



Simpson's Paradox

January 10th, 2022

What is Simpson's Paradox (SP)?

- A phenomenon in probability and statistics in which a trend appears in several groups of data but disappears or reverses when the groups are combined
- Admission figures for the fall of 1973 showed that men applying were more likely than women to be admitted
- Women tended to apply to more competitive departments with lower rates of admission, whereas men tended to apply to less competitive departments with higher rates of admission (such as engineering)

Smokers in Whickham

- Between 1972 and 1974, a survey carried out in Whickham, a mixed urban and rural district near Newcastle upon Tyne, United Kingdom
- A follow-up study was conducted 20 years later
- 1,314 who were classified either as current smokers or as never having smoked
- A 20-year survival status was determined for all the women in the original survey.

Survival Rates

- 24% (139/582) of smokers died
- 31% (230/732) of nonsmokers died
- A significant protective effect of smoking?

	<i>Smoker</i>		<i>Total</i>
	<i>Yes</i>	<i>No</i>	
<i>Dead</i>	139	230	369
<i>Alive</i>	443	502	945
	582	732	1,314

$\chi^2 = 9.12$ on 1 df; $P = .0025$.
Odds ratio = .68 (95% confidence limits .53–.88).

Age!

- Few of the older women (over 65 at the original survey) were smokers, but many of them had died by the time of follow-up
- Why? Possible explanations
 - Smoking wasn't as popular for that age group
 - those who had smoke are less likely to survive to be seen in the original study

	Age group													
	18–24 Smoker		25–34 Smoker		35–44 Smoker		45–54 Smoker		55–64 Smoker		65–74 Smoker		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
$\hat{\psi}$	2.30		0.75		2.40		1.44		1.61		1.15		—	

NOTE: For each age group the odds ratio ($\hat{\psi}$) is given for a smoker dying relative to a nonsmoker.

The Prisoner's Dilemma

- Two members of a criminal organization are arrested and imprisoned
- They can't communicate
- The prosecutors lack sufficient evidence to convict the pair on the principal charge, but they have enough to convict both on a lesser charge
- The prosecutors offer each prisoner a bargain. Each prisoner is given the opportunity either to betray the other by testifying that the other committed the crime, or to cooperate with the other by remaining silent
- The possible outcomes are:
 1. If A and B each betray the other, each of them serves two years in prison
 2. If A betrays B but B remains silent, A will be set free and B will serve three years in prison
 3. If A remains silent but B betrays A, A will serve three years in prison and B will be set free
 4. If A and B both remain silent, both will serve only one year in prison (on the lesser charge).

The Prisoner's Dilemma

- From A's perspective
 - B stays silent, best choice is to betray (defect)
 - B betrays (defects), best choice is to betray (defect)

A \ B	B stays silent	B betrays
A stays silent	-1, -1	-3, 0
A betrays	0, -3	-2, -2

There is a game theory course in the Econ Dept (ECN 122) that can be used as Computer/Math Elective credit for the CS major

“Real Life”

- People tend to be kind to others in those situations in which people tend to be kind to them
- There are some situations that encourage cooperation (within a group)
 - Nice situations
- There are some situations where people do not tend to co-operate (between groups)
 - Nasty situations
- Because all players are influenced by whether they are in nice or nasty situation, their behavior will be correlated
- Overall, people will tend both to cooperate, and be cooperated with, in nice situations; but the reverse in nasty situations

Reinforcement Learning

- Reward positive outcomes
- Punish undesired outcomes
- Learn based on experience
- Reinforcement learning methods are based on the average amount of payoff that each behavior receives

Reinforcement Learning

- Average reward is higher for defecting
 - Reinforcement learning will teach agent to defect

	Person 1	Person 2
CC	5	5
CD	0	7
DC	7	0
DD	3	3

C Avg	2.5
D Avg	5

Simpson's Paradox

- If agents' responses are positively correlated, this means that CC and DD outcomes are the most common
- When D is played, it is typically associated with a low outcome (where both prisoners defect)
- When C is played, the pay-off is the reasonably high CC pay-off
- C is associated on average with somewhat less than the CC pay-off
- D is associated on average with somewhat more than the DD pay-off
- Average reward is higher for cooperating

	Person 1	Person 2	Biased Rounds
CC	5	5	50
CD	0	7	10
DC	7	0	10
DD	3	3	50

C Avg	4.16666667
D Avg	3.66666667

Simpson's Paradox

- Average payoff for C is 4.17
- Average payoff for D is 3.67
- Reinforcement learner will play C
- "Rational" move is still always to play D
 - You will always do better playing D

	Person 1	Person 2	Biased Rounds
CC	5	5	50
CD	0	7	10
DC	7	0	10
DD	3	3	50

C Avg	4.16666667
D Avg	3.66666667

Sources

Chater, N., Vlaev, I., & Grinberg, M. (2008). A new consequence of Simpson's paradox: Stable cooperation in one-shot prisoner's dilemma from populations of individualistic learners. *Journal of Experimental Psychology: General*, 137(3), 403.

Appleton, D. R., French, J. M., & Vanderpump, M. P. J. (1996). Ignoring a covariate: An example of Simpson's paradox. *The American Statistician*, 50(4), 340-341. Retrieved from <https://www.proquest.com/scholarly-journals/ignoring-covariate-example-simpsons-paradox/docview/228470284/se-2?accountid=14505>

https://en.wikipedia.org/wiki/Prisoner%27s_dilemma

https://en.wikipedia.org/wiki/Simpson's_paradox