UCLA STAT 13

Introduction to Statistical Methods for the Life and Health Sciences

•**Instructor:** **Ivo Dinov,**
Asst. Prof. of Statistics and Neurology

•**Teaching Assistants:** **Chris Barr & Ming Zheng**

University of California, Los Angeles, Fall 2004

http://www.stat.ucla.edu/~dinov/courses_students.html

Chapter 4: Probabilities and Proportions

- Where do probabilities come from?
- Simple probability models
- probability rules
- Conditional probability
- Statistical independence

Let's Make a Deal Paradox – aka, Monty Hall 3-door problem

- This paradox is related to a popular television show in the 1970's. In the show, a contestant was given a choice of three doors/cards of which one contained a prize (diamond). The other two doors contained gag gifts like a chicken or a donkey (clubs).

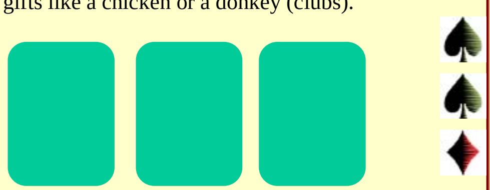

Let's Make a Deal Paradox.

- After the contestant chose an initial door, the host of the show then revealed an empty door among the two unchosen doors, and asks the contestant if he or she would like to switch to the other unchosen door. The question is should the contestant switch. Do the odds of winning increase by switching to the remaining door?

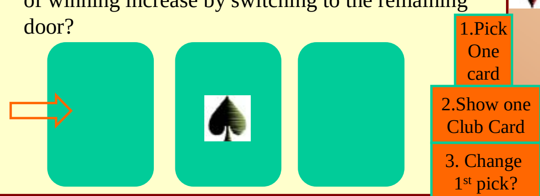

Let's Make a Deal Paradox.

- The *intuition* of most people tells them that each of the doors, the chosen door and the unchosen door, are equally likely to contain the prize so that there is a 50-50 chance of winning with either selection? This, however, is **not the case**.
- The probability of winning by using the switching technique is 2/3, while the odds of winning by not switching is 1/3. The easiest way to explain this is as follows:

Let's Make a Deal Paradox.

- The probability of picking the wrong door in the initial stage of the game is 2/3.
- If the contestant picks the wrong door initially, the host must reveal the remaining empty door in the second stage of the game. Thus, if the contestant switches after picking the wrong door initially, the contestant will win the prize.
- The probability of winning by switching then reduces to the probability of picking the wrong door in the initial stage which is clearly 2/3.
- **Demos:**
- file:///C:/Ivo.dir/UCLA_Classes/Applets.dir/SOCR/Prototype1.1/classes/TestExperiment.html
- C:\Ivo.dir\UCLA_Classes\Applets.dir\StatGames.exe

Long run behavior of coin tossing

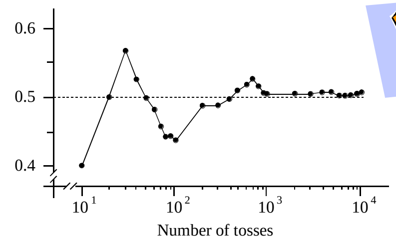

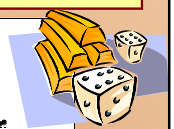

Figure 4.1.1 Proportion of heads versus number of tosses for John Kerrich's coin tossing experiment.

Definitions …

- The law of averages about the behavior of coin tosses – the relative proportion (relative frequency) of heads-to-tails in a coin toss experiment becomes more and more stable as the number of tosses increases. The law of averages applies to relative frequencies not absolute counts of #H and #T.
- Two widely held misconceptions about what the law of averages about coin tosses:
  - Differences between the actual numbers of heads & tails becomes more and more variable with increase of the number of tosses – a seq. of 10 heads doesn't increase the chance of a tail on the next trial.
  - Coin toss results are fair, but behavior is still unpredictable.

Coin Toss Models

- Is the coin tossing model adequate for describing the sex order of children in families?
  - This is a rough model which is not exact. In most countries rates of B/G is different; form 48% … 52%, usually. Birth rates of boys in some places are higher than girls, however, female population seems to be about 51%.
  - Independence, if a second child is born the chance it has the same gender (as the first child) is slightly bigger.

Data from a "random" draw

366 cylinders (for each day in the year) for the US Vietnam war draft. The N-th drawn number, corresp. to one B-day, indicating order of drafting.

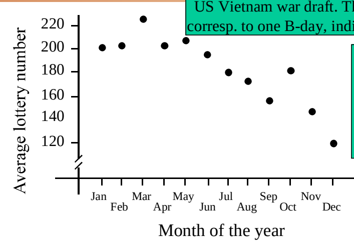

So, people born later in the year tend to have lower lottery numbers and a bigger chance of actually being drafted.

