# Section 2: Estimation, Confidence Intervals and Testing Hypothesis 

Carlos M. Carvalho<br>The University of Texas at Austin<br>McCombs School of Business

http://faculty.mccombs.utexas.edu/carlos.carvalho/teaching/

## A First Modeling Exercise

- I have US\$ 1,000 invested in the Brazilian stock index, the IBOVESPA. I need to predict tomorrow's value of my portfolio.
- I also want to know how risky my portfolio is, in particular, I want to know how likely am I to lose more than $3 \%$ of my money by the end of tomorrow's trading session.
- What should I do?


## IBOVESPA - Data



- As a first modeling decision, let's call the random variable associated with daily returns on the IBOVESPA $X$ and assume that returns are independent and identically distributed as

$$
X \sim N\left(\mu, \sigma^{2}\right)
$$

- Question: What are the values of $\mu$ and $\sigma^{2}$ ?
- We need to estimate these values from the sample in hands ( $n=113$ observations)...
- Let's assume that each observation in the random sample $\left\{x_{1}, x_{2}, x_{3}, \ldots, x_{n}\right\}$ is independent and distributed according to the model above, i.e., $x_{i} \sim N\left(\mu, \sigma^{2}\right)$
- An usual strategy is to estimate $\mu$ and $\sigma^{2}$, the mean and the variance of the distribution, via the sample mean $(\bar{X})$ and the sample variance $\left(s^{2}\right) \ldots$ (their sample counterparts)

$$
\begin{gathered}
\bar{X}=\frac{1}{n} \sum_{i=1}^{n} x_{i} \\
s^{2}=\frac{1}{n-1} \sum_{i=1}^{n}\left(x_{i}-\bar{X}\right)^{2}
\end{gathered}
$$

For the IBOVESPA data in hands,
BOVESPA


- The red line represents our "model", i.e., the normal distribution with mean and variance given by the estimated quantities $\bar{X}$ and $s^{2}$.
- What is $\operatorname{Pr}(X<-3)$ ?


## Annual Returns on the US market...

Assume I invest some money in the U.S. stock market. Your job is to tell me the following:

- what is my expected one year return?
- what is the standard deviation (volatility)?
- what is the probability my investment grow by $10 \%$ or more?
- What happens in 20 years if I invest $\$ 1$ today on the market?


## Models, Parameters, Estimates...

In general we talk about unknown quantities using the language of probability... and the following steps:

- Define the random variables of interest
- Define a model (or probability distribution) that describes the behavior of the RV of interest
- Based on the data available, we estimate the parameters defining the model
- We are now ready to describe possible scenarios, generate predictions, make decisions, evaluate risk, etc...


## Oracle vs SAP Example (understanding variation)

## RESEARCH NOTE

## "SAP customers are 20\% less profitable than their industry peers"

- Nucleus Research Study, March 2006, based on an analysis of 81 publicly traded SAP customers.


## Don't SAP Your Profits.

Get Results With Oracle Applications.

## ORACLE

## Oracle vs. SAP

- Do we "buy" the claim from this add?
- We have a dataset of 81 firms that use SAP...
- The industry ROE is $15 \%$ (also an estimate but let's assume it is true)
- We assume that the random variable $X$ represents ROE of SAP firms and can be described by

$$
X \sim N\left(\mu, \sigma^{2}\right)
$$

|  | $\bar{X}$ | $s^{2}$ |
| :---: | :---: | :---: |
| SAP firms | 0.1263 | 0.065 |

- Well, $\frac{0.12}{0.15} \approx 0.8$ ! I guess the add is correct, right?
- Not so fast...


## Oracle vs. SAP

- Let's assume the sample we have is a good representation of the "population" of firms that use SAP...
- What if we have observed a different sample of size 81 ?


## Oracle vs. SAP

- Selecting a random, with replacement, from the original 81 samples I get a new $\bar{X}=0.09 \ldots$ I do it again, and I get $\bar{X}=0.155 \ldots$ and again $\bar{X}=0.132 \ldots$



## data sample $\swarrow \quad \downarrow$



.

## Oracle vs. SAP

- After doing this 1000 times... here's the histogram of $\bar{X} \ldots$ Now, what do you think about the add?

