ngoodman@
stanford.edu

Social cognition, communication, and the language of thought

Noah D. Goodman
Stanford University

IPAM graduate summer school
July 14, 2011

Statistics and composition

Thought is useful in an uncertain world

Thought is productive: "the infinite use of finite means"

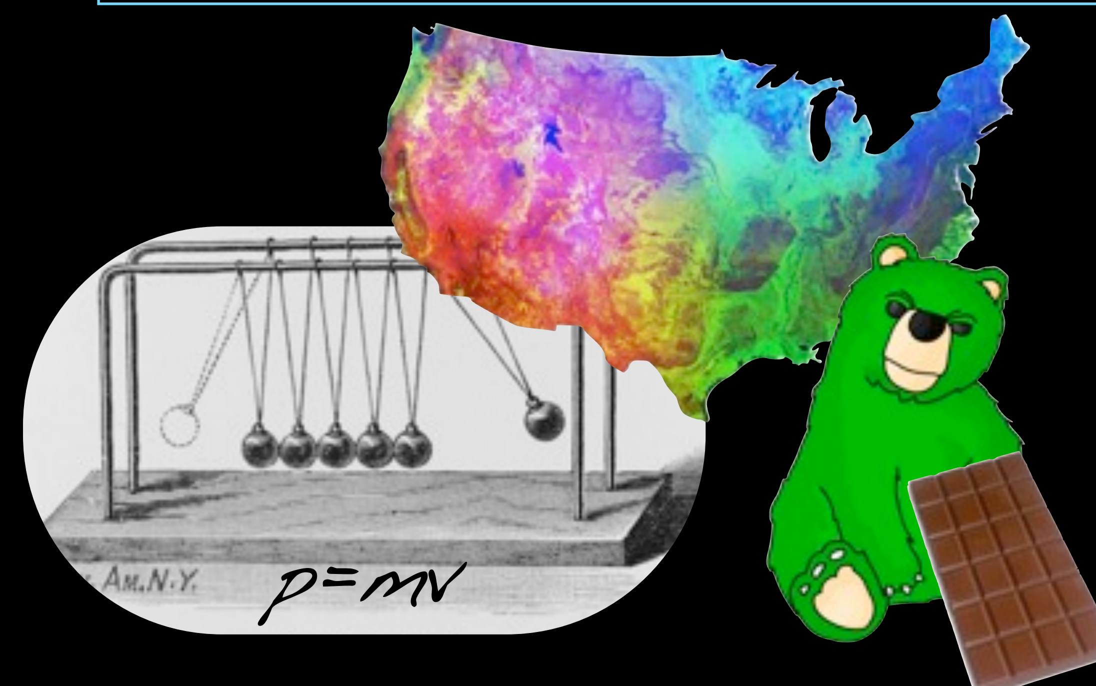

Probabilistic inference

Generative models

Compositional representations

Statistics and composition

Probabilistic language of thought hypothesis

Thought is useful in an uncertain world

Thought is productive: "the infinite use of finite means"

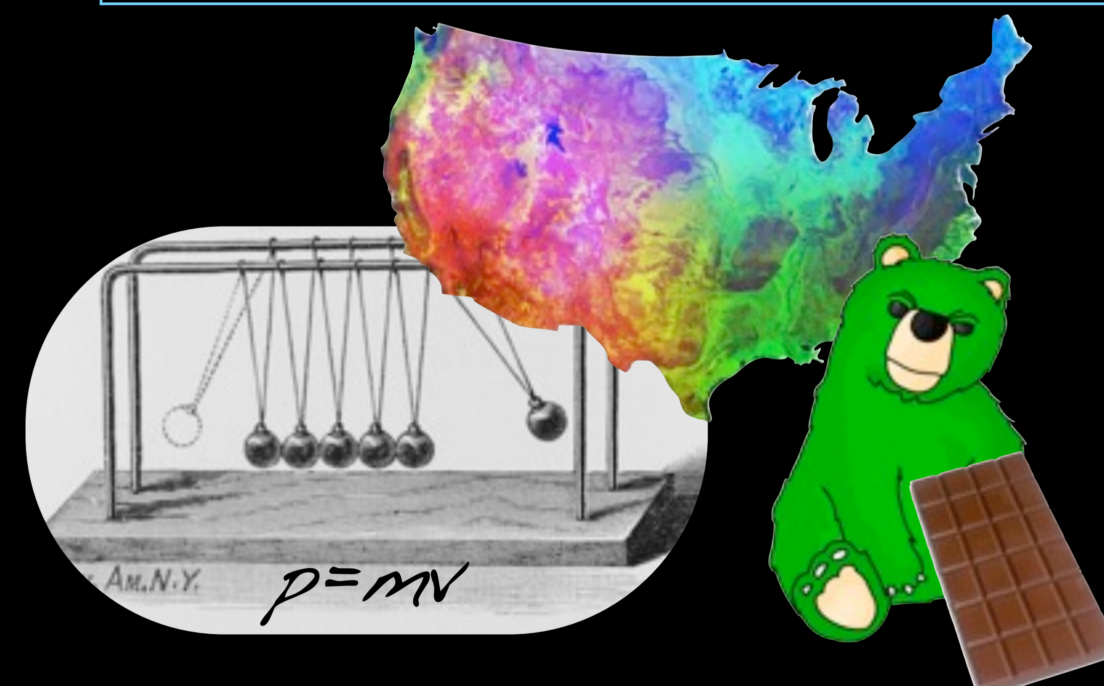

Probabilistic inference

Generative models

Compositional representations

PLoT

- The probabilistic language of thought hypothesis:
  - Mental representations are compositional,
  - Their meaning is probabilistic,
  - They encode generative knowledge,
- Hence, they support thinking and learning by probabilistic inference.

PLoT

- The probabilistic language of thought hypothesis:
  Mental representations are functions in a **stochastic process** calculus (e.g. ψλ-calculus / Church).
  - Intuitive framework theories.
  - Flexible reasoning and language use.
  - Learning structured concepts.

Outline

- Theory of mind and learning from others' actions.
- Multi-agent reasoning: coordination games.
- Communicating with natural signs: intuitive pedagogy.
- Communicating with arbitrary signs: natural language.

Bob's box

Goodman, Baker, Tenenbaum (2009)
Goodman & Baker (in prep.)

Bob's box

- Bob has a box with two buttons and a light.

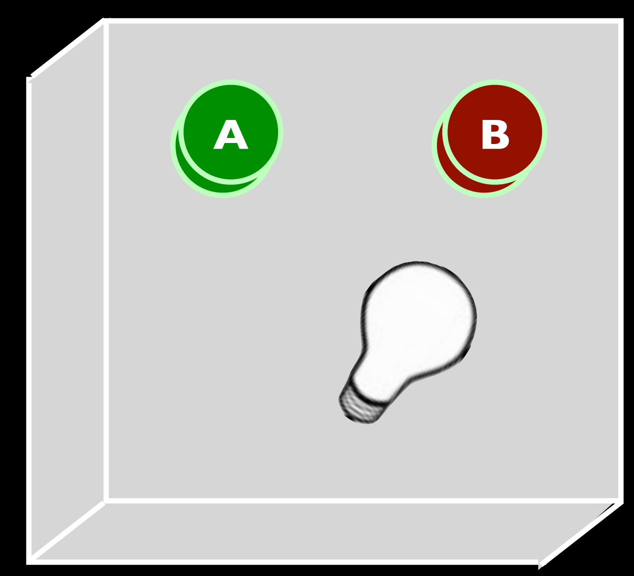

Goodman, Baker, Tenenbaum (2009)
Goodman & Baker (in prep.)

Bob's box

- Bob has a box with two buttons and a light.
- He presses both buttons, and the light comes on.

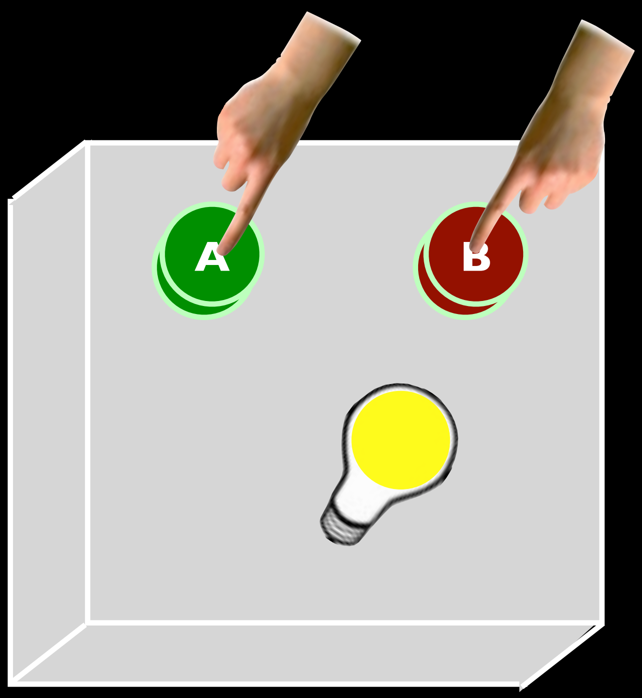

Goodman, Baker, Tenenbaum (2009)
Goodman & Baker (in prep.)

Bob's box

- Bob has a box with two buttons and a light.
- He presses both buttons, and the light comes on.
- How does the box work?

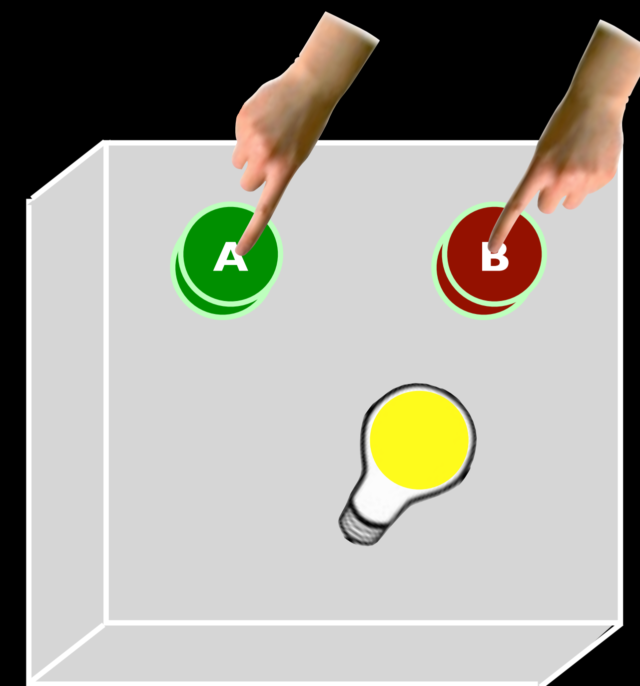

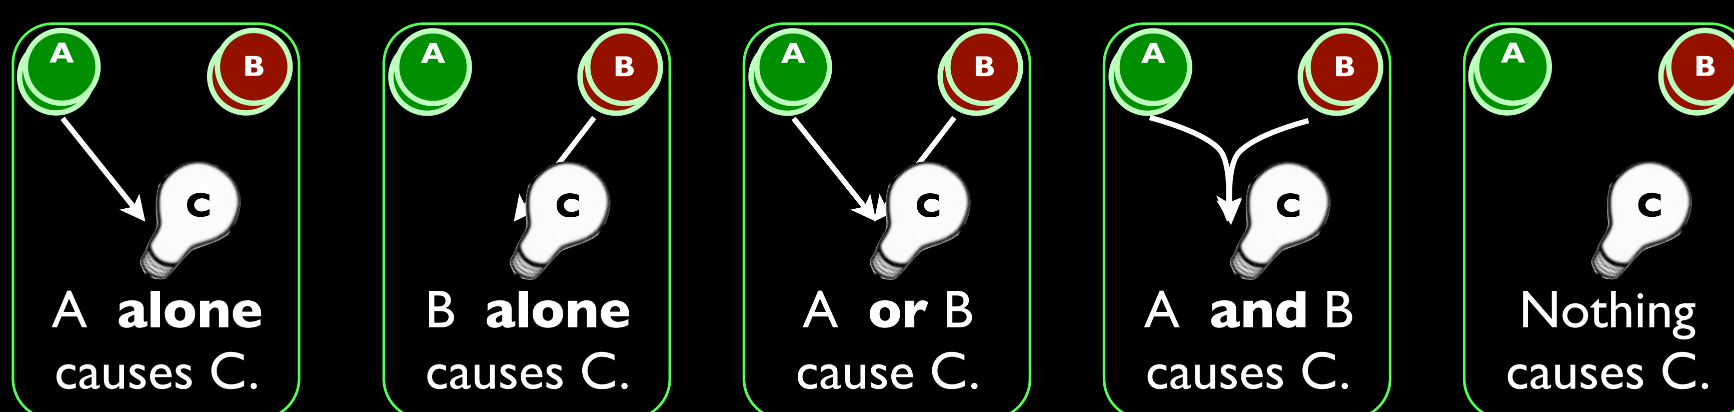

Goodman, Baker, Tenenbaum (2009)
Goodman & Baker (in prep.)

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                  state
                  action))
```

Causal learning models

Causal-only

```
(query
 (define world-cs (cs-prior))
 (define action (uniform))
 (define outcome (world-cs
                  state
                  action))
 world-cs
 (and (press-A action)
      (press-B action)
      (light-on outcome)))
```

- Given actions and outcomes, infer most likely causal structure. E.g. Griffiths & Tenenbaum (2005)

Causal learning models

Causal-only

```
(query
 (define world-cs (cs-prior))
 (define action (uniform))
 (define outcome (world-cs
                  state
                  action))
 world-cs
 (and (press-A action)
      (press-B action)
      (light-on outcome)))
```

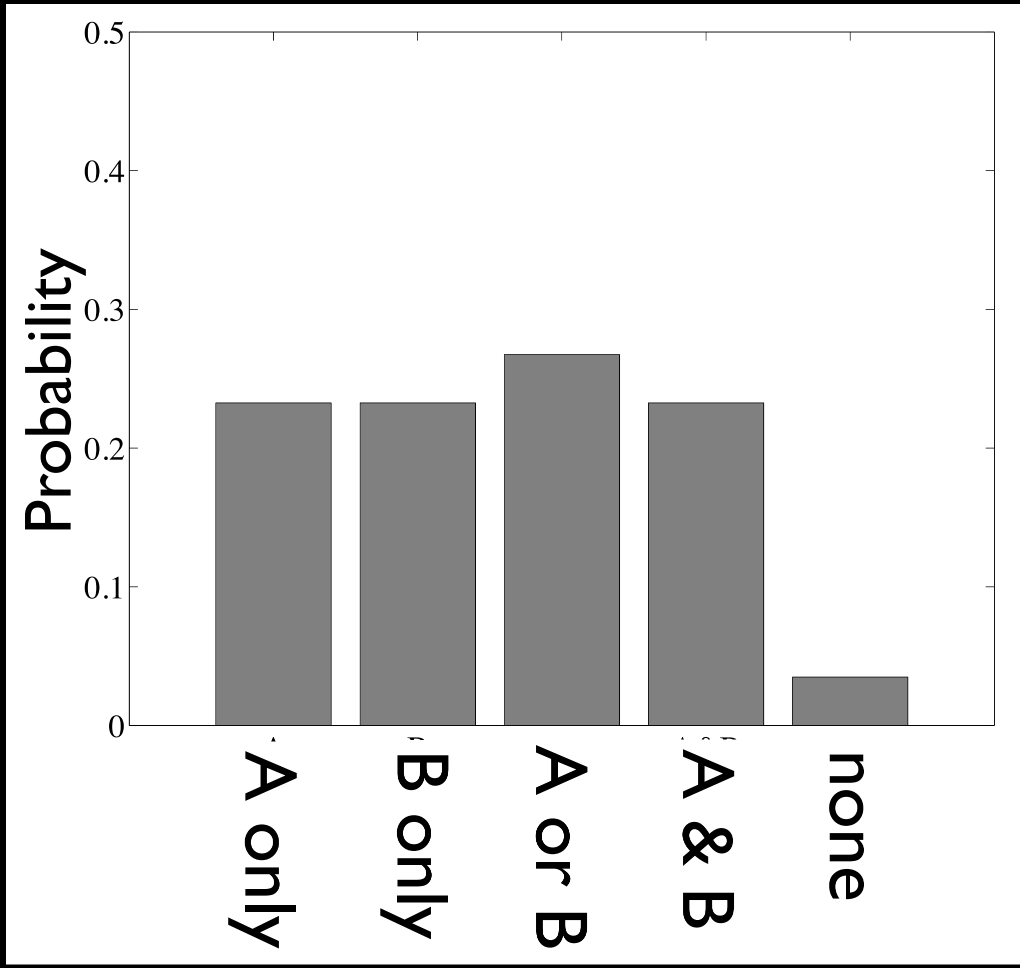

- Given actions and outcomes, infer most likely causal structure. E.g. Griffiths & Tenenbaum (2005)
- Predicts weak inferences (confounded evidence).

Explaining actions

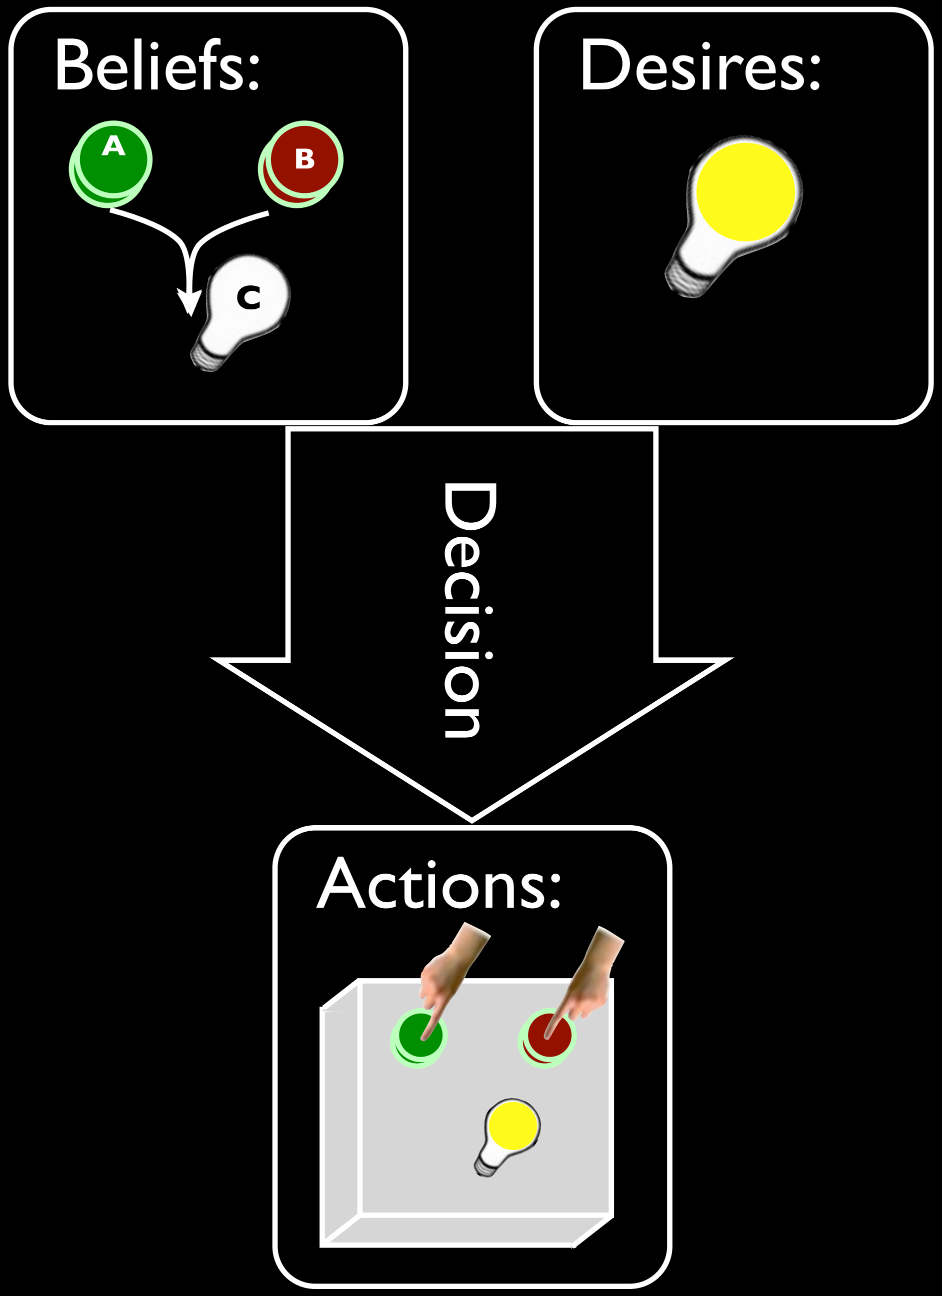

See: Goodman, et al 2009; Baker, et al 2009.

Explaining actions

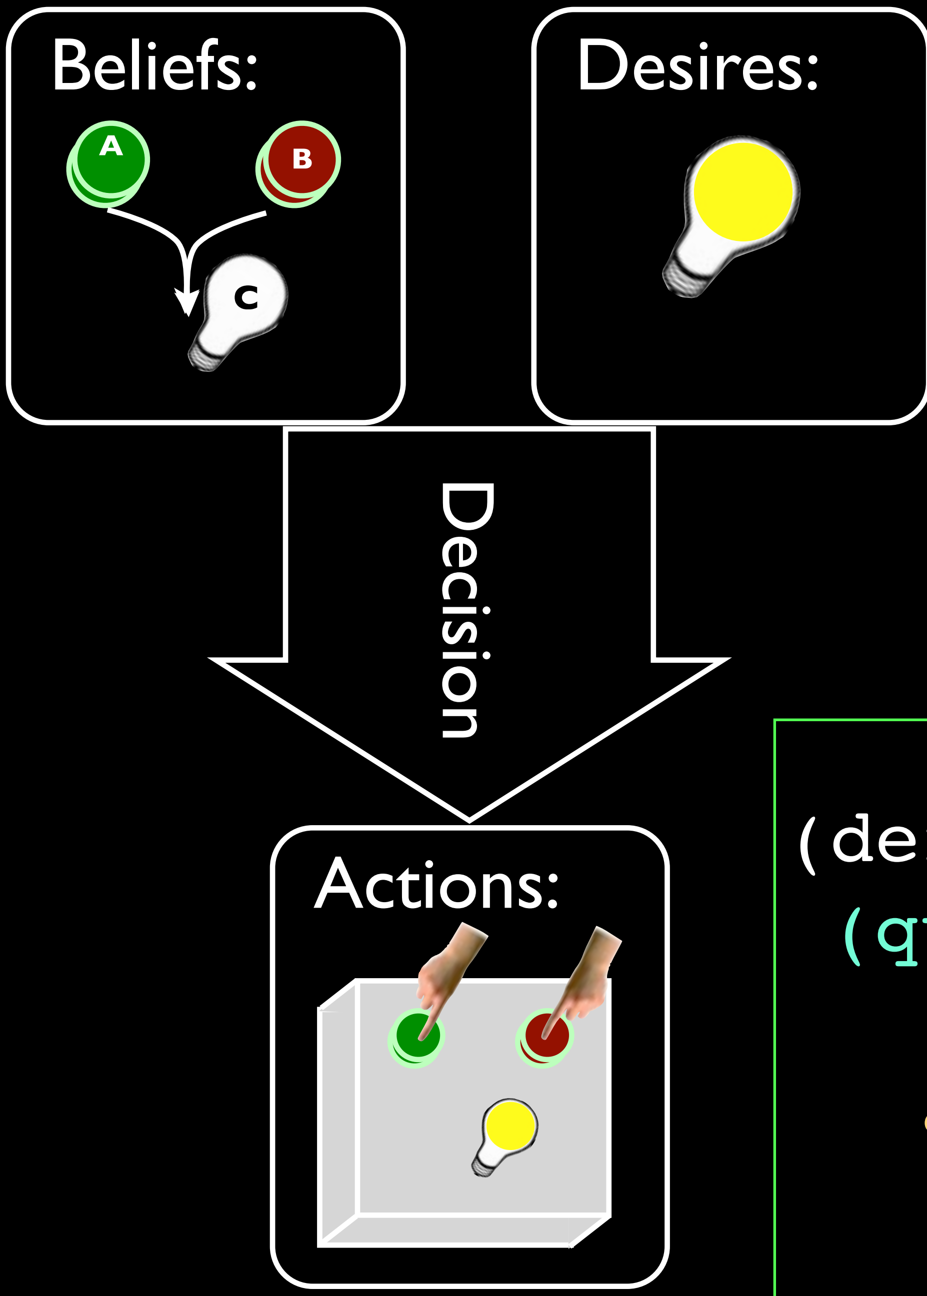

Rational action as inference:

```
(define (decide state causal-model goal?)
 (query
   (define action (action-prior))
   action
   (goal?
     (causal-model state action)))))
```

See: Goodman, et al 2009; Baker, et al 2009.

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                 init-state
                 action))
```

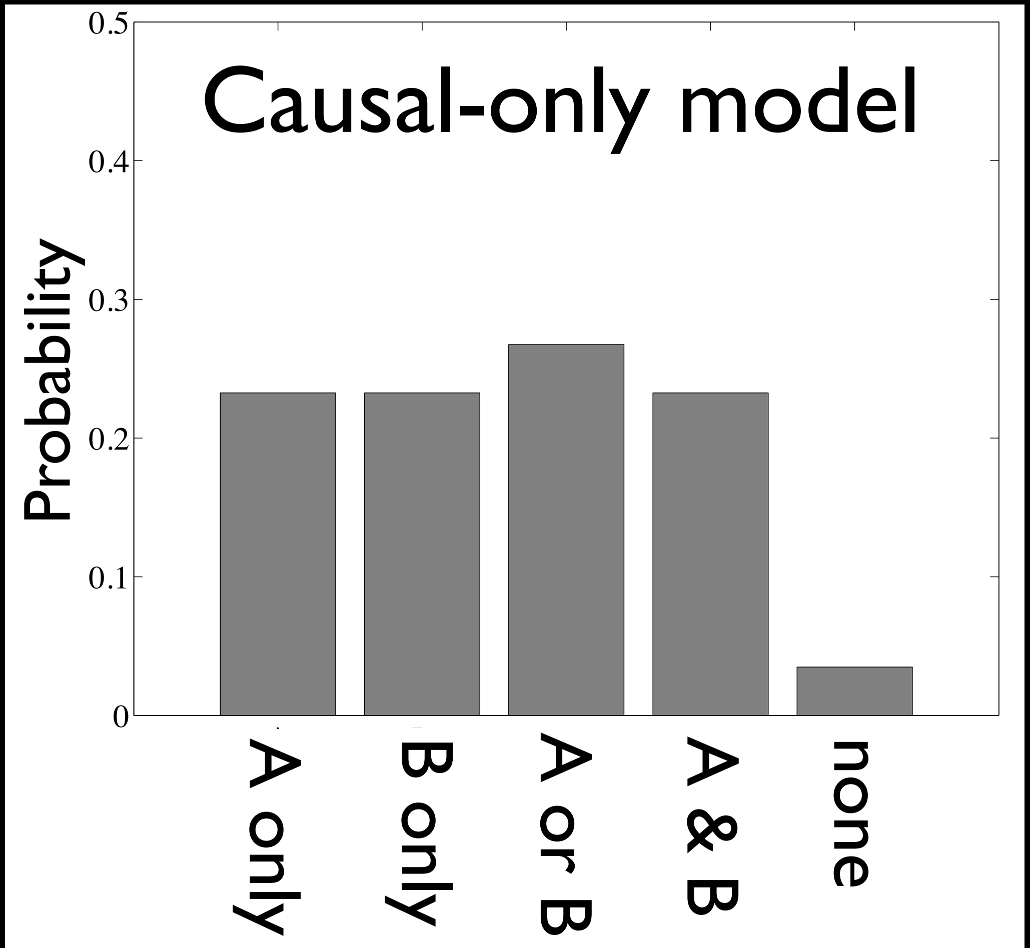

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                state
                cs-belief
                goal?))
(define outcome (world-cs
                 state
                 action))
```

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                 init-state
                 action))
```

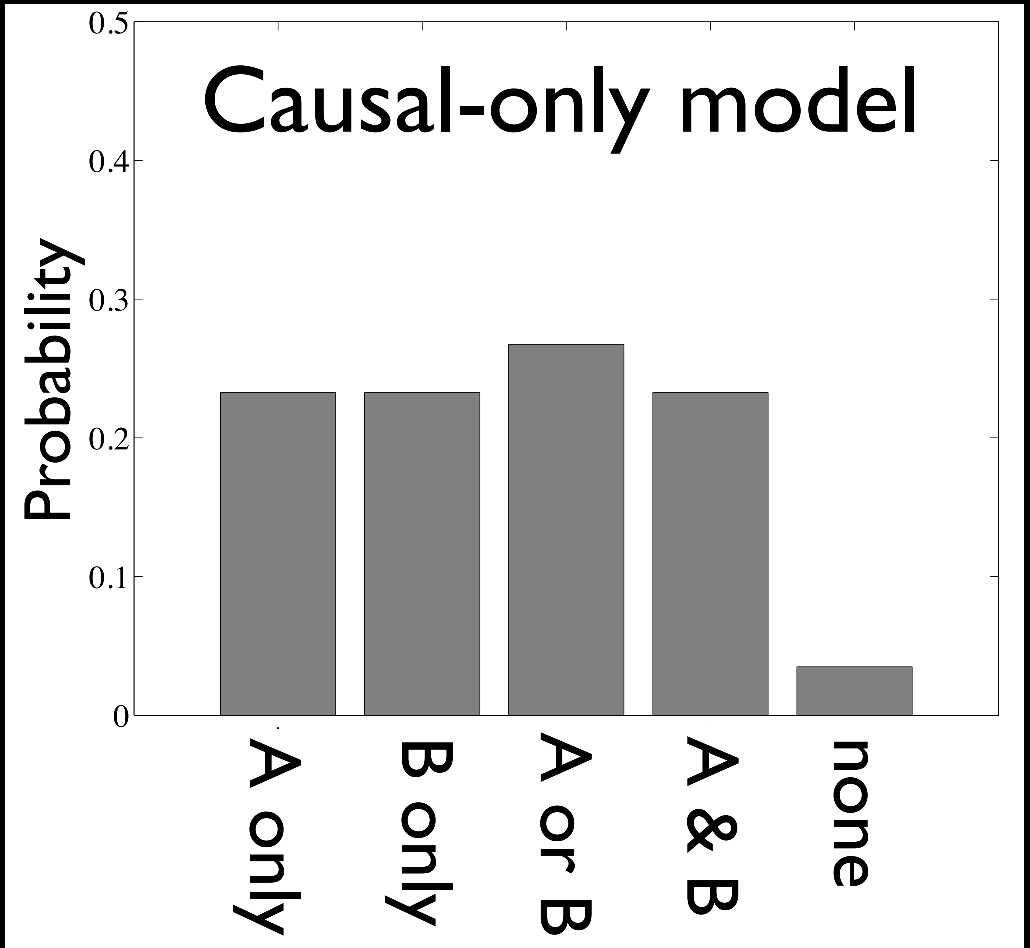

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                state
                cs-belief
                goal?))
(define outcome (world-cs
                 state
                 action))
```

Rational agent assumption

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                  init-state
                  action))
```

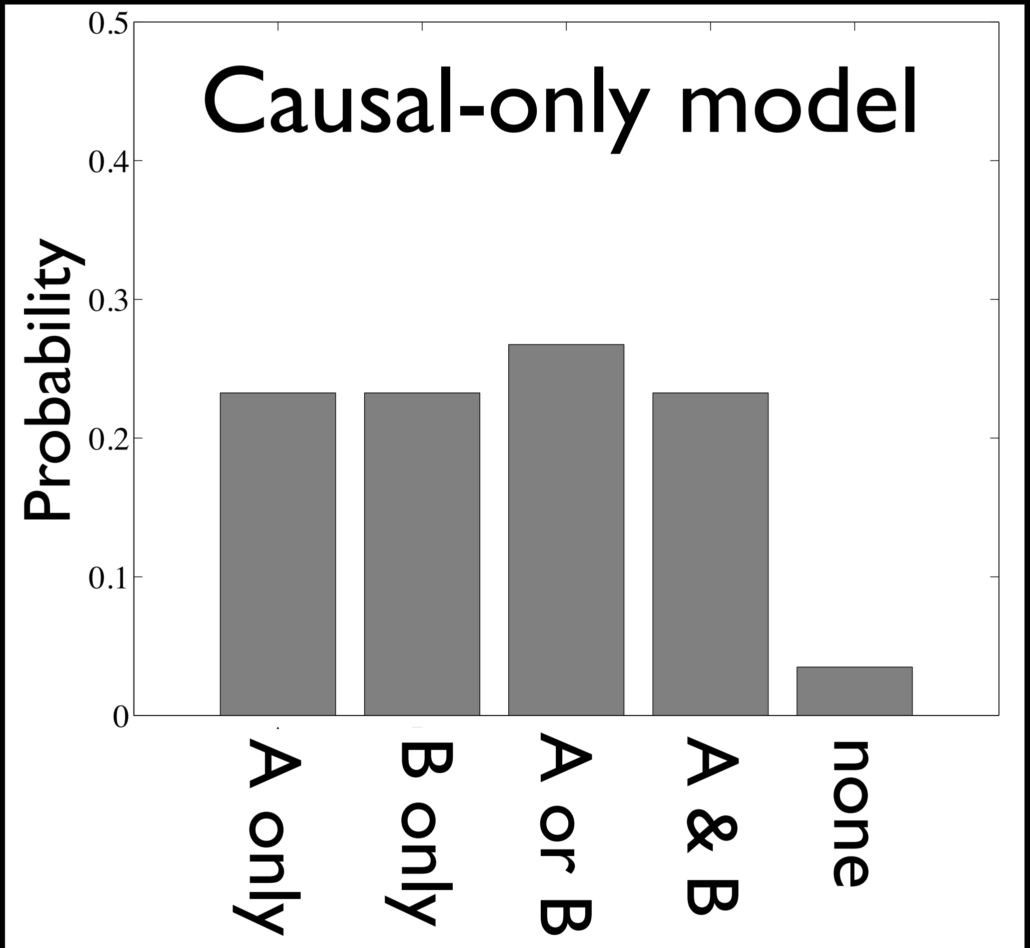

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                 state
                 cs-belief
                 goal?))
(define outcome (world-cs
                  state
                  action))
```

Knowledgeable agent assumption

Rational agent assumption

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                  init-state
                  action))
```

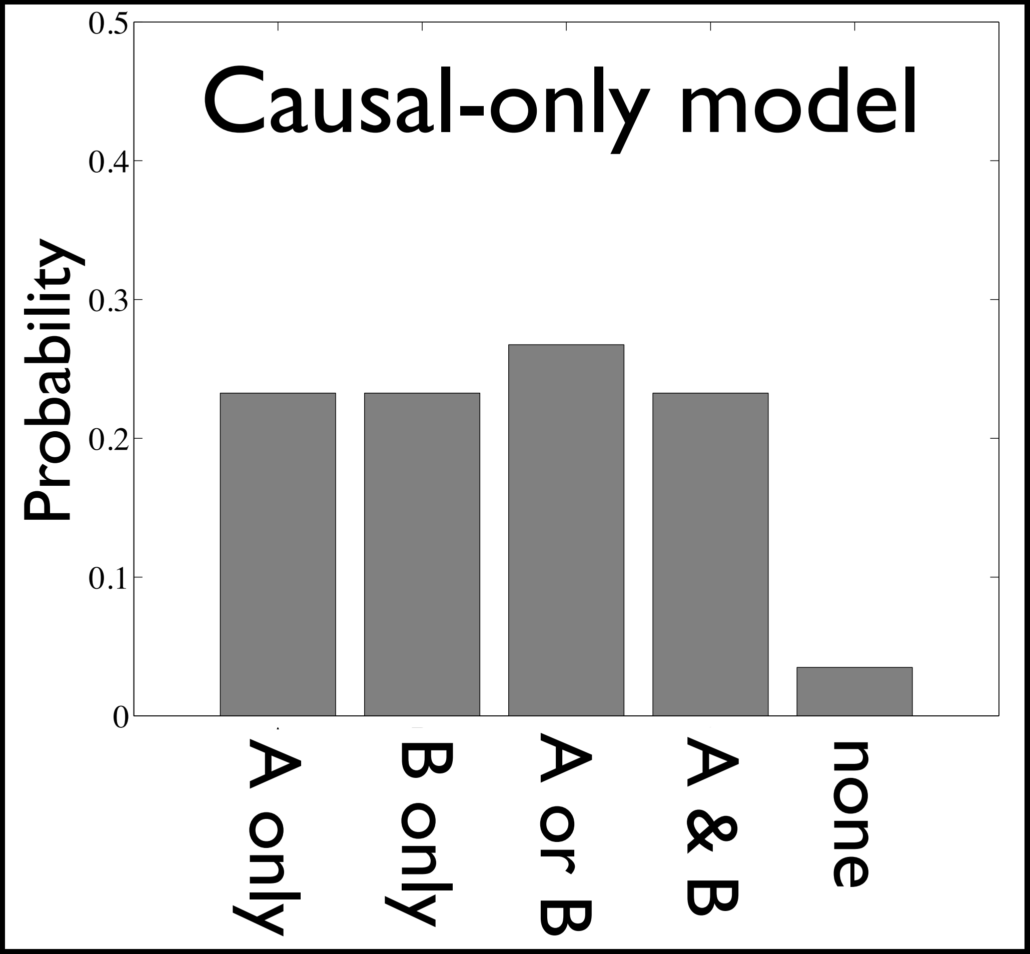

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                 state
                 cs-belief
                 goal?))
(define outcome (world-cs
                  state
                  action))
```

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                  init-state
                  action))
```

