The scientific process as cumulative

1 September 2021

Modern Research Methods

Molly Lewis

Course Website: https://cumulativescience.netlify.com/

Business

- Qs about syllabus?
- Survey
  - Geared toward particular area of research?
  - Group vs individual work?
  - Lots of support for learning R?
- Lab Friday
  - Bring laptop
  - Have R, RStudio and tidyverse installed
  - Directions on website – we'll help you in lab if stuck

Last Time: Cumulative Science

The Scientific Process

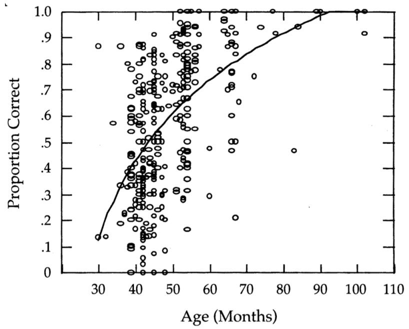

THEORY 1

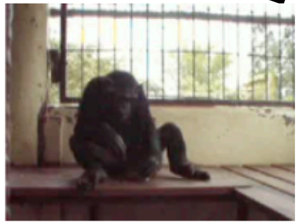

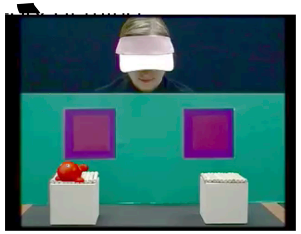

THEORY 2

Today: An introduction to cumulative science tools

How do kids learn the meaning of new words?

Graduate Student, Molly

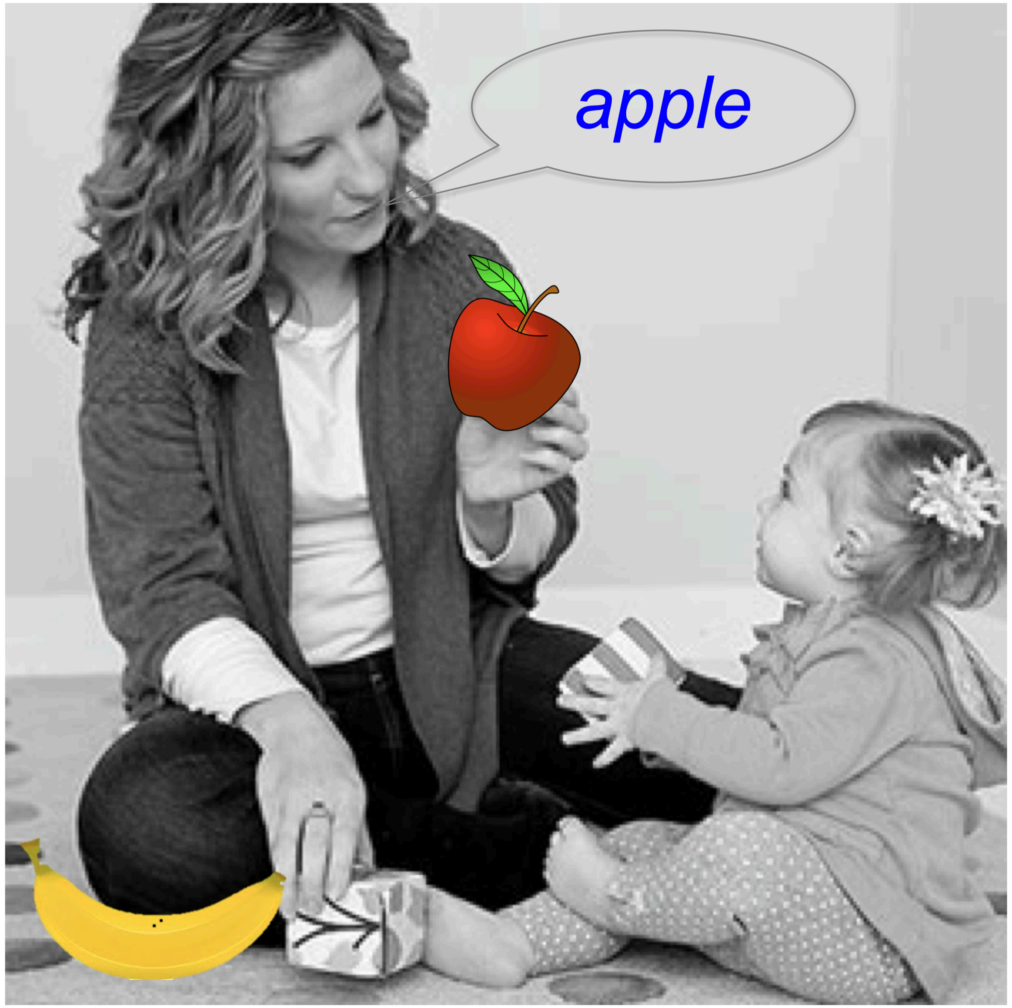

There are infinite possible meanings in the local environment when a child hears a new word, how to figure out right one?

