CSE 512 - Data Visualization

Uncertainty

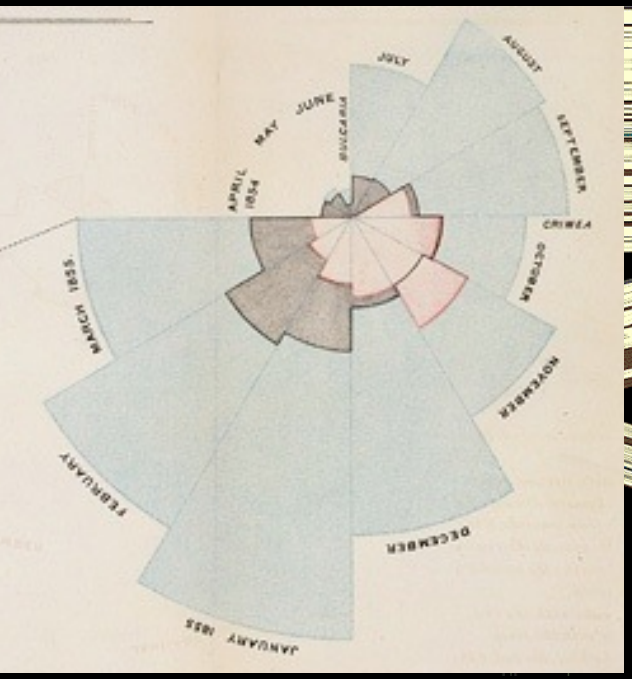

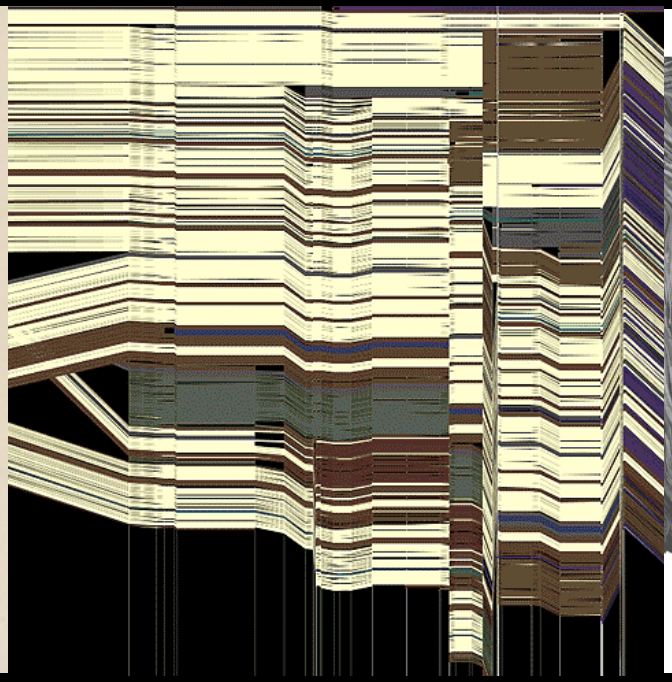

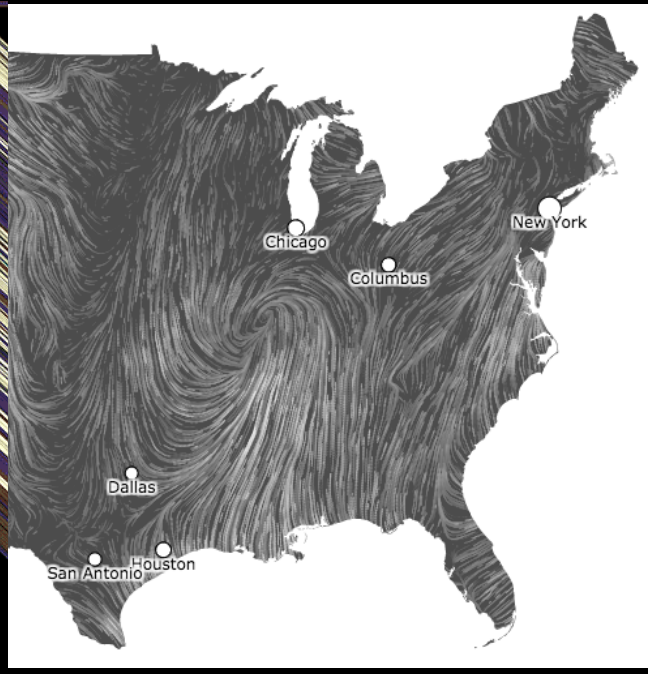

Leilani Battle University of Washington
(with significant material from Michael Correll)

Questions To Answer

What Does Uncertainty Mean?

How Should I Visualize It?

What Can Go Wrong?

What we talk about when we talk about “uncertainty”…

Things “Uncertainty” Can Mean

Doubt

Risk

Variability

Error

Lack of Knowledge

Hedging

…

Uncertainty Vis Pipeline

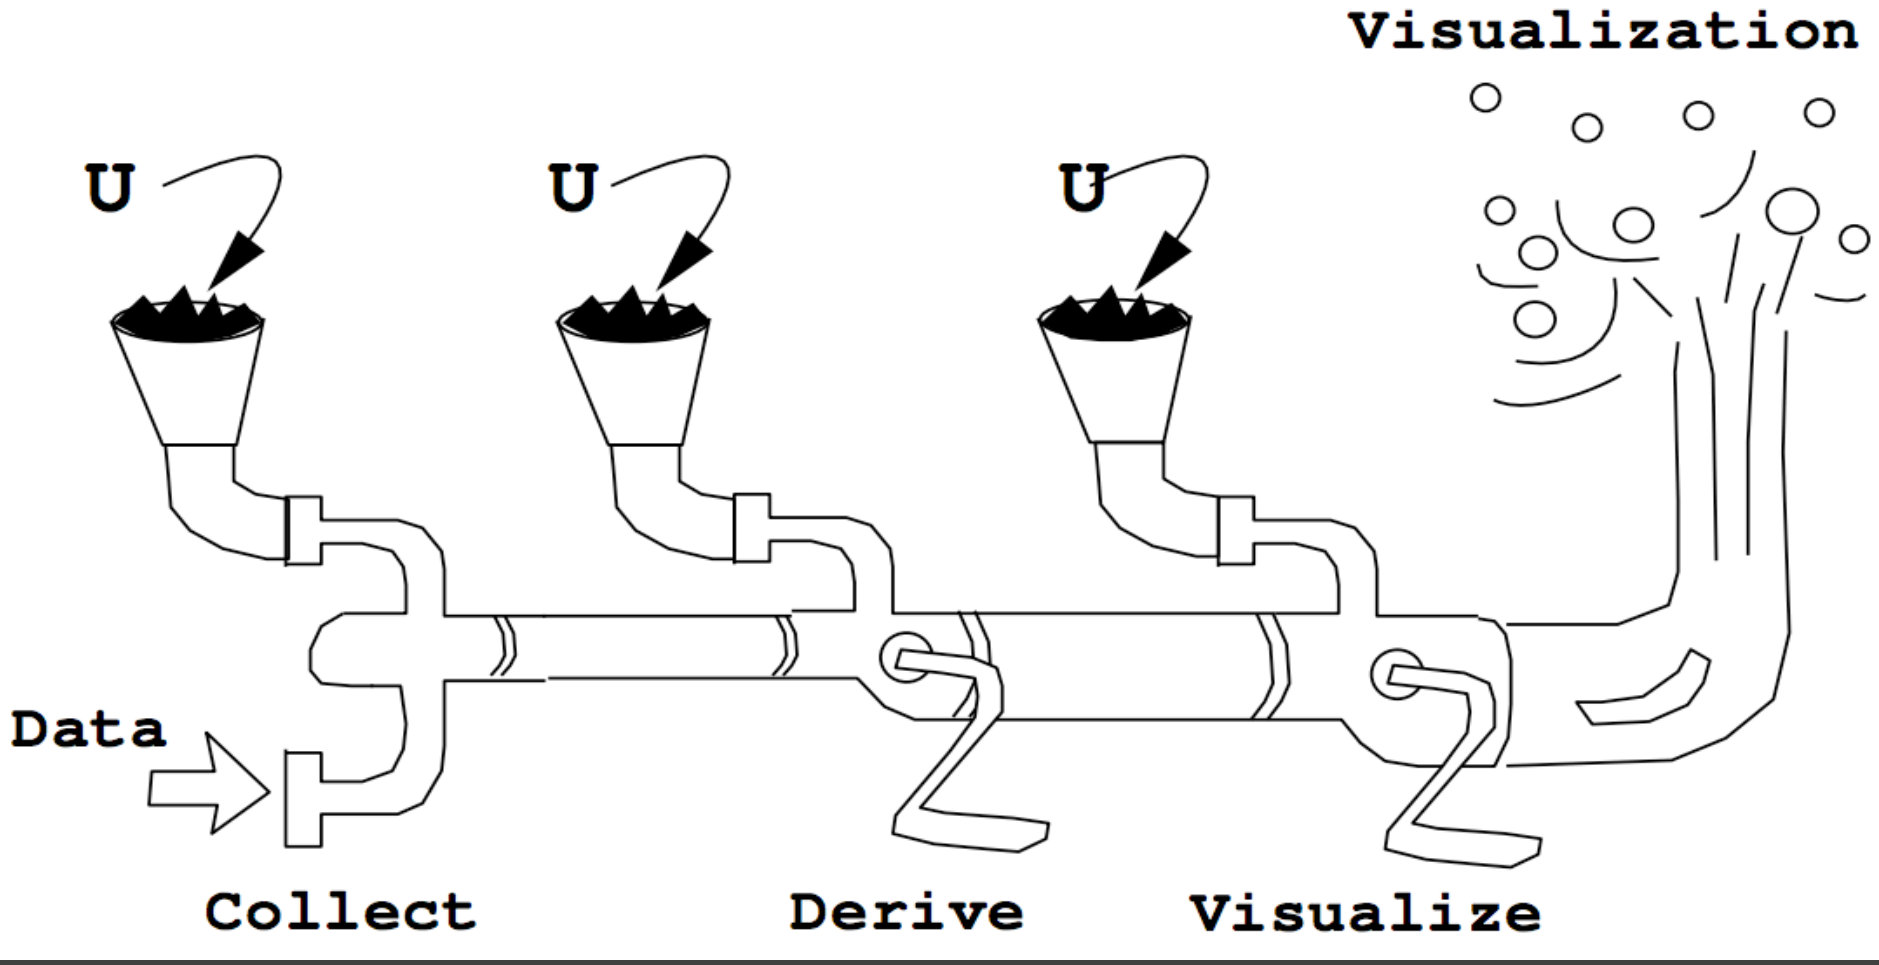

Pang et al. Approaches to Uncertainty Visualization. The Visual Computer, 1997.

A Bar Chart

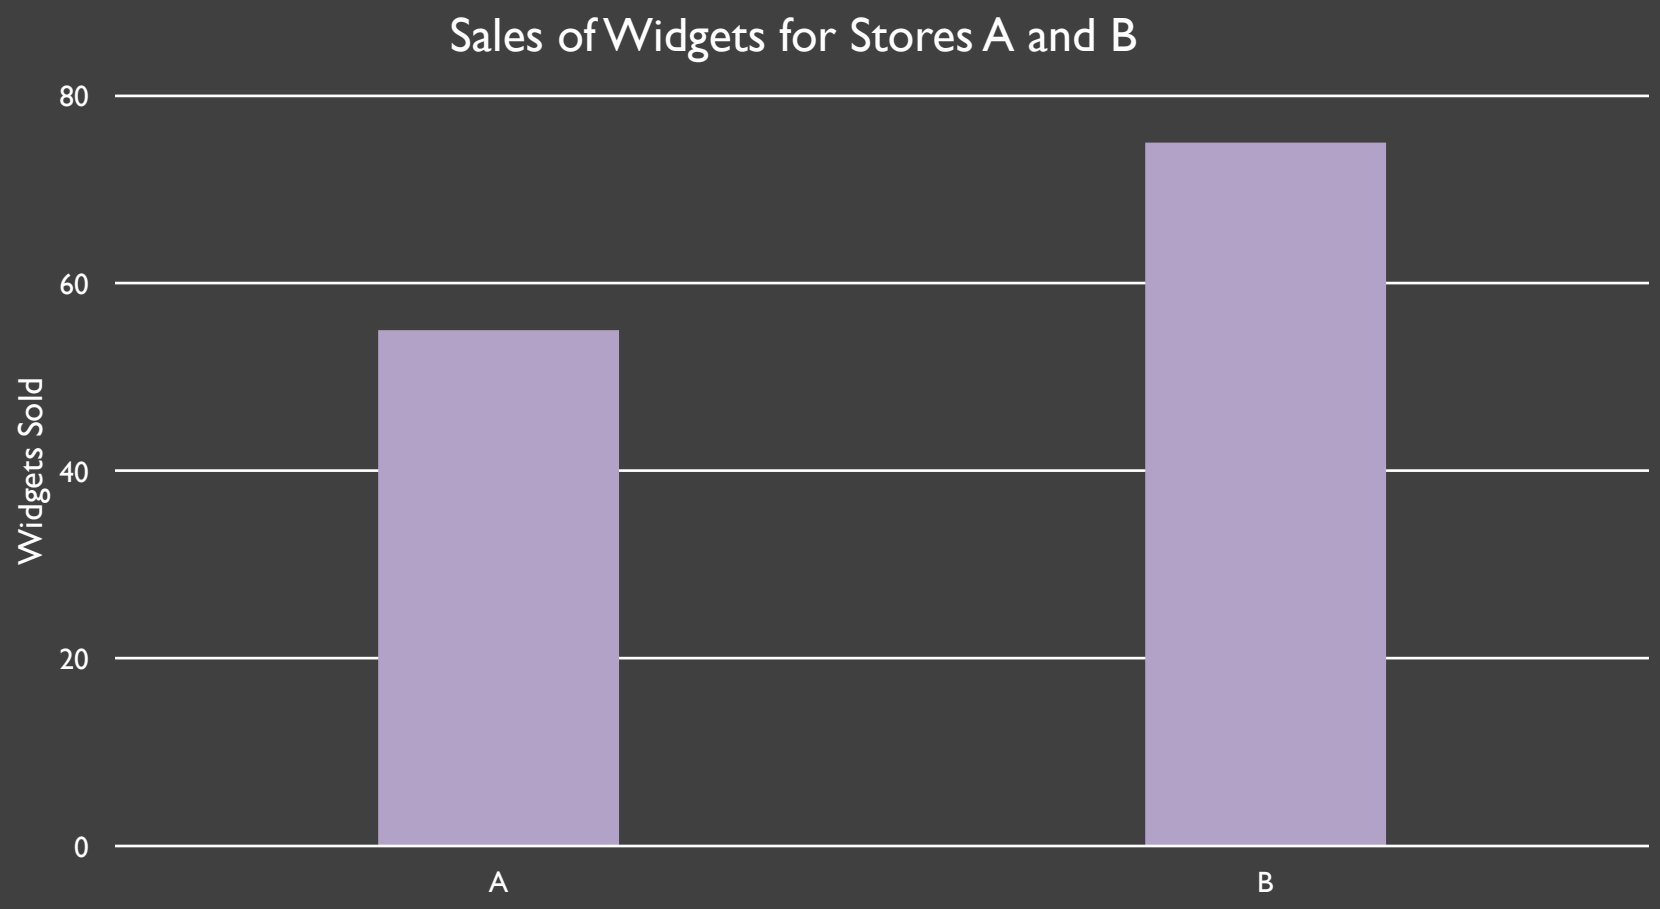
Sales of Widgets for Stores A and B
Widgets Sold
80
60
40
20
0
A
B

Measurement Uncertainty

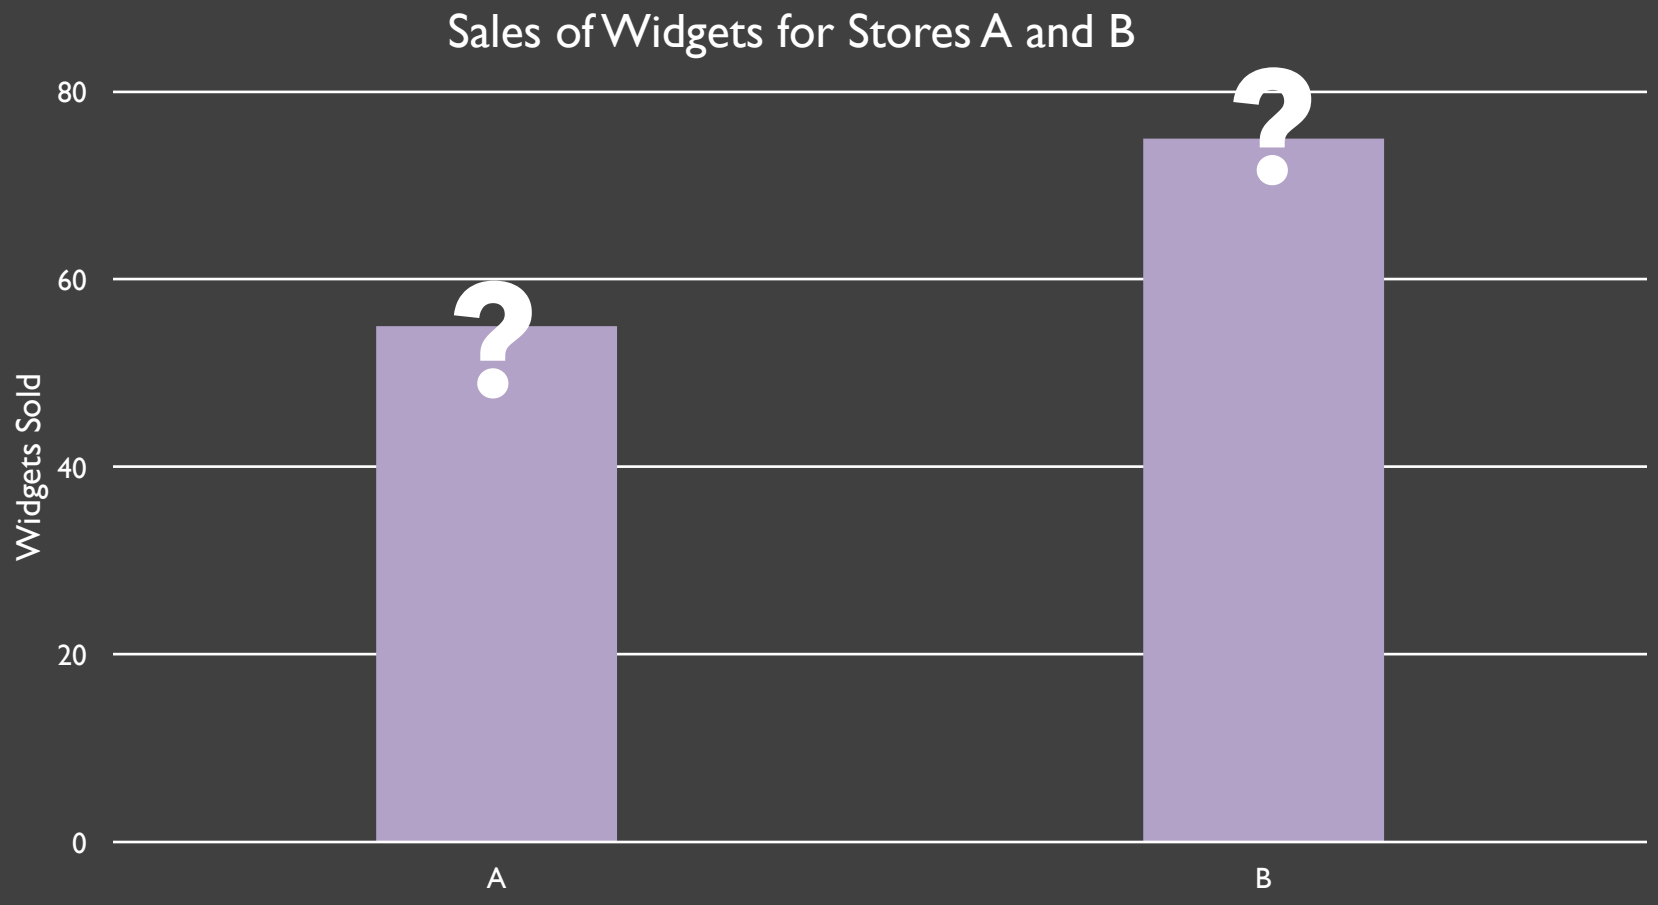

Forecast Uncertainty

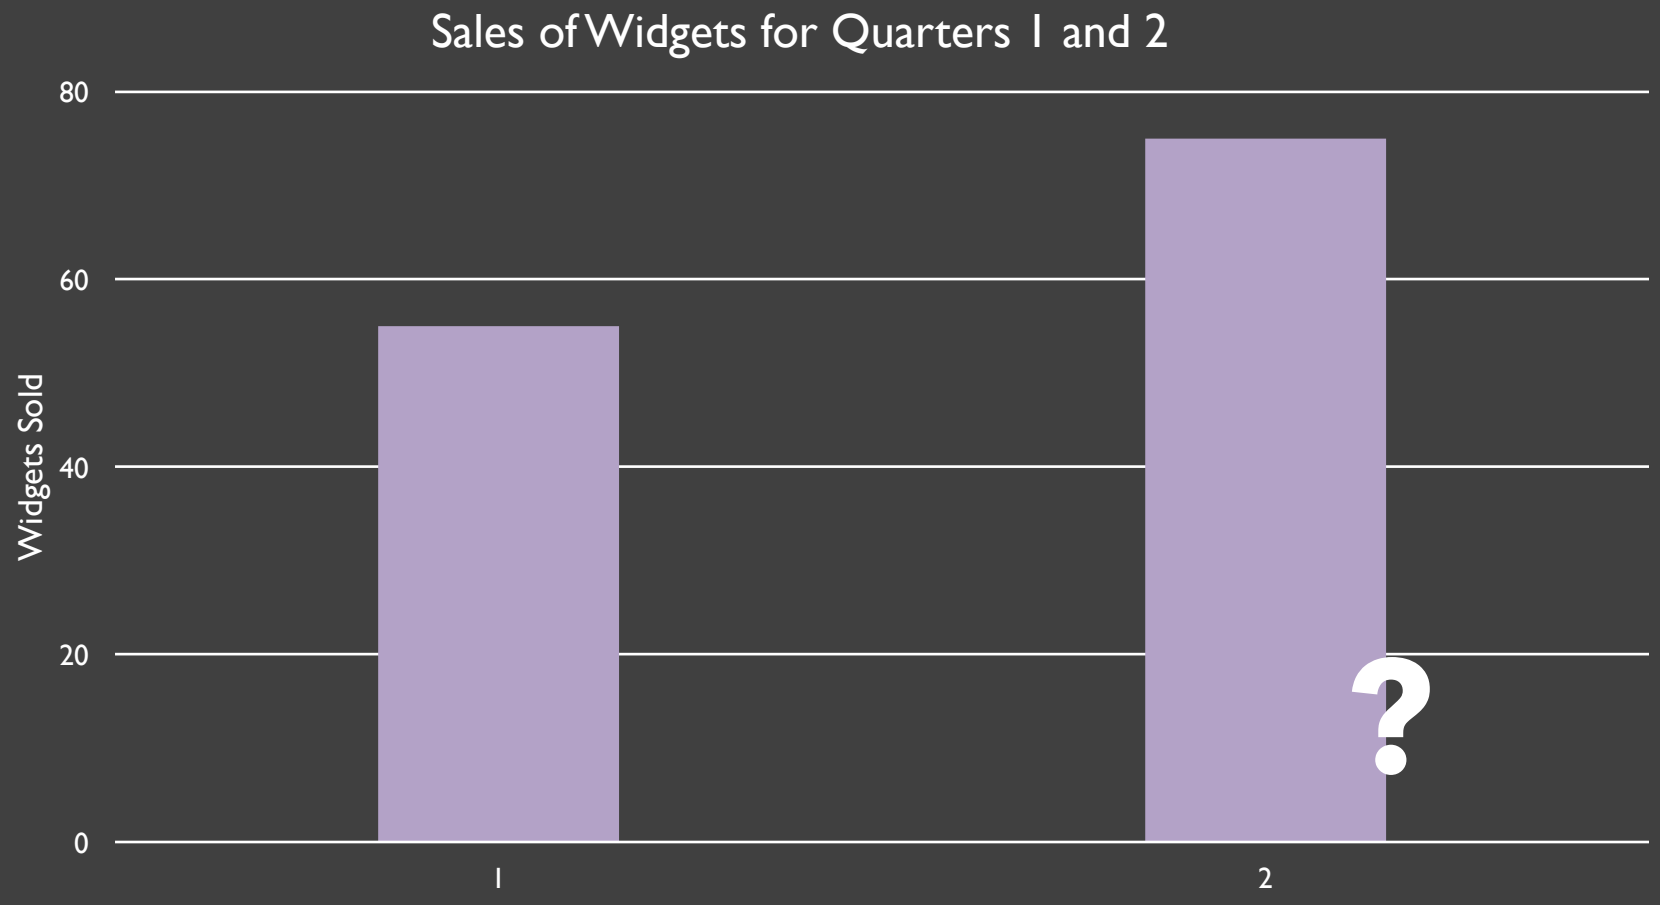

Decision Uncertainty

Uncertainty Sources

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

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

Forecast Uncertainty: “We’re not sure what will happen to the data next”

Decision Uncertainty: “We’re not sure what to do with the data”

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.

Should I Bring an Umbrella?

Decision Uncertainty

"50% Chance of Rain"

Types of Error

	I False Positive
II False Negative	

Model Uncertainty

"50% Chance of Rain"

Model Uncertainty

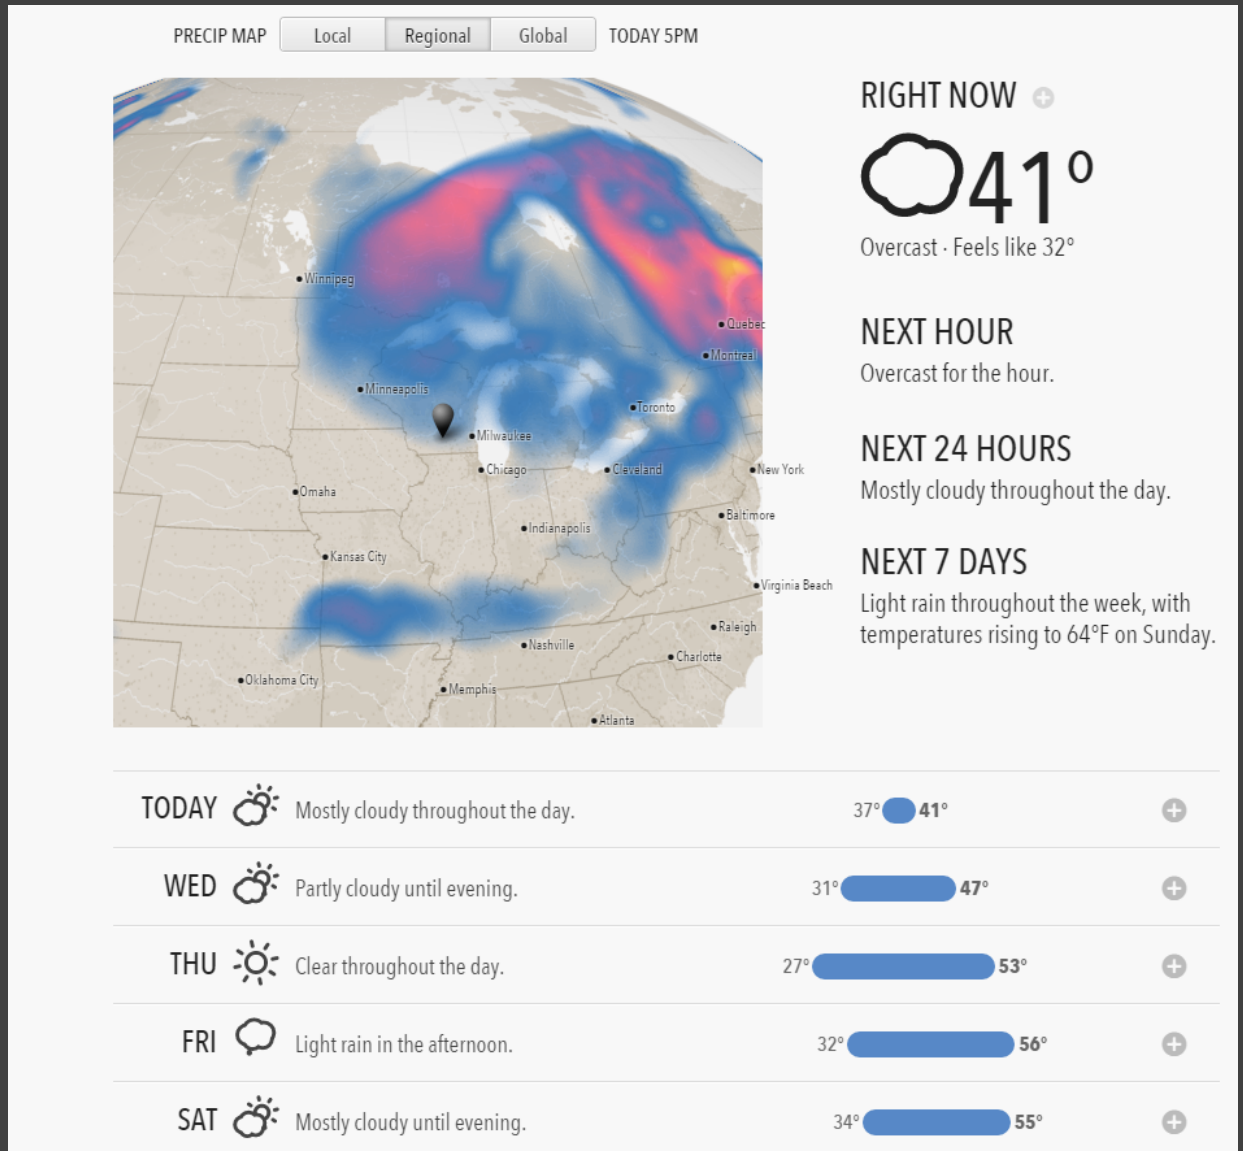

Measurement Uncertainty

Precision

Measurement Uncertainty

Precision

Measurement Uncertainty

Precision

Measurement Uncertainty

Precision

Accuracy

Measurement Uncertainty

Precision

Accuracy

Measurement Uncertainty

Precision

Accuracy

Measurement Uncertainty

Precision

Accuracy

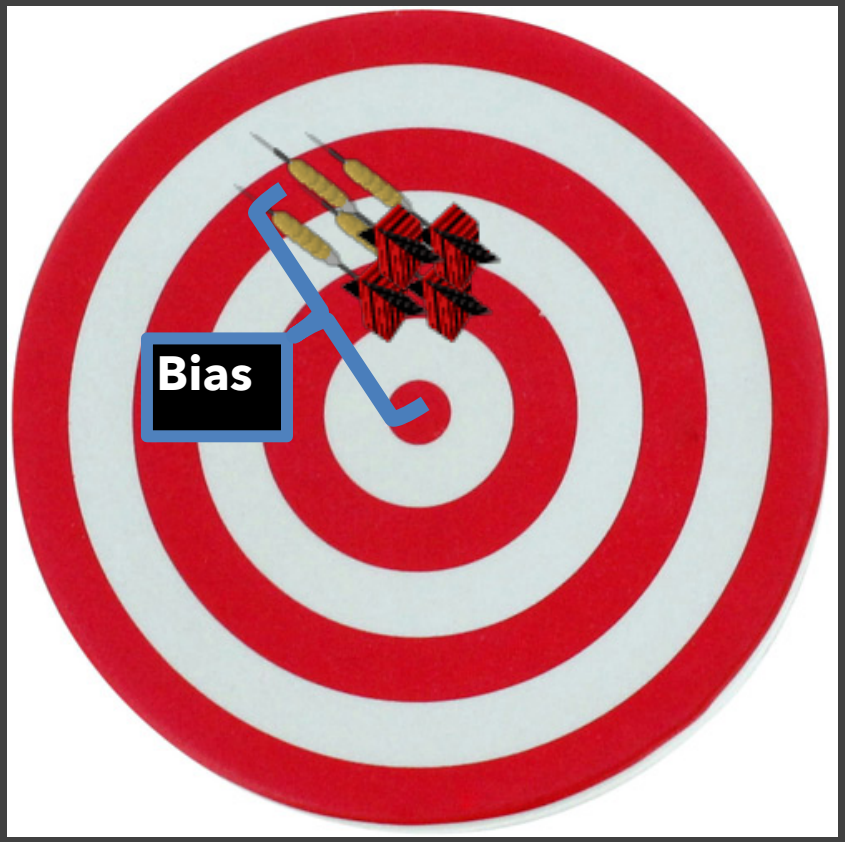

What Does Uncertainty Mean?

Any one of a number of potentially interconnected quantitative, qualitative, or factors that affect the quality, reliability, or utility of your data or data-driven decisions. Anything that can cause you to be unsure about your data or how to use it.

What Does Uncertainty Mean?

Any one of a number of potentially interconnected quantitative, qualitative, or factors that affect the quality, reliability, or utility of your data or data-driven decisions. Anything that can cause you to be unsure about your data or how to use it.

