Towards a Topological Theory of Knowledge, Inquiry and Correlations

Or: how to use questions to answer other questions

Alexandru Baltag, ILLC, Amsterdam

To Know is to Know the Answer

Propositional Knowledge ("knowledge that") is just a special case of knowing the answer to a question:
"knowledge wh" (whether, what, where, who, how etc.)

Schaffer: " All knowledge involves a question. To know is to know the answer".

Hintikka: interrogative model of knowledge. "Socratic epistemoloy".

Kelly: learning-theoretic (and topological) approach. "Reliable inquiry" as the foundation of epistemology.

What is an answer and when do we know it?

What does it mean to ”know the answer” to a question?

What is the relationship between this and the usual propositional ”knowledge that” operator?

Can we know the answer without knowing the question?

Can we know the answer without knowing that it is the answer to THIS question?

Knowledge Whether and Partitional Questions

Partitional Questions (Groenendijk and Stokhoff): A question is a partition $(P_1, \ldots, P_n)$ of the state space W.

$$KP_i = P_i \wedge K?(P_1, \ldots, P_n)$$

a Knows P_i means: P_i is true and a Knows WHETHER P_1 or $\ldots P_n$.

Special Case: Binary Questions $?P := (P, W = P)$.

Usually interpreted as a "contrastive" theory of knowledge, in the tradition of the "Relevant Alternatives" approach to knowledge (Lewis etc).

Relevant distinctions

A different interpretation: the "relevant distinctions" approach. Cf. and my joint paper with Rachel Boddy and Sonja Smets, in our contribution to the recent "Hintikka: In Memoriam" volume .

Answers that are "too refined" are not appropriate: an over-informative answer is NOT a correct answer!

E.g. typically, the complete answer to a question is not a full specification of the actual world.

Also, in a observational/experimental setting, the actual world might never be knowable: it is NOT a feasible answer.

Knowledge is limited by the underlying questions and by what can be observed or learned: so not every information is knowable.

Beyond Partitions and Downward Closure

Let us first stick with propositional questions:

the ones whose answers are propositions (=sets of possible worlds, usually captured by sentences in a formal language).

To formalize this in general, we need a notion of question that goes beyond partitions.

Inquisitive semantics and inquisitive logics (Ciardelli and F. Roelofsen) generalize the partitional model of questions, but in the WRONG direction (from an epistemic perspective): answers are assumed to be close downwards.

Instead, closure under finite conjunctions is the natural requirement.

Propositional Questions as Topological Bases

Given a a set W of possible worlds, a question (or an "interrogative agenda") is a family $Q \subseteq \mathcal{P}(W)$ of non-empty sets (called (feasible) answers), closed under non-empty intersections:

$$A, B \in Q, A \cap B \neq \emptyset \;\Rightarrow\; (A \cap B) \in Q.$$

In other words, a question is a special case of topological basis.

For every world w and question Q, we denote by

$Q(w) = \{A \in Q : w \in A\}$ the set of true (feasible) answers at w.

Intuitively: if A and B are true (and feasible) answers to Q (at a given world w, then the conjunction $A \cap B$ is also a true (and feasible) answer (at w).

Example: Binary Questions $?_P$

For a proposition $P \subseteq W$, the binary question $?P$ is given by:

$$?P := \{P, W \setminus P\}$$

Determined Questions

In principle, if we have an infinite set $A_1, \ldots, A_n \ \ldots, \in Q$ of true (and feasible) answers, their conjunction $\bigcap_{i \in N} A_i$ is also true. BUT... is it feasible??

In some contexts, we do want to consider this as a feasible answer, so we may require closure under arbitrary non-empty intersections (i.e. Alexandroff topology). We call such questions determined. The other (un-determined) questions are called open-ended.

The advantage is that determined questions have a unique "complete answer" in any world w, namely

$$Q_w := \bigcap \{A \in Q : w \in A\}.$$

This is the strongest answer in Q that is true at w.

The complete answer

Given a proposition $P \subseteq W$ and a determined question Q, the proposition "P is the complete answer to Q" (at the actual world) is denoted by $Q : P$.