But, it gets even harder…

“dax”

Proposal in the literature

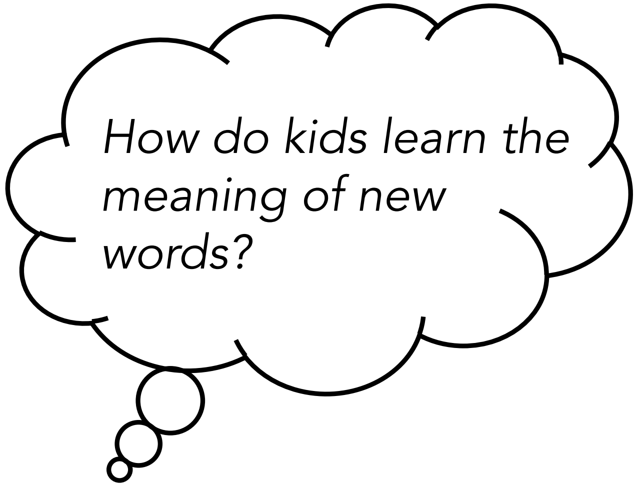

Psychological Review
2007, Vol. 114, No. 2, 245–272

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.2.245

Word Learning as Bayesian Inference

Fei Xu
University of British Columbia

Joshua B. Tenenbaum
Massachusetts Institute of Technology

The authors present a Bayesian framework for understanding how adults and children learn the meanings of words. The theory explains how learners can generalize meaningfully from just one or a few positive examples of a novel word's referents, by making rational inductive inferences that integrate prior knowledge about plausible word meanings with the statistical structure of the observed examples. The theory addresses shortcomings of the two best known approaches to modeling word learning, based on deductive hypothesis elimination and associative learning. Three experiments with adults and children test the Bayesian account's predictions in the context of learning words for object categories at multiple levels of a taxonomic hierarchy. Results provide strong support for the Bayesian account over competing accounts, in terms of both quantitative model fits and the ability to explain important qualitative phenomena. Several extensions of the basic theory are discussed, illustrating the broader potential for Bayesian models of word learning.

Keywords: word learning, Bayesian inference, concepts, computational modeling

Let's try it out

P("dax" means dog) =
P("dax" means dalmation) =

"dax"

P("dax" means dalmation) =
P("dax" means dog) =

If I'm picking examples from the dalmation category, I'm more likely to pick three dalmations

If I'm picking examples from the dog category, it would be <u>really</u> unlikely to pick three dalmations

It would be a "suspicious coincidence"!

Xu and Tenenbaum (2007)

The Size Principle

dax

dalmatian? (Subordinate)

dog? (Basic)

animal? (Superordinate)

In general, more exemplars make the more specific category more likely.

dax

dax

dax

dalmatian (Subordinate)

dog (Basic)

animal (Superordinate)

Xu and Tenenbaum (2007)

Testing the suspicious coincidence effect

Here is a rab.

Can you give Mr. Frog all the other rabs?

To give a rab, click on it below. When you have given all the rabs, click the Next button.