Figure 4.3.1 Average lottery numbers by month. Replotted from data in Fienberg [1971].

Types of Probability

- Probability models have two essential components (*sample space*, the space of all possible outcomes from an experiment; and a list of *probabilities* for each event in the sample space). Where do the outcomes and the probabilities come from?
- Probabilities from models – say mathematical/physical description of the sample space and the chance of each event. Construct a fair die tossing game.
- Probabilities from data – data observations determine our probability distribution. Say we toss a coin 100 times and the observed Head/Tail counts are used as probabilities.
- Subjective Probabilities – combining data and psychological factors to design a reasonable probability table (e.g., gambling, stock market).

Sample Spaces and Probabilities

- When the relative frequency of an event in the past is used to estimate the probability that it will occur in the future, what assumption is being made?
  - The underlying process is stable over time;
  - Our relative frequencies must be taken from large numbers for us to have confidence in them as probabilities.
- All statisticians agree about how probabilities are to be combined and manipulated (in math terms), however, not all agree what probabilities should be associated with for a particular real-world event.
- When a weather forecaster says that there is a 70% chance of rain tomorrow, what do you think this statement means? (Based on our past knowledge, according to the barometric pressure, temperature, etc. of the conditions we expect tomorrow, 70% of the time it did rain under such conditions.)

Sample spaces and events

- A ***sample space***, *S*, for a random experiment is the set of all possible outcomes of the experiment.
- An ***event*** is a *collection of outcomes*.
- An event ***occurs*** if any outcome making up that event occurs.

The complement of an event

- The complement of an event *A*, denoted $\bar{A}$, occurs *if and only if A* does not occur.

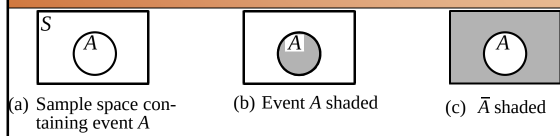

Figure 4.4.1 An event *A* in the sample space *S*.

Combining events – *all statisticians agree on*

- "***A* or *B***" contains all outcomes in *A* **or** *B* (or both).
- "***A* and *B***" contains all outcomes which are **in both** *A* and *B*.

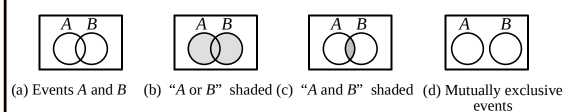

Figure 4.4.2 Two events.

From *Chance Encounters* by C.J. Wild and G.A.F. Seber, © John Wiley & Sons, 2000.

Mutually exclusive events cannot occur at the same time.

Probability distributions

- Probabilities always lie between 0 and 1 and they sum up to 1 (across all simple events) .
- ***pr(A)*** can be obtained by adding up the probabilities of all the outcomes in *A*.

$$pr(A) = \sum_{\substack{E \text{ outcome} \\ \text{in event } A}} pr(E)$$

Job losses in the US

TABLE 4.4.1 Job Losses in the US (in thousands) for 1987 to 1991

	Reason for Job Loss			
	Workplace moved/closed	**Slack work**	**Position abolished**	**Total**
Male	1,703	1,196	548	3,447
Female	1,210	564	363	2,137
Total	2,913	1,760	911	5,584

Job losses cont.

	Workplace moved/closed	Slack work	Position abolished	Total
Male	1,703	1,196	548	3,447
Female	1,210	564	363	2,137
Total	2,913	1,760	911	5,584

TABLE 4.4.2 Proportions of Job Losses from Table 4.4.1

	Reason for Job Loss			
	Workplace moved/closed	**Slack work**	**Position abolished**	**Row totals**
Male	.305	.214	.098	.617
Female	.217	.101	.065	.383
Column totals	.552	.315	.163	1.000

Review

- What is a sample space? What are the two essential criteria that must be satisfied by a possible sample space? (completeness – every outcome is represented; and uniqueness – no outcome is represented more than once.
- What is an event? (collection of outcomes)
- If A is an event, what do we mean by its complement, $\bar{A}$? When does $\bar{A}$ occur?
- If A and B are events, when does ***A* or *B*** occur? When does ***A* and *B*** occur?

Properties of probability distributions

- A sequence of number $\{p_1, p_2, p_3, \dots, p_n\}$ is a probability distribution for a sample space $S = \{s_1, s_2, s_3, \dots, s_n\}$, if $pr(s_k) = p_k$, for each 1<=k<=n. The two essential properties of a probability distribution $p_1, p_2, \dots, p_n$?

$$p_k \ge 0; \quad \sum_k p_k = 1$$

- How do we get the probability of an event from the probabilities of outcomes that make up that event?
- If all outcomes are distinct & equally likely, how do we calculate $pr(A)$? If $A = \{a_1, a_2, a_3, \dots, a_9\}$ and $pr(a_1)=pr(a_2)=\dots=pr(a_9)=p$; then

$$pr(A) = 9 \times pr(a_1) = 9p.$$

Example of probability distributions

- Tossing a coin twice. *Sample space* S={HH, HT, TH, TT}, for a fair coin each outcome is equally likely, so the probabilities of the 4 possible outcomes should be identical, p. Since, p(HH)=p(HT)=p(TH)=p(TT)=p and

$$p_k \ge 0; \quad \sum_k p_k = 1$$

- p = ¼ = 0.25.

