Interactions I

David Bamman
Info 202: Information Organization and Retrieval

Nov. 16, 2016

Interactions

- Interactions include any activity, function, or service supported by or enabled with respect to the resources in a collection or with respect the collection as a whole
- Whenever we select a resource for inclusion in an organizing system, describe it, or arrange it according to an organizing principle, we (should) have an interaction in mind
- TDO calls this “organizing on the way in”

We design organizing systems because we have some interaction in mind

Generic Interactions

- Finding a resource that you know exists
- Identifying a resource to make sure you have the one you were looking for
- Selecting a resource from a set of candidates in a collection
- Obtaining the resource if what you have at this point is just a resource description


SE RUEGA NO TOCAR

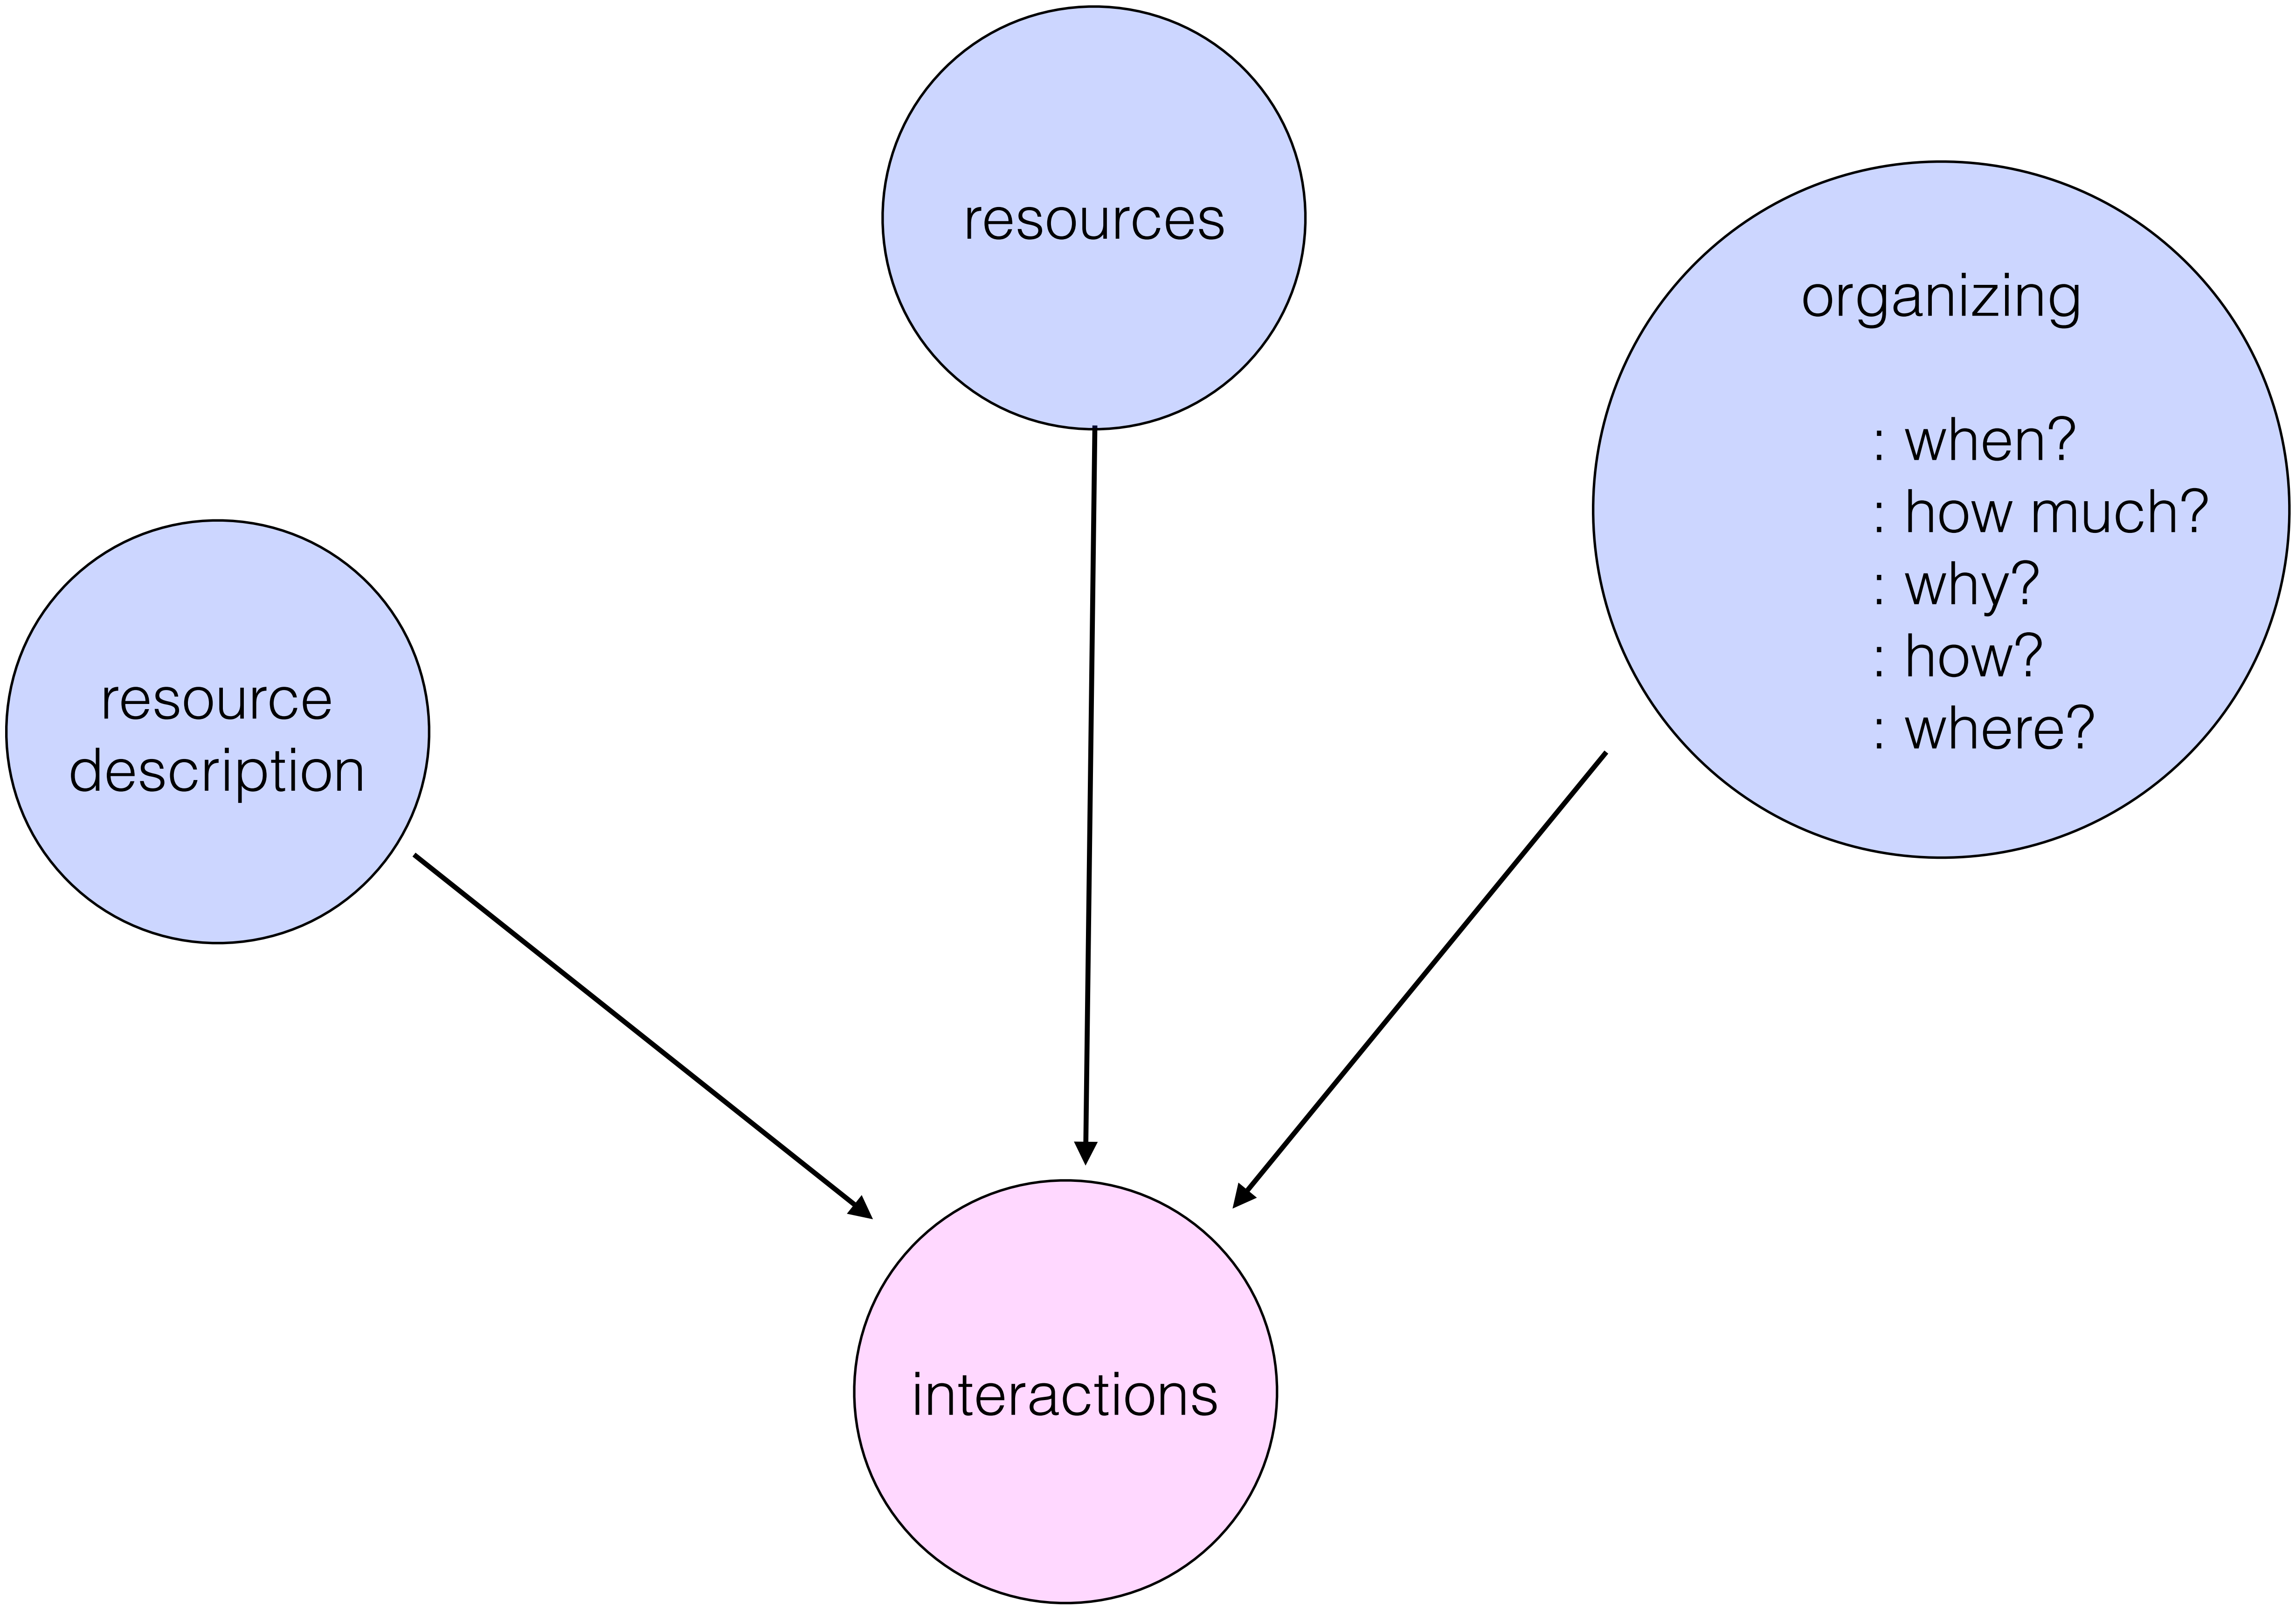
resources
organizing
: when?
: how much?
: why?
: how?
: where?
resource
description
interactions

Interactions

- The design of interactions operates within constraints imposed by:
    - Resources
    - Organizing system
    - Producers
    - Users
    - Technical, legal, policy, social requirements

Interactions

- If a collection of resources isn't already organized to support our desired interactions, computational techniques can re-organize them (if they are digital) or their digital resource descriptions

- … by indexing, searching, sorting, transforming, translating, labeling, categorizing, combining, summarizing… (IR, ML, NLP)

- This doesn't always mean that the collection was organized poorly; it might just be organized for different interactions than the ones we now want

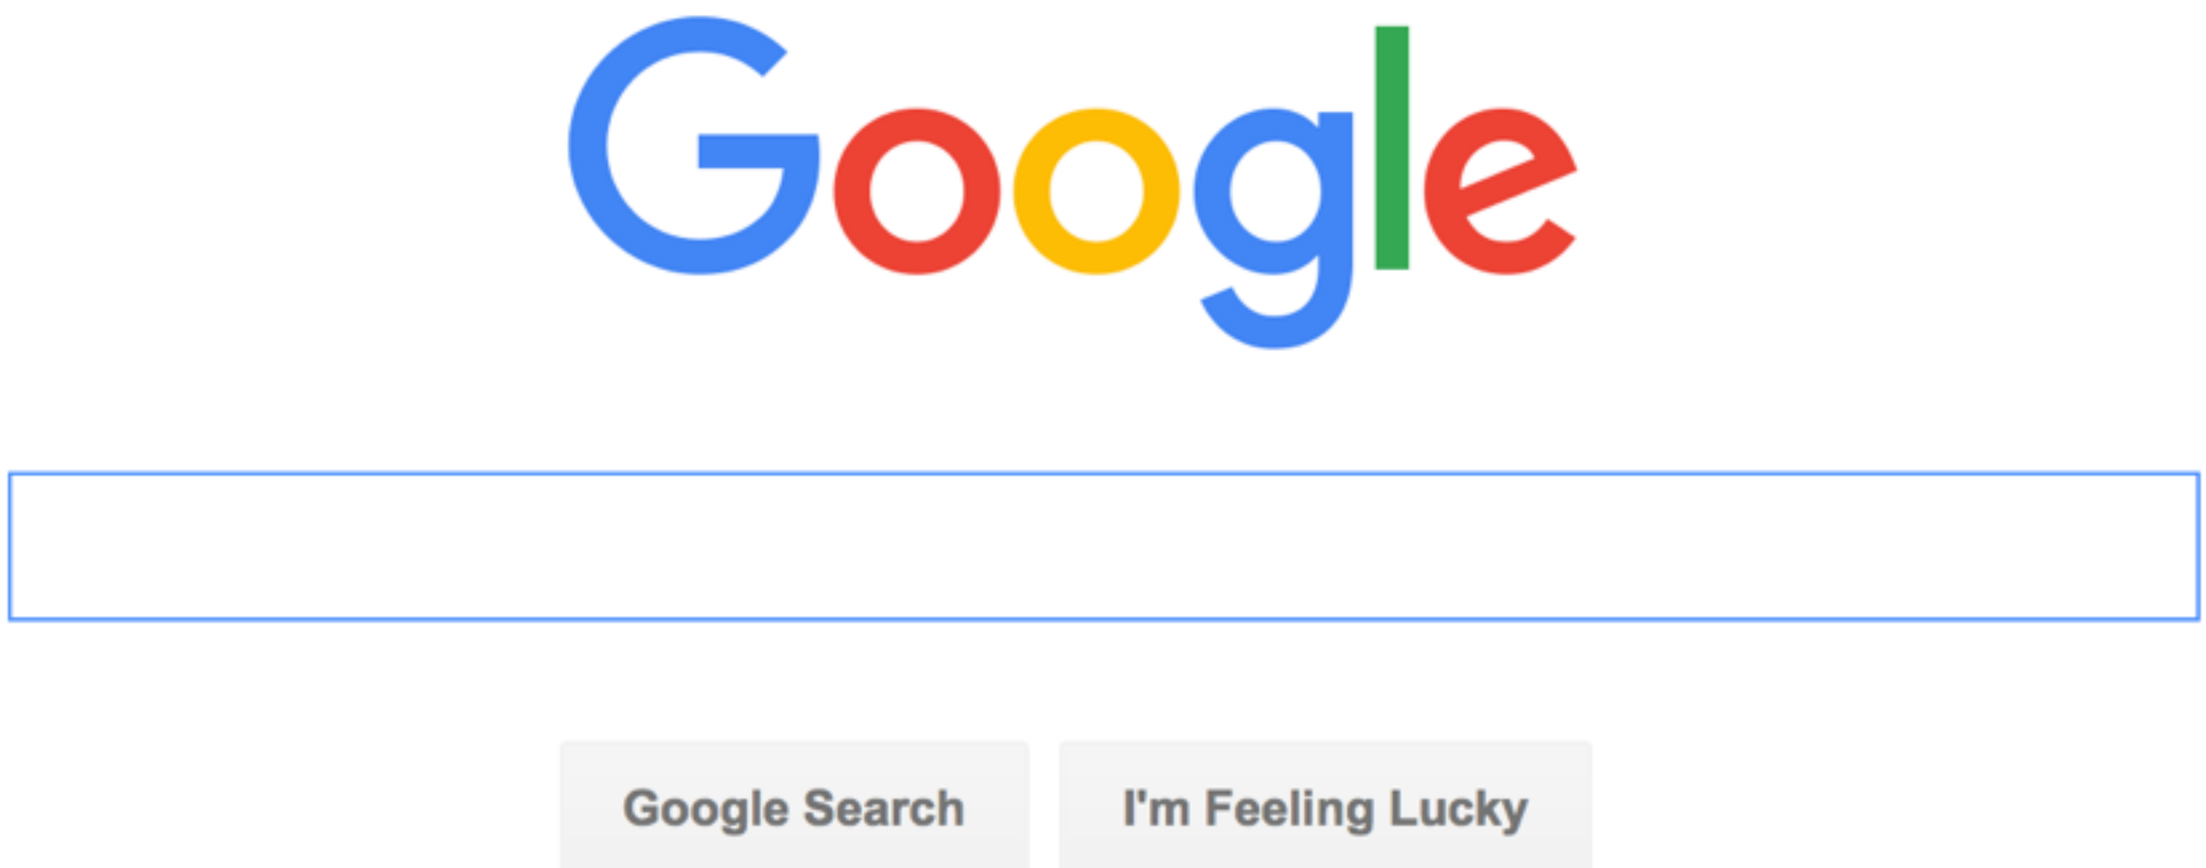
Google
Google Search
I'm Feeling Lucky

Interactions

- With general organizing systems, we often organize resources to support a range of interactions, each of which must be discoverable by users.
- In seeing data analysis as an instance of organizing, we can prioritize the design of a single interaction.

Data science as an organizing system

- The selection of data
- The description of data
- Leveraging relationships between data points
- To enable interactions: classification, prediction, recommendation, inference, hypothesis testing

Case study: prediction

Information organization here involves selecting data and describing it to enable an interaction: prediction

- what is being organized?
- why is it being organized?
- how much is it being organized?
- when is it being organized?
- how (or by whom) is it being organized?
- where is it being organized?

Classification

- Assigning resources to categories

— as an additional resource description (to support some other interaction)

— as an end goal itself (where the interaction = prediction)

Inference

- If our end goal is to infer the most likely value for something we don't observe, our interaction = inference.

- Inference could reason over the most likely likely category (classification), real value (regression), or our degree of belief in a hypothesis.

- Not simply the act of picking the most likely value, but also of quantifying our uncertainty about that value.

Hypotheses

hypothesis
The average income in two sub-populations is different
Web design A leads to higher CTR than web design B
Self-reported location on Twitter is predictive of political preference
Male and female literary characters become more similar over time

Null hypothesis

- A claim, assumed to be true, that we'd like to test (because we think it's wrong)

hypothesis	H_0
The average income in two sub-populations is different	The incomes are the same
Web design A leads to higher CTR than web design B	The CTR are the same
Self-reported location on Twitter is predictive of political preference	Location has no relationship with political preference
Male and female literary characters become more similar over time	There is no difference in M/F characters over time

Hypothesis testing

- If the null hypothesis were true, how likely is it that you'd see the data you see?

Example

- Hypothesis: Berkeley residents tend to be politically liberal

- H_0: Among all N registered {Democrat, Republican} primary voters, there are an equal number of Democrats and Republicans in Berkeley.

$$\frac{\#dem}{N} = \frac{\#rep}{N} = 0.5$$

Example

- If we had access to the party registrations (and knew the population), we would have our answer.
- Instead, we only have access to samples
- Any sample has variability simply as a function of random noise.