LOTS OF THINGS

Visualizing Distributions

Distribution Visualizations

Strip Plot

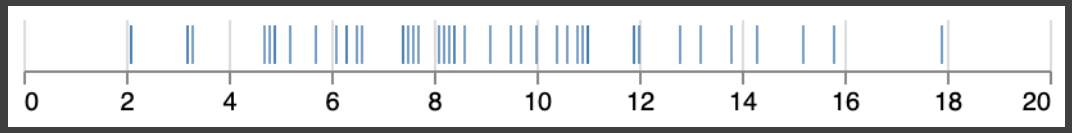

Jittered Plot

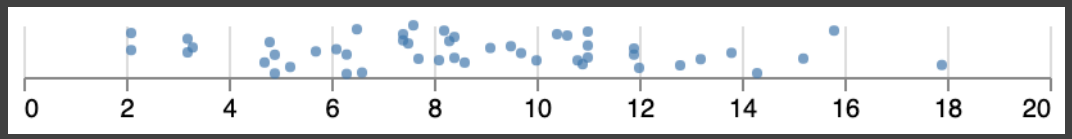

Box Plot

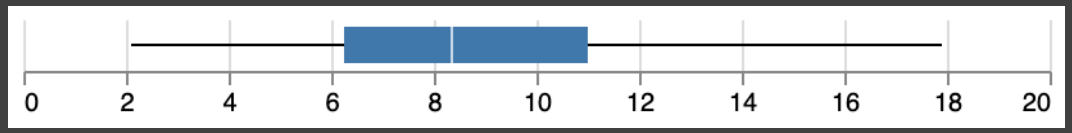

Dot Plot

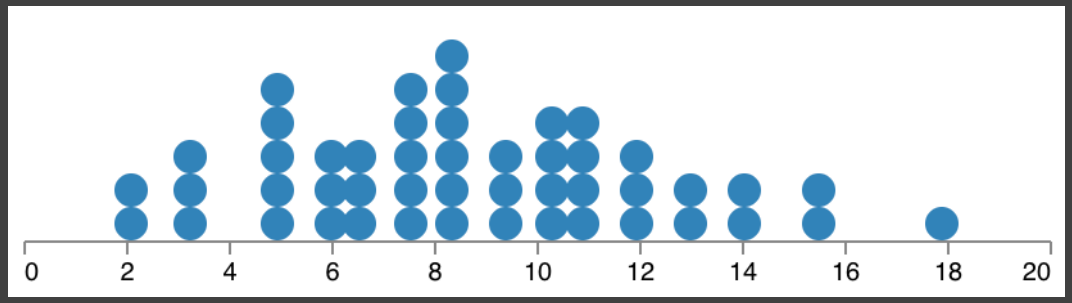

Distribution Visualizations

Histogram

bin size = 2

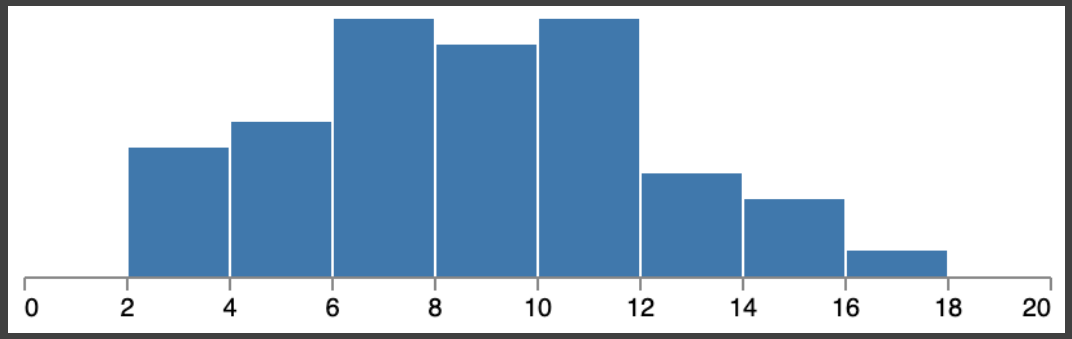

Density Plot

kde, **σ** = 0.5

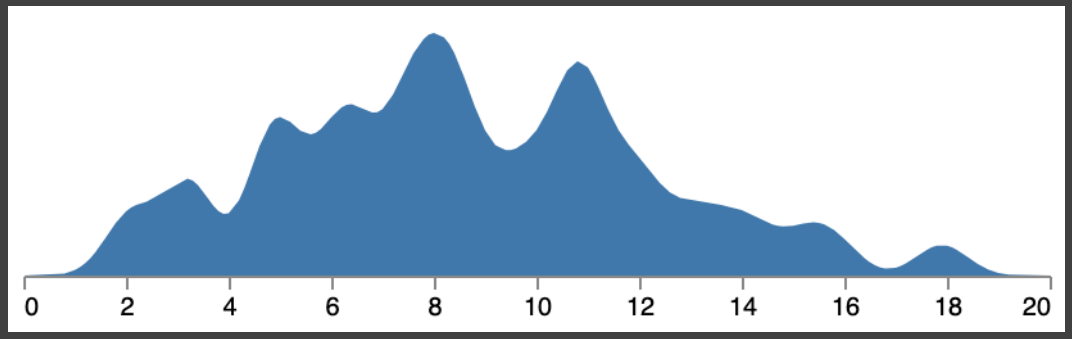

Violin Plot

kde, **σ** = 0.5

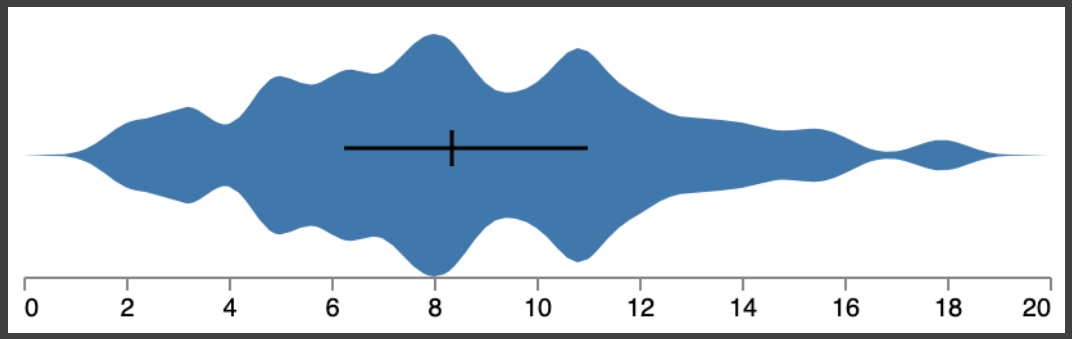

Identical boxplots, different distributions

Boxplots are great. They show medians and ranges and enable comparison of different groups. However, boxplots can be misleading. Different datasets can have the same descriptive statistics (left), but quite different underlying distributions (middle). Therefore, it is crucial to visualize the distribution in addition to descriptive statistics. Violin plots with integrated boxplots are great for this.

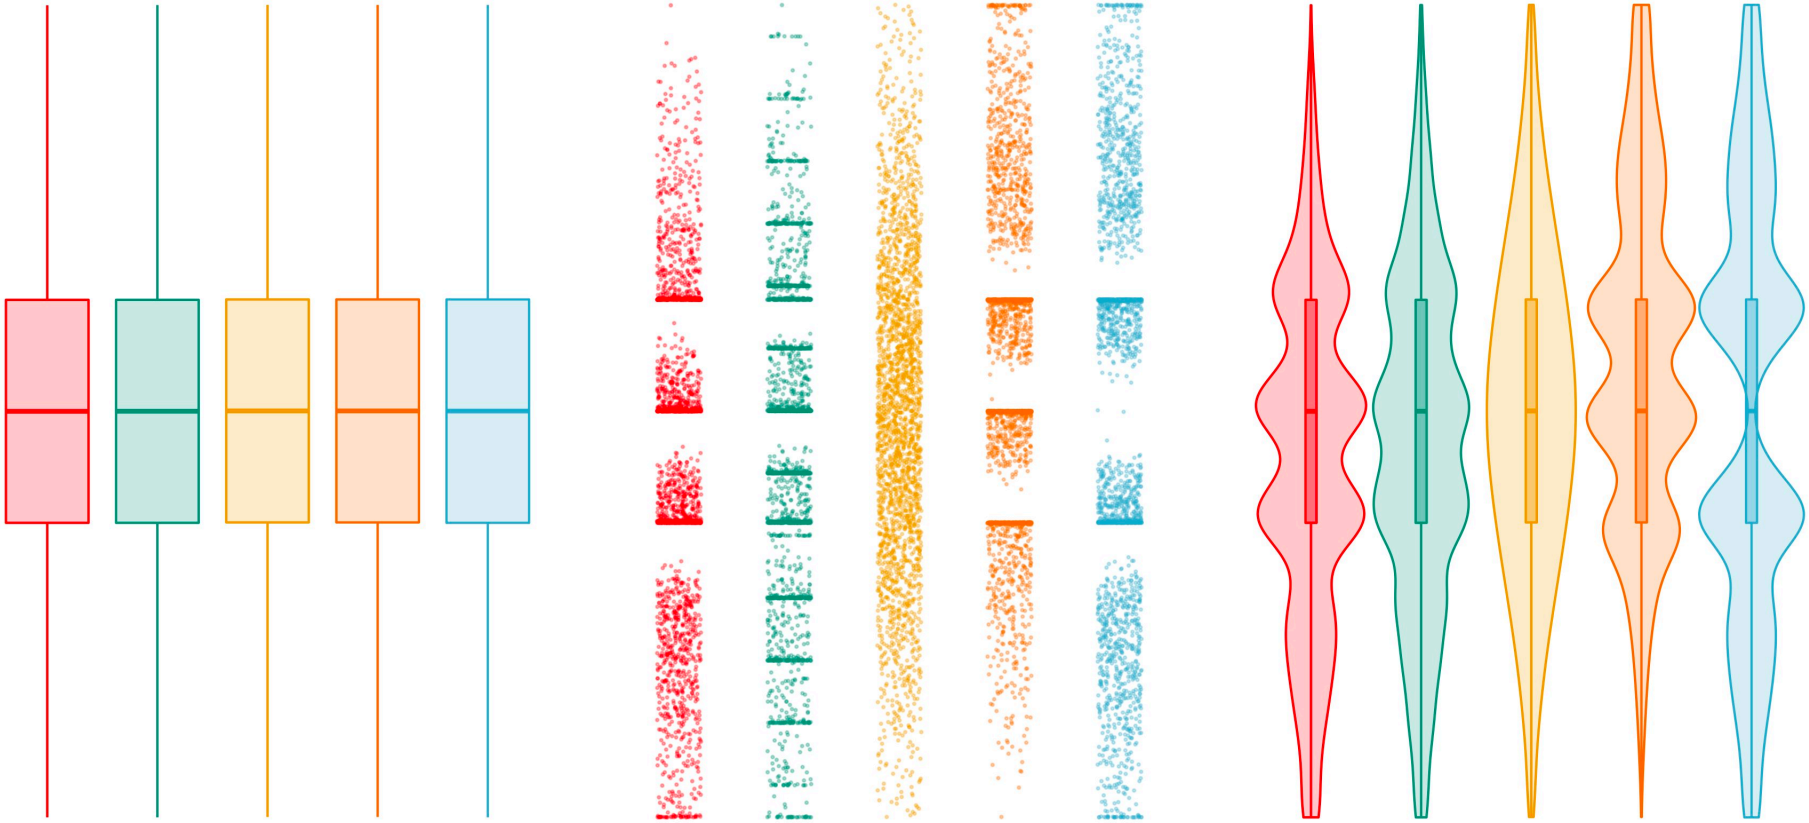

@MBozukova | #30DayChartChallenge | Day 27 - educational
Data source: {datasauRus} R package

Now in 2D! Heatmaps, Contours

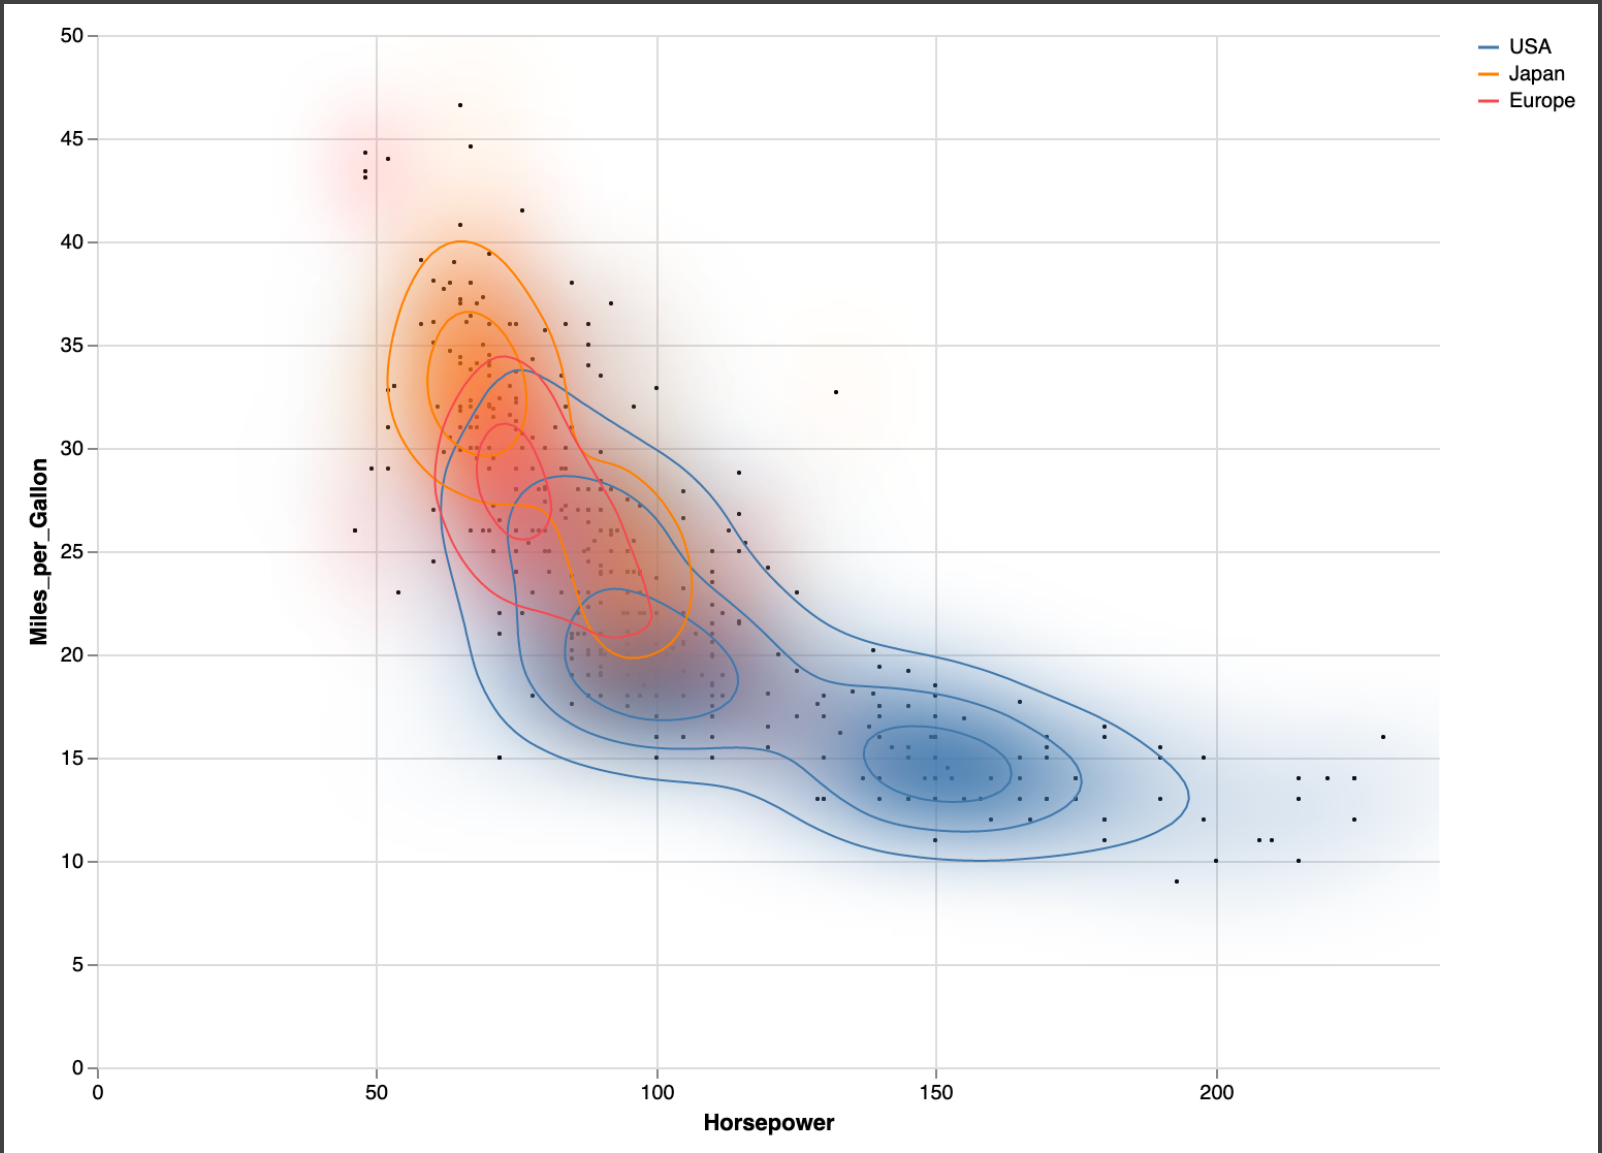

Quantified Uncertainty

Error Bars

Standard Deviation (σ)
Standard Error (σ / √n)
1.5 * IQR (Interquartile Range)
Confidence Intervals
… and so on

Confidence Intervals

What does a 95% confidence interval indicate?

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

Wrong!

Rather, given an infinite number of independent 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

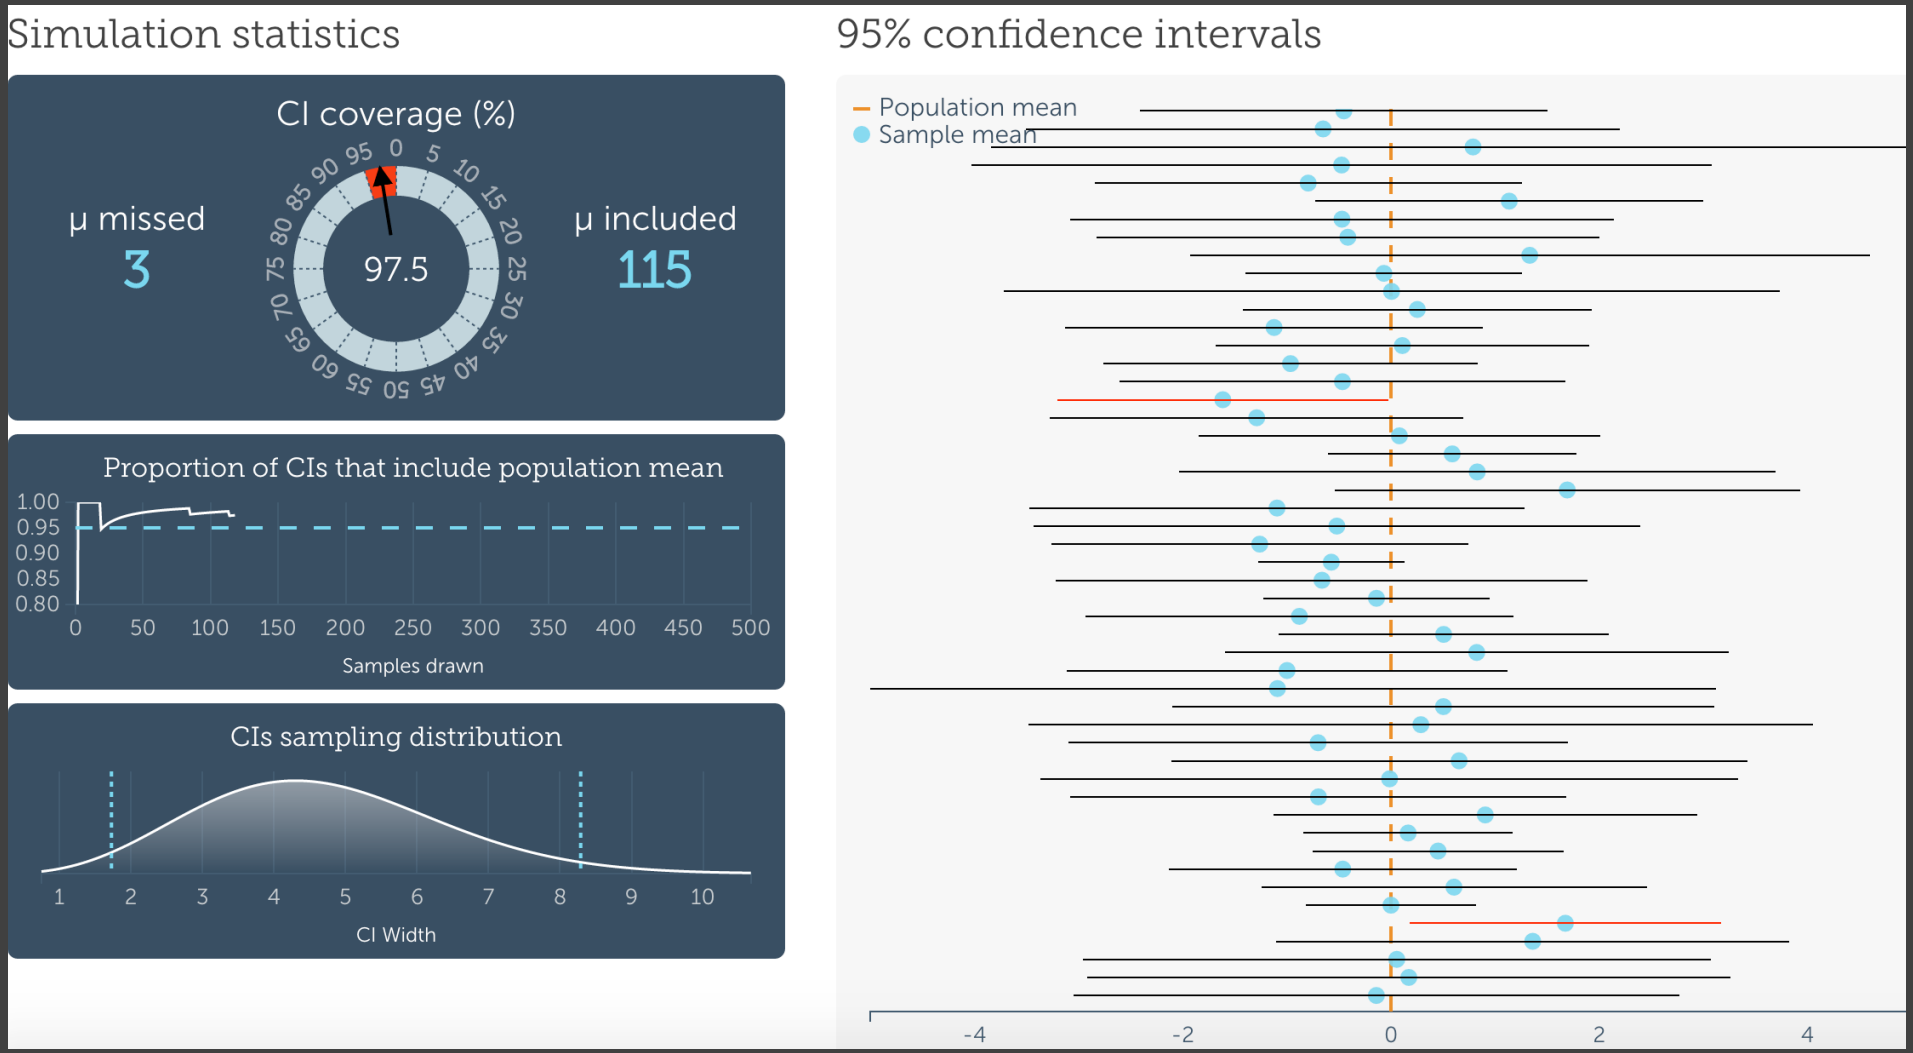

Regression Coefficients

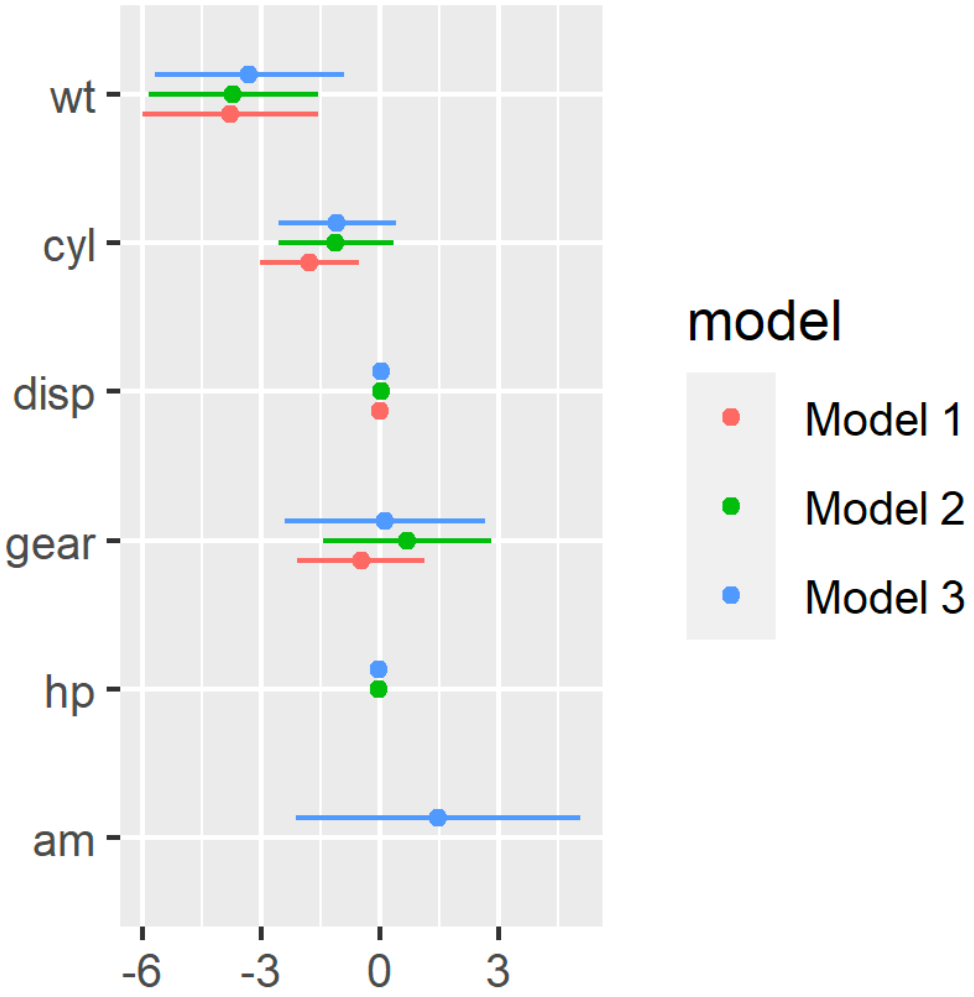

95% CIs for regression model parameters.

Here, we compare fitted parameters from 3 different models. Not all predictors are included in all models.

Visual comparison: does the CI overlap 0?

Error Bars

The mean treatment effect is higher than than the placebo.

Is this difference in means *statistically significant?*

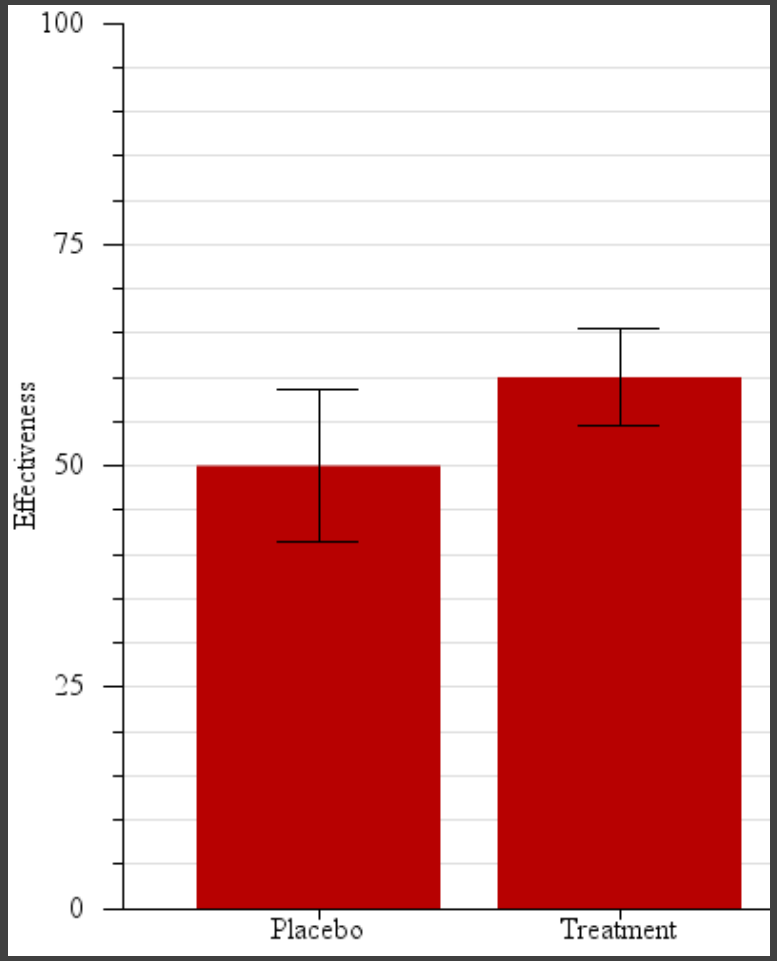

Guess the p-value…

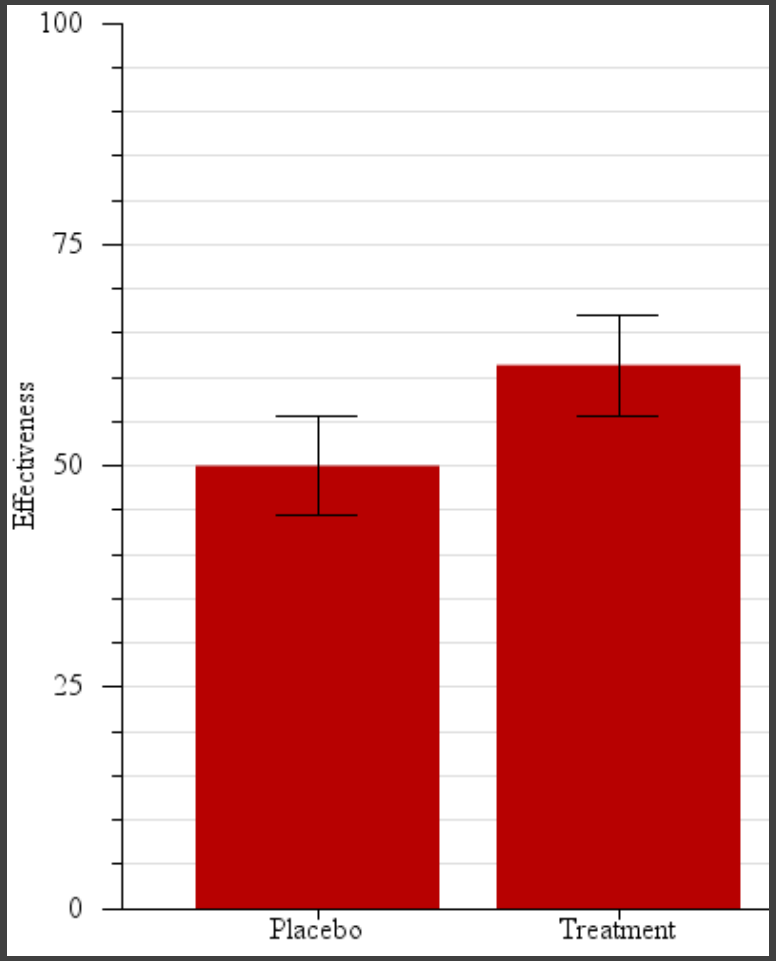

Error bars depict
95% Conf. Interval

Guess the p-value…

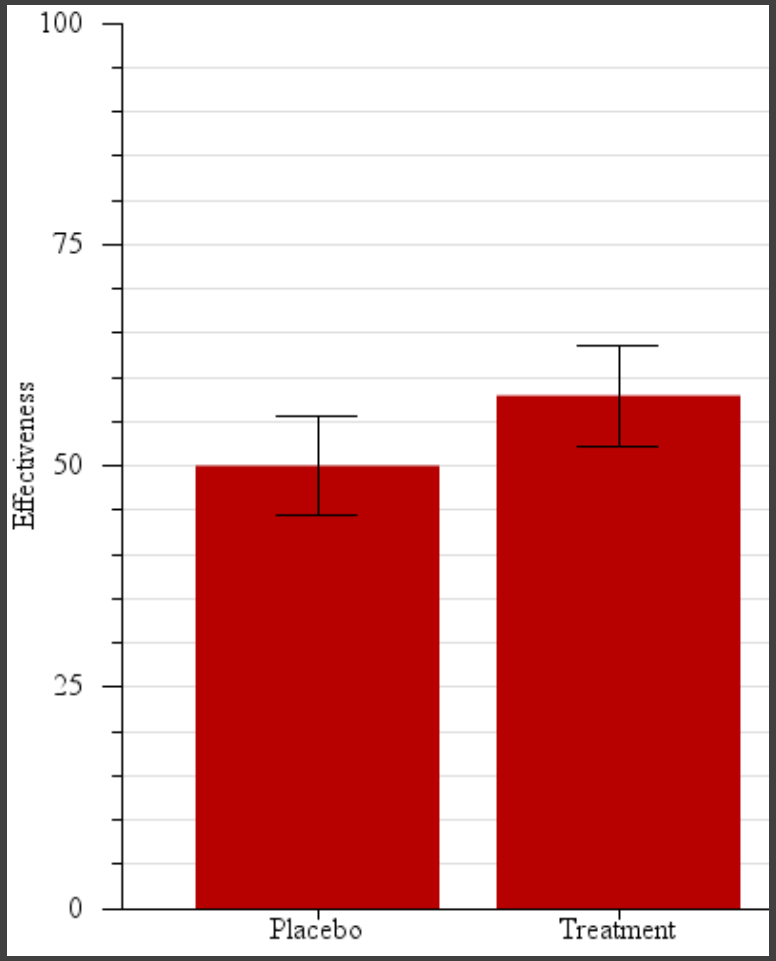

Error bars depict
95% Conf. Interval

Guess the p-value…

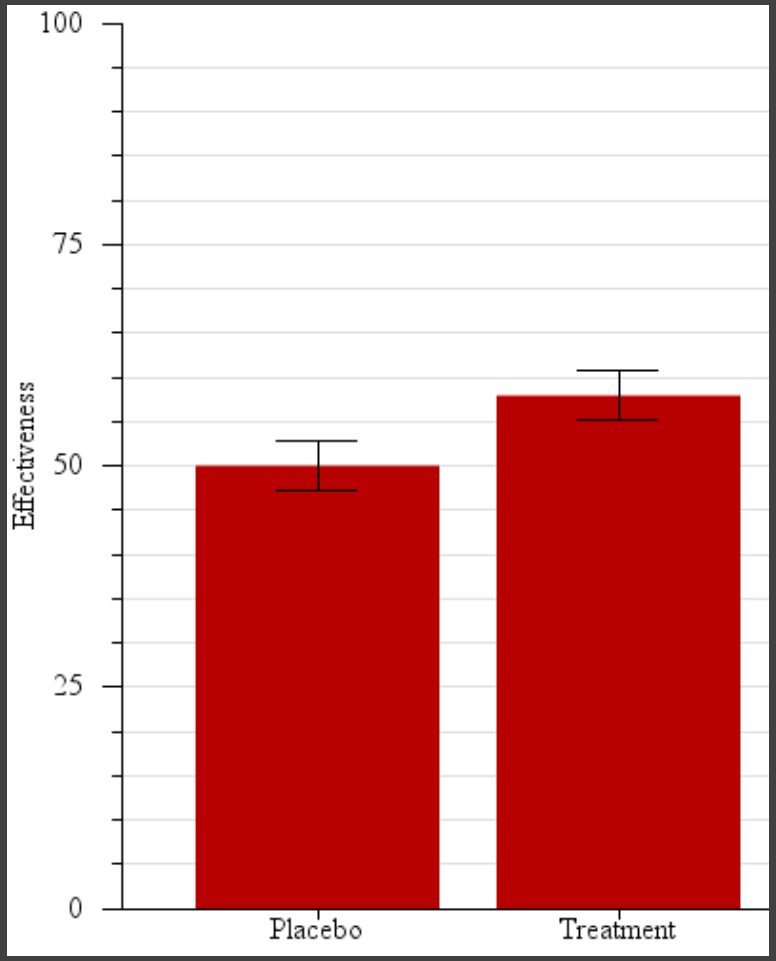

Error bars depict *standard error*

Inference by Eye

95% CIs

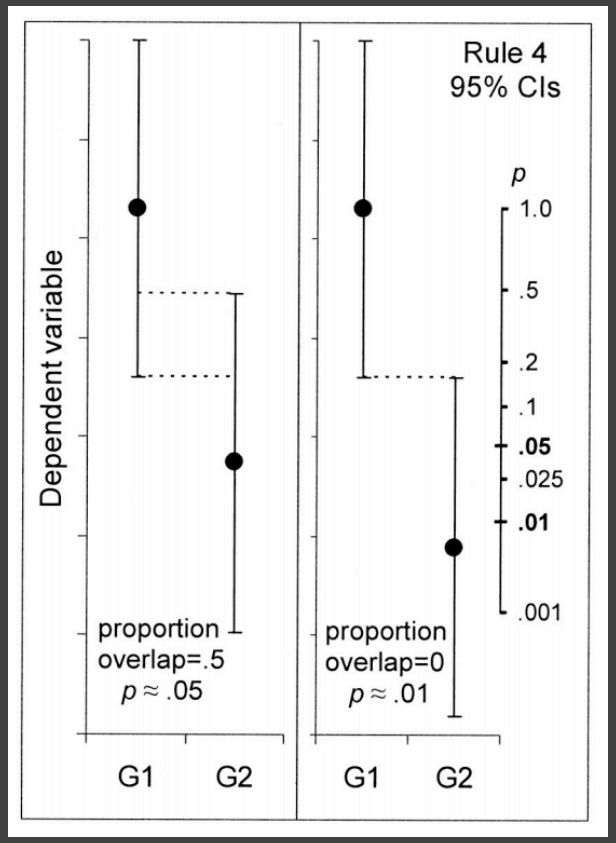

Standard Error

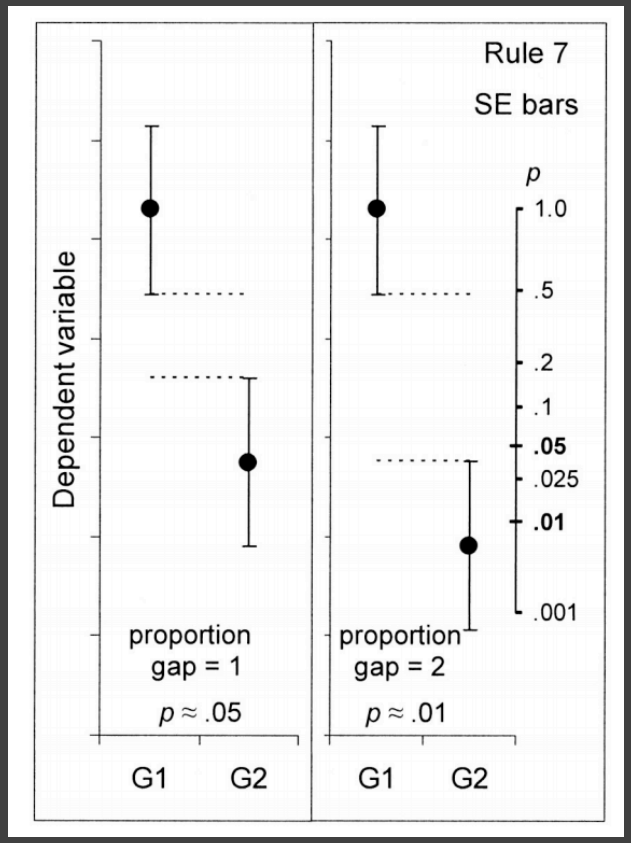

Cumming, Geoff and Finch, Sue. Inference by eye: confidence intervals and how to read pictures of data. American Psychologist, 2005.