Proportion vs. Probability

- How do the concepts of a proportion and a probability differ? A proportion is a partial description of a real population. The probabilities give us the chance of something happening in a random experiment. Sometimes, proportions are identical to probabilities (e.g., in a real population under the experiment *choose-a-unit-at-random*).
- See the ***two-way table of counts*** (*contingency table*) on Table 4.4.1, slide 19. E.g., *choose-a-person-at-random* from the ones laid off, and compute the chance that the person would be a male, laid off due to position-closing. We can apply the same rules for manipulating probabilities to proportions, in the case where these two are identical.

Rules for manipulating Probability Distributions

For mutually exclusive events,

pr(*A* or *B*) = pr(*A*) + pr(*B*)

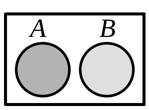

From *Chance Encounters* by C.J. Wild and G.A.F. Seber, © John Wiley & Sons, 2000.

Descriptive Table

pr(*Wild in* **and** *Seber in*) → 0.5; pr(*Seber in*) → 0.6; pr(*Wild in*) → 0.7

		Seber In	Seber Out	Total
Wild	In	0.5	?	0.7
	Out	?	?	?
	Total	0.6	?	1.00

Algebraic Table

	B	$\bar{B}$	Total
A	pr(A and B)	pr(A and $\bar{B}$)	pr(A)
$\bar{A}$	pr($\bar{A}$ and B)	pr($\bar{A}$ and $\bar{B}$)	pr($\bar{A}$)
Total	pr(B)	pr($\bar{B}$)	1.00

Availability of the Textbook authors to students

.5	?	.7
?	?	?
.6	?	1.00

→

.5	?	.7
?	?	.3
.6	.4	1.00

→

.5	.2	.7
?	?	.3
.6	.4	1.00

→

.5	.2	.7
.1	.2	.3
.6	.4	1.00

TABLE 4.5.1
Completed Probability Table

Wild	Sebe In	Sebe Out	Total
In	.5	.2	.7
Out	.1	.2	.3
Tota	.6	.4	1.0

Unmarried couples

Select an unmarried couple *at random* – the table proportions give us the probabilities of the events defined in the row/column titles.

TABLE 4.5.2 Proportions of Unmarried Male-Female Couples Sharing Household in the US, 1991

	Female				
Male	**Never Married**	**Divorced**	**Widowed**	**Married to other**	**Total**
Never Married	0.401	.111	.017	.025	.554
Divorced	.117	.195	.024	.017	.353
Widowed	.006	.008	.016	.001	.031
Married to other	.021	.022	.003	.016	.062
Total	.545	.336	.060	.059	1.000

Review

- If A and B are mutually exclusive, what is the probability that both occur? (0) What is the probability that at least one occurs? (sum of probabilities)
- If we have two or more mutually exclusive events, how do we find the probability that at least one of them occurs? (sum of probabilities)
- Why is it sometimes easier to compute *pr(A)* from $pr(A) = 1 - pr(\bar{A})$? (The complement of the even may be easer to find or may have a known probability. E.g., a random number between 1 and 10 is drawn. Let A ={a number less than or equal to 9 appears}. Find pr(A) = 1 – pr($\bar{A}$)). probability of $\bar{A}$ is pr({10 appears}) = 1/10 = 0.1. Also Monty Hall 3 door example!

Melanoma – type of skin cancer – an example of laws of conditional probabilities

TABLE 4.6.1: 400 Melanoma Patients by Type and Site

	Site			
Type	**Head and Neck**	**Trunk**	**Extremities**	**Row Totals**
Hutchinson's melanomic freckle	22	2	10	34
Superficial	16	54	115	185
Nodular	19	33	73	125
Indeterminant	11	17	28	56
Column Totals	68	106	226	400

Contingency table based on Melanoma histological type and its location

Conditional Probability

The ***conditional probability*** of ***A*** occurring ***given*** that B occurs is given by

$$pr(A \mid B) = \frac{pr(A \text{ and } B)}{pr(B)}$$

Suppose we select one out of the 400 patients in the study and we want to find the probability that the cancer is on the *extremities given that* it is of type *nodular*: P = 73/125 = P(C. on Extremities | Nodular)

$$\frac{\#\text{nodular patients with cancer on extremities}}{\#\text{nodular patients}}$$

Multiplication rule- what's the percentage of Israelis that are poor and Arabic?

$$pr(A \text{ and } B) = pr(A \mid B)pr(B) = pr(B \mid A)pr(A)$$

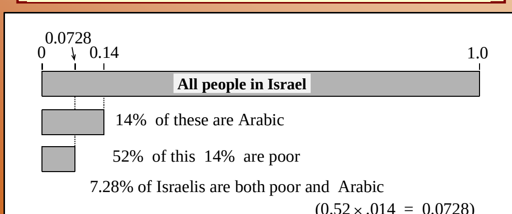

Figure 4.6.1 Illustration of the multiplication rule.

