"Simple" Sensitivity Analyses for Matched Samples

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

Goal of a Formal Sensitivity Analysis

To replace a general qualitative statement that applies in all observational studies…

"the association we observe between treatment and outcome does not imply causation"
"hidden biases can explain observed associations"

… with a quantitative statement that is specific to what is observed in a particular study…

"to explain the association seen in a particular study, one would need a hidden bias of a particular magnitude."

If the association is strong, the hidden bias needed to explain it would be large. If a study is free of hidden bias (main example: a carefully randomized trial), this means that any two units (patients, subjects, whatever) that appear similar in terms of their observed covariates actually have the same chance of assignment to treatment. There is *hidden bias* if two units with the same observed covariates have different chances of receiving the treatment. A sensitivity analysis asks: How would inferences about treatment effects be altered by hidden biases of various magnitudes? How large would these differences have to be to alter the qualitative conclusions of the study?

The Sensitivity Parameter, Γ

Suppose we have two units (subjects, patients), say, j and k, with the same observed covariate values $\mathbf{x}$ but different probabilities p of treatment assignment (possibly due to some unobserved covariate), so that $\mathbf{x}_{[j]} = \mathbf{x}_{[k]}$ but that possibly $p_{[j]} \neq p_{[k]}$.

Units j and k might be matched to form a matched pair in our attempt to control overt bias due to the covariates $\mathbf{x}$. The odds that units j and k receive the treatment are, respectively,

$\dfrac{p_{[j]}}{1 - p_{[j]}}$ and $\dfrac{p_{[k]}}{1 - p_{[k]}}$, and the odds ratio is the ratio of these odds.

Imagine that we knew that this odds ratio for units with the same $\mathbf{x}$ was at most some number Γ, so that $\Gamma \geq 1$. That is,

$$\frac{1}{\Gamma} \leq \frac{p_{[j]}\left(1 - p_{[k]}\right)}{p_{[k]}\left(1 - p_{[j]}\right)} \leq \Gamma$$

We call Γ the sensitivity parameter, and it forms the basis for our sensitivity analyses. If $\Gamma = 1$, then $p_{[j]} = p_{[k]}$ whenever $\mathbf{x}_{[j]} = \mathbf{x}_{[k]}$, so the study would be free of hidden bias, and standard statistical methods designed for randomized trials would apply.

If $\Gamma = 2$, then two units who appear similar, who have the same set of observed covariates **x**, could differ in their odds of receiving the treatment by as much as a factor of 2, so that one could be twice as likely as the other to receive the treatment.

In other words, Γ is a measure of the **degree of departure** from a study that is free of hidden bias. A sensitivity analysis will consider possible values of Γ and show how the inference might change. A study is *sensitive* if values of Γ close to 1 could lead to inferences that are very different from those obtained assuming the study is free of hidden bias. A study is *insensitive* if extreme values of Γ are required to alter the inference.

<u>Scenario 1</u>.   A Binary Outcome – A Sensitivity Analysis for McNemar's Test

Exposure: Heavy Smoker vs. Nonsmoker;
Outcome: Death due to Lung Cancer (no censoring)

Suppose we paired 1000 heavy smokers to 1000 nonsmokers on the basis of a series of baseline characteristics (without using propensity methods, but that doesn't matter here).

Totally fake data follow:

	Heavy Smoker Dies from Lung Cancer	Heavy Smoker Doesn't Die from Lung Cancer	Total
Nonsmoker Dies from Lung Cancer	175	**12**	187
Nonsmoker Doesn't Die from Lung Cancer	**110**	703	813
Total	285	715	1000

Of the $S = 1000$ matched pairs, suppose that there were 122 pairs in which exactly one person died of lung cancer. Of these, there were 12 pairs in which the nonsmoker died of lung cancer, and 110 pairs in which the heavy smoker died of lung cancer. In other words, 122 of the pairs are discordant for death from lung cancer.