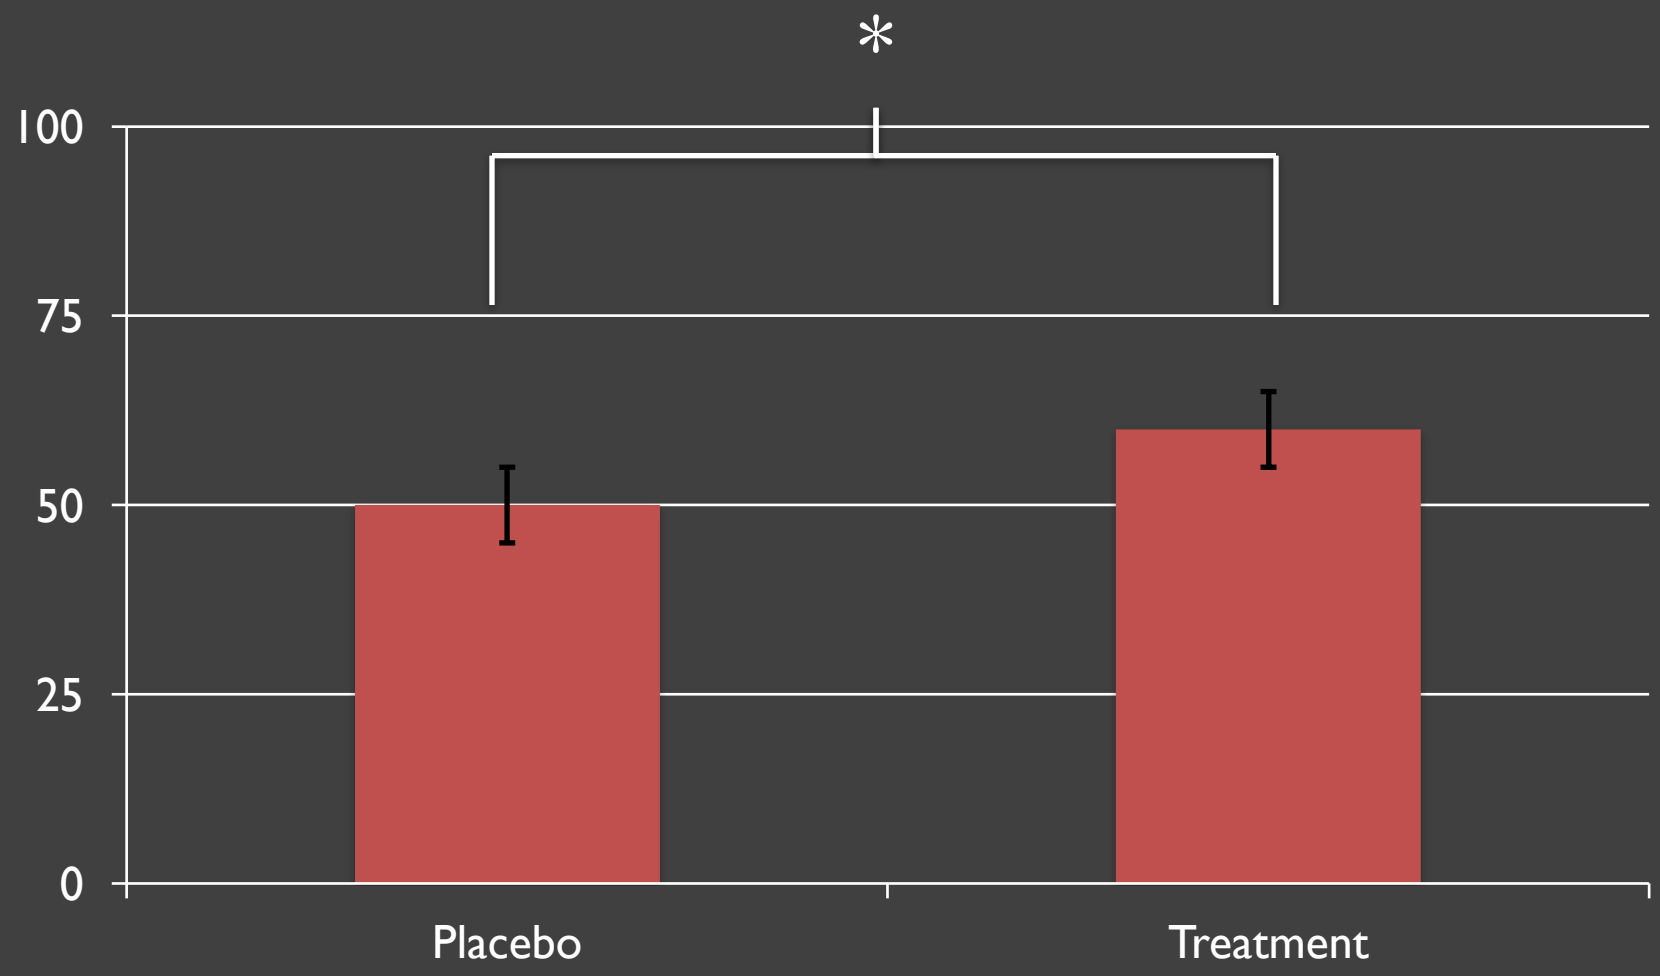
*
100
75
50
25
0
Placebo
Treatment

Misplaced Emphasis?

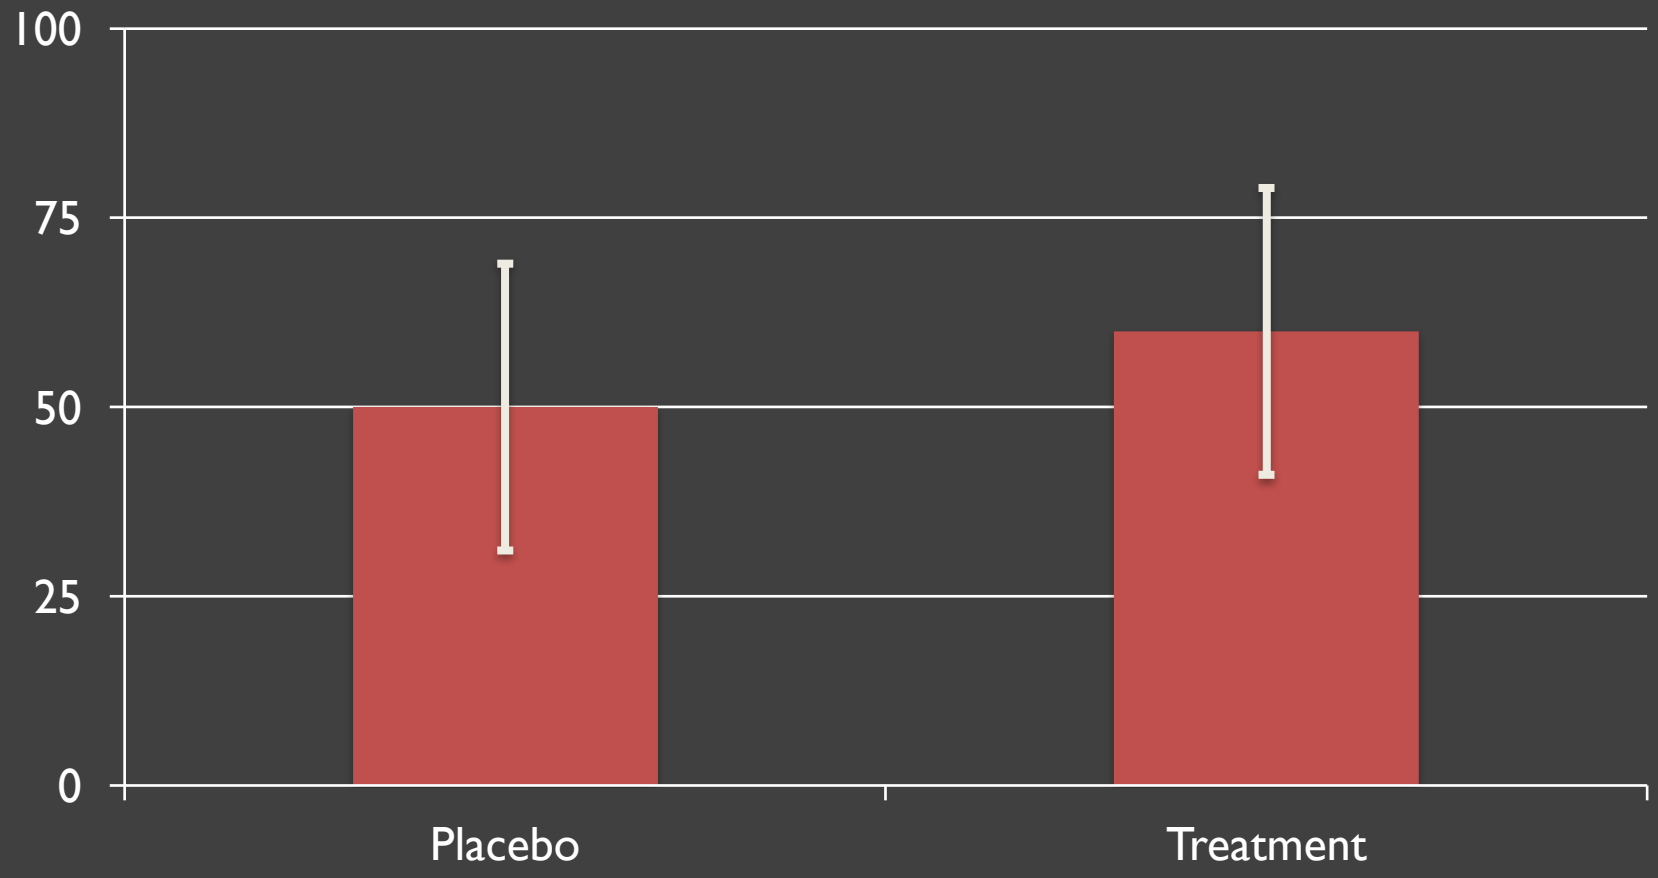

Misplaced Emphasis?

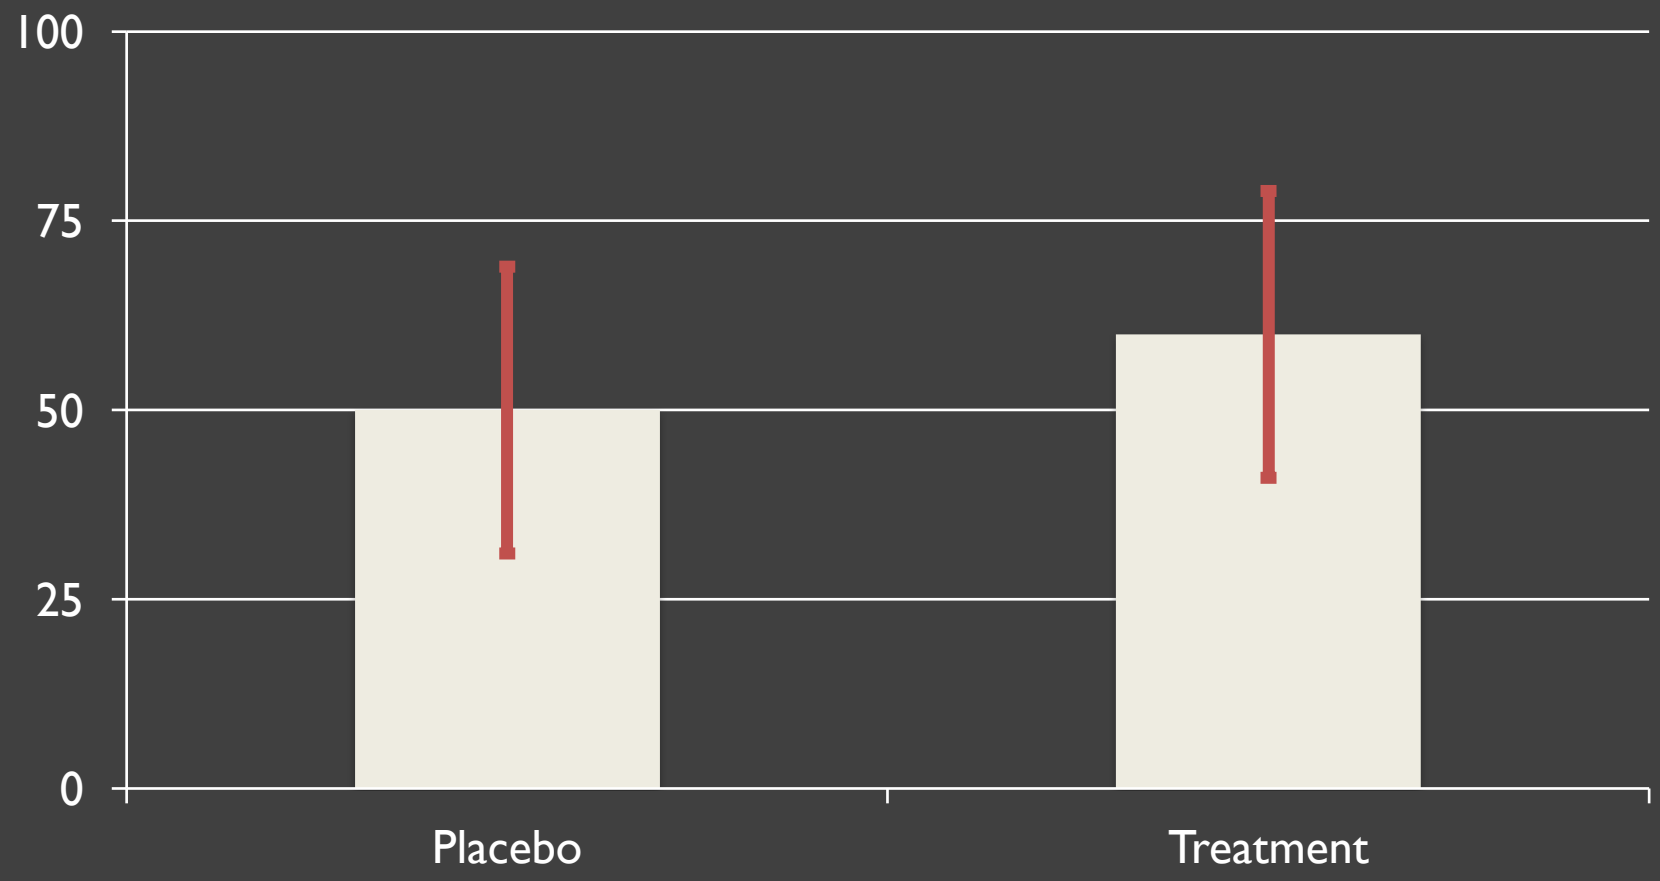

Within-the-Bar Bias

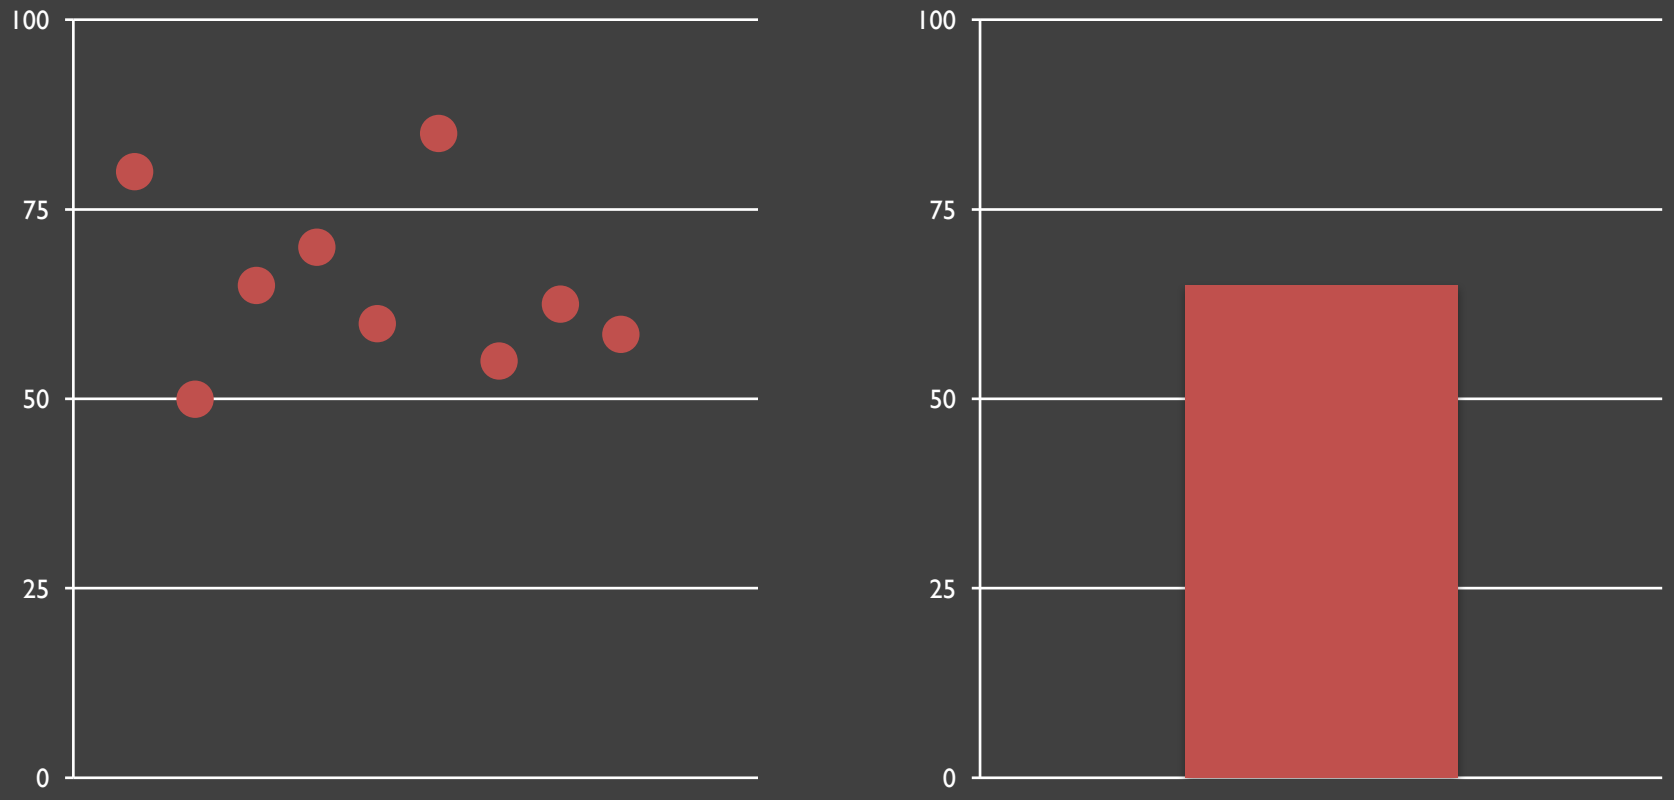

Newman & Scholl. (2012) “Bar graphs depicting averages are perceptually misinterpreted: the within-the-bar bias.”

