Computational Cognitive Science

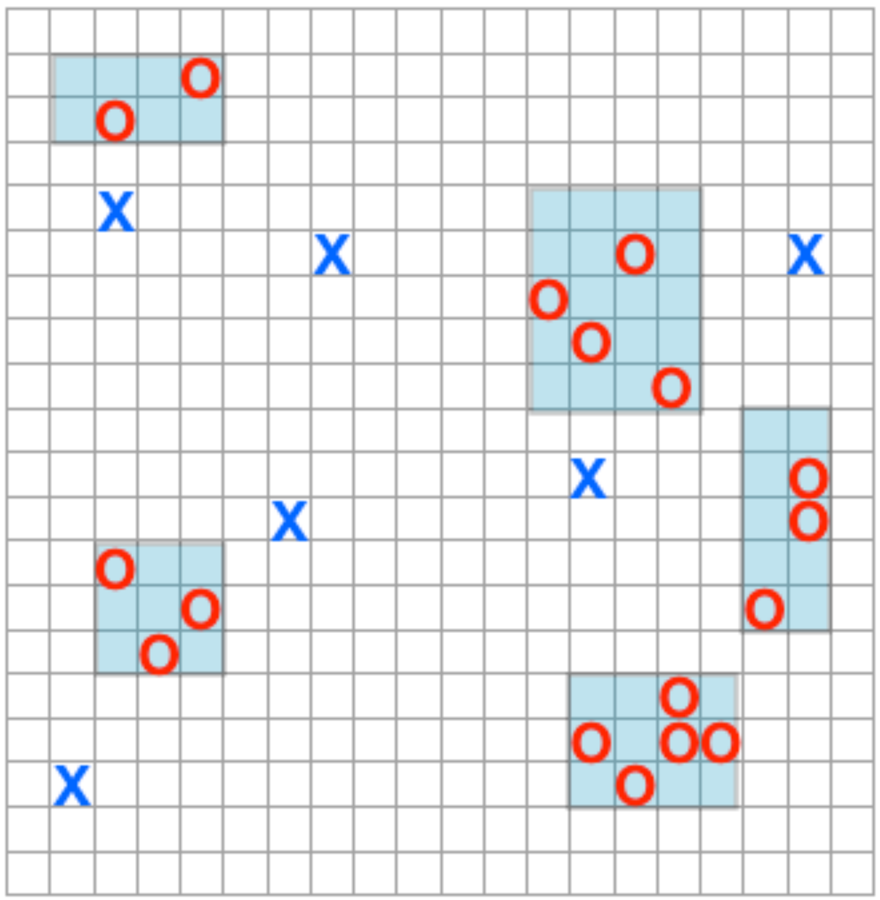

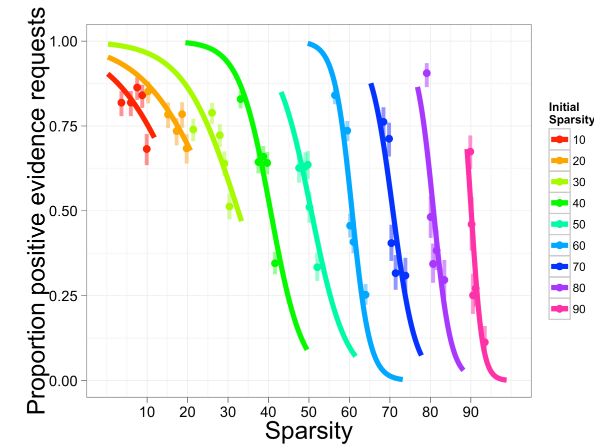

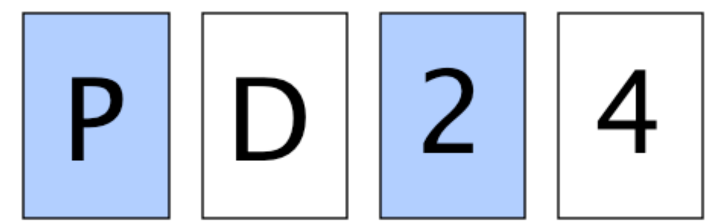

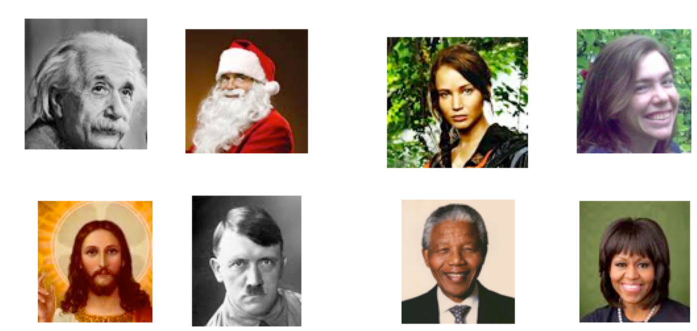

Lecture 22: What kind of information are people sensitive to?

Last time

- There is a difference between strong sampling and pedagogical sampling

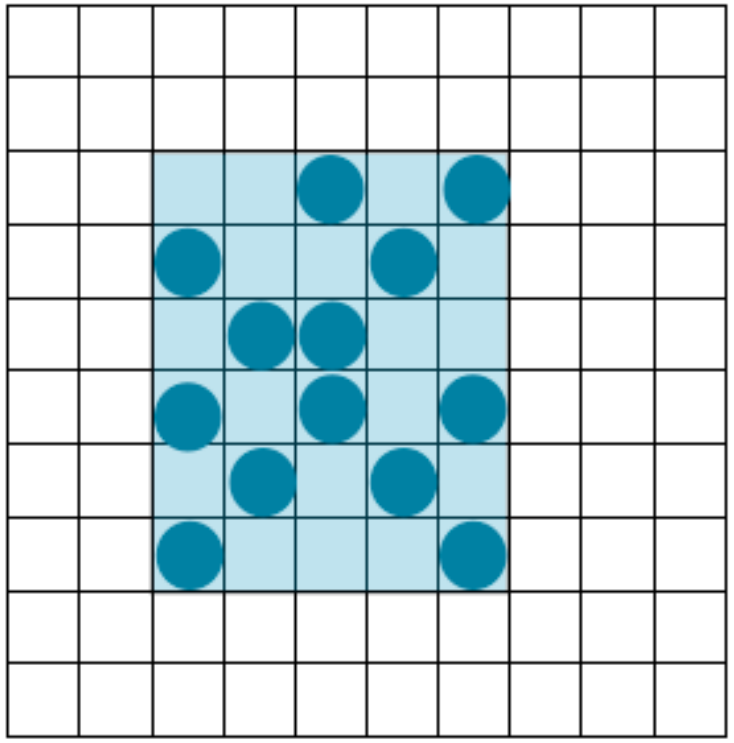

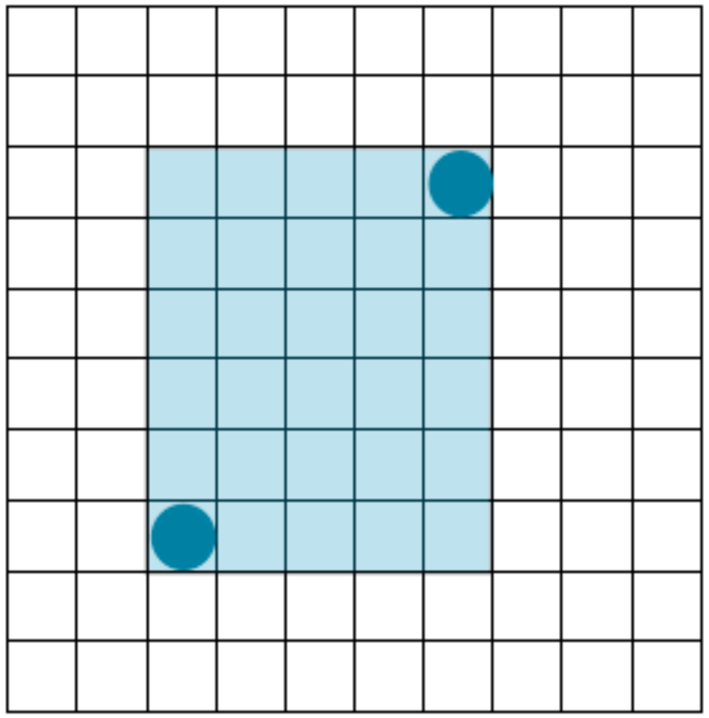

Last time

- There is a difference between strong sampling and pedagogical sampling
- People are sensitive to this difference

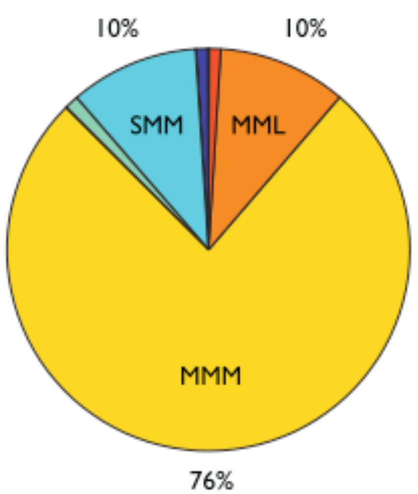

(a) Strong sampling predictions

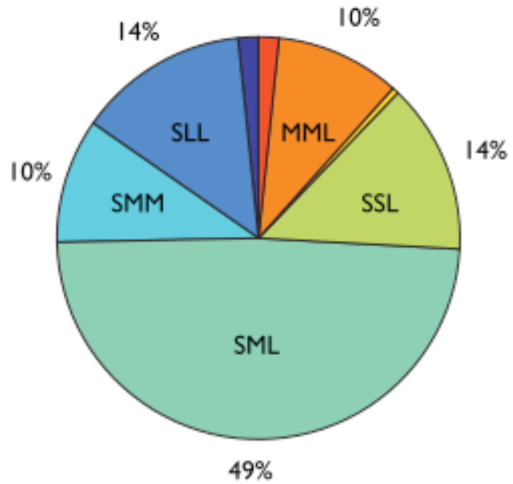

(b) Pedagogical model predictions

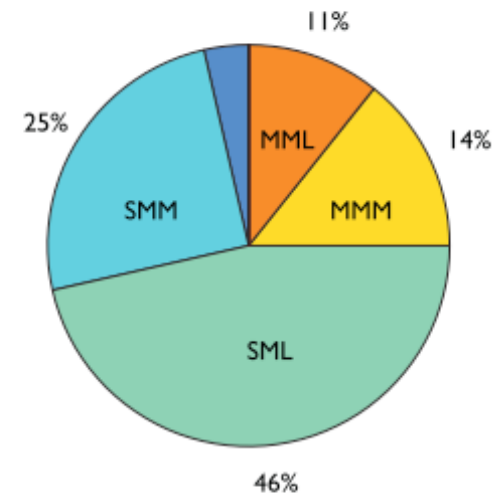

(c) Experiment results

Last time

People are sensitive to non-obvious aspects of the data. Another is how informative different datapoints are. What about that?

Plan

- Lecture 1 [short]: Another kind of sampling: pedagogical
  - Model for pedagogical sampling
  - Sensitivity to pedagogical data
  - Double-edged sword of pedagogy
- Lecture 2 [long]: Sensitivity to the informativeness of the data
  - Confirmation bias and the positive test strategy (PTS)
  - When is positive evidence most useful?
  - When do people use a PTS?
  - Are people sensitive to evidence utility in general?

Sensitivity to informativeness of data

Some of the research in the last lecture suggests that people are generally sensitive to non-obvious aspects of where data come from...

.. but other research suggests that people are surprisingly *irrational* when it comes to weighing the utility of data, and knowing what is informative or useful

Guess the number rule

I'm thinking of a number rule. You should give me triples that you want to test in order to figure out the number rule. I'll let you know whether they follow it or not.

2 4 6

Guess the number rule

2 4 6

- People usually have hypotheses like “increasing by two” and guess things like “12 14 16” or “3 5 7”
- The problem is, these guesses are *confirmatory* only.

They aren’t like: 5 8 13

which would be false if the hypothesis were true

- As a result, most people don’t get the true number rule, which is “increasing numbers”

Confirmation bias can lead to error

- Especially when the correct hypothesis is a *superset* of the one you are entertaining

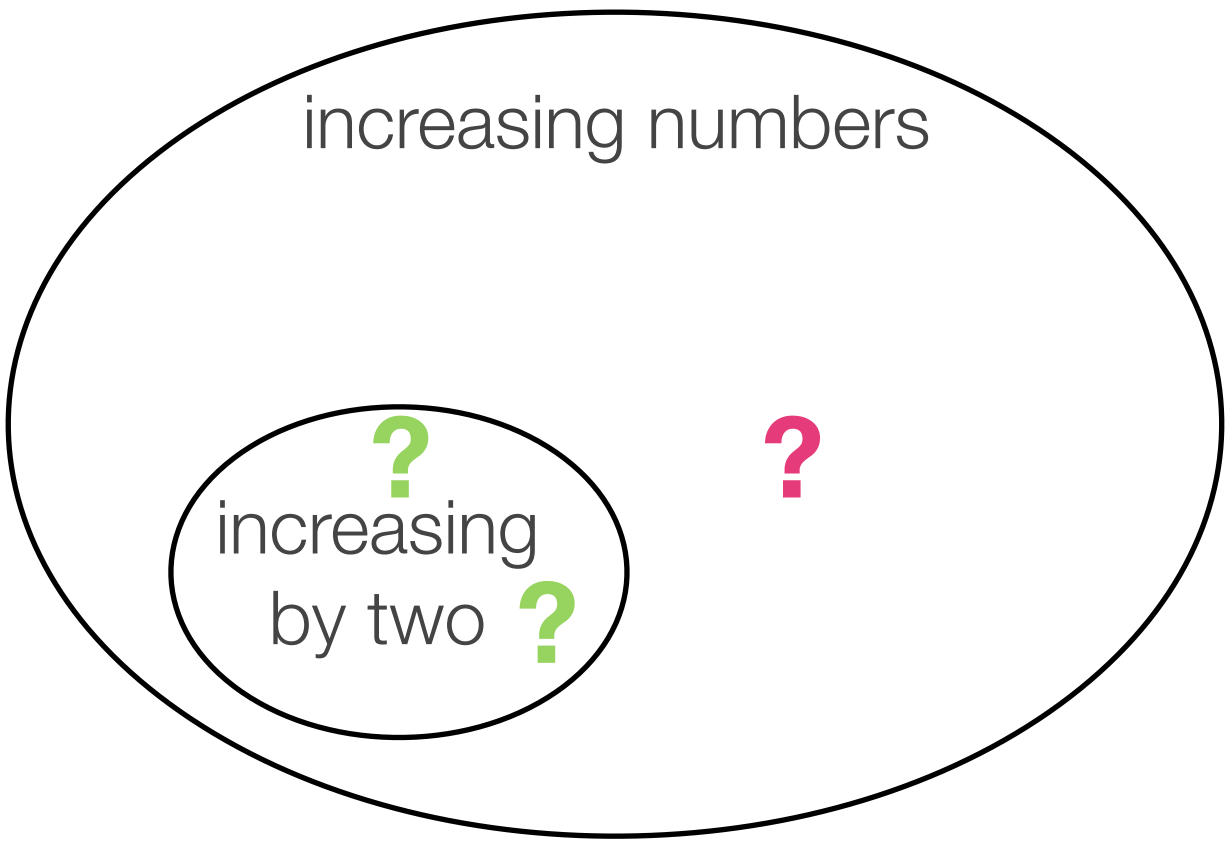

Confirmation bias can lead to error

- Especially when the correct hypothesis is a *superset* of the one you are entertaining

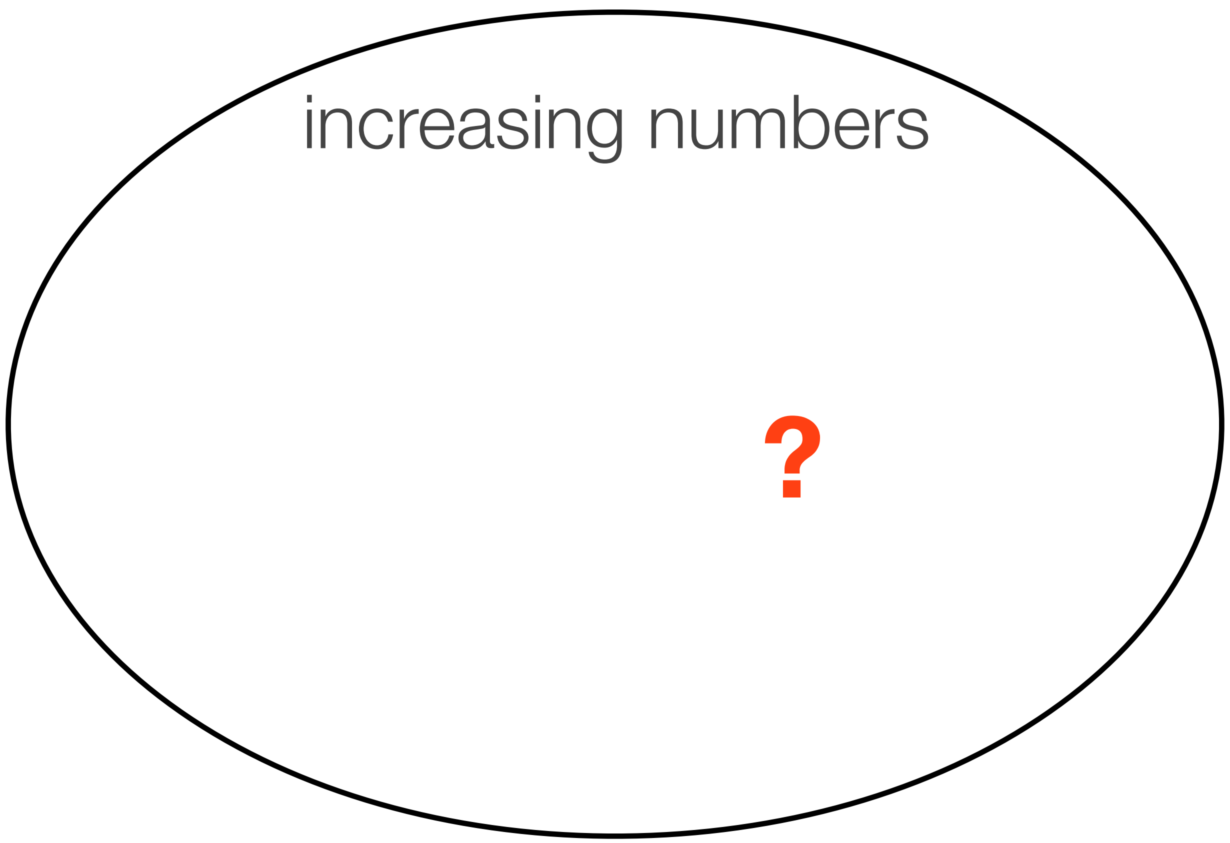

Confirmation bias crops up in many places

- You *should* flip the following ones

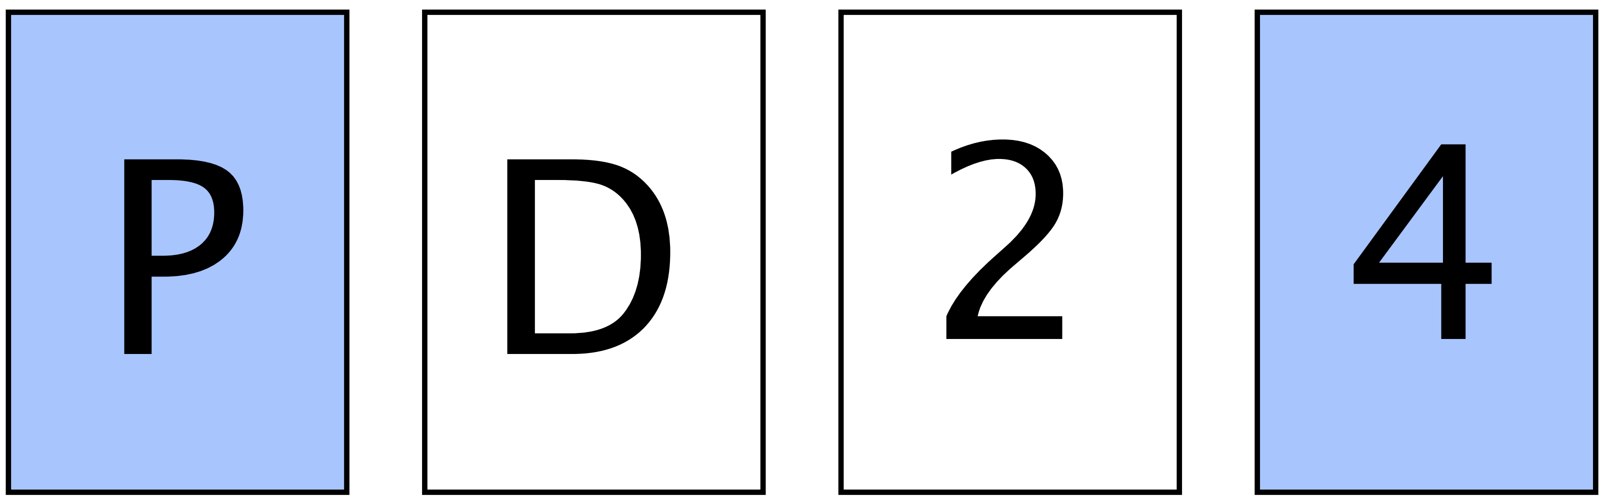

Cards have a letter and a number
Which do you flip to test the hypothesis

"If P then 2"

Confirmation bias crops up in many places

- Most people flip:

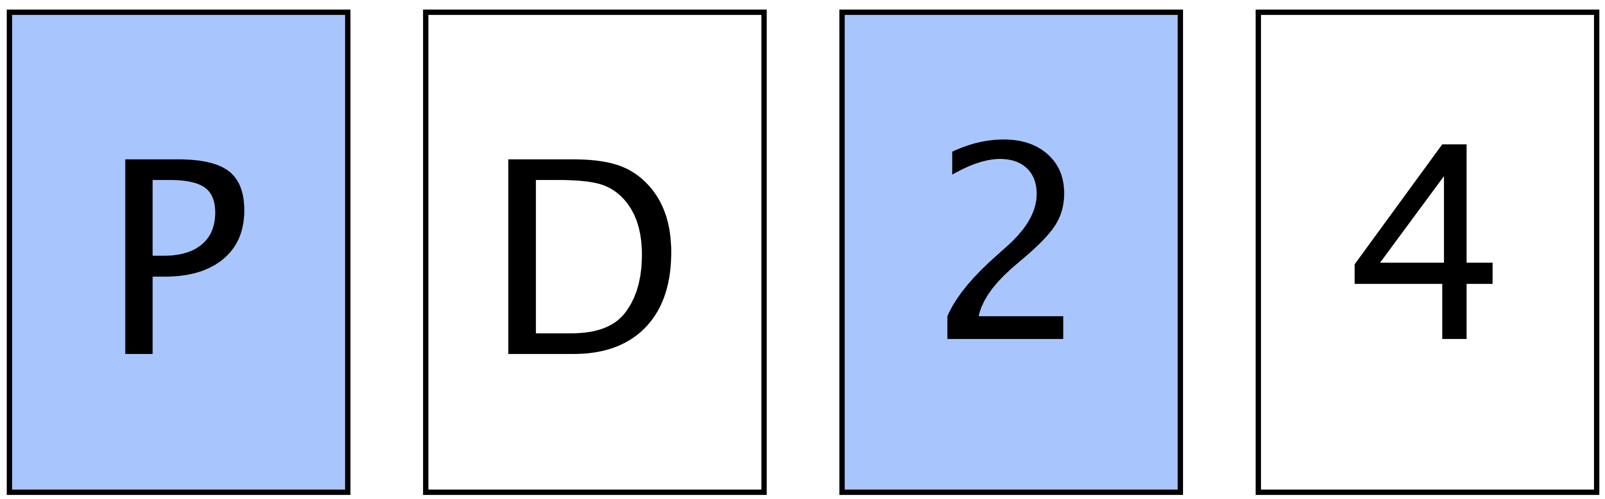

Cards have a letter and a number
Which do you flip to test the hypothesis

"If P then 2"

Confirmation bias crops up in many places

They are not trying to *falsify* their hypothesis!

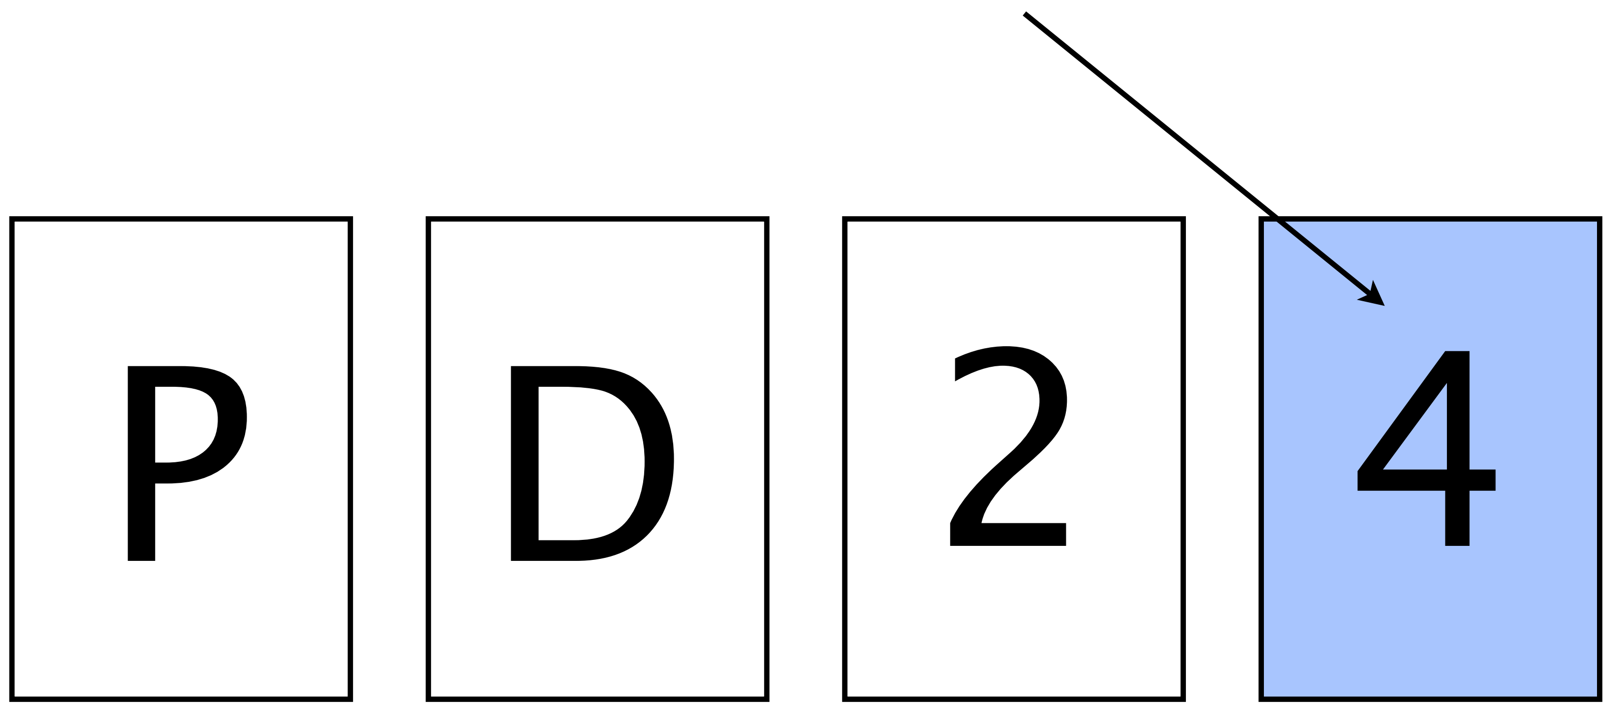

Cards have a letter and a number
Which do you flip to test the hypothesis

"If P then 2"

Two kinds of strategies

Positive test strategy (an aspect of confirmation bias): generate an item that you think will be true and ask about it

If it is true, this is evidence for; if it's not, it eliminates the hypothesis

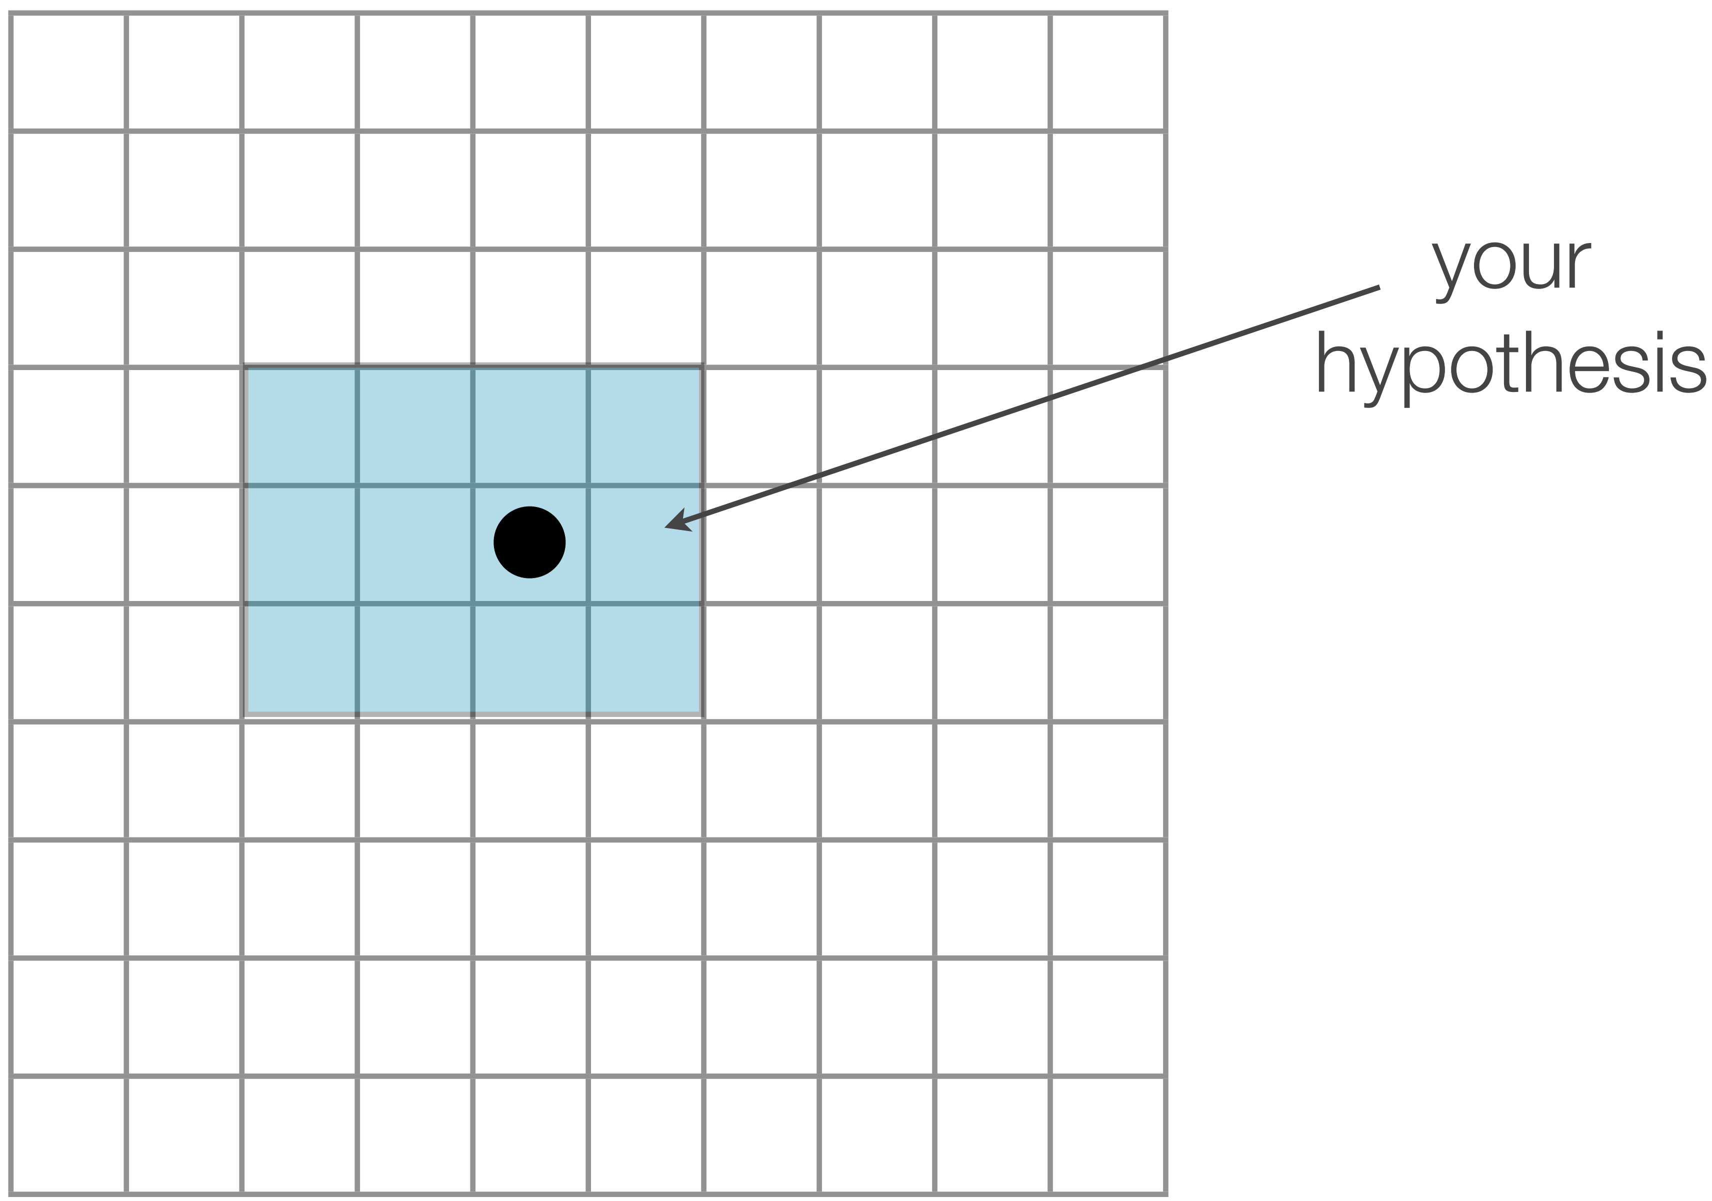

Two kinds of strategies

Negative test strategy: generate an item that you think will *not* be true and ask about it

If it is false, this is evidence for; if it true, it eliminates the hypothesis

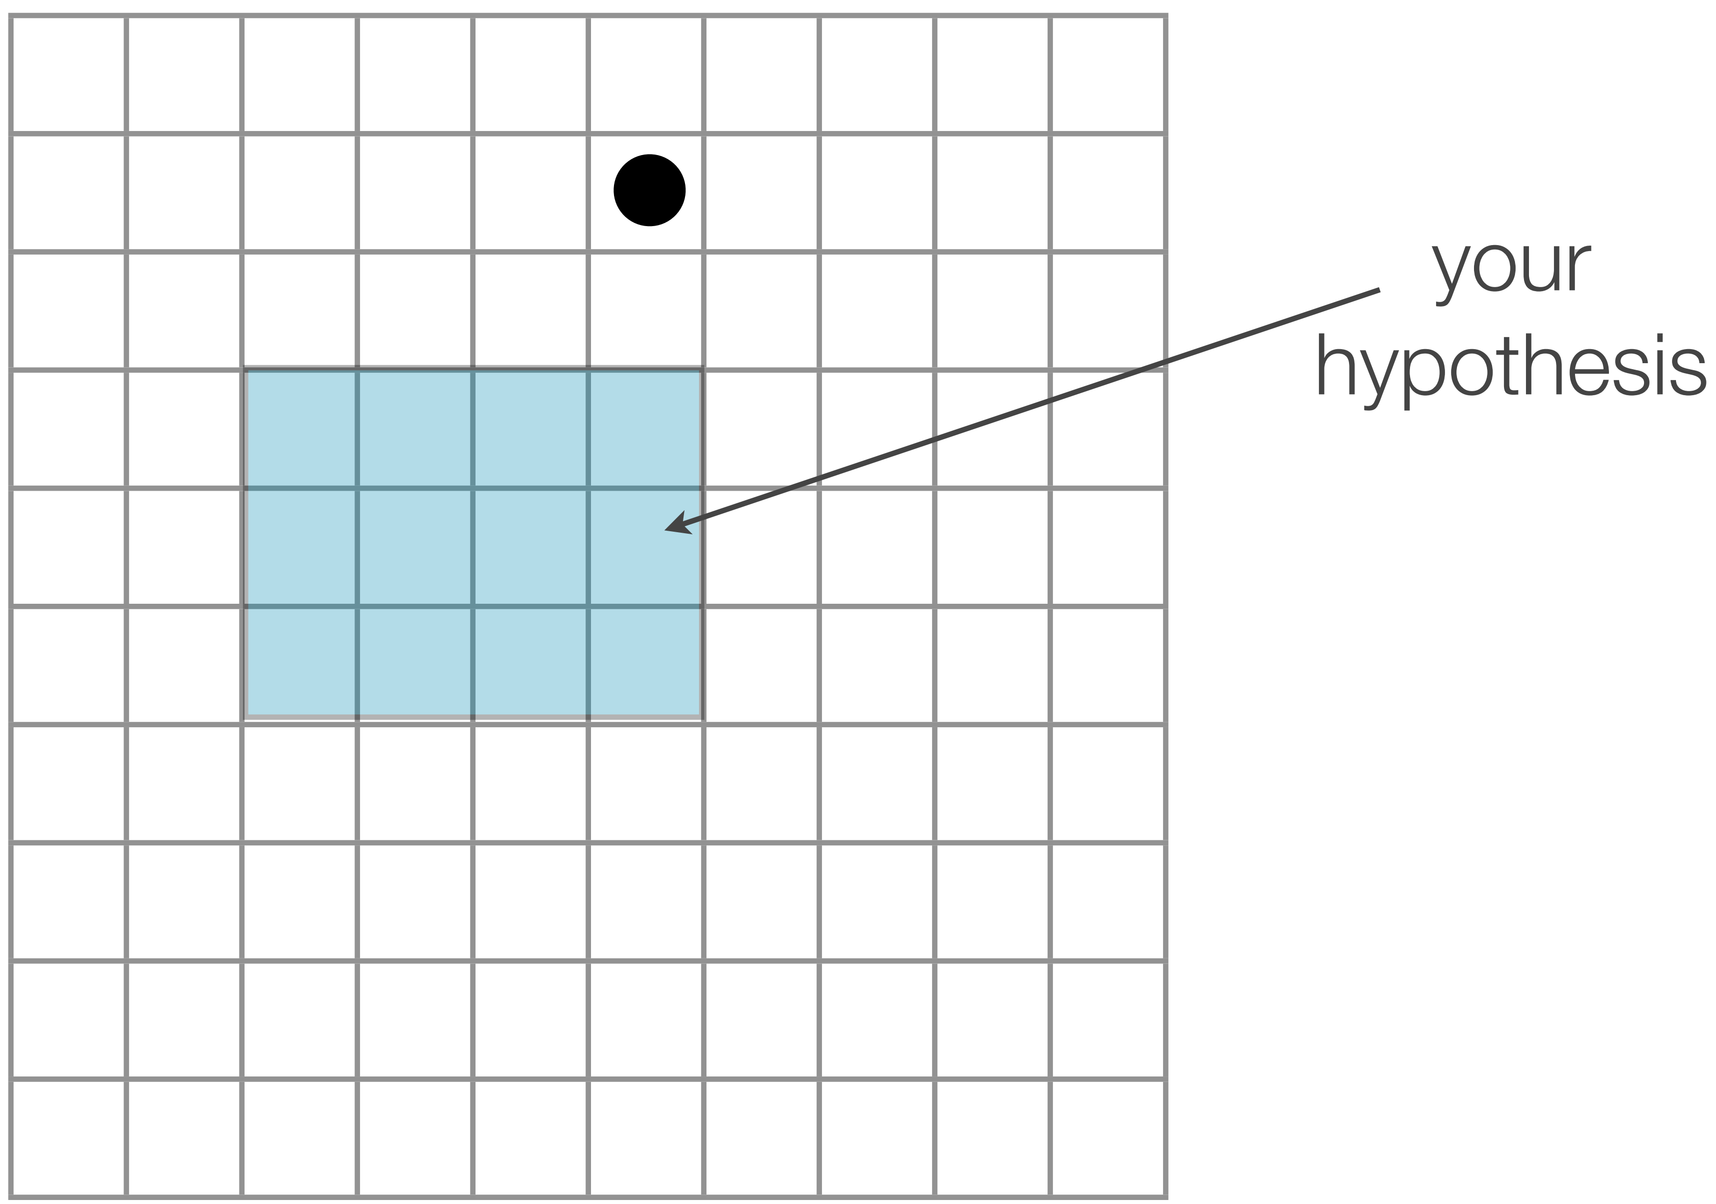

This can apply quite widely

Number rules

INFORMATION SO FAR :
rule true for {9, 27, 45, 81}
rule false for {3, 2, 51}

HYPOTHESIS: rule is “multiples of 9”

POSITIVE TEST: ask if {18} is true.
NEGATIVE TEST: ask if {21} is true

Scientific Laws

INFORMATION SO FAR :
rule true for {experiment1, experiment2}
rule false for {experiment3}

HYPOTHESIS: rule is theory of relativity

POSITIVE TEST: one kind of experiment
NEGATIVE TEST: another kind of exp’t

Language

INFORMATION SO FAR :
rule true for {“run!”, “he ate”}
rule false for {“blaarg”}

HYPOTHESIS: some grammar **G**

POSITIVE TEST: try a sentence from **G**
NEGATIVE TEST: try sentence not in **G**

Categories

INFORMATION SO FAR :
rule true for {cat, dog, parrot}
rule false for {tiger}

HYPOTHESIS: rule is “pets”

POSITIVE TEST: ask if {goldfish} is true.
NEGATIVE TEST: ask if {mosquito} is true

Why do people use positive tests?