Each participant saw some "1 example" trials, and some "3 example" trials

Children and adults make this inference

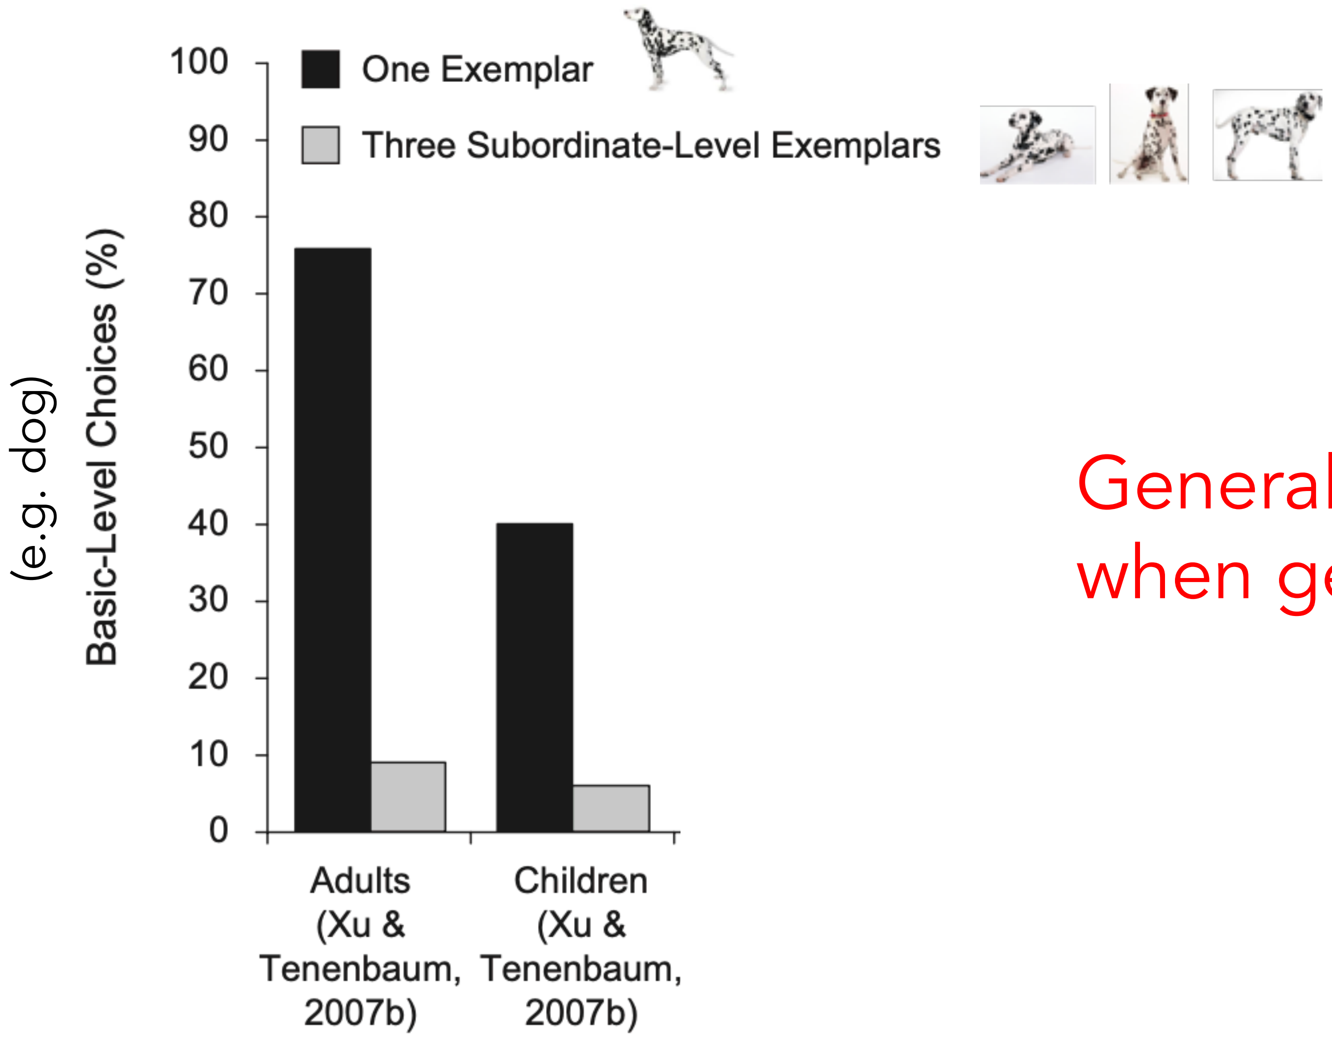

Generalize word <u>more</u> narrowly when get more examples!

Xu and Tenenbaum (2007)

A theory of how children could learn the meaning of new words at **multiple levels of abstraction**.

“dax”

NUMBER of examples of word meaning provides information

Psychological Review
2007, Vol. 114, No. 2, 245–272

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.2.245

Word Learning as Bayesian Inference

Fei Xu
University of British Columbia

Joshua B. Tenenbaum
Massachusetts Institute of Technology

The authors present a Bayesian framework for understanding how adults and children learn the meanings of words. The theory explains how learners can generalize meaningfully from just one or a few positive examples of a novel word’s referents, by making rational inductive inferences that integrate prior knowledge about plausible word meanings with the statistical structure of the observed examples. The theory addresses shortcomings of the two best known approaches to modeling word learning, based on deductive hypothesis elimination and associative learning. Three experiments with adults and children test the Bayesian account’s predictions in the context of learning words for object categories at multiple levels of a taxonomic hierarchy. Results provide strong support for the Bayesian account over competing accounts, in terms of both quantitative model fits and the ability to explain important qualitative phenomena. Several extensions of the basic theory are discussed, illustrating the broader potential for Bayesian models of word learning.

Keywords: word learning, Bayesian inference, concepts, computational modeling

2007

Psychological Review
2007, Vol. 114, No. 2, 245–272

0033-295X/07/$12.00 DOI: 10.1037/0033-295X.114.2.245

Word Learning as Bayesian Inference

Fei Xu
University of British Columbia

Joshua B. Tenenbaum
Massachusetts Institute of Technology

The authors present a Bayesian framework for understanding how adults and children learn the meanings of words. The theory explains how learners can generalize meaningfully from just one or a few positive examples of a novel word's referents, by making rational inductive inferences that integrate prior knowledge about plausible word meanings with the statistical structure of the observed examples. The theory addresses shortcomings of the two best known approaches to modeling word learning, based on deductive hypothesis elimination and associative learning. Three experiments with adults and children test the Bayesian account's predictions in the context of learning words for object categories at multiple levels of a taxonomic hierarchy. Results provide strong support for the Bayesian account over competing accounts, in terms of both quantitative model fits and the ability to explain important qualitative phenomena. Several extensions of the basic theory are discussed, illustrating the broader potential for Bayesian models of word learning.

Keywords: word learning, Bayesian inference, concepts, computational modeling

Google Scholar

word learning as bayesian inference

Articles

About 194,000 results (**0.23** sec)

Any time
Since 2020
Since 2019
Since 2016
Custom range...

Word learning as **Bayesian inference**.
F Xu, JB Tenenbaum - Psychological review, 2007 - psycnet.apa.org
The authors present a Bayesian framework for understanding how adults and children learn the meanings of words. The theory explains how learners can generalize meaningfully from just one or a few positive examples of a novel word's referents, by making rational inductive …
Cited by 795 Related articles All 19 versions Web of Science: 314

2011

Wait a minute!

Research Article

aps ASSOCIATION FOR PSYCHOLOGICAL SCIENCE

Psychological Science
22(8) 1049–1057

DOI: 10.1177/0956797611413934
http://pss.sagepub.com
SAGE

Learning Words in Space and Time: Probing the Mechanisms Behind the Suspicious-Coincidence Effect

John P. Spencer[1], Sammy Perone[1], Linda B. Smith[2], and Larissa K. Samuelson[1]
[1]Department of Psychology and Delta Center, University of Iowa, and [2]Department of Psychological and Brain Sciences, Indiana University

Your theory is wrong…

“The striking finding that led Xu and Tenenbaum (2007b) to this conclusion—broader generalization from a single instance than from three (nearly identical) instances—**is also consistent with mechanistic accounts** couched in terms of memories and representations for learning events. […] In the case of the suspicious-coincidence effect, two such task factors may be particularly critical: The fact that the exemplars **are simultaneously visible** in the task space and that they are nearly identical instances in **close spatial proximity.** ” – Spencer, et al. (2011)

Your theory predicts that it’s just NUMBER of examples but other things might matter too.