Within-the-Bar Bias

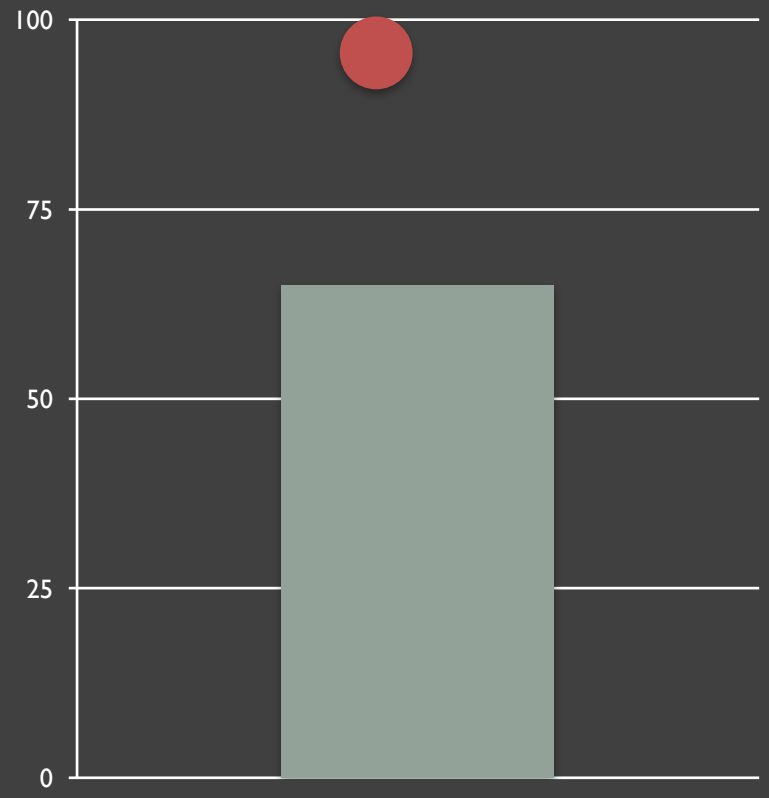

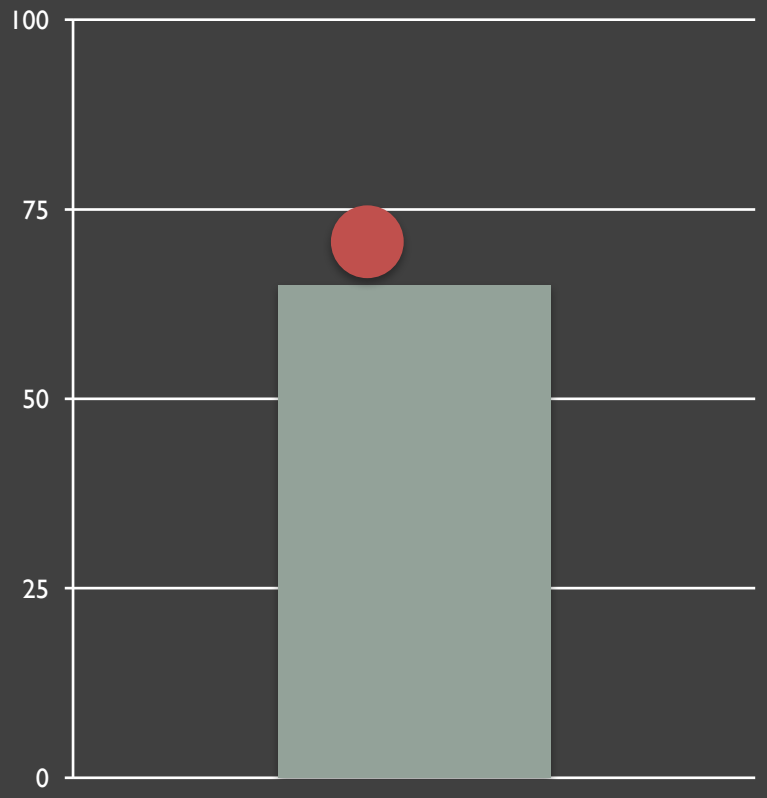

Within-the-Bar Bias

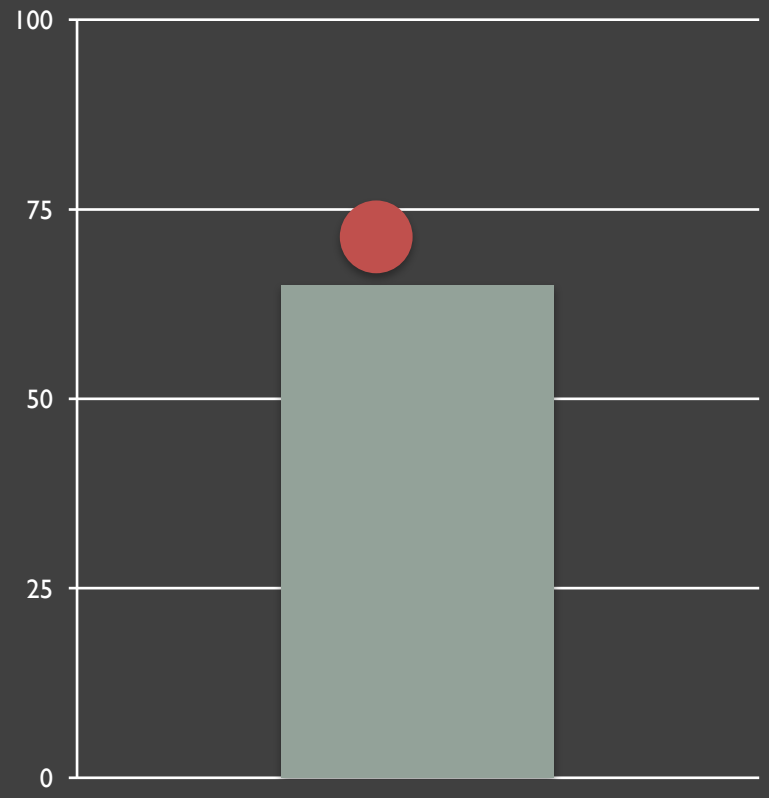

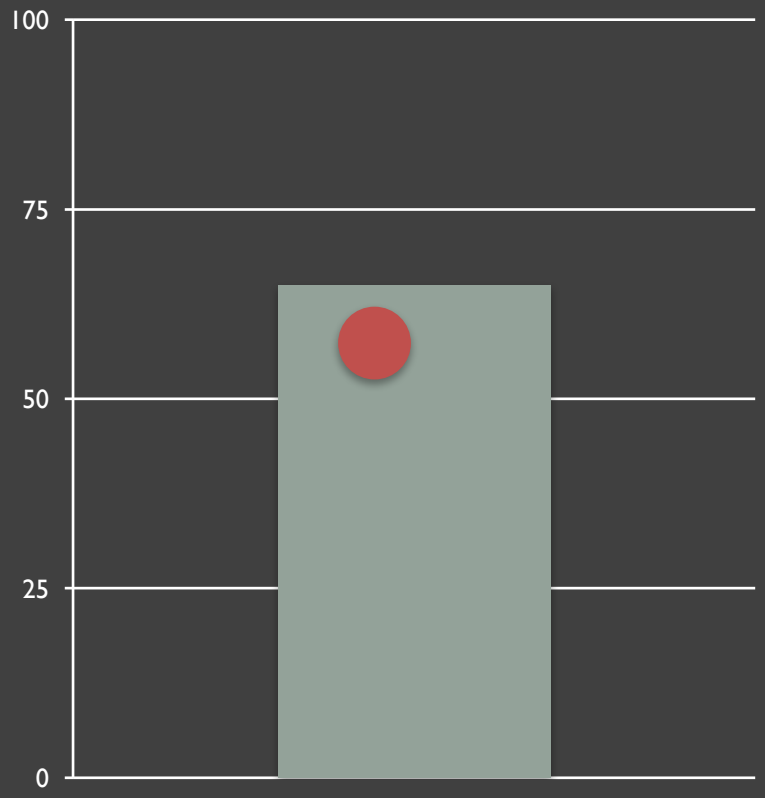

Within-the-Bar Bias

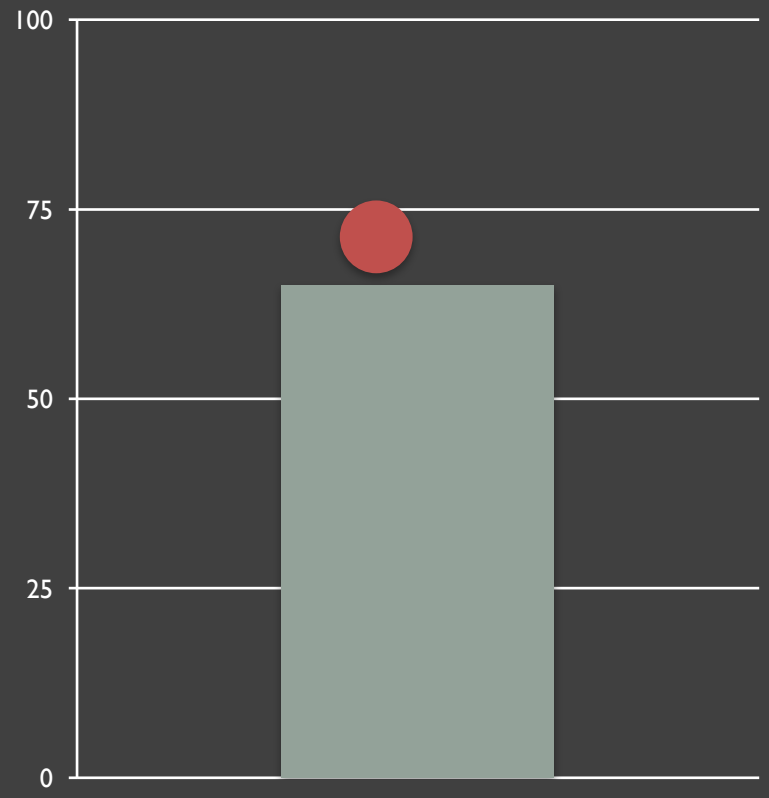

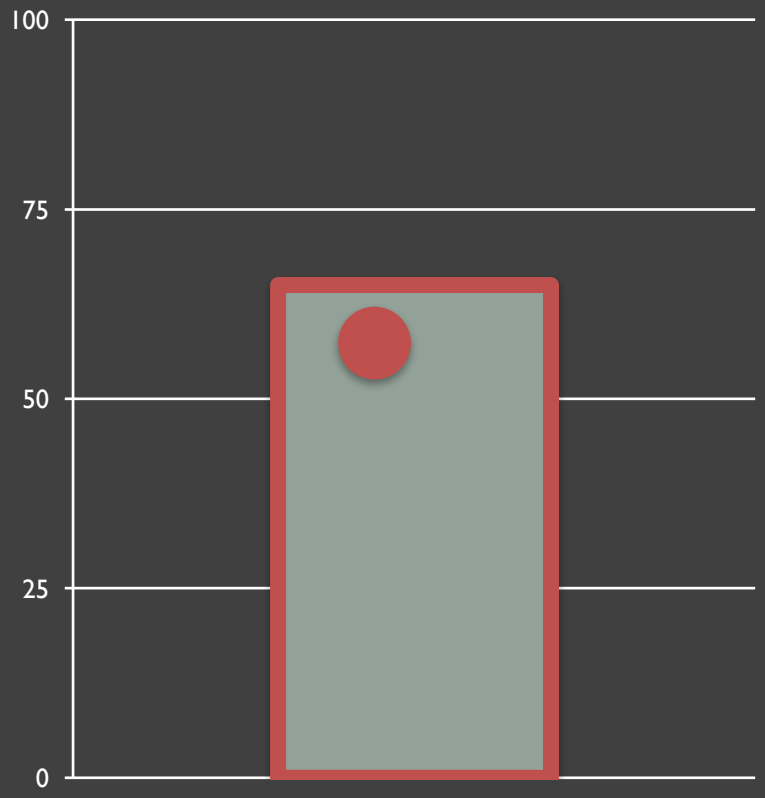

Alternatives to Error Bars

Gradient Plot

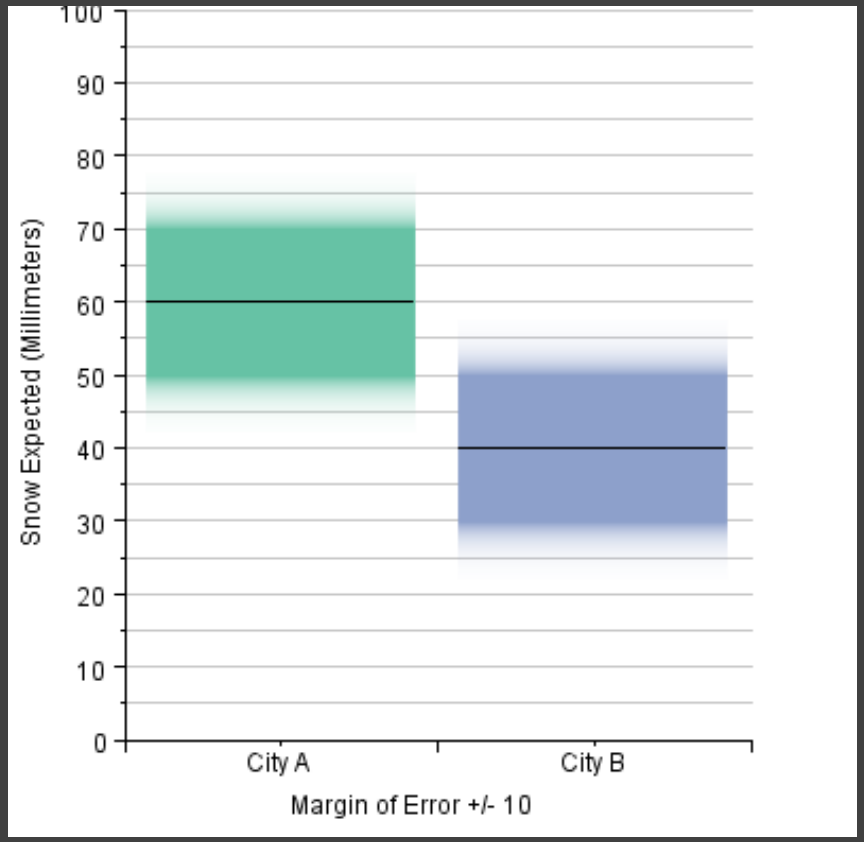

Violin Plot

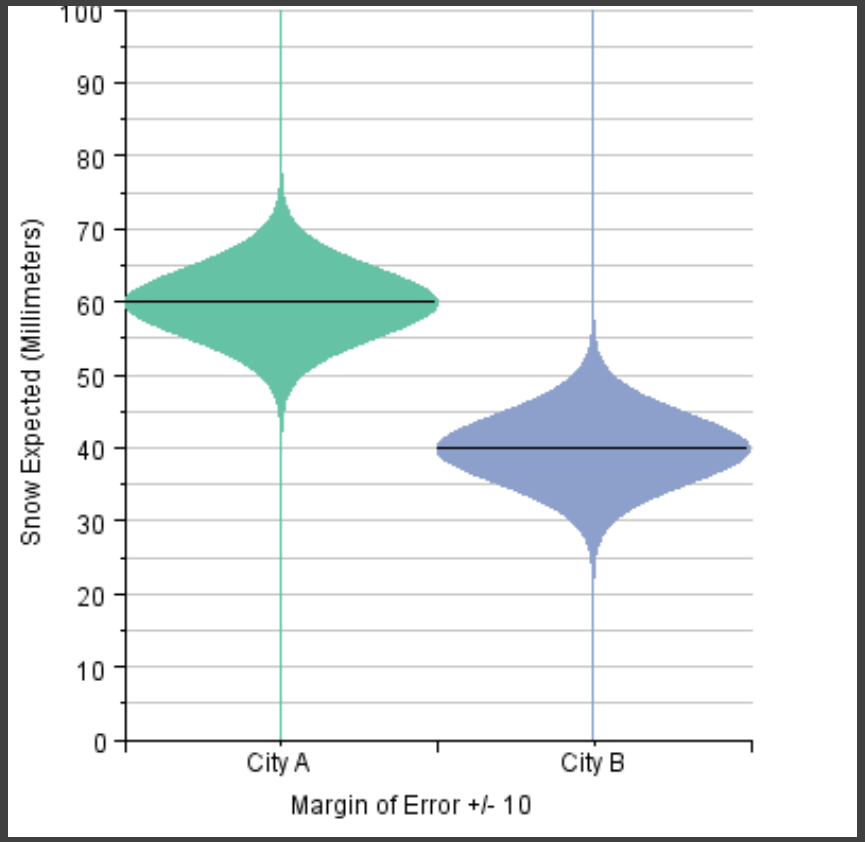

For inference tasks, focus on the **uncertainty** not the point estimate!

Encoding Uncertainty

Uncertainty Vis Pipeline

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

SNAP

Data Map

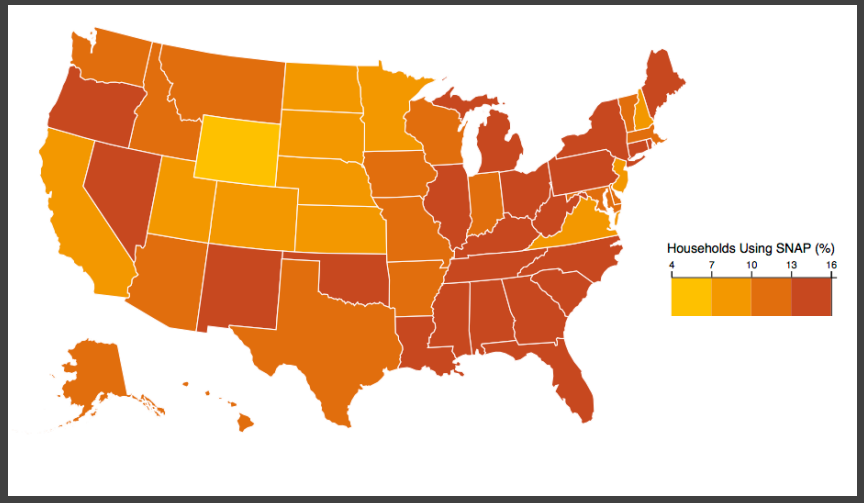

SNAP

Data Map

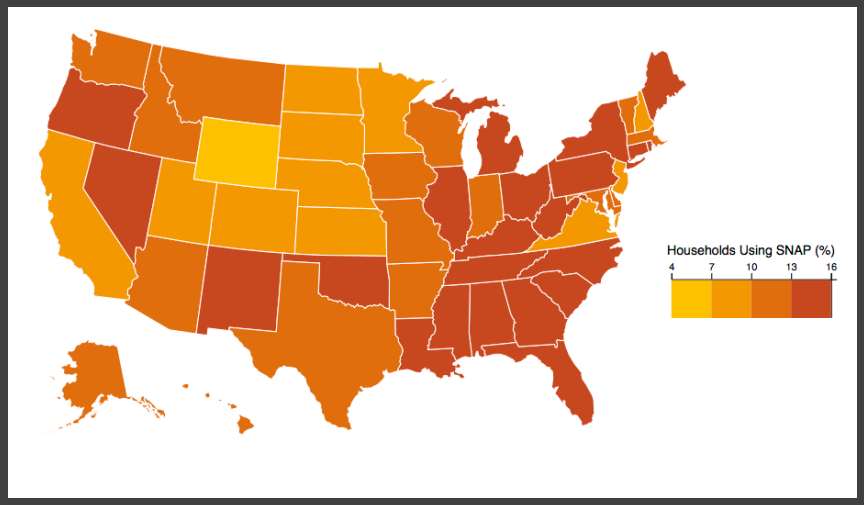

Uncertainty Map

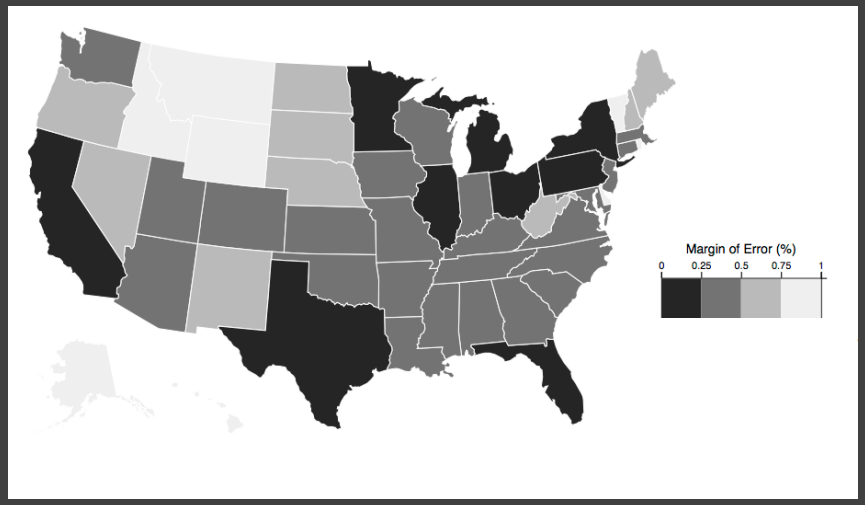

Uncertainty Vis Pipeline

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

Uncertainty Vis Pipeline

1) Quantify uncertainty
2) Choose a free visual variable
3) Encode uncertainty with the variable
4) Unify the Data Map and Uncertainty Map

How to Unify?

Data Map

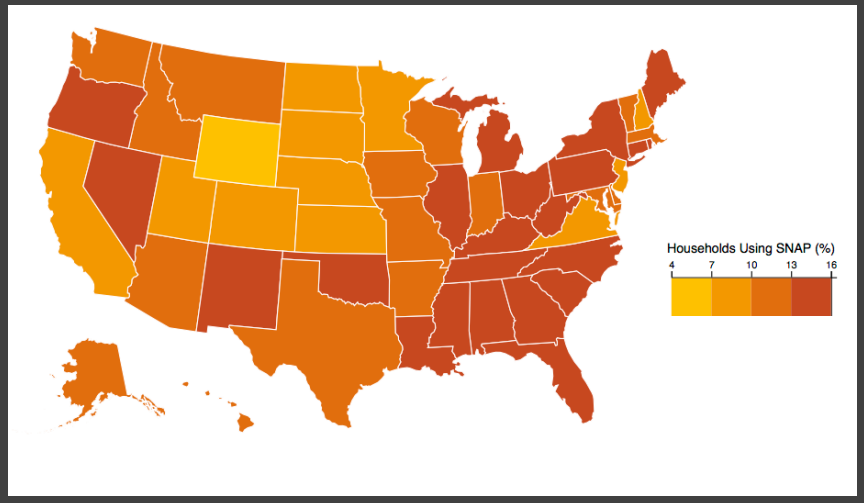

Uncertainty Map

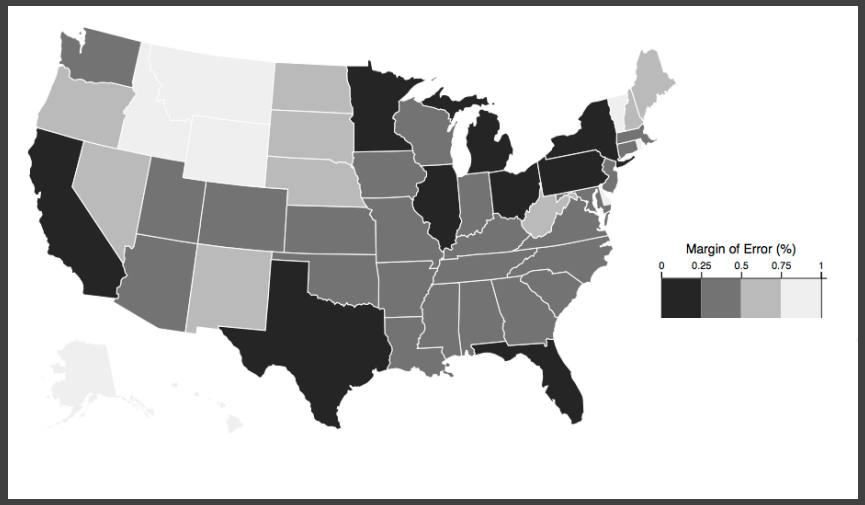

Juxtaposition

Data Map

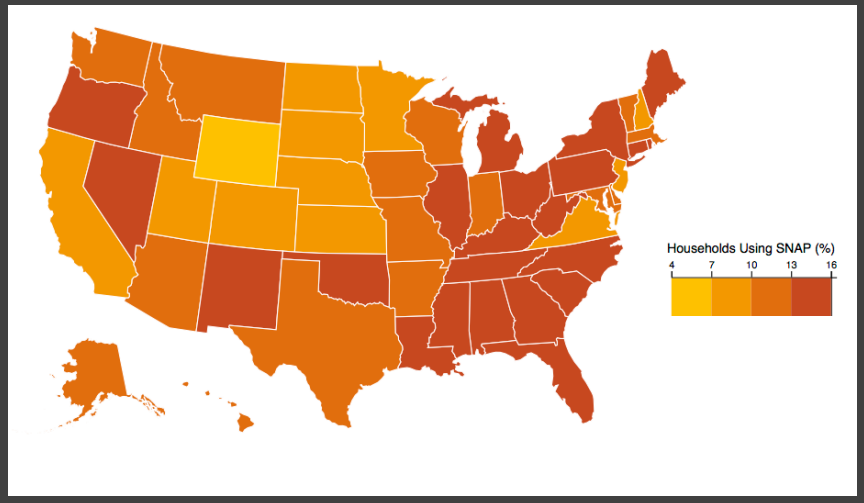

Uncertainty Map

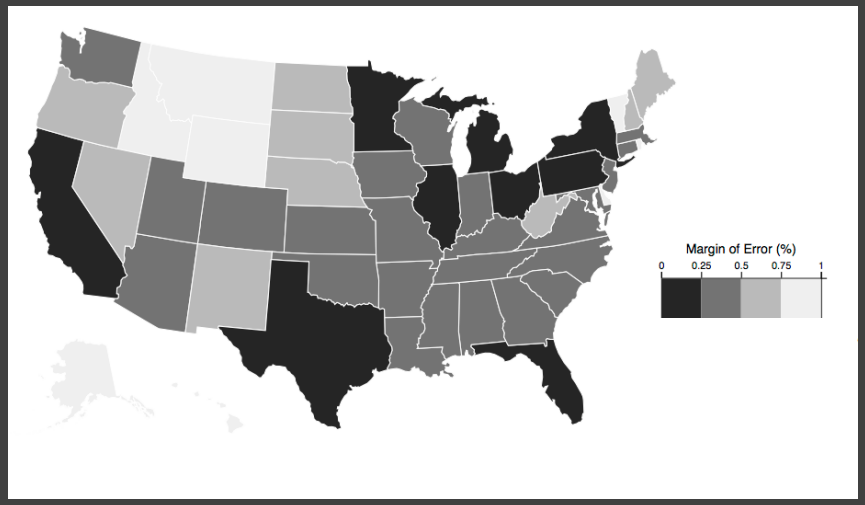

Superposition

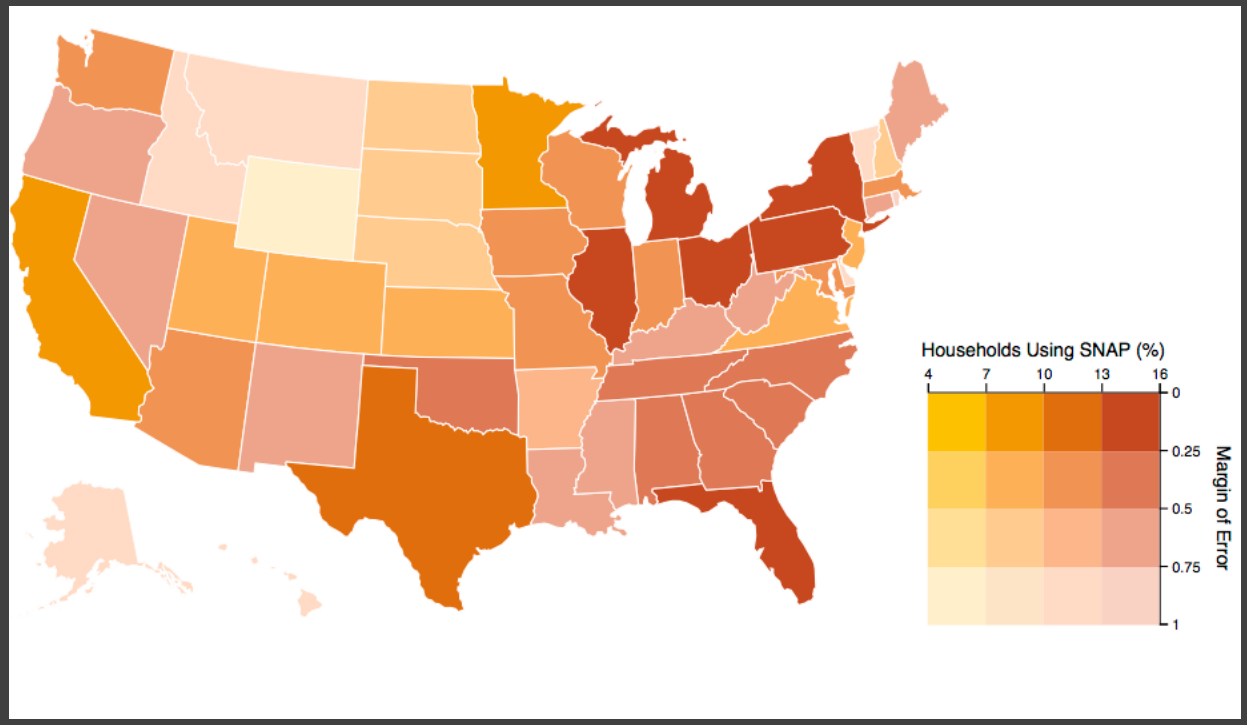

Superposition

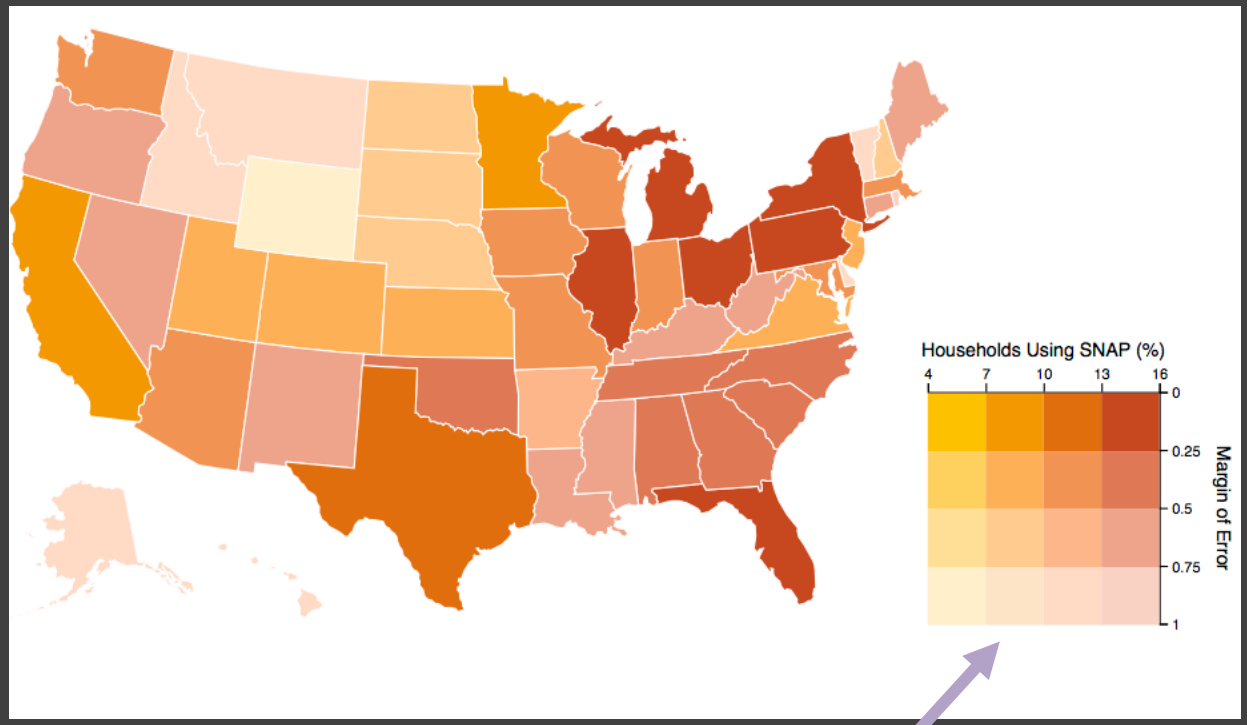

Bivariate Map

Superposition

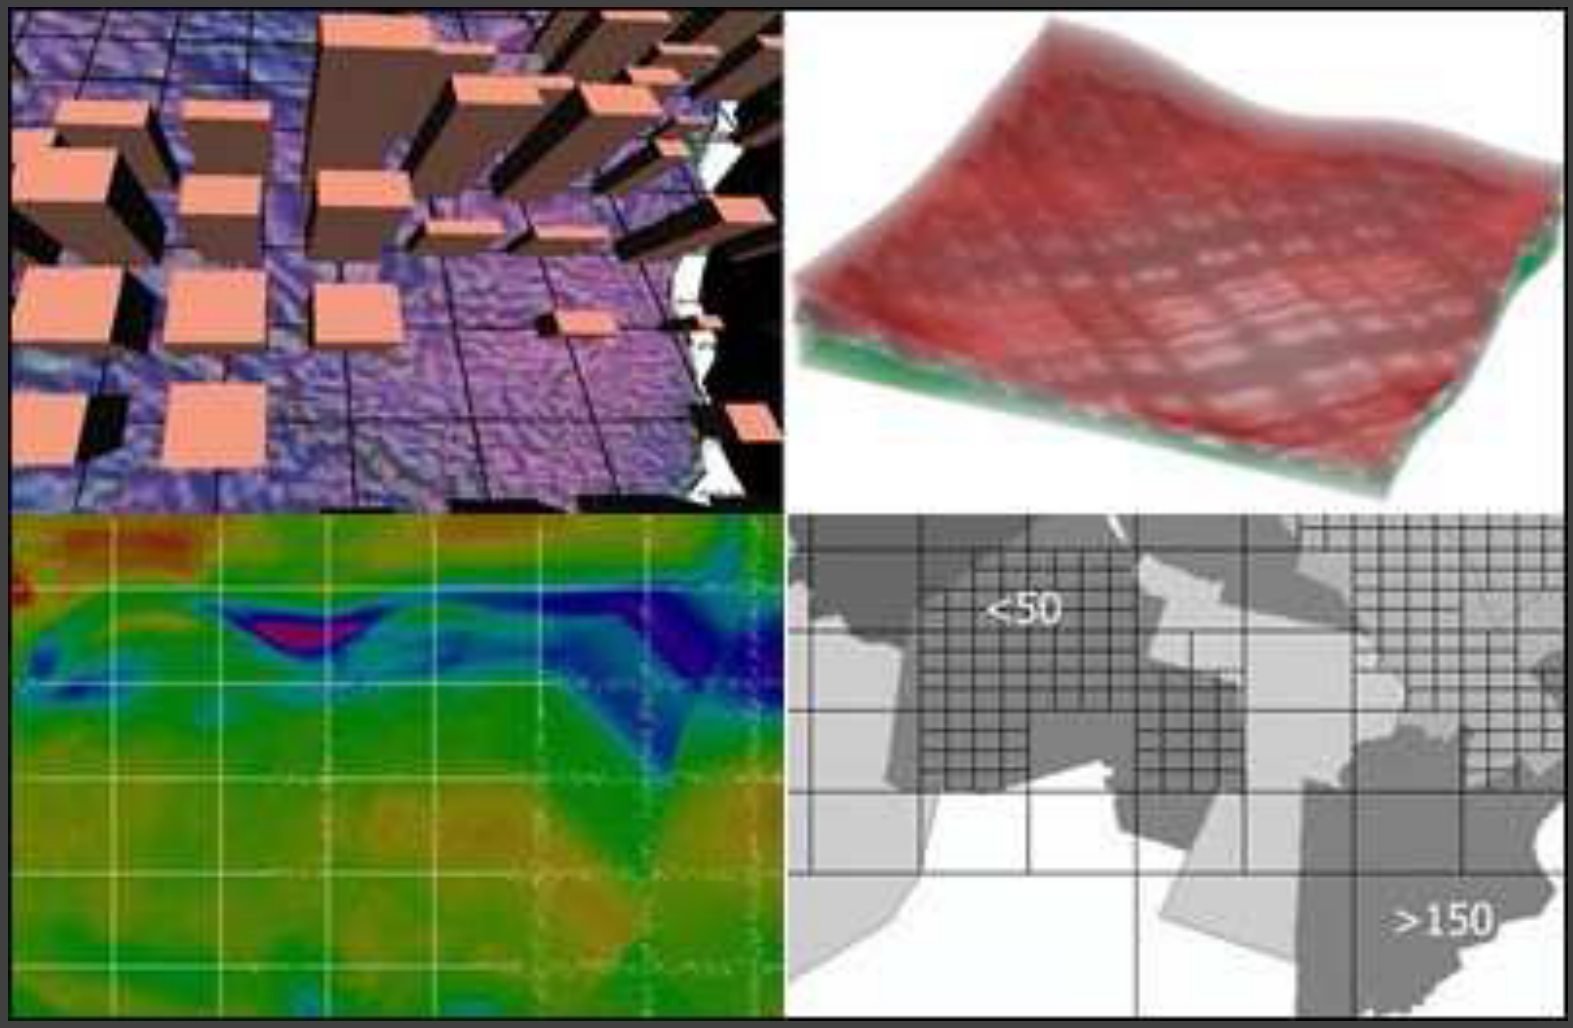

Griethe, Henning and Schumann, Heidrun. The Visualization of Uncertain Data: Methods and Problems. SimVis, 2006.

Uncertainty Vis Pipeline

1) Quantify uncertainty
2) Choose a free visual variable
3) Encode uncertainty with the variable
4) Unify the Data Map and Uncertainty Map

Uncertainty Vis Pipeline

1) Quantify uncertainty
2) Choose a free **visual variable**
3) Encode uncertainty with the variable
4) Unify the Data Map and Uncertainty Map

Semiotics of Uncertainty

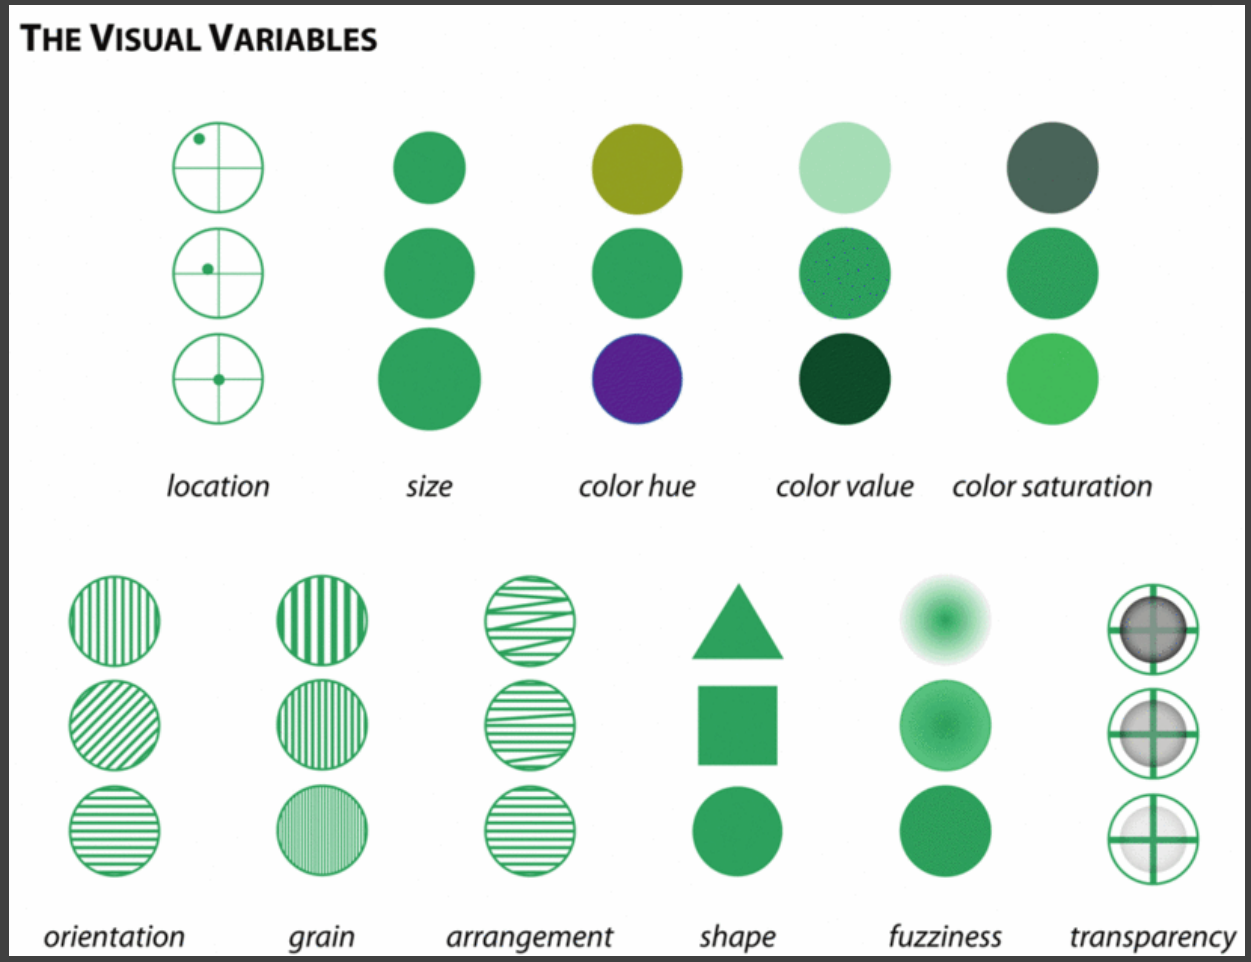

MacEachren et al. Visual Semiotics & Uncertainty Visualization: An empirical study. IEEE VIS, 2012.