Are they stupid?

Not necessarily stupid!

There are actually a number of situations in which a positive test strategy increases the expected information gain of your data

Plan

- Lecture 1 [short]: Another kind of sampling: pedagogical
  - Model for pedagogical sampling
  - Sensitivity to pedagogical data
  - Double-edged sword of pedagogy
- ➡ Lecture 2 [long]: Sensitivity to the informativeness of the data
  - Confirmation bias and the positive test strategy (PTS)
  - ➡ When is positive evidence most useful?
  - When do people use a PTS?
  - Are people sensitive to evidence utility in general?

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

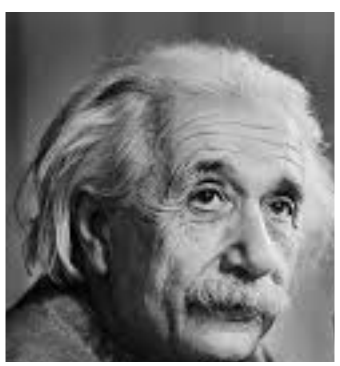



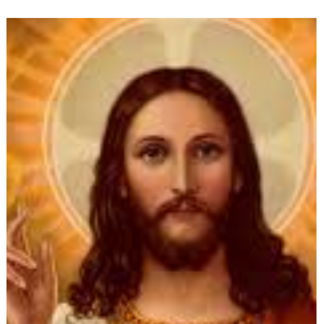

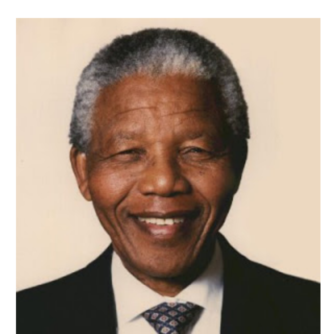
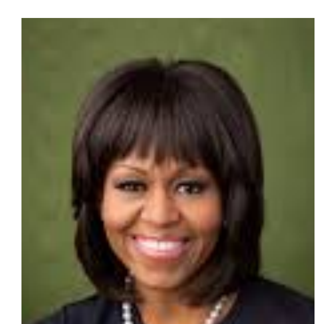

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

do they have facial hair?

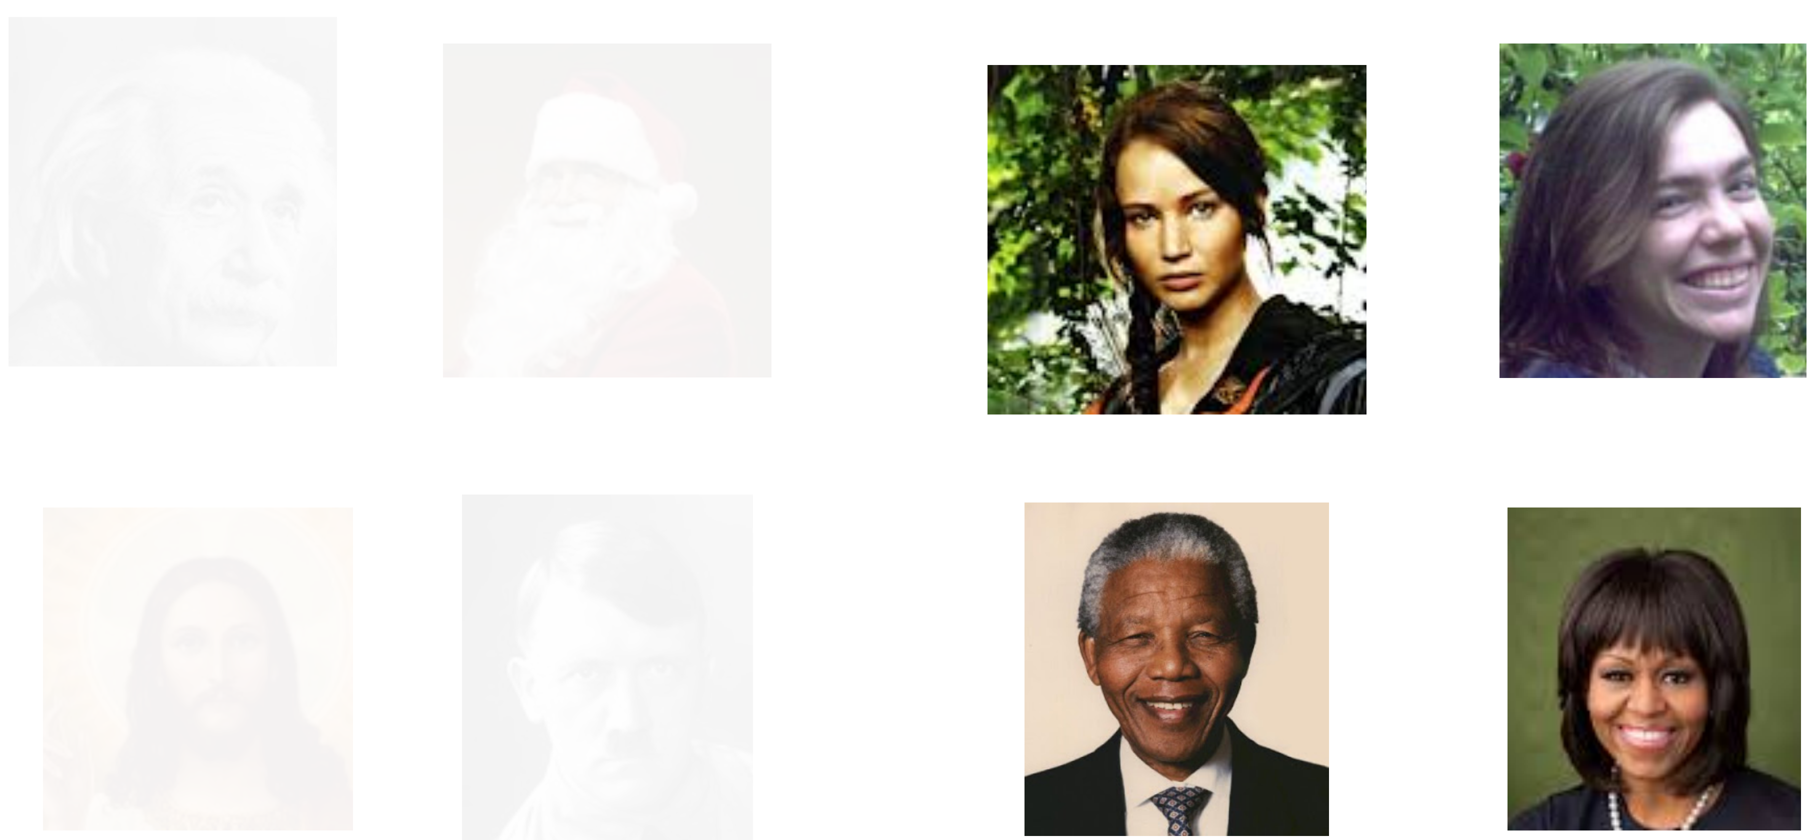

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

are they black?

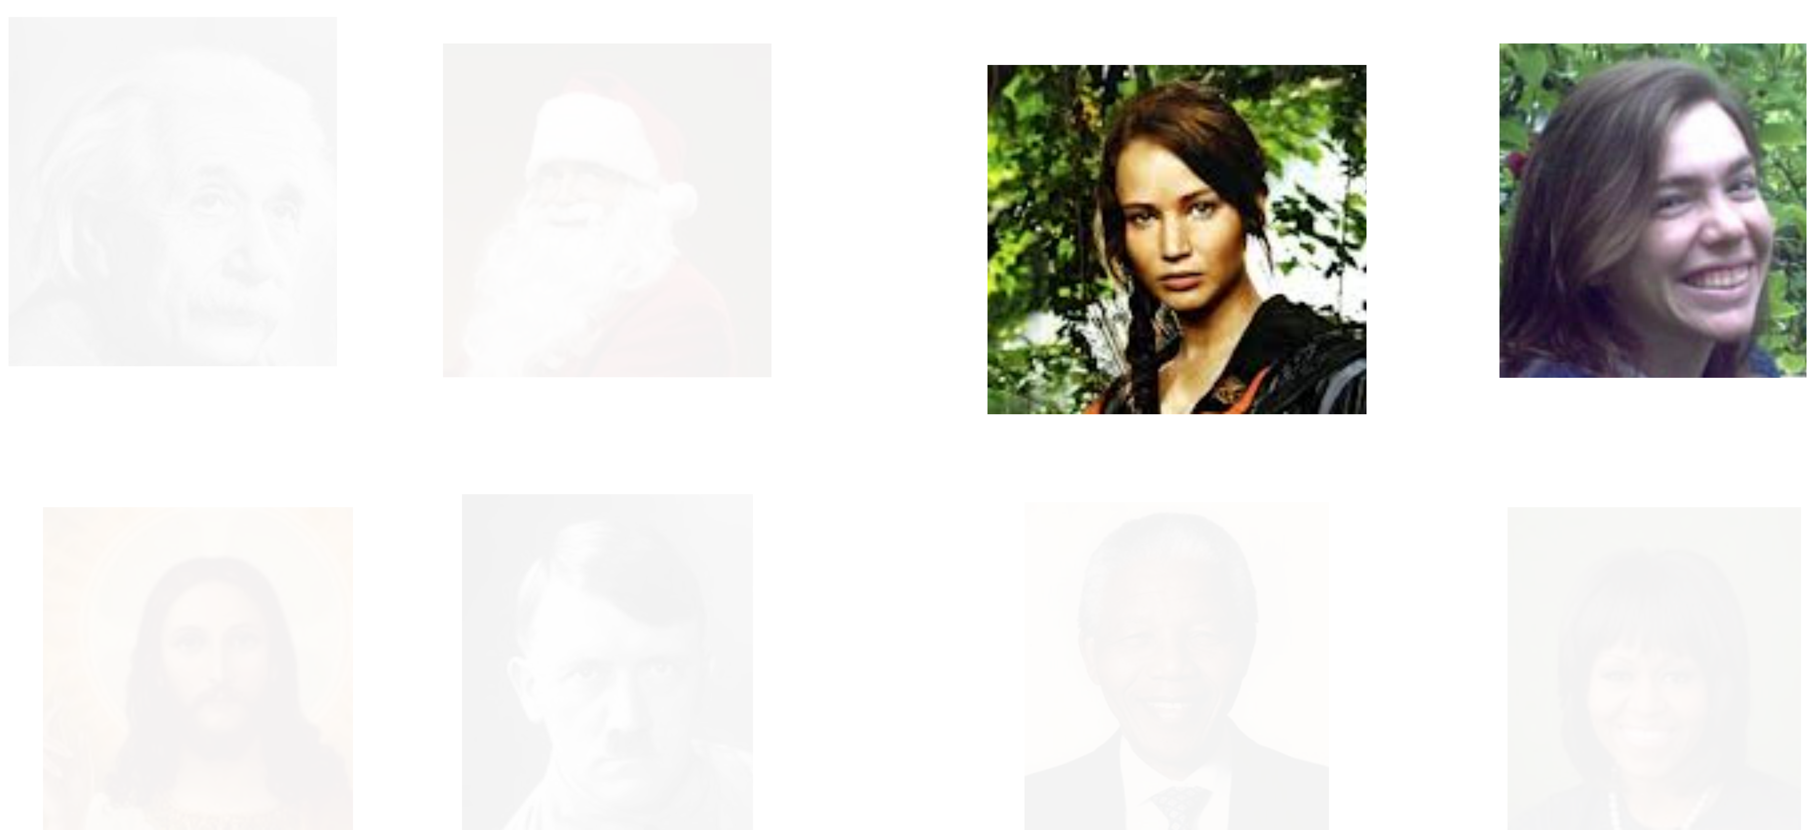

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

are they super awesome?

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

are they super awesome?

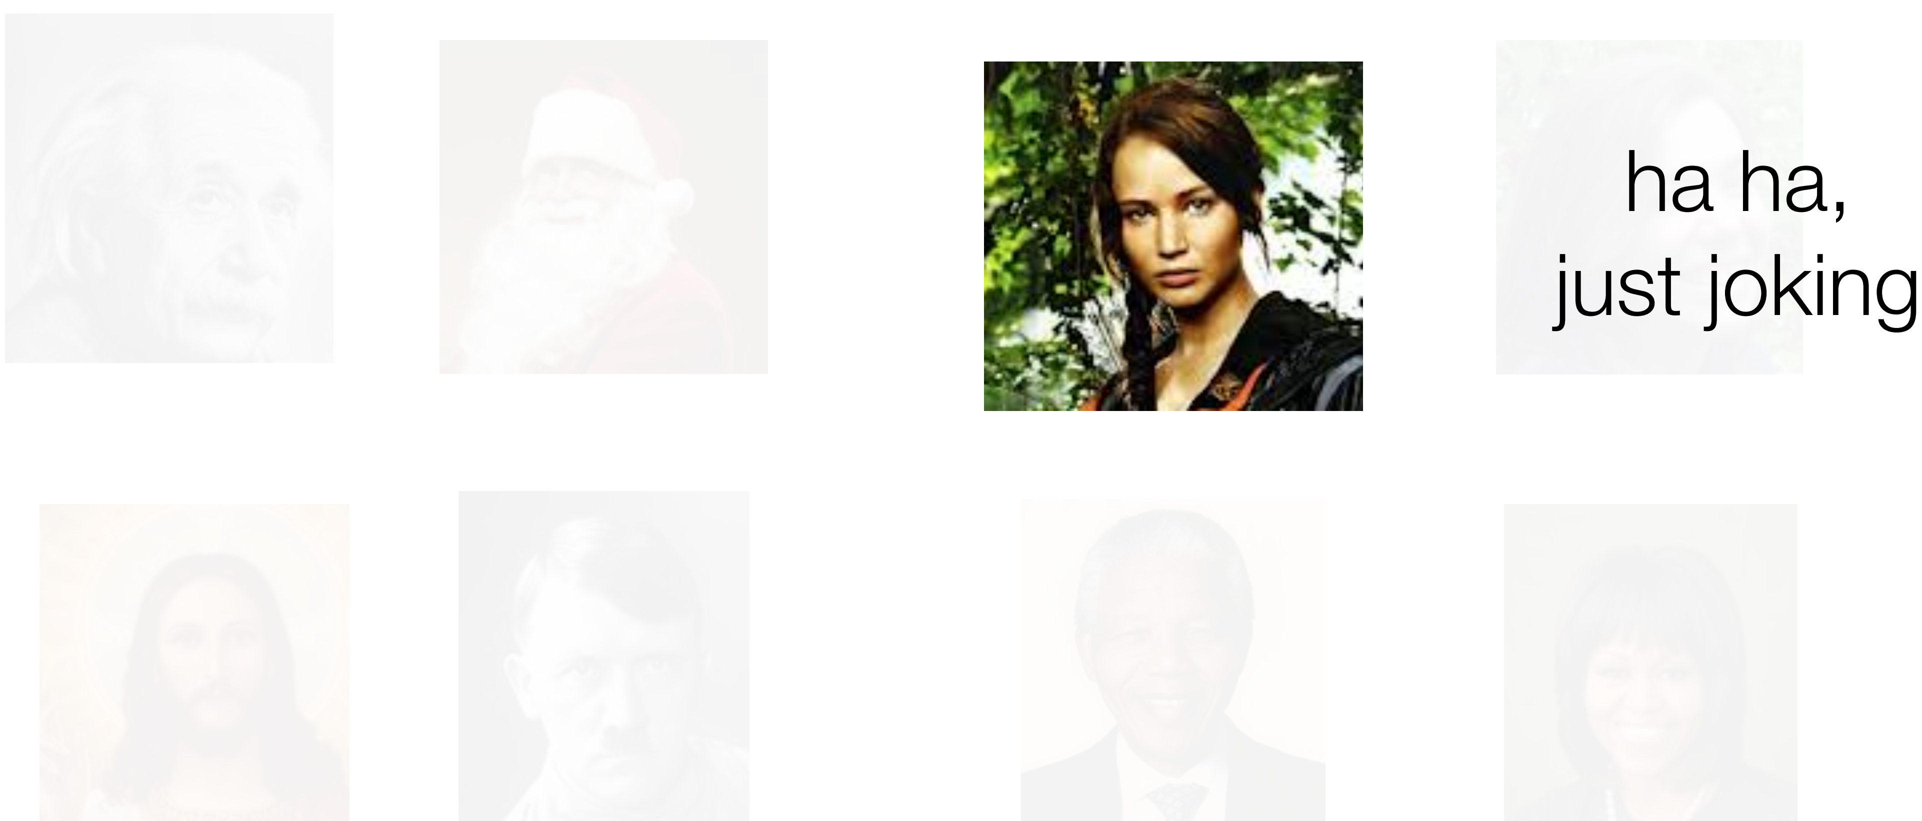

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

bisecting each time got us to the correct answer in three queries

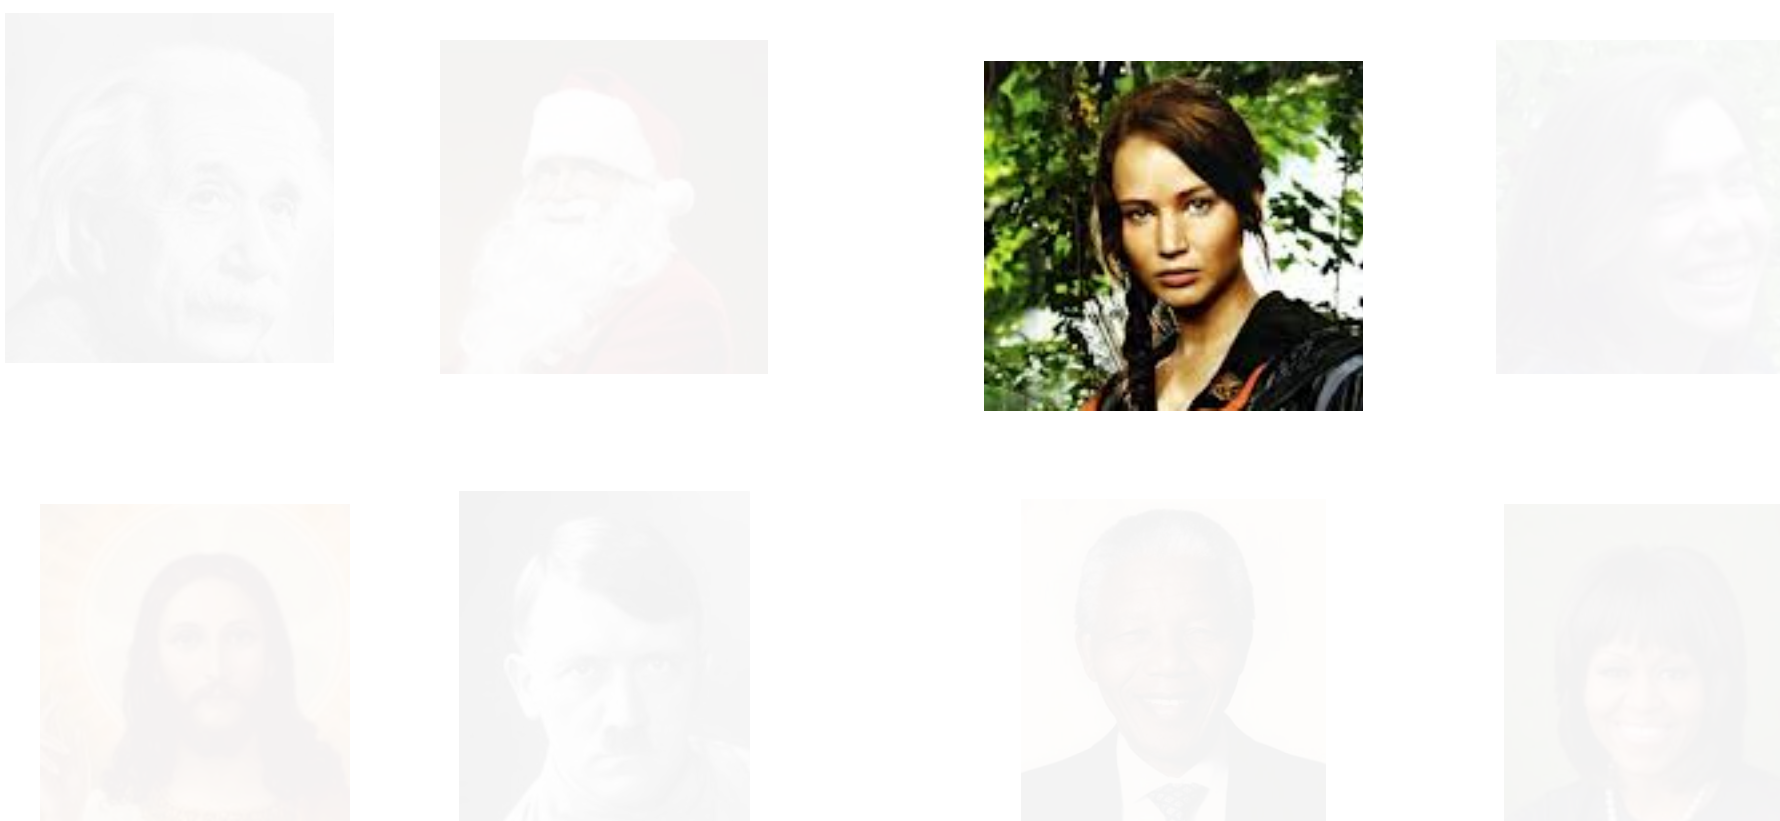

Optimal strategy

- When you don't know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

do they have a beard?

might be lucky...

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

do they have a beard?

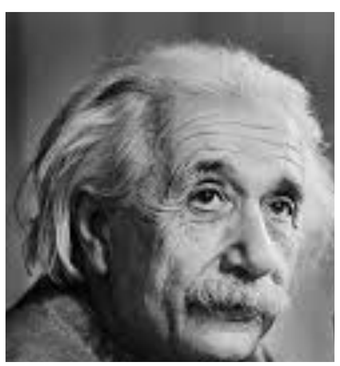

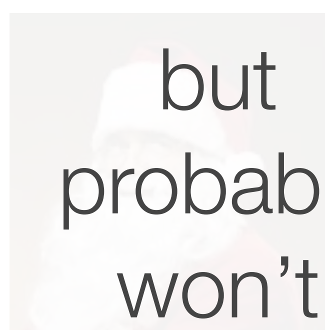

but probably won’t

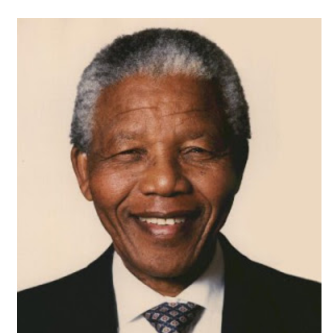

Optimal strategy

- When you don’t know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

do they teach CCS?

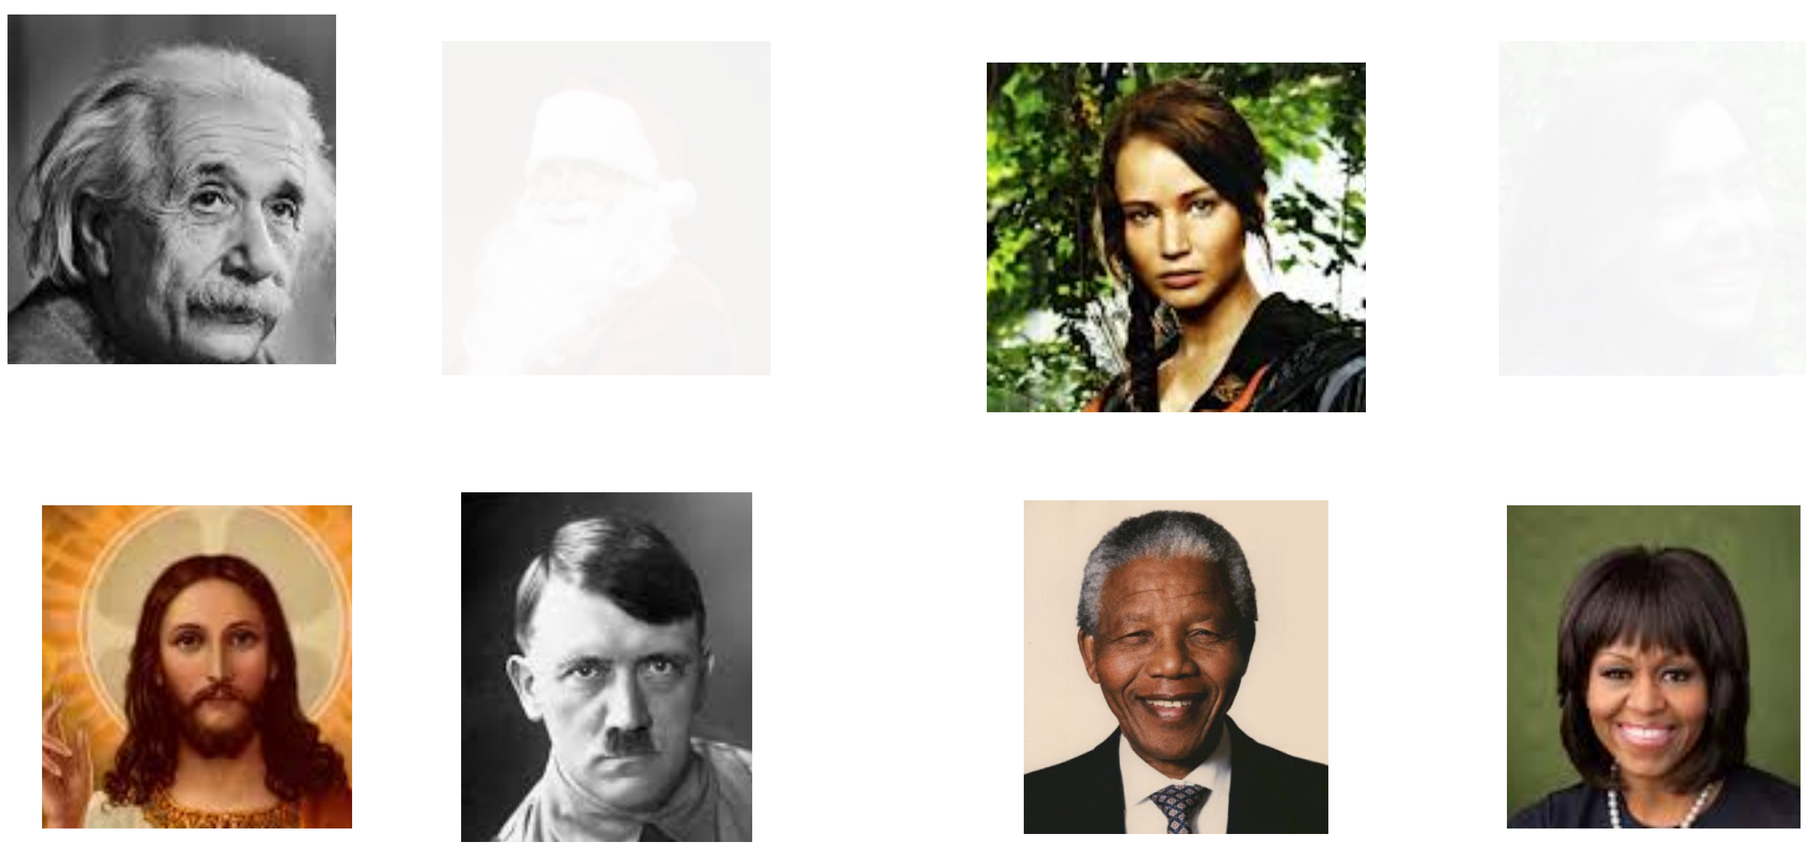

Optimal strategy

- When you don't know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

do they use a bow and arrow?

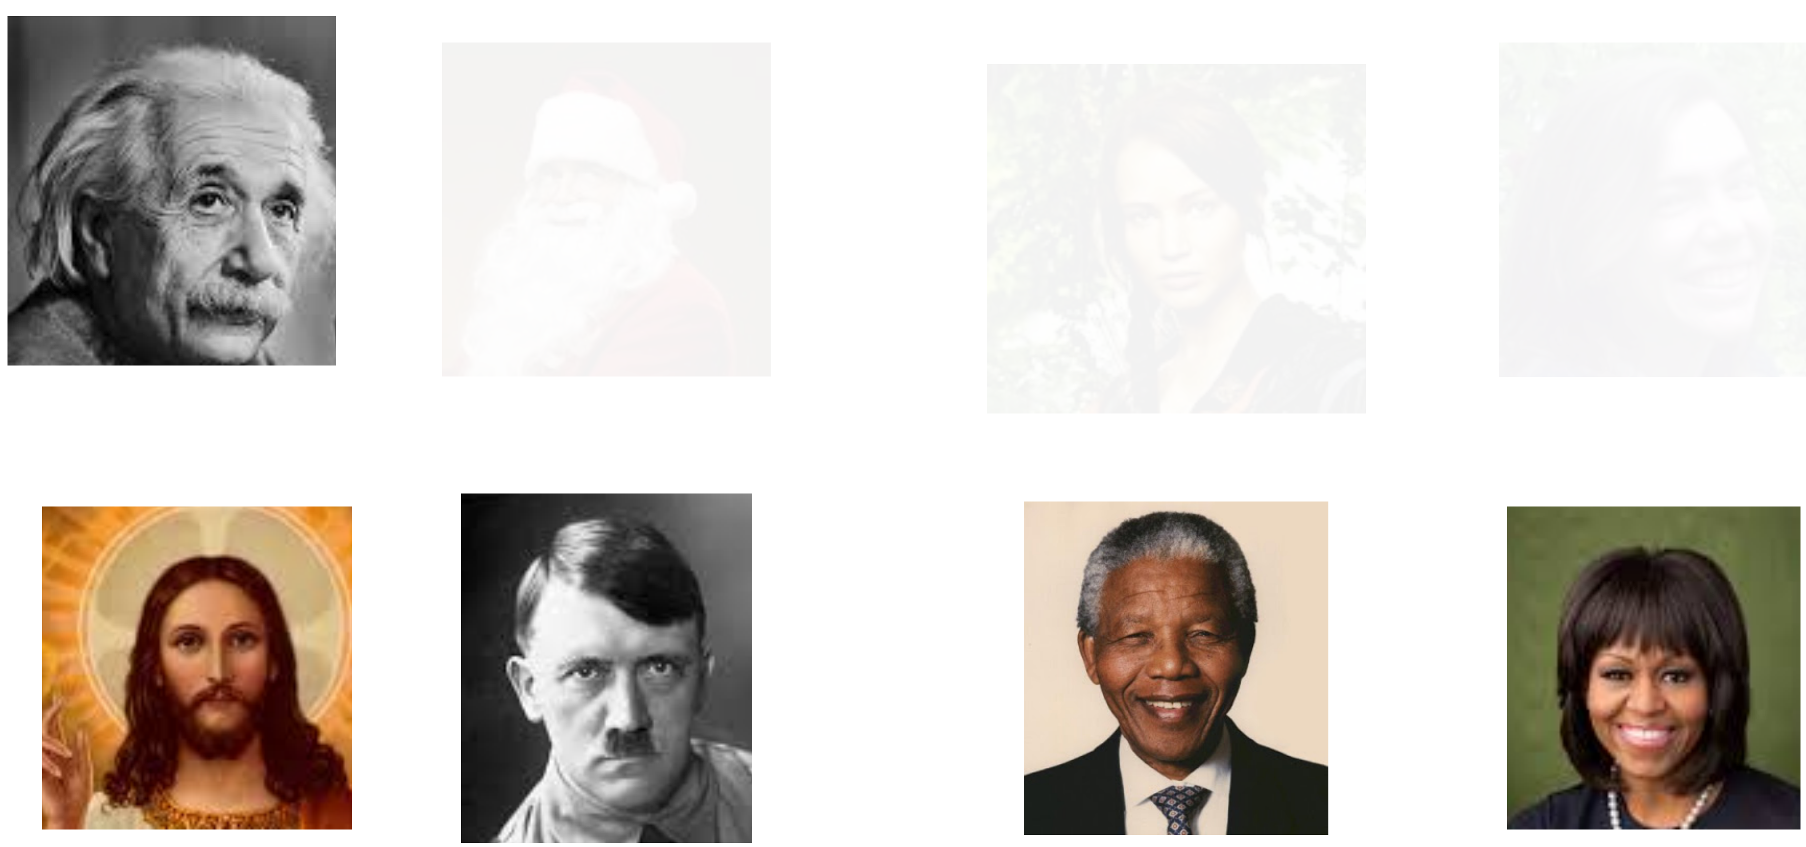

Optimal strategy

- When you don't know what the correct answer is, you gain the most information by asking queries that eliminate 50% of the possibilities

Twenty questions

... and so forth

Optimal strategy

When trying to determine which rule is correct out of an entire hypothesis space, the bisection strategy is theoretically optimal

To learn the true rule fastest, one seeks to minimize the expected entropy of the distribution over the m possible rules.

Straightforward to show that this depends on the number $m(x)$ of not-yet-falsified rules that are consistent with the query x:

$$E[H(\mathcal{O} \mid x)] = \frac{m(\neg x) \ln m(\neg x) + m(x) \ln m(x)}{m}$$

minimized when $m(x) = m(\neg x) = m/2$

Optimal strategy

When trying to determine which rule is correct out of an entire hypothesis space, the bisection strategy is theoretically optimal

To learn the true rule fastest, one seeks to minimize the expected entropy of the distribution over the m possible rules.

Straightforward to show that this depends on the number $m(x)$ of not-yet-falsified rules that are consistent with the query x:

There is math for this:

$$\begin{aligned}
E[H(\mathcal{O}\,|\,x)] &= -E_{\mathcal{O}(x)}\left[\sum_i \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x))\ln \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x))\right] \\
&= -\Pr(\mathcal{O}(x)=0)\left(\sum_i \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x)=0)\ln \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x)=0)\right) \\
&\quad -\Pr(\mathcal{O}(x)=1)\left(\sum_i \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x)=1)\ln \Pr(\mathcal{O}\rightarrow r_i\,|\,\mathcal{O}(x)=1)\right) \\
&= -\frac{m(\neg x)}{m}\left(\sum_{i|r_i(x)=0}\frac{1}{m(\neg x)}\ln\frac{1}{m(\neg x)}\right)-\frac{m(x)}{m}\left(\sum_{i|r_i(x)=1}\frac{1}{m(x)}\ln\frac{1}{m(x)}\right) \\
&= \frac{m(\neg x)\ln m(\neg x)+m(x)\ln m(x)}{m}
\end{aligned}$$

Optimal strategy

When trying to determine which rule is correct out of an entire hypothesis space, the bisection strategy is theoretically optimal

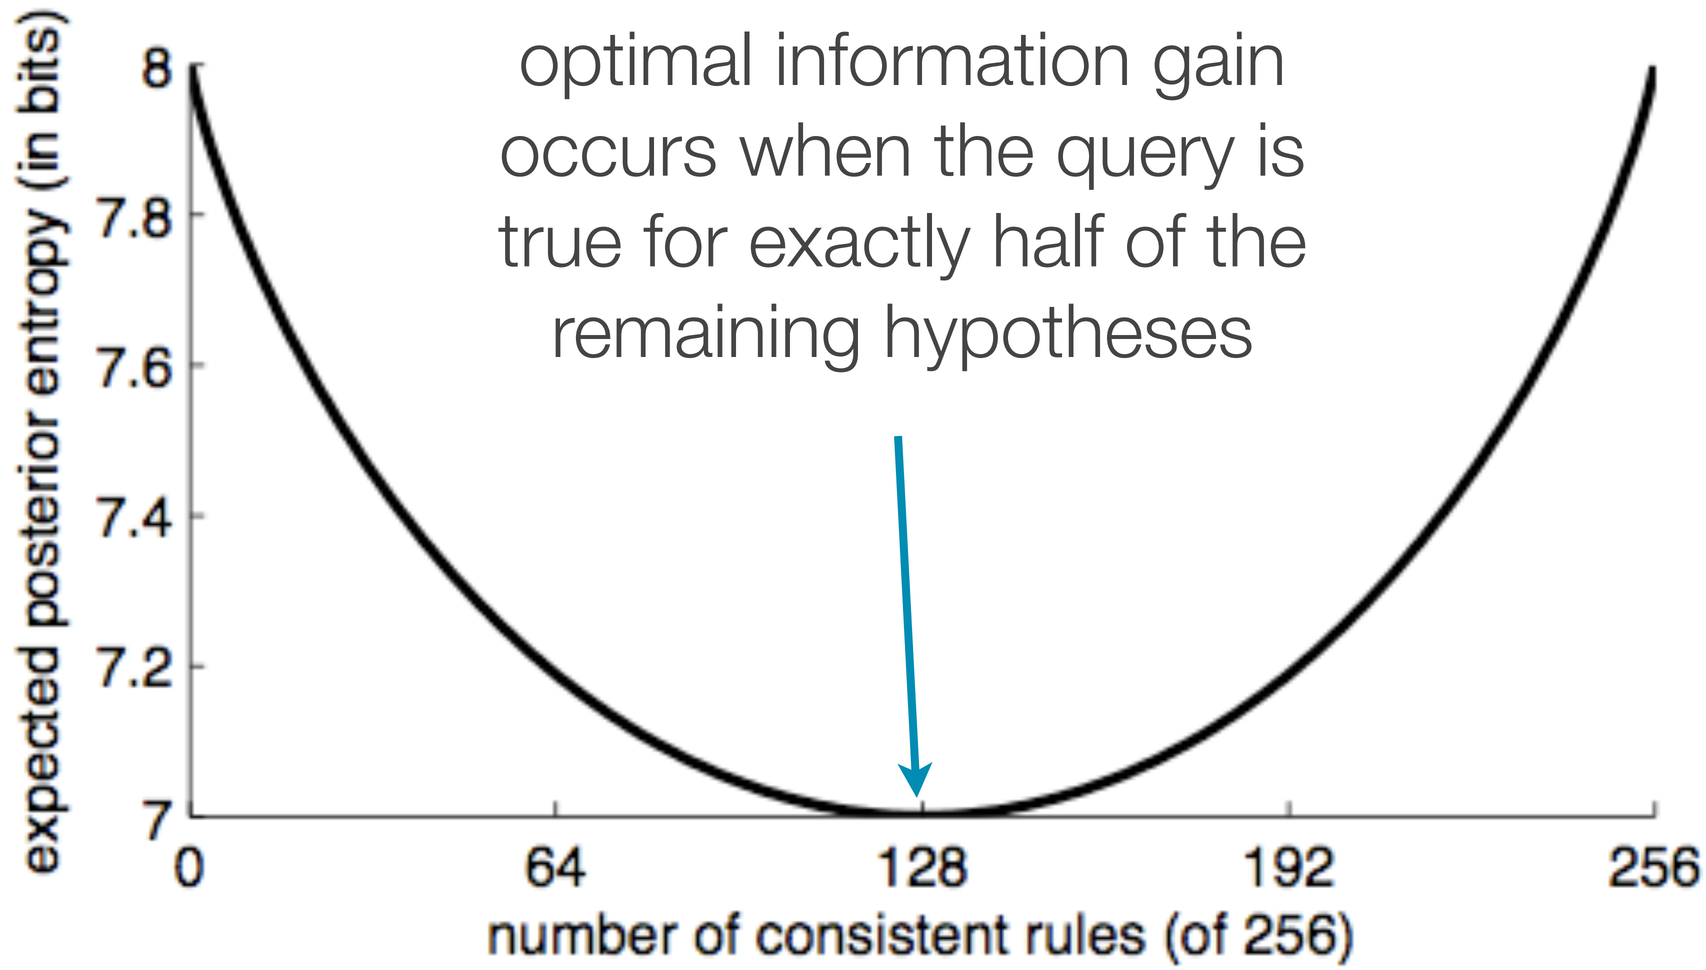

What if you can't enumerate your hypotheses?