Formally:

$$Q : P := \{w \in W : Q_w = P\}$$

But in general, in an infinite set of worlds, the complete answer might not be a feasible (i.e. knowable or observable) answer.

Why Not Closure under Complementation?

Why not assume that Q is closed under complements?

The negation of a feasible (knowable) answer might not be feasible.

Certain properties are observable by experiments only if they are true.
If they are false, their falsehood is not observable.

This can happen even in a finite state space.

Partitional Questions

Partitional questions are a special case of determined questions, namely the one in which every two distinct answers $A, B \in Q$ (with $A \neq B$) are disjoint ($A \cap B = \emptyset$).

Partitional questions are in a sense, not only determined, but "fully determined", since there is only one true answer.

Knowing an answer A to a partitional question Q implies knowing that A is the answer to Q.

But this is NOT true in general (for non-partitional questions).

Examples

EXAMPLE 1: "What is the height of that tree?"

(a) Idealized, conceptual interpretation: full accuracy.

$$Q = \{\{x\} : 0 < x < \infty\}, \quad \text{so } Q_x = \{x\}.$$

This is a fully determined (partitional) question.

(b) Experimental interpretation: "'measure the tree's height" with some possible error.

$$Q' = \{(x-\epsilon, x+\delta) : 0 < x, \epsilon, \delta < \infty\}, \quad \text{so } Q'(x) = \{(x-\epsilon, x+\delta) : 0 < \epsilon, \delta < \infty\}$$

This is an open-ended question: the conjunction of all the answers (=the singleton $\{x\}$) is NOT a feasible answer.

ANOTHER EXAMPLE: observable properties

In an empirical setting, we can think of each observable property as a possible answer to the question: "What is the evidence?" (or rather "what evidence is out there?").

Each of the answers is feasible, since we can observe it (if true). The same goes for any finitely many of them (hence closure under finite intersections).

$W = N = \{0, 1, 2, \ldots\}$, $Q = \{P_0, P_1, P_2, \ldots\}$, where $P_x = \{n \in N : n \geq x\}$.

At any world $x \neq 0$, the complete answer is P_x, but this is consistent with all worlds $y < x$, for which the complete answer (P_y) is different.

So, even if we observe the complete answer, we DON'T KNOW that it is THE COMPLETE answer!

ANOTHER EXAMPLE: The Epistemic Question

Let $?a$ be the question "What does a know?".

One can argue that this IS a question in our sense (assuming that knowledge is closed under finite conjunctions).

Treating $?a$ as a question in our sense amounts to adopting the topological semantics for K_a (as interior in a certain topology).

One might even argue that this is a determined question: hyper-rational agents may know the conjunction of all they know.

Treating $?a$ as a determined question amounts to adopting the relational $S4$ semantics for K_a (as Kripke modality for a certain preorder); i.e. an 'Alexandroff topology'.

Not always partitional!

But assuming that $?a$ as a "fully determined" (i.e. partitional) questions amounts to assuming the relational $S5$ semantics (based on equivalence relations).

This only fits "infallible" knowledge (absolutely certain, fully introspective).

An agent who is not negatively introspective may believe (that he knows) things that he doesn't actually know.

Such an agent a knows the (complete) answer to $?a$ (since he knows all the propositions that he knows!), but he will typically not know that this is the (complete) answer to $?a$ (since he thinks he knows many other things).

The partitional refinement of Q

For every question Q, there is an associated partitional question $?Q$, called the partitional refinement of Q. Its equivalence relation is

$$w \sim_Q w' \text{ iff } Q(w) = Q(w').$$

The unique (complete) answer to $?Q$ at w is given by:

$$(?Q)_w := \{w' : w \sim_Q w'\}$$

The question $?Q$ is given by the associated partition:

$$?Q = \{(?Q)_w : w \in W\}$$

The partitional refinement $?Q$ IS a refinement of Q: every answer to Q is a union of answers to $?Q$.

$$A \in Q \text{ implies } A = \bigcup_{w \in A} (?Q)_w$$

For instance, $??a$ is the question: "what does a know and what doesn't she know?". The complete answer at w is the set

$$(??a)_w := \{w' \in W :?_a(w) =?_a(w')\}$$

of all the worlds in which a knows the same things that she knows at w.