Histogram of sample mean


## Sampling Distribution of Sample Mean

Consider the mean for an iid sample of $n$ observations of a random variable $\left\{X_{1}, \ldots, X_{n}\right\}$

It turns out

$$
\bar{X} \sim N\left(\mu, \frac{\sigma^{2}}{n}\right)
$$

This is called the sampling distribution of the mean...

## Sampling Distribution of Sample Mean

- The sampling distribution of $\bar{X}$ describes how our estimate would vary over different datasets of the same size $n$
- It provides us with a vehicle to evaluate the uncertainty associated with our estimate of the mean...
- It turns out that $s^{2}$ is a good proxy for $\sigma^{2}$ so that we can approximate the sampling distribution by

$$
\bar{X} \sim N\left(\mu, \frac{s^{2}}{n}\right)
$$

- We call $\sqrt{\frac{s^{2}}{n}}$ the standard error of $\bar{X} \ldots$ it is a measure of its variability... I like the notation

$$
s_{\bar{X}}=\sqrt{\frac{s^{2}}{n}}
$$

## Back to the Oracle vs. SAP example

Back to our simulation...

Histogram of sample mean


## Confidence Intervals

$$
\bar{X} \sim N\left(\mu, s_{\bar{X}}^{2}\right)
$$

so...

$$
(\bar{X}-\mu) \sim N\left(0, s_{\bar{X}}^{2}\right)
$$

right?

- What is a good prediction for $\mu$ ? What is our best guess?? $\bar{x}$
- How do we make mistakes? How far from $\mu$ can we be?? $95 \%$ of the time $\pm 2 \times s_{\bar{X}}$
- $\left[\bar{X} \pm 2 \times s_{\bar{X}}\right]$ gives a $95 \%$ range of plausible values for $\mu \ldots$ this is called the $95 \%$ Confidence Interval for $\mu$.


## Oracle vs. SAP example... one more time

In this example, $\bar{X}=0.1263, s^{2}=0.065$ and $n=81 \ldots$ therefore, $s_{\bar{X}}^{2}=\frac{0.065}{81}$ so, the $95 \%$ confidence interval for the ROE of SAP firms is

$$
\begin{gathered}
{\left[\bar{X}-2 \times s_{\bar{X}} ; \bar{X}+2 \times s_{\bar{X}}\right]} \\
=\left[0.1263-2 \times \sqrt{\frac{0.065}{81}} ; 0.1263+2 \times \sqrt{\frac{0.065}{81}}\right] \\
=[0.069 ; 0.183]
\end{gathered}
$$

- Is 0.15 a plausible value? What does that mean?


## Back to the Oracle vs. SAP example

Back to our simulation...

Histogram of sample mean


## Let's revisit the US stock market example from before...

Let's run a simulation based on our results...

## Estimating Proportions... another modeling example

Your job is to manufacture a part. Each time you make a part, it is defective or not. Below we have the results from 100 parts you just made. $Y_{i}=1$ means a defect, 0 a good one.
How would you predict the next one?


There are 18 ones and 82 zeros.

In this case, it might be reasonable to model the defects as iid...

We can't be sure this is right, but, the data looks like the kind of thing we would get if we had iid draws with that p!!!

If we believe our model, what is the chance that the next 10 are good?
$.82^{10}=0.137$.

## Standard Error of Proportions

We used the proportion of defects in our sample to estimate $p$, the true, long-run, proportion of defects.

Could this estimate be wrong?!! Let $\hat{p}$ denote the sample proportion.

The standard error associated with the sample proportion as an estimate of the true proportion is:

$$
s_{\hat{p}}=\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}
$$

## Standard Error of Proportions

We estimate the true $p$ by the observed sample proportion of 1 's, $\hat{p}$.

The (approximate) 95\% confidence interval for the true proportion is:

$$
\hat{p} \pm 2 s_{\hat{p}} .
$$

## Defects:

In our defect example we had $\hat{p}=.18$ and $n=100$.
This gives

$$
s_{\hat{p}}=\sqrt{\frac{(.18)(.82)}{100}}=.04
$$

The confidence interval is $.18 \pm .08=(0.1,0.26)$

## Polls: yet another example...

## (Read chapter 10 of "Naked Statistics" if you have a chance)

If we take a relatively small random sample from a large population and ask each respondent yes or no with yes $\approx Y_{i}=1$ and no
$\approx Y_{i}=0$, where $p$ is the true population proportion of yes.
Suppose, as is common, $n=1000$, and $\hat{p} \approx .5$.
Then,

$$
s_{\hat{p}}=\sqrt{\frac{(.5)(.5)}{1000}}=.0158
$$

The standard error is .0158 so that the $\pm$ is .0316 , or about $\pm 3 \%$.
(Sounds familiar?!)