Signal vs. Noise

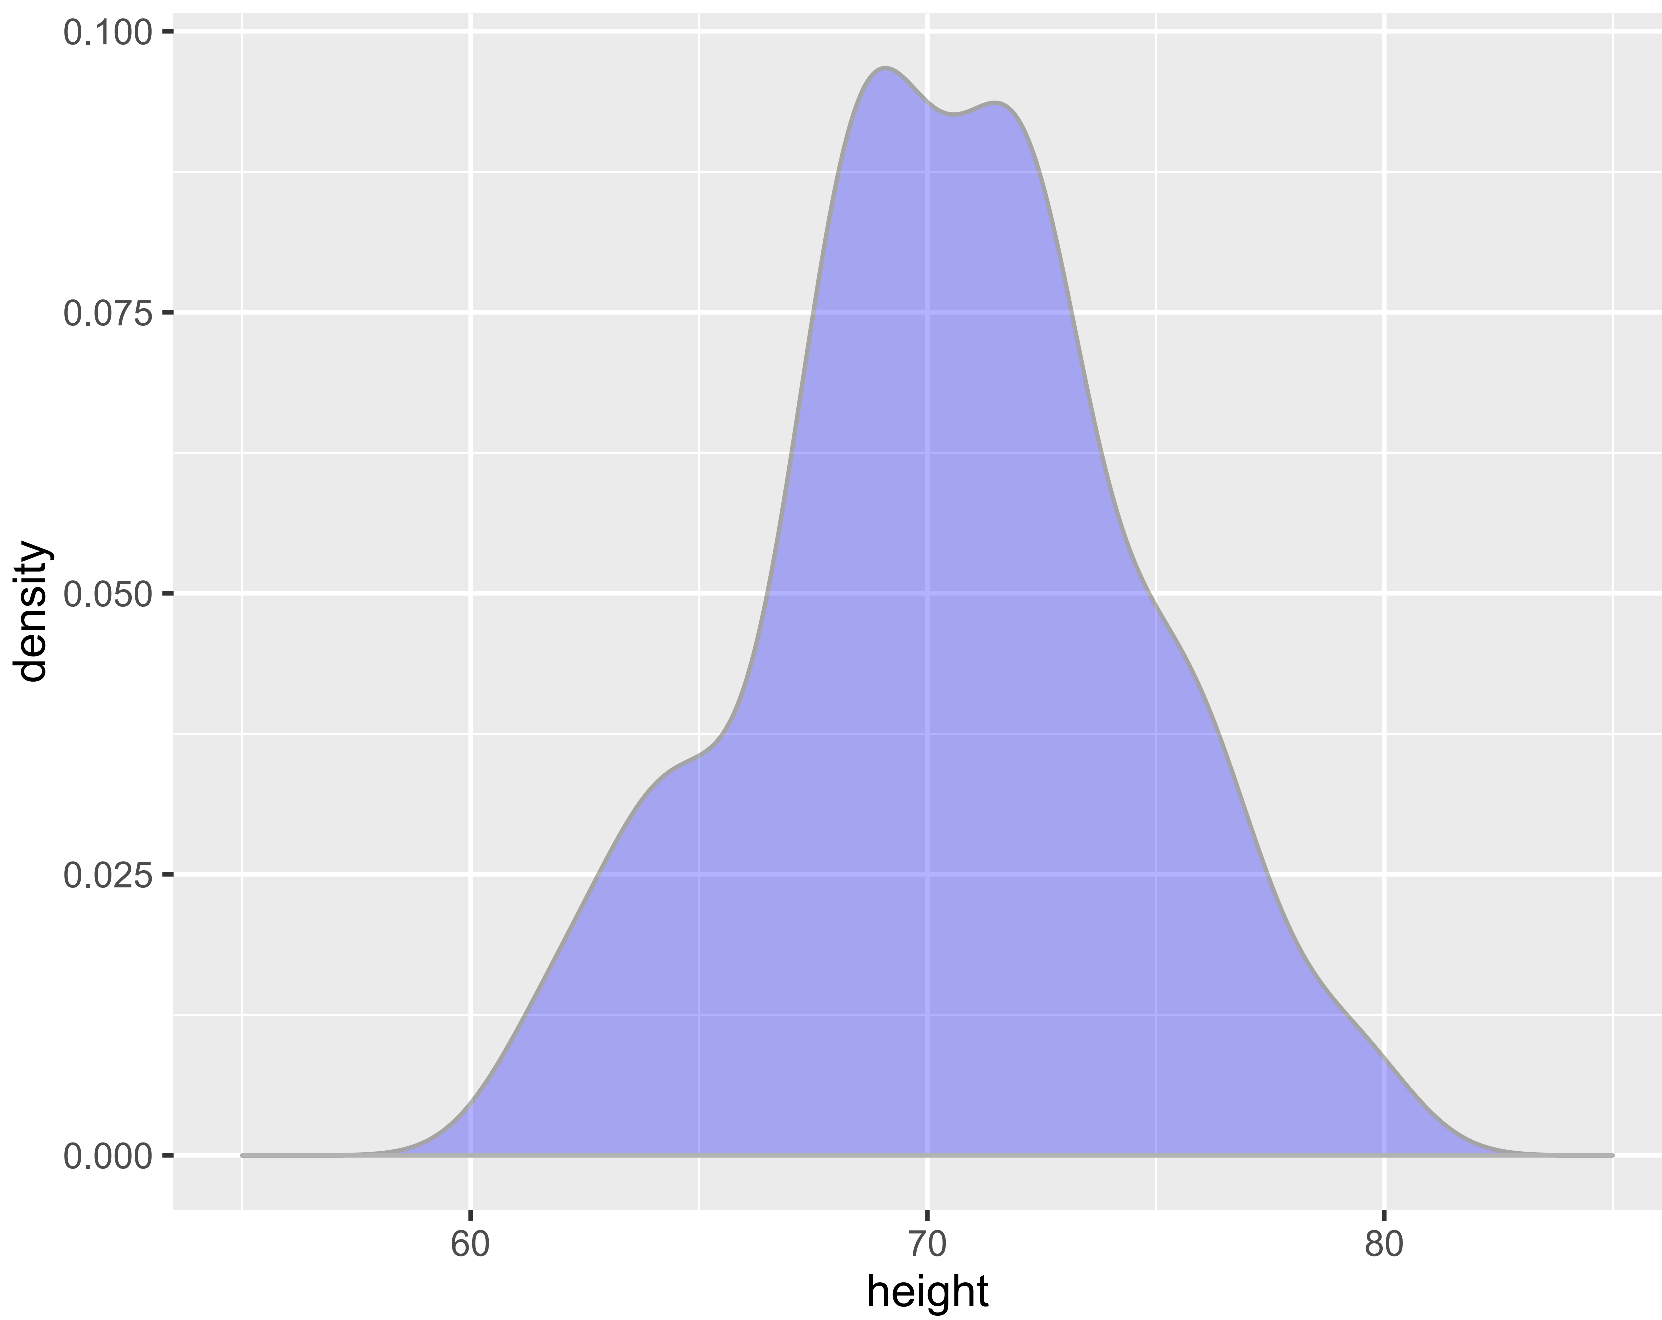

Signal vs. Noise

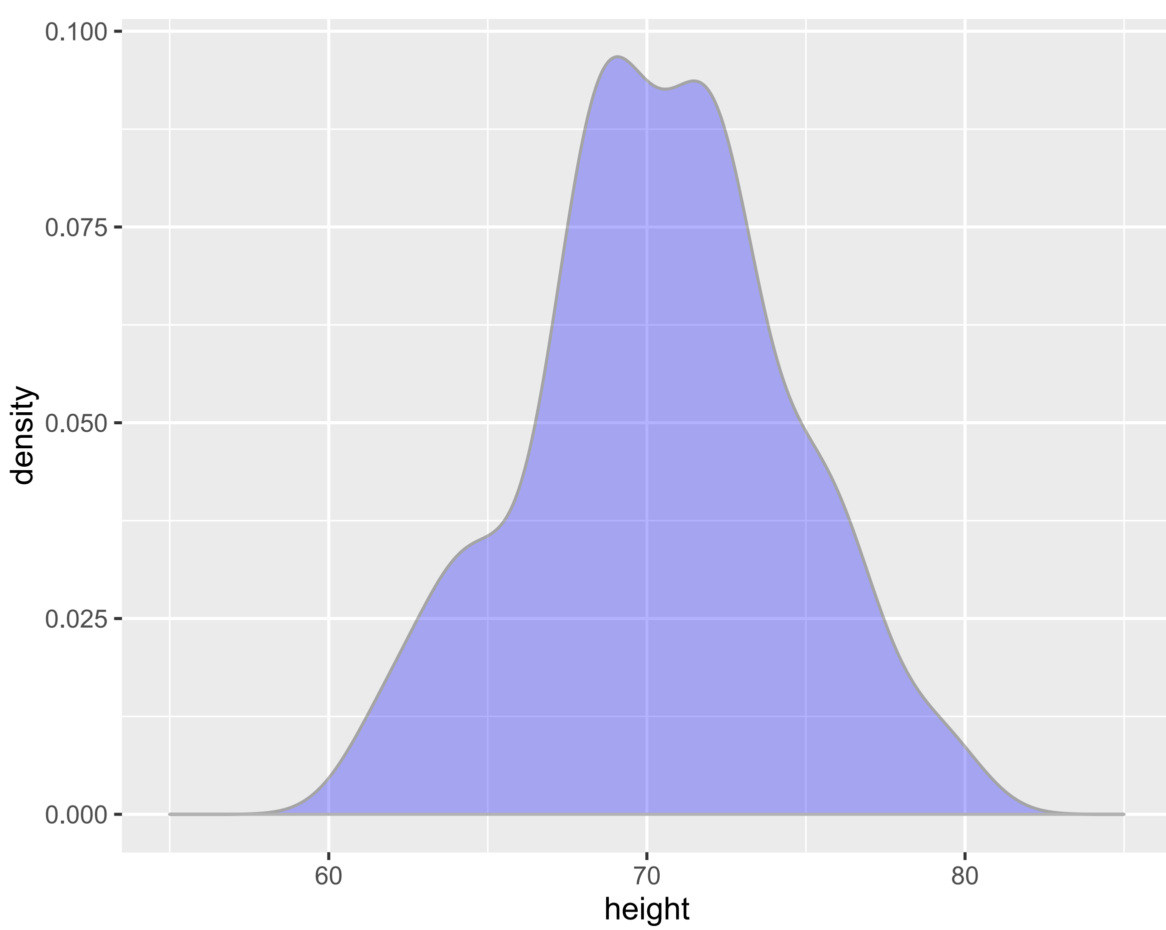

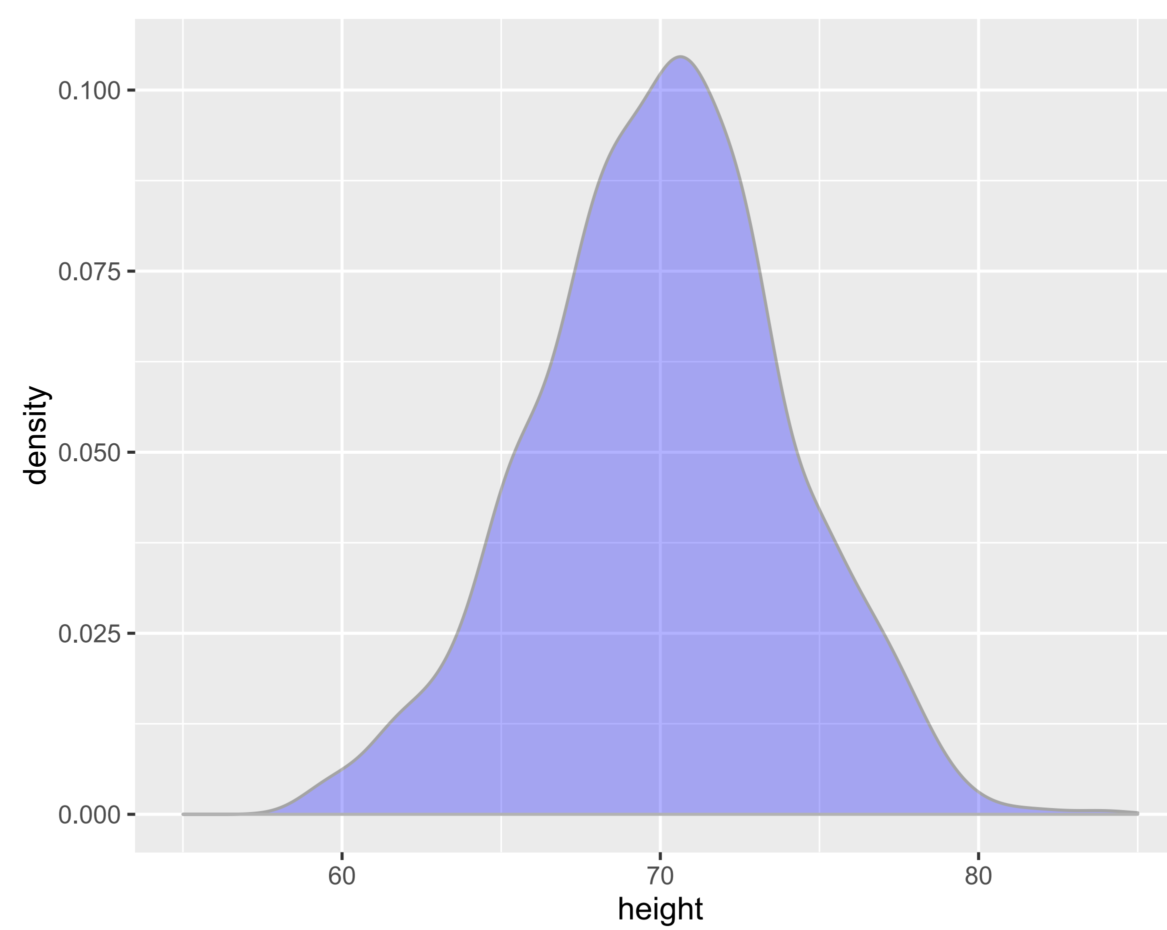

Signal vs. Noise

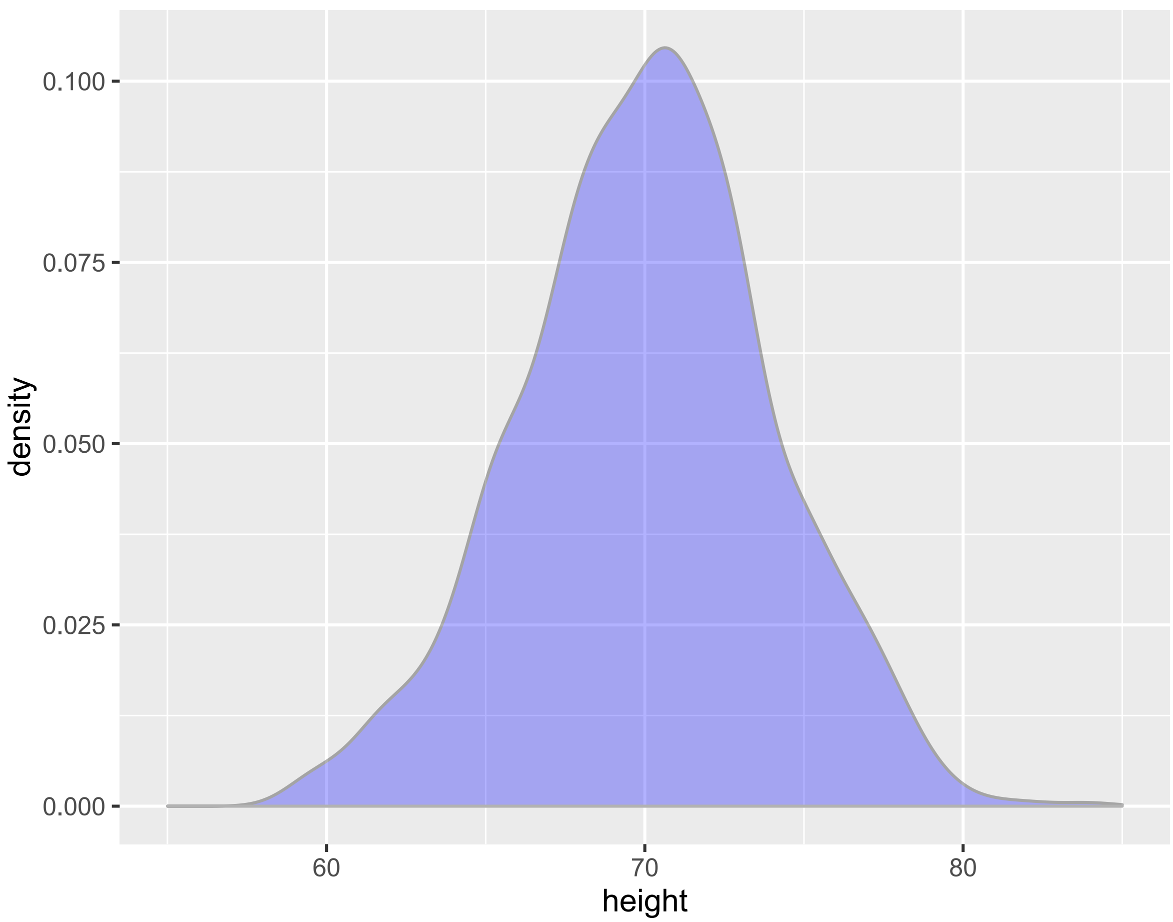

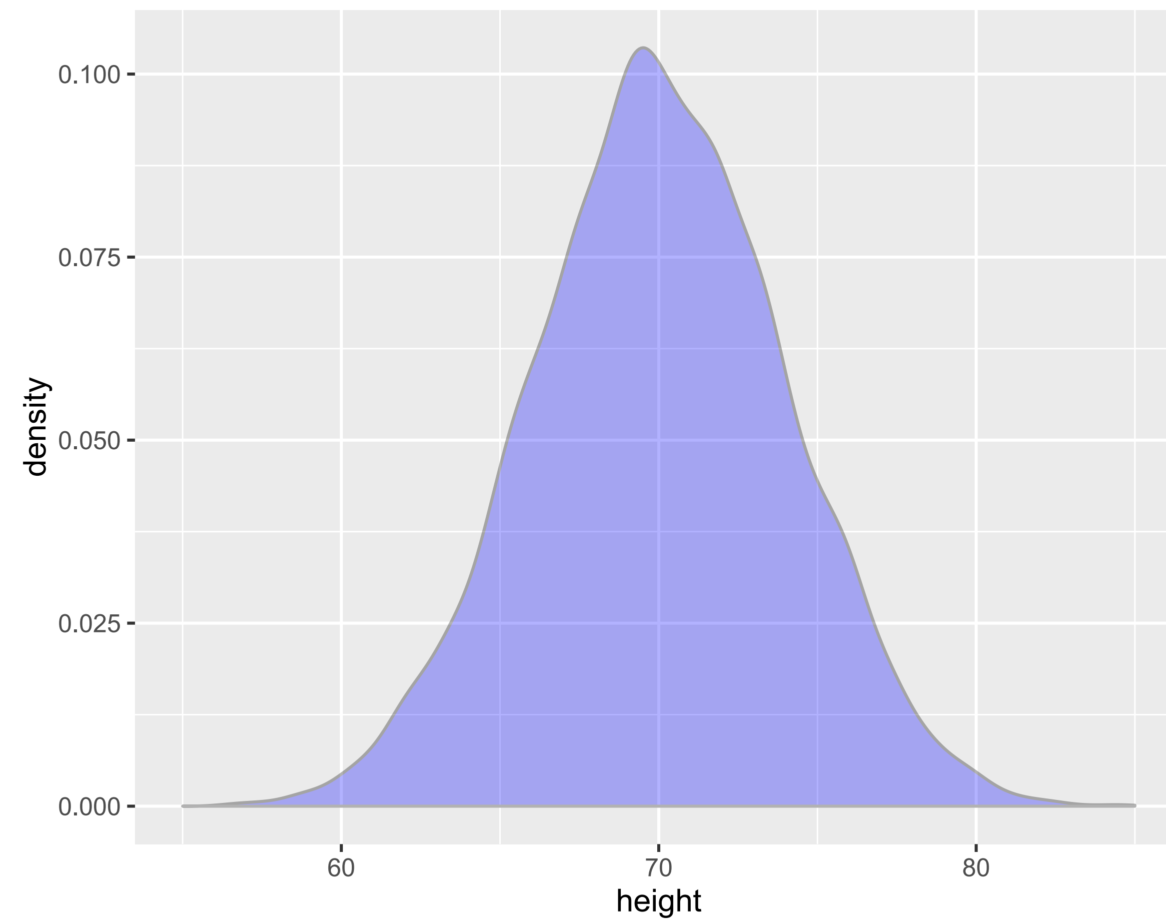

Signal vs. Noise

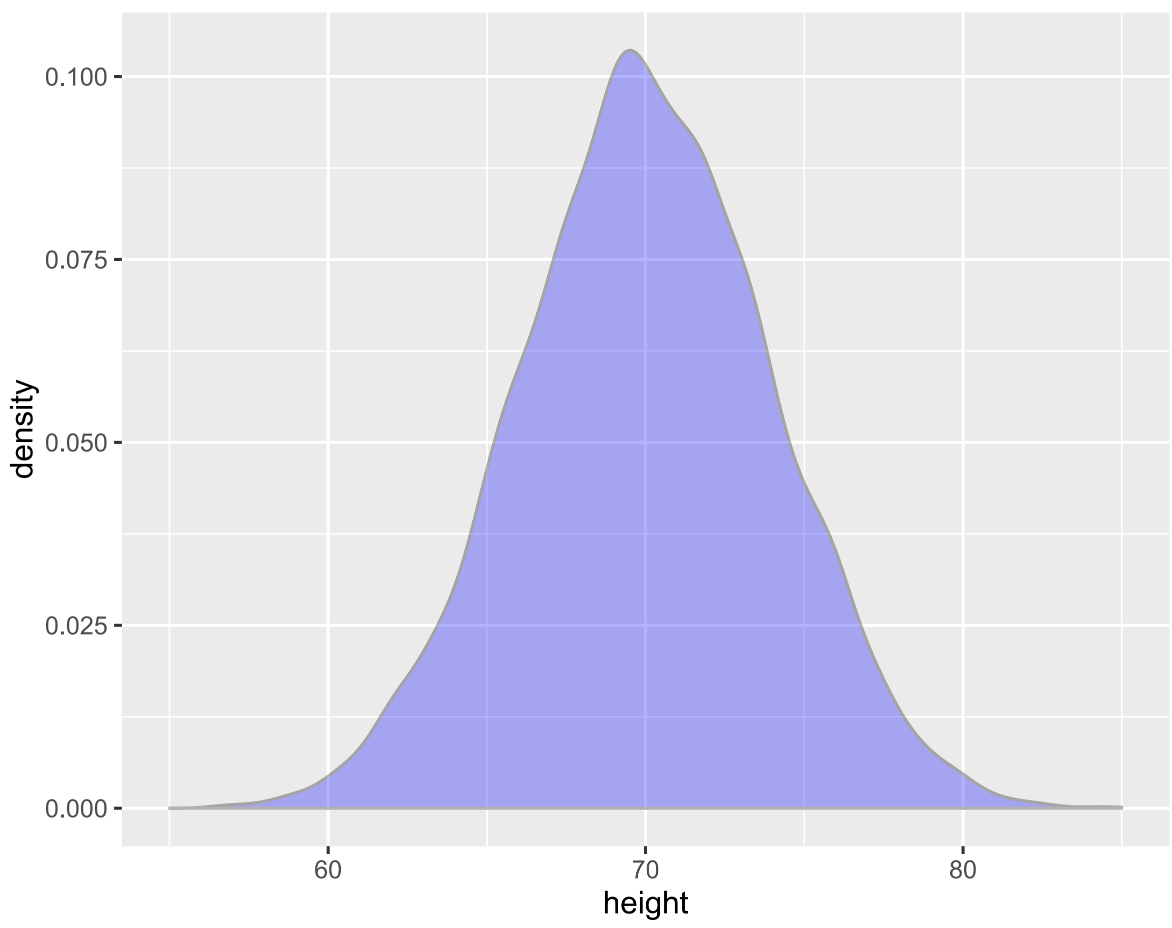

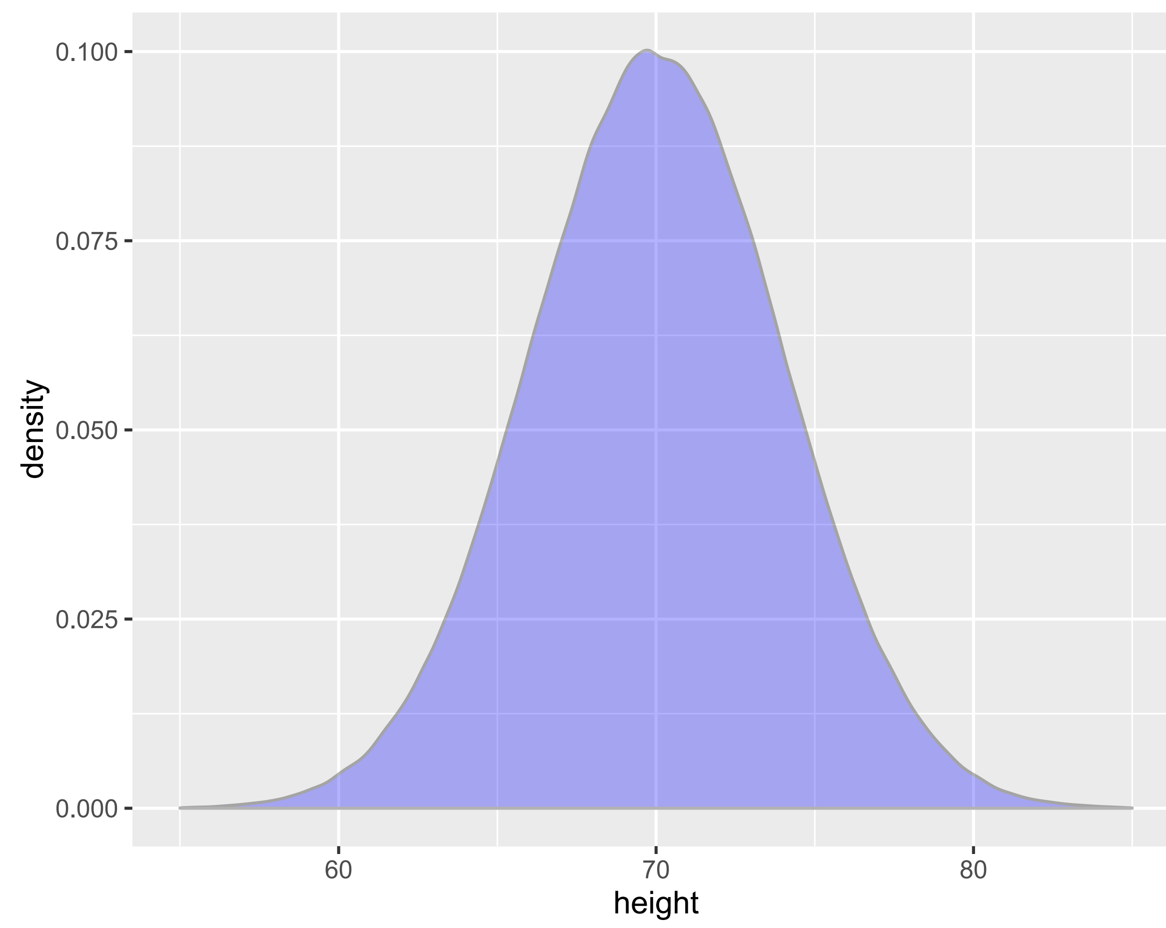

Signal vs. Noise

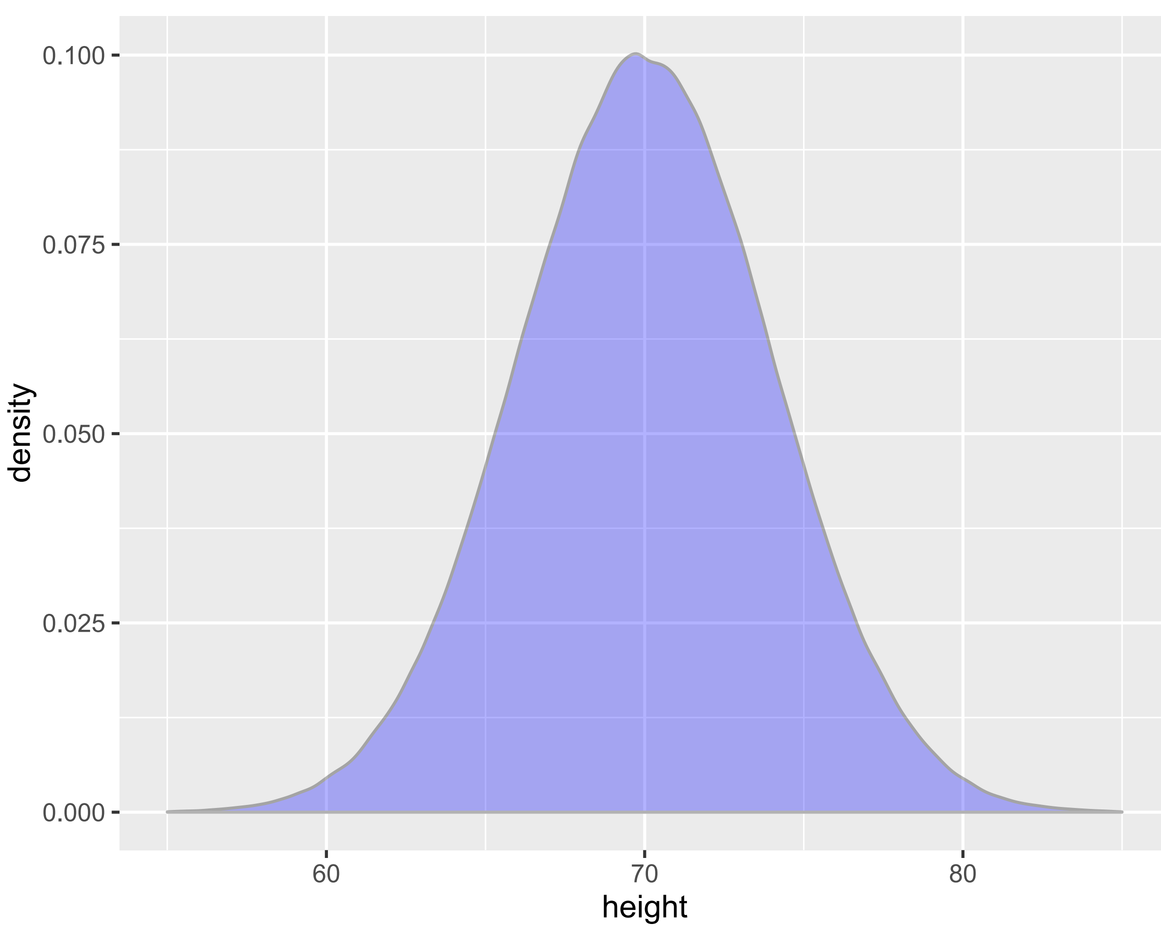

