

Pragmatics, probabilities & psychologists: A Bayesian perspective on some reasoning problems

> Dani Navarro School of Psychology University of New South Wales

> > compcogscisydney.com



Contributors, coauthors, collaborators

Keith Ransom Wouter Voorspoels Steph Banner Ally Tingey Jessica Lee

Amy Perfors Brett Hayes Drew Hendrickson Peter Lovibond Michelle Keshwa Matthew Welsh Matt Dry Michael Lee Chris Donkin

Kristy Martire Ben Newell Wai Keen Vong Lauren Kennedy Steve Langsford Candy Liu Anastasia Ejova Rachel Stephens **Gert Storms** Pat Shafto **Baxter Eaves** Charles Kemp Nancy Briggs **Titia Benders** Sean Tauber





Australian Government

Australian Research Council





What are the connections between human reasoning and statistical inference?

What should we do with this sample of evidence?



The problem of inductive generalisation



What factors shape our inductive inferences?



What factors shape our inductive inferences?



Reasoners consider hypotheses





The sample rules out some and not others...



Inductive generalisation is based on hypotheses consistent with the sample







Epistemic vigilance: Statistical reasoning about untrustworthy data

These birds have plaxium blood





Does this bird have plaxium blood?





I This is silly, but "it's all made up" is <u>absolutely</u> a legitimate sampling assumption Does this bird have plaxium blood?

The price of inductive freedom is epistemic vigilance



Shafto, Eaves, Navarro & Perfors (2012) Developmental Science Mascaro & Sperber (2009)

The price of inductive freedom is epistemic vigilance



Shafto, Eaves, Navarro & Perfors (2012) Developmental Science Mascaro & Sperber (2009)

Why epistemic vigilance?

People will try to "mislead with a half truth" if the listener is naive...





Ransom, Voorspoels, Perfors & Navarro (2017)





Ransom, Voorspoels, Perfors & Navarro (2017) Everyday reasoning about the world is intertwined with social reasoning about other people



Why are you telling me this? Why are you telling me this?

Where did you hear this?

Do you even *know* what you're talking about?



What do you want me to do with this information?

What does all this buy us? Taking a hint from a helpful teacher

Ransom, Perfors & Navarro (2016). *Cognitive Science*









Sampling mechanism:



Prediction:



Adding positive instances has minimal effect if they're too similar to things I already know about

Adding positive instances from the same category conveys *intent*, and drives attention to that category



Previous experience? (filler trials)

	Helpful cover story, filler trials imply helpful	
Cover story?	Neutral cover story, filler trials imply helpful	Neutral cover story, filler trials imply random
		Random cover story, filler trials imply random









-0.2













Ι



Knowledge about animal categories (theory of the world) creates structural differences between the different arguments



The sampling model (theory of the context) describes how "adding more data" can have different effects across conditions and arguments


Using negative evidence to take hints from helpful teachers

Voorspoels, Navarro, Perfors, Ransom & Storms (2015). *Cognitive Psychology*

Positive evidence



Mozart produces alpha waves in the brain

This seems helpful!

Negative evidence

This... not so much



The sound of a falling rock does not

see Hempel (1945), Good (1960), etc

Okay, we start by telling people that Mozart does produce alpha waves...



+Mozart





+Mozart

... and they reason sensibly



+Mozart

-Metallica

Adding Metallica as a negative example has a modest, sensible effect on inferences about Nirvana



Um.



+Mozart



-Falling rock







Negative evidence is interpreted as marking the category boundary



Bayesian reasoners with a random sampling assumption do *not* produce the effect





Bayesian reasoners with a helpful sampling assumption *do* produce the effect



What does it mean to be "helpful" anyway?





Prediction:

If the negative evidence is perceived as a helpful hint we should continue to get the effect

If it is construed as an arbitrary fact, the effect should vanish



Here's the experimental results:

Hint Arbitrary







A⁺ A⁺ C⁻

Superficially useless information can have a huge effect when it is <u>deemed</u> to be **helpful**

WTF is this "falling rocks" thing? It must be **relevant** somehow, so...

Extension: Negative evidence, fear conditioning & inductive reasoning

(work in progress!)

Fear conditioning*





Negative evidence along the same dimension ("near" CS-)





CS- decreases generalisation on this side



CS- increases generalisation on this side



What happens when the "far" CS- has no value on the blue-green dimension?





CS- increases generalisation across the whole dimension





These are essentially the same design



Near negative



We needed a fancy sampling assumption for this

What about this?