Sequential presentation of exemplars

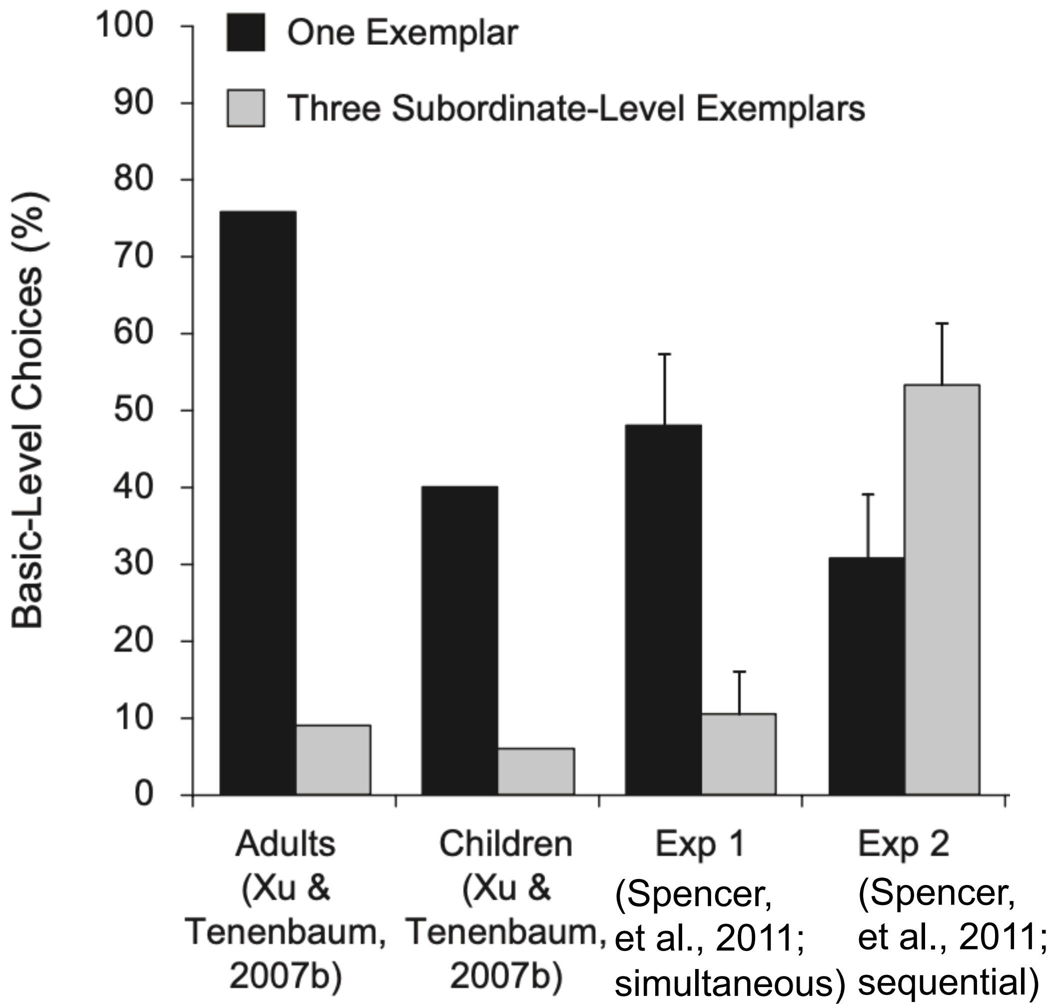

...sequential makes the effect reverse??

Resolving the conflict in this literature

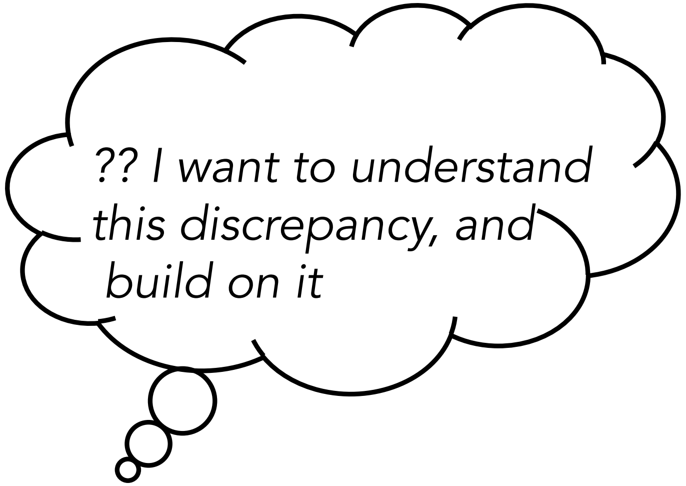

2018

Did a <u>replication</u> of both studies.

REPLICATE = Repeat a study with the same population, hypothesis, experimental design, and analysis plan and get same result (Patil, et al. 2016)

Replicating previous results

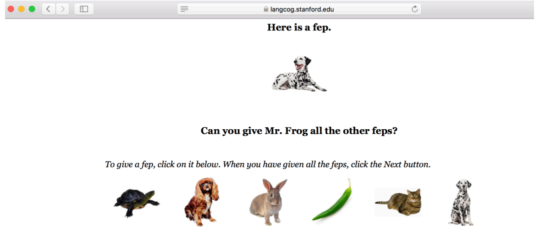

- Stimuli and code from original experiments weren't available so I had to implement using Javascript and HTML (https://tinyurl.com/ry3tvyz)
- Cleaned and analyzed data in R
- Before I ran my study, I pre-registered experimental code/analysis plan (https://osf.io/wgvcw) - why?
- Conducted a replication of these studies online using a large sample (*N* = 600) of participants

Reproducibility

REPRODUCE = Repeat procedure (e.g. experimental code, analysis code) and get same result

- All my code is available online so that other researchers can **reproduce** my experiment and analysis
- Website called Github (https://github.com/)
- https://github.com/mllewis/XTMEM

GitHub

A methodological difference between two studies…

- Xu and Tenenbaum (2007) – 1 trial first, then 3 trial
- Spencer et al (2011) – 3 trial first then 1 trial

- Might this matter? Who knows – I'll test both.

What did I find?

