Understanding Uncertainty

David W Scott

Noah Harding Professor of Statistics, Rice University

2011 RUSMP Fall Networking Conference

Houston, Texas

September 24, 2011

The scientific process has many components

- experimentation
- modeling
- data collection
- design of experiments
- curve fitting
- parameter estimation
- uncertainty quantification

Uncertainty Quantification

The fundamental result of statistics is:

Uncertainty may be reduced by averaging.

This simple statement applies in a range of applications

a. Given a random sample $\{x_1, x_2, \ldots, x_n\}$,

$$\text{Average}(\bar{x}) = \text{Average}(x_1)$$

but

$$\text{Variance}(\bar{x}) = \frac{1}{n}\text{Variance}(x_1)$$

b. Reconstructing a crystalline sample takes advantage of the periodic structure and Fourier techniques. For example, x-ray crystallography begins with as pure a crystal as available.

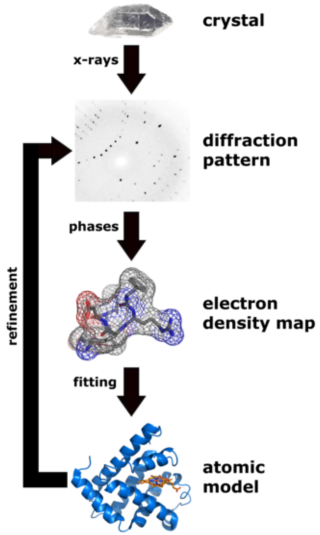

Thomas Splettstoesser, Heidelberg University

c. Given thousands of images of flash-frozen viruses at random angles, Wah Chiu (Baylor College of Medicine) has shown how to reconstruct an individual virus:

   1. isolate each virus image
   2. cluster the images into 50-60 groups
   3. average images within a cluster
   4. try to figure out the angle of each cluster
   5. use known symmetry to aid reconstruction

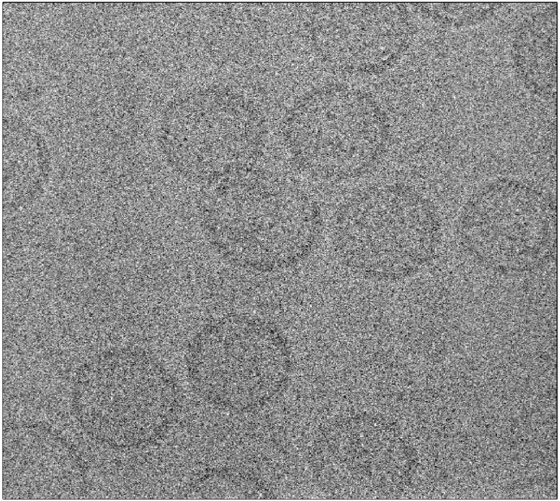

Cryo-electron microscope images of P22 virus.

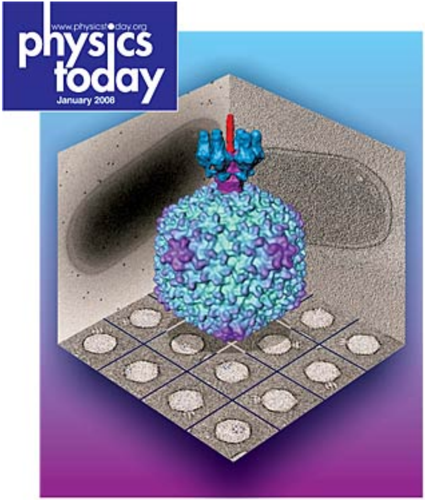

Focus on phages

Statistical Uncertainty

- A tool for understanding uncertainty is simulation
- These computer experiments can imitate any “real” process imagined
- By replication, the accuracy (or, equivalently, the uncertainty) may be observed

Example 1.

- At some point during the Fall of 2008, you heard the following poll result:

47%	*Obama*
53%	*Other*

(*"Other"* may break down as 43% McCain, 10% undecided, ignoring other candidates.)

- A Gallup poll might report their findings based on 1200 phone interviews
- So a *single* computer simulation would involve flipping a biased coin 1200 times, and counting the number of "heads"
- Repeat the simulation a large number of times (1000 here) and accumulate the results in a frequency chart (histogram)

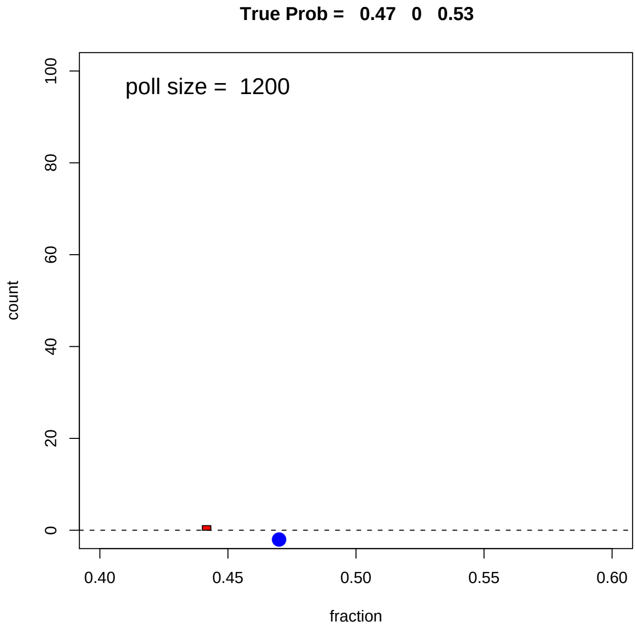

True Prob = 0.47 0 0.53
poll size = 1200
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

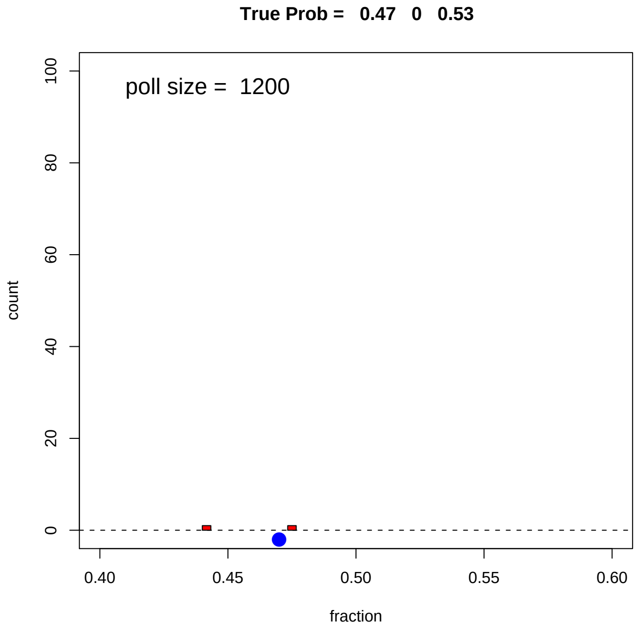
True Prob = 0.47 0 0.53
poll size = 1200
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

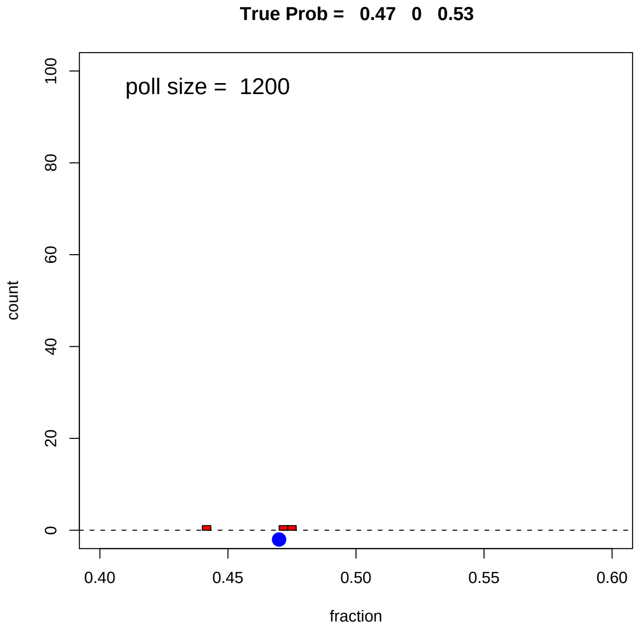
True Prob = 0.47 0 0.53
poll size = 1200
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

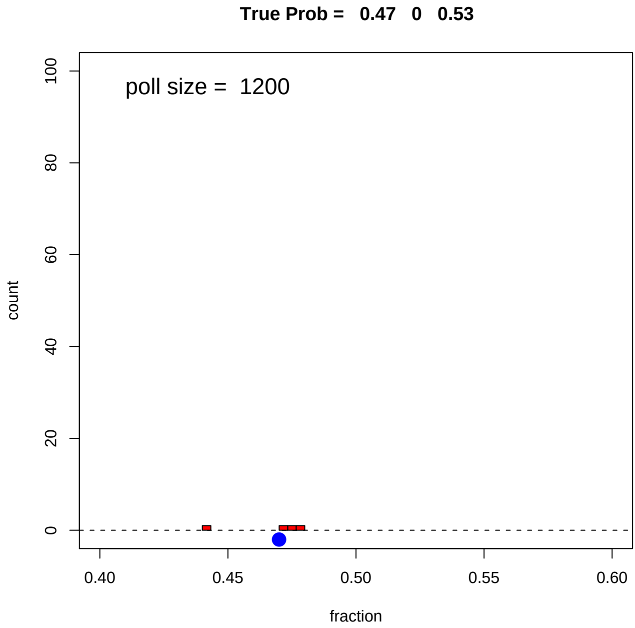
True Prob = 0.47 0 0.53
poll size = 1200
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