Interrogative Order

Every question Q induces a preorder relation $\leq_Q$ (reflexive and transitive) on W, called the interrogative order, and given by

$$w \leq_Q s \text{ iff } Q(w) \subseteq Q(s),$$

i.e. all answers true at w are also true at s.

This interrogative order is directly related to the partitional refinement of Q:

$$w \sim_Q s \text{ iff } w \leq_Q s \leq_Q w.$$

For determined questions, the following are equivalent:

1. $w \leq_Q s$
2. $s \in Q_w$
3. $Q_s \subseteq Q_w$

Cumulating Questions

Given questions Q, Q', we can ask both questions simultaneously: this is the cumulated question $Q \sqcap Q'$, given by

$$Q \sqcap Q' = \{A \cap A' : A \in Q, A' \in Q', A \cap A' \neq \emptyset\}$$

This means that

$$(Q \sqcap Q')(w) = \{A \cap A' : A \in Q, A' \in Q, w \in A \cap A'\}$$

For determined questions:

$$(Q \sqcap Q')_w = Q_w \cap Q'_w.$$

In principle, we can cumulate an infinite family of questions $Q_1 \sqcap Q_2 \sqcap \ldots Q_n \sqcap \ldots$.

Question Dependency ("Solvability")

Knowing the answer to a question Q might be enough to know the answer to question Q', in which case we write

$$K_Q Q'$$

and say that Q solves Q'.

For DETERMINED questions, this means: the complete answer to Q entails the complete answer to Q'.

$$w \in K_Q Q' \text{ iff } Q_w \subseteq Q'_w \text{ iff } \forall s (w \leq_Q s \to w \leq_{Q'} s).$$

In particular, for FULLY determined (partitional) questions, we have:

$$w \in K_Q Q' \text{ iff } \forall s (w \sim_Q s \to w \sim_{Q'} s),$$

i.e. all worlds that agree with w on (the answer to) Q also agree on (the answer to) Q'.

But for OPEN-ENDED ones, we need to be more careful with our definition:

$K_Q Q'$ every answer to Q' is entailed by some answer to Q.

$$w \in K_Q Q' \text{ iff } \forall A' \in Q'(w) \exists A \in Q_w : A \subseteq A'.$$

This is somewhat related to "interrogative implication" in Inquisitive Semantics, and to dependency formulas in Dependence Logic.

Note that we always have

$$K_{Q \sqcap Q'} Q,$$

$$K_{Q \sqcap Q'} Q',$$

$$K_{?Q} Q.$$

Knowing the Answer

An agent a knows (all) the answer(s) to Q iff $?a$ solves Q:

$$K_a Q := K_{?a} Q$$

An agent a knows the answer to Q given the answer to Q' if $?a \sqcap Q'$ solves Q:

$$K_a^{Q'} Q := K_{?a \sqcap Q'} Q.$$

Knowledge of a proposition $P \subseteq W$ ("knowledge that") is a special case:

$$K_a P := P \cap K_a ?P$$

Examples

Review our examples.

Is the height of the tree knowable?

In our third example (of observable properties), can the real world (number x) be known given the observable evidence?

Let Ockham be a special agent who “knows” all upp-closed sets (i.e. considers higher numbers more plausible than lower numbers).

Can Ockham know the real world knowable given the observable evidence?

Question Equivalence and the Complete Answer

Two questions are equivalent if knowing one's answer is equivalent to knowing the other's answer:

$$Q \simeq Q' \;:=\; K_Q Q' \cap K_{Q'} Q$$

For determined questions, we can express the fact that $P \subseteq W$ is the COMPLETE ANSWER to Q (at w), by saying that P is true and Q is equivalent to $?P$ (at w):

$$Q_w = P \;\text{ iff }\; w \in P \text{ and } w \in (Q \simeq ?P)$$

i.e.

$$Q : P = P \cap (Q \simeq ?P)$$

Knowing the Answer Without Knowing it's the Answer

An agent can know the answer to Q without knowing that it's the answer to Q.