Multiplication rule- what's the percentage of Israelis that are poor and Arabic?

$$pr(A \text{ and } B) = pr(A \mid B)pr(B) = pr(B \mid A)pr(A)$$

	Jews	Arabs	Total
Poor		0.0728	?
Not-poor			
Total	0.86	0.14	1.0

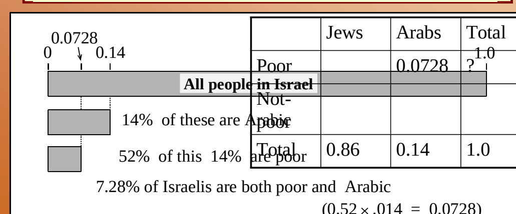

Figure 4.6.1 Illustration of the multiplication rule.

Review, Fri., Oct. 12, 2001

$pr(A \text{ and } B) = pr(A \mid B)pr(B) = pr(B \mid A)pr(A)$

$pr(A) = 1 - pr(\bar{A})$

1. Proportions (partial description of a real population) and probabilities (giving the chance of something happening in a random experiment) may be identical – under the experiment *choose-a-unit-at-random*
2. Properties of probabilities.

$\{p_k\}_{k=1}^{N}$ define probabilities $\Leftrightarrow p_k \ge 0; \ \sum_k p_k = 1$

A tree diagram for computing conditional probabilities

Suppose we draw 2 balls at random one at a time *without replacement* from an urn containing **4 black** and **3 white** balls, otherwise identical. What is the probability that the ***second ball is black***? Sample Spc?

P({*2-nd ball is black*}) =
P({*2-nd is black*} &{*1-st is black*}) +
P({2-nd is black} &{*1-st is white*}) =
4/7 x 3/6 + 4/6 x 3/7 = 4/7.

Mutually exclusive

A tree diagram

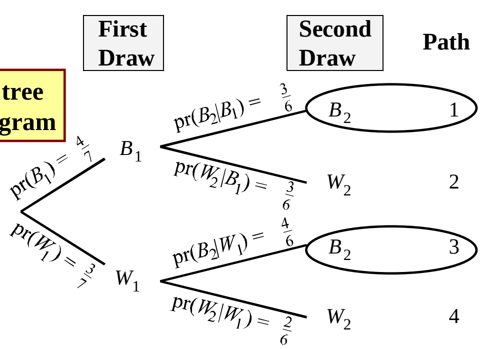

Figure 4.6.2 Tree diagram for a sampling problem.

Tree diagram for poverty in Israel

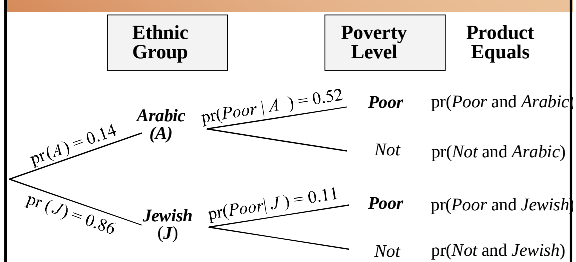

2-way table for poverty in Israel

pr(*Poor* and *Arabic*) = pr(*Poor*|*Arabic*) × pr(*Arabic*) [= 52% of 14%]

pr(*Poor* and *Jewish*) = pr(*Poor*|*Jewish*) × pr(*Jewish*) [= 11% of 86%]

		Ethnicity: Arabic	Ethnicity: Jewish	Total
Poverty	Poor	.52 × .14	.11 × .86	?
	Not poor	?	?	?
	Total	.14	.86	1.00

pr(*Arabic*) = .14 pr(*Jewish*) = .86

Figure 4.6.4 Proportions by Ethnicity and Poverty.

P(A & B) = P(A | B) x P(B),
P(A | B) = P(A & B) / P(B)
P(A & B) = P(B & A) = P(B | A) x P(A).
P(A | B) = [P(B | A) x P(A)] / P(B).

