Causal thinking – what did cognitive psychological research find?

York Hagmayer

University of Göttingen

Linguistic Perspectives on Causation Workshop

Hebrew University Jerusalem, 2017

In a nutshell

- Causal reasoning serves many functions and comes in many forms.
- Causal reasoning builds upon a number of different notions of cause and causation. These notions mirror philosophical conceptions.
- In cognitive psychology causal reasoning is usually studied by running experiments in which participants are given a specific task and well controlled, novel and artificial data.
- Only a few studies within psychology tried to investigate causal reasoning by looking at causal questions and the ensuing answers and arguments in everyday life. Such studies could provide interesting new insights.
- To run and analyze such studies a collaboration of linguists and psychologists is needed.

Functions of causal reasoning

Epistemic

- To explain
- To acquire knowledge

Pragmatic

- To predict
- To decide
- To regulate emotion and motivation
- To assign guilt and responsibility

(cf. Danks, 2014)

Forms of causal reasoning – Induction of causal knowledge

Inferring generic causal knowledge from observed data (and pre-existing causal knowledge)

- Induction of causal structure
  - Does A cause B?
  - Is the relation of A und B due to a common cause C?

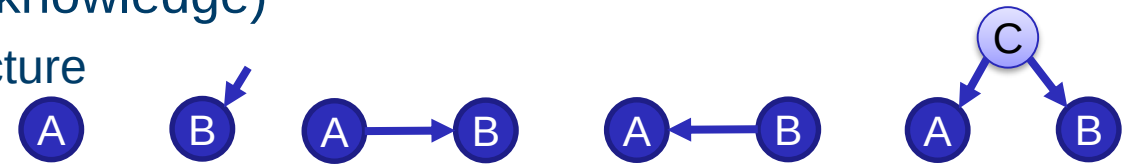

- Induction of causal strength, causal power, …
  - How strongly does A affect B?
  - How often does A generate B when A is present and no other generative causes of B are present?

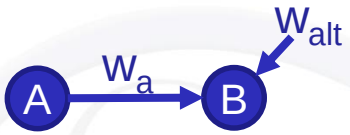

Function: Epistemic = Acquisition of causal knowledge

Dominant theories: Causal Bayes nets (e. Tenenbaum et al., 2011)
Cue-based induction of causal models (e.g. Lagnado et al.,2007)

Forms of causal reasoning – Inferences from causal knowledge

Drawing inferences from generic causal knowledge

- Predicting consequences of actual and hypothetical events
  - How likely will B result if A happens?
- Diagnosing the causes of events
  - B happens, how likely was cause A present?
- Predicting the results of interventions in a causal system
  - How likely will B result if A is generated through an intervention?
  - How likely would B result if A would be prevented from happening?

Note: These inferences concern types of events, not particular instances

Functions: Pragmatic

Dominant theories: Causal Bayes nets (e.g., Rehder & Hastie, 2003; Meder et al., 2014)
Causal model theory (e.g., Meder et al., 2008; Waldmann & Hagmayer, 2005)

Forms of causal reasoning – Inferences about singular / actual causation

Drawing inferences about the causal relation between particular instances

- Causal attribution (e.g., Cheng & Novick, 2005; Stephan & Waldmann, 2016)
  - An instance of B happened. What caused b?
    Inference 1: Likelihood of cause A being present.
    Inference 2: Likelihood that A actually caused b.
- Actual causation (e.g., Icard, 2017; Halpern & Hitchcock, 2011)
  - Instances of a and b happened. A is a generic cause of B.
    How likely did a actually cause b?
- Causal selection (e.g., Kominsky, 2015)
  - Instances a, b and c happened. A and B are known causes of C.
    Did a or b cause c?

Functions: Epistemic = explanation
Pragmatic = Assignment of blame, regulation of emotion

Forms of causal reasoning – Others

- Inferring causation from bivariate time-series
  (e.g., Rottman & Keil, 2012)
- Inferring actual causation from a single observation (perception of causality)
  (e.g., Michotte, 1946; White, 2006)
- Perceiving and judging the dynamic interaction of causes/forces on a single occasion (e.g., Wolff, 2007)
- Inferring actual causation from narratives (e.g., Pennington & Hastie, 1992)
- …

Notions of causes and causation – A psychological plurality (cf. Schulz & Gopnik, 2007; Waldmann, 2016)

Causes …

- … explain things
- … precede their effect
- … are contiguous to their effect
- … are similar to their effects
- … covary with their effect
- … make their effects more or less likely when present
- … change their effects, when intervened on
- … are necessary and/or sufficient for their effects
- … have counterfactual dependence with their effects
- … are connected to their effects through a mechanism

⇒ Depending on the task and available information people seem to rely on different notions

Notions of causes and causation – A psychological plurality

Which causes do explain? (cf. Lombrozo & Vasilyeva, 2016)

- Efficient causes: generators and inhibitors
- Material causes
- Formal causes
- Final causes: Functions, purposes
- Intentional causes: Intentions, desires

⇒ Most current theoretical models of causal reasoning focus on efficient causes (events), accounts for causal explanation also consider intentional and functional causes

Empirical research

Traditional experimental research in cognitive psychology

Predictions from theoretical model(s) of causal cognition are tested by well-controlled studies using artificial stimuli

- Theories are often derived or inspired by philosophical theories or models from computer sciences (e.g., Causal Bayes nets)
- Participants have little or no knowledge and previous experience
- Information given is controlled and restricted
- A specific task is assigned to participants
- A particular question is asked and response options are fixed

Exemplar experiment on causal attribution by Cheng & Novick (1990)

Task: Explain what caused a singular target event,

e.g. “Jane had fun washing the dishes on this occasion“

Manipulated: Information on probabilistic dependencies provided to participants in addition to target event

Problem 3 (LLH):

Adam thinks that narcissus (a flower) smells nice on this occasion.