Indeed, although

$$(Q : P \wedge K_a Q) \rightarrow K_a P$$

is always the case, we also have that in general

$$(Q : P \wedge K_a Q) \not\rightarrow K_a(Q : P)$$

EXAMPLE: A non-negatively-introspective agent a knows the answer to $?_a$ without knowing it is the (complete) answer.

The Logic of Epistemic Questions: SYNTAX

LEQ has a twofold syntax, consisting of a set $\mathcal{L}$ of propositional formulas φ and a set $\mathcal{Q}$ of question expressions Q, defined by double recursion:

$$\varphi ::= \quad p \quad | \quad \top \quad | \quad \neg\varphi \quad | \quad \varphi \wedge \varphi \quad | \quad K_Q Q$$

$$Q ::= \quad ?x \quad | \quad ?a \quad | \quad ?\varphi \quad | \quad Q \sqcap Q \quad | \quad ?Q$$

where $a \in \mathcal{A}$ are agents and $?x \in Var$ are variables denoting some primitive questions.

The set of basic questions consists of all epistemic questions $?a$ and all primitive questions $?x$ (for all a and all x).

SEMANTICS

Models $M = (W, \| \bullet \|)$, where:

- W is a set of possible worlds;
- $\| \bullet \|$ is a valuation function, mapping
  - atomic sentences $p \in P$ into propositions over W ($\|p\| \subseteq W$),
  - basic question variables $?x \in Var$ and knowledge questions $?a$ into questions over W.

SEMANTICS– continued

We extend the valuation to an interpretation map $\| \bullet \|_M$, mapping

- formulas $\varphi \mapsto \|\varphi\|_M \subseteq W$ into propositions, and
- question expressions Q into questions $\|Q\|$ over W:

$$\|\top\| = W$$

$$\|\neg\varphi \to \psi\| = W \setminus \|\varphi\|$$

$$\|\varphi \wedge \psi\| = \|\varphi\| \cap \|\psi\|$$

$$\|?\varphi\| = ?\|\varphi\|$$

$$\|Q \sqcap Q'\| = \|Q\| \sqcap \|Q'\|$$

$$\|?Q\| = ?\|Q\|$$

Abbreviations

$$K_Q\varphi \;=\; \varphi \wedge K_Q?\varphi \quad \text{(the issue modality } Q\varphi \text{ of van Benthem and Minica)}$$

$$K\varphi \;=\; K_{?\top}\varphi \quad \text{(global modality)}$$

$$K_a Q \;=\; K_{?a} Q \quad \text{(knowledge of answer)}$$

$$K_a^{Q'} Q \;=\; K_{?a \sqcap Q'} Q \quad \text{(conditional knowledge of answer)}$$

$$K_a\varphi \;=\; K_{?a}\varphi \quad \text{(knowledge that)}$$

$$K_a^{Q}\varphi \;=\; K_{?a \sqcap Q}\varphi \quad \text{(conditional knowledge that)}$$

$$D_G\varphi \;=\; K_{?a_1 \sqcap ?a_2 \sqcap ?a_n}\varphi, \quad \text{for } G = \{a_1, \ldots, a_n\} \quad \text{(distributed knowledge)}$$

$$K_a b \;=\; K_{?a}?b \quad (a \text{ knows all that } b \text{ knows})$$

$$Q \simeq Q' \;:=\; K_Q Q' \wedge K_{Q'} Q \quad \text{(question equivalence)}$$

$$Q : \varphi \;:= \varphi \wedge (Q \simeq ?\varphi) \quad (\varphi \text{ is the complete answer to } Q)$$

$$K_a{}^{\text{only}}\,\varphi \;:= ?a : \varphi \quad \text{(``only knowing'' } \varphi)$$

Axiomatizations

A model is determined if the interpretations of all basic questions $?a$ and $?x$ are determined questions.