## Example: Salary Discrimination

Say we are concerned with potential salary discrimination between males and females in the banking industry... To study this issue, we get a sample of salaries for both 100 males and 150 females from multiple banks in Chicago. Here is a summary of the data:

|  | average | std. deviation |
| :--- | :---: | :---: |
| males | $150 k$ | $30 k$ |
| females | $143 k$ | $15 k$ |

What do we conclude? Is there a difference FOR SURE?

## Example: Salary Discrimination

Let's compute the confidence intervals:
males:

$$
\left(150-2 \times \sqrt{\frac{30^{2}}{100}} ; 150+2 \times \sqrt{\frac{30^{2}}{100}}\right)=(144 ; 156)
$$

females:

$$
\left(143-2 \times \sqrt{\frac{15^{2}}{150}} ; 143+2 \times \sqrt{\frac{15^{2}}{150}}\right)=(140.55 ; 145.45)
$$

## Example: Google Search Algorithm

Google is testing a couple of modifications in its search algorithms... they experiment with 2,500 searches and check how often the result was defined as a "success". Here's the data from this experiment:

| Algorithm | current | $\bmod 1$ | $\bmod 2$ |
| :--- | :---: | :---: | :---: |
| success | 1755 | 1850 | 1760 |
| failure | 745 | 650 | 740 |

The probability of success is estimated to be $\hat{p}=0.702$ for the current algorithm, $\hat{p}_{A}=0.74$ for modification (A) and $\hat{p}_{B}=0.704$ for modification (B) .

## Example: Google Search Algorithm

Let's compute the confidence intervals and check if these modifications are REALLY working... current:
$\left(.702-2 \times \sqrt{\frac{.702 *(1-.702)}{2500}} ; .702+2 \times \sqrt{\frac{.702 *(1-.702)}{2500}}\right)=(0.683 ; 0.720)$
$\bmod (A):$
$\left(.740-2 \times \sqrt{\frac{.740 *(1-.740)}{2500}} ; .740+2 \times \sqrt{\frac{.740 *(1-.740)}{2500}}\right)=(0.723 ; 0.758)$
$\bmod (B):$
$\left(.704-2 \times \sqrt{\frac{.704 *(1-.704)}{2500}} ; .704+2 \times \sqrt{\frac{.704 *(1-.704)}{2500}}\right)=(0.686 ; 0.722)$
What do we conclude?

## Standard Error for the Difference in Means

It turns out there is a more precise way to address these comparisons problems (for two groups)...

We can compute the standard error for the difference in means:

$$
s_{\left(\bar{X}_{a}-\bar{X}_{b}\right)}=\sqrt{\frac{s_{X_{a}}^{2}}{n_{a}}+\frac{s_{X_{b}}^{2}}{n_{b}}}
$$

or, for the difference in proportions

$$
s_{\left(\hat{p}_{a}-\hat{p}_{b}\right)}=\sqrt{\frac{\hat{p}_{a}\left(1-\hat{p}_{a}\right)}{n_{a}}+\frac{\hat{p}_{b}\left(1-\hat{p}_{b}\right)}{n_{b}}}
$$

## Confidence Interval for the Difference in Means

We can then compute the
confidence interval for the difference in means:

$$
\left(\bar{X}_{a}-\bar{X}_{b}\right) \pm 2 \times s_{\left(\bar{X}_{a}-\bar{X}_{b}\right)}
$$

or, the confidence interval for the difference in proportions

$$
\left(\hat{p}_{a}-\hat{p}_{b}\right) \pm 2 \times s_{\left(\hat{p}_{a}-\hat{p}_{b}\right)}
$$

## Let's revisit the examples... Salary Discrimination

$$
s_{\left(\bar{X}_{\text {males }}-\bar{x}_{\text {females }}\right)}=\sqrt{\frac{30^{2}}{100}+\frac{15^{2}}{150}}=3.24
$$

so that the confidence interval for the difference in means is:

$$
(150-143) \pm 2 \times 3.24=(0.519 ; 13.48)
$$

What is the conclusion now?