If the study were a randomized experiment, or if it was an observational study free of any hidden bias (neither of which are true), then we'd use McNemar's test to compare the 110 lung cancer deaths among smokers to a binomial distribution with 122 trials and probability of "success" ½ to yield a significance level of $p < 0.0001$. In R, we'd have…

```
> Temp <- matrix(c(175,12,110,703), nrow=2)
> Temp
        [,1] [,2]
   [1,]  175  110
   [2,]   12  703
> mcnemar.test(Temp, correct=F)
McNemar's Chi-squared test
data:  Temp
McNemar's chi-squared = 78.7213, df = 1, p-value < 2.2e-16
```

"Simple" Sensitivity Analyses for Matched Samples

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

In the absence of hidden bias, there would be strong evidence that smoking causes lung cancer. **How much hidden bias would need to be present to alter this conclusion?**

Let T = # of the 122 discordant pairs in which the heavy smoker died of lung cancer. Here T = 110.

Here are the key formulae for this scenario:

Any particular value of the sensitivity parameter Γ determines the values for p^+ and p^-, specifically, $p^+ = \Gamma / (1+ \Gamma)$ and $p^- = 1 / (1 + \Gamma)$. We then use these p^+ and p^- values to determine the bounds on the p value for the McNemar statistic for varying values of Γ.

The formulae are identical for the two bounds, except for the p^+ and p^- substitution…

$$\text{Upper Bound for } p \text{ value is } \sum_{a=110}^{122} \binom{122}{a} \left(p^+\right)^a \left(1-p^+\right)^{122-a}$$

$$\text{Lower Bound for } p \text{ value is } \sum_{a=110}^{122} \binom{122}{a} \left(p^-\right)^a \left(1-p^-\right)^{122-a}$$

While these formulae look troublingly complicated at first, they are just the binomial probabilities of obtaining a value of T = 110 or higher assuming a binomial distribution with n = 122 trials and common probability $p = p^+$ (for upper bound) or $p = p^-$ (for lower bound). Happily, Excel, R, SAS (or any other useful statistical software) can find these two probabilities. Below, I'm using an Excel sheet to demonstrate (which is a little imprecise – as is the case for most of Excel's statistical functions, but lands you in the right ballpark and is easy to use.) The first part of the sheet (I'll show you the whole thing shortly) is shown below…

	A	B	C	D	E	F
1	Sensitivity Analysis for McNemar's Test: Simplied Formula					
2	Section 4.3.2. of Rosenbaum PR (2002) Observational Studies, 2nd Edition.					
3	INSERT VALUES (IN RED) IN CELLS HIGHLIGHTED IN YELLOW.					
4	Two-By-Two Table	Treated, outcome = Yes	Treated, outcome = No			
5	Control, outcome = Yes	175	12	187		
6	Control, outcome = No	110	703	813		
7		285	715	1000		
8						
9	Computed Summaries					
10	# of Pairs	1000	# of matched pairs (overall)			
11	# of Discordant Pairs	122	# of matched pairs in which exactly one has the outcome			
12	Test Statistic	110	# of discordant pairs where Treated has outcome			

In this case, I've already typed in the correct values for the 2 by 2 table based on outcomes and exposures in the matched pairs.
- Remember that you should have one observation in this 2x2 table for each **pair** of subjects. Here, we have 1000 pairs.

"Simple" Sensitivity Analyses for Matched Samples

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

- We have 122 discordant pairs (i.e. pairs in which exactly one of the two pair members has outcome = Yes) – note that the discordant pairs are just the off-diagonal cells here.
- The test statistic, which is the number of pairs in which the Treated's outcome is Yes but the Control's outcome is No, is 110.

The remainder of the sheet does the actual set of sensitivity analysis calculations. Here's the whole thing…

In the **Sensitivity Analysis** section, the spreadsheet automatically completes the calculations for gamma values between 1.0 (i.e. no hidden bias) and 6.0, stepping by 0.5. We see, for instance, the *p* value assuming no hidden bias is 0, to four decimal places (this is obtained from the gamma = 1.0 row). The sensitivity analysis tips over significance at the two-tailed α = 0.05 level somewhere between Γ = 5.0 and Γ = 5.5.