A model is fully determined (partitional) if these interpretations are partitional questions.

I have complete axiomatizations over the class of determined models and the class of partitional models (but NOT yet for open-ended questions!).

Examples of Axioms and Rules

The $S4$ axioms and rules for $K_Q\varphi$;

Determinacy Rule:

From $Q : p \to \varphi$, infer φ, provided that p does not occur in Q or φ.

(This says that every determined question has a complete answer.)

Cumulation Axiom:

$$(Q : \varphi \wedge Q' : \psi) \;\to\; (Q \sqcap Q') : (\varphi \wedge \psi)$$

Binary Questions:

$$\neg\varphi \to ?\varphi : \neg\varphi$$

Partitional Refinement:

$$Q : \varphi \to ?Q : (Q : \varphi), \text{ and } S5 \text{ axioms for } K_{?Q}$$

One-Step Dynamics: Learning the Answer

The complete answer to Q is publicly announced: $!Q$

It only makes sense for DETERMINED questions.

The model M changes to $M^{!Q}$, where the new epistemic questions are

$$\|?a\|^{!Q} := \|?a\| \cup \|Q\|, i.e. :$$

$$\|?a^{!Q}\|(w) := \|?a\|(w) \cup \|Q\|(w)$$

and so agent a's new knowledge is

$$(?a)_w^{!Q} := (?a)_W \cap Q_w.$$

SEMANTICS: $\|\langle !Q \rangle \varphi\|_M = \|\varphi\|_{M^{!Q}}$

" φ is achieved after learning the complete answer to Q "

Gradual Dynamics: Learning Answers

For OPEN-ENDED questions, all we can do is to gradually learn answers.

So M^Q is now a "process", i.e. a family of models

$$M^Q = \{M^A : A \in Q\},$$

where each M^A has A as set of worlds, and all the components (valuation, interpretation of questions) are restricted to A.

$$\text{SEMANTICS:} \quad \|\langle Q\rangle\varphi\|_M = \bigcup_{A\in Q} \|\varphi\|_{M^A}.$$

" φ may be achieved by learning some (enough) answers to Q ".

NOTE: $\langle !Q\rangle$ is NOT a special case of $\langle Q\rangle$. We rather have:

$\langle !Q\rangle\varphi = \langle Q\rangle[Q]\varphi$ ("eventually φ, after learning enough answers"), (where $[Q]\varphi := \neg\langle Q\rangle\neg\varphi$).

Conditional Knowledge Pre-encodes Knowledge

For basic questions $Q \in \{?x : x \in Var\} \cup \{?b : b \in Agents\}$, we have the validity

$$K_a^Q Q' \to \langle Q \rangle K_a Q'.$$

The general version is a bit more complicated, because of Moore sentences.

GENERALIZATION: Learning Different Questions

I will do it only for determined questions:

it is publicly announced that a is privately told the (complete) answer to some question Q_a, b the answer to some question Q_b etc.

This action $!(a : Q_a, b : Q_b, \ldots)$ changes any model M to $M^{!(a:Q_a,\ldots)}$, where the new epistemic questions are

$$\|?a\|^{!(a:Q_a,b:Q_b,\ldots)} := \|?a\| \cup \|Q_a\|$$

$$\|?b\|^{!(a:Q_a,b:Q_b,\ldots)} := \|?a\| \cup \|Q_b\|$$

etc, so that e.g.

$$(?a)_w^{!(a:Q_a,b:Q_b,\ldots)} = (?a)_w \cap (Q_a)_w$$

EXAMPLE: Cheryl's birthday!

FURTHER GENERALIZATION: Question Refinement

Learning is just a special case of question refinement

$$!(Q_1 := Q_1 \sqcap Q'_1, Q_2 := Q_2 \sqcap Q'_2, \ldots),$$

by which certain basic questions
$Q_1, Q_2, \ldots \in \{?x : x \in Var\} \cup \{?b : b \in Agents\}$ are separately but simultaneously refined with (possibly more complex) question expressions $Q'_1, Q'_2, \ldots$, such that in the new model

$$\|Q_1\|^{!(Q_1:=Q_1\sqcap Q'_1, Q_2:=Q_2\sqcap Q'_2, \ldots)} := \|Q_1\| \cup \|Q'_1\|$$

$$\|Q_2\|^{!(Q_1:=Q_1\sqcap Q'_1, Q_2:=Q_2\sqcap Q'_2, \ldots)} := \|Q_2\| \cup \|Q'_2\|$$

etc.