## Let's revisit the examples... Google Search

Let's look at the difference between the current algorithm and modification B...

$$
s_{\left(\hat{\rho}_{\text {current }}-\hat{p}_{\text {new }}\right)}=\sqrt{\frac{0.702 * 0.298}{2500}+\frac{0.704 * 0.296}{2500}}=0.0129
$$

so that the confidence interval for the difference in means is:

$$
(0.702-0.704) \pm 2 \times 0.0129=(-0.0278 ; 0.0238)
$$

What is the conclusion now?

## The Bottom Line...

- Estimates are based on random samples and therefore random (uncertain) themselves
- We need to account for this uncertainty!
- "Standard Error" measures the uncertainty of an estimate
- We define the " $95 \%$ Confidence Interval" as

$$
\text { estimate } \pm 2 \times \text { s.e. }
$$

- This provides us with a plausible range for the quantity we are trying to estimate.


## The Bottom Line...

- When estimating a mean the $95 \%$ C.I. is

$$
\bar{X} \pm 2 \times s_{\bar{X}}
$$

- When estimating a proportion the $95 \%$ C.I. is

$$
\hat{p} \pm 2 \times s_{\hat{p}}
$$

- The same idea applies when comparing means or proportions


## Testing

Suppose we want to assess whether or not $\mu$ equals a proposed value $\mu^{0}$. This is called hypothesis testing.

Formally we test the null hypothesis:
$H_{0}: \mu=\mu^{0}$
vs. the alternative
$H_{1}: \mu \neq \mu^{0}$

## Testing

That are 2 ways we can think about testing:

1. Building a test statistic... the t-stat,

$$
t=\frac{\bar{X}-\mu^{0}}{s_{\bar{X}}}
$$

This quantity measures how many standard deviations the estimate $(\bar{X})$ from the proposed value $\left(\mu^{0}\right)$.

If the absolute value of $t$ is greater than 2 , we need to worry (why?)... we reject the hypothesis.

## Testing

2. Looking at the confidence interval. If the proposed value is outside the confidence interval you reject the hypothesis.

Notice that this is equivalent to the t-stat. An absolute value for $t$ greater than 2 implies that the proposed value is outside the confidence interval... therefore reject.

This is my preferred approach for the testing problem. You can't go wrong by using the confidence interval!

## Testing (Proportions)

- The same idea applies to proportions... we can compute the t-stat testing the hypothesis that the true proportion equals $p^{0}$

$$
t=\frac{\hat{p}-p^{0}}{s_{\hat{p}}}
$$

Again, if the absolute value of $t$ is greater than 2 , we reject the hypothesis.

- As always, the confidence interval provides you with the same (and more!) information.
(Note: In the proportion case, this test is sometimes called a z-test)


## Testing (Differences)

- For testing the difference in means:

$$
t=\frac{\left(\bar{X}_{a}-\bar{X}_{b}\right)-d^{0}}{s_{\left(\bar{X}_{a}-\bar{X}_{b}\right)}}
$$

- For testing a difference in proportions:

$$
t=\frac{\left(\hat{p}_{a}-\hat{p}_{b}\right)-d^{0}}{s_{\left(\hat{p}_{a}-\hat{p}_{b}\right)}}
$$

In both cases $d^{0}$ is the proposed value for the difference (we often think of zero here... why?)

Again, if the absolute value of $t$ is greater than 2 , we reject the hypothesis.

## Testing... Examples

Let's recap by revisiting some examples:

- What hypothesis were we interested in the Oracle vs. SAP example? Use a t-stat to test it...
- Using the t-stat, test whether or not the Patriots are cheating in their coin tosses
- Use the t-stat to determine whether or not males are paid more than females in the Chicago banking industry
- What does the t-stat tells you about Google's new search algorithm?


## The Importance of Considering and Reporting Uncertainty

In 1997 the Red River flooded Grand Forks, ND overtopping its levees with a 54 -feet crest. $75 \%$ of the homes in the city were damaged or destroyed!

It was predicted that the rain and the spring melt would lead to a 49 -feet crest of the river. The levees were 51-feet high.

The Water Services of North Dakota had explicitly avoided communicating the uncertainty in their forecasts as they were afraid the public would loose confidence in their abilities to predict such events.

## The Importance of Considering and Reporting Uncertainty

It turns out the prediction interval for the flood was $49 \mathrm{ft} \pm 9 \mathrm{ft}$ leading to a $35 \%$ probability of the levees being overtopped!!