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                 state
                 cs-belief
                 goal?))
(define outcome (world-cs
                  state
                  action))
```

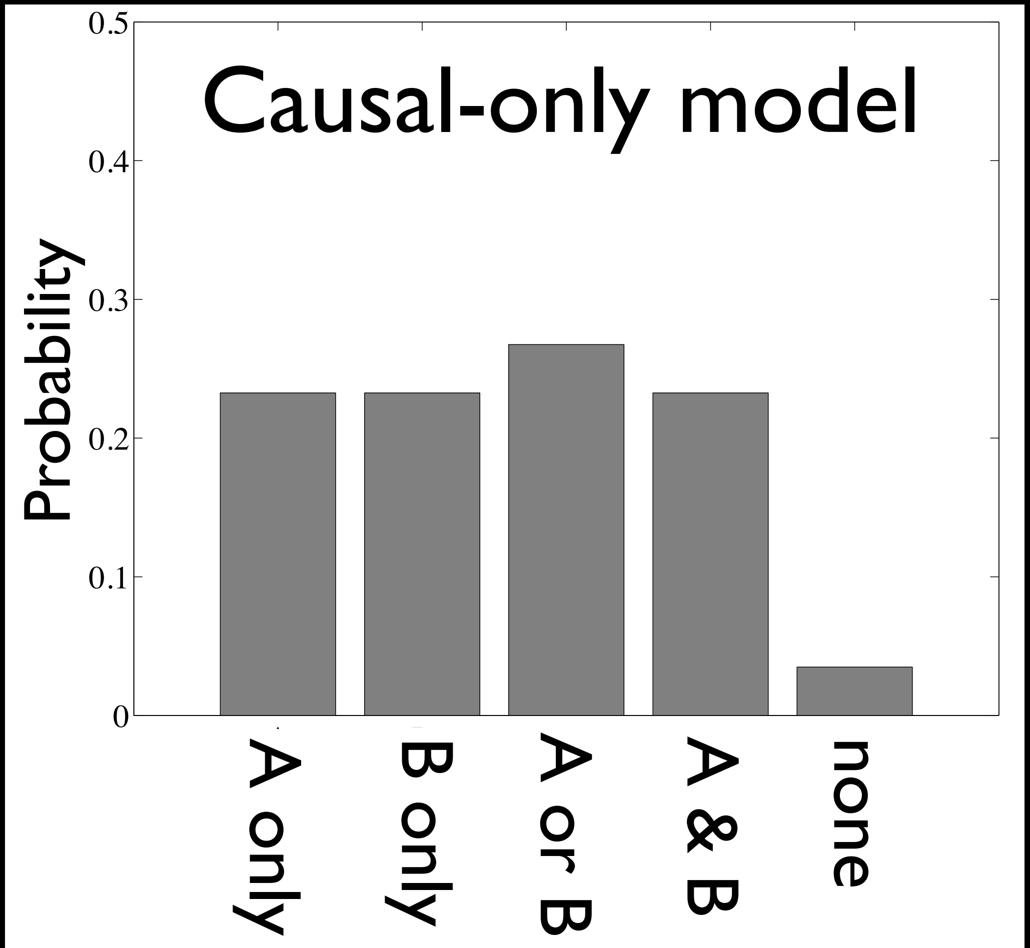

Causal learning models

Causal-only

```
(define world-cs (cs-prior))
(define action (uniform))
(define outcome (world-cs
                  init-state
                  action))
```

Social & causal

```
(define world-cs (cs-prior))
(define goal? (goal-prior))
(define cs-belief world-cs)
(define action (decide
                 state
                 cs-belief
                 goal?))
(define outcome (world-cs
                  state
                  action))
```

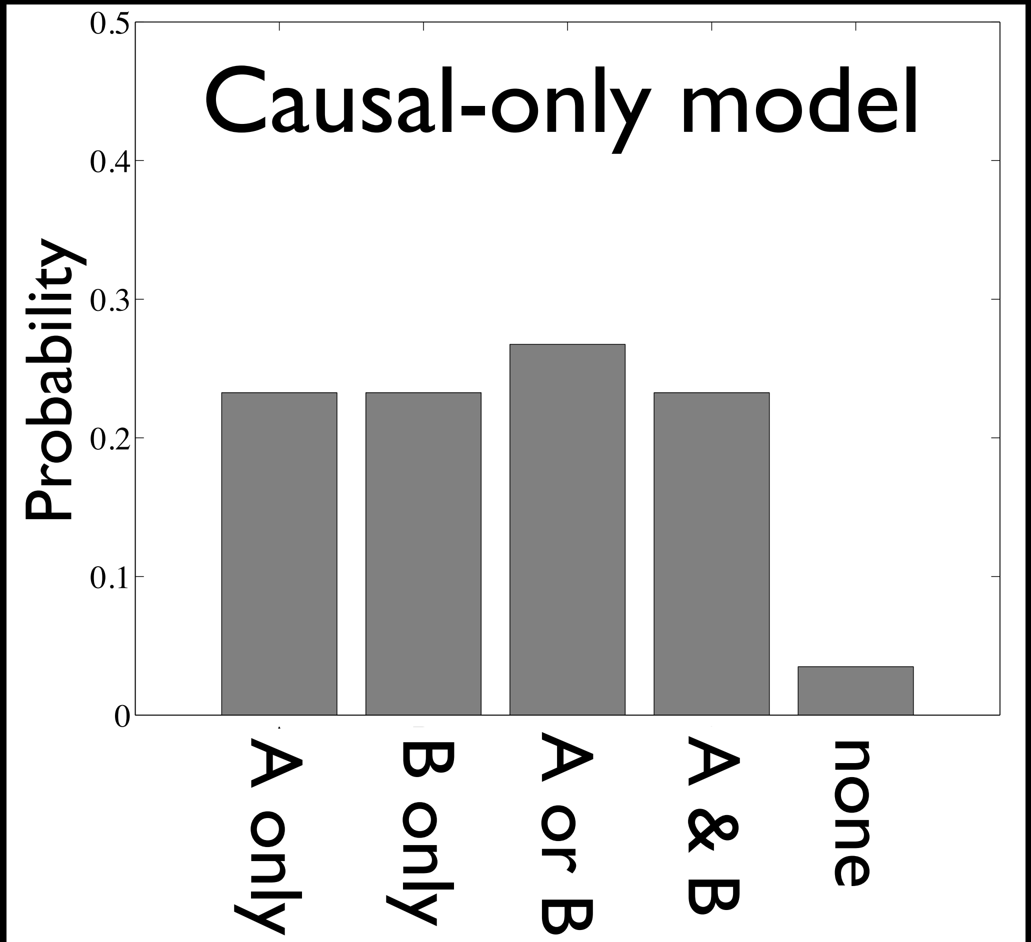

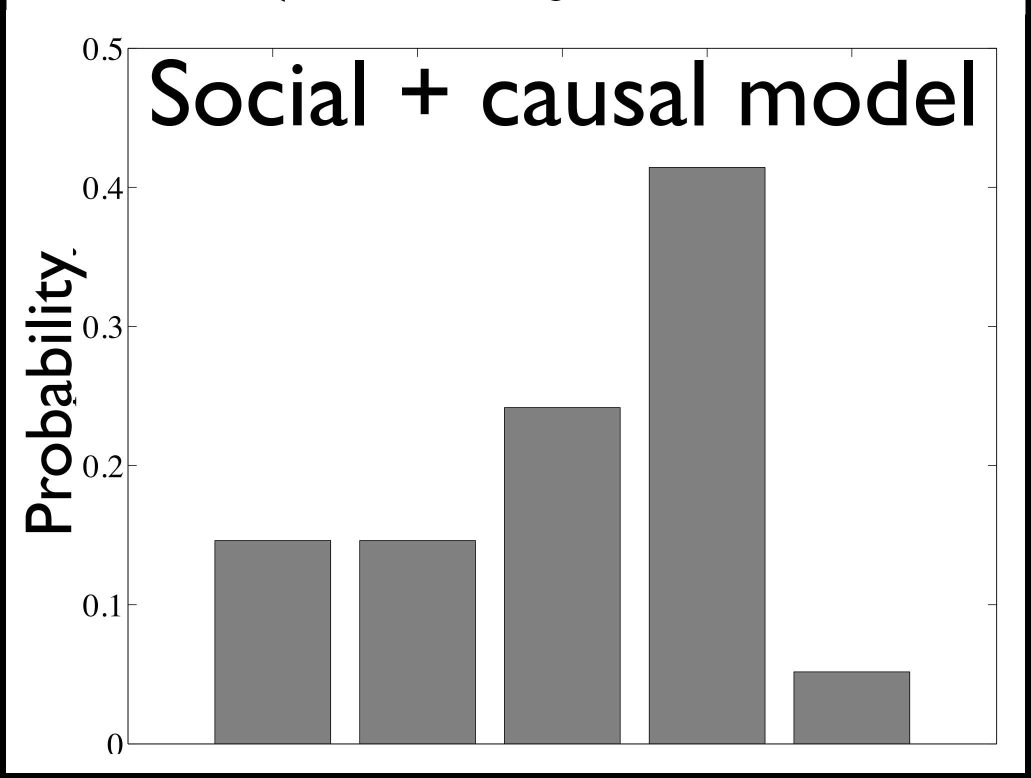

Expt 1: social vs physical

You work at a genetically-engineered plants nursery, and one of your coworkers is tending to some almost-dead flowers that you haven't seen before.

(Social cond.:)
Your coworker pours a yellow liquid and a blue liquid on the flowers.

By the end of the day, the flowers are growing again.

What causes the flowers to grow?

Expt 1: social vs physical

You work at a genetically-engineered plants nursery, and one of your coworkers is tending to some almost-dead flowers that you haven't seen before.

(Social cond.:)
Your coworker pours a
A yellow liquid and a blue liquid B
on the flowers.

By the end of the day, the flowers are growing again.

What causes the flowers to grow?

A only	_10$_
B only	_10$_
A or B	_20$_
A & B	_40$_
neither	_5$___

Expt 1: social vs physical

You work at a genetically-engineered plants nursery, and one of your coworkers is tending to some almost-dead flowers that you haven't seen before.

(Social cond.:)
Your coworker pours a
A yellow liquid and a blue liquid B
on the flowers.

(Physical cond.:)
A small earthquake knocks over a yellow liquid and a blue liquid, which pour on the flowers.

By the end of the day, the flowers are growing again.

What causes the flowers to grow?

A only	_10$_
B only	_10$_
A or B	_20$_
A & B	_40$_
neither	_5$__

Expt 1: social vs physical

You work at a genetically-engineered plants nursery, and one of your coworkers is tending to some almost-dead flowers that you haven't seen before.

(Social cond.:)
Your coworker pours a A yellow liquid and a blue liquid B on the flowers.

(Physical cond.:)
A small earthquake knocks over a yellow liquid and a blue liquid, which pour on the flowers.

By the end of the day, the flowers are growing again.

What causes the flowers to grow?

A only	_10$_
B only	_10$_
A or B	_20$_
A & B	_40$_
neither	_5$__

- 9 different cover stories, 3 domains.

Expt 1: social vs physical

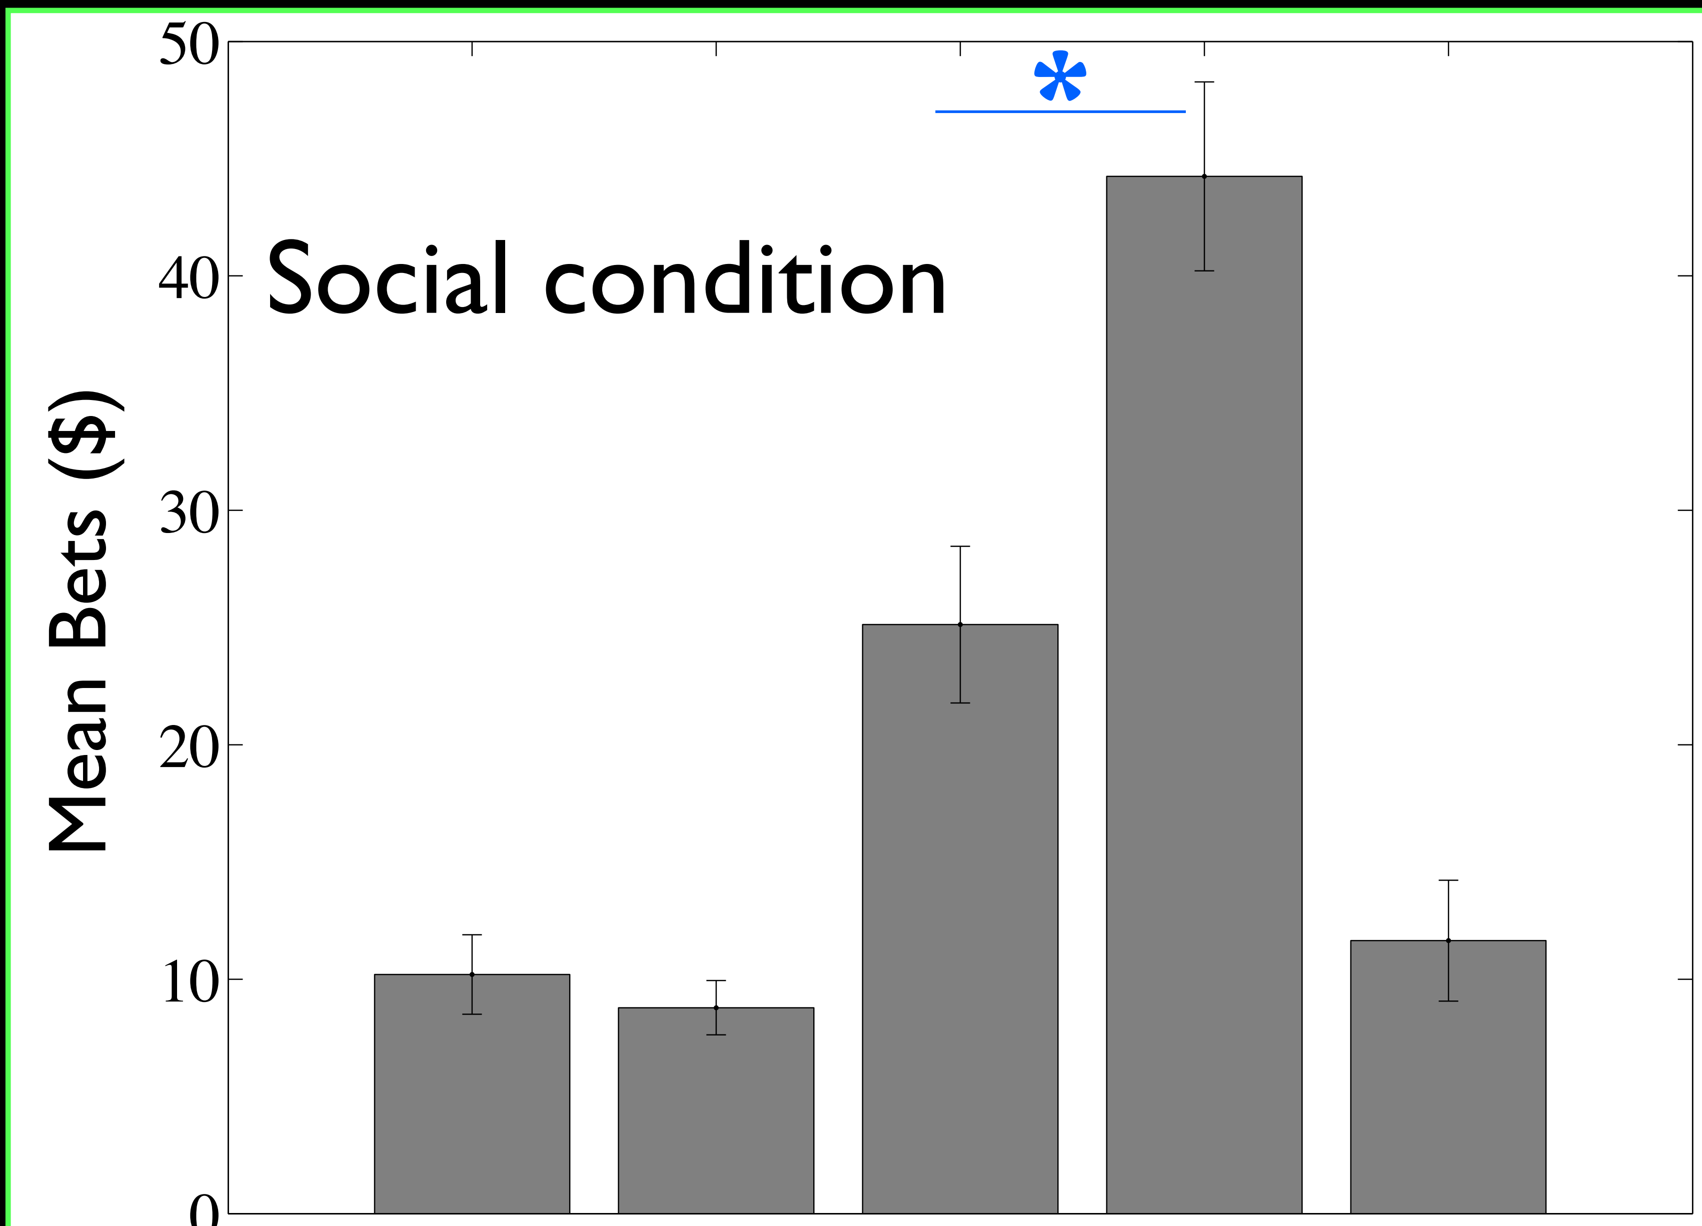

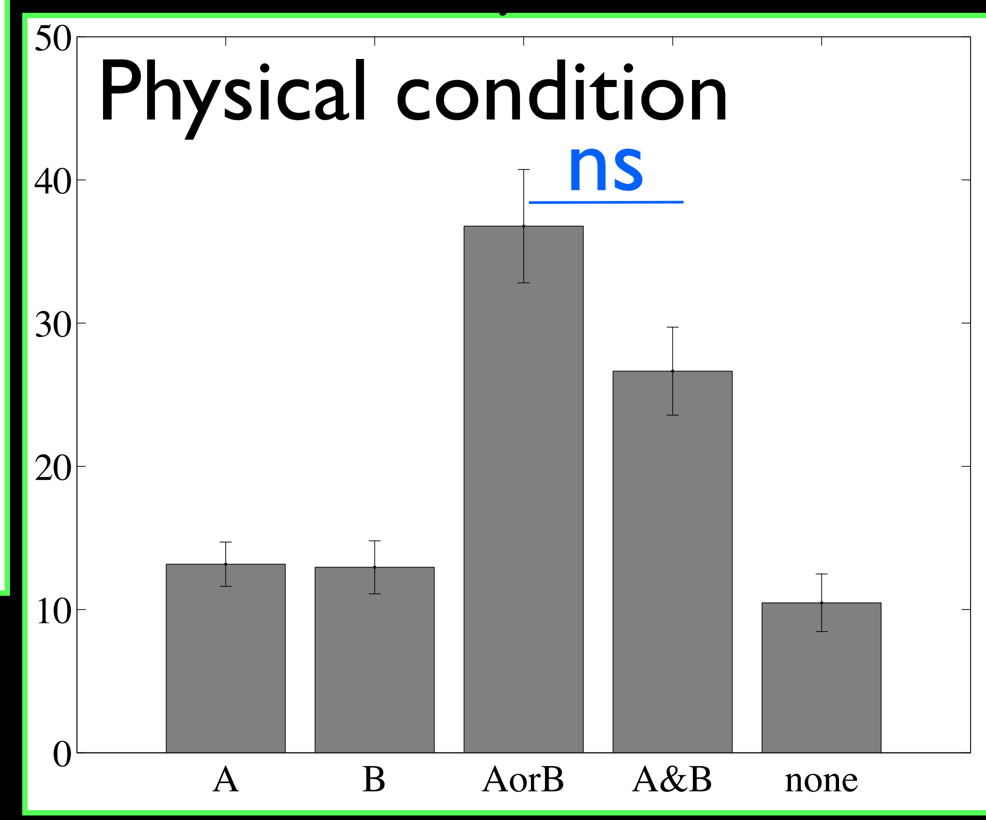

N=15

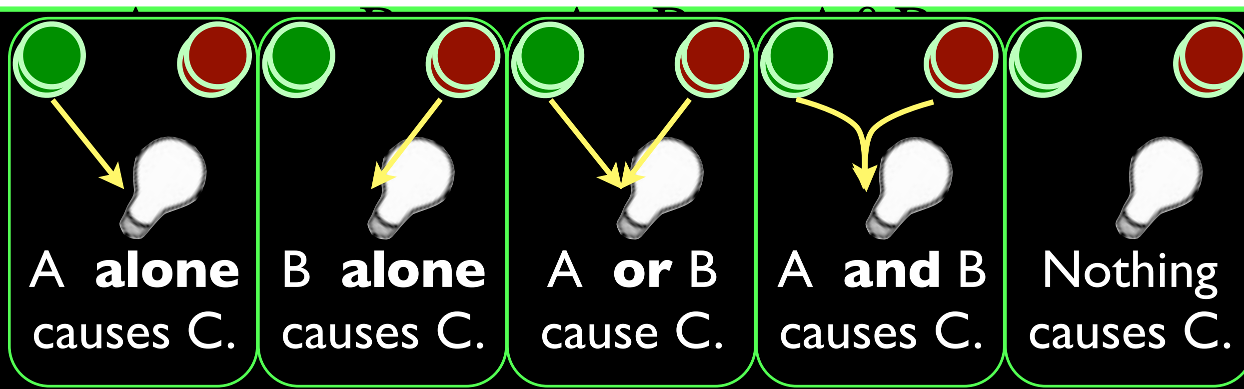

Expt 1: social vs physical

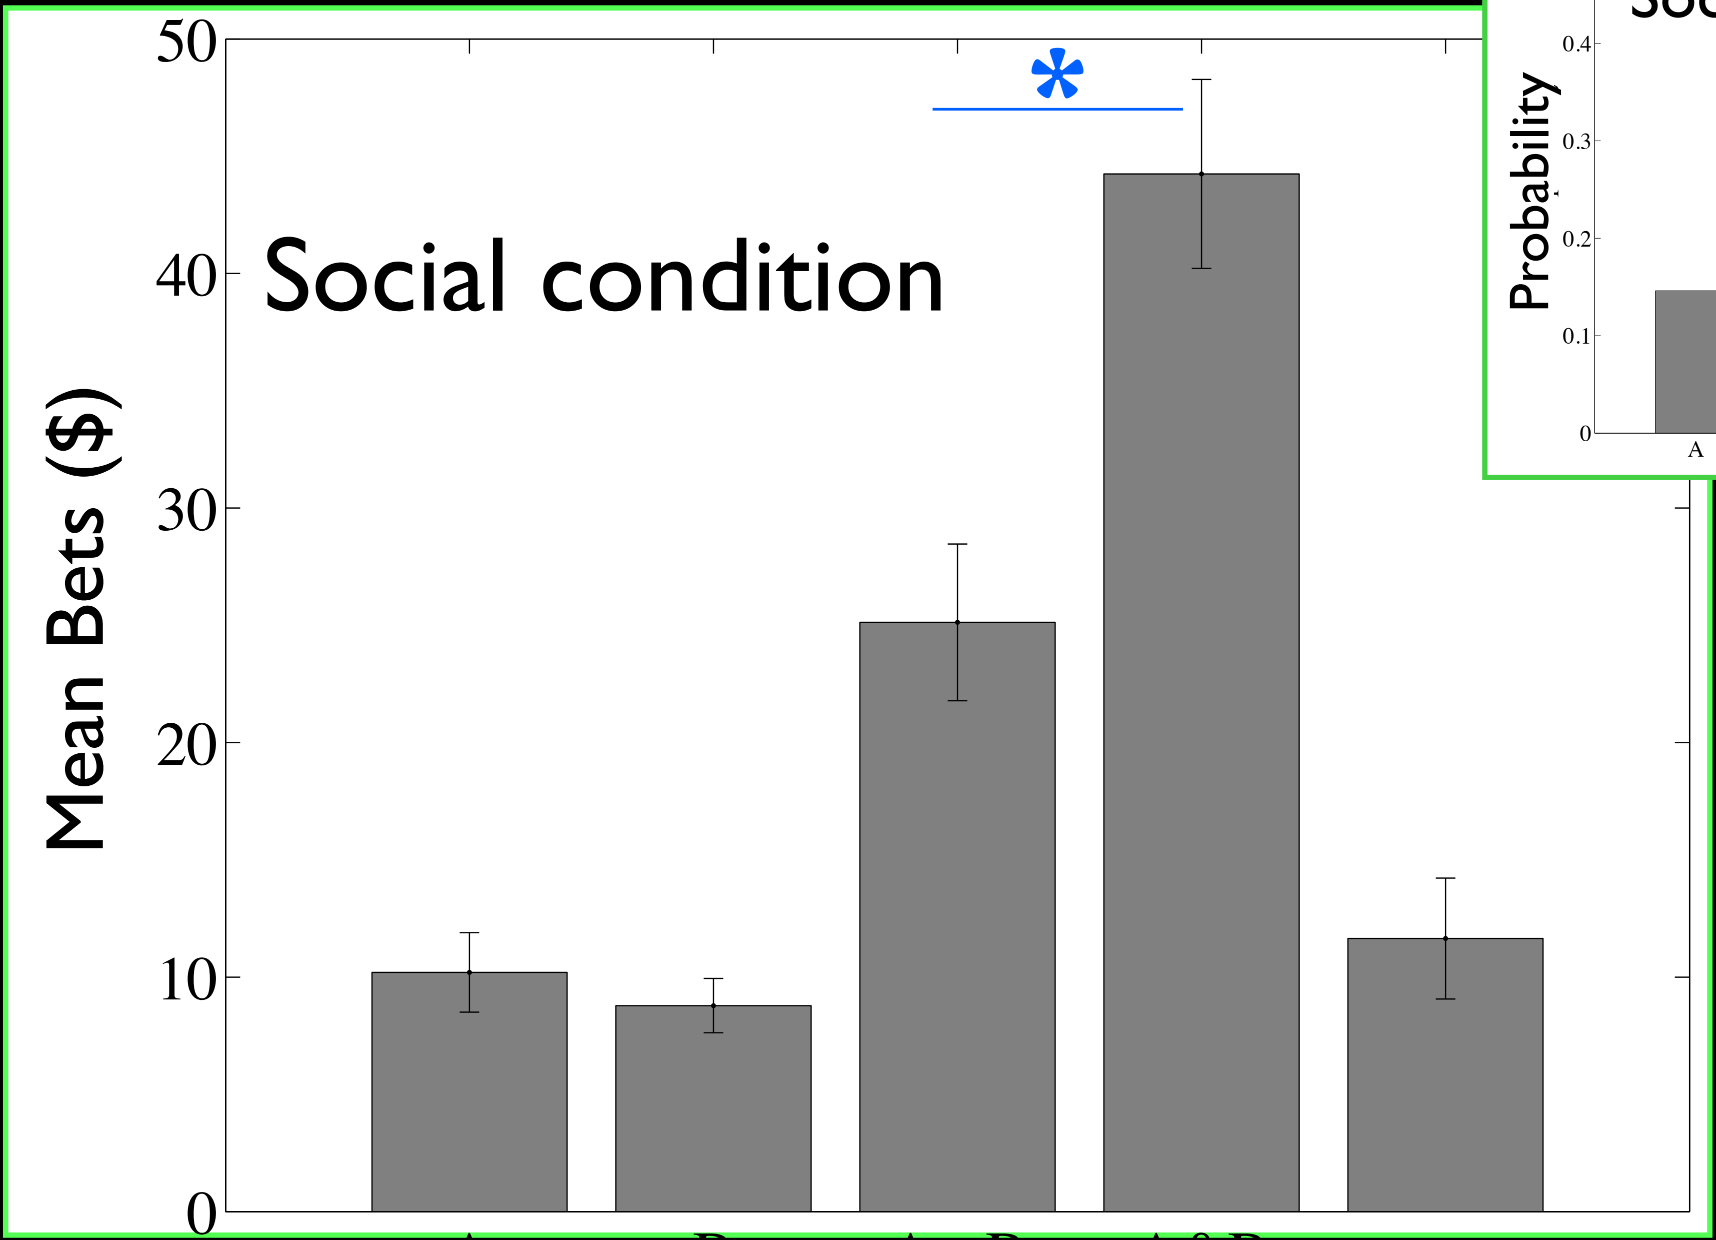

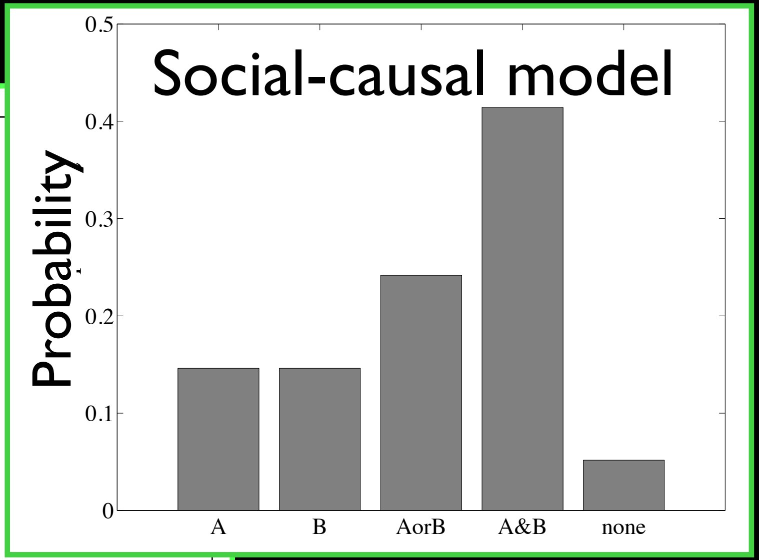

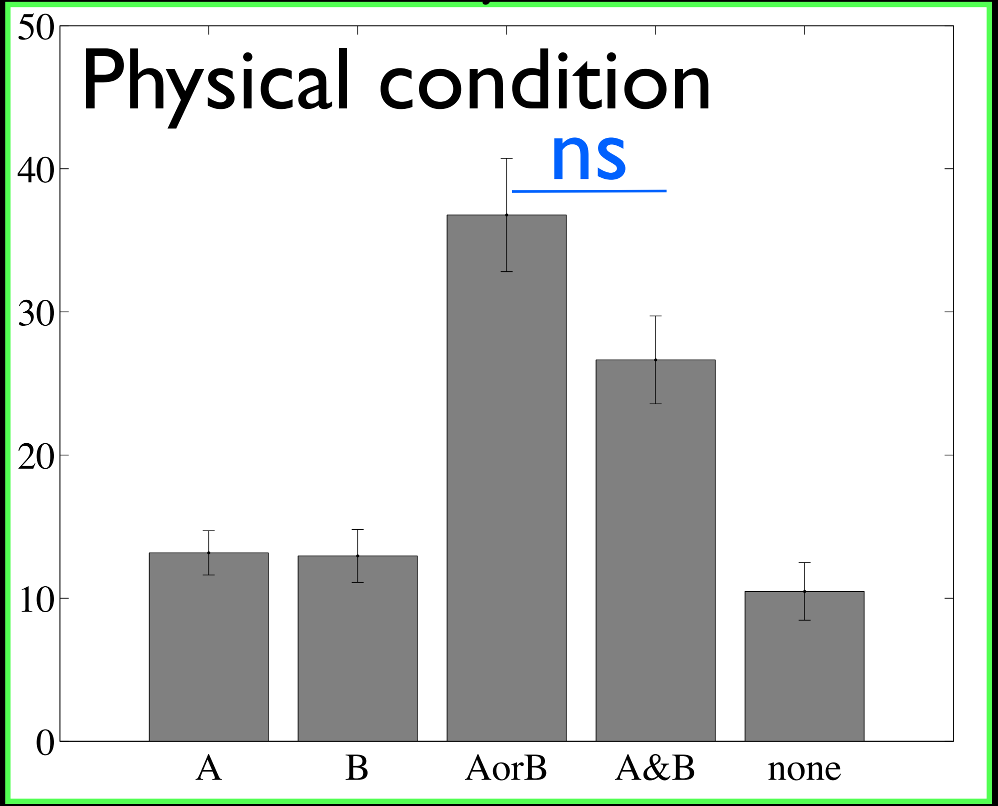

N=15

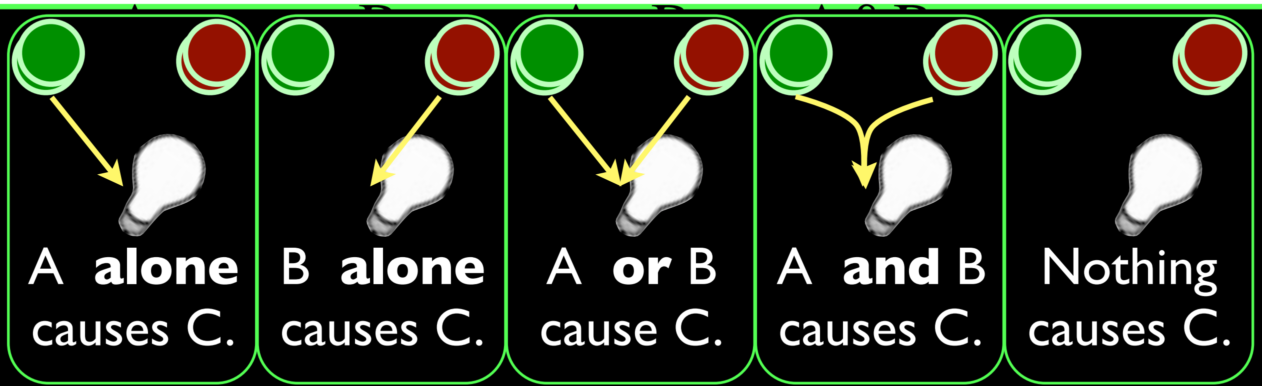

Expt 2: intentional vs accidental

- Controlling for agency.
  - Elicit intentionality judgements.
  - Median-split:

“While reaching for a notebook, your coworker accidentally knocks over a yellow liquid and a blue liquid, which pour on the flowers.”

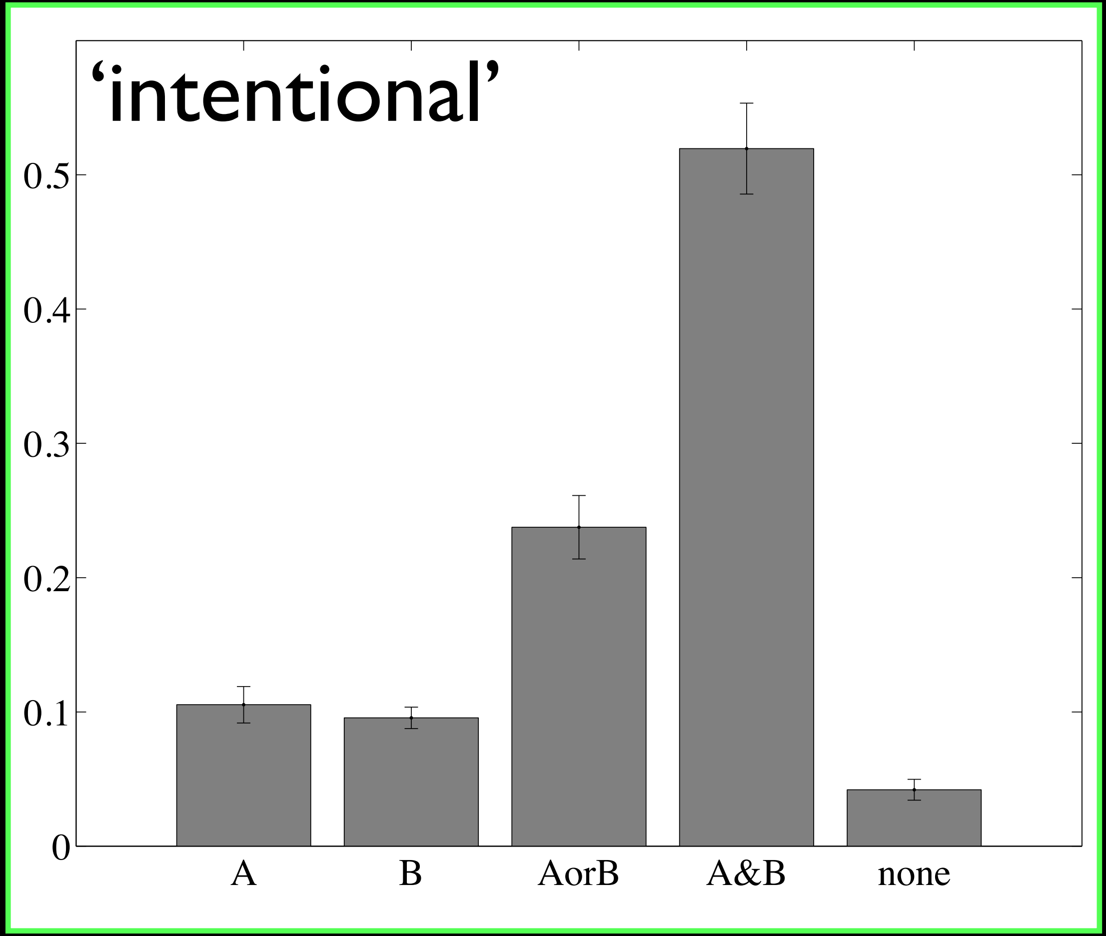

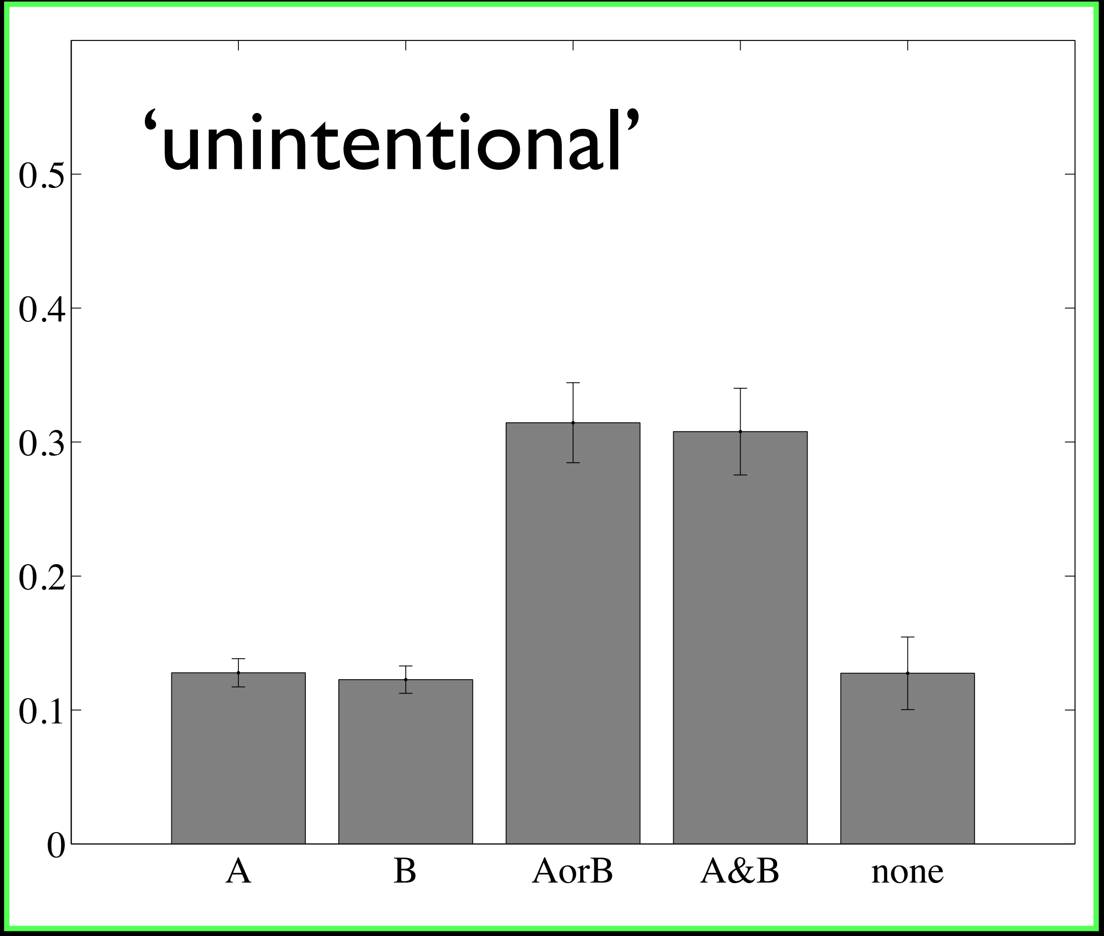

N=17

Expt 3: prior knowledge

- Effect of prior knowledge. (And in-lab replication.)

"Your coworker drinks a yellow liquid and pours a blue liquid on the flowers."

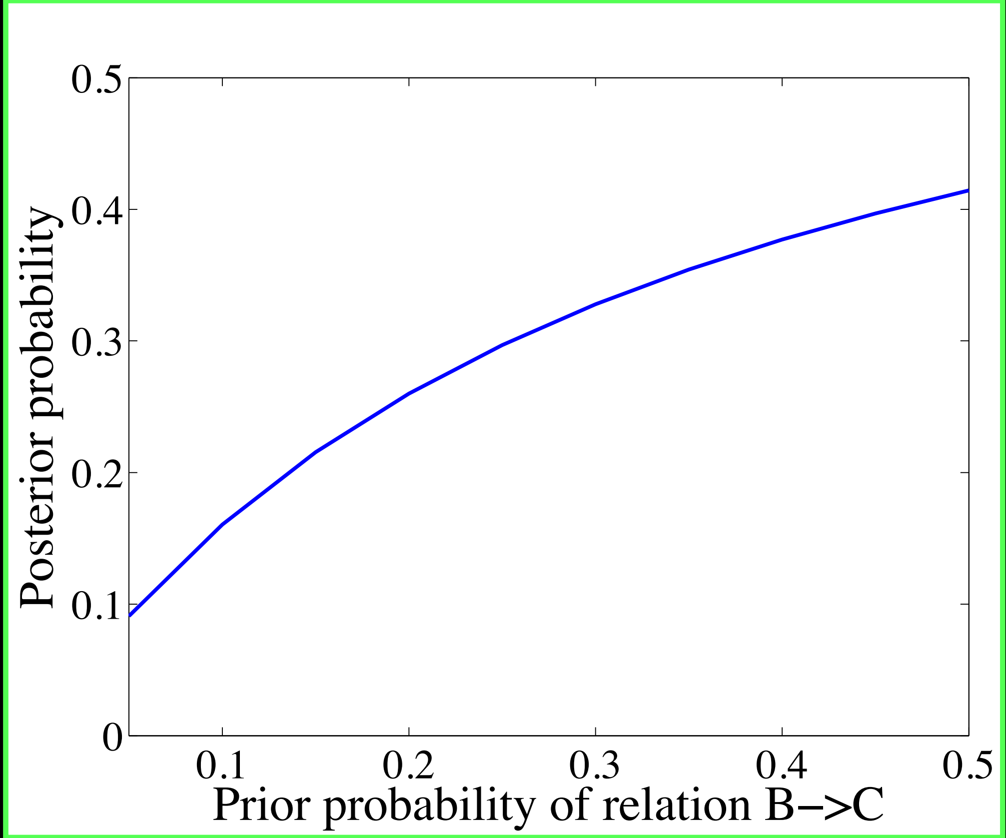

- Elicit prior plausibility judgements.
- Mean bet on 'A&B' vs. prior:

N=15

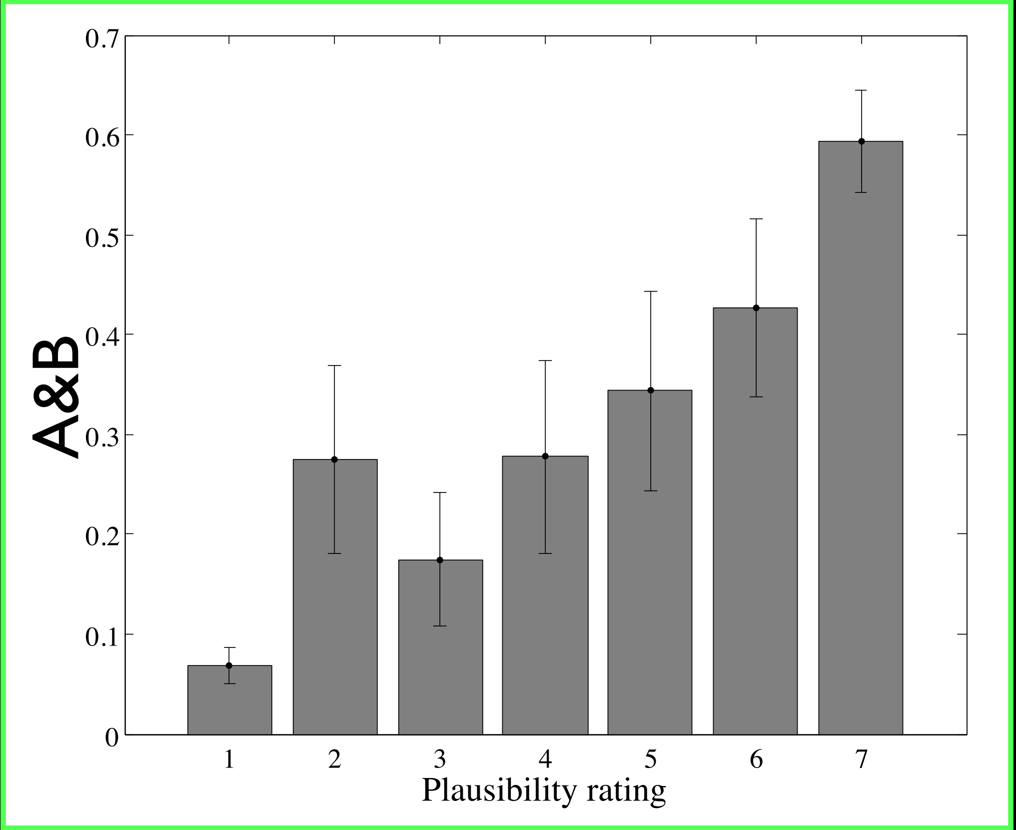

Goals?

- In this example goals were simple state features..