FURTHER GENERALIZATION: Question Change

We can further generalize this to question change (substitution, by dropping the cumulative nature of the question expressions on the right:

$$!(Q_1 := Q'_1, Q_2 := Q'_2, \ldots).$$

Applying this to epistemic questions, we can obtain non-monotonic forms of knowledge change:

$$!(?a := Q_a, ?_b := Q_b, \ldots).$$

This can represent various forms of forgetting ("agent a forgets everything except the answer to Q_a "), or mixtures of learning and forgetting.

FURTHER GENERALIZATION: Event Models

This can be used to represent more private forms of learning, forgetting, question change, raising of new issues etc.

As well as learning the answer to a question (from a given list), without getting to know the question!

The End? Far from it!

We saw that one can know the answer to a question without knowing that is THE (complete) answer to THAT question.

But can we know an answer without knowing that it is AN answer?

For the special case of FULLY DETERMINED questions, the answer is unique, so knowing an answer is the same in this case as knowing the answer.

Do exist fully determined questions, for which we can know the answer without knowing that is (a, the) answers to THIS question?

Well, this is not possible in our current models!

But... maybe the models are wrong!

The answer is 42. But what is the question?

According to The Hitchhiker's Guide to the Galaxy, the "Answer to the Ultimate Question of Life, the Universe, and Everything" is 42.

This was calculated by an enormous supercomputer named Deep Thought over a period of 7.5 million years.

Unfortunately, no one really knows what the question is.

Deep Thought claims that just asking " the Ultimate Question of Life, the Universe, and Everything" is too vague!

'If you'd care to make this question precise', he says, 'you'd see that the answer is obviously 42 '.

Earth has been constructed in order to compute the precise Question. Earth's biosphere is a super-super-computer designed to search for the question, over billions of years.

Possible clarifications

Here are some possible ways the Ultimate Question more precise, encoded as a question about the value of a variable:

$x :=$ the number of illustrations in Alice's Adventures in Wonderland;
$y :=$ my office number at ILLC (without floor number or building code);
$z :=$ Yde Venema's office number at ILLC;
$w :=$ Johan van Benthem's office number at ILLC;
$v :=$ the perfect score in the USA Mathematical Olympiad (USAMO);

After hearing Deep Thought's answer, you know the value to (at least) one of these variables:

$$K\{x, y, z, w, v\}$$

But which one?

This, my friends, is of course another question!

The Earth's Biosphere is still working on it!

But to spare you more billions of years of waiting, I hereby give you:

THE FULL ANSWER:

all of the above EXCEPT z.

(Yde's office number is not 42.)

So the Question about Life, the Universe and Everything is NOT about Sonja's office!

But all the others are serious contenders for the Ultimate Question.

Knowing What

Answers are not necessarily propositions, propositional knowledge is not enough.

$$
\begin{aligned}
\text{Questions} &= \text{Variables} \\
\text{Answers} &= \text{Values}
\end{aligned}
$$

Plaza, Wang and Fan, van Eijck and Gattinger: Knowing What - knowing the value of a variable. Kx: the agent knows the (current) value of x (in the actual world).

Knowledge-that is a special case of knowing-what: use of special Boolean variables $?_\varphi$, encoding the truth value of φ. Then

$$K\varphi = \varphi \wedge K?_\varphi.$$

→ To paraphrase Quine:

"To Know is to Know the Value of a Variable"

Feasible Answers

In an empirical context, the exact value may not be knowable. We need to model the feasible answers.

A wh- question is a map $x : W \to X$ from the set of possible worlds to some topological space (X, τ_x) (with the designated basis).