To isolate the correct value of the sensitivity parameter to greater detail, you can use the insert gamma value below cell (row 29) and twiddle the value of gamma until the upper bound for the *p* value is just under 0.05. This case implies a Γ threshold of about 5.43.

The conclusion here (stating specifically that $\Gamma > 5$ but < 6) would be something like:
To attribute the higher rate of *death from lung cancer* to an unobserved covariate rather than to an effect of *smoking*, that unobserved covariate would need to produce more than a <u>fivefold</u> increase in the odds of *smoking*, and it would need to be a near perfect predictor of *lung cancer*.

As we shall see by comparison in later examples, this is a high degree of insensitivity to hidden bias. In many other studies, biases smaller than $\Gamma = 5$ could explain the association between treatment and outcome.

> **Now suppose that there were 20 pairs (instead of 110) where the treated patient died but the control patient did not. What is the impact of this change?**

Scenario 2. A Continuous Outcome – A Sensitivity Analysis for the Wilcoxon Signed Rank Test

Exposure: Parent working (or not) in a factory where lead was used to make batteries;
Outcome: level of lead found in the child's blood (in µg/dl of whole blood)

Morton et al. (1982) studied lead in the blood of 33 kids (from different families) whose parents worked in a factory where lead was used in making batteries. The covariate **x** was two-dimensional, recording age and neighborhood of residence. They matched each exposed child to one control child of the same age and neighborhood whose parents were employed in other industries not using lead.

Lead in Children's Blood (µg/dl)

Pair	Exposed	Control	Difference	Rank	Pair	Exposed	Control	Difference	Rank
1	38	16	22	22	18	10	13	-3	5.5
2	23	18	5	8	19	45	9	36	30
3	41	18	23	23.5	20	39	14	25	26
4	18	24	-6	9.5	21	22	21	1	2.5
5	37	19	18	21	22	35	19	16	18.5
6	36	11	25	26	23	49	7	42	31
7	23	10	13	14	24	48	18	30	28
8	62	15	47	32	25	44	19	25	26
9	31	16	15	17	26	35	12	23	23.5
10	34	18	16	18.5	27	43	11	32	29
11	24	18	6	9.5	28	39	22	17	20
12	14	13	1	2.5	29	34	25	9	12.5
13	21	19	2	4	30	13	16	-3	5.5
14	17	10	7	11	31	73	13	60	33
15	16	16	0	1	32	25	11	14	15.5
16	20	16	4	7	33	27	13	14	15.5
17	15	24	-9	12.5					

The Wilcoxon signed rank statistic for S matched pairs is computed by ranking the absolute value of the differences within each pair from 1 to S, and then summing the ranks of the pairs where the exposed unit had a higher response than the matched control.

"Simple" Sensitivity Analyses for Matched Samples

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

Pair	Exposed	Control	Difference	Rank	Pair	Exposed	Control	Difference	Rank
1	38	16	22	22	18	10	13	-3	5.5
2	23	18	5	8	19	45	9	36	30
3	41	18	23	23.5	20	39	14	25	26
4	18	24	-6	9.5	21	22	21	1	2.5
5	37	19	18	21	22	35	19	16	18.5
6	36	11	25	26	23	49	7	42	31
7	23	10	13	14	24	48	18	30	28
8	62	15	47	32	25	44	19	25	26
9	31	16	15	17	26	35	12	23	23.5
10	34	18	16	18.5	27	43	11	32	29
11	24	18	6	9.5	28	39	22	17	20
12	14	13	1	2.5	29	34	25	9	12.5
13	21	19	2	4	30	13	16	-3	5.5
14	17	10	7	11	31	73	13	60	33
15	16	16	0	1	32	25	11	14	15.5
16	20	16	4	7	33	27	13	14	15.5
17	15	24	-9	12.5					

In this case[1], we sum up the ranks associated with the shaded pairs above (i.e. 22 + 8 + 23.5 + 21 + …) and conclude that the Wilcoxon signed rank statistic for the differences in children's lead levels is 527.