```
(define (goal? state)(light? state))
```

- This can be extended to more complex goals.
  - For example social goals

```
(define (helped? state) ((friends-goal state) state))
```

  (see Tenenbaum tomorrow)...

  - What happens when two agents have goals involving each other (and know this)?

Outline

- Theory of mind and learning from others' actions.
- Multi-agent reasoning: coordination games.
- Communicating with natural signs: intuitive pedagogy.
- Communicating with arbitrary signs: natural language.

Coordination

Alice and Bob arrange to meet at "the bar".
Each later realizes they didn't agree on *which* bar.
They must guess where to meet.

- Coordination games (Schelling, 1960; Clark, 1996; etc) involve partners reasoning about each other, without any communication.
- Model this as social cognition?

Coordination

```
(define (location)
  (if (flip .55) ‘good-bar ‘bad-bar))

(define (bob)
  (location))

(define (alice)
  (location))
```

Coordination

```
(define (location)
  (if (flip .55) 'good-bar 'bad-bar))

(define (bob)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice))))

(define (alice)
  (location))
```

Where to go,
To meet Alice?

Coordination

```
(define (location)
  (if (flip .55) ‘good-bar ‘bad-bar))

(define (bob)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice))))

(define (alice)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (bob))))
```

Where to go,
To meet Alice?

Where to go,
To meet Bob?

Coordination

```
(define (location)
  (if (flip .55) 'good-bar 'bad-bar))

(define (bob)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice))))

(define (alice)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (bob))))
```

Where to go,
To meet Alice?

Where to go,
To meet Bob?

Coordination

```
(define (location)
  (if (flip .55) ‘good-bar ‘bad-bar))

(define (bob depth)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice (- depth 1)))))

(define (alice depth)
  (query
    (define go-to (location))
    go-to
    (or (= depth 0)
        (equal? go-to (bob depth)))))
```

Coordination

```
(define (location)
  (if (flip .55) ‘good-bar ‘bad-bar))

(define (bob depth)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice (- depth

(define (alice depth)
  (query
    (define go-to (location))
    go-to
    (or (= depth 0)
         (equal? go-to (bob depth
```

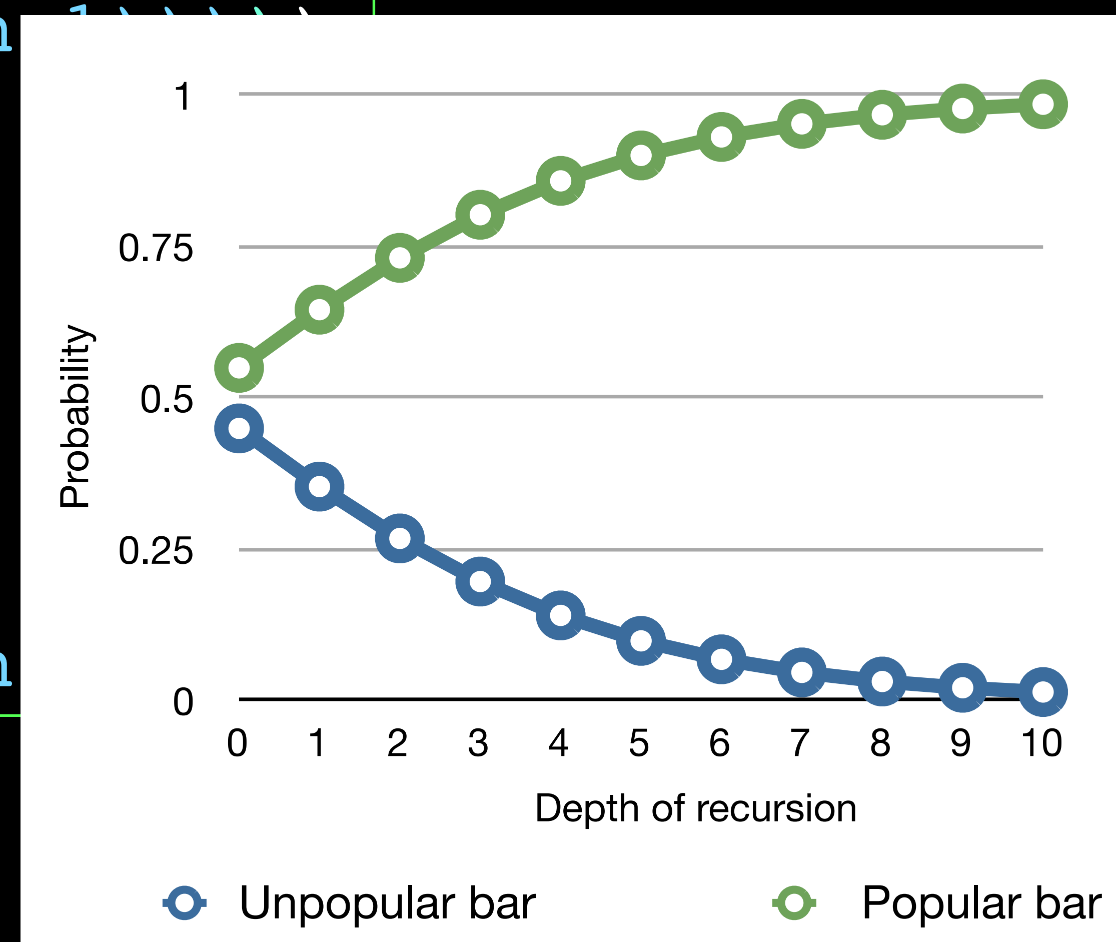

Coordination

```
(define (location)
  (if (flip .55) 'good-bar 'bad-bar))

(define (bob)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice))))

(define (alice)
  (query
    (define go-to (location))
    go-to
    (or (flip 0.2)
        (equal? go-to (bob)))))
```

Where to go,

(according to prior, or)

To meet Bob?

Coordination

```
(define (location)
  (if (flip .55) 'good-bar 'bad-bar))

(define (bob)
  (query
    (define go-to (location))
    go-to
    (equal? go-to (alice))))

(define (alice)
  (query
    (define go-to (location))
    go-to
    (or (flip 0.2)
        (equal? go-to (bob)))))
```

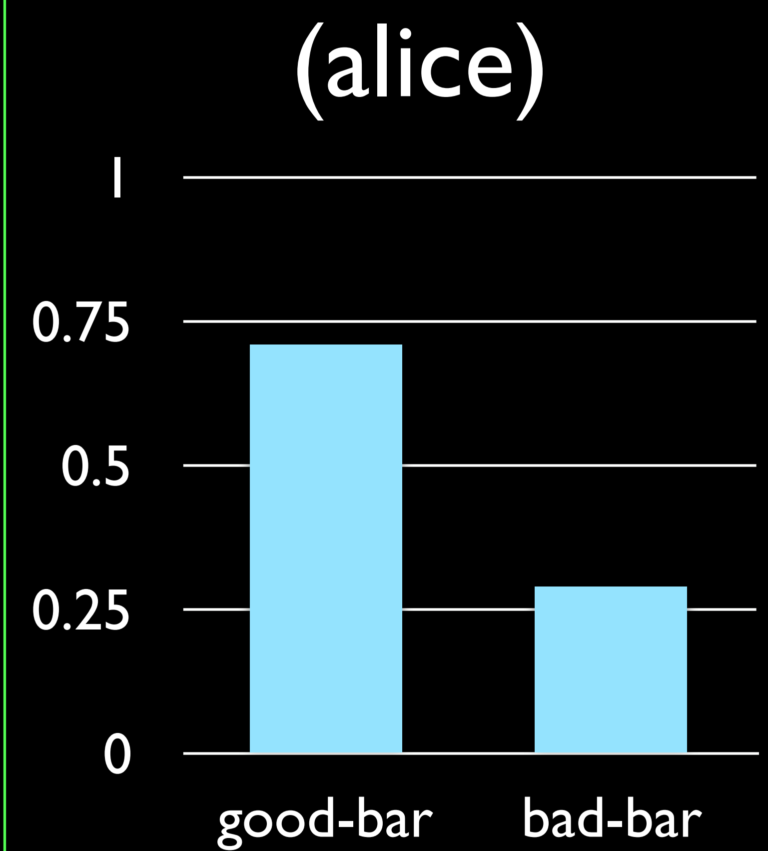

Where to go,

(according to prior, or)

To meet Bob?

Outline

- Theory of mind and learning from others' actions.
- Multi-agent reasoning: coordination games.
- Communicating with natural signs: intuitive pedagogy.
- Communicating with arbitrary signs: natural language.

Communication

- We can't always rely on simple coordination to arrange things -- there are too many options.
  - What if Alice and Bob hadn't said they'd meet at "the bar"?
- Instead we pass *signs* that help us to coordinate.
- *Natural signs* have meaning in the world.

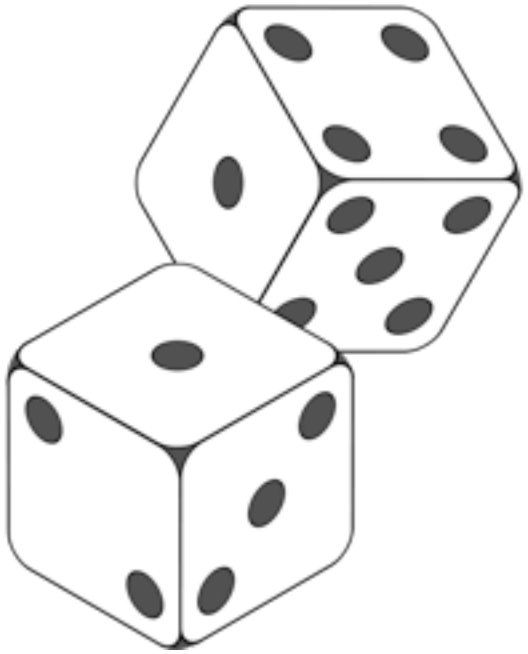

A:

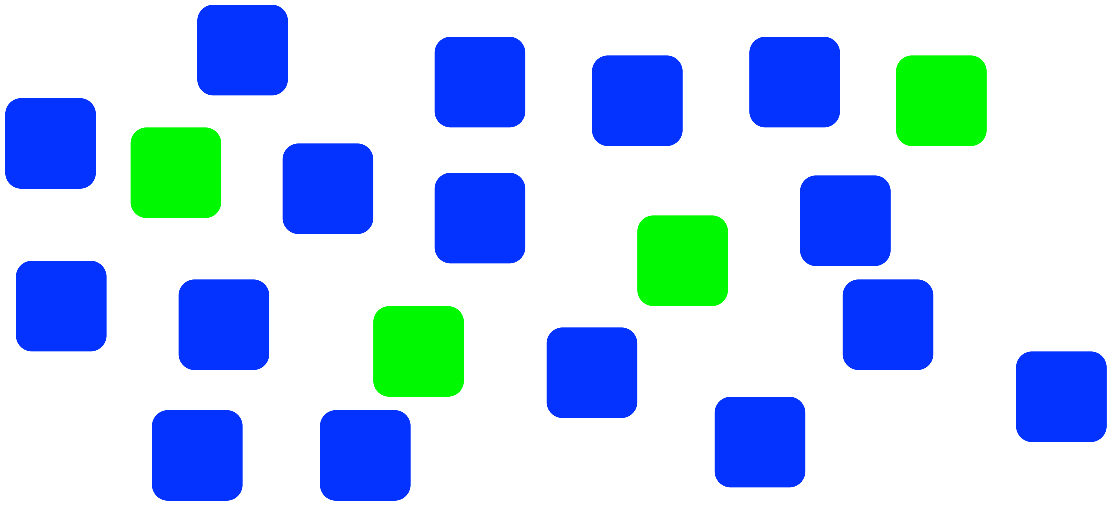

B:

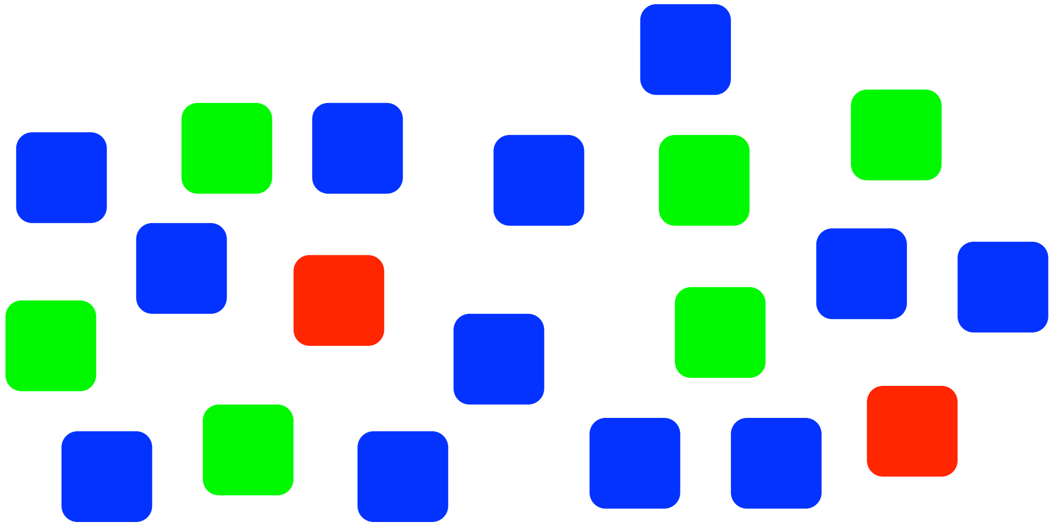

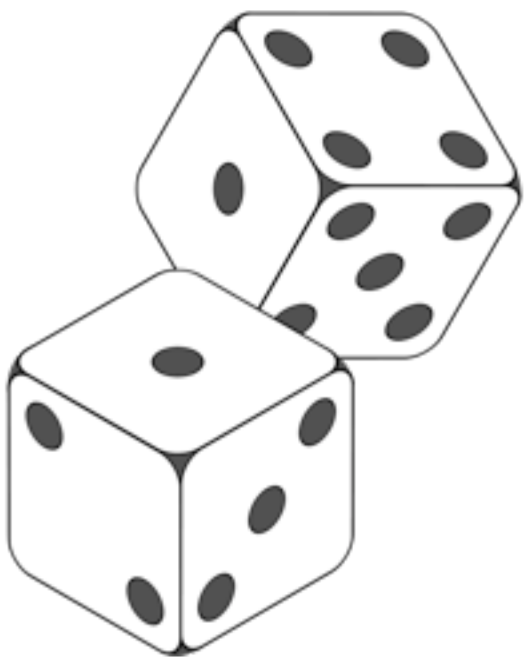

A:

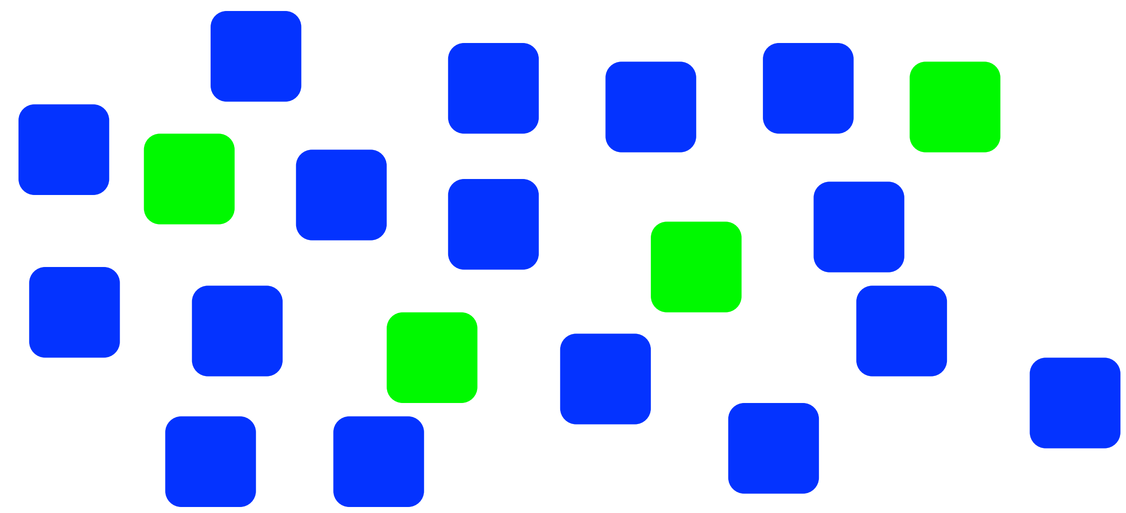

B:

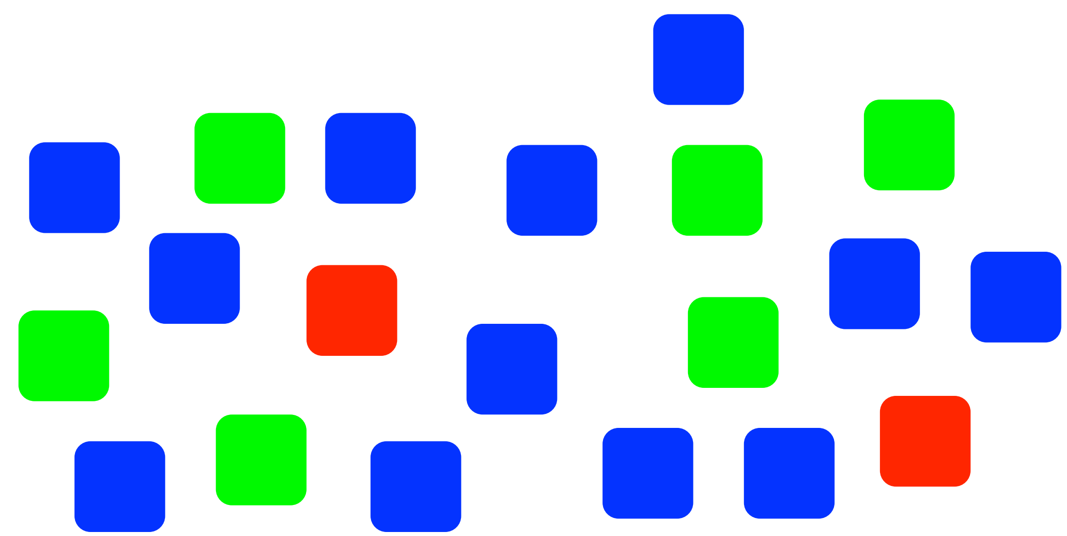

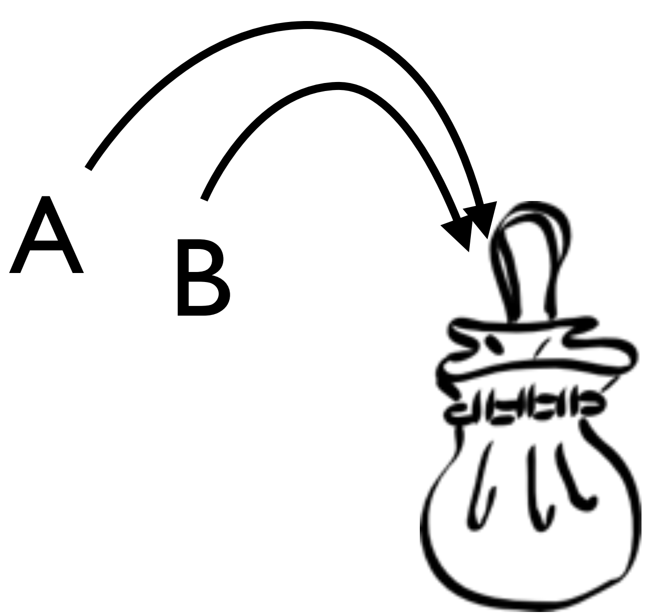

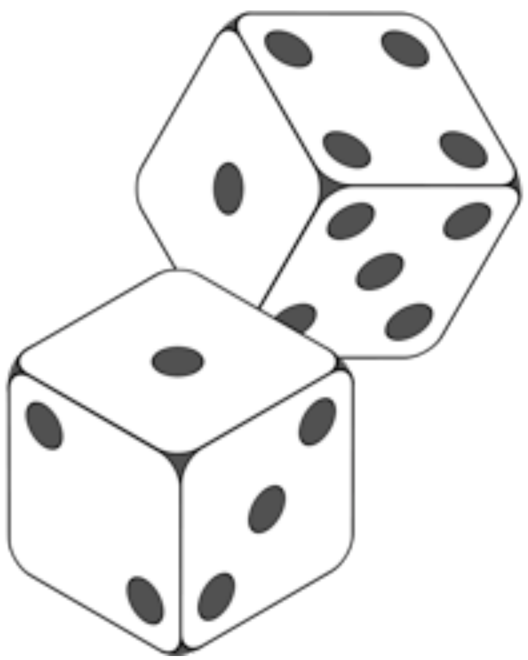

A:

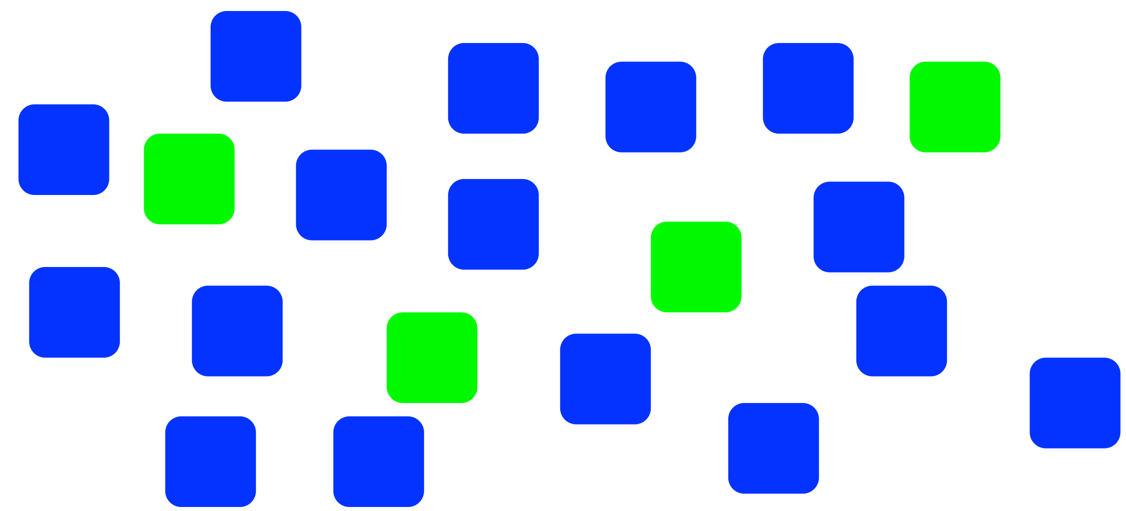

B:

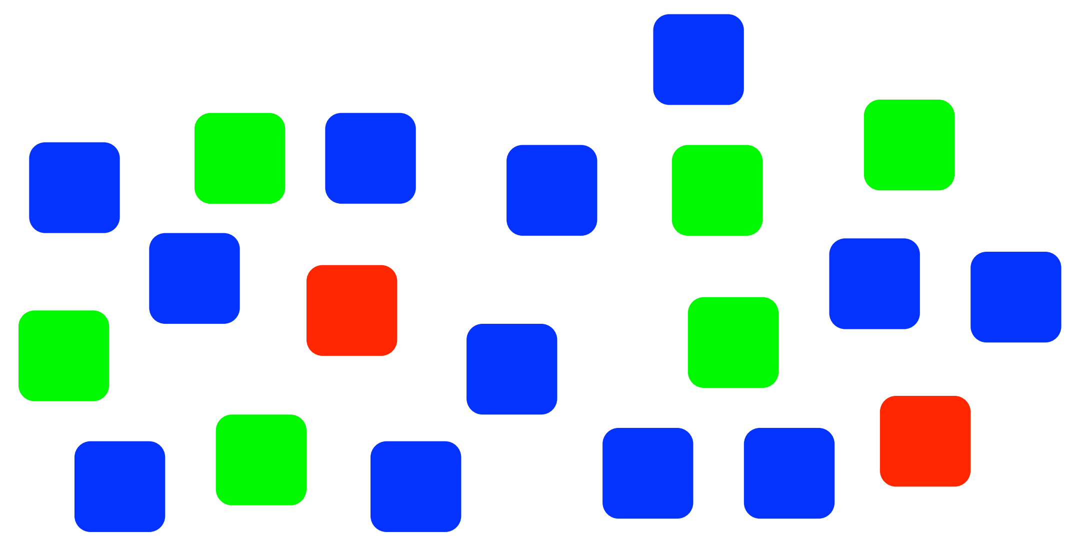

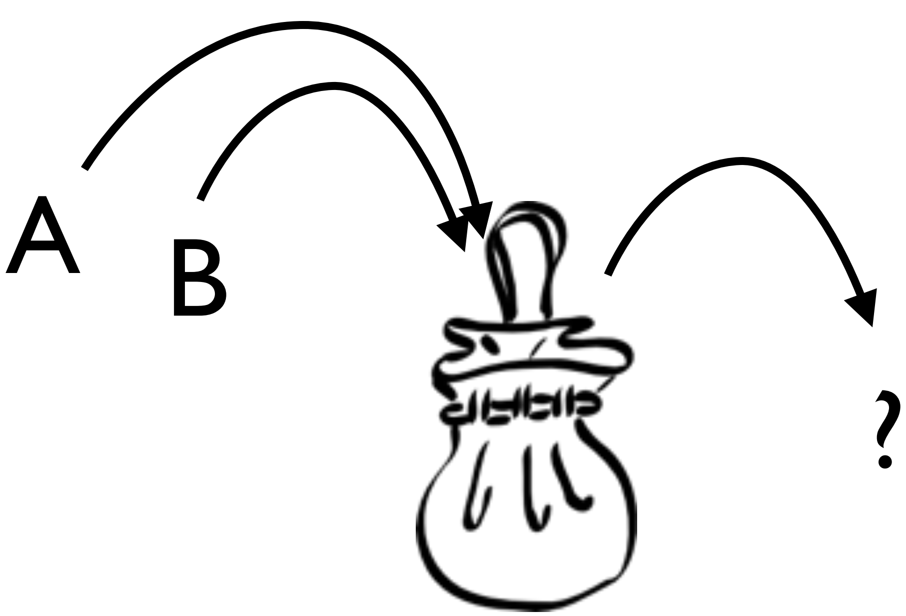

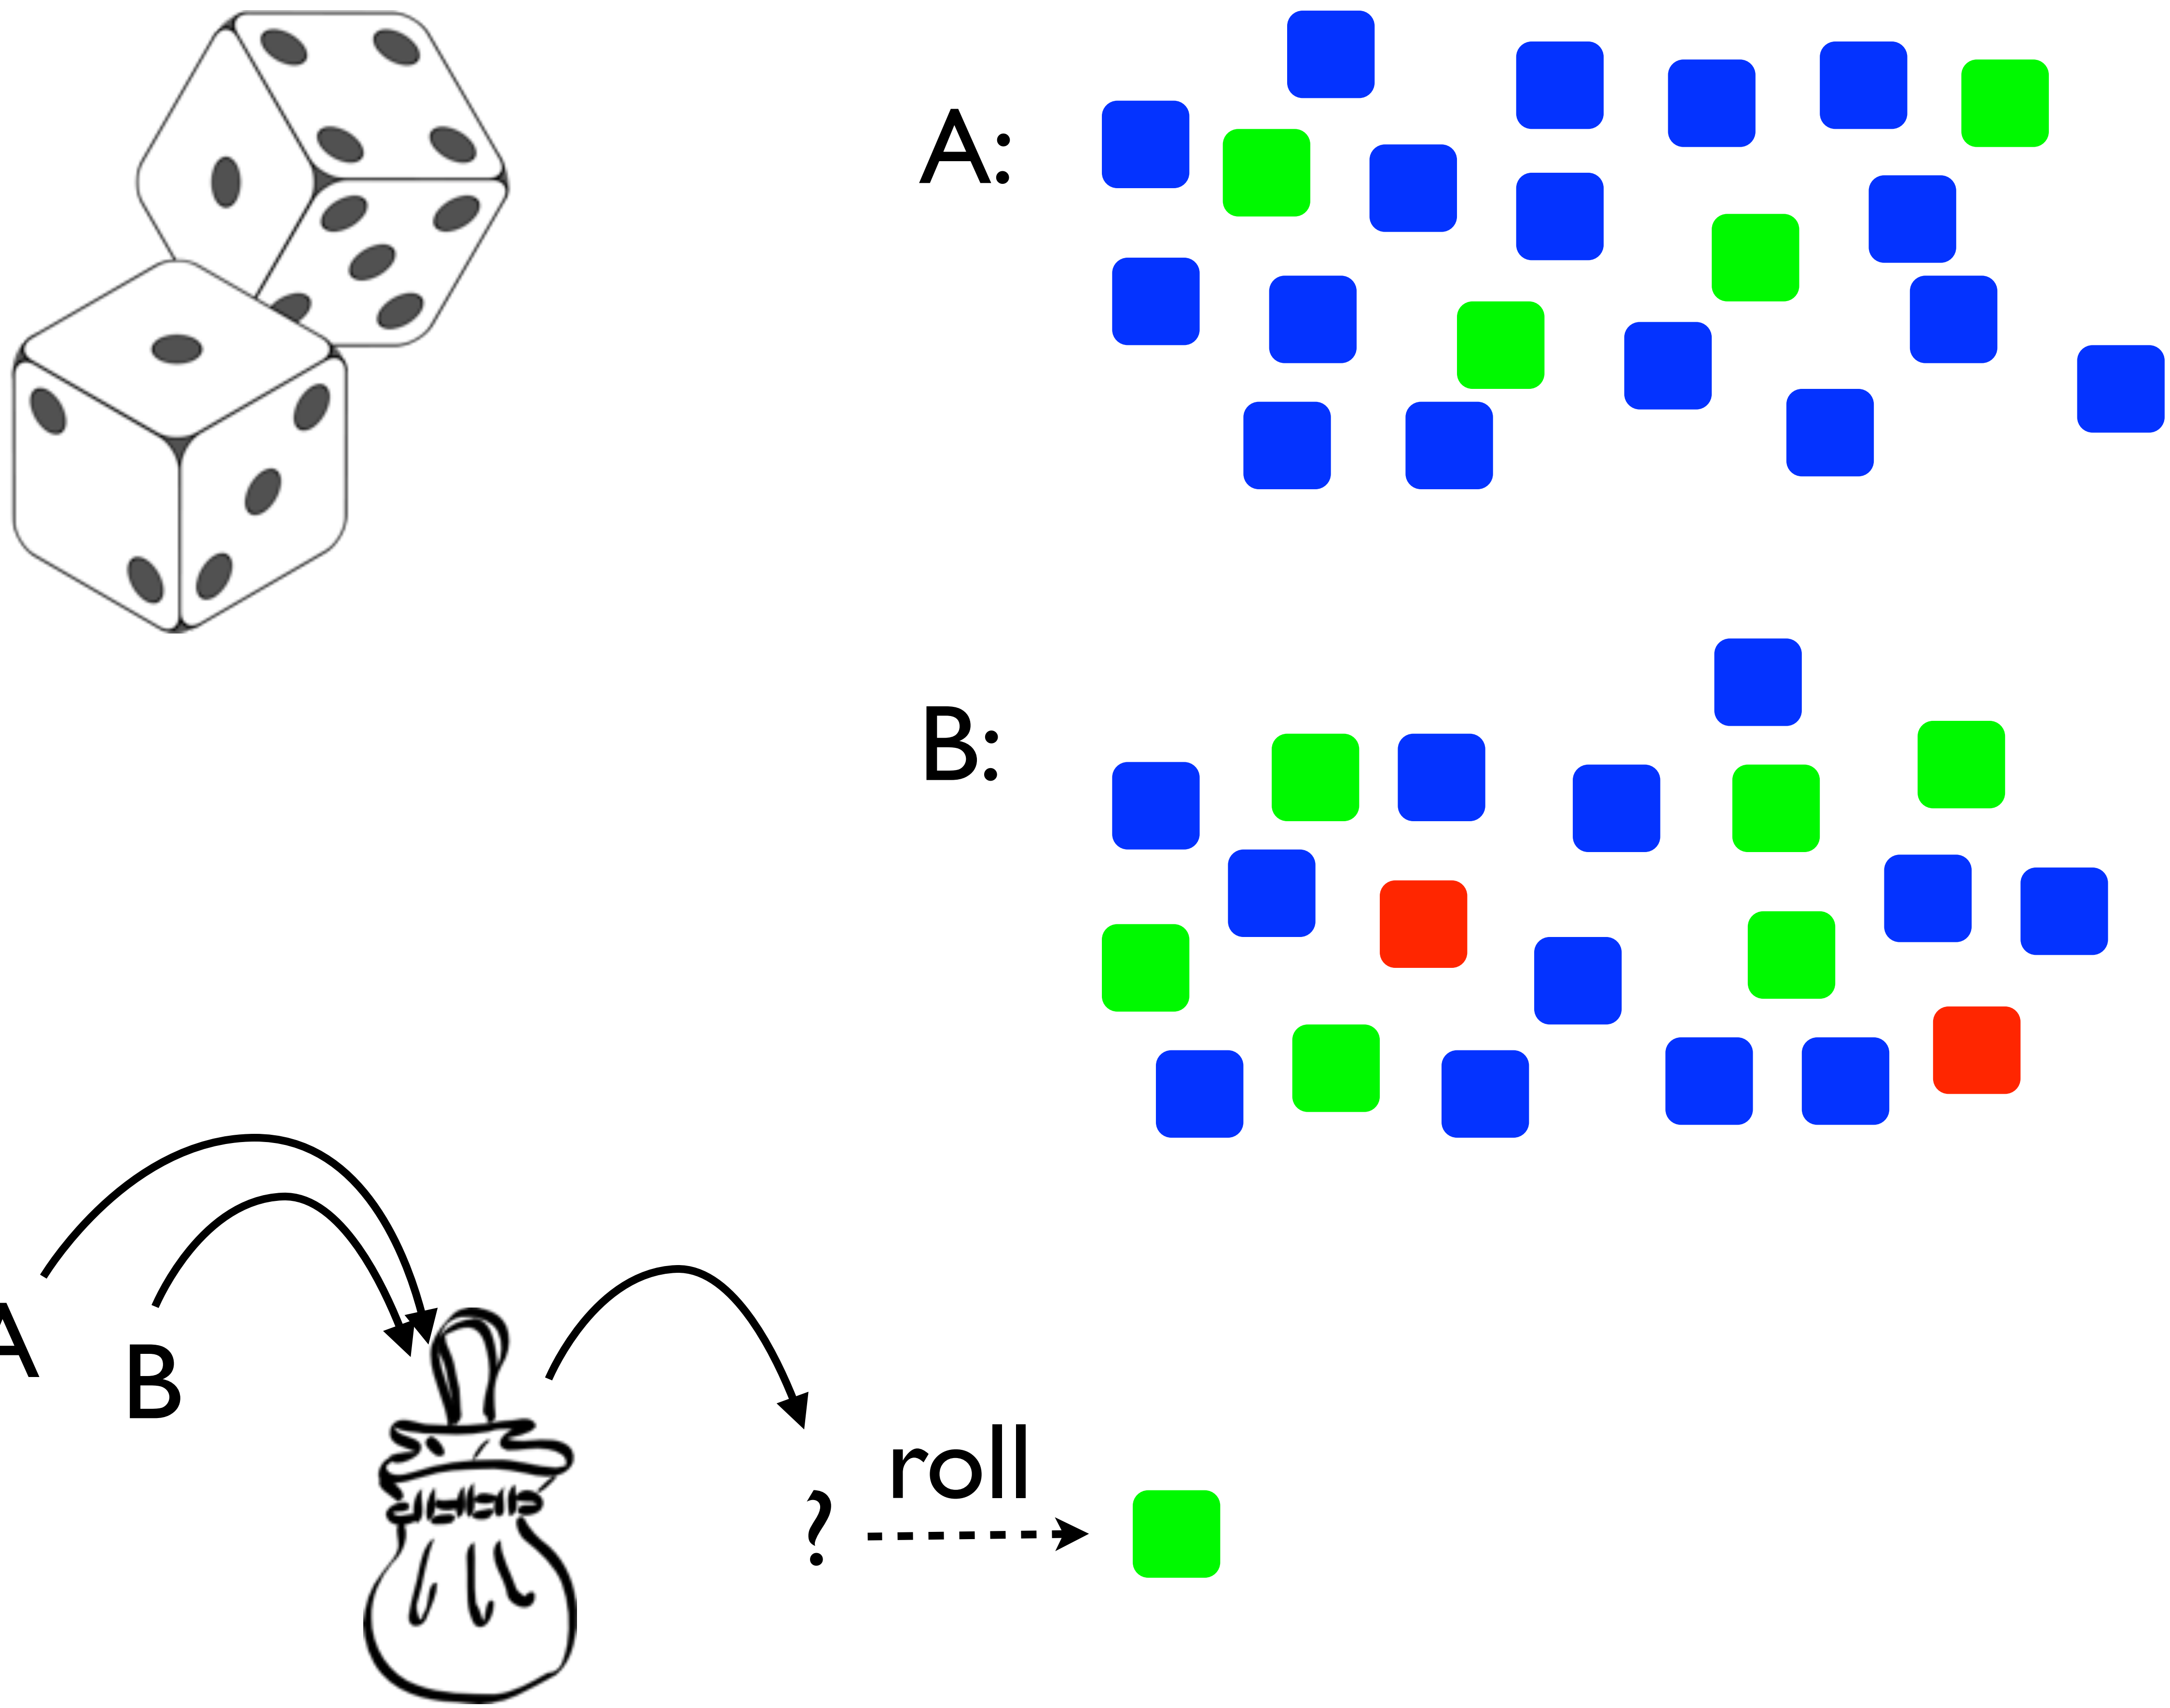

A:
B:
A
B
?
roll

Observation

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,

If it came up green?

Observation

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

(Observer)

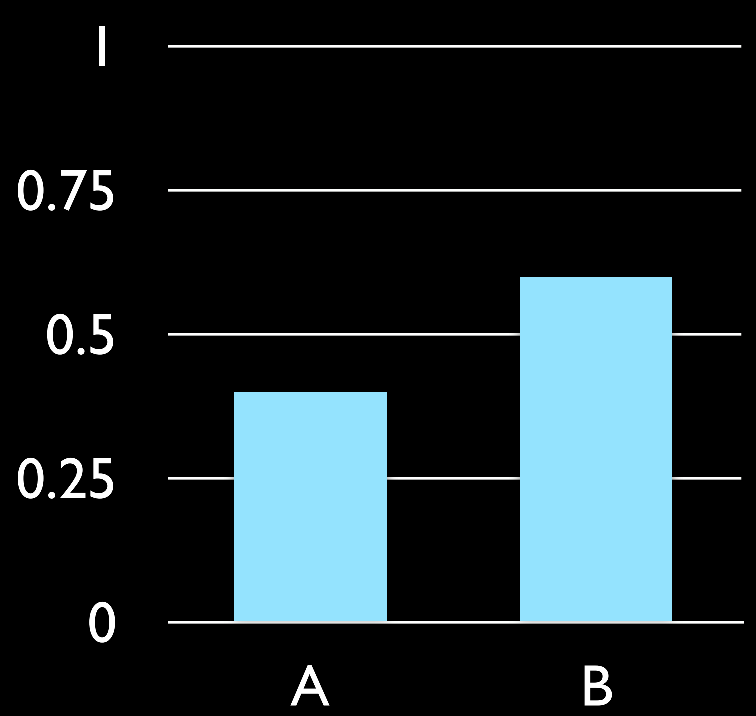

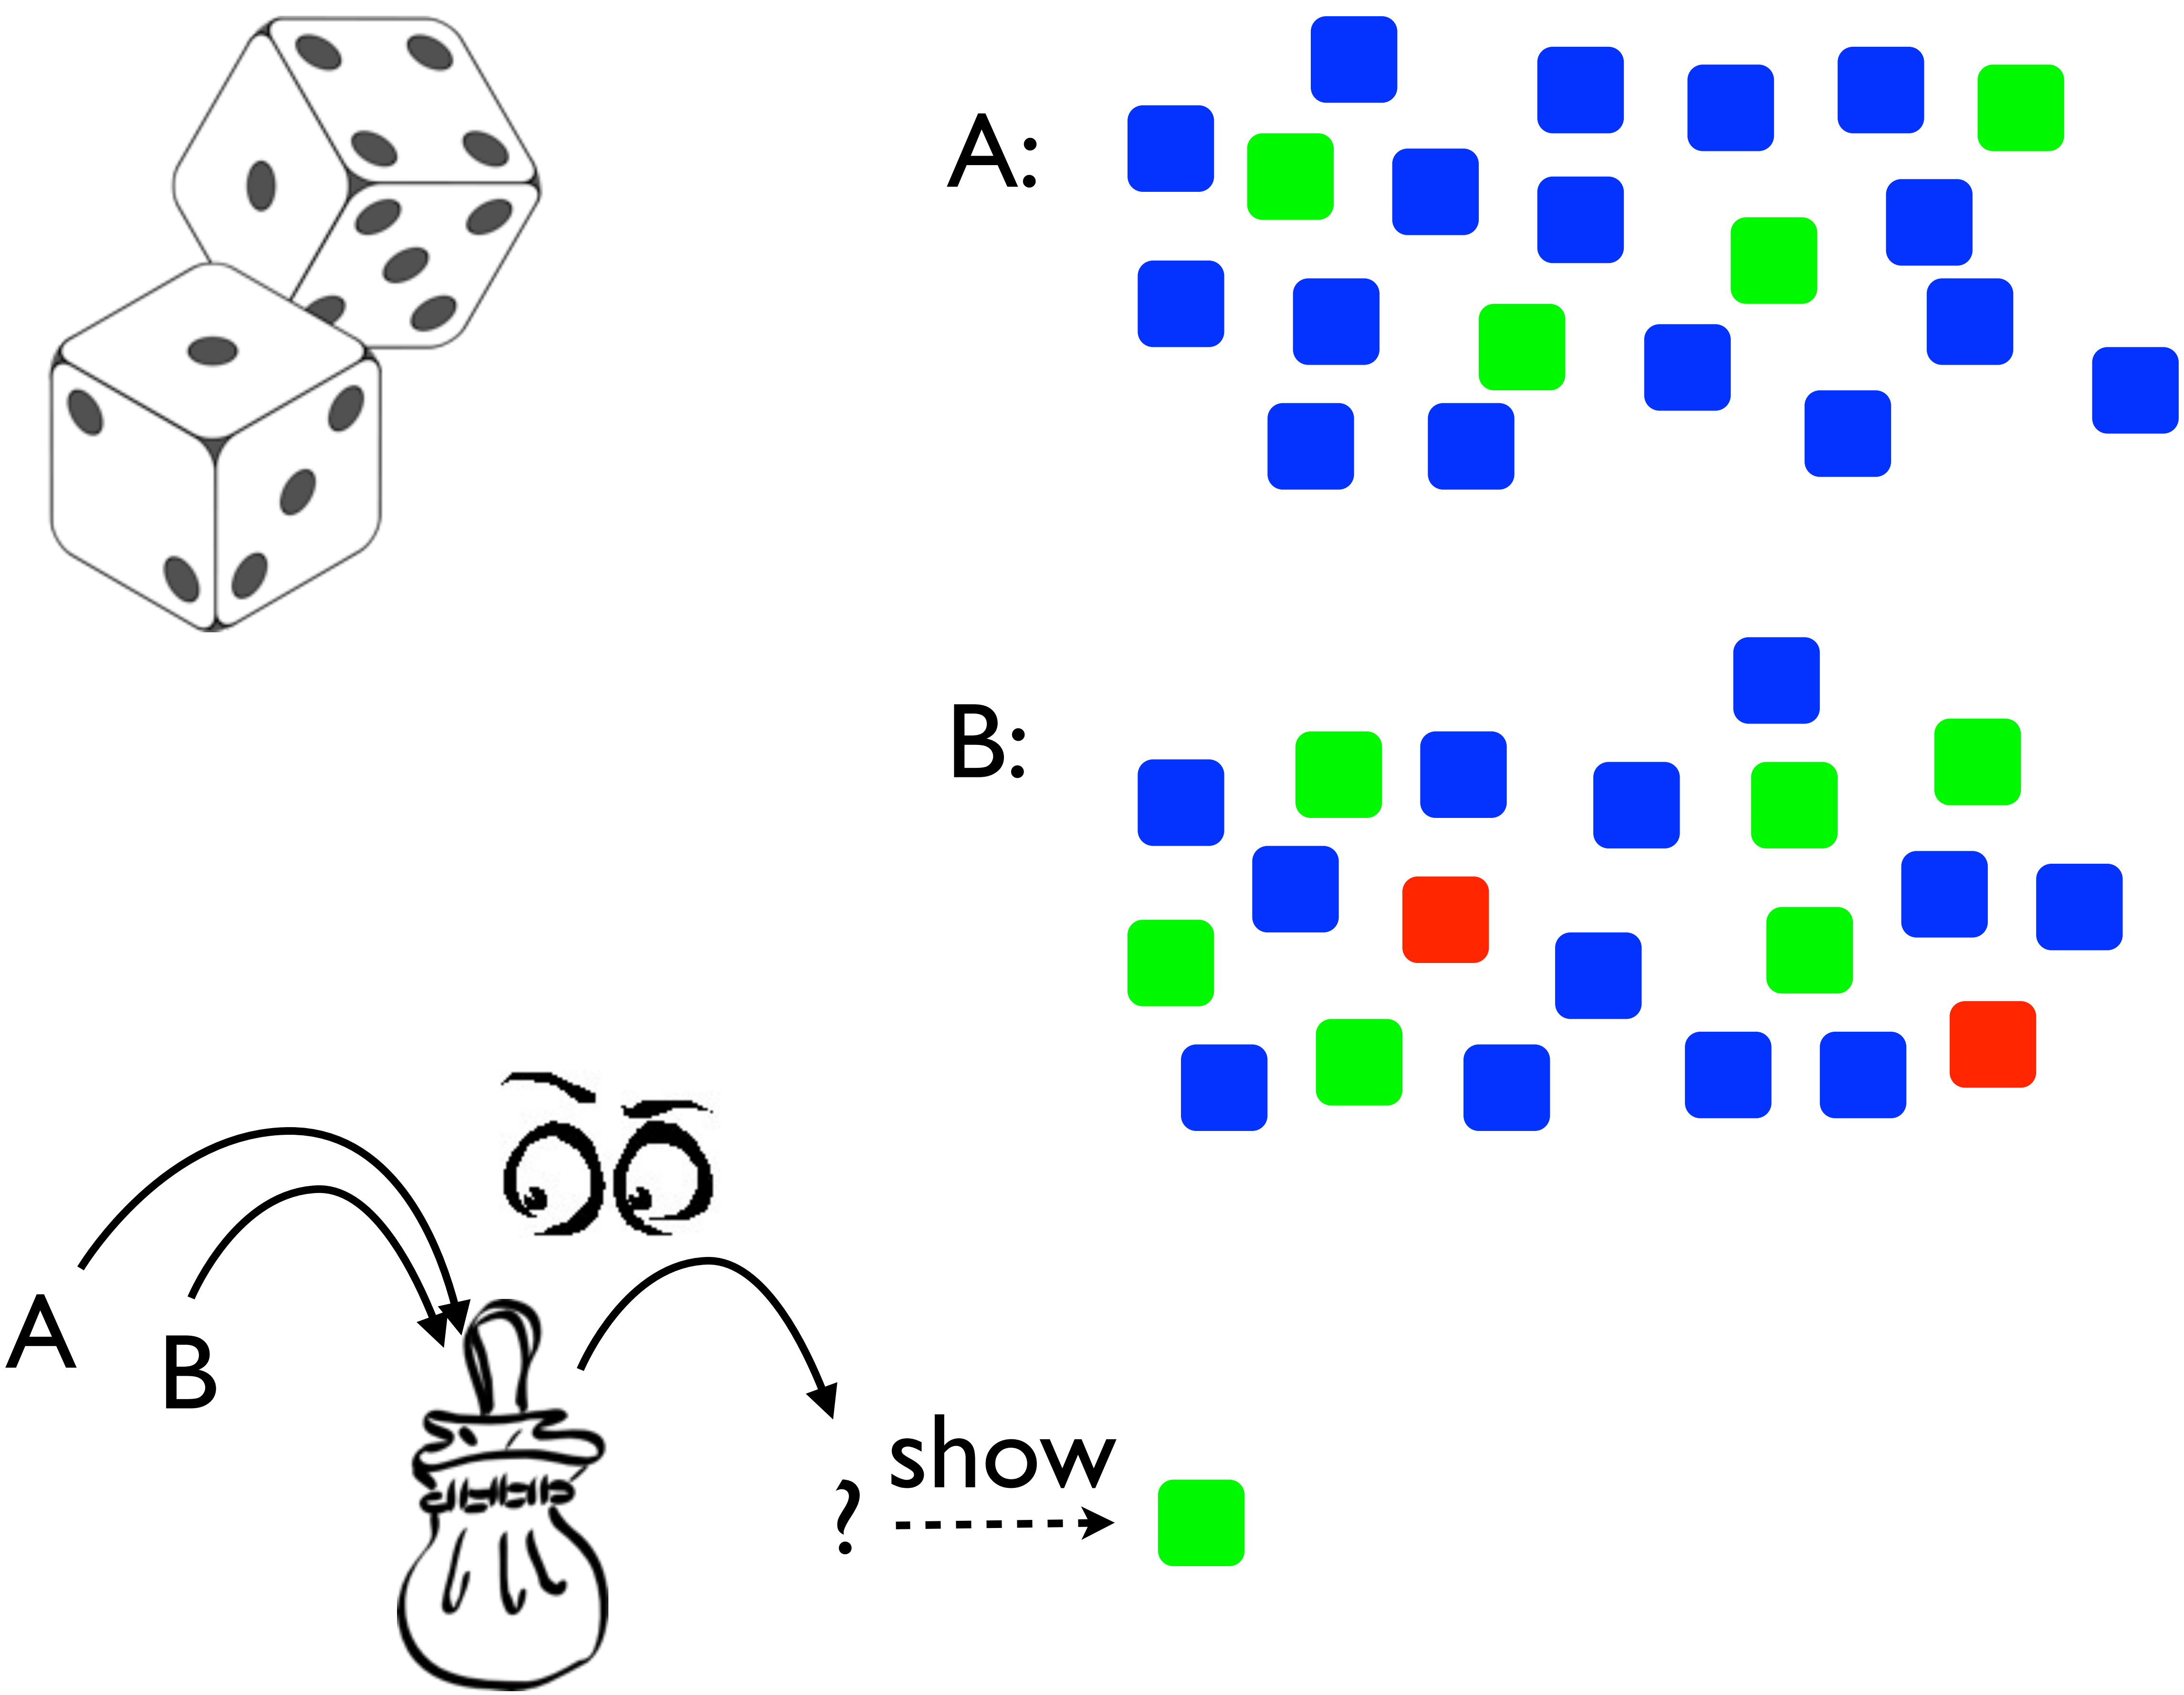
A:
B:
A
B
?
show

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,

If it came up green?

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

Which die is it,

If it came up green?

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

Which die is it,
If it came up green?

What should I show,

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

What should I show,
So that the observer will
infer the correct die?

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

What should I show,
So that the observer will
infer the correct die?

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (equal? side (speaker die))))
```