1. In fact, Adam has always thought that narcissus smells nice.
2. But nobody else has ever thought that narcissus smells nice.
3. Everyone has always thought that all other flowers smell nice.

Finding: Participants are highly sensitive to probabilistic dependencies. They seem to compute statistical main effects and interactions of person, stimulus, and occasion to decide.

Study by Ahn et al. (1995) on causal attribution

- Task: “Write down any question you would want to have answered in order to identify the causes of the event. Write down the most likely cause.”
- Events/actions to be explained:
  Partially the same as in Cheng & Novick (1990)
  e.g. “Jane said the morning prayer on this occasion.”

Study by Ahn et al. (1995) on causal attribution

Findings

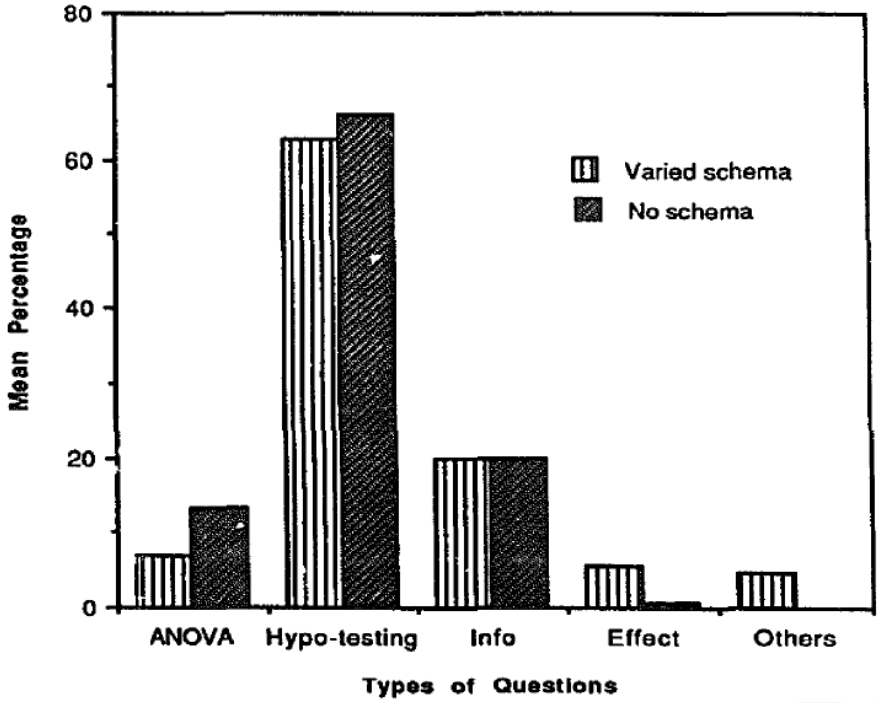

⇒ Participants rarely asked for probabilistic dependencies, even when prompted by the instruction

Traditional experimental research may miss important aspects of causal reasoning

- Tasks assigned to participants and causal questions asked might not be the ones they ask themselves in everyday life.

- Information provided may not be the information participants would search for themselves in order to complete a task.

Study 1 – Everyday causal questions

(with Marzieh Asghari & Alex Voss)

Our aim was to find out …

- … which causal questions people ask.
- … what they want to have explained, when they ask for an explanation.

Study 1: Everyday causal questions

Study 1a: Questions posted on the internet

- Analysis of all questions that were judged by at least 150 people to be a "good question" on *gutefrage.net* by April 1st 2015
- N = 1445 questions,
  N = 380 causal questions
  (most other questions inquired about some trivia facts or how to questions)

Study 1b: Causal questions arising throughout a day

- 50 students (mean age 24, 80% female)
- Task: "Text all causal questions that occur to you during one day"
- N = 612 questions,
  N = 404 causal questions
  (most other questions inquired about action choice "shall I study or eat first?", were meta-questions concerning the study "why do they conduct this study?", or expressions of frustration "why is the stoplight red again?")

Study 1 – Everyday causal questions
Typology of questions

Explanation

- Single instance: Why X on this particular occasion?
  X ∈ {event, state, action, feature, causal relation, other)
- Generic type: Why X in general?

Causation

- Single instance: Did A cause X on this occasion? (singular/actual causation)
- Generic type: Does A cause X in general?

Prognosis

- Single instance: X is present, what will result on this occasion?
- Generic type: If X is present, what will result?

Resulting utility

- Is X useful? How useful is X?

Intervention planning

- Given X, what can be done to achieve Y?

Study 1: Everyday causal questions
Examples for questions

Study 1a: Questions posted on the internet (answers.yahoo.com)

Explanation

- “Why does my 7 year old hoard litter, scraps, random things?”
- “Why are potatoes shaped like an oval?”
- “How can body mists cool you down during summer?”

Causation

- “Can hitting your nipple a lot cause cancer?”
- “Does surgery hurt?”

Intervention

- “How can I lose weight fast with barely any work?”

Study 1b: Causal questions arising throughout a day (translated)

Explanation

- “Why do I offer things I do not want to provide?”
- “Why does studying calm me down?”

Causation

- “Did too much stress cause my headache?”
- “Does my brother miss me when I think of him?”

