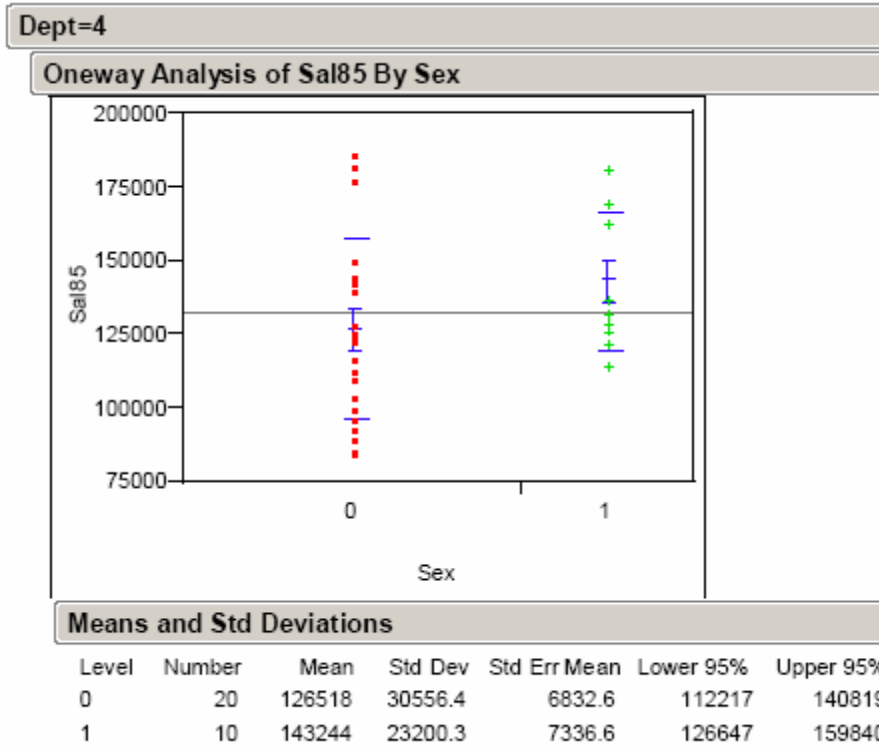


Figure 7 Analysis of Salary Levels for Department 4 (Pediatrics)




Figure 8 Analysis of Salary Levels for Department 5 (Medicine)




Figure 9 Analysis of Salary Levels for Department 6 (Surgery)





Figure 10 Comparison of Salaries across Departments



Contingency Analysis of Sex By Dept

Contingency Table

Dept	Sex		Count
	0	1	
1	20	30	50
	18.87	19.35	
2	20	20	40
	18.87	12.90	
3	11	10	21
	10.38	6.45	
4	20	10	30
	18.87	6.45	
5	30	50	80
	28.30	32.26	
6	5	35	40
	4.72	22.58	
	106	155	261

Figure 11 Percentage of Men and Women in each Department