Indeed, most people are in a situation in which:

(a) they must find the correct hypothesis out of an entire space of possibilities

(b) they are unable to evaluate all hypotheses at once

What is the best strategy in this case?

The answer depends on the notion of **sparsity**

- Sparse rules capture a minority of entities:
  - Few animals are PETS
  - Few numbers are DIVISIBLE BY 10
- Nonsparse rules are the opposite:
  - Most animals are MOTILE
  - Most numbers are COMPOSITE
- The sparsity assumption:
  - “Good rules tend to be sparse”

The answer depends on the notion of **sparsity**

When rules are sparse, “false” is the most common response to a random query

- PTS: “false” is strongly informative
- NTS: “true” is strongly informative

PTS makes the <u>typical</u> case informative, but NTS makes the <u>atypical</u> case informative

The answer depends on the notion of sparsity

Intuitively: PTS helps to overcome the bias in the world for "no" responses

helps, on average, to eliminate more of the hypotheses

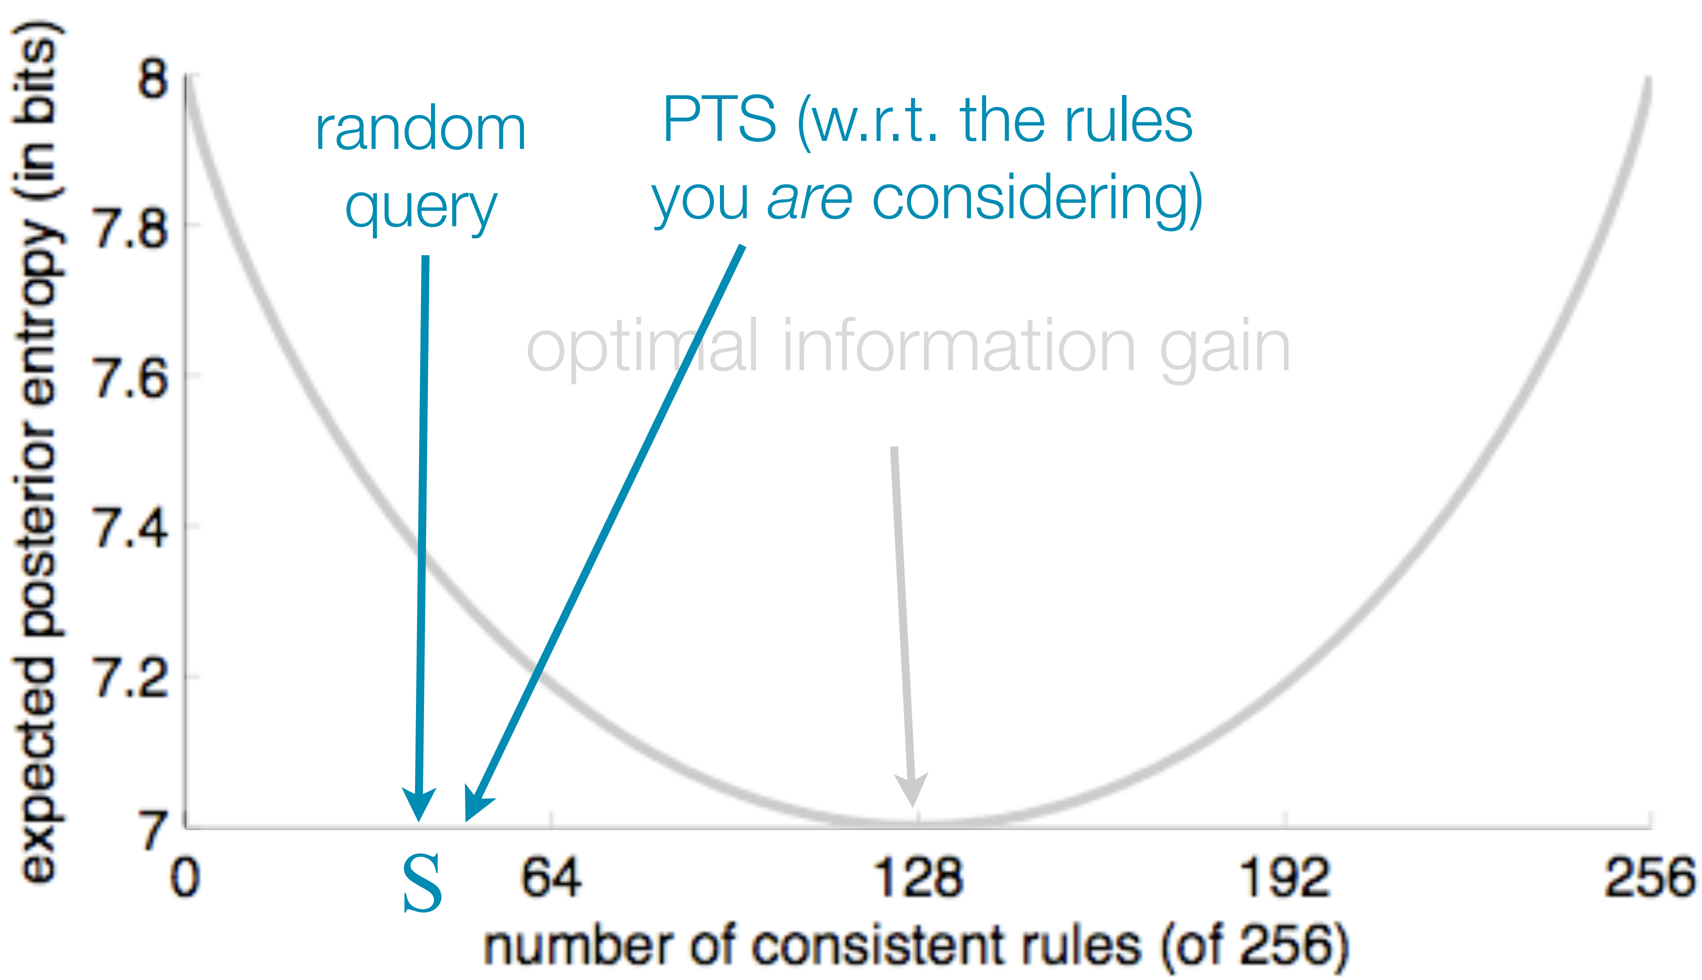

Simulations support this intuition

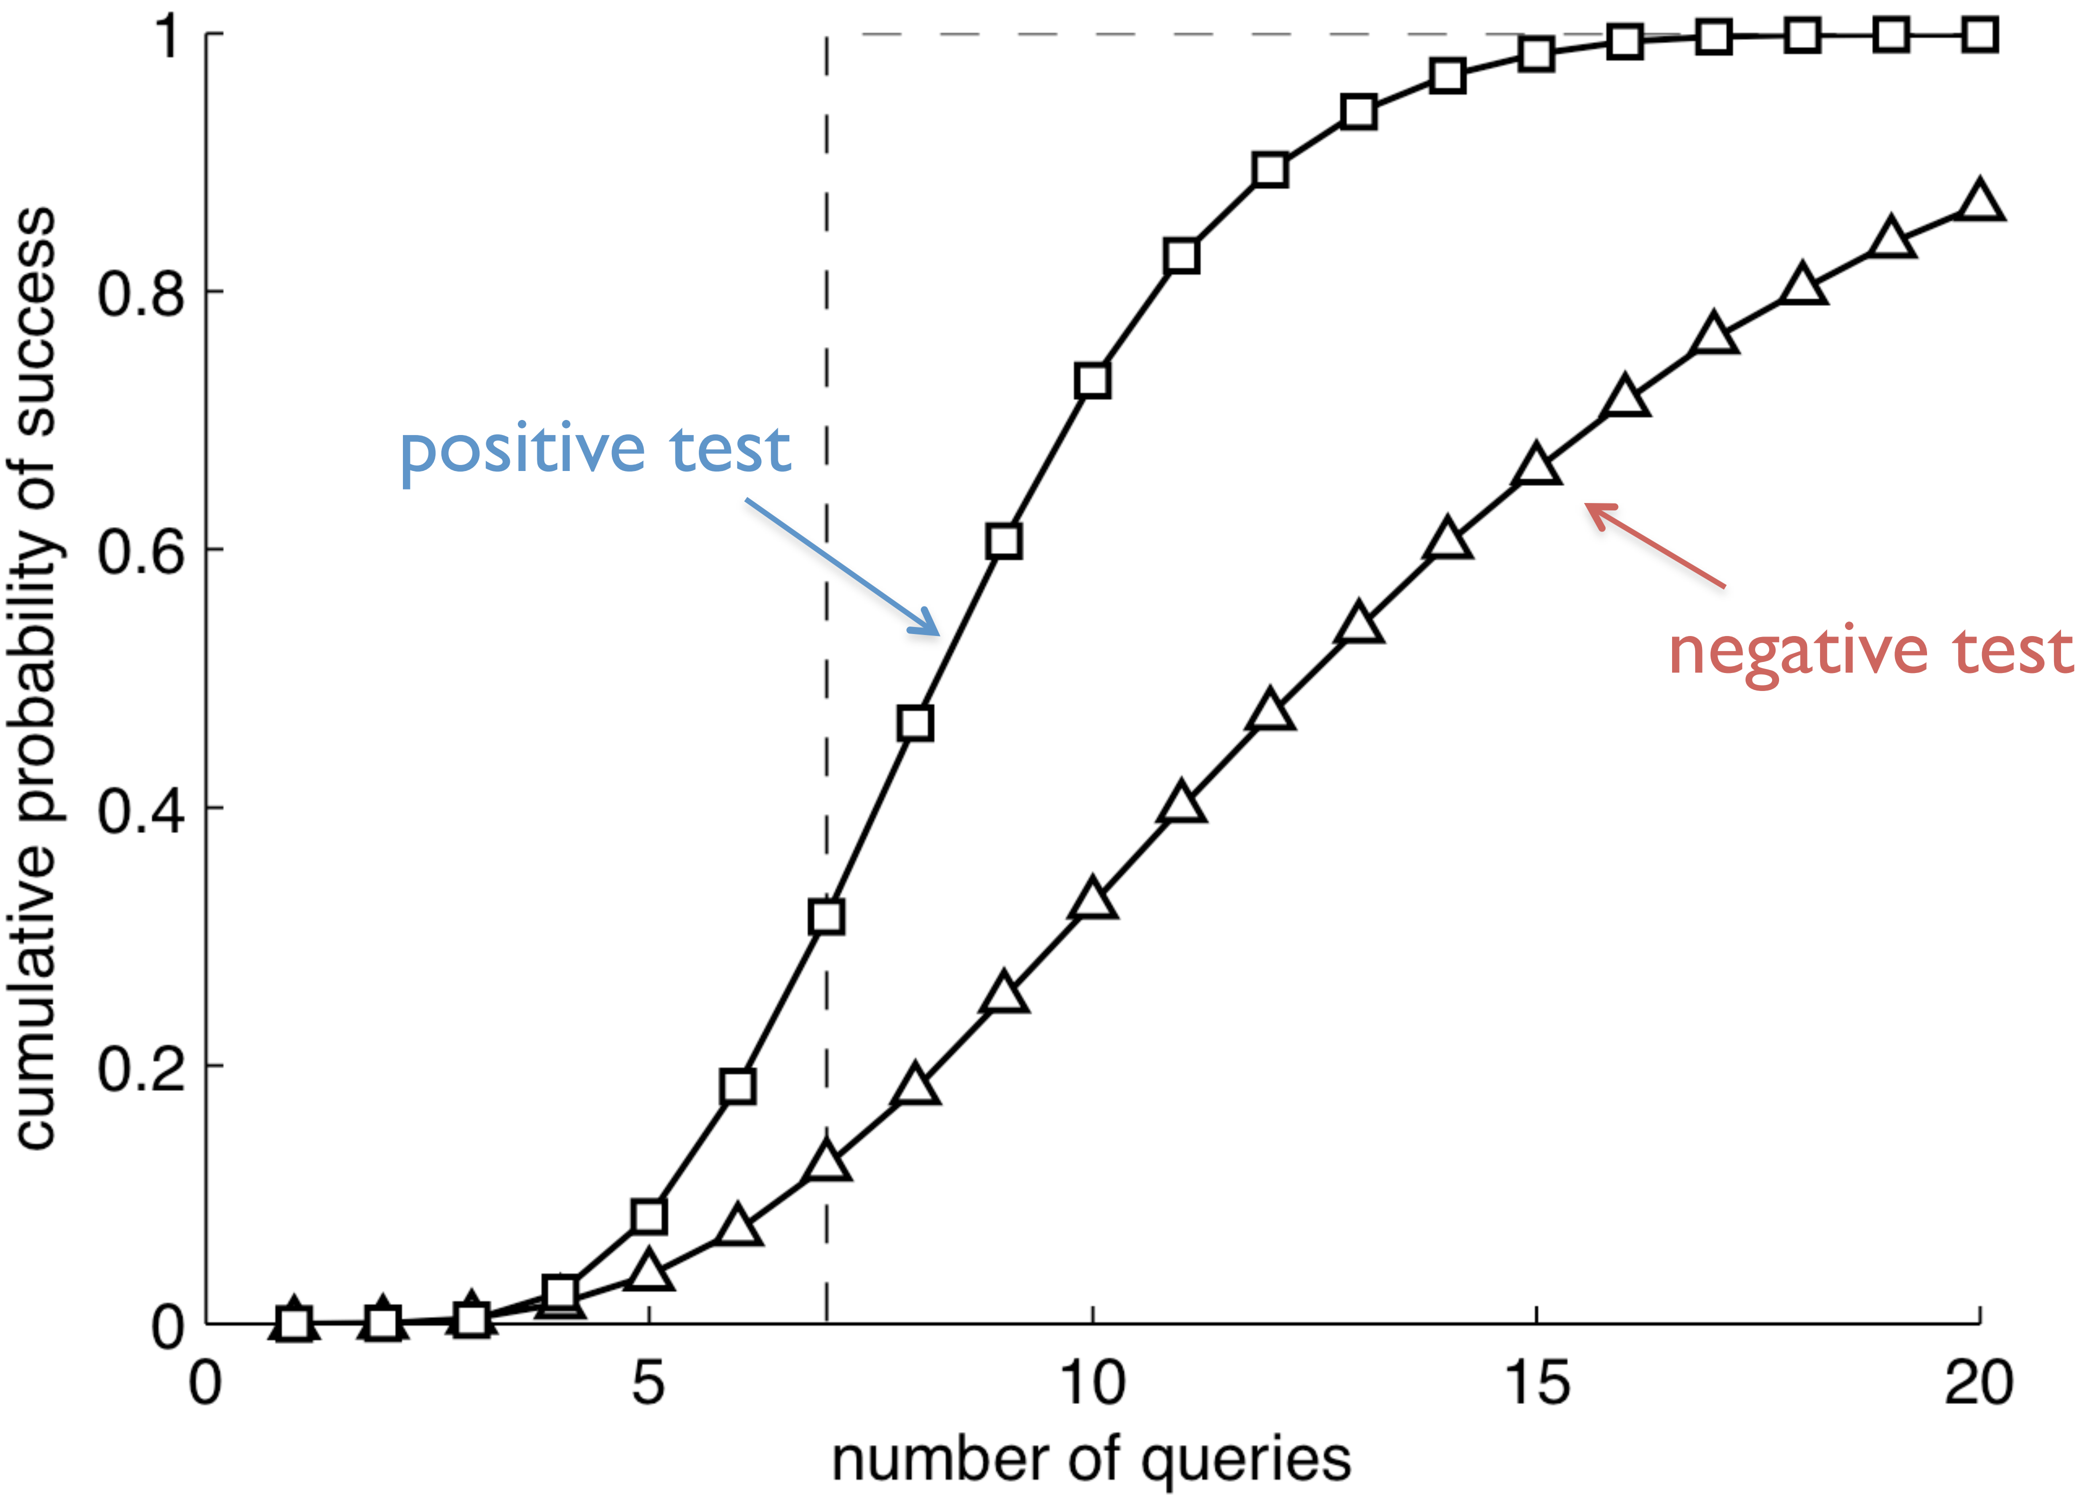

So PTS makes sense when hypotheses are sparse!

- Positive test strategy is more effective at identifying the correct rule out of an entire space of hypotheses when
  - only a few hypotheses can be considered at once
  - hypotheses in general tend to be sparse

- This is because, when hypotheses are sparse, the world has a bias to answering “no”
  - This rules out relatively few hypotheses (and does worse than the optimal bisection strategy)
  - The best way to counteract that bias is to ask questions that would yield “yes” for at least a few hypotheses (i.e., the ones you are considering)

So PTS makes sense when hypotheses are sparse!

Are people sensitive to this?

In other words, do they use a PTS when the hypotheses are sparse, but stop using one if they aren't?

Test: Game of battleships

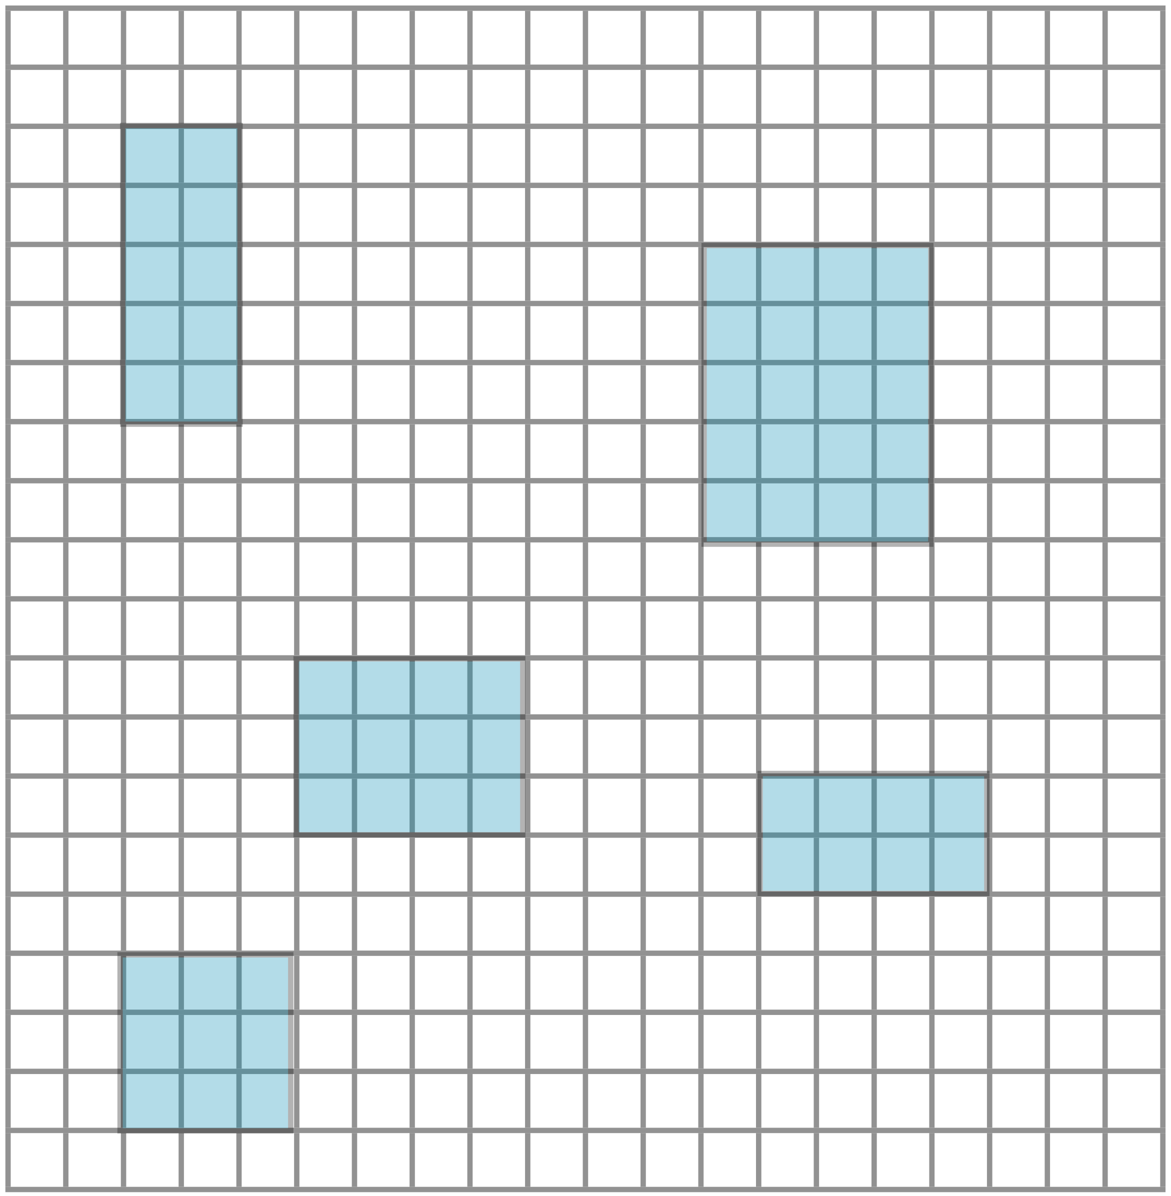

Where are the hidden ships?

Test: Game of battleships

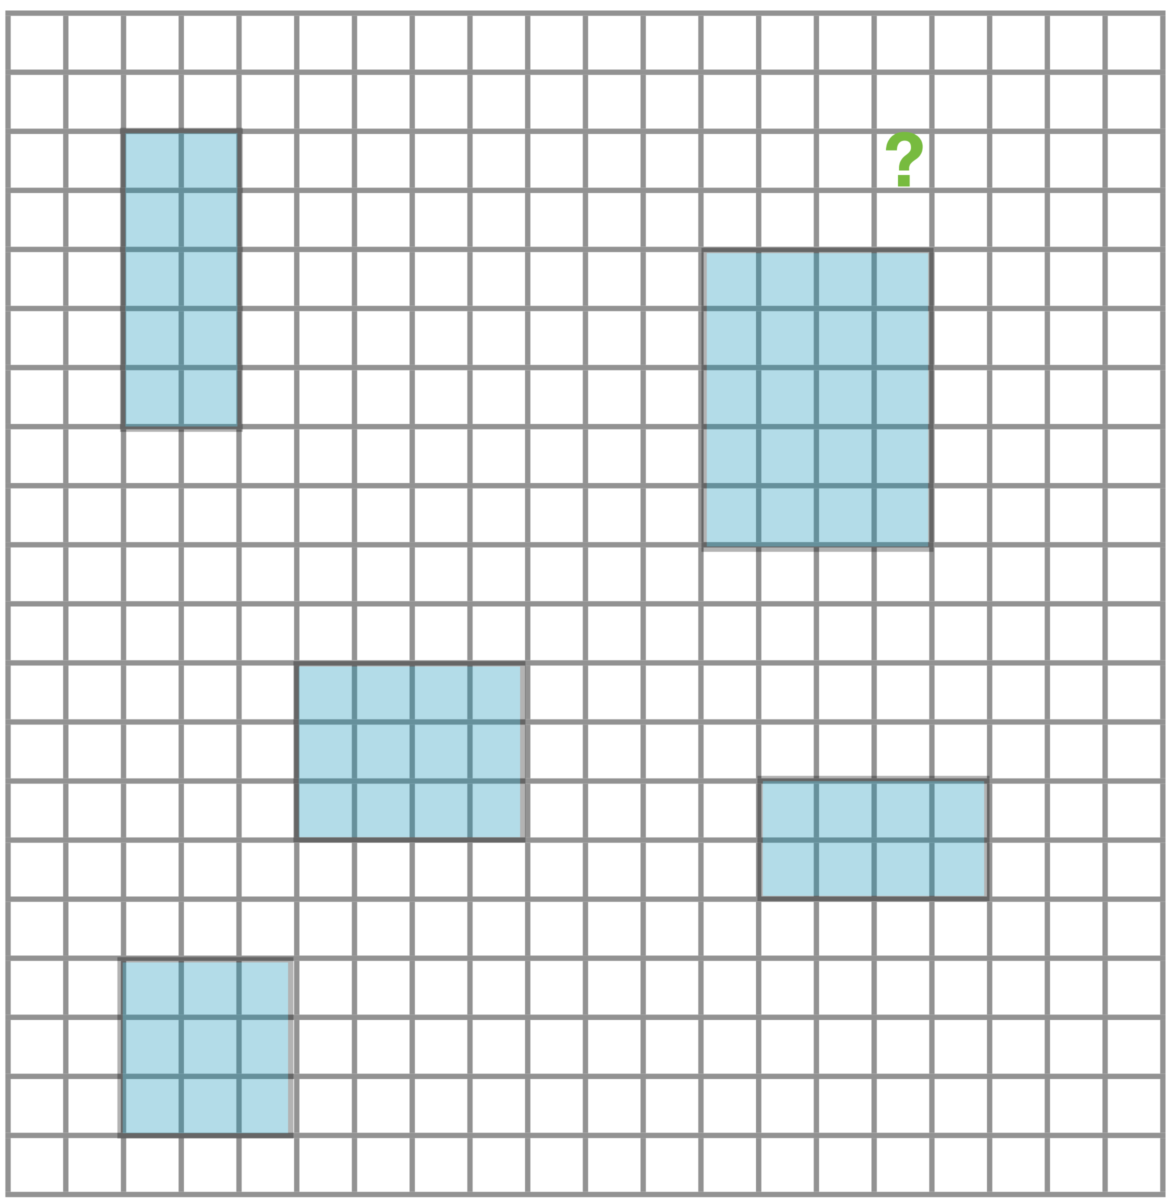

Normally, in the game, you ask about specific points

Test: Game of battleships

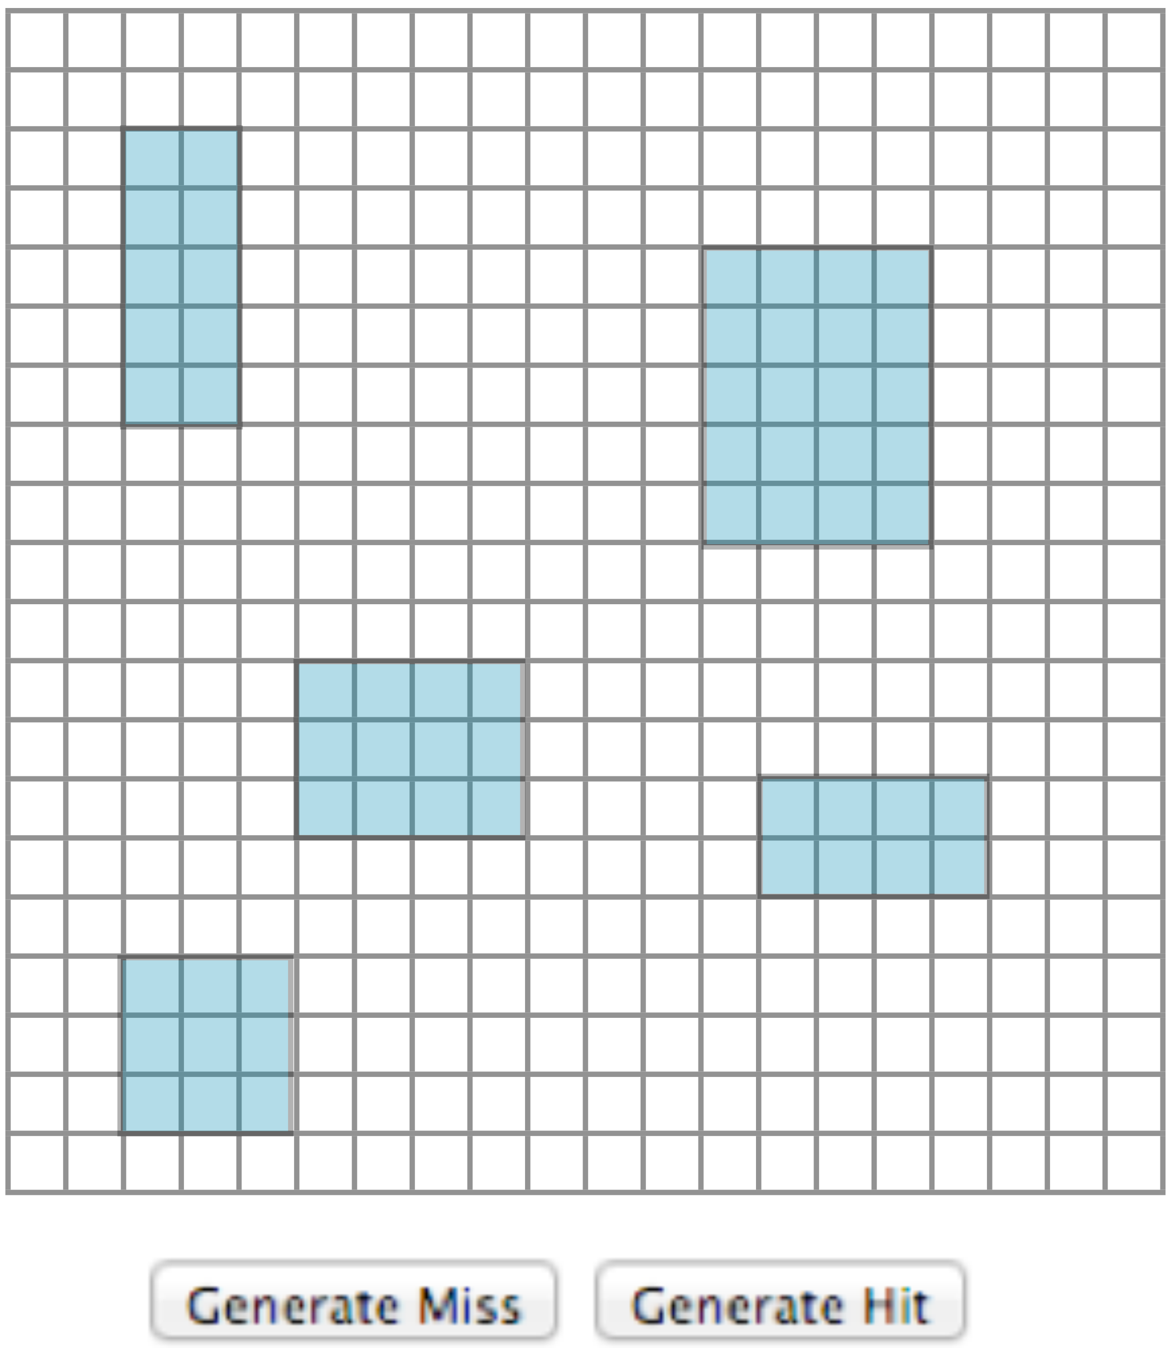

We can directly measure people's reliance on positive or negative tests by making them ask for those...

Test: Game of battleships

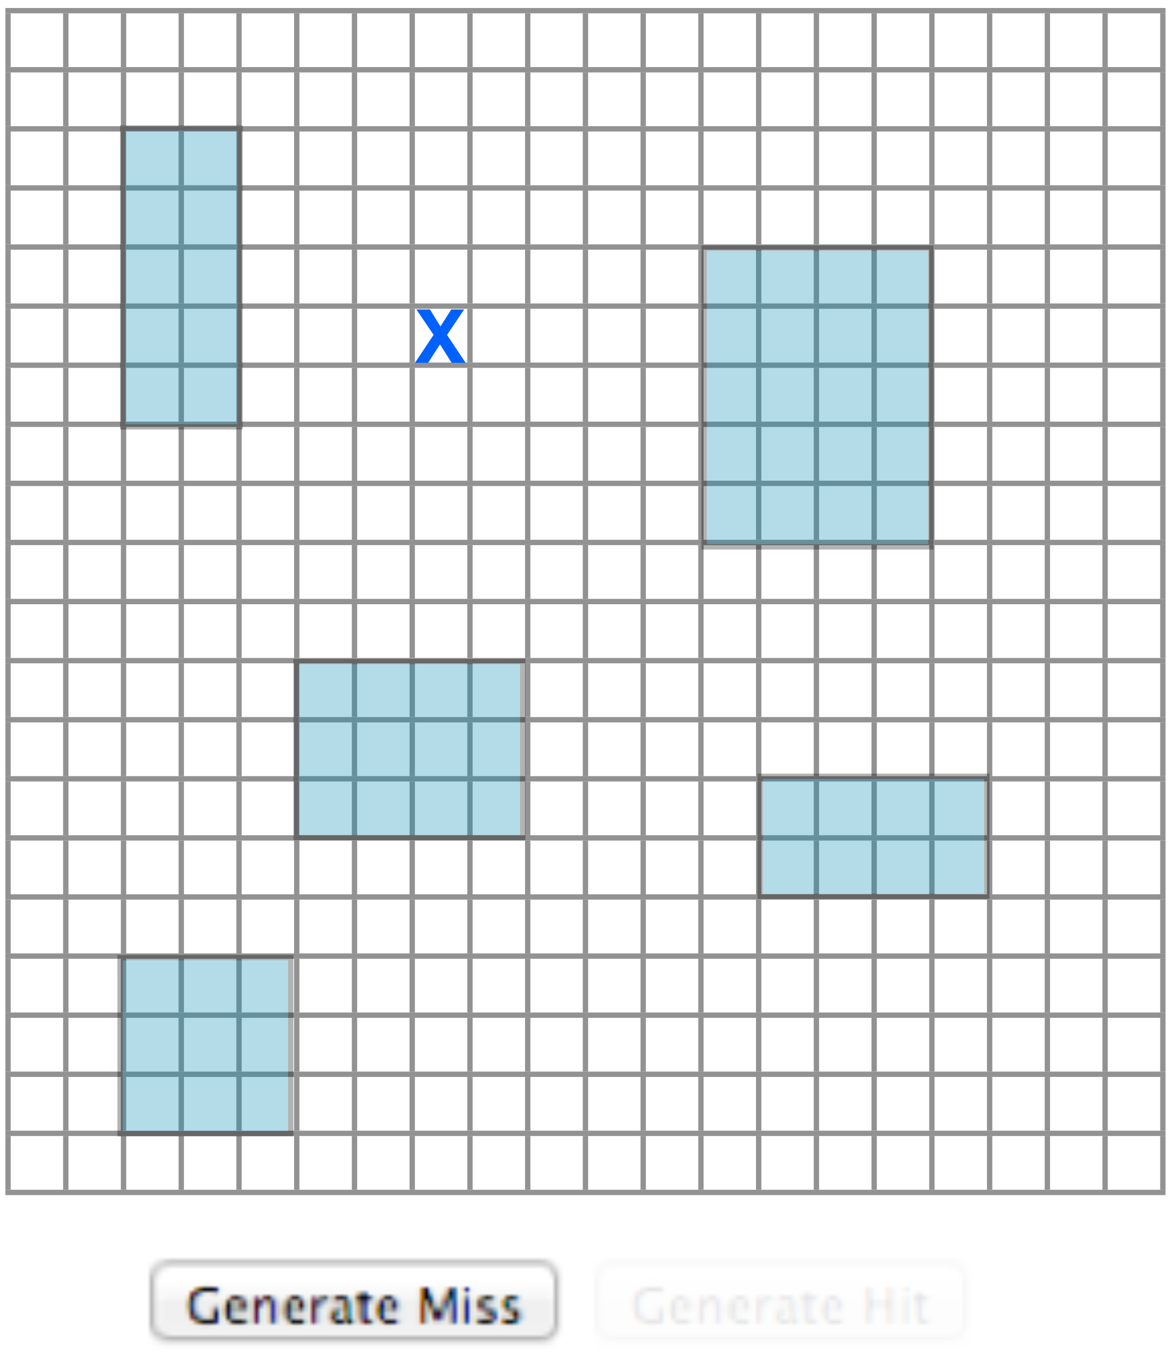

We can directly measure people's reliance on positive or negative tests by making them ask for those...

Test: Game of battleships

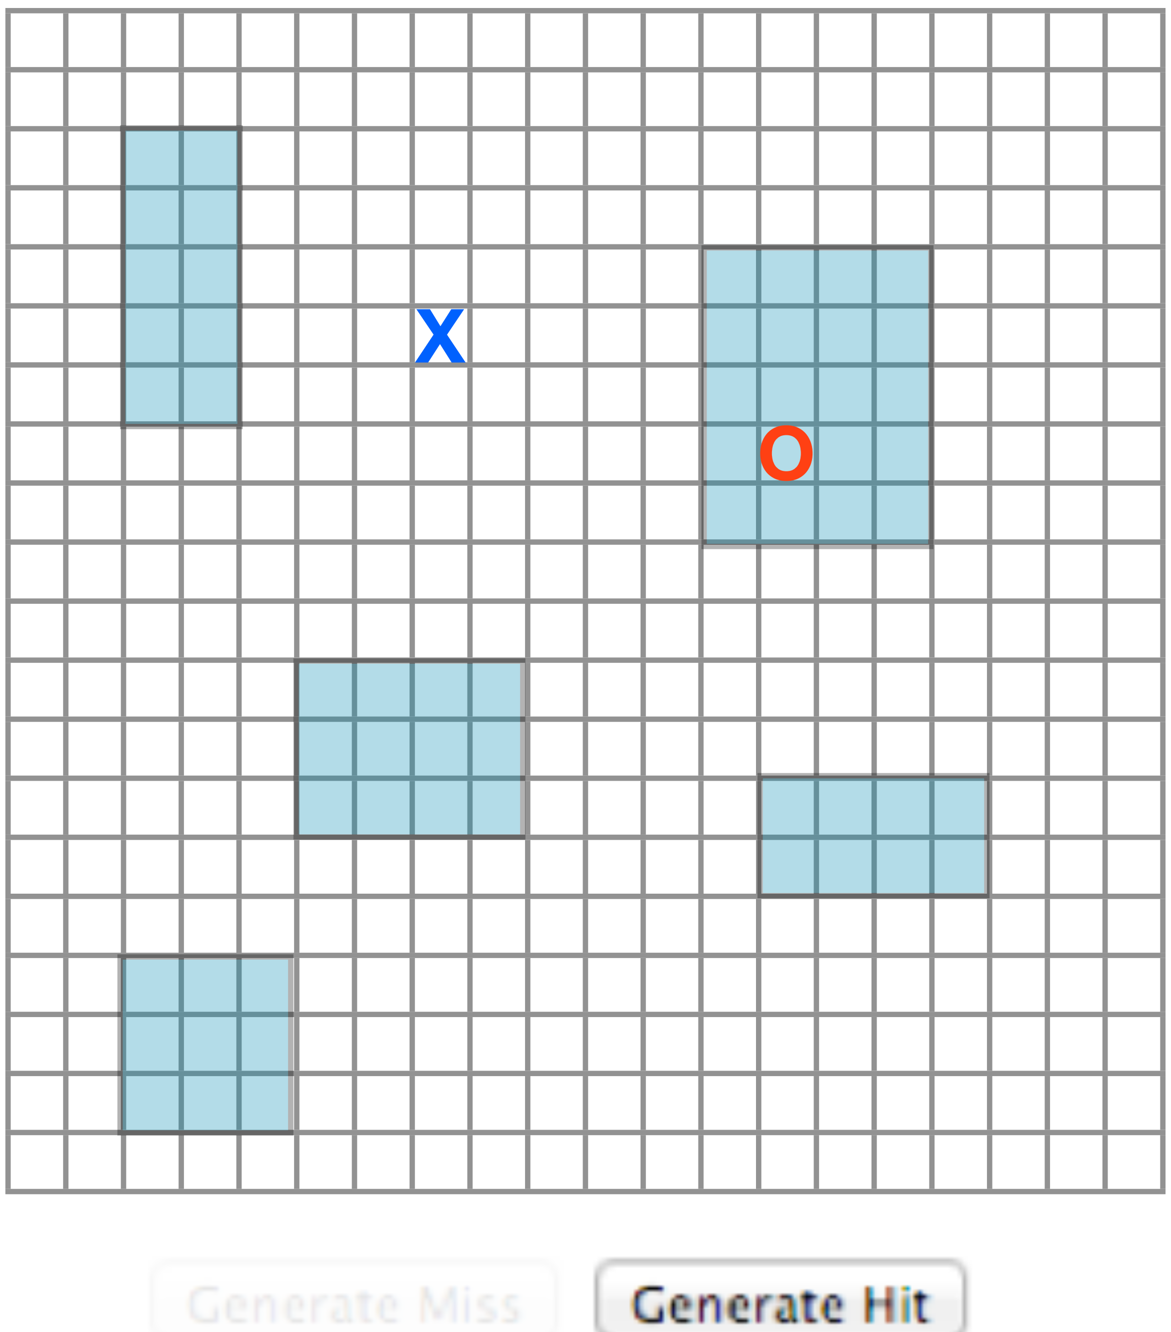

We can directly measure people's reliance on positive or negative tests by making them ask for those...