Intervention

- „What can I do prevent future headaches?“

Study 1: Everyday causal questions
Types of Questions

Study 1a: Questions posted on the internet (gutefrage.net)

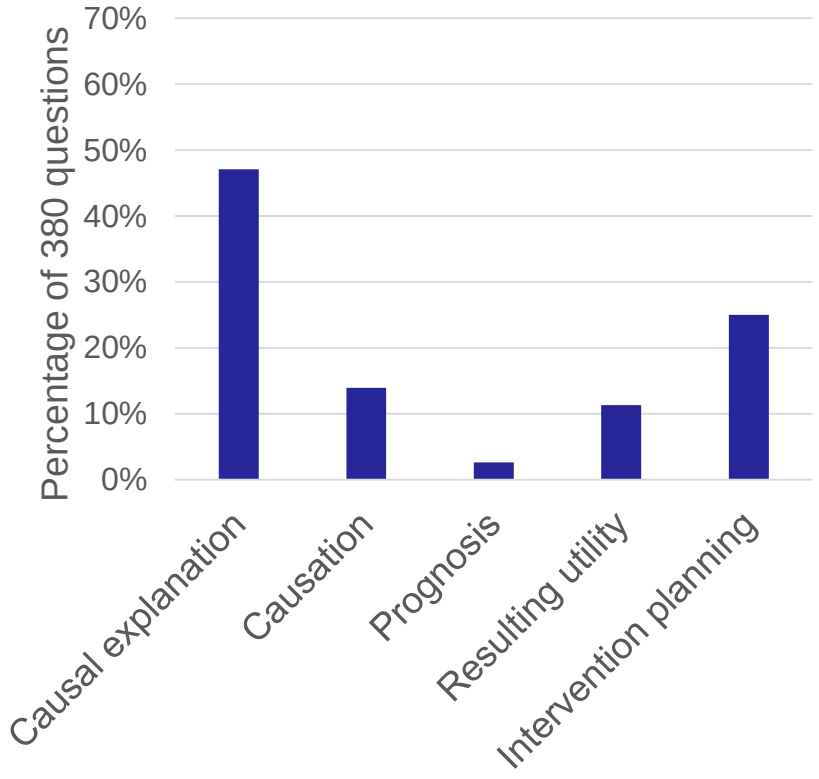

Study 1b: Causal questions arising throughout a day

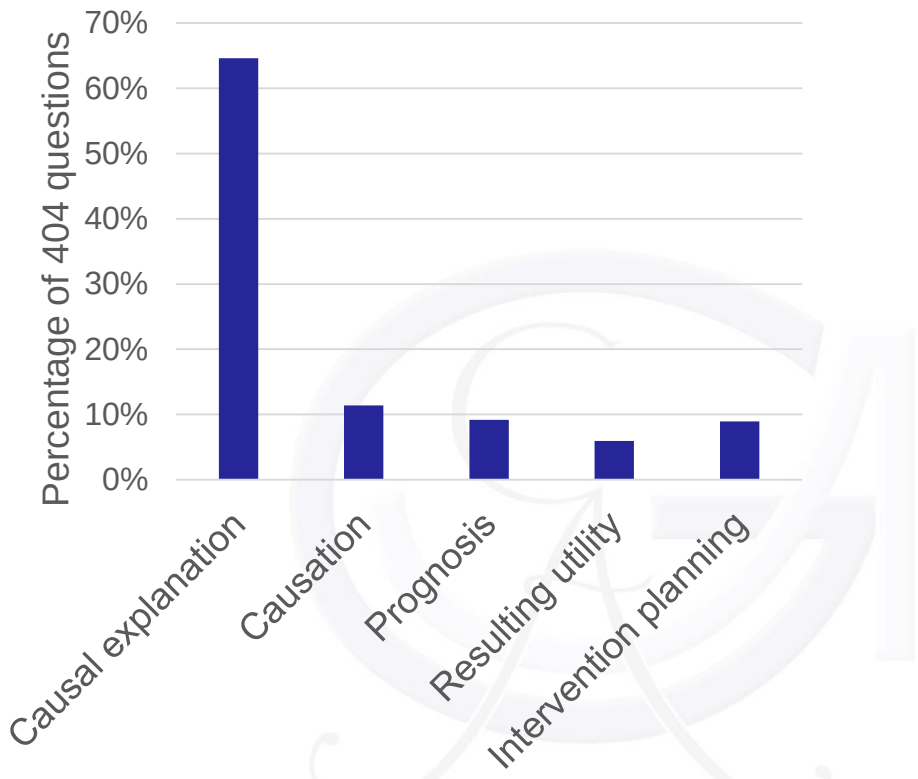

Study 1: Everyday causal questions
What do people want to have explained?