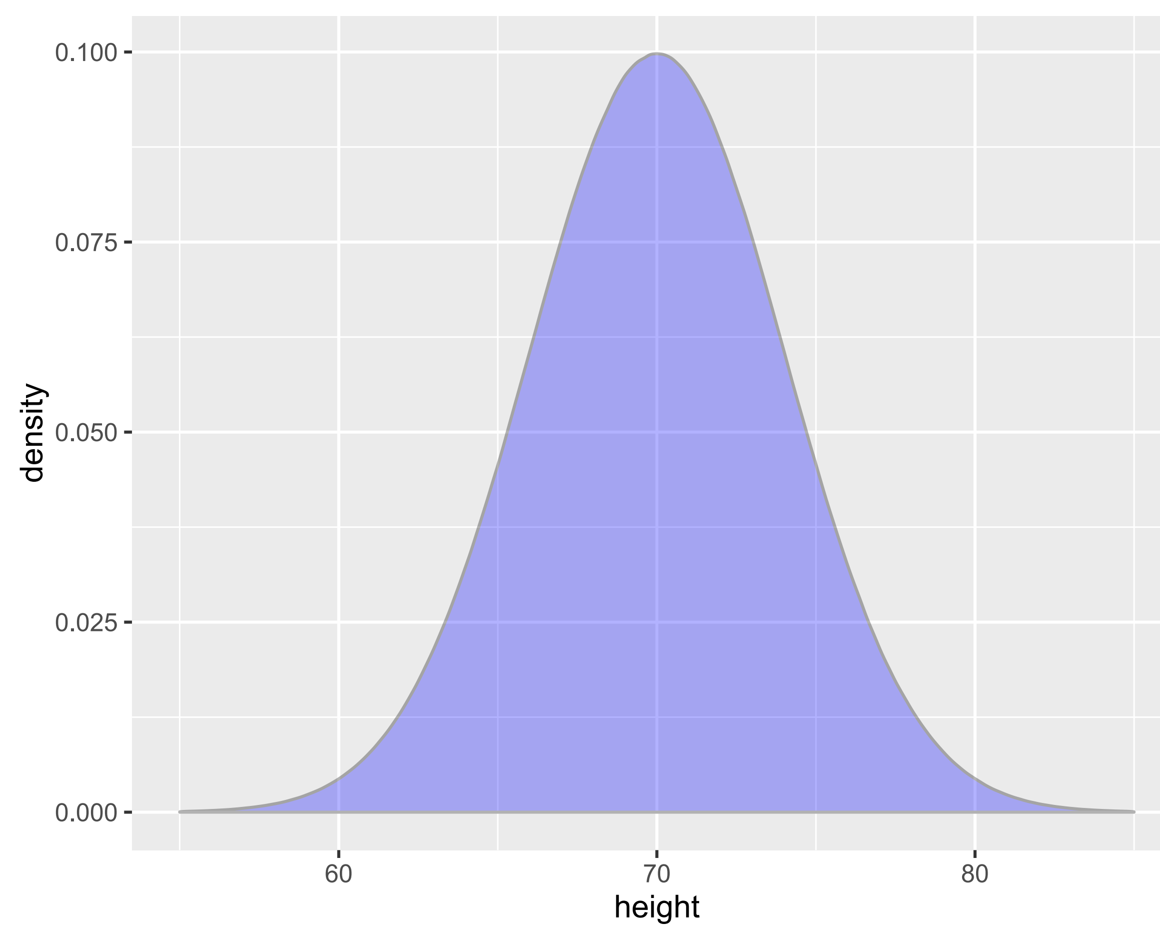

Signal vs. Noise

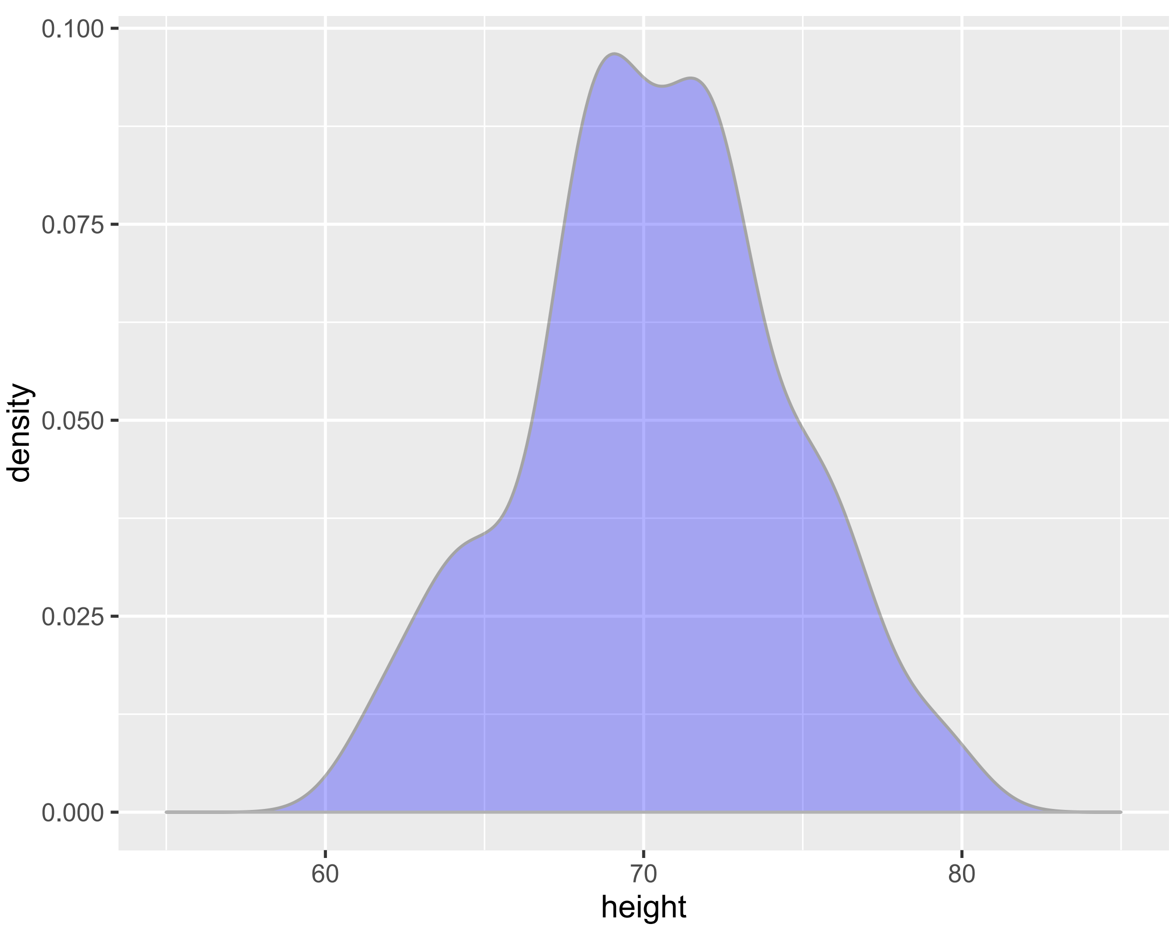

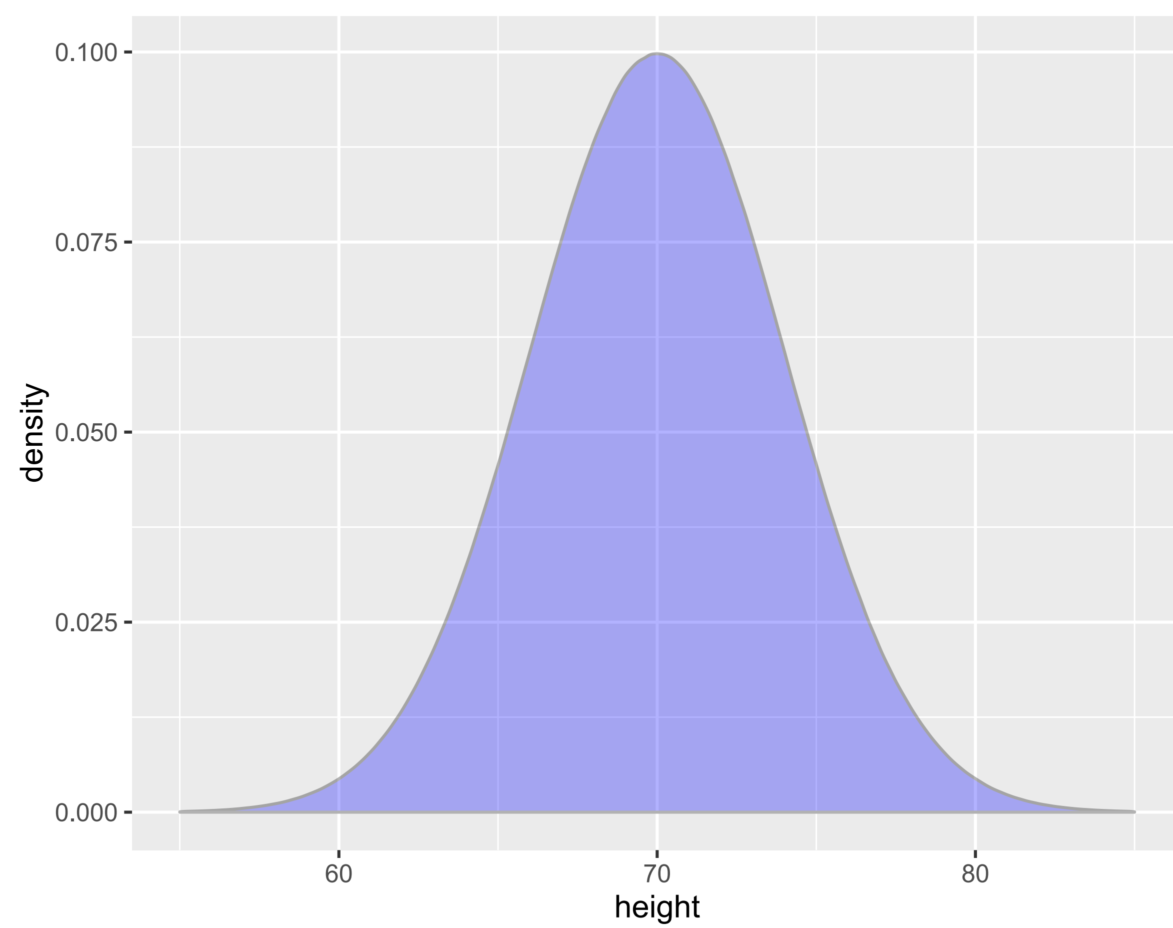

The data we see (left) is a combination of some signal (right) and random noise

Example

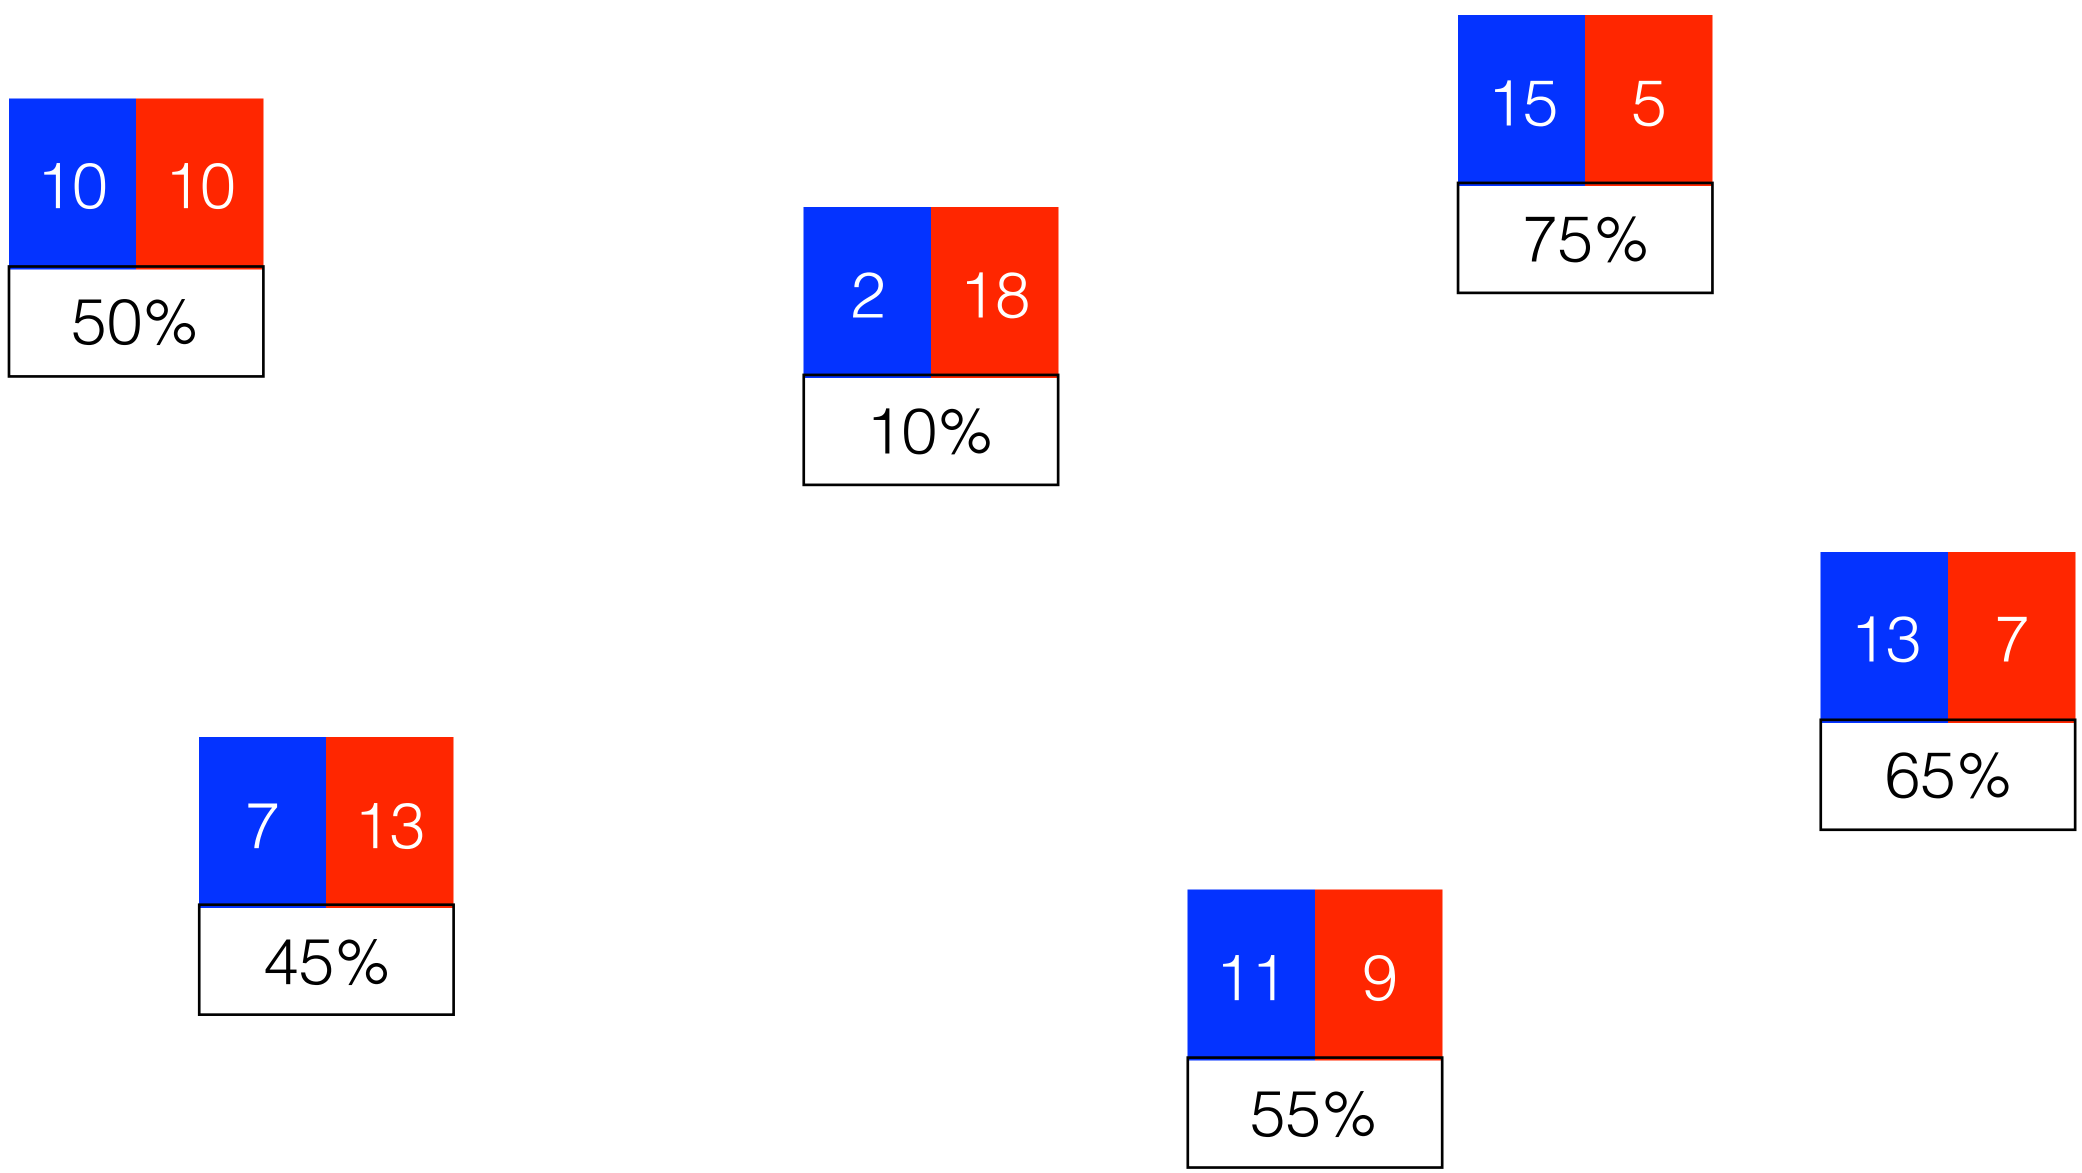

Hypothesis testing

- Hypothesis testing measures our confidence in what we can infer about a null from a sample.