It turns out that if there are no ties among the absolute differences and no zero differences and no hidden bias, then the expectation and variance of this test statistic is known, and we appeal to a standard normal distribution to evaluate our test statistic.

$$\text{Here, } E(T) = \frac{S(S+1)}{4} = \frac{33(34)}{4} = 280.5 \text{ and}$$

$$\text{var}(T) = \frac{S(S+1)(2S+1)}{24} = \frac{33(34)(67)}{24} = 3132.25, \text{ so that } SD(T) = \sqrt{3132.25} = 55.97$$

We then calculate a Z statistic, as $Z = \frac{T - E(T)}{SD(T)} = \frac{527 - 280.5}{55.97} = 4.40$, and we can

compare this to a standard normal table[2] to get a two-tailed p value < 0.0001. Alternatively, we could simply let R, or SAS, or some other package complete the calculations.

At any rate, **if the study was free of hidden bias**, this would constitute strong evidence ($p < 0.0001$) of an effect of parental exposure to lead on children's lead levels.

---

[1] Note that the easier thing to do here would be to subtract off the total of the ranks associated with the unshaded pairs (those in which the exposed child DID NOT have more lead in their blood), and then subtract this from the total sum of the integers from 1 to S, which is always just S(S+1)/2. Note that 9.5 + 1 + 12.5 + 5.5 + 5.5 = 34, and 33*34/2 = 561, so that we again get 527 for the Wilcoxon statistic.

[2] To do this in Excel, the formula you need is **=2*(1-NORMSDIST(Z))** where Z is the value of interest.

Now, of course, there were some tied differences (in which average ranks were used) and there was one zero difference. Thus, the null expectation and variance expressions above are a little bit off, but that difference doesn't affect the conclusion here at all.

We've seen that if the study were free of hidden bias, that is, if $\Gamma = 1$, then there would be strong evidence that parents' occupational exposures to lead increased the level of lead in their children's blood. The sensitivity analysis we'll conduct now asks how this conclusion might be changed by hidden biases of various magnitudes.

To establish this sensitivity analysis, we need the following key formulae:

$$E\left(T^+\right) = \frac{p^+ S\left(S+1\right)}{2} \text{ and } var\left(T^+\right) = p^+\left(1-p^+\right)\frac{S\left(S+1\right)\left(2S+1\right)}{6}, \text{ where } p^+ = \frac{\Gamma}{\Gamma+1}$$

With $p^- = \frac{1}{\Gamma+1}$ in place of p^+, we get the same expectation and variance for T^- ...

$$E\left(T^-\right) = \frac{p^- S\left(S+1\right)}{2} \text{ and } var\left(T^-\right) = p^-\left(1-p^-\right)\frac{S\left(S+1\right)\left(2S+1\right)}{6}, \text{ where } p^- = \frac{1}{\Gamma+1}$$

Note that if $\Gamma = 1$, then $p^+ = p^- = \frac{1}{2}$, and the expressions revert to the standard formula for the signed rank test shown above.

As an example, suppose $\Gamma = 2$, so that matched children might differ in their odds of exposure to lead by a factor of 2 due to hidden bias.

$$\text{We have } p^- = \frac{1}{\Gamma+1} = \frac{1}{3}, E\left(T^-\right) = \frac{\frac{1}{3}\left(33\right)\left(33+1\right)}{2} = 187 \text{ and}$$

$$var\left(T^-\right) = \frac{1}{3}\left(1-\frac{1}{3}\right)\frac{33\left(33+1\right)\left(66+1\right)}{6} = 2784.22, \text{ so that } SD\left(T^-\right) = \sqrt{2784.22} = 52.77$$

$$\text{and } p^+ = \frac{\Gamma}{\Gamma+1} = \frac{2}{3}, \text{ so } E\left(T^+\right) = \frac{\frac{2}{3}\left(33\right)\left(33+1\right)}{2} = 374, \text{ and}$$

$$var\left(T^+\right) = \frac{2}{3}\left(1-\frac{2}{3}\right)\frac{33\left(33+1\right)\left(66+1\right)}{6} = 2784.22, \text{ so that } SD\left(T^+\right) = \sqrt{2784.22} = 52.77$$

$SD\left(T^-\right) = SD\left(T^+\right)$, as will be true for any particular Γ, so we can calculate either alone.