Study 1a: Questions posted on the internet

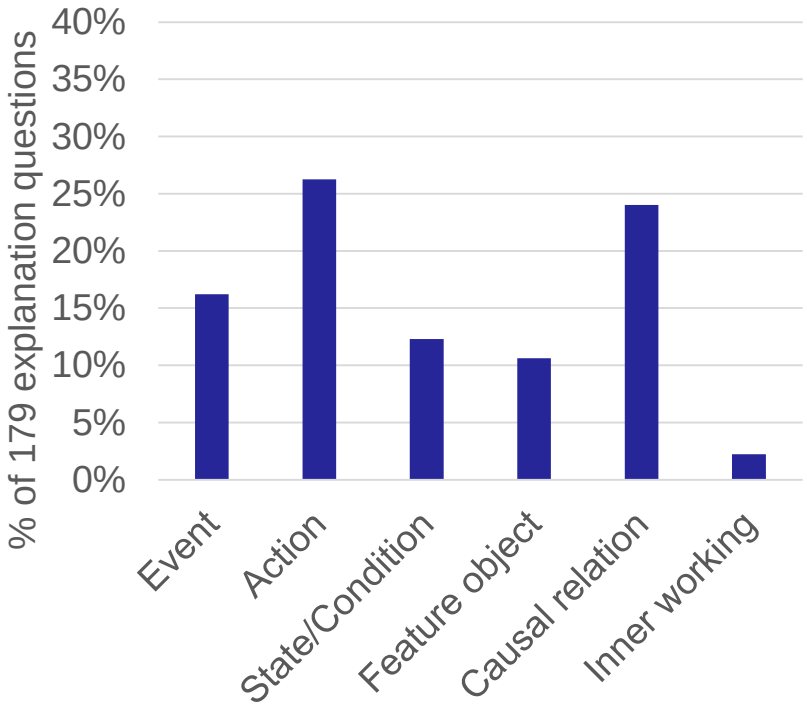

Study 1b: Causal questions arising throughout a day

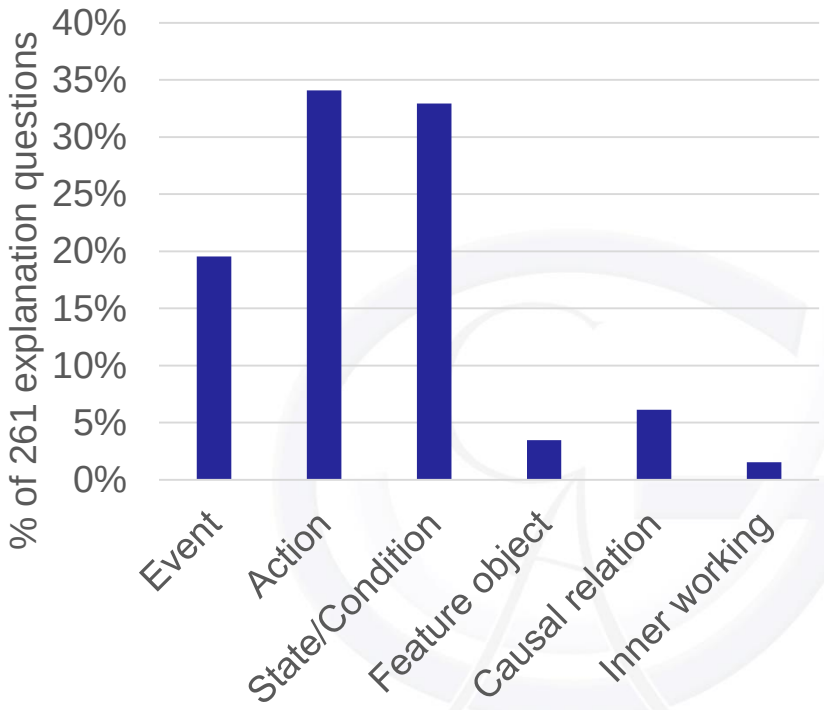

Study 1: Everyday causal questions
What do people want to have explained?