Bayesian reasoning with random sampling produces the wrong pattern



(aside: compare to animal results, Switalski et al 1966)

Bayesian reasoning that assumes an intentional* sampling process works*



Taking the *wrong* hint because your teacher is a jerk

(another work in progress!)

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Which is more probable?(a) Linda is a bank teller(b) Linda is a feminist bank teller

The social/pragmatic account



The social/pragmatic account





Social / pragmatic context

Navarro, Tingey, Perfors & Keshwa (in prep)



Random / disconnected fact condition

Navarro, Tingey, Perfors & Keshwa (in prep)


The "taboo" task

Generate a description that implies but does not openly state that "Linda is a feminist":



Linda is 31 and has had a rough upbringing, growing up with an abusive father which restricted her mother and her freedom. This upbringing was what made her decide to major in sociology and psychology within university. She has strong views on politics and other similar matters that affect men and women. She regularly attends rallies and protests on the weekend.

The "taboo" task

Generate a description that implies but does not openly state that "Paula is a bank teller":



Paula is 30, and loves buying clothes even at her age of 30. She is in contact with money so much that she has been able to calculate the exact change given before the cashier has given it to her. Her skills in counting are ingrained within her brain that she cannot turn it off, due to years dealing with cash

The "taboo" task

Generate a description that implies but does not openly state that "Brenda is a feminist & bank teller":



Brenda is 32 years old, methodical, logical, and passionate about her beliefs. She is very good with both people and numbers and is often able to spot errors. She is trusted by her friends to handle the money when planning an overseas trip. She is also a very individual woman and looks up to celebrities such as Emma Watson

Several different versions

"Feminist / Bank Teller", "Engineer / Jazz Musician", "Introvert / Chef", "Journalist / Anxious Person", "Painter / Accountant", "Extrovert / Statistician", "Pacifist / Boxer", "Butcher / Empath", "Writer / Mechanic"

"Mind reading" task:

Isabelle is 41 years old and is very bright and good with numbers. Her creative flair has always been a passion although until recently she didn't act on it. As a women with two professions she works extremely hard and ensures that her conflicting logical and free spirited natures are harmonious in all aspects of life.

Which of the following do you believe the writer was trying to communicate when they wrote this description:

- Isabelle is a painter
- Isabelle is an accountant
- Isabelle is both a painter and an accountant
- None of the above



Standard conjunction task:

Ryan is 26. He spends his spare time unwinding and sitting on a couch at the end of the day reading or watching a movie. He has a small but tight knit group of friends. He likes talking to them individually and dislikes group outings.

How likely is it that this person belongs to each of the following categories? Please give an estimate of the probability from 0 to 100% for each category (0 being impossible and 100 being certain).





In progress: social vs random vignettes



In progress: social vs random vignettes



More tensions between social and random sampling: variations on the Monty Hall Dilemma

(yet another work in progress!)

The Monty Hall dilemma





A suitably constrained host:

	If A is correct	If B is correct	If C is correct
Host opens A	0%		
Host opens B		0%	
Host opens C			0%

Host won't open the prize door



A suitably constrained host:

	If A is correct	If B is correct	If C is correct
Host opens A	0%	0%	0%
Host opens B		0%	
Host opens C			0%

Nor will they open the door you chose (A)



A suitably constrained host:

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	50%		100%
Host opens C	50%	100%	

Otherwise random



A suitably constrained host: a Bayesian reason to switch

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	50%	0%	100%
Host opens C			

* this is the correct solution to the original problem as stated by vos Savant



An indifferent host chooses randomly

	If A is correct	If B is correct	If C is correct
Host opens A	33%	33%	33%
Host opens B	33%	33%	33%
Host opens C	33%	33%	33%

An indifferent host: a Bayesian reason for indifference

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	33%	33%	33%
Host opens C			



A malicious host who never offers a bet when your choice was wrong!