X =the set of possible exact answers/values

τ_x=the family of possible feasible answers (=observable approximations of the exact value)

The (basic) open neighborhoods of $x(w)$ are the (correct) feasible answers at w.

$\tau_x(x(w))$ =the family of feasible answers at w.

Propositionalization

Any wh- question $x : W \to (X, \tau_x)$ induces a propositional question (=topological basis) $\tilde{x}$ on W, given by

$$\tilde{x} = \{x^{-1}(U) : U \in \tau_x\}$$

$\tilde{x}$ is called the propositionalization of question x. Its answers are propositions of the form "O is an answer to question x ".

Cumulating Questions

Given questions $x : W \to (X, \tau_x)$, $y : W \to (Y, \tau_y)$, we can ask both questions simultaneously:
this is the cumulated question $x \sqcap y$, given by

Set of exact values: $X \times Y$

Exact answer at w: $x(w), y(w))$

Feasible answers: the product topology.

(Product basis: finite intersections of $A \times B$, where A, B are in the corresponding designated bases, for x and respectively for y.)

Learning the value of a variable

Suppose we learn (an) approximate value of a variable x. What will be known after that?

To predict this dynamics we need to pre-encode it using conditional forms:

namely, conditional knowledge (of x given y, or given $y_1, \ldots, y_n$)

$$K_{y_1,\ldots,y_n} x$$

and as a special case conditional knowledge of a proposition given a variable

$$K_x \varphi.$$

Knowing Correlations

"To Know is to Know a Dependence between Variables"

Cf. an "Epistemic" version of Dependence Logic and/or erotetic-inquisitive logics.

Dynamics: Knowledge acquisition = learning correlations (with the goal of eventually tracking causal relationships in the actual world).

Further Generalization: Knowledge-as-Correlation

A general logic of informational correlations, which can express the fact that a situation (tuple of values of certain variables) carries information about another situation.

This logic subsumes distributed knowledge, "knowing at least as much" as another agent, "knowing only" phi, interrogative modalities etc.

Knowledge is just a special case: the agent's epistemic state carries information about some other situation.

"Knowing IS a Dependence between Variables"

Cf. Situation Theory (Barwise), the information-correlations throry of knowledge (Dretske), "information as range" (cf van Benthem and Martinez).

Question Dependency ("Solvability")

Knowing how to answer a question x might be enough to know how to answer question y, in which case we write

$$K_x y$$

and say that x solves y.

Formally: $K_x y$ holds iff every correct feasible answer to y is entailed by some correct feasible answer to x.

$$w \in K_x y \text{ iff } \forall O \in \tau_y(y(w)) \exists U \in \tau_x(x(w)) : x^{-1}(U) \subseteq y^{-1}(O).$$

For $K_{x_1,\ldots,x_n} y$ we just use the definition for $K_{x_1 \sqcap \ldots x_n} y$ (using the cumulative question with the product topology).

This is a "continuous" version of the semantics of Dependence Logic.

Knowing the Answer

An agent a knows (all) the answer(s) to x iff $?a$ solves x:

$$K_a x := K_{?a} y$$

An agent a potentially knows y given the answer to x if $?a \sqcap x$ solves y:

$$K_a^x y := K_{?a \sqcap x} y.$$

Knowledge of a proposition $P \subseteq W$ ("knowledge that") is a special case:

$$K_a P := P \cap K_a ?P$$

More generally

$$K_{x_1,\ldots,x_n} P := P \cap K_{x_1,\ldots,x_n} ?P$$

Topologically

$$K_x P = Int_{\tilde{x}} P.$$

Locality

Knowing a correlation $K_{x_1,\ldots,x_n}y$ is a "local" version of (the topological variant of) the Dependence Logic statement $=(y, x_1, \ldots, x_n)$.

Its "'locality" is due to the fact it is dependent on the actual world.

Epistemically, this is expressed by its non-introspective nature:

$$K_x y \not\to K_x K_x y$$

If given the "right" value of x (in the actual world), the agent would be able to compute the "right" value of y. But this might not be true for other values of x (in other worlds), and so she might not know.