Replication of Xu and Tenenbaum (2007)

Replication of Spencer et al. (2011)

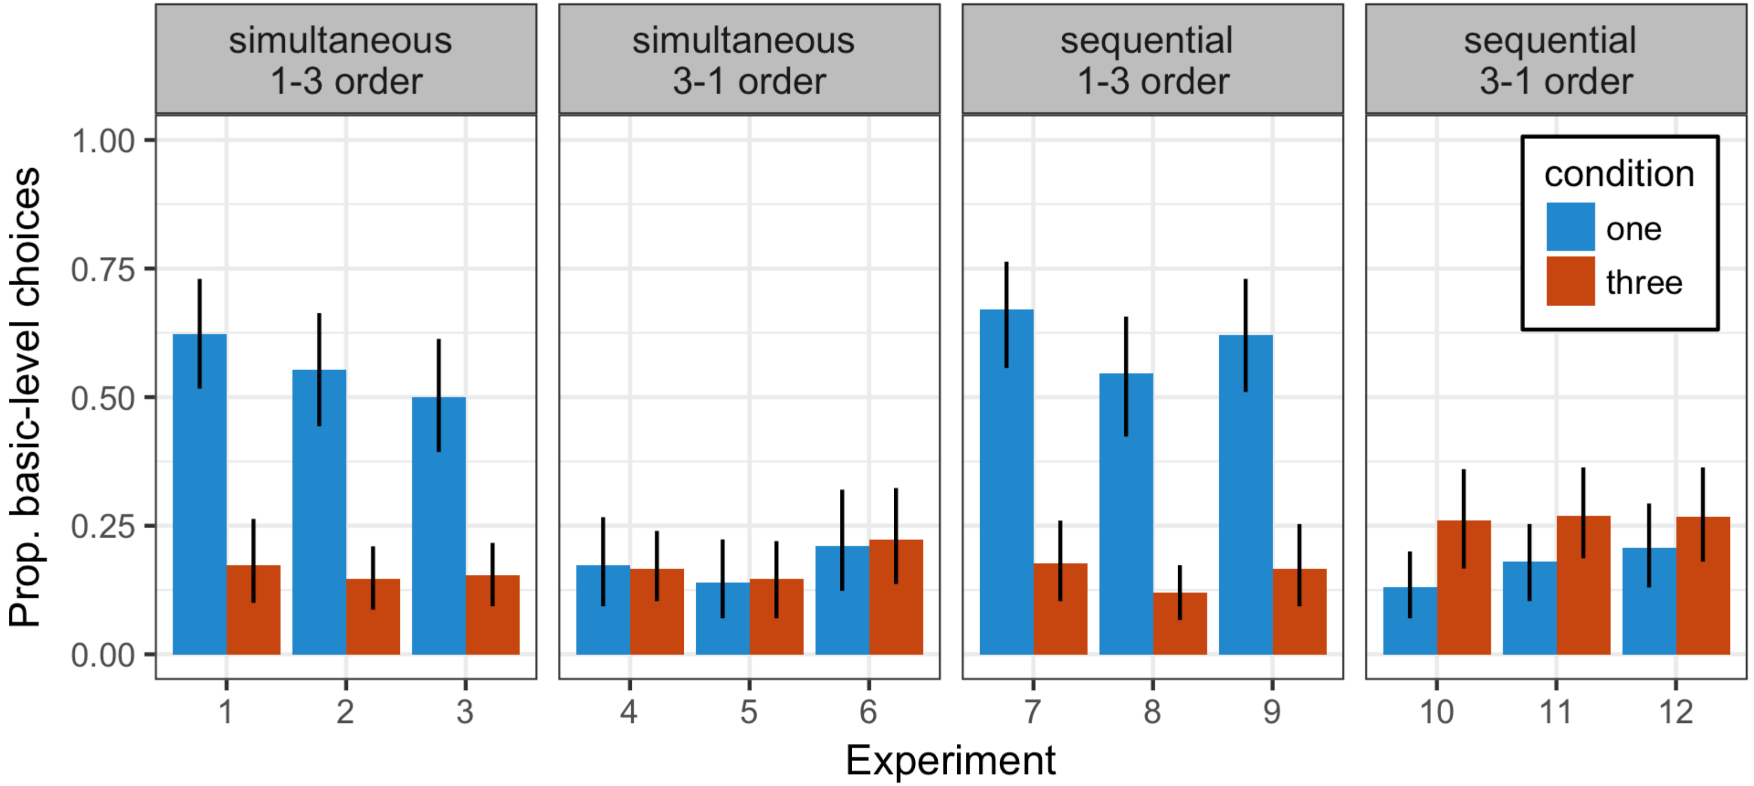

(Lewis & Frank, *Psych. Science*, 2018)

Trial order matters!

- Only see the suspicious coincidence effect in the 1-3 ordering
- How can we test this idea?
- Effect sizes and meta-analysis