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	p%		
Host opens C			
Host does not open a door		100%	100%



A malicious host with discretion: a Bayesian reason to stay

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	p%	0%	0%
Host opens C			
Host does not open a door			



A helpful host with discretion:

	If A is correct	If B is correct	If C is correct
Host opens A	0%	0%	0%
Host opens B		0%	
Host opens C			0%
Host does not open a door			



A helpful host with discretion:

	If A is correct	If B is correct	If C is correct
Host opens A	0%	0%	0%
Host opens B		0%	
Host opens C			0%
Host does not open a door	100%	0%	0%

A helpful host with discretion: A Bayesian reason to switch

	If A is correct	If B is correct	If C is correct
Host opens A			
Host opens B	0%	0%	100%
Host opens C			
Host does not open a door			

Probability of switching



Perfors, Navarro, Benders & Donkin (in limbo)

People (incorrectly?) view the original MHD as most similar to the malicious version





Perfors, Navarro, Benders & Donkin (in limbo)

Possibly people are treating MHD as a "social reasoning" problem, and thinking that the host is malicious?



<u>Can</u> people be sensitive to conditional sampling without requiring a social component?

(also in progress, but almost finished)

Most of these effects rely on sampling by people



This problem can be solved using social cognition

Maybe this is <u>all</u> social reasoning?

Sampling across spatial locations







Eurasian magpie



Australian magpie

Not social cognition

Sampling across time



Not social cognition

You are currently classifying predators according to whether they pose a threat to humans. *Your team*, working at *this location recently* collected 200 observations and found that 50 (25%) of them met this criterion. This week, you have made another 4 observations, of which 3 (75%) met the above criterion. What proportion of predators in the area do you estimate pose a threat to humans?



Welsh & Navarro (2012). Organisational Behavior and Human Decision Processes

Let's make this a little more sneaky...



20 small birds with plaxium blood (SP+)

Hayes, Banner & Navarro (2017)



Category sampling: select items based on category membership (i.e. small birds)

Hayes, Banner & Navarro (2017)



Property sampling: select items based on possession of the property (i.e. plaxium blood)

Hayes, Banner & Navarro (2017)



Mean Projection Score

Lawson & Kalish (2009)
Hypotheses a reasoner might consider



Hypotheses consistent with the data



Category sampling



Frame explains absence of LP+ and LP-

Hypothesis must account for absence of SP-

Category sampling



2 of 3 hypotheses allow LP+ ... so generalisation to large birds is very plausible

Property sampling



Frame explains absence of SP- and LP-

Hypothesis must account for absence of LP+

Property sampling



No remaining hypotheses allow LP+... so generalisation to large birds is very implausible

Replication of L&K 2009



Explicit negative evidence (actual LP-) attenuates value of *implicit* negative evidence (no LP+)





A toy model





People pay attention to *mechanistic* constraints on sampling processes (not just social cues), and this shapes our reasoning in a sensible way



More extensions?

Choice: What drives people's active sampling?



with Sean Tauber and Ben Newell

Law: Evidence sampling and expertise in the courtroom



Martire, Edmonds, Navarro & Newell (2017) Martire, Growns & Navarro (under review)

Society: Trust-based sampling via selforganising social networks (fake news...)



Development: Exploratory versus goaldirected sampling by preschoolers





2 Chapter 1. Probability Models

servations that are mutually independent and identically distributed (IID), or X might be some general quantity. The set of possible values for X is the armyle space and is often denoted as X. The members P_{θ} of the parametric family will be distributions over this space X. If X is continuous or discrete, then densities or probability mass functions¹ exist. We will denote the density or mass function for P_{θ} by $f_{X|\Theta}(\cdot|\theta)$. For example, if X is a single random variable with continuous distribution, then

$$P_{\theta}(a < X \leq b) = \int_{a}^{b} f_{X|\Theta}(x|\theta) dx$$



¹Using the theory of measures (see Appendix A) we will be able to dispense with the distinction between densities and probability mass functions. They will both be special cases of a more general type of "density."



On the origins of data and the rationality* of human reasoning





People are smart. Limited, but smart.

"Common sense" reasoning is infuriatingly cunning, and requires people to learn from complex data sources (e.g., other people)



We need to disentangle facts from agendas

with Amy Perfors and Pat Shafto





We need to detect trickery

too many collaborators to list





٩Þ	∽⊳	୩	⊲þ	
Yun	Dax	Huk	???	
Whic	Yun Dax Huk ??? Which category does this belong to? Yun Dax Huk New			
Y	fun Da	ax Hu	uk New	

We need to know when to reject the rules we're given

with Charles Kemp





¶►	9Þ		qÞ-	
Yun	Dax	Huk	???	
Whic	h catego	ny does i	this belong	to?
Y	ùn D	ax. Hi	ik New	

We need to read the intention of potentially malicious agents



too many collaborators to list

Common sense reasoning requires uncommonly rich statistical models



Thanks!