Test: Game of battleships

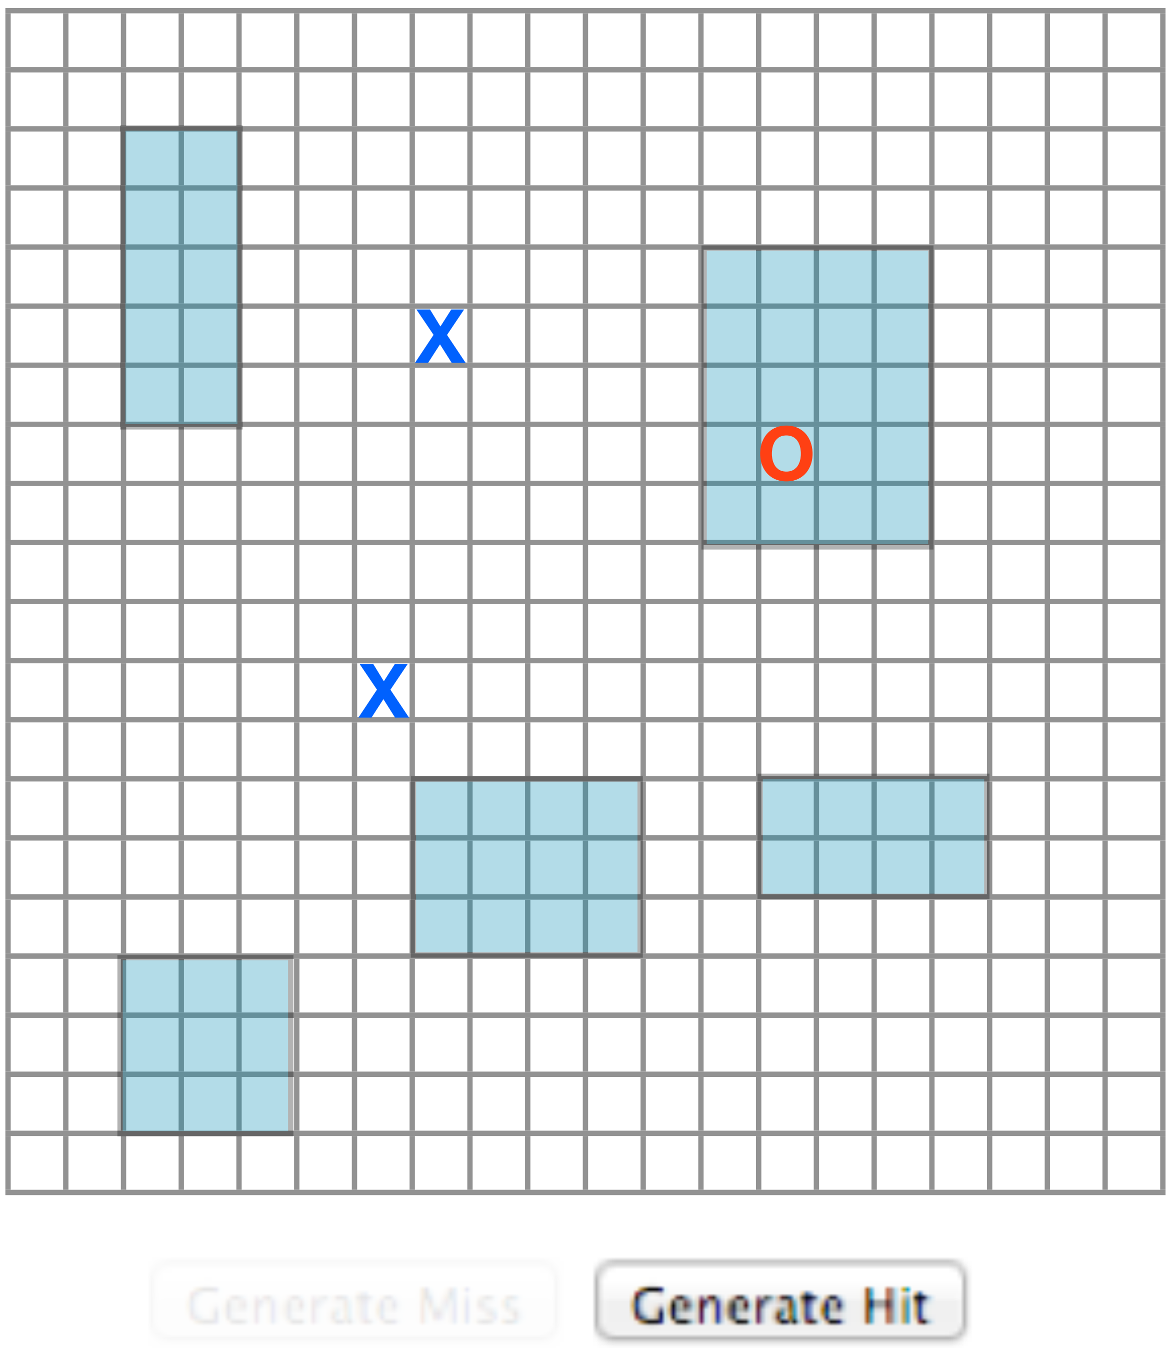

If one of the points is inconsistent with their ships, they have to move them around

Test: Game of battleships

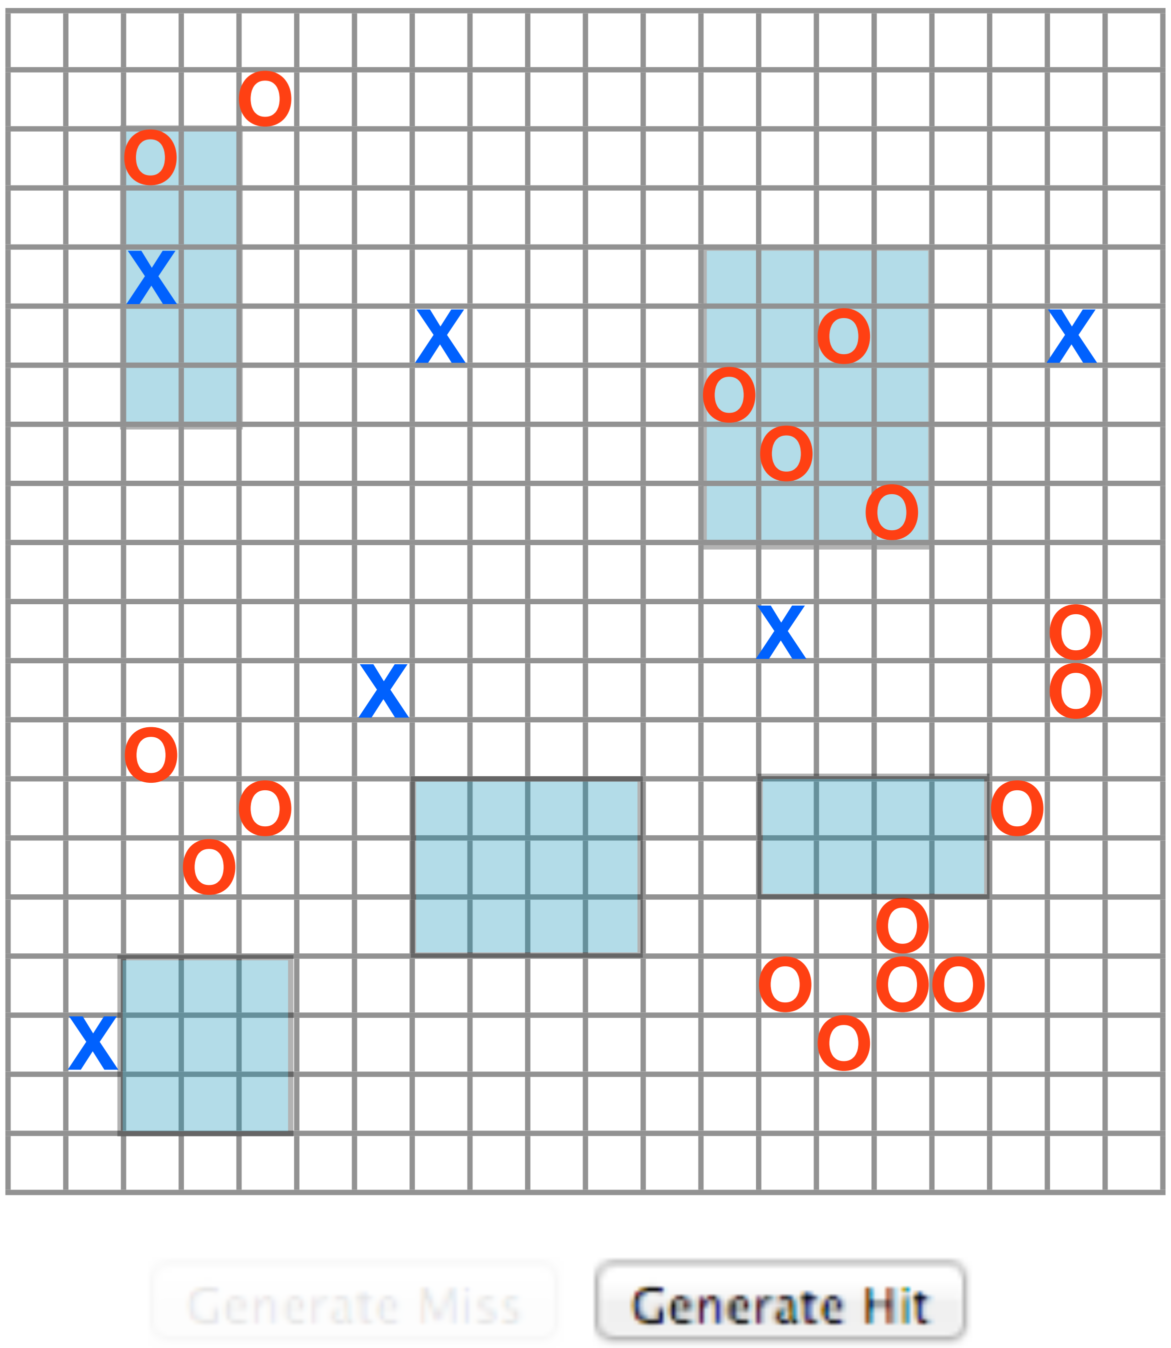

Keep generating points until they think they know where the ships are

Test: Game of battleships

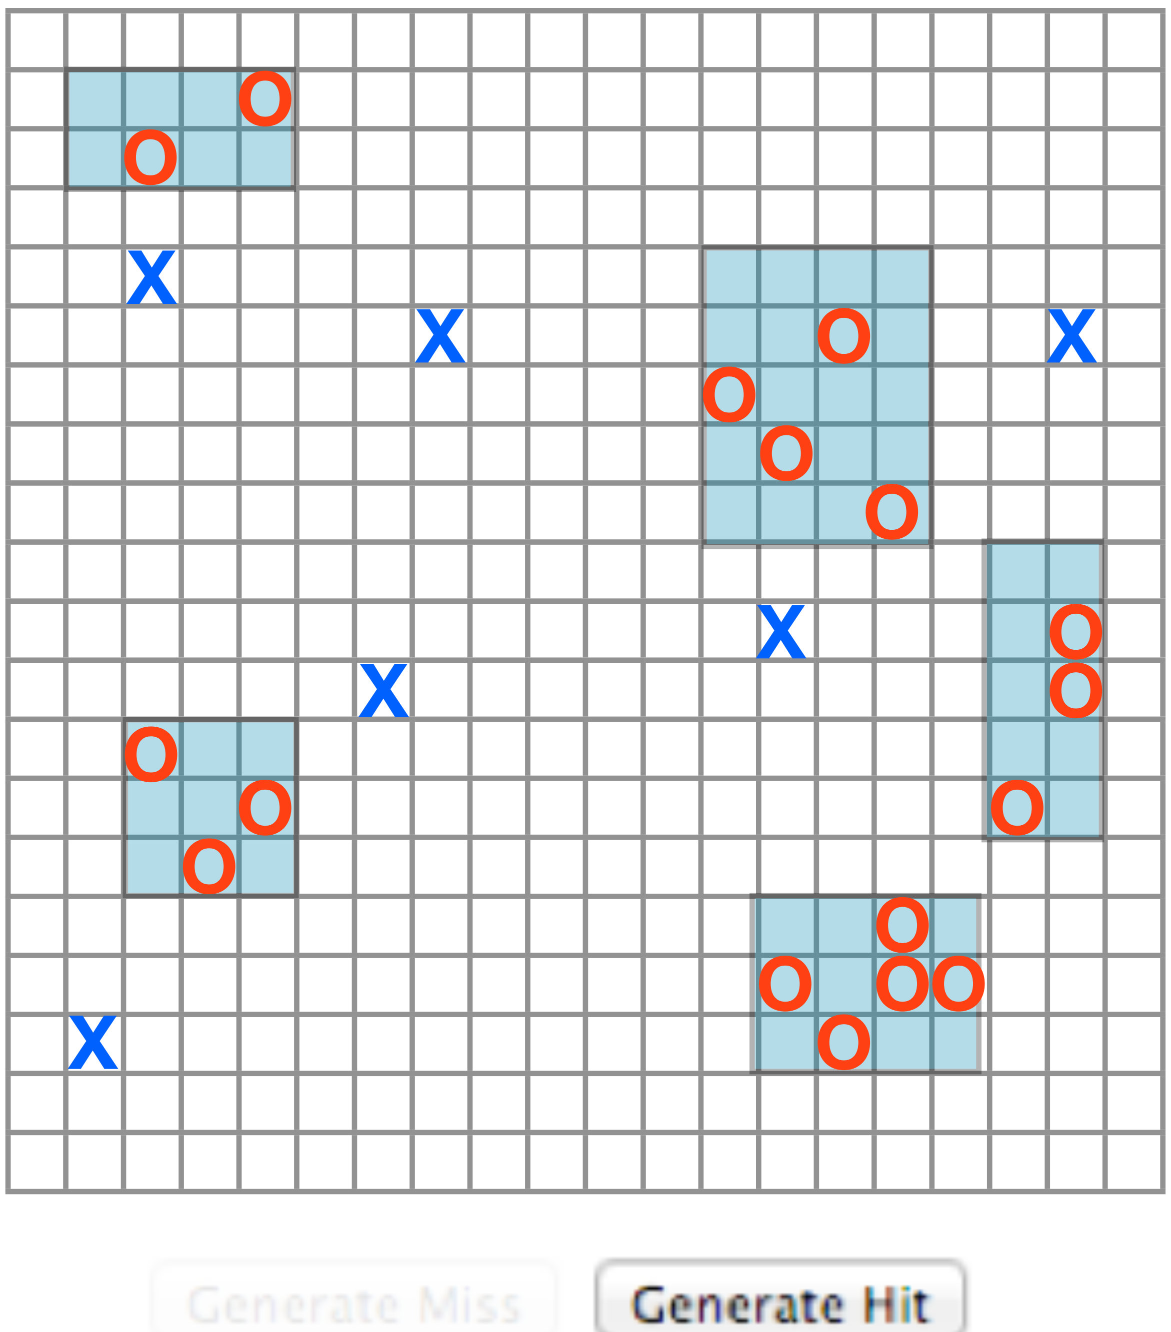

Keep generating points until they think they know where the ships are

Test: Game of battleships

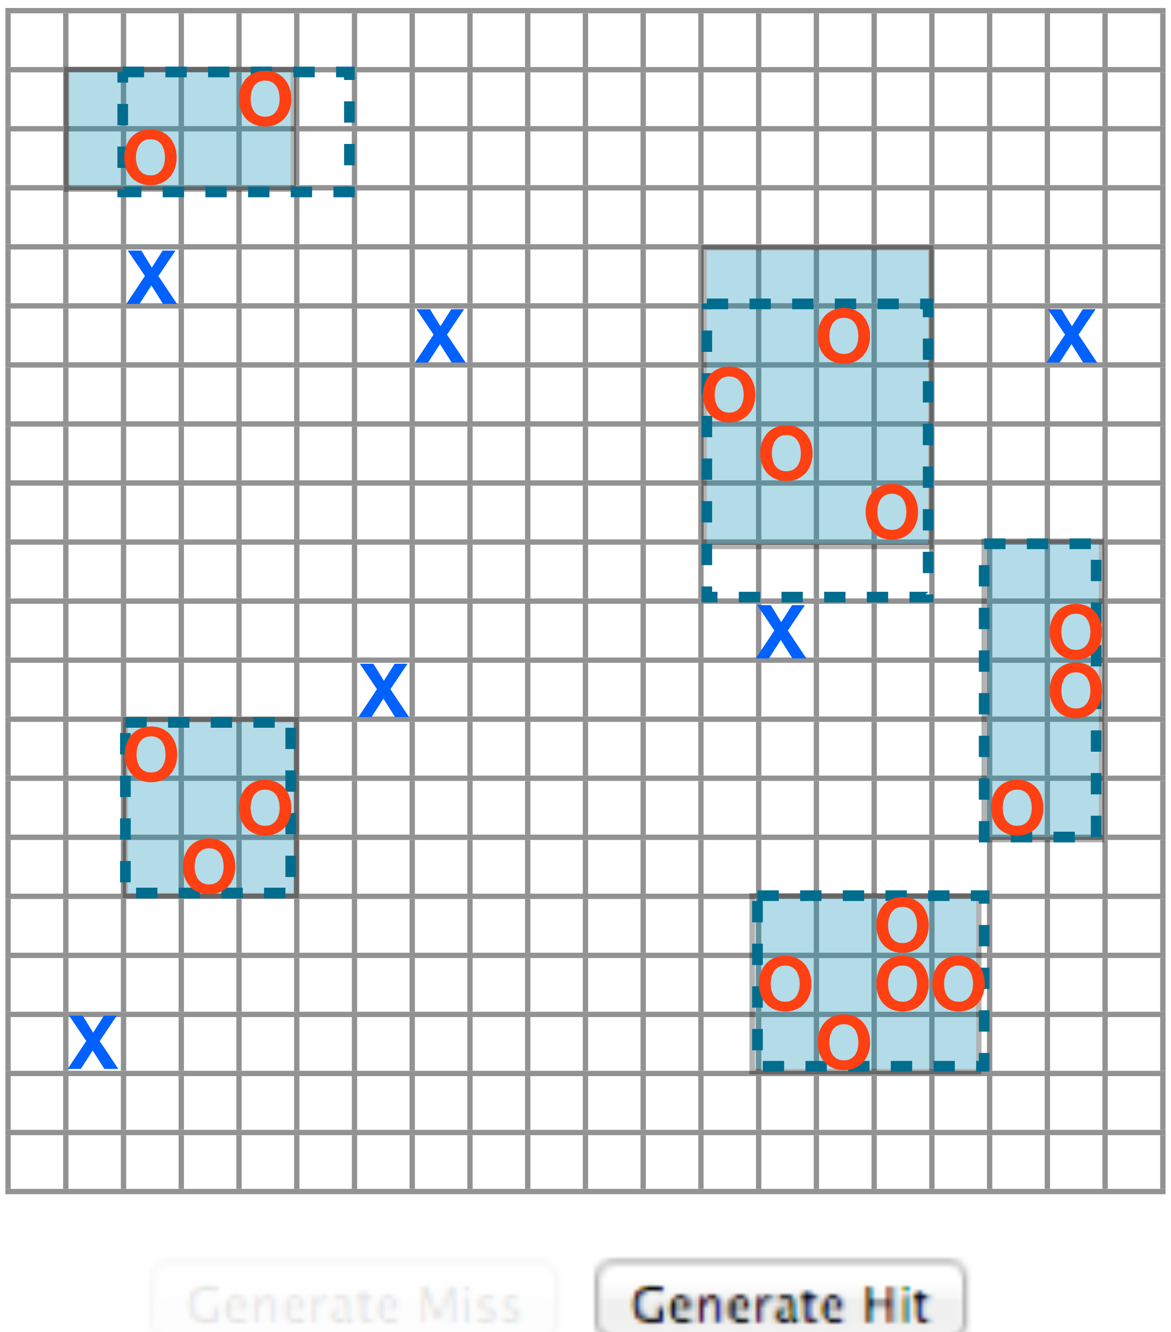

Get a score based on how much overlap they have with the true correct hypotheses

Changing the ships changes the sparsity!

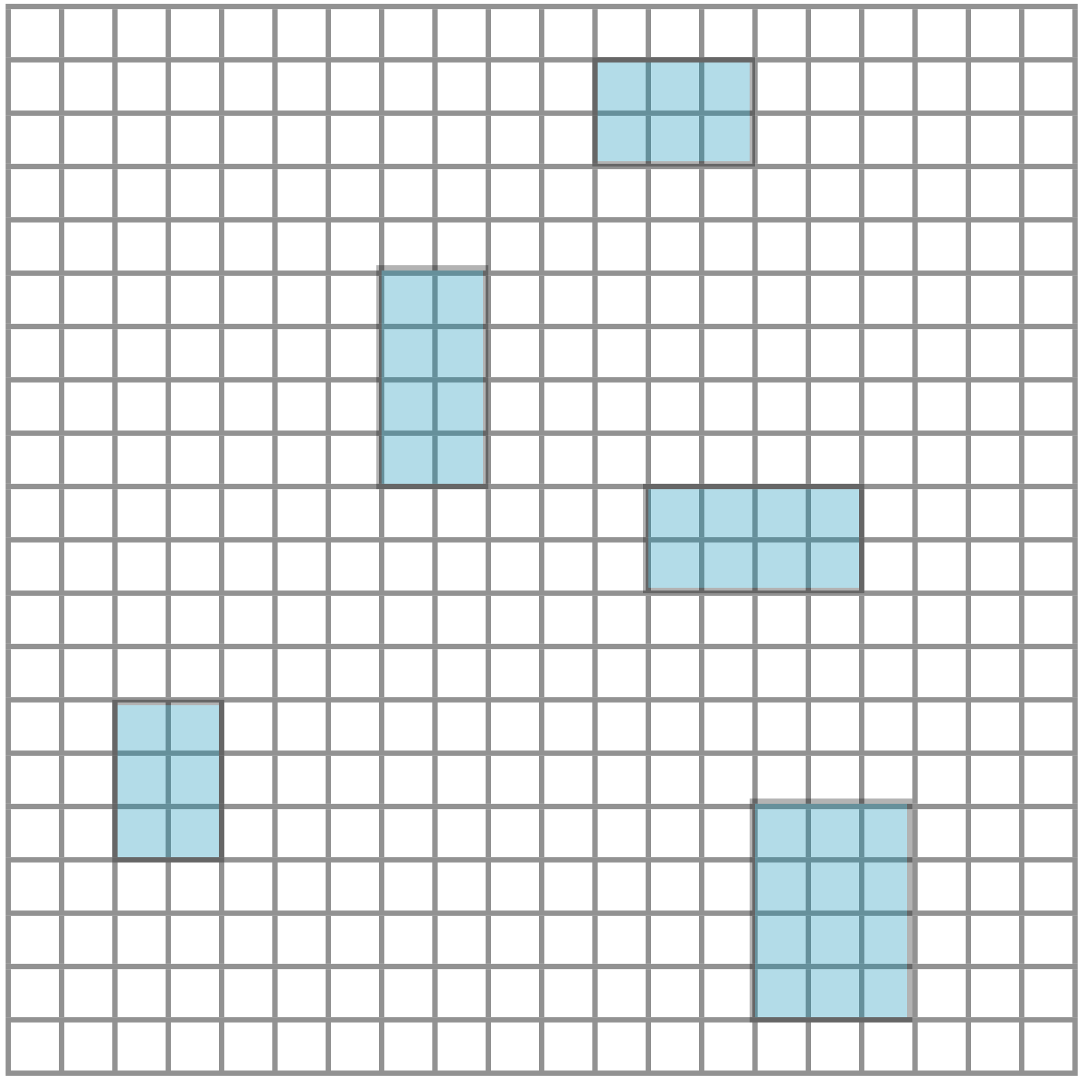

10% sparsity: the hypotheses cover 10% of the possible data points in the world

Changing the ships changes the sparsity!

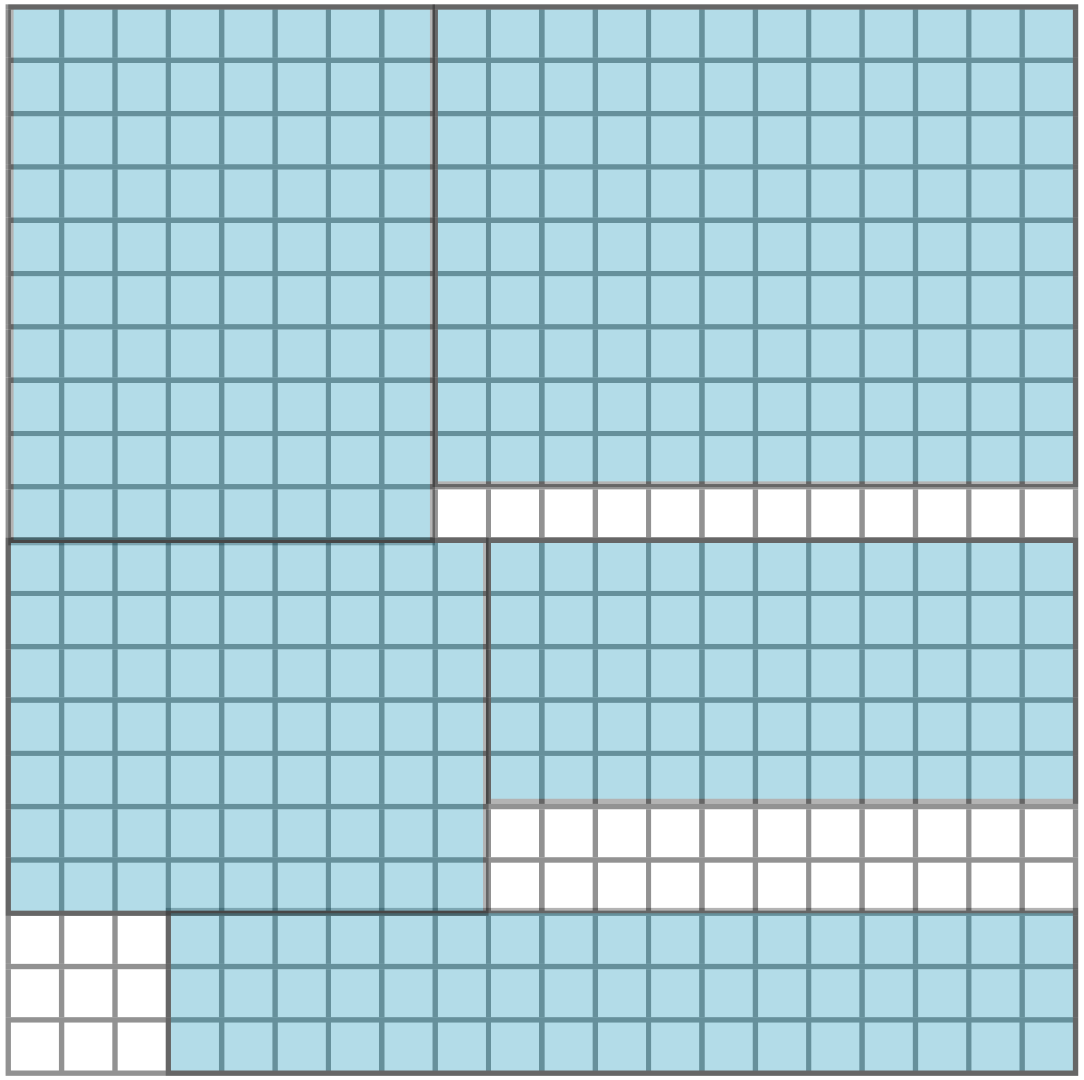

90% sparsity: the hypotheses cover 90% of the possible data points in the world

Changing the ships changes the sparsity!

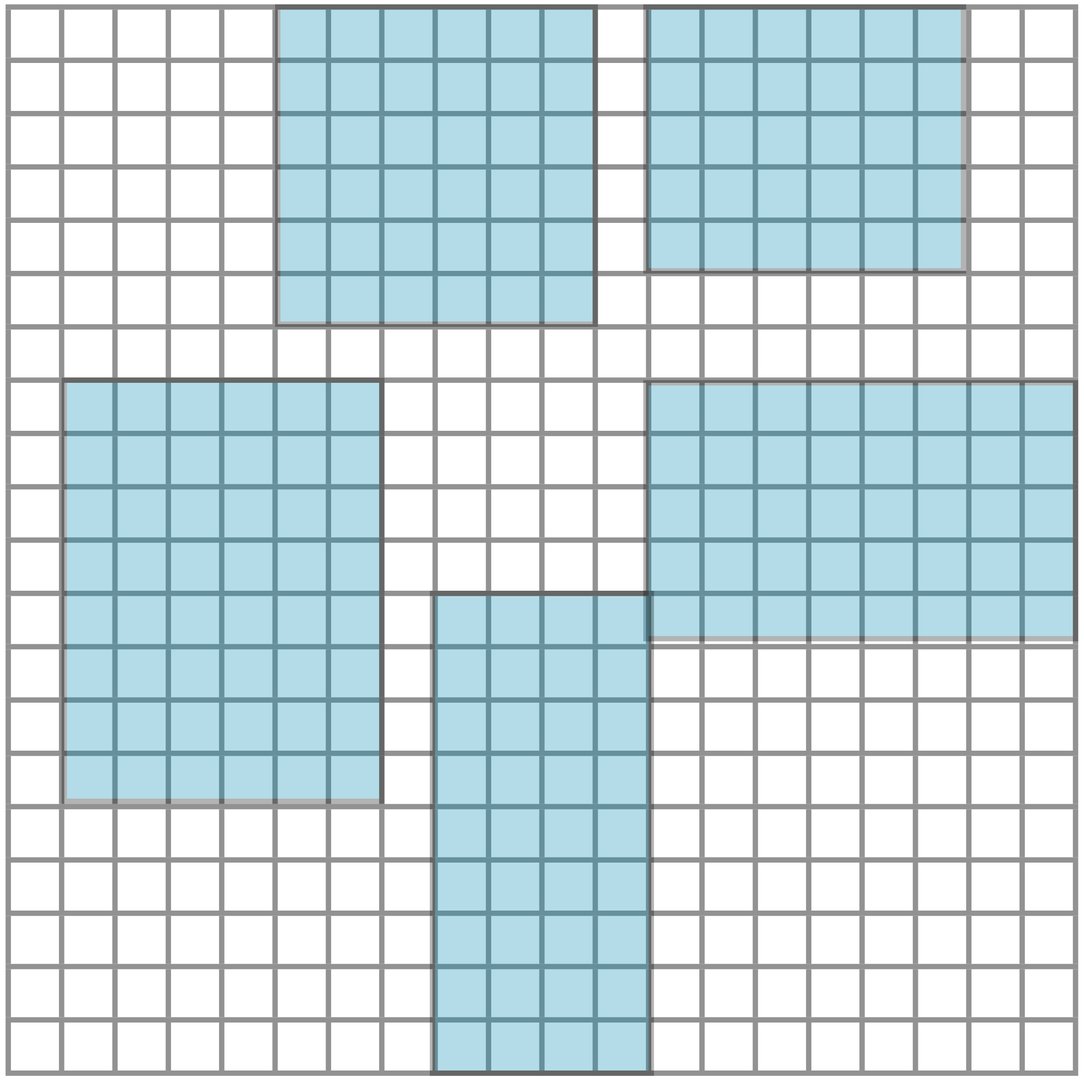

50% sparsity: the hypotheses cover 50% of the possible data points in the world

Sparsity governs which data is most useful

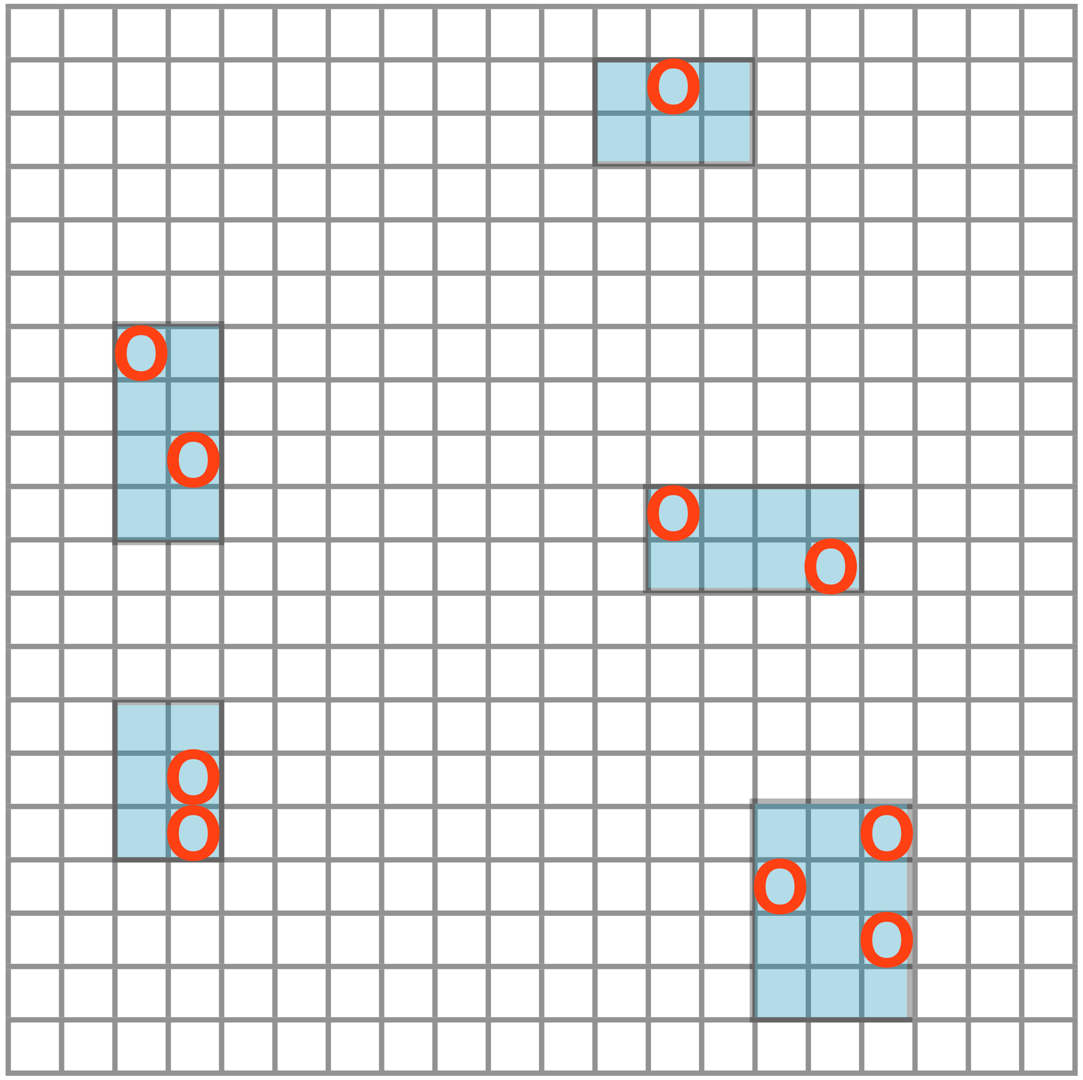

As predicted, when sparsity is low a positive test (asking for a hit) is much more informative

Sparsity governs which data is most useful

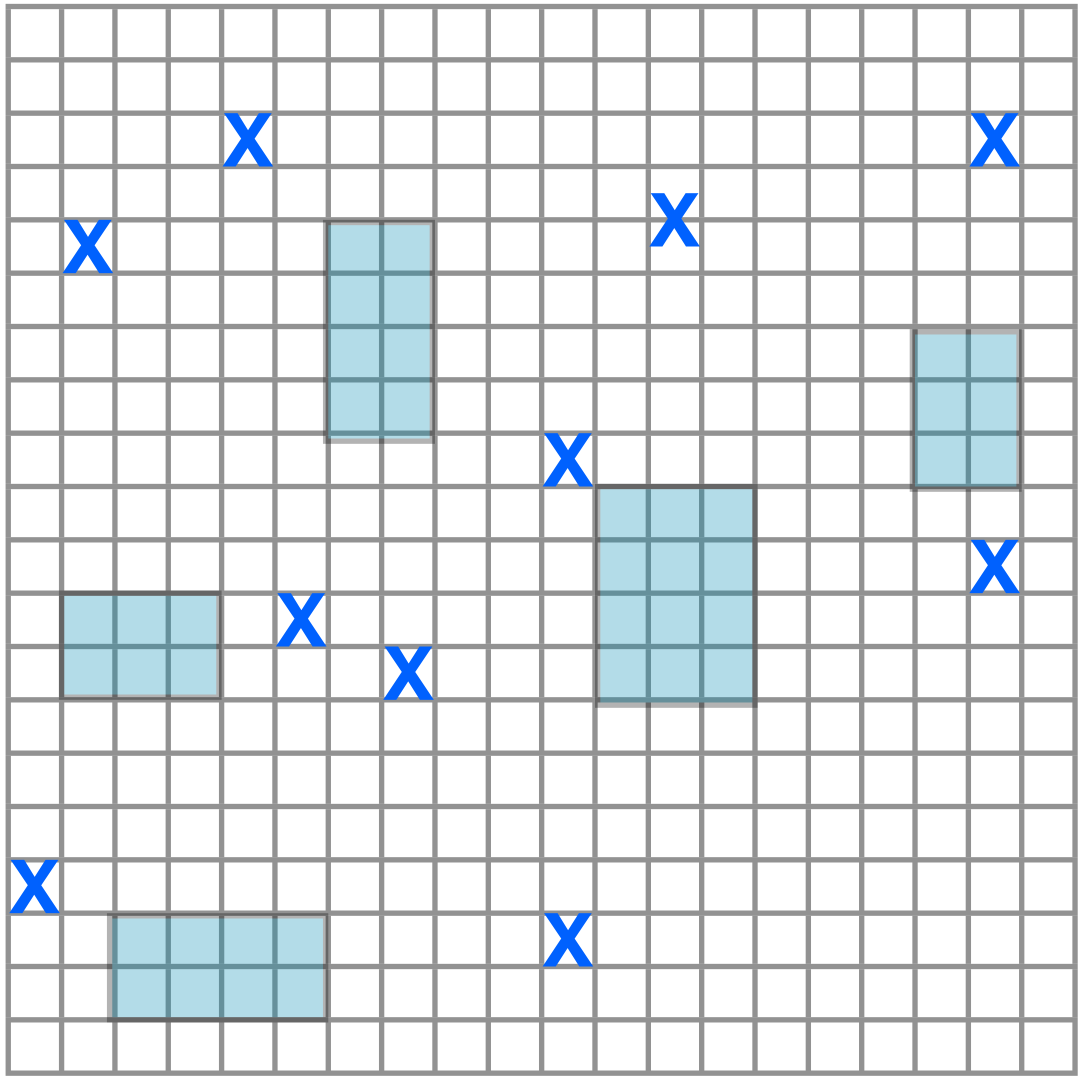

As predicted, when sparsity is low a positive test (asking for a hit) is much more informative