Should we take the point prediction (49ft) or the interval as an input for a decision problem?

In general, the distribution of potential outcomes are very relevant to help us make a decision

## The Importance of Considering and Reporting <br> Uncertainty

The answer seems obvious in this example (and it is!)... however, you see these things happening all the time as people tend to underplay uncertainty in many situations!
"Why do people not give intervals? Because they are embarrassed!"
Jan Hatzius, Goldman Sachs economists talking about economic forecasts...

Don't make this mistake! Intervals are your friend and will lead to better decisions!

## Building Portfolios

- Let's assume we are considering 2 investment opportunities

1. IBM stocks
2. ALCOA stocks

- How should we start thinking about this problem?


## Building Portfolios

Let's first learn about the characteristics of each option by assuming the following models:

- IBM $\sim N\left(\mu_{I B M}, \sigma_{I B M}^{2}\right)$
- ALCOA $\sim N\left(\mu_{\text {Alcoa }}, \sigma_{\text {Alcoa }}^{2}\right)$

After observing some return data we can came up with estimates for the means and variances describing the behavior of these stocks

## Building Portfolios



| IBM | ALCOA |
| :---: | :---: |
| $\bar{X}_{I B M}=12.5$ | $\bar{X}_{\text {Alcoa }}=14.9$ |
| $s_{I B M}=10.5$ | $s_{\text {Alcoa }}=14.0$ |
| $\operatorname{corr}(I B M, A L C O A)=0.33$ |  |

## Covariance

Measure the direction and strength of the linear relationship between $Y$ and $X$

$$
\operatorname{Cov}(Y, X)=\frac{\sum_{i=1}^{n}\left(Y_{i}-\bar{Y}\right)\left(X_{i}-\bar{X}\right)}{n-1}
$$



- $s_{y}=15.98, s_{x}=9.7$
- $\operatorname{Cov}(X, Y)=125.9$

How do we interpret that?

## Correlation

Correlation is the standardized covariance:

$$
\operatorname{corr}(X, Y)=\frac{\operatorname{cov}(X, Y)}{\sqrt{s_{x}^{2} s_{y}^{2}}}=\frac{\operatorname{cov}(X, Y)}{s_{x} s_{y}}
$$

The correlation is scale invariant and the units of measurement don't matter: It is always true that $-1 \leq \operatorname{corr}(X, Y) \leq 1$.

This gives the direction $(-$ or + ) and strength $(0 \rightarrow 1)$ of the linear relationship between $X$ and $Y$.

## Correlation

$$
\operatorname{corr}(Y, X)=\frac{\operatorname{cov}(X, Y)}{\sqrt{s_{x}^{2} s_{y}^{2}}}=\frac{\operatorname{cov}(X, Y)}{s_{x} s_{y}}=\frac{125.9}{15.98 \times 9.7}=0.812
$$



## Correlation






## Building Portfolios

- How about combining these options? Is that a good idea? Is it good to have all your eggs in the same basket? Why?
- What if I place half of my money in ALCOA and the other half IBM...

In order to answer this question we need to understand how IBM and Alcoa move together!

## Building Portfolios

We need to understand the behavior of the weighted sum (linear combinations) of two random variables...

Let $X$ and $Y$ be two random variables:

- $E(a X+b Y)=a E(X)+b E(Y)$
- $\operatorname{Var}(a X+b Y)=a^{2} \operatorname{Var}(X)+b^{2} \operatorname{Var}(Y)+2 a b \times \operatorname{Cov}(X, Y)$
where $\operatorname{Cov}(X, Y)$ is the covariance between $X$ and $Y$.


## Back to Building Portfolios

- So, by using what we learned about the means, variances and covariance, we get to:

$$
\begin{aligned}
E(P) & =0.5 \bar{X}_{I B M}+0.5 \bar{X}_{A l c o a} \\
\operatorname{Var}(P) & =0.5^{2} * s_{I B M}^{2}+0.5^{2} * s_{A / c o a}^{2}+2 * 0.5 * 0.5 * \operatorname{Cov}(I B M, \text { Alcoa })
\end{aligned}
$$

- $E(P)$ and $\operatorname{Var}(P)$ refer to the estimated mean and variance of our portfolio


## Building Portfolios

Here are the results for different combinations of Alcoa and IBM...