Study 1a: Questions posted on the internet

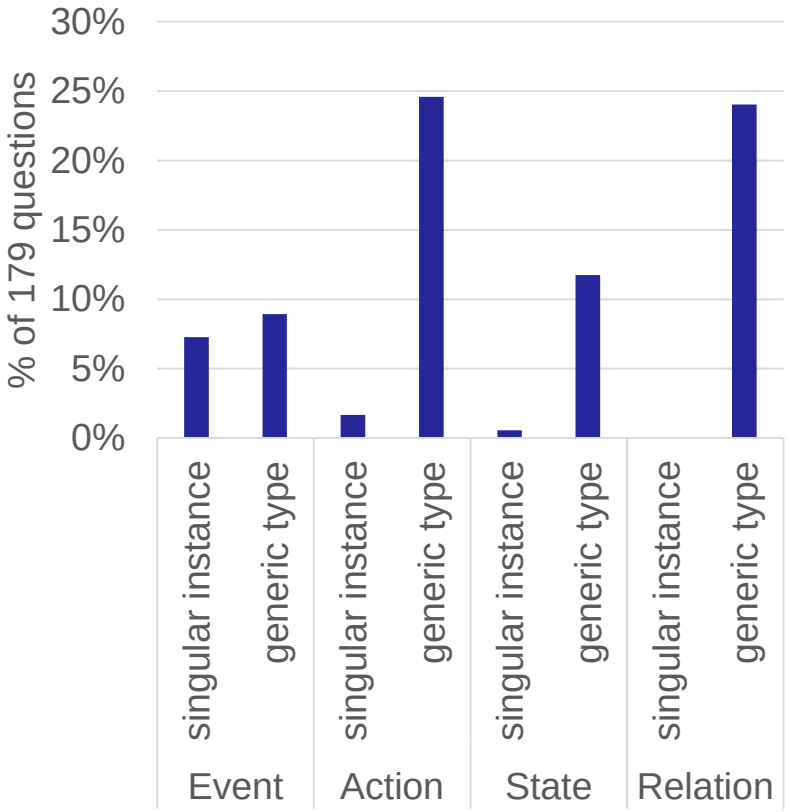

Study 1b: Causal questions arising throughout a day

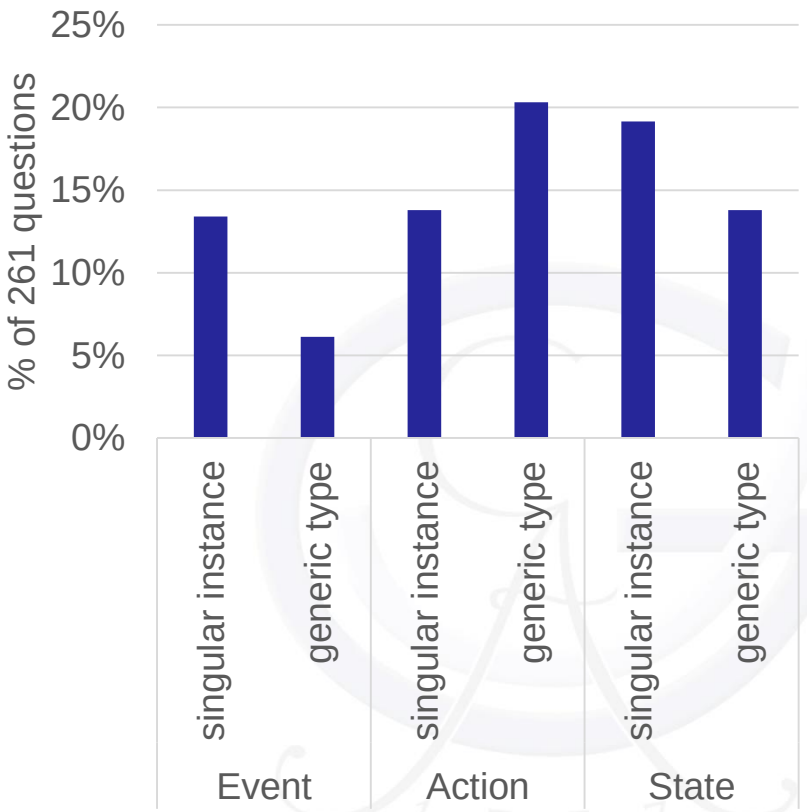

Study 1 – Summary

- Causal questions reflect the functions of causal reasoning.
- A majority of questions concern causal explanation.
- Questions for explanations not only concern single instances of events or actions (as investigated in experimental research)

  ⇒ Some forms of explanation (e.g. of causal relations) may be overlooked
- Even when a specific instance is given, reasoners ask for generic causes (≈30% in Study 1b).

  ⇒ Some questions seem to be asked to support rather than replace causal reasoning

Study 2 – Asking questions to provide an explanation
(with Neele Engelmann)

- In Study 1 only a few people asked for information that may help them to provide an explanation themselves. But cognitive psychologists are interested in the causal reasoning underlying causal explanation.

Our aims of Study 2 were …

- … to find out what reasoners want to know when they are asked to provide an explanation for familiar vs. unfamiliar events and for a singular token event vs. a type of event.
- … to investigate whether information search conforms to theoretical models of causal attribution and explanation.

Study 2 – Task and Materials

Task: Your task will be to provide an explanation. Before you do, you may ask any kind of question to obtain information that may help you.

Unfamiliar token event
„Although being healthy normally, Tina suffers from the disease Krokuritasis today.“

Familiar token event
„Although being healthy normally, Tina suffers from stomach ache today.“

Unfamiliar type of event
„Although being healthy normally, people suffer from the disease Krokuritasis“

Familiar type of event
„Although being healthy normally, people suffer from stomach ache.“

⇒ Neither potential causes nor their presence is known

⇒ Potential causes are known, but their presence is unknown

Study 2 – Predictions from theoretical models

Caveat

- No cognitive-psychological theory makes explicit predictions about information search my means of asking questions.
- But they require certain knowledge on behalf of the reasoner and/or they make claims about what a reasoner would consider in learning and inference.

⇒ Assumption: When asking questions, people intend to acquire missing knowledge or knowledge needed to derive the relevant inference

Study 2 – Predictions from theoretical models

(1) To identify potential causes, ask for …

- … contiguous factors (theories of causal perception, cue-based models of causal induction)
- … covariations, factors on which the explanandum is probabilistically dependent (attribution theories, Causal Bayes nets, cue-based models)
- … potential causes directly

(2) To establish presence or infer likelihood of potential causes, ask for …

- … structure and strength of all causal relations (Causal Bayes nets)
- … presence directly

Study 2 – Predictions from theoretical models

(3) To establish actual causation for a token instance, ask for …

- … counterfactual dependence (all counterfactual accounts)
- … presence of a mechanism connecting present cause and explanandum (mechanistic accounts)
- … sufficiency and necessity (Causal Bayes nets)

(4) To estimate causation for a type of event, ask for …

- … sufficiency (probability of), causal power (Causal Bayes nets)
- … necessity (probability of)

Study 2 – Results (N=116 participants, N=1727 questions)