Sparsity governs which data is most useful

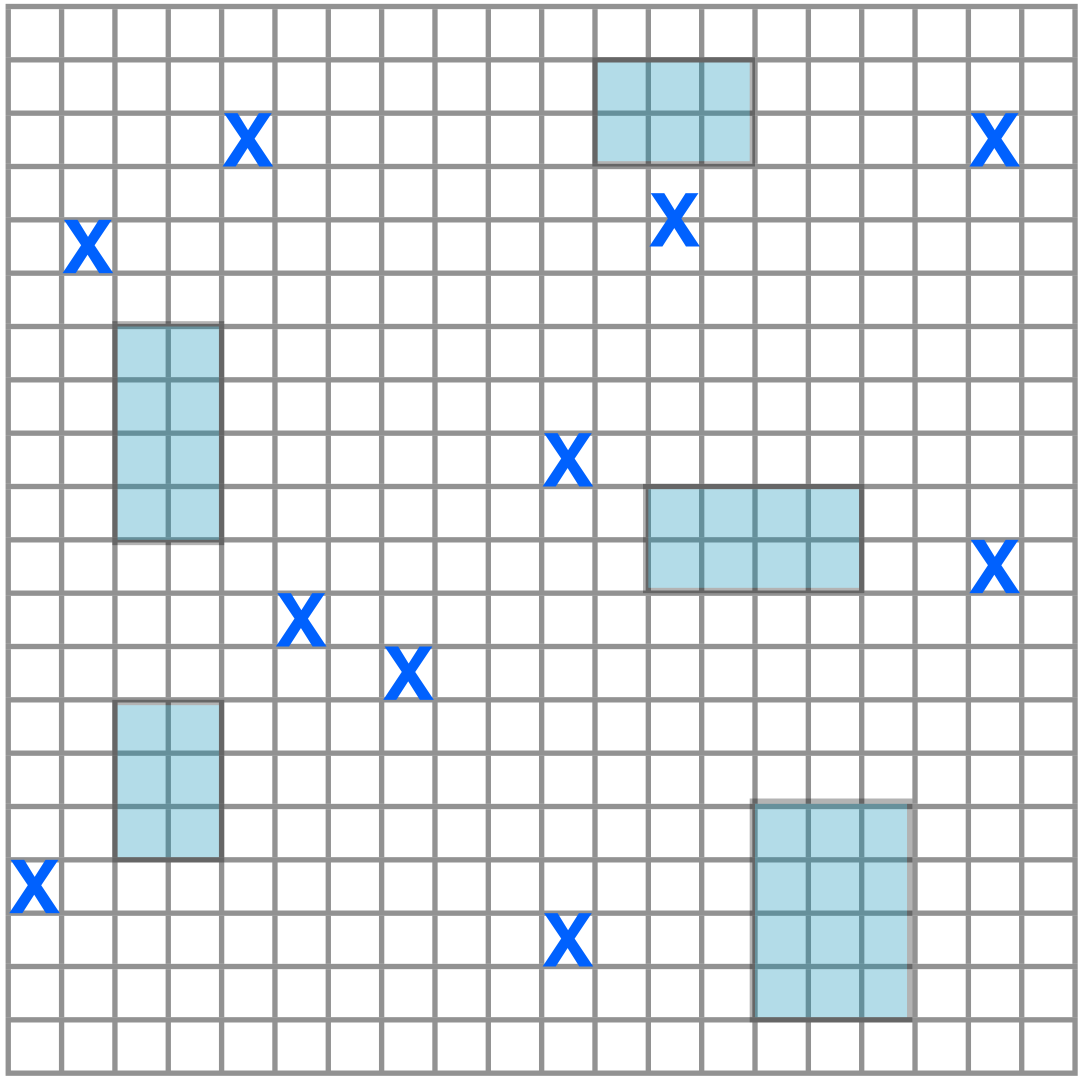

As predicted, when sparsity is low a positive test (asking for a hit) is much more informative

Sparsity governs which data is most useful

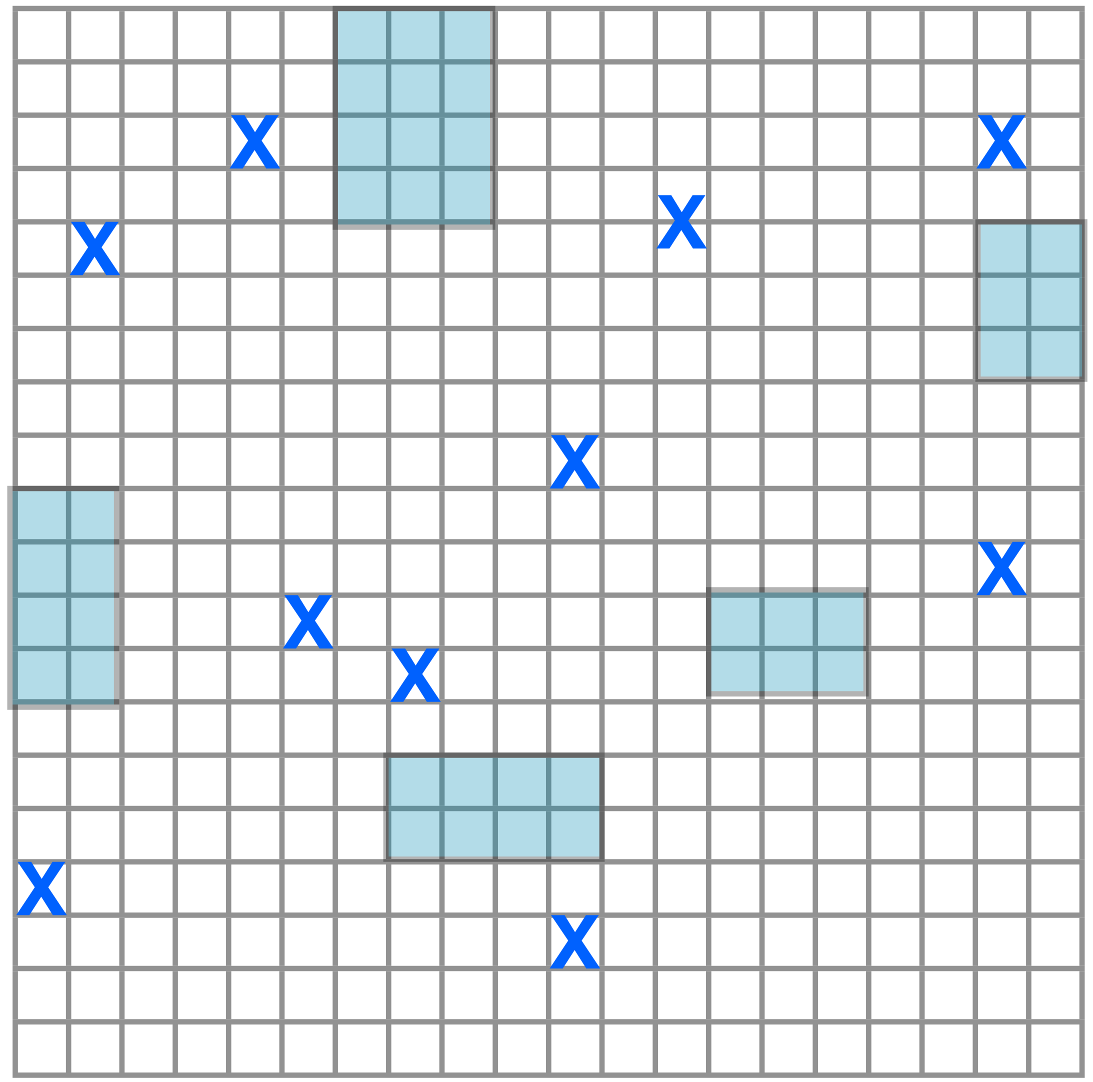

As predicted, when sparsity is low a positive test (asking for a hit) is much more informative

Sparsity governs which data is most useful

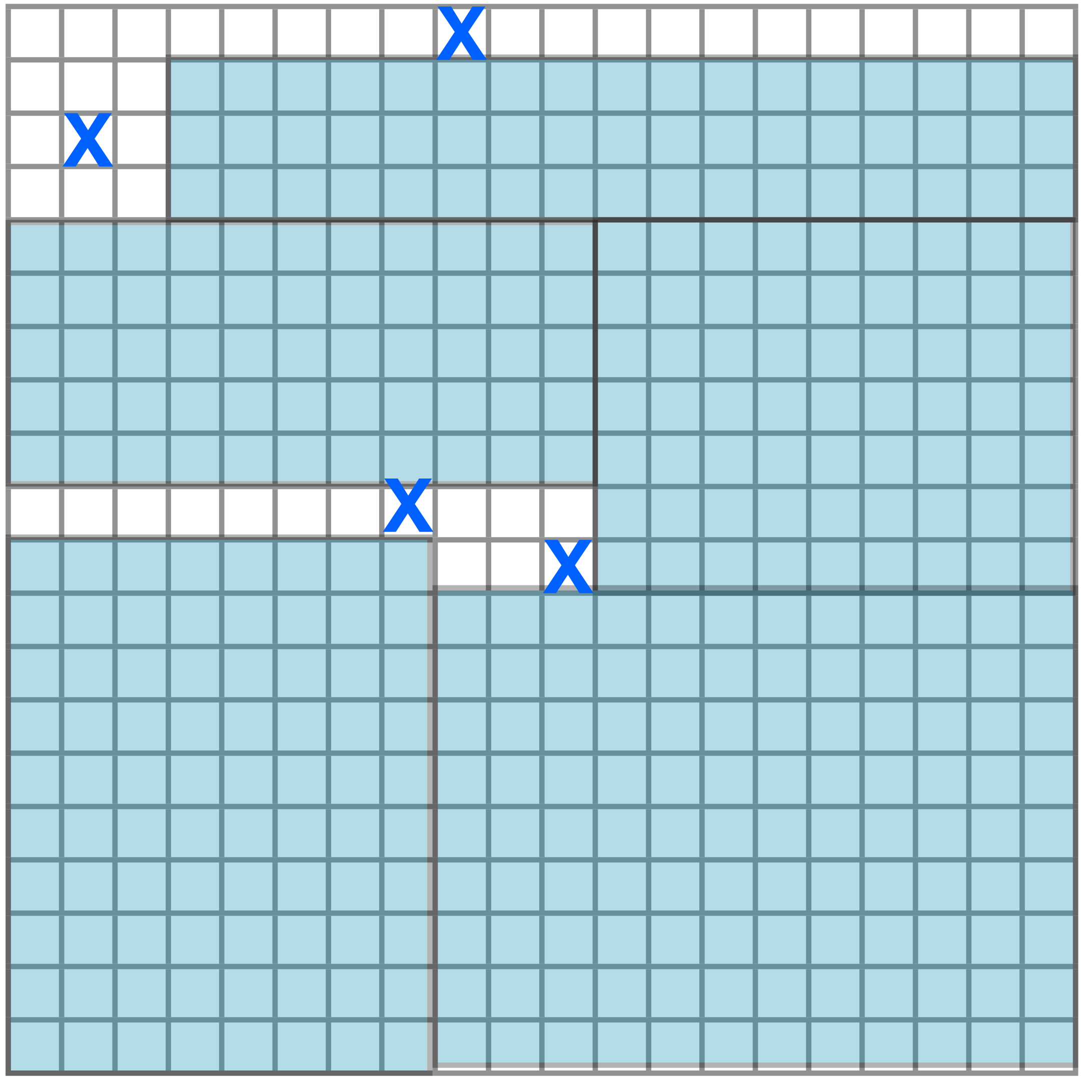

But when sparsity is high asking for a *miss* is much more informative

Sparsity governs which data is most useful

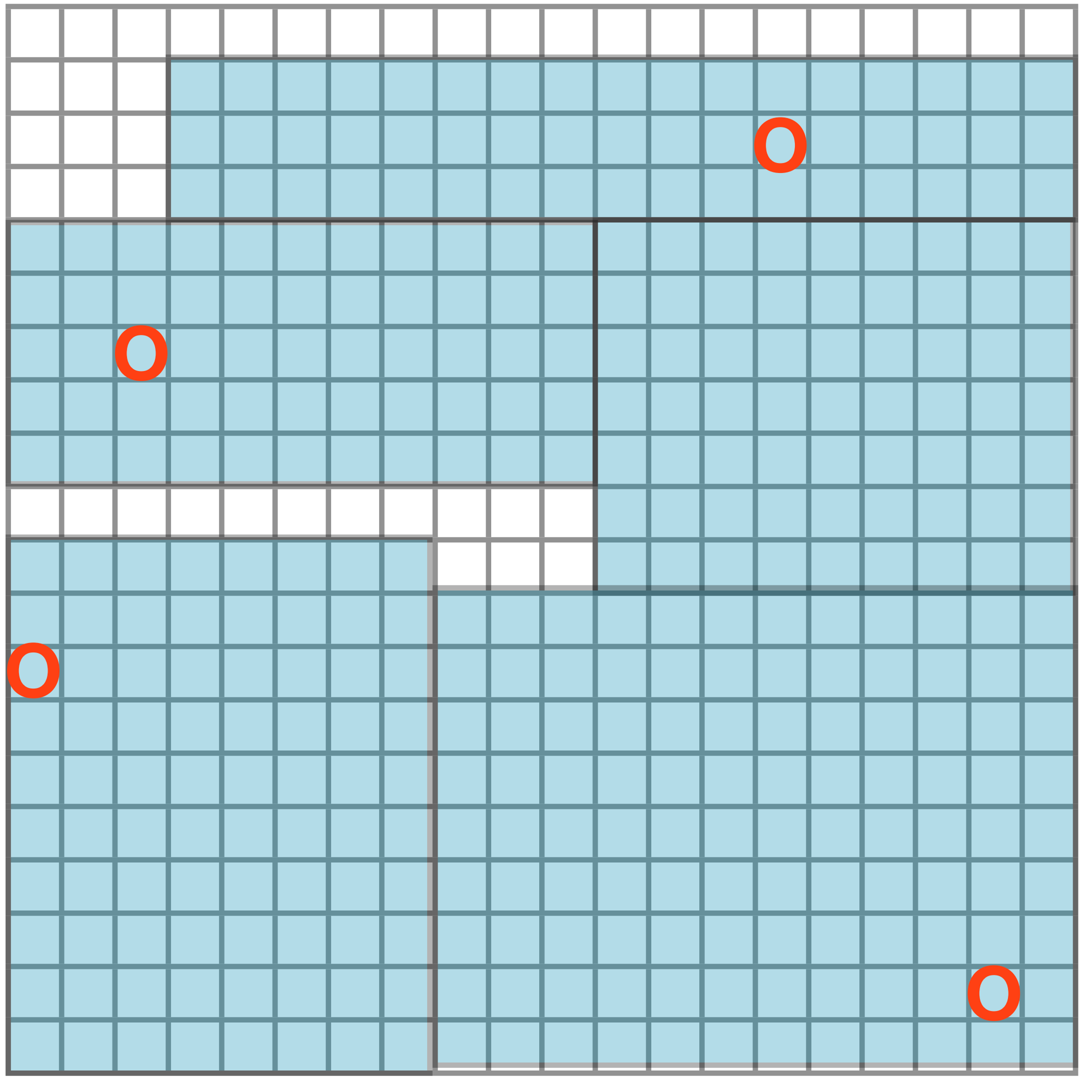

But when sparsity is high asking for a *miss* is much more informative

Sparsity governs which data is most useful

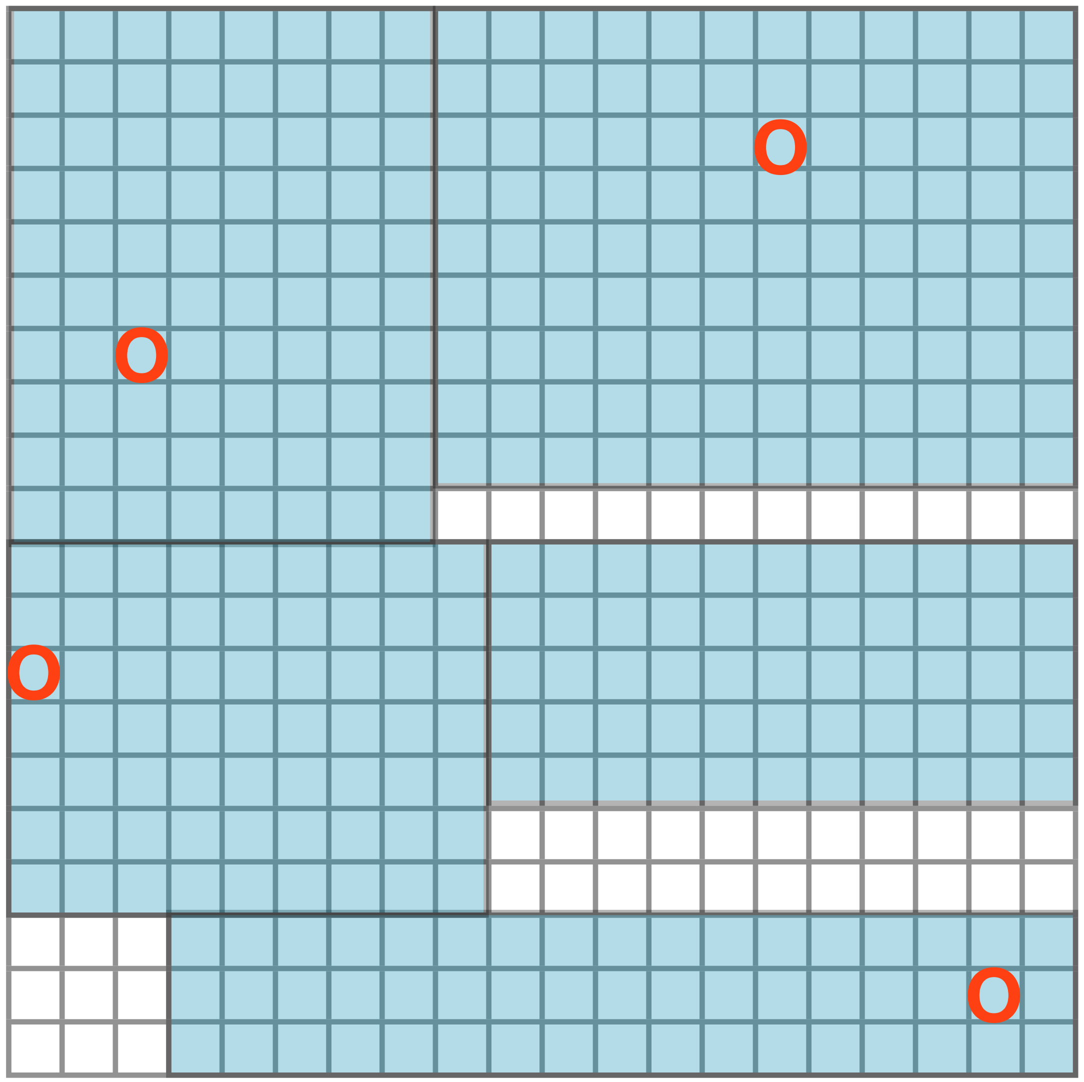

But when sparsity is high asking for a *miss* is
much more informative

Sparsity governs which data is most useful

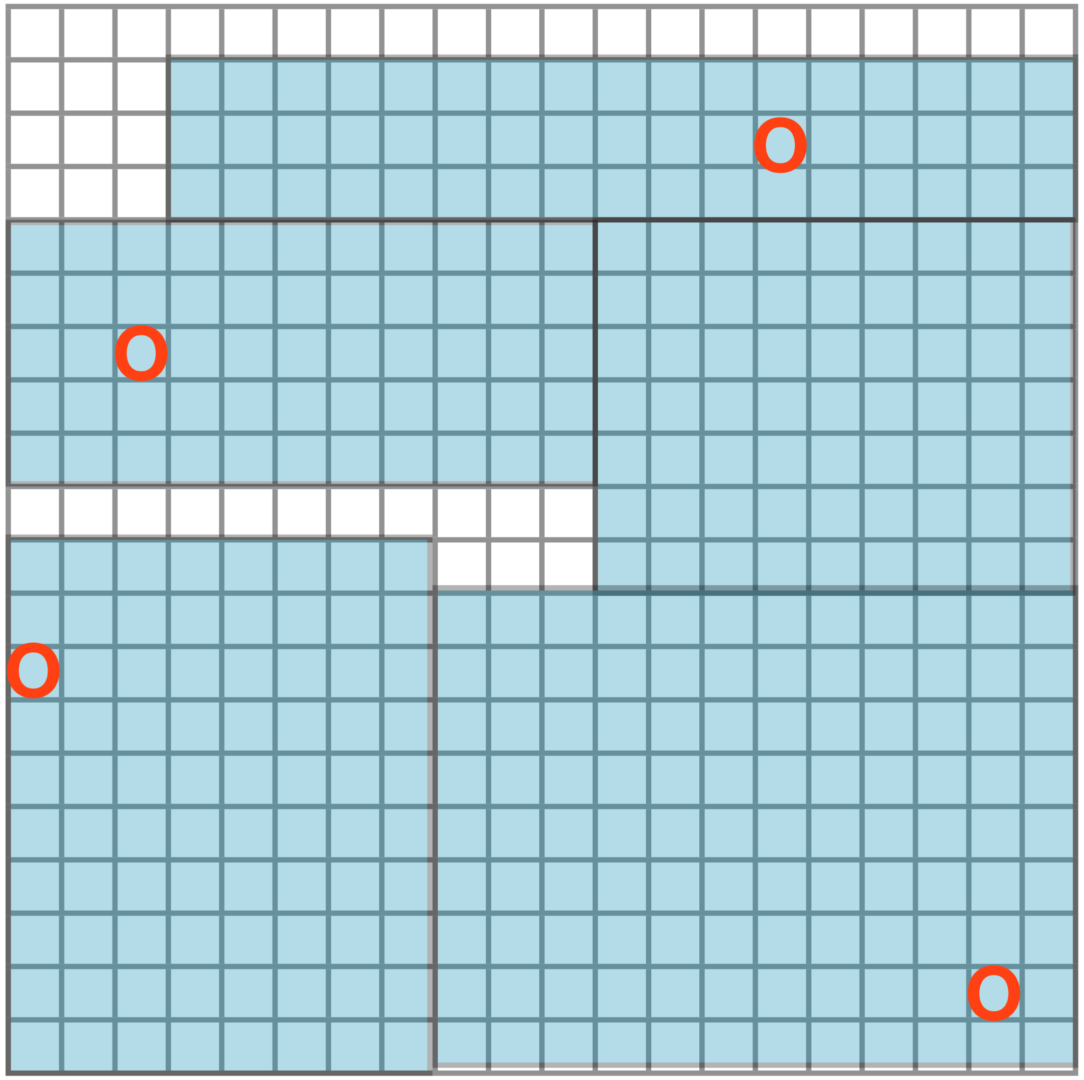

But when sparsity is high asking for a *miss* is much more informative

Sparsity governs which data is most useful

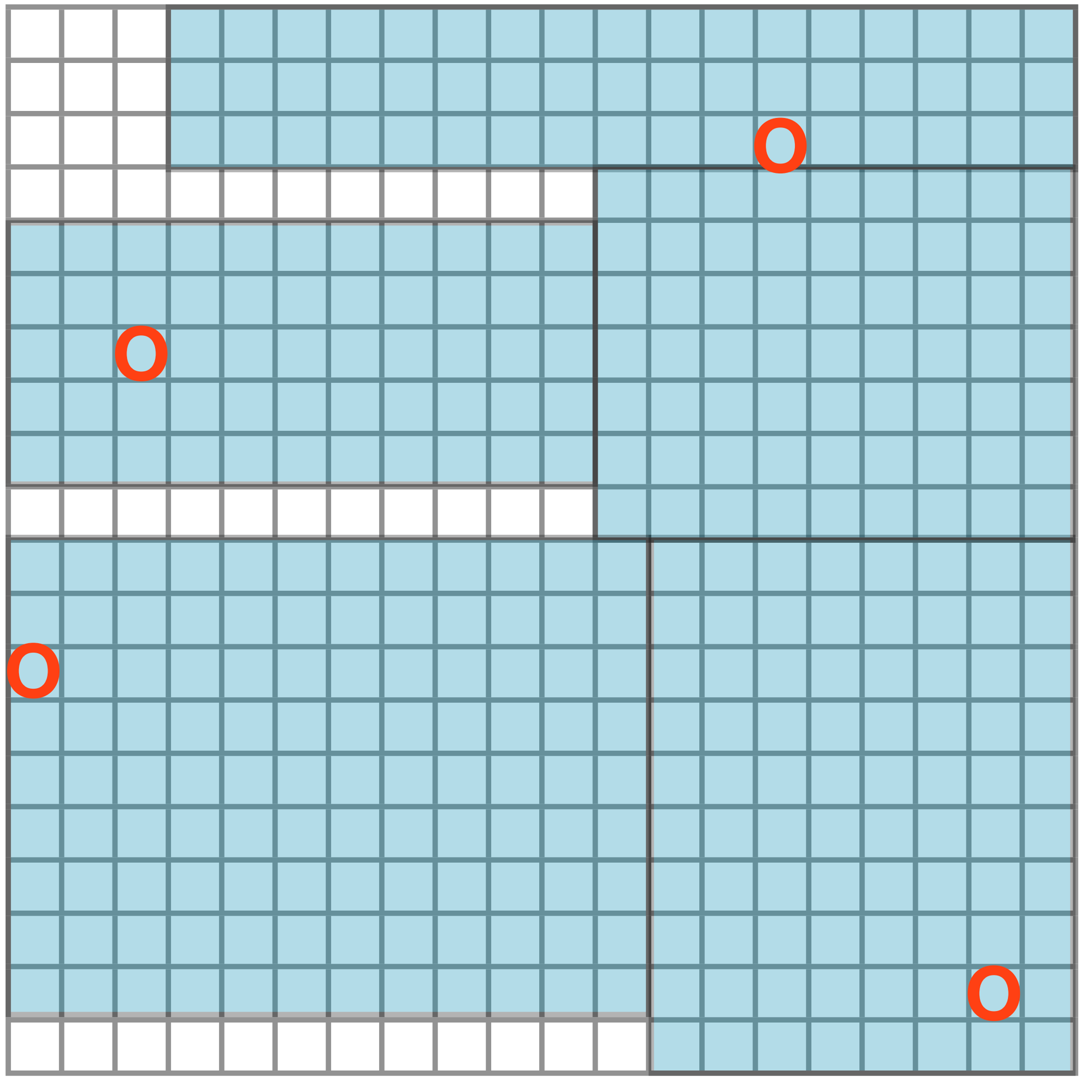

But when sparsity is high asking for a *miss* is much more informative

Do people change their strategy with sparsity?

yes; they ask for more hits (positive evidence) when the hypothesis space is sparse

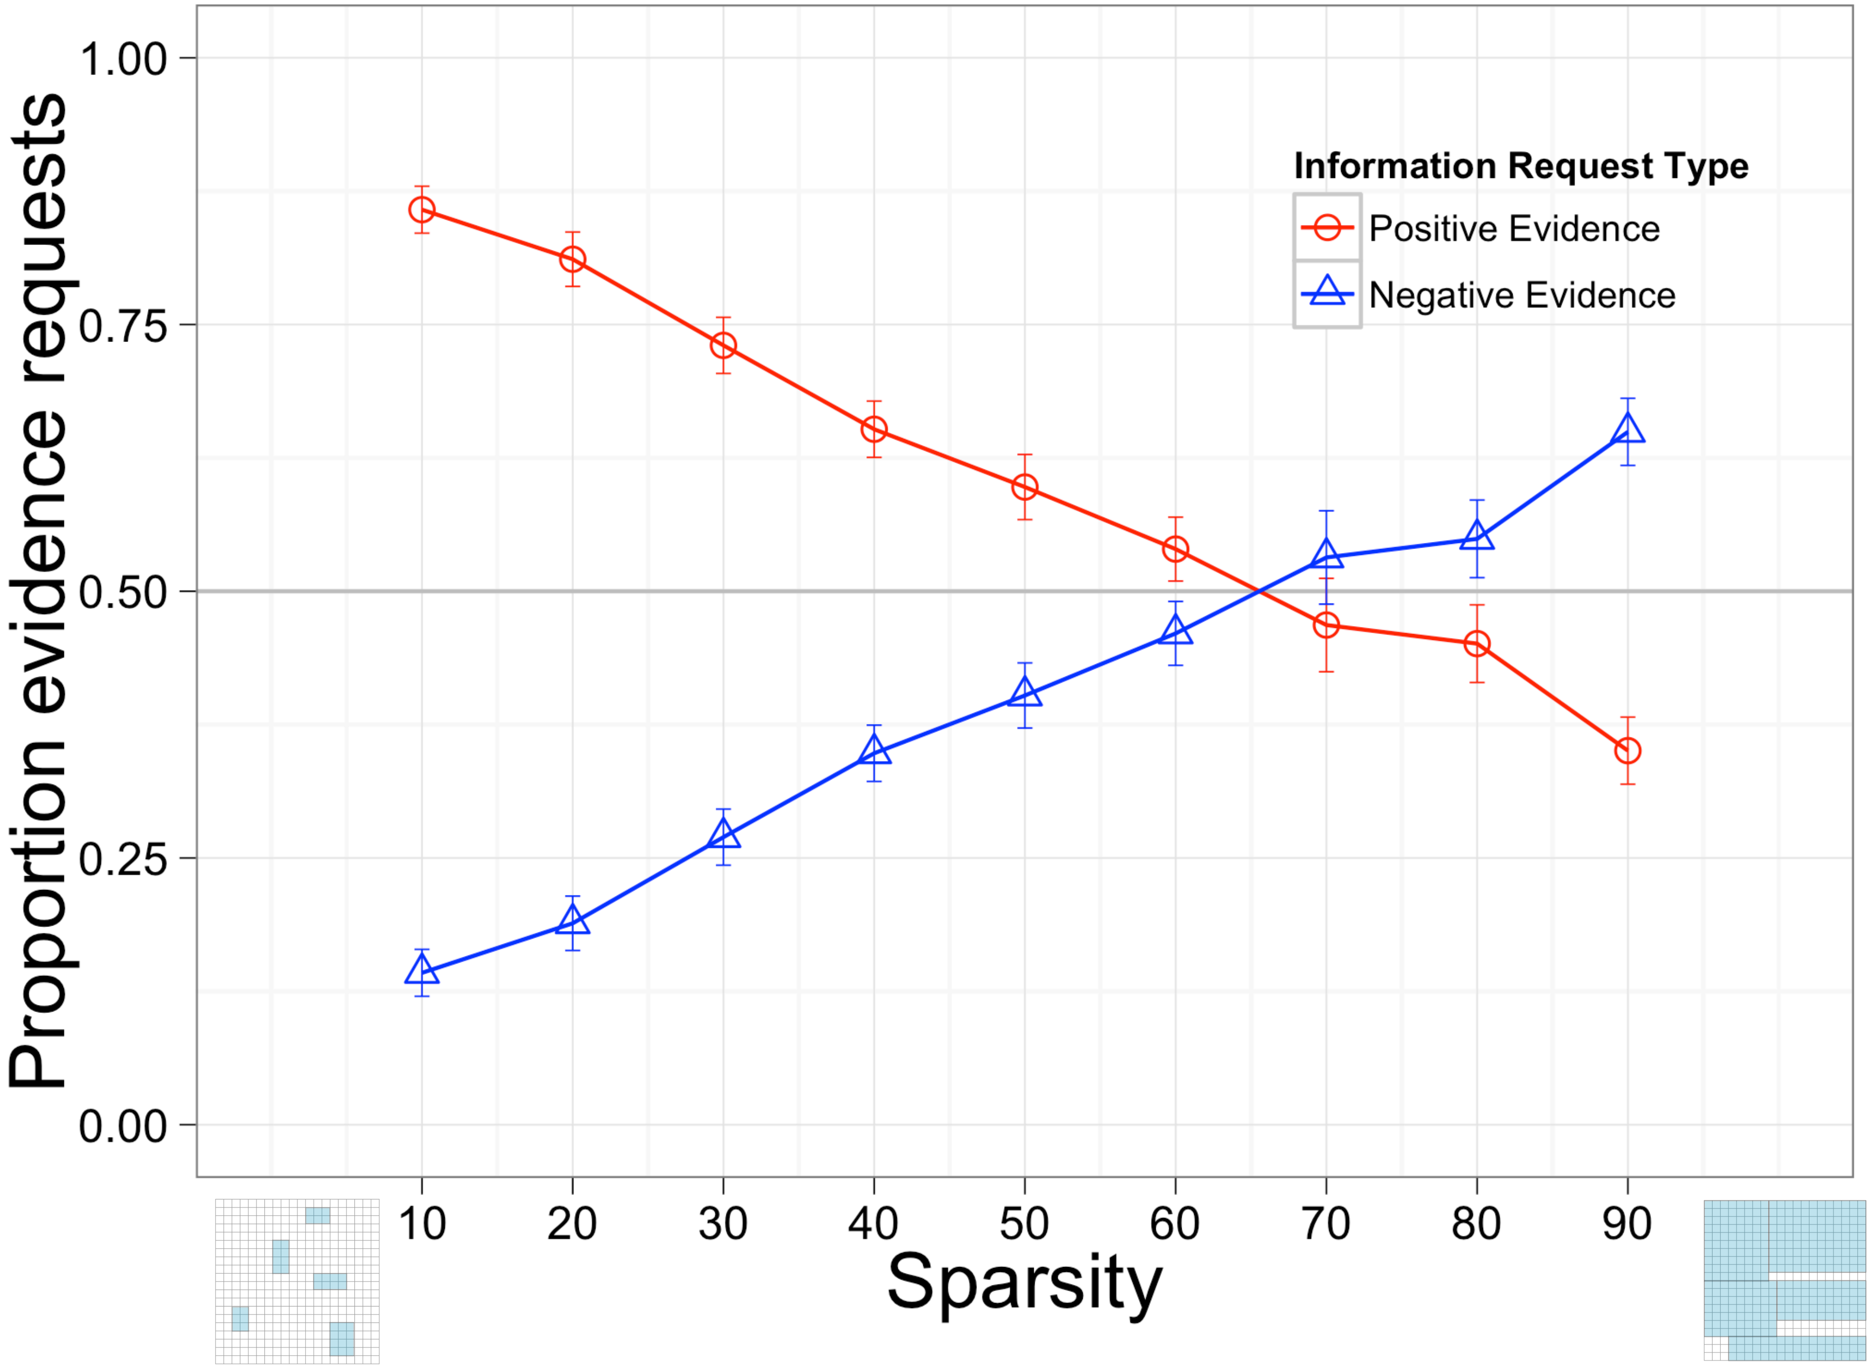

Do people change their strategy with sparsity?

yes; they ask for more hits (positive evidence) when the hypothesis space is sparse

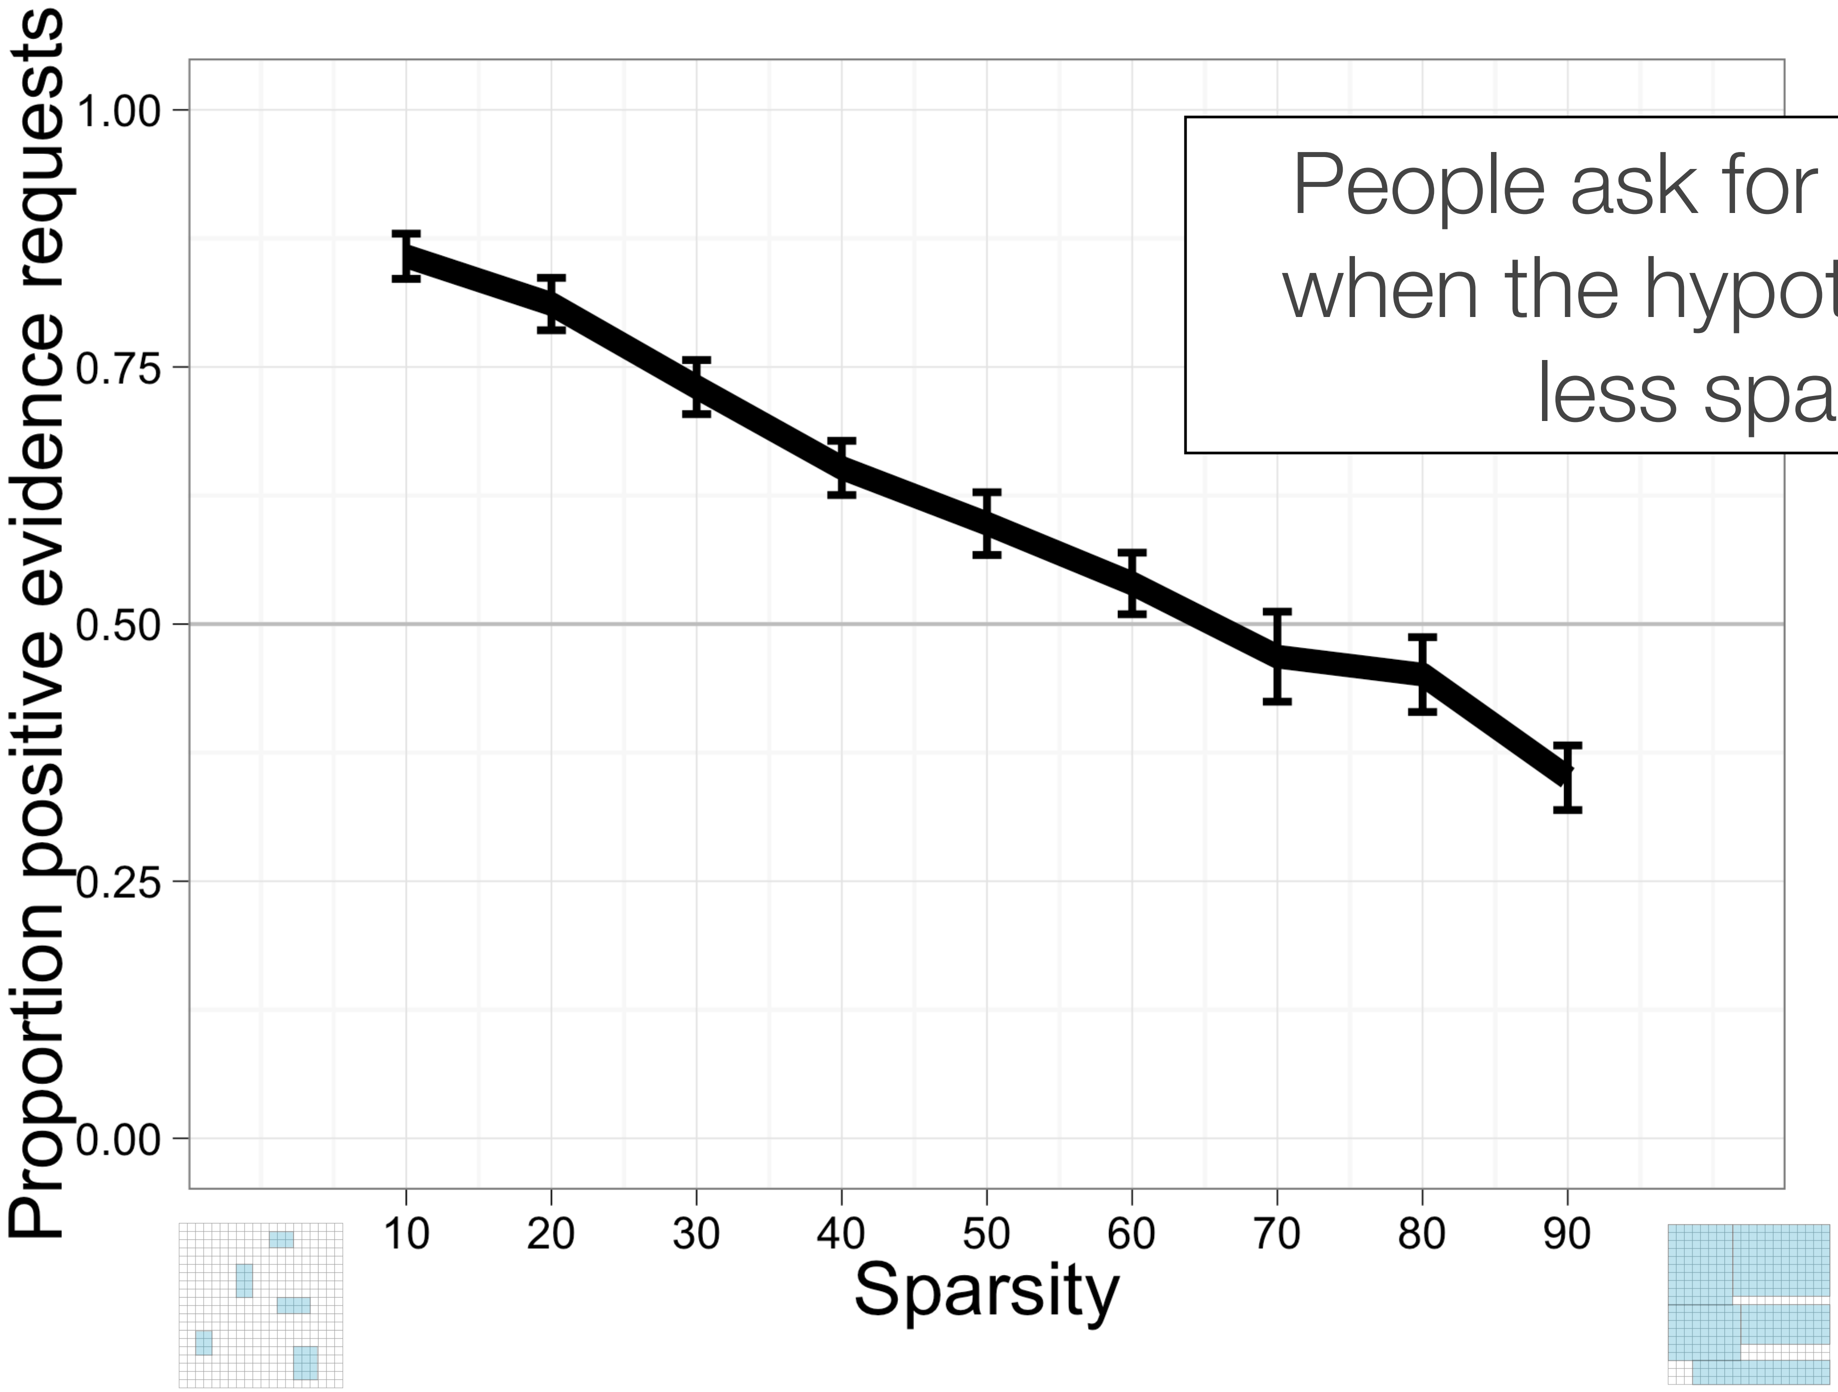

People ask for fewer hits when the hypotheses are less sparse

Do people change their strategy with sparsity?