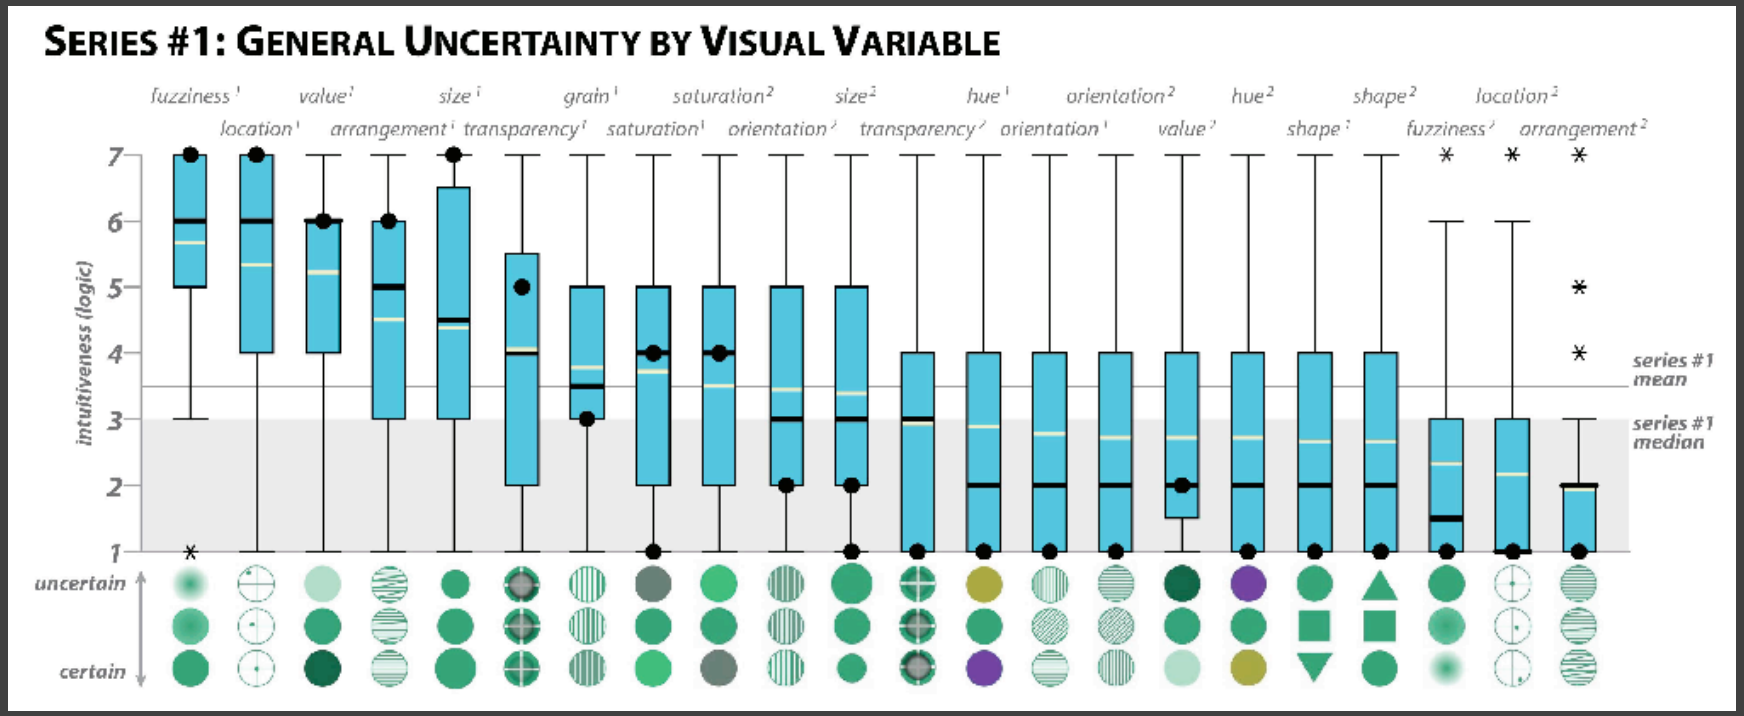
SERIES #1: GENERAL UNCERTAINTY BY VISUAL VARIABLE
fuzziness 1
location 1
value 1
arrangement 1
size 1
transparency 1
grain 1
saturation 1
saturation 2
orientation 2
size 2
transparency 2
hue 1
orientation 1
orientation 2
value 2
hue 2
shape 1
shape 2
fuzziness 2
location 2
arrangement 2
intuitiveness (logic)
7
6
5
4
3
2
1
series #1 mean
series #1 median
uncertain
certain

Fuzziness Juxtaposition

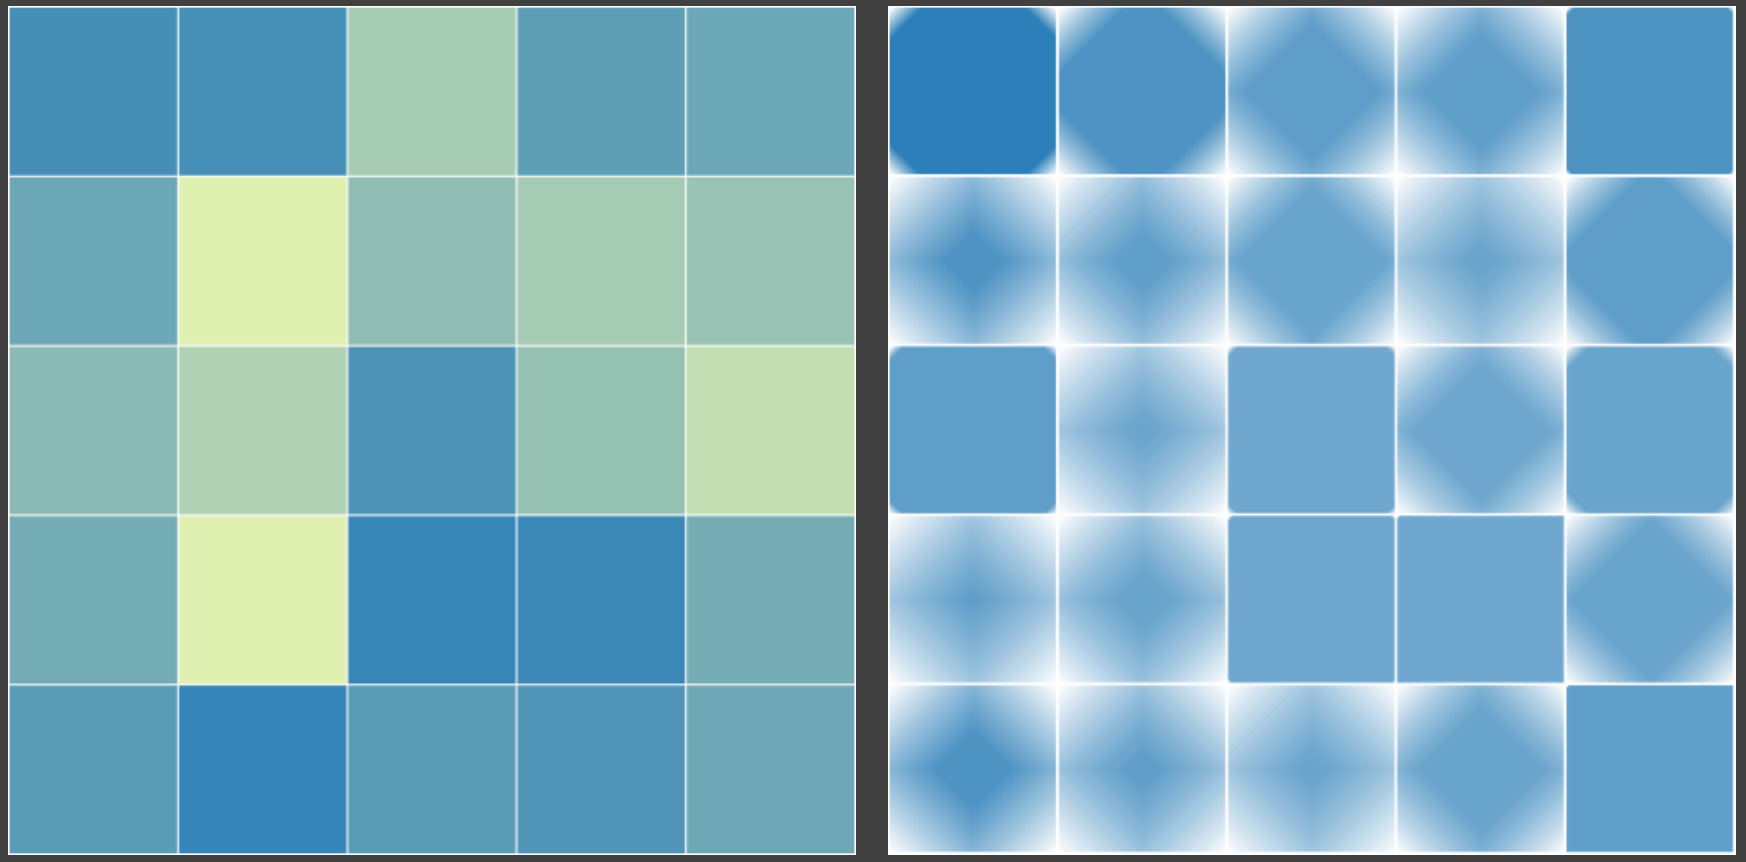

Fuzziness Superposition

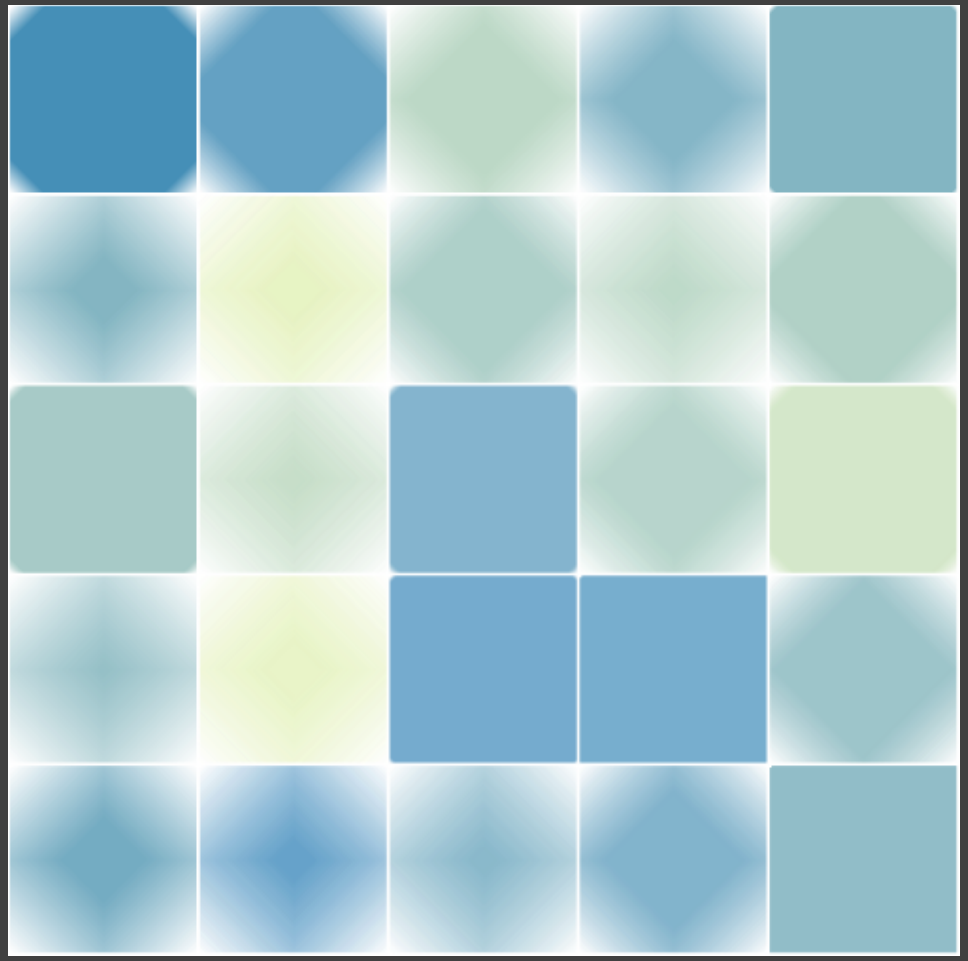

Size Juxtaposition

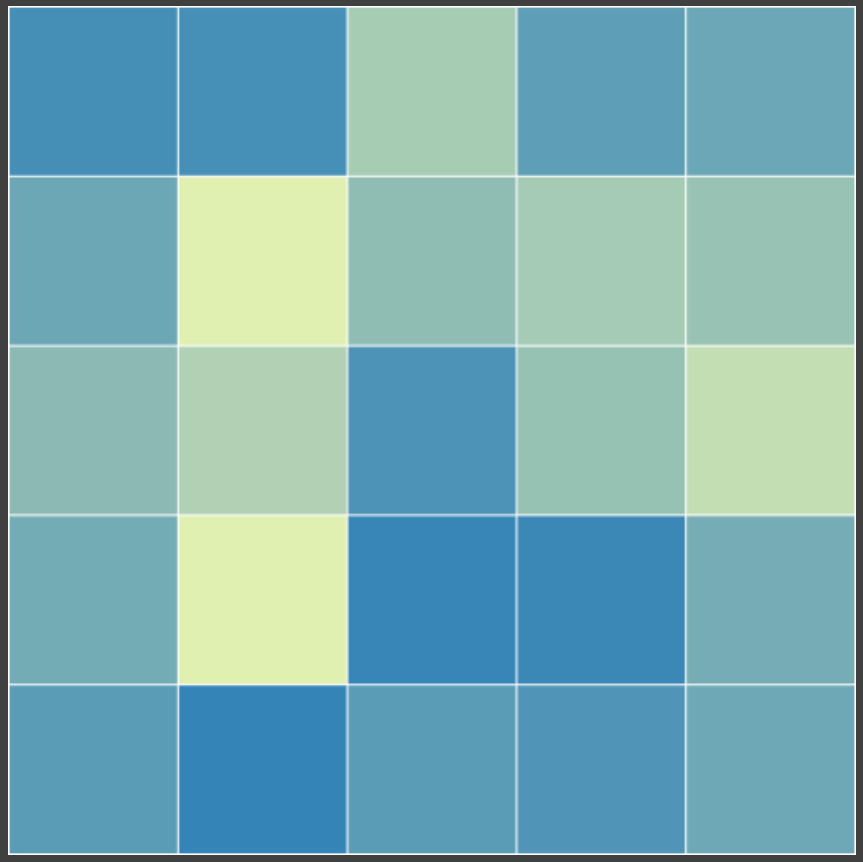

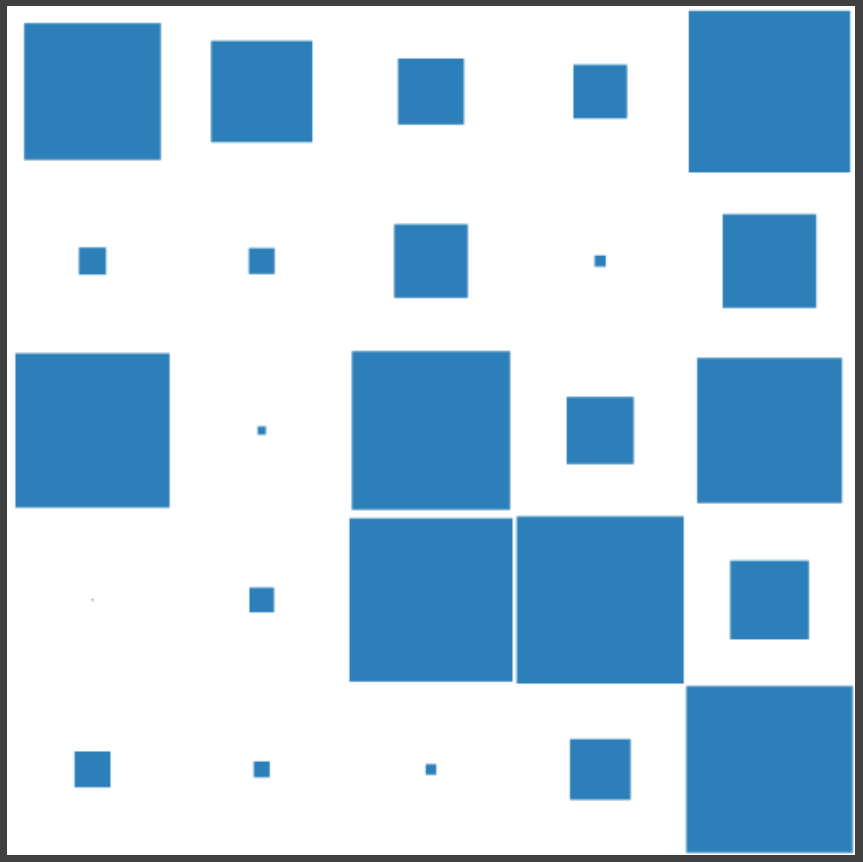

Size Superposition

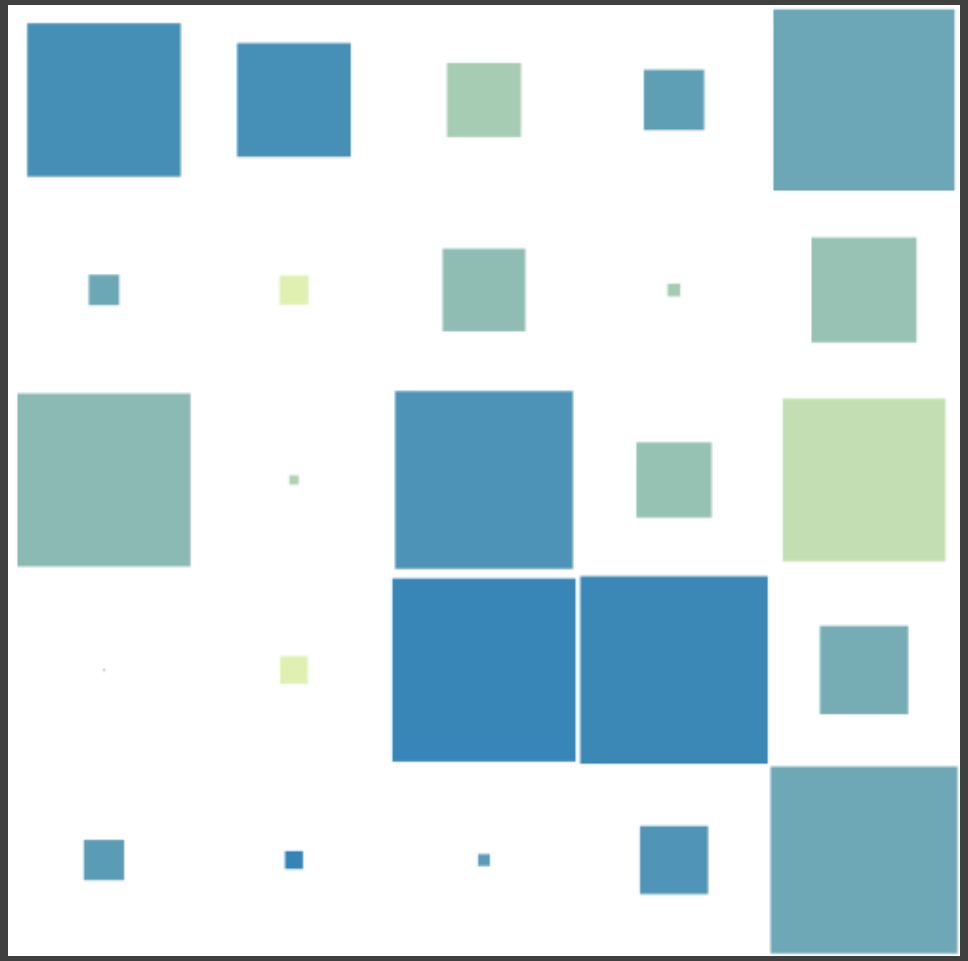

“Sketchiness”

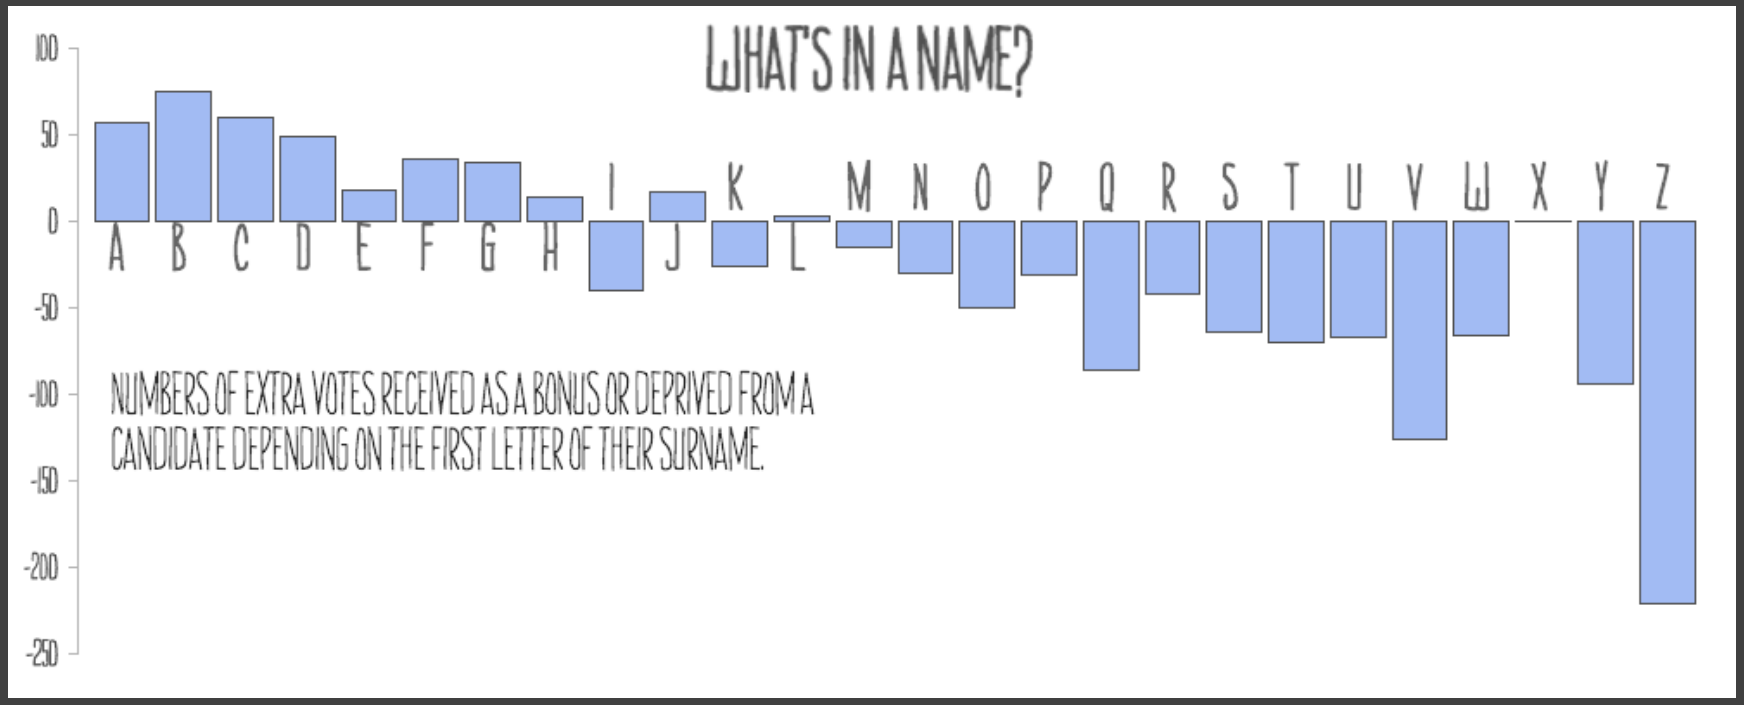

Wood, Jo et al. Sketchy rendering for information visualization. IEEE VIS, 2012.

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

“Sketchiness”

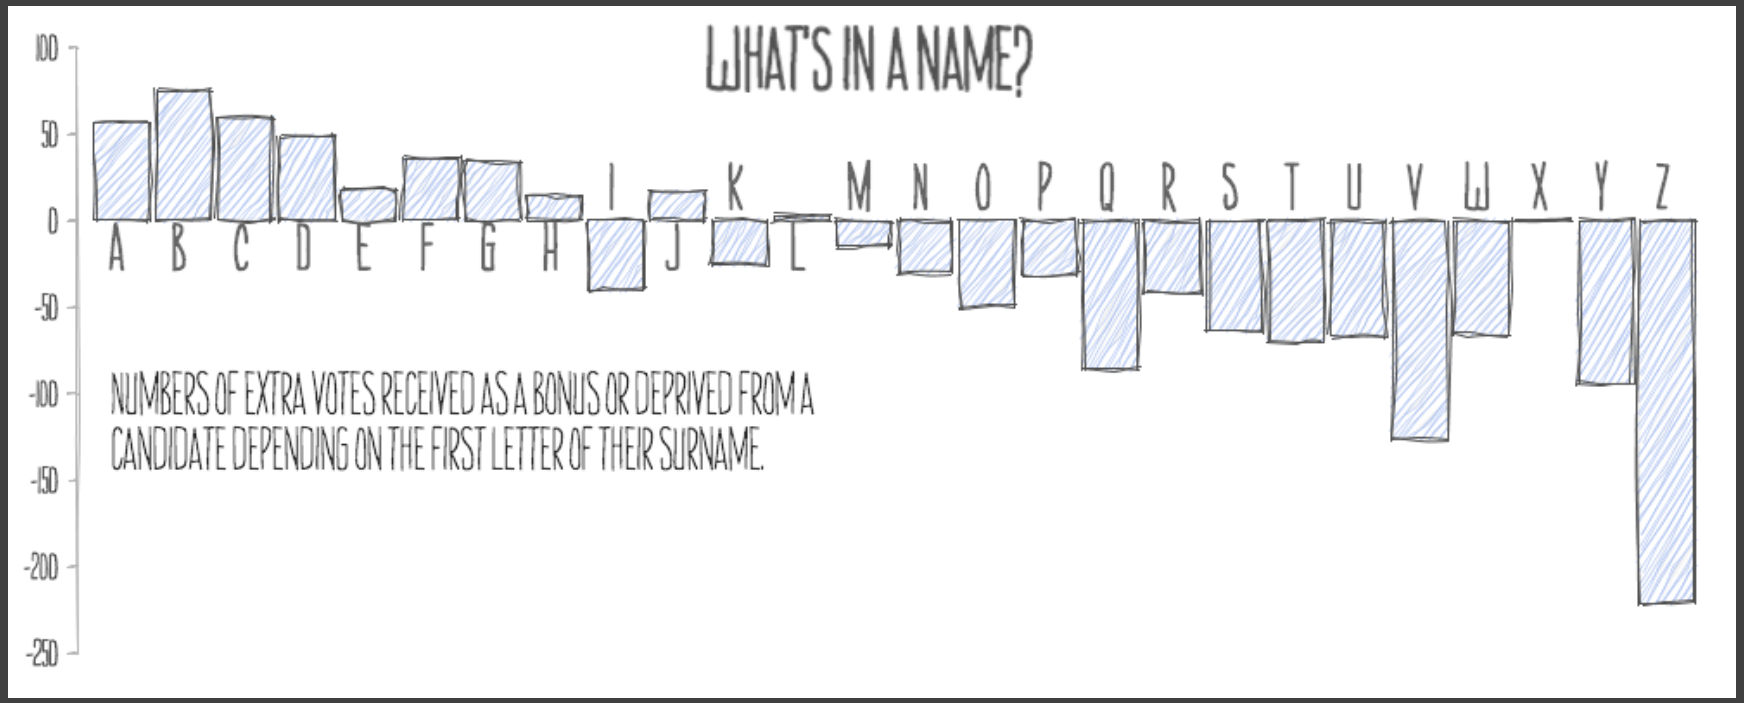

Wood, Jo et al. Sketchy rendering for information visualization. IEEE VIS, 2012.

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

Perceived Data Quality

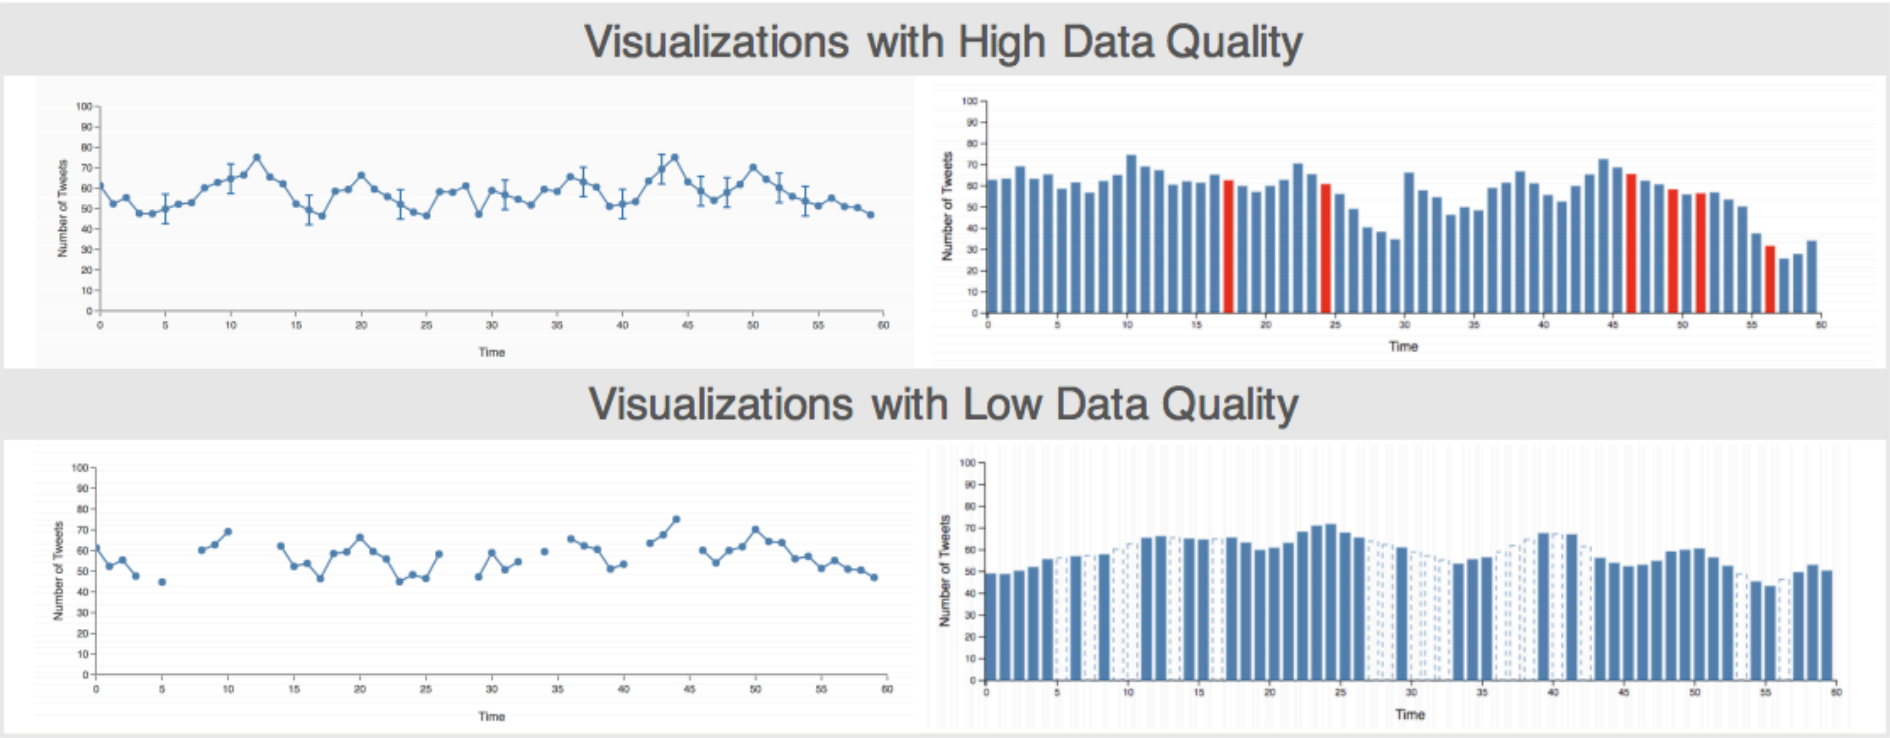

Song and Szafir. Where's My Data? Evaluating Visualizations with Missing Data. IEEE VIS, 2018.

Perceived Data Quality

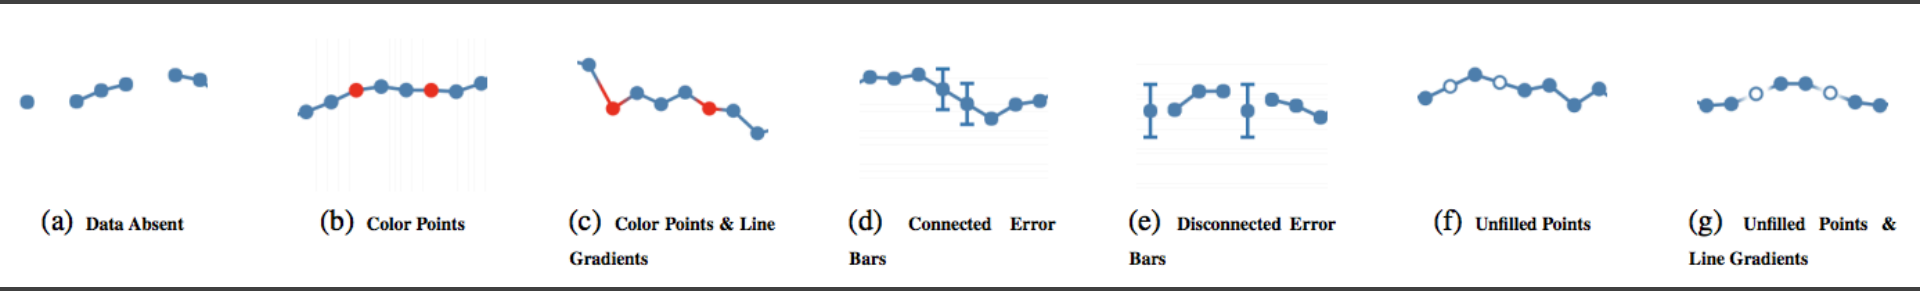

Error & perceived quality decrease with more missing data

Highlighting missing data increased perceived quality

Linear interpolation led to highest perceived quality

Absent data led to lower perceived quality, credibility, confidence

Limited evidence for accuracy bias from imputation methods

Song and Szafir. Where's My Data? Evaluating Visualizations with Missing Data. IEEE VIS, 2018.

Encoding Uncertainty

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

However, intuitive variables may not always be accurately interpreted!