Example

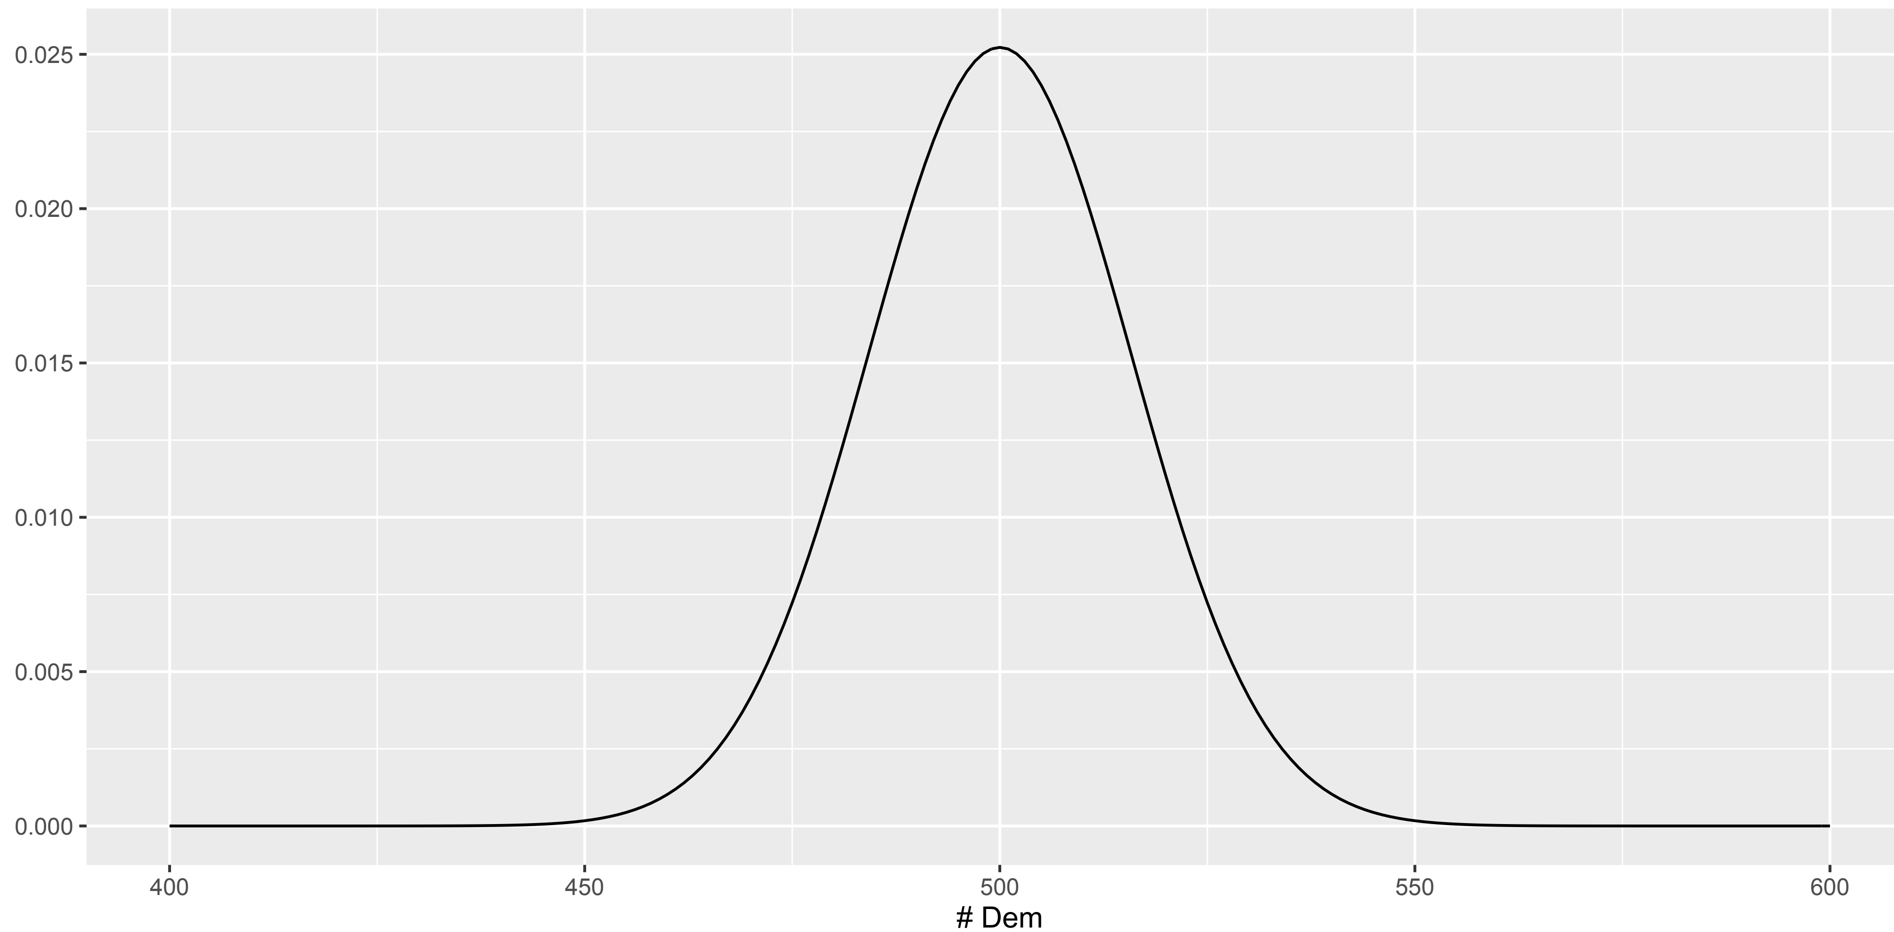

Binomial probability distribution for number of democrats in n=1000 with p = 0.5

Example

At what point is a sample statistic unusual enough to reject the null hypothesis?

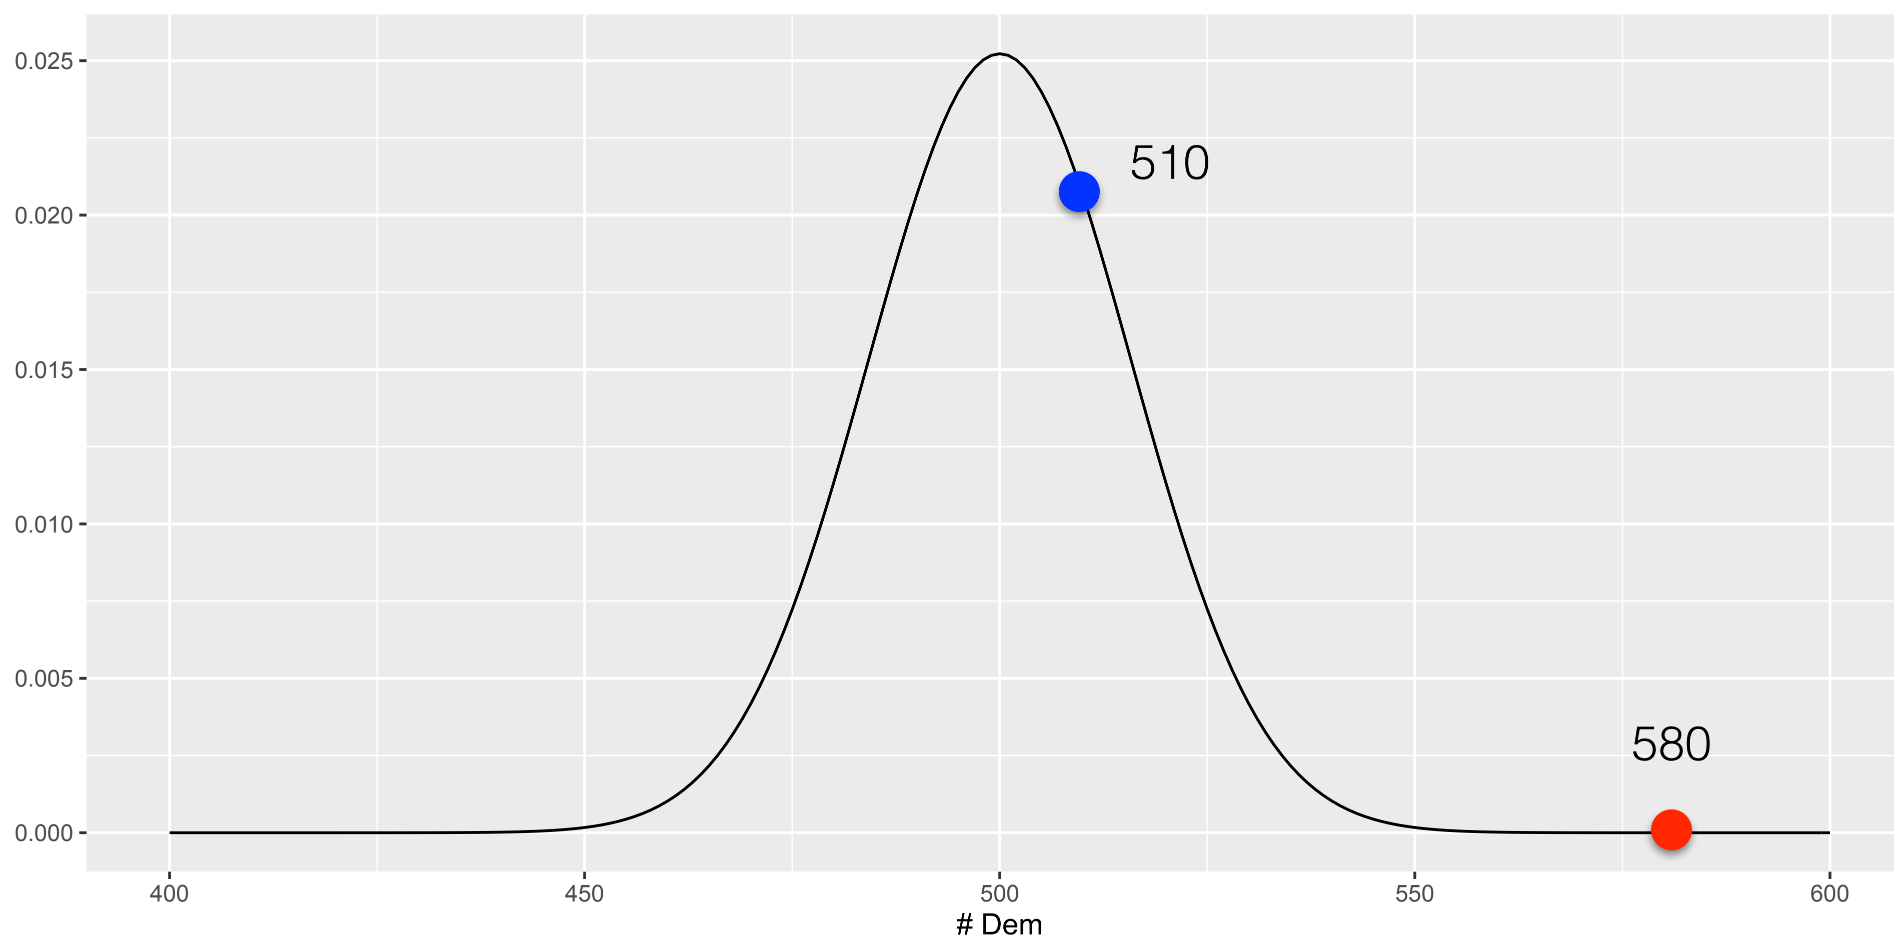

Example

- The form we assume for the null hypothesis lets us quantify that level of surprise.

- We can do this for many parametric forms that allows us to measure $P(X \leq x)$ for some sample of size n; for large n, we can often make a normal approximation.

Z score

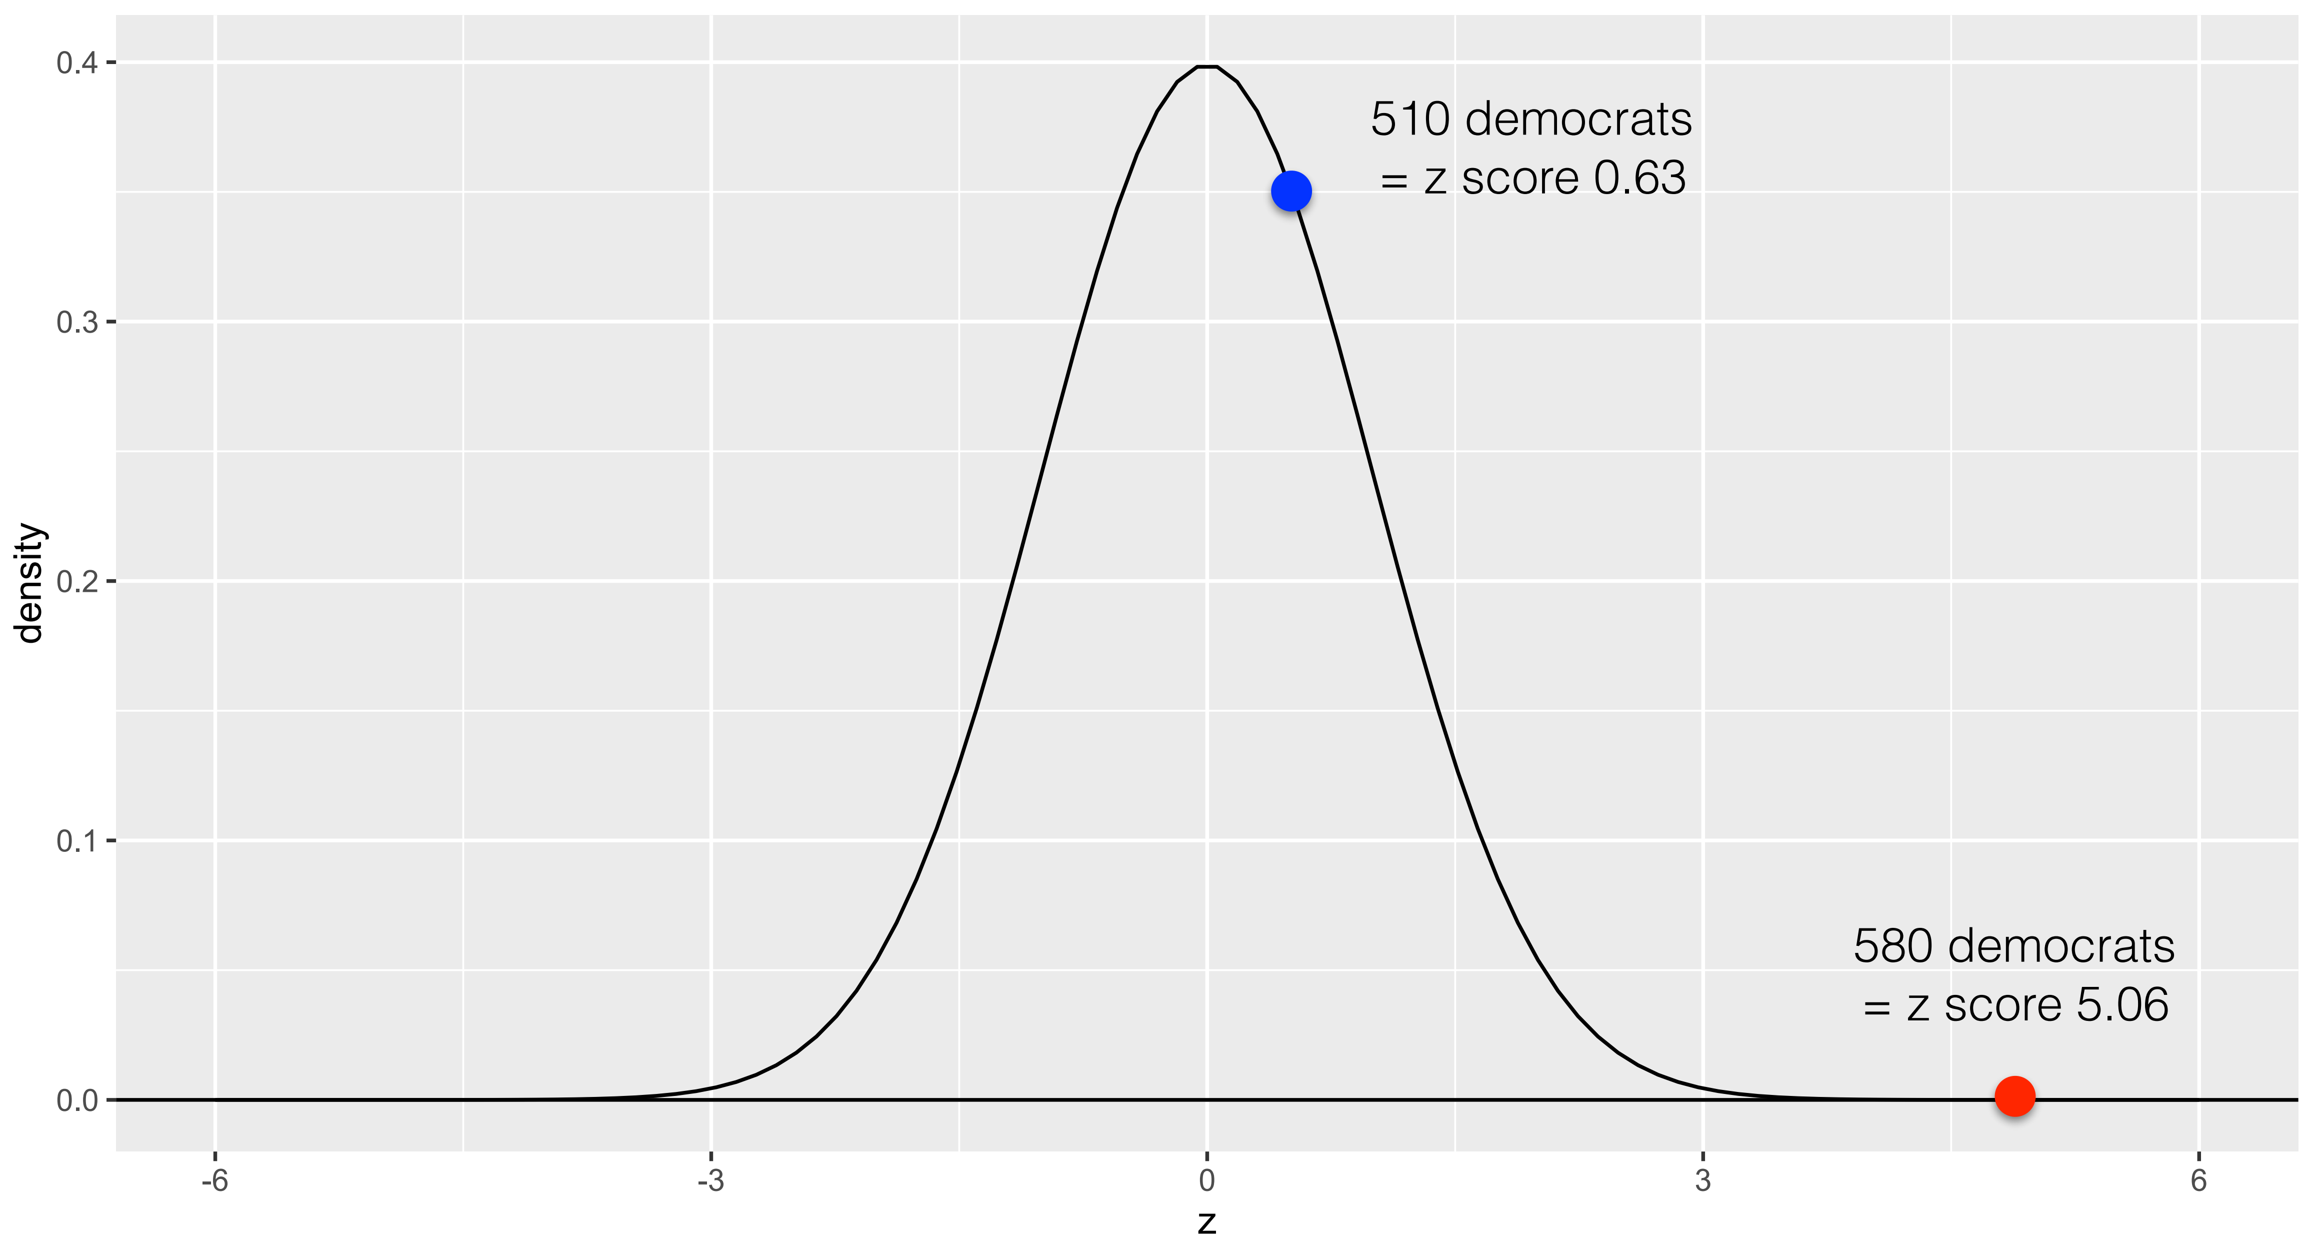

0.4
0.3
0.2
0.1
0.0
density
-6
-3
0
3
6
z
510 democrats
= z score 0.63
580 democrats
= z score 5.06