Unfamiliar token event

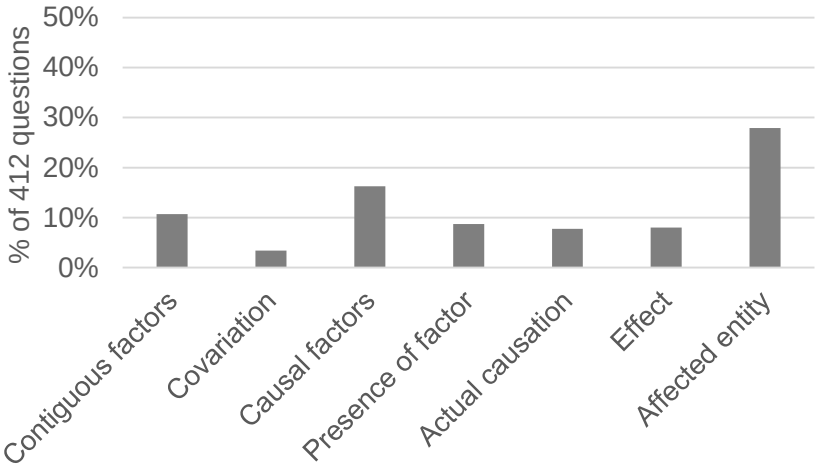

Familiar token event

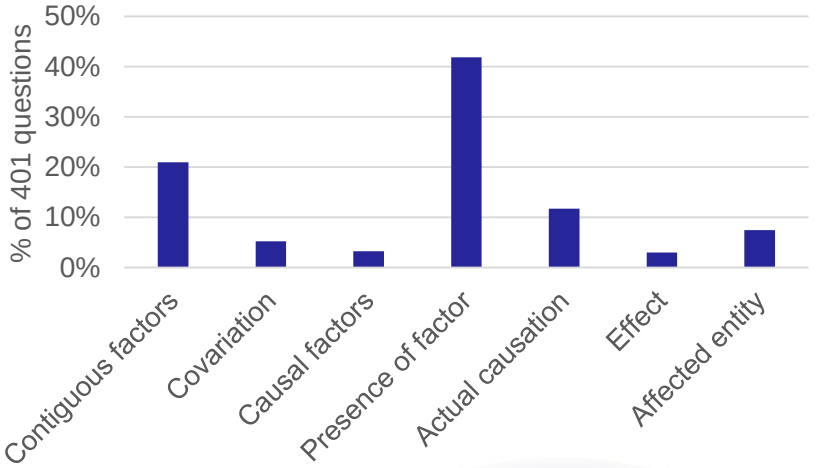

Unfamiliar type of event

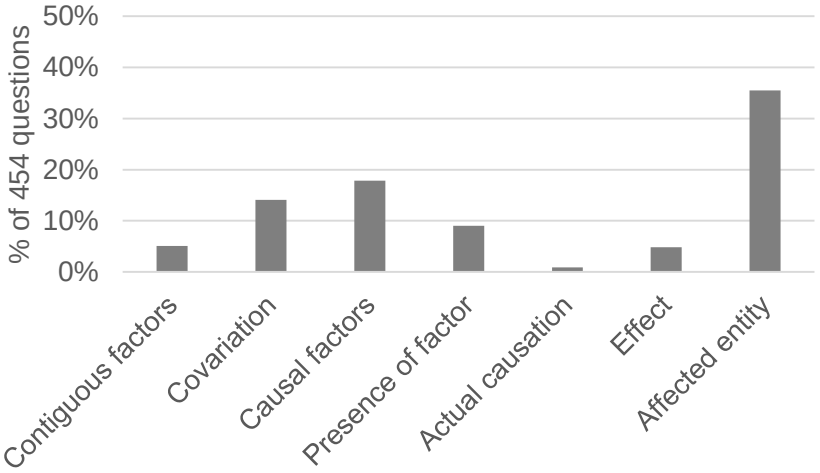

Familiar type of event

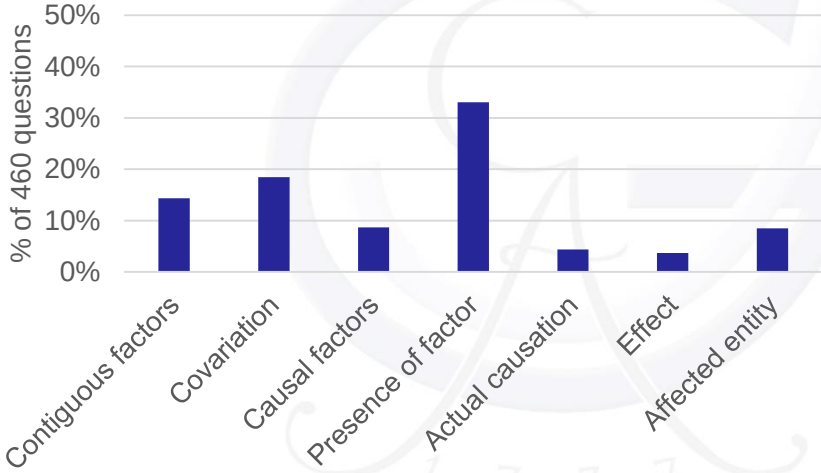

Study 2 – Summary

The questions people put forward, when asked to explain, depend on …

- … whether they have respective background causal knowledge.
  If they do not, they ask questions, which would allow for a categorization of the explanandum and the affected entity.
  If they do, they tend to ask for the presence of potential causes.
- … whether they have to explain a single token instance or a generic type of event.
  Given a single token, they ask for the presence of factors or contiguous factors. They ask for actual causation, but very often.
  Given a type of event, they ask similar questions, but more questions about covarying factors. They do not specifically ask for sufficiency, causal strength, or necessity.

A causal model hypothesis on asking questions

- People represent their causal knowledge about a domain as a qualitative causal model, which represents the relevant factors and their assumed causal relations.
- They ask causal questions,
  1. when they are uncertain about the causal model (its factors and/or relations) and have an epistemic need
  2. when they instantiate a causal model for a particular occasion to derive inferences for epistemic or pragmatic purposes and are uncertain which cause is present or which intervention is effective.
- They ask questions that allow for a categorization of the explanandum, when they do not know which causal model to apply.
- When asking a question to another person they will provide information about their causal model, what they know for certain and what they are uncertain about to allow responders to take this into account.

⇒ Hypothesis is consistent with findings from Study 2

⇒ But it has to be tested more directly

To test the causal model hypothesis on asking questions

- The causal model of people asking questions has to be known
- The occasion for which questions can be asked has to be controlled
- The questions and especially the information provided in the questions has to be analyzed and compared to the causal model
- The answers by responders and the ensuing inferences on behalf of the inquirer have to be analyzed

A first try with postings on BBC's Gardener's World and student's explanations of acne showed that …

- … there is probably a lot of conversational pragmatics involved.
- … the inferences / arguments have to be linguistically analyzed.

Conclusions

- Causal reasoning serves many functions and comes in many different forms.
- Experimental research on causal reasoning in cognitive psychology provides many important insights on how people use given information to answer specified questions. It also allows to directly test respective theories. But findings may only partially reflect how people reason causally in everyday life.
- By looking at the causal questions people ask and the ensuing answers and arguments novel insights may be generated.
- Looking a questions, answers, and arguments requires linguistic analyses, which means collaborations will be essential.