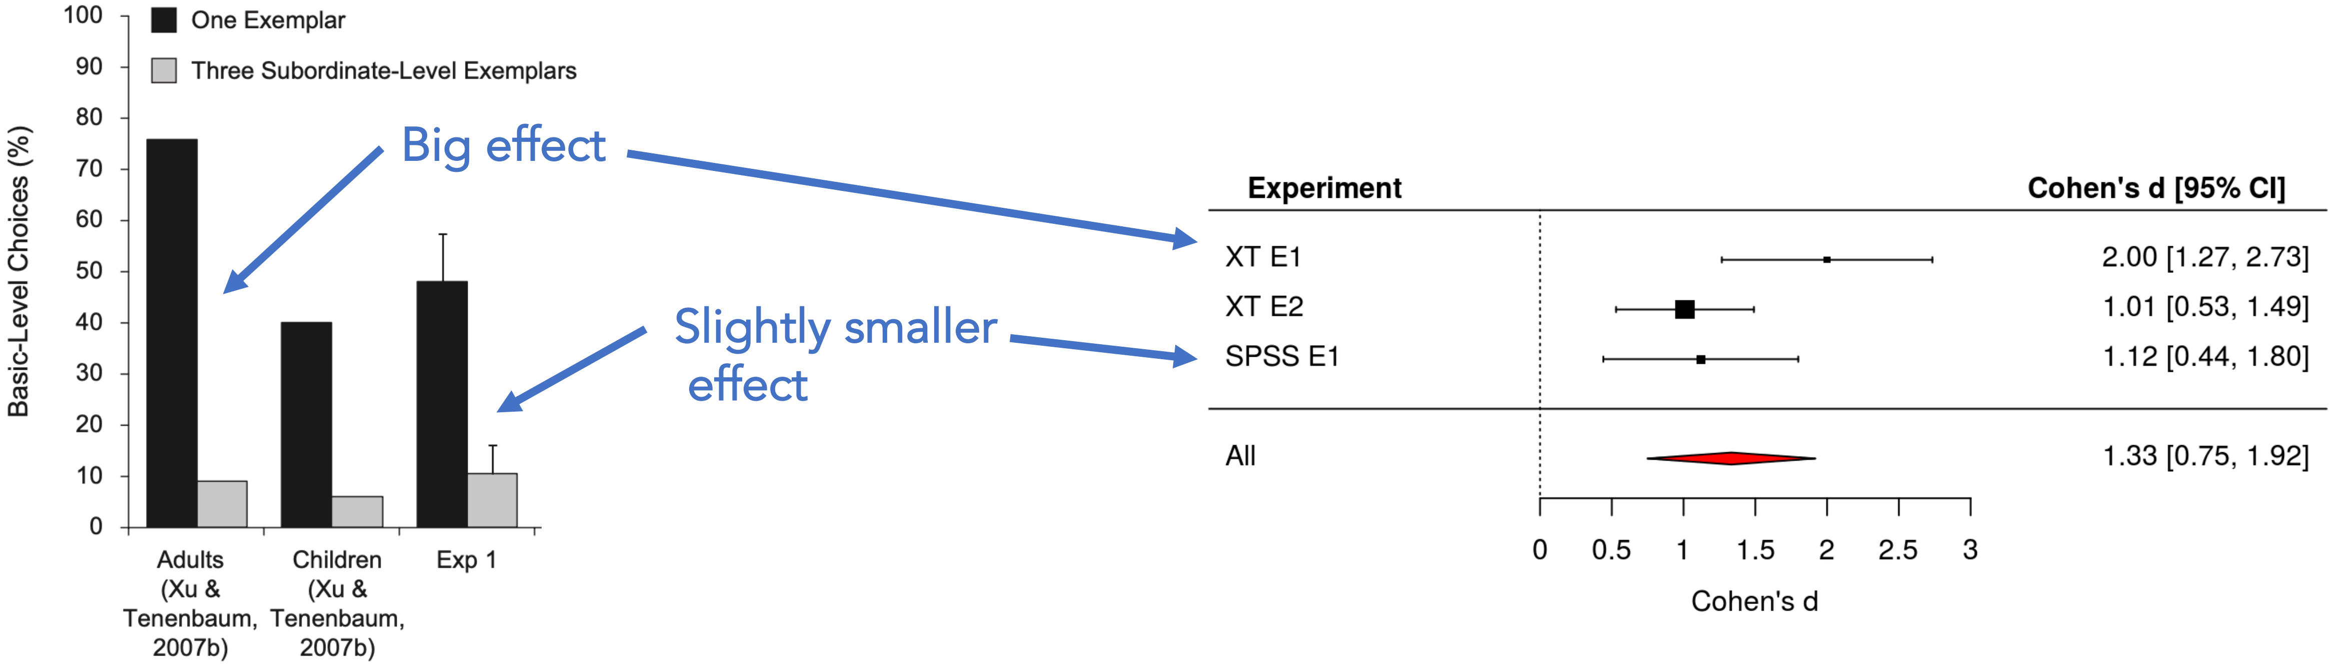

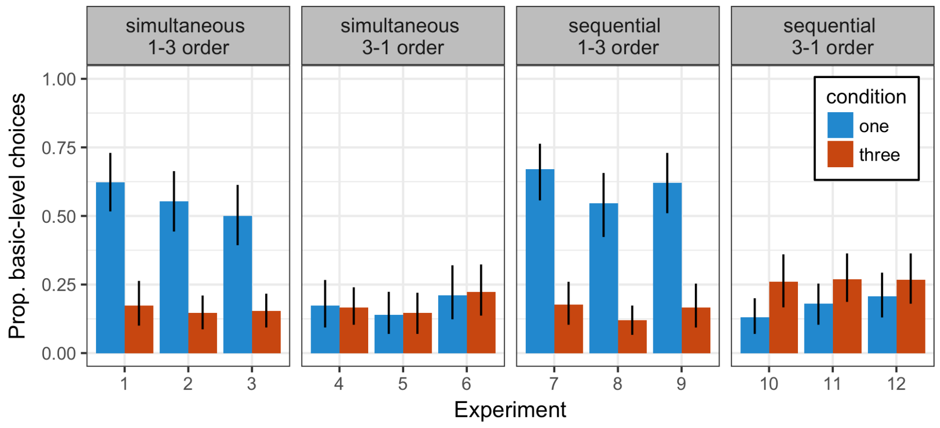

simultaneous 1-3 order
simultaneous 3-1 order
sequential 1-3 order
sequential 3-1 order
condition
one
three
Prop. basic-level choices
1.00
0.75
0.50
0.25
0.00
1
2
3
4
5
6
7
8
9
10
11
12
Experiment