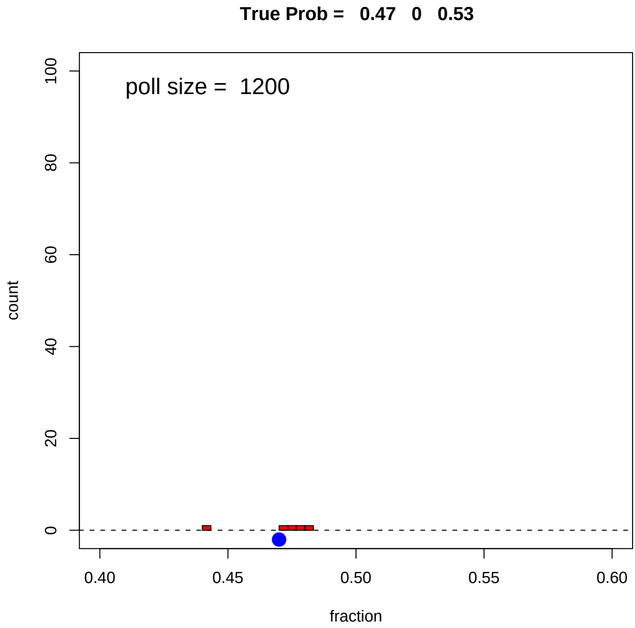
True Prob = 0.47 0 0.53
poll size = 1200
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

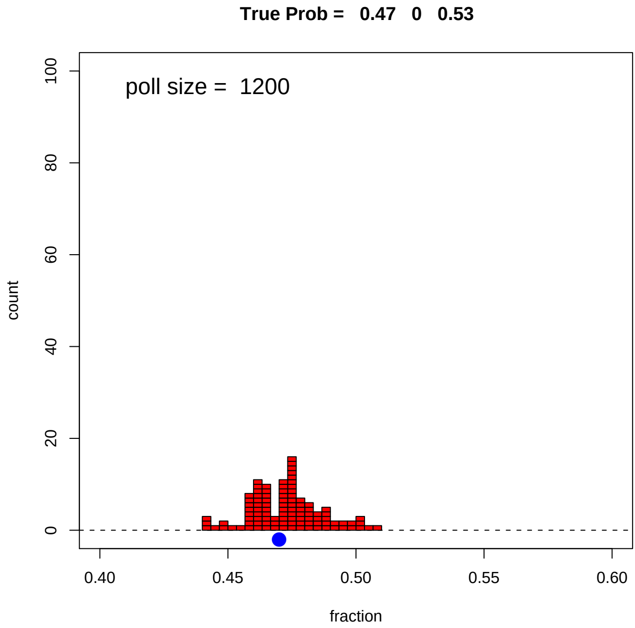
True Prob = 0.47 0 0.53
poll size = 1200
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

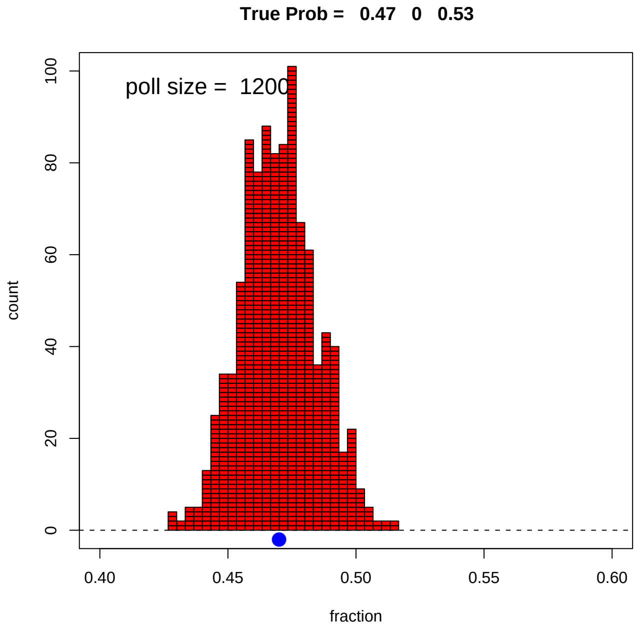
True Prob = 0.47 0 0.53
poll size = 1200
count
fraction
0.40
0.45
0.50
0.55
0.60
0
20
40
60
80
100

- This simulation suggests (visually) that the uncertainty is about ± 3 points
- Now it turns out that it is basically OK to ascribe this same level of uncertainty to a Gallup poll
- The *New York Times "Polling Standards"* includes:

> *UNDERSTANDING THE MARGIN OF SAMPLING ERROR*
>
> *A typical nationwide telephone poll of 1,000 respondents has a margin of sampling error of plus or minus three percentage points. This means that in 19 cases out of 20, overall results based on such samples will differ by no more than three percentage points in either direction from what would have been obtained by seeking out all American adults.*

- Examine the simulation again
- Question: What is the likelihood Obama will get more than 50% of the vote?
- The visualization suggests not

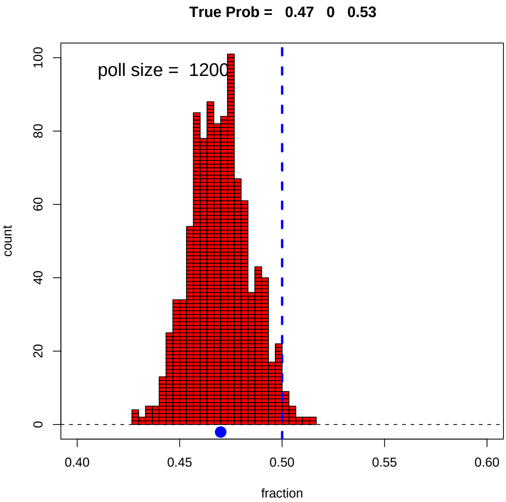

- Examine the simulation again
- Question: What is the likelihood Obama will get more than 50% of the vote?
- The visualization suggests not
- **What about the undecided voters???**

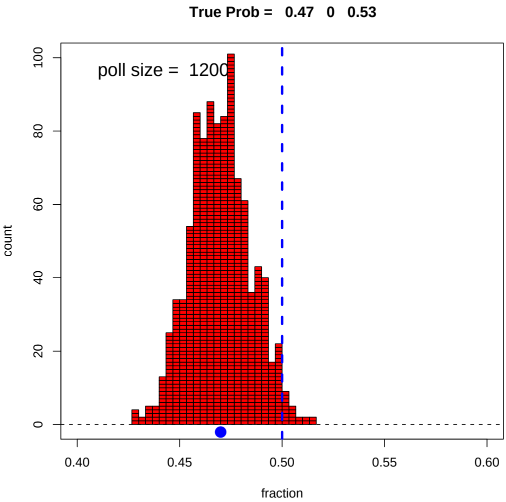

The Undecideds

- From the simulation point of view, we can “force” an answer from such individuals
- We would not use the same biased (47%) coin for the undecideds
- Seems obvious that these folks are truly on the fence: therefore, 50-50

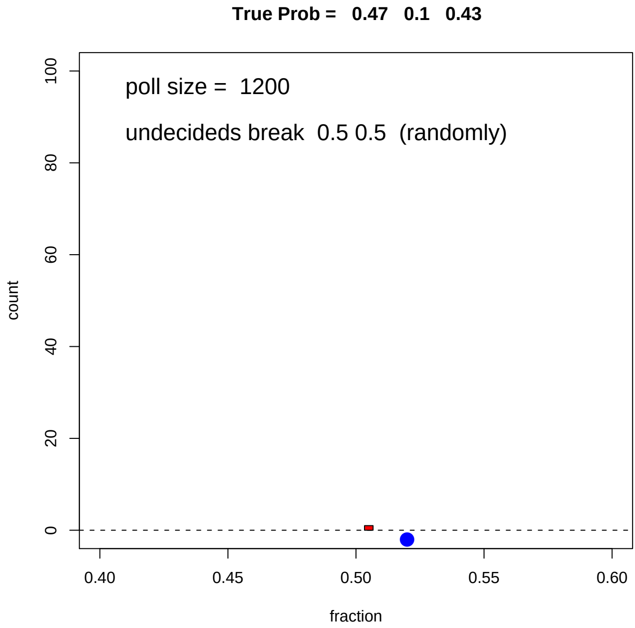
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.5 0.5 (randomly)
count
100
80
60
40
20
0
0.40
0.45
0.50
0.55
0.60
fraction

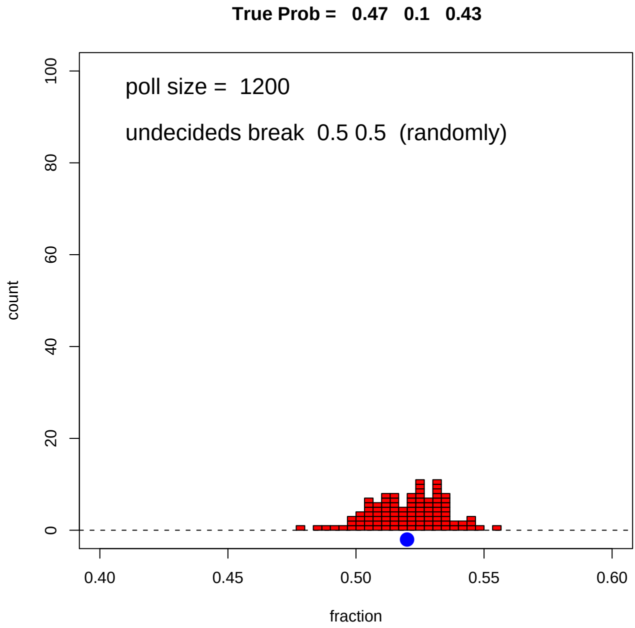

True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.5 0.5 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