Now, the standard normal deviates (Z-scores) are:

$$Z^- = \frac{T - E\left(T^-\right)}{SD\left(T^-\right)} = \frac{527 - 187}{52.77} = 6.44 \text{ and } Z^+ = \frac{T - E\left(T^+\right)}{SD\left(T^+\right)} = \frac{527 - 374}{52.77} = 2.90$$

And the relevant p value range is therefore from less than 0.0001 (associated with 6.44) to 0.0037 (associated with 2.90). A hidden bias of size $\Gamma = 2$ is insufficient to explain the observed difference between exposed and control children.

"Simple" Sensitivity Analyses for Matched Samples

Yes, of course, I have a spreadsheet to do the sensitivity analysis calculation here…

Microsoft Excel - sensitivityspreadsheet.xls

File Edit View Insert Format Tools Data Window Help Adobe PDF   binomial distribution

Arial   10   B I U $ %

A1   Sensitivity Analysis for McNemar's Test: Simpled Formula

	A	B	C	D	E	F	G	H
1	Sensitivity Analysis for McNemar's Test: Simpled Formula							
2	Section 4.3.3. of Rosenbaum PR (2002) Observational Studies, 2nd Edition.							
3	INSERT VALUES (IN RED) IN CELLS HIGHLIGHTED IN YELLOW.							
4								
5	Data							
6	Total # of Pairs (S)	33						
7	Wilcoxon Test Statistic, T	527	sum of ranks of absolute differences for pairs where exposed had higher value than control					
8								
9	Sensitivity Analysis							
10	Gamma Values	2-tail P value (lower bound)	2-tail P value (upper bound)	P+	P-	E(T+)	E(T-)	SD(T+)
11	1.0	0.0000	0.0000	0.500	0.500	280.50	280.50	55.97
12	1.5	0.0000	0.0005	0.400	0.600	224.40	336.60	54.84
13	2.0	0.0000	0.0037	0.333	0.667	187.00	374.00	52.77
14	2.5	0.0000	0.0125	0.286	0.714	160.29	400.71	50.57
15	3.0	0.0000	0.0284	0.250	0.750	140.25	420.75	48.47
16	3.5	0.0000	0.0514	0.222	0.778	124.67	436.33	46.53
17	4.0	0.0000	0.0807	0.200	0.800	112.20	448.80	44.77
18	4.5	0.0000	0.1152	0.182	0.818	102.00	459.00	43.17
19	5.0	0.0000	0.1538	0.167	0.833	93.50	467.50	41.71
20	5.5	0.0000	0.1952	0.154	0.846	86.31	474.69	40.39
21	6.0	0.0000	0.2388	0.143	0.857	80.14	480.86	39.17
22								
23	Insert Gamma Value Below	2-tail P value (lower bound)	2-tail P value (upper bound)	P+	P-	E(T+)	E(T-)	SD(T+)
24	4.286	0.0000	0.0999	0.189	0.811	106.13	454.87	43.84
25	Stop when value for the upper bound of the P value (cell C24) is just below desired two-tailed significance level							
26								

Binary Outcome - McNemar \ Continuous Outcome - Signed Rk / Sheet3 /

Ready   NUM

The tipping point for the sensitivity parameter is a little over 4.25. To explain away the observed association between parental exposure to lead and child's lead level, a hidden bias or unobserved covariate would need to increase the odds of exposure by more than a factor of $\Gamma = 4.25$. The association cannot be attributed to small hidden biases, but it is somewhat more sensitive to bias than the study of heavy smokers in Scenario 1.

Note that all you need to insert here is the number of matched pairs and the Wilcoxon test statistic. If you have a big data set, getting the Wilcoxon statistic will be the most time-consuming issue…

Scenario 3.   A Censored Survival Outcome – A Sensitivity Analysis

Exposure: Either Normal or Low Serum Potassium Level in the DIG trial (HF Pts)
 Outcome: All cause mortality during the follow-up period (i.e. there is censoring)

In this case, we're doing a secondary data analysis of the DIG trial, to study heart failure patients with either normal or low serum potassium levels. We identified 1187 matched pairs of a normal potassium and a low potassium HF patient with similar baseline characteristics at admission.