Tests

- We will define “unusual” to equal the most extreme areas in the tails

least likely 10%

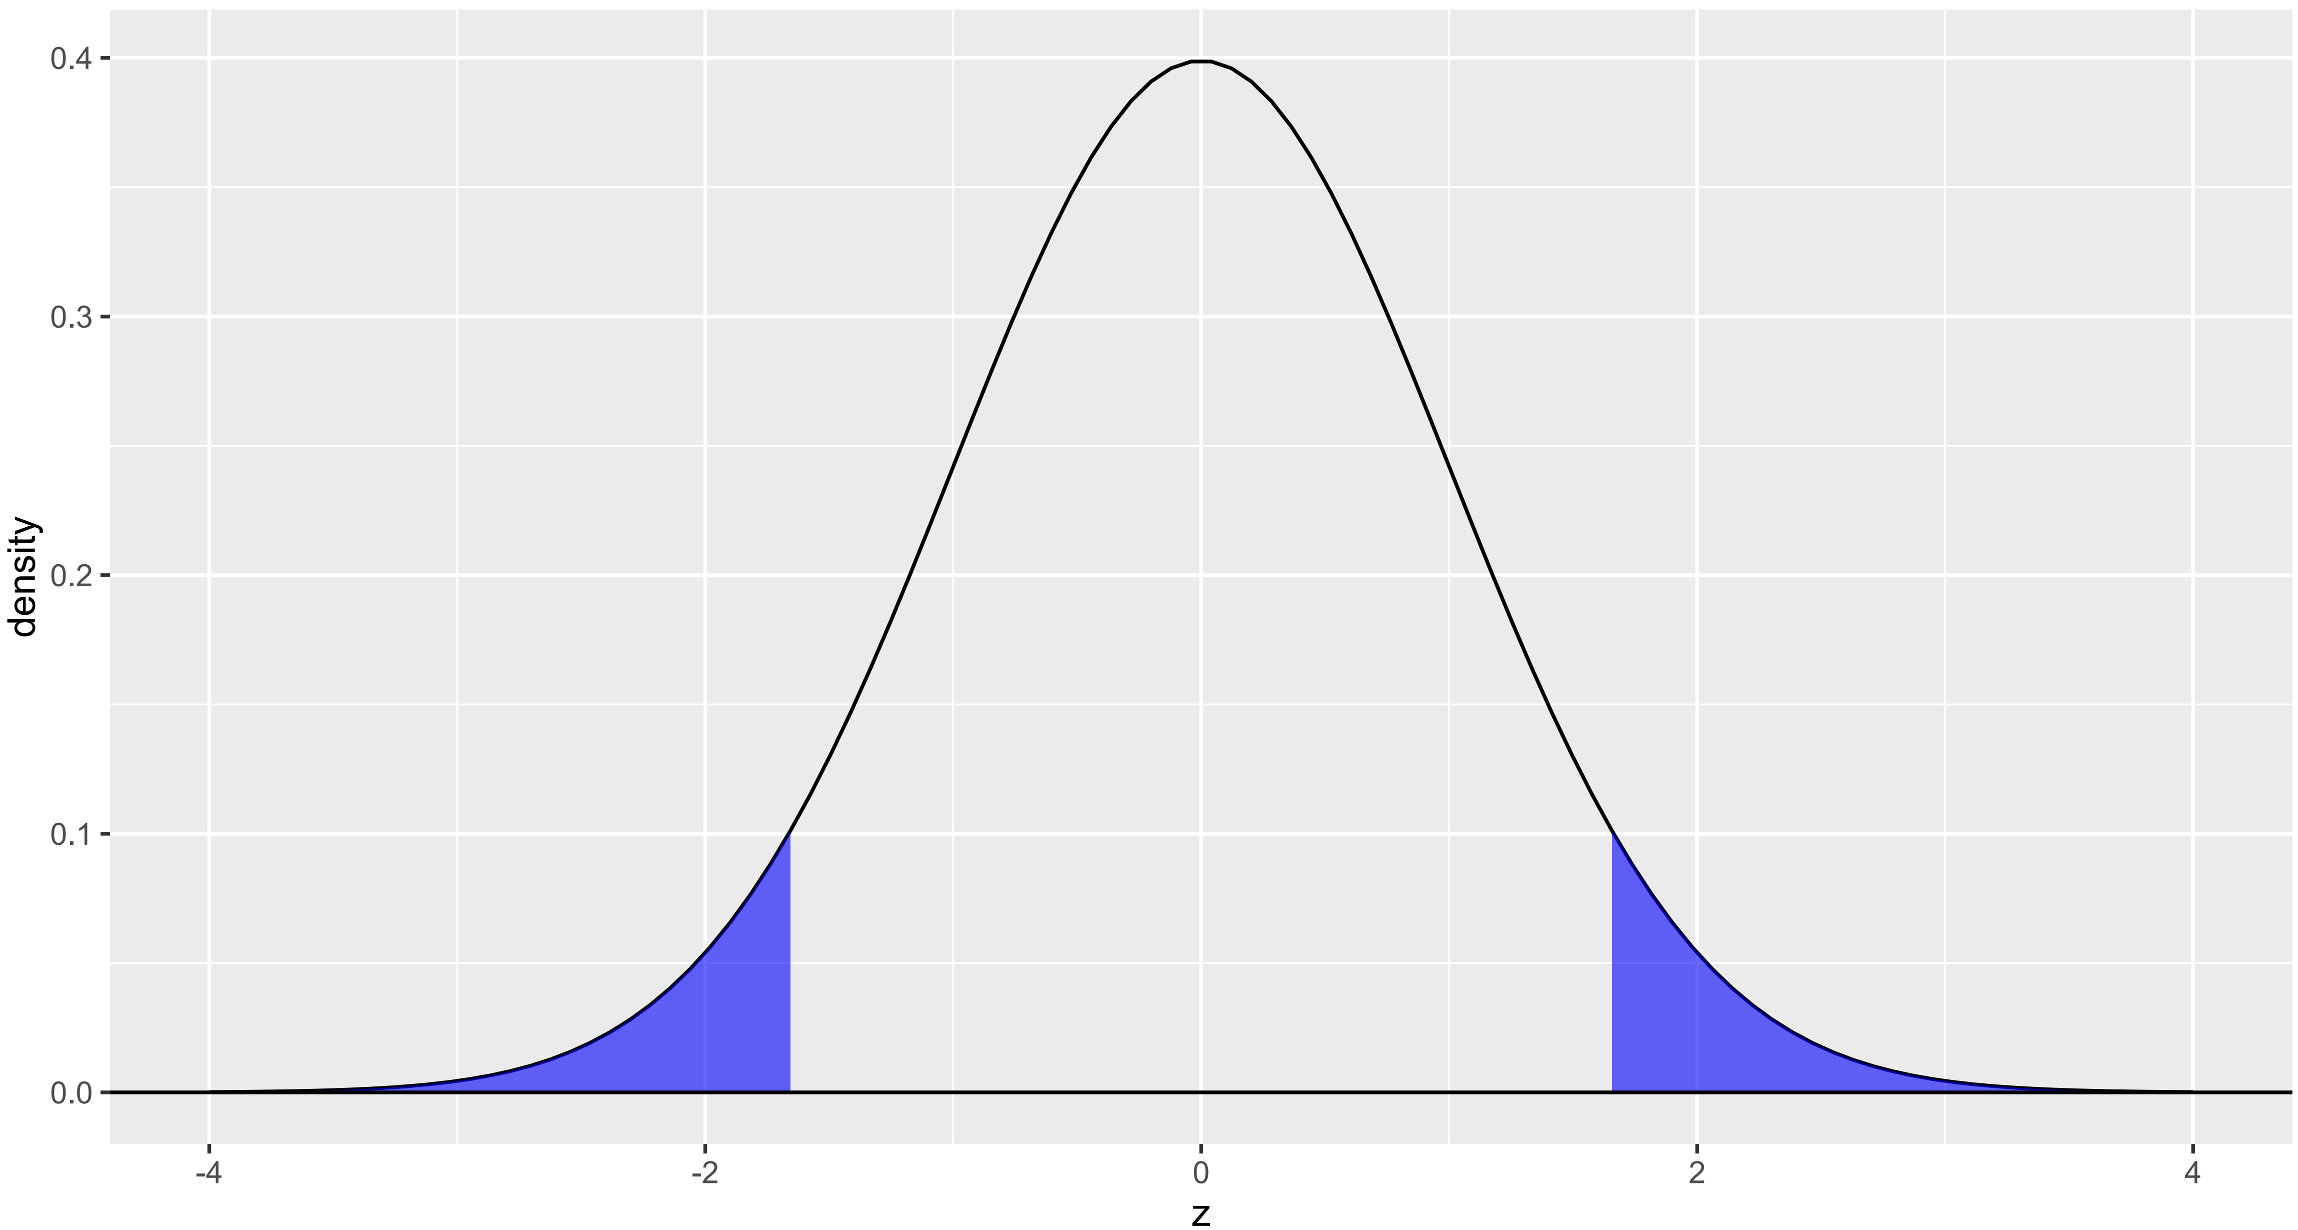

least likely 5%

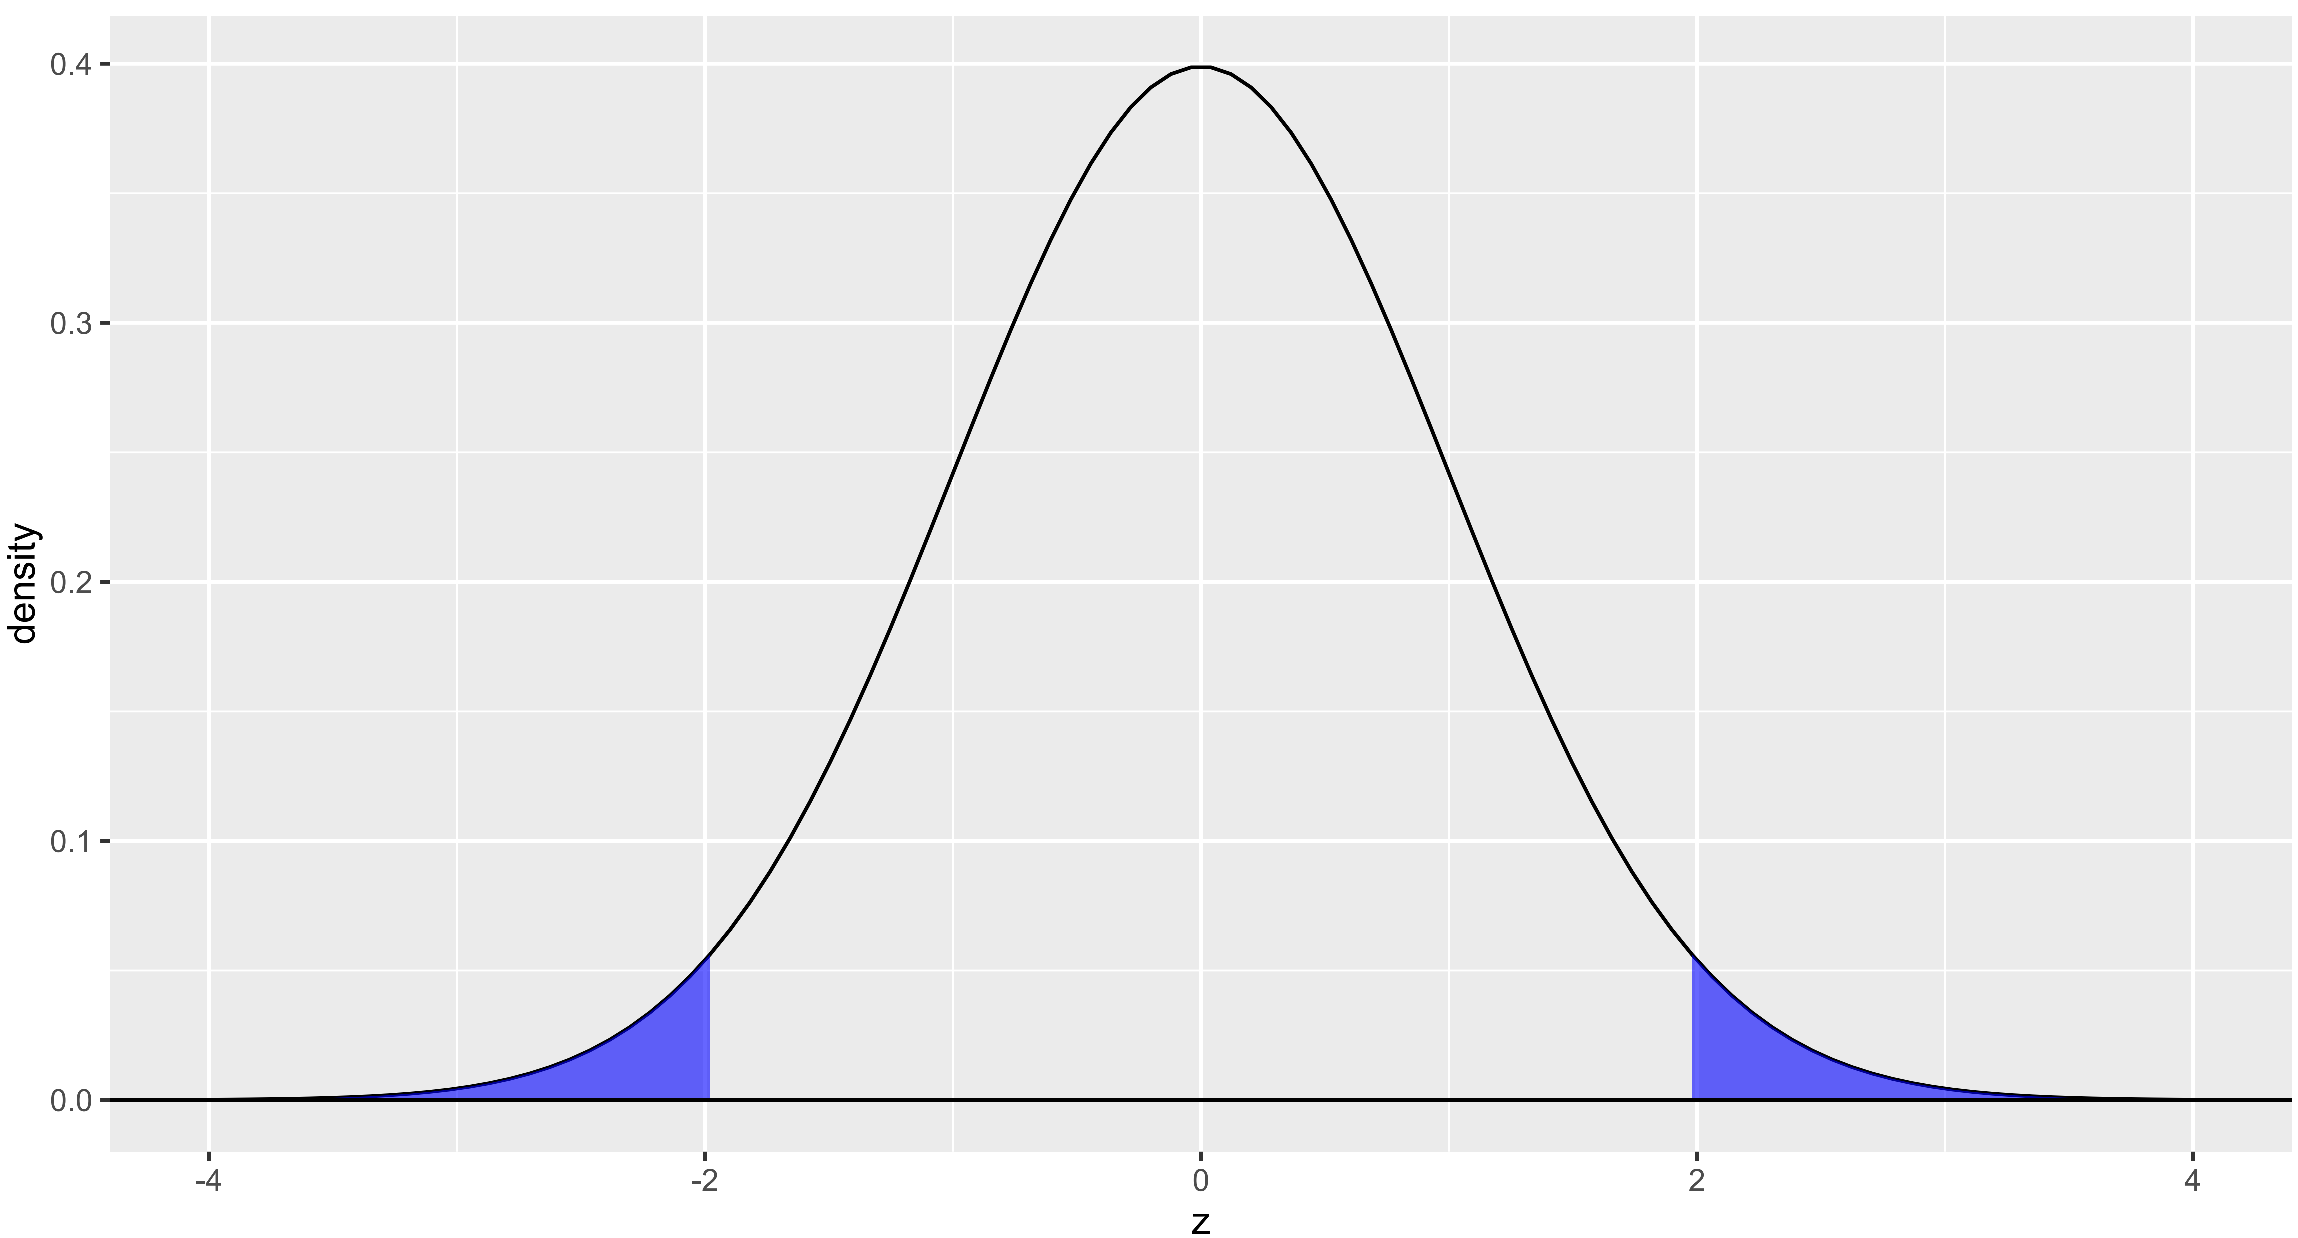

least likely 1%

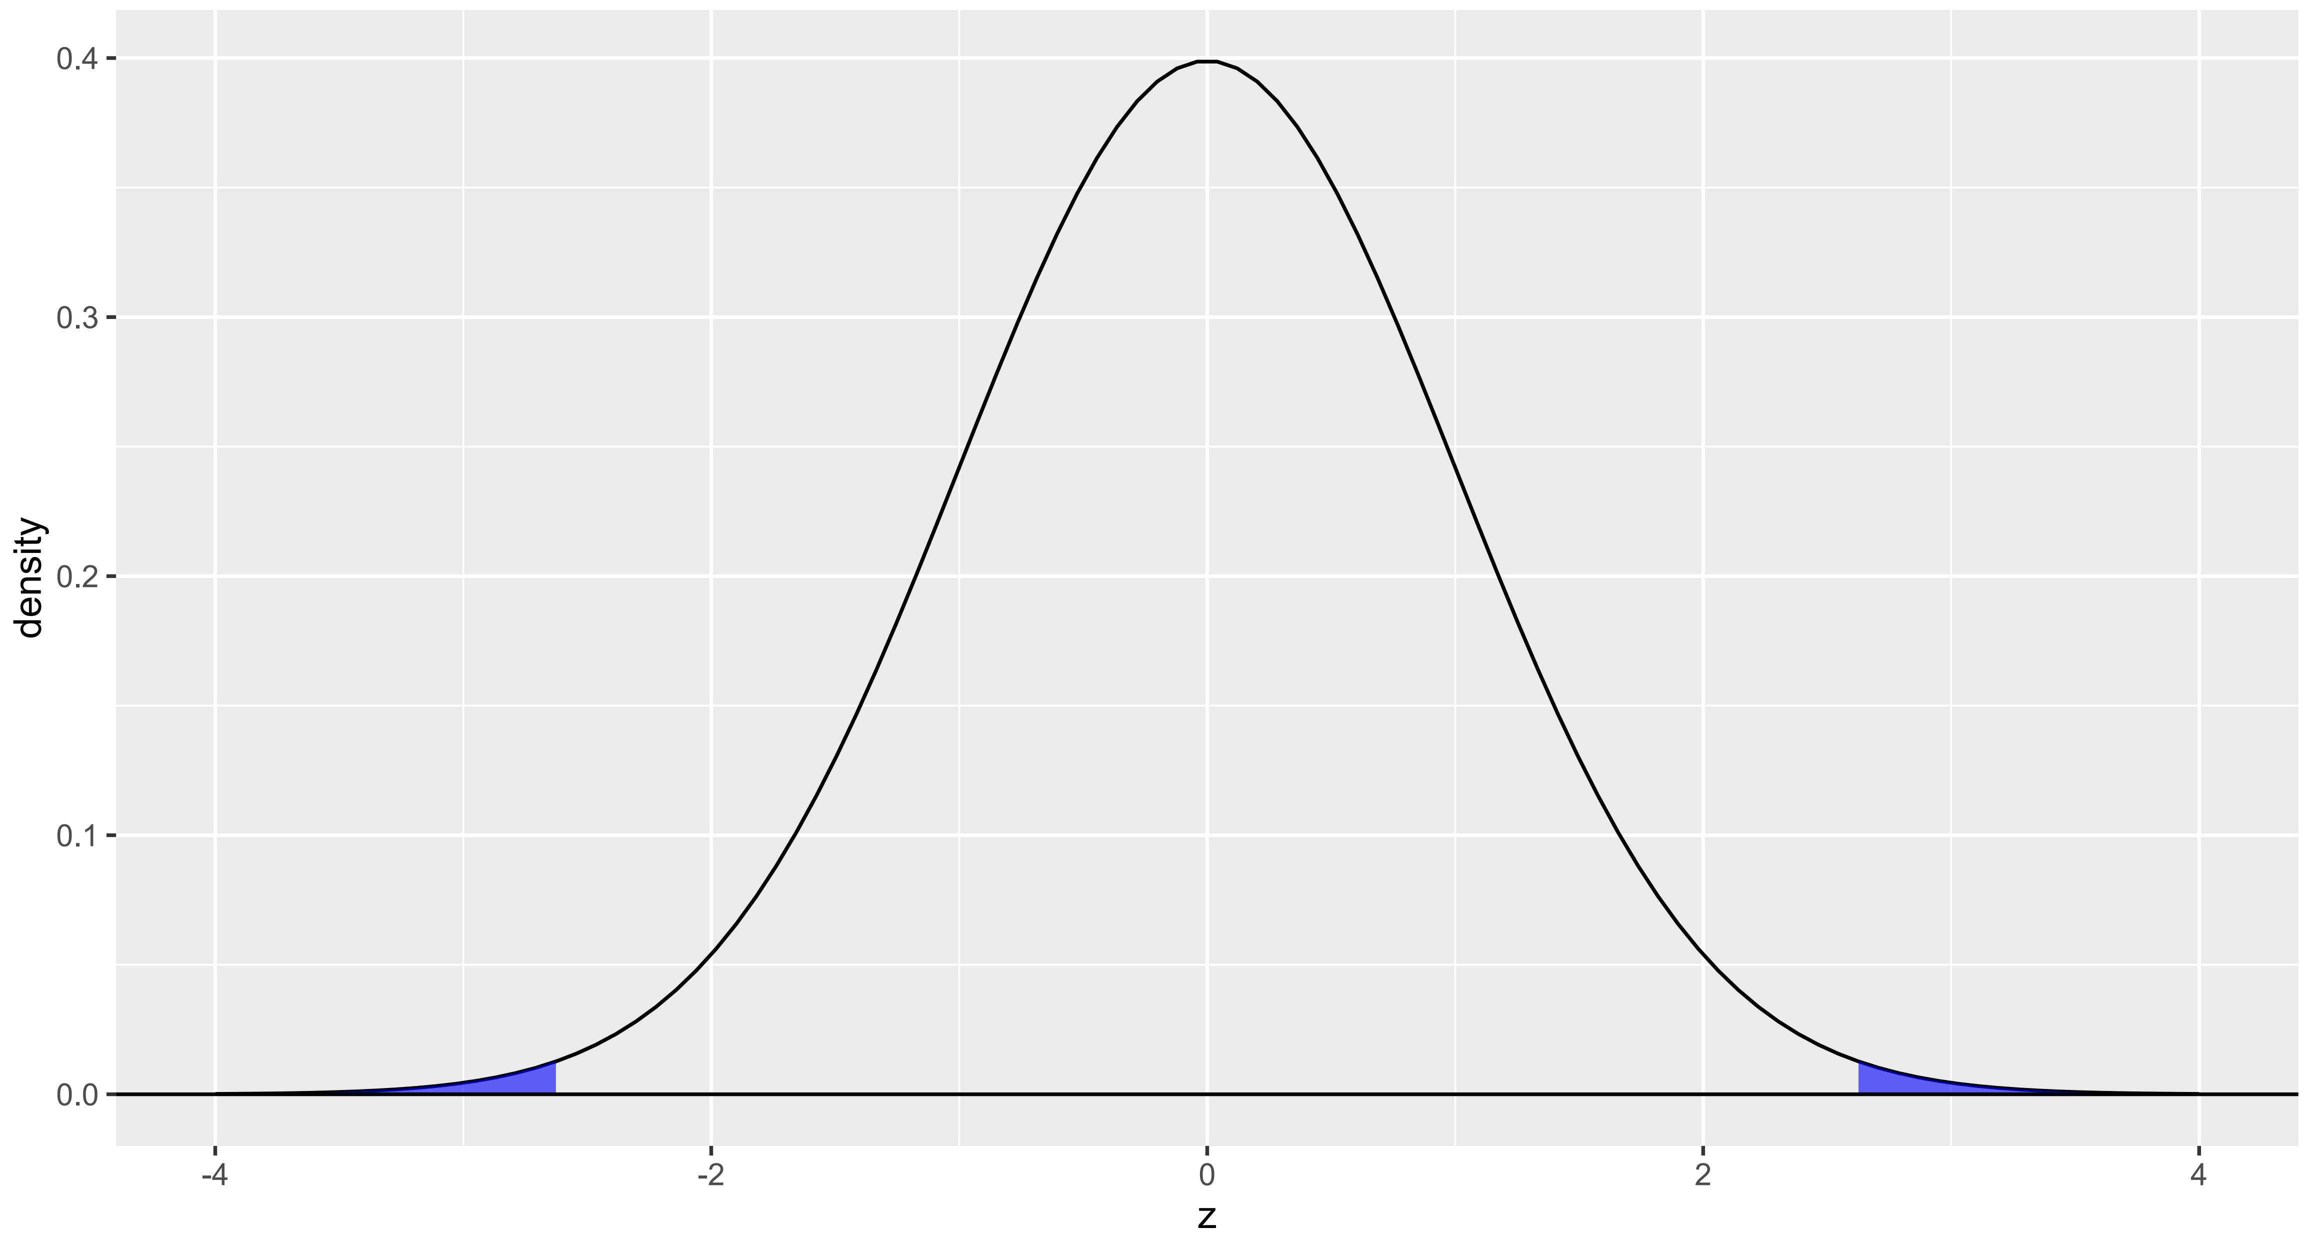

Insight

- The goal of this kind of work is insight.