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

What should I show,
So that the observer will
infer the correct die?

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (equal? side (speaker die))))
```

Which die is it,

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (observer side))))
```

What should I show,
So that the observer will
infer the correct die?

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (equal? side (speaker die))))
```

Which die is it,
If the speaker chose
this side?

Communication

```
(define (observer side)
  (query
   (define die (die-prior))
   die
   (equal? side (roll die))))
```

Which die is it,
If it came up green?

```
(define (speaker die)
  (query
   (define
   side
   (equal?
```

```
(define (l
  (query
   (define
   die
   (equal?
```

now,
ver will
ct die?

it,
chose

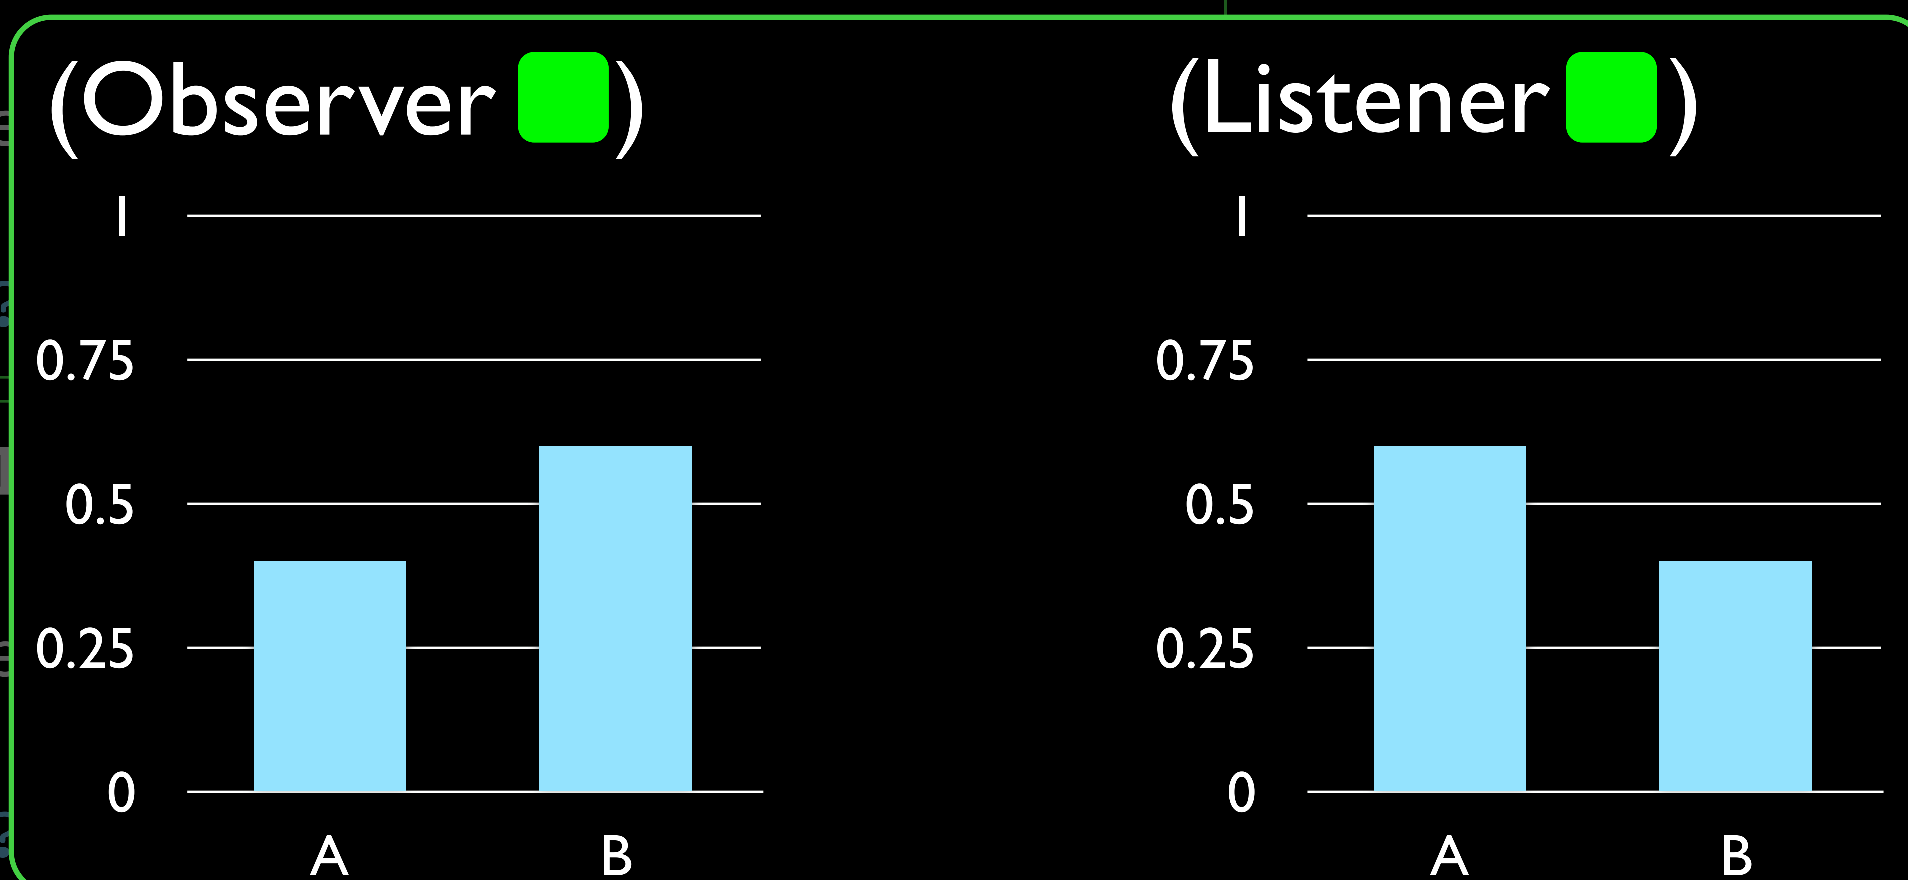

Communication

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (listener side))))
```

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (if (flip 0.2)
       (equal? side (roll die))
       (equal? side (speaker die)))))
```

Communication

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (listener side))))
```

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (if (flip 0.2)
       (equal? side (roll die))
       (equal? side (speaker die)))))
```

If this side is likely

(or)

the speaker chose this side?

Communication

```
(define (speaker die)
  (query
   (define side (roll die))
   side
   (equal? die (listener side))))
```