yes; they ask for more hits (positive evidence) when the hypothesis space is sparse

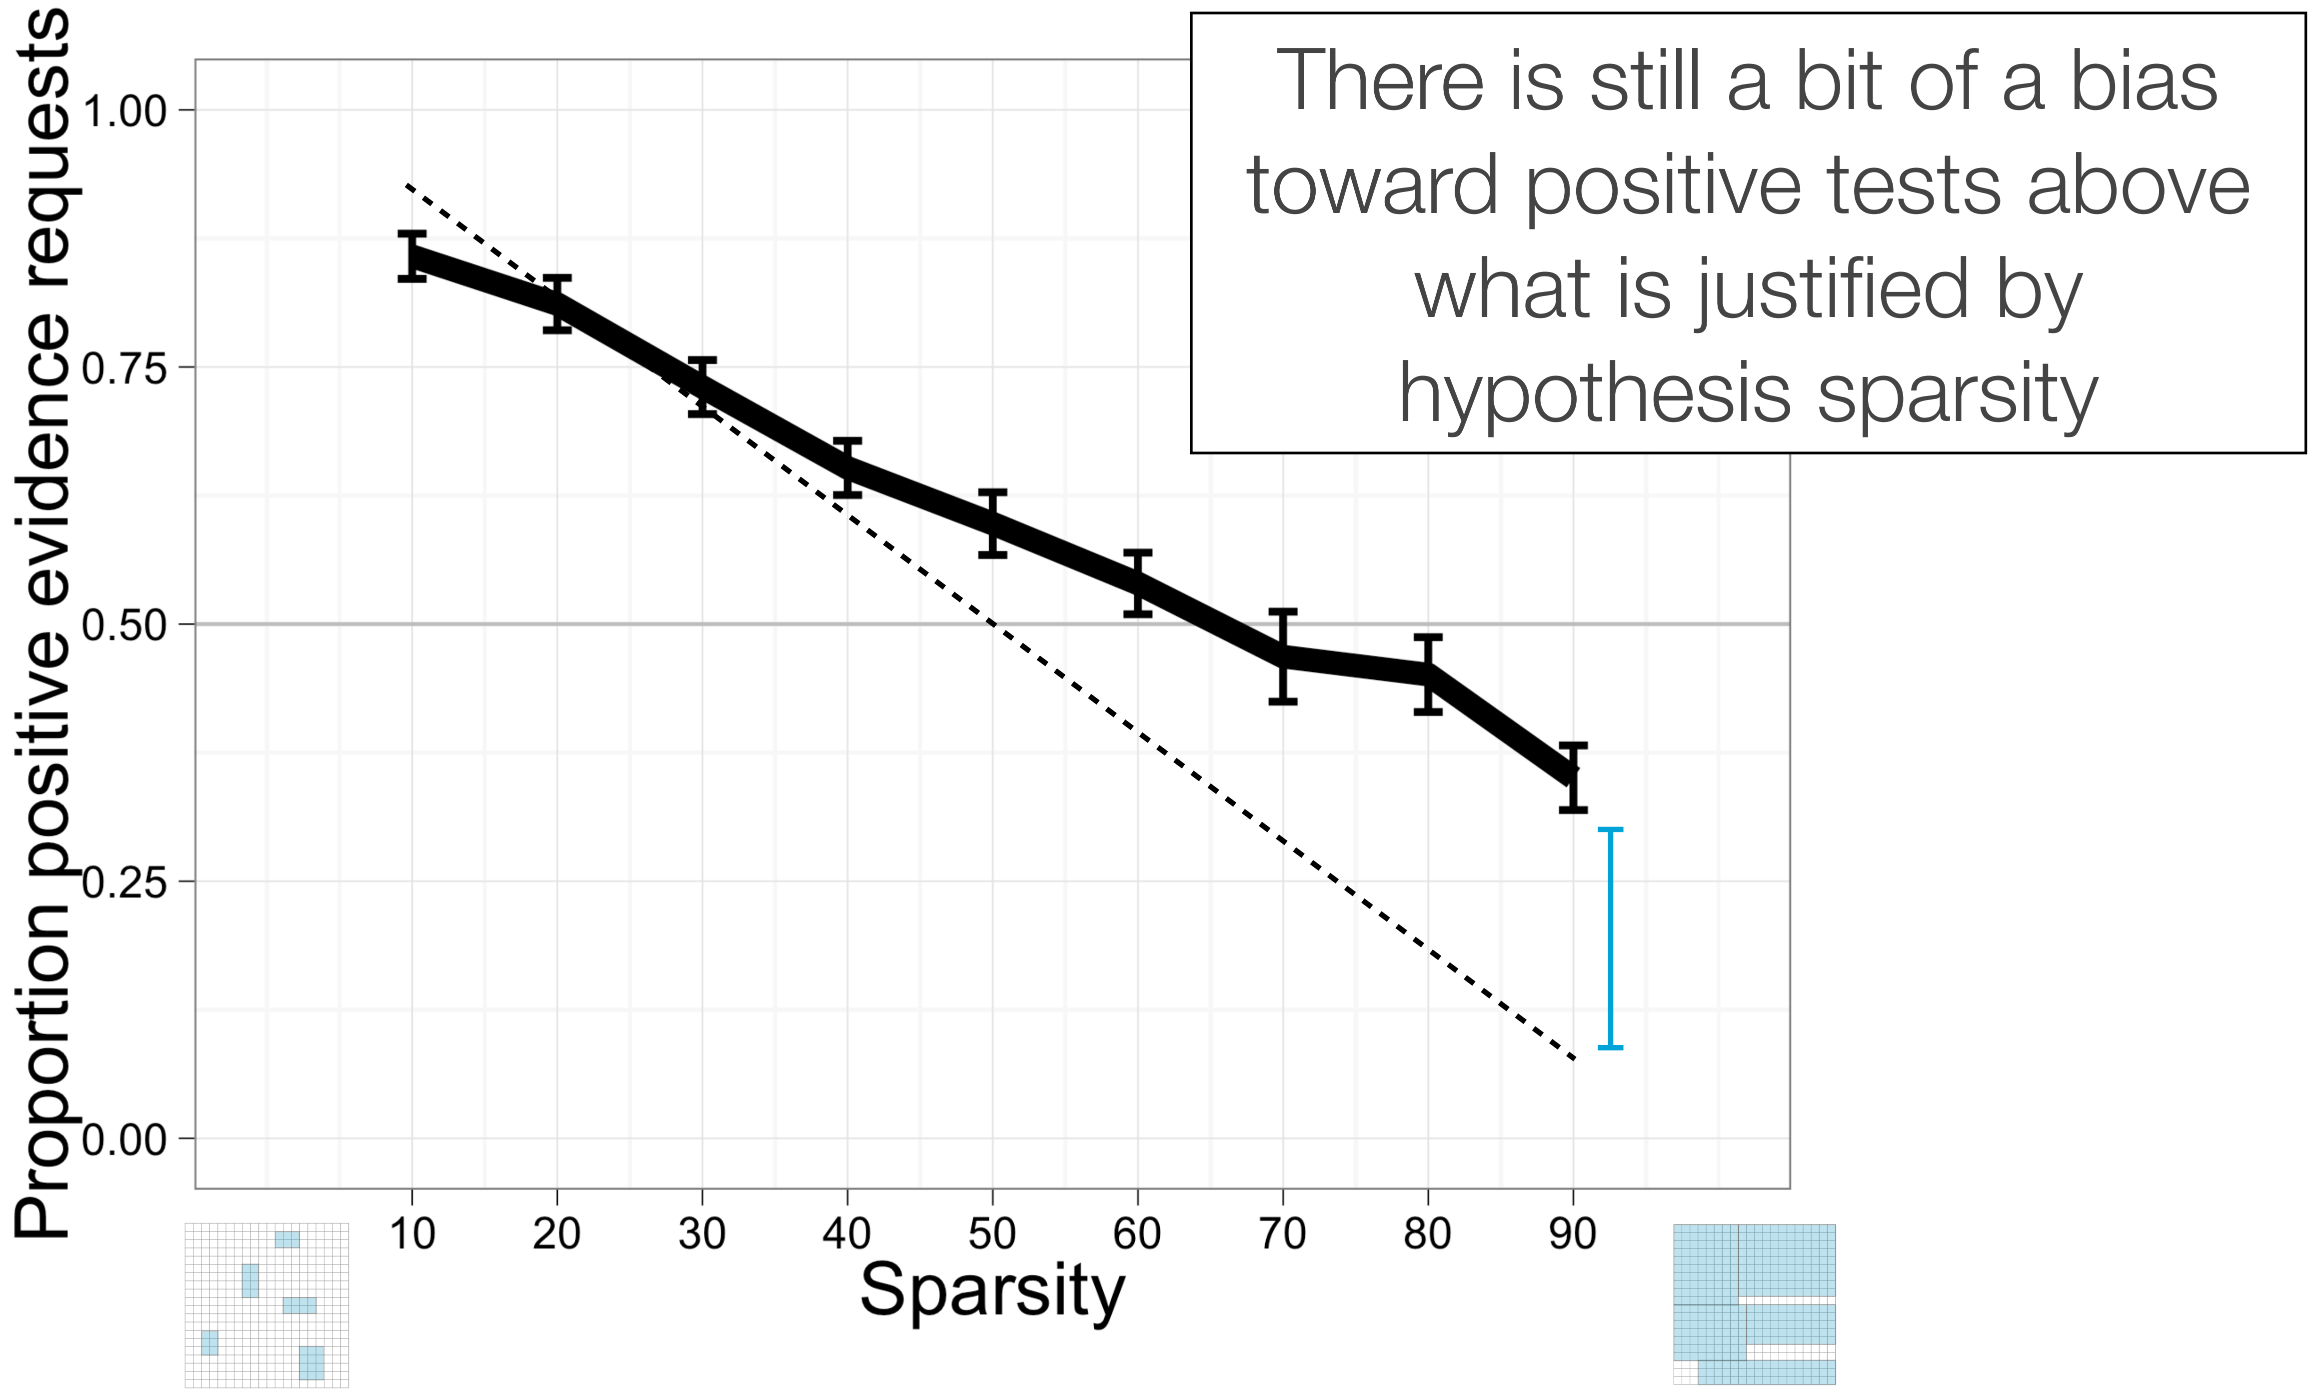

Do people change their strategy with sparsity?

yes; they ask for more hits (positive evidence) when the hypothesis space is sparse

Is this just a heuristic, or are people actually sensitive to the utility / informativeness of data in general?

Sensitivity to utility in general

- One way to test this is to see if people change what they ask for as they go on

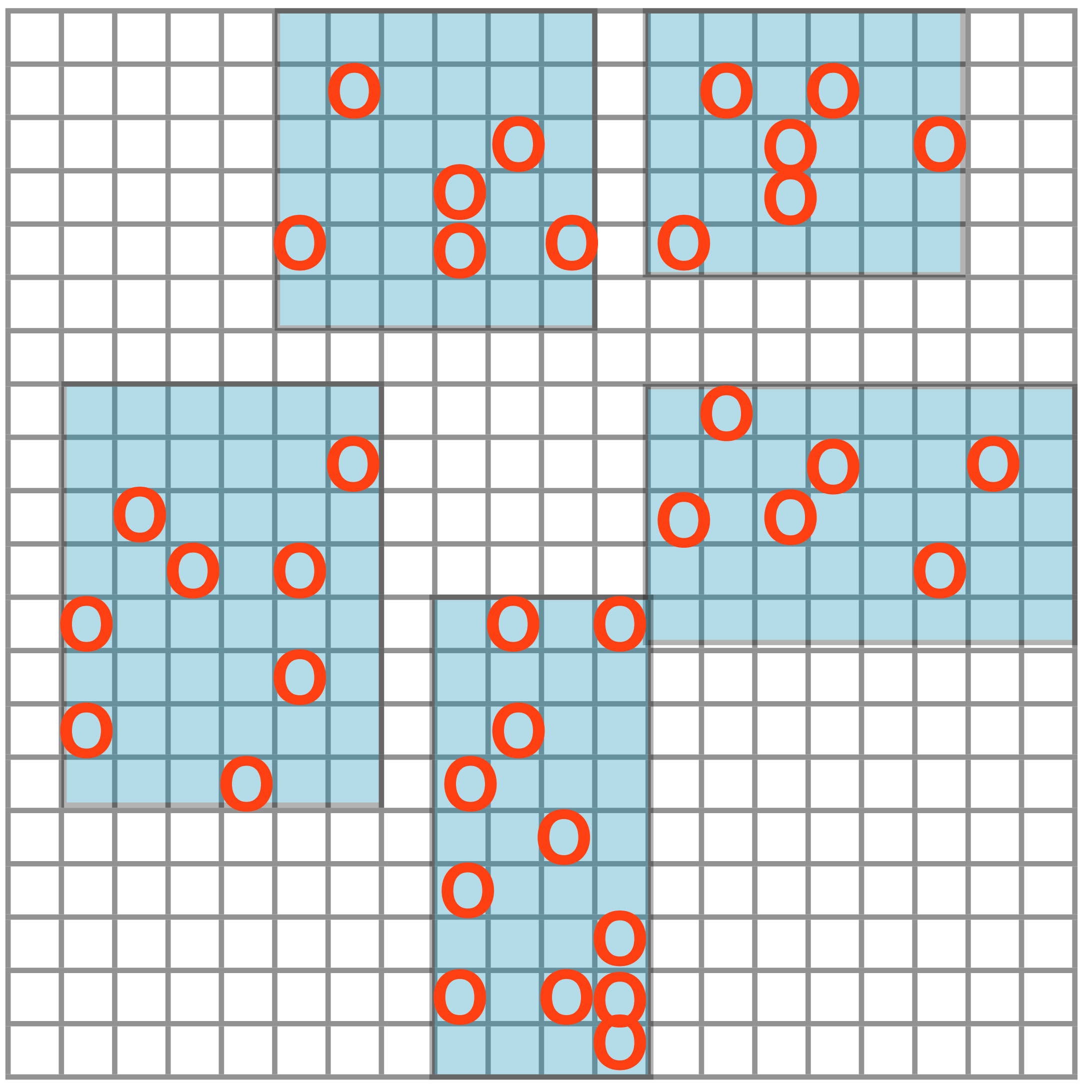

Additional data points change the sparsity of the remaining space

Sensitivity to utility in general

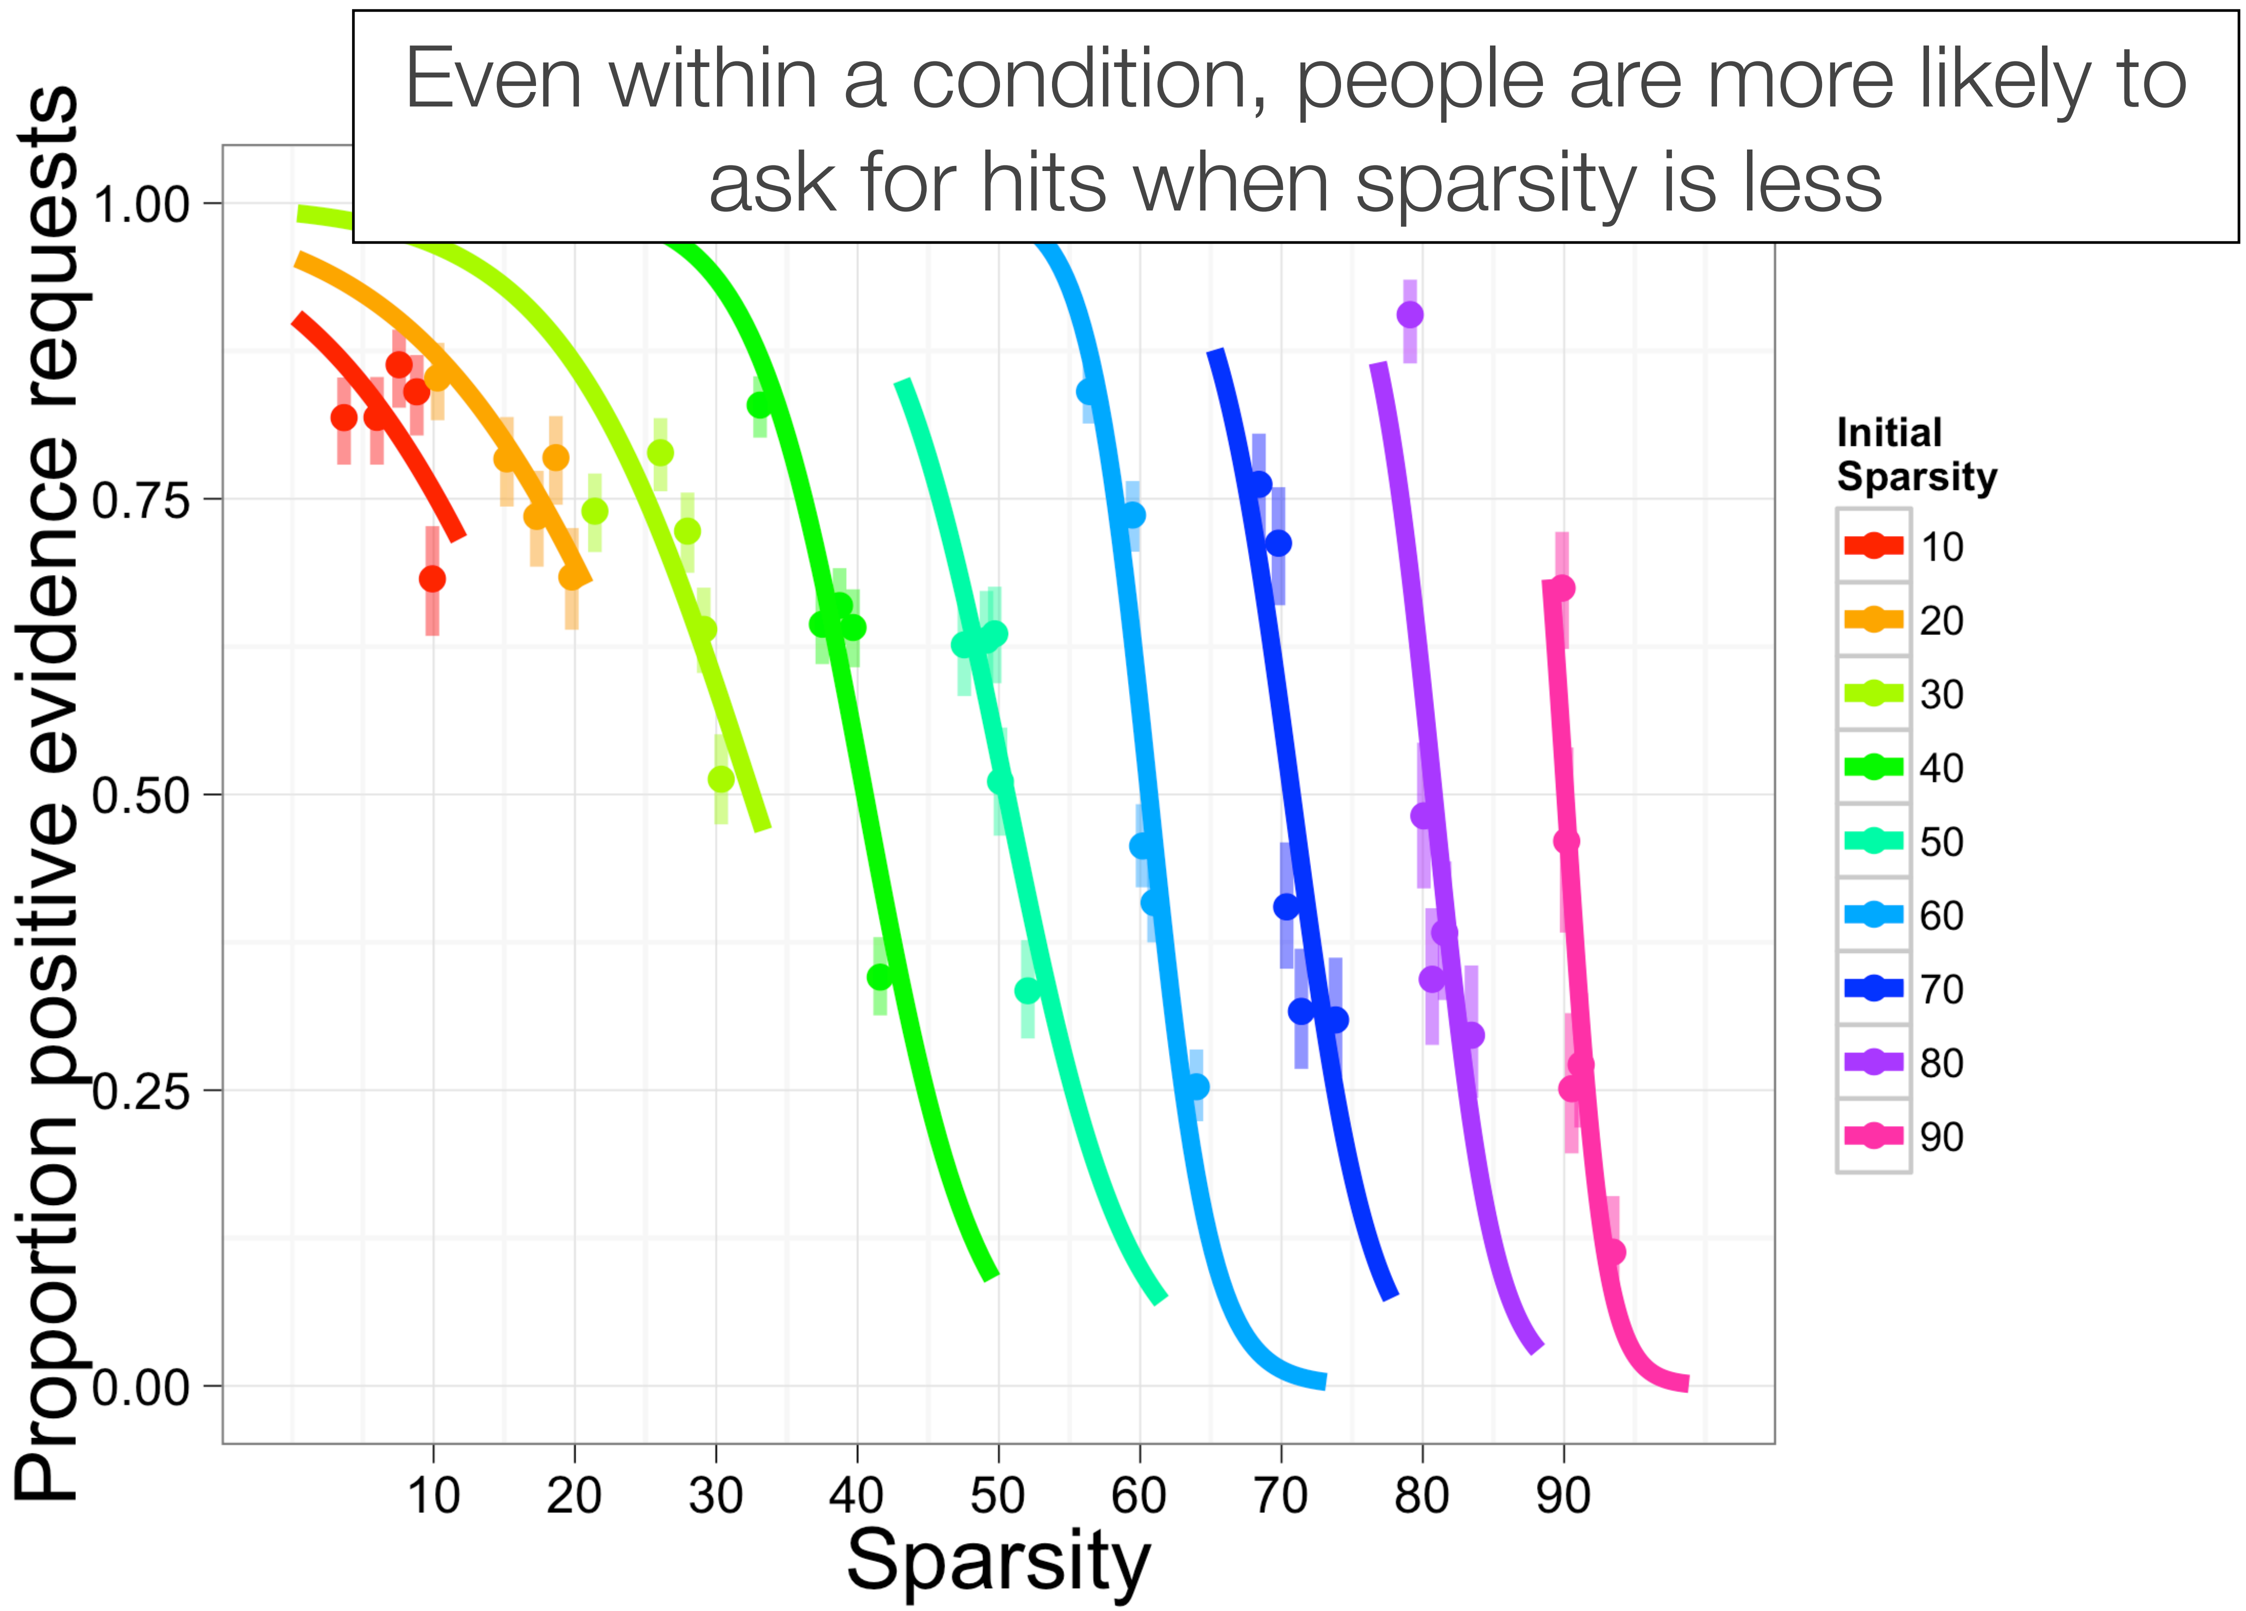

Sensitivity to utility in general

- Information requests are influenced by the sparsity of the hypotheses
  - initial sparsity
  - action-by-action changes in sparsity

- People seem to be making ‘reasonable’ requests

- This is just an indirect measure of sensitivity to information -- what if we let them ask for specific points directly, and look at how valuable they are?