- Statistical hypothesis tests are part of that process, but so are exploratory analysis, etc.

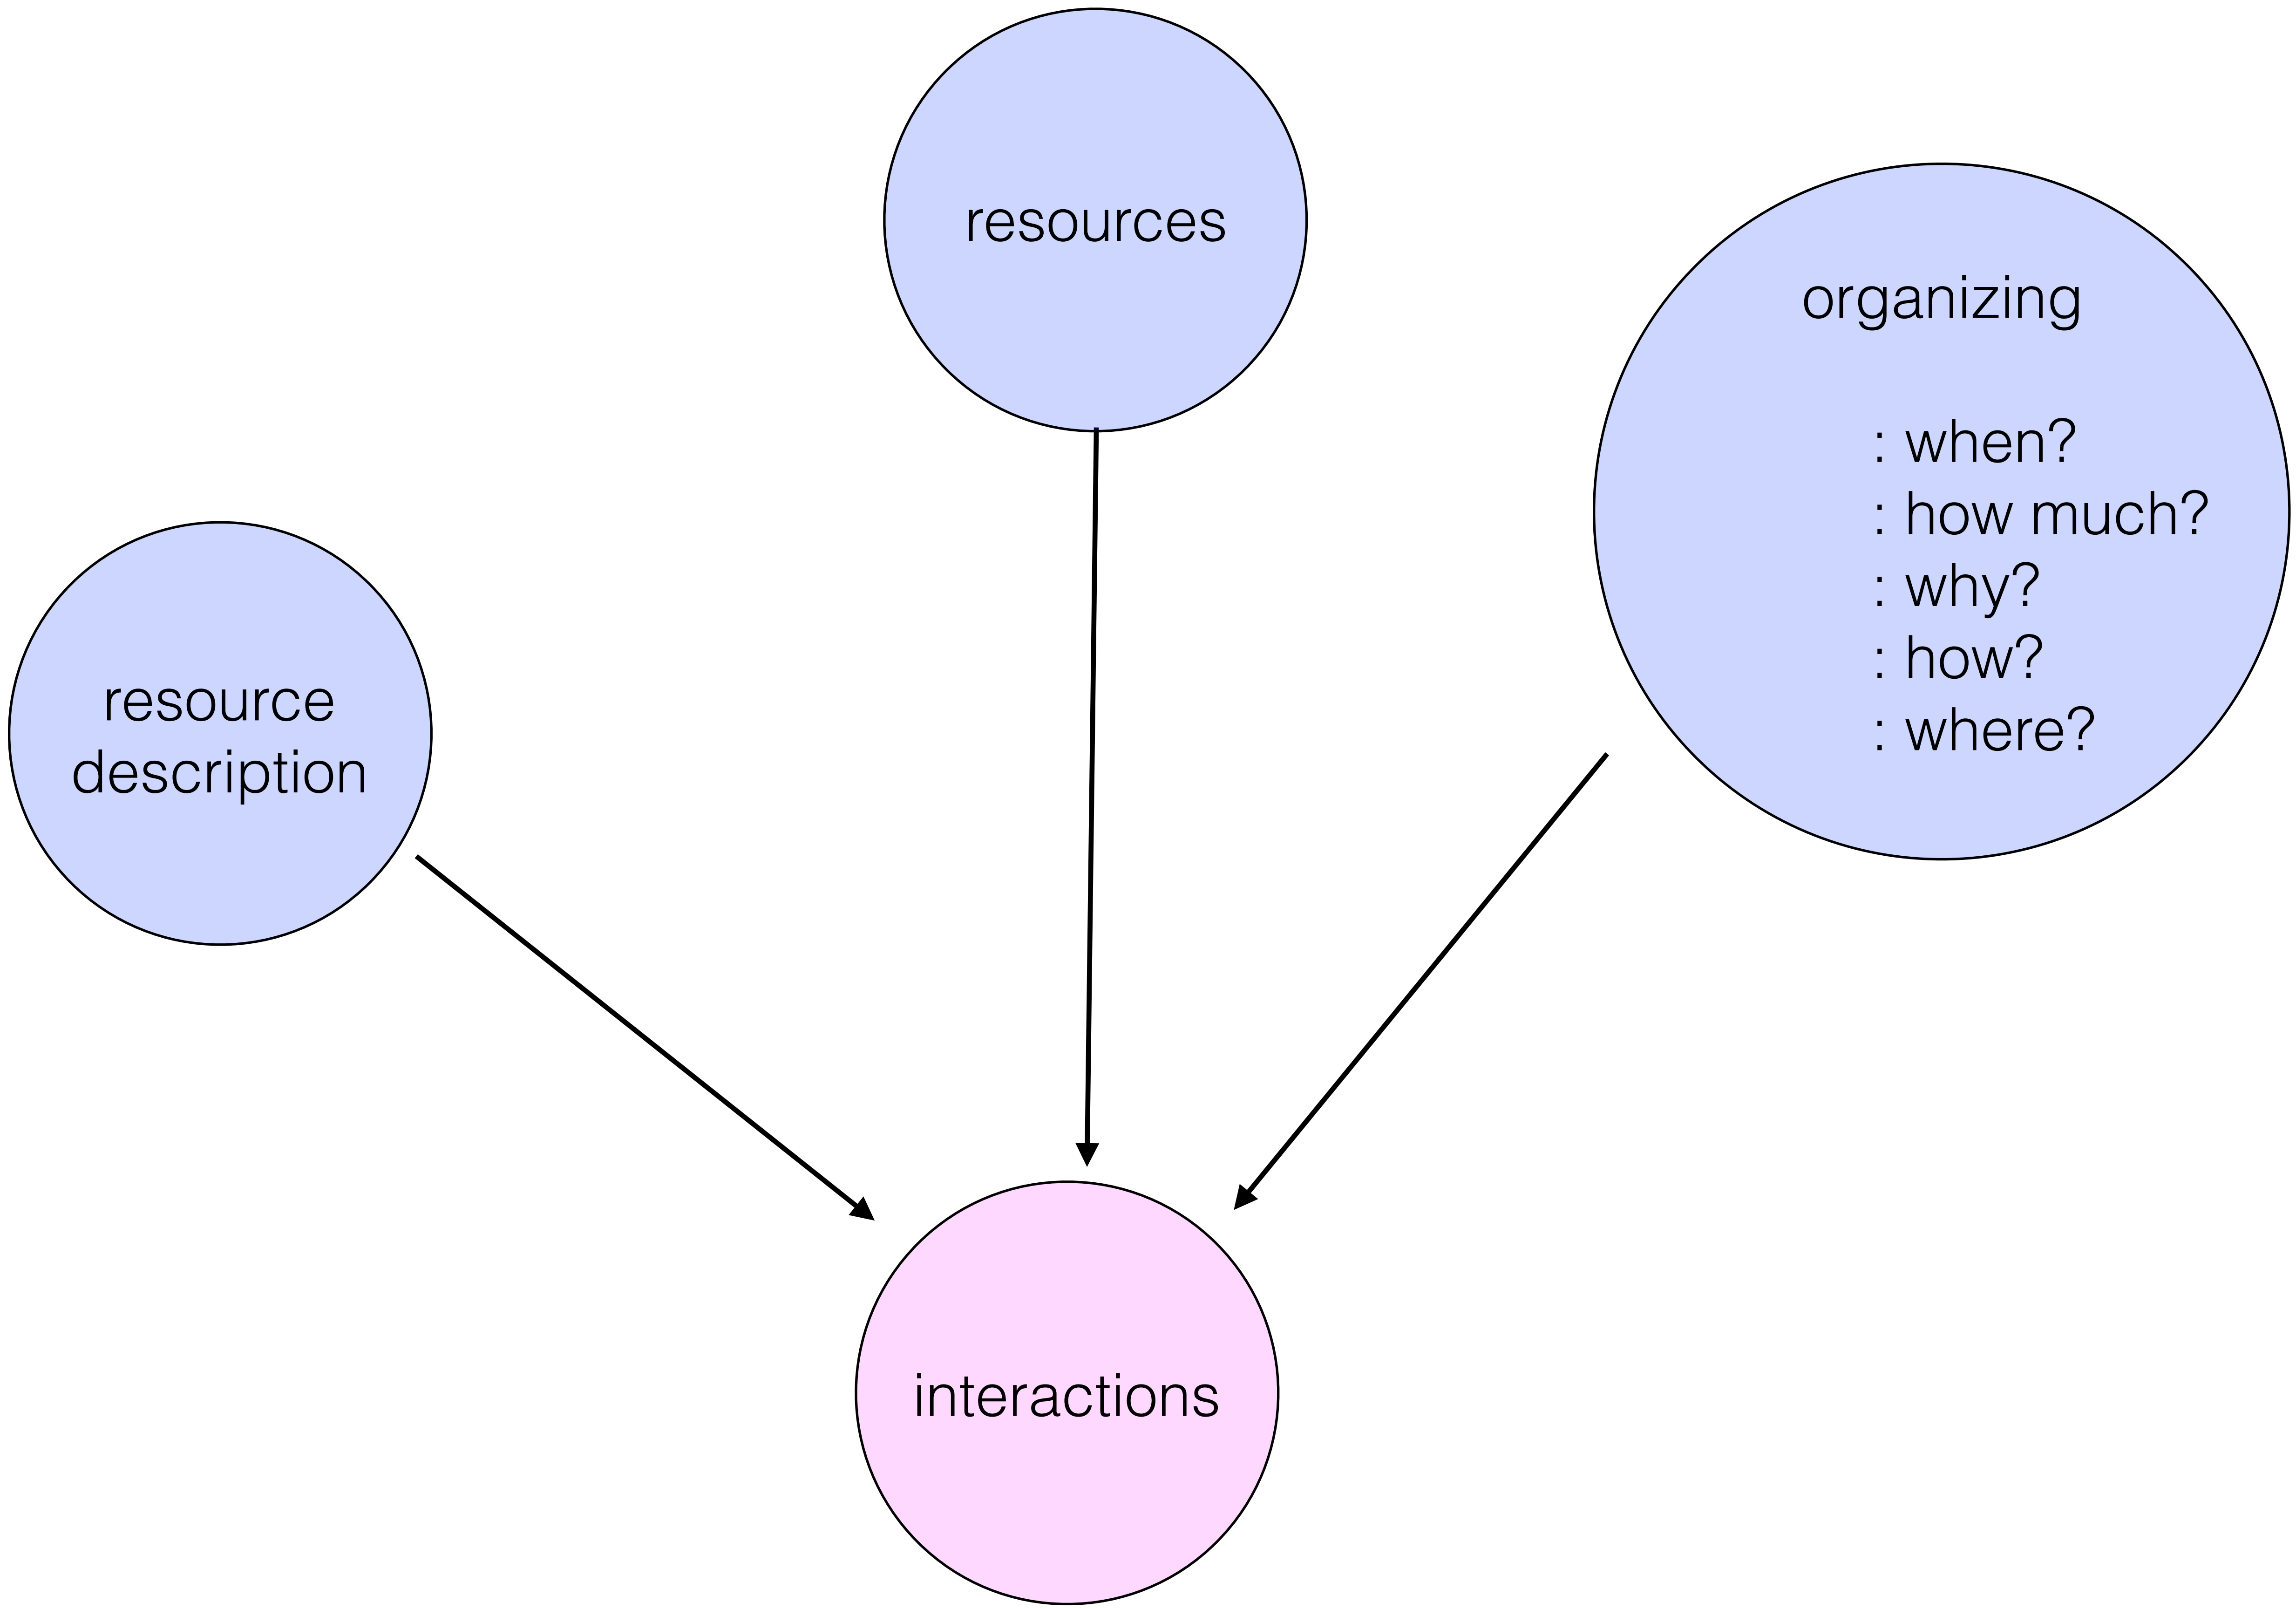
resources
organizing
: when?
: how much?
: why?
: how?
: where?
resource
description
interactions

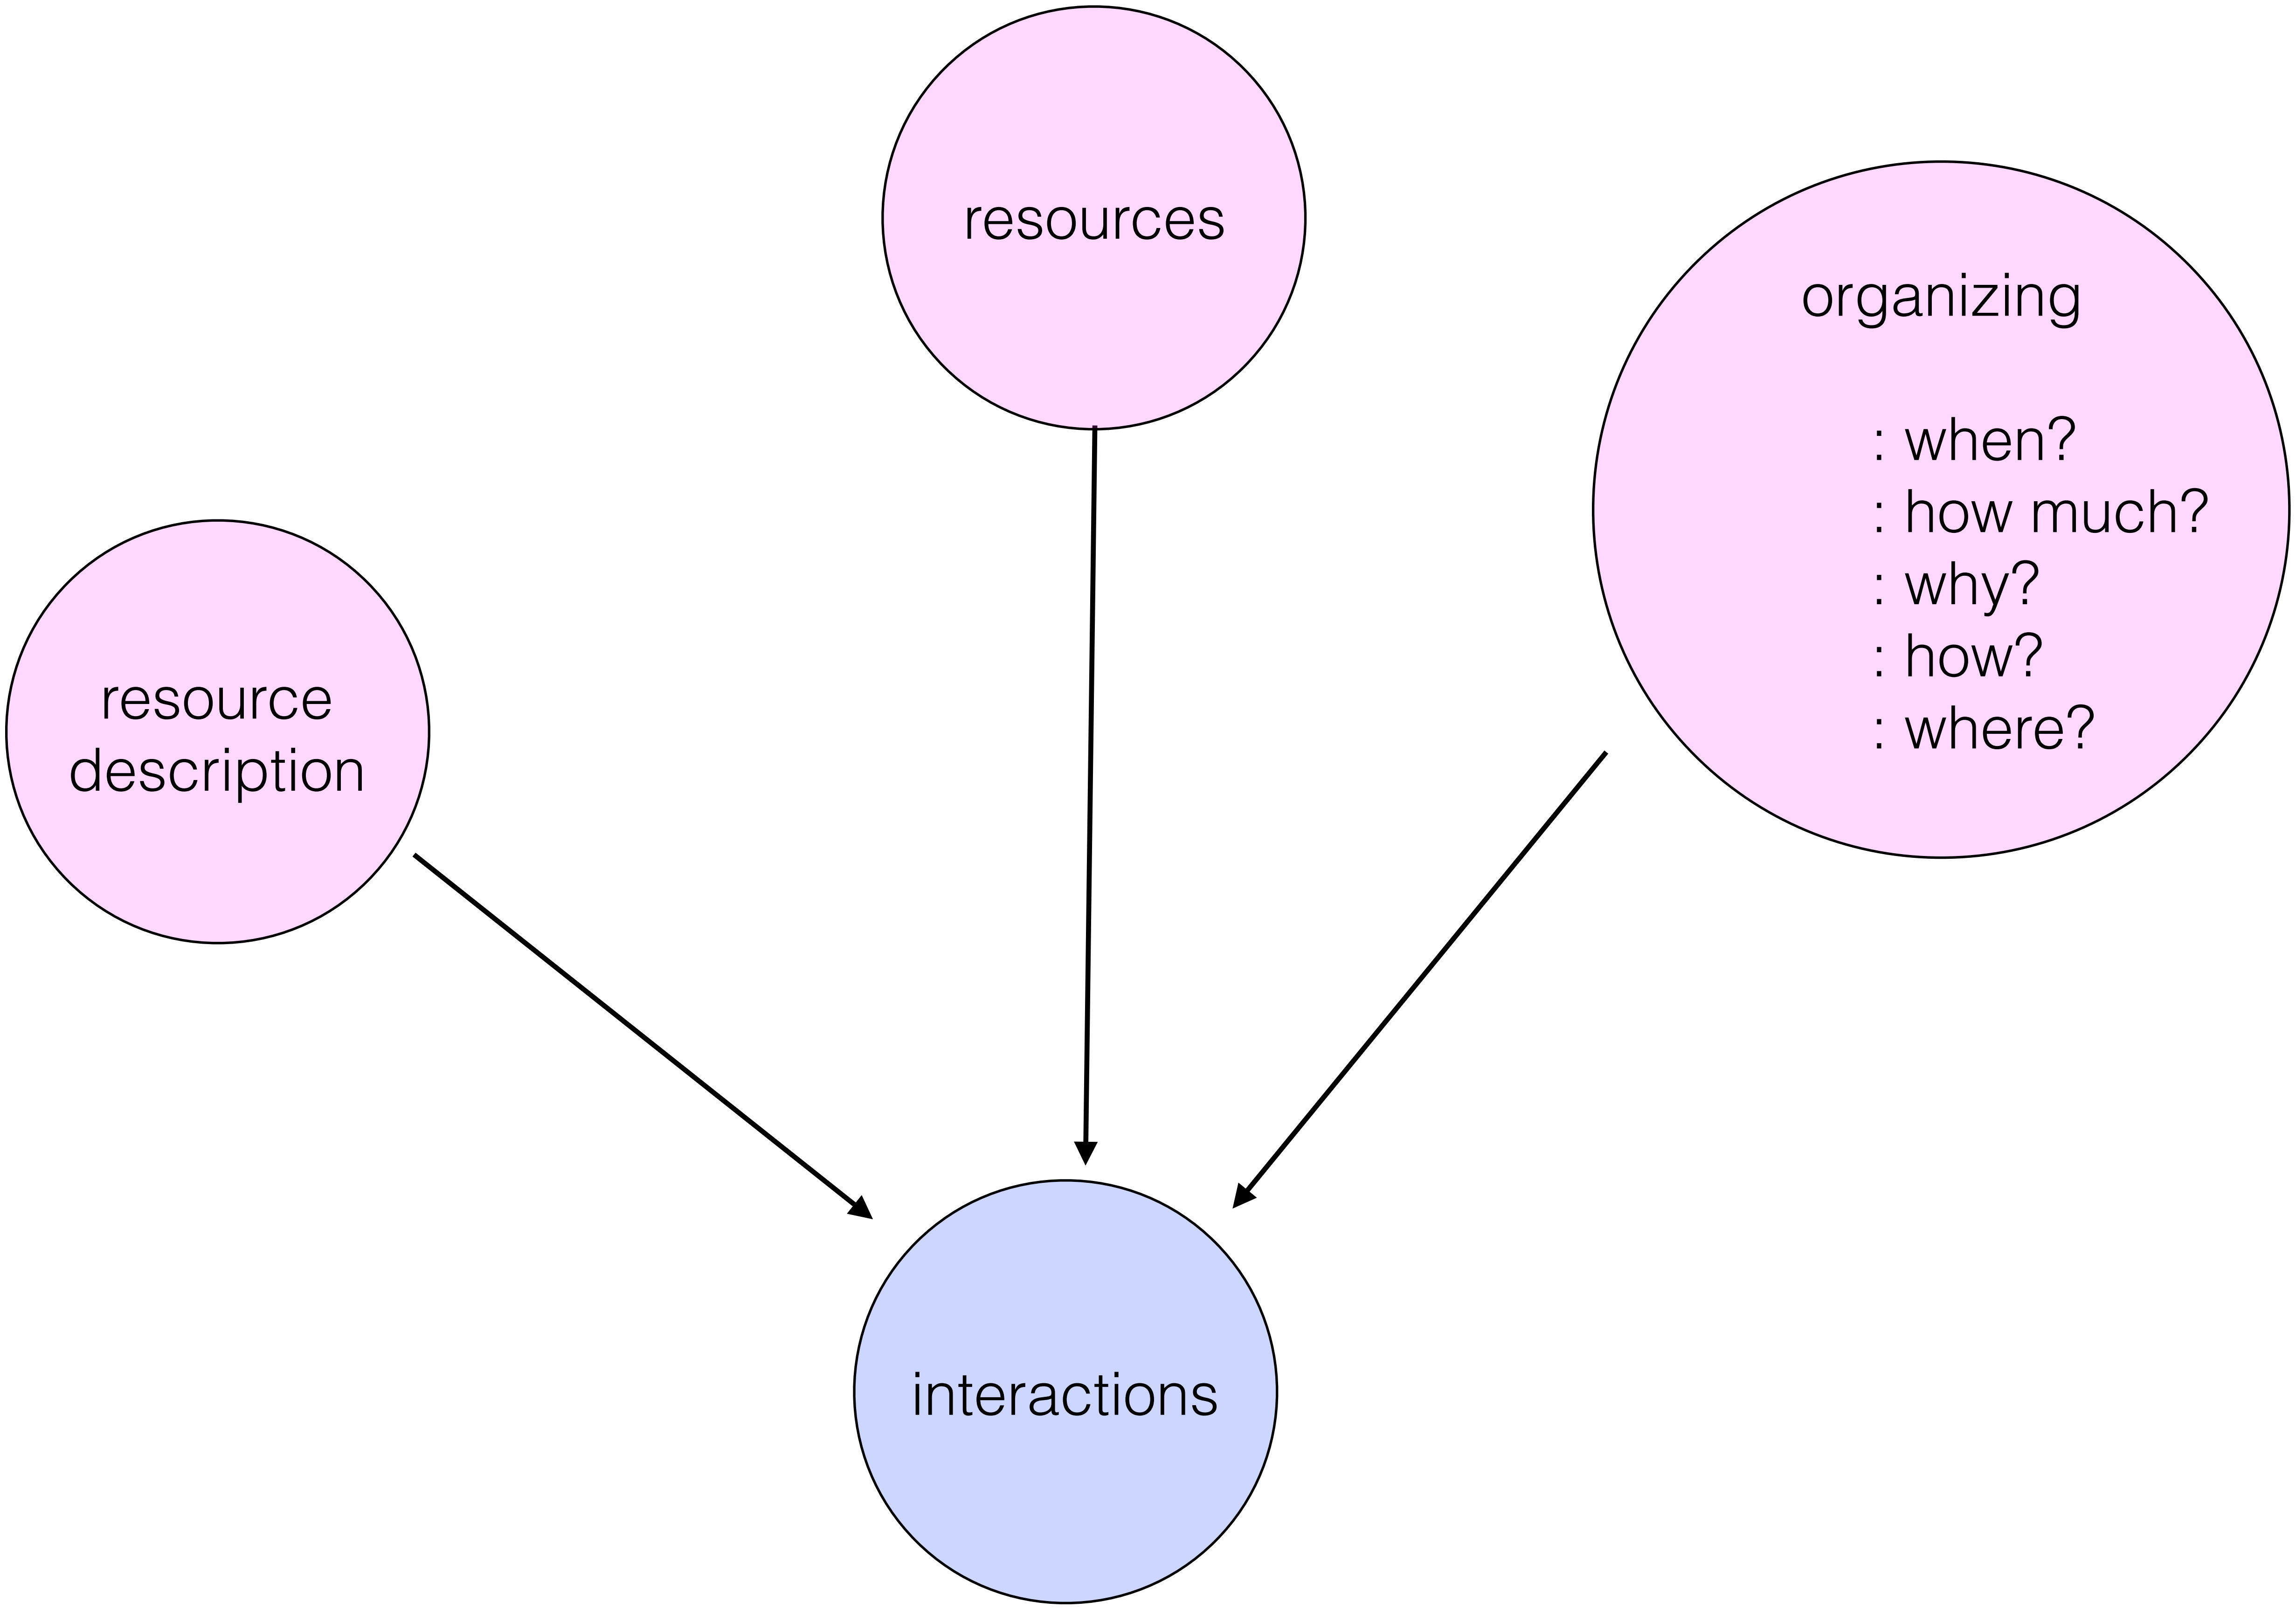
resources
organizing
: when?
: how much?
: why?
: how?
: where?
resource description
interactions