```
(define (listener side)
  (query
   (define die (die-prior))
   die
   (if (flip 0.2)
       (equal? side (roll die))
       (equal? side (speaker die)))))
```

(listener)

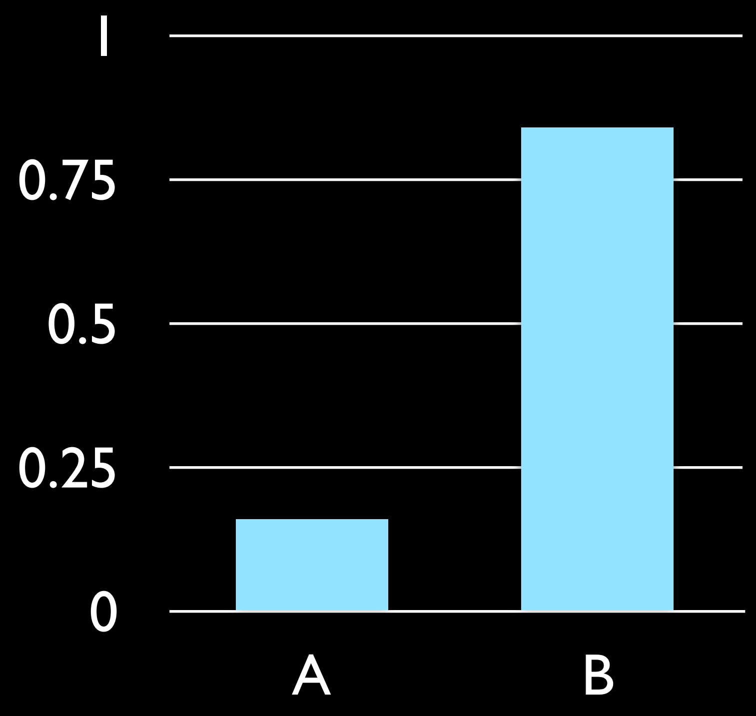

If this side is likely

(or)

the speaker chose this side?

Intuitive pedagogy

- This has been proposed as a model of natural pedagogy.
  (Shafto & Goodman, 2008;
  Shafto, Goodman, Griffiths, under review)
- Teaching games: have a *teacher* try to convey a *hypothesis* by sending *examples* to a *student*.

The Rectangle Game

Cf. Tenenbaum, 1999

The Rectangle Game

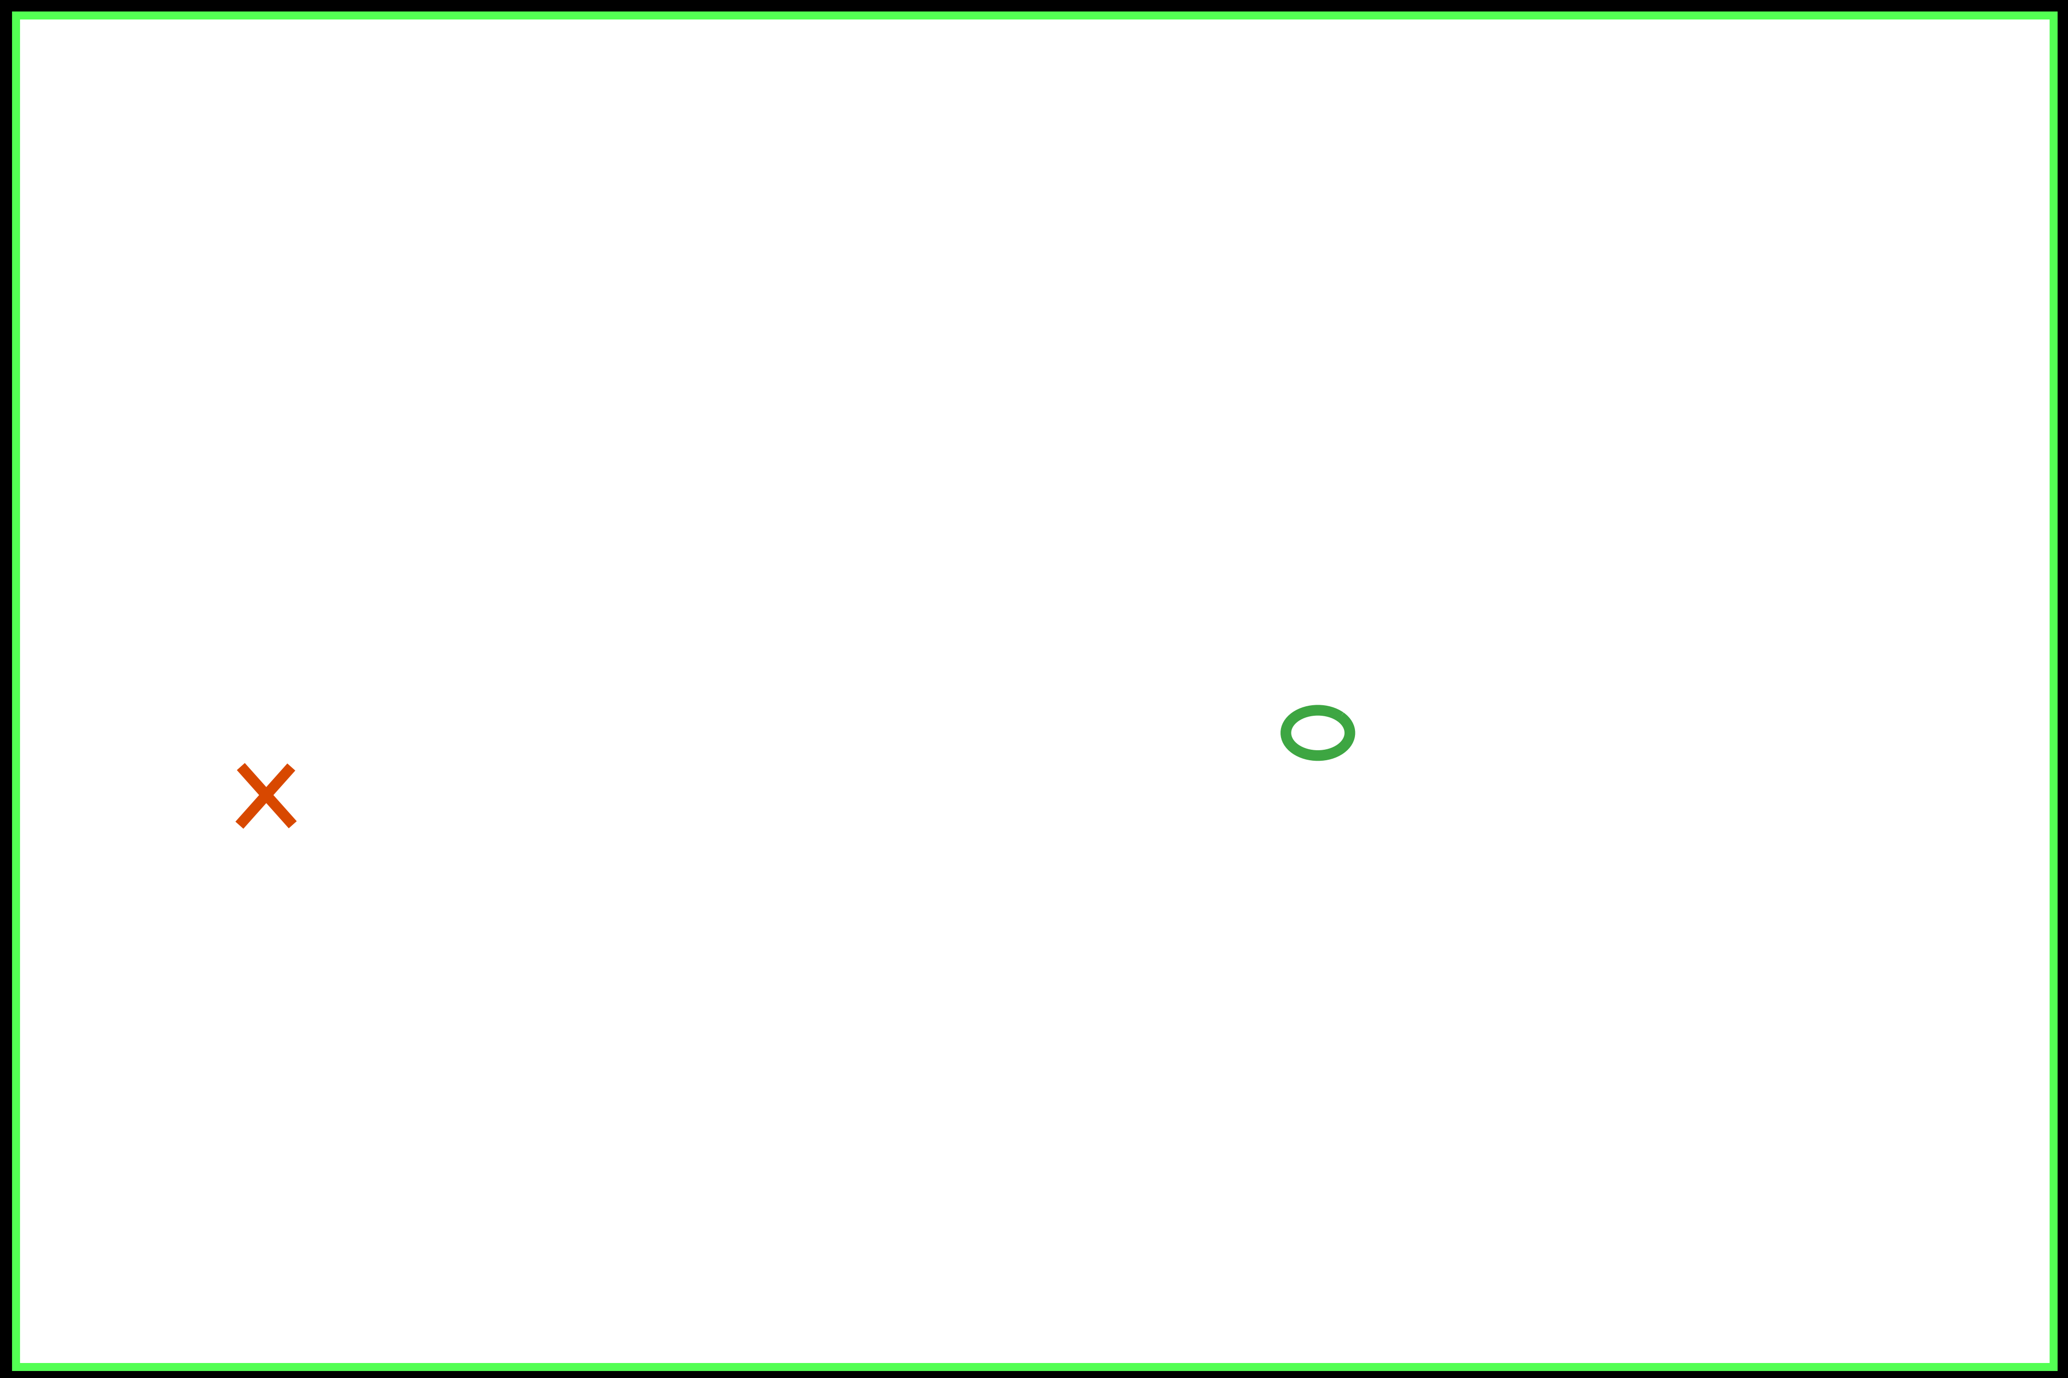

Cf. Tenenbaum, 1999

The Rectangle Game

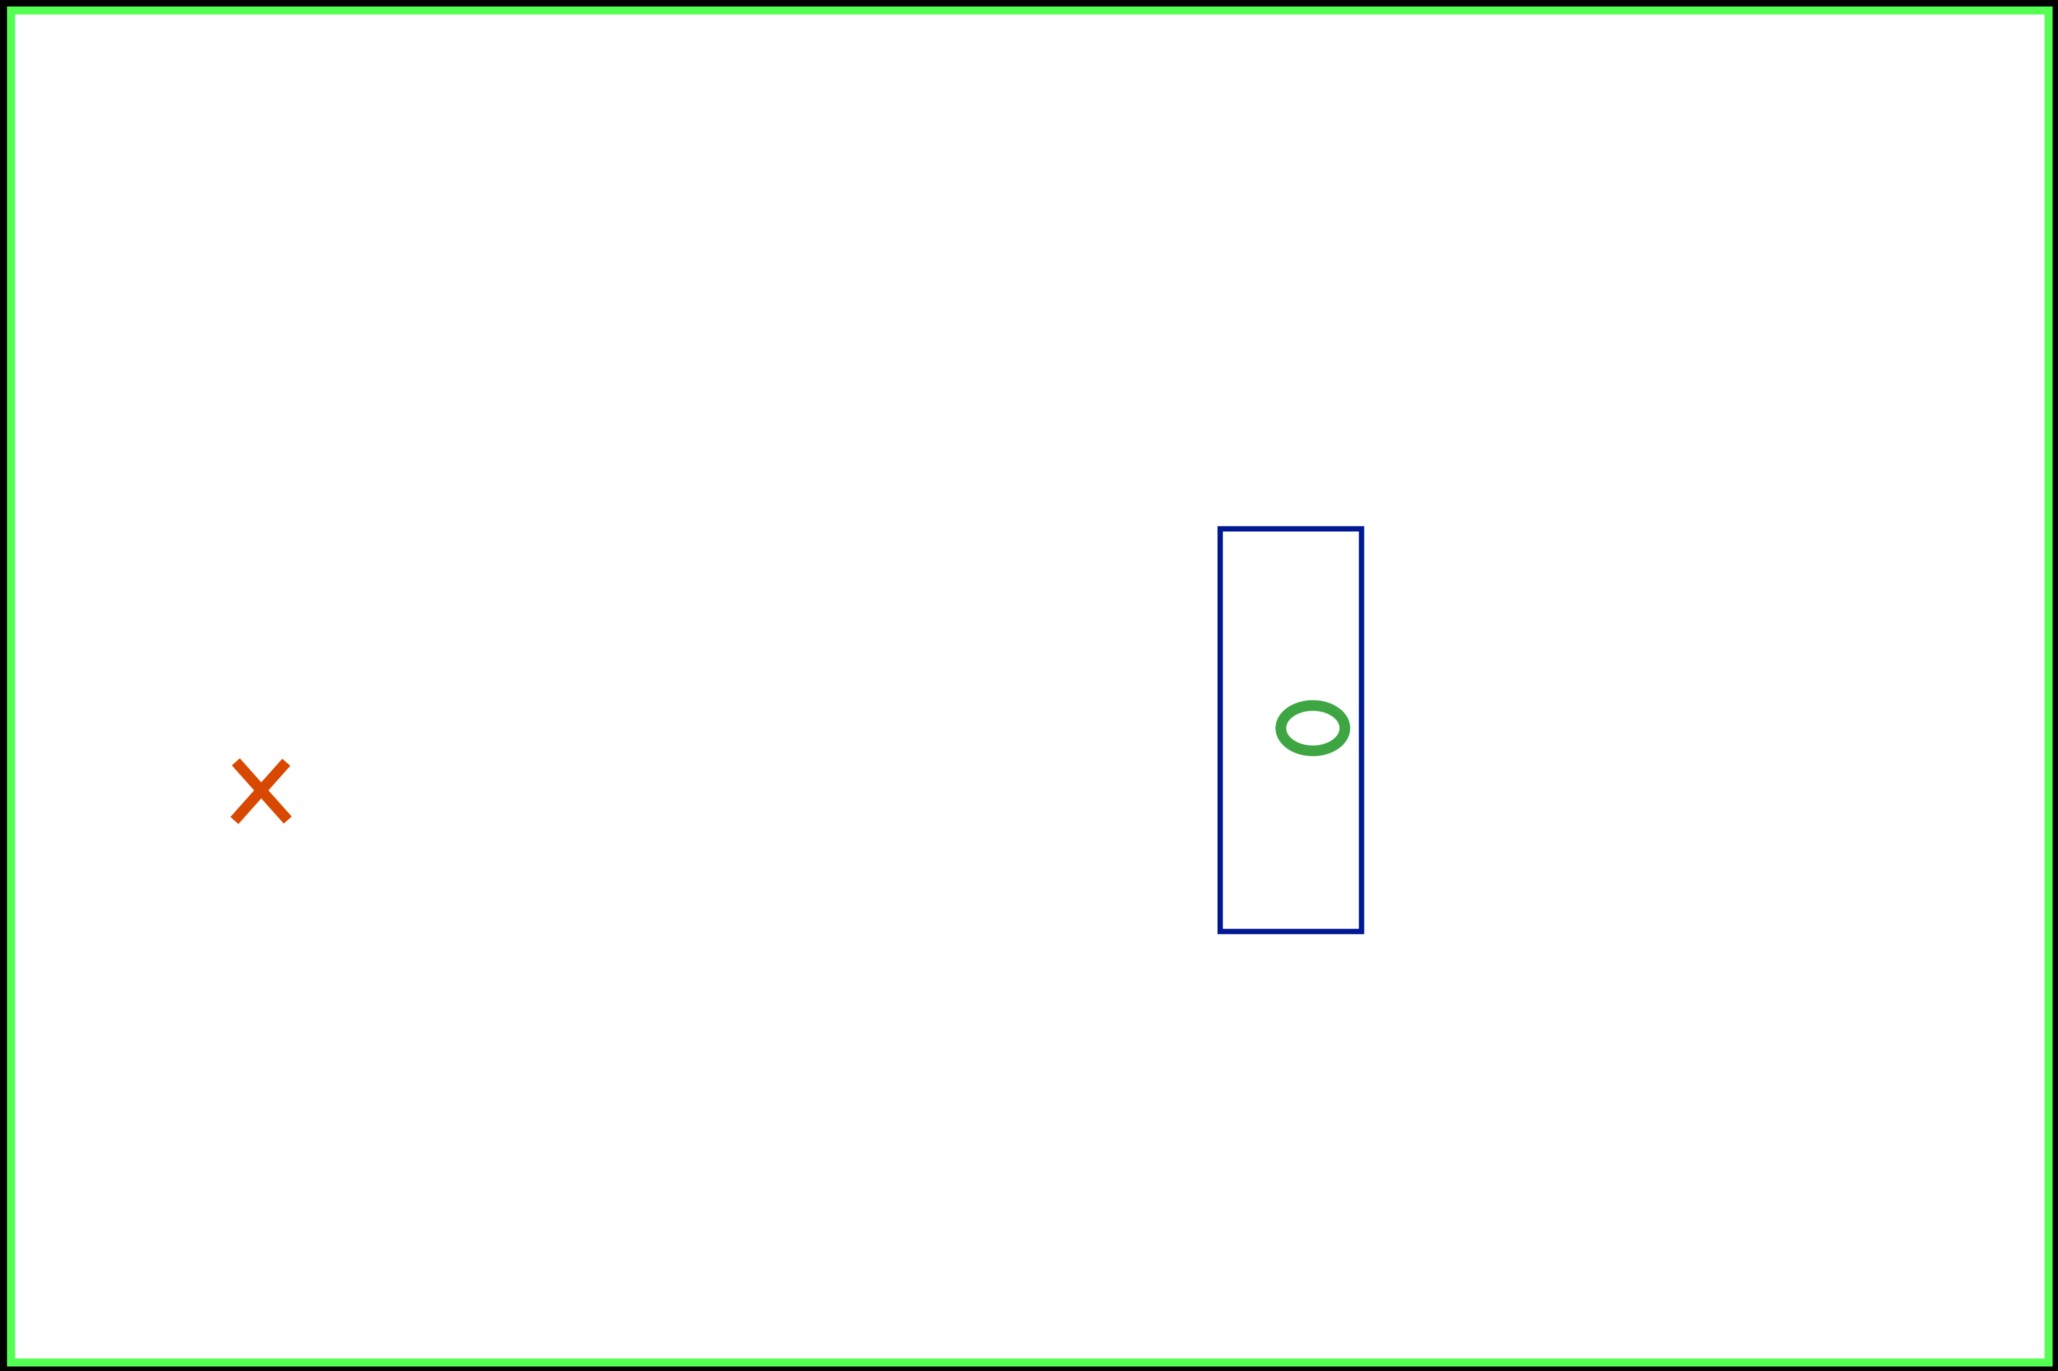

Cf. Tenenbaum, 1999

The Rectangle Game

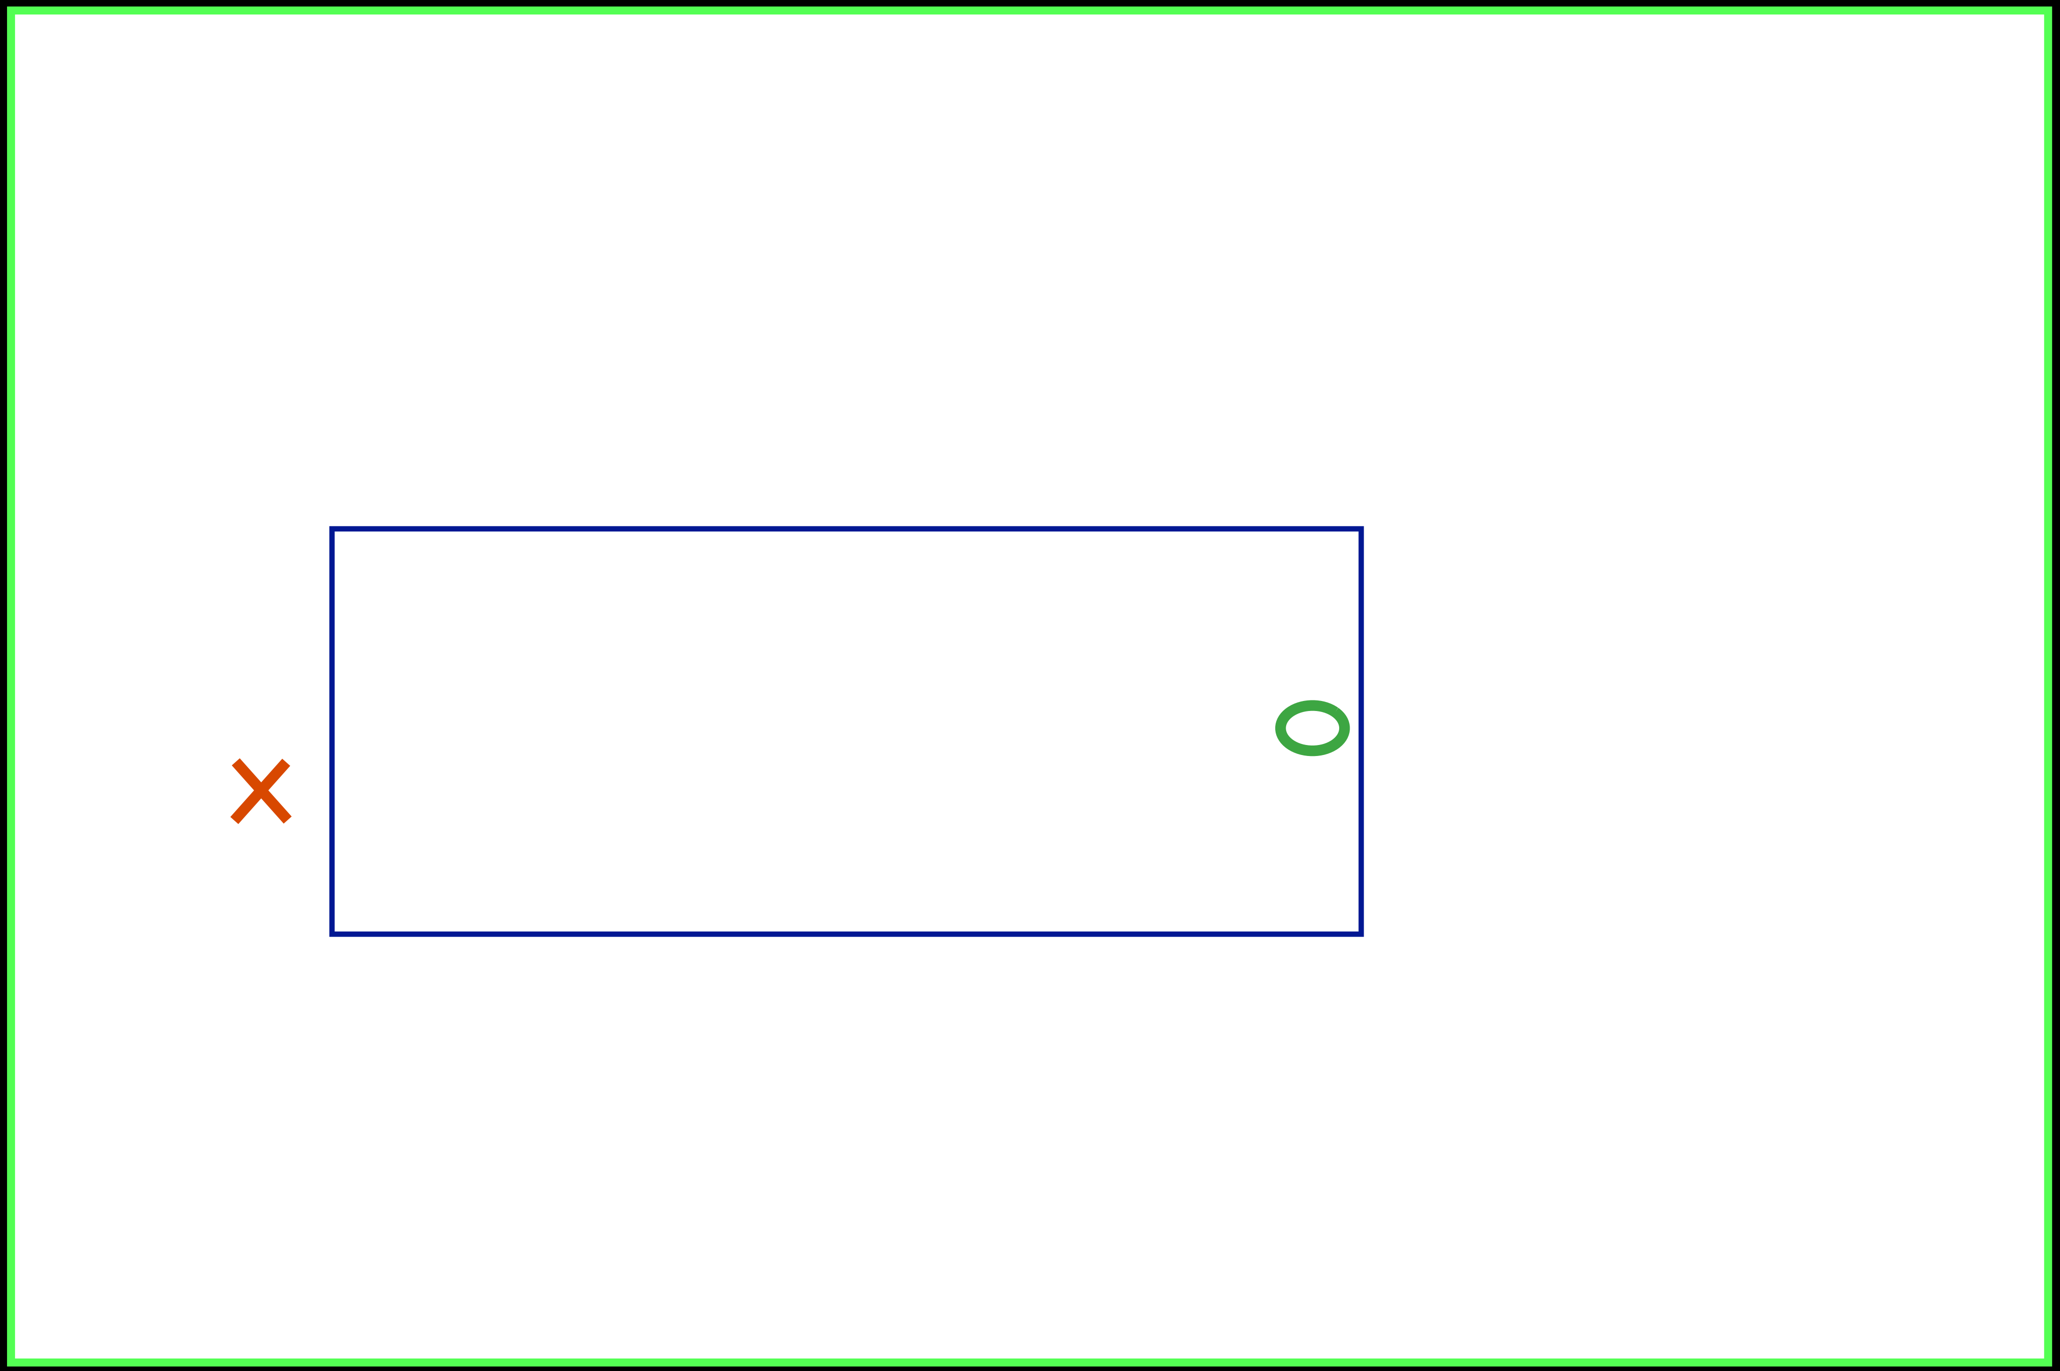

Cf. Tenenbaum, 1999

Learning Results

Examples: | Model best guess: | Learner guesses:

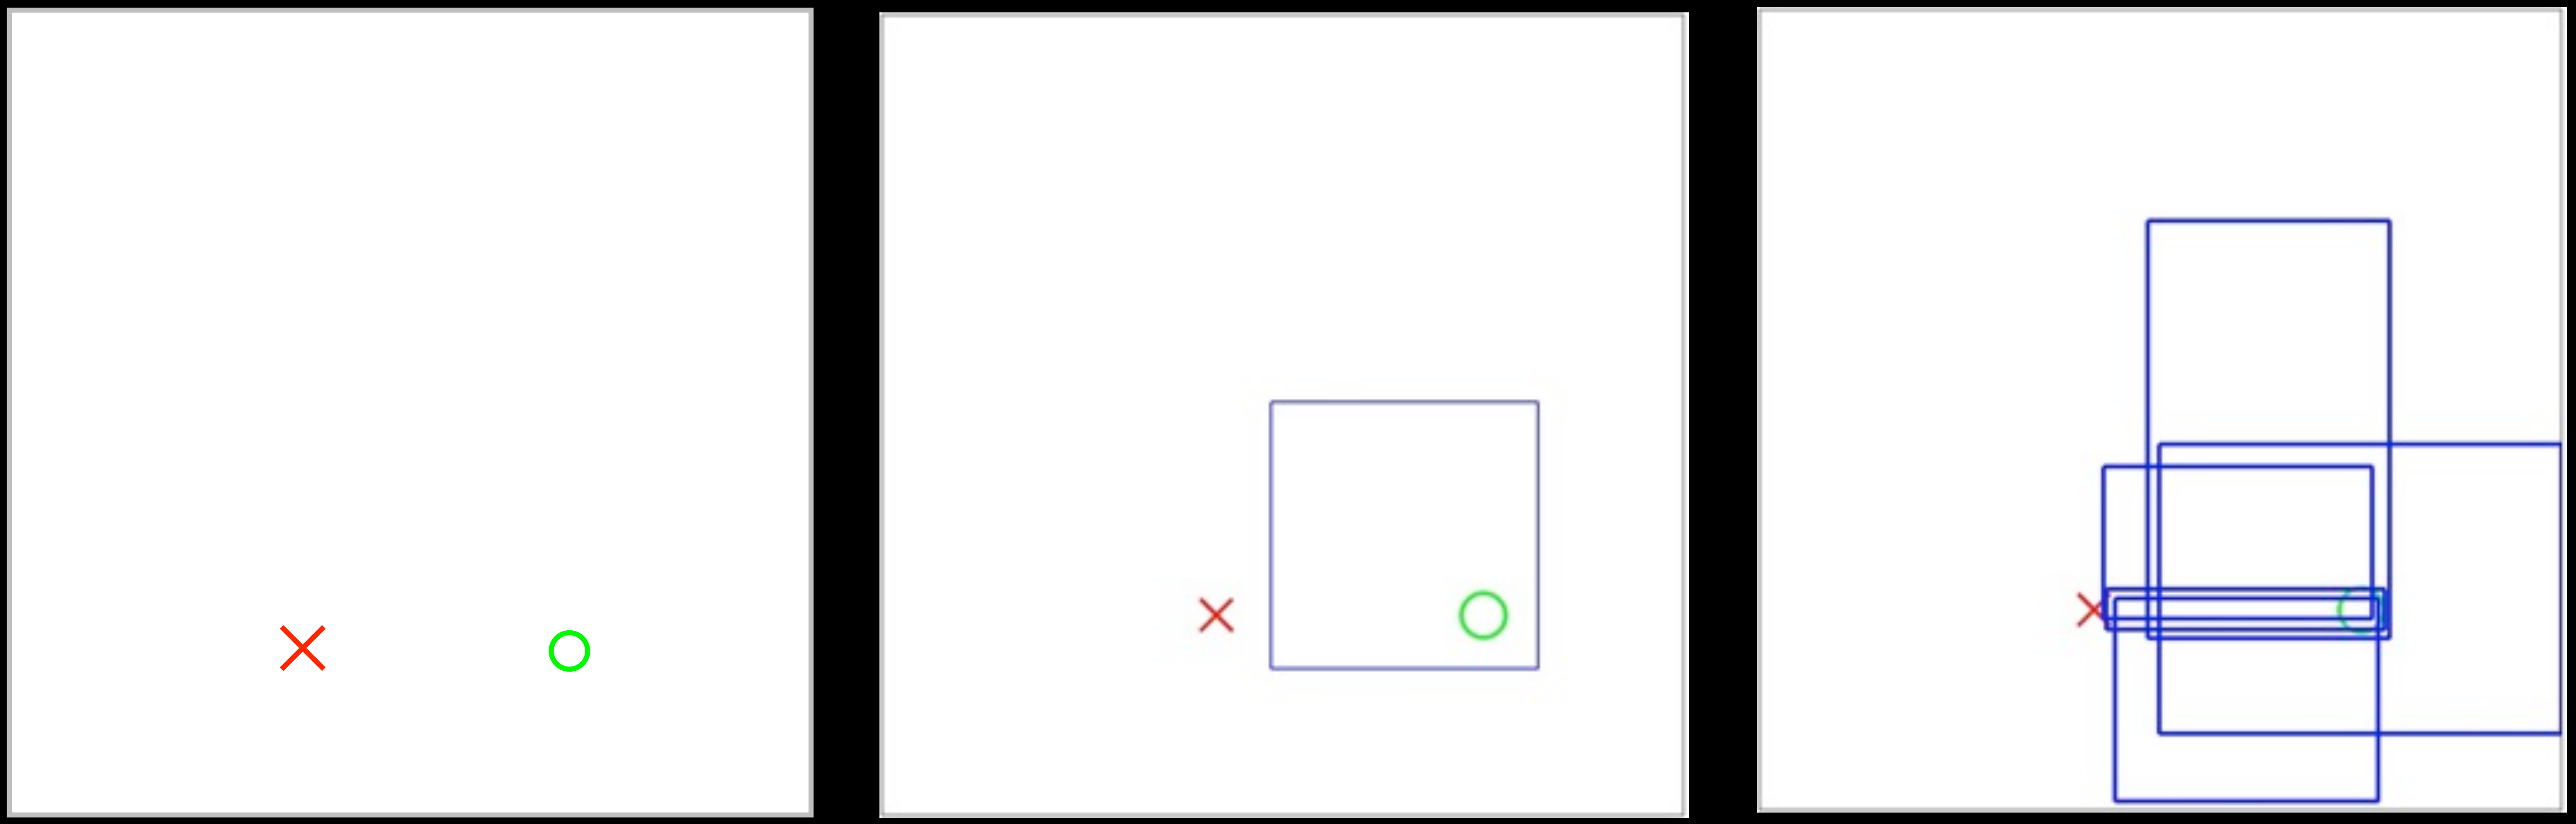

Shafto & Goodman (2008)

Learning Results

- Where did learners tend to draw rectangles, relative to examples?
- Bin examples in grid wrt rectangle.

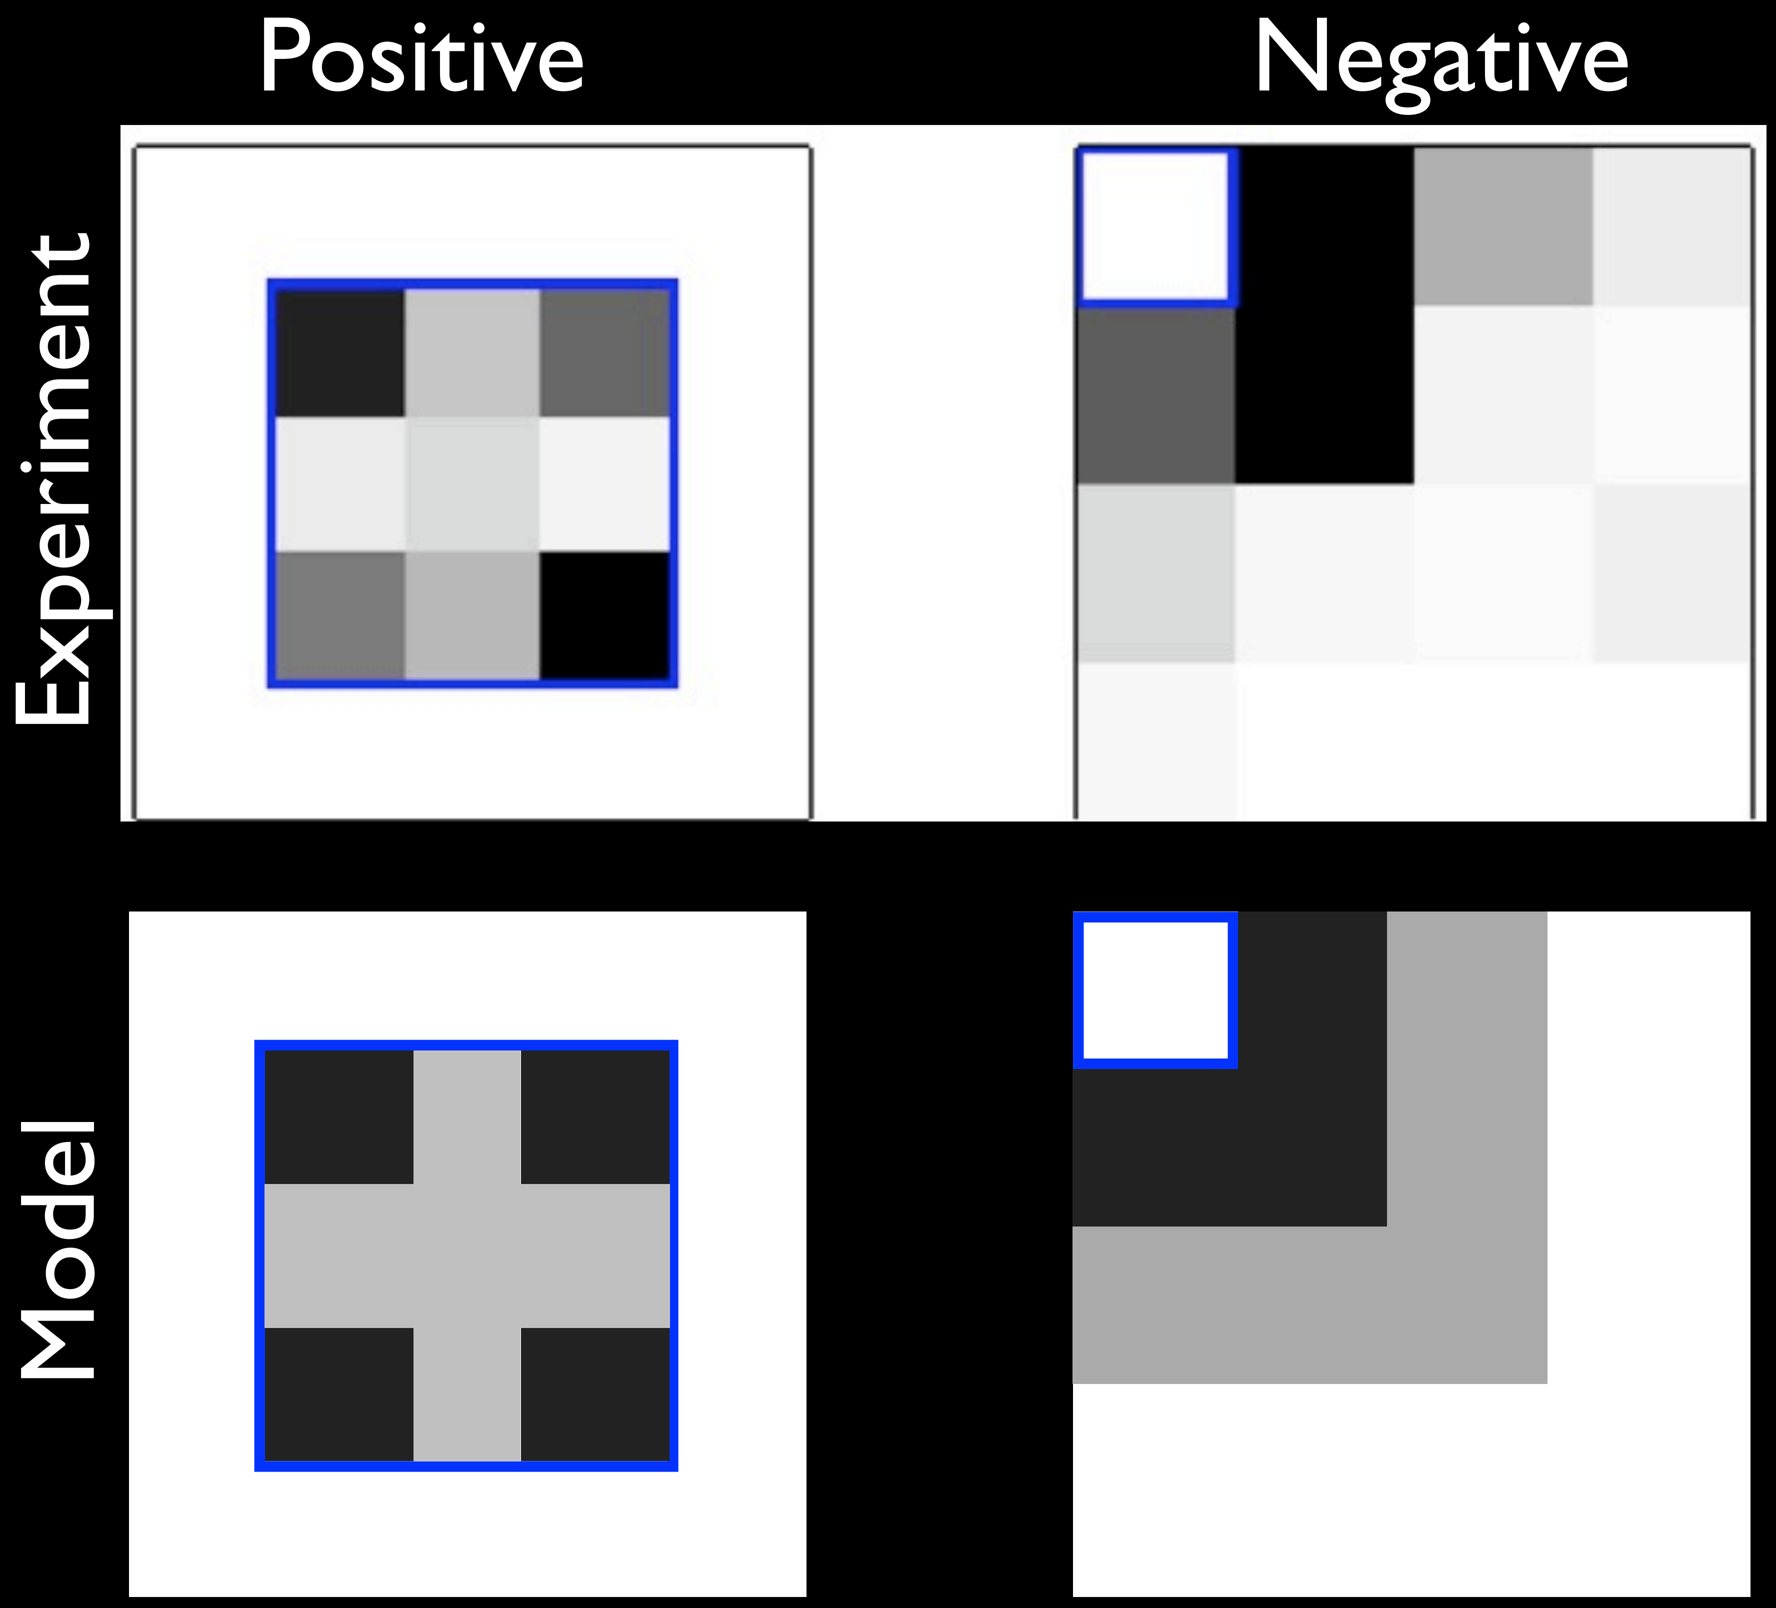

Control: random game

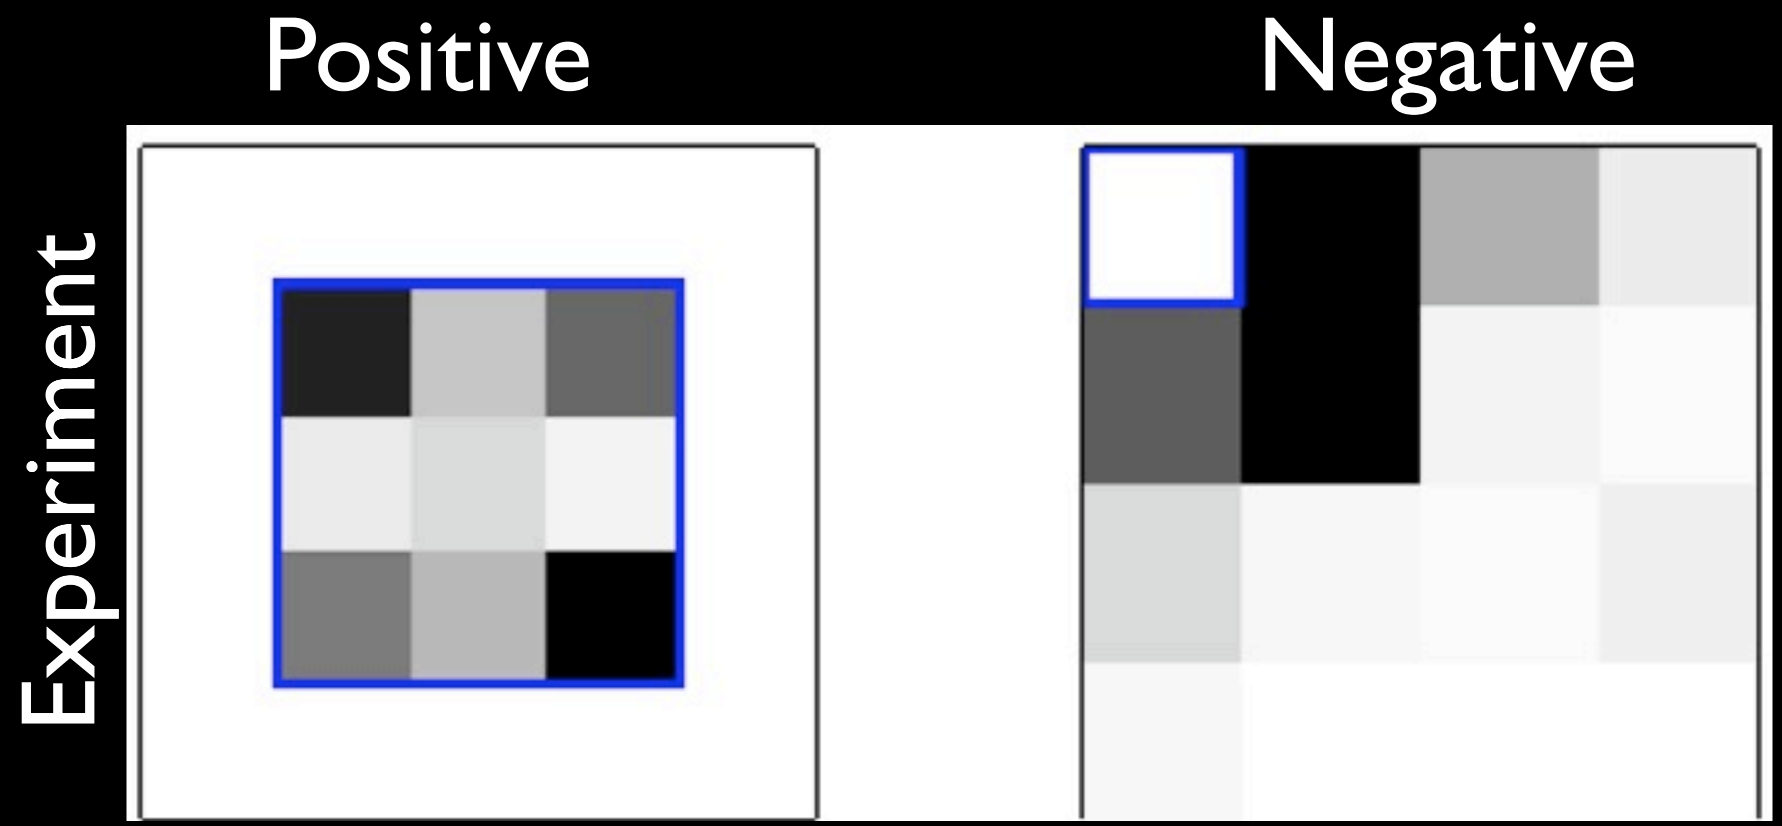

- Learner chooses the points to be labeled.
- Locations are now un-informative.

Control: random game

- Learner chooses the points to be labeled.
- Locations are now un-informative.

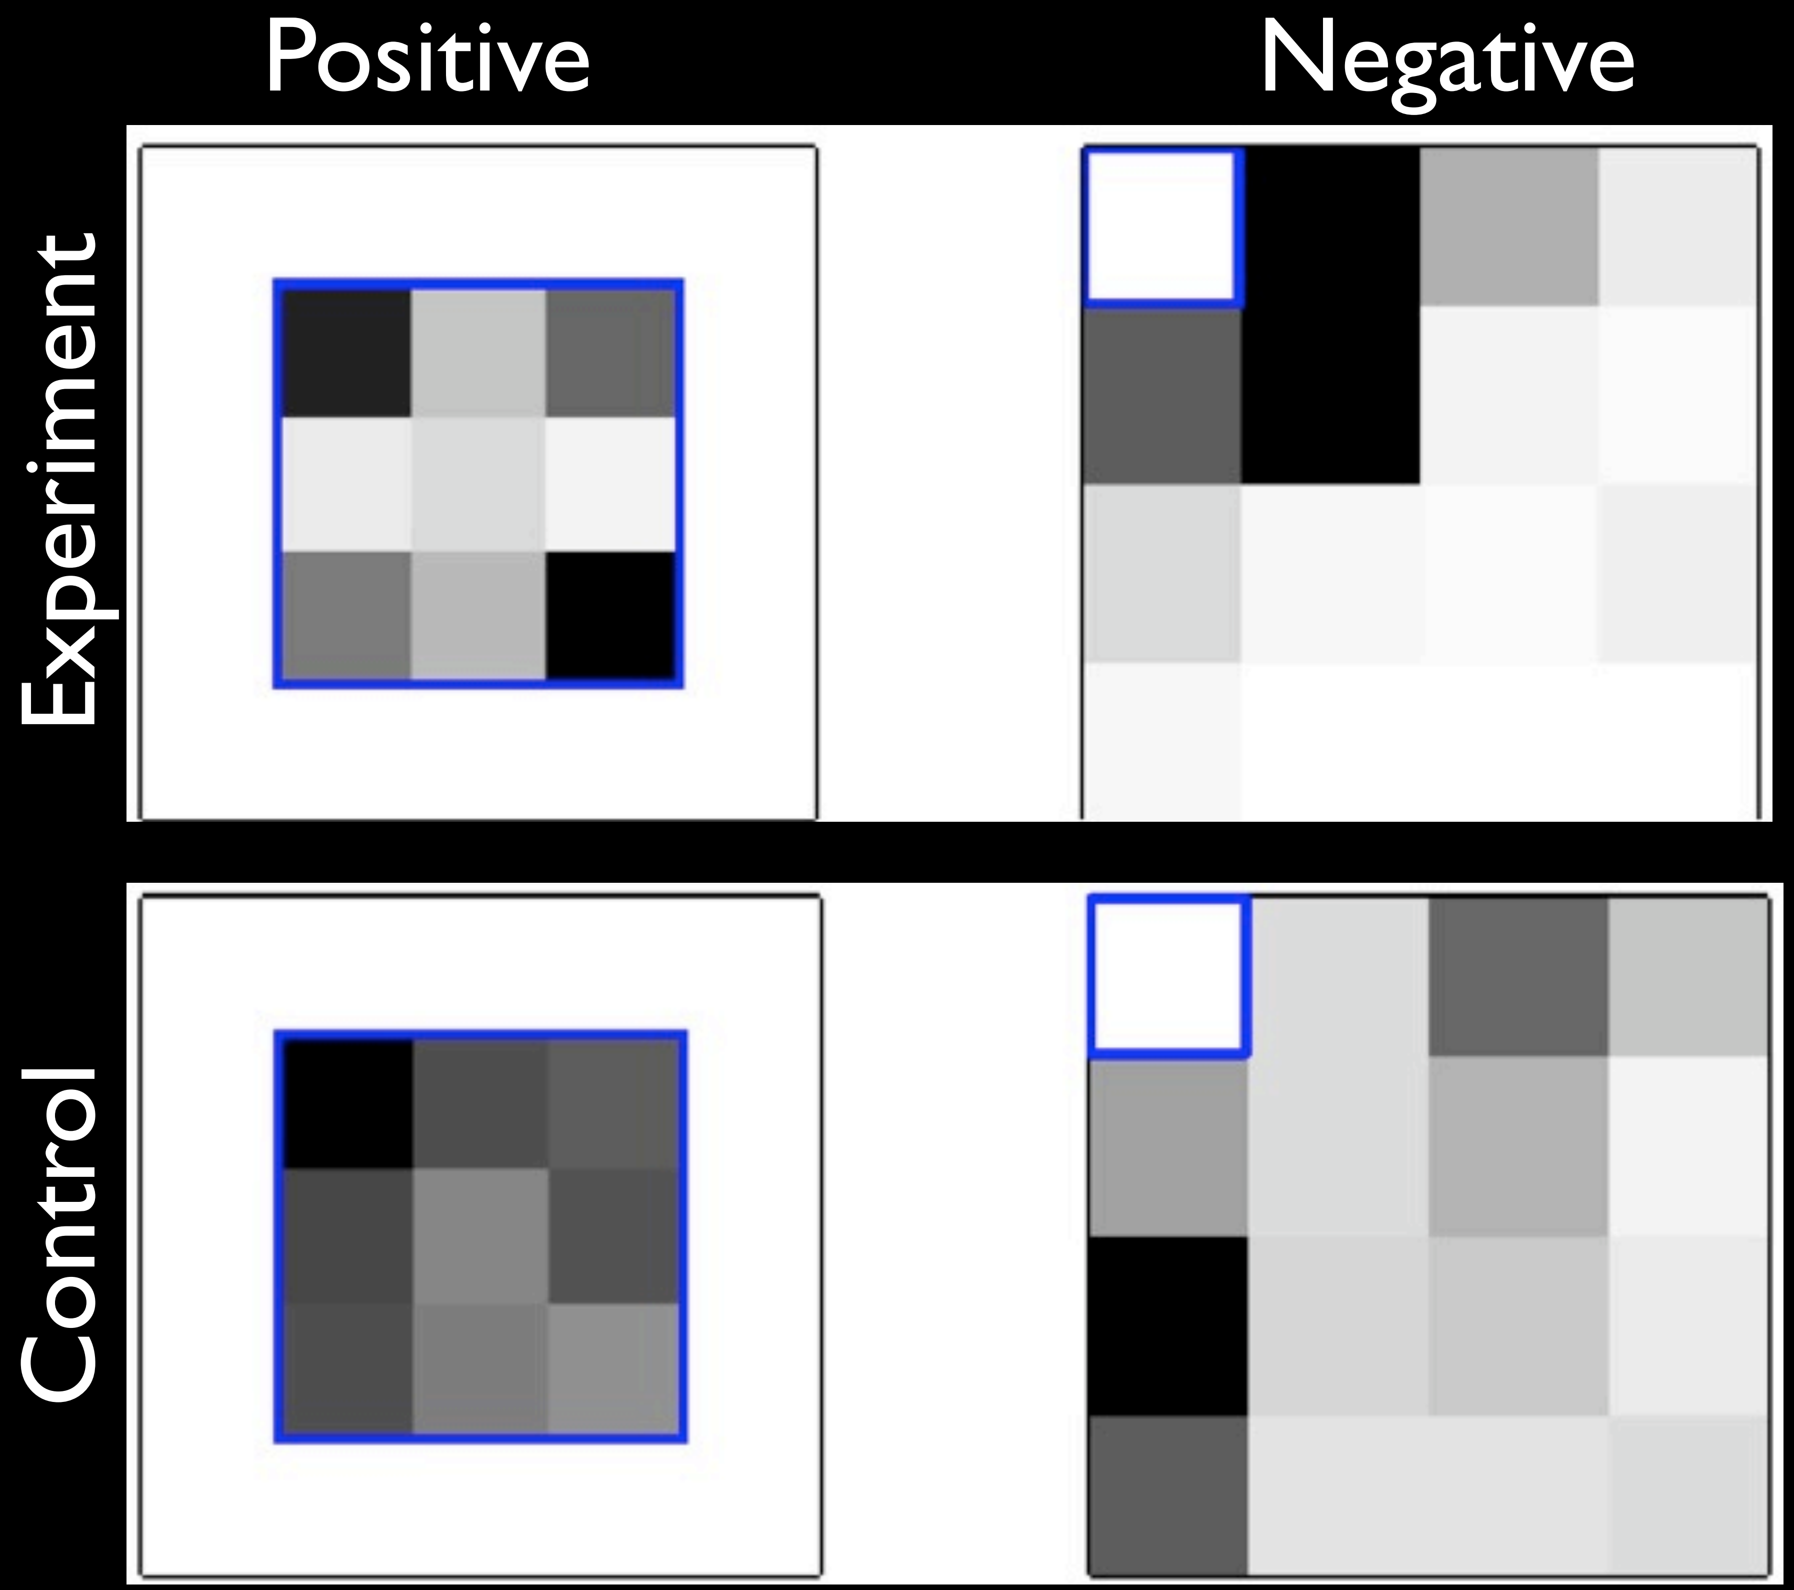

The double-edged sword

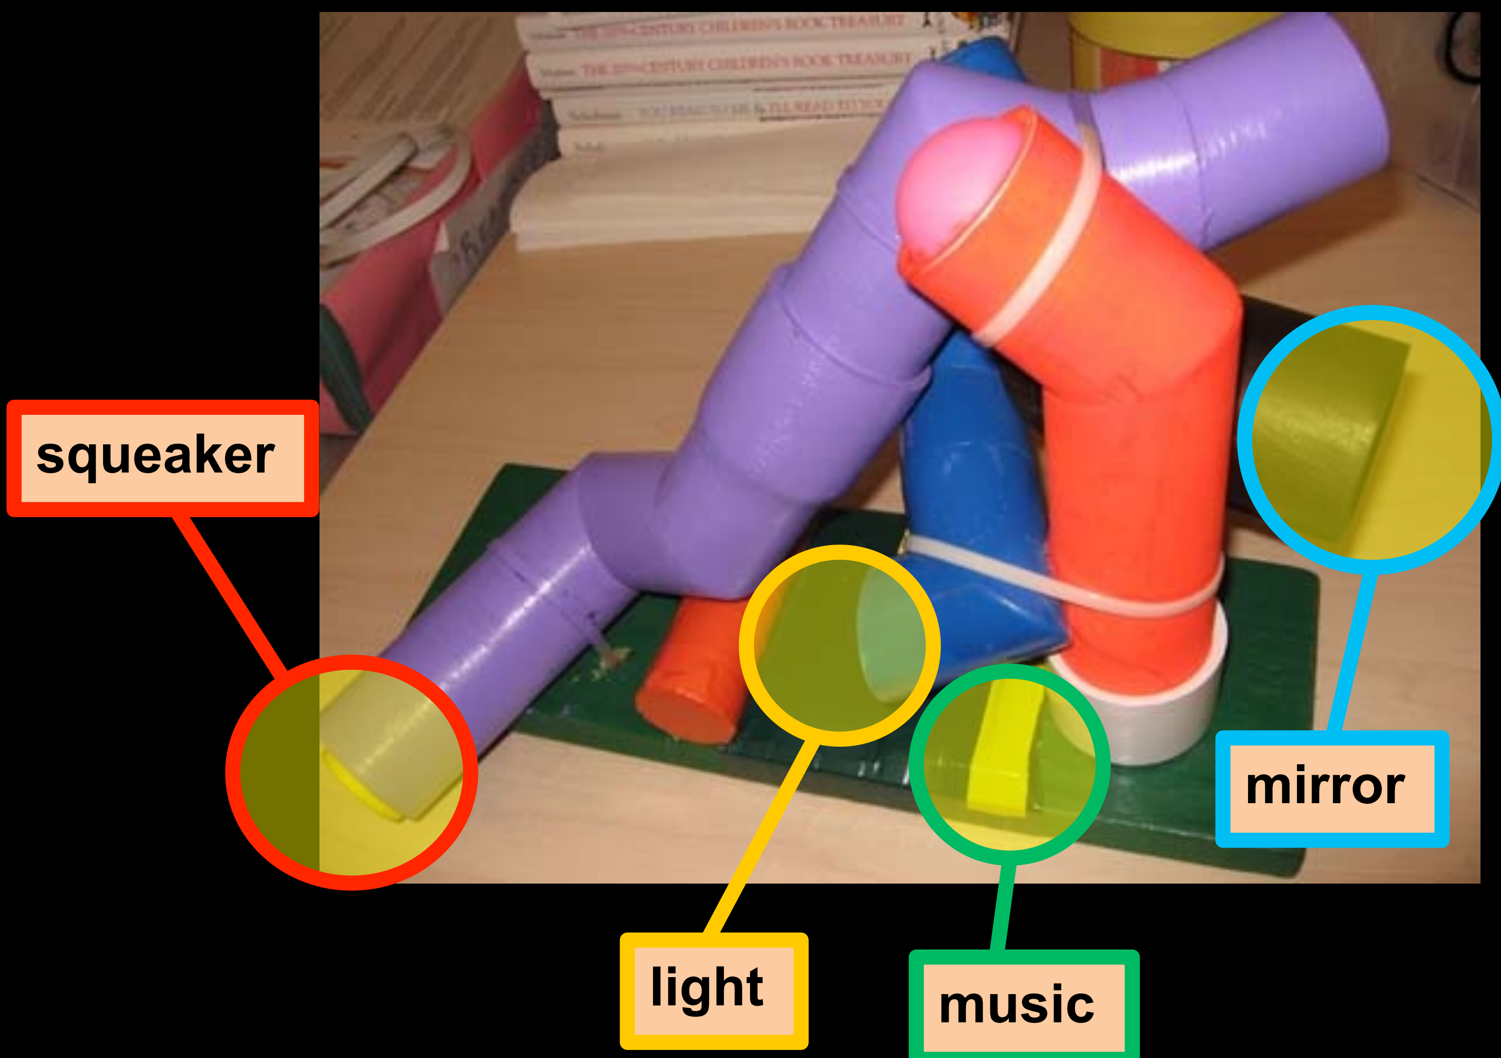

DIRECT INSTRUCTION	ACCIDENTAL	BASELINE
"Watch this, I'm going to show you my toy."	"Look at this neat toy I found here."	"Look at this neat toy that I have."
[intentionally pull tube]	[accidentally pull tube]	[rotate toy for child]
"Wow, see that?"	"Wow, see that?"	"Wow, see that?"

"Now it's your turn, why don't you go ahead and play for a little bit and see if you can figure out how it works."

Bonawitz*, Shafto*, Gweon, Goodman, Spelke, & Schulz (2011).

The double-edged sword

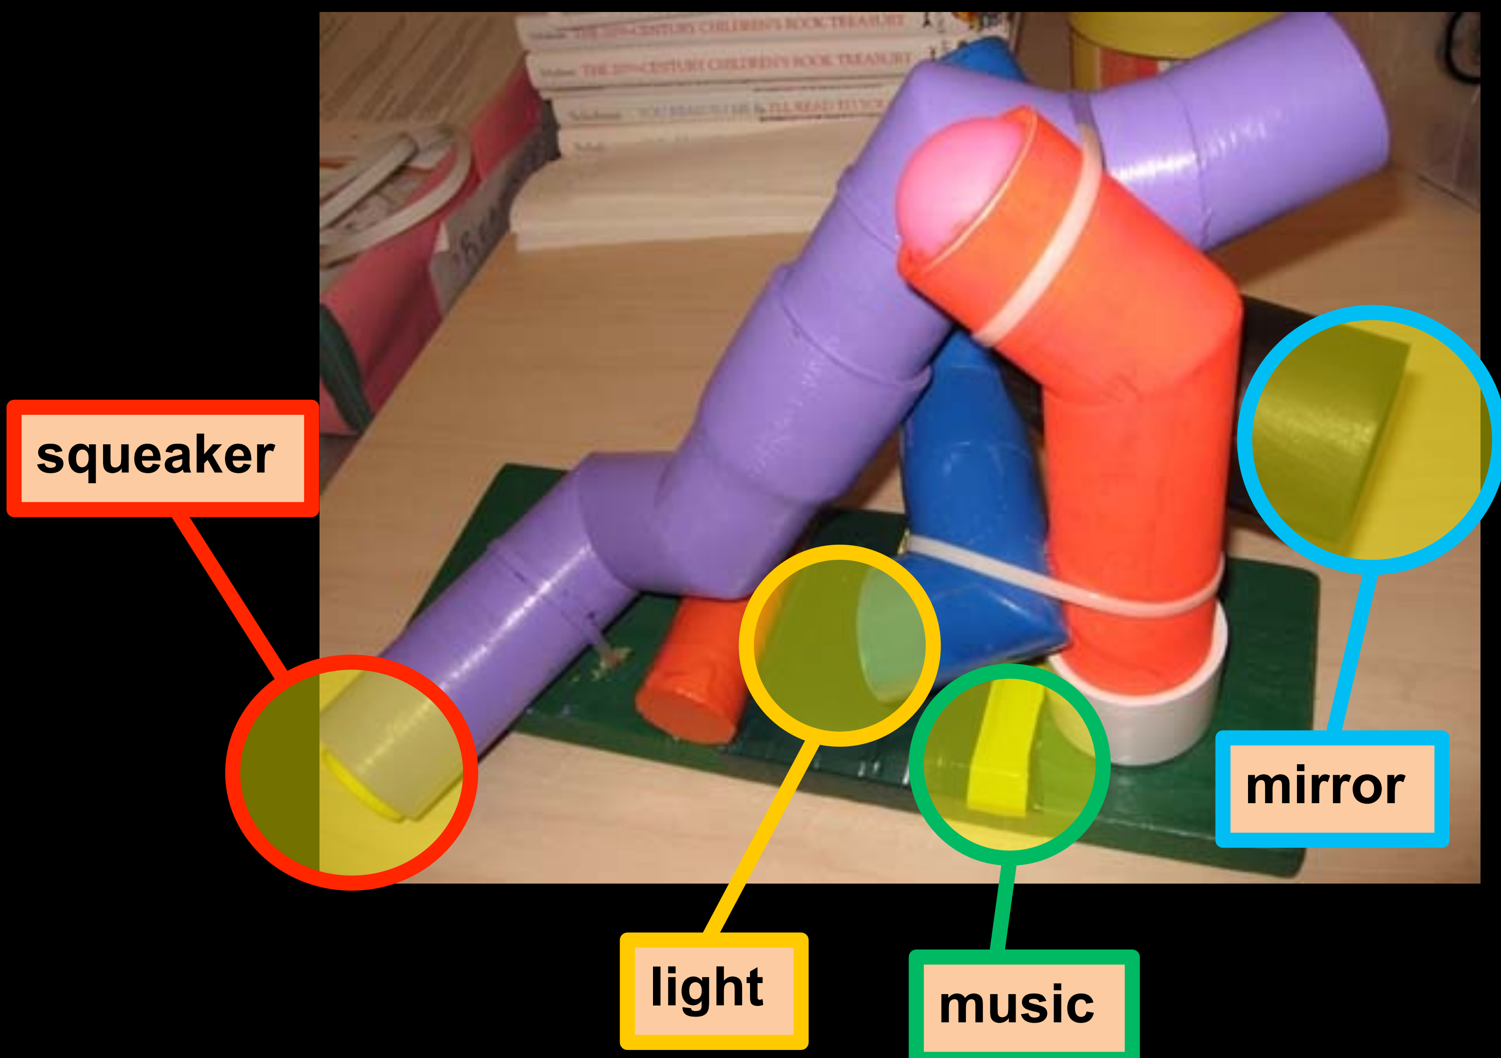

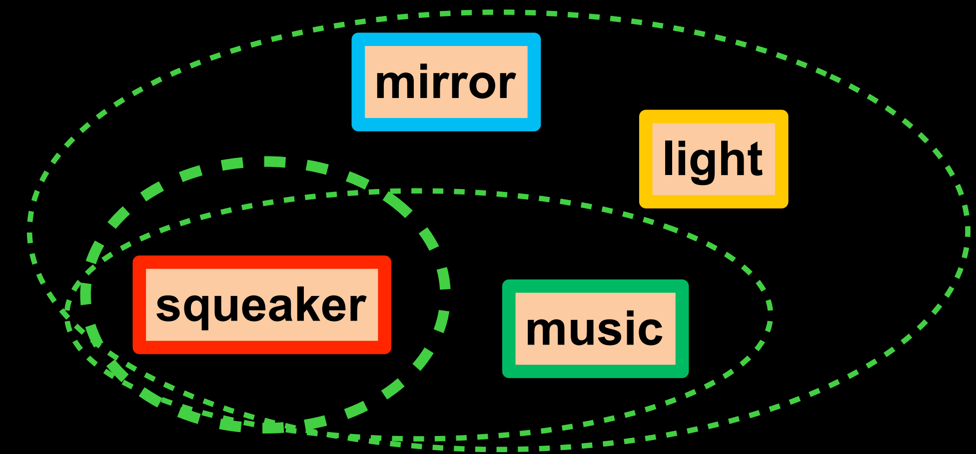

Prediction: generalize narrowly after direct instruction, because teacher *could have* shown more functions.

DIRECT INSTRUCTION	ACCIDENTAL	BASELINE
"**Watch this,** I'm going to show you **my** toy." [**intentionally pull tube**] "Wow, see that?"	"Look at this neat toy I found here." [**accidentally pull tube**] "Wow, see that?"	"Look at this neat toy that I have." [**rotate toy for child**] "Wow, see that?"

"Now it's your turn, why don't you go ahead and play for a little bit and see if you can figure out how it works."

Bonawitz*, Shafto*, Gweon, Goodman, Spelke, & Schulz (2011).

The double-edged sword

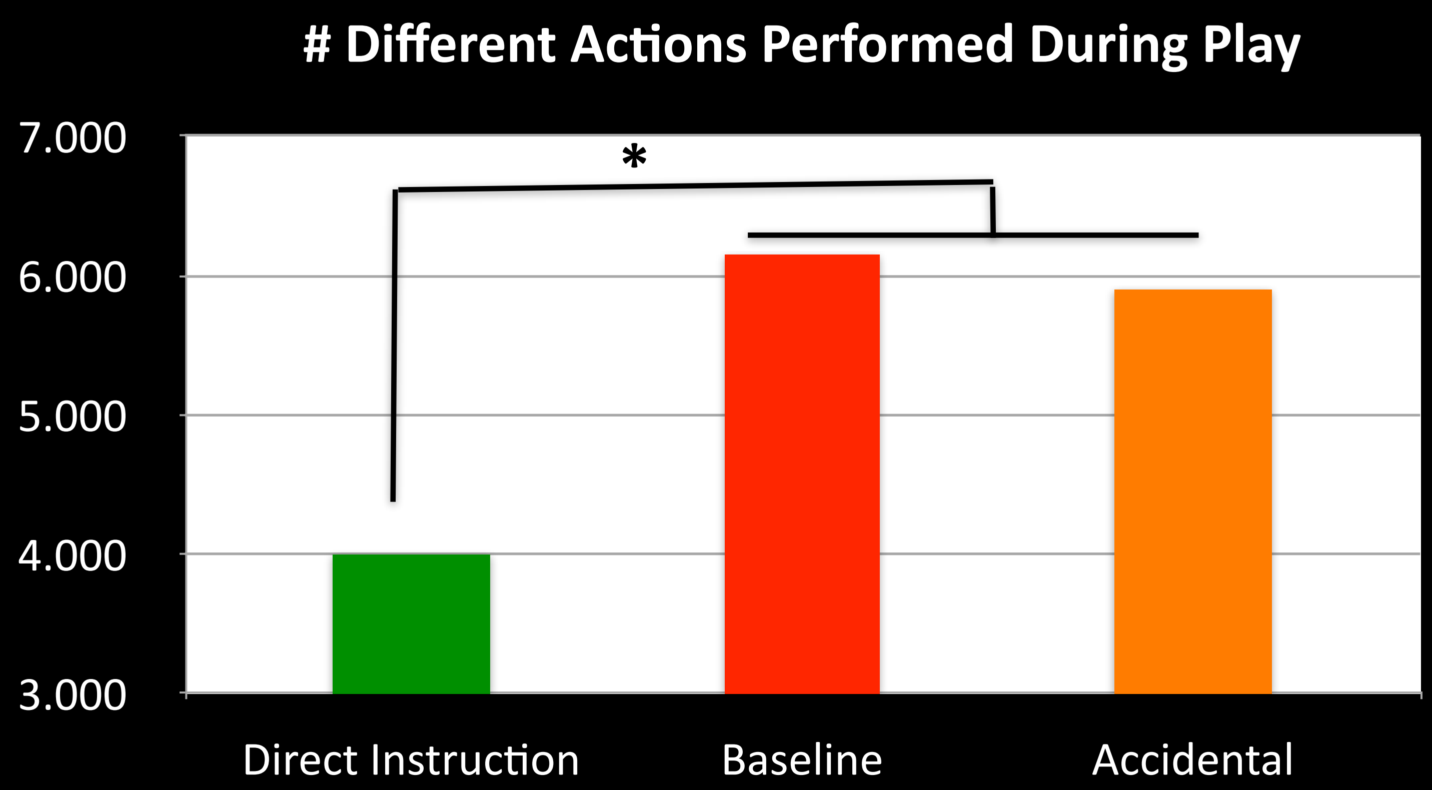

DIRECT INSTRUCTION	ACCIDENTAL	BASELINE
"**Watch this,** I'm going to show you **my** toy."	"Look at this neat toy I found here."	"Look at this neat toy that I have."
[**intentionally pull tube**]	[**accidentally pull tube**]	[**rotate toy for child**]
"Wow, see that?"	"Wow, see that?"	"Wow, see that?"

"Now it's your turn, why don't you go ahead and play for a little bit and see if you can figure out how it works."

Bonawitz*, Shafto*, Gweon, Goodman, Spelke, & Schulz (2011).

Outline

- Theory of mind and learning from others' actions.
- Multi-agent reasoning: coordination games.
- Communicating with natural signs: intuitive pedagogy.
- Communicating with arbitrary signs: natural language.

Communication

- We can’t rely on simple coordination to arrange things.
- Instead we pass *signs* that help us to coordinate.
- *Natural signs* have meaning in the world.
- *Arbitrary signs* only have conventional (literal) meaning.

Language

- Speaker chooses an *utterance*.
- Each utterance has a *literal meaning*:
  - For now, a truth-function: a predicate on states of the world.

Goodman & Stuhlmueller (in prep)

Language