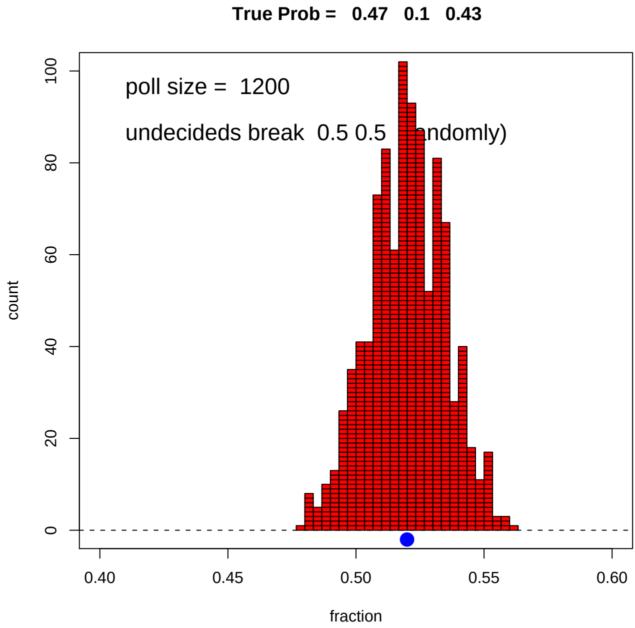
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.5 0.5 (randomly)
count
fraction
0 20 40 60 80 100
0.40 0.45 0.50 0.55 0.60

- Uncertainty is about the same
- Obama very likely to get at least 50%

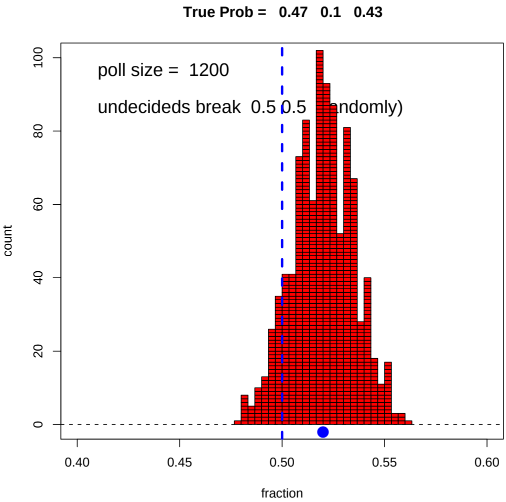

- Uncertainty is about the same
- Obama very likely to get at least 50%
- **So why report undecideds in any case?**

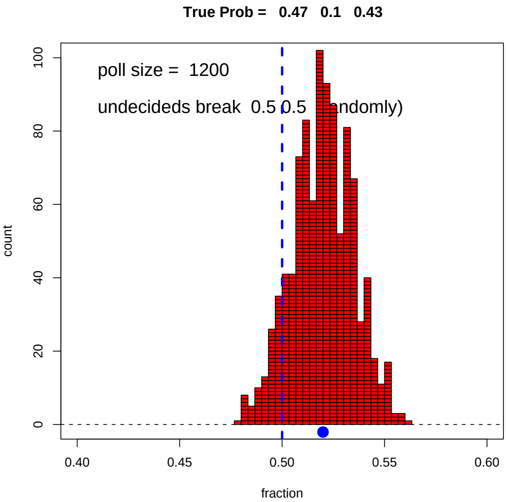

The Undecideds Revisited

- Treating undecideds as 50-50 is not the only choice

The Undecideds Revisited

- Treating undecideds as 50-50 is not the only choice
- Recall the common claim of an "October surprise"
- Perhaps undecideds are waiting for a last-minute reason to vote for Obama (or not)

The Undecideds Revisited

- Treating undecideds as 50-50 is not the only choice

- Recall the common claim of an "October surprise"
- Perhaps undecideds are waiting for a last-minute reason to vote for Obama (or not)

- *Model the October surprise as a 70-30 split, or as a 30-70 split, randomly*

The Undecideds Revisited

- Treating undecideds as 50-50 is not the only choice
- Recall the common claim of an "October surprise"
- Perhaps undecideds are waiting for a last-minute reason to vote for Obama (or not)
- *Model the October surprise as a 70-30 split, or as a 30-70 split, randomly*
- *Or as a 80-20 or 20-80 split, randomly*

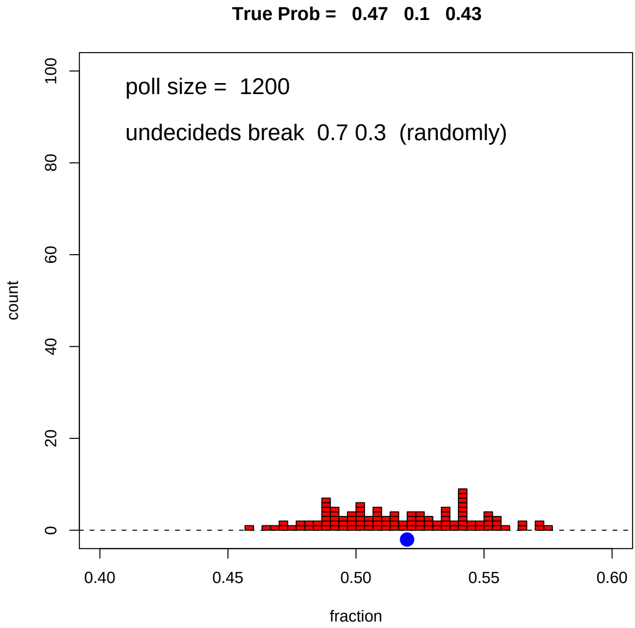
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.7 0.3 (randomly)
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

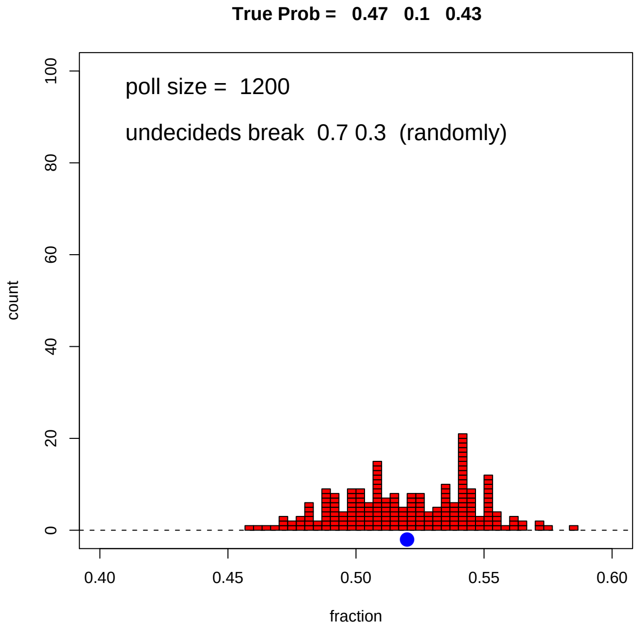
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.7 0.3 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

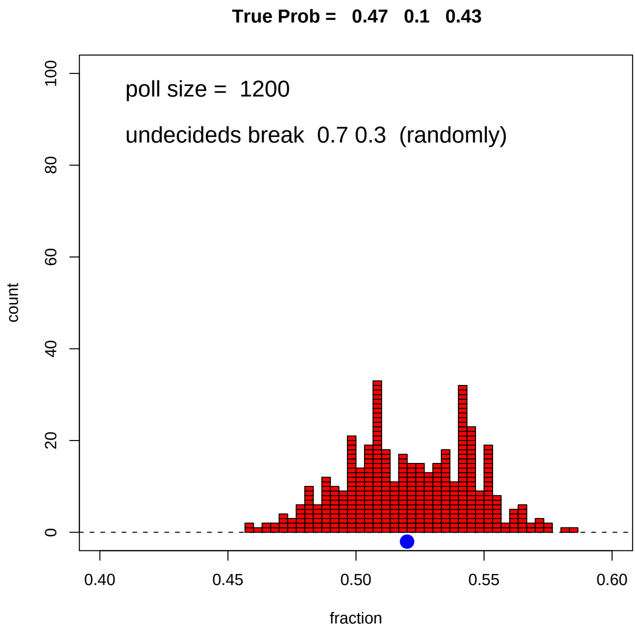
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.7 0.3 (randomly)
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

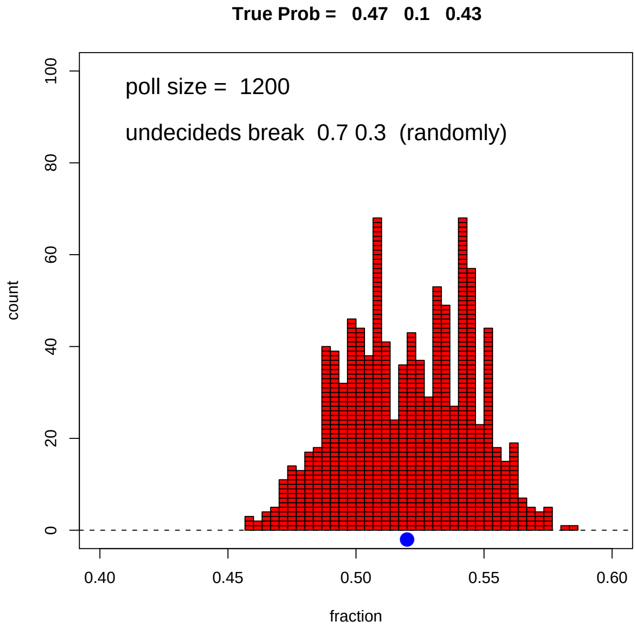
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.7 0.3 (randomly)
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