Sensitivity to utility in general

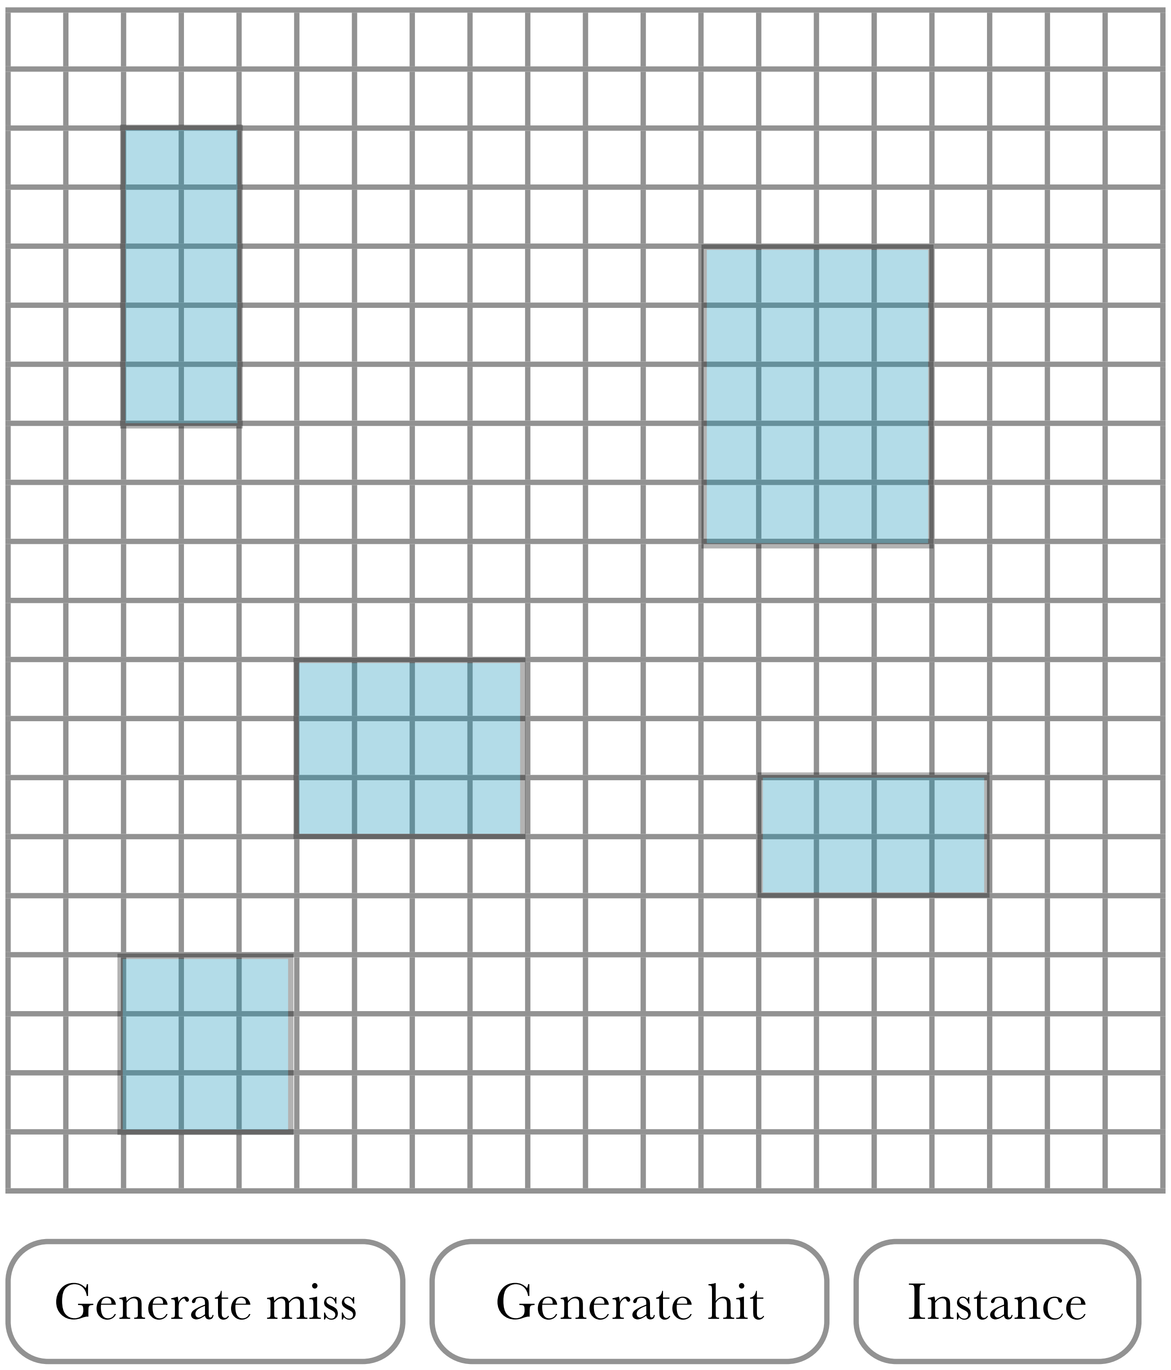

Sensitivity to utility in general

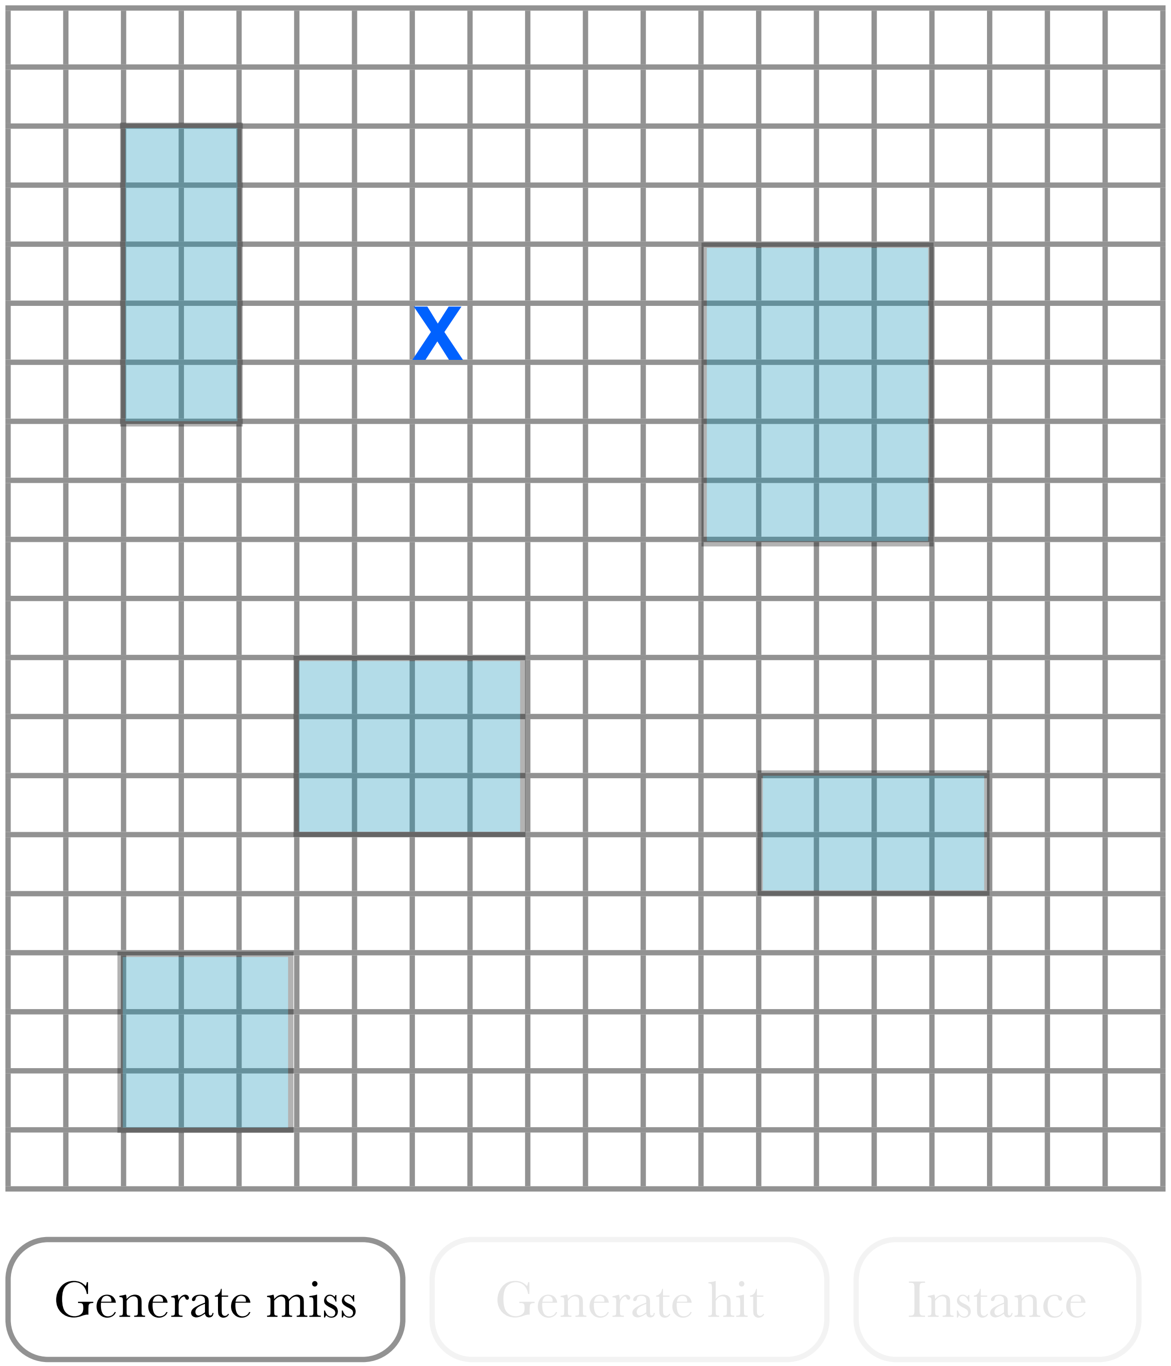

Sensitivity to utility in general

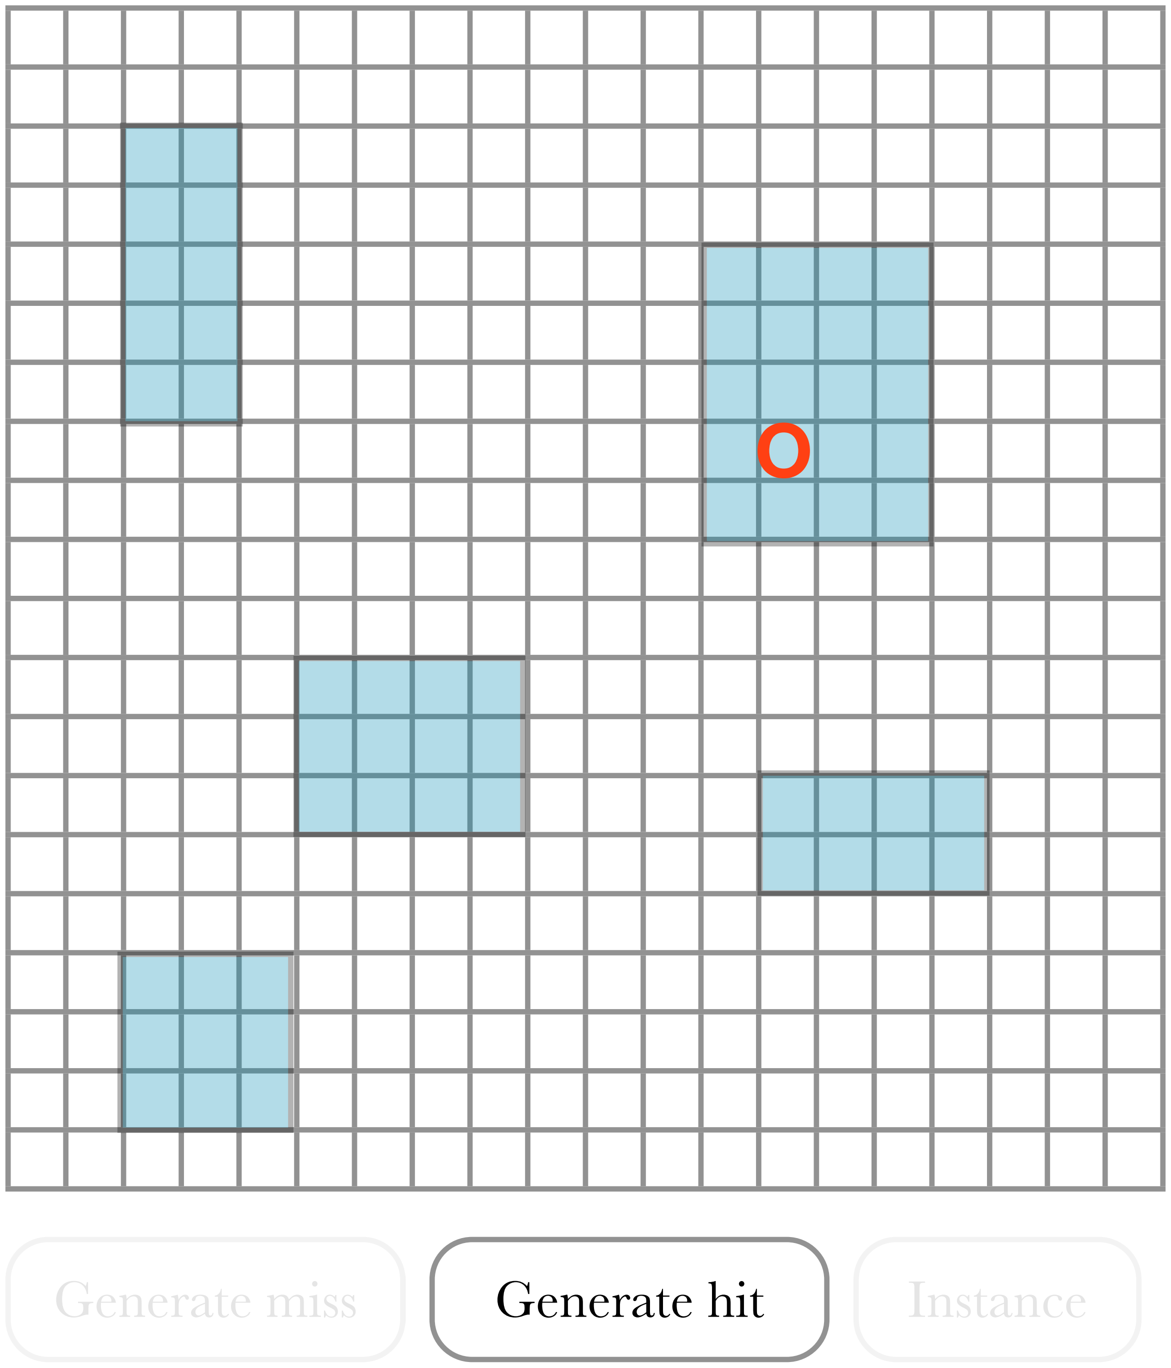

Sensitivity to utility in general

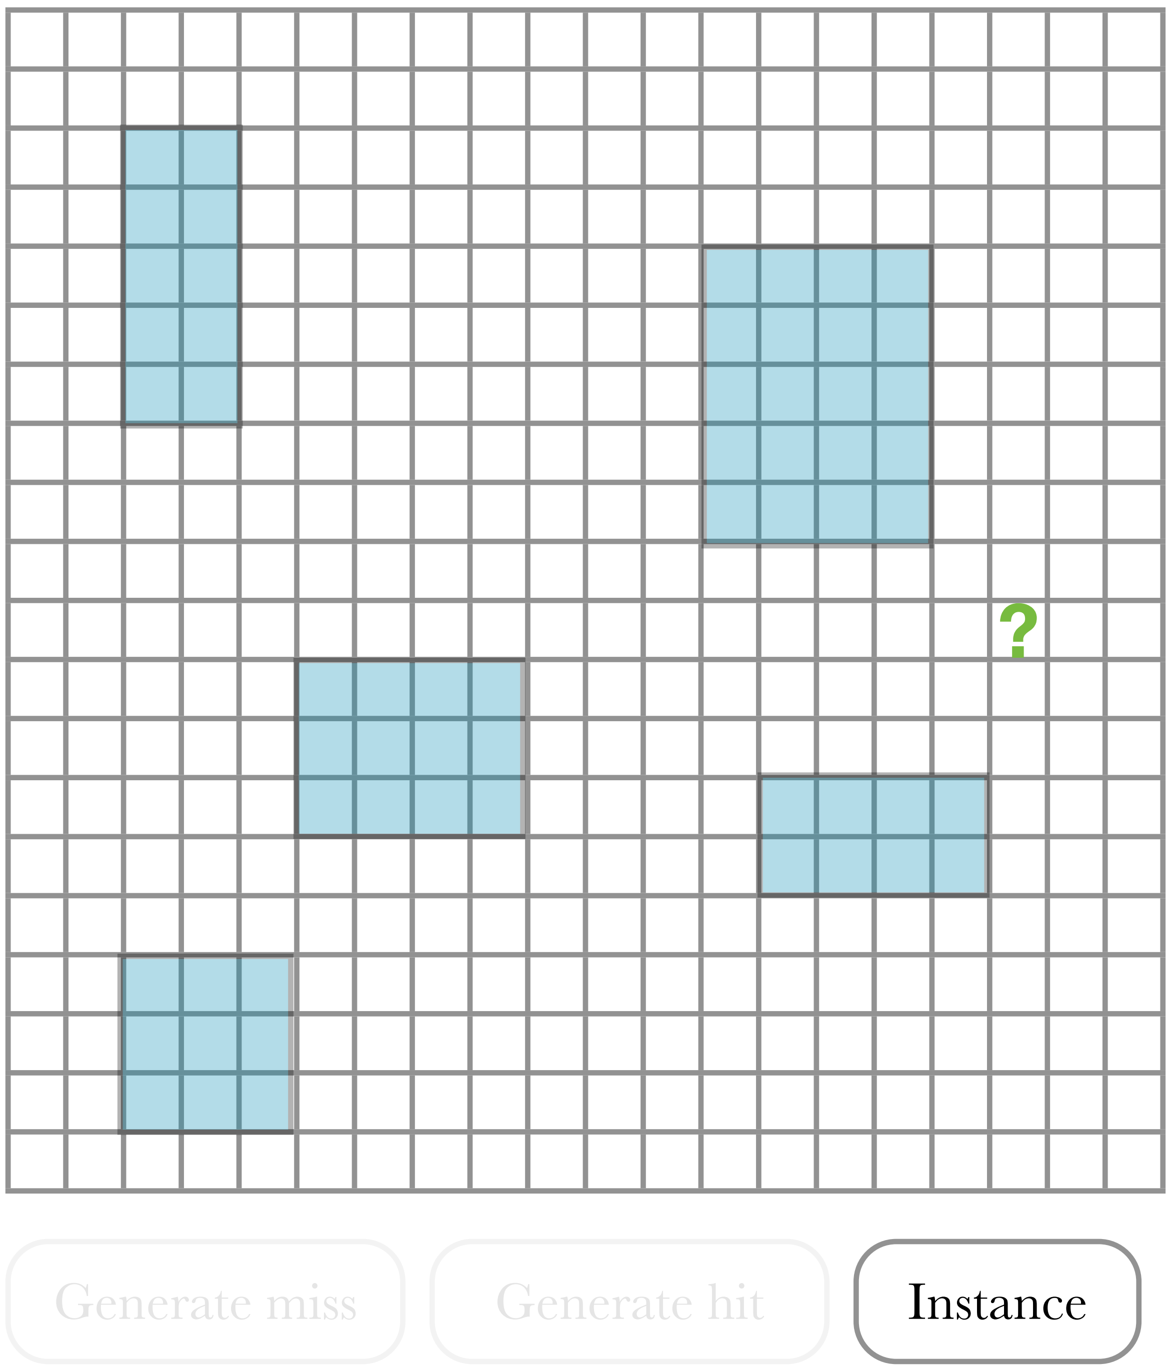

Sensitivity to utility in general

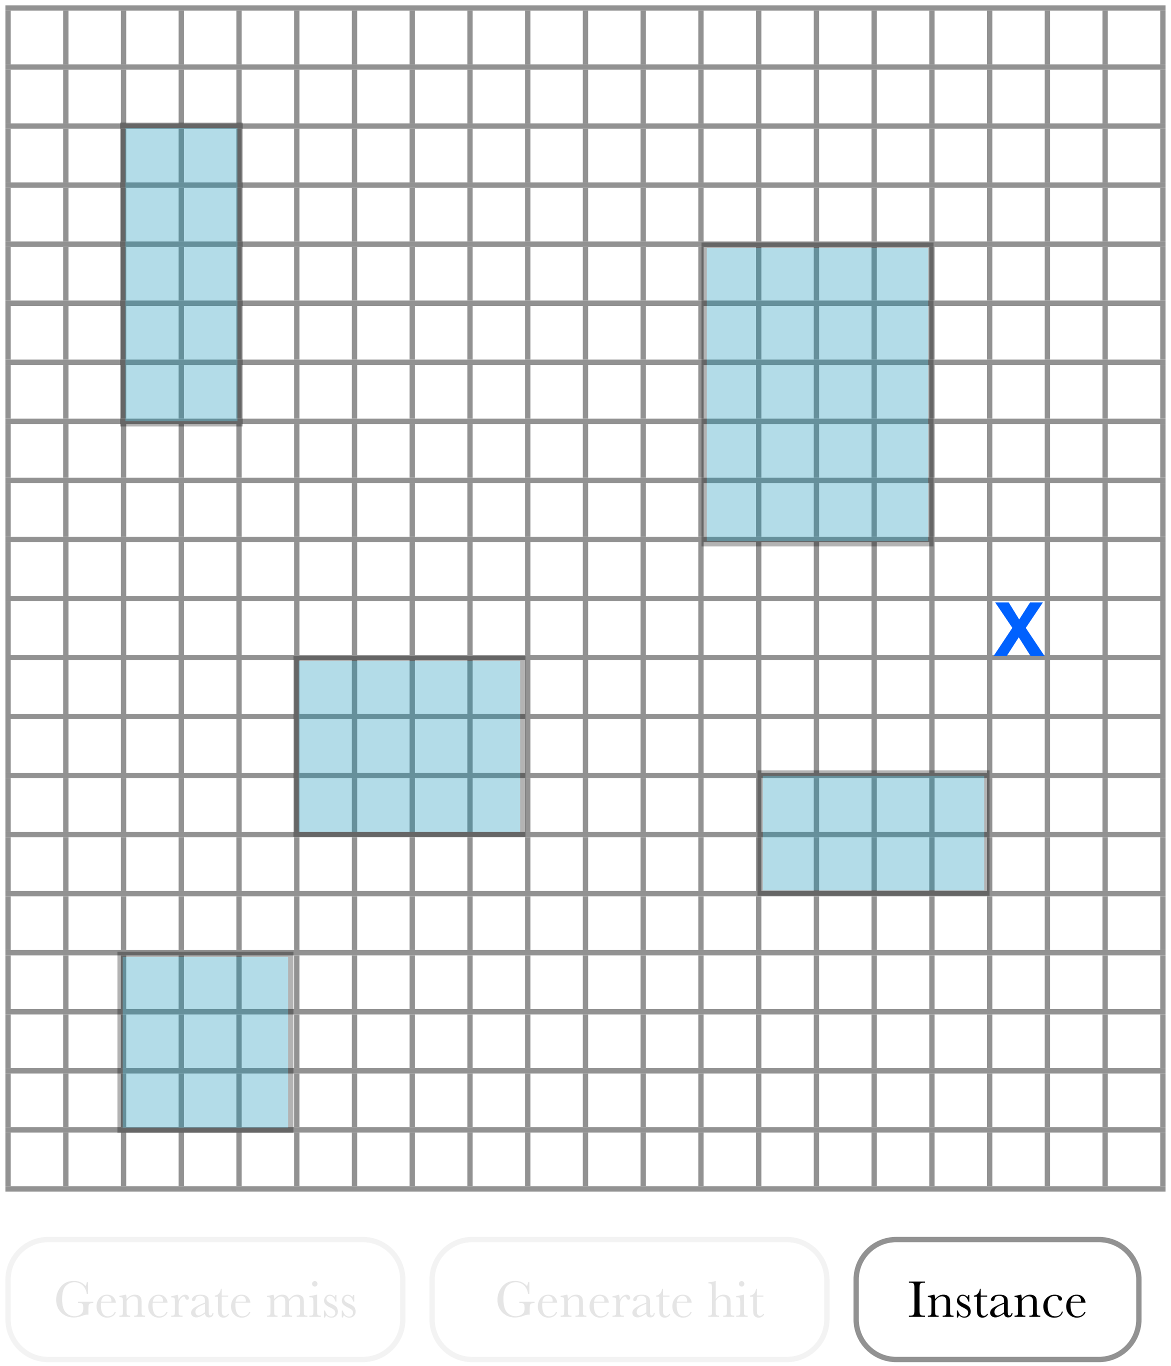

Sensitivity to utility in general

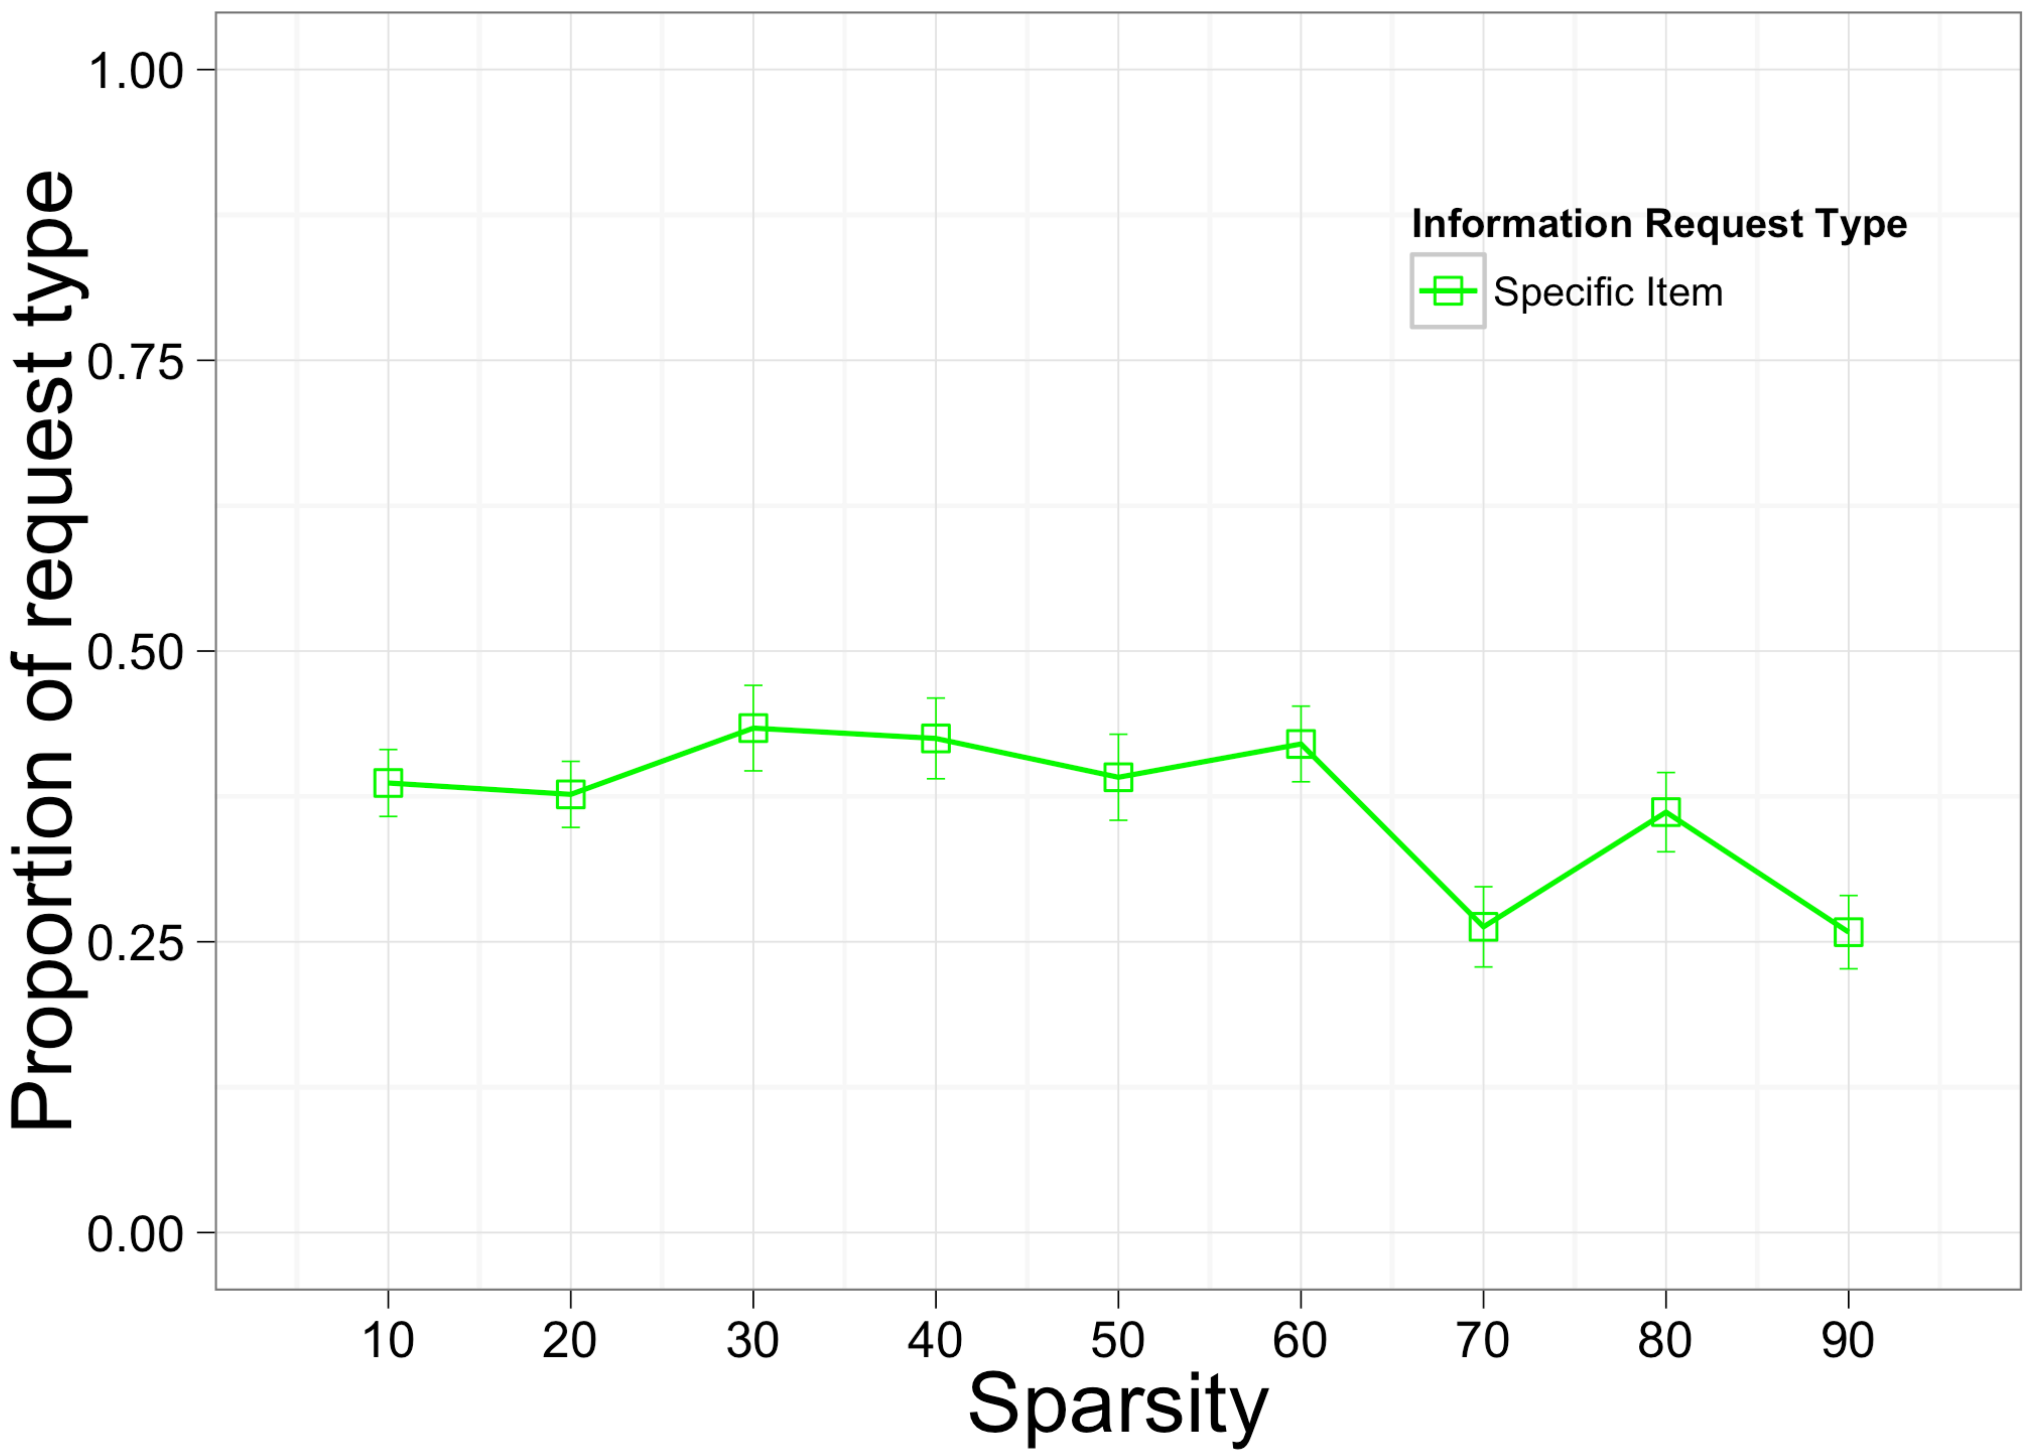

Sensitivity to utility in general

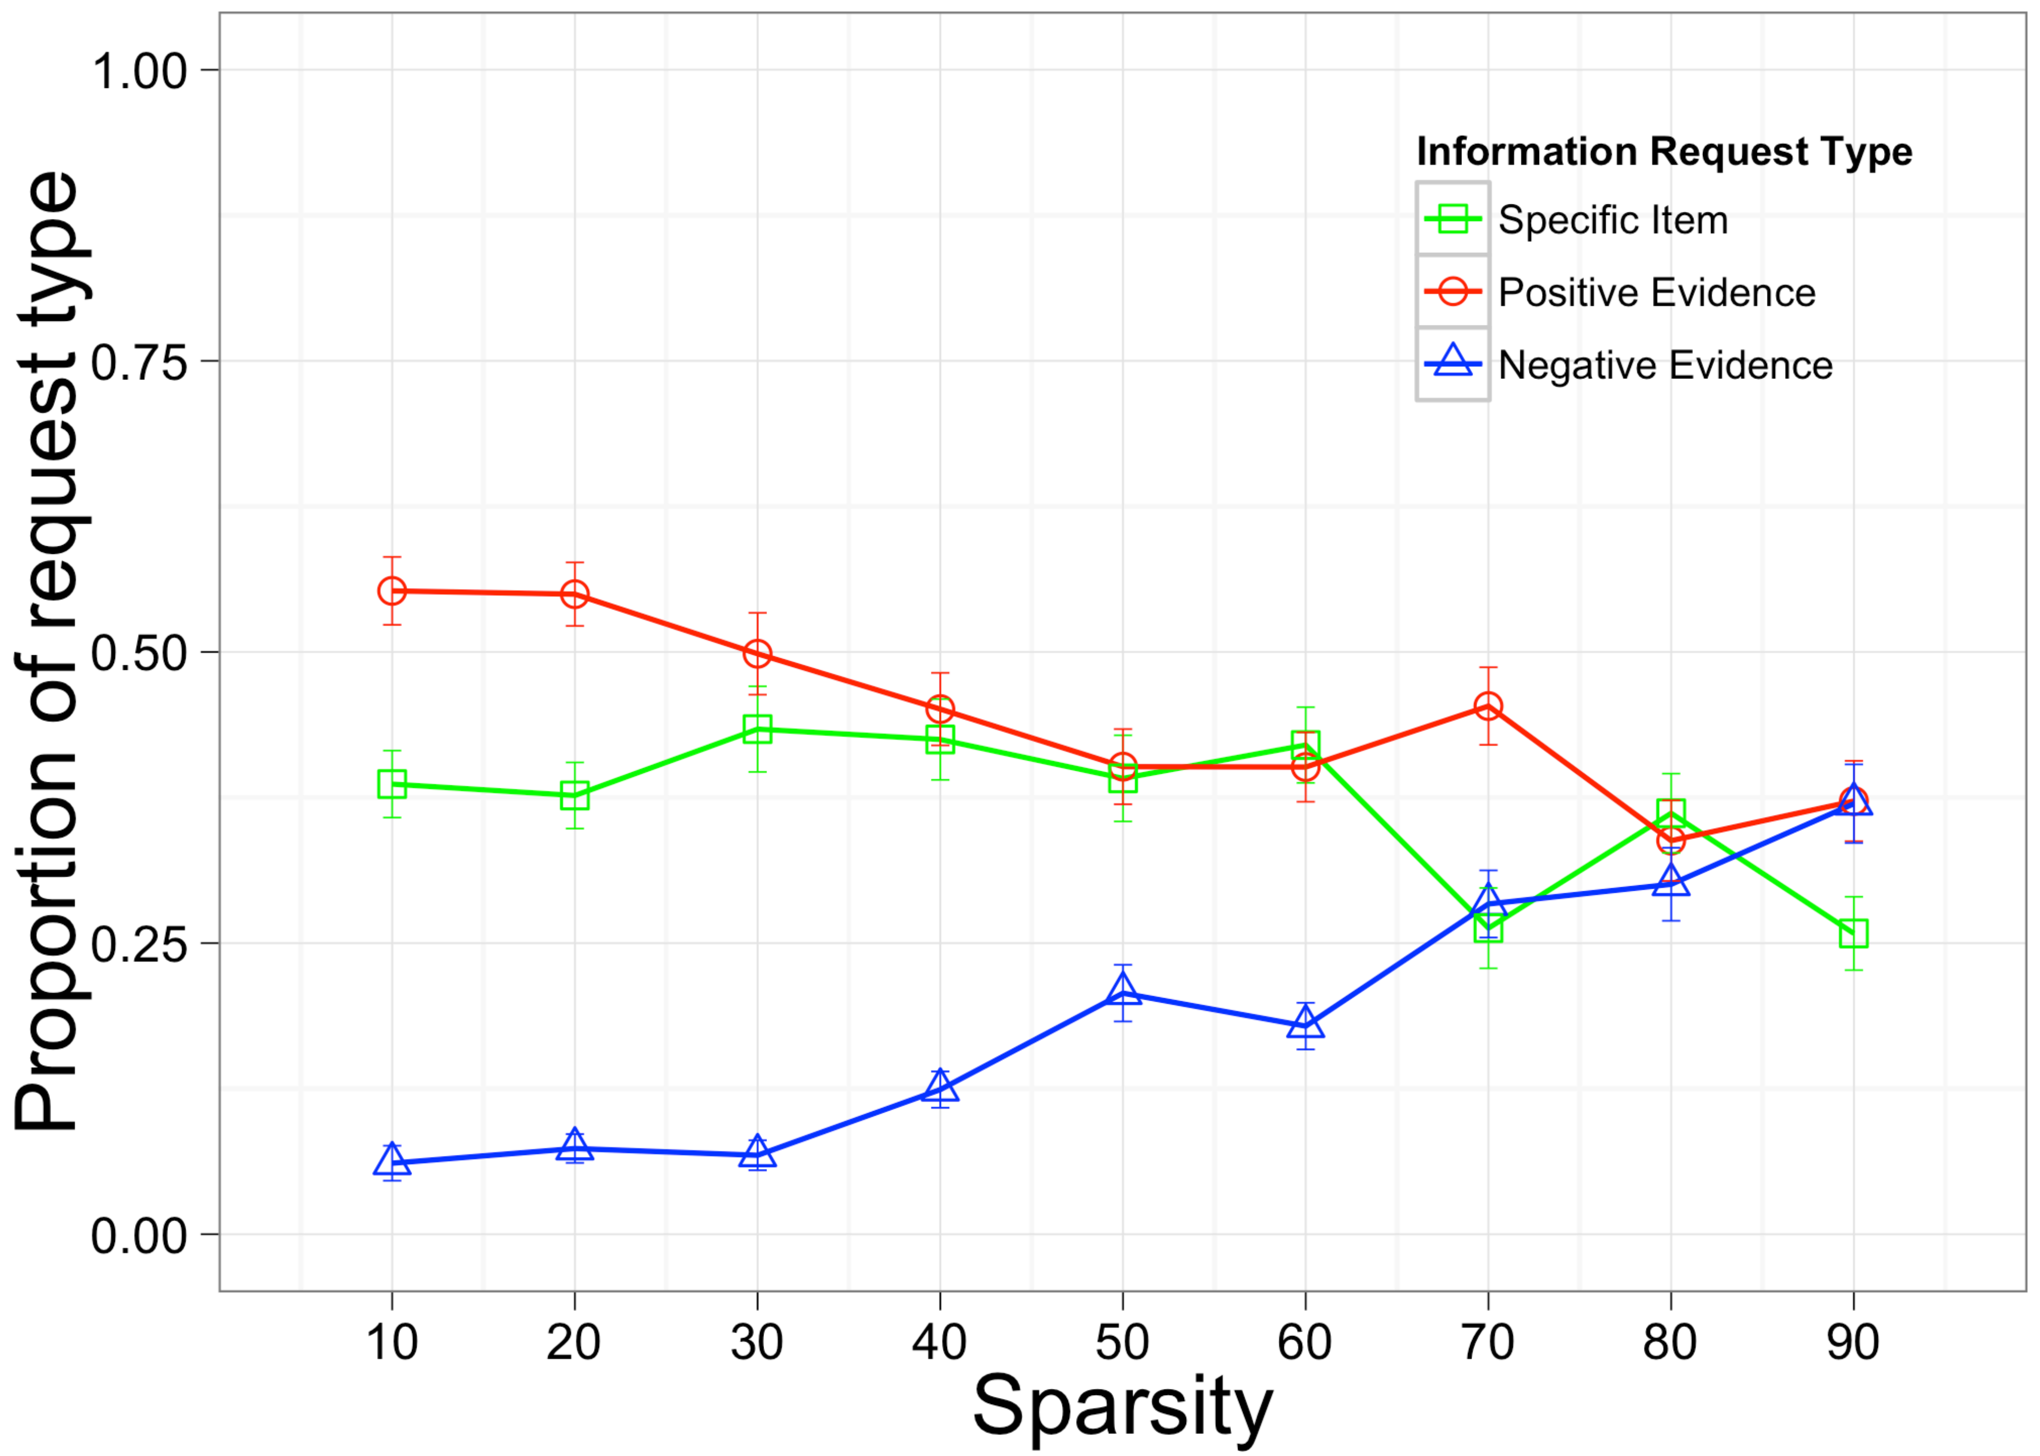

Sensitivity to utility in general

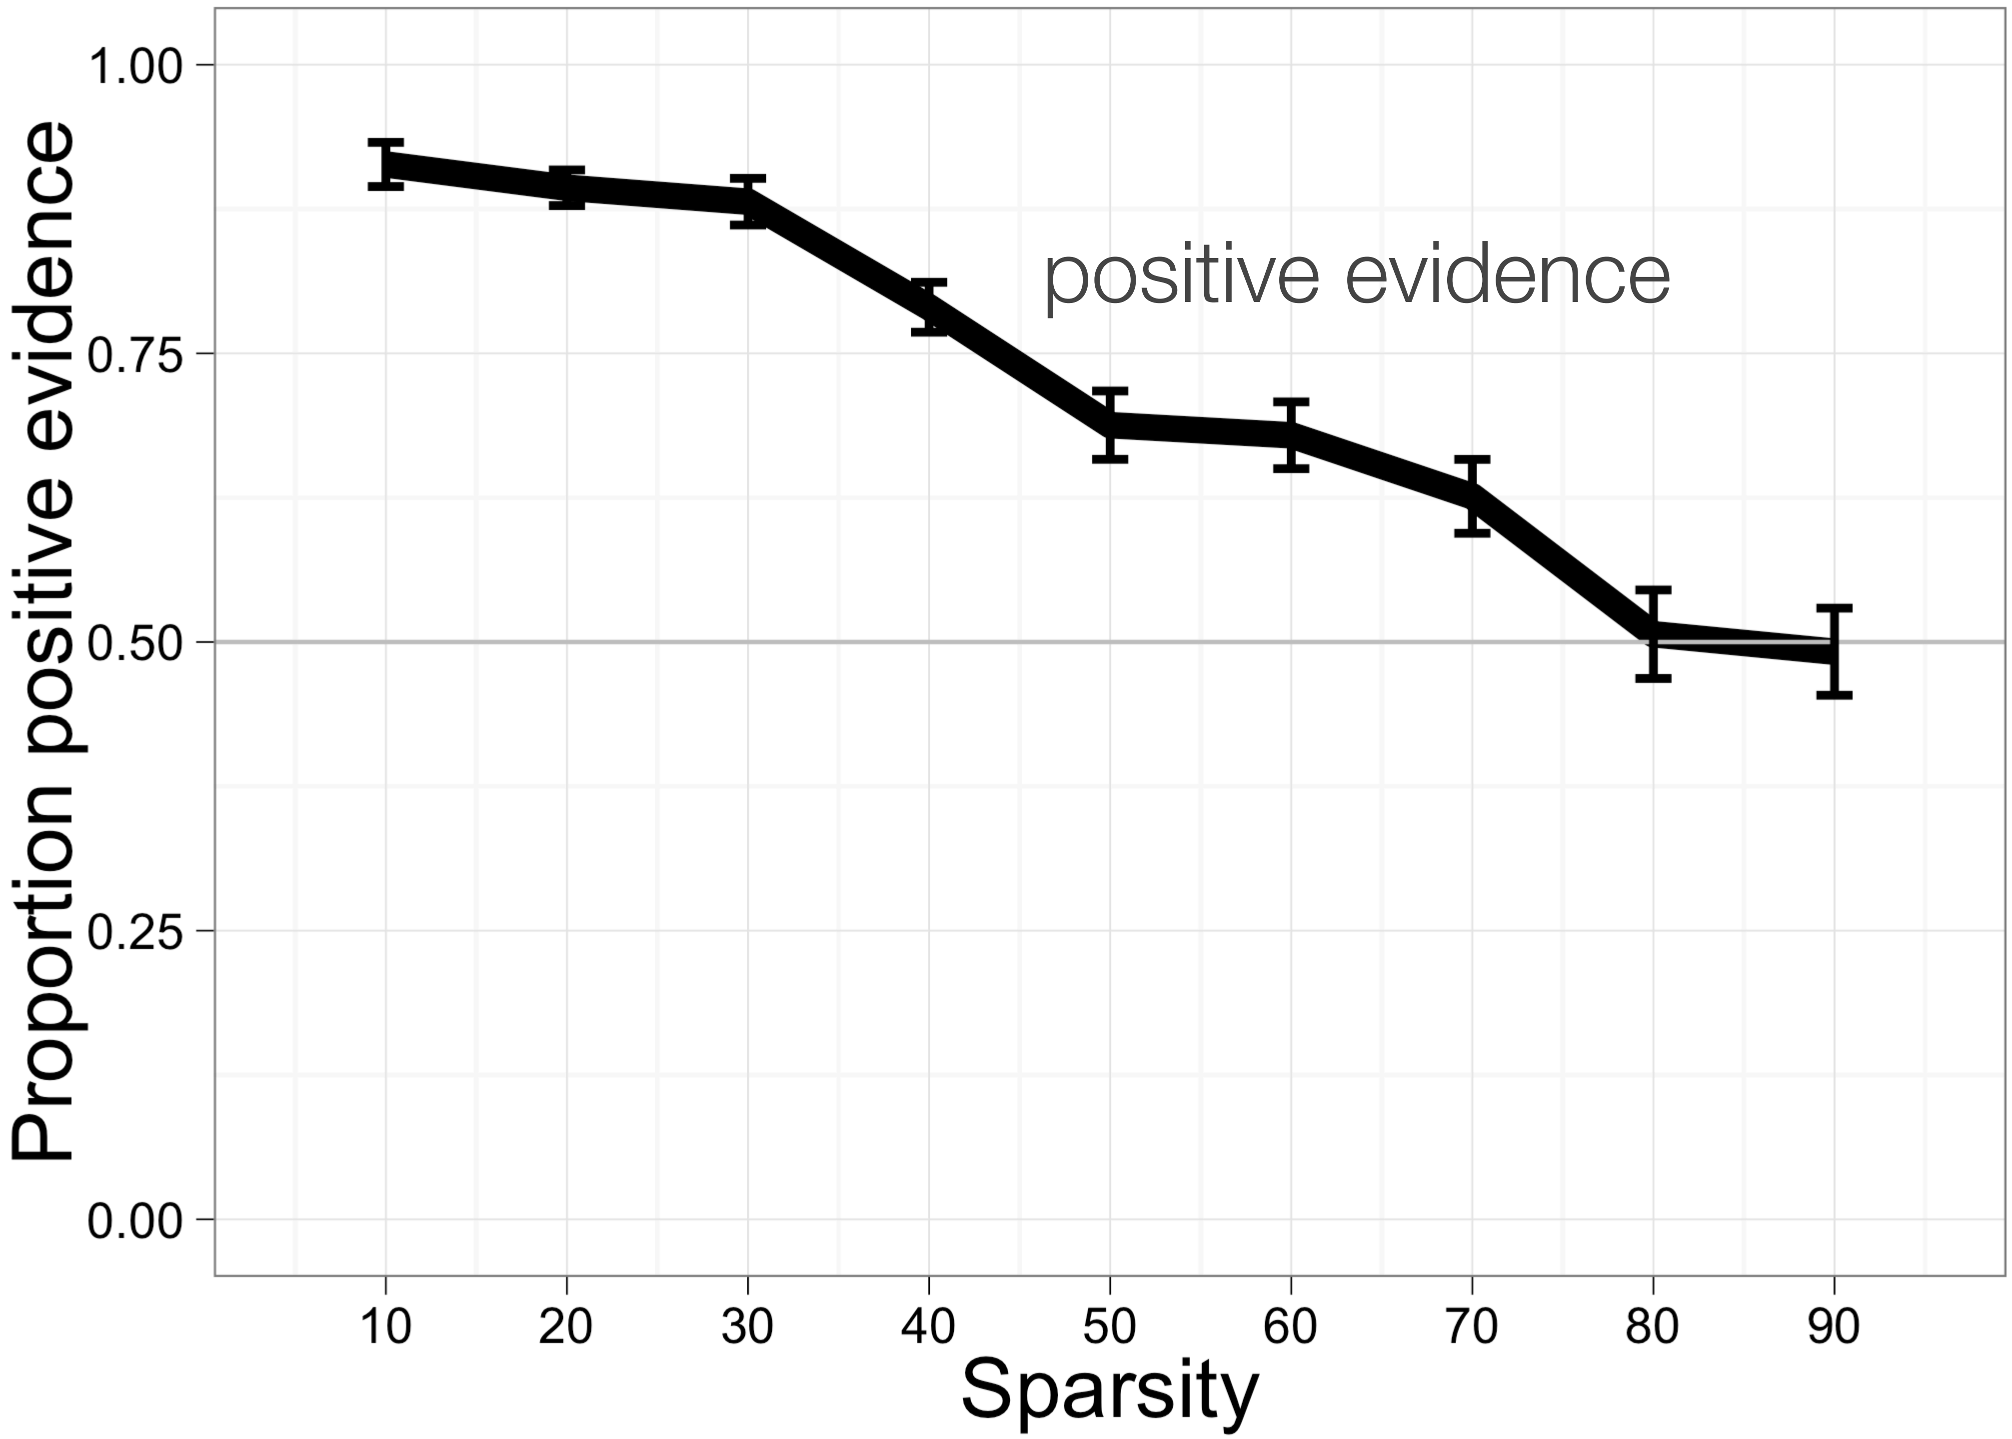

Sensitivity to utility in general

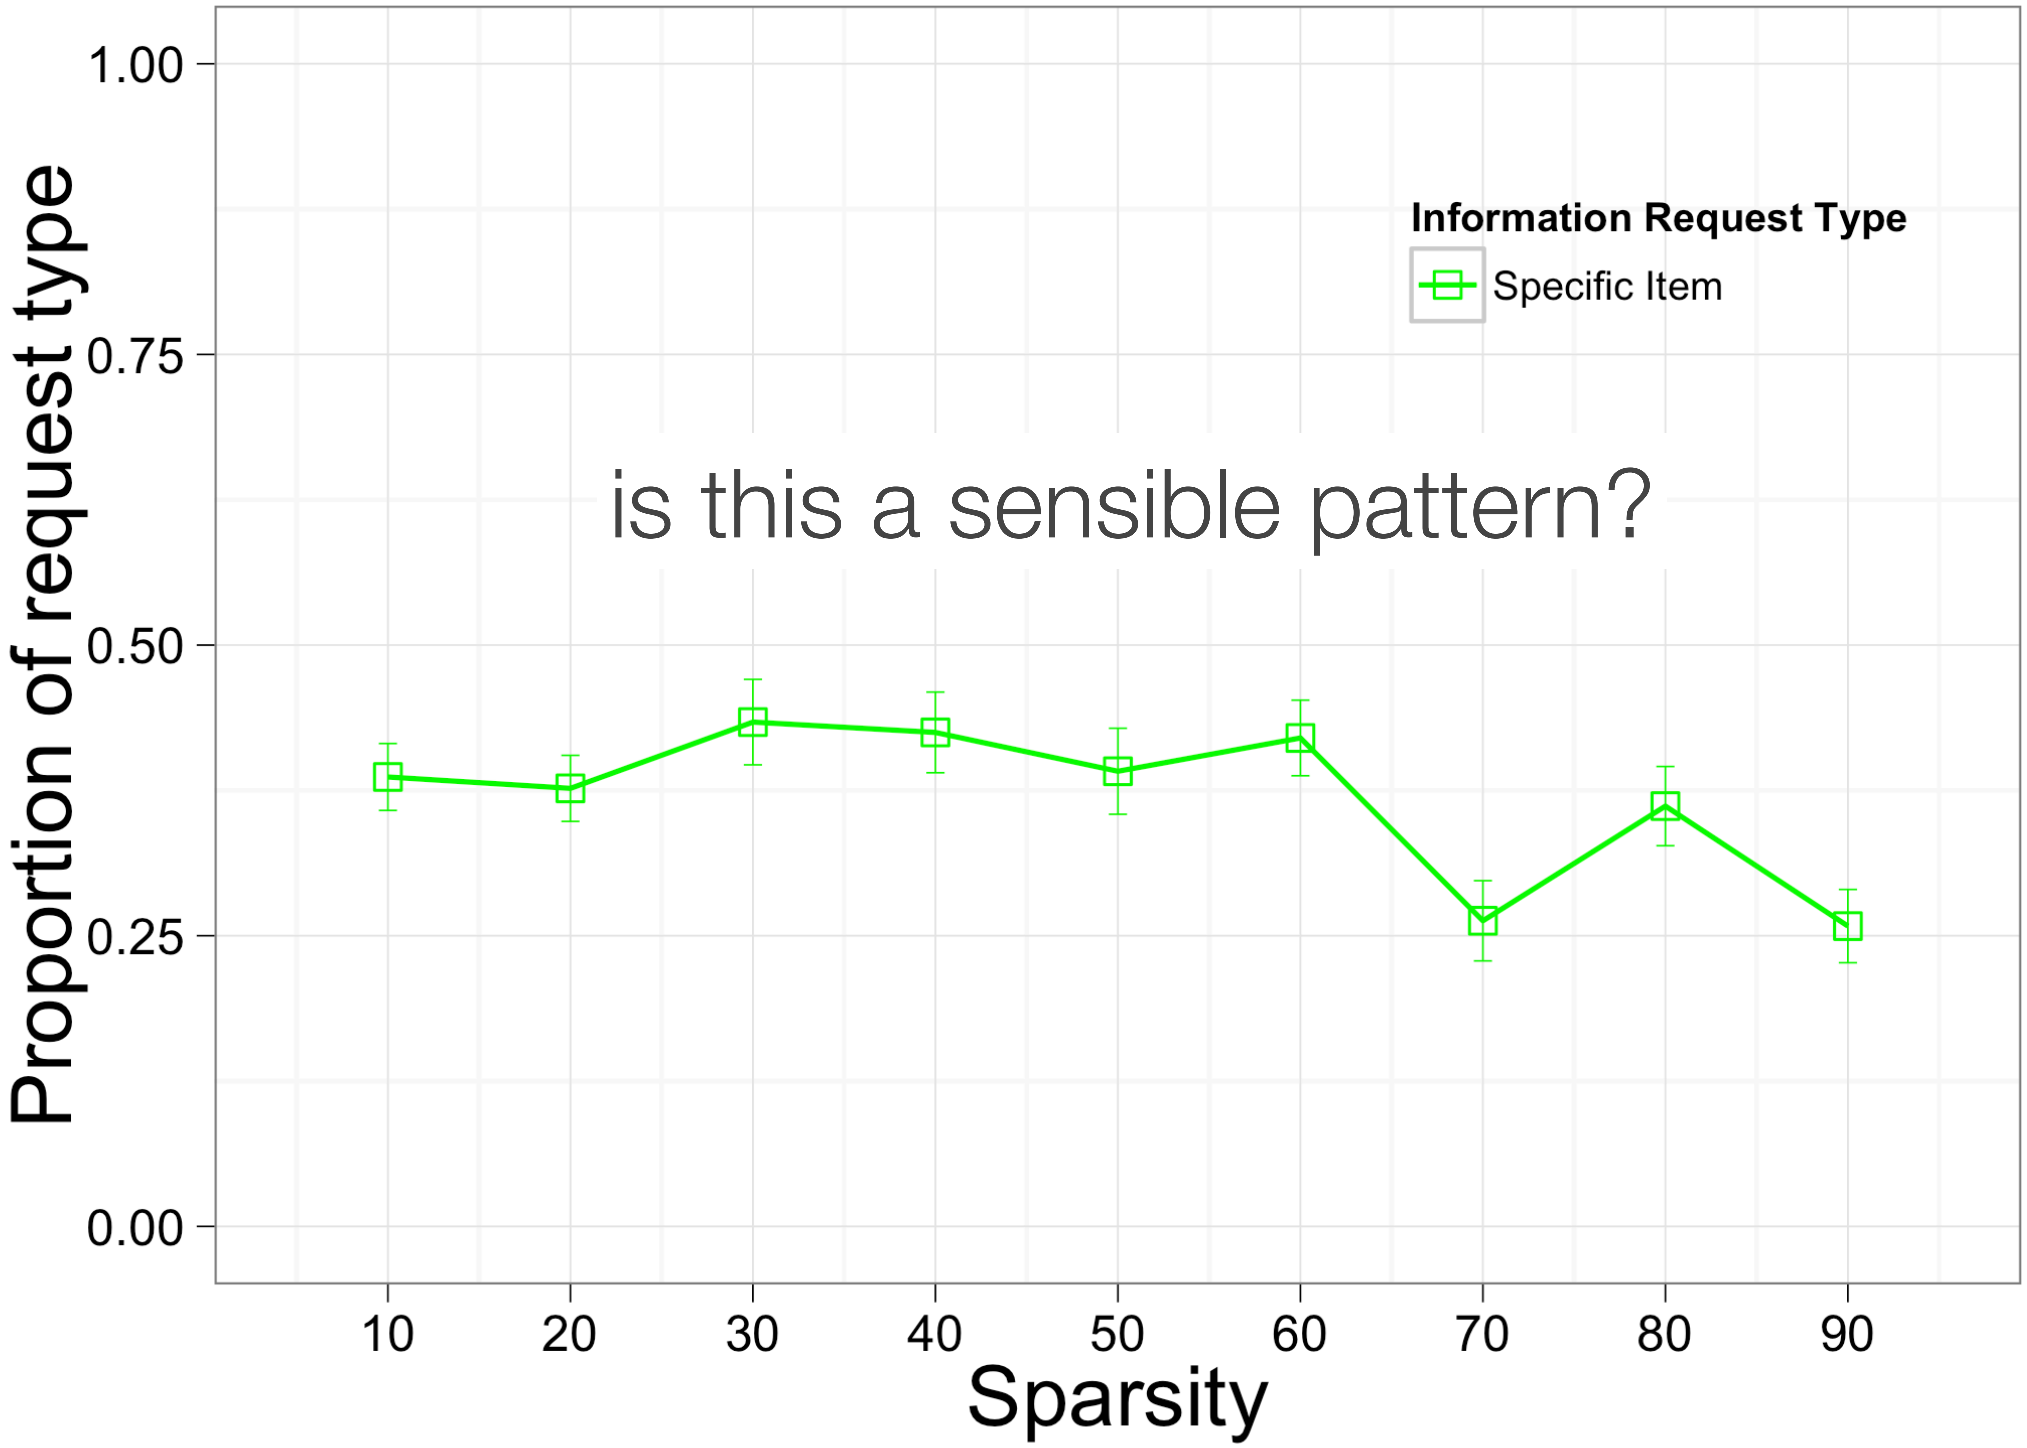

Choosing useful options

Do people choose ‘good’ request options?

1. Are some requests better than others?

2. Are people more likely to select more useful request options?

Choosing useful options

Do people choose ‘good’ request options?

➡ 1. Are some requests better than others?

2. Are people more likely to select more useful request options?

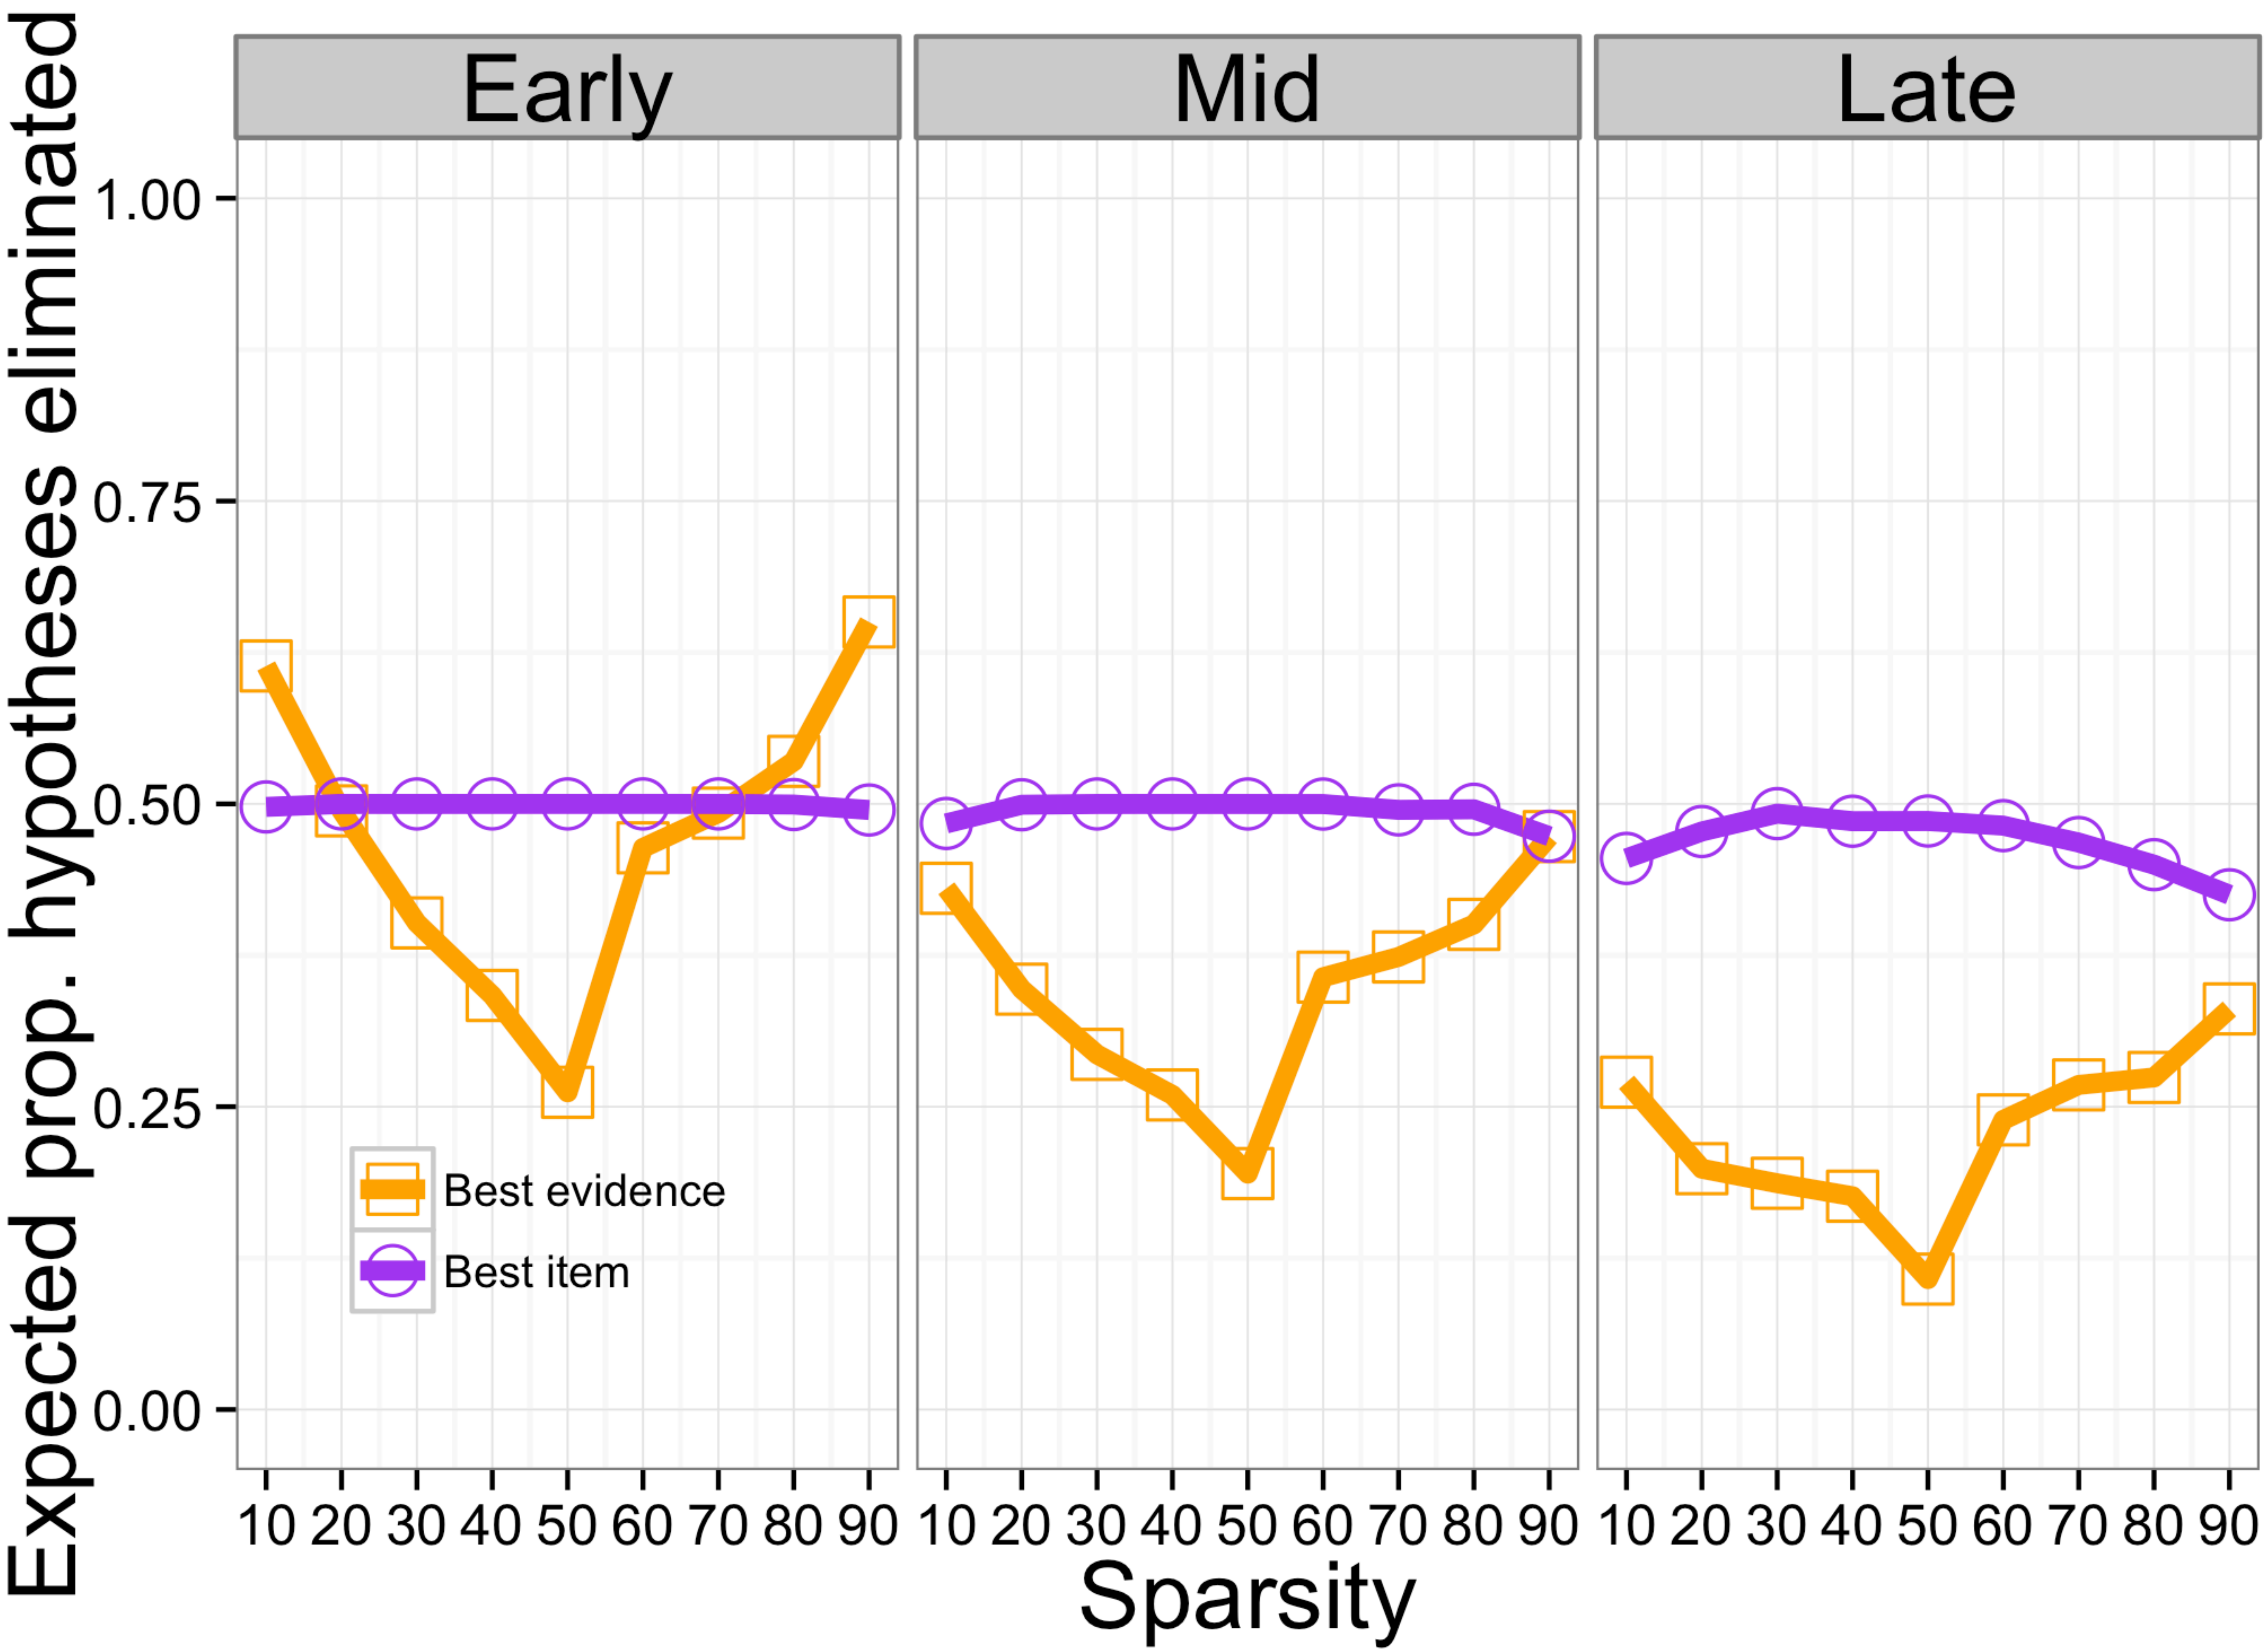

Early
Mid
Late
Expected prop. hypotheses eliminated
1.00
0.75
0.50
0.25
0.00
Best evidence
Best item
10 20 30 40 50 60 70 80 90
10 20 30 40 50 60 70 80 90
10 20 30 40 50 60 70 80 90
Sparsity

Choosing useful options

Do people choose ‘good’ request options?

1. Are some requests better than others?
YES

➡ 2. Are people more likely to select more useful request options?

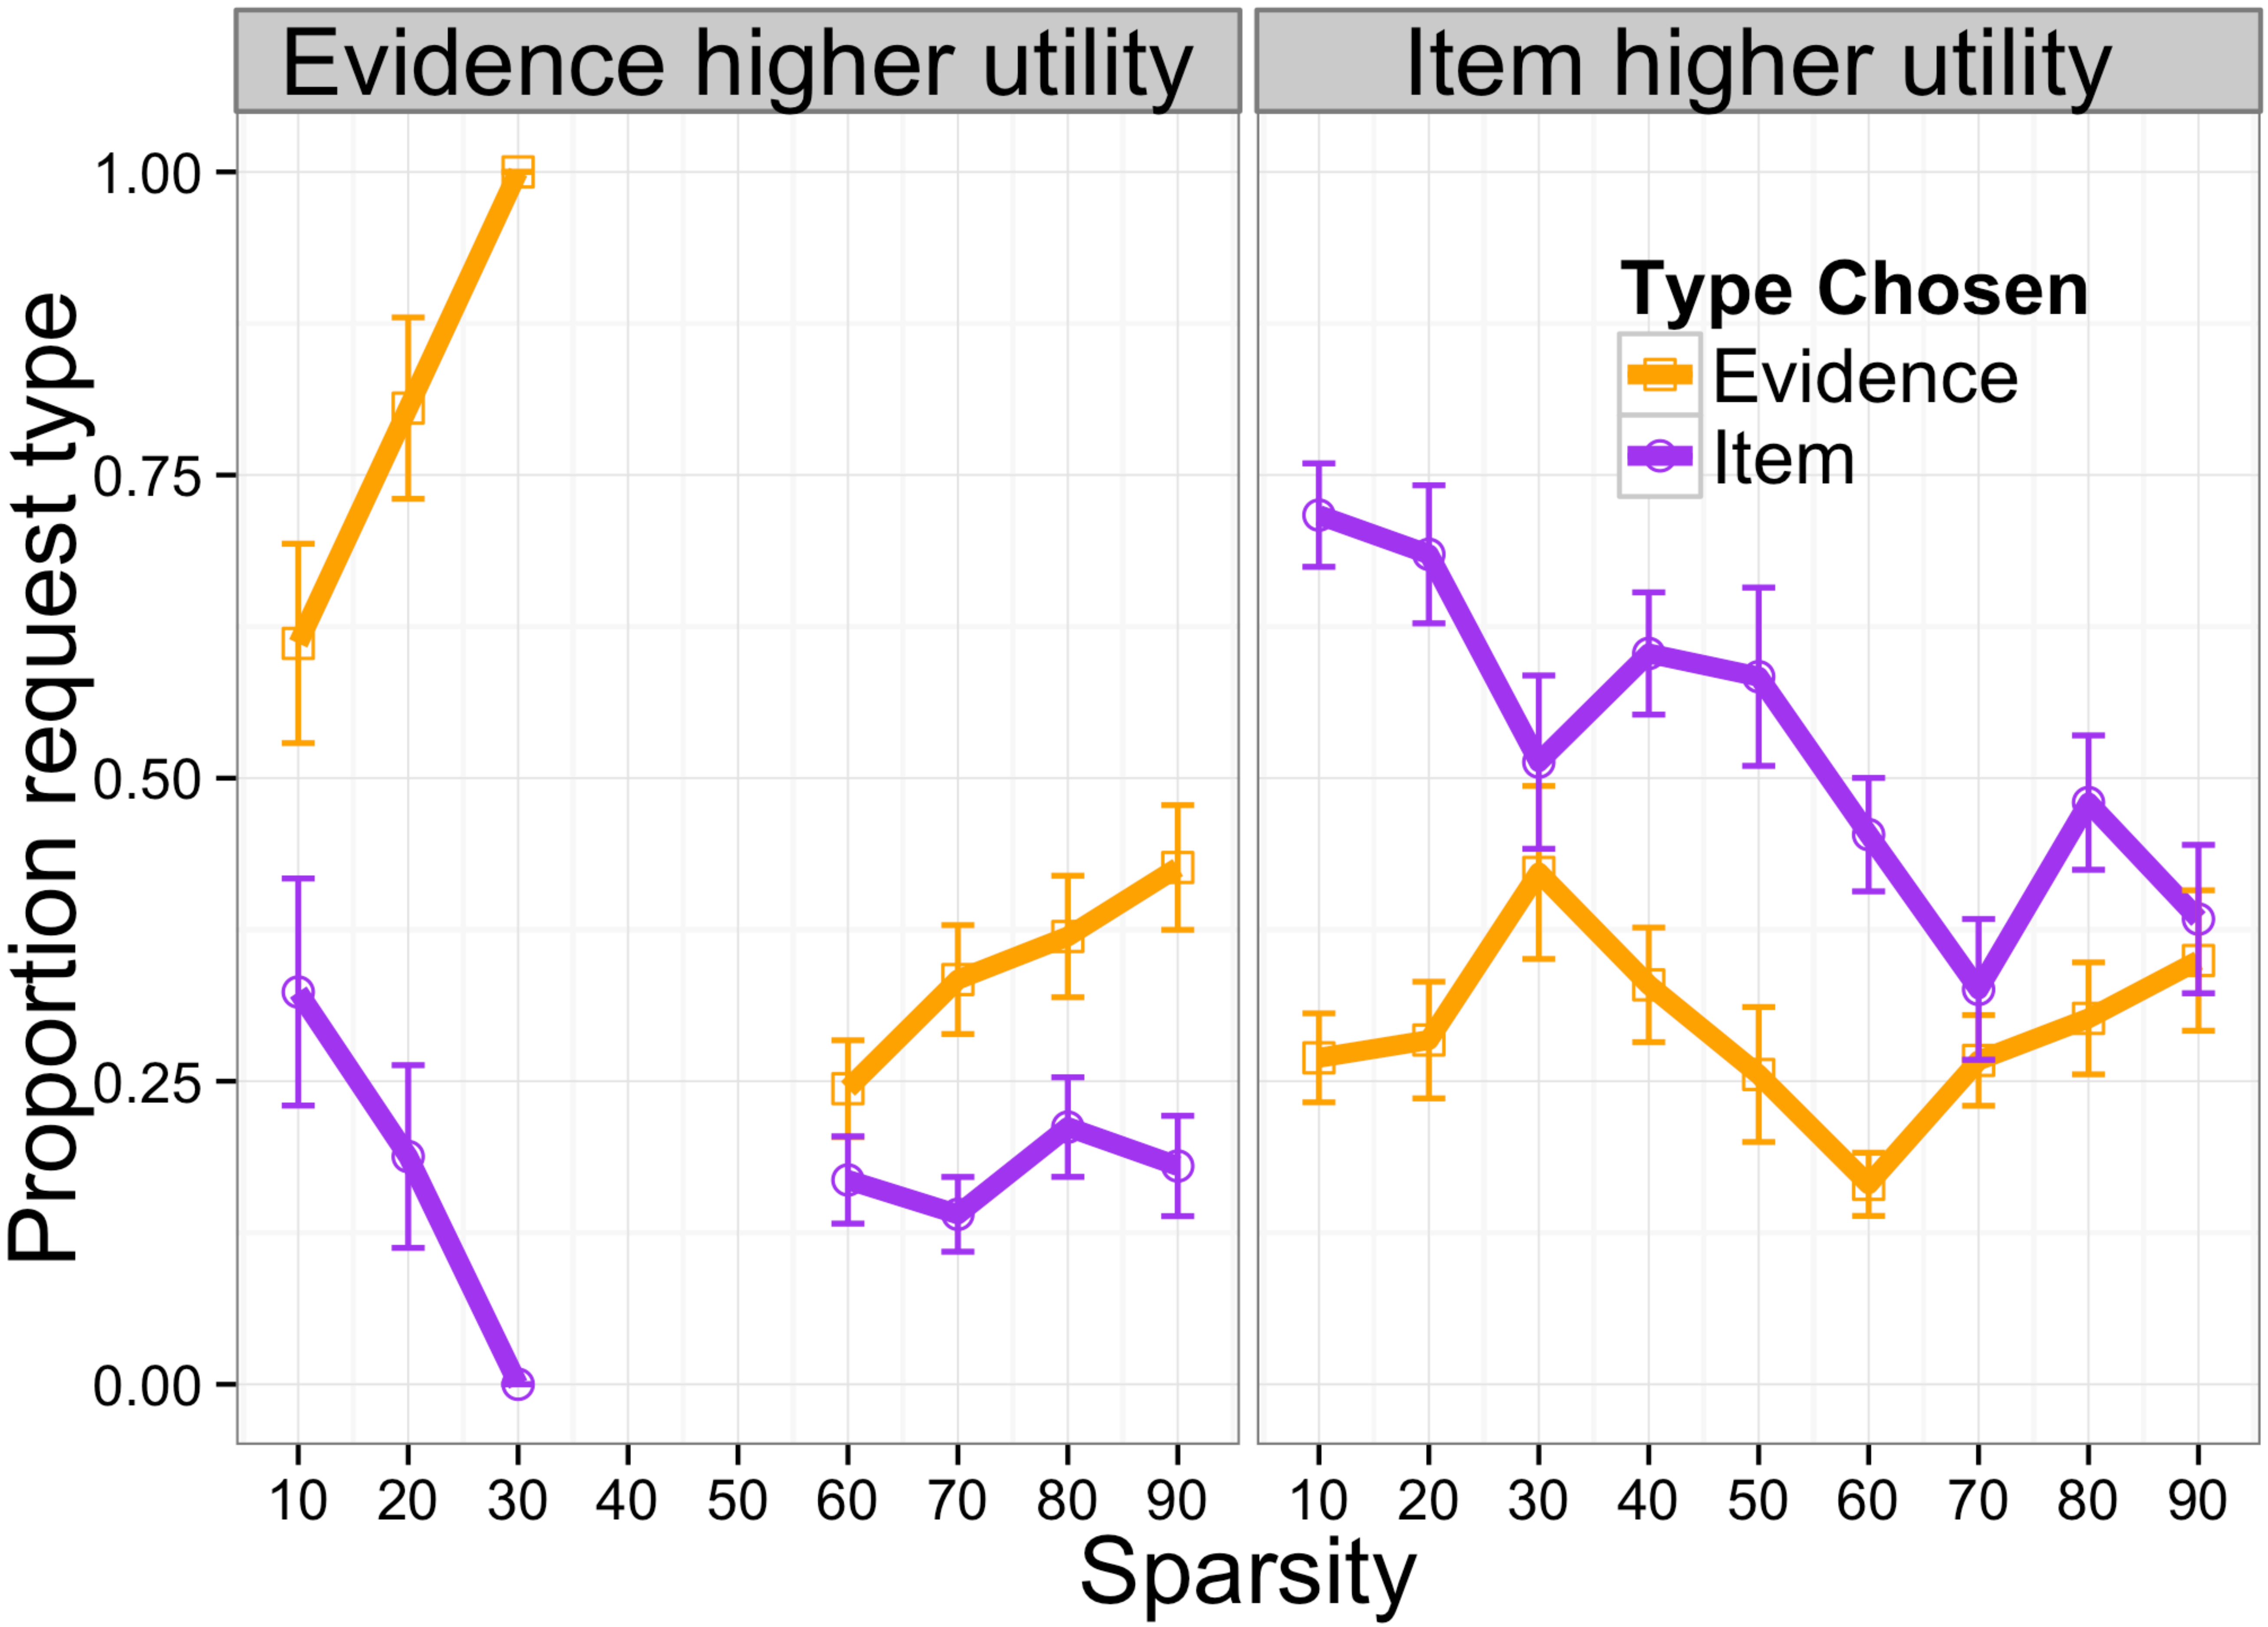

Evidence higher utility
Item higher utility
Type Chosen
Evidence
Item
Proportion request type
1.00
0.75
0.50
0.25
0.00
10
20
30
40
50
60
70
80
90
Sparsity

Choosing useful options

Do people choose ‘good’ request options?

1. Are some requests better than others?
YES

➡ 2. Are people more likely to select more useful request options?
YES

Summary

- People show a bias for positive tests in a wide variety of situations

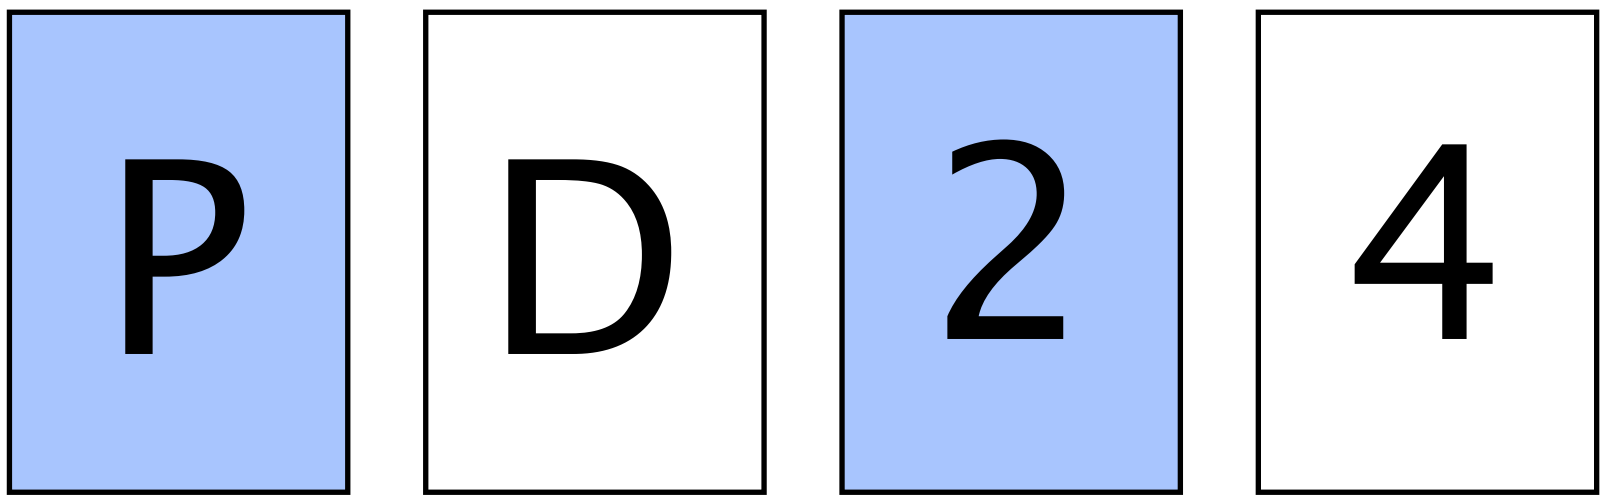

Summary