```
(define (speaker state)
  (query
   (define words (sentence-prior))
   words
   (equal? state (listener words))))
```

Language

```
(define (speaker state)
  (query
   (define words (sentence-prior))
   words
   (equal? state (listener words))))
```

```
(define (listener words)
  (query
   (define state (state-prior))
   state
   (if (flip laziness)
       (words state)
       (equal? words (speaker state)))))
```

Language

```
(define (speaker state)
  (query
   (define words (sentence-prior))
   words
   (equal? state (listener words))))
```

```
(define (listener words)
  (query
   (define state (state-prior))
   state
   (if (flip laziness)
       (words state)
       (equal? words (speaker state)))))
```

If the words are true?

(or)

If the speaker chose these words?

Scalar implicature

Literal meanings:

"**none** of the circles are red"

"**some** of the circles are red"

"**all** of the circles are red"

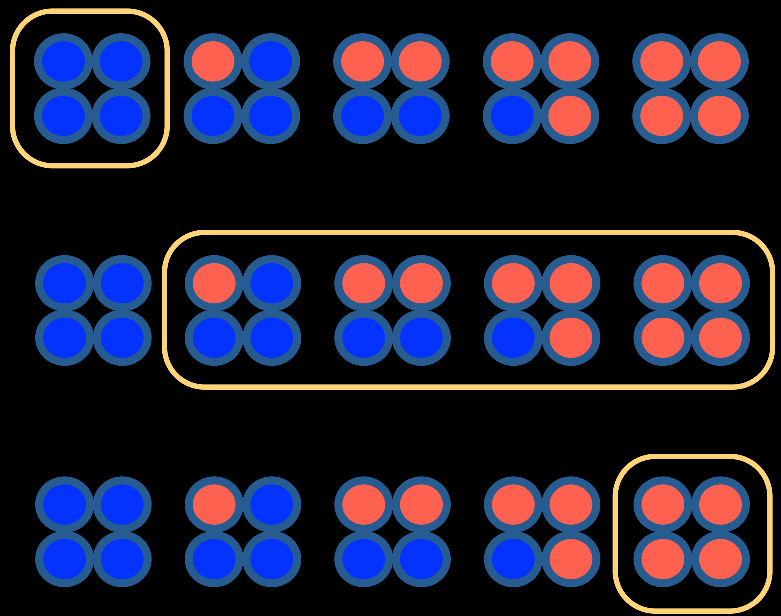

Scalar implicature

Literal meanings:

"**none** of the circles are red"

"**some** of the circles are red"

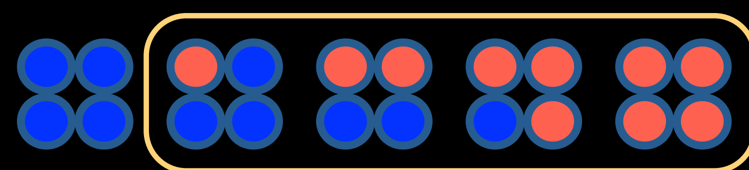

"**all** of the circles are red"

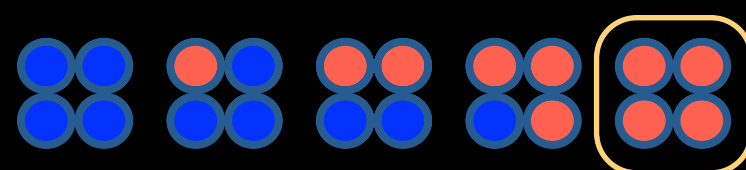

Scalar implicature

Literal meanings:

“**none** of the circles are red”

“**some** of the circles are red”

“**all** of the circles are red”

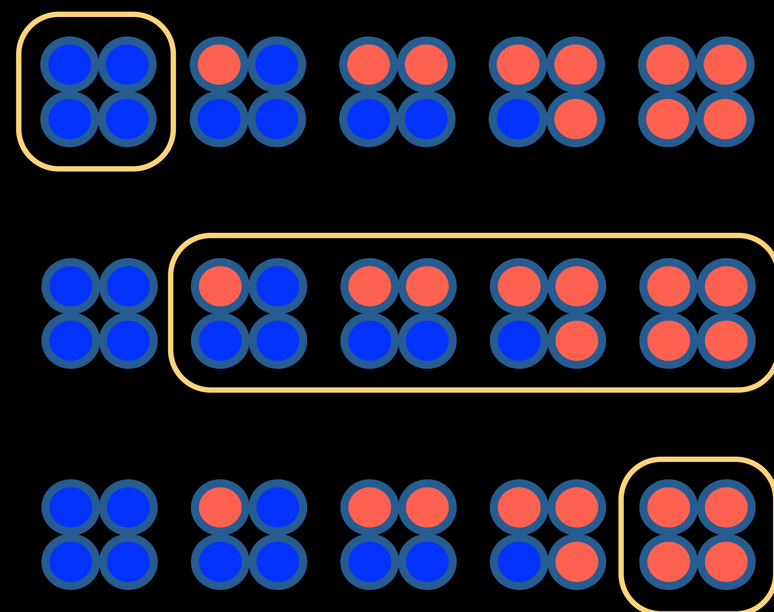

`(listener “some..”):`

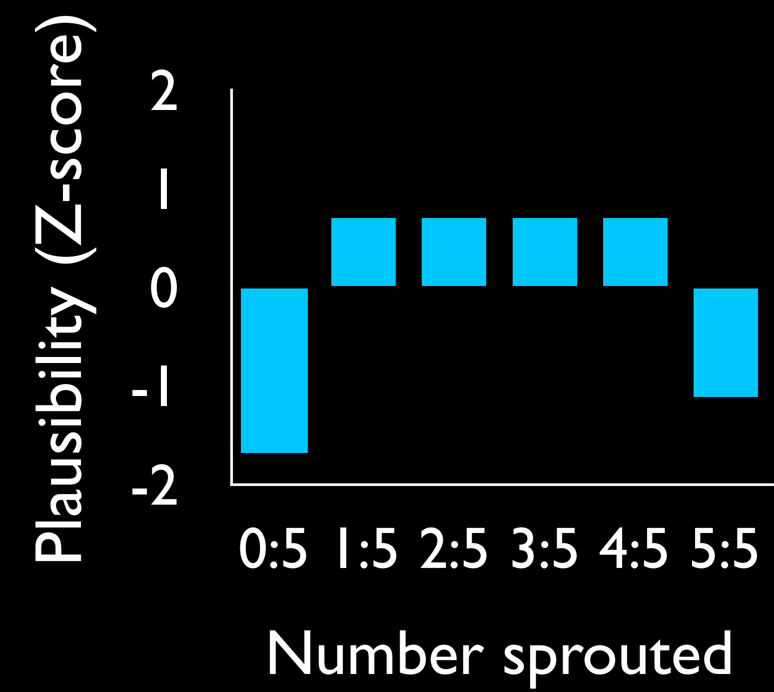

Scalar implicature

Literal meanings:

"**none** of the circles are red"

"**some** of the circles are red"

"**all** of the circles are red"

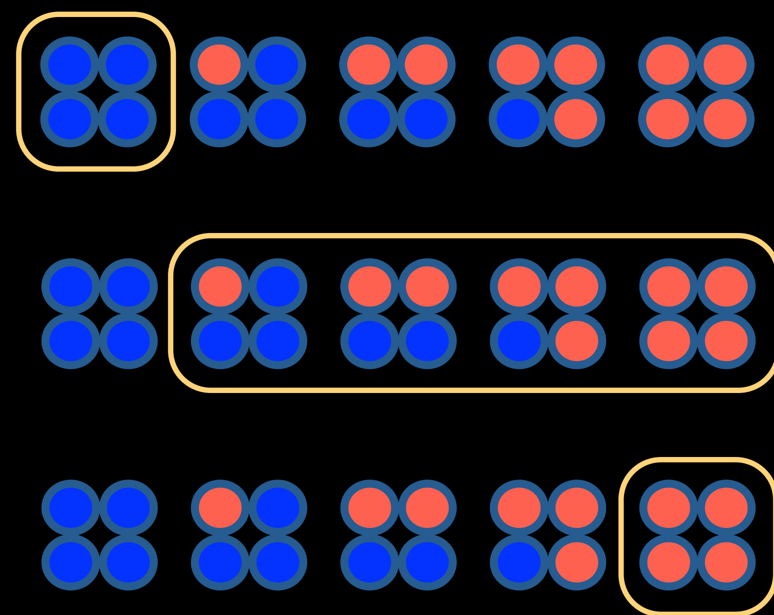

`(listener "some..")`:

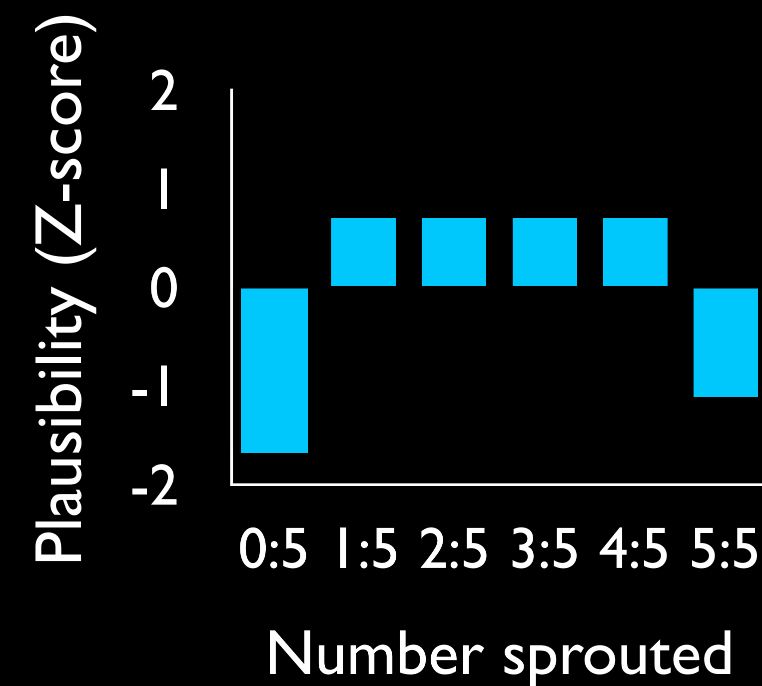

(Plants usually sprout.)

Plants usually sprout:

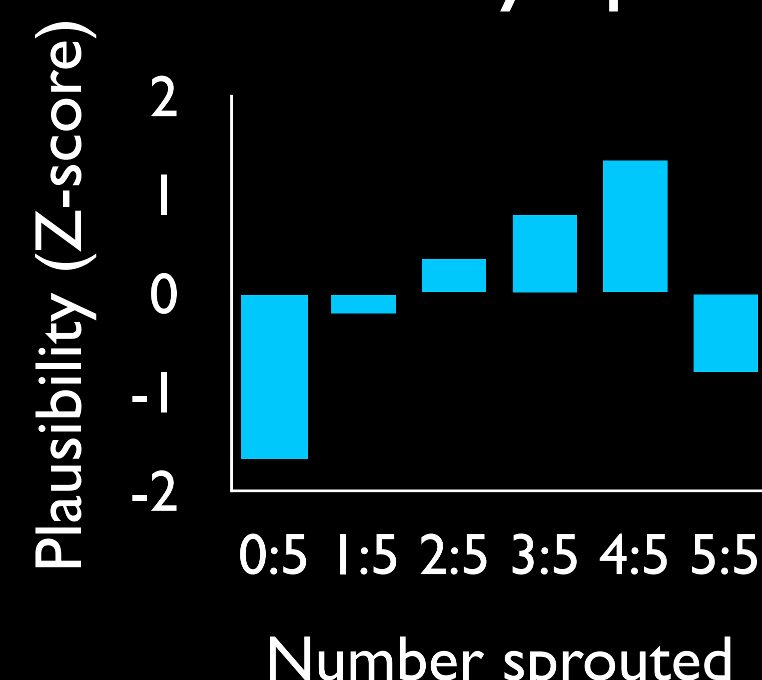

Scalar implicature

(Plants usually sprout.)

Scalar implicature

- Speaker has only partial knowledge of world state.
- Listeners knows that.

(Plants usually sprout.)

Scalar implicature

- Speaker has only partial knowledge of world state.
- Listeners knows that.

```
(define (speaker state access)
  (query
   (define words (sentence-prior))
   words
   (equal? (belief state access)
           (listener words access))))
```

(Plants usually sprout.)

Scalar implicature

- Speaker has only partial knowledge of world state.
- Listeners knows that.