2-way table for poverty in Israel cont.

pr(*Poor* and *Arabic*) = pr(*Poor*|*Arabic*) × pr(*Arabic*) [= 52% of 14%]

pr(*Poor* and *Jewish*) = pr(*Poor*|*Jewish*) × pr(*Jewish*) [= 11% of 86%]

		Ethnicity: Arabic	Ethnicity: Jewish	Total
Poverty	Poor	.52 × .14	.11 × .86	?
	Not poor	?	?	?
	Total	.14	.86	1.00

pr(*Arabic*) = .14 pr(*Jewish*) = .86

TABLE 4.6.3 Proportions by Ethnicity and Poverty

		Ethnicity: Arabic	Ethnicity: Jewish	Total
Poverty	Poor	.0728	.0946	.1674
	Not Poor	.0672	.7654	.8326
	Total	.14	.86	1.00

Conditional probabilities and 2-way tables

- Many problems involving conditional probabilities can be solved by constructing two-way tables
- This includes *reversing the order of conditioning*

P(A & B) = P(A | B) x P(B) = P(B | A) x P(A)

Classes vs. Evidence Conditioning

- **Classes**: healthy(NC), cancer
- **Evidence**: positive mammogram (pos), negative mammogram (neg)
- If a woman has a positive mammogram result, what is the probability that she has breast cancer?

$$P(class \mid evidence) = \frac{P(evidence \mid class) \times P(class)}{P(evidence)}$$

$$P(cancer) = 0.01$$

$$P(pos \mid cancer) = 0.8$$

$$P(positive) = 0.107$$

$$P(cancer \mid pos) = ?$$

Proportional usage of *oral contraceptives* and their rates of failure

We need to complete the two-way contingency table of proportions

pr(*Failed* and *Oral*) = pr(*Failed* | *Oral*) × pr(*Oral*) [= 5% of 32%]

pr(*Failed* and *IUD*) = pr(*Failed* | *IUD*) × pr(*IUD*) [= 6% of 3%]

		Method					
		Steril.	Oral	Barrier	IUD	Sperm.	Total
Outcome	Failed	0 × .38	.05 × .32	.14 × .24	.06 × .03	.26 × .03	?
	Didn't	?	?	?	?	?	?
	Total	.38	.32	.24	.03	.03	1.00

pr(*Steril.*) = .38 pr(*Barrier*) = .24 pr(*IUD*) = .03

Oral contraceptives cont.

pr(*Failed* and *Oral*) = pr(*Failed* | *Oral*) × pr(*Oral*) [= 5% of 32%]

pr(*Failed* and *IUD*) = pr(*Failed* | *IUD*) × pr(*IUD*) [= 6% of 3%]

		Method					
		Steril.	Oral	Barrier	IUD	Sperm.	Total
Outcome	Failed	0 × .38	.05 × .32	.14 × .24	.06 × .03	.26 × .03	?
	Didn't	?	?	?	?	?	?
	Total	.38	.32	.24	.03	.03	1.00

pr(*Steril.*) = .38 pr(*Barrier*) = .24 pr(*IUD*) = .03

TABLE 4.6.4 Table Constructed from the Data in Example 4.6.8

		Method					
		Steril.	**Oral**	**Barrier**	**IUD**	**Sperm.**	**Total**
Outcome	Failed	0	.0160	.0336	.0018	.0078	.0592
	Didn't	.3800	.3040	.2064	.0282	.0222	.9408
	Total	.3800	.3200	.2400	.0300	.0300	1.0000

Remarks …

- In *pr(A | B)*, how should the symbol " | " is read *given that*.
- How do we interpret the fact that: *The event A always occurs when B occurs*? What can you say about pr(*A* | *B*)?

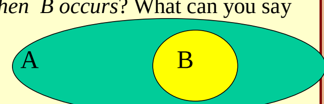

- When drawing a probability tree for a particular problem, how do you know *what events* to use for the first fan of branches and which events to use for the subsequent branching? (at each branching stage condition on all the info available up to here. E.g., at first branching use all simple events, no prior is available. At 3-rd branching condition of the previous 2 events, etc.).

TABLE 4.6.5 Number of Individuals Having a Given Mean Absorbance Ratio (MAR) in the ELISA for HIV Antibodies

MAR	Healthy Donor		HIV patients	
<2	202	} 275	0	} 2 False-Negatives (FNE)
2 - 2.99	73	Test cut-off	2	
3 - 3.99	15	} False-positives	7	
4 - 4.99	3		7	
5 - 5.99	2		15	
6 -11.99	2		36	
12+	0		21	
Total	297		88	

Power of a test is: 1-P(FNE)= 1-P(Neg|HIV) ~ 0.976

Adapted from Weiss et al.[1985]

HIV cont.

pr(*HIV* and *Positive*) = pr(*Positive*|*HIV*) × pr(*HIV*) [= 98% of 1%]

pr(*Not HIV* and *Negative*) = pr(*Negative*|*Not HIV*) × pr(*Not HIV*) [= 93% of 99%]

		Test result Positive	Negative	Total	
Disease status	HIV	.98 × .01	?	.01	← pr(*HIV*) = .01
	Not HIV	?	.93 × .99	.99	← pr(*Not HIV*) = .99
	Total	?	?	1.00	

Figure 4.6.6 Putting HIV information into the table.