Model Visualization

Polling Data

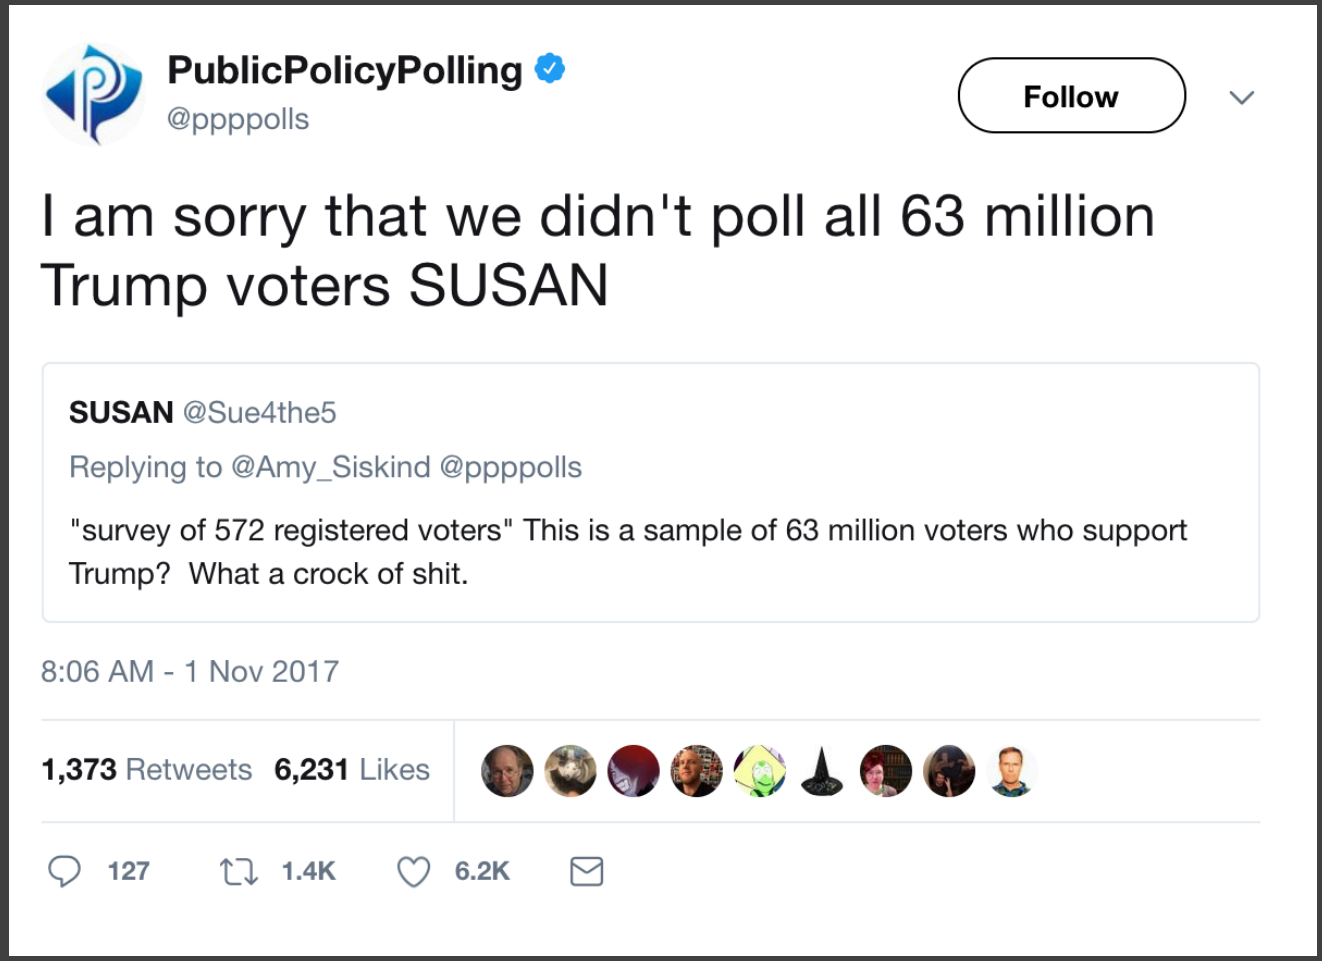

PublicPolicyPolling ✓
@ppppolls

Follow

I am sorry that we didn't poll all 63 million Trump voters SUSAN

> **SUSAN** @Sue4the5
> Replying to @Amy_Siskind @ppppolls
> "survey of 572 registered voters" This is a sample of 63 million voters who support Trump? What a crock of shit.

8:06 AM - 1 Nov 2017

1,373 Retweets **6,231** Likes

127 1.4K 6.2K

The NYT Needle

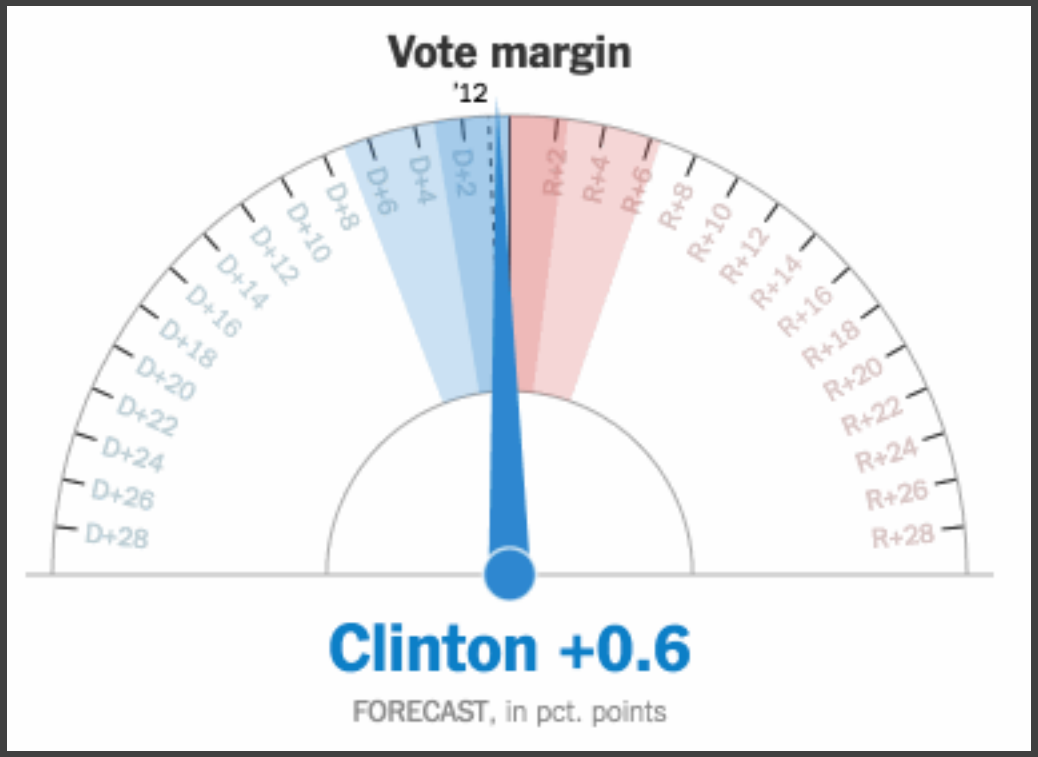

ELECTION RESULTS BY SEARCHLIGHT.

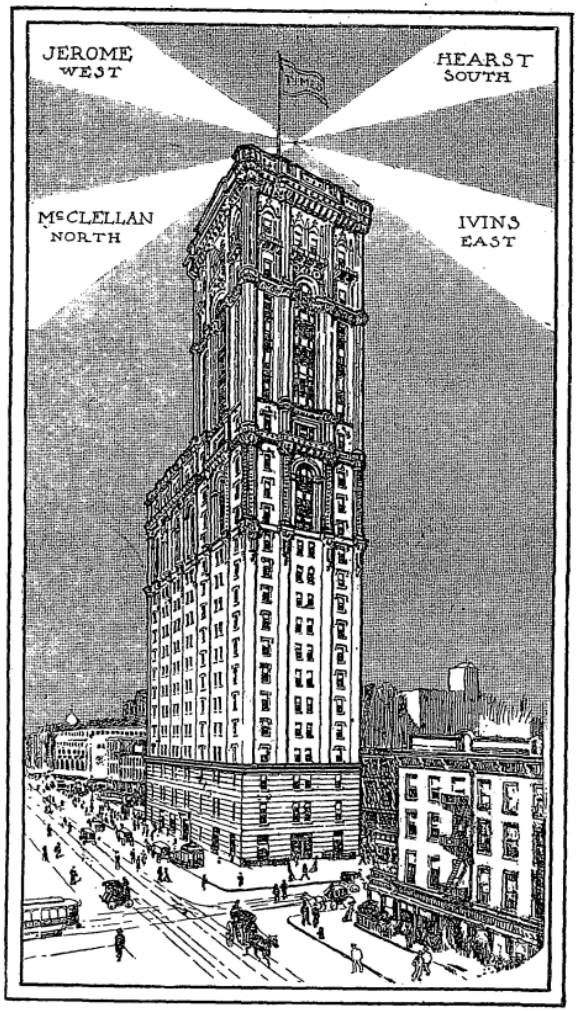

The Times Election Searchlight Code.

News Will Be Flashed from the Tower of The Times Building on Tuesday Night.

The results of the election next Tuesday night will be flashed by electric light from the tower of the Times Building, so that for miles around people will be able to tell which of the candidates has won.

This will be entirely separate and distinct from the elaborate bulletin service which THE TIMES will also maintain. To display the detailed bulletins so that the crowds can see them easily and comfortably, a stereopticon machine will be set up in the triangle north of the Times Building and the bulletins displayed on canvas stretched from the north side of the building. There will be a similar service at the Harlem office of THE TIMES, 129 West 125th Street.

The electric signals from the tower of the Times Building will be flashed from a point 365 feet above the street level. A steady light to the north will show that McClellan has been elected; a steady light to the east will indicate Ivins's election, and a steady light to the south will indicate that Hearst has won.

Jerome's election will be indicated by a steady light to the west. A light to the north, waving from east to west, will indicate Osborne's election. A light to the south, waving from east to west, will indicate Shearn's election.

ELECTION RESULTS BY SEARCHLIGHT.

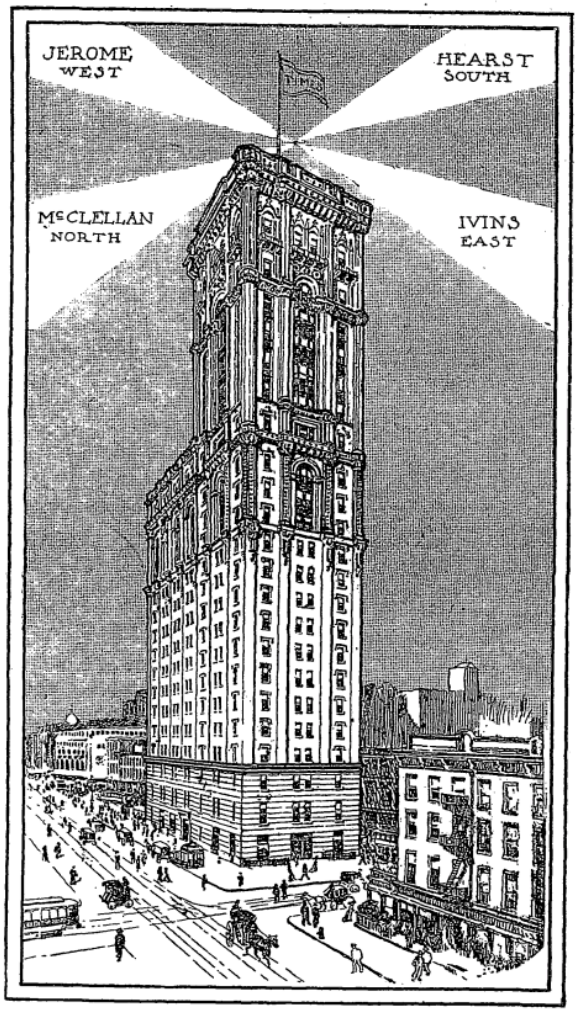

The Times Election Searchlight Code.

News Will Be Flashed from the Tower of The Times Building on Tuesday Night.

The results of the election next Tuesday night will be flashed by electric light from the tower of the Times Building, so that for miles around people will be able to tell which of the candidates has won.

This will be entirely separate and distinct from the elaborate bulletin service which THE TIMES will also maintain. To display the detailed bulletins so that the crowds can see them easily and comfortably, a stereopticon machine will be set up in the triangle north of the Times Building and the bulletins displayed on canvas stretched from the north side of the building. There will be a similar service at the Harlem office of THE TIMES, 129 West 125th Street.

The electric signals from the tower of the Times Building will be flashed from a point 365 feet above the street level. A steady light to the north will show that McClellan has been elected; a steady light to the east will indicate Ivins's election, and a steady light to the south will indicate that Hearst has won.

Jerome's election will be indicated by a steady light to the west. A light to the north, waving from east to west, will indicate Osborne's election. A light to the south, waving from east to west, will indicate Shearn's election.

Election Bulletins
BY BOMBS.

TUESDAY NIGHT
THE TRIBUNE
will send up from the roof of the
GREAT NORTHERN HOTEL

hourly, shells containing blue and red stars—exactly on the hour—at 7, 8, 9, 10, 11 p. m. 12 midnight, 1 and 2 a. m. Wednesday morning, unless election is decided earlier, in which case twelve bombs will be sent up in rapid succession. Blue to indicate McKinley's election. Red to indicate Bryan's election.

SIX BOMBS EVERY HOUR.

The first bomb sent up, if blue, indicates the returns in **COOK COUNTY** at that hour are favorable to McKinley; if red, favorable to Bryan.

After sixty seconds two bombs will be sent up in rapid succession, and will indicate, if blue, that returns from **ILLINOIS** favor McKinley; if red, Bryan.

After sixty seconds more three bombs will be sent up in rapid succession, and if blue will indicate that at that hour returns from the **entire country** favor McKinley; if red, Bryan. Each bomb bursts high in the air, scattering a shower of stars.

Value-Suppressing Uncertainty Palette

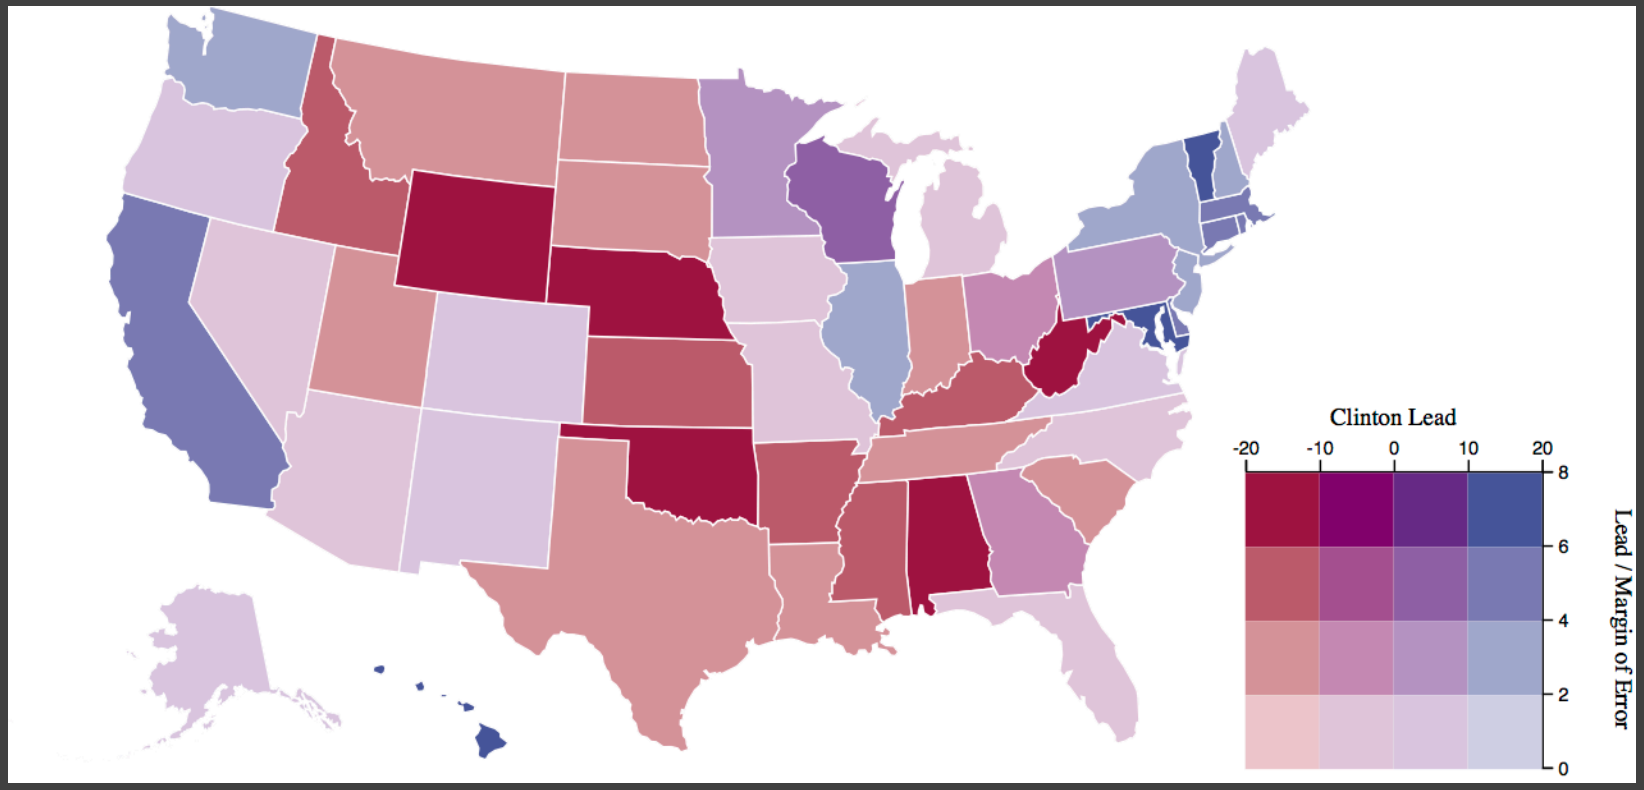

Bivariate Map

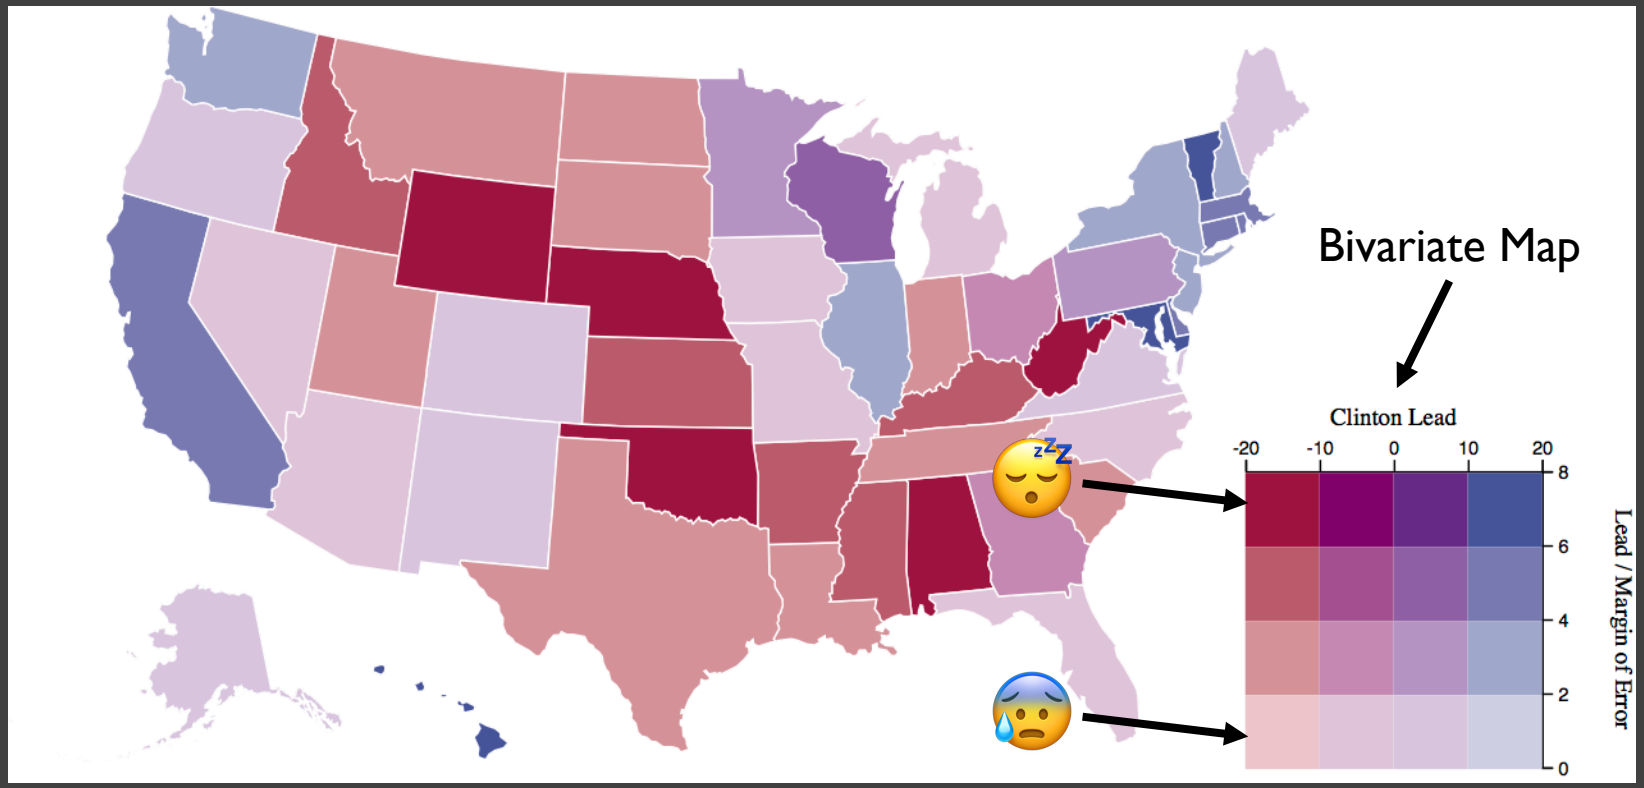

Value-Suppressing Uncertainty Palette

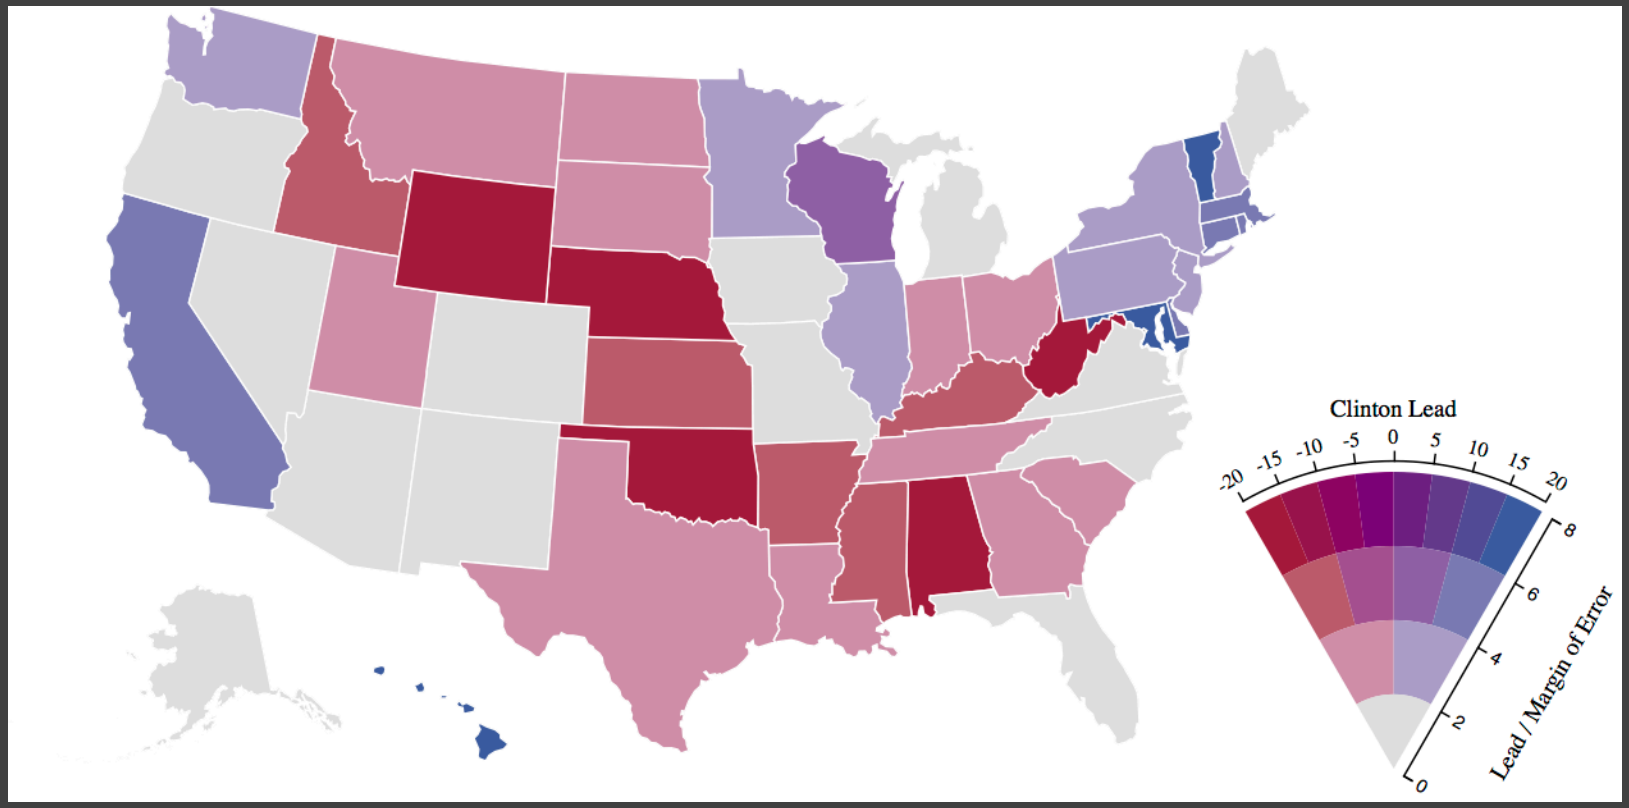

Correll, Moritz & Heer. "Value-Suppressing Uncertainty Palettes." CHI 2018.

Value-Suppressing Uncertainty Palette

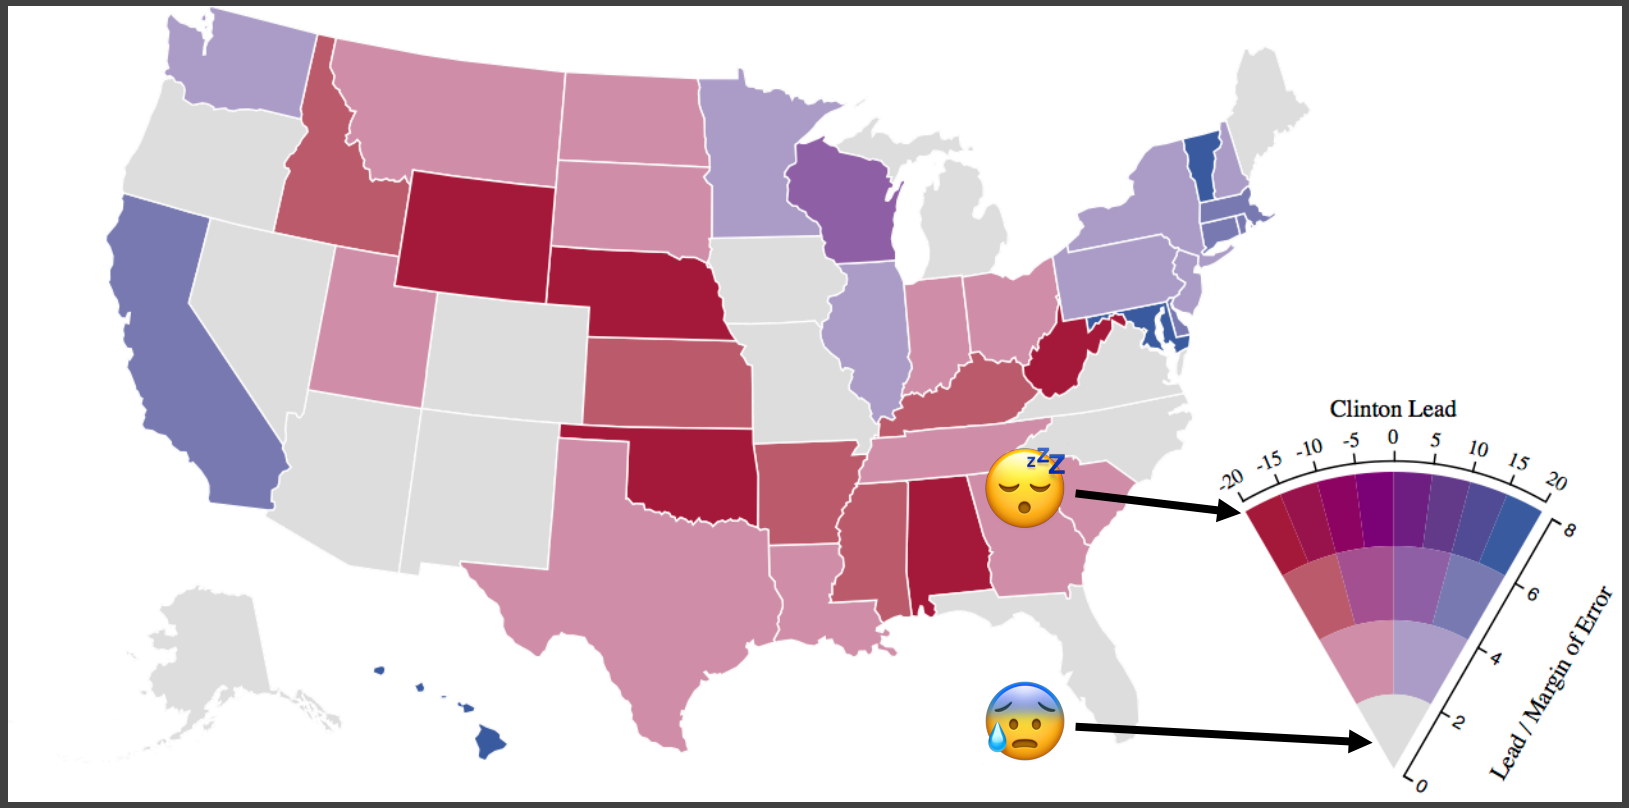

Correll, Moritz & Heer. "Value-Suppressing Uncertainty Palettes." CHI 2018.

Model Visualization

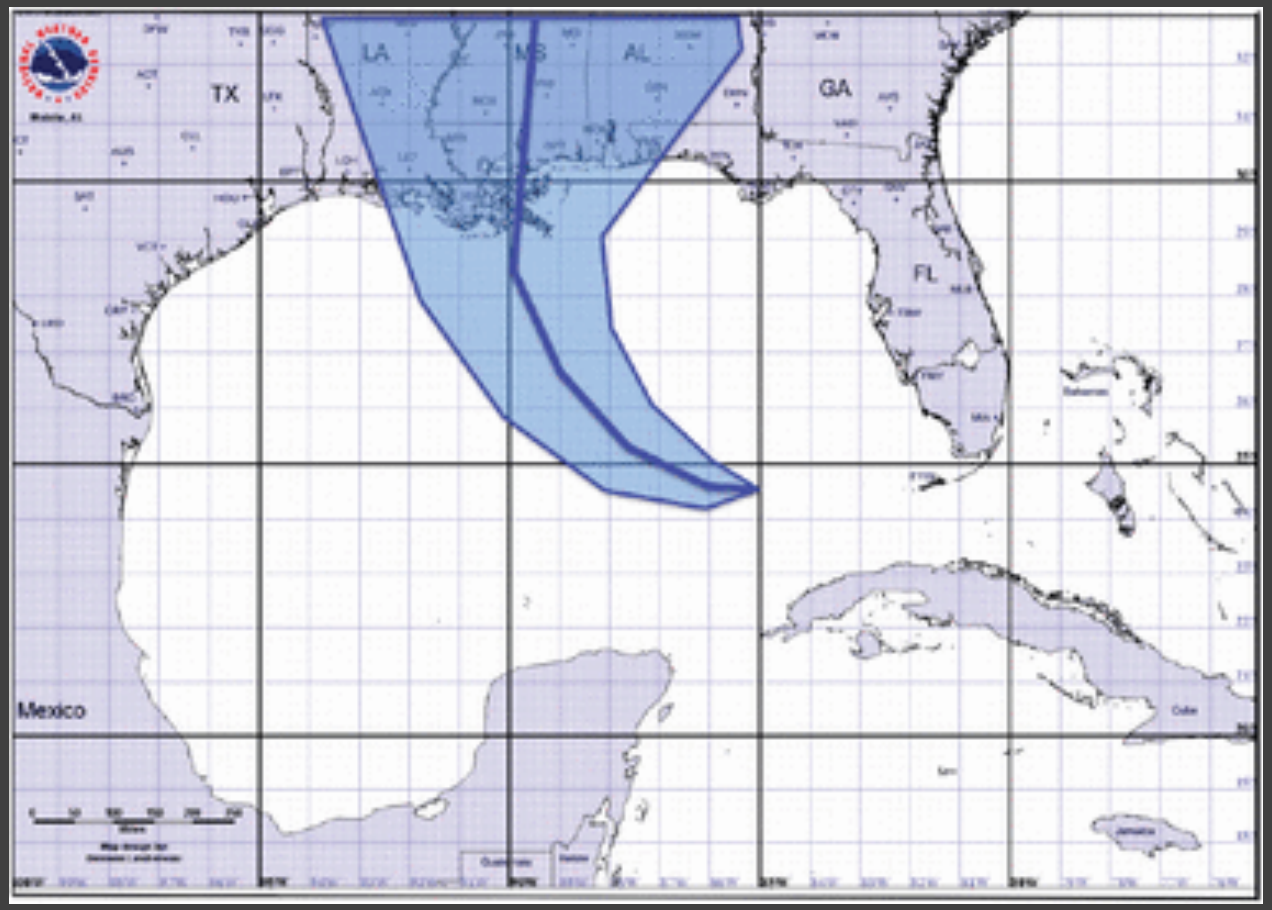

Cox, Jonathan and House, Donald and Lindell, Michael. Visualising uncertainty in predicted hurricane tracks. International Journal for Uncertainty Quantification, 2013.