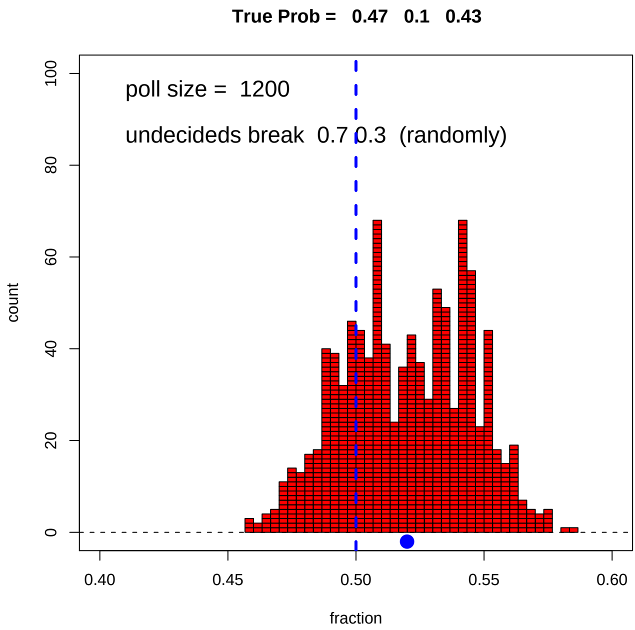

True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.7 0.3 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

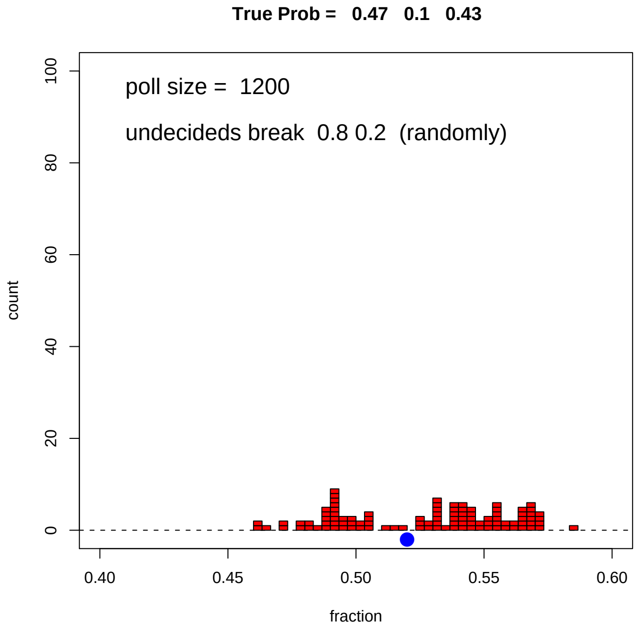
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.8 0.2 (randomly)
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

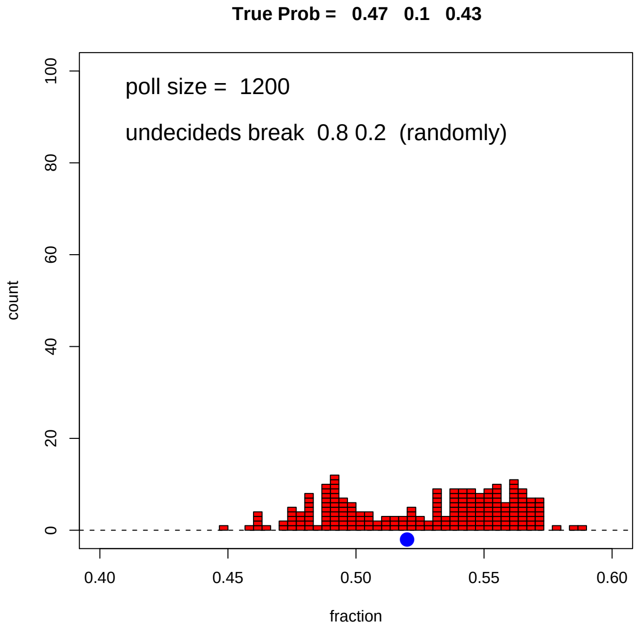
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.8 0.2 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

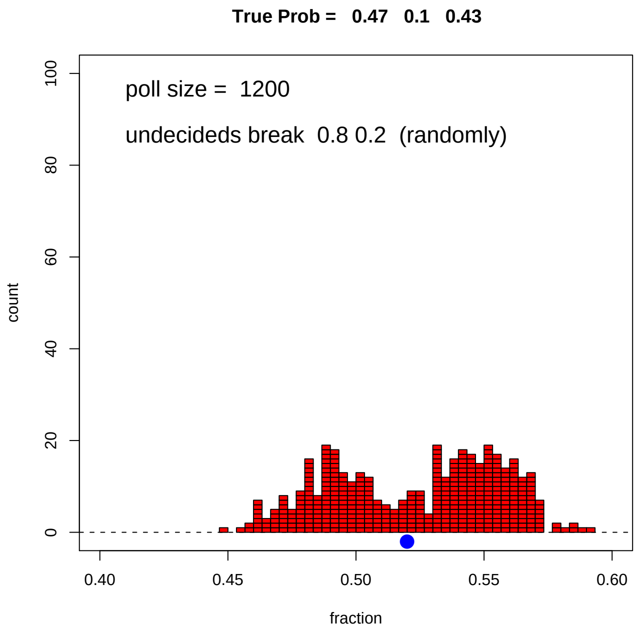
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.8 0.2 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

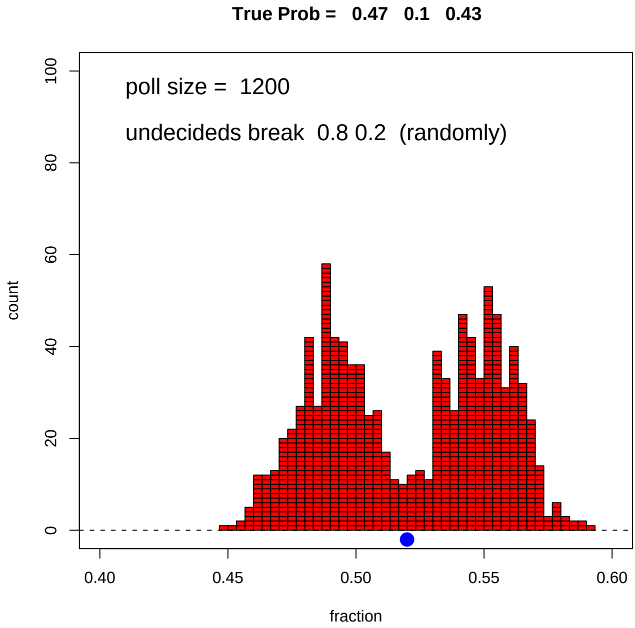
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.8 0.2 (randomly)
count
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60
fraction

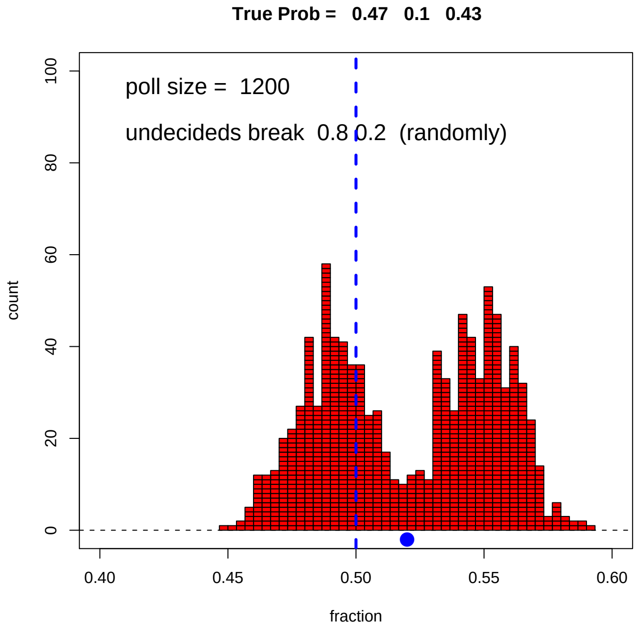
True Prob = 0.47 0.1 0.43
poll size = 1200
undecideds break 0.8 0.2 (randomly)
count
fraction
0
20
40
60
80
100
0.40
0.45
0.50
0.55
0.60

- The October-surprise theory is probably accurate
- Some believe McCain's *support* of the stimulus bill was exactly the October surprise, but in Obama's favor

- The October-surprise theory is probably accurate
- Some believe McCain's *support* of the stimulus bill was exactly the October surprise, but in Obama's favor

- One more simulation, focusing on cell phones

- The October-surprise theory is probably accurate
- Some believe McCain's *support* of the stimulus bill was exactly the October surprise, but in Obama's favor

- One more simulation, focusing on cell phones

- For a discussion of why cell phones are excluded from surveys: *www.pollster.com/blogs/cell_phones_and_political_surv.php*

- The October-surprise theory is probably accurate
- Some believe McCain's *support* of the stimulus bill was exactly the October surprise, but in Obama's favor

- One more simulation, focusing on cell phones

- For a discussion of why cell phones are excluded from surveys: *www.pollster.com/blogs/cell_phones_and_political_surv.php*

- (While we enjoy teaching closed-form expressions for uncertainty, simulation is much easier for more realistic models)

What is the Effect of Cell Phones?