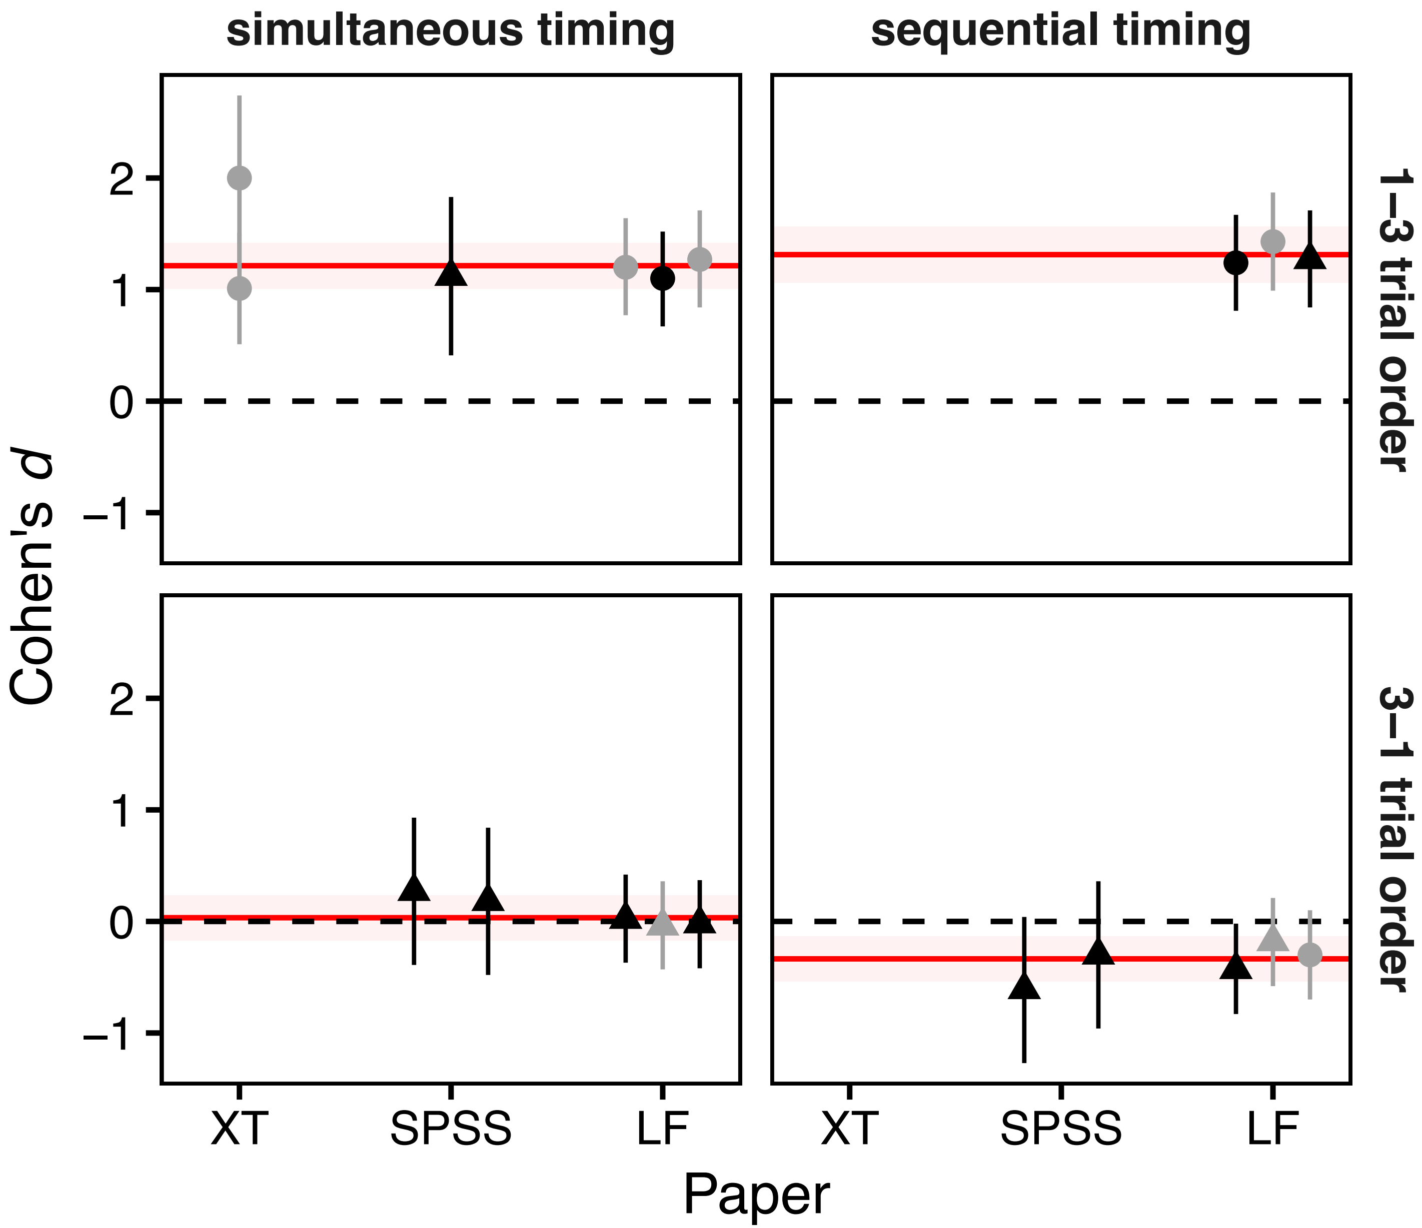

simultaneous timing
sequential timing
1–3 trial order
3–1 trial order
Cohen's d
2
1
0
−1
XT
SPSS
LF
Paper

Peer Review Process

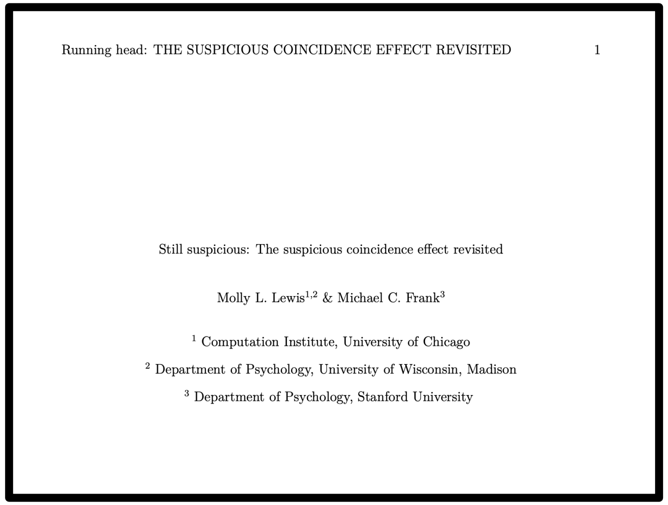

Running head: THE SUSPICIOUS COINCIDENCE EFFECT REVISITED 1

Still suspicious: The suspicious coincidence effect revisited

Molly L. Lewis[1,2] & Michael C. Frank[3]