Model Visualization

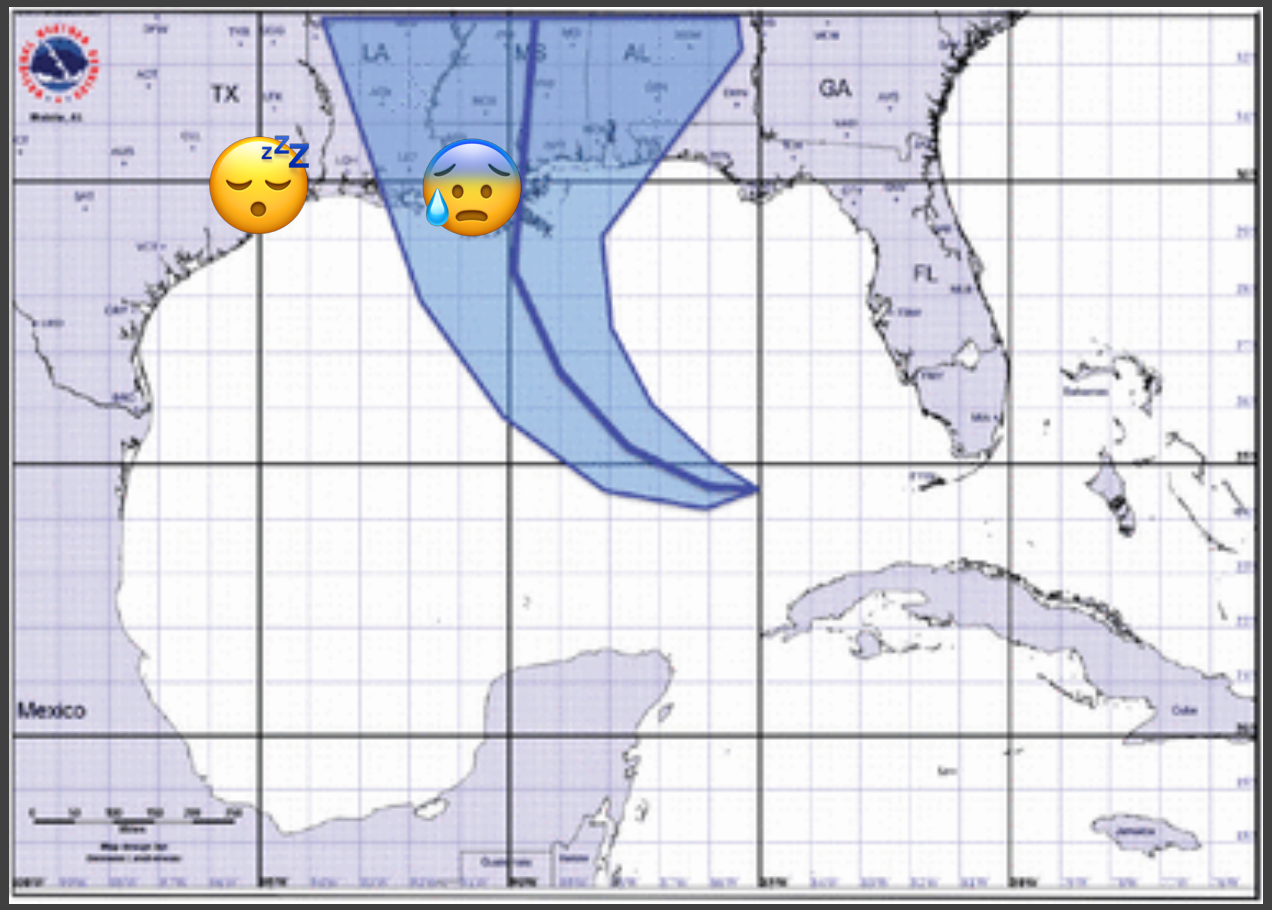

Cox, Jonathan and House, Donald and Lindell, Michael. Visualising uncertainty in predicted hurricane tracks. International Journal for Uncertainty Quantification, 2013.

Model Visualization

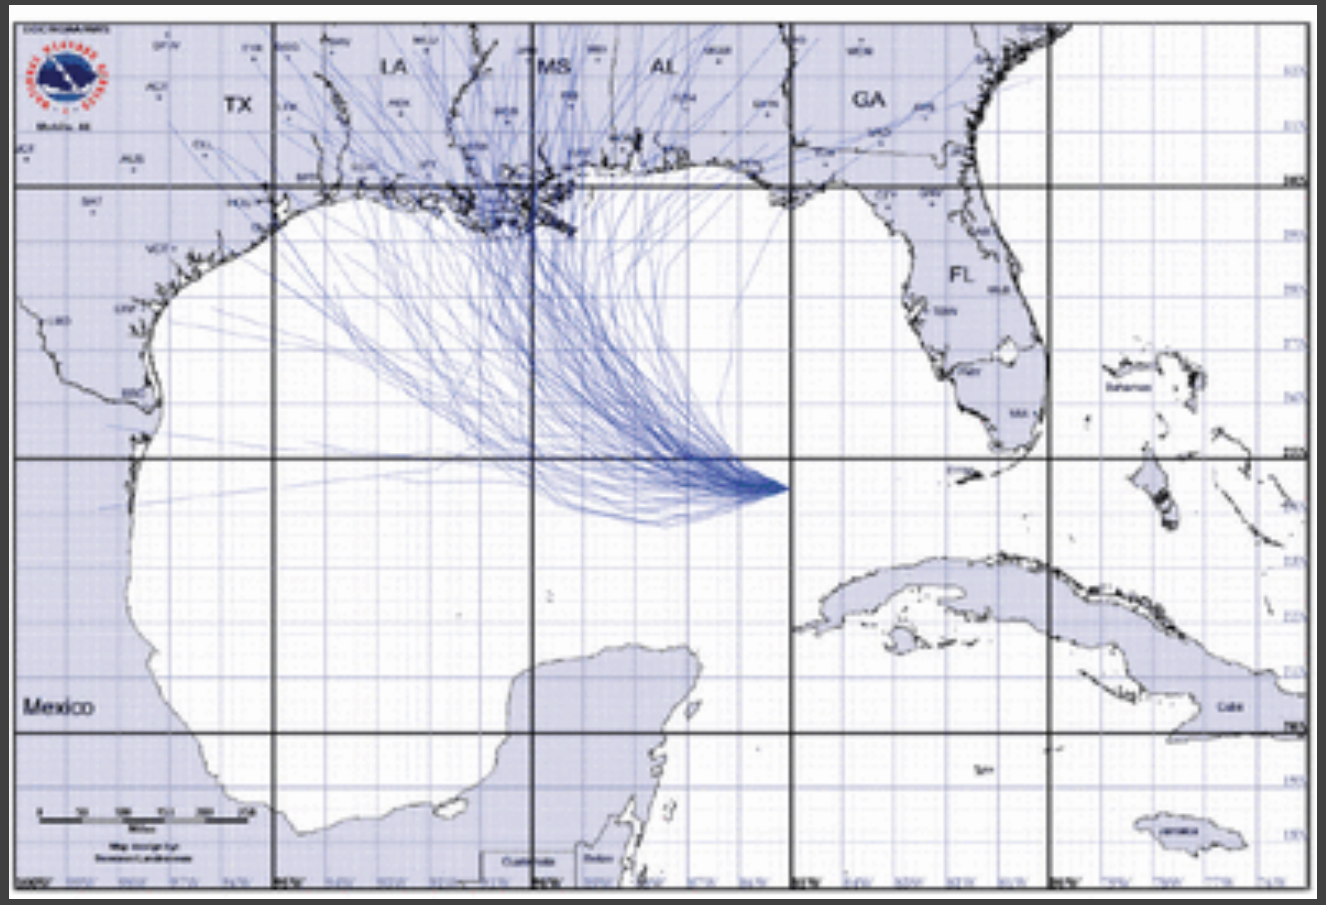

Model Visualization

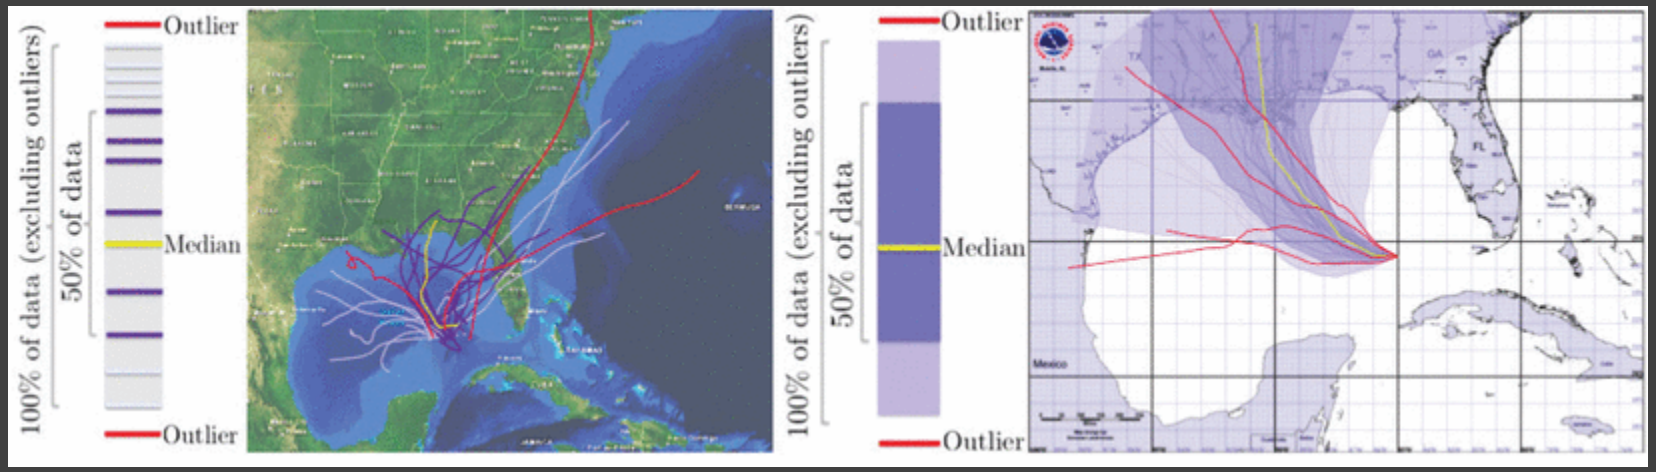

M. Mirzargar, R. Whitaker and R. Kirby. Curve Boxplot:
Generalization of Boxplot for Ensembles of Curves. IEEE VIS 2014.

Predicted Bus Arrival Times

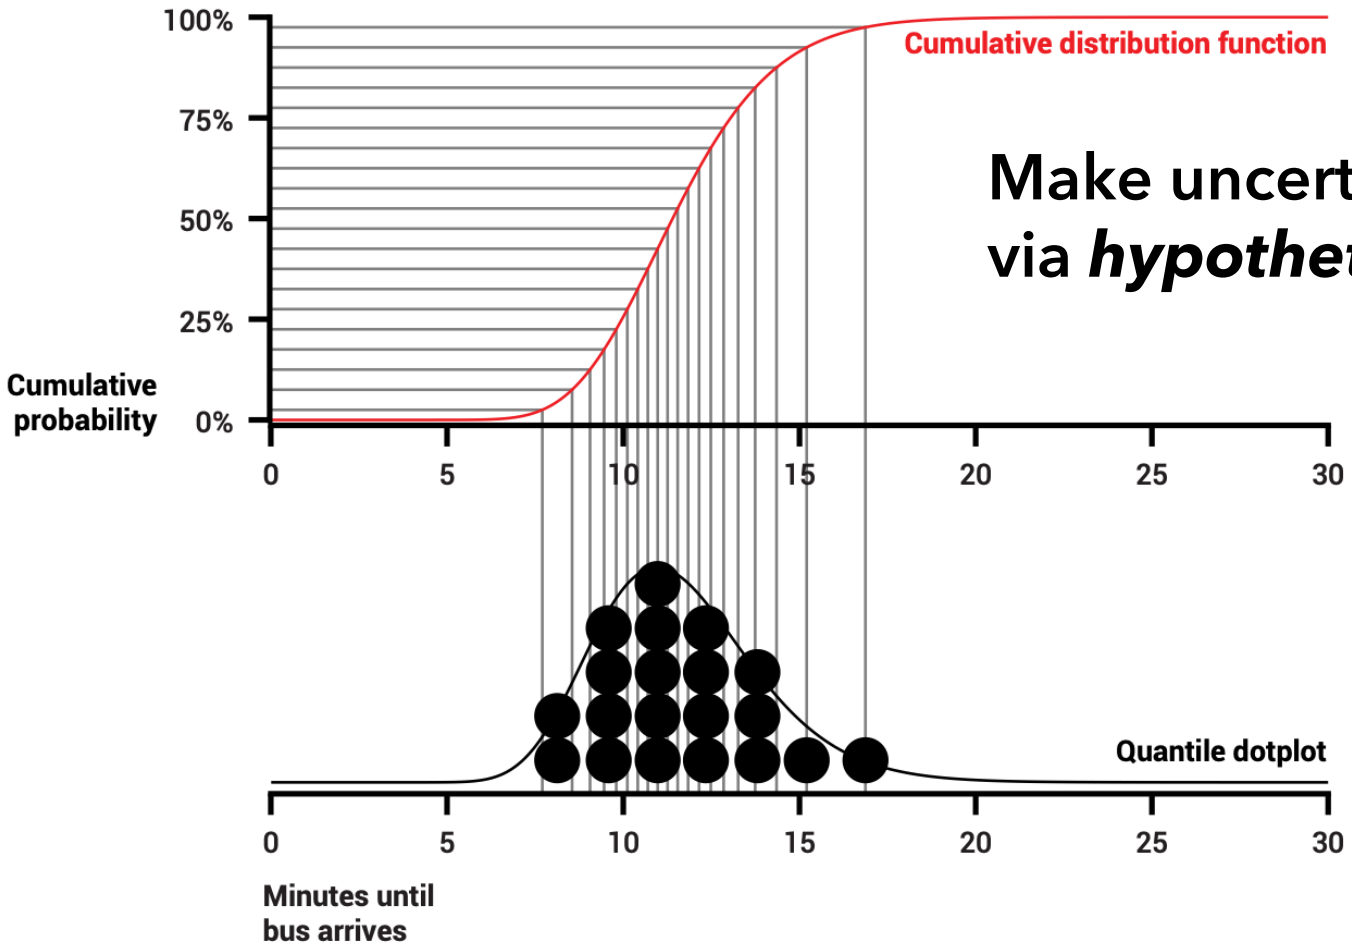

Make uncertainty concrete via ***hypothetical outcomes***.

[Kay et al. 2016]

Predicted Bus Arrival Times

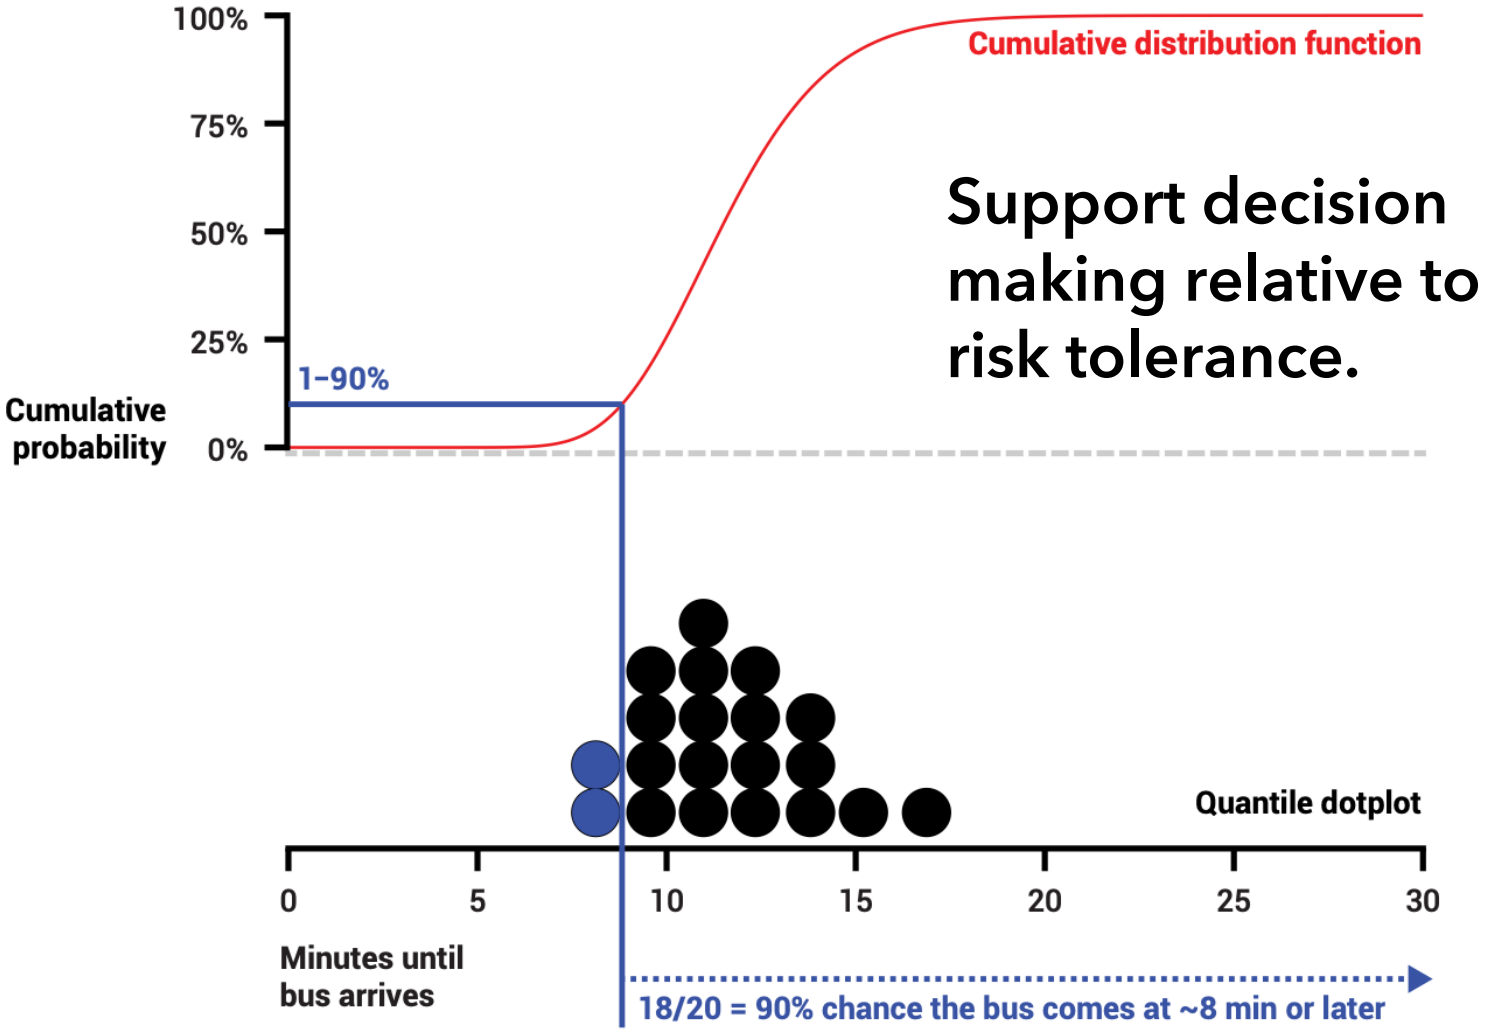

Support decision making relative to risk tolerance.

[Kay et al. 2016]

Predicted Bus Arrival Times

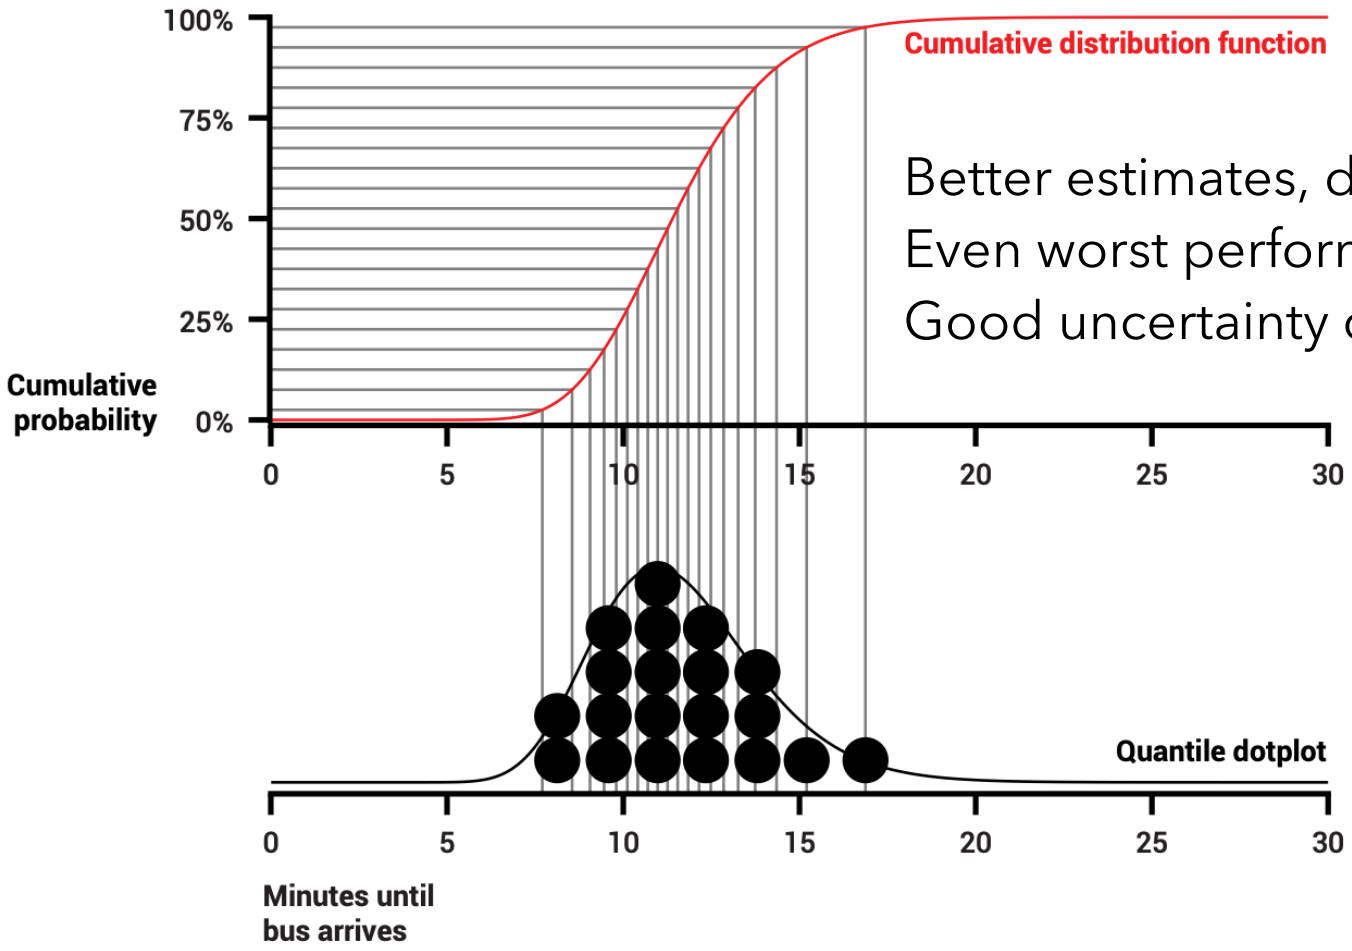

Better estimates, decisions with time.
Even worst performers improve.
Good uncertainty displays possible!

[Kay et al. 2016]

Hypothetical Outcome Plots

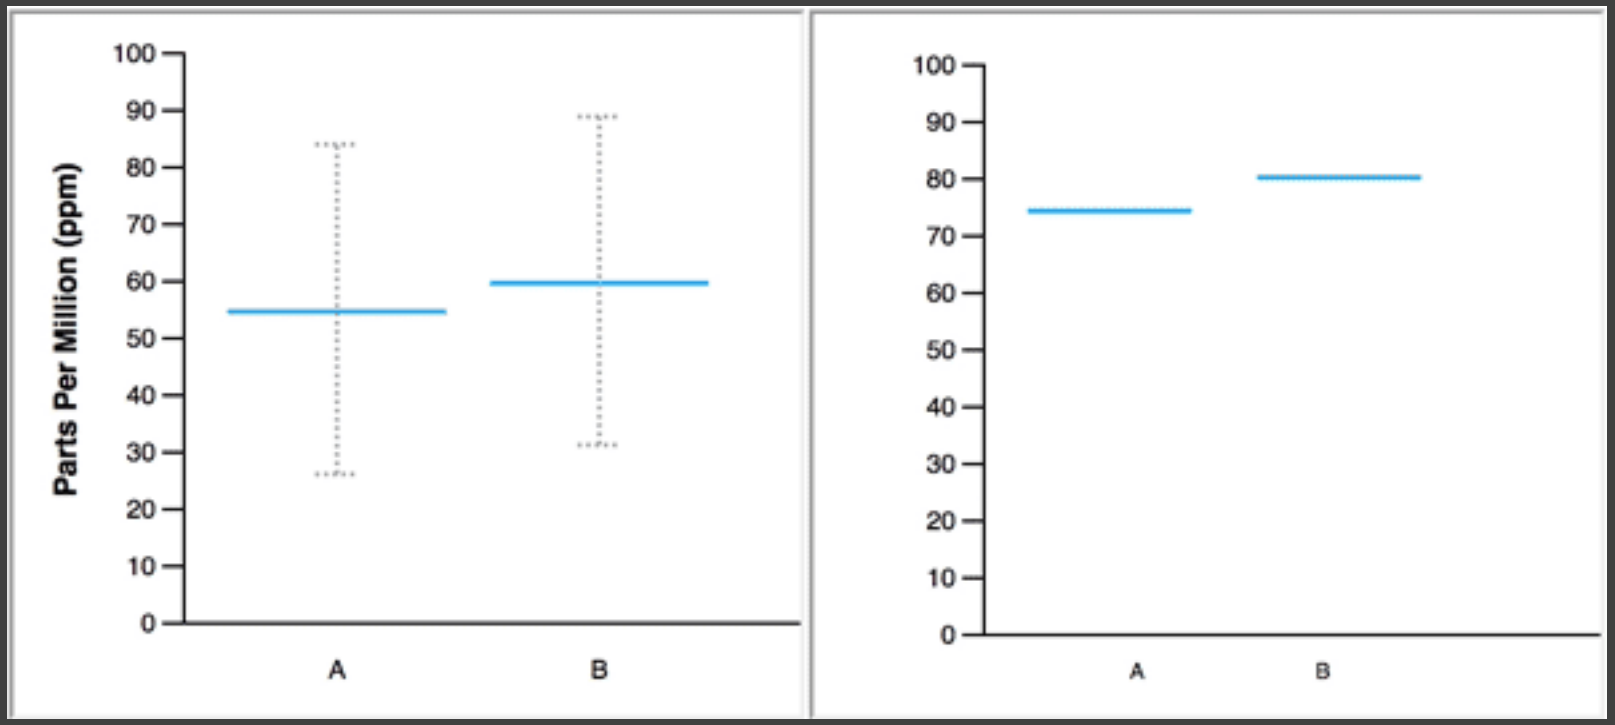

Life Expectancy

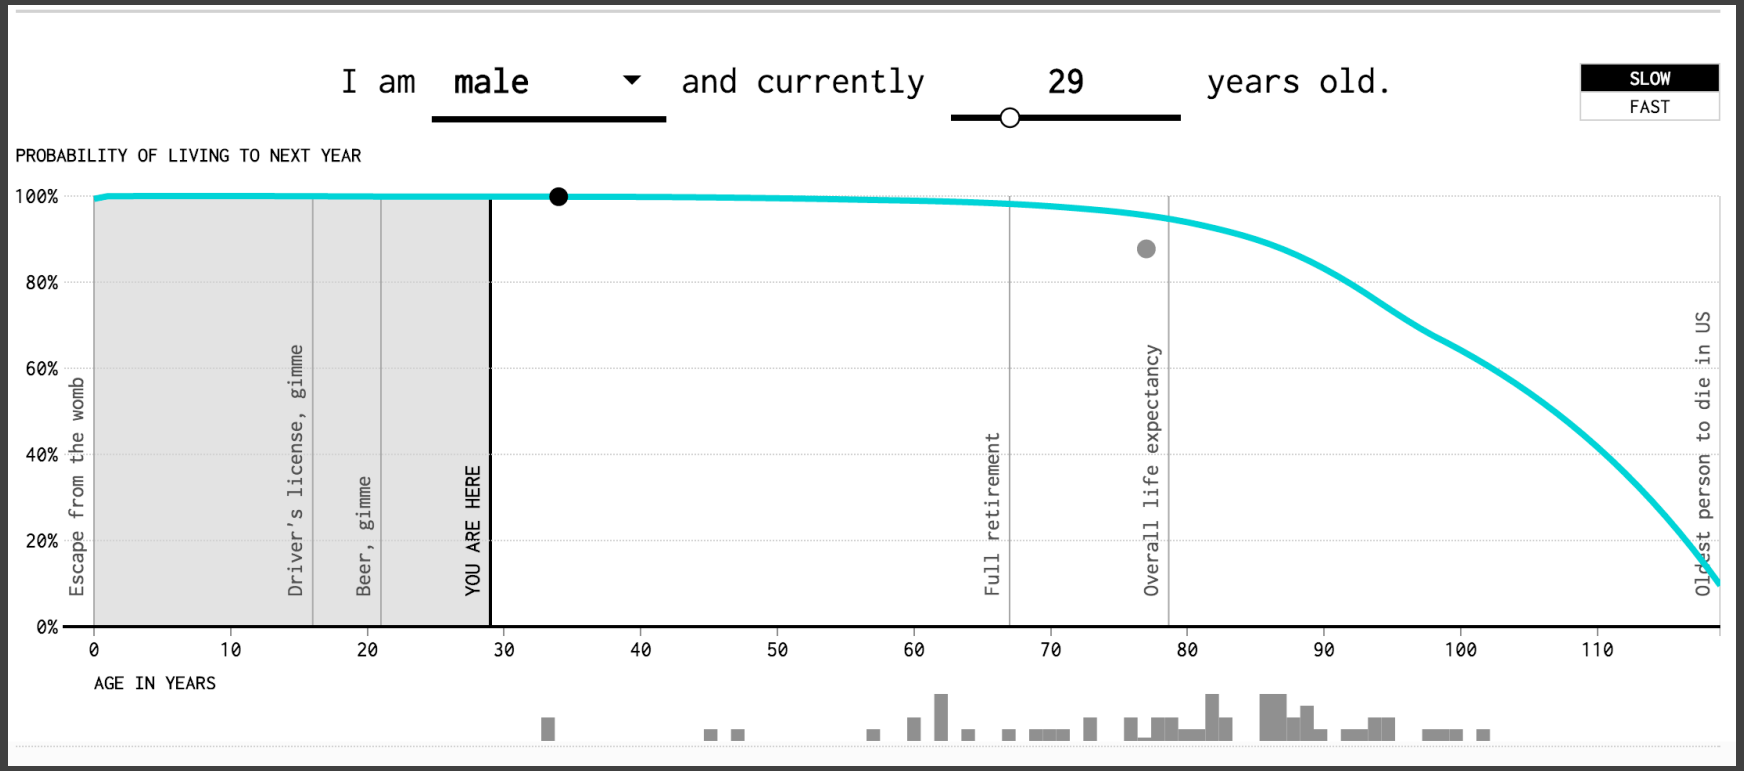

Model Visualization

Building models is necessary to quantify uncertainty

It is important to communicate the variability in model outcomes

Dynamic or ensemble displays can help communicate complex models

How Should I Visualize Uncertainty?

Choose an appropriate visual variable based on the domain, literacy, and expertise of your audience. Be mindful that any display of uncertainty inherently increases the complexity of your visualization, and that there is a preference/performance gap.

How Should I Visualize Uncertainty?

Choose an appropriate visual variable based on the domain, literacy, and expertise of your audience. Be mindful that any display of uncertainty inherently increases the complexity of your visualization, and that there is a preference/performance gap.

IT DEPENDS

Administrivia

Final Project Schedule

Proposal	Wed, May 18
Milestone	Fri, May 27
Poster	Wed, June 1
Deliverables	Tue, June 7

Logistics

Final project description posted online

Groups of up to 4 people

Select topics and form groups now!

Potential project ideas shared on edstem!

Reminders!

Virtual Guest Lecture This Thursday 5/19!

Final Project Proposal Due Wed 5/18, 11:59pm
https://courses.cs.washington.edu/courses/cse512/22sp/fp.html

Three Peer Evaluations Due Tue 5/17, 11:59pm
https://courses.cs.washington.edu/courses/cse512/22sp/a3-peereval.html

What Can Go Wrong?

Which Stock To Buy?

Company A

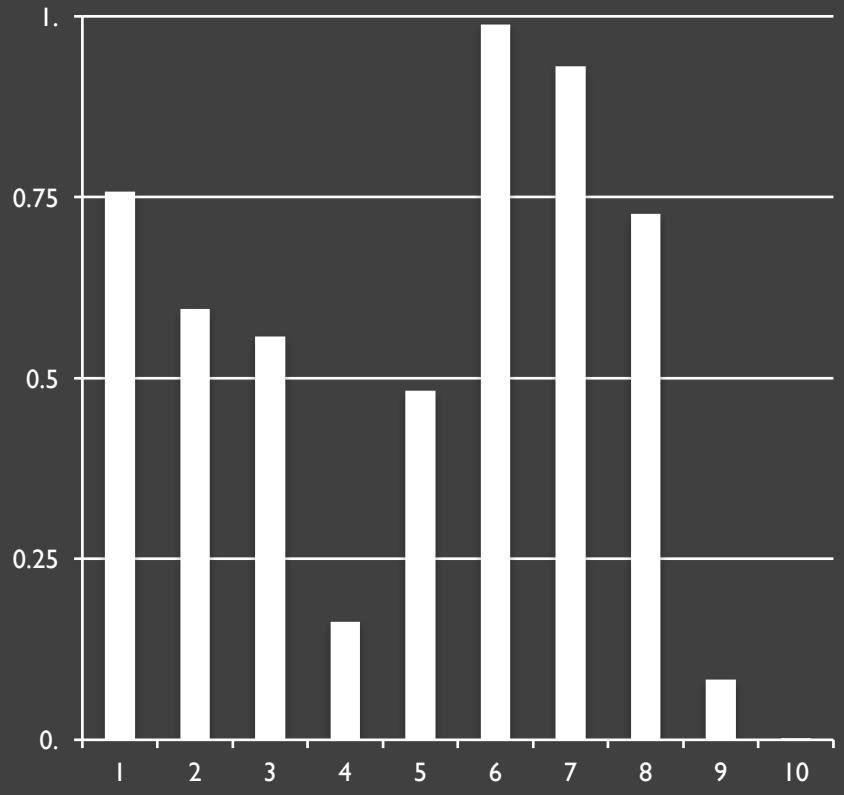

Company B

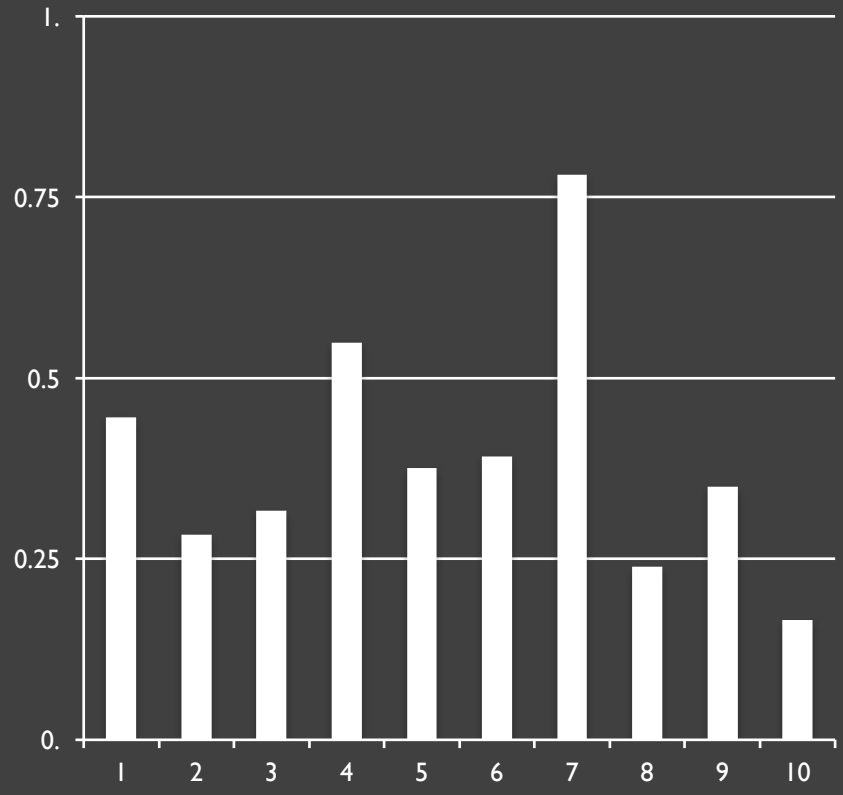

Neither!