- A Harris poll (4/08) showed 89% of adults have a cell phone (up from 77% in 12/06)
- 20% have no land line
- 14% only use a cell phone
- These 14% of voters are not equally distributed by age:

18 – 29	49%
30 – 39	22%
40 – 49	13%
50 – 64	11%
65–	6%

- Obama's supporters are also not distributed equally by age
- An 4/08 Gallup poll found the Obama/McCain supporters broke

18 – 29	57%	37%
30 – 49	46%	46%
50 – 64	44%	47%
65–	35%	51%

- Let us take as our model the cell phone Harris poll combined with the actual CNN 2008 election exit poll numbers

Age	*Fraction*	*Obama*	*McCain*	*Other*	*Cell Only*
18 – 24	10%	66%	32%	2%	49%
25 – 29	8%	66%	31%	3%	49%
30 – 39	18%	54%	44%	2%	22%
40 – 49	21%	49%	49%	2%	13%
50 – 64	27%	50%	49%	1%	11%
65–	16%	45%	53%	2%	6%

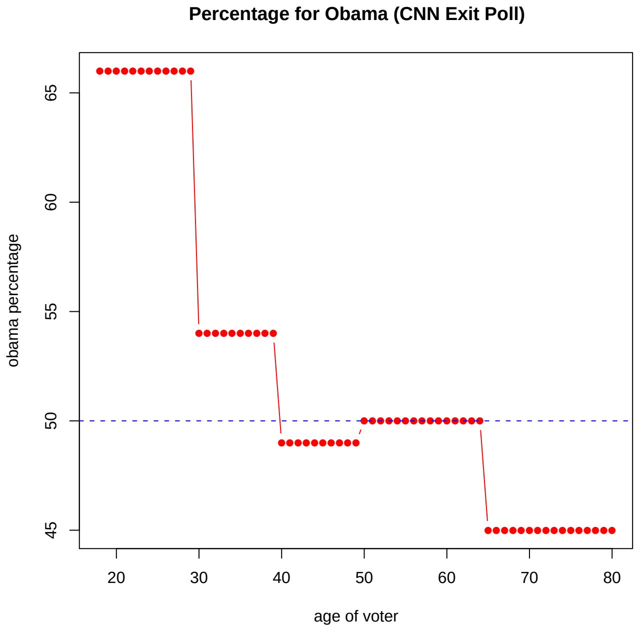

Percentage for Obama (CNN Exit Poll)
obama percentage
65
60
55
50
45
20
30
40
50
60
70
80
age of voter

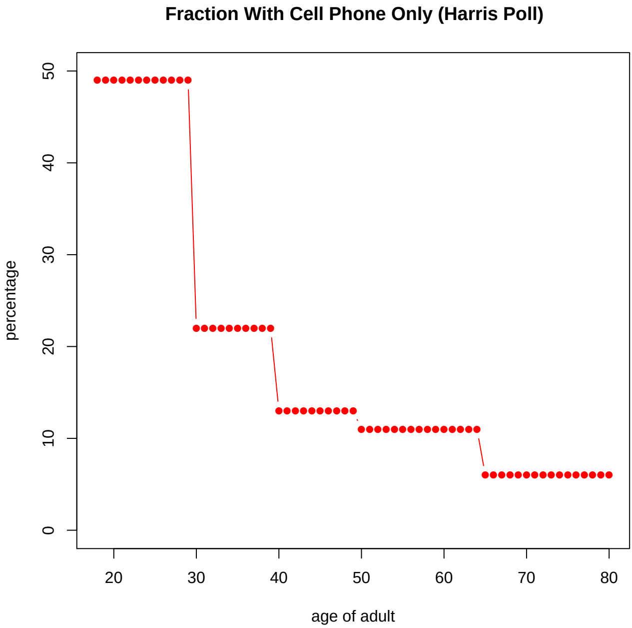

Fraction With Cell Phone Only (Harris Poll)
percentage
age of adult
0
10
20
30
40
50
20
30
40
50
60
70
80

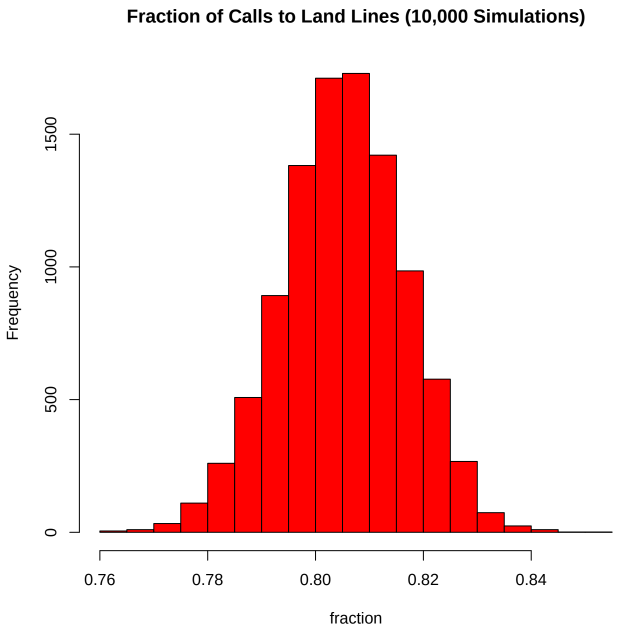
Fraction of Calls to Land Lines (10,000 Simulations)
Frequency
0
500
1000
1500
0.76
0.78
0.80
0.82
0.84
fraction

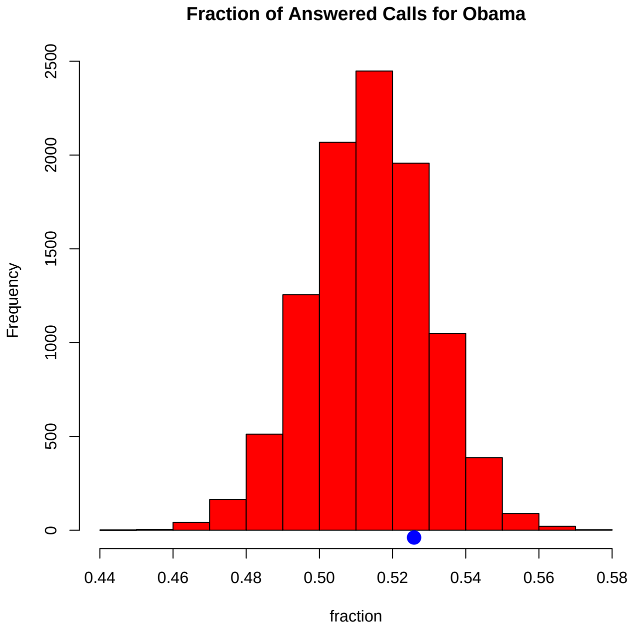
Fraction of Answered Calls for Obama
Frequency
0
500
1000
1500
2000
2500
0.44
0.46
0.48
0.50
0.52
0.54
0.56
0.58
fraction

- This simulation suggests (visually) that the uncertainty is not only somewhat greater than ± 3 points, but that it underestimates Obama's support (biased) by a point

- This simulation suggests (visually) that the uncertainty is not only somewhat greater than ± 3 points, but that it underestimates Obama's support (biased) by a point
- We can easily extend (complicate?) the simulation to account for other realistic polling realities (gender? race?)

- This simulation suggests (visually) that the uncertainty is not only somewhat greater than ± 3 points, but that it underestimates Obama's support (biased) by a point
- We can easily extend (complicate?) the simulation to account for other realistic polling realities (gender? race?)
- The *New York Times "Polling Standards"* also includes this paragraph:

> *UNDERSTANDING THE MARGIN OF SAMPLING ERROR*
>
> *The margin of sampling error is the only quantifiable error in a typical random sample telephone poll, but there are other errors too. The refusal rate, question order, interviewer techniques and question wording are all additional sources of error and bias in polls.*

- This simulation suggests (visually) that the uncertainty is not only somewhat greater than ± 3 points, but that it underestimates Obama's support (biased) by a point
- We can easily extend (complicate?) the simulation to account for other realistic polling realities (gender? race?)
- The *New York Times "Polling Standards"* also includes this paragraph:

> *UNDERSTANDING THE MARGIN OF SAMPLING ERROR*
>
> *The margin of sampling error is the only quantifiable error in a typical random sample telephone poll, but there are other errors too. The refusal rate, question order, interviewer techniques and question wording are all additional sources of error and bias in polls.*

- (No mention of cell phones)

Example 2: Is It Real?

Study this graphic from the first page of USA Today (10/13/06)

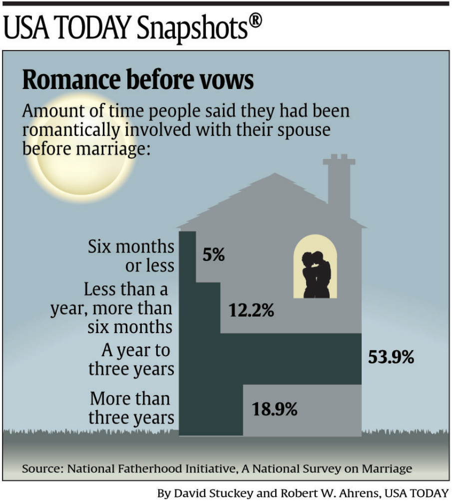

- Quality of USA Today graphics used to be error-prone
- Still tends towards junk art
- Low data-to-ink ratio (Tufte)

- Quality of USA Today graphics used to be error-prone
- Still tends towards junk art
- Low data-to-ink ratio (Tufte)

- Question here:

  Is there any structure apparent from such compressed data?