HIV – reconstructing the contingency table

pr(*HIV* and *Positive*) = pr(*Positive*|*HIV*) × pr(*HIV*) [= 98% of 1%]

pr(*Not HIV* and *Negative*) = pr(*Negative*|*Not HIV*) × pr(*Not HIV*) [= 93% of 99%]

		Test result Positive	Negative	Total	
Disease status	HIV	.98 × .01	?	.01	← pr(*HIV*) = .01
	Not HIV	?	.93 × .99	.99	← pr(*Not HIV*) = .99
	Total	?	?	1.00	

TABLE 4.6.6 Proportions by Disease Status and Test Result

		Test Result Positive	Negative	Total
Disease Status	HIV	.0098	.0002	.01
	Not HIV	.0693	.9207	.99
	Total	.0791	.9209	1.00

Proportions of HIV infections by country

TABLE 4.6.7 Proportions Infected with HIV

Country	No. AIDS Cases	Population (millions)	pr(HIV)	Having | Test pr(HIV | Positive)
United States	218,301	252.7	0.00864	0.109
Canada	6,116	26.7	0.00229	0.031
Australia	3,238	16.8	0.00193	0.026
New Zealand	323	3.4	0.00095	0.013
United Kingdom	5,451	57.3	0.00095	0.013
Ireland	142	3.6	0.00039	0.005

Statistical independence

- Events *A* and *B* are *statistically independent* if knowing whether *B* has occurred gives no new information about the chances of *A* occurring,

  i.e. if **pr(*A* | *B*) = pr(*A*)**
- Similarly, P(*B* | *A*) = P(*B*), since

  P(B|A)=P(B & A)/P(A) = P(A|B)P(B)/P(A) = P(B)
- If *A* and *B* are *statistically independent*, then

$$\text{pr}(A \text{ and } B) = \text{pr}(A) \times \text{pr}(B)$$

Example using independence

There are many genetically based blood group systems. Two of these are: *Rh* blood type system (Rh+ and Rh-) and the *Kell* system (K+ and K-). For Europeans the following proportions are experimentally obtained.

Table 4.7.1 Blood Type Data

.92×.81, pr(*RH*+) = .81, .08×.81, .08×.19

	K+	*K*-	Total
Rh+	?	?	.81
Rh-	?	?	.19
Total	.08	.92	1.00

pr(*K*+) = .08, .92 × .19

→

	K+	*K*-	Total
Rh+	.0648	.7452	.81
Rh-	.0152	.1748	.19
Total	.08	.92	1.00

How can we fill in the inside of the two-way contingency table? It is known that anyone's blood type in one system is *independent* of their type in another system.

P(Rh+ and K+) = P(Rh+) x P(K+)= 0.81 x 0.08 = 0.0648

People vs. Collins

TABLE 4.7.2 Frequencies Assumed by the Prosecution

Yellow car	$\frac{1}{10}$	Girl with blond hair	$\frac{1}{3}$
Man with mustache	$\frac{1}{4}$	Black man with beard	$\frac{1}{10}$
Girl with ponytail	$\frac{1}{10}$	Interracial couple in car	$\frac{1}{1000}$

- The first occasion where a conviction was made in an American court of law, largely on statistical evidence, 1964. A woman was mugged and the offender was described as a wearing dark cloths, with blond hair in a pony tail who got into a yellow car driven by a black male accomplice with mustache and beard. The suspect brought to trial were picked out in a line-up and fit all of the descriptions. Using the *product rule for probabilities* an expert witness computed the chance that a random couple meets these characteristics, as 1:12,000,000.

Nuclear reactor safety

Probability of worst consequence = *Probability of initiating event* × *Probability of safety system failure* × *Probability of worst weather* × *Probability of highest population density*

- The calculations that the chance of 1:1,000,000,000 per reactor/per year that containment of a possible nuclear core meltdown would fail, releasing volatile products into the atmosphere, were based on *faulty tree analysis*. Many probability statements were made which were later critiqued on the basis of:
- Individual prob's were purely subjective and had no data backing.
- There were unfounded assumptions of independence in chains of events.
- *Fault tree analysis* may consider failure chances only from an anticipated cause.

Summary

- What does it mean for two events *A* and *B* to be *statistically independent*?
- Why is the working rule under independence, *P(A and B) = P(A) P(B)*, just a special case of the multiplication rule *P(A & B) = P(A | B) P(B)* ?
- *Mutual independence* of events $A_1, A_2, A_3, \ldots, A_n$ *if and only* if $P(A_1 \& A_2 \& \ldots \& A_n) = P(A_1)P(A_2)\ldots P(A_n)$
- What do we mean when we say two human characteristics are *positively associated*? *negatively associated*? (blond hair – blue eyes, pos.; black hair – blue eyes, neg.assoc.)