The main complicating issue here is that many of the patients were censored (they dropped out of the study, or they survived to the end of the follow-up period). I'll spare you the formulae this time. The tricky part is conceptually simple, but computationally

"Simple" Sensitivity Analyses for Matched Samples

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

time-consuming. You need to count the number of pairs with a clear "winner" (subject who survives longer) and then determine the "winner" for each of those pairs.

Pair	Exposed Subject		Control Subject		Clear Winner?
	Survival Time	Censored?	Survival Time	Censored?	
1	200	No	100	No	Exposed
2	200	Yes	100	No	Exposed
3	200	No	100	Yes	Unknown
4	200	Yes	100	Yes	Unknown
5	100	No	200	No	Control
6	100	Yes	200	No	Unknown
7	100	No	200	Yes	Control
8	100	Yes	200	Yes	Unknown
9	100	Either	100	Either	Unknown

In this set of 9 pairs, we have four clear winners, 2 Exposed and 2 Control.

The inputs for the spreadsheet below are:
- The Number of Pairs in which the "Winner" (longer survival time) can be conclusively determined.
- # of pairs with a clear "winner" in which the patient with normal potassium outlived the patient with low potassium

In the specific scenario, there were 440 pairs with a clear "winner" in terms of all-cause mortality, and in 335 of these pairs the "winner" was the patient with normal potassium.

The tipping point for the sensitivity parameter is a little over $\Gamma = 2.5$.

CRSP 500 – Spring 2008 – Thomas E. Love, Ph. D., Instructor

Exercises for Discussion in Class (1 & 2 from Rosenbaum Chapter 4)

1. Many drugs used to treat cancer are quite harsh, and there is the possibility that these drugs can harm hospital workers who are exposed by accident. Kevekordes, Gebel, Hellwig, Dames and Dunkelberg (1998) studied this possibility when a "malfunction of a safety hood result[ed] in air flowing from the hood along the arms of the person preparing infusions of antineoplastic drugs." They studied 10 nurses who may have experienced substantial exposures, matching each nurse to a control based on gender, age, and intensity of smoking. They measured genetic damage using the cytokinesis block micronucleus test, reporting mean micronuclei/10^3 binucleate lymphocytes (mm/10^3), as follows. Perform an appropriate sensitivity analysis.

Pair	Ages	Smoking	Exposed	Control
1	37/37	NonS	20	11
2	24/25	S	10	9
3	33/32	S	22	19
4	29/30	NonS	13	9
5	23/23	NonS	13	7
6	28/29	NonS	14	11
7	25/24	NonS	12	6
8	38/40	S	21	23
9	32/32	NonS	9	4
10	33/34	S	21	14

2. Starting with 10,872 death certificates with the diagnosis of sporadic motor neuron disease (MND), Graham, Macdonald and Hawkes (1997) examined their birth certificates and found that 70 of these cases with MND had a living twin free of MND. These 70 twin pairs formed the basis for a case-referent study. Because little is known about the causes of MND, they examined "many variables" as potential causes in their explanatory study. The strongest association they found was with "carrying out car or vehicle maintenance." There were 16 twin pairs discordant for this variable, and 14/16 had an exposed case, while 2/16 had a referred referent, yielding an estimated odds ratio of 14/2 = 7 in the absence of hidden bias. Do a sensitivity analysis for the significance level using an appropriate test.

3. In the study described previously as part of Scenario 3, Ahmed, Love et al. studied the impact of low vs. normal serum potassium on two additional outcomes, specifically cardiovascular mortality, and heart failure mortality. Of the pairs with "clear winners" in terms of cardiovascular mortality, 274 had normal potassium winners, while 77 had low potassium winners. In terms of heart failure mortality, the pairs with winners broke down as 135 for normal potassium and 27 for low potassium. What conclusions can be drawn?