Is the Time-to-Marriage Normal?

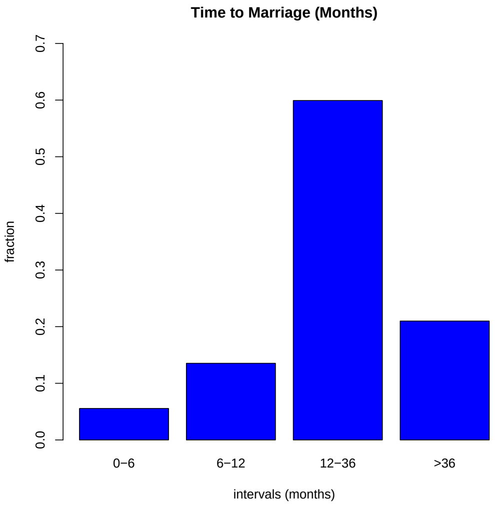

Is the Time-to-Marriage Normal?

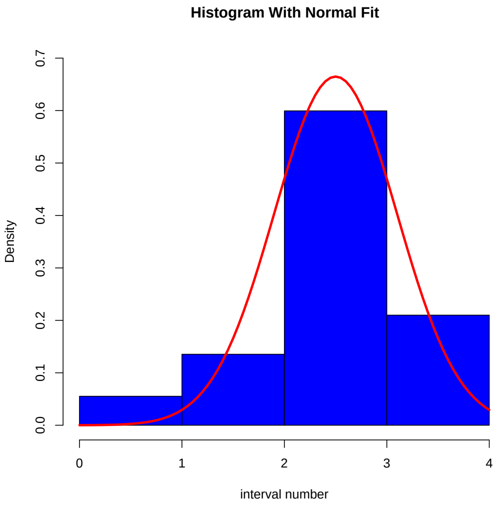

- Look at the numbers from the chart again

Age Range(Months)	*Fraction*
$0 - 6$	5.0%
$6 - 12$	12.2%
$12 - 36$	53.9%
$36 -$	18.9%

- Look at the numbers from the chart again

Age Range(Months)	*Fraction*
0 – 6	5.0%
6 – 12	12.2%
12 – 36	53.9%
36 –	18.9%

- *These only add up to 90%!*

- Using the original numbers from the survey cited

Age Range	*Number*	*Fraction*	(*Chart*)
0 – 6	181	15.00%	(5.0%)
6 – 12	147	12.18%	(12.2%)
12 – 36	651	53.94%	(53.9%)
36 –	228	18.89%	(18.9%)

- Using the original numbers from the survey cited

Age Range	*Number*	*Fraction*	(*Chart*)
0 – 6	181	15.00%	(5.0%)
6 – 12	147	12.18%	(12.2%)
12 – 36	651	53.94%	(53.9%)
36 –	228	18.89%	(18.9%)

- The error is entirely in the first interval

- Using the original numbers from the survey cited

Age Range	*Number*	*Fraction*	(*Chart*)
0 – 6	181	15.00%	(5.0%)
6 – 12	147	12.18%	(12.2%)
12 – 36	651	53.94%	(53.9%)
36 –	228	18.89%	(18.9%)

- The error is entirely in the first interval
- (The original survey gave these percentages — USA Today just copied the mistake)

Is the Time-to-Marriage Normal? (corrected data)

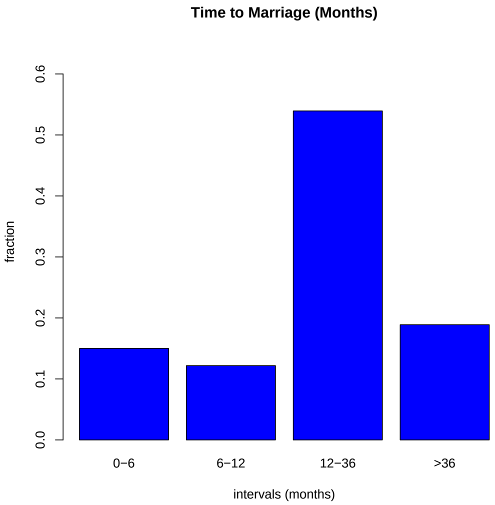

Is the Time-to-Marriage Normal? (corrected data)

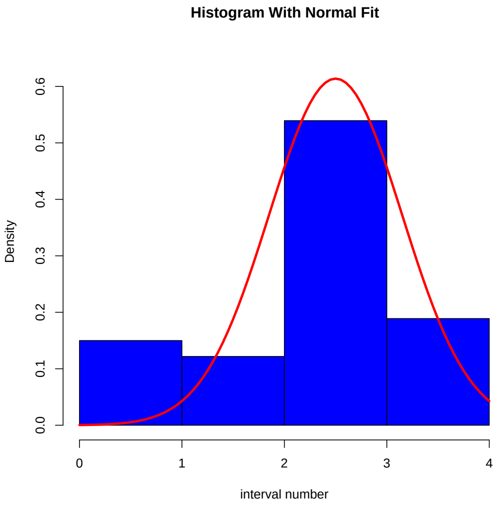

How To Handle the 3rd Interval?

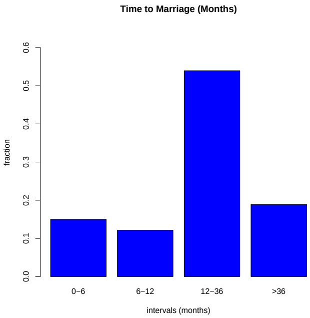

How To Handle the 3rd Interval?

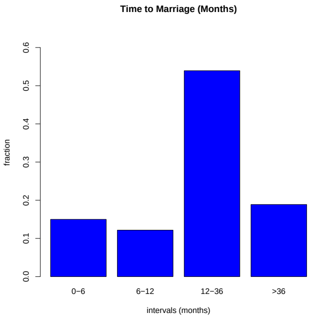

- The third bin is 4 times wider than the first two

How To Handle the 3rd Interval?

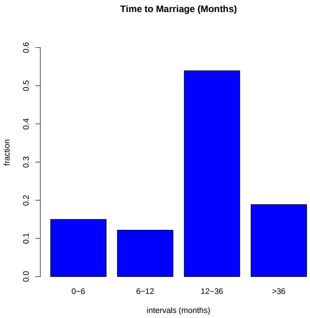

- The third bin is 4 times wider than the first two
- So split into 4 intervals — divide the count equally

Is the Time-to-Marriage Uniform?

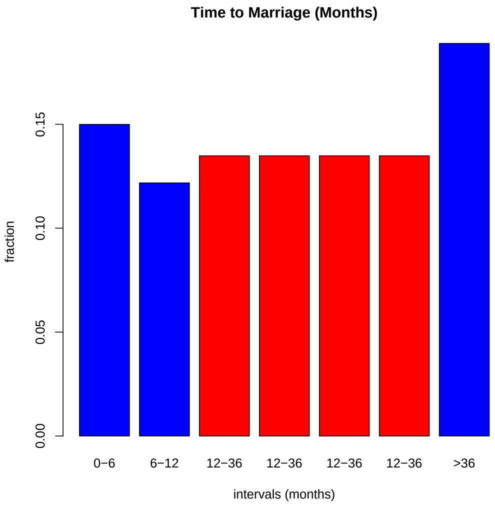

Is the Time-to-Marriage Uniform?

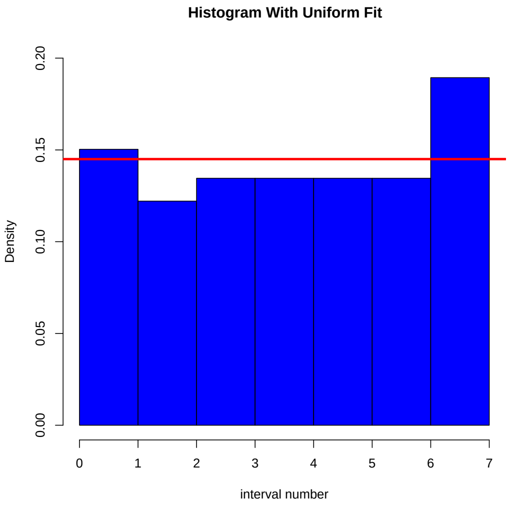

How To Handle the 4th Interval?

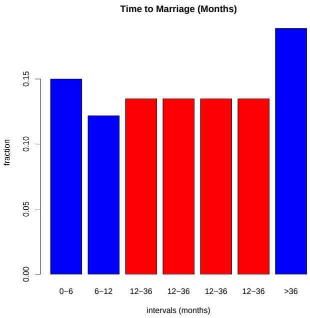

How To Handle the 4th Interval?

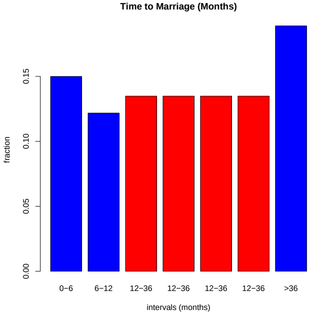

- The width of the fourth bin is indeterminate

How To Handle the 4th Interval?

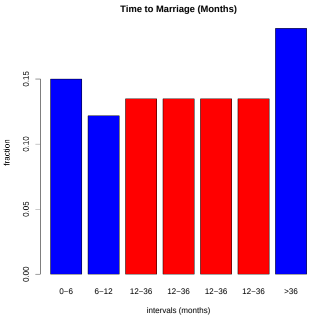

- The width of the fourth bin is indeterminate
- So split into 6 intervals — divide the count equally

Is the Time-to-Marriage Exponential?