```
(define (speaker state access)
  (query
   (define words (sentence-prior))
   words
   (equal? (belief state access)
           (listener words access))))
```

```
(define (belief state access)
 ...for each object,
      if access, then true state,
        else draw from prior...)
```

(Plants usually sprout.)

Scalar implicature

(Plants usually sprout.)

Scalar implicature

split in two?

(Plants usually sprout.)

Scalar implicature

Model:

Full knowledge

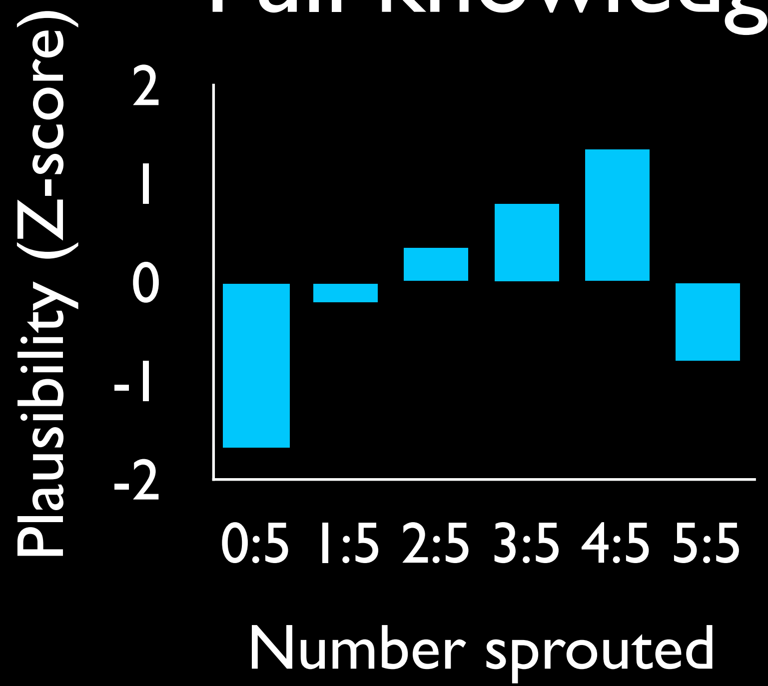

split in two?

(Plants usually sprout.)

Scalar implicature

Model:

Full knowledge — Partial knowledge

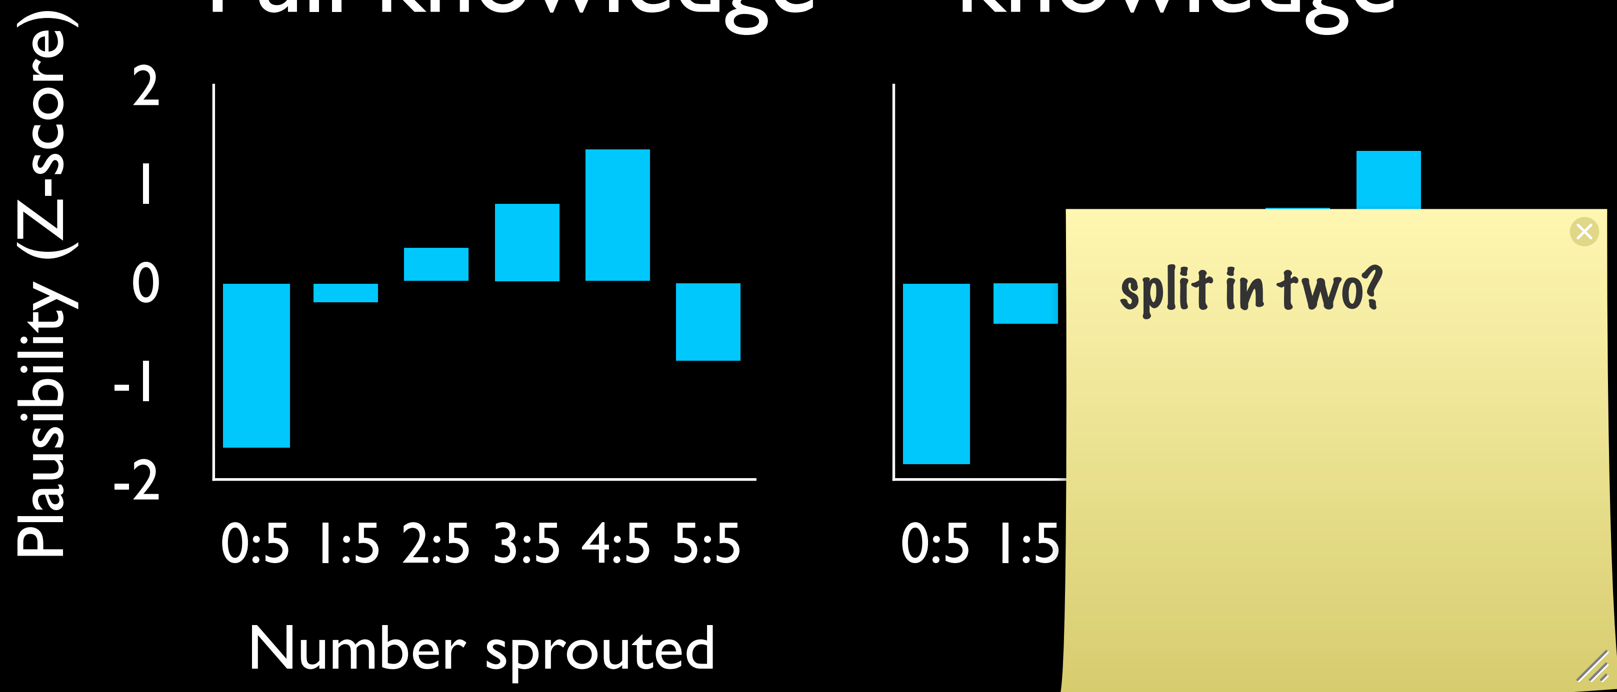

split in two?

(Plants usually sprout.)

Scalar implicature

Model:

Full knowledge | Partial knowledge

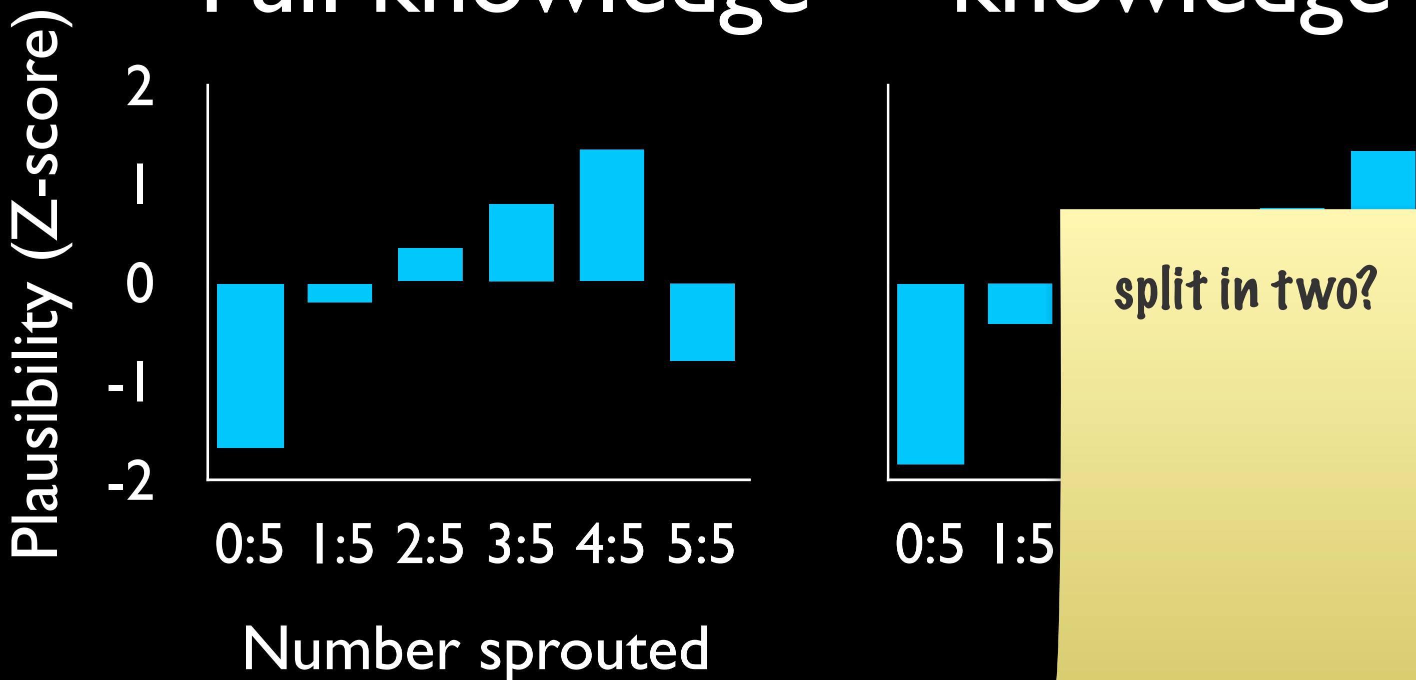

split in two?

- Three scenarios, varying information access.

Corendula is a type of plant that, when put into water, almost always sprouts within a day. Two ...ys ago, botanist Jim put ***five*** Corendula ...eds into water. He just got back to his plants ...lay, has looked at ***two of the*** seeds and ...ys: "Some of my five seeds have sprouted."

- Elicit likelihood ratings.

(Plants usually sprout.)

Scalar implicature

Model:

Full knowledge

Partial knowledge

split in two?

- Three scenarios, varying information access.

Corendula is a type of plant that, when put into water, almost always sprouts within a day. Two ys ago, botanist Jim put ***five*** Corendula eds into water. He just got back to his plants ay, has looked at ***two of the*** seeds and s: "Some of my five seeds have sprouted."

- Elicit likelihood ratings.

Human:

(Plants usually sprout.)

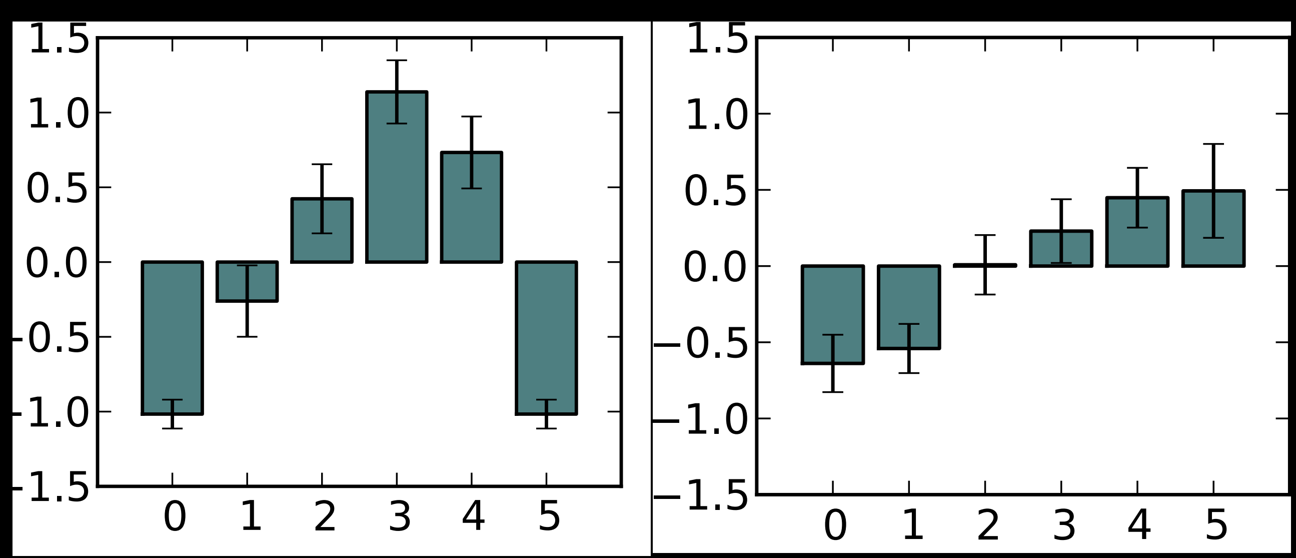

N=15

Horn’s principle

- Horn’s principle of division of pragmatic labor:
  “(un)marked expressions typically get an (un)marked interpretation” (e.g., Van Rooy, 2004)
- What does this mean? Does it follow from social reasoning models?

With Leon Bergen, Roger Levy, Andreas Stuhlmueller.

Horn’s principle example

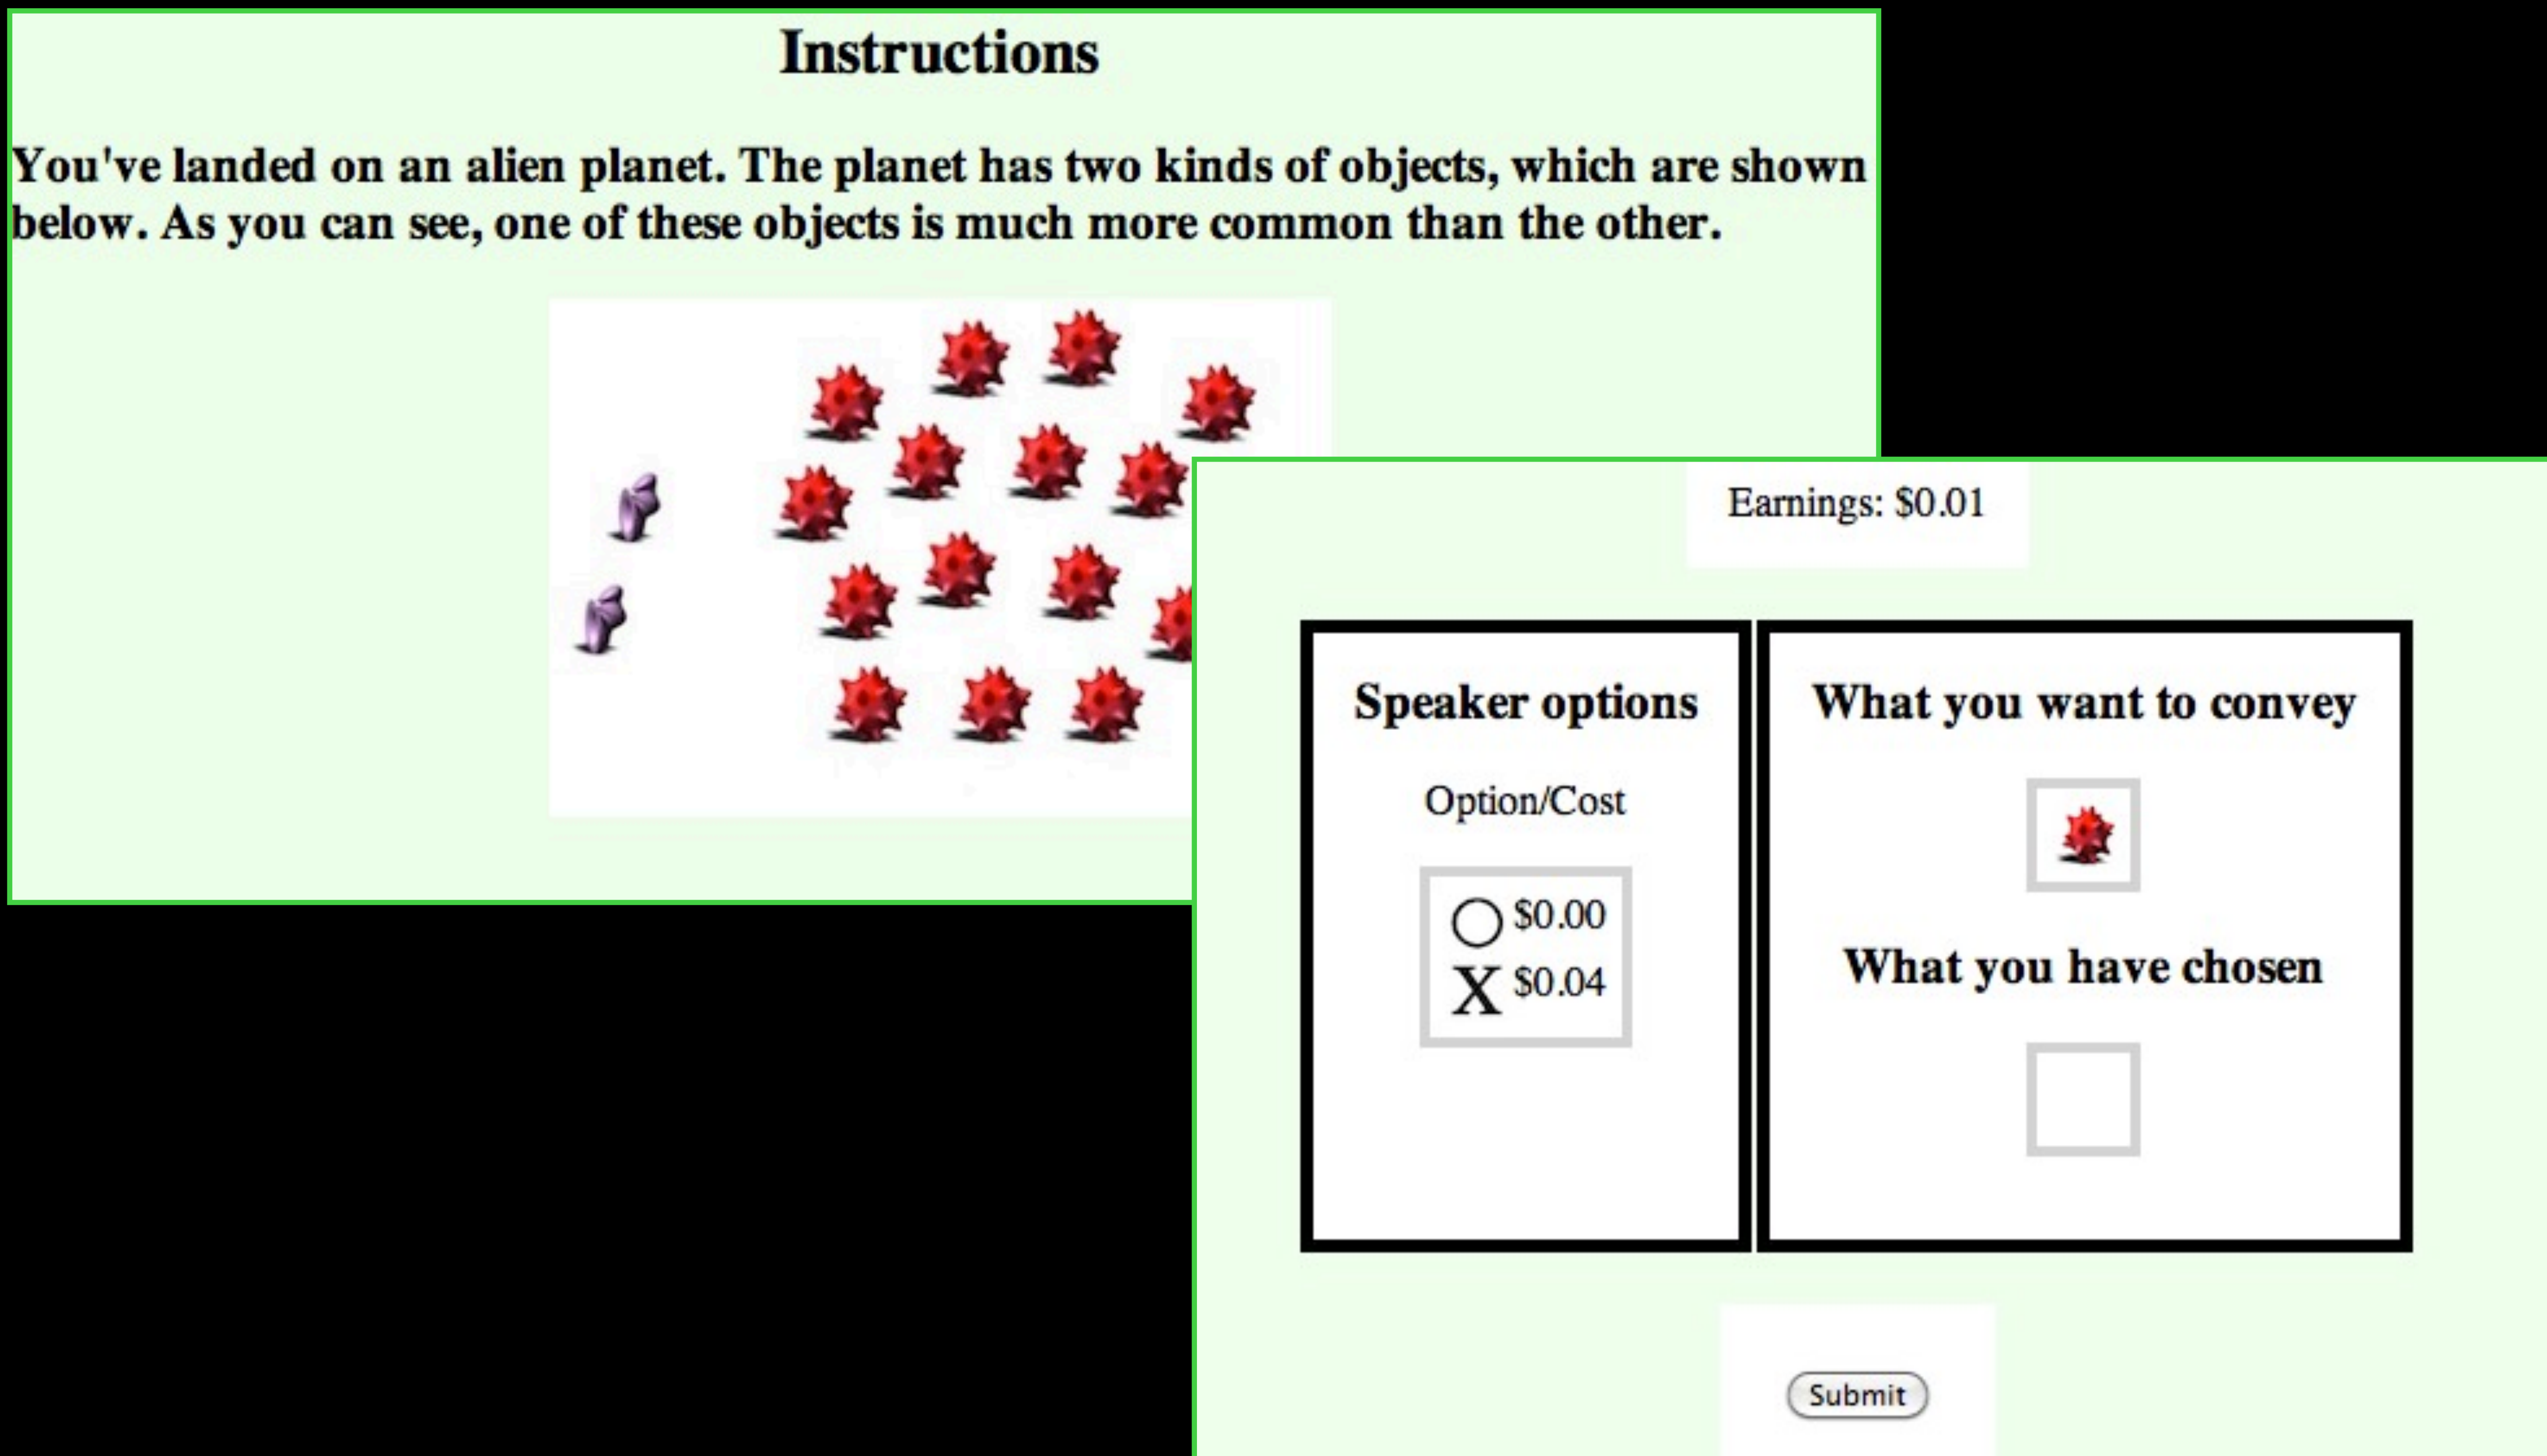

Horn's principle example

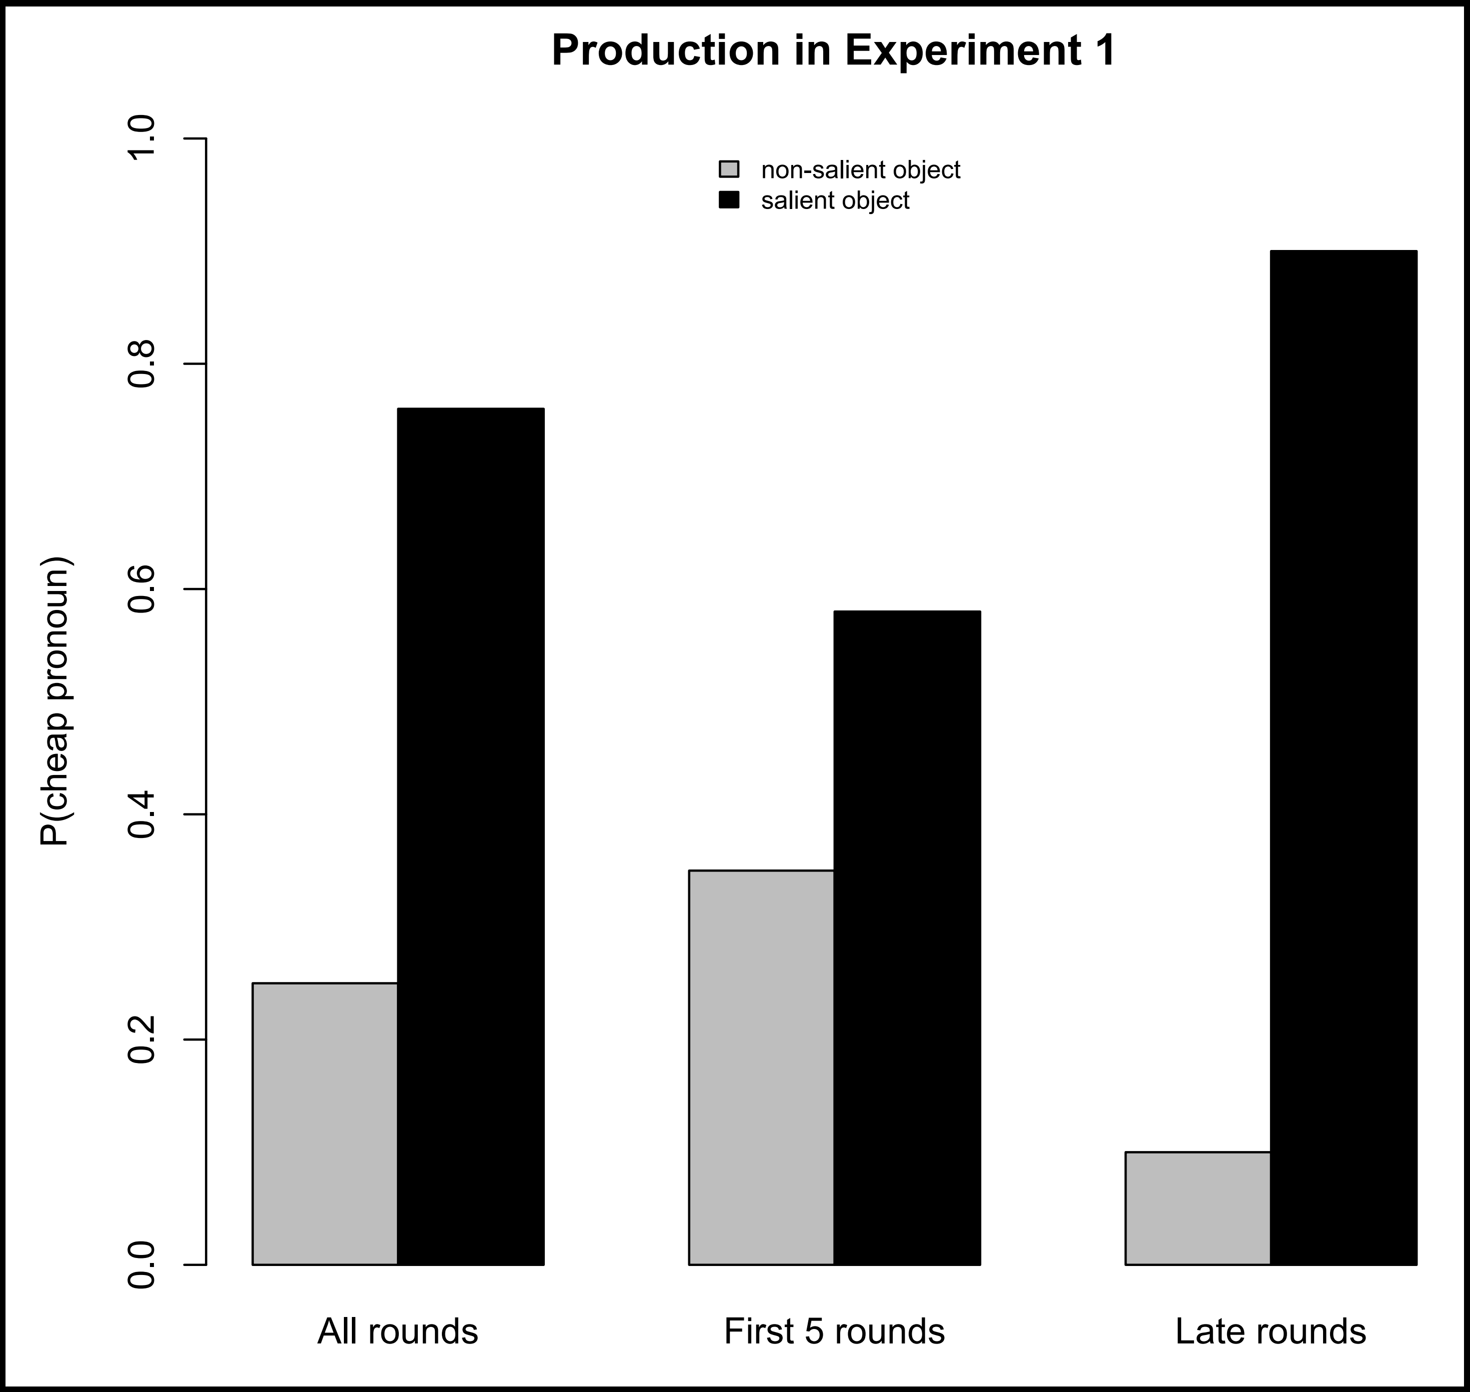

The cheap pronoun is used for the common object.

Modeling Horn's principle

- Basic model doesn't work.
  - Cf. non-informative equilibria in signaling games.
- Be more optimal? Nope.
- Select whole strategies.
- Other options....

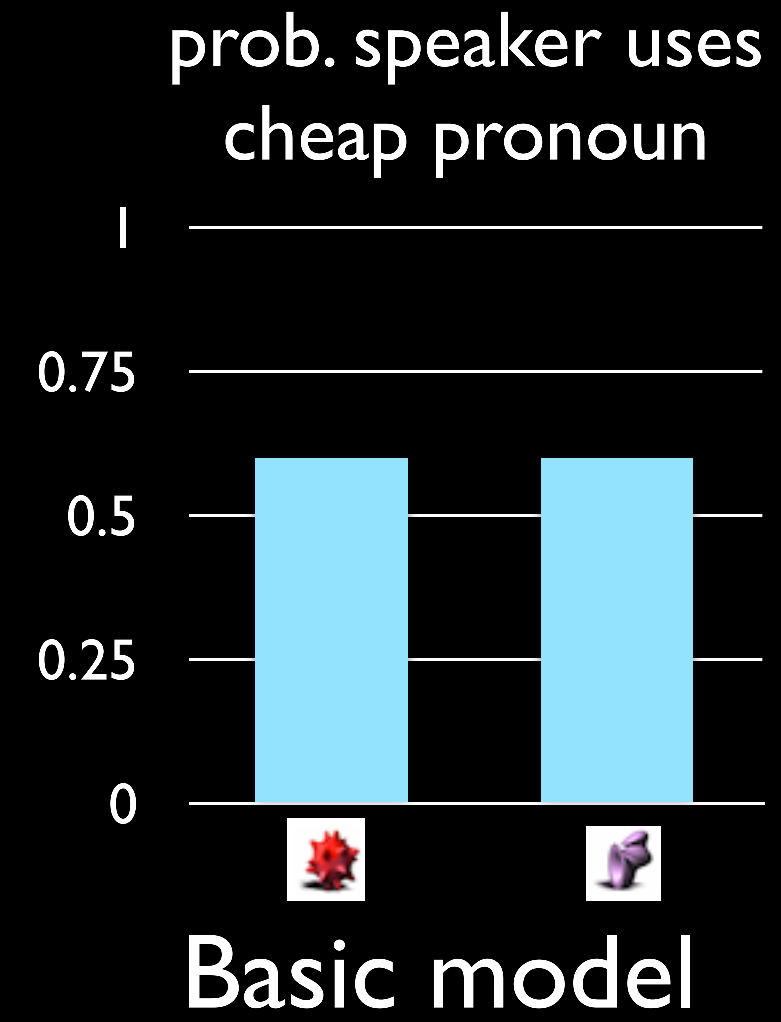

Modeling Horn’s principle

- Basic model doesn’t work.
  - Cf. non-informative equilibria in signaling games.
- Be more optimal? Nope.
- Select whole strategies.
- Other options….

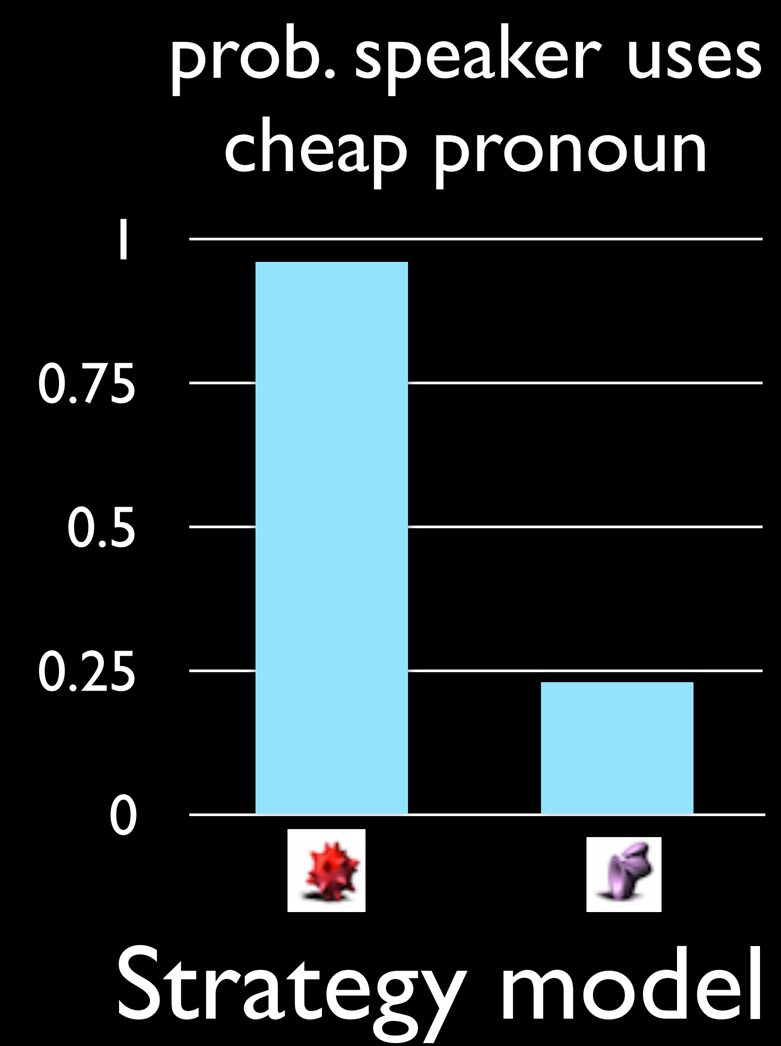

Application to implicature

- Puzzle: why does “some or all” not have a “not all” implicature?
  - If “some or all” implies ignorance,
  - the earlier result explains why the potential “not all” implicature is canceled.
  - Ignorance follows from Horn’s principle (as in the previous simple example)!
    - “some or all” more complex than “some”,
    - knowledge more common than ignorance.

Knowledge inference

Your friend Jim is slowly unwrapping ten candies in another room. After a moment he says to you "**all** of the candies are chocolate."

Do you think Jim knows exactly how many of the candies are chocolate?

- 1 (Definitely knows)
- 2
- 3
- 4
- 5 (Definitely does **not** know)

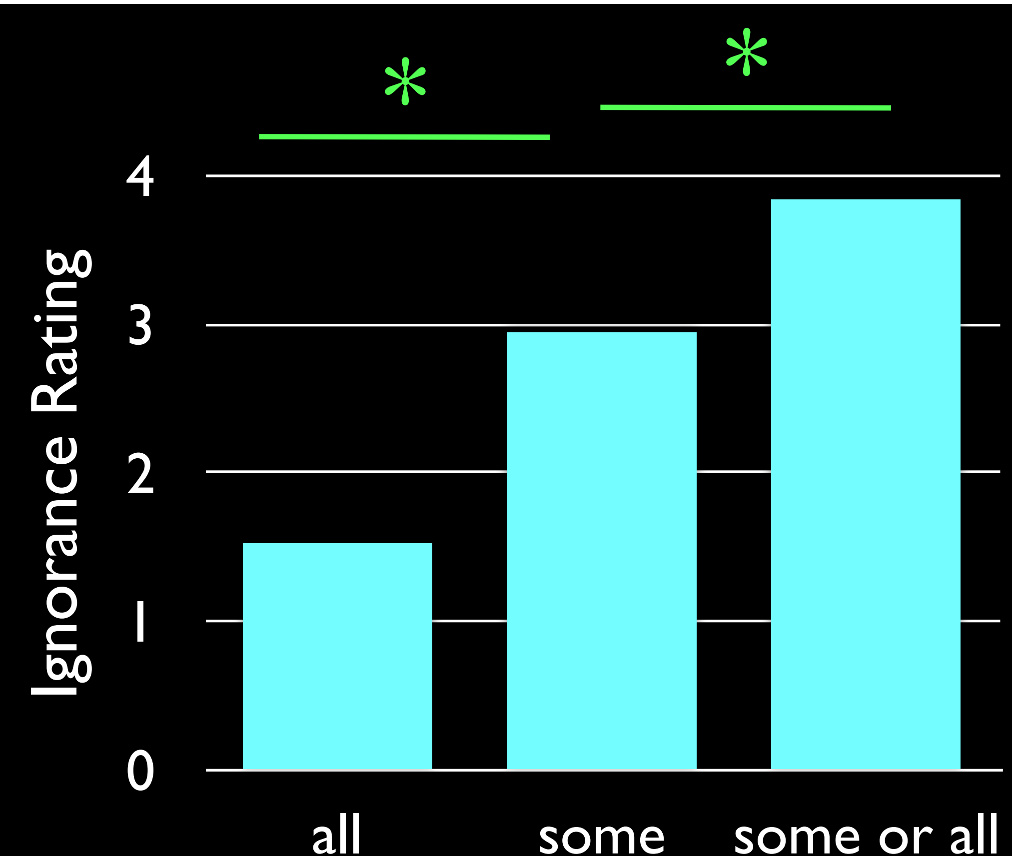

N=19

Word learning

words: "blue rings"
objects: rings, big bird

words: "and green rings"
objects: rings, big bird

words: "and yellow rings"
objects: rings, big bird

words: "Bigbird! Do you want to hold the rings?"
objects: big bird

In any one situation, children hear many words and see many objects.

Referential word learning

- Bayesian inference to learn word-object mappings?
- Words come from people...