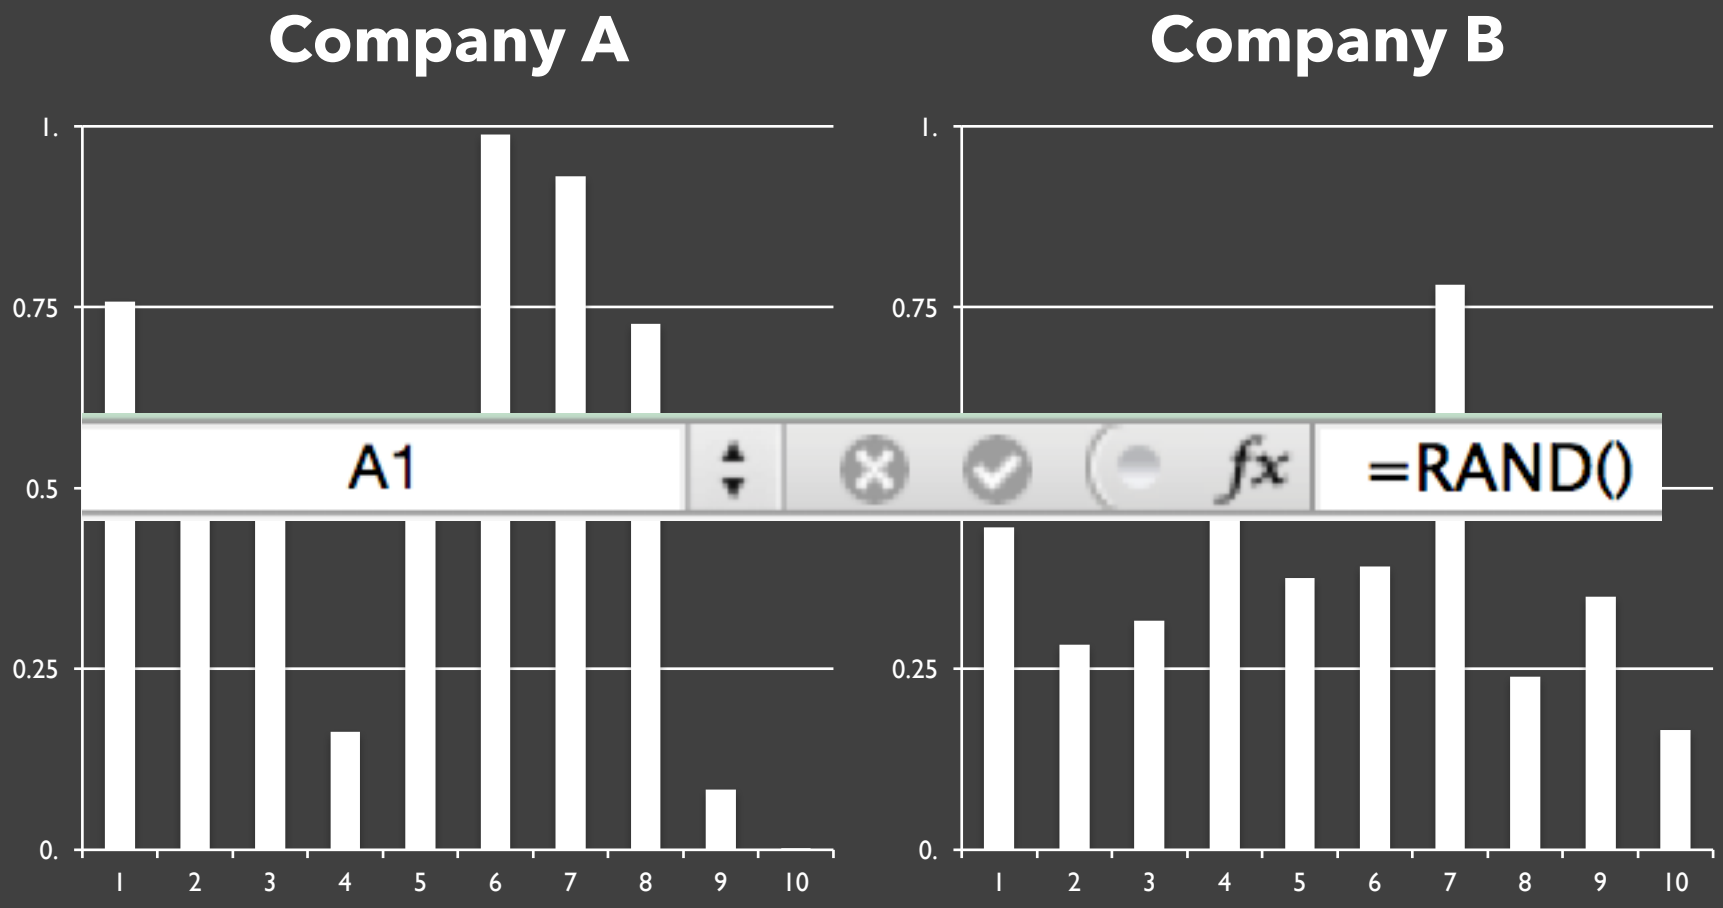

Pareidolia

Jobs Reports

If the economy actually added 150,000 jobs last month, it would be possible to see any of these headlines:

The jobs number is just an estimate, and it comes with uncertainty.

Under 55,000 jobs	55,000 to 110,000	110,000 to 140,000	160,000 to 190,000	190,000 to 245,000	245,000+
4% chance	19% chance	19% chance	19% chance	19% chance	4% chance

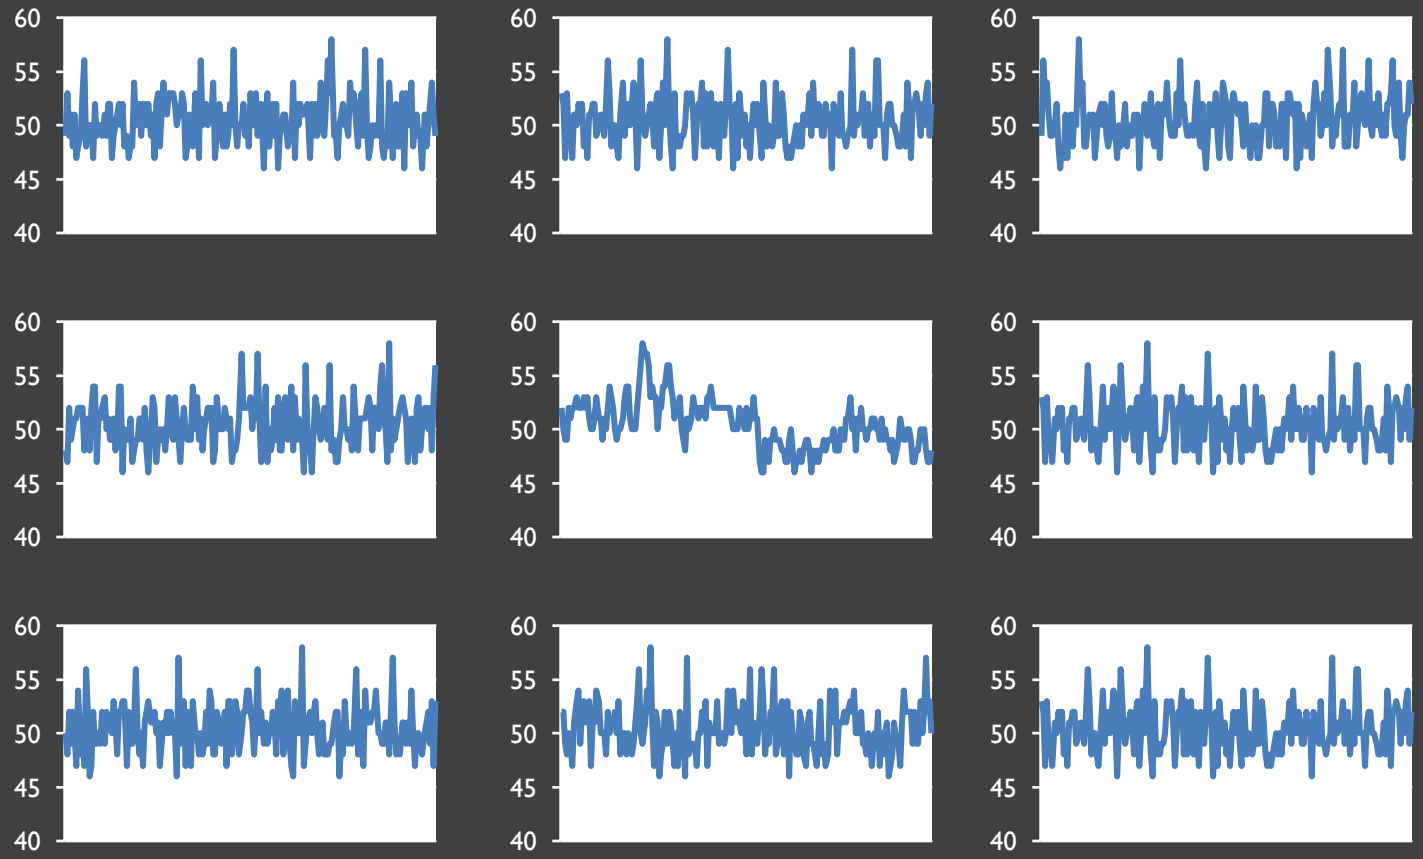
60
55
50
45
40

Visual Lineups

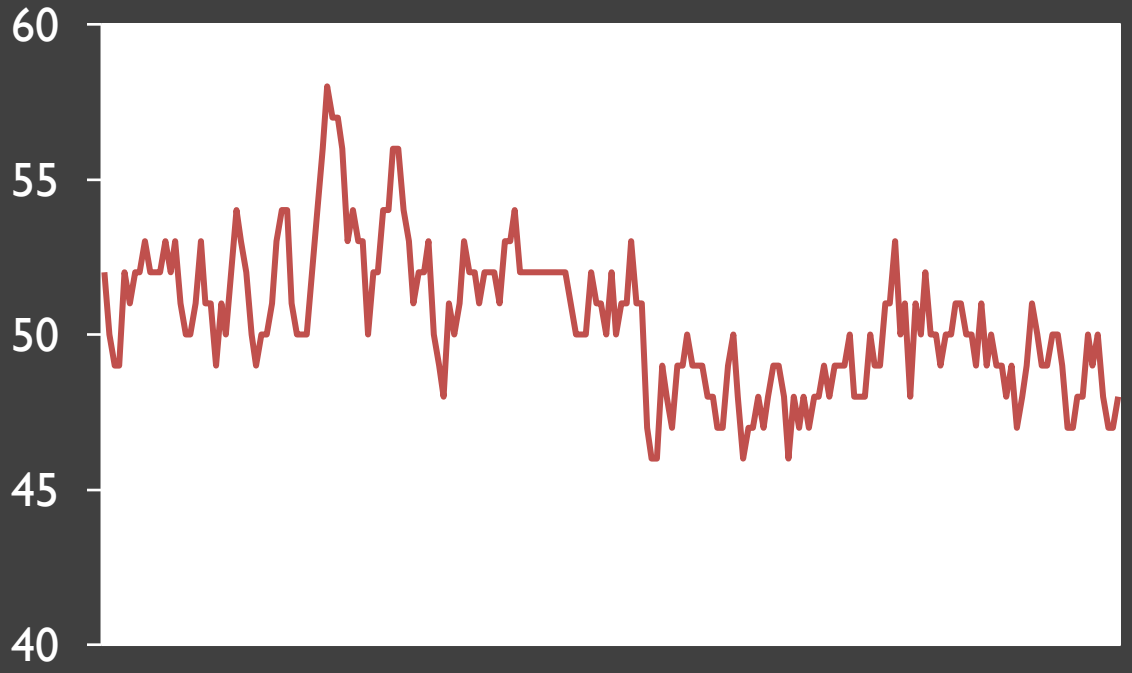

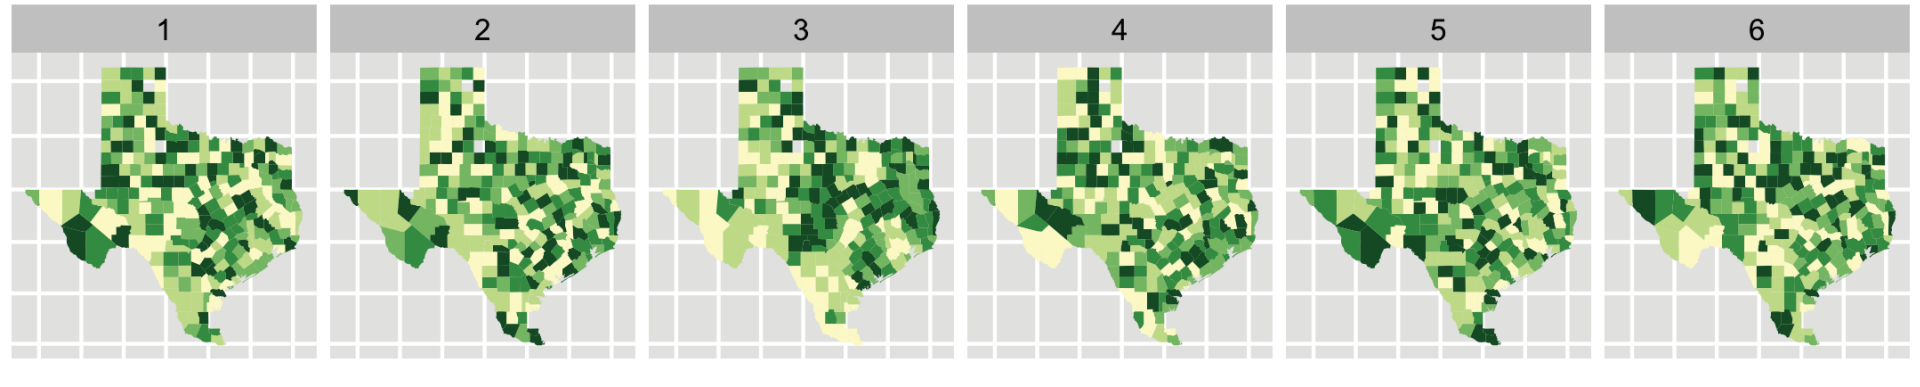

Choropleth maps of cancer deaths in Texas.

One plot shows a real data sets. 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.

Hadley Wickham et al. "Graphical inference for Infovis." IEEE transactions on visualization and computer graphics 16.6 (2010): 973–9.

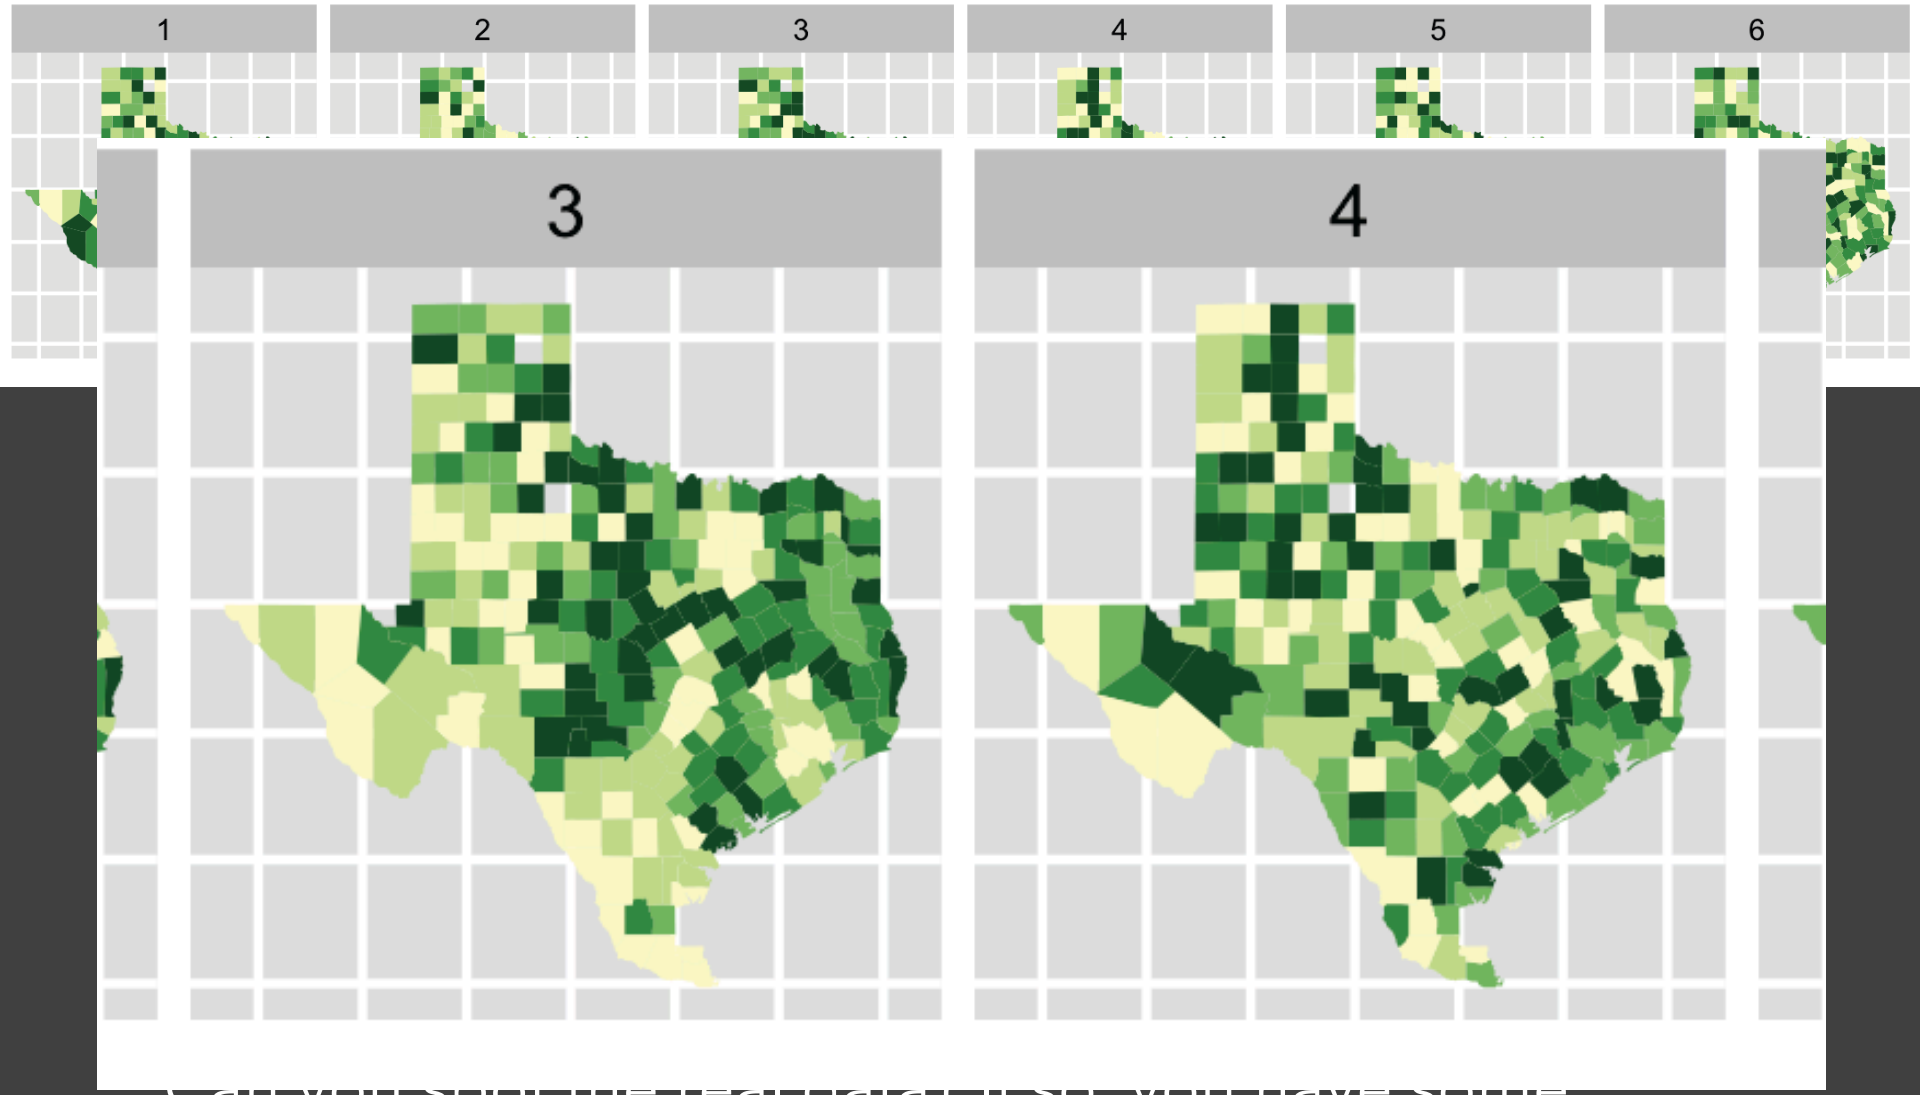

Can you spot the real data? If so, you have some evidence of spatial dependence in the data.

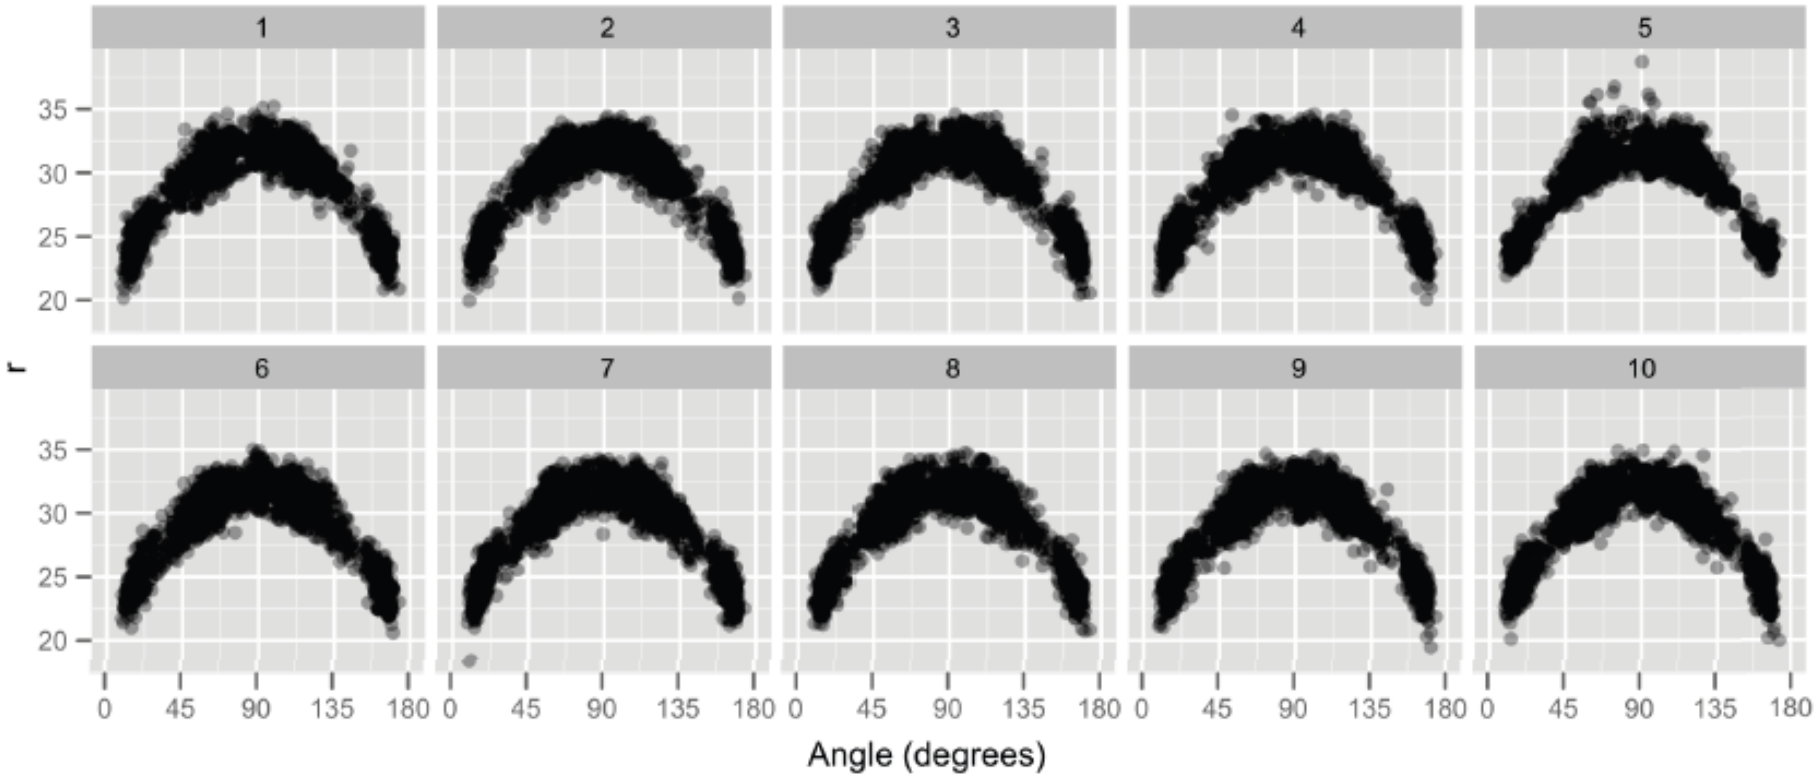

Distance vs. angle for 3 point shots by the LA Lakers.

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

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

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

$$P(+) = P(+ \wedge C)P(C) + P(+ \wedge \sim C)P(\sim C)$$

Bayes' Law

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

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

$P(+) = P(+ \wedge C)P(C) + P(+ \wedge \sim C)P(\sim C)$

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

$P(+) = 0.107$

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

Problems

People are bad at this.

People who should be good at this are bad at it.

How you present the problem affects how bad people are at it.

Base Rate Fallacy

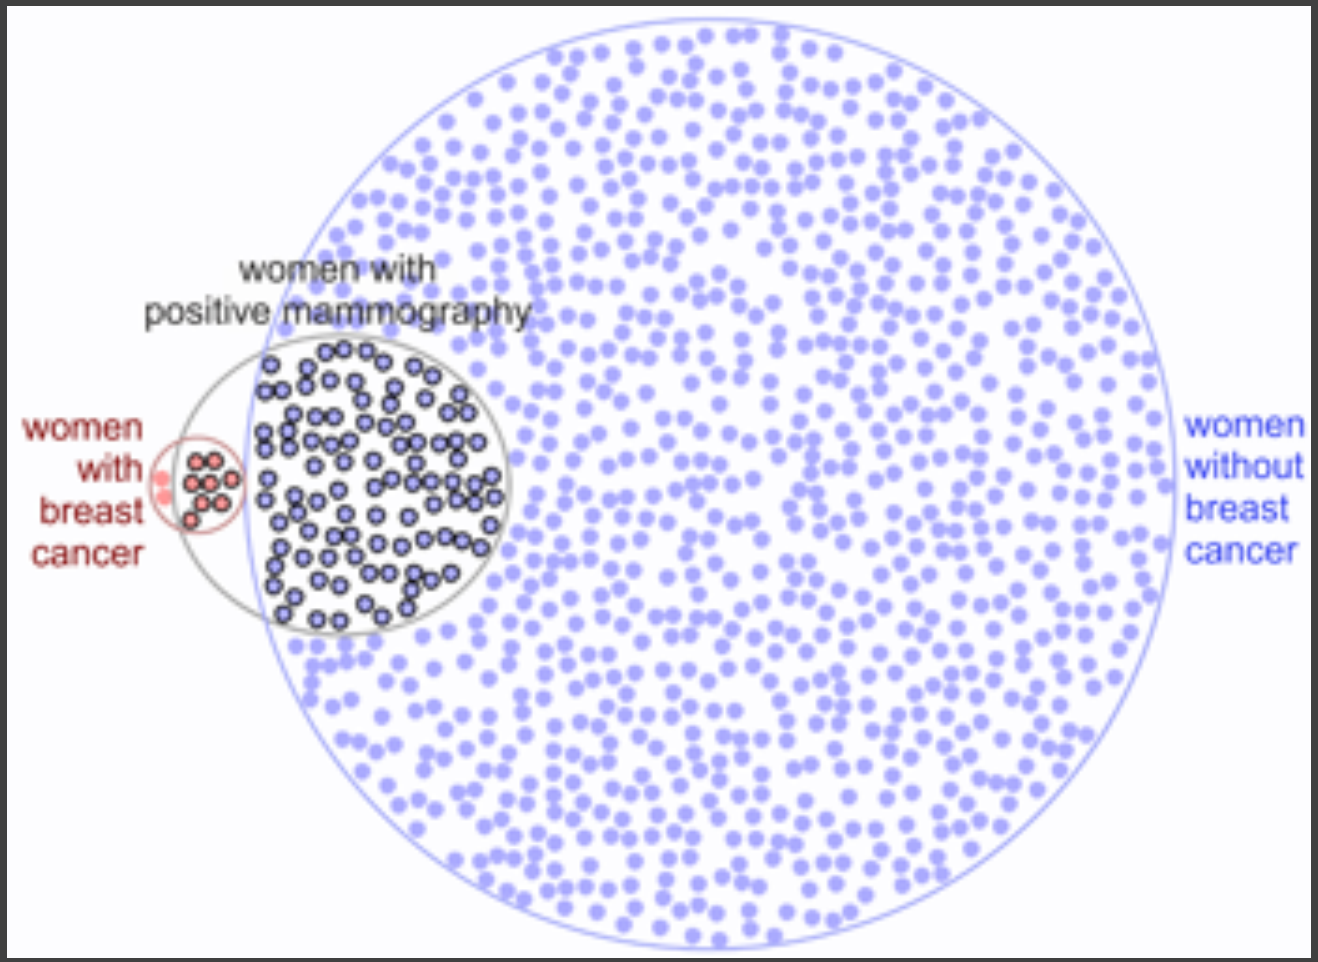

Luana Micallef, Pierre Dragicevic, and Jean-Daniel Fekete. “Assessing the Effect of Visualizations on Bayesian Reasoning Through Crowdsourcing.” VIS 2012.

How To Present Probabilities

Frequency

Probability P(A) = 0.6

Percentage 60% chance of A

Natural 3 out of 5 times, A happens.

Alvitta Ottley, et al. “Improving Bayesian reasoning: the effects of phrasing, visualization, and spatial ability.” VIS 2016.

Quantile Dot Plots

Less Error

More Error

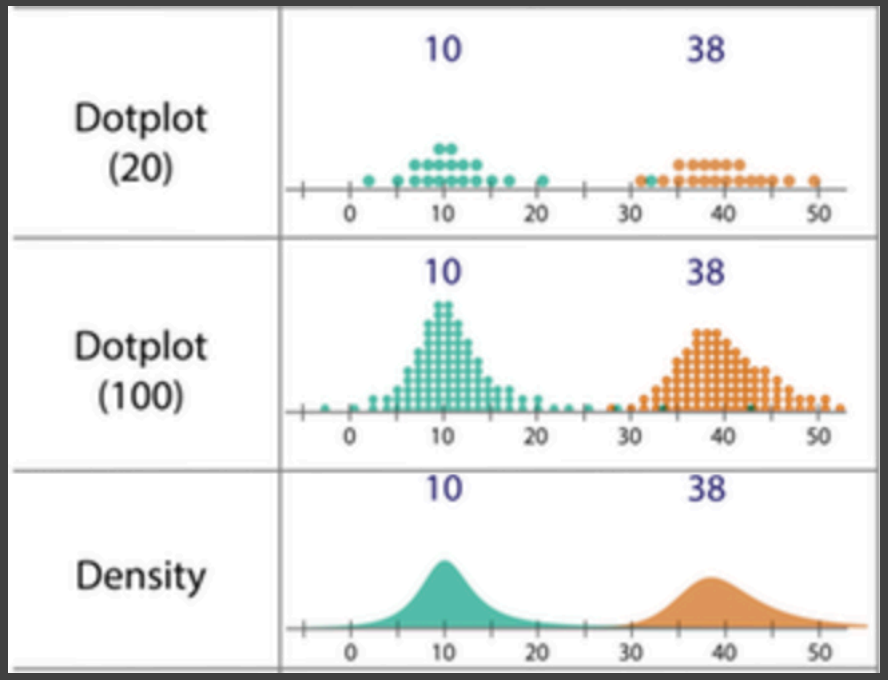

Matthew Kay, Tara Kola, Jessica Hullman, Sean Munson. "When(ish) is My Bus? User-centered Visualizations of Uncertainty in Everyday, Mobile Predictive Systems." CHI 2016.

What Can Go Wrong?

Uncertainty can be difficult to understand, and require a statistical background and high numeracy. Additionally, cognitive and perceptual biases can result in people making poor or error-prone decisions from uncertain data.

What Can Go Wrong?

Uncertainty can be difficult to understand, and require a statistical background and high numeracy. Additionally, cognitive and perceptual biases can result in people making poor or error-prone decisions from uncertain data.

A LOT

Questions To Answer

What Does Uncertainty Mean?

How Should I Visualize It?

What Can Go Wrong?

Questions To Answer

What Does Uncertainty Mean?

LOTS OF THINGS

How Should I Visualize It?

IT DEPENDS

What Can Go Wrong?

A LOT

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.