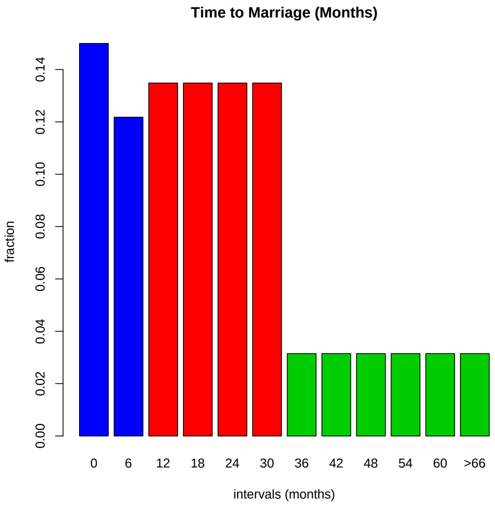

Is the Time-to-Marriage Exponential?

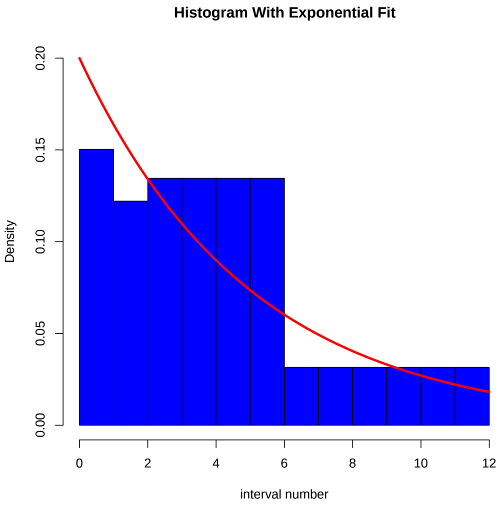

A New Density Estimator

- Consider a fine histogram that
  1. exactly matches the 4 interval proportions
  2. minimizes $\int f''(x)^2 dx$ (discrete approximation)
  3. is as wide as possible and nonnegative (4th bin problem)

A New Density Estimator

- Consider a fine histogram that
  1. exactly matches the 4 interval proportions
  2. minimizes $\int f''(x)^2 dx$ (discrete approximation)
  3. is as wide as possible and nonnegative (4th bin problem)

- Reference: Scott, D.W. and Scott, W.R. (2008), "Smoothed Histograms for Frequency Data on Irregular Intervals," *The American Statistician*, 62, 256–261

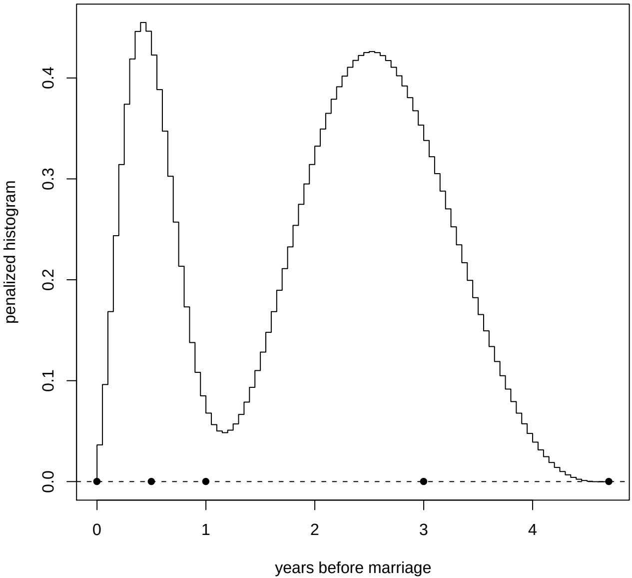

penalized histogram
0.0
0.1
0.2
0.3
0.4
0
1
2
3
4
years before marriage

Is the Time-to-Marriage Really Bimodal?

- Simulation again
- Start with the original interval counts

181	147	651	228

Is the Time-to-Marriage Really Bimodal?

- Simulation again
- Start with the original interval counts

181	147	651	228

- Use these empirical proportions to generate another (multinomial) sample of size 1207

185	161	649	212

Is the Time-to-Marriage Really Bimodal?

- Simulation again
- Start with the original interval counts

181 147 651 228

- Use these empirical proportions to generate another (multinomial) sample of size 1207

185 161 649 212

- Compute the new histogram and see if it is bimodal

Is the Time-to-Marriage Really Bimodal?

- Simulation again
- Start with the original interval counts

181	147	651	228

- Use these empirical proportions to generate another (multinomial) sample of size 1207

185	161	649	212

- Compute the new histogram and see if it is bimodal
- cf. bootstrapping (repeat 10 times)

Is the Time-to-Marriage Really Bimodal?

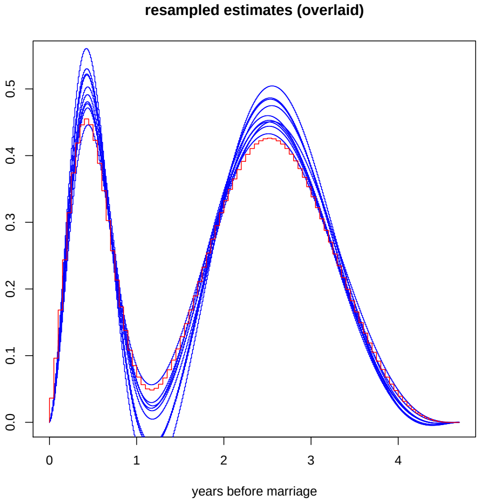

- Simulation not only useful for usual uncertainty measures
- Also useful for assessing the veracity of (apparent) features in the data

- Simulation not only useful for usual uncertainty measures
- Also useful for assessing the veracity of (apparent) features in the data

- Data errors creep in

- Simulation not only useful for usual uncertainty measures
- Also useful for assessing the veracity of (apparent) features in the data

- Data errors creep in
- How can one be assured no theoretical errors in a simulation code?
- How can one be assured no bugs in a simulation code?

- Simulation not only useful for usual uncertainty measures
- Also useful for assessing the veracity of (apparent) features in the data

- Data errors creep in
- How can one be assured no theoretical errors in a simulation code?
- How can one be assured no bugs in a simulation code?

- Answer: ——–-

Example 3: How To Visualize Uncertainty in Models with Many Parameters?

- Tables or graphs or ?

Example 3: How To Visualize Uncertainty in Models with Many Parameters?

- Tables or graphs or ?
- Consider multiple regression (with normal errors)

$$y = \theta_0 + \theta_1 x_1 + \theta_2 x_2 + \cdots + \theta_p x_p + \sigma z$$

Example 3: How To Visualize Uncertainty in Models with Many Parameters?

- Tables or graphs or ?
- Consider multiple regression (with normal errors)

$$y = \theta_0 + \theta_1 x_1 + \theta_2 x_2 + \cdots + \theta_p x_p + \sigma z$$

- Collect n instances of this model; least-squares solution is

$$\hat{\theta} = (X^t X)^{-1} X^t Y$$

and

$$\text{variance}(\hat{\theta}) \approx \hat{\sigma}^2 (X^t X)^{-1}$$

- variance($\hat{\theta}$) is a matrix: therefore, $\hat{\theta}_k$ and $\hat{\theta}_\ell$ are correlated, and the diagonal elements contain the variance($\hat{\theta}_k$)

Estimation of $\boldsymbol{\theta}_p \in \Re^p$

- how to understand uncertainty of $\hat{\boldsymbol{\theta}}_p$?
  - interpreting individual parameters $\hat{\boldsymbol{\theta}}_p^{(i)}$
  - stability related to n and collinearity
- beyond pairwise correlations, hard to explain/understand
- visualization helpful/better than analytics?
- sensitivity analysis often mentioned

Estimation of $\boldsymbol{\theta}_p \in \Re^p$

- how to understand uncertainty of $\hat{\boldsymbol{\theta}}_p$?
  - interpreting individual parameters $\hat{\boldsymbol{\theta}}_p^{(i)}$
  - stability related to n and collinearity
- beyond pairwise correlations, hard to explain/understand
- visualization helpful/better than analytics?
- sensitivity analysis often mentioned

- consider examples from regression

- parameter-by-parameter confidence intervals

$$\beta_k \in \hat{\beta}_k \pm t_{.975}\sqrt{\hat{\Sigma}_{kk}}$$

  - too limited
  - changing one variable at a time
  - ("the effect of the k-th variable, *all other things being equal* ...")

- parameter-by-parameter confidence intervals

$$\beta_k \in \hat{\beta}_k \pm t_{.975}\sqrt{\hat{\Sigma}_{kk}}$$

  - too limited
  - changing one variable at a time
  - ("the effect of the k-th variable, *all other things being equal* ...")
- simultaneous confidence region given by (hyper-ellipse)

$$(\beta - \hat{\beta})^t(\hat{\sigma}^2(X^tX)^{-1})^{-1}(\beta - \hat{\beta}) \leq \frac{n \cdot p}{n-p-1} f_{p,n-p-1}(0.95)$$

- parameter-by-parameter confidence intervals

$$\beta_k \in \hat{\beta}_k \pm t_{.975}\sqrt{\hat{\Sigma}_{kk}}$$

  - too limited
  - changing one variable at a time
  - ("the effect of the k-th variable, *all other things being equal ...*")

- simultaneous confidence region given by (hyper-ellipse)

$$(\beta - \hat{\beta})^t(\hat{\sigma}^2(X^tX)^{-1})^{-1}(\beta - \hat{\beta}) \leq \frac{n \cdot p}{n-p-1} f_{p,n-p-1}(0.95)$$

- example: law school admissions data ($n = 15$, Efron)

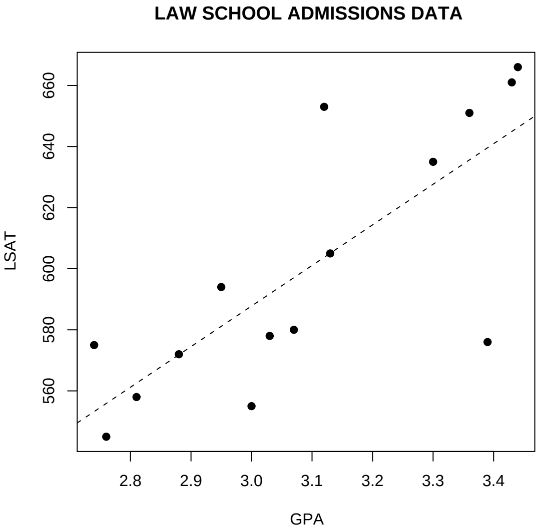
LAW SCHOOL ADMISSIONS DATA
LSAT
560
580
600
620
640
660
2.8
2.9
3.0
3.1
3.2
3.3
3.4
GPA

LAW SCHOOL ADMISSIONS DATA

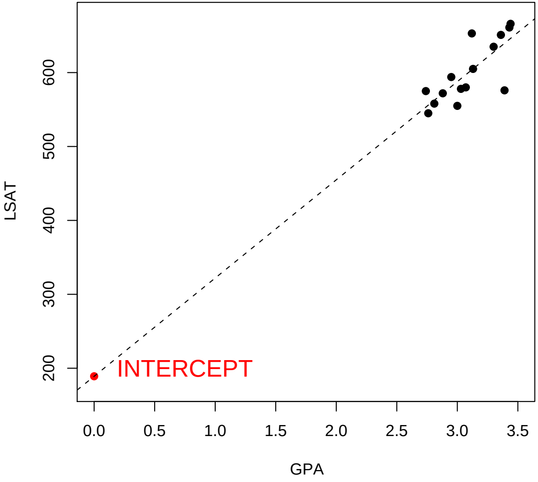

- look at confidence interval for intercept, $\hat{\beta}_0$
- look at confidence interval for slope, $\hat{\beta}_1$
- bivariate confidence interval (rectangle) — too conservative

- look at confidence interval for intercept, $\hat{\beta}_0$
- look at confidence interval for slope, $\hat{\beta}_1$
- bivariate confidence interval (rectangle) — too conservative
- should be an ellipse (compare)

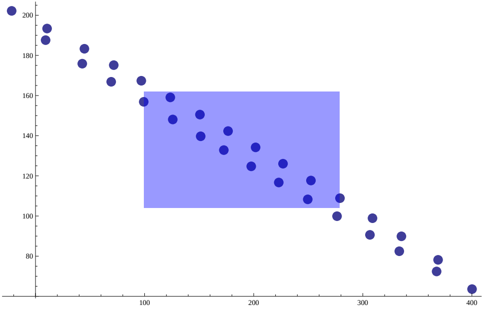

- (see Mathematica animations — prg2.nb)

- visualization improvement?
    - change focus from parameters
    - choose a “smooth path” through the confidence ellipse (rather than on boundary)
    - examine the corresponding *model visualization*

- visualization improvement?
  - change focus from parameters
  - choose a "smooth path" through the confidence ellipse (rather than on boundary)
  - examine the corresponding *model visualization*
- facilitates understanding the impact of parameter uncertainty on model and model predictions

- visualization improvement?
    - change focus from parameters
    - choose a "smooth path" through the confidence ellipse (rather than on boundary)
    - examine the corresponding *model visualization*
- facilitates understanding the impact of parameter uncertainty on model and model predictions
- eigenvectors of $\hat{\Sigma}$ convenient "smooth path" in $\Re^p$
    - $\lambda_1 = 8879.3$ and $\lambda_2 = 4.4$ (data uncentered/unscaled)
    - $\lambda_1 = 0.0285$ and $\lambda_2 = 0.0266$ (data standardized)
- centering and scaling do affect perception (always center/standardize)

- for $p > 2$, eigenvector strategy highly effective

- for $p > 2$, eigenvector strategy highly effective
- how to visualize regression when $p > 2$?
- use alternative parallel plot (for additive model)
    - ironically, changes focus back to the variables
    - everything standardized, so easy to compare effects

- for $p > 2$, eigenvector strategy highly effective
- how to visualize regression when $p > 2$?
- use alternative parallel plot (for additive model)
  - ironically, changes focus back to the variables
  - everything standardized, so easy to compare effects
- look at the stackloss data (3 predictors + int), predicting stack.loss
  - $\lambda_2/\lambda_1 = 0.314$
  - $\lambda_3/\lambda_1 = 0.197$
  - $\lambda_4/\lambda_1 = 0.097$

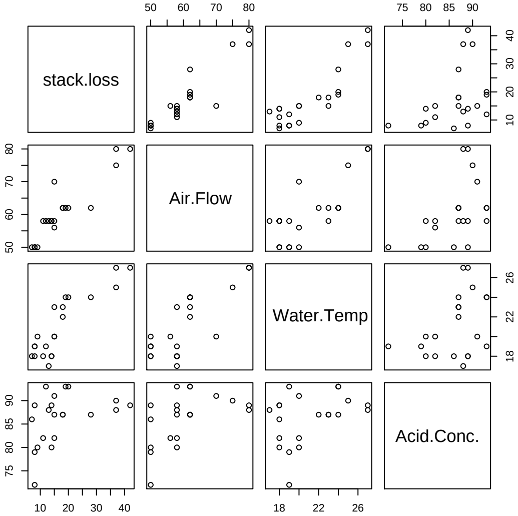
stack.loss
Air.Flow
Water.Temp
Acid.Conc.
50
60
70
80
75
80
85
90
40
30
20
10
80
70
60
50
26
24
22
18
90
85
80
75
10
20
30
40
18
22
26

- finally, we will look at the transformed Boston Housing data (13 predictors + int), predicting "median house value" ($R^2 = 0.77$)
    - $\lambda_2/\lambda_1 = 0.695$
    - $\lambda_3/\lambda_1 = 0.624$
    - $\lambda_4/\lambda_1 = 0.548$
    - $\lambda_5/\lambda_1 = 0.385$
    - $\vdots$
    - $\lambda_{14}/\lambda_1 = 0.013$
- (Mathematica notebook prg4.nb)

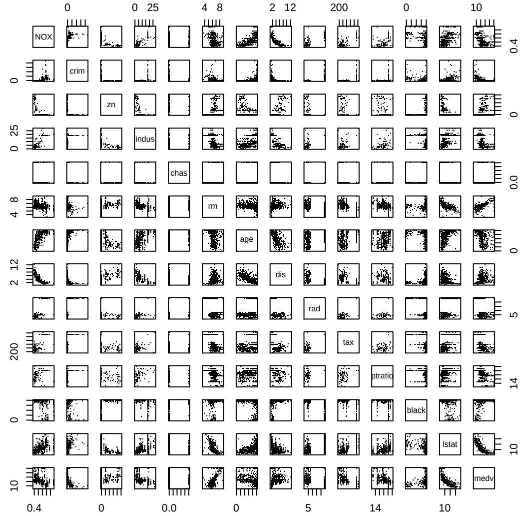

NOX
crim
zn
indus
chas
rm
age
dis
rad
tax
ptratio
black
lstat
medv

- hard to look at correlation matrix and “see” higher-dimensional collinearities

- eigenvectors sorted by the “most active” set of coefficients

Concluding Remarks

- estimates without uncertainty bad practice

Concluding Remarks

- estimates without uncertainty bad practice
- parameter-by-parameter uncertainty misleading (tables)

Concluding Remarks

- estimates without uncertainty bad practice
- parameter-by-parameter uncertainty misleading (tables)
- visualization helps understand complex parameter interactions
  - can visualize the parameters
  - or can visualize the resulting model as the parameters are (smoothly) varied within the hyperellipse

Concluding Remarks

- estimates without uncertainty bad practice
- parameter-by-parameter uncertainty misleading (tables)
- visualization helps understand complex parameter interactions
  - can visualize the parameters
  - or can visualize the resulting model as the parameters are (smoothly) varied within the hyperellipse
- modern computer simulations (economy, weather,...) have thousands of parameters — uncertainty poorly understood

Concluding Remarks

- estimates without uncertainty bad practice
- parameter-by-parameter uncertainty misleading (tables)
- visualization helps understand complex parameter interactions
  - can visualize the parameters
  - or can visualize the resulting model as the parameters are (smoothly) varied within the hyperellipse
- modern computer simulations (economy, weather,...) have thousands of parameters — uncertainty poorly understood
- help practicing scientists understand the same lesson; convince scientists to reserve some computation for uncertainty

Concluding Remarks

- estimates without uncertainty bad practice
- parameter-by-parameter uncertainty misleading (tables)
- visualization helps understand complex parameter interactions
  - can visualize the parameters
  - or can visualize the resulting model as the parameters are (smoothly) varied within the hyperellipse
- modern computer simulations (economy, weather,...) have thousands of parameters — uncertainty poorly understood
- help practicing scientists understand the same lesson; convince scientists to reserve some computation for uncertainty
- *Thank you.*

http://www.stat.rice.edu/∼scottdw