- People show a bias for positive tests in a wide variety of situations
- This bias, called the positive test strategy, may be sensible if most hypotheses in the space are sparse (as is typical for most hypotheses in the real world)

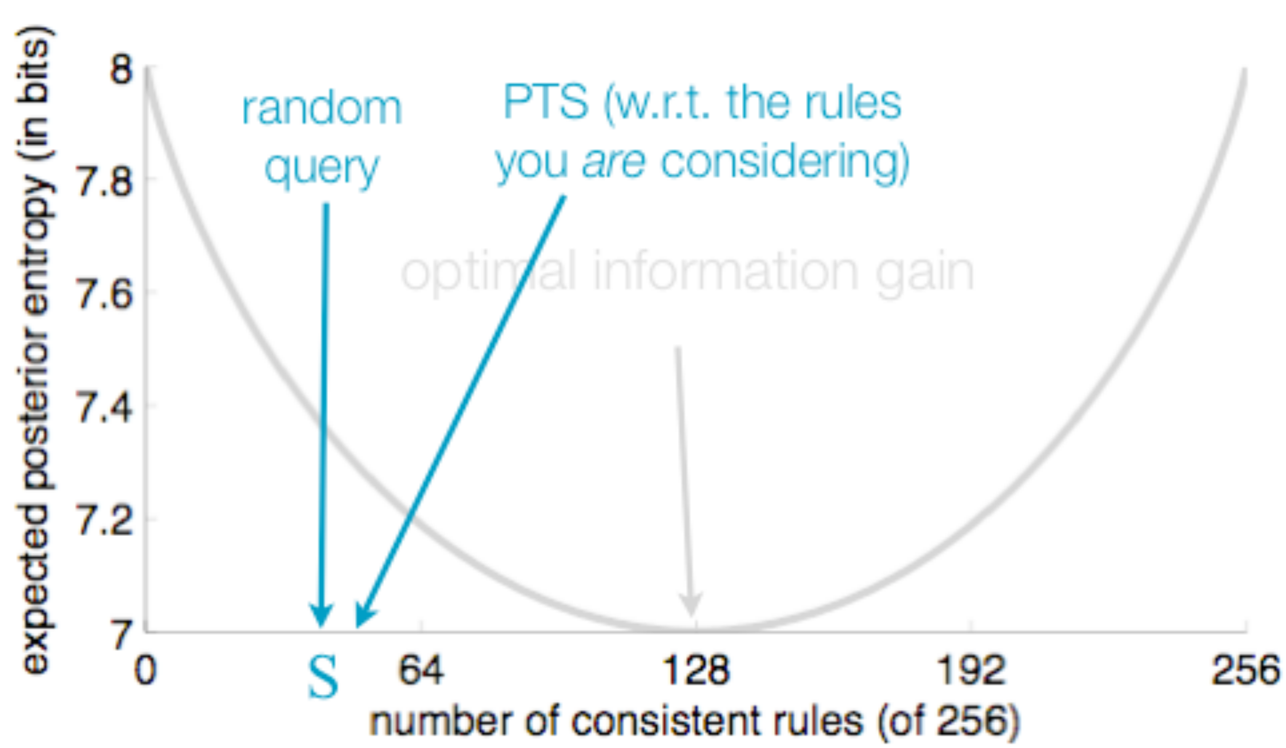

Summary

- People show a bias for positive tests in a wide variety of situations
- This bias, called the positive test strategy, may be sensible if most hypotheses in the space are sparse (as is typical for most hypotheses in the real world)
- People indeed are sensitive to sparsity when asking for different kinds of evidence

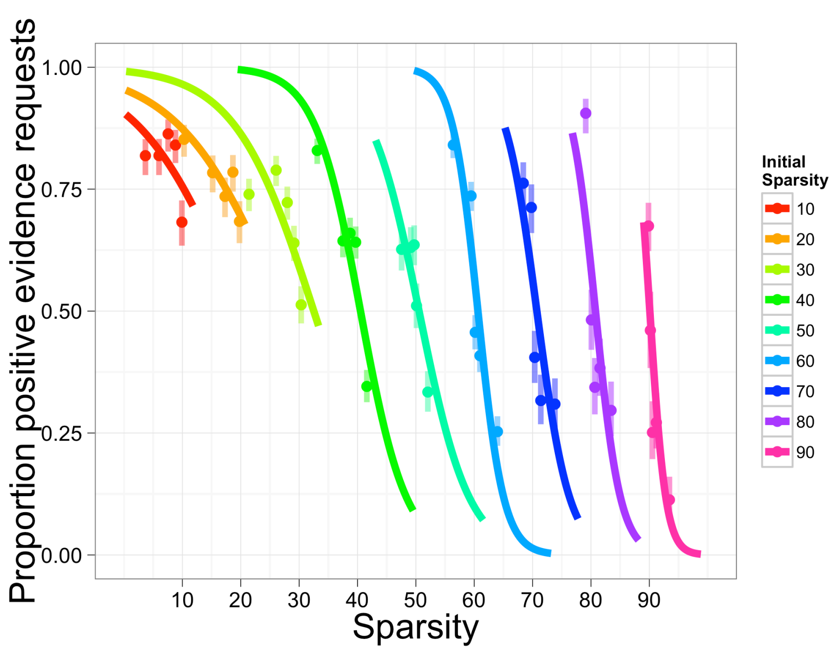

Summary

- People show a bias for positive tests in a wide variety of situations
- This bias, called the positive test strategy, may be sensible if most hypotheses in the space are sparse (as is typical for most hypotheses in the real world)
- People indeed are sensitive to sparsity when asking for different kinds of evidence
- This isn't just a heuristic -- they are also more sensitive more broadly to informativeness of data

Overall moral of this week

People are often sensitive to surprisingly subtle probabilistic reasoning -- often if they look like they're doing it "wrong", it's because they're doing something more sensible

Going forward: how do people make decisions on the basis of the inferences they make?

Additional references (not required)

- Hendrickson, A., Perfors, A., & Navarro, D. (2014). Adaptive information source selection during hypothesis testing. In P. Bello, M. Guarani, M. McShane, and B Scassellati (eds). *Proceedings of the 36th Annual Conference of the Cognitive Science Society*.

- Klayman, J. & Ha, Y-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review 94*: 211-228.

- Navarro, D., & Perfors, A. (2011). Hypothesis generatin, the positive test strategy, and sparse categories. *Psychological Review 118*: 120-134.

- Wason, P. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology 20*: 273-281.