Capturing Lewis's "Elusive Knowledge"

Zhaoqing Xu

Department of Philosophy, Peking University
zhaoqingxu@gmail.com

September 22, 2011

1. Introduction
2. Philosophical Background
   - Dretske's Relevant Alternatives Theory
   - Lewis's Elusive Knowledge
3. Formalization
4. Interaction with Belief
5. Dynamics
6. Conclusion

"Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge, it goes away."

— David Lewis, *Elusive Knowledge*, p.550.

Example [Dretske, 1970]:

- Do you know what this animal is?

Figure: Zebra? Or cleverly disguised mule?

Skeptical Paradox:

- (a) We know it is a zebra;

Skeptical Paradox:

- (a) We know it is a zebra;
- (b) We do not know it is not a cleverly disguised mule;

Skeptical Paradox:

- (a) We know it is a zebra;
- (b) We do not know it is not a cleverly disguised mule;
- (c) If we know something is a zebra, then we know it is not a cleverly disguised mule.

Skeptical Paradox:

- (a) We know it is a zebra;
- (b) We do not know it is not a cleverly disguised mule;
- (c) If we know something is a zebra, then we know it is not a cleverly disguised mule.

- (a') I know I have hands;
- (b') I do not know I am not handless and deceived (by evil demon or scientist);
- (c') If I know I have hands, then I know that I am not handless and deceived.

Solutions:

- Mooreans do *modus ponens* from (a) and (c), thus reject (b);

Solutions:

- Mooreans do *modus ponens* from (a) and (c), thus reject (b);
- Skeptics do *modus tollens* from (b) and (c), thus reject (a);

Solutions:

- Mooreans do *modus ponens* from (a) and (c), thus reject (b);
- Skeptics do *modus tollens* from (b) and (c), thus reject (a);
- Fred Dretske has another story to tell. He accepts both (a) and (b) but rejects (c), since he denies the closure principle of knowledge (that if one knows that P and knows that P implies Q, then one knows that Q) [Dretske, 1970, 1981];

Solutions:

- Mooreans do *modus ponens* from (a) and (c), thus reject (b);
- Skeptics do *modus tollens* from (b) and (c), thus reject (a);
- Fred Dretske has another story to tell. He accepts both (a) and (b) but rejects (c), since he denies the closure principle of knowledge (that if one knows that P and knows that P implies Q, then one knows that Q) [Dretske, 1970, 1981];
- However, David Lewis has still another story to tell. "Knowledge *is* closed under implication. ... Dretske gets the phenomenon right,...; it is just that he misclassifies what he sees. He thinks it is a phenomenon of logic, when really it is a phenomenon of pragmatics." [Lewis, 1996, p.564]

Dretske's RAT [1970; 1981]:

Dretske identifies a spectrum of sentential operators: *fully penetrating*, *semi-penetrating*, and *non-penetrating*.

- An operator O is *fully penetrating* if whenever P entails Q, $O(P)$ entails $O(Q)$; e.g., "it is true that", " it is a fact that", "it is necessary that", etc.

Dretske's RAT [1970; 1981]:

Dretske identifies a spectrum of sentential operators: *fully penetrating*, *semi-penetrating*, and *non-penetrating*.

- An operator O is *fully penetrating* if whenever P entails Q, $O(P)$ entails $O(Q)$; e.g., "it is true that", " it is a fact that", "it is necessary that", etc.
- *Non-penetrating* operators fails "to penetrate to some of the most elementary logical consequences of a proposition" [Dretske, 1970, p.1008]; e.g., "it is strange that", "it was a mistake that", "it was accidental that", etc.

RAT (Cont'):

- Dretske's thinks epistemic operators are *semi-penetrating*: (i) they are not non-penetrating, since "it seems ... fairly obvious that if someone knows that P and (ii) Q ... he thereby knows that Q" and "if he knows that P is the case, he knows that P or Q is the case" [ibid, p.1009]; (iii)they are not fully-penetrating, since Dretske denies the general closure principle of knowledge.

RAT (Cont'):

- Dretske's thinks epistemic operators are *semi-penetrating*: (i) they are not non-penetrating, since "it seems ... fairly obvious that if someone knows that P and (ii) Q ... he thereby knows that Q" and "if he knows that P is the case, he knows that P or Q is the case"[ibid, p.1009]; (iii)they are not fully-penetrating, since Dretske denies the general closure principle of knowledge.
- Formally, Dretