CSE 442 - Data Visualization

Uncertainty

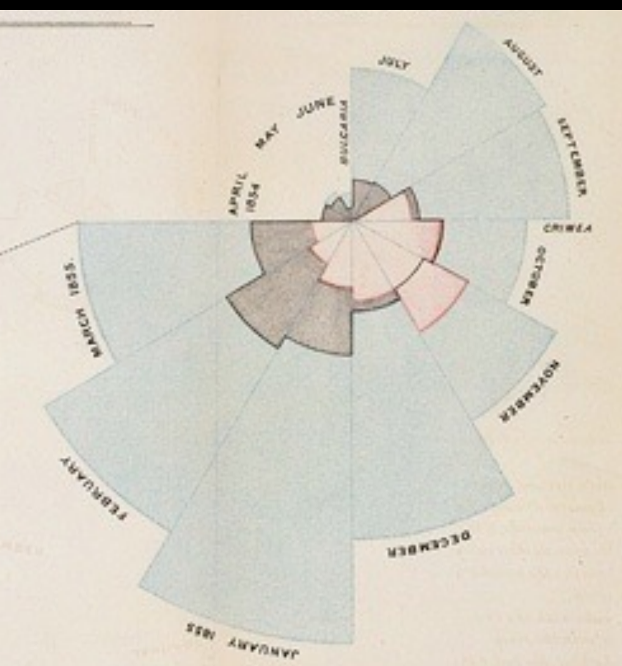

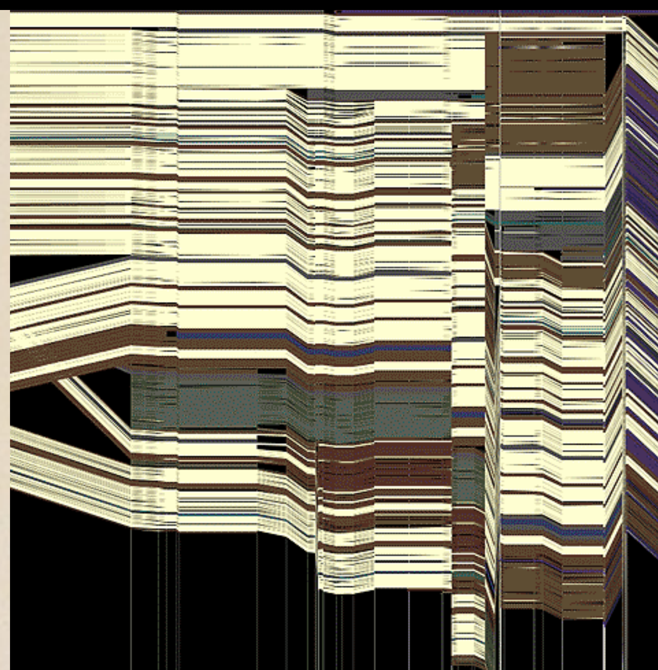

Jeffrey Heer University of Washington

(with significant material from Michael Correll)

The Visualization Pipeline

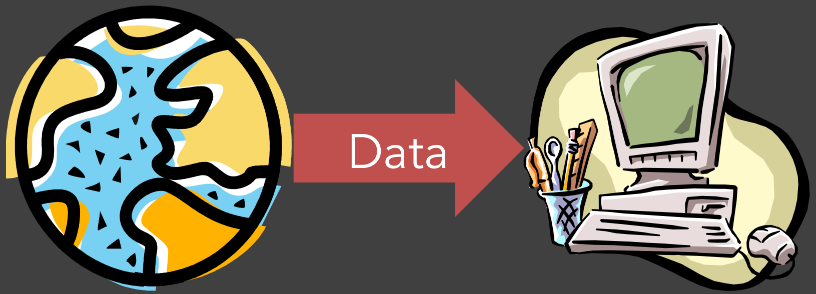

Collection/Curation

The Visualization Pipeline

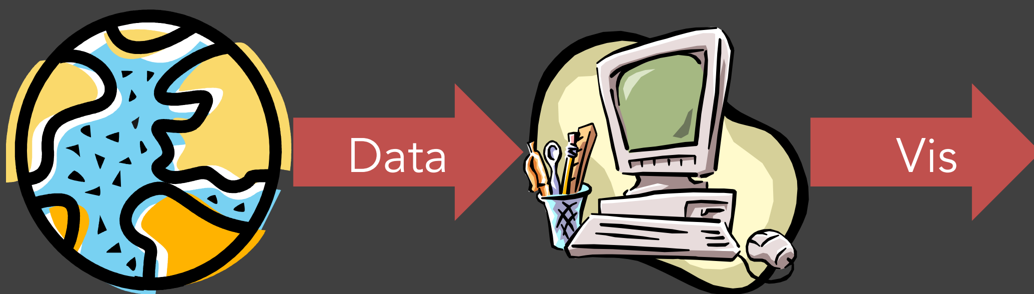

Collection/Curation

Design

The Visualization Pipeline

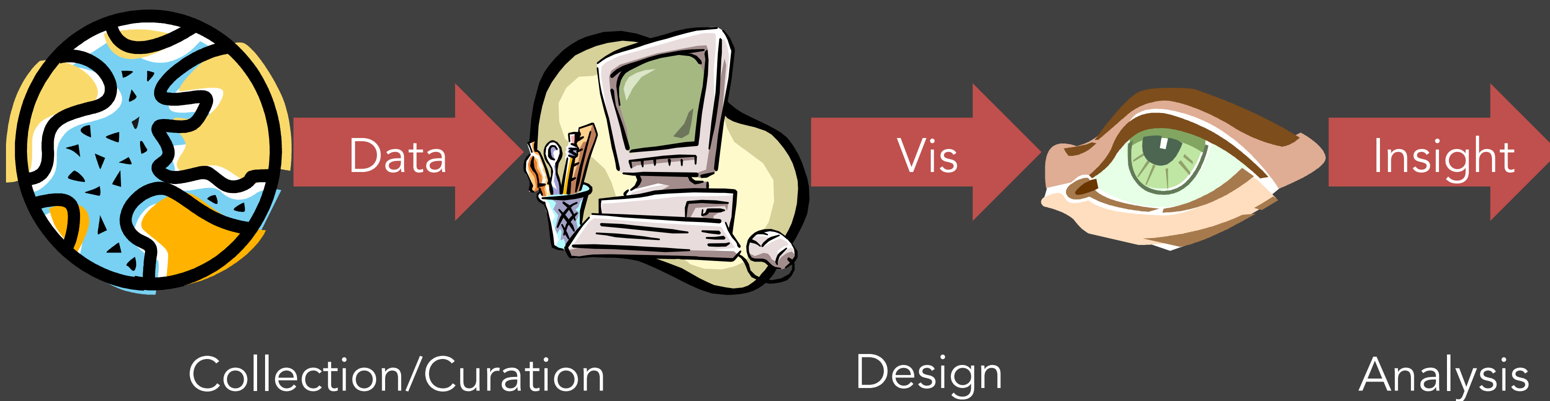

The Visualization Pipeline?

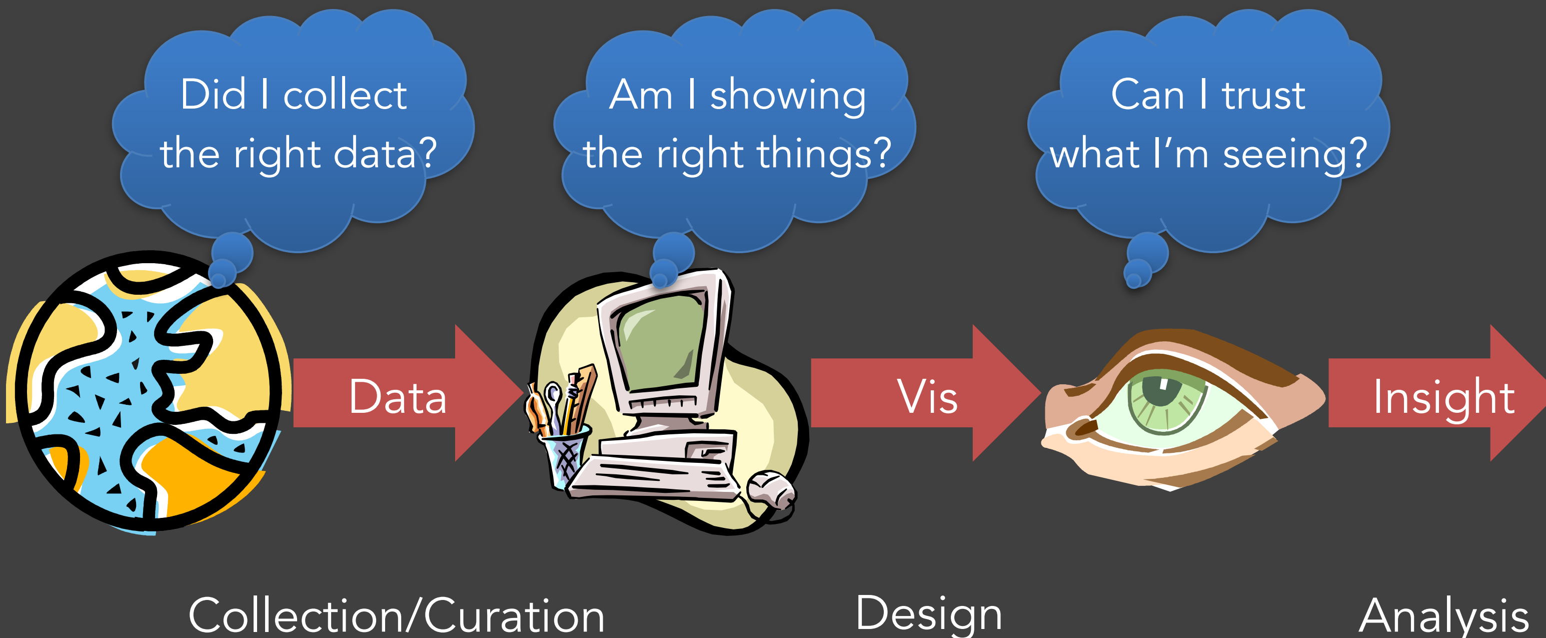

Unknown Unknowns

Things “Uncertainty” Can Mean

Doubt

Risk

Variability

Error

Lack of Knowledge

Hedging

…

Uncertainty Visualization

There are different **types** and **sources** of uncertainty.

We can **quantify** or **model** our uncertainty.

The visual presentation of uncertainty can **clash** with cognitive and perceptual biases.

Terminology

Terminology

Aleatory Uncertainty

Epistemic Uncertainty

Type I error

Type II error

Precision

Bias

Aleatory Uncertainty

Aleatory Uncertainty

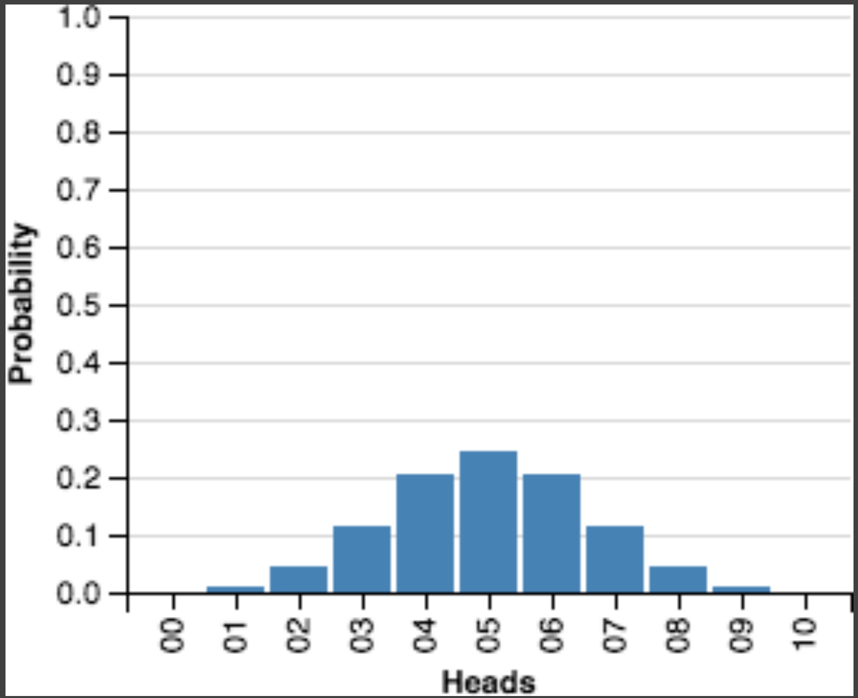

Aleatory Uncertainty

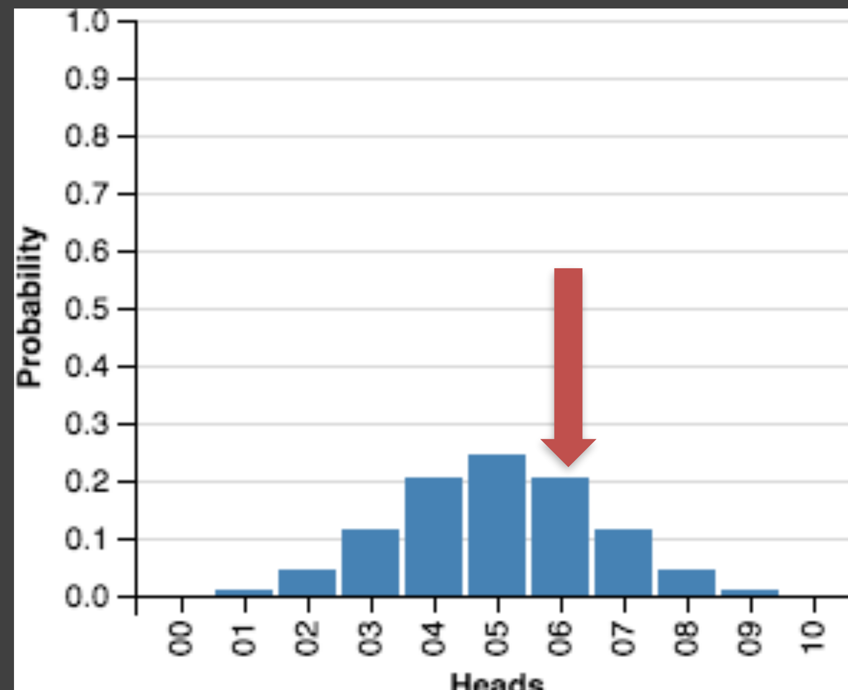

Aleatory Uncertainty

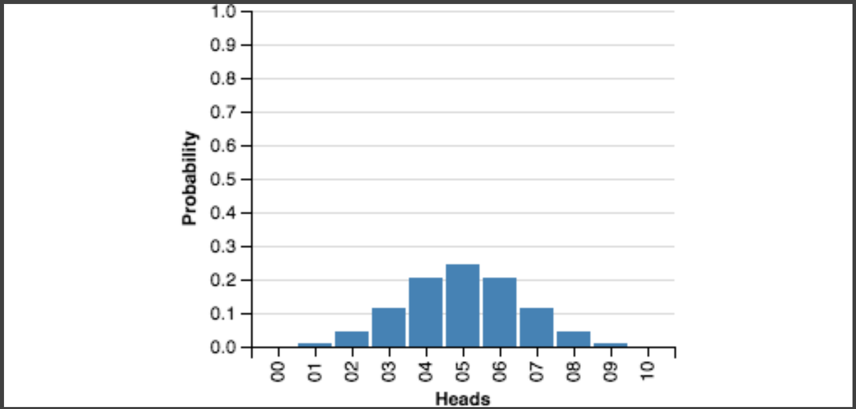

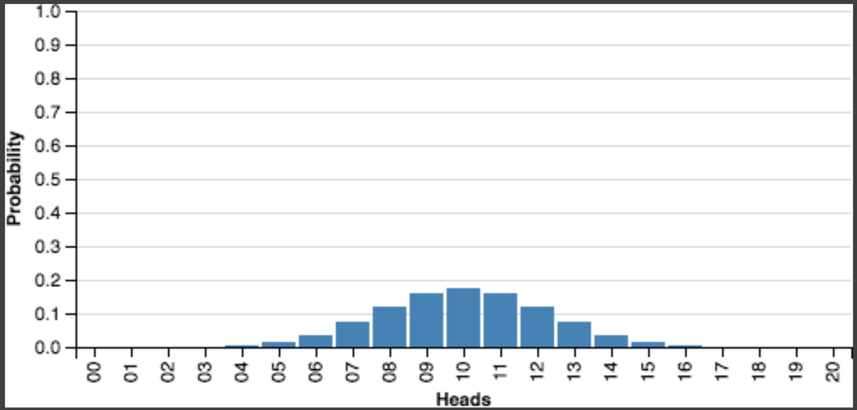

John Edmund Kerrich

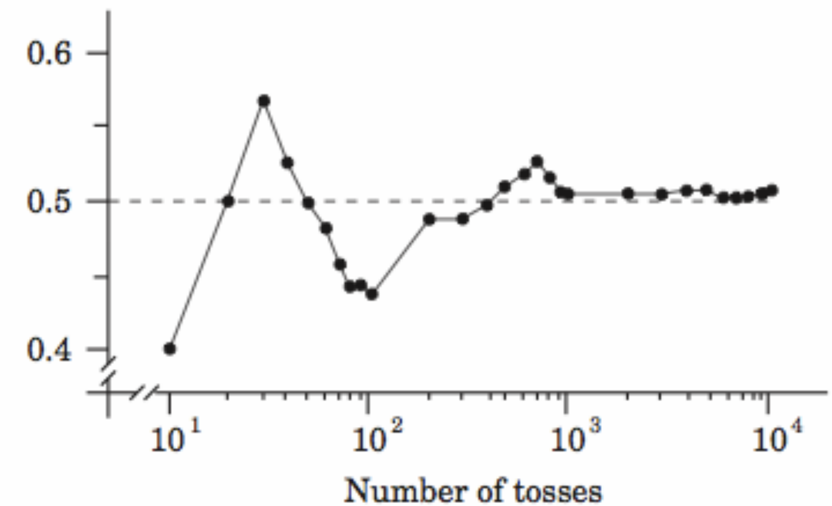

FIGURE 4.1.1 Proportion of heads versus number of tosses for John Kerrich's coin-tossing experiment.

Epistemic Uncertainty

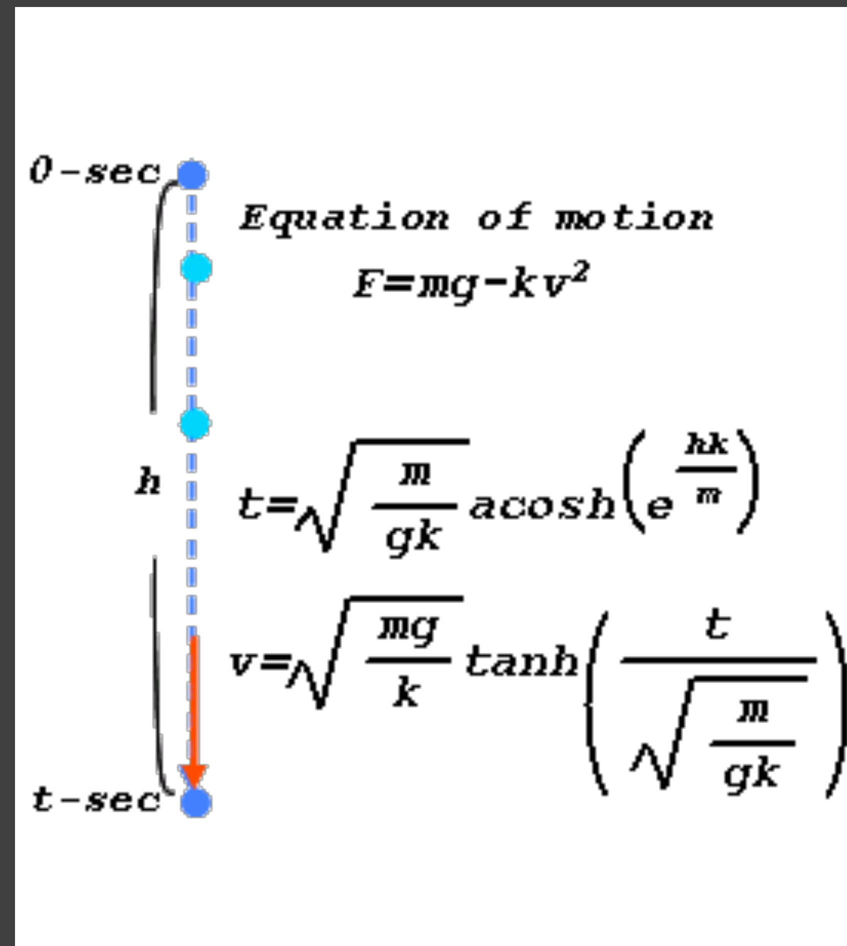

Uncertainty Types

Aleatory

Variability: things that we don't know (but can reason about the likelihood of).

Epistemic

Things we could in principle know for certain, but have not measured.

Should I Bring an Umbrella?

Decision Uncertainty

"50% Chance of Rain"

Risk and Error

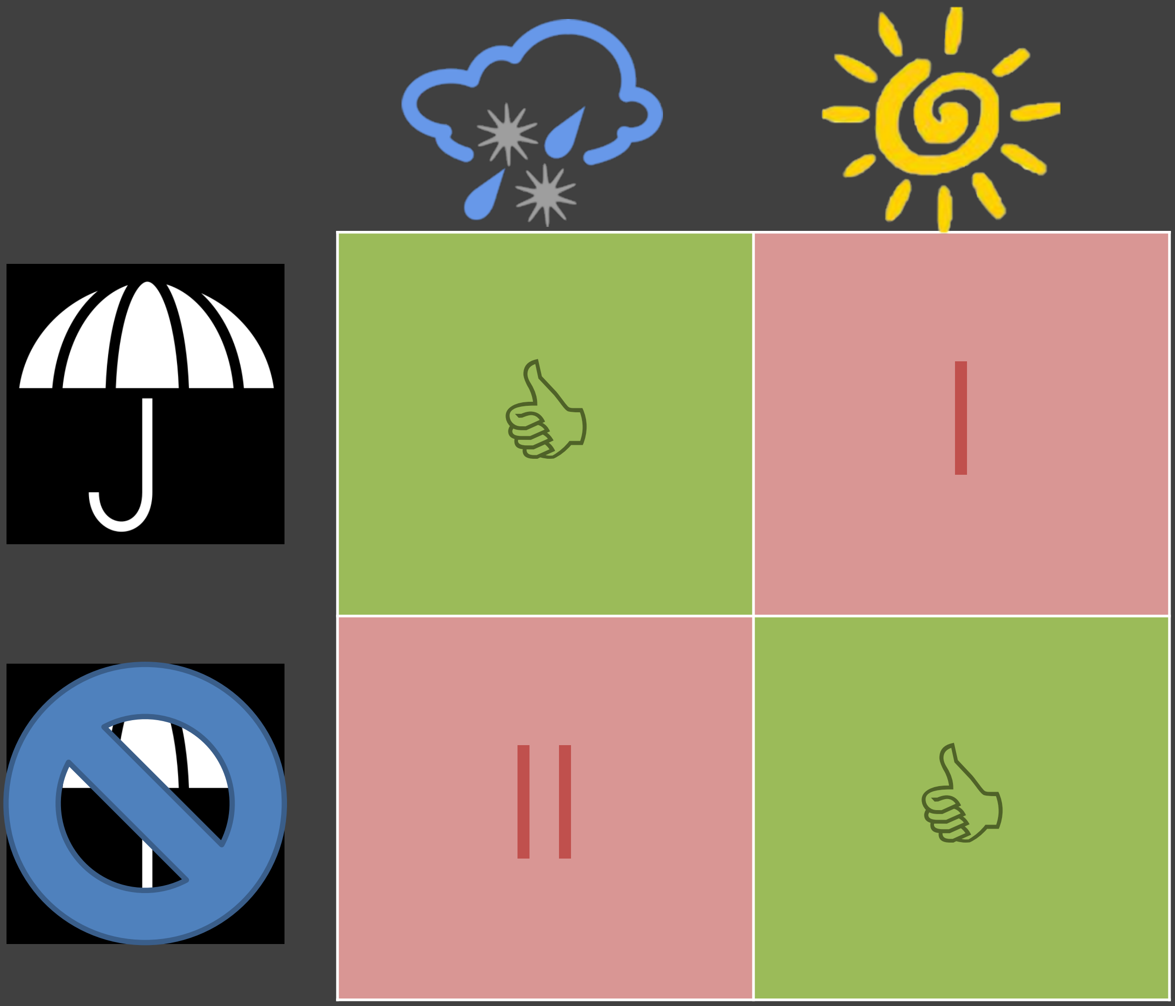

Type I and Type II

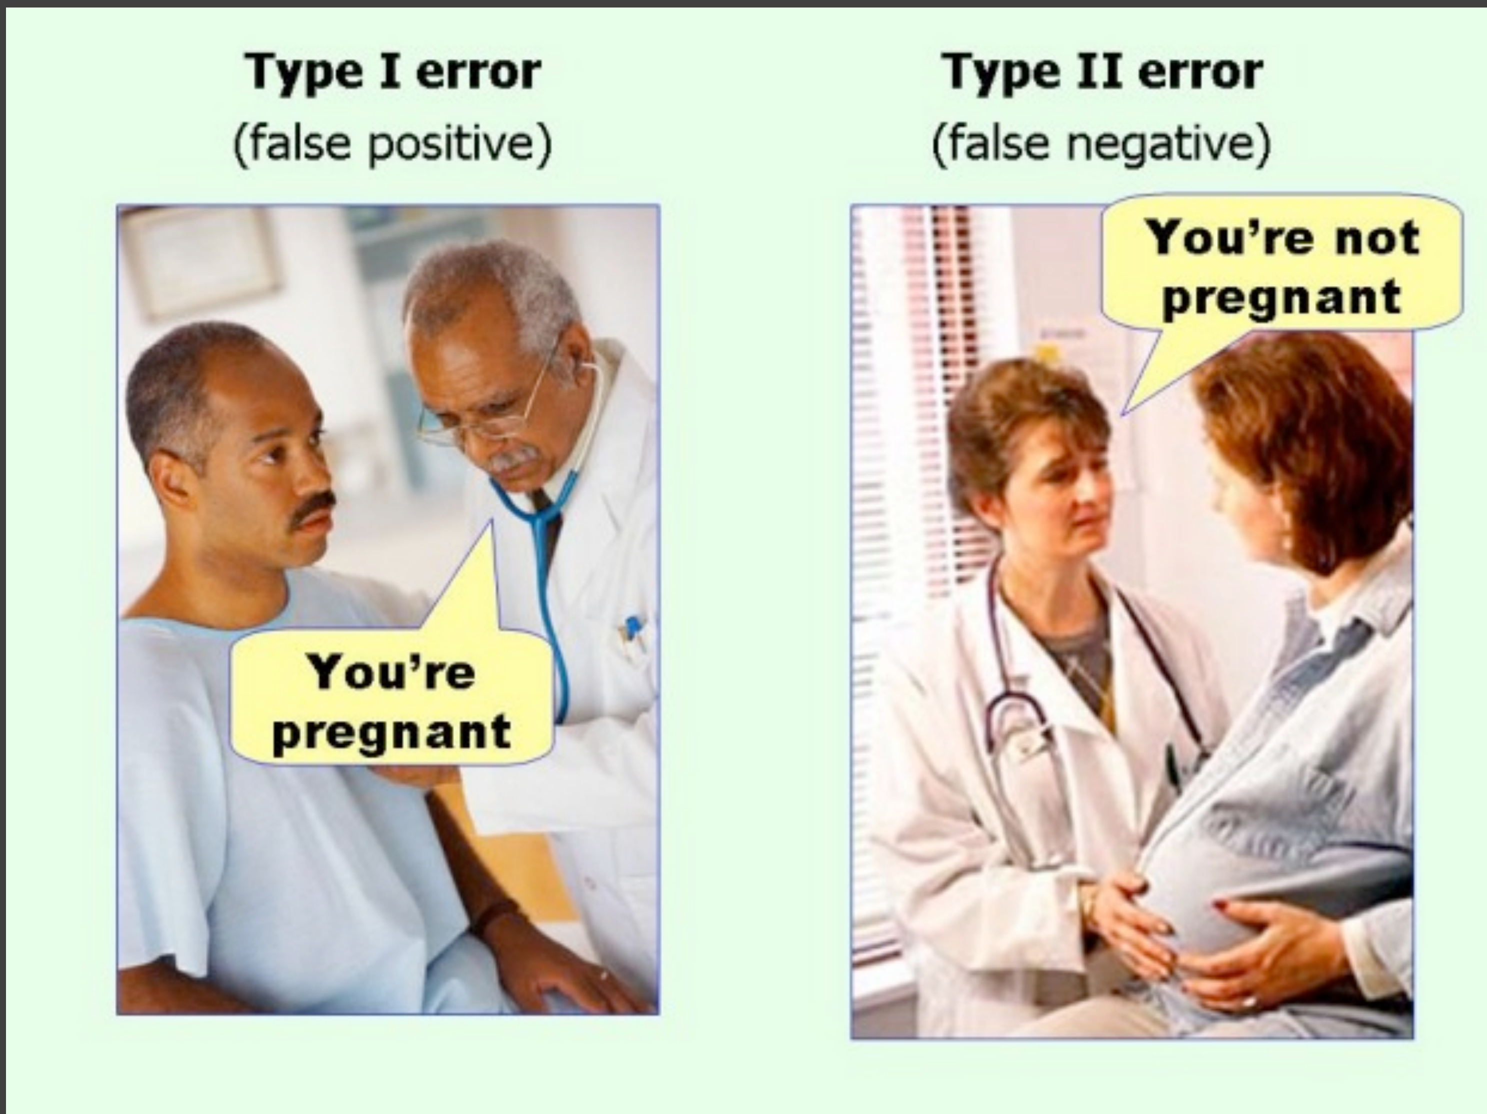

Model Uncertainty

“50% Chance of Rain”

Model Uncertainty

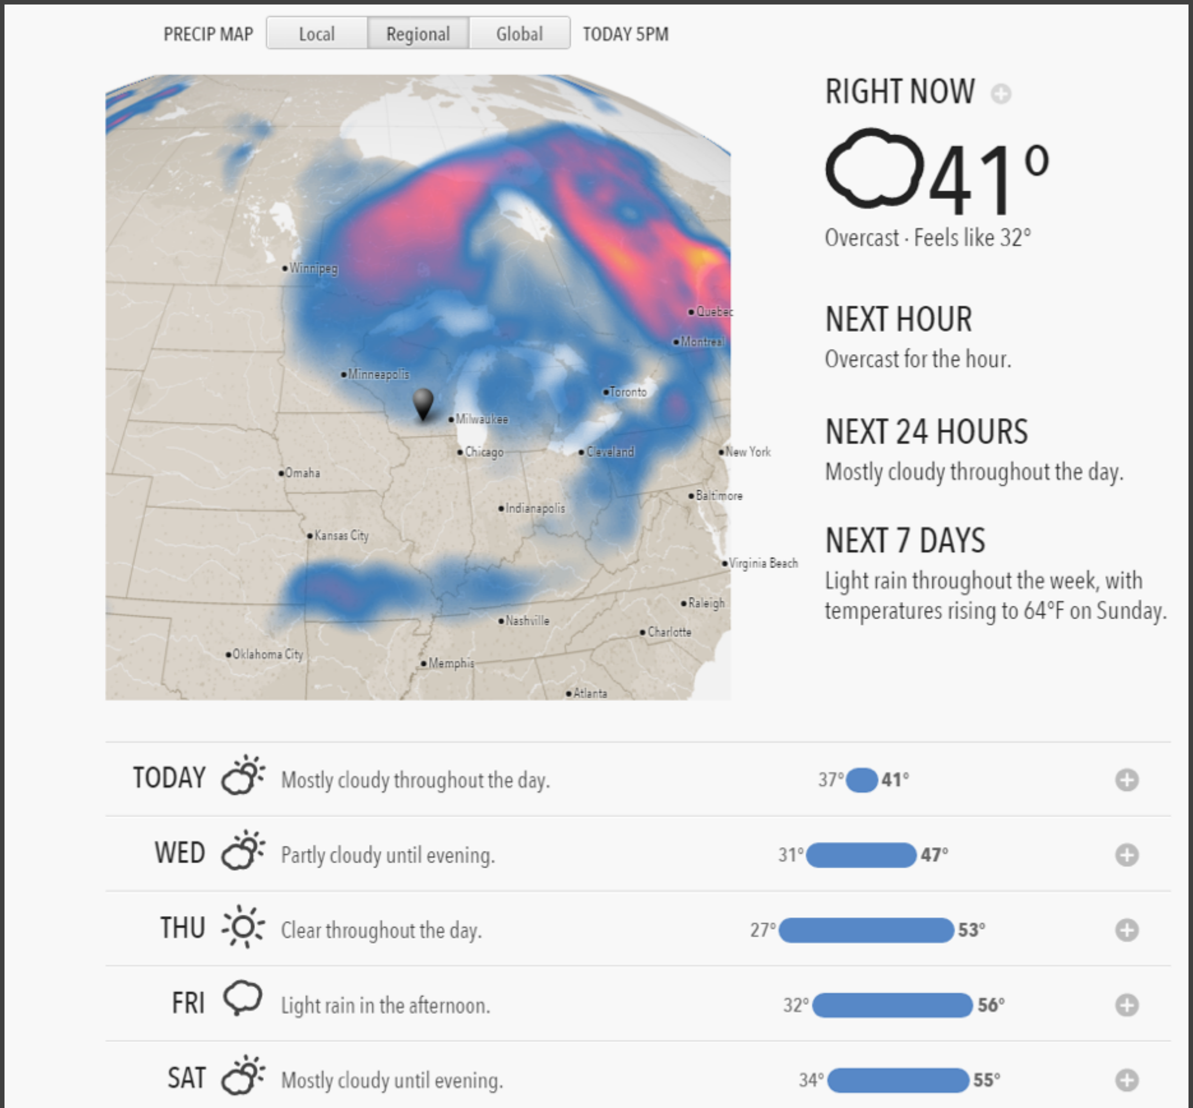

Measurement Uncertainty

Measurement Uncertainty

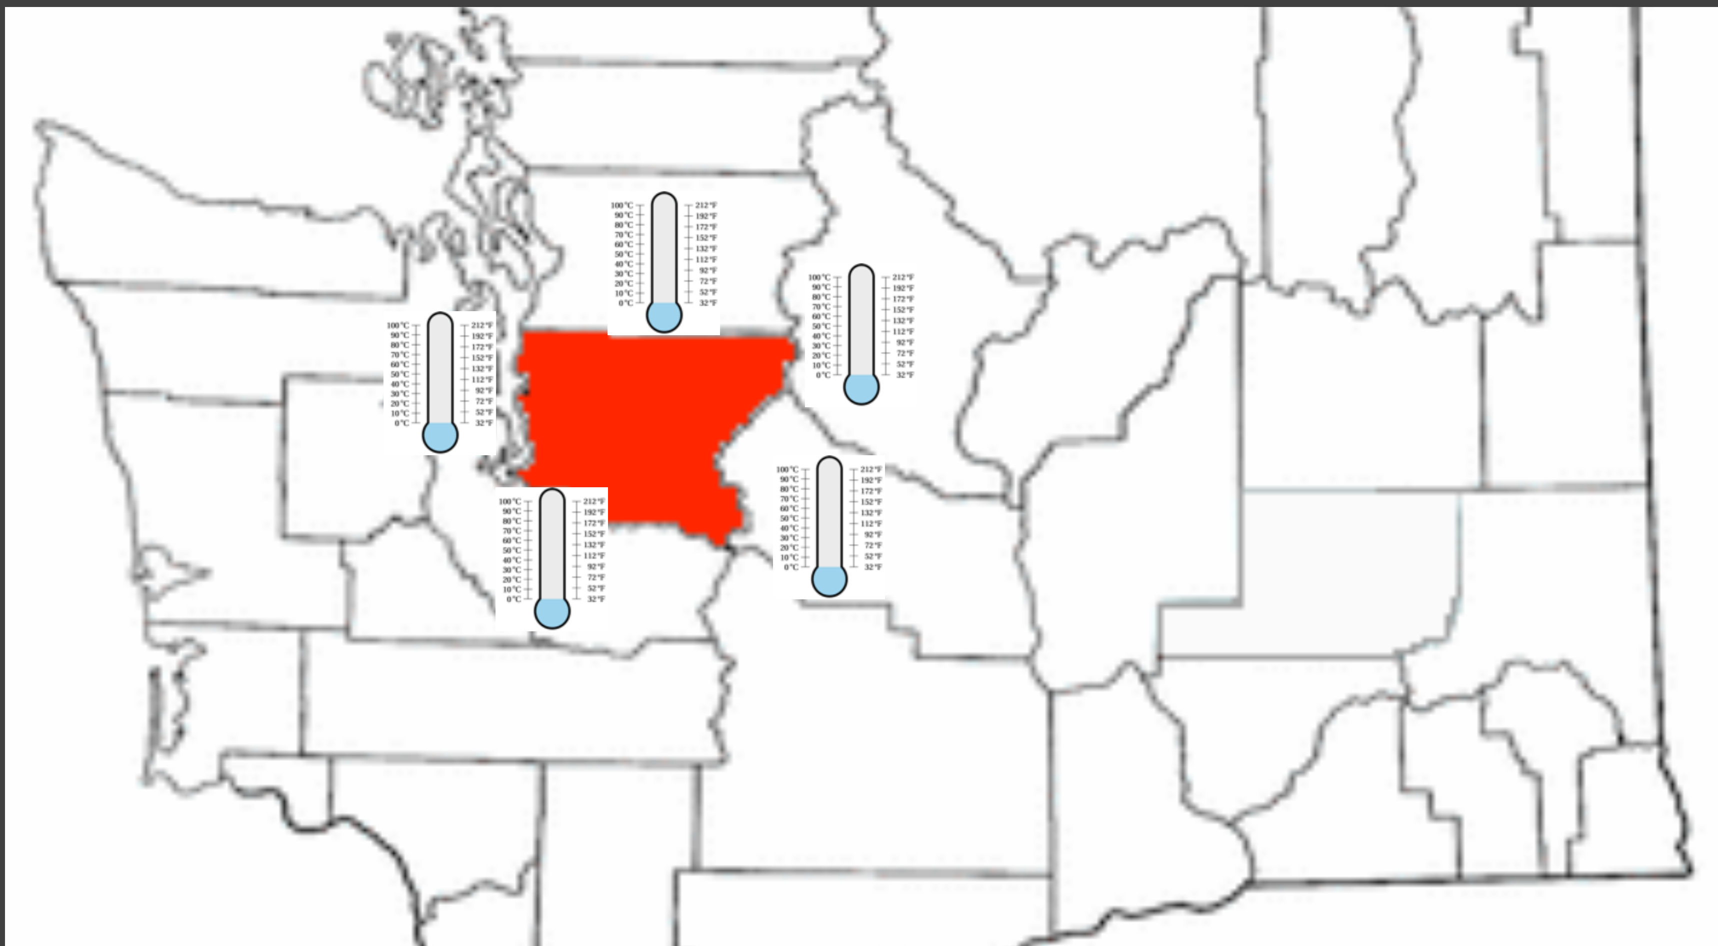

Measurement Uncertainty

Accuracy

Measurement Uncertainty

Accuracy

Measurement Uncertainty

Accuracy

Precision

Measurement Uncertainty

Accuracy

Precision

Measurement Uncertainty

Accuracy

Precision

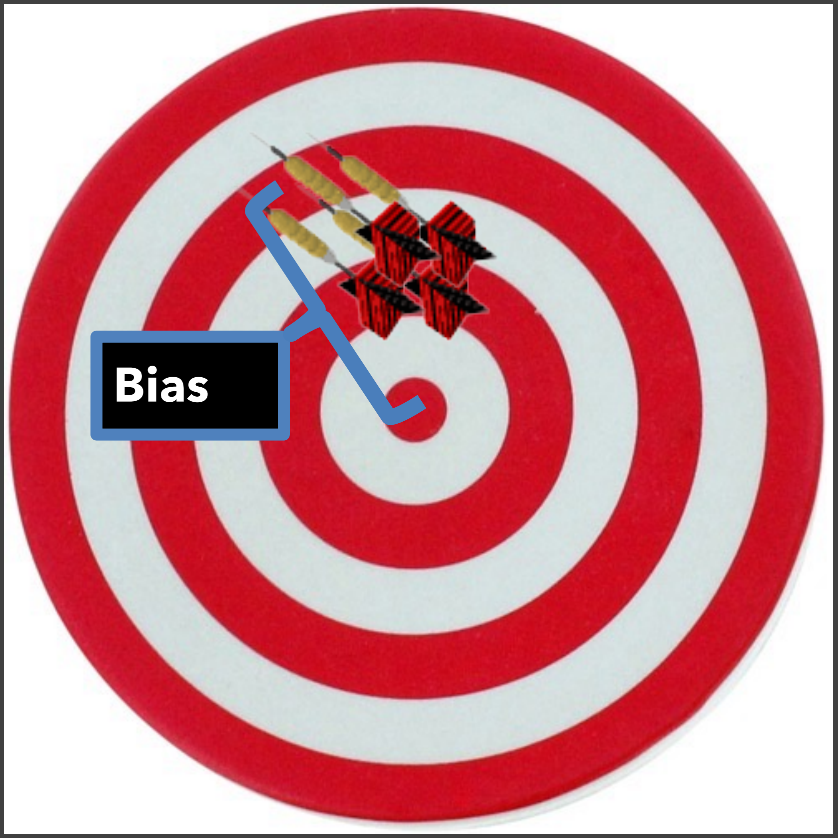

Should you take this $4 bet?

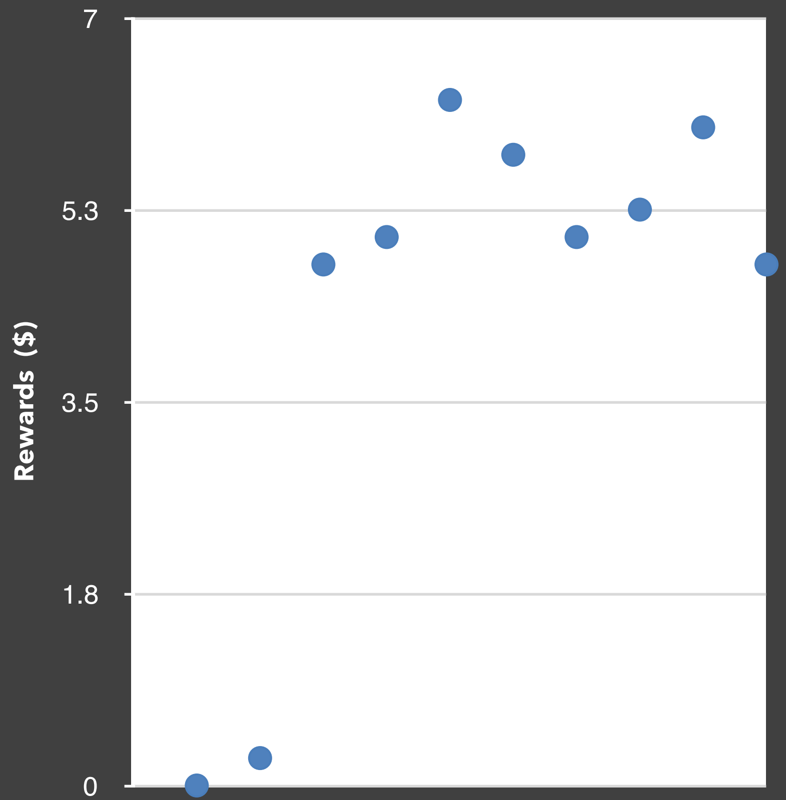

Should you take this $4 bet?

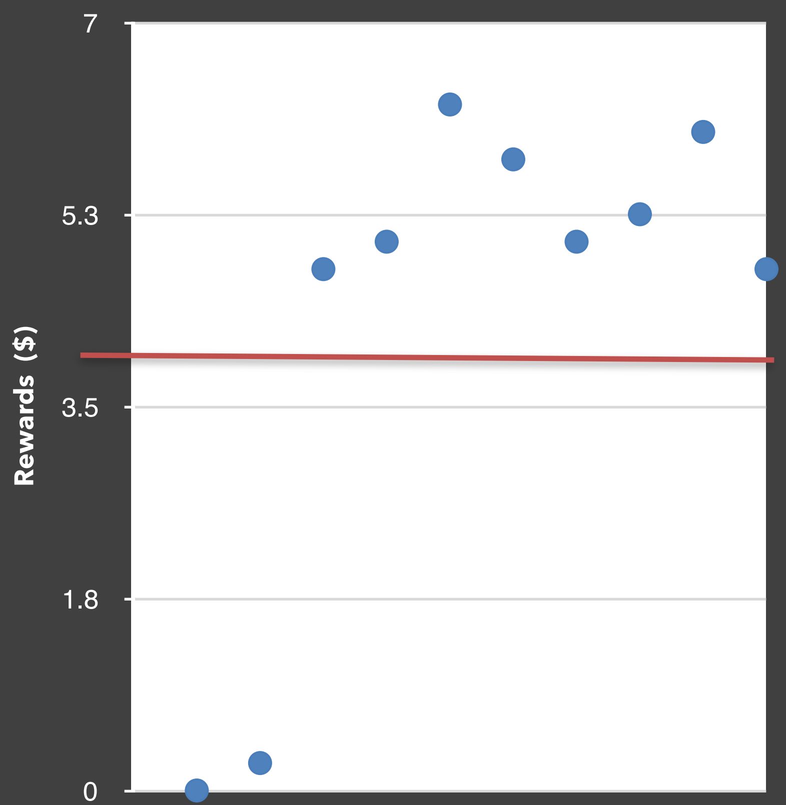

Should you take this $4 bet?

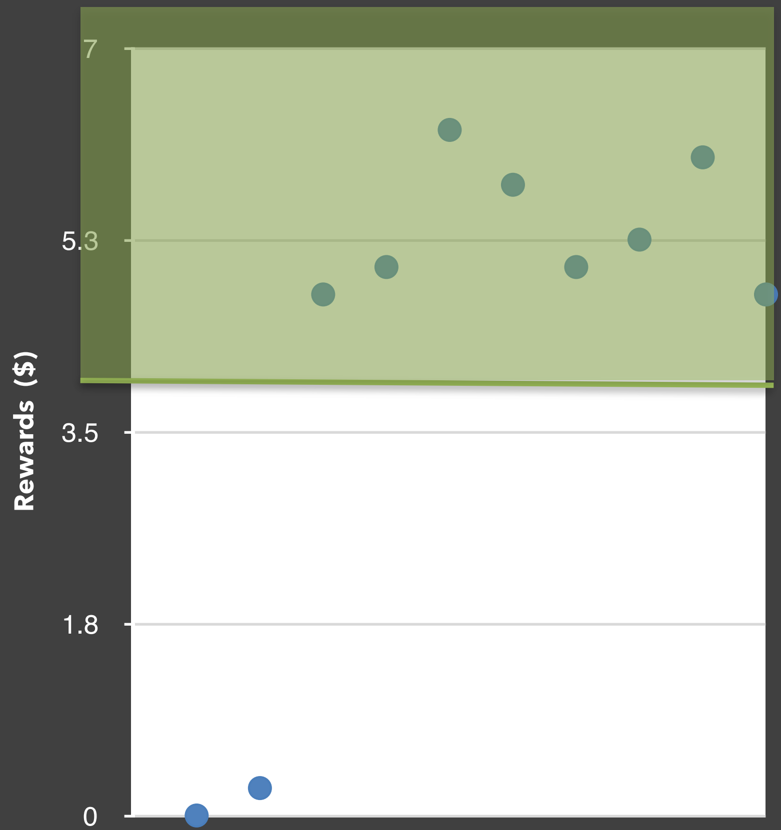

Should you take this $4 bet?

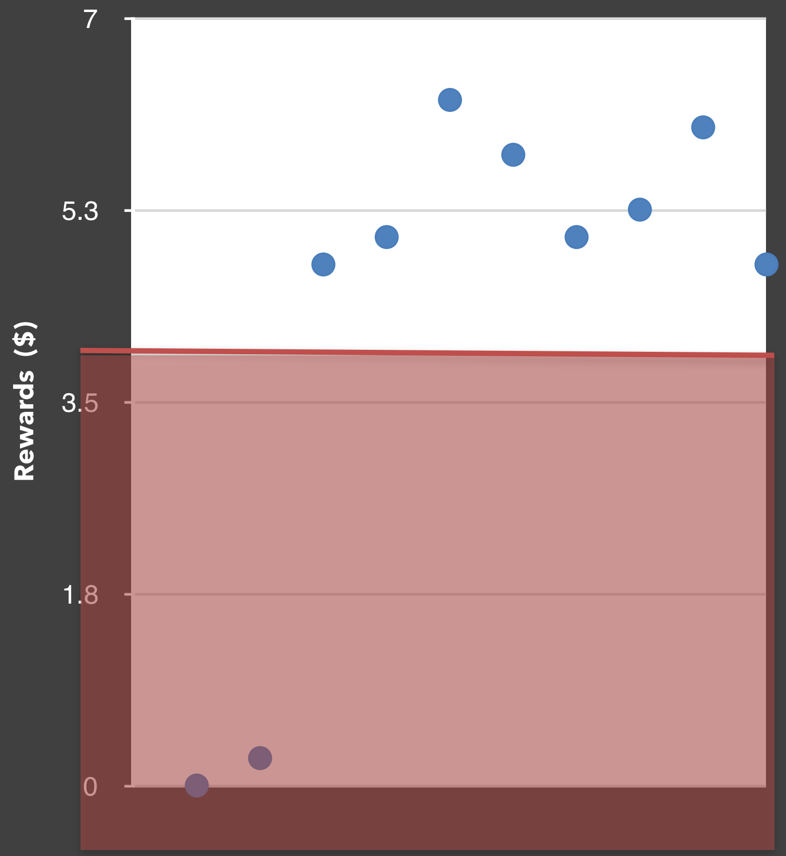

Expected Value

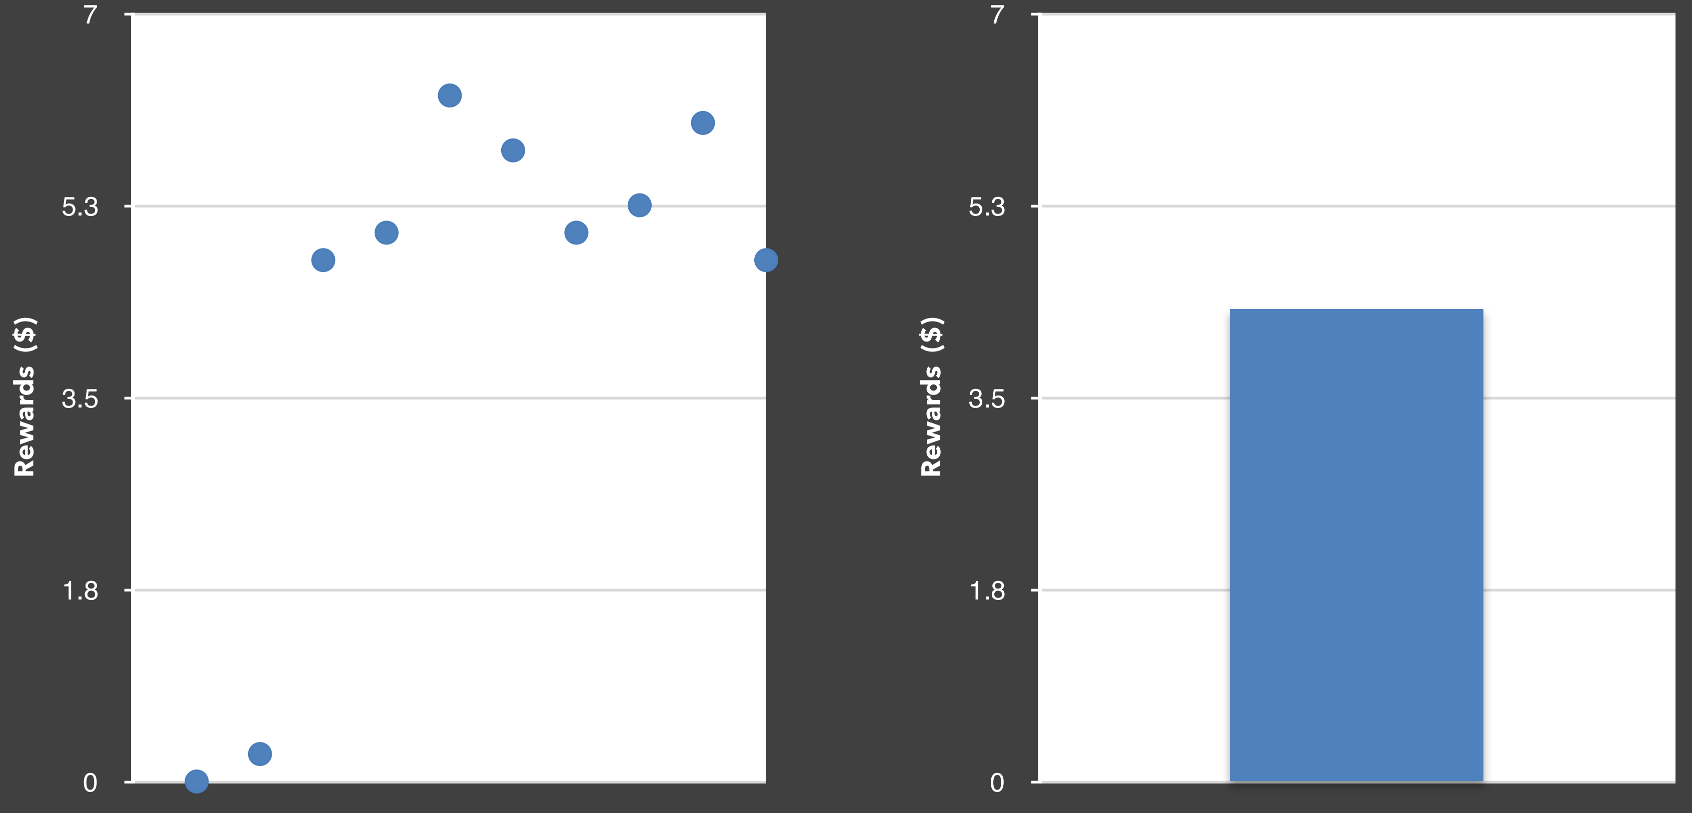

Mean And Error

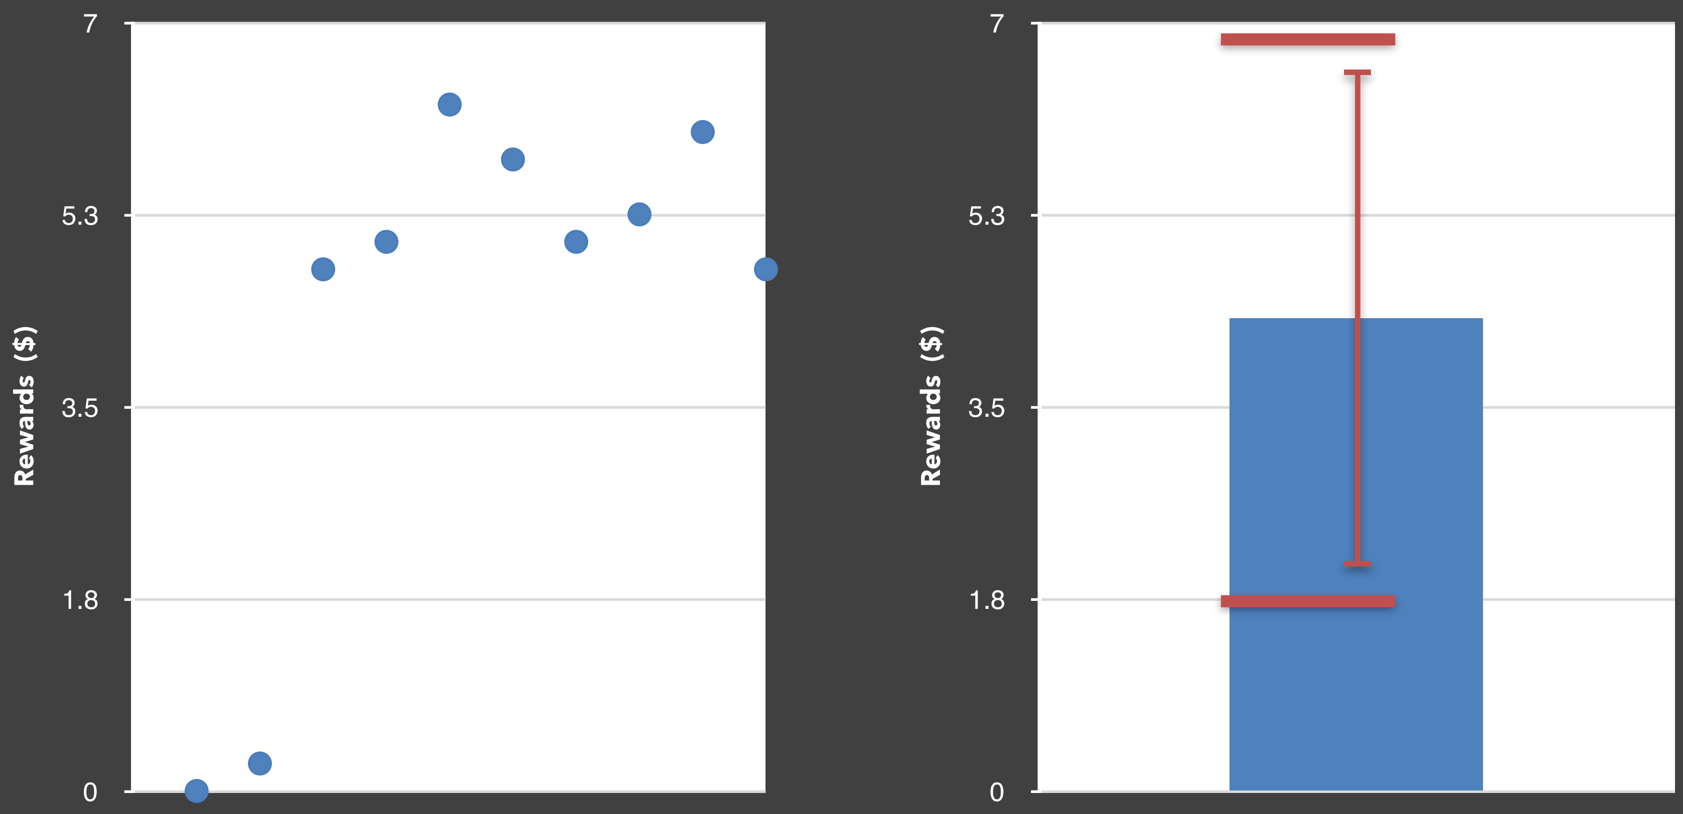

Statistical Inference

Assuming bet returns are normally distributed.

M = 4.14

SD = 2.33

n = 10

P(μ>4) = **0.95**

∎ Take the bet

Statistical Inference

Assuming bet returns are normally distributed.

MODEL

$M = 4.14$

$SD = 2.33$

$n = 10$

$P(\mu>4) = \mathbf{0.95}$

∎ Take the bet

Statistical Inference

Assuming bet returns are normally distributed. MODEL

$M = 4.14$

$SD = 2.33$ MEASUREMENT

$n = 10$

$P(\mu > 4) = \mathbf{0.95}$

∎ Take the bet

Statistical Inference

Assuming bet returns are normally distributed.

MODEL

$M = 4.14$

$SD = 2.33$

$n = 10$

MEASUREMENT

$P(\mu>4) =$ **0.95**

∎ Take the bet

DECISION

Uncertainty Sources

Measurement Uncertainty: “We’re not sure what the data are”

Model Uncertainty: “We’re not sure how the data fit together”

Decision Uncertainty: “We’re not sure what to do now that we have the data”


1
0.75
0.5
0.25
0

Measurement Uncertainty

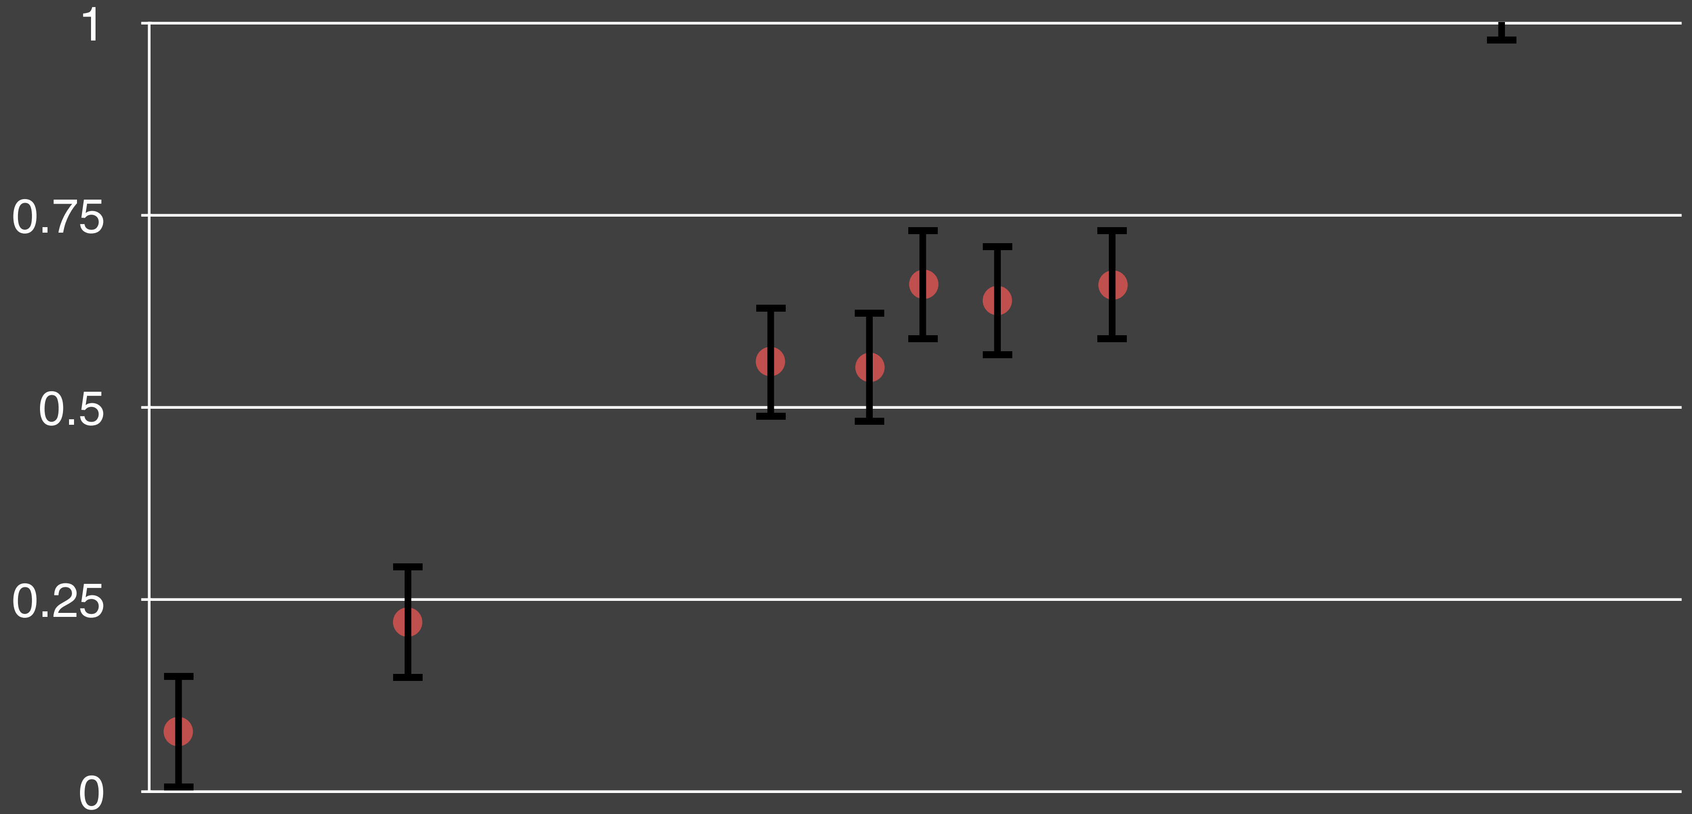

Model Uncertainty

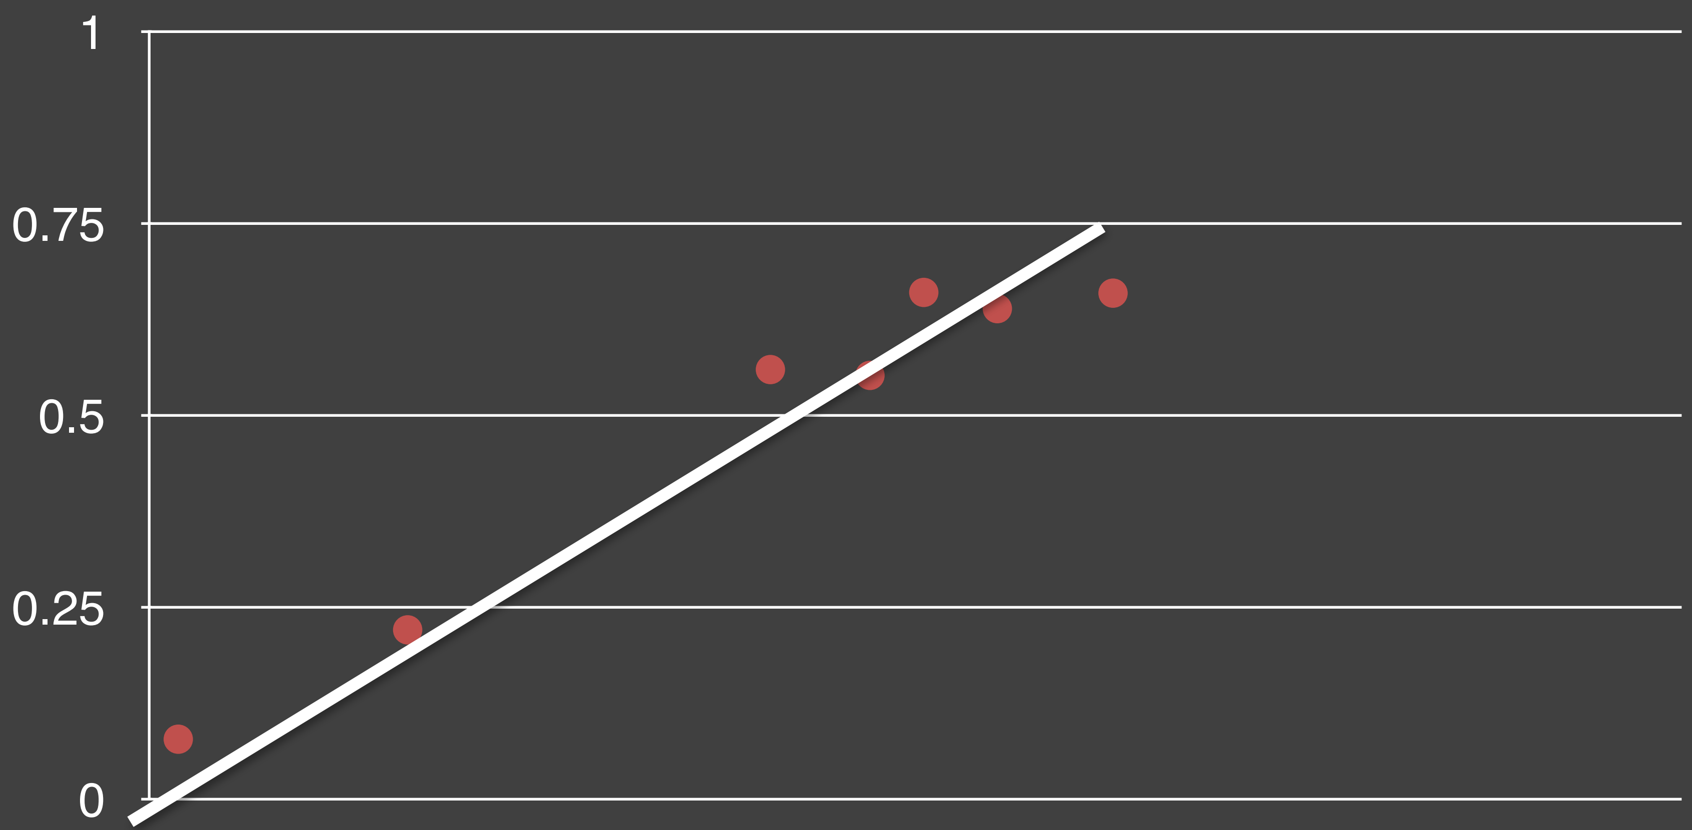

Model Uncertainty

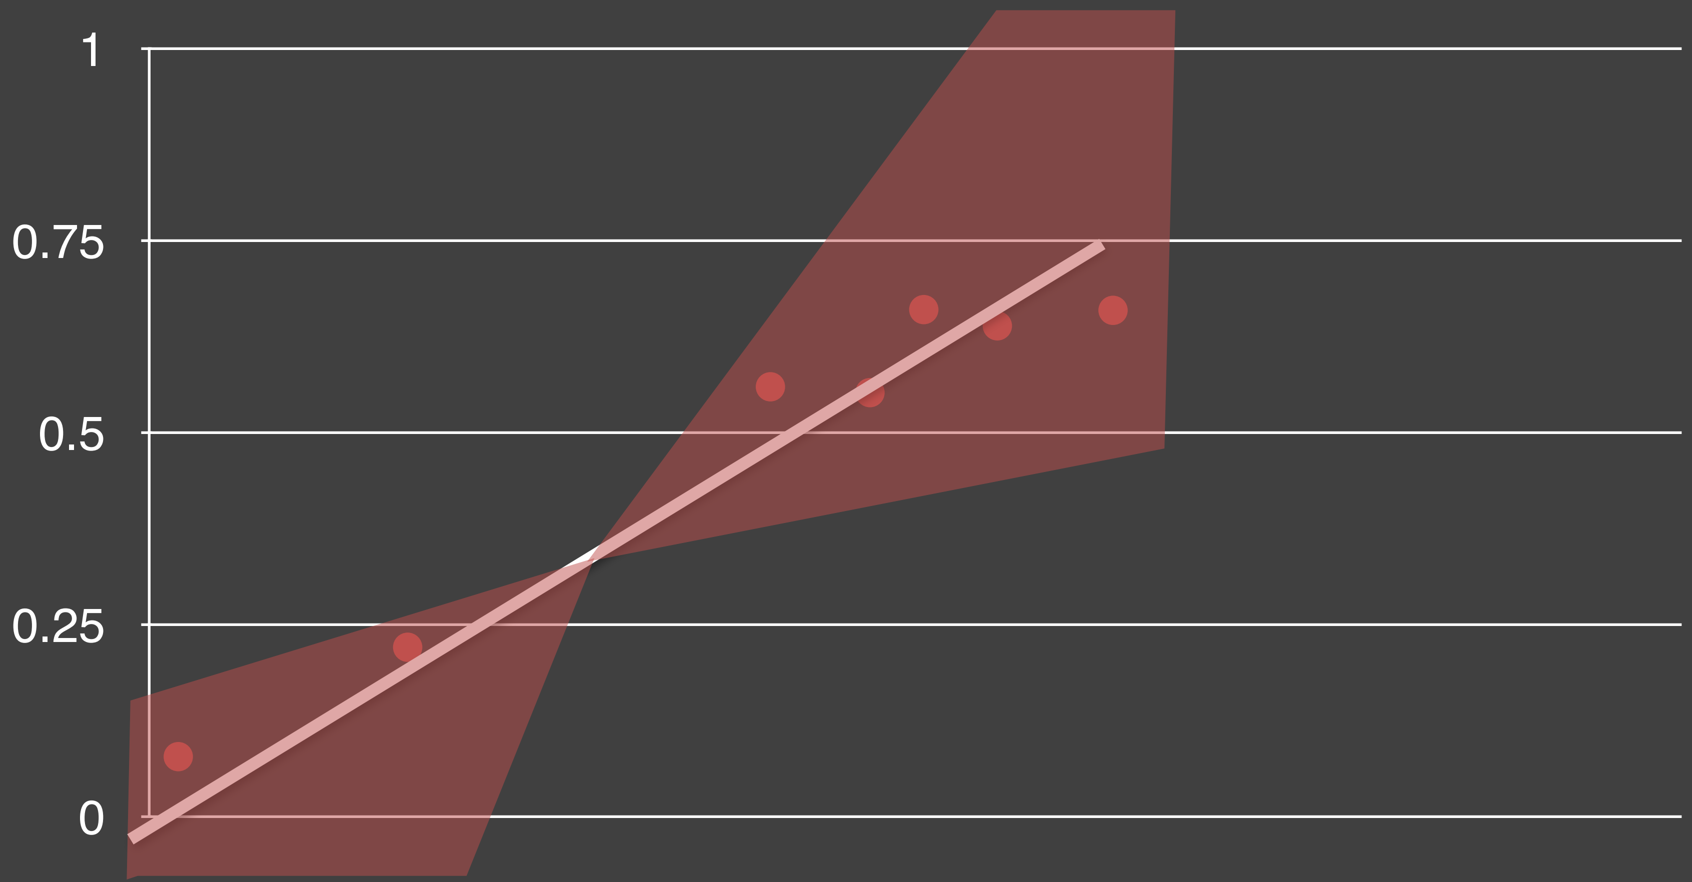

Decision Uncertainty

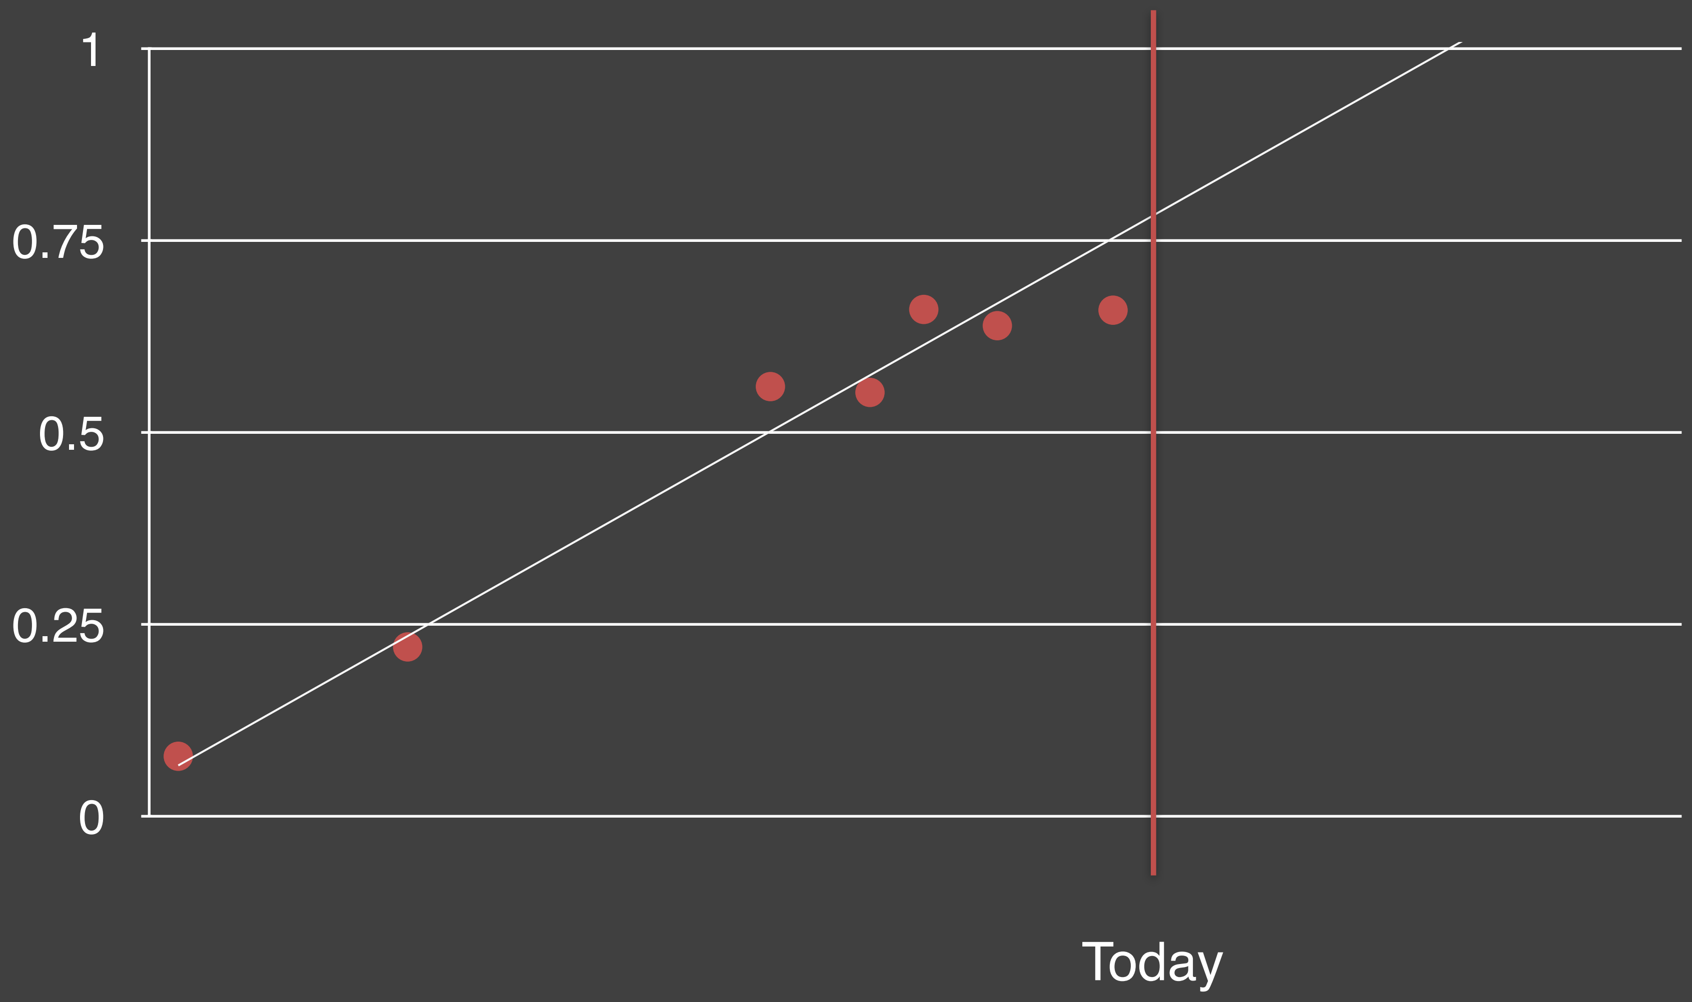

Decision Uncertainty

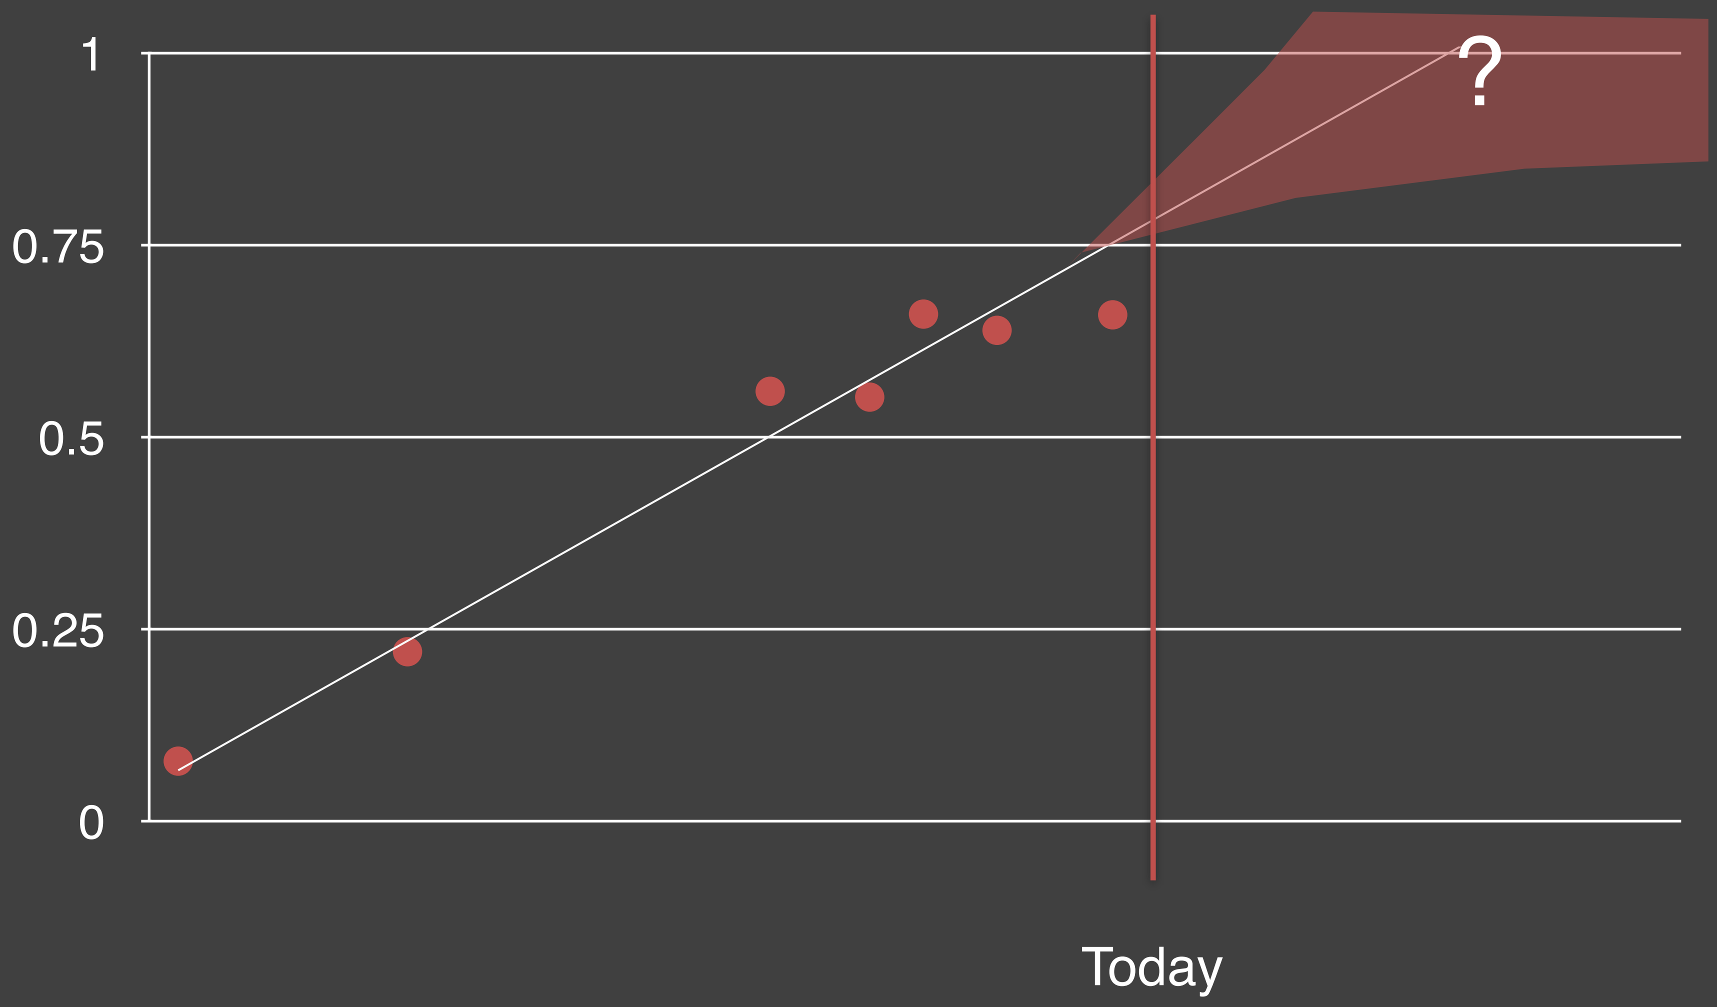

Uncertainty Vis Pipeline

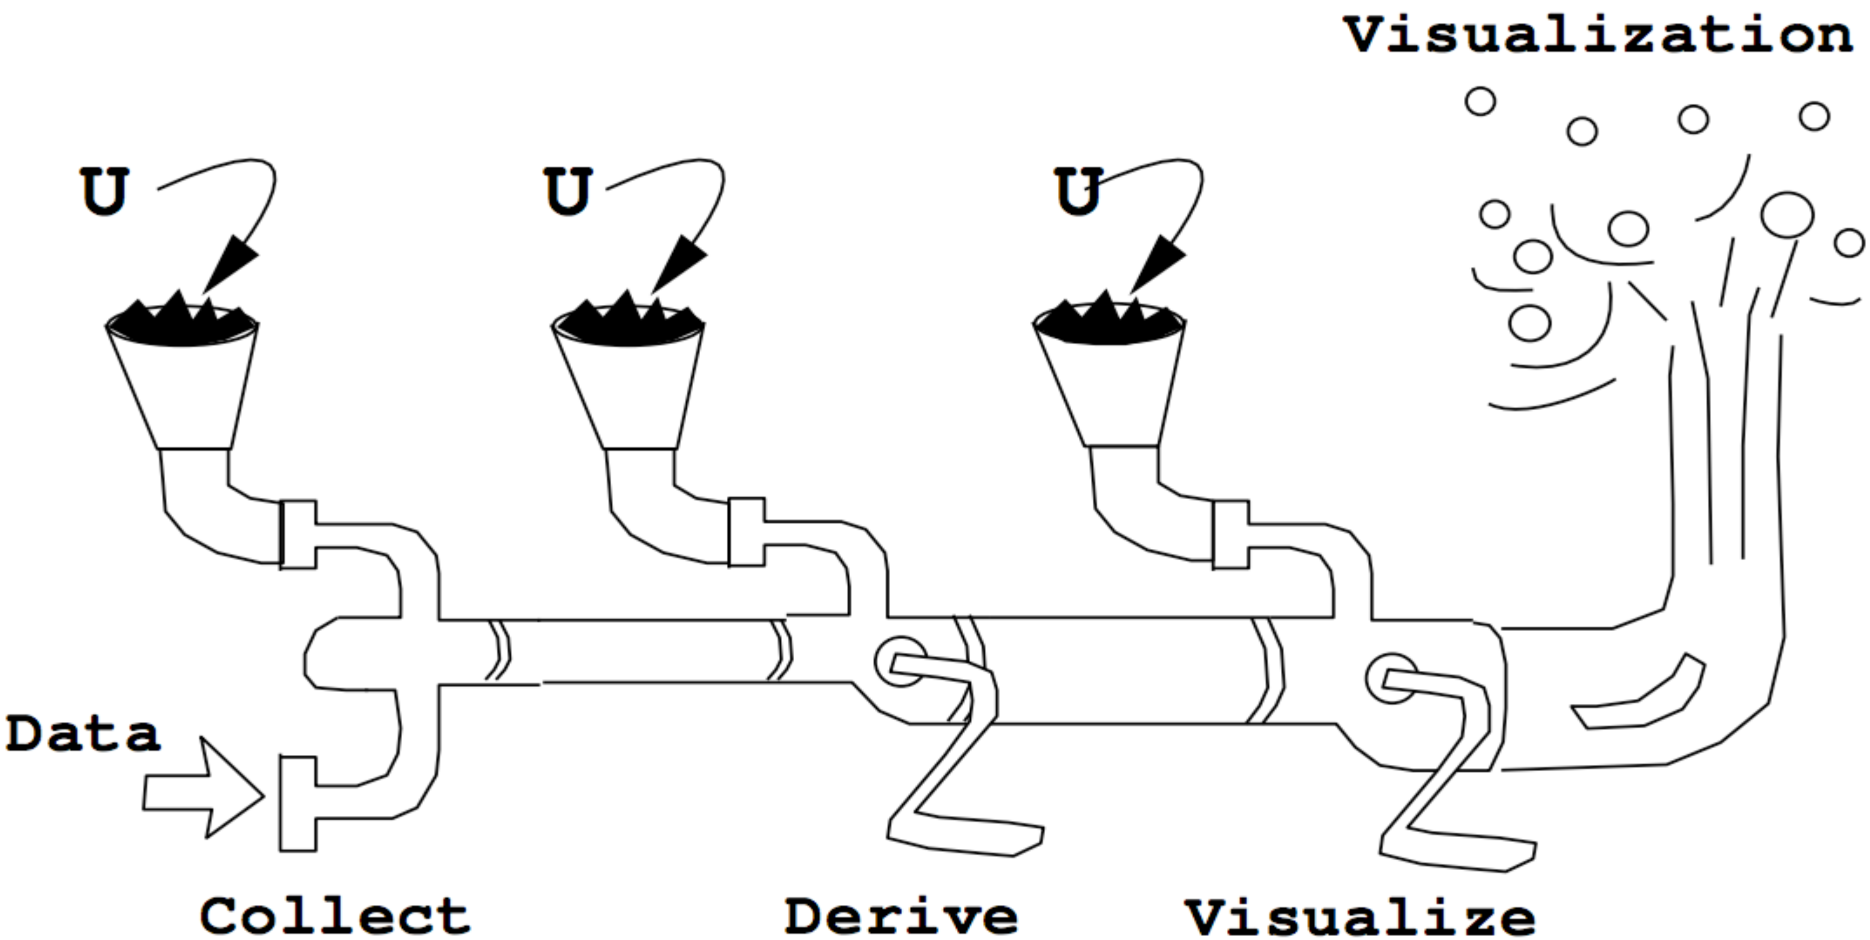

Pang et al. Approaches to Uncertainty Visualization. The Visual Computer, 1997.

Uncertainty Vis Pipeline

1) Quantify Uncertainty
2) Choose a free visual variable
3) Encode uncertainty with the variable

Data Map

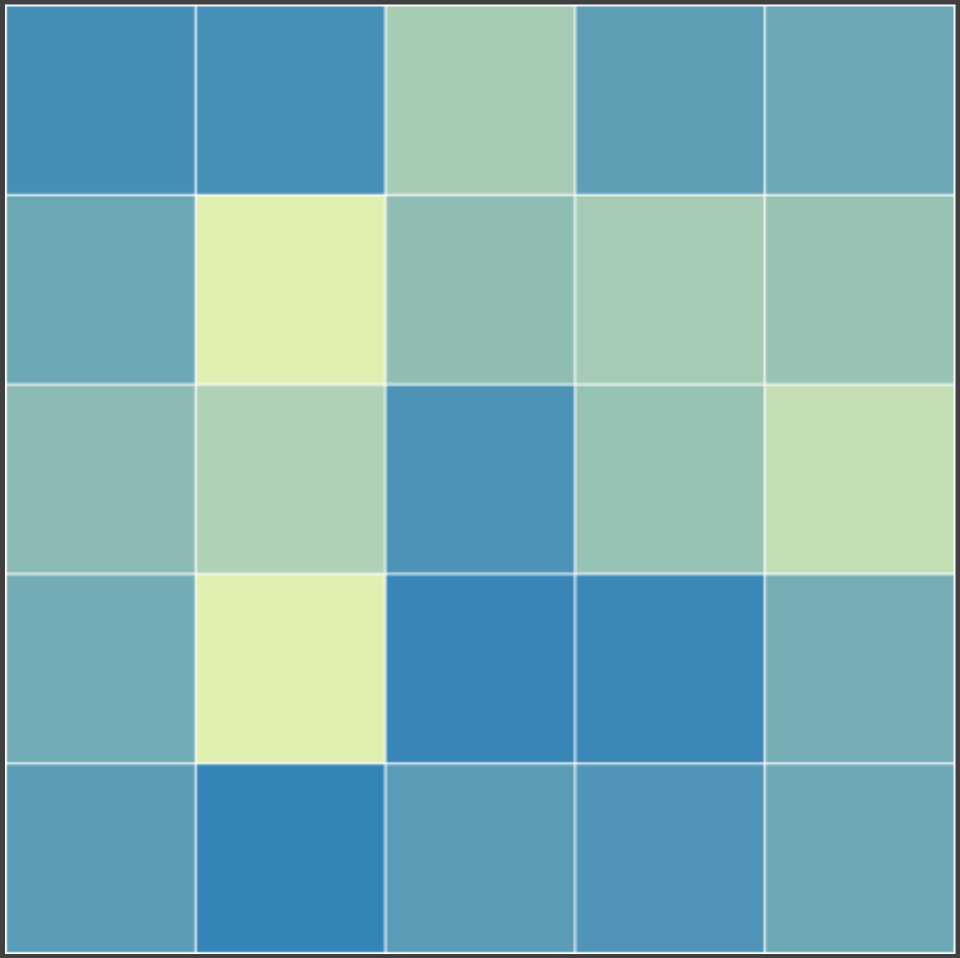

Uncertainty Map

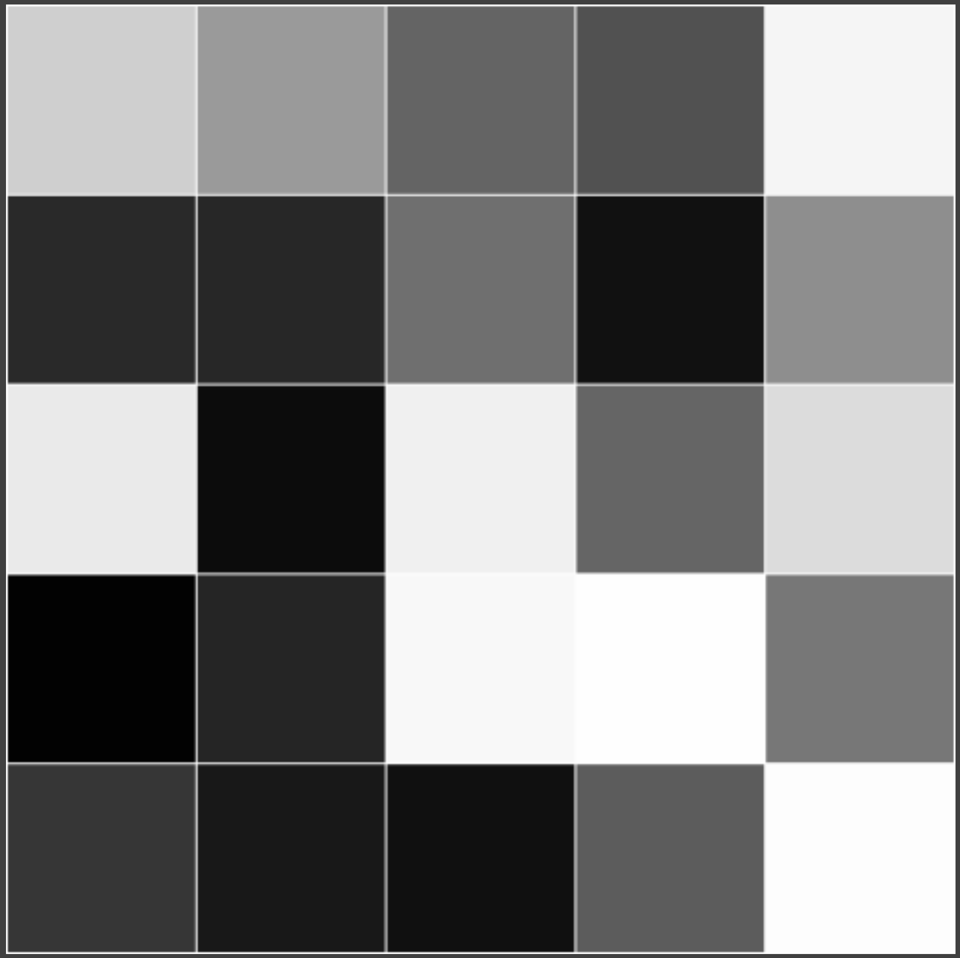

Juxtaposition

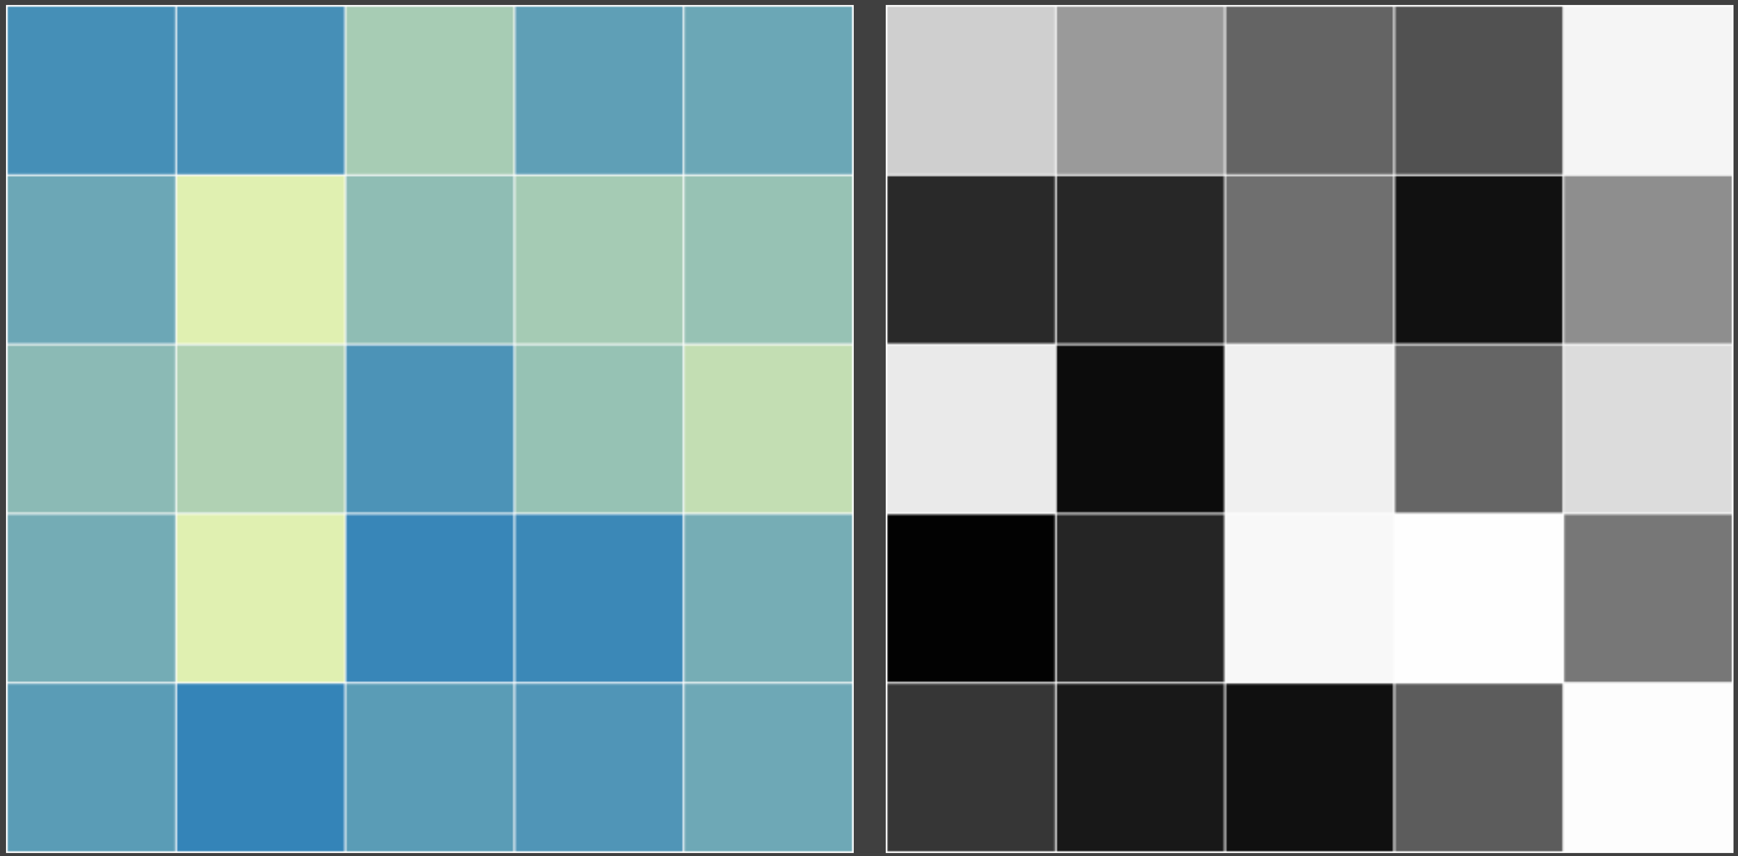

Juxtaposition

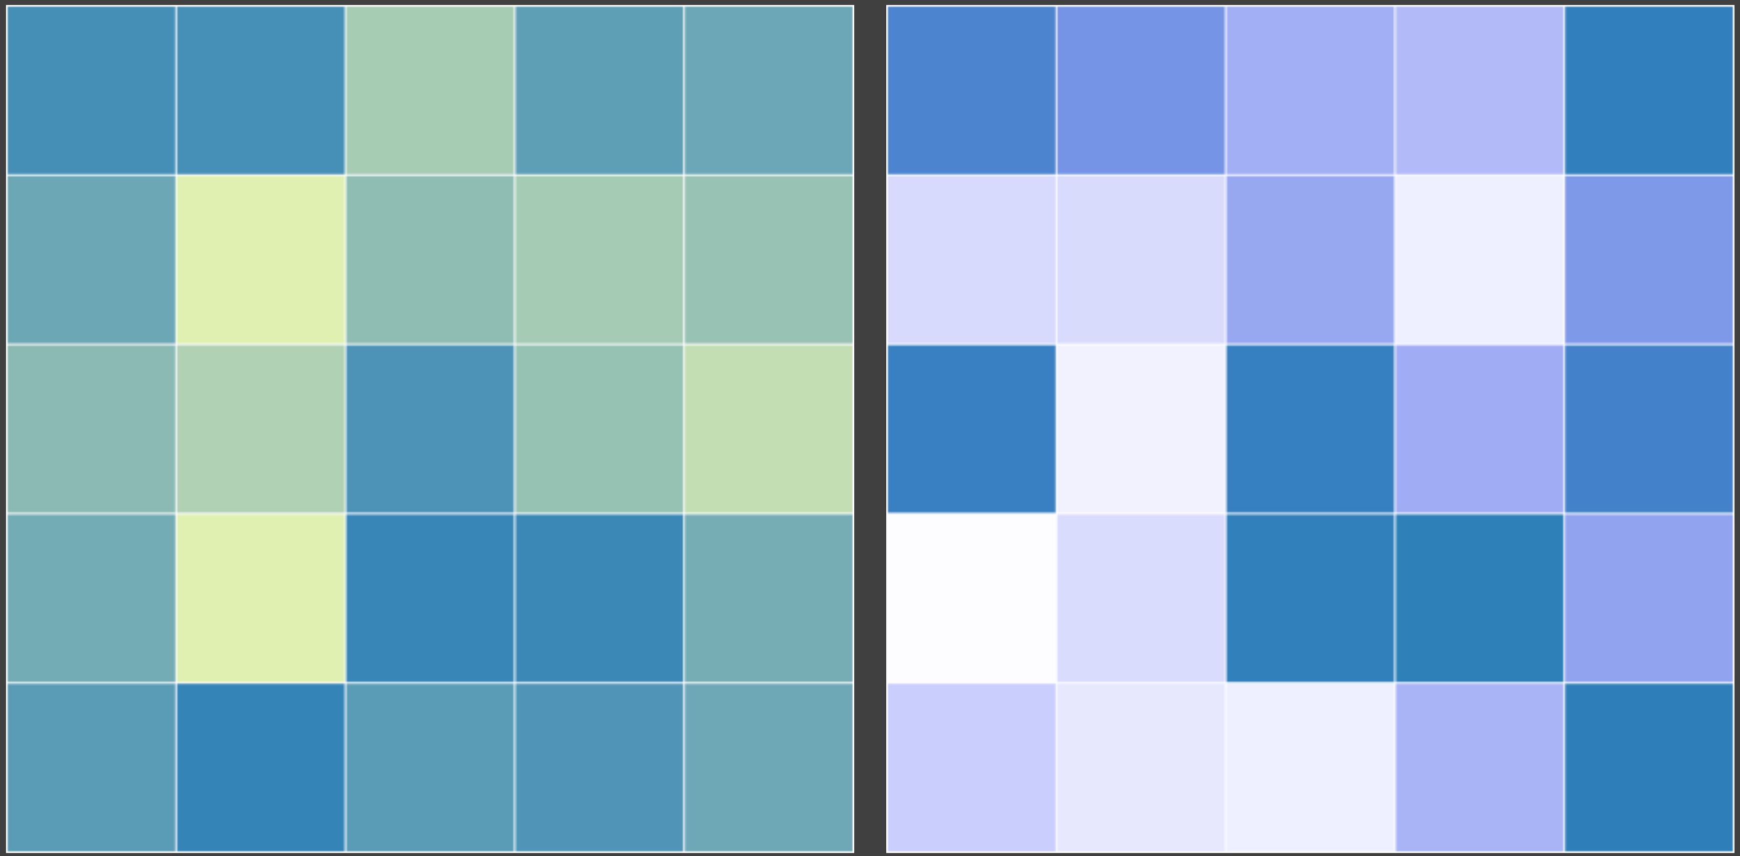

Juxtaposition

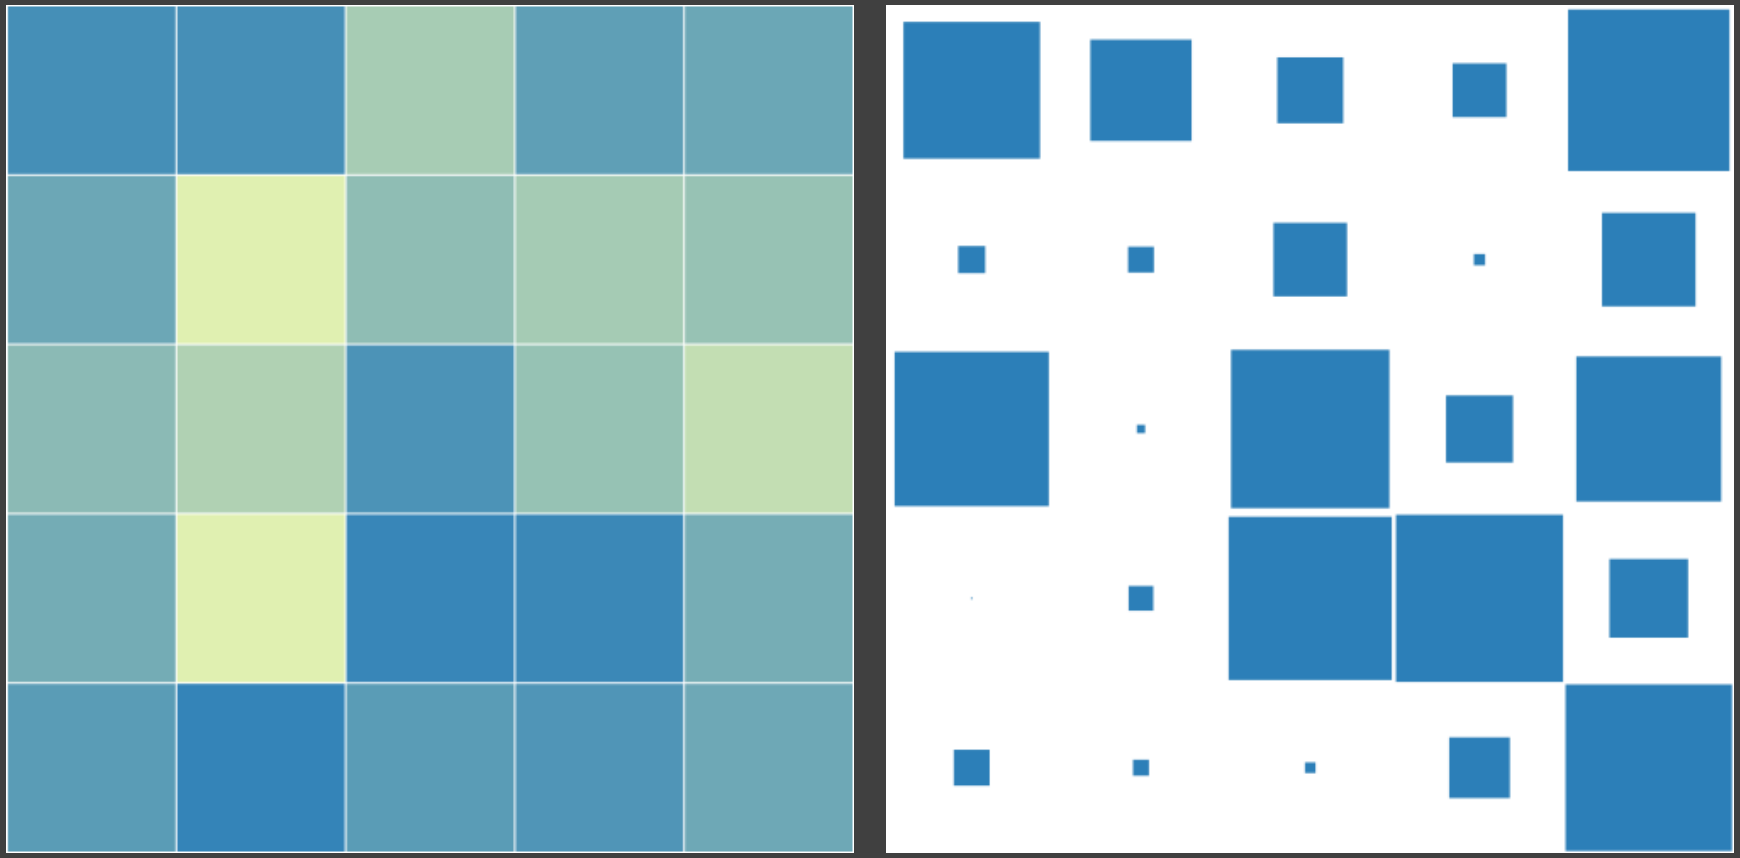

Juxtaposition

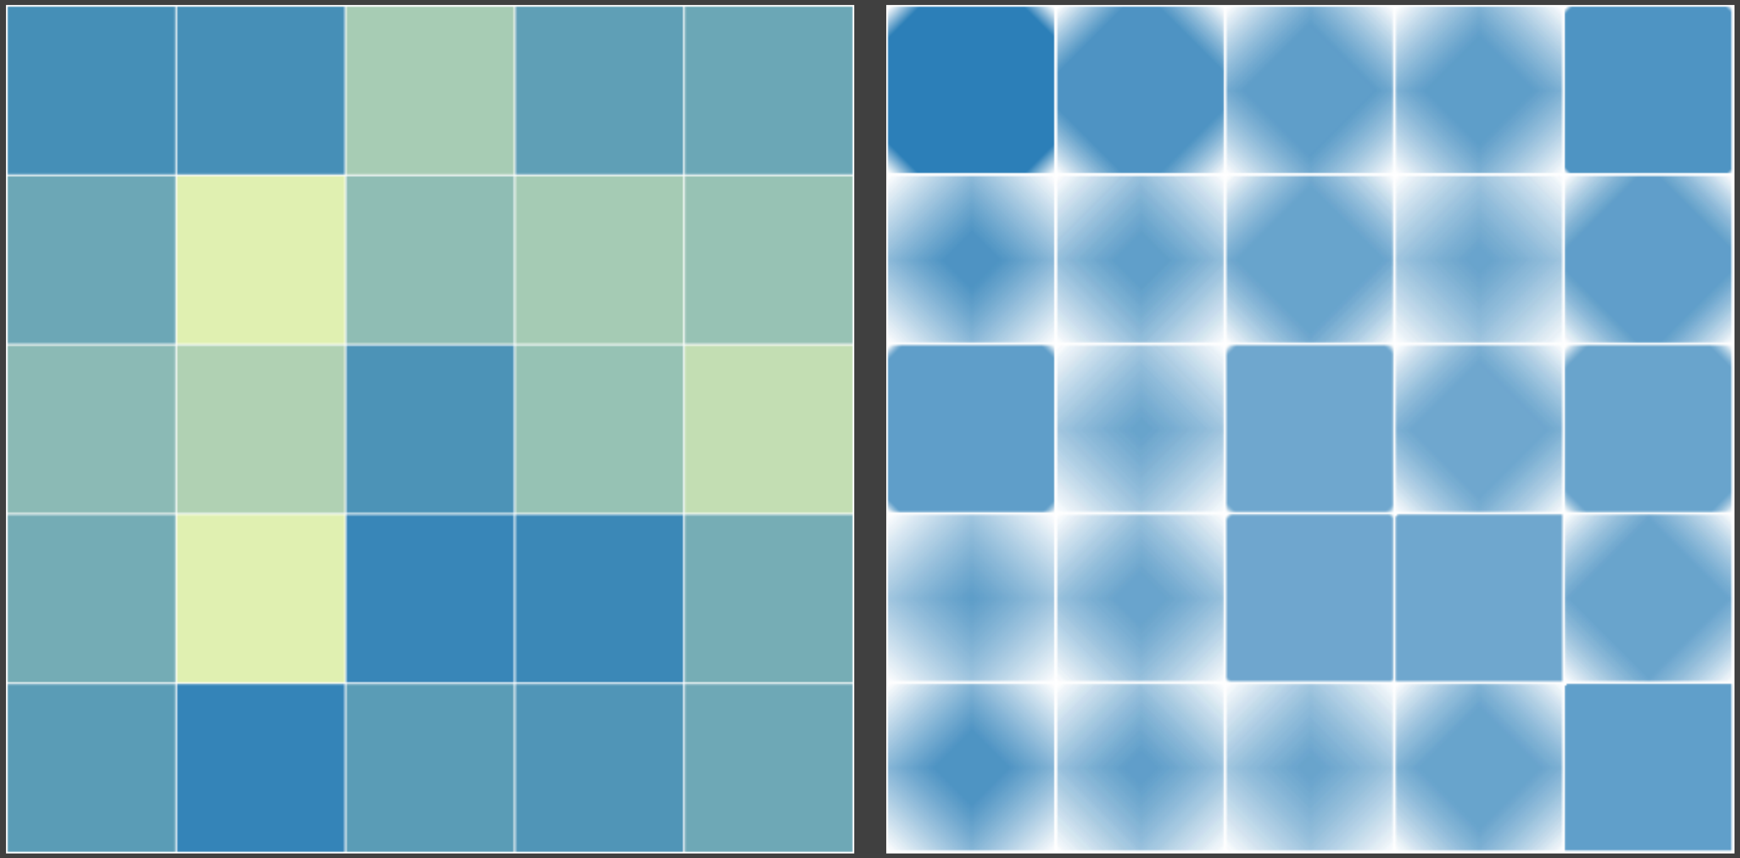

Superposition

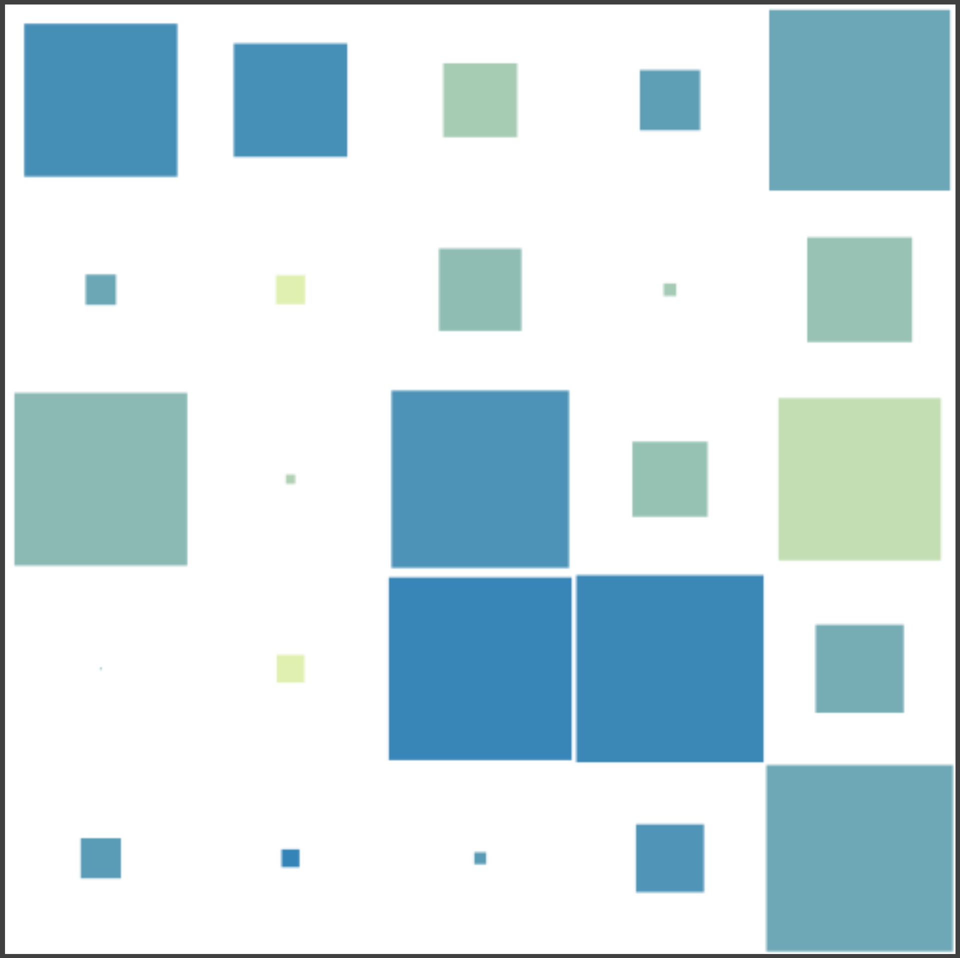

Superposition

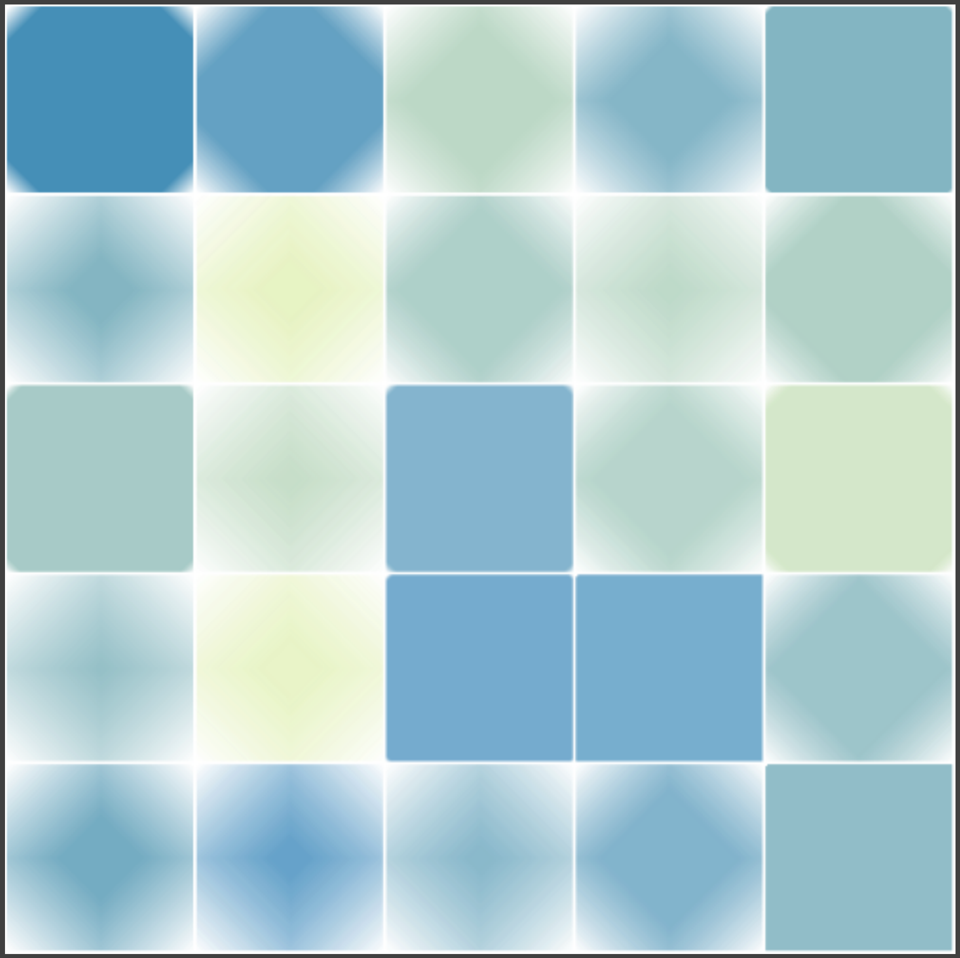

Superposition

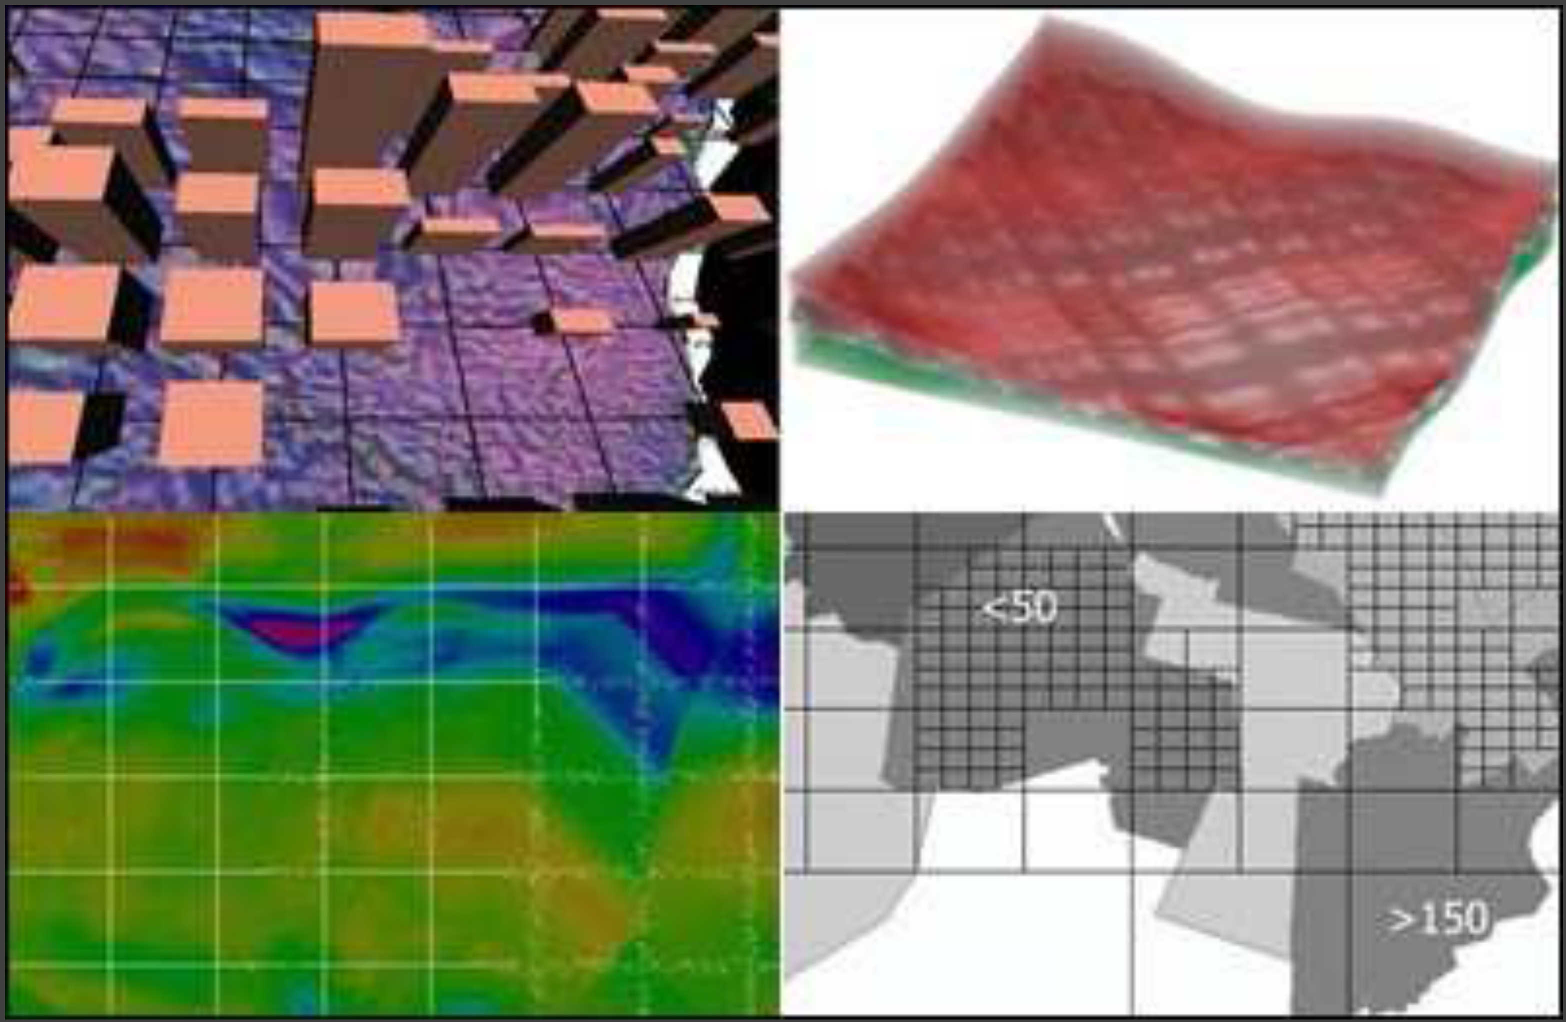

Griethe, Henning and Schumann, Heidrun. The Visualization of Uncertain Data: Methods and Problems. SimVis, 2006.

Uncertainty Vis Pipeline?

1) Quantify Uncertainty
2) Choose a free visual variable
3) Encode uncertainty with the variable

Design Decisions:

How to unify data and uncertainty map(s)?

Semiotics of Uncertainty

Visual Variables for Uncertainty

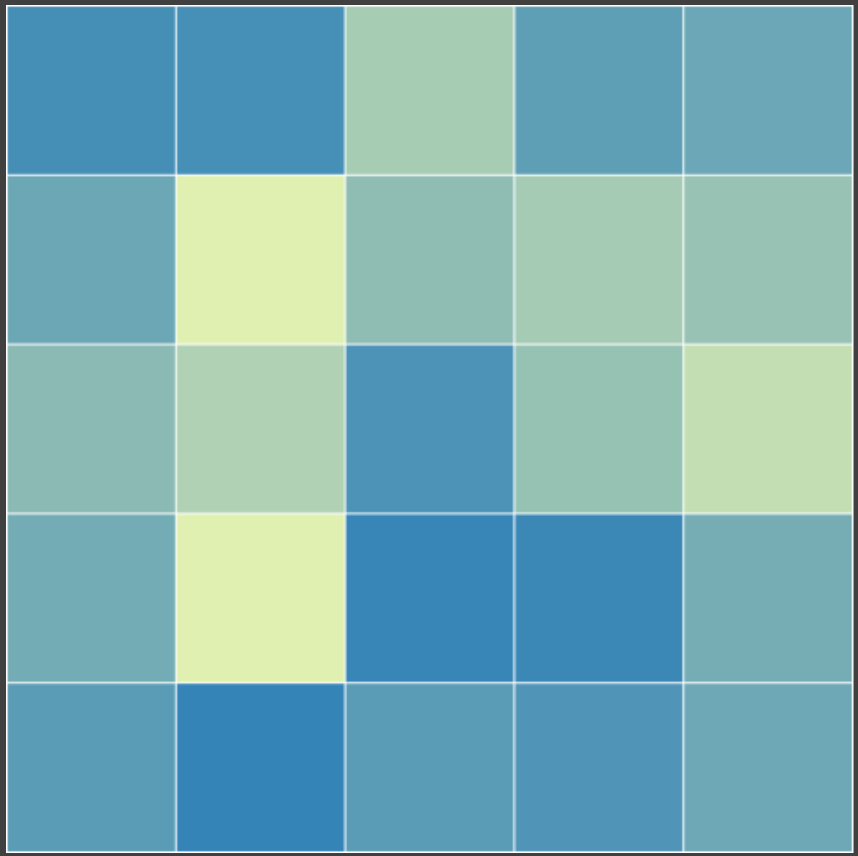

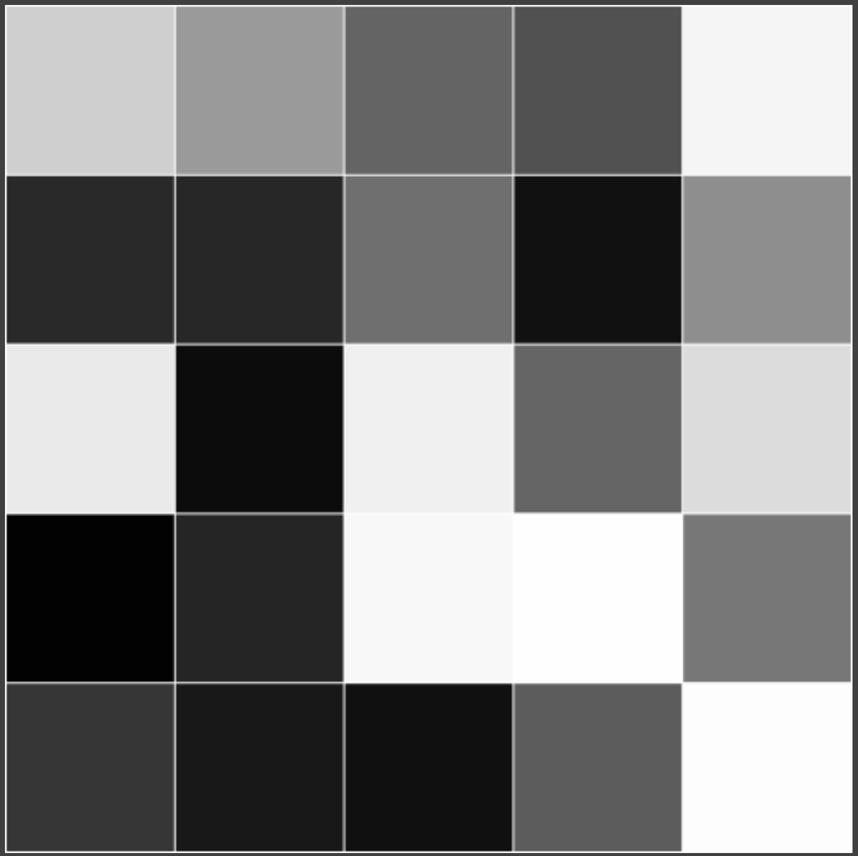

Value

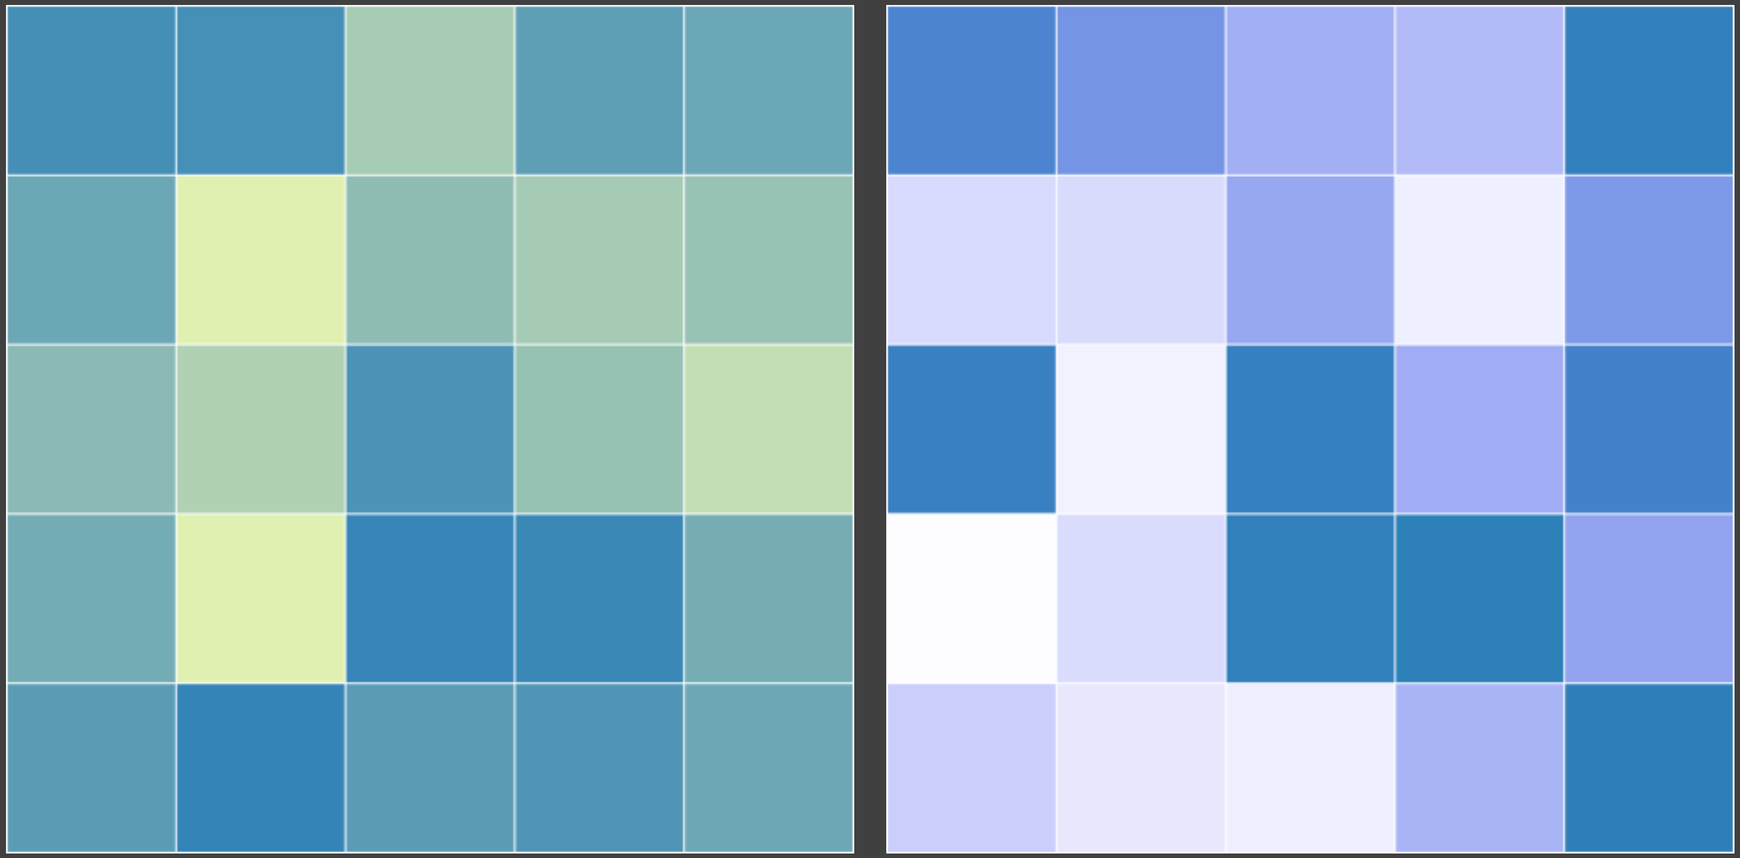

Size

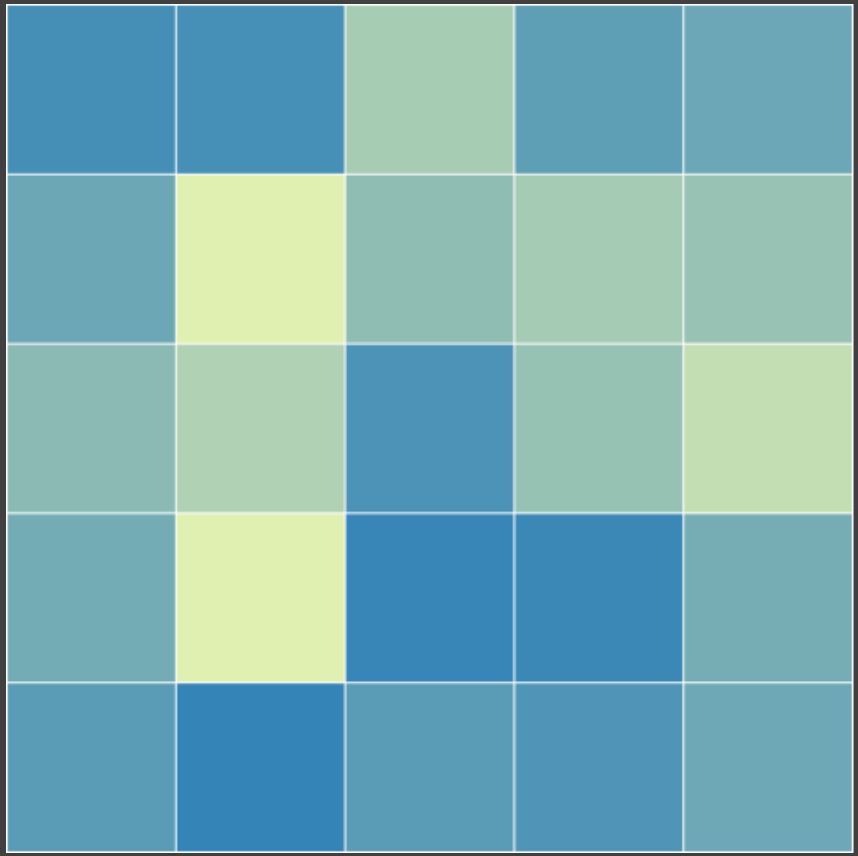

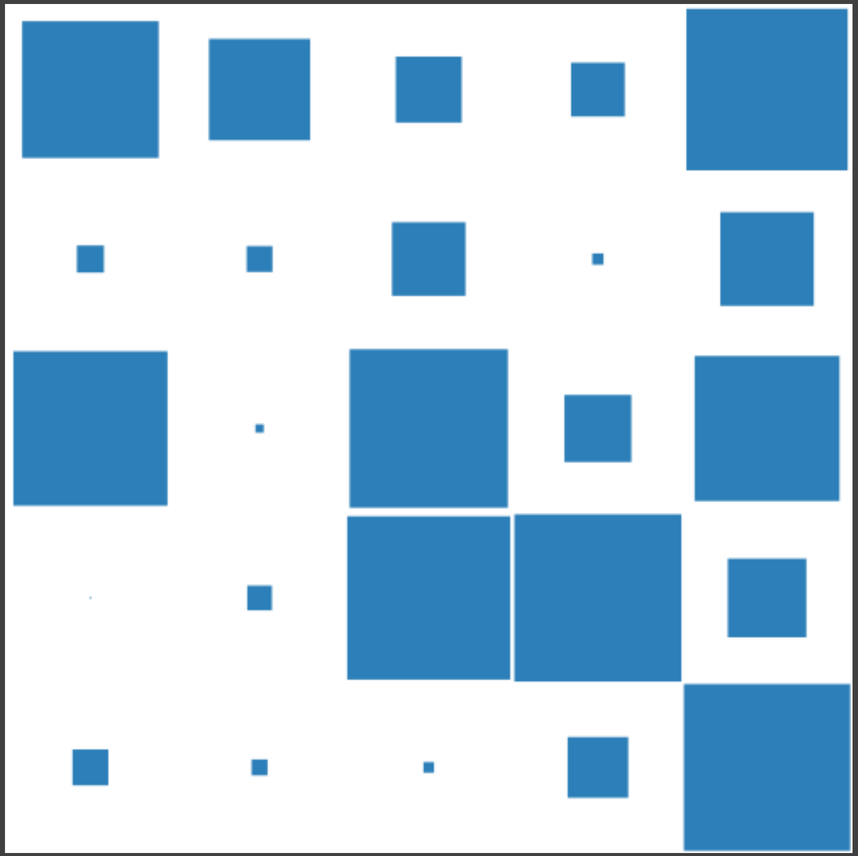

Fuzziness

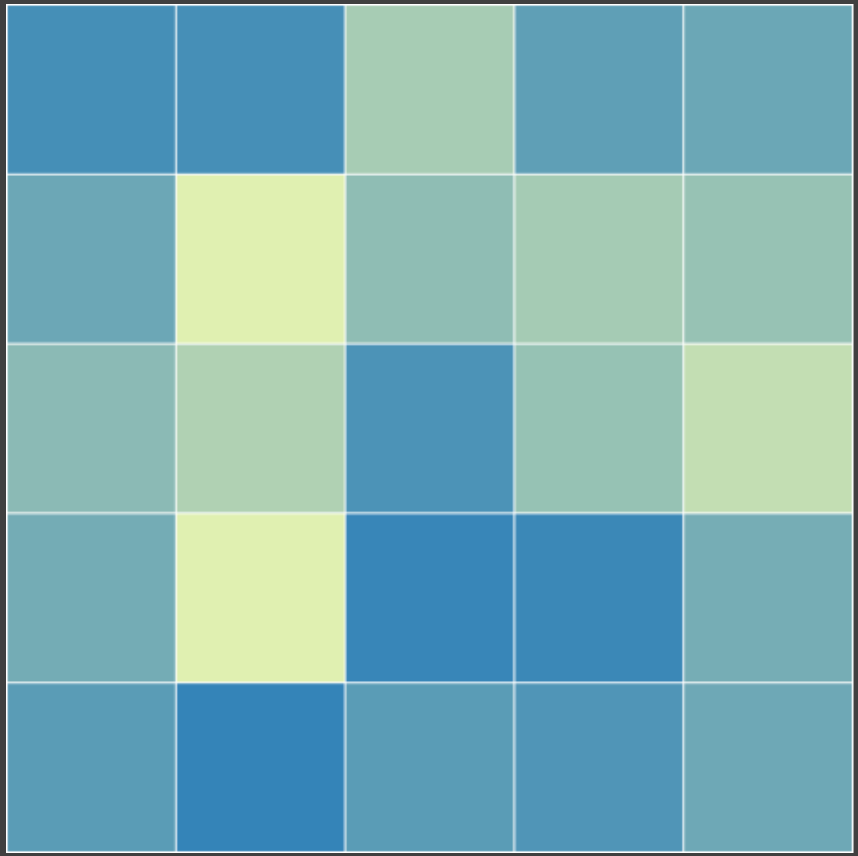

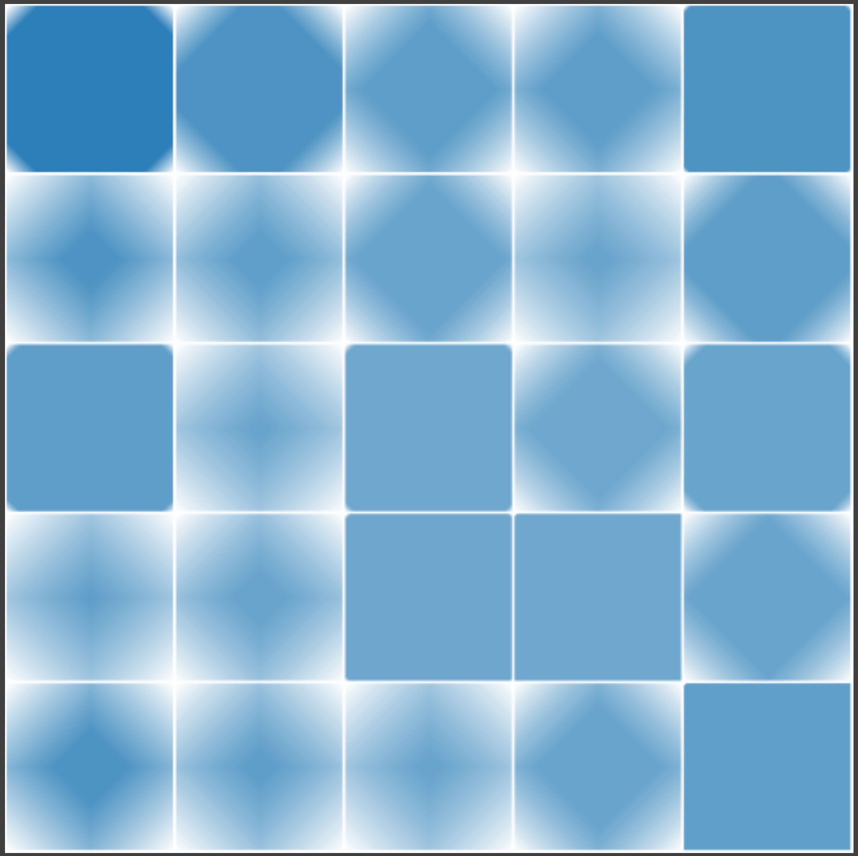

Semiotics of Uncertainty

Semiotics of Uncertainty

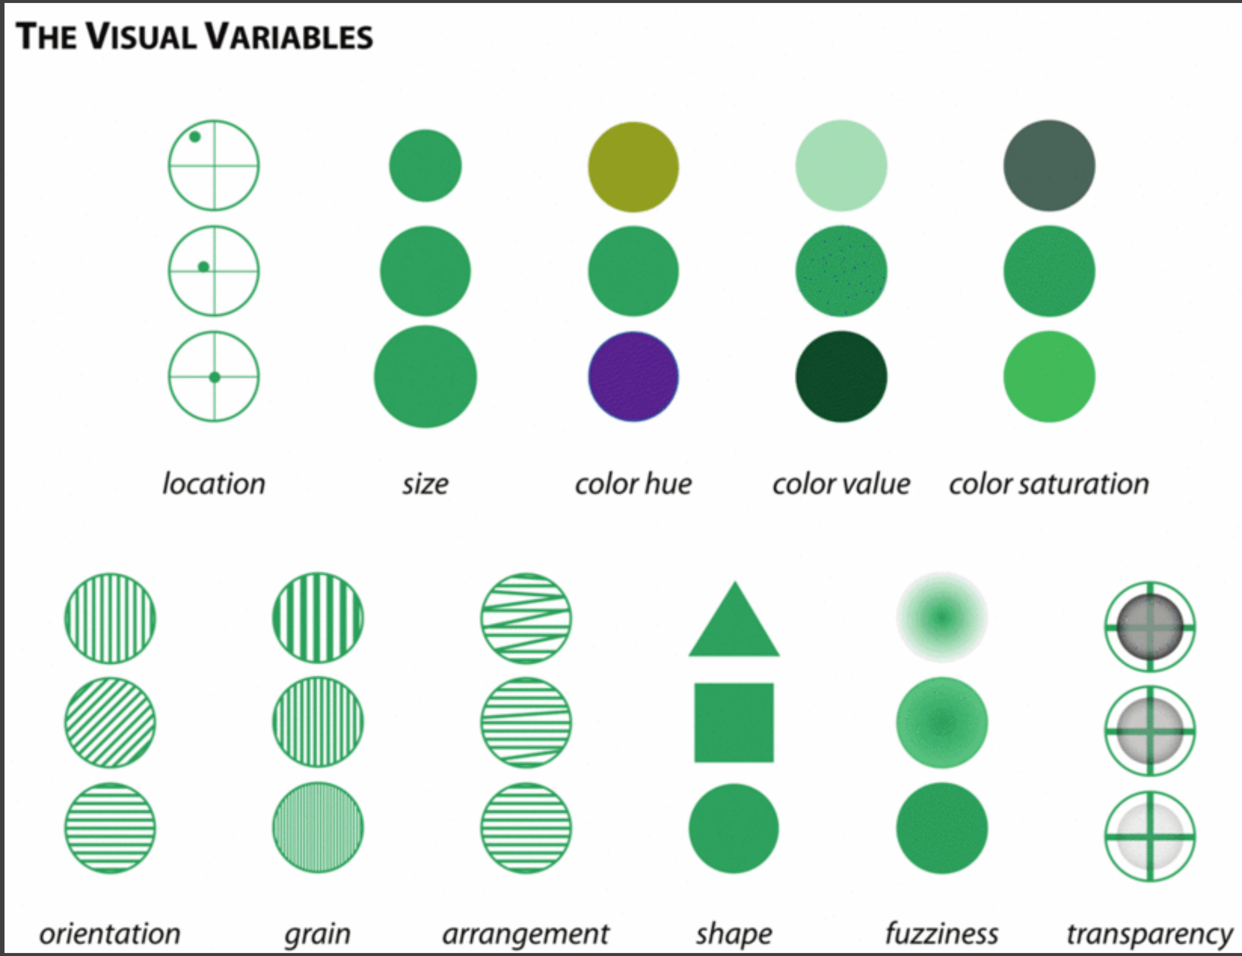

MacEachren et al. Visual Semiotics & Uncertainty Visualization: An empirical study. IEEE VIS, 2012.

Series #1: General Uncertainty by Visual Variable

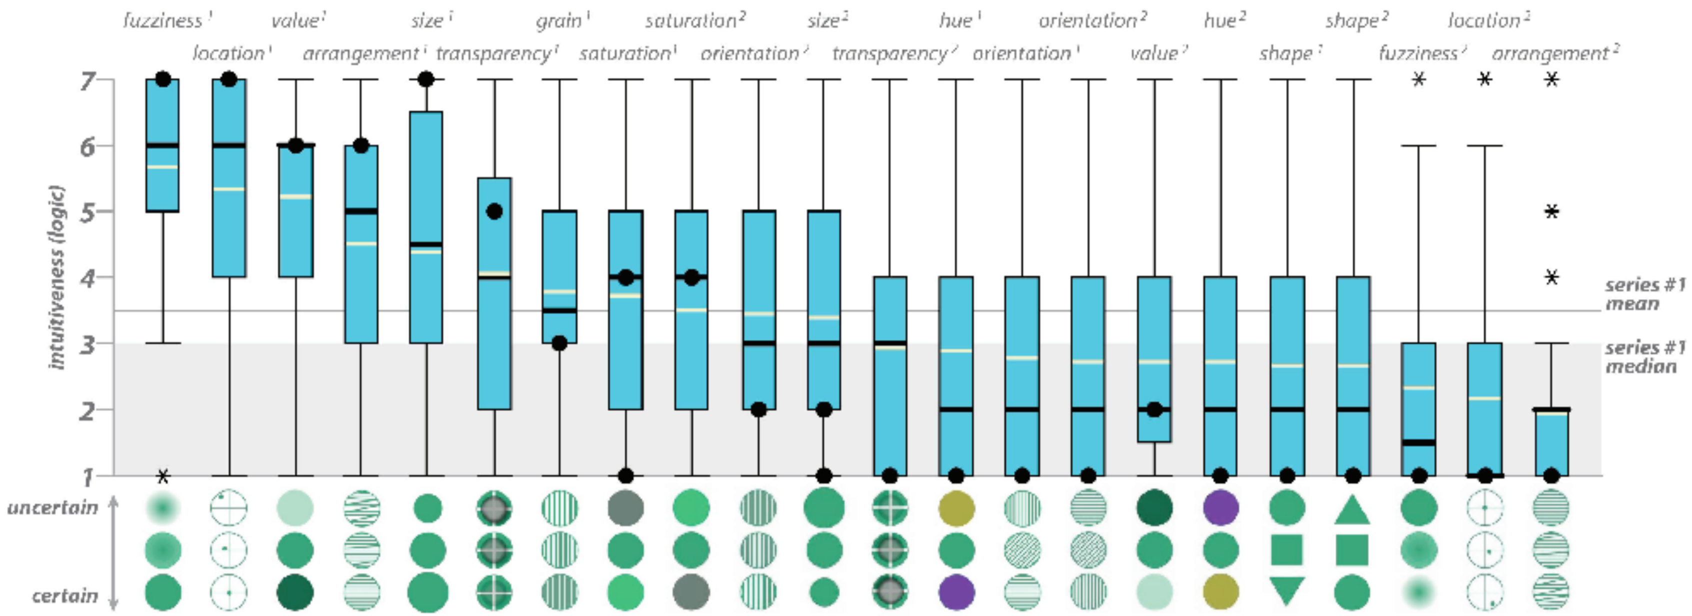

"Sketchiness"

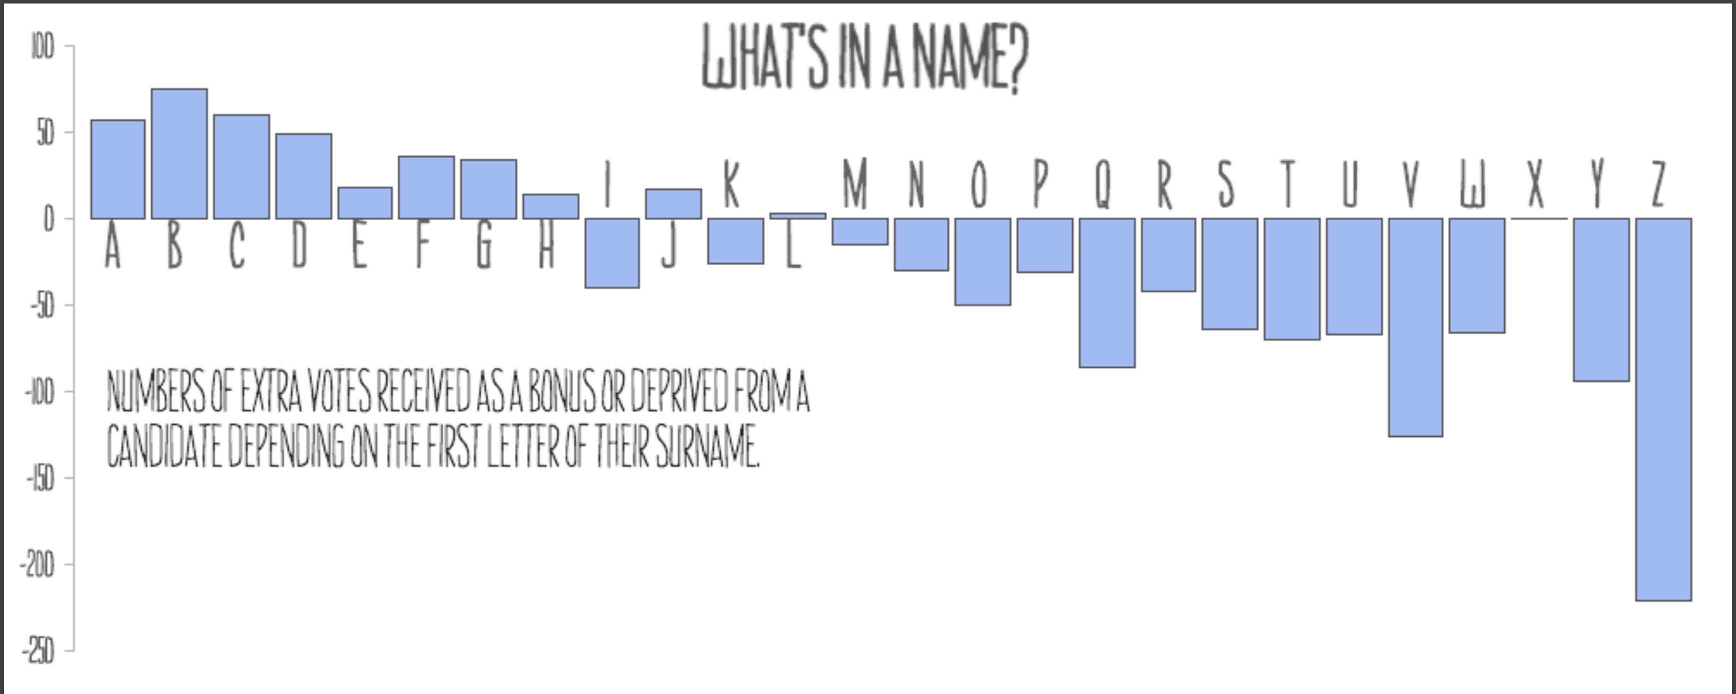

Wood, Jo et al. Sketchy rendering for information visualization. IEEE VIS, 2012.

Boukhelifa, Nadia et al. Evaluating skrtchiness as a visual variable for the depiction of qualitative uncertainty. IEEE VIS, 2012.

"Sketchiness"

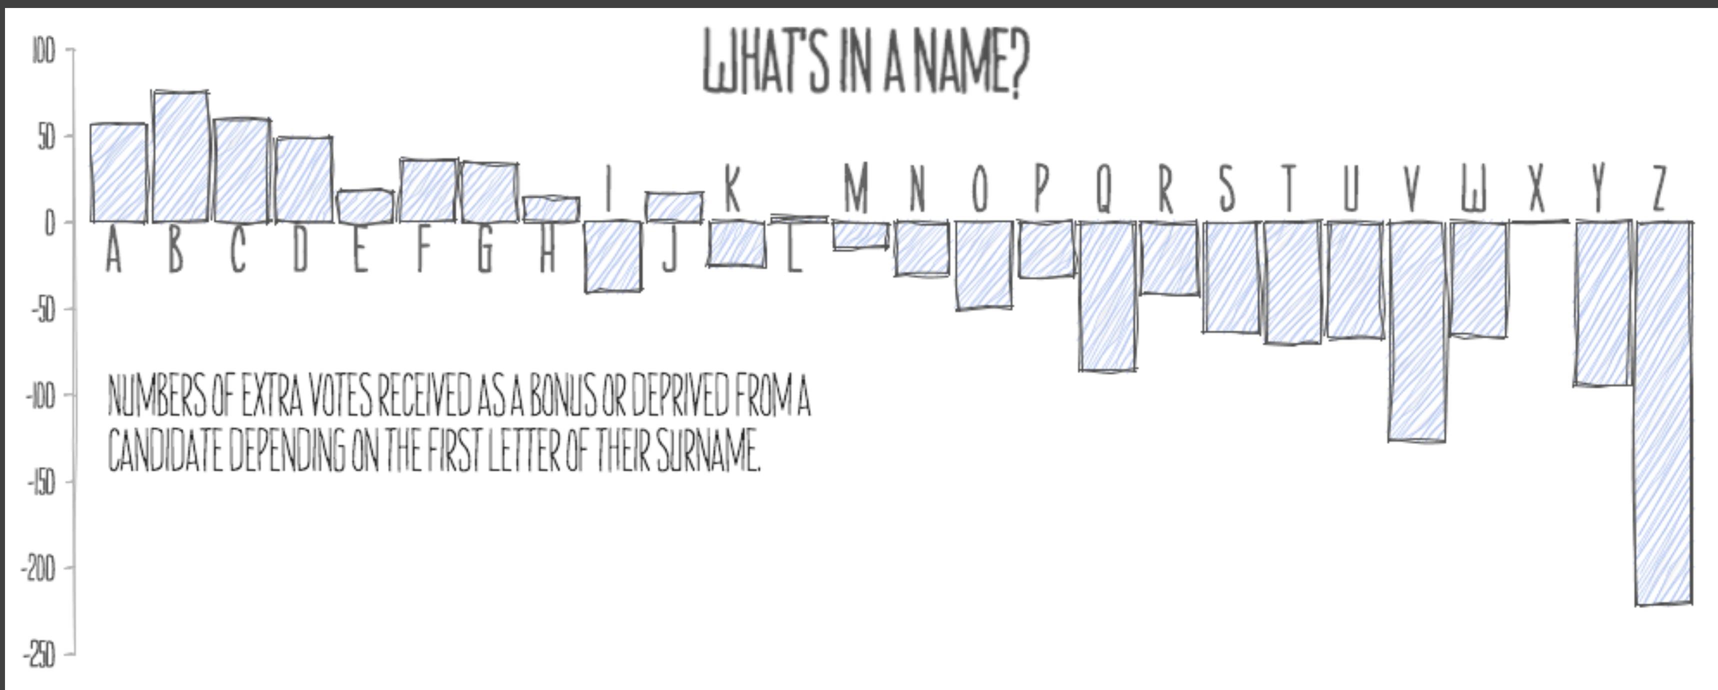

Wood, Jo et al. Sketchy rendering for information visualization. IEEE VIS, 2012.

Boukhelifa, Nadia et al. Evaluating skrtchiness as a visual variable for the depiction of qualitative uncertainty. IEEE VIS, 2012.

Encoding Uncertainty

Some visual variables (like fuzziness) have a **semiotic connection** to uncertainty.

However, “intuitive” variables may not always be effective / accurately interpreted!

Error Bars

p-values

The probability of results at least as extreme as the observed results, given some null hypothesis: $\mathbf{p = P(D \mid H_0)}$

If $p < \alpha$ (usually 0.05), then the result is considered to be *statistically significant*.

Error Bars

The mean treatment effect is higher than than the placebo.

Is this difference in means *statistically significant?*

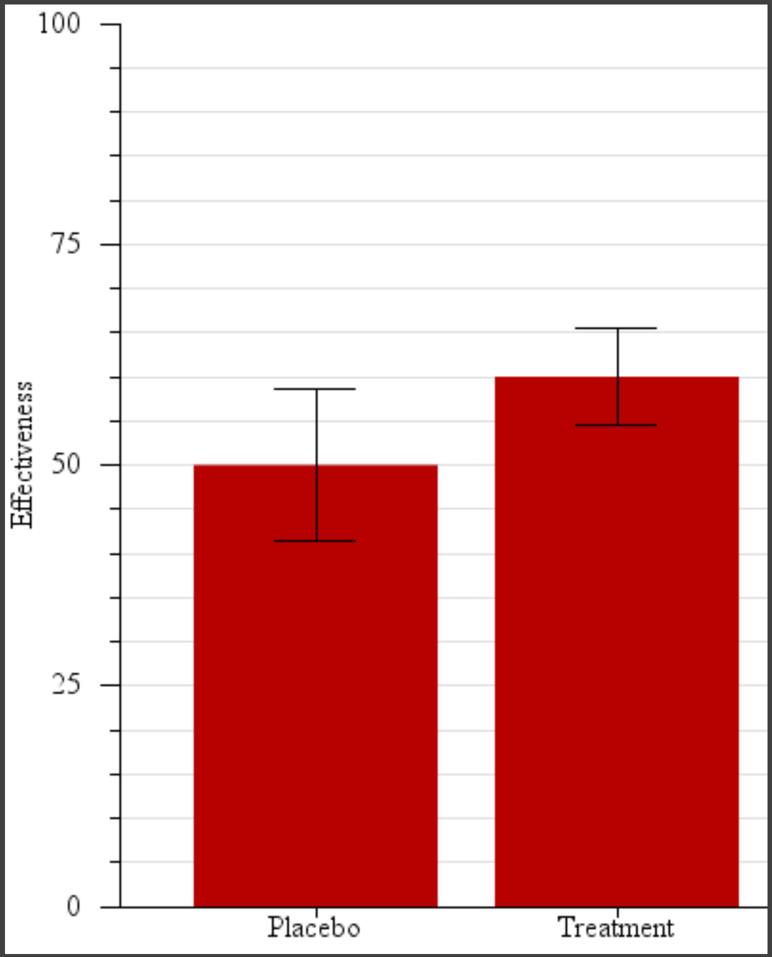

Error Bars

Standard Deviation?

Standard Error ($\sigma/\sqrt{n}$)?

T-Confidence Interval?

Z-Confidence Interval?

Bootstrapped Interval?

Min/Max?

1.5 * IQR (Q3-Q1)?

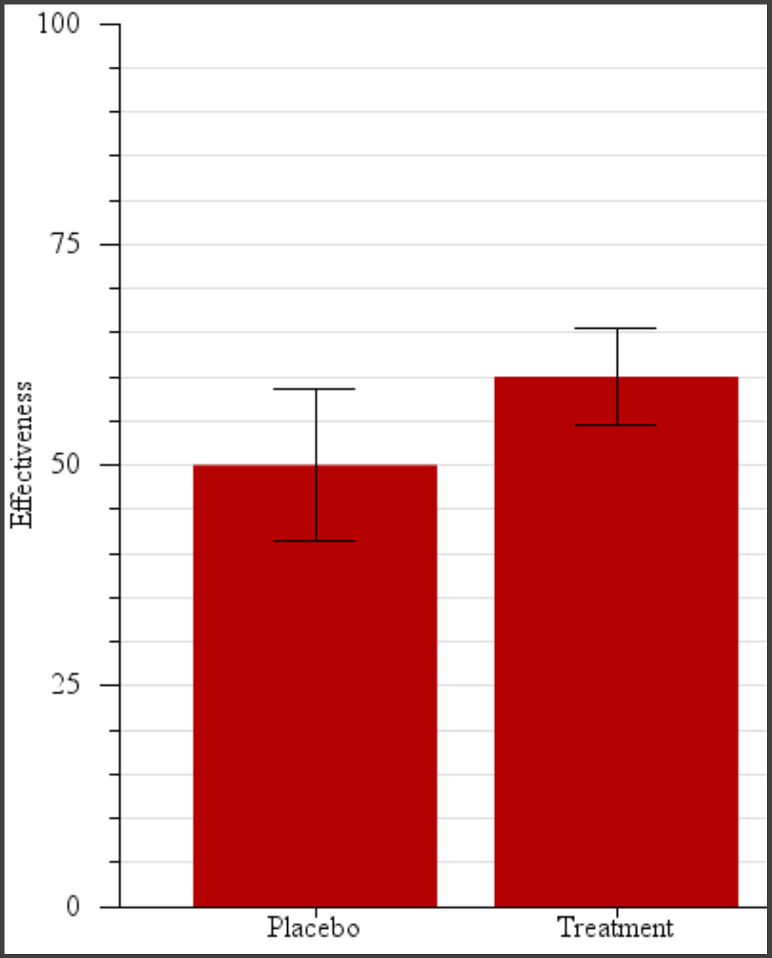

Guess the p-value…

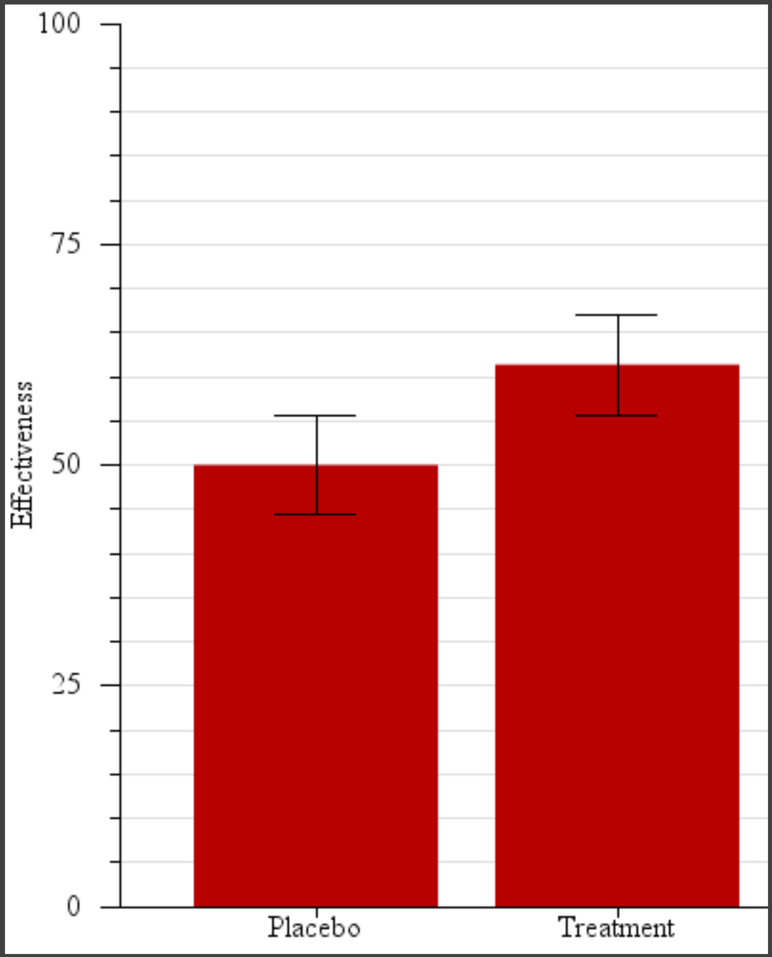

Error bars depict
95% Conf. Interval

Guess the p-value…

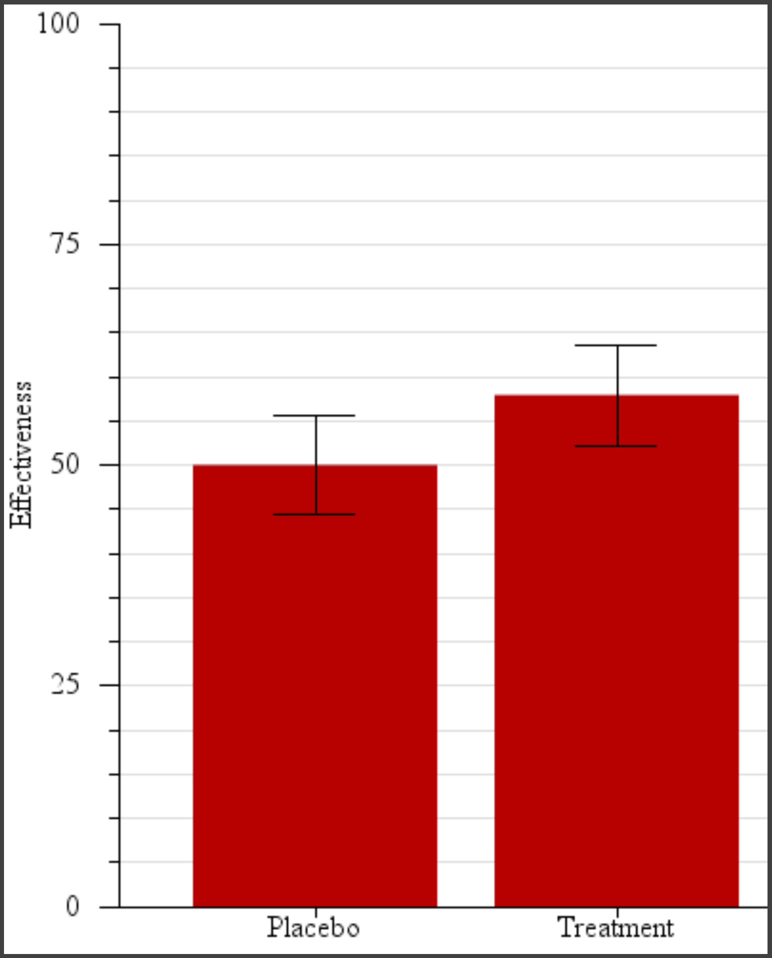

Error bars depict
95% Conf. Interval

Guess the p-value...

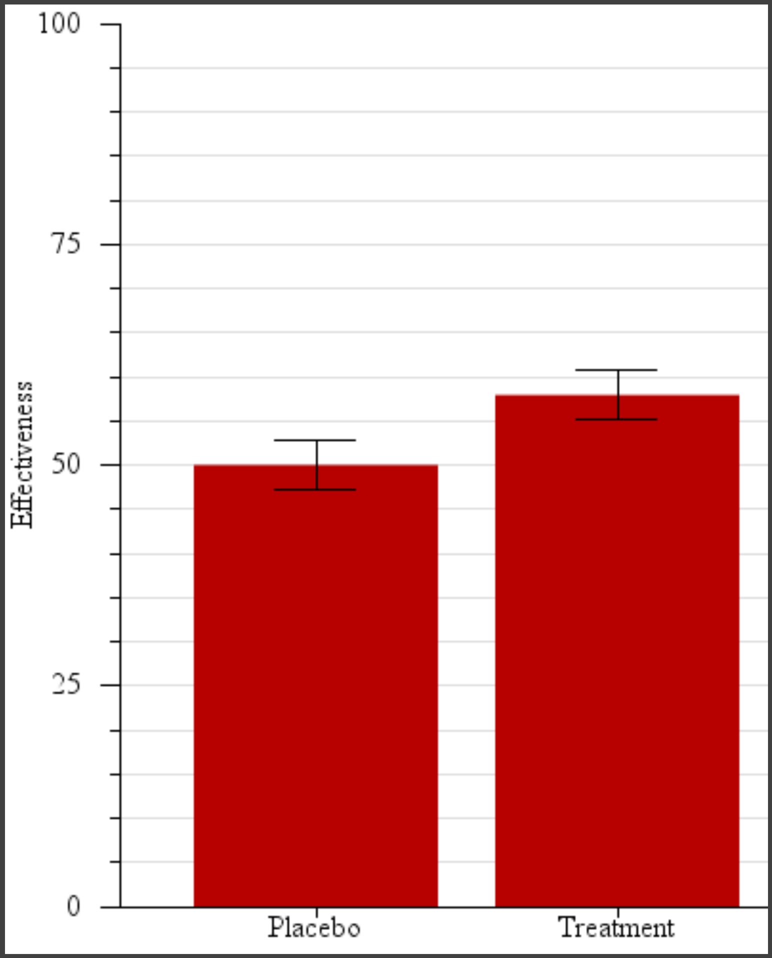

Error bars depict
standard error

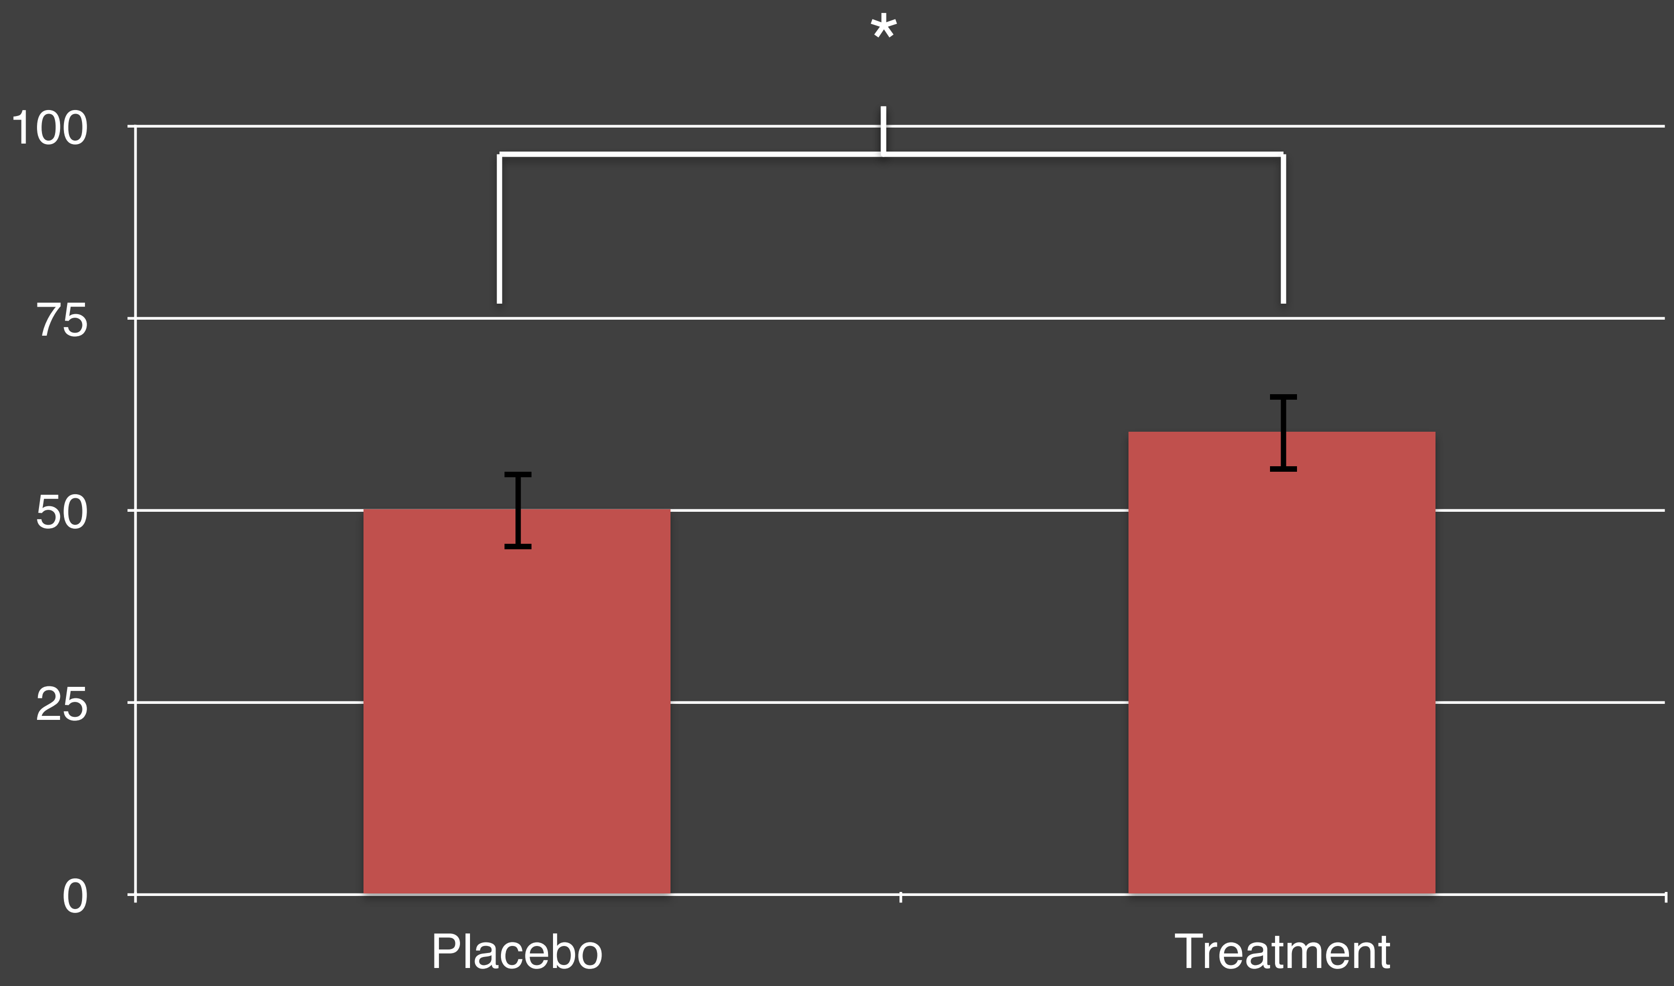
*
100
75
50
25
0
Placebo
Treatment

Inference by Eye

95% CIs

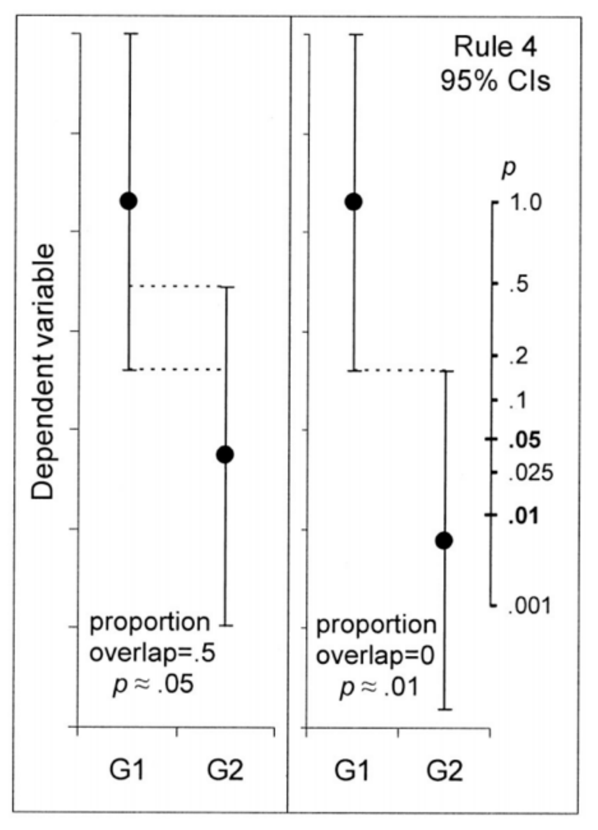

Standard Error

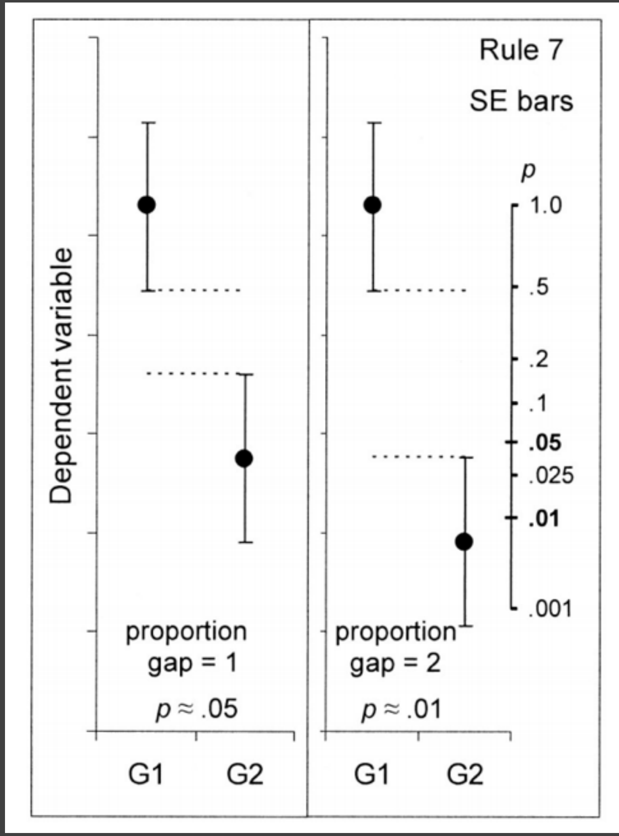

Cumming, Geoff and Finch, Sue. Inference by eye: confidence intervals and how to read pictures of data. American Psychologist, 2005.

Confidence Intervals

What does a 95% confidence interval indicate?

One "folk" interpretation: there is a 95% chance that the population mean is within the interval.

Wrong!

True answer: Given an infinite number of indep. experiments, 95% of the confidence intervals generated will contain the true population mean.

"Confidence" concerns the procedure, not the data. (Though see Bayesian *credible intervals*…)

Confidence Intervals

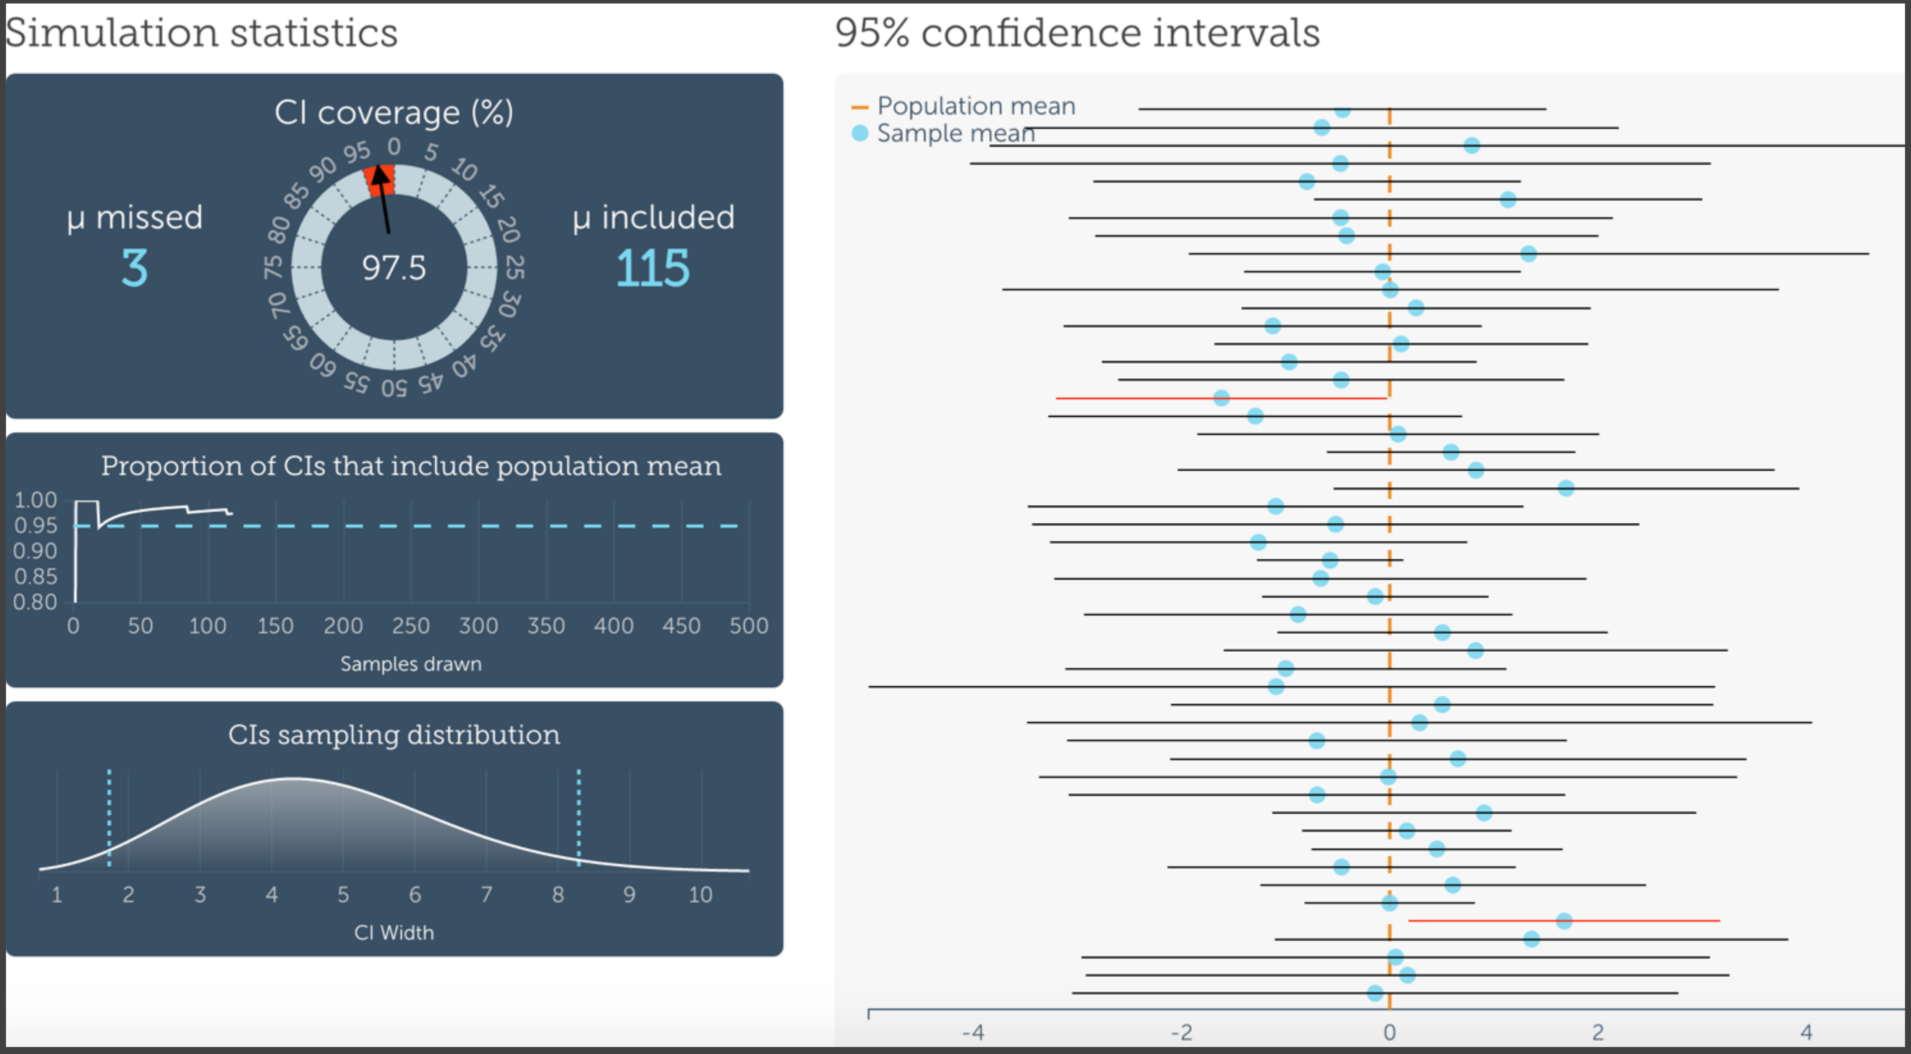

Error Bar Visualization

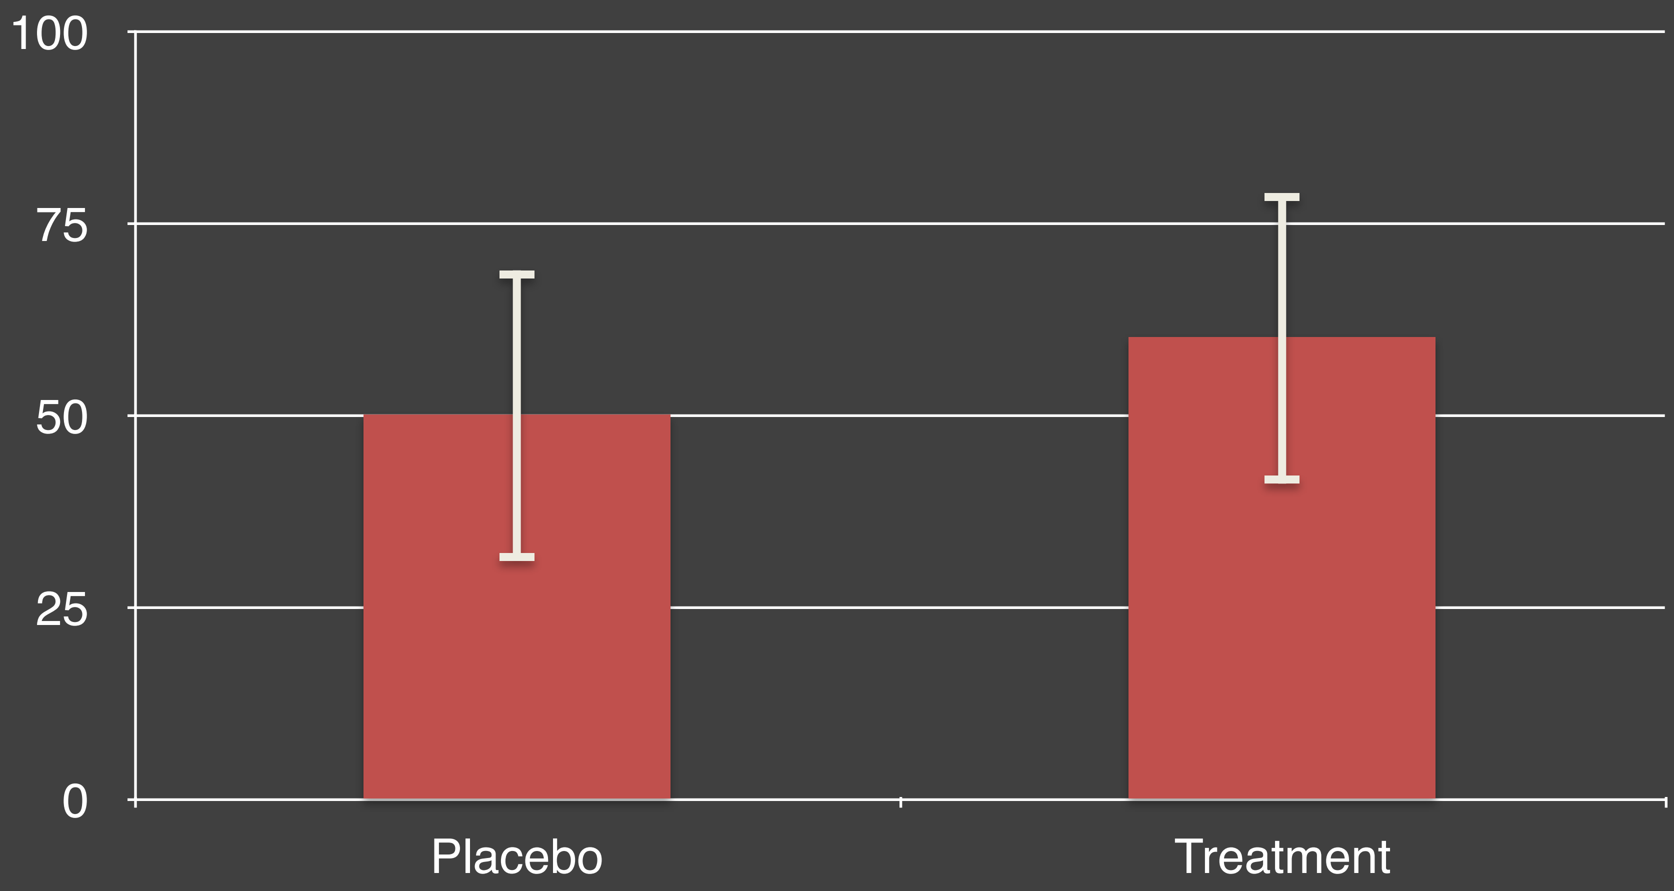
100
75
50
25
0
Placebo
Treatment

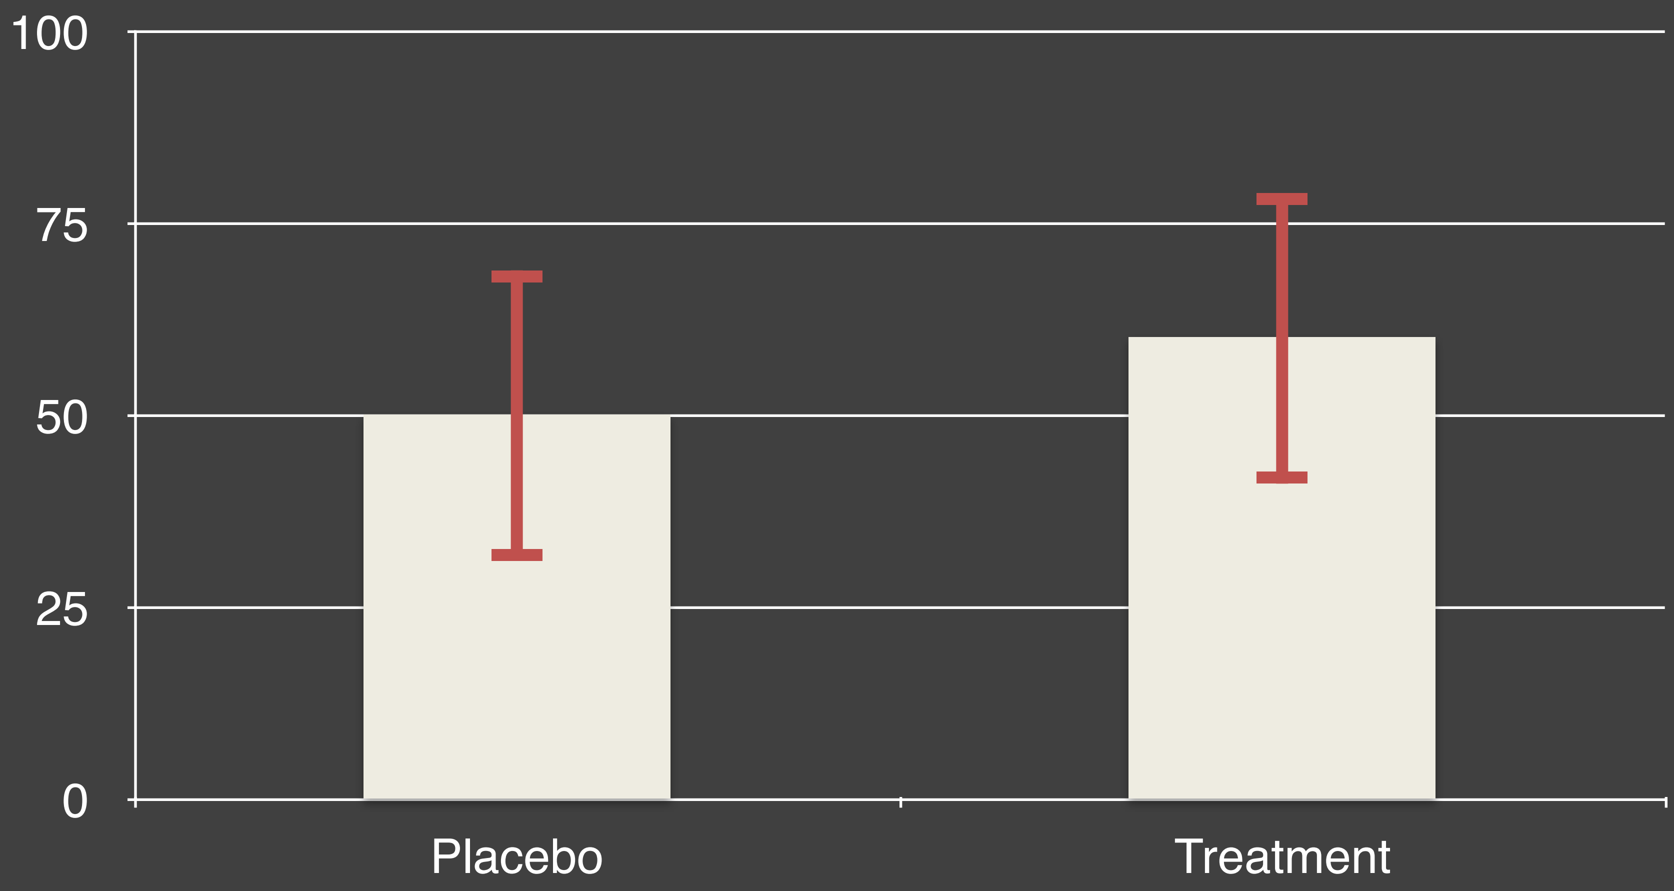
100
75
50
25
0
Placebo
Treatment

Within-the-bar bias

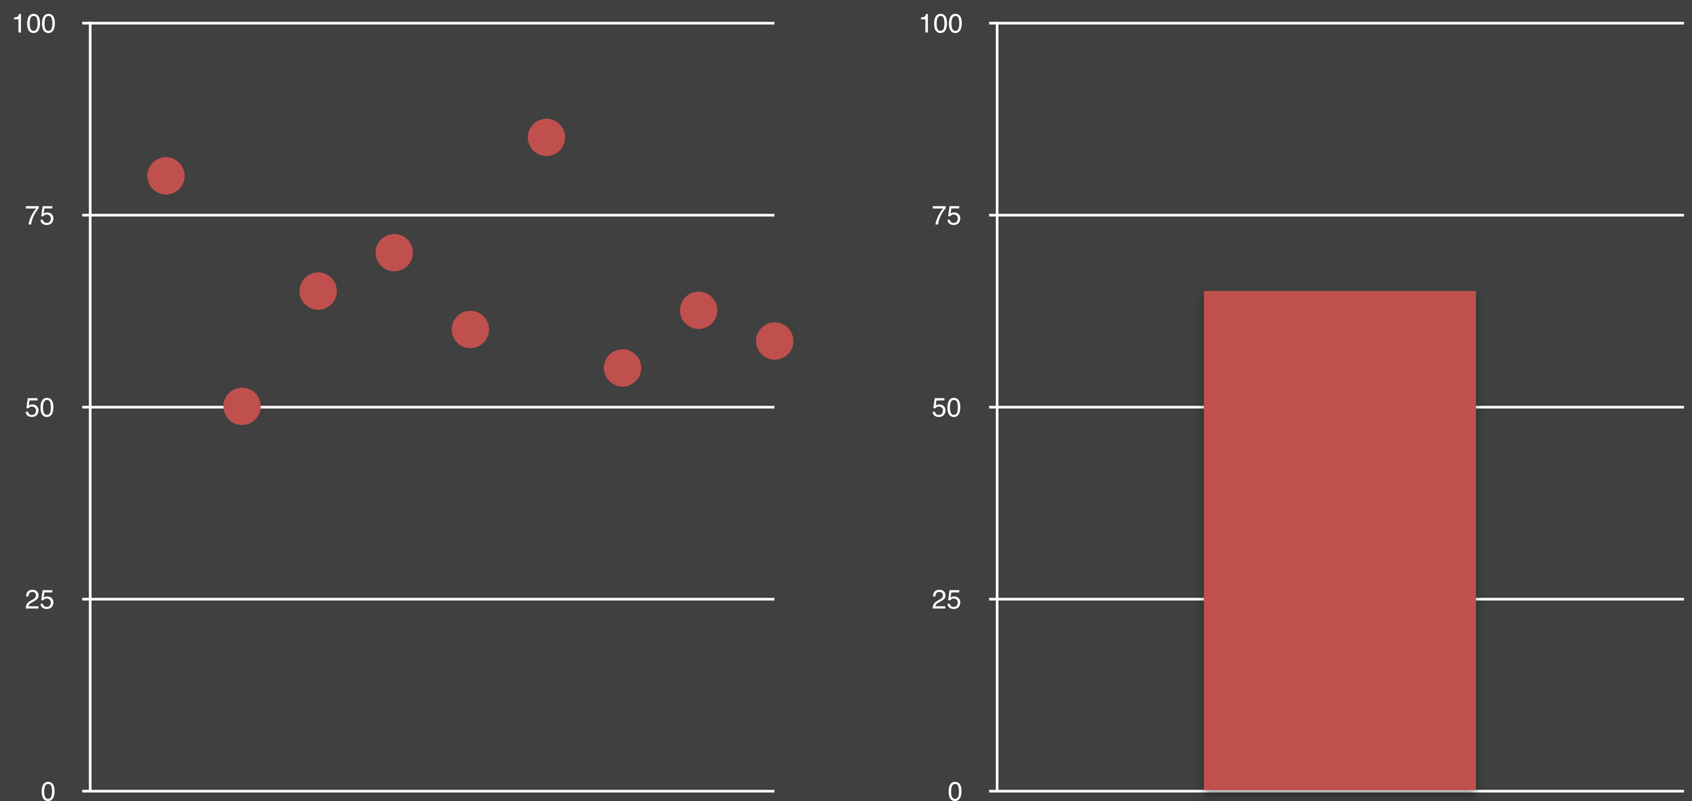

Newman & Scholl. (2012) "Bar graphs depicting averages are perceptually misinterpreted: the within-the-bar bias."

Within-the-bar bias

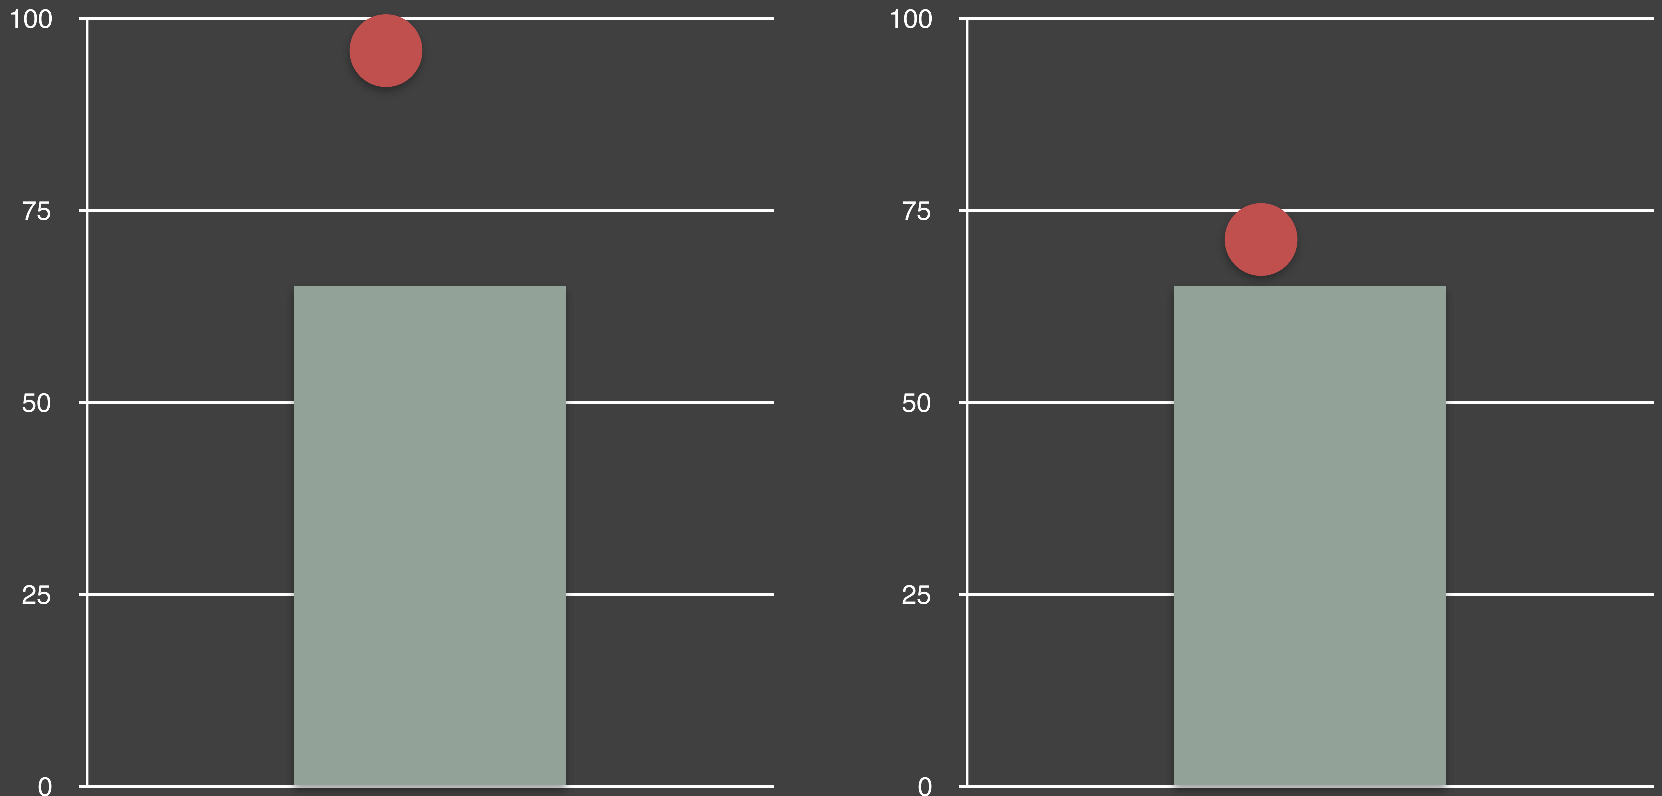

Within-the-bar bias

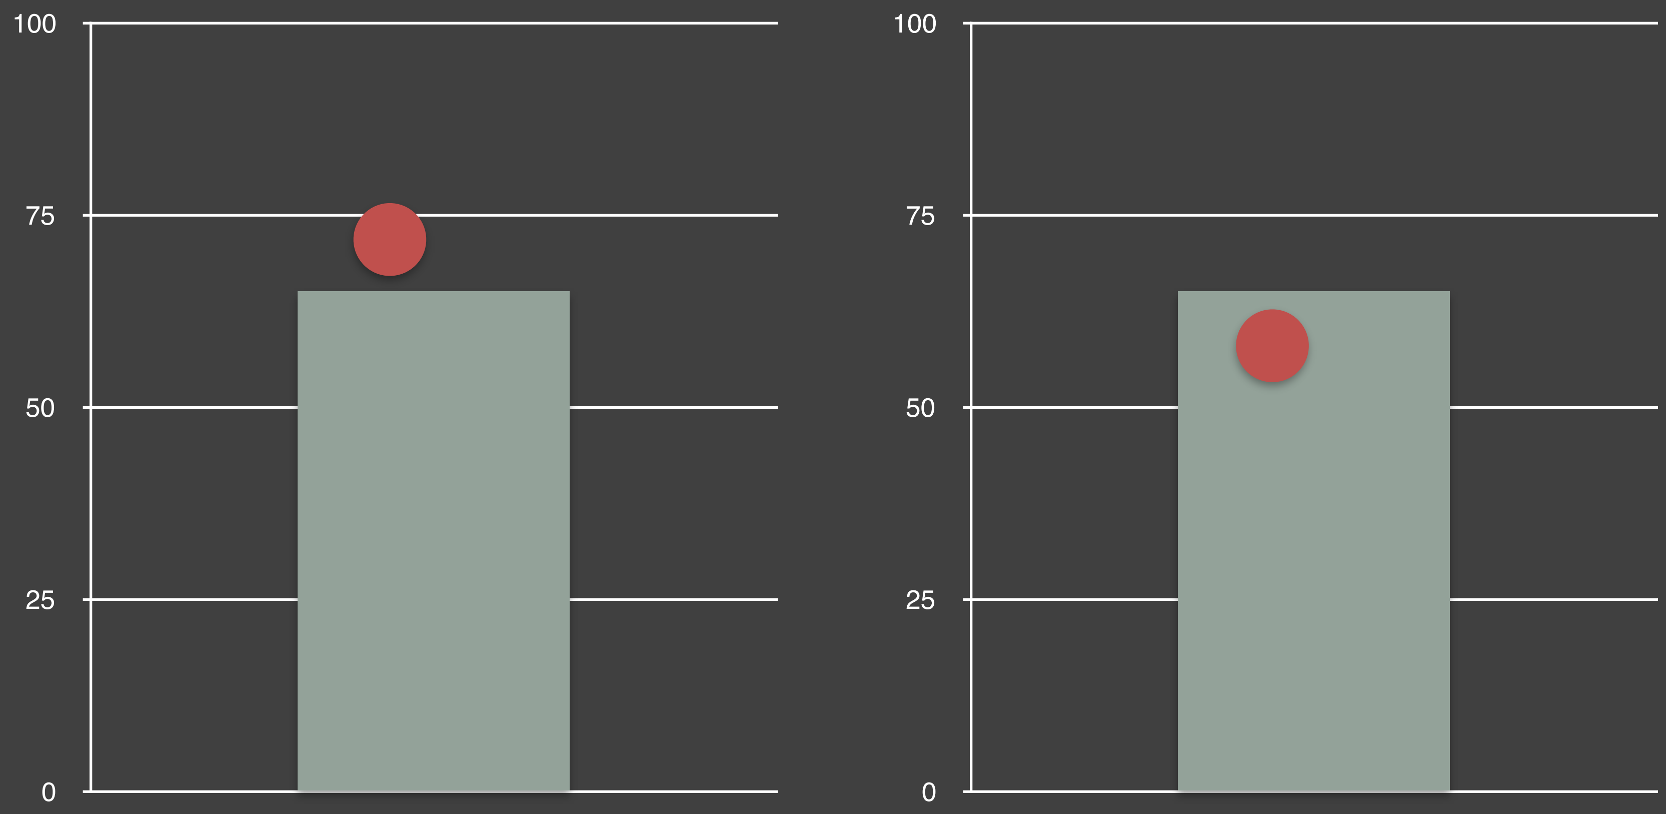

Within-the-bar bias

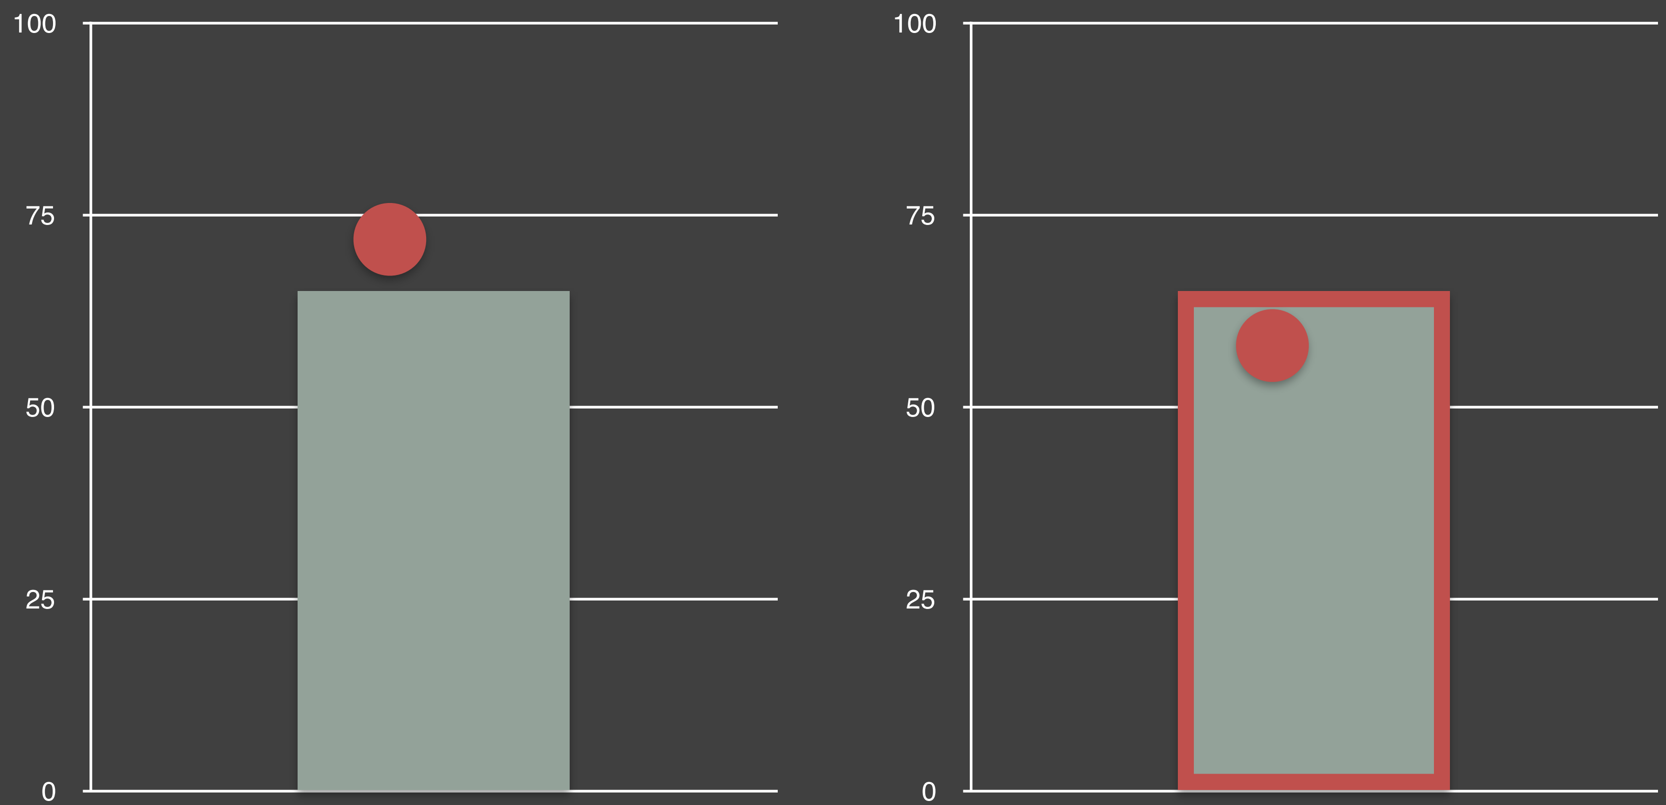

Alternatives

Gradient Plot

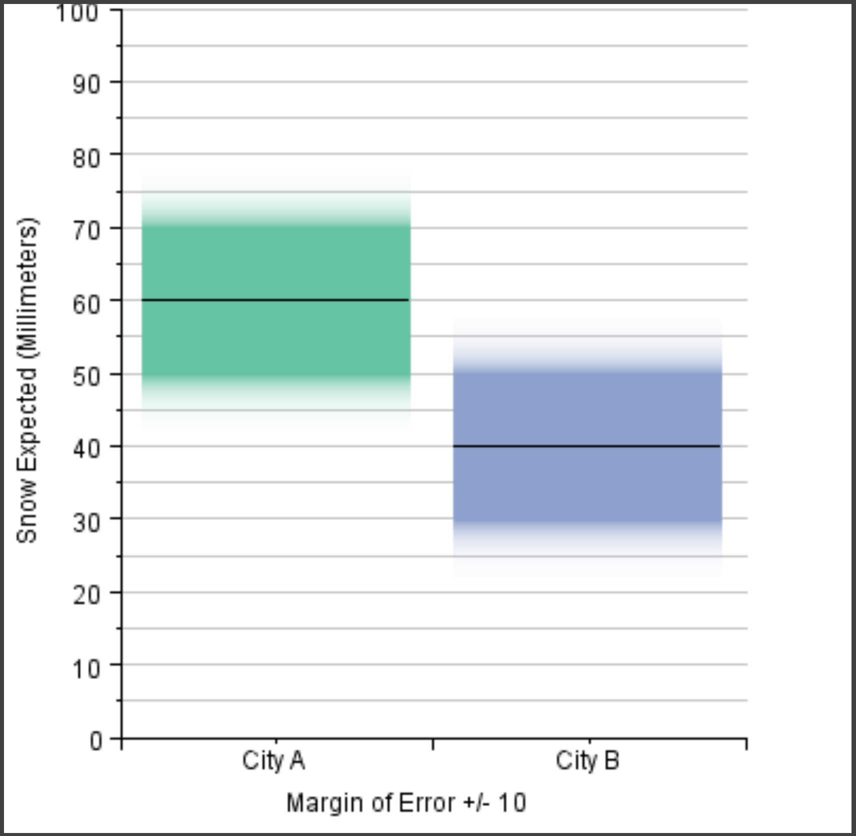

Violin Plot

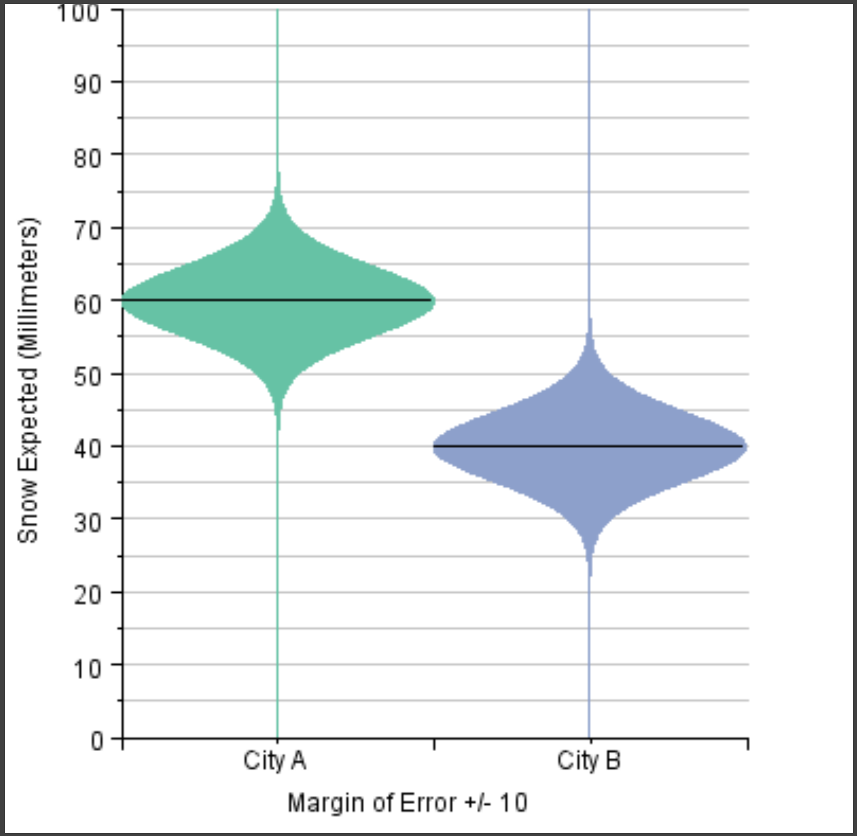

Model Visualization

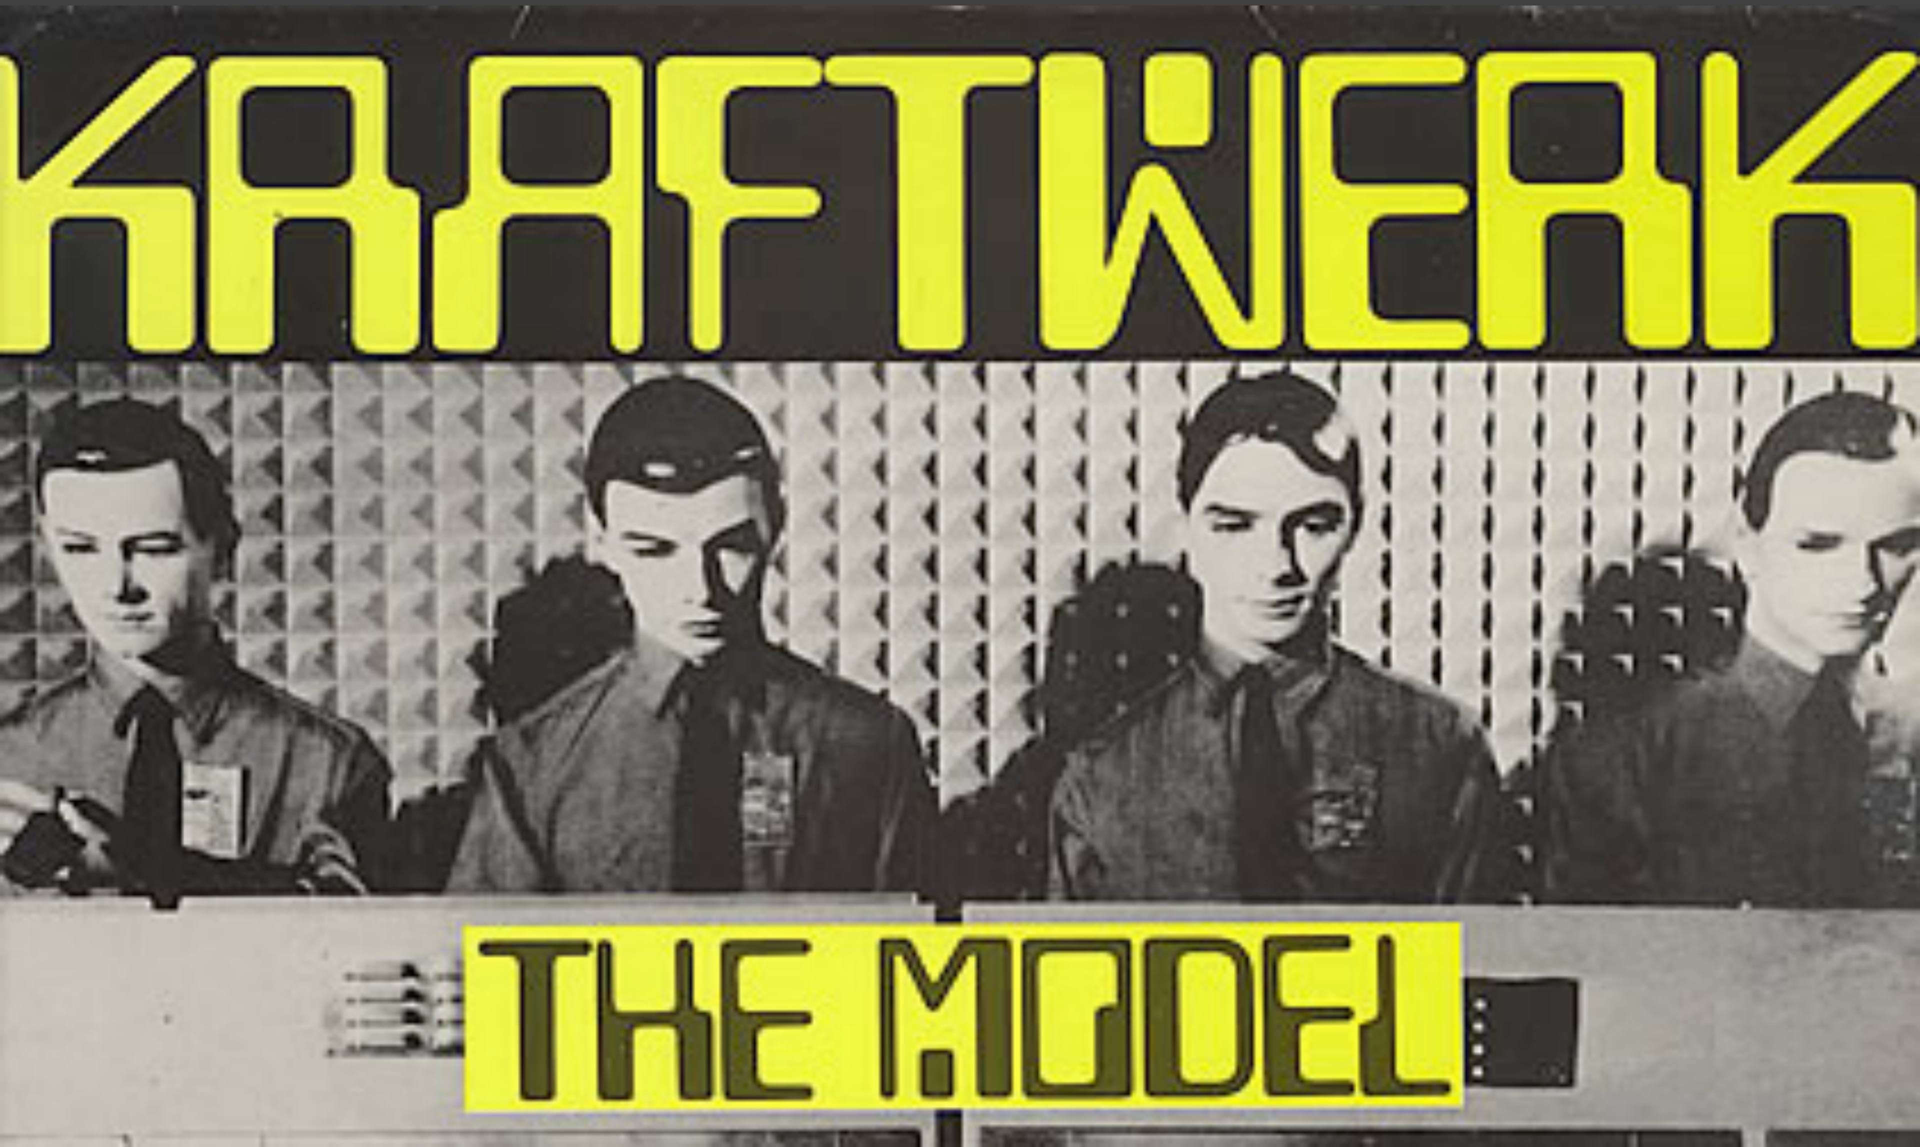

Polling Data

Candidate A is ahead of Candidate B in the polls, with 55% of the likely voters*

Polling Data

Candidate A is ahead of Candidate B in the polls, with 55% of the likely voters*

*poll of 100 people, margin of error +/-5

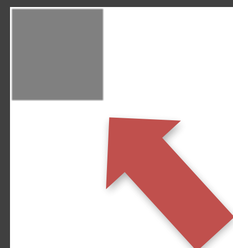

A Likely Voter

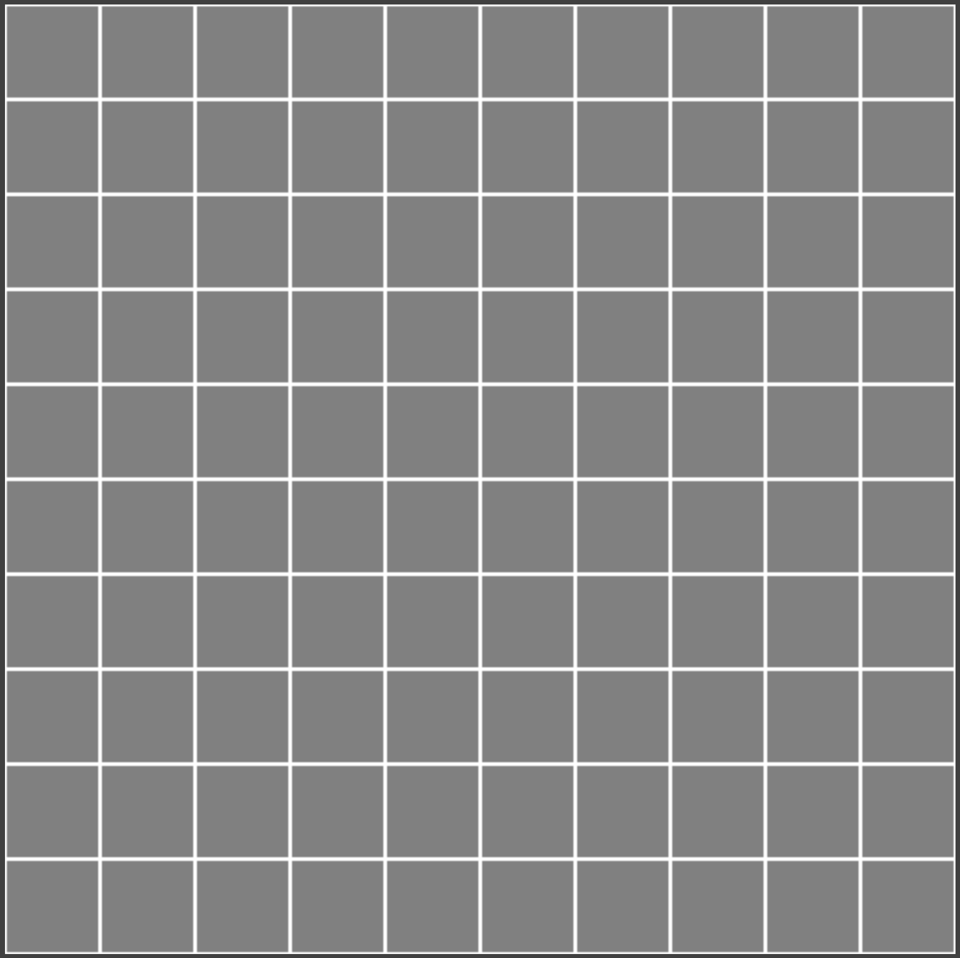

Poll

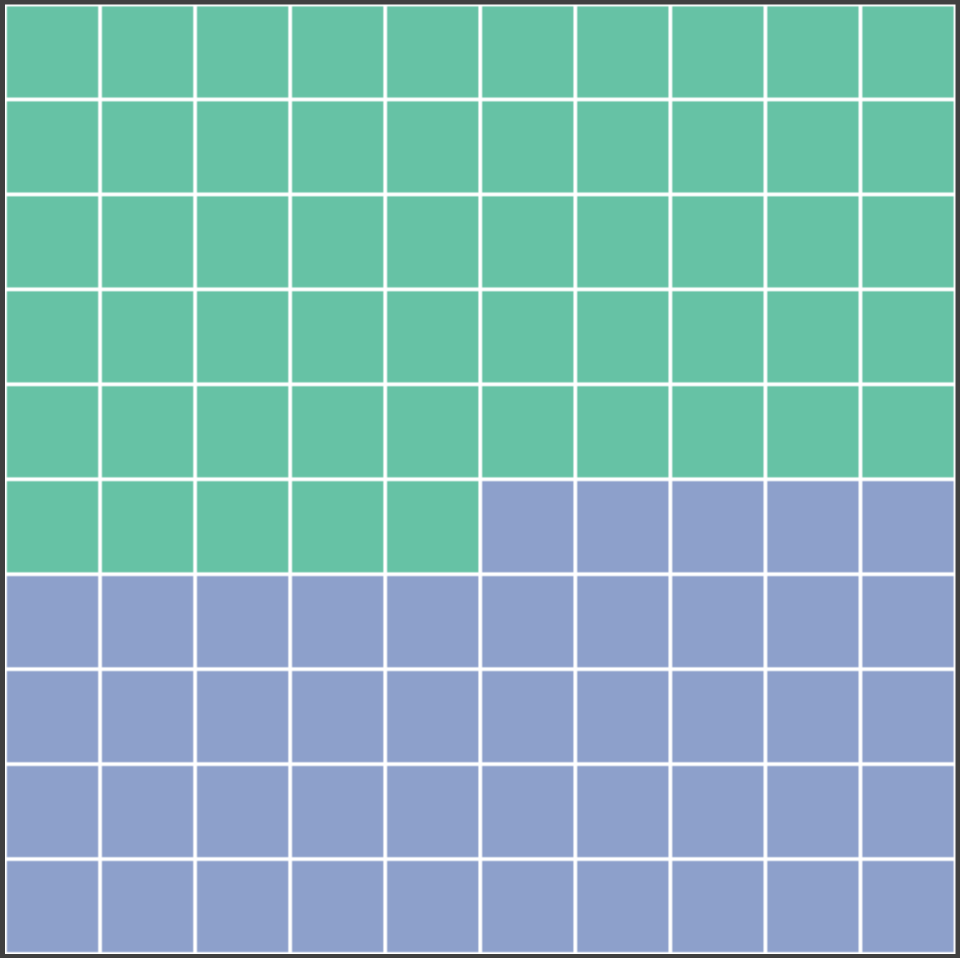

Poll

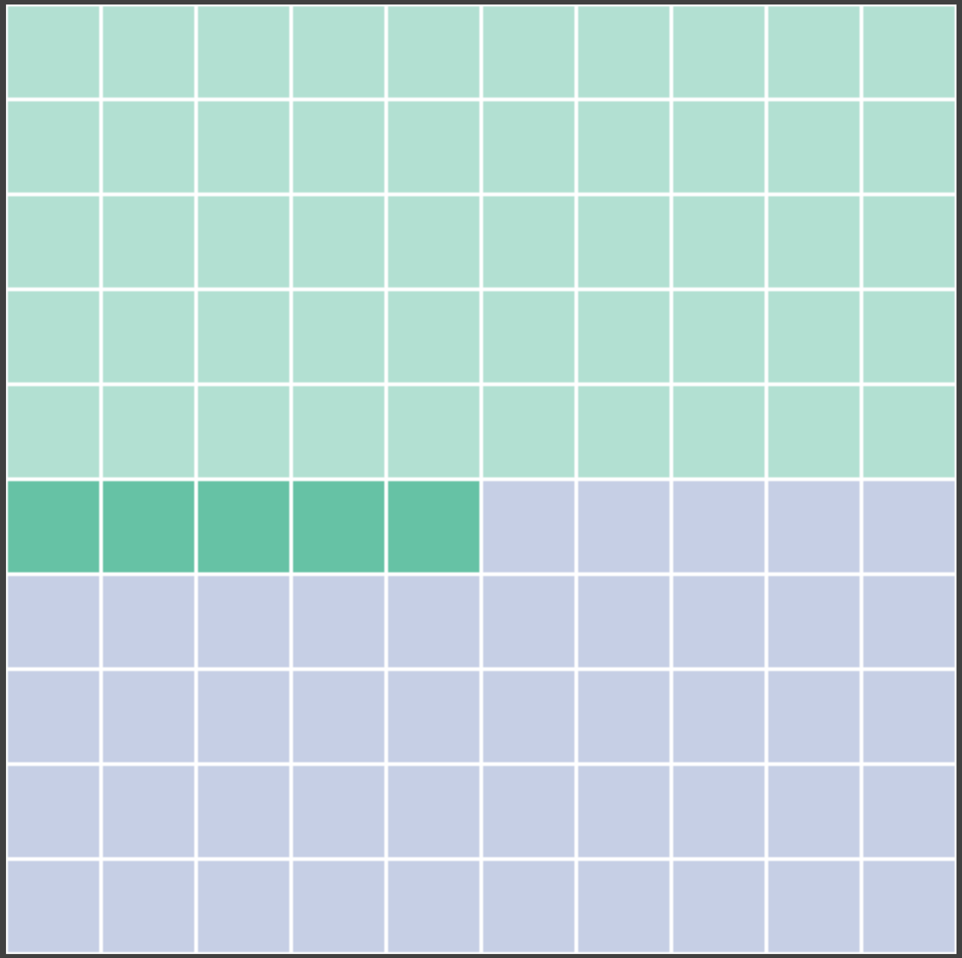

Actual Election?

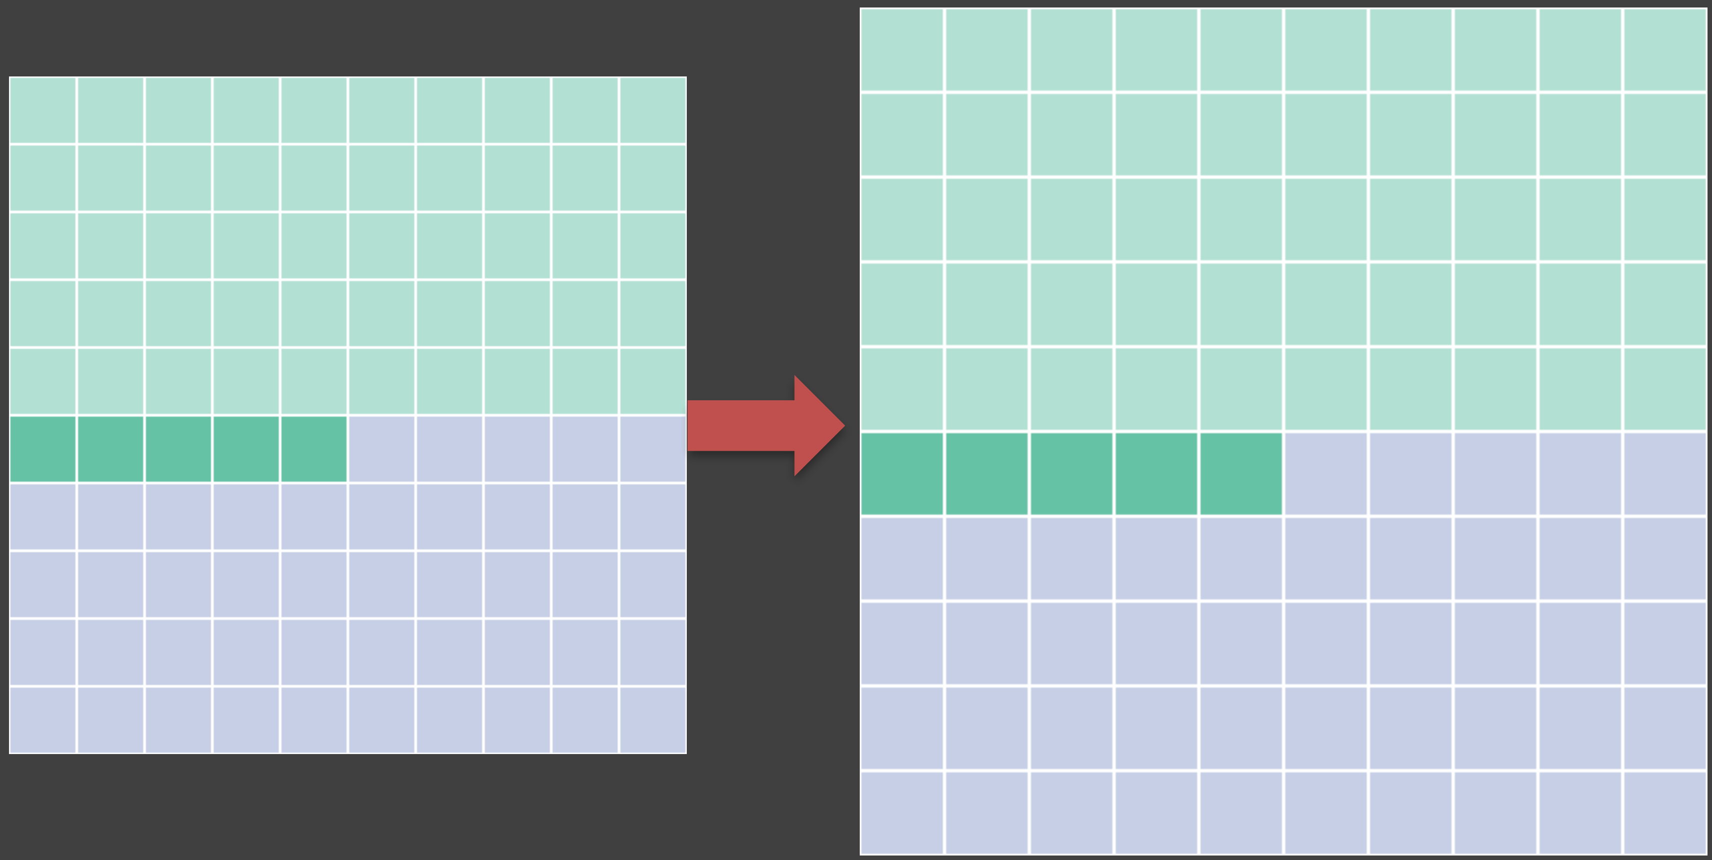

Actual Election?

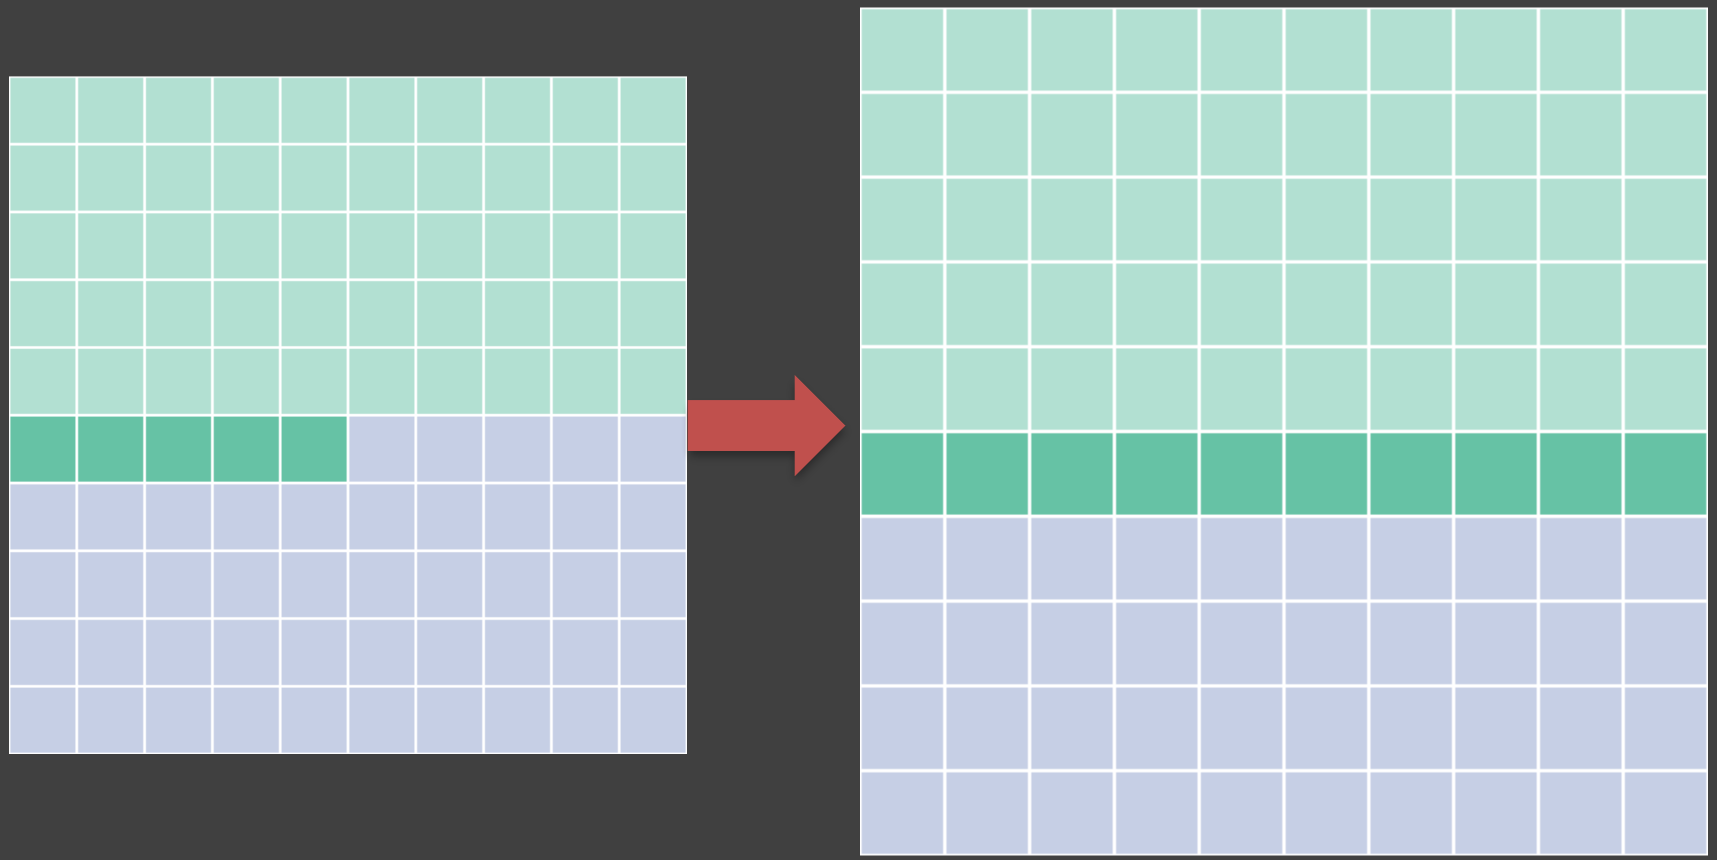

Actual Election?

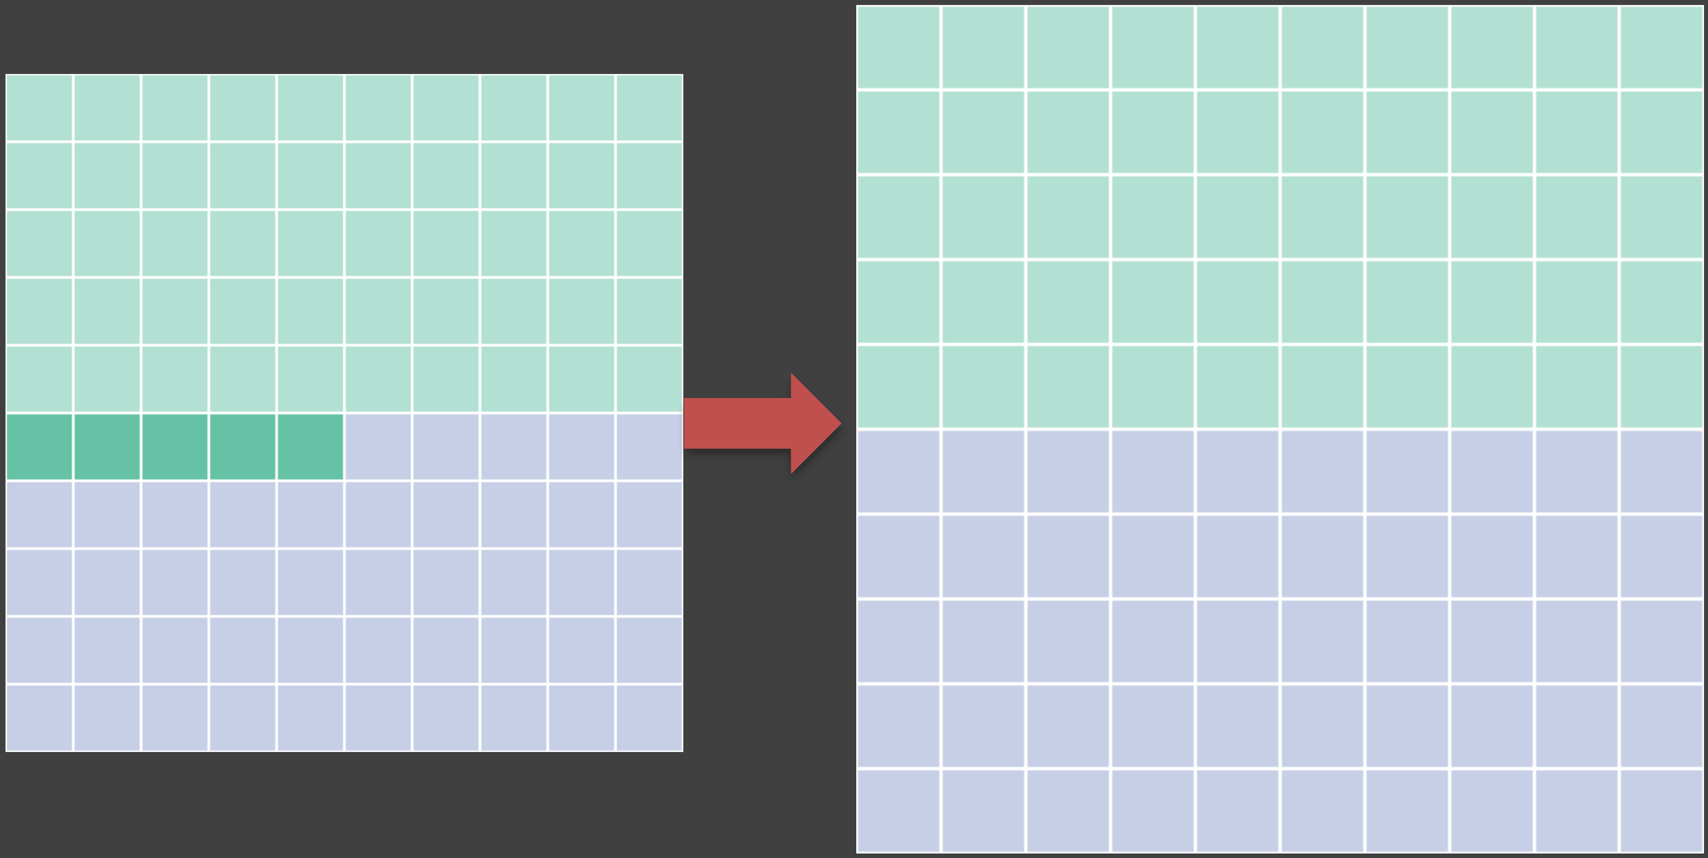

Actual Election?

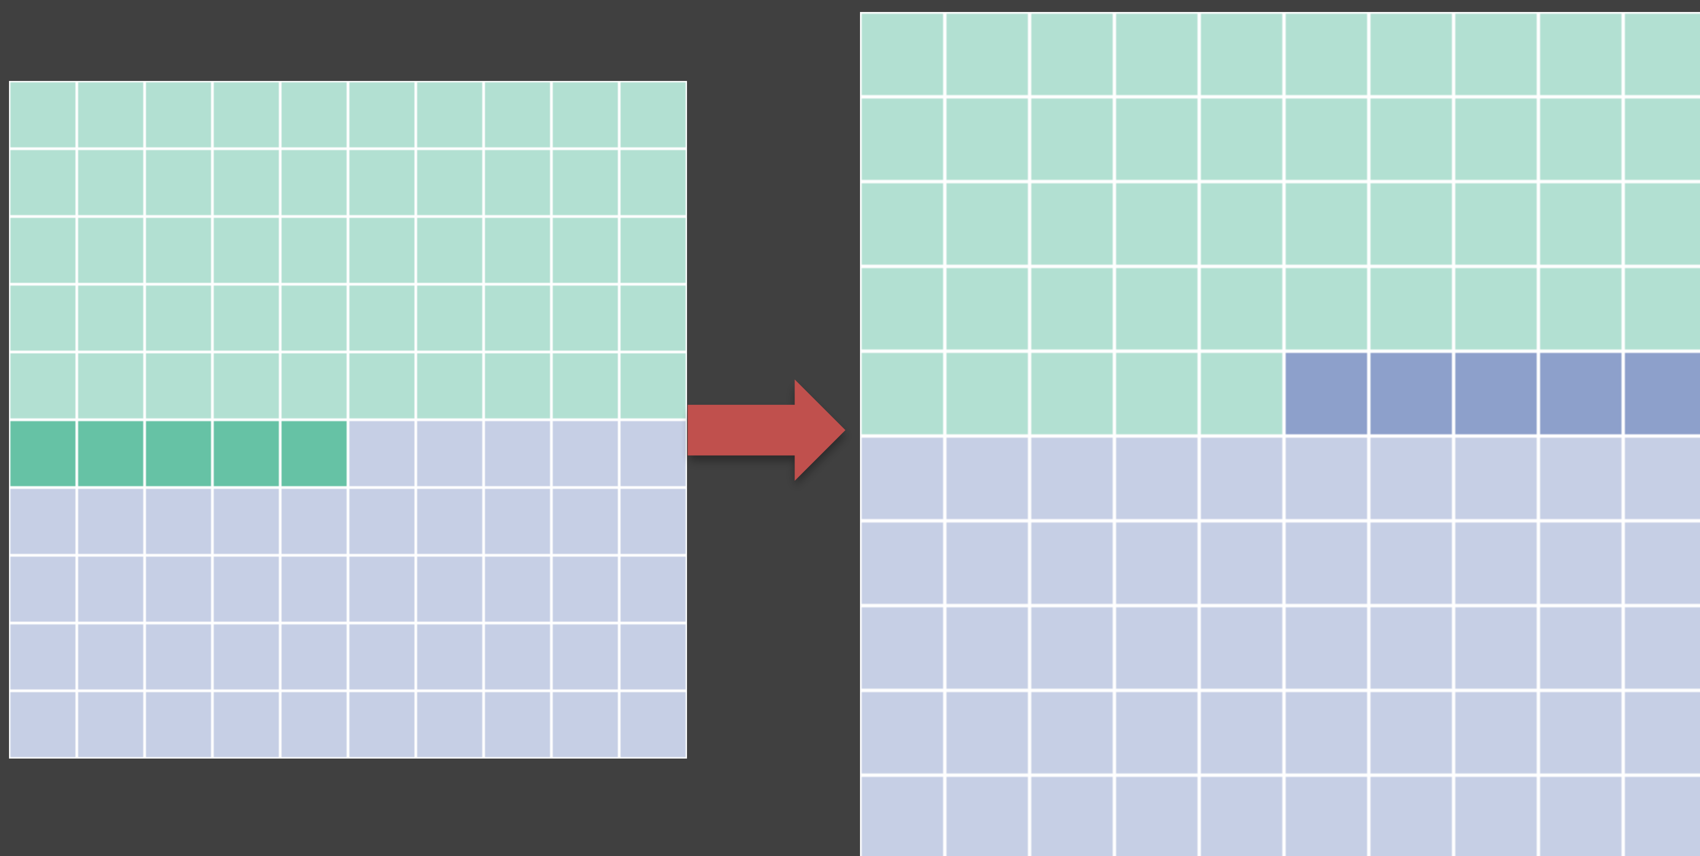

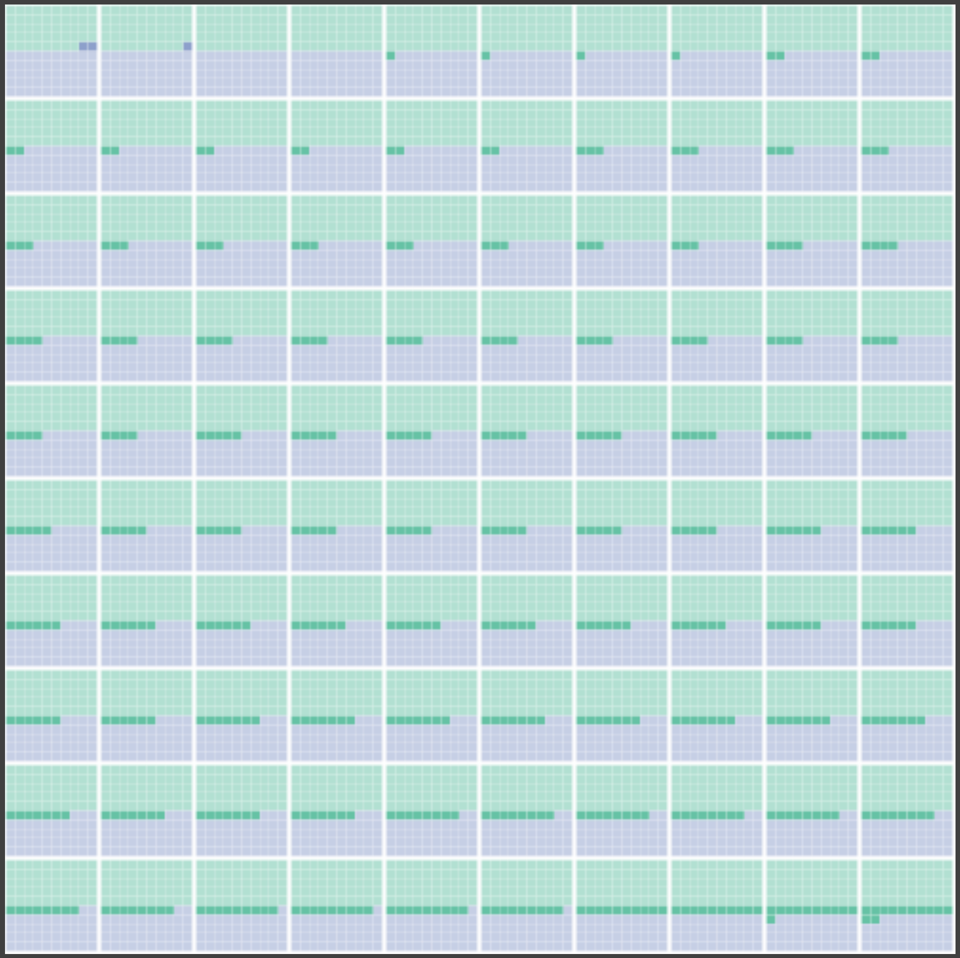

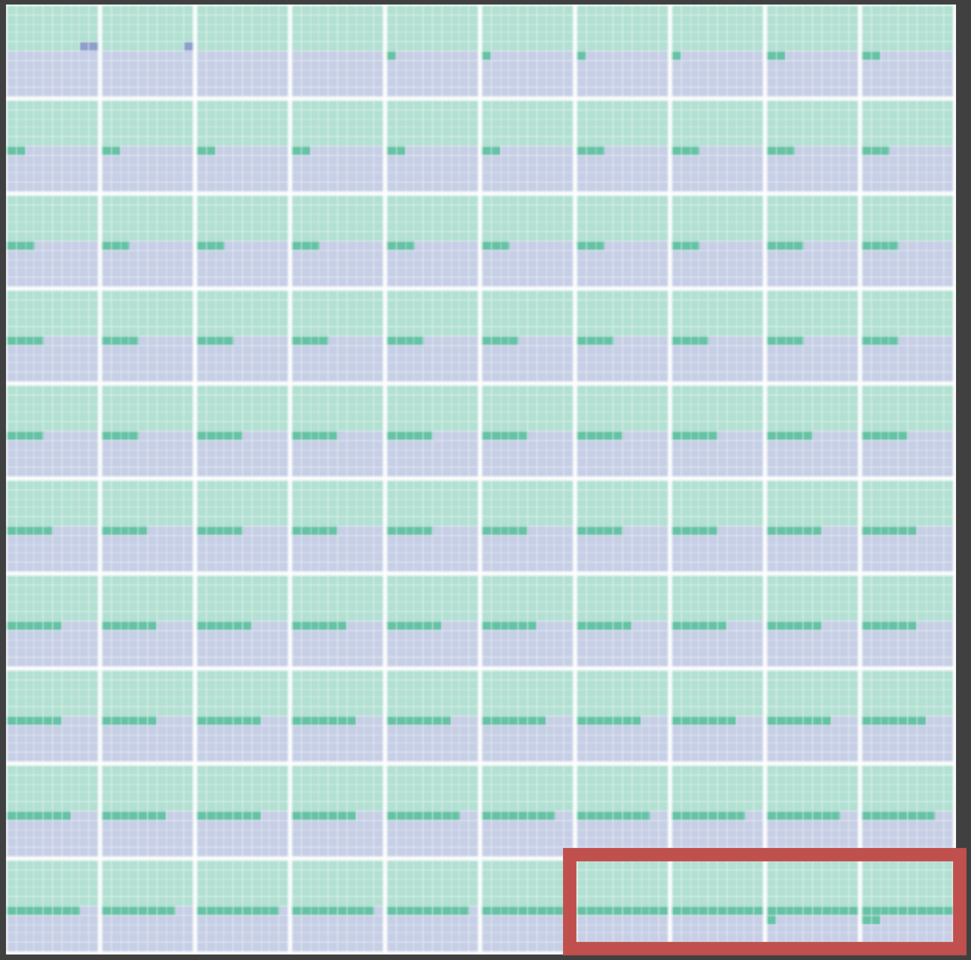

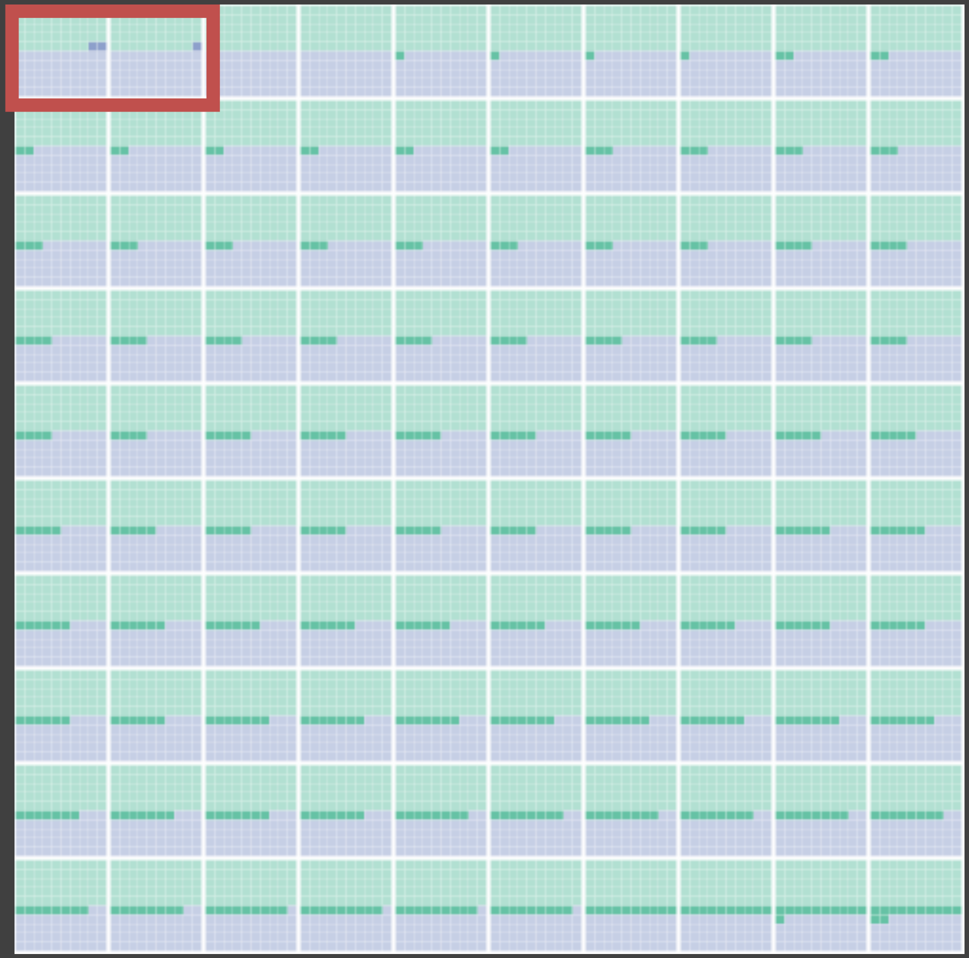

Pangloss Plot

Candidate A is ahead of Candidate B in the polls, with 55% of the likely voters*

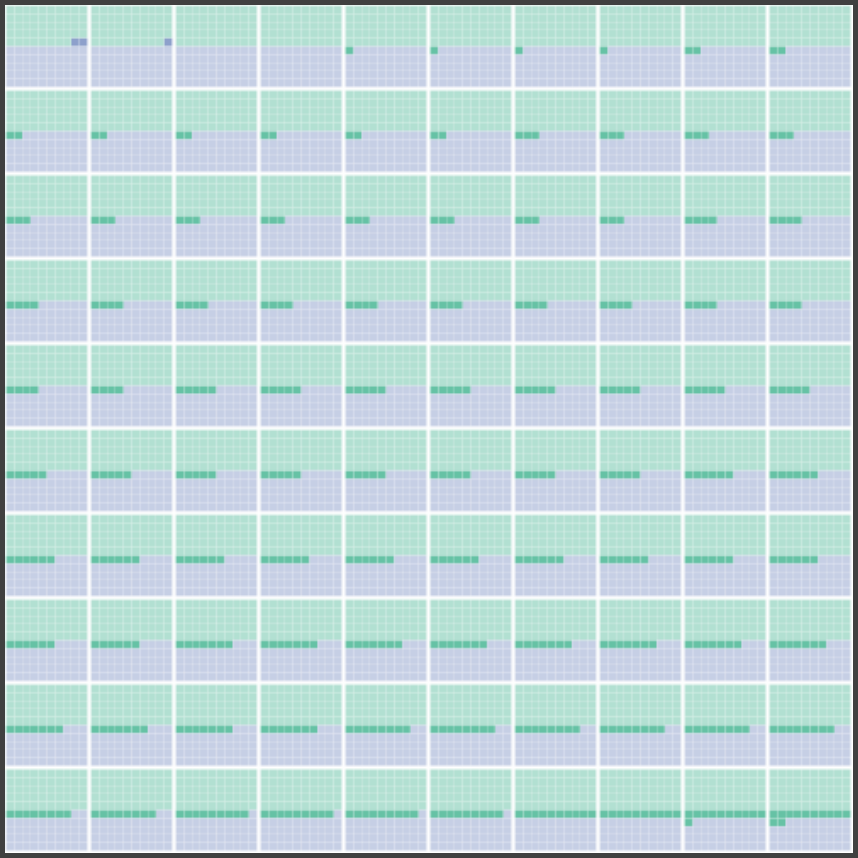

*poll of 100 people, margin of error +/-5

Pangloss Plot

Romney is ahead of Obama in the polls, with 51% of the likely voters*

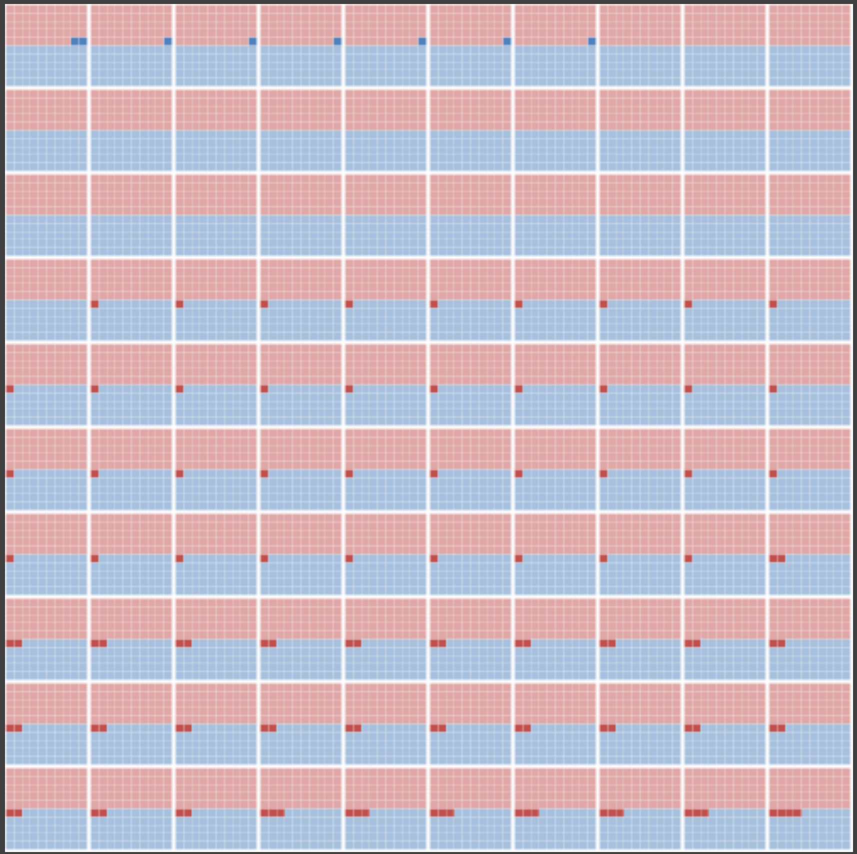

*poll of 3,117 people, margin of error +/-2

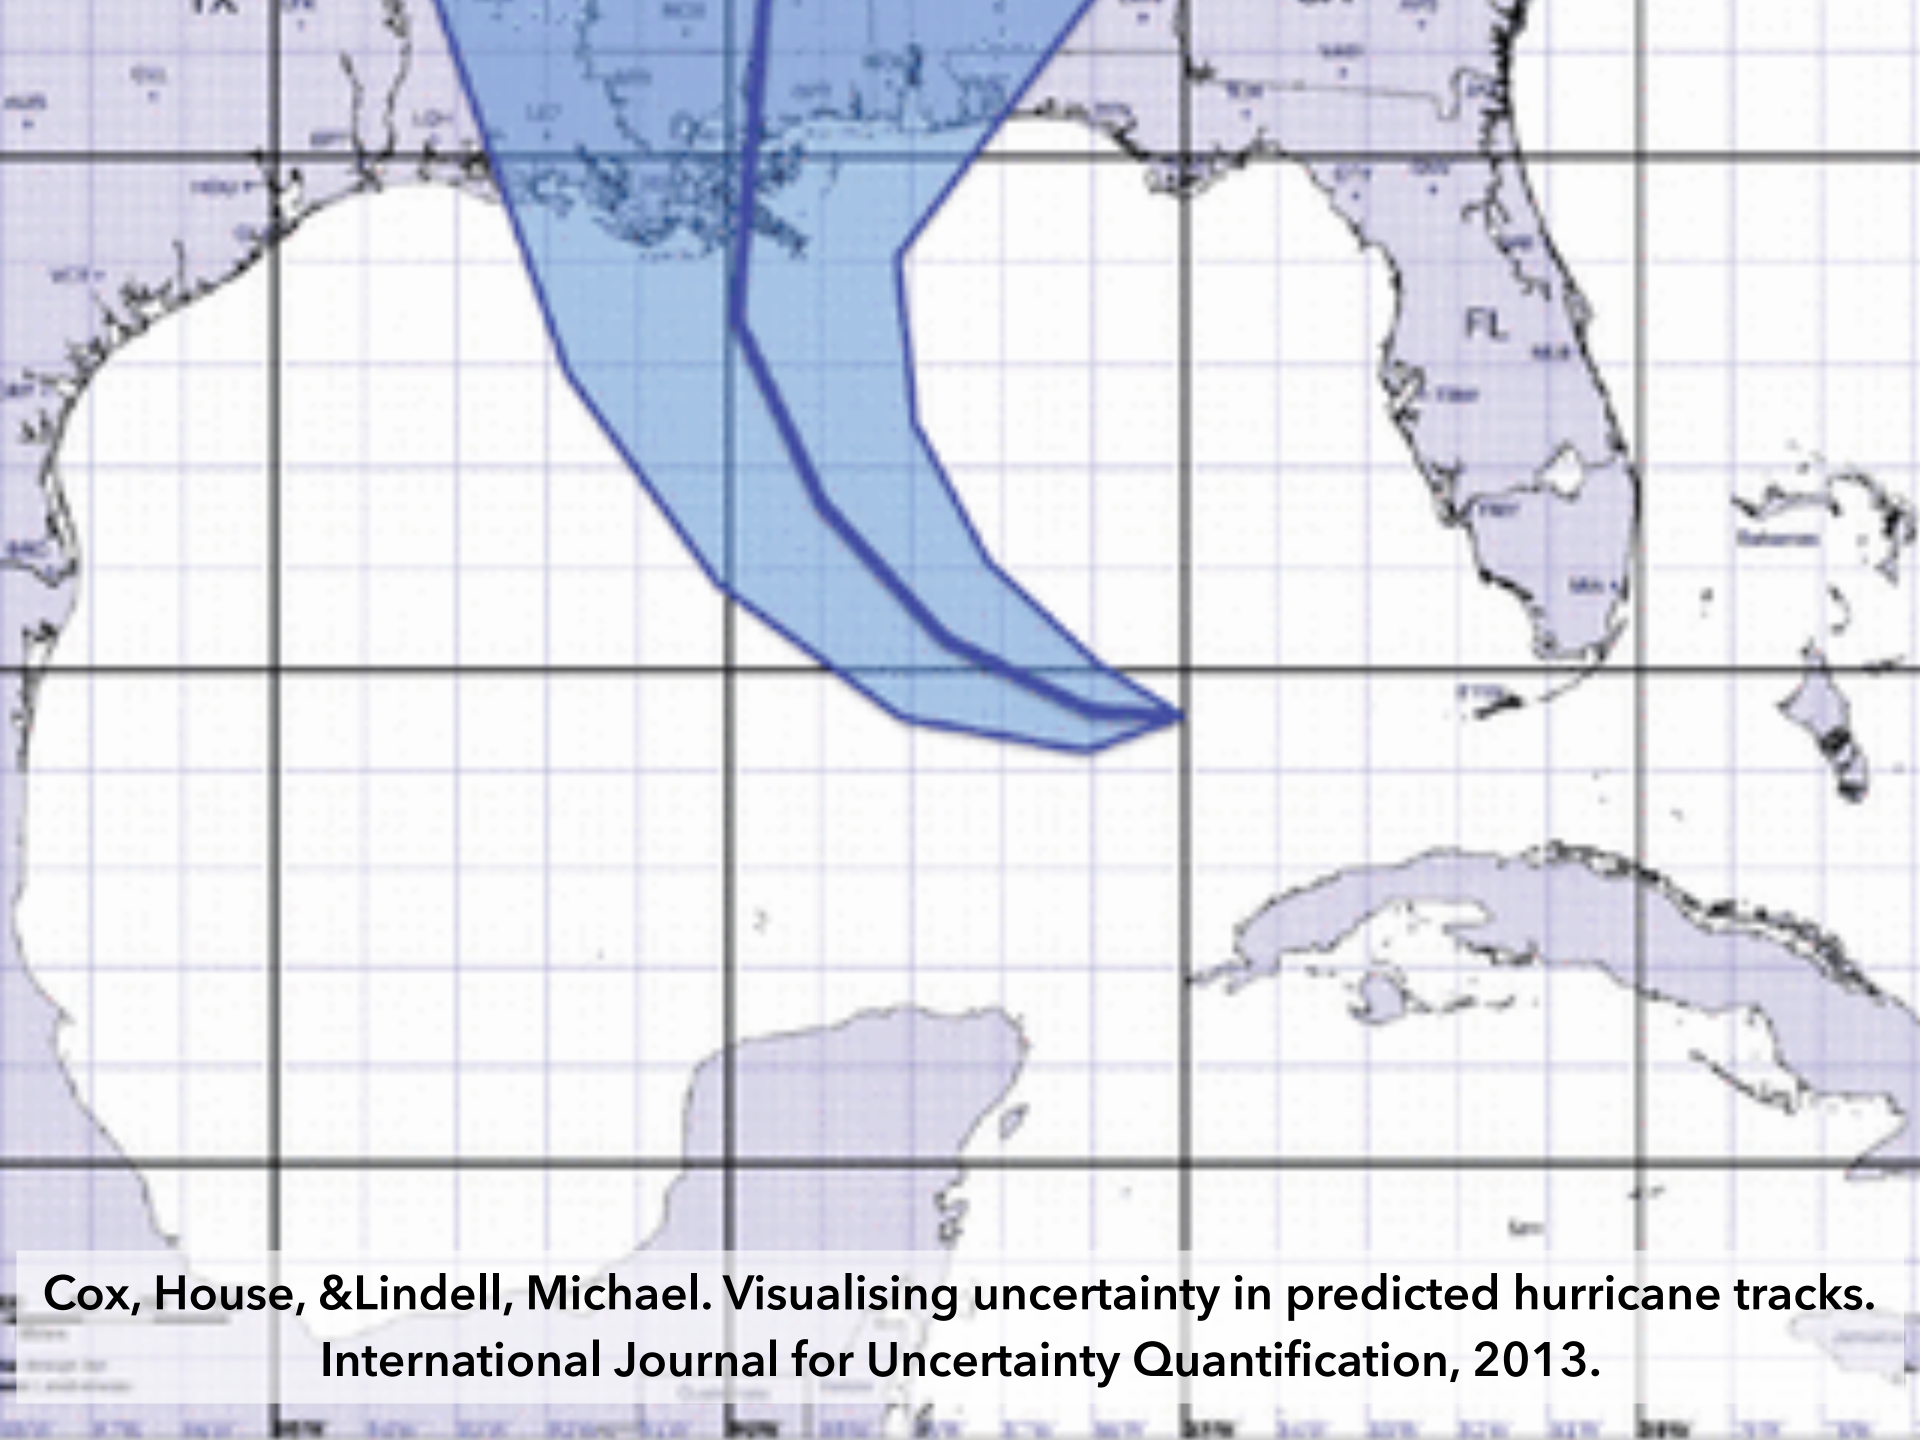

Cox, House, &Lindell, Michael. Visualising uncertainty in predicted hurricane tracks.
International Journal for Uncertainty Quantification, 2013.

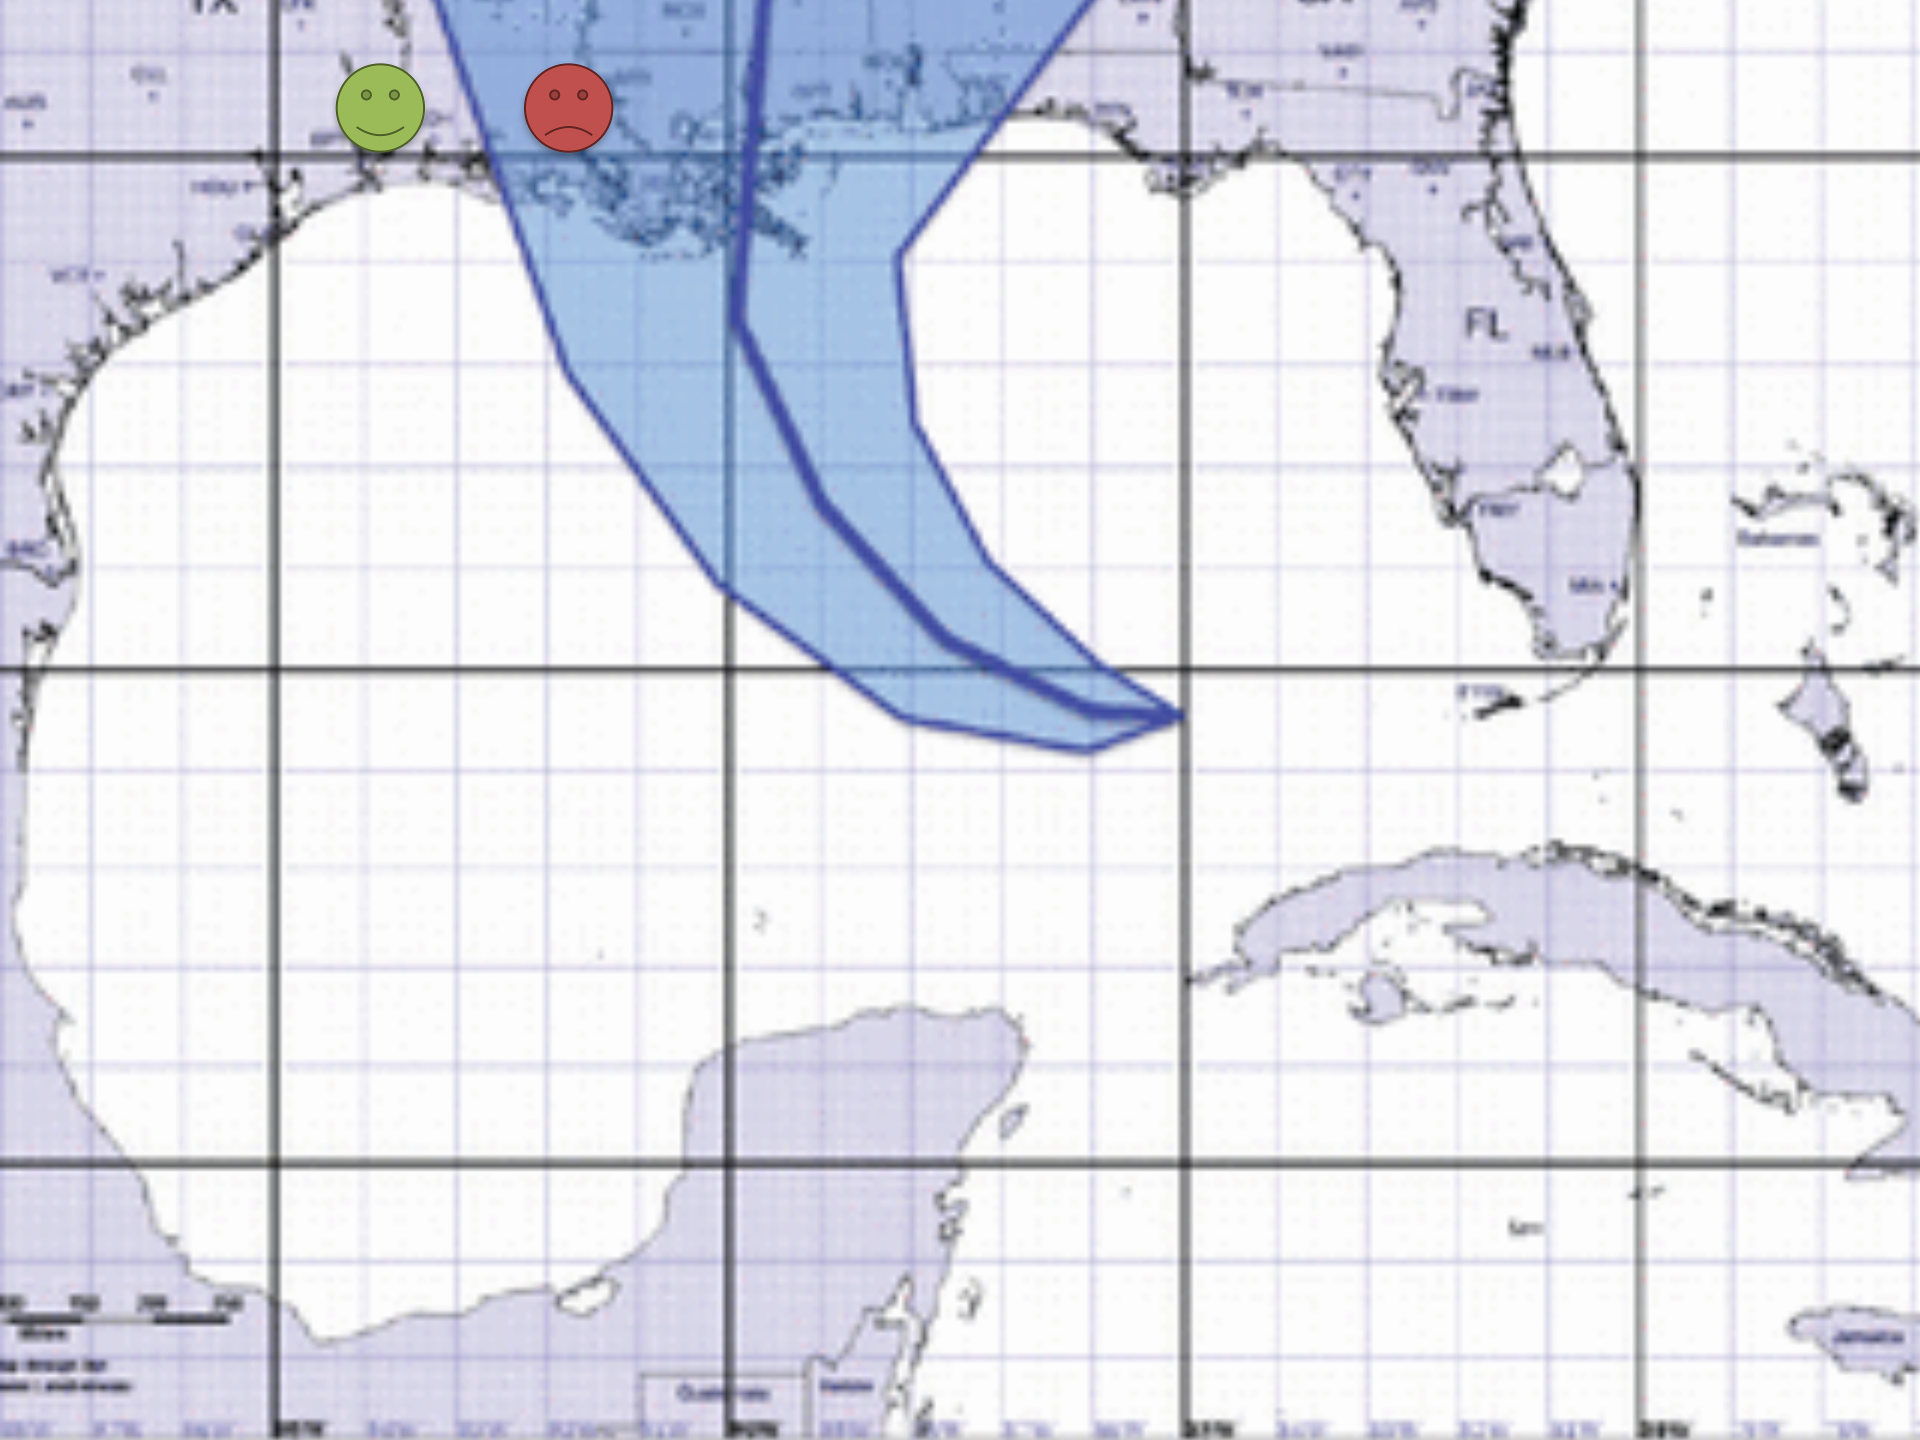
FL

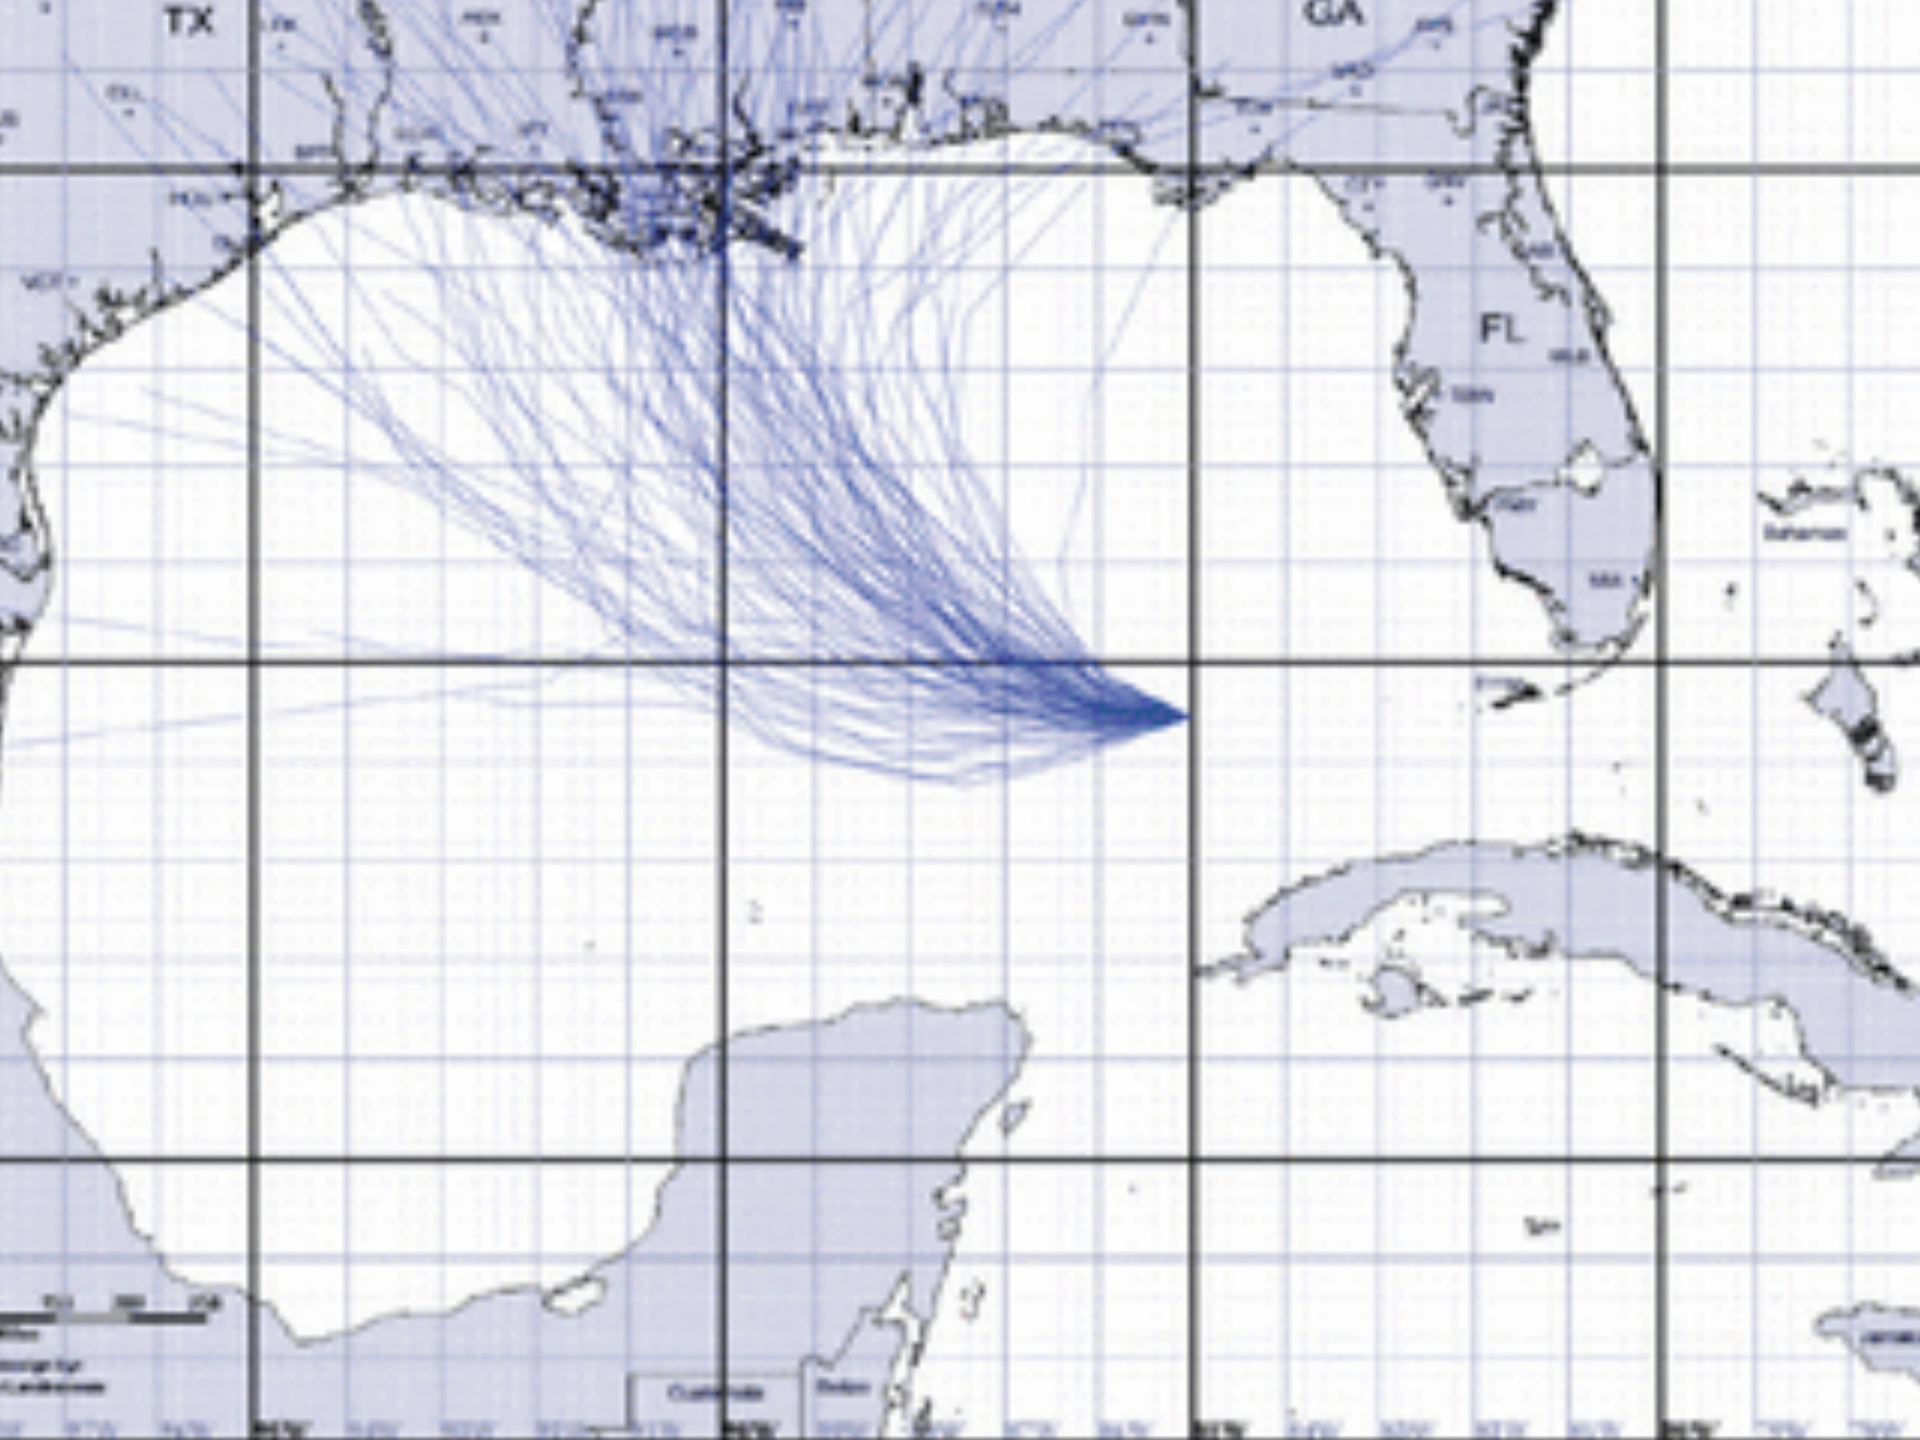
TX
GA
FL

Model Visualization

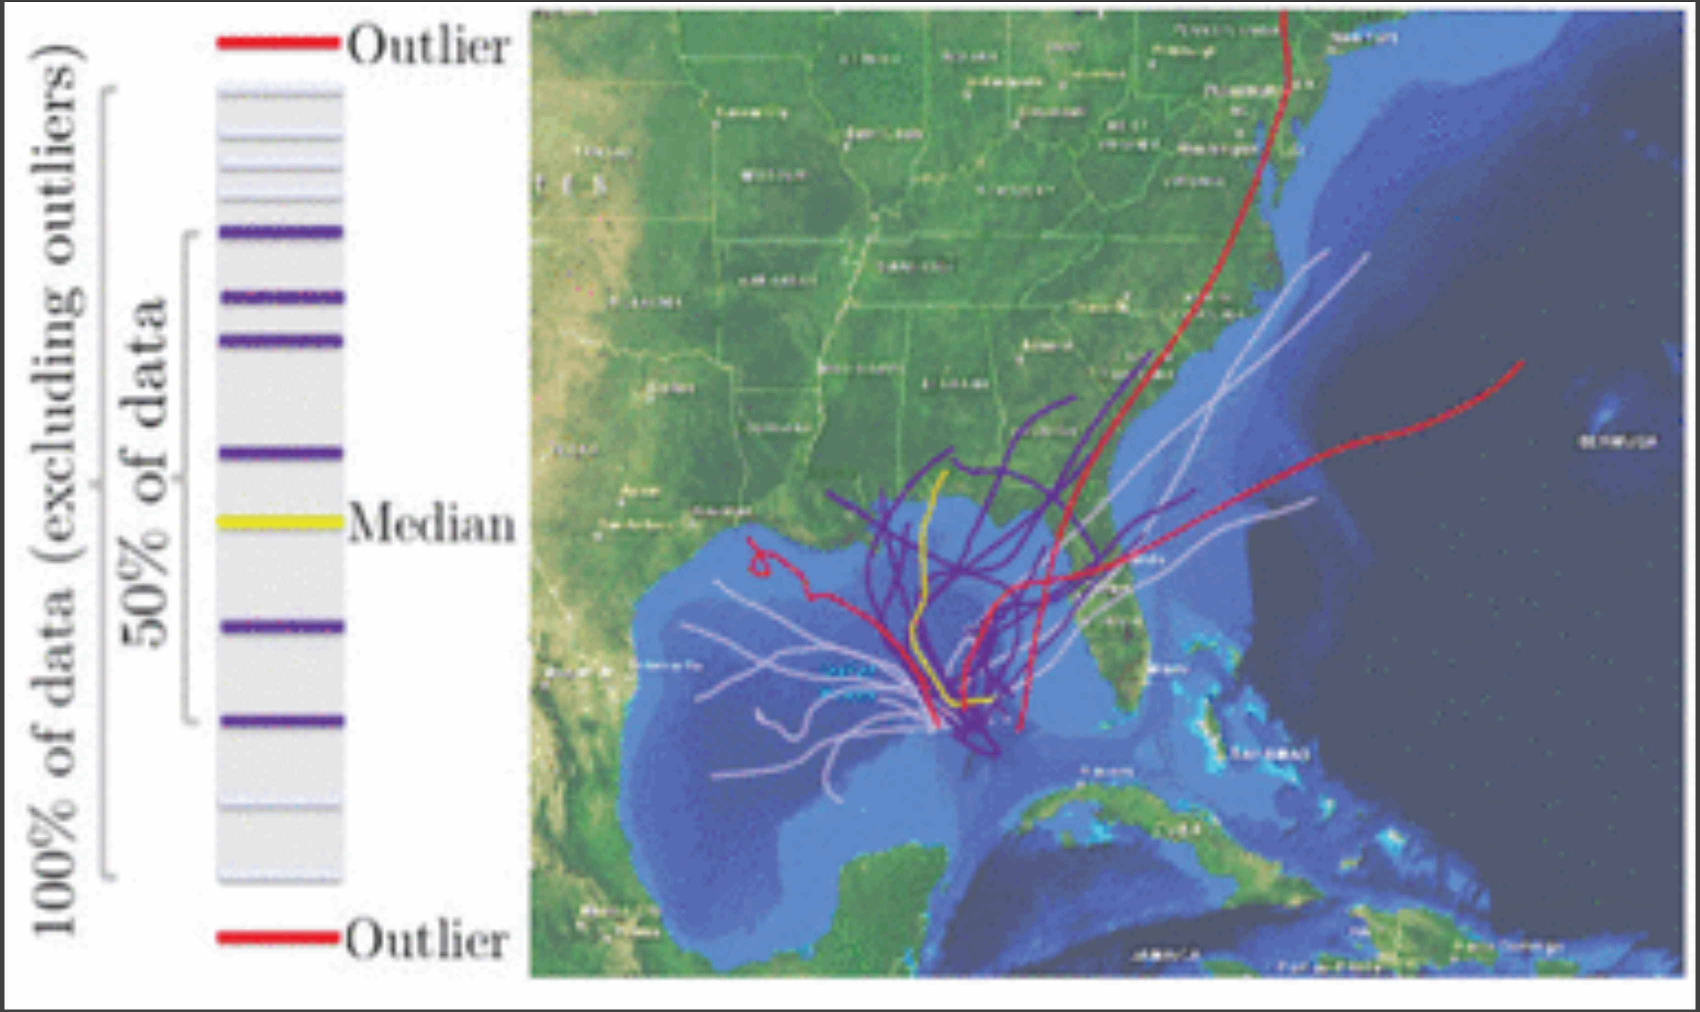

M. Mirzargar, R. Whitaker and R. Kirby. Curve Boxplot: Generalization of Boxplot for Ensembles of Curves. IEEE VIS 2014.

Model Visualization

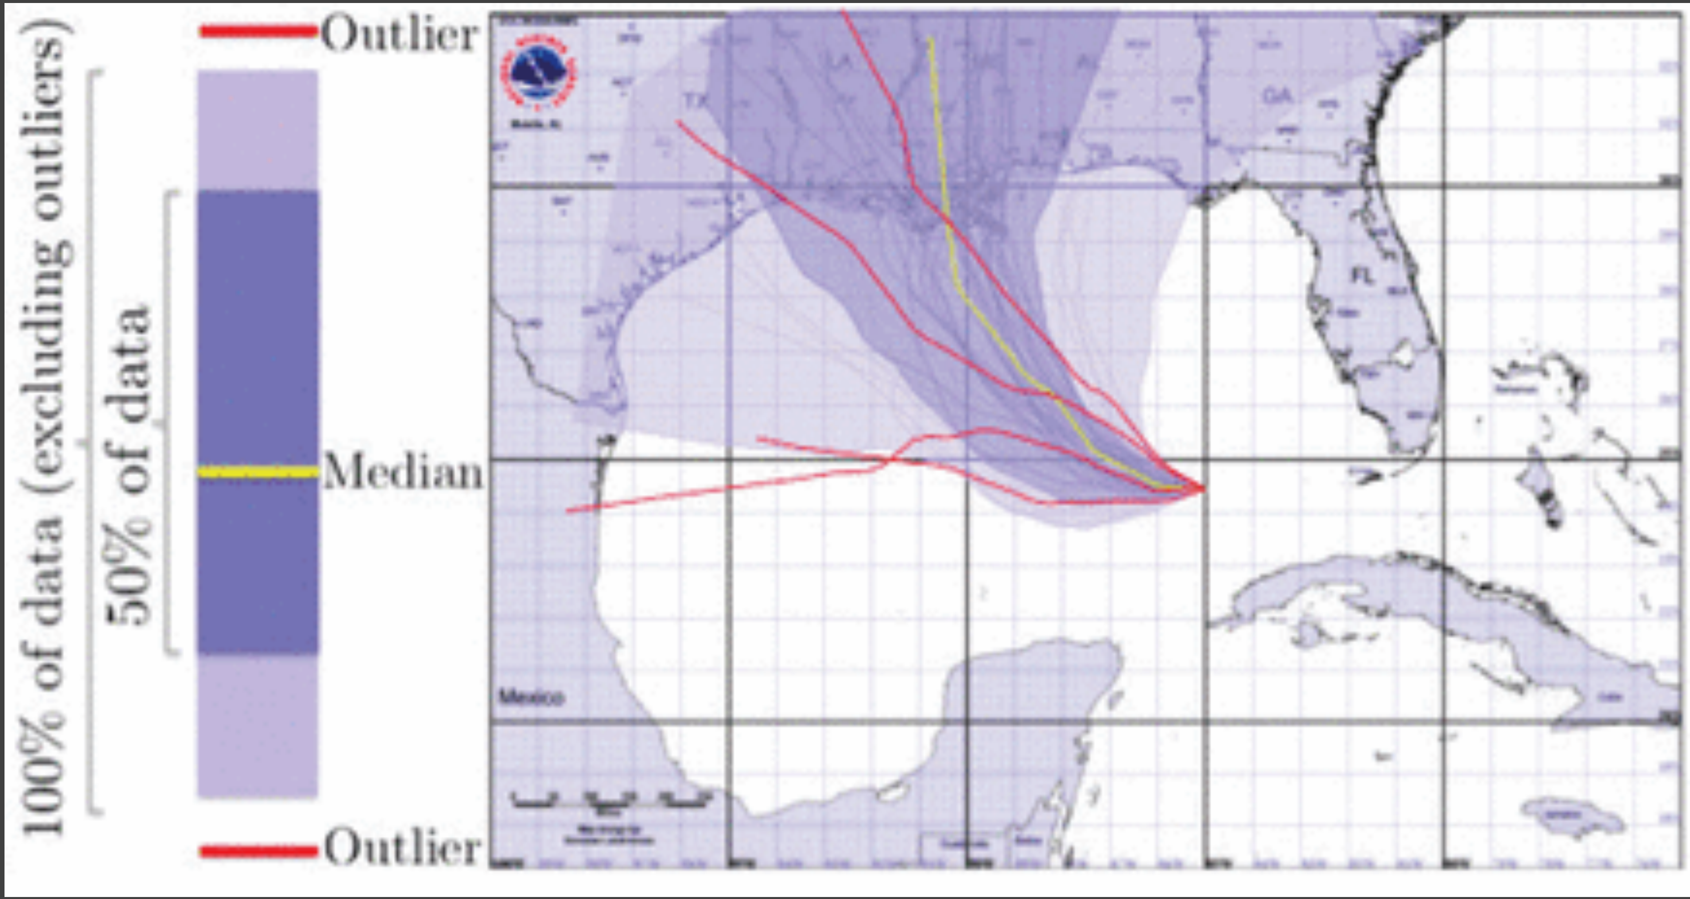

M. Mirzargar, R. Whitaker and R. Kirby. Curve Boxplot: Generalization of Boxplot for Ensembles of Curves. IEEE VIS 2014.

Life Expectancy

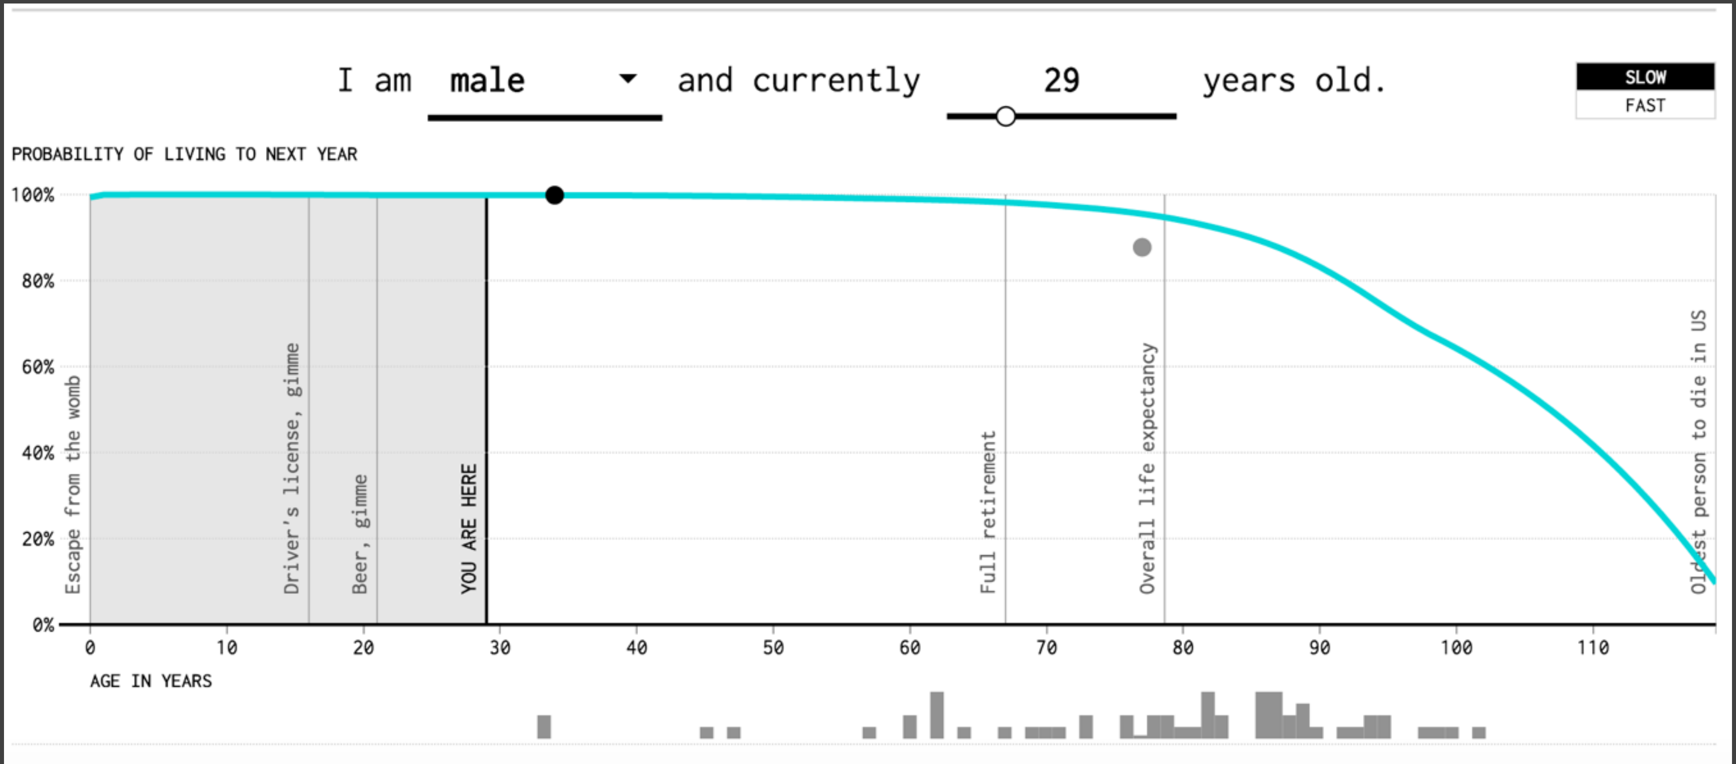

Gun Deaths

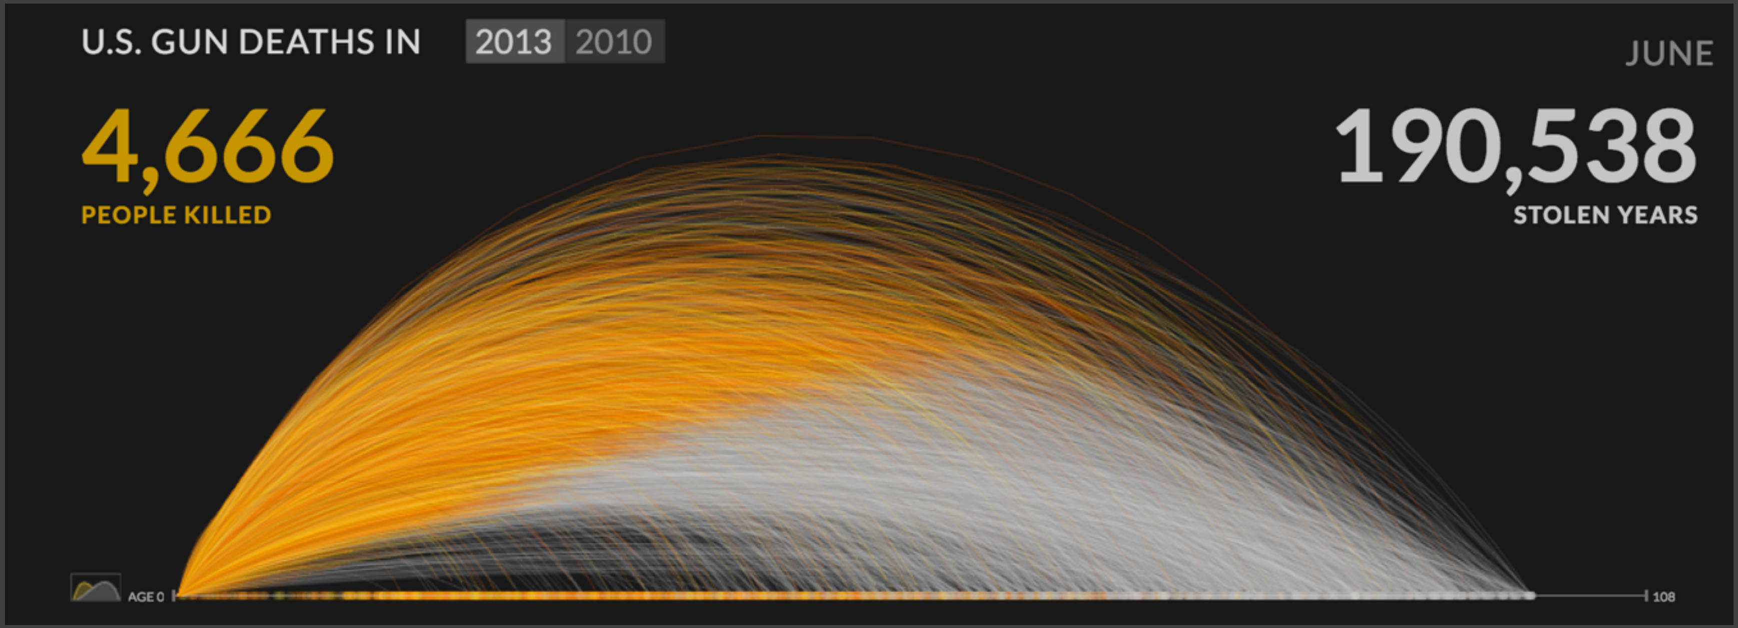
U.S. GUN DEATHS IN
2013
2010
JUNE
4,666
PEOPLE KILLED
190,538
STOLEN YEARS
AGE 0
108

Model Visualization

Building models is necessary to quantify uncertainty

It is important to communicate the variability in model outcomes

Dynamic displays can help communicate complex models

Cognitive Biases

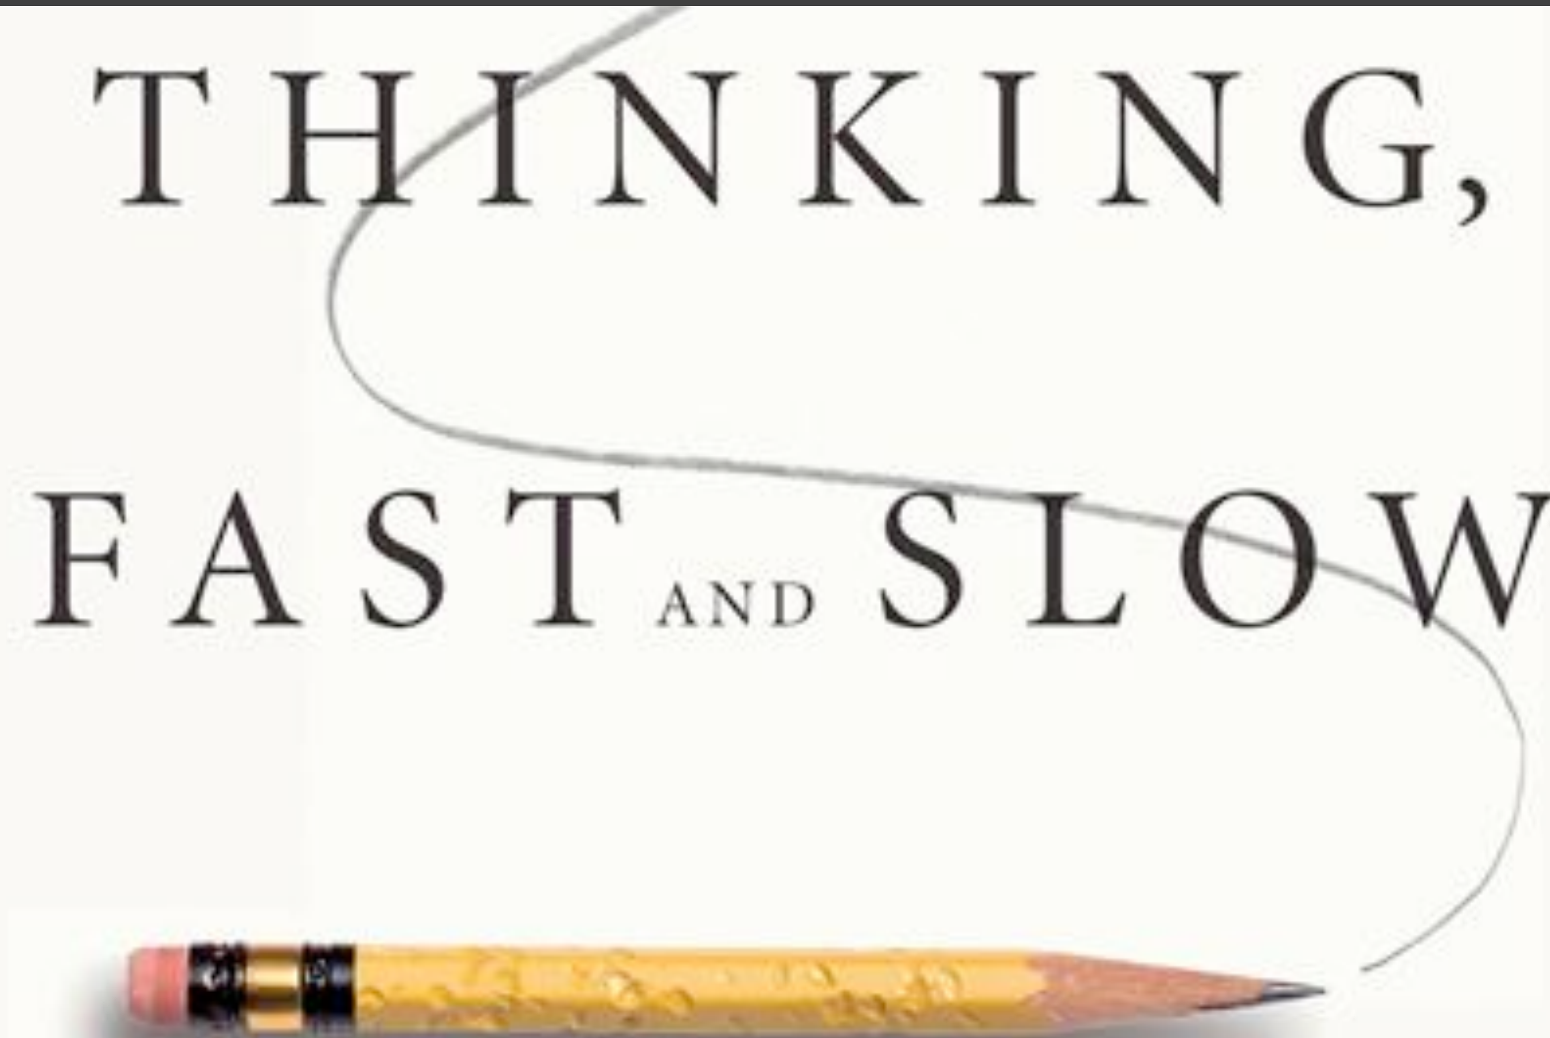

“A large pharmaceutical company has recently developed a new drug to boost peoples’ immune function. It reports that trials it conducted demonstrated a drop of forty percent (from eighty seven to forty seven percent) in occurrence of the common cold. It intends to market the new drug as soon as next winter, following FDA approval.”

Persuaded by Nothing

"A large pharmaceutical company has recently developed a new drug to boost peoples' immune function. It reports that trials it conducted demonstrated a drop of forty percent (from eighty seven to forty seven percent) in occurrence of the common cold. It intends to market the new drug as soon as next winter, following FDA approval."

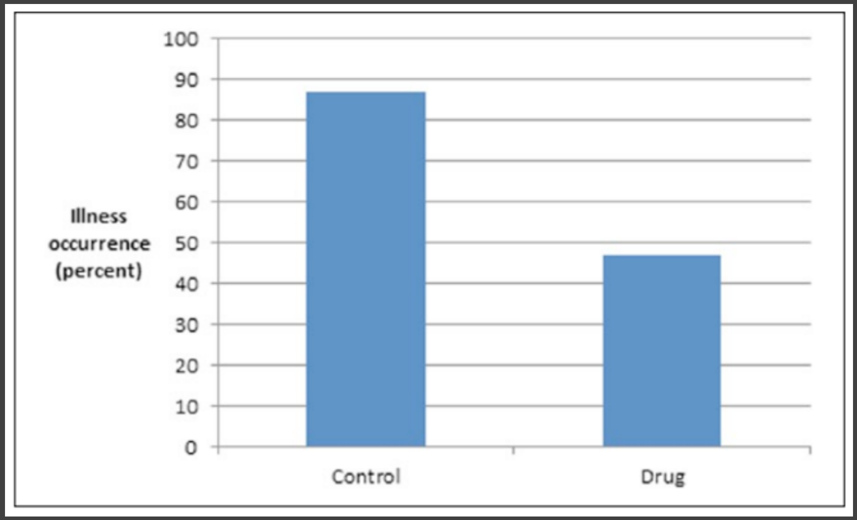

Tal & Wansink. (2016) Blinded with science: Trivial graphs and formulas increase ad persuasiveness and belief in product efficacy.

Which Stock To Buy?

Company A

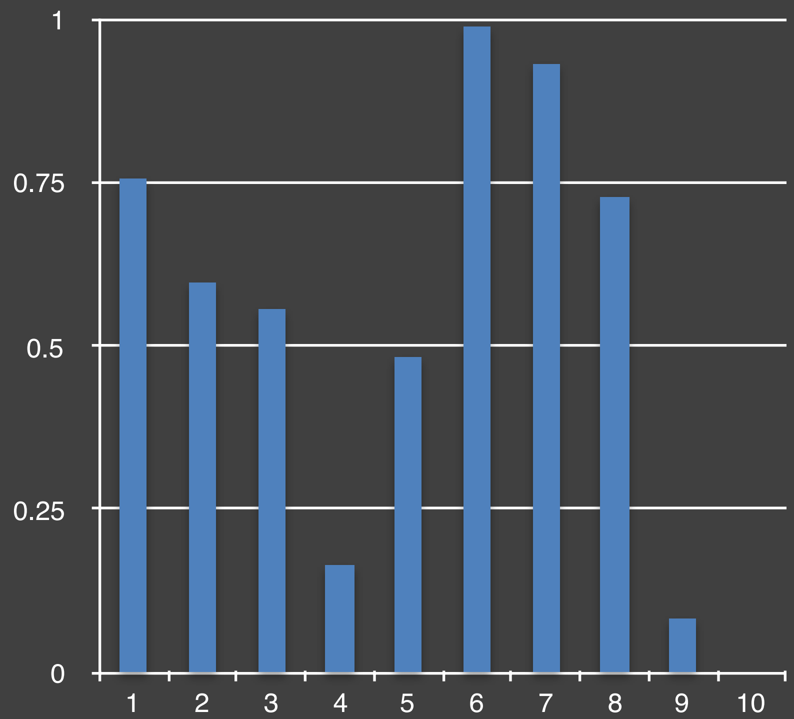

Company B

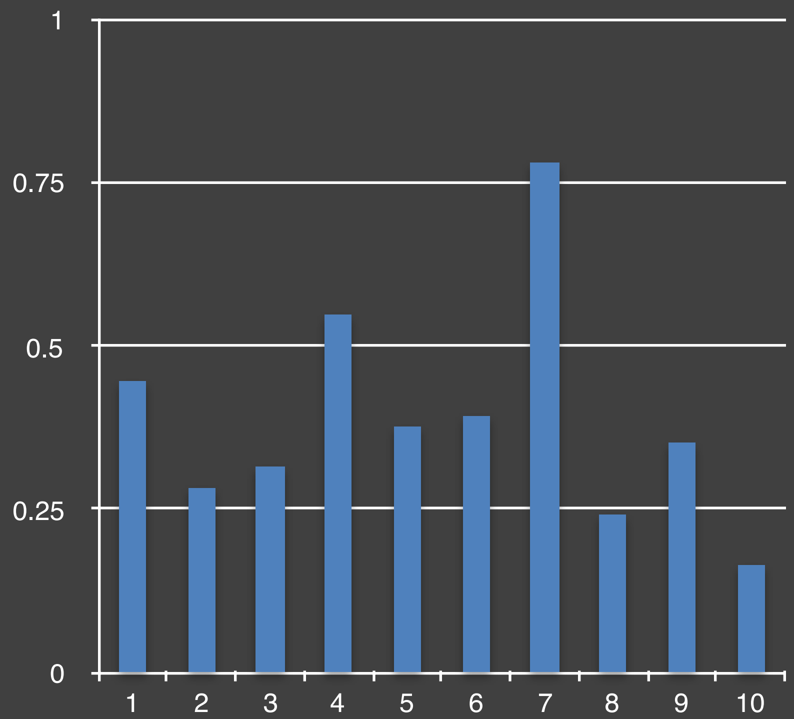

Neither!

Pareidolia

Jobs Reports

Graphical Inference: Visual Lineups

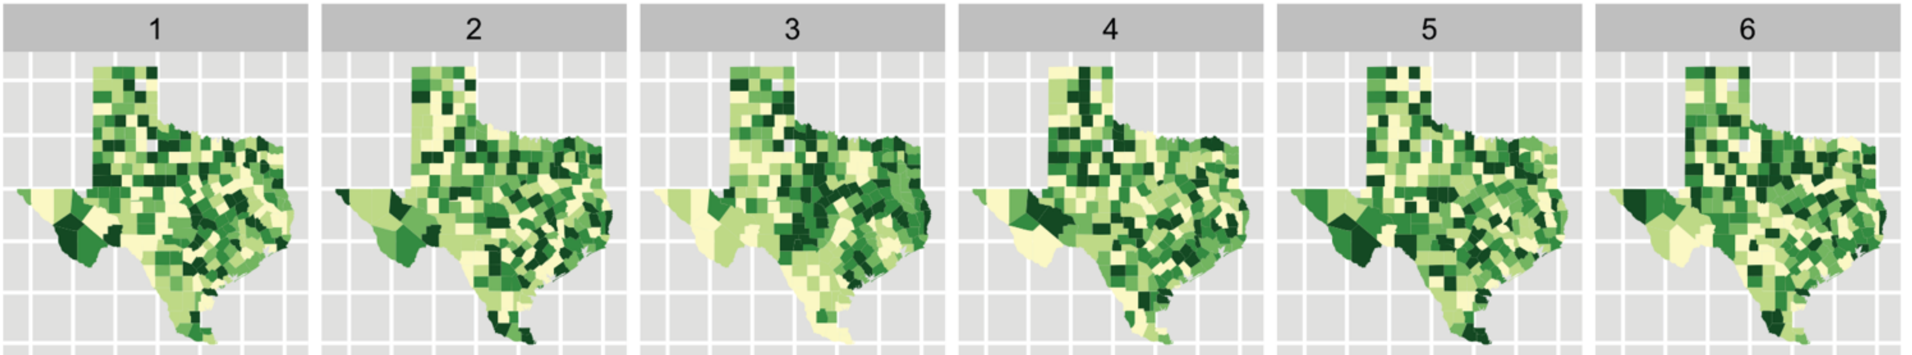

Choropleth maps of cancer deaths in Texas.

One plot shows a real data set. The others are simulated under the null hypothesis of spatial independence.

Can you spot the real data? If so, you have some evidence of spatial dependence in the data.

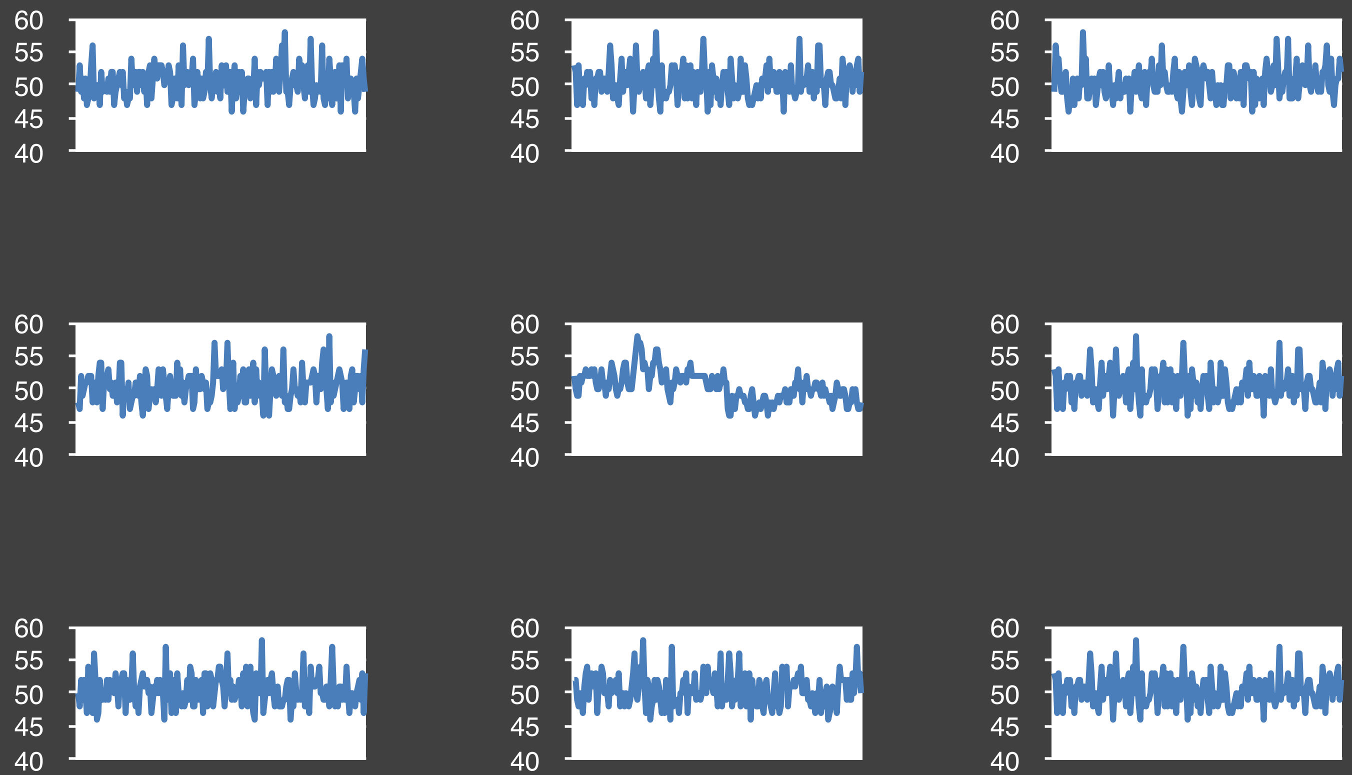

60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40
60
55
50
45
40

Visual Lineups

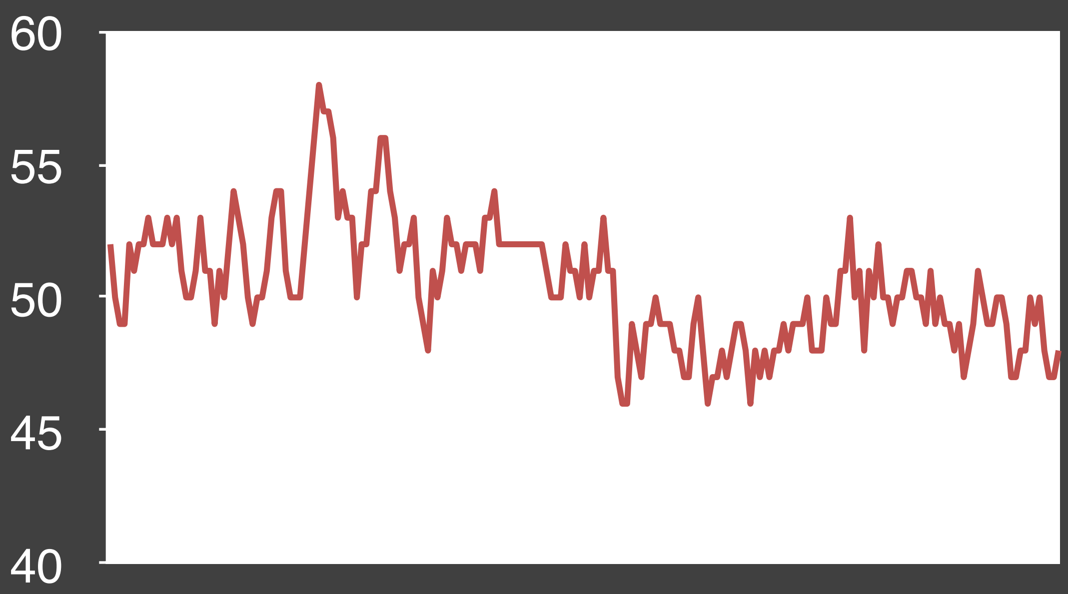

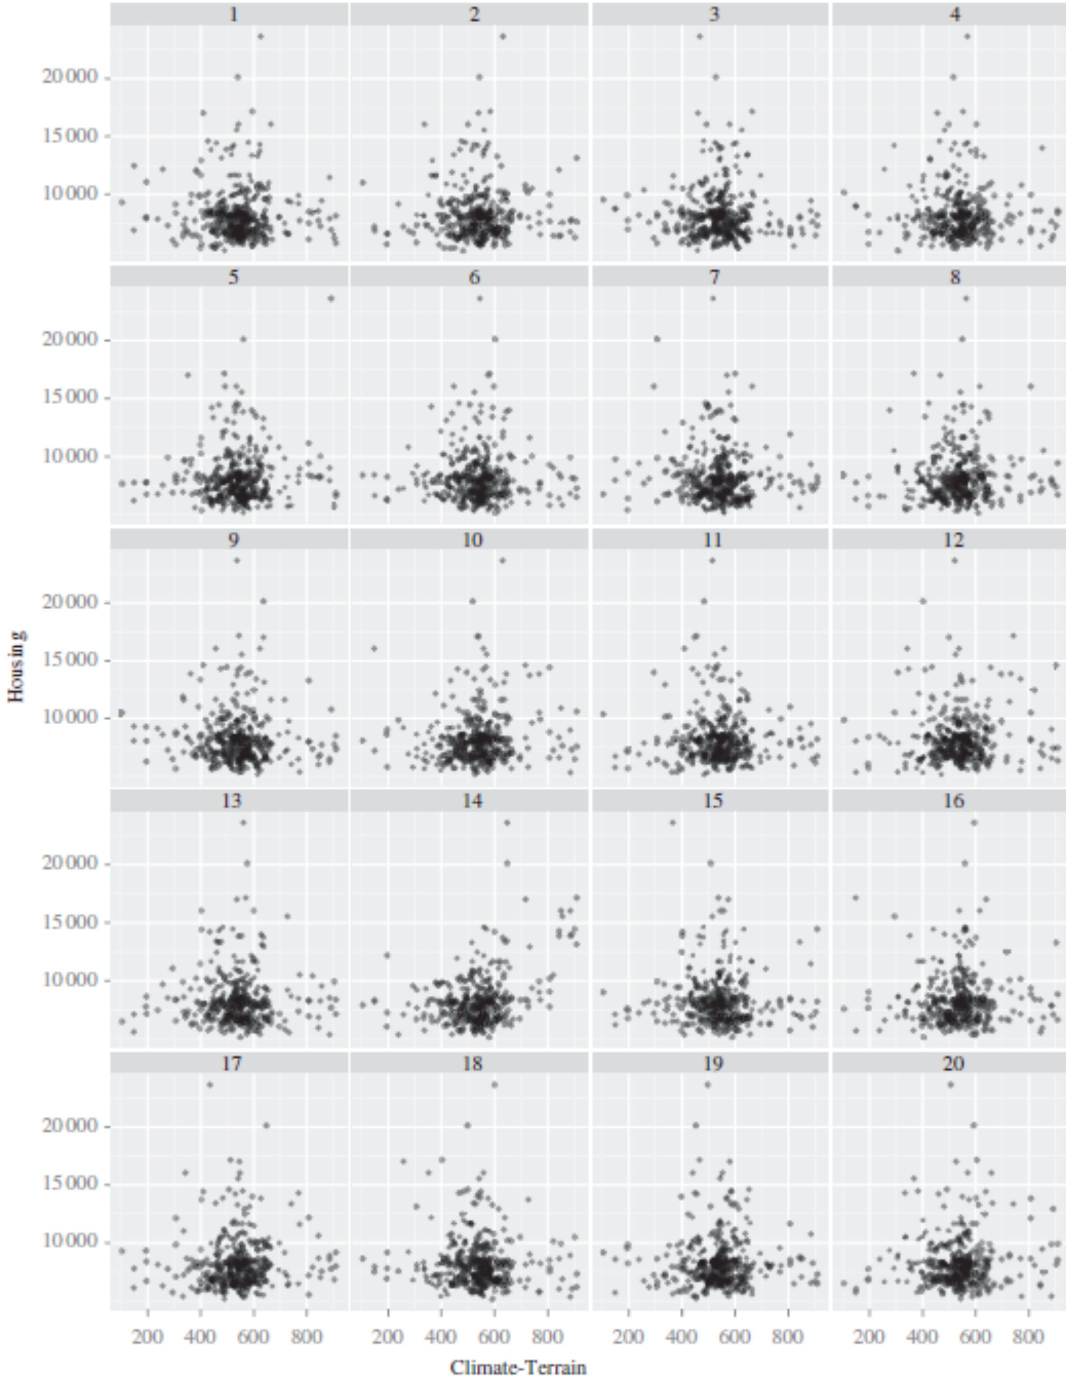

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
Housing
20 000
15 000
10 000
200
400
600
800
Climate-Terrain

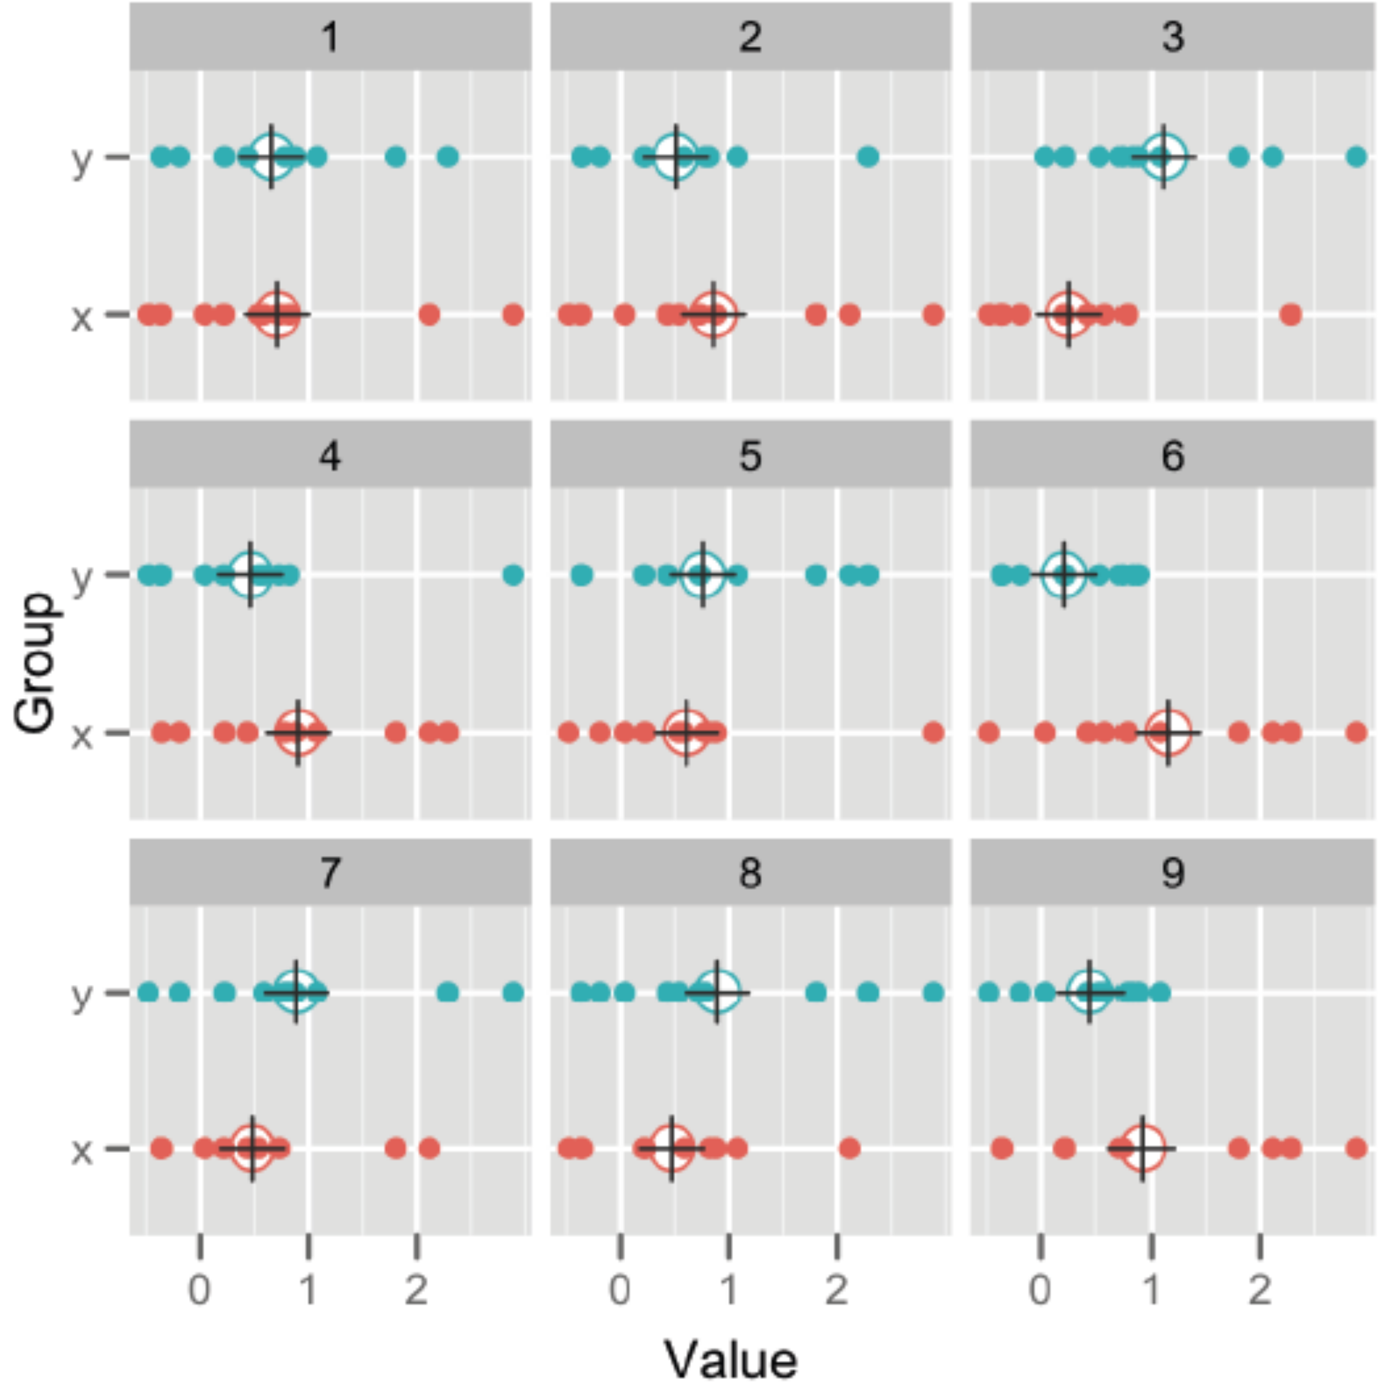

1
2
3
4
5
6
7
8
9
y
x
Group
0
1
2
Value

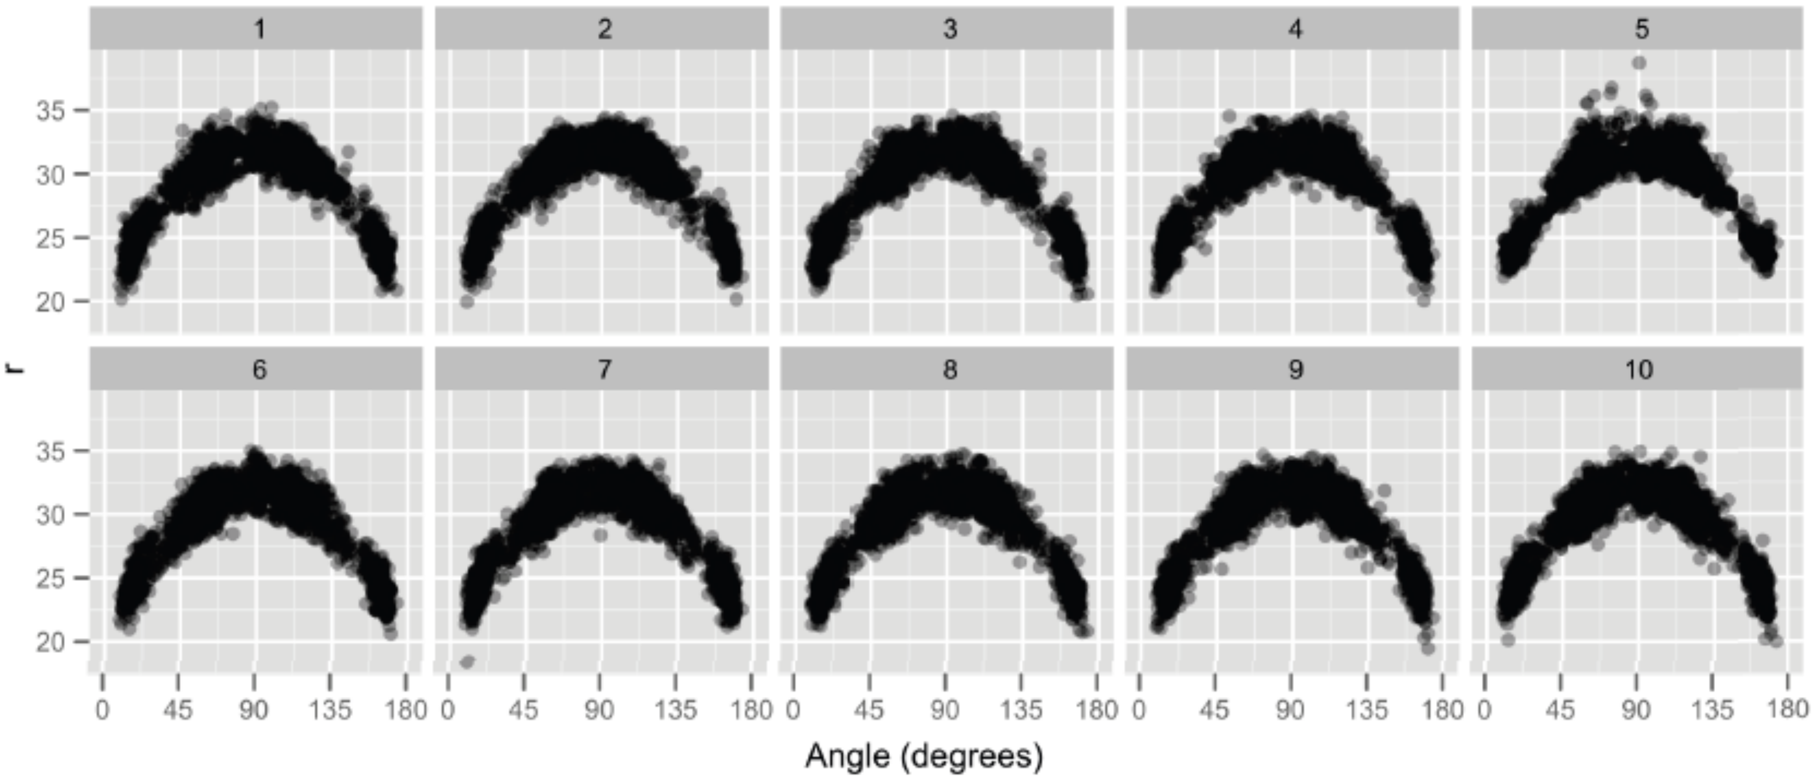

Distance vs. angle for 3 point shots by the LA Lakers.

One plot is the real data. The others are generated according to a null hypothesis of quadratic relationship.

Negative Results

People tend to analyze patterns and make decisions, even if there is “nothing to see.”

Negative or null results can correspond to weak and non-robust visual patterns across a model space.

Base Rate Fallacy

1% of 40 year old women have breast cancer

The probability a mammogram will detect breast cancer is 80%

The probability of a false positive is 10%.

If a 40 year old woman gets a positive result, what is the probability she has breast cancer?

Bayes' Law

$$P(A|B) = P(B|A)\, P(A) \,/\, P(B)$$

Bayes' Law

$$P(A|B) = P(B|A)\ P(A) / P(B)$$

$$P(\text{Cancer} \mid +\text{Test}) = P(+\text{Test}|\text{Cancer})\ P(\text{Cancer}) / P(+\text{Test})$$

Bayes' Law

$$P(A|B) = P(B|A)\, P(A) / P(B)$$

$$P(\text{Cancer} \mid +\text{Test}) = P(+\text{Test}|\text{Cancer})\, P(\text{Cancer}) / P(+\text{Test})$$

$$P(+) = P(+ \mid C)\, P(C) + P(+ \mid {\sim}C)\, P({\sim}C)$$

Bayes' Law

$$P(A|B) = P(B|A)\, P(A) \,/\, P(B)$$

$$P(\text{Cancer} \mid +\text{Test}) = P(+\text{Test}|\text{Cancer})\, P(\text{Cancer}) \,/\, P(+\text{Test})$$

$$P(+) = P(+ \mid C)\, P(C) + P(+ \mid \sim C)\, P(\sim C)$$

$$P(+) = 0.8 * 0.01 + 0.1 * 0.99$$

$$P(+) = 0.107$$

$$P(C \mid +) = 0.8 * 0.01 \,/\, 0.107 \approx \mathbf{0.075}$$

Base Rate Fallacy

Micallef, Dragicevic, & Fekete. (2012) Assessing the Effect of Visualizations on Bayesian Reasoning Through Crowdsourcing

Bayesian Surprise Maps

Can we compensate for some biases?

Idea: Draw attention to what is “surprising” in the data based on models of expectation.

In choropleth maps, we might expect:

- High correlation with population density
- Higher variability due to smaller samples

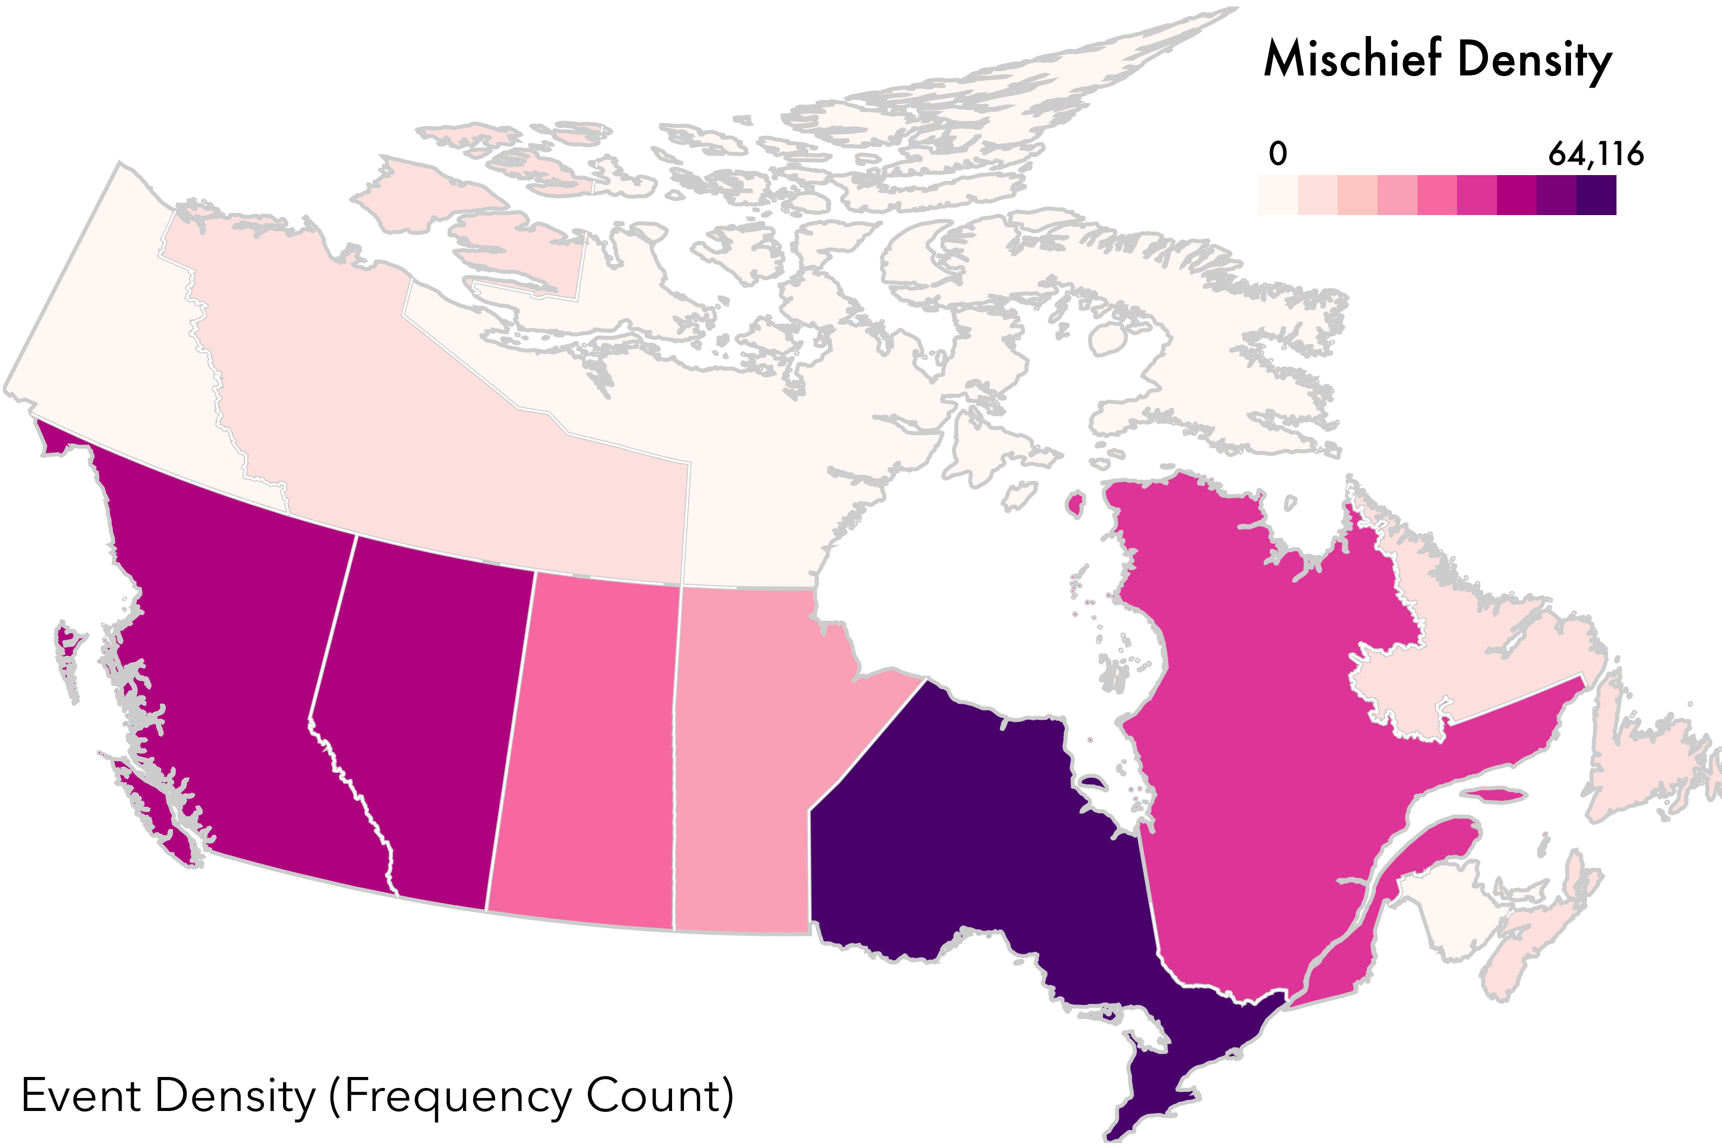

Event Density (Frequency Count)

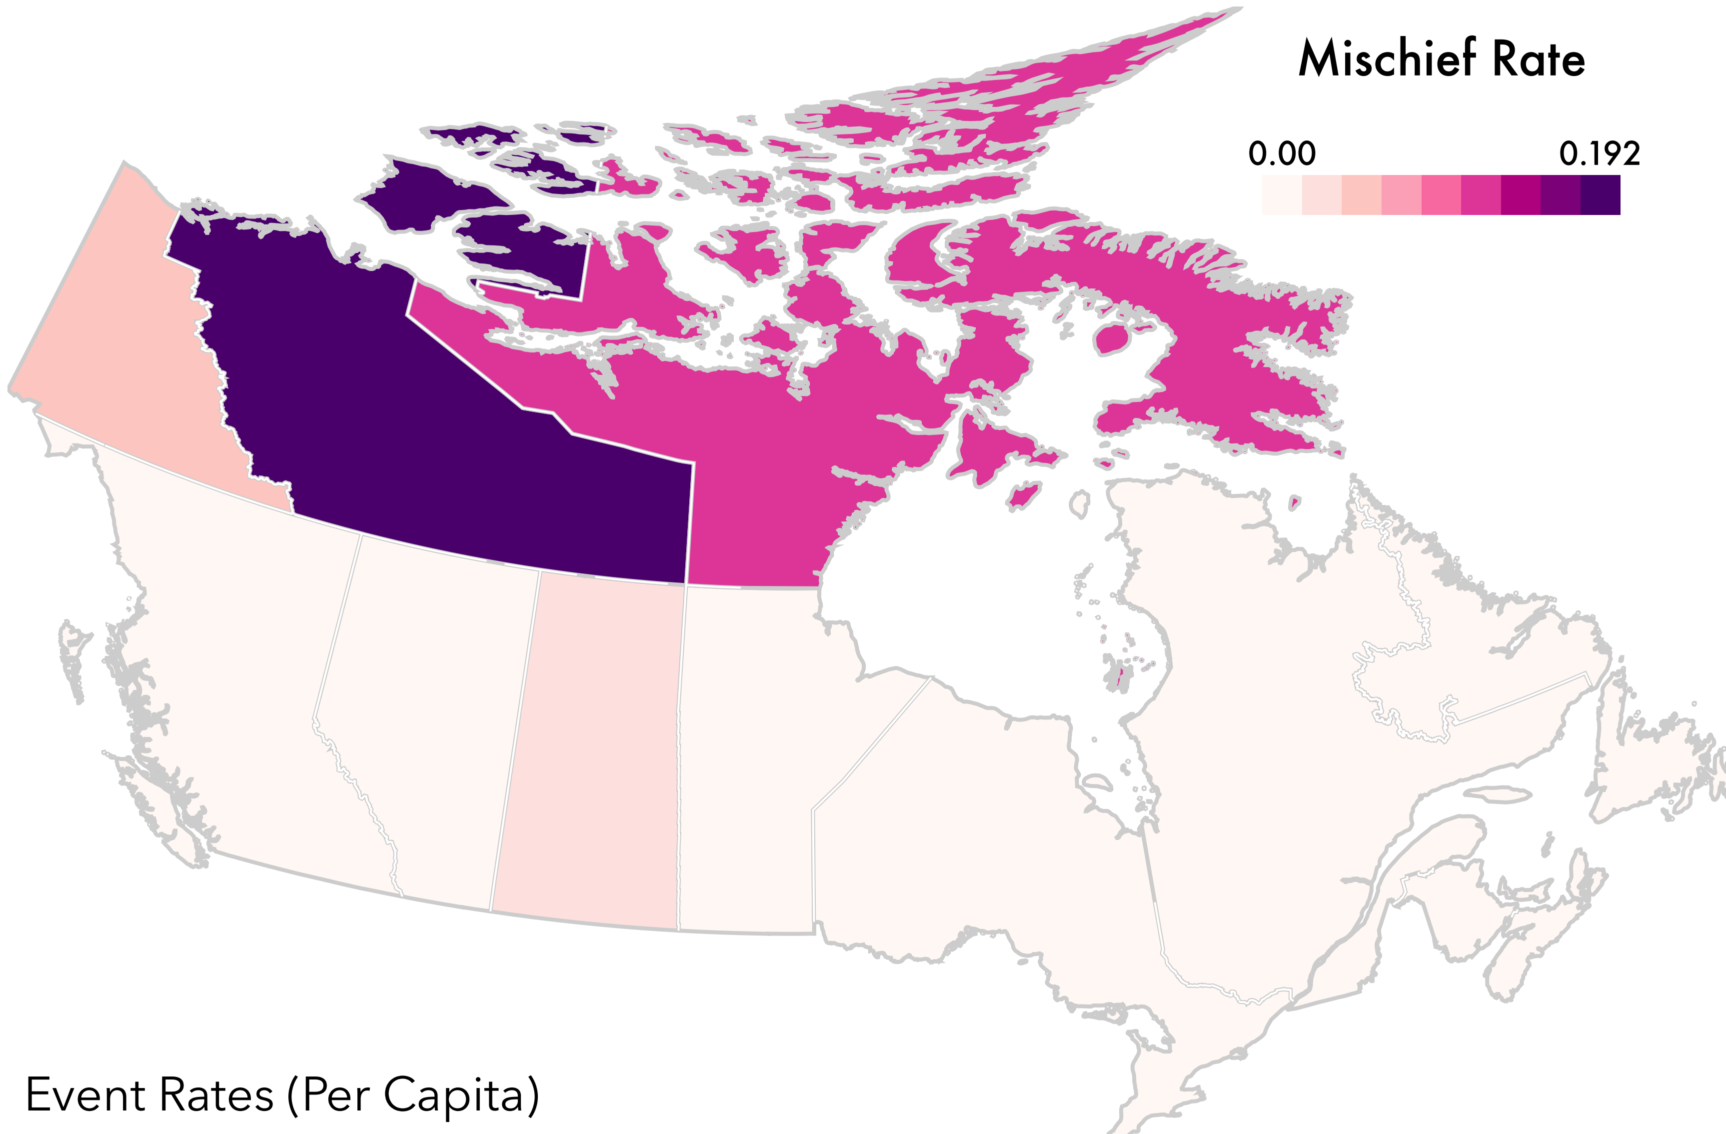
Mischief Rate
0.00
0.192
Event Rates (Per Capita)

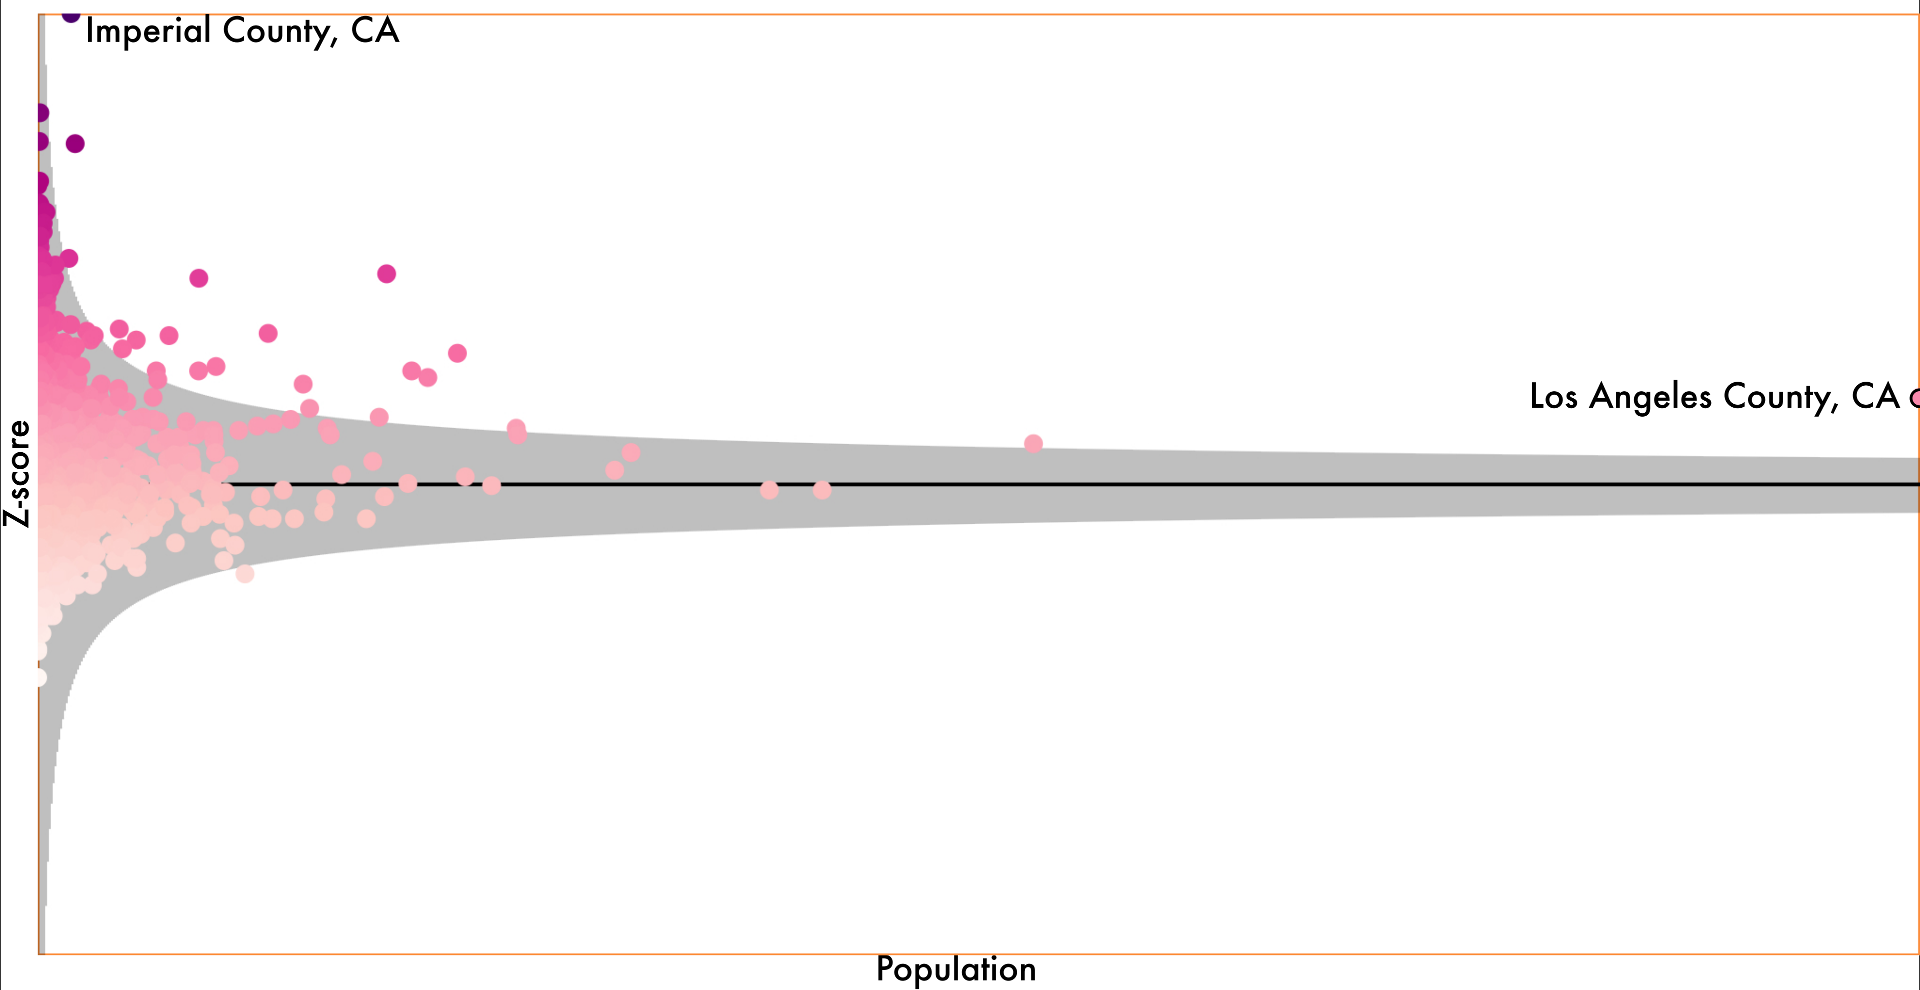

De Moivre Funnel: Effect Size vs. Sample Size

Variability Decreases with Sample Size - Standard Error = $\sigma / \sqrt{N}$

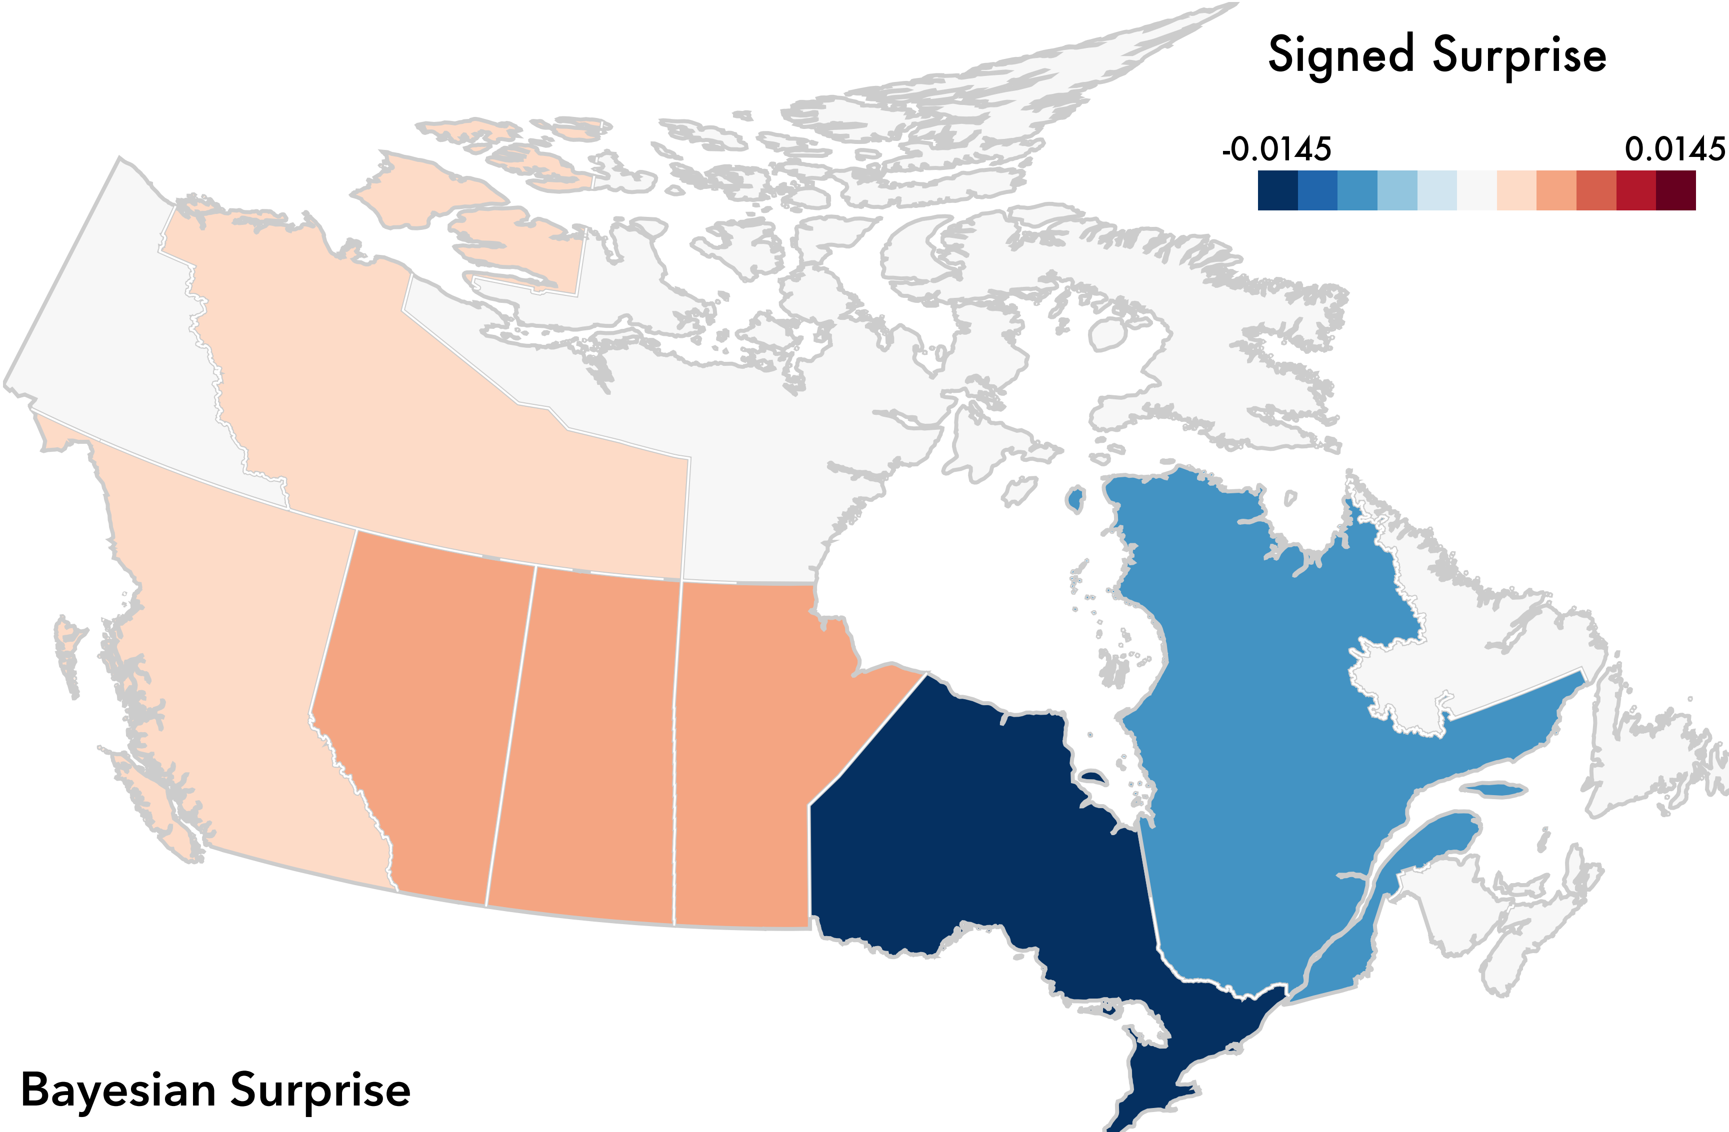

Bayesian Surprise

Modes: Population Density + De Moivre Funnel

Bayesian Surprise Maps

1. Start with set of models & prior beliefs for each.
2. Observe data, update beliefs over models by applying Bayes' Law.
3. Measure discrepancy between prior and posterior models (e.g., using Kullback-Leibler Divergence). The degree of change is our measure of surprise.

Provides a *saliency measure* of what to focus on.

Can be applied over time to data streams.

Conclusion

There are different **types** and **sources** of uncertainty associated with data.

We can **quantify** or **model** our uncertainty.

The visual presentation of uncertainty can **clash** with cognitive and perceptual biases.