- so model word generation via the (unknown) intention of the speaker.

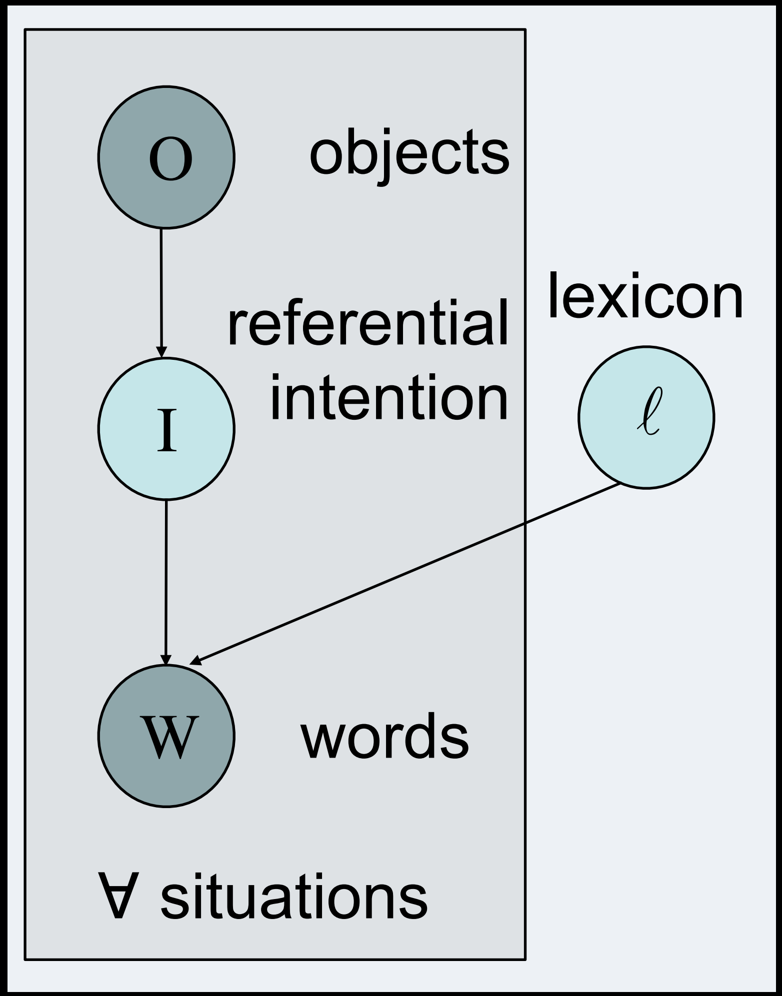

Frank, Goodman, Tenenbaum (2009)

Corpus Results

Best Lexicon:

Word	Object
bear	**bear**
bigbird	**bird**
bird	**duck**
birdie	**duck**
book	**book**
bottle	bear
bunnies	**bunny**
bunnyrabbit	**bunny**
hand	**hand**
hat	**hat**
hiphop	mirror
kittycat	**kitty**
lamb	**lamb**
laugh	cow
meow	baby
mhmm	hand
mirror	**mirror**
moocow	**cow**
oink	**pig**
on	ring
pig	**pig**
put	ring
ring	**ring**
sheep	**sheep**

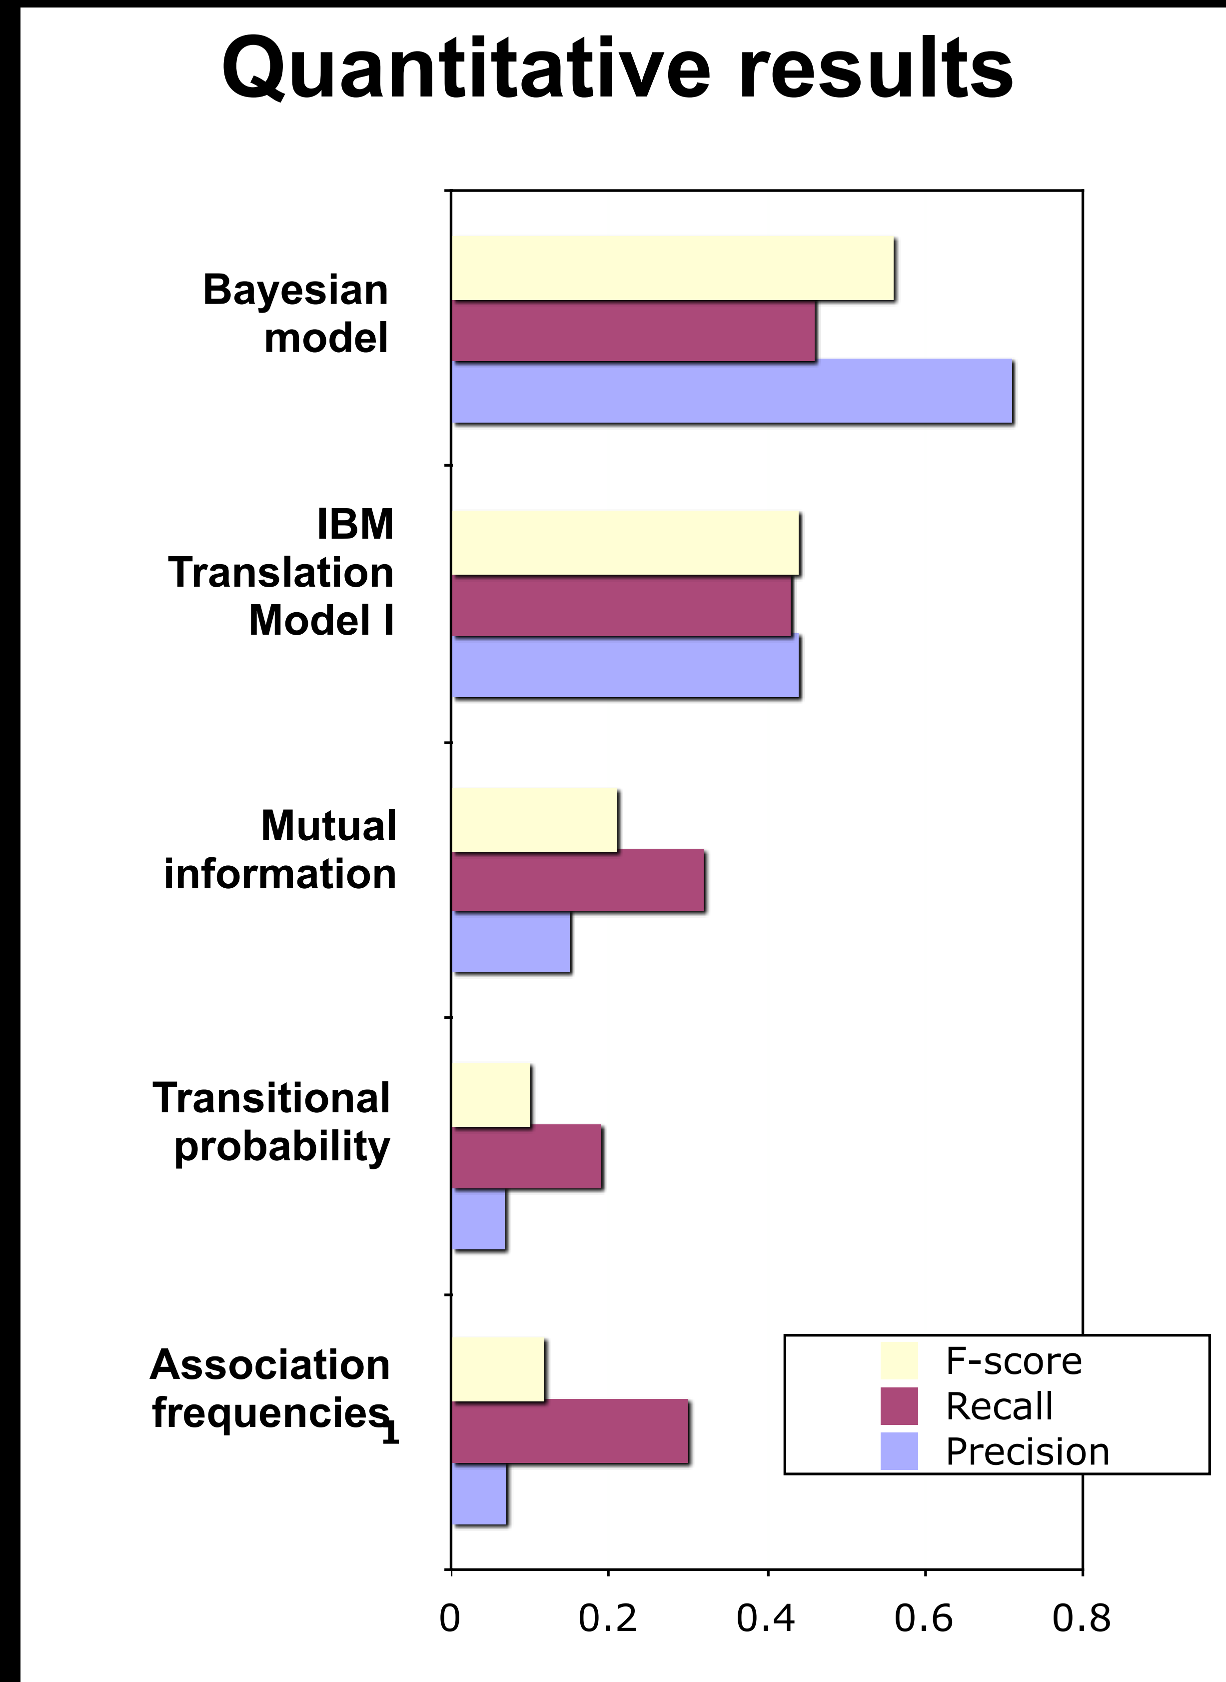

Corpus Results

Best Lexicon:

Word	Object
bear	**bear**
bigbird	**bird**
bird	**duck**
birdie	**duck**
book	**book**
bottle	bear
bunnies	**bunny**
bunnyrabbit	**bunny**
hand	**hand**
hat	**hat**
hiphop	mirror
kittycat	**kitty**
lamb	**lamb**
laugh	cow
meow	baby
mhmm	hand
mirror	**mirror**
moocow	**cow**
oink	**pig**
on	ring
pig	**pig**
put	ring
ring	**ring**
sheep	**sheep**

Quantitative results

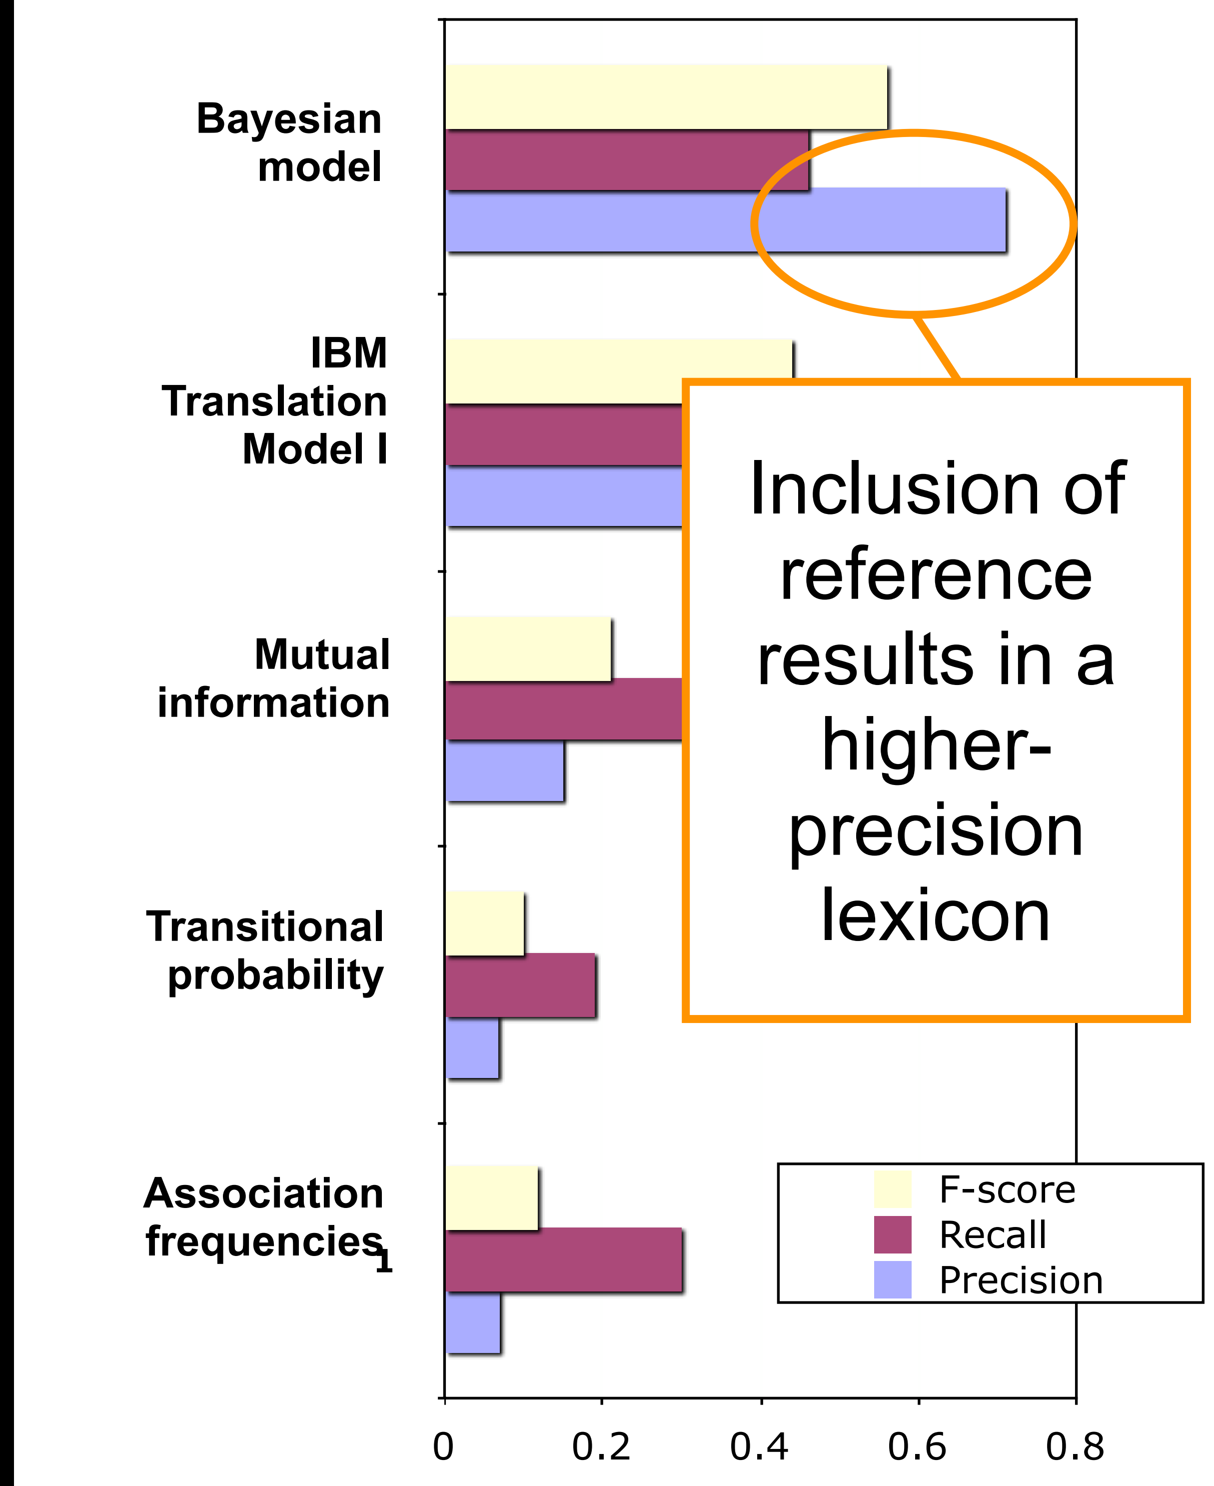

Mutual exclusivity

- Mutual exclusivity:
  A novel word is mapped to a novel object.
- This follows for free from *explaining away* and the *size principle*:
  - Conditioned on the situation, BIRD-dax and NOVEL-dax mappings are dependent.
  - BIRD-dax is unlikely because BIRD has never occurred with dax before.

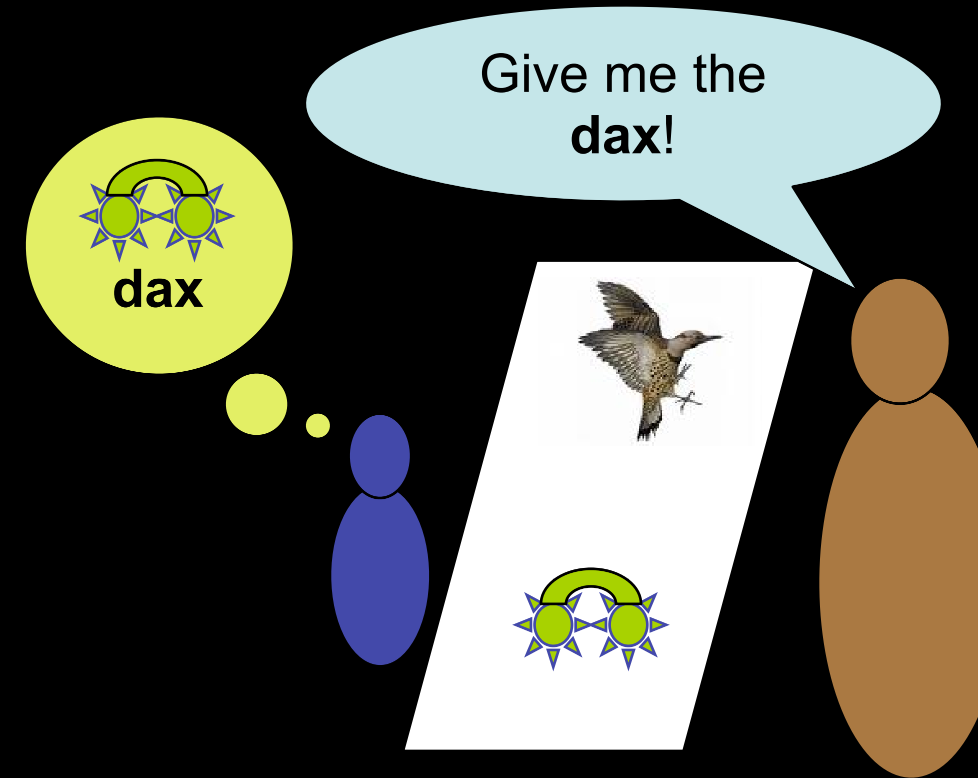

Ad-hoc implicature

- Children fail standard scalar implicature until 5 or 6yrs.
(E.g. Papafragou & Musolino, 2003; Noveck, 2001)

Stiller, Goodman, Frank (2011)

Ad-hoc implicature

- Children fail standard scalar implicature until 5 or 6yrs. (E.g. Papafragou & Musolino, 2003; Noveck, 2001)

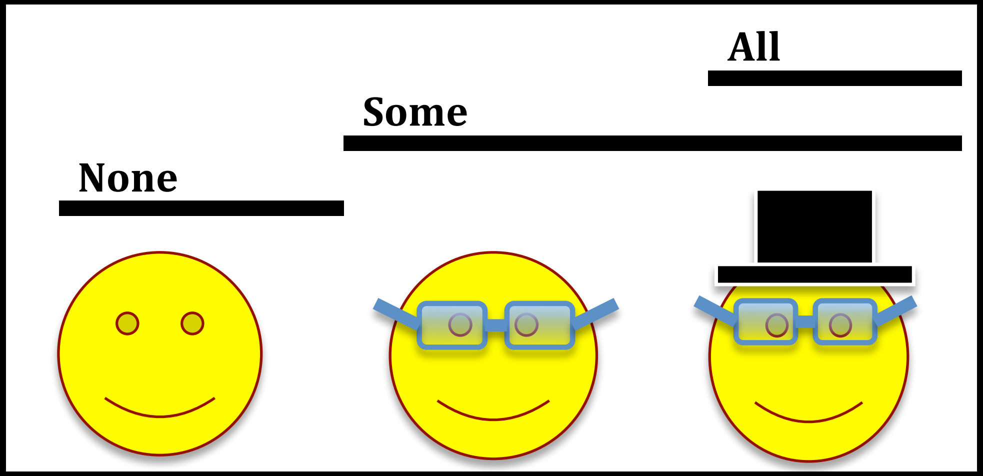

“**My friend has glasses**.”
“Can you show me my friend?”

Stiller, Goodman, Frank (2011)

Ad-hoc implicature

- Children fail standard scalar implicature until 5 or 6yrs. (E.g. Papafragou & Musolino, 2003; Noveck, 2001)

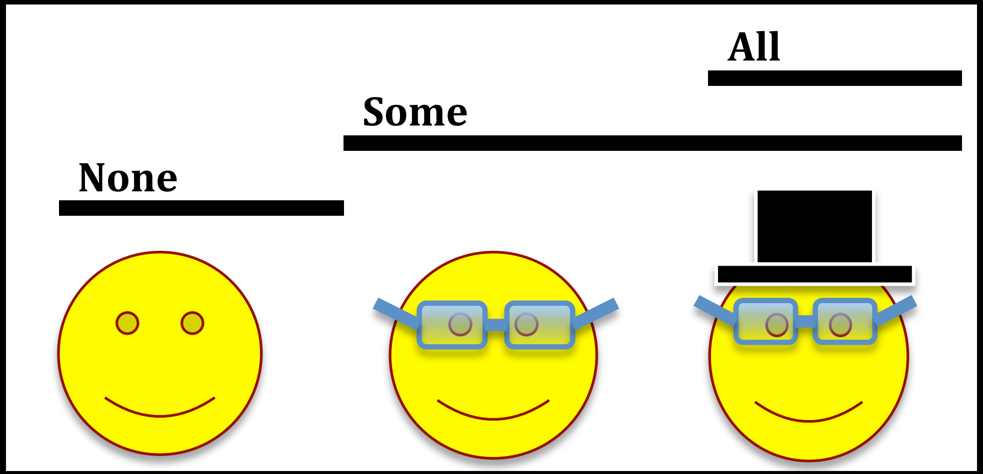

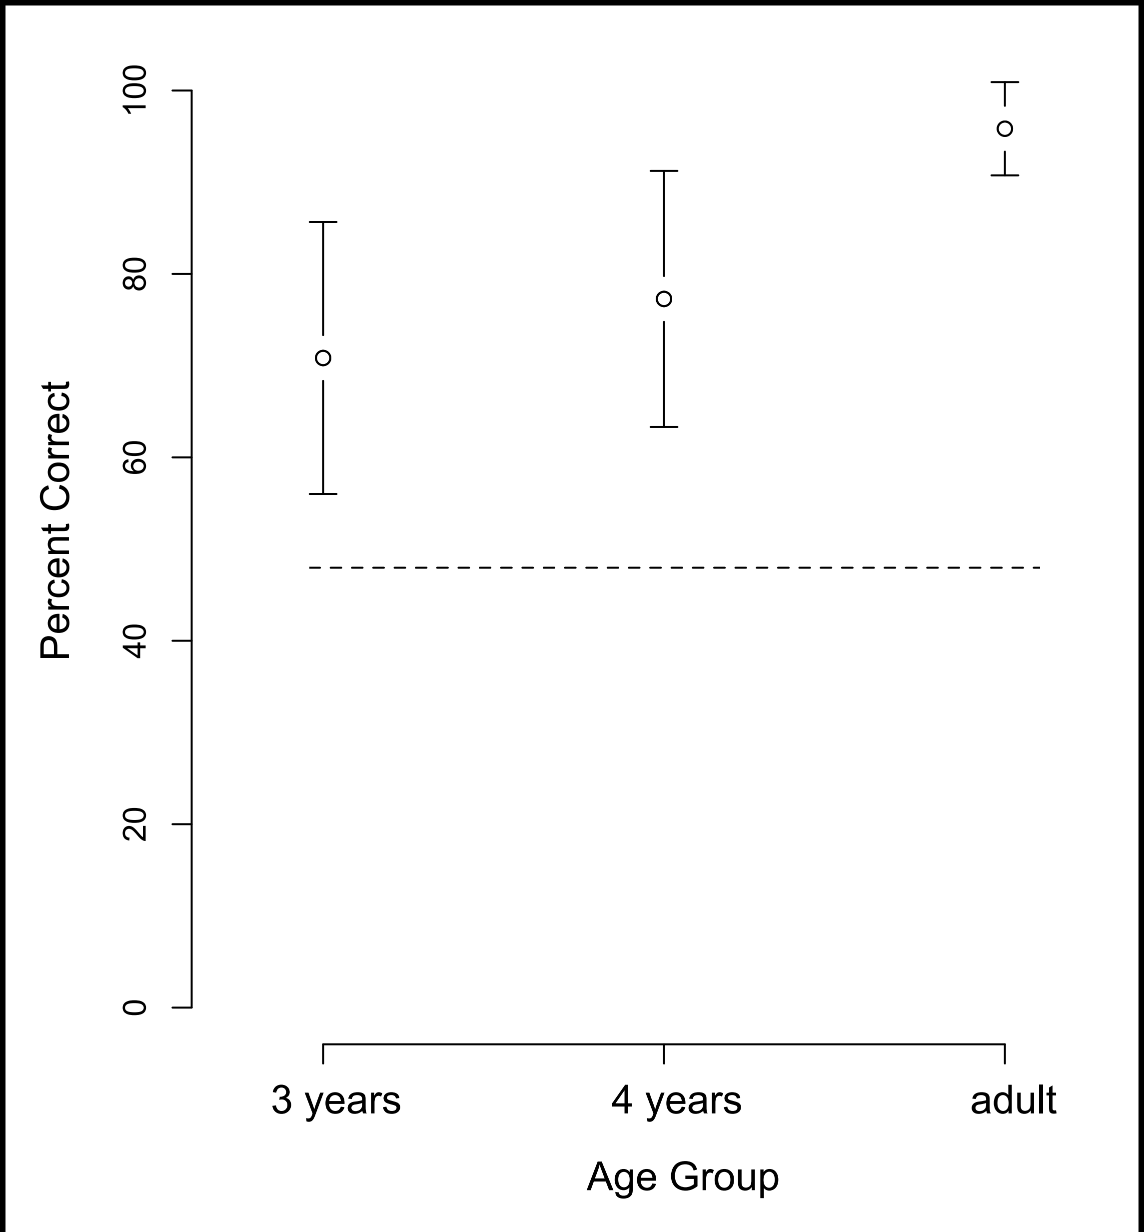

"My friend has glasses."
"Can you show me my friend?"

Stiller, Goodman, Frank (2011)

Ad-hoc implicature

- Children fail standard scalar implicature until 5 or 6yrs. (E.g. Papafragou & Musolino, 2003; Noveck, 2001)

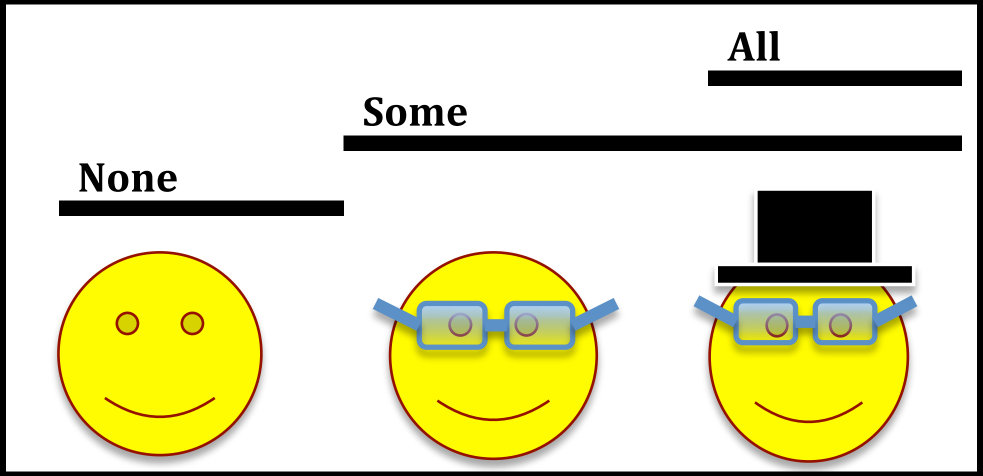

"**My friend has glasses.**"
"Can you show me my friend?"

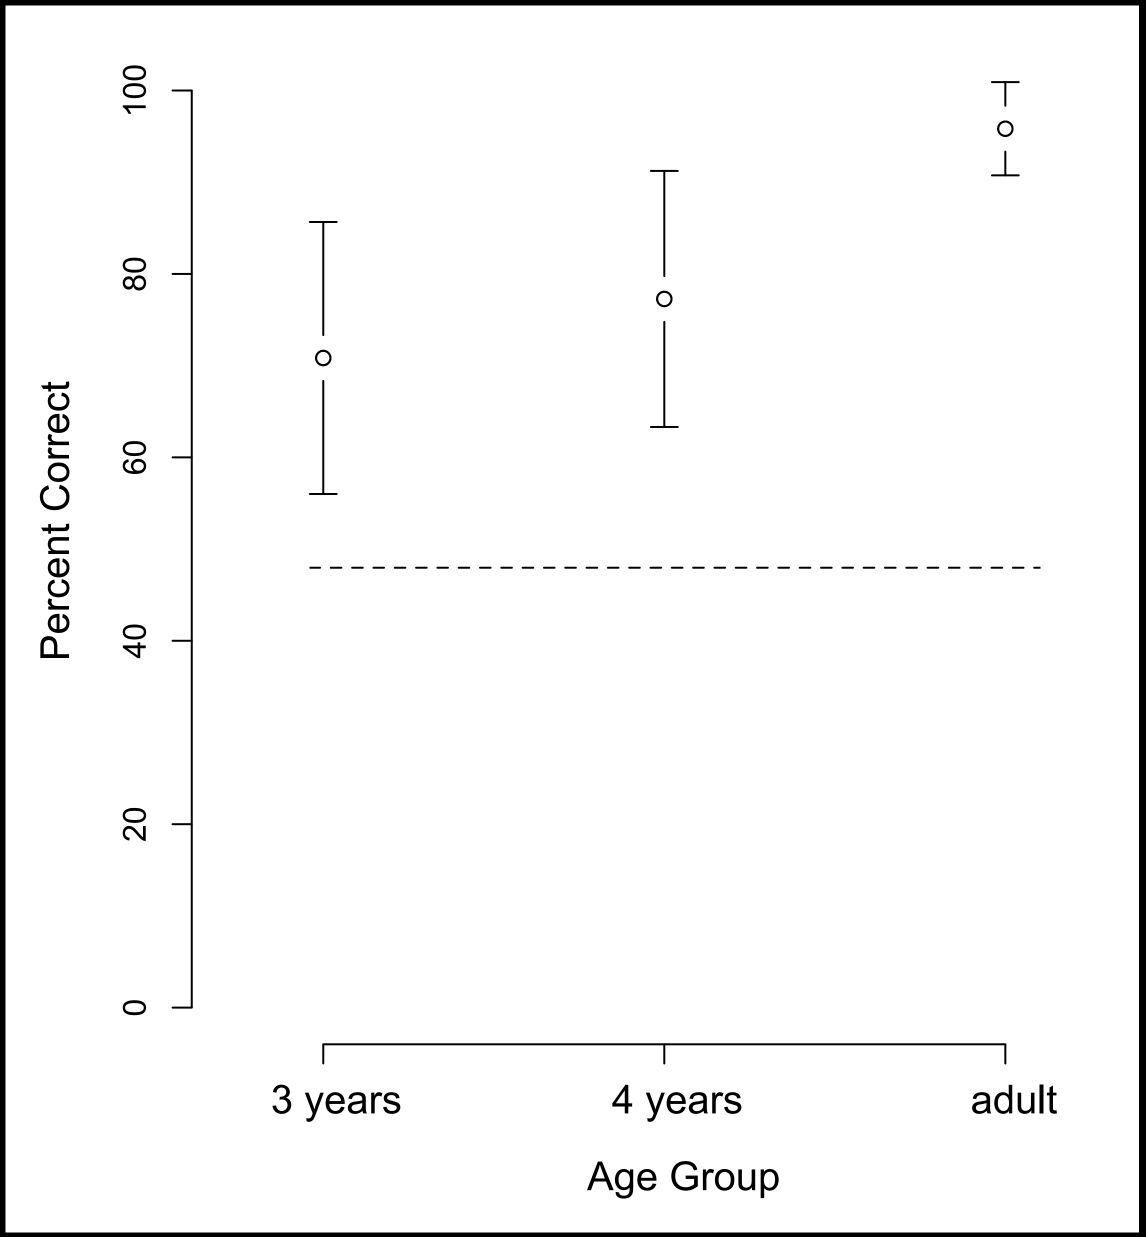

Children can do ad-hoc implicature!

Stiller, Goodman, Frank (2011)

Implicature for learning

Does daxy mean square or dotted?

Frank, Goodman, Lai, Tenenabaum (2009)
Frank & Goodman (in prep)

Implicature for learning

Does daxy mean square or dotted?

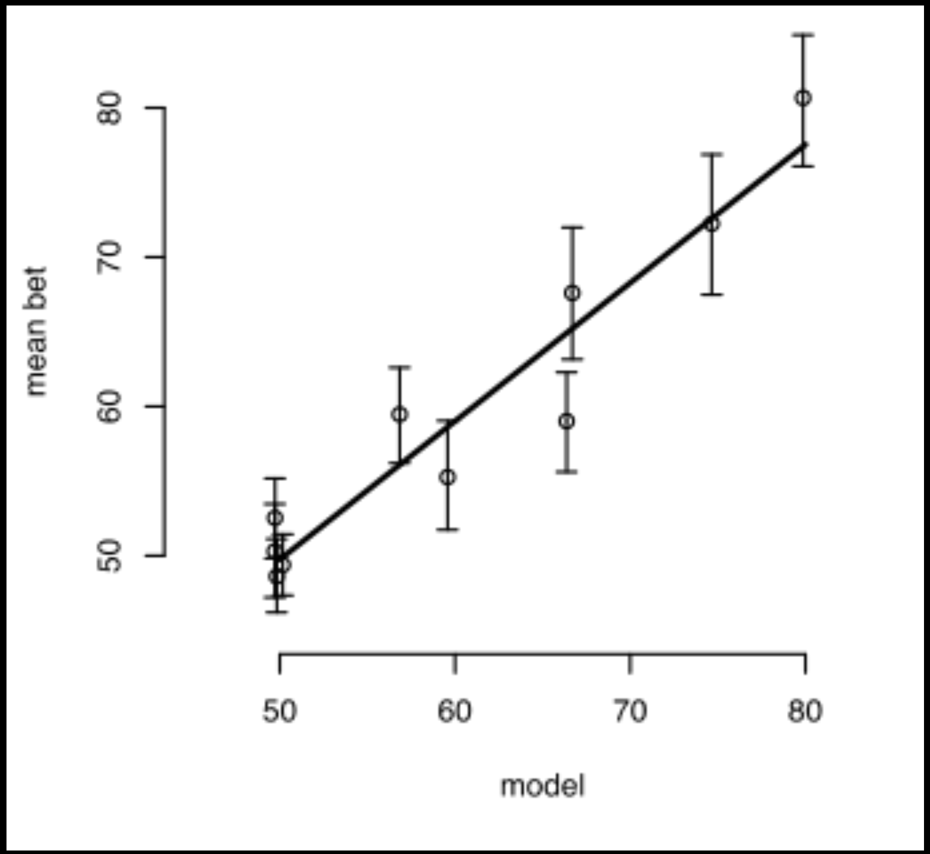

Adults can infer word meanings by way of implicature.
(Predicted by model.)

Frank, Goodman, Lai, Tenenabaum (2009)
Frank & Goodman (in prep)

Implicature for learning

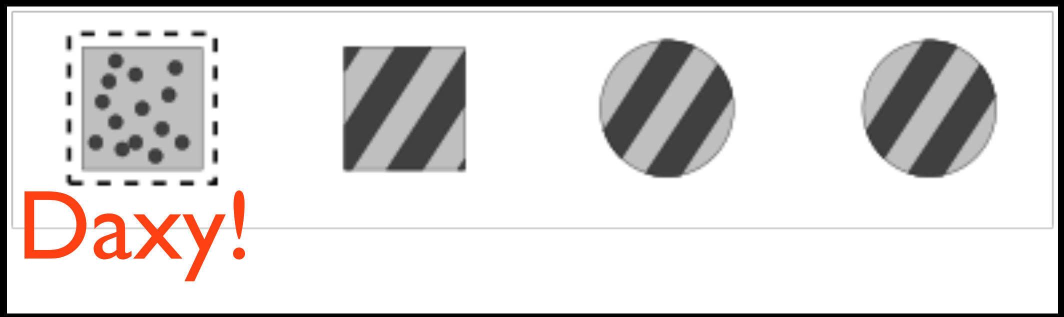

Does daxy mean square or dotted?

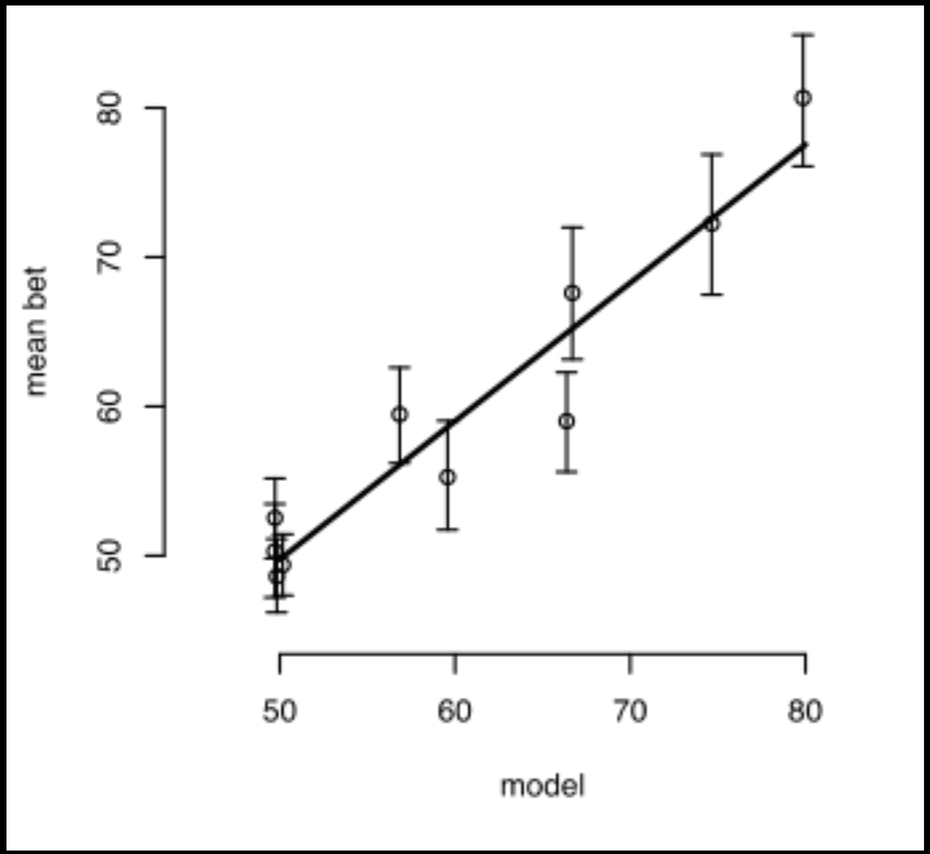

Adults can infer word meanings by way of implicature.
(Predicted by model.)

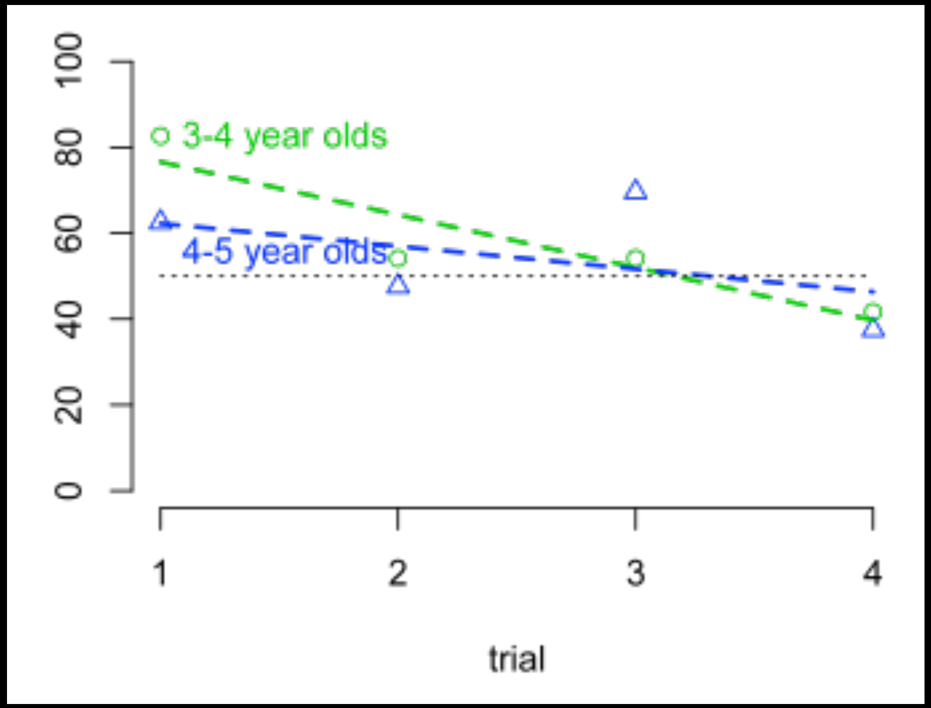

So can kids!

Frank, Goodman, Lai, Tenenabaum (2009)
Frank & Goodman (in prep)

Semantics...

- What are literal meanings?
  - Conditioning statements used to update prior distribution.
- How are they built compositionally?
  - How does formal semantics change when moving from λ-calculus to $\psi\lambda$-calculus?
- How do non-literal meanings arise?
  - From interactions with an intuitive theory of mind.

Semantics...

- What are literal meanings?
  - Conditioning statements used to update prior distribution.
- How are t … itionally?
  - How does … hange when moving fro … λ-calculus?
- How do non-literal meanings arise?
  - From interactions with an intuitive theory of mind.

A current focus of Stanford CoCoLab!

The end