- How much data do you need to be able to assert your belief in an outcome with confidence? (Statistical power)

- What relationship does the data have to the outcome?

- How do you interpret the measurements you make?

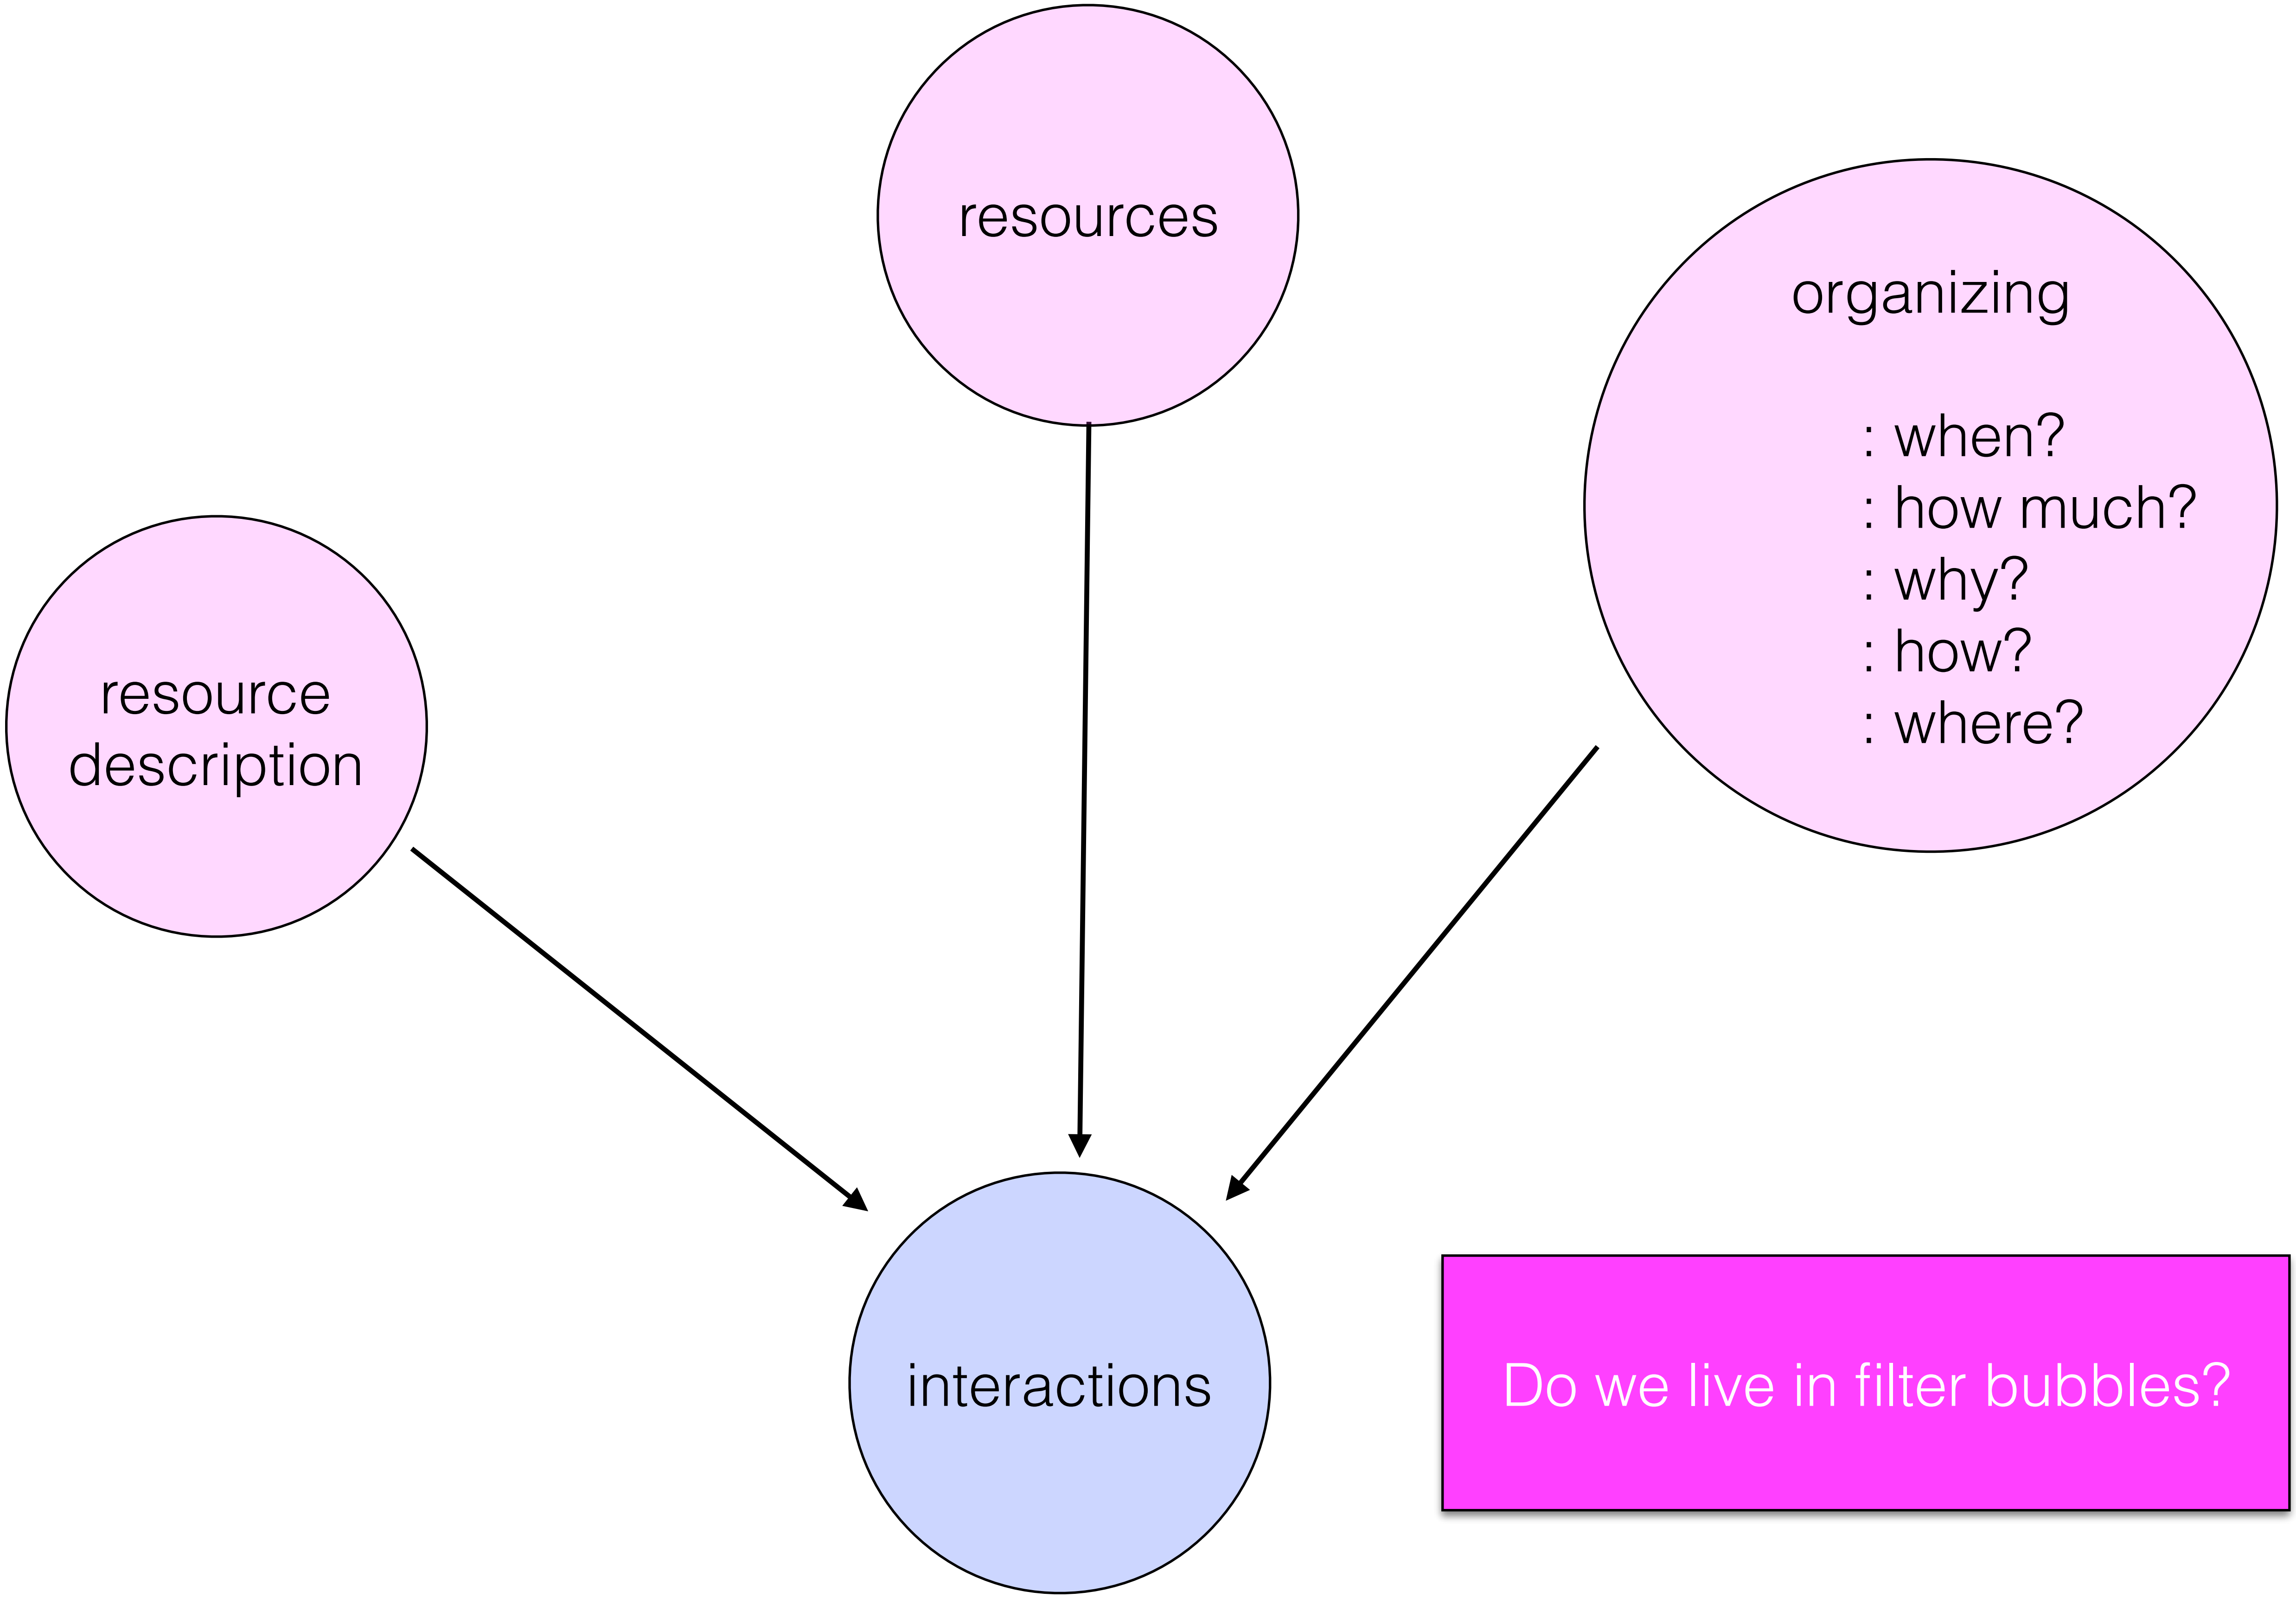
resources
organizing
: when?
: how much?
: why?
: how?
: where?
resource description
interactions
Do we live in filter bubbles?

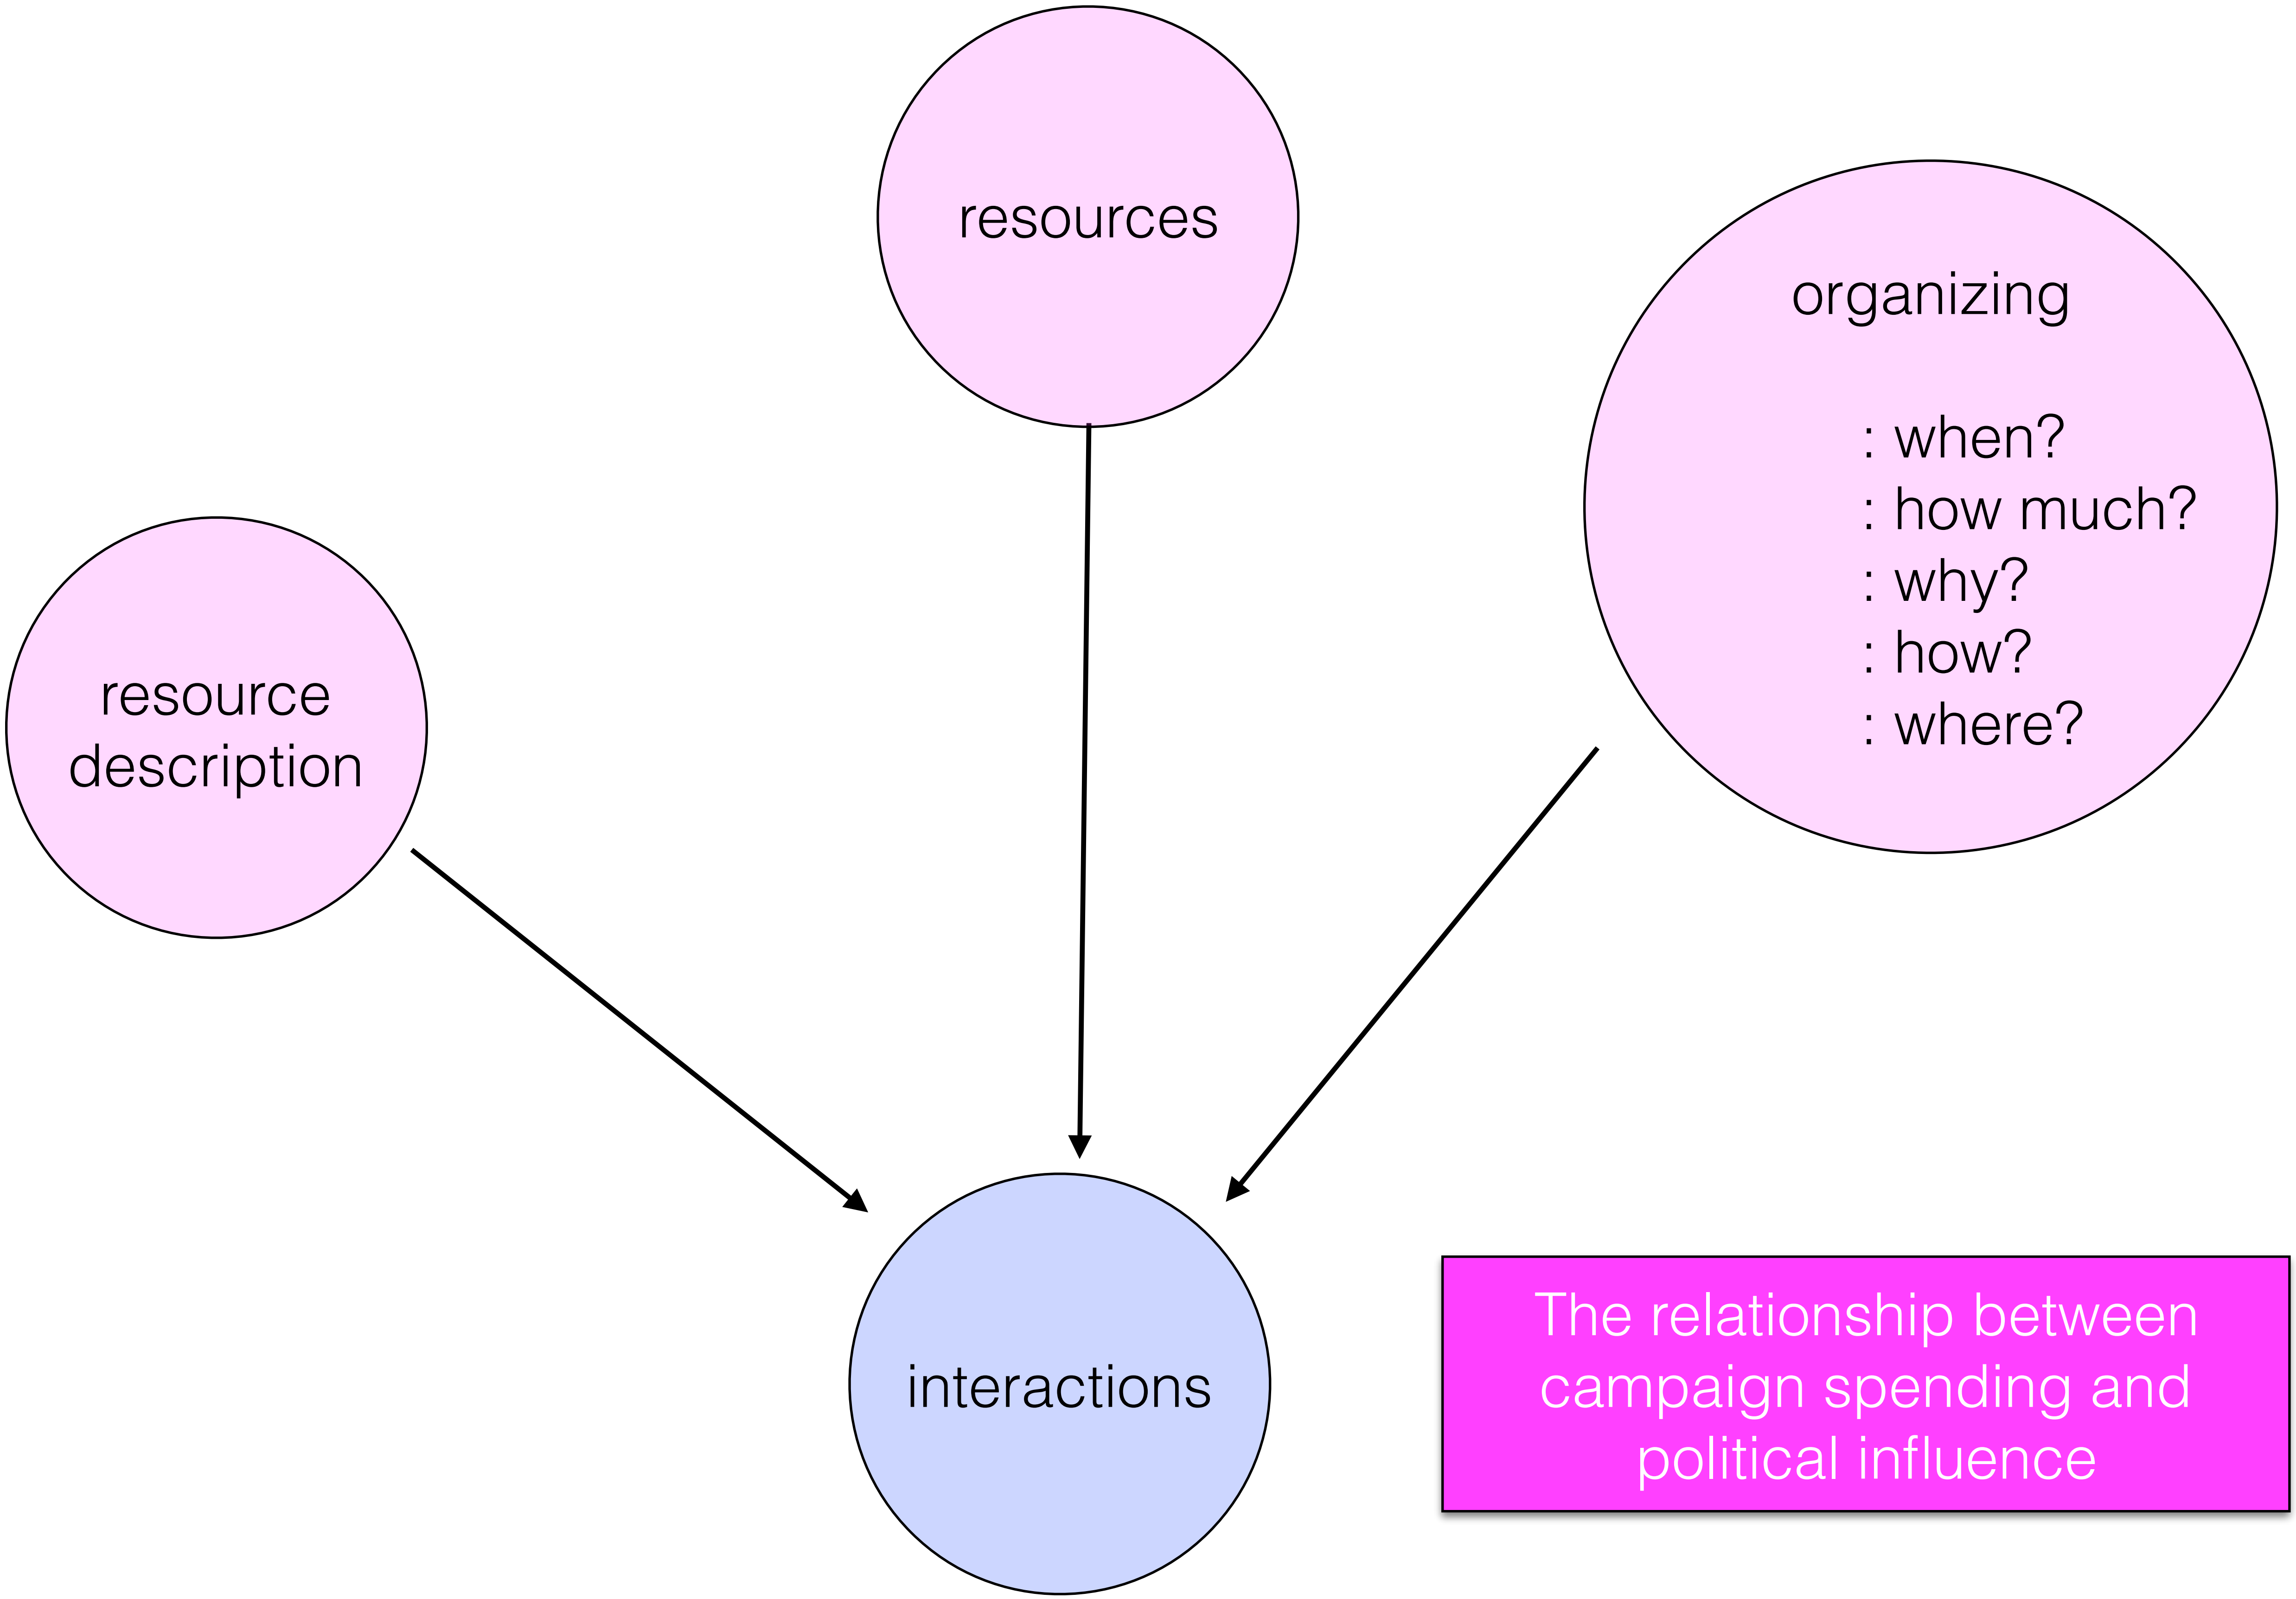
resources
organizing
: when?
: how much?
: why?
: how?
: where?
resource description
interactions
The relationship between campaign spending and political influence