Review

- What happens to the calculated P(*A* and *B*) if we treat positively associated events as independent? if we treat negatively associated events as independent?

  (Example, let B={A + {b}}, A & B are pos-assoc'd, P(A&B)=P(A)[P(A)+P({b})], under indep. assump's. However, P(A&B)=P(B|A)P(A)=1 x P(A) > P(A)[P(A)+P({b})], underestimating the real chance of events. If A & B are neg-assoc'd → A & comp(B) are pos-assoc'd. In general, this may lead to answers that are grossly too small or too large …)
- Why do people often treat events as independent? When can we trust their answers?(Easy computations! Not always!)

Summary of ideas

- The ***probabilities*** people quote come from 3 main sources:
  - (i) ***Models*** (idealizations such as the notion of equally likely outcomes which suggest probabilities by symmetry).
  - (ii) ***Data*** (e.g.relative frequencies with which the event has occurred in the past).
  - (iii) ***subjective feelings*** representing a degree of belief
- A simple probability model consists of a sample space and a probability distribution.
- A sample space, *S*, for a random experiment is the set of all possible outcomes of the experiment.

Summary of ideas cont.

- A list of numbers $p_1, p_2, \ldots$ is a ***probability distribution*** for $S = \{s_1, s_2, s_3, \ldots\}$, provided
  - all of the p_i's lie between 0 and 1, and
  - they add to 1.
- According to the probability model, p_i is the probability that outcome s_i occurs.
- We write $p_i = P(s_i)$.

Summary of ideas cont.

- An ***event*** is a collection of outcomes
- An event ***occurs*** if any outcome making up that event occurs
- The probability of event A can be obtained by adding up the probabilities of all the outcomes in A
- If all outcomes are equally likely,

$$\text{pr}(A) = \frac{\text{number of outcomes in } A}{\text{total number of outcomes}}$$

Summary of ideas cont.

- The ***complement*** of an event *A*, denoted $\bar{A}$, occurs if *A* does not occur
- It is useful to represent events diagrammatically using ***Venn diagrams***
- ***A union*** of events, ***A* or *B*** contains all outcomes in *A* or *B* (including those in both). It occurs if at least one of *A* or *B* occurs
- ***An intersection*** of events, ***A and B*** contains all outcomes which are in ***both*** *A* and *B*. It occurs only if both *A* and *B* occur
- ***Mutually exclusive*** events cannot occur at the same time

Summary of ideas cont.

- The ***conditional probability*** of *A* occurring ***given*** that *B* occurs is given by

$$\text{pr}(A \mid B) = \frac{\text{pr}(A \text{ and } B)}{\text{pr}(B)}$$

- Events *A* and *B* are ***statistically independent*** if knowing whether *B* has occurred gives no new information about the chances of *A* occurring, i.e. if $P(A \mid B) = P(A)$ → $P(B|A)=P(B)$.
- If events are physically independent, then, under any sensible probability model, they are also statistically independent
- Assuming that events are independent when in reality they are not can often lead to answers that are grossly too big or grossly too small

Formula Summary

- For discrete sample spaces, pr(*A*) can be obtained by adding the probabilities of all outcomes in *A*
- For equally likely outcomes in a finite sample space

$$\text{pr}(A) = \frac{\text{number of outcomes in } A}{\text{total number of outcomes}}$$

Formula summary cont.

- $\text{pr}(S) = 1$
- $\text{pr}(\bar{A}) = 1 - \text{pr}(A)$
- If *A* and *B* are mutually exclusive events, then $\text{pr}(A \text{ or } B) = \text{pr}(A) + \text{pr}(B)$

  (here "or" is used in the inclusive sense)
- If $A_1, A_2, \ldots, A_k$ are mutually exclusive events, then

  $\text{pr}(A_1 \text{ or } A_2 \text{ or } \ldots \text{ or } A_k) = \text{pr}(A_1)+\text{pr}(A_2)+\ldots+\text{pr}(A_k)$

Formula summary cont.

Conditional probability

- Definition:

$$\text{pr}(A \mid B) = \frac{\text{pr}(A \text{ and } B)}{\text{pr}(B)}$$

- Multiplication formula:

$$\text{pr}(A \text{ and } B) = \text{pr}(B|A)\text{pr}(A) = \text{pr}(A|B)\text{pr}(B)$$

Formula summary cont.