Introspective version

The introspective version of knowledge of a correlation is

$$K_x K_x y$$

or more generally

$$K_{x_1,\ldots,x_n} K_{x_1,\ldots,x_n} y.$$

Topological Characterization: $w \in K_{\overrightarrow{x}} K_{\overrightarrow{x}} y$ holds iff there is some neighborhood in the product topology $U \in \tau_{\overrightarrow{x}}(\overrightarrow{x}(w))$ and some continuous function $f : U \to y[W]$ s.t. $f \circ \overrightarrow{x} = y$ holds on $dom(y)$.

Continuity of the dependency = knowability of the correlation.

An Example

Newton is asked: "Can you tell the position of the Moon in year 3000?"

$y : W \to R$, y(w)= the position of the Moon in 3000 in world w, with the standard topology on R.

Answer: "I don't know".

"But if I give you the current positions, velocities and masses of all bodies in the solar system, can you answer the previous question?"

$x_1, \ldots, x_n$: variables encoding all the initial conditions.

Answer: "Yes, I can".

$$\neg K_n y \wedge K_n^{x_1,\ldots,x_n} y$$

Newton versus Ockham

$$K_n K_n^{x_1,\ldots,x_n} y$$

This essentially encodes that Newton knows the Law of Gravitation.

How about Ockham?

Well, he doesn't know it. But given enough evidence e (pairs of input-outputs), he will come to know it.

$$K_o K_o^{e,x_1,\ldots,x_n} y$$

This is specific to Ockham: his knowledge $?_o$ is given by an Alexandrov topology= plausibility order, in which worlds governed by polynomial gravity laws of lower degree are more plausible than the ones of higher degree (or with non-linear laws).

Other agents might not know y after seeing the same evidence.

Logic

Syntax: given question variables x (and strings of variables $\overrightarrow{x}$), atomic propositions p and answer-labels U.

$$\varphi ::= \quad p \quad | \quad Q : U \quad | \quad K_{\overrightarrow{Q}} Q \quad | \quad \neg\varphi \quad | \quad \varphi \wedge \varphi$$

$$Q ::= \quad x \quad | \quad ?\varphi$$

Models: as above. Valuations map p's into $|p| \subseteq W$, x's into $|x| : W \to (X, \tau_x)$, U's into $|U| \in \tau_x$.

Semantics: $w \models Q : U$ iff $|U| \in \tau_x(|Q|(w))$, and obvious in rest.

Abbreviations

$$K_{\overrightarrow{Q}}\varphi \ := \ \varphi \wedge K_{\overrightarrow{Q}}?\varphi$$

$$K\varphi \ := \ K_{?\top}\varphi \ \text{(global modality)} \ .$$

$$\overrightarrow{Q} : \overrightarrow{U} \ := \ \bigwedge_i Q_i : U_i.$$

Axiomatization

$S5$ axioms and rules for global modality K.

$$\Big(\overrightarrow{x} : \overrightarrow{O} \wedge K(\overrightarrow{x} : \overrightarrow{O} \to \varphi)\Big) \to K_{\overrightarrow{x}}\varphi$$

From $\Big(\overrightarrow{Q} : \overrightarrow{O} \wedge K(\overrightarrow{Q} : \overrightarrow{O} \to \varphi)\Big) \to \eta$, infer $K_{\overrightarrow{Q}}\varphi \to \eta$, provided $\overrightarrow{O}$ are "fresh".

$$\big(K_{\overrightarrow{x}}y \wedge y : O\big) \to K_{\overrightarrow{x}}(y : O)$$

From $(\eta \wedge Q : O) \to K_{\overrightarrow{Q}}(y : O)$, infer $\eta \to K_{\overrightarrow{Q}}y$, provided O is fresh.

$$(\varphi \leftrightarrow ?\varphi : \varphi) \wedge (\neg\varphi \leftrightarrow ?\varphi : \neg\varphi)$$

Question Substitution: from $\varphi(x)$ infer $\varphi(Q)$.