Thank you for your interest!

I’m very much looking forward to learn from you!

Study 3 – BBC Gardener's World

All questions on daffodils by June 1st 2015

Study 3 – BBC Gardener's World

Questions asking for explanations (N=13)

- Most questions (8) concern „daffodil blindness“

  „Residents in the Parish planted about 1000 daffodil bulbs, this was a few years ago, up until now they have bloomed well, this year we have noticed that almost 1/2 of them blind,could it have been the hot summer, or not so cold winter, any suggestions please.”

  Answer 1 of 8: “Are they maybe becoming congested? It might be time to split the clumps up. Or they might have been attacked by narcissus fly.”

- Some questions asking for intervention (N=12) require an explanation (N=7) because they concern daffodil blindness or weak growth

  „Daffs etc that come up blind. I have read both to bin without exception and to leave them for another chance. Which advice should I follow please? Looks this year as tho' some established and some first year bulbs are not going to flower.“

  Answer 2 of 6: ”Depends on why they are blind. If like us you are inflicted with Narcissus root fly, then dig them up and burn. (…).”

Study 3 – Causal Model for Daffodils

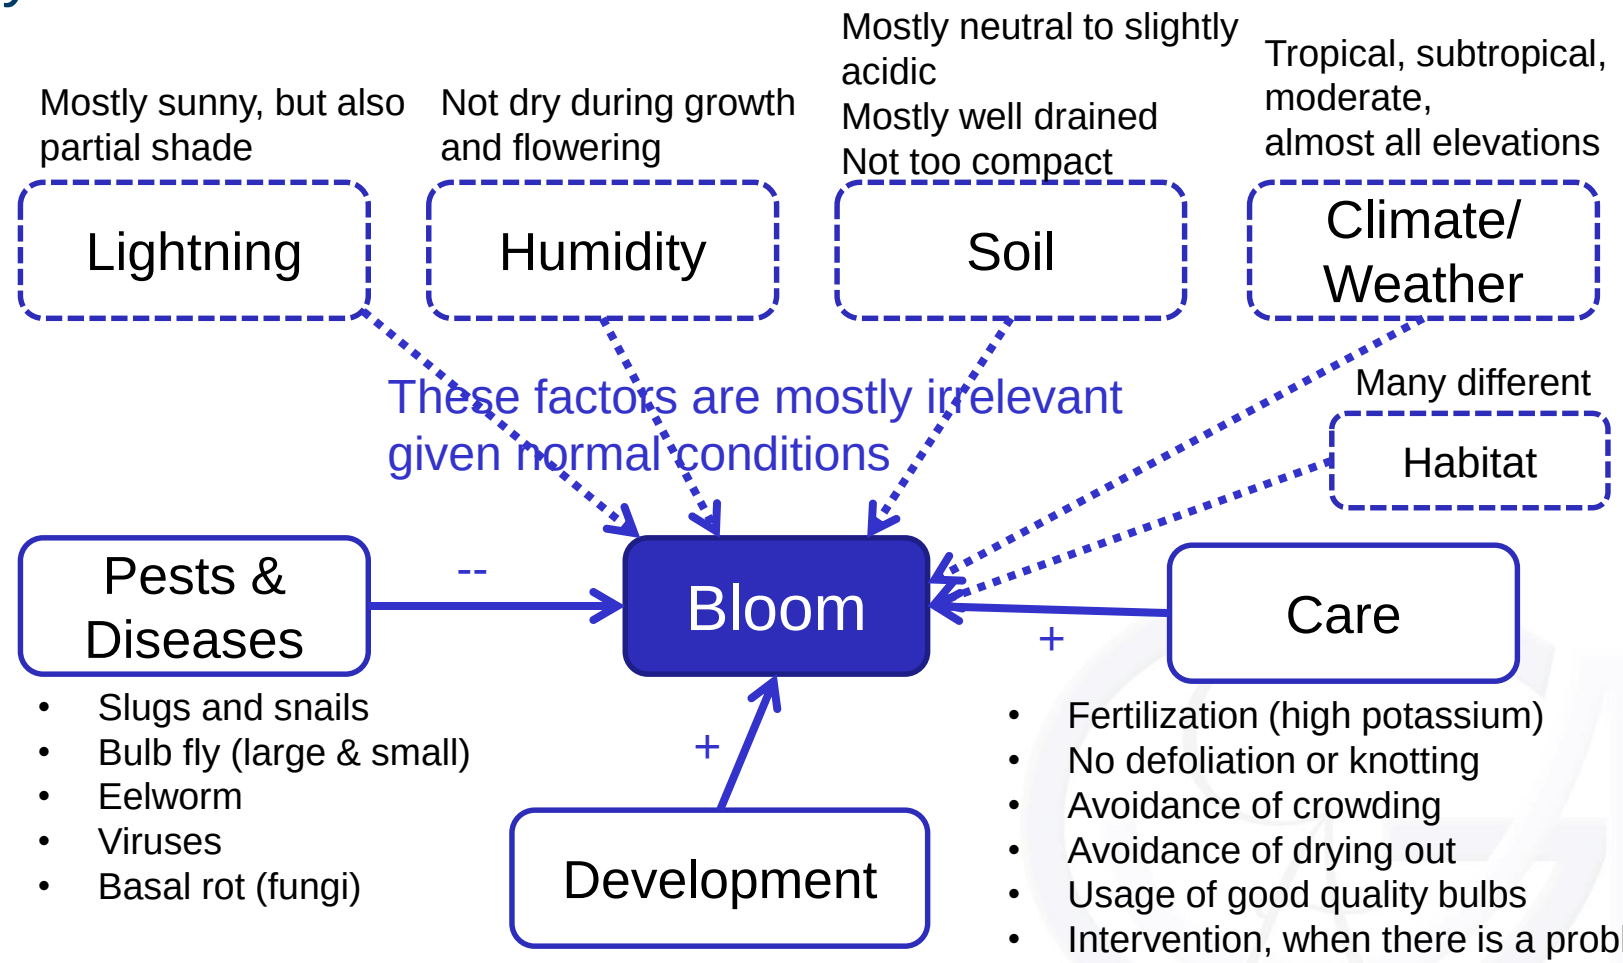

Daffodils study: Intervention

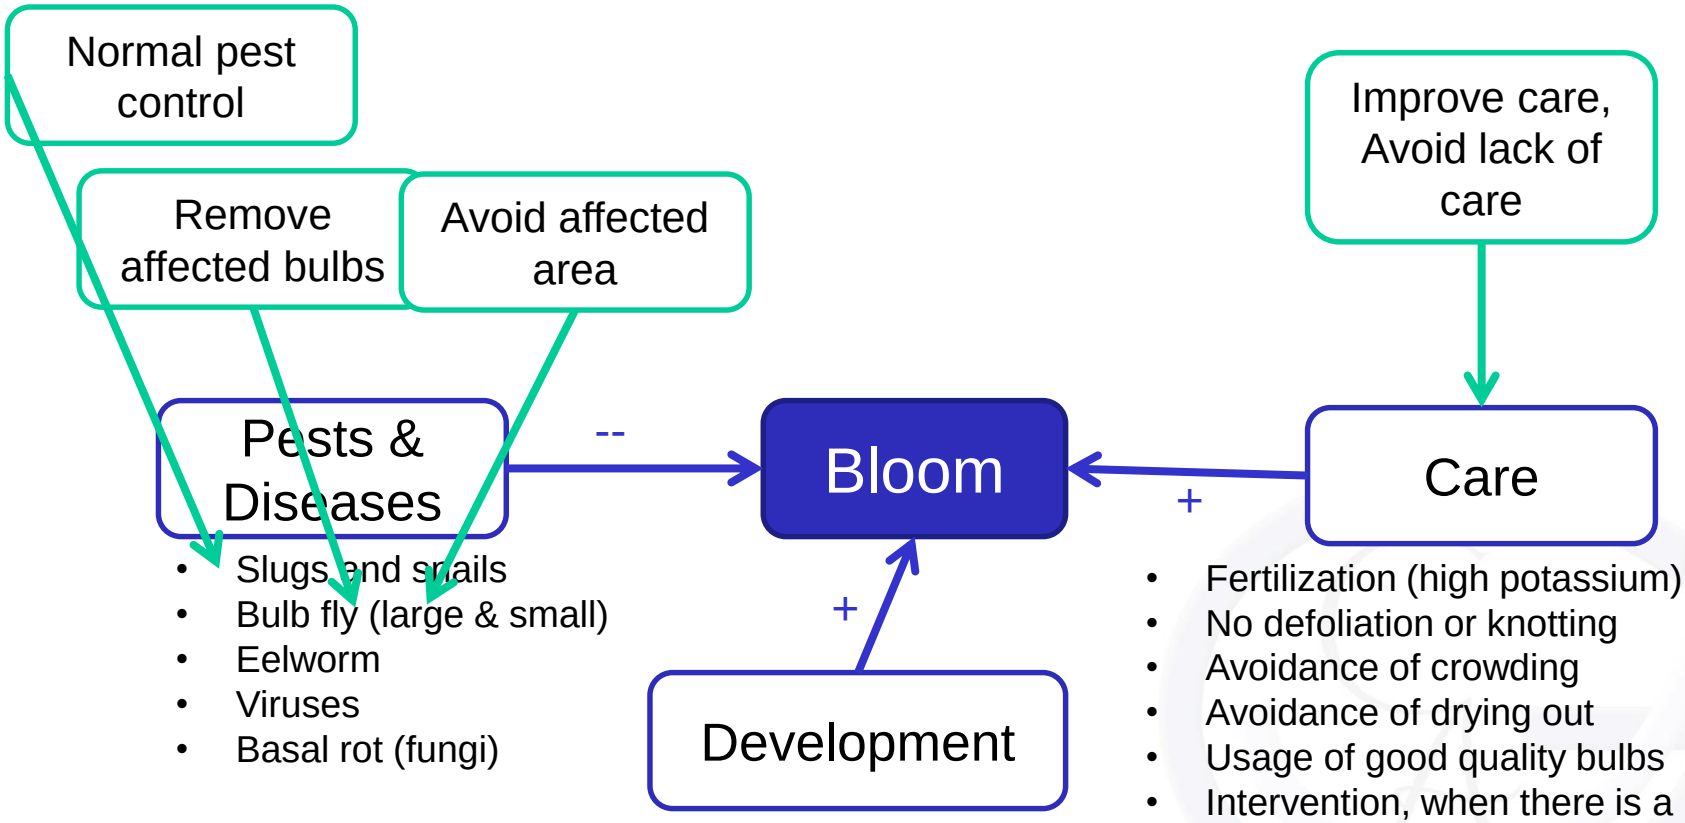

Study 3 – Findings wrt. explanation questions

- Inquirers provide a lot of information on factors that are known causes of the explanandum (3.8 on average). Many also make suggestions about the actual cause.
- Responders mention even more causes (2.4 on average) and inquire about specific aspects of these causes (e.g., type of potting soil used)
- Inquirers often just ask for an explanation. But sometimes they point out what they are uncertain about.
  - Presence of a causal factor (e.g. a pest like root fly)
  - Actual causation (e.g., late frost)
- Responders only partially address what the inquirer is uncertain about.

⇒ A linguistic analysis would be very helpful for a better understanding