Multiplication Rule under independence:

- If *A* and *B* are independent events, then

$$\text{pr}(A \text{ and } B) = \text{pr}(A)\,\text{pr}(B)$$

- If $A_1, A_2, \ldots, A_n$ are mutually independent,

$\text{pr}(A_1 \text{ and } A_2 \text{ and } \ldots \text{ and } A_n) = \text{pr}(A_1)\,\text{pr}(A_2) \ldots \text{pr}(A_n)$

Law of Total Probability

- If $\{A_1, A_2, \ldots, A_n\}$ are a partition of the sample space (mutually exclusive and $UA_i=S$) then for any event *B*

$$P(B) = P(B|A_1)P(A_1) + P(B|A_2)P(A_2) + \ldots + P(B|A_n)P(A_n)$$

Ex:

$P(B) = P(B|A_1)P(A_1) + P(B|A_2)P(A_2)$

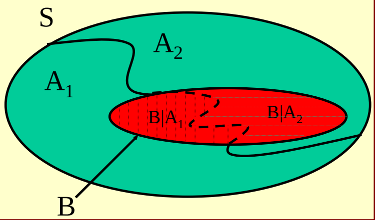

Bayesian Rule

- If $\{A_1, A_2, \ldots, A_n\}$ are a non-trivial partition of the sample space (mutually exclusive and $UA_i=S$, $P(A_i)>0$) then for any non-trivial event and *B* ($P(B)>0$)

$$P(A_i \mid B) = P(A_i \cap B) / P(B) = [P(B \mid A_i) \times P(A_i)] / P(B) =$$

$$= \frac{P(A_i \mid B) \times P(A_i)}{\sum_{k=1}^{n} P(B \mid A_k) P(A_k)}$$

Bayesian Rule

$$P(A_i) = \frac{P(A_i \mid B) \times P(A_i)}{\sum_{k=1}^{n} P(B \mid A_k) P(A_k)}$$

D **= the test person has the disease.**

T **= the test result is positive.**

Ex: (Laboratory blood test) **Assume:**

P(positive Test| Disease) = 0.95

P(positive Test| no Disease)=0.01

P(Disease) = 0.005

Find:

P(Disease | positive Test)=?

P(D | T) = ?

$$P(D \mid T) = \frac{P(D \cap T)}{P(T)} = \frac{P(T \mid D) \times P(D)}{P(T \mid D) \times P(D) + P(T \mid D^c) \times P(D^c)}$$

$$= \frac{0.95 \times 0.005}{0.95 \times 0.005 + 0.01 \times 0.995} = \frac{0.00475}{0.02465} = 0.193$$

Classes vs. Evidence Conditioning

- **Classes**: healthy(NC), cancer
- **Evidence**: positive mammogram (pos), negative mammogram (neg)
- If a woman has a positive mammogram result, what is the probability that she has breast cancer?

$$P(class \mid evidence) = \frac{P(evidence \mid class) \times P(class)}{\sum_{classes} P(evidence \mid class) \times P(class)}$$

$P(cancer) = 0.01$

$P(pos \mid cancer) = 0.8$

$P(pos \mid healthy) = 0.1$

$P(cancer \mid pos) = ?$

P(C|P)=P(P|C)xP(C)/[P(P|C)xP(C)+ P(P|H)xP(H)]

P(C|P)=0.8x0.01 / [0.8x0.01 + 0.1x0.99] = ?

Bayesian Rule (different data/example!)

True Disease State

Test Results	No Disease	Disease	Total
Negative	OK (0.98505)	False Negative II (0.00025)	0.9853
Positive	False Positive I (0.00995)	OK (0.00475)	0.0147
Total	0.995	0.005	1.0

$$P(T \cap D^C) = P(T \mid D^C) \times P(D^C) = 0.01 \times 0.995 = 0.00995$$

Power of Test = $1 - P(T^C \mid D) = 0.99975$

Sensitivity: **TP/(TP+FN)** = 0.00475/(0.00475+ 0.00025)= 0.95

Specificity: **TN/(TN+FP)** = 0.98505/(0.98505+ 0.00995) = 0.99

Examples – Birthday Paradox

- The Birthday Paradox: In a random group of N people, what is the change that at least two people have the same birthday?
- E.x., if N=23, P>0.5. Main confusion arises from the fact that in real life we rarely meet people having the same birthday as us, and we meet more than 23 people.
- The reason for such high probability is that any of the 23 people can compare their birthday with any other one, not just you comparing your birthday to anybody else's.
- There are N-Choose-2 = 20*19/2 ways to select a pair or people. Assume there are 365 days in a year, P(one-particular-pair-same-B-day)=1/365, and
- P(one-particular-pair-failure)=1-1/365 ~ 0.99726.
- For N=20, 20-Choose-2 = 190. E={No 2 people have the same birthday is the event all 190 pairs fail (have different birthdays)}, then P(E) = P(failure)^{190} = 0.99726^{190} = 0.59.
- Hence, P(at-least-one-success)=1-0.59=0.41, quite high.
- Note: for N=42 → P>0.9 …