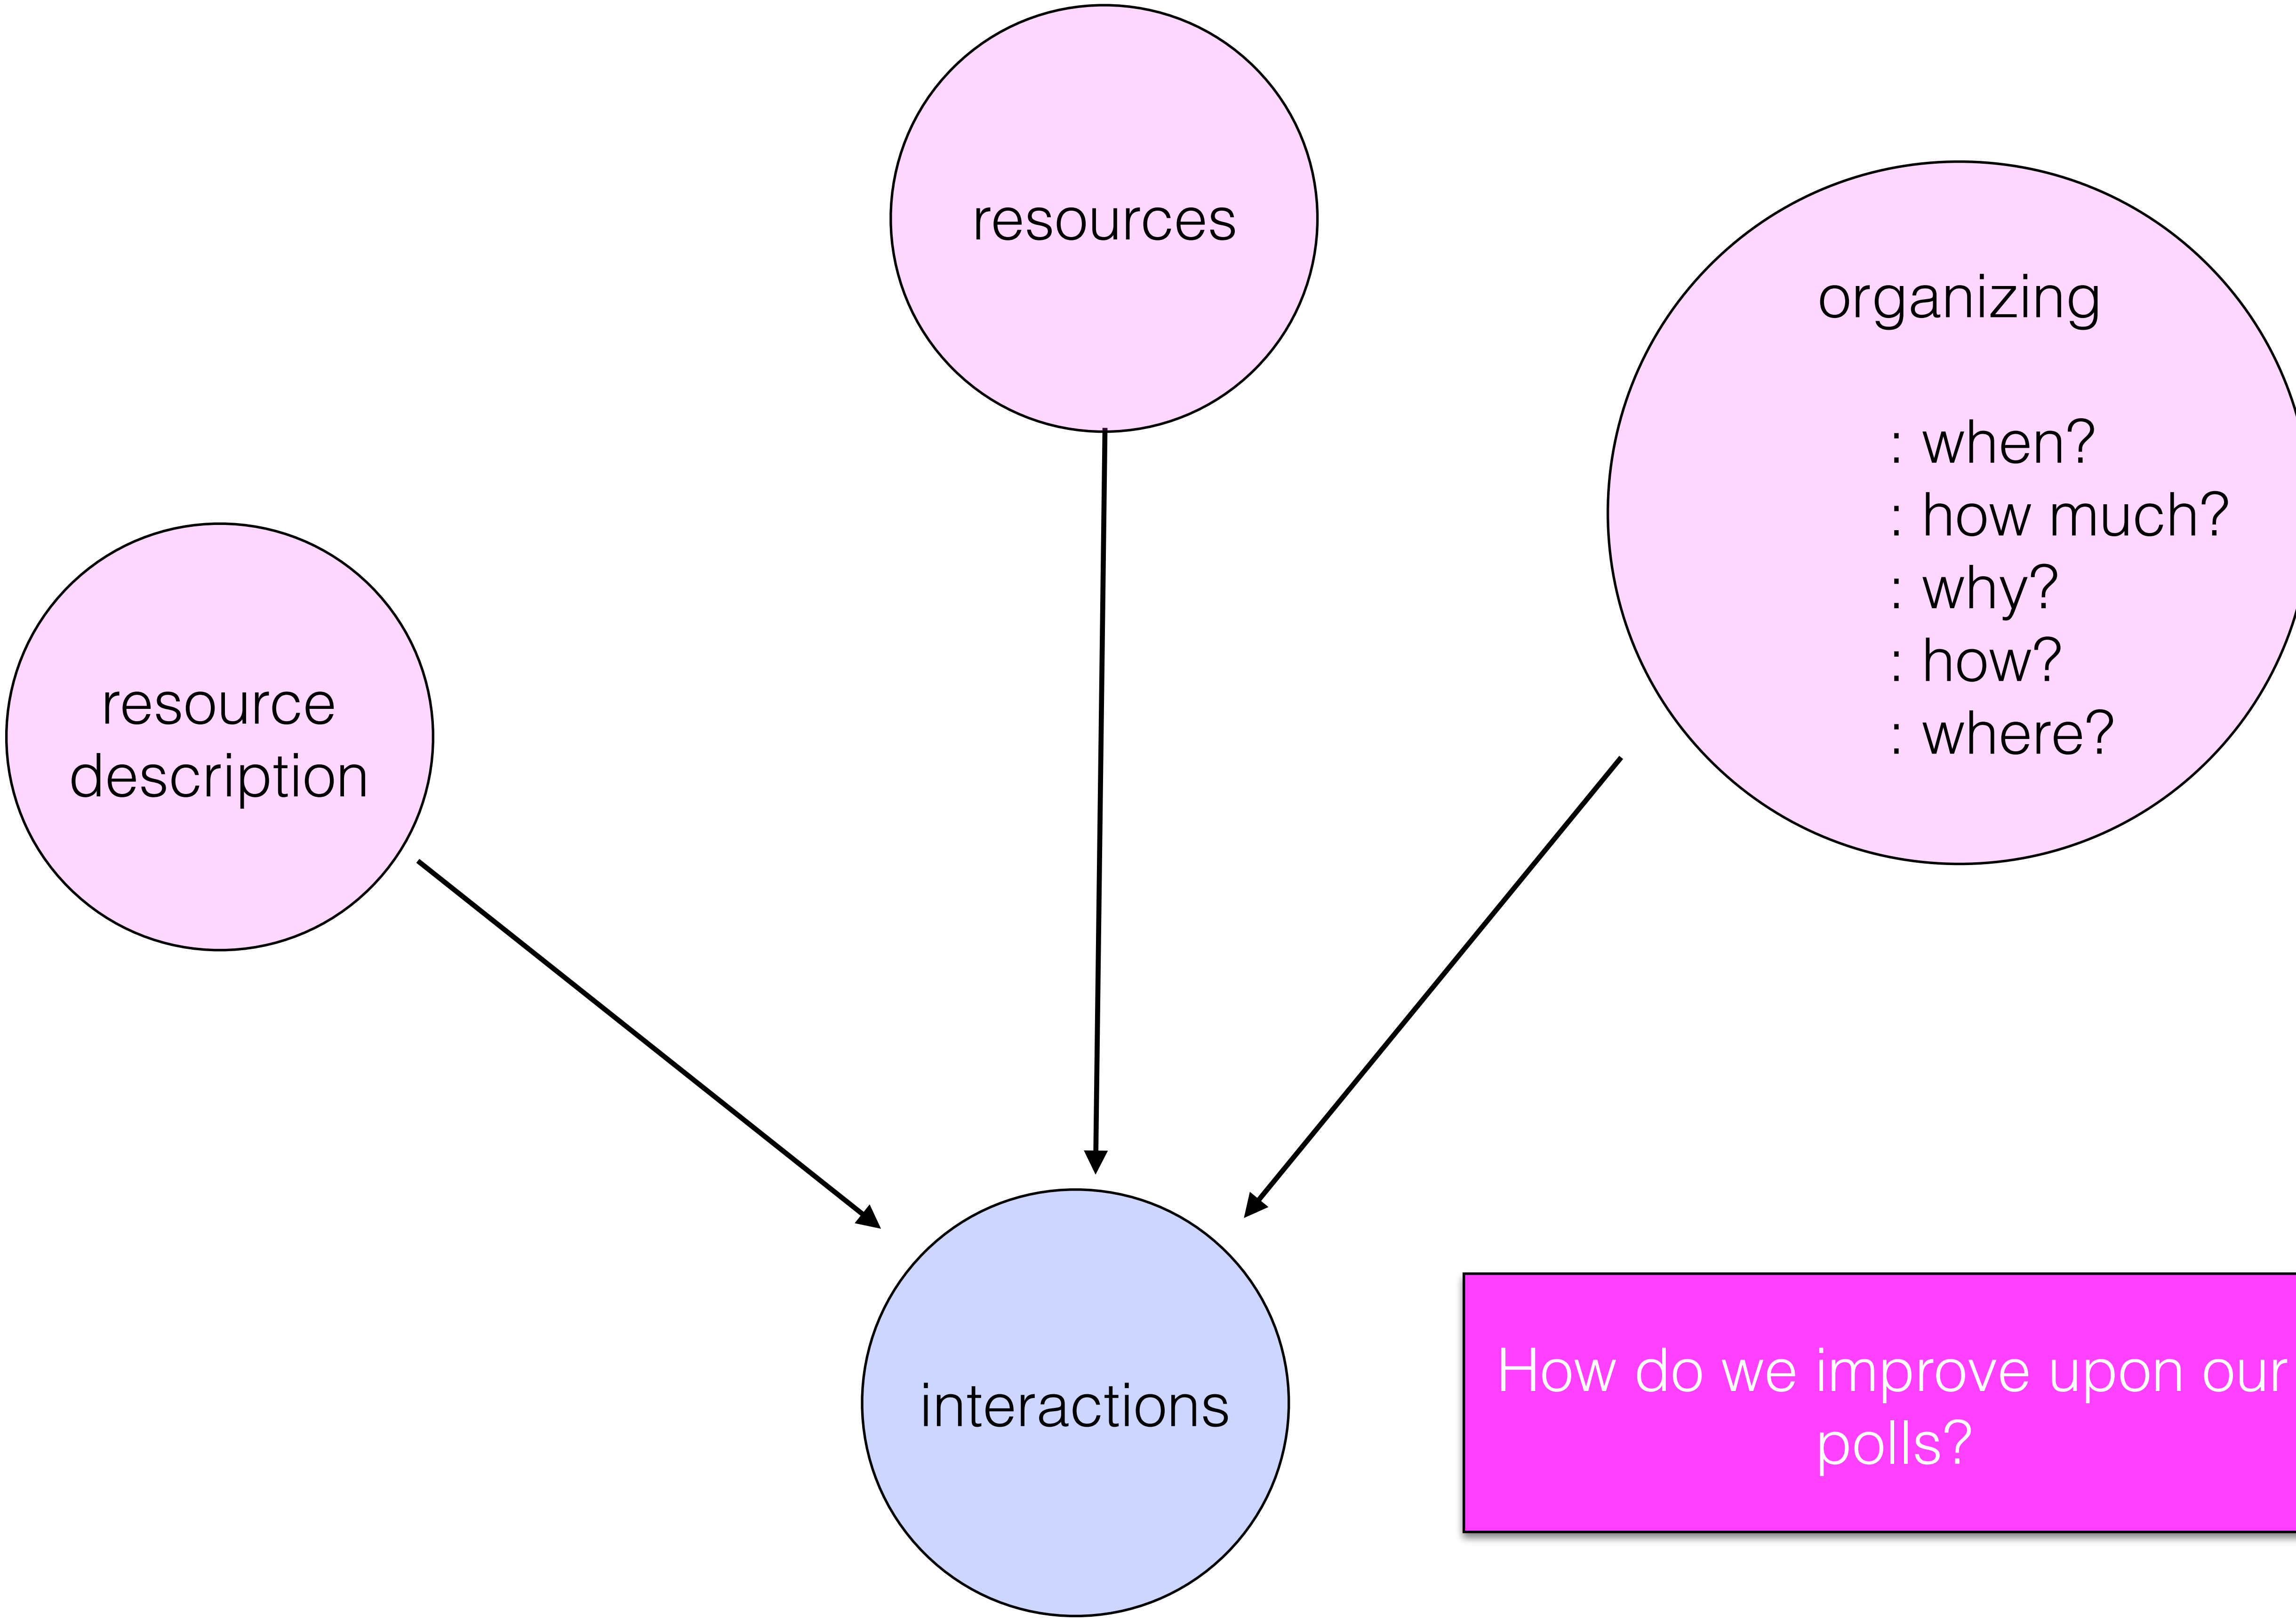

How do we improve upon our polls?

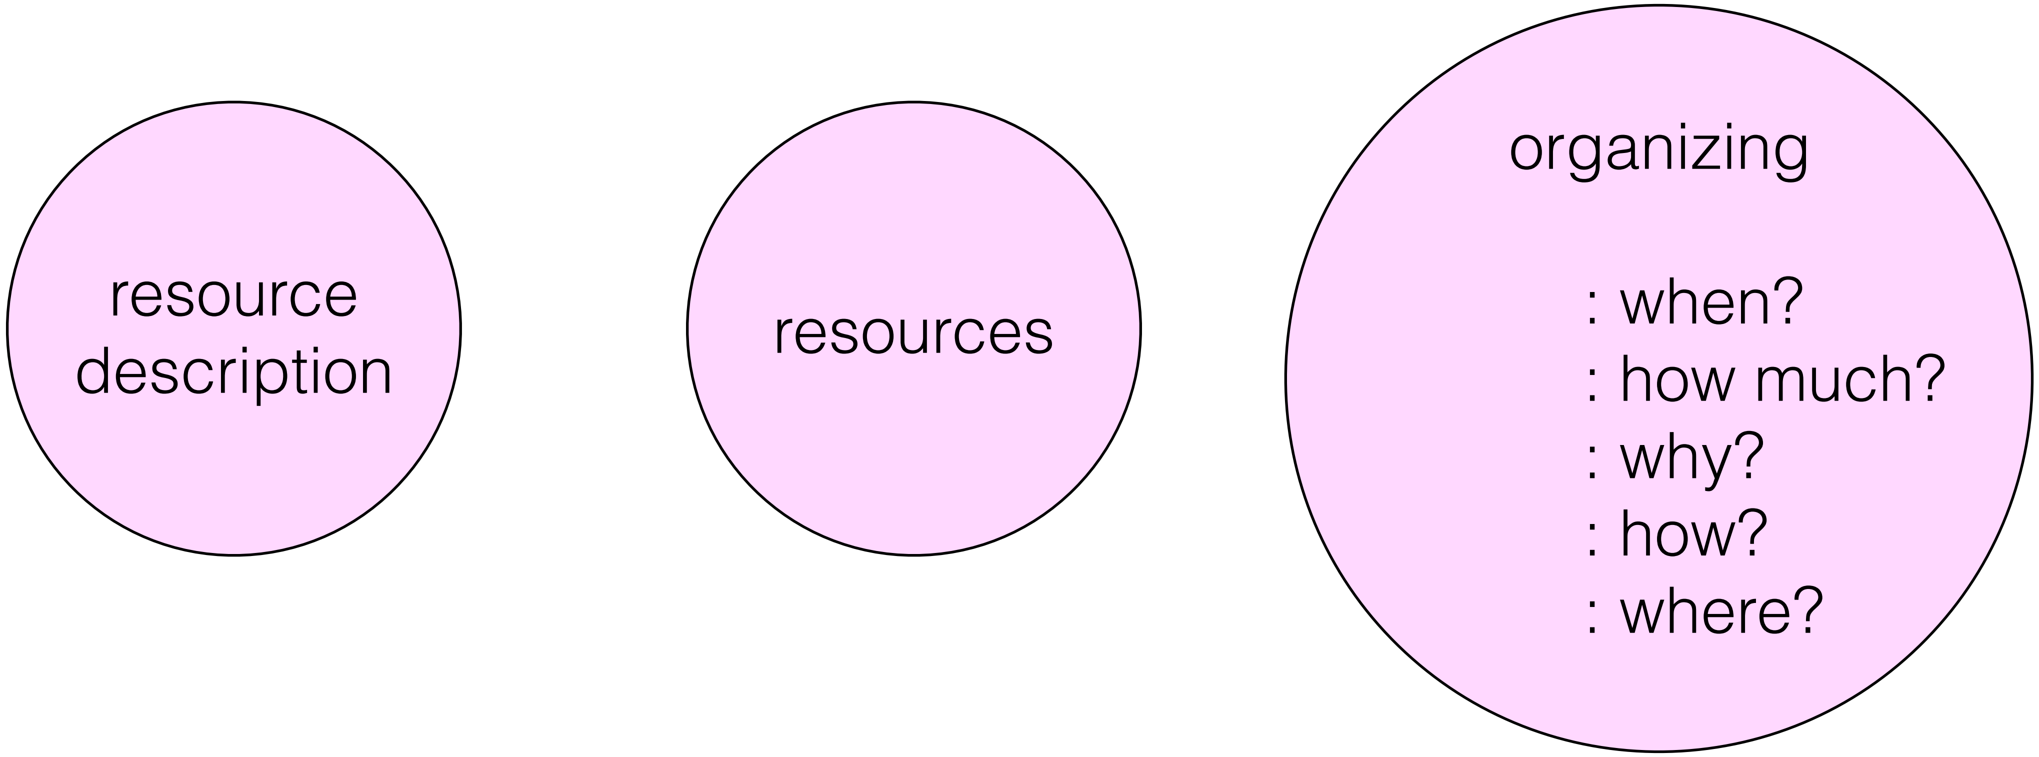

Question
Do we live in filter bubbles?
How do we measure the relationship between campaign spending and political influence?
How do we improve upon our polls?

Evaluating interactions

- Efficiency (timeliness/cost-effectiveness)
- Effectiveness (accuracy/relevance)
- Satisfaction (user)

Evaluating inferences

- Validity of results
    - Face validity (is patently untrue?)
    - Social validity (does it make a difference?)
    - Structural validity (is there meaningful theory for why the results come out how they did?)
    - Semantic validity (is the concept we're thing we're evaluating the one we actually are?)