[1] Computation Institute, University of Chicago
[2] Department of Psychology, University of Wisconsin, Madison
[3] Department of Psychology, Stanford University

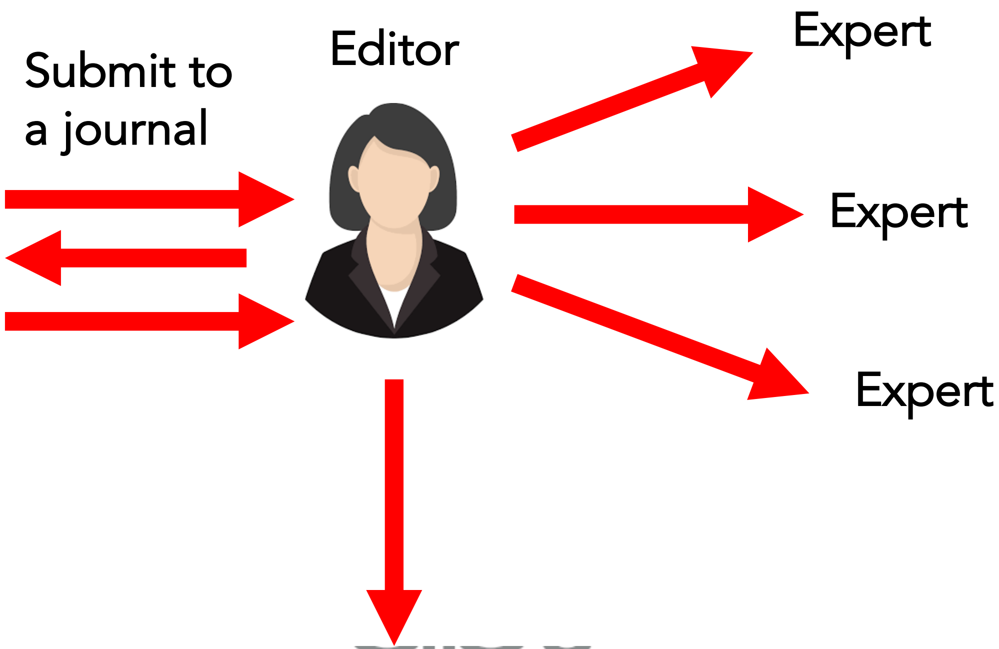

ASSOCIATION FOR PSYCHOLOGICAL SCIENCE

Preregistered Direct Replication

Still Suspicious: The Suspicious-Coincidence Effect Revisited

Psychological Science
2018, Vol. 29(12) 2039–2047

DOI: 10.1177/0956797618794931
www.psychologicalscience.org/PS
SAGE

Molly L. Lewis[1,2] and Michael C. Frank[3]
[1]Computation Institute, University of Chicago; [2]Department of Psychology, University of Wisconsin–Madison; and [3]Department of Psychology, Stanford University

Tools I used in this project

- Data analysis and visualization in R
- Preregistration
- Replications
- Reproducible workflows (e.g. Github)
- Effect sizes and meta-analysis

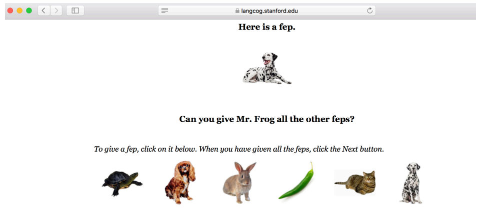

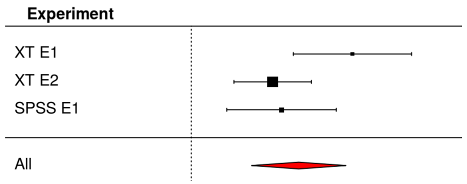

- In this class, you will learn about all of these tools
- You will not master any of them, but my goal is to introduce them to you so you can have the ability to learn more

Next Time: Introduction to R (Lab)

- Bring laptop, install R and R Studio

Chapter 3 Getting started with R

> *Robots are nice to work with.*
>
> –Roger Zelazny[13]

In this chapter I'll discuss how to get started in R. I'll briefly talk about how to download and install R, but most of the chapter will be focused on getting you started typing R commands. Our goal in this chapter is not to learn any statistical concepts: we're just trying to learn the basics of how R works and get comfortable interacting with the system. To do this, we'll spend a bit of time using R as a simple calculator, since that's the easiest thing to do with R. In doing so, you'll get a bit of a feel for what it's like to work in R. From there I'll introduce some very basic programming ideas: in particular, I'll talk about the idea of defining *variables* to store information, and a few things that you can do with these variables.

tidyverse part of the tidyverse 1.3.1.9000 Reference

Welcome to the Tidyverse

Hadley Wickham, Mara Averick, Jennifer Bryan, Winston Chang, Lucy D'Agostino McGowan, Romain François, Garrett Grolemund, Alex Hayes, Lionel Henry, Jim Hester, Max Kuhn, Thomas Lin Pedersen, Evan Miller, Stephan Milton Bache, Kirill Müller, Jeroen Ooms, David Robinson, Dana Paige Seidel, Vitalie Spinu, Kohske Takahashi, Davis Vaughan, Claus Wilke, Kara Woo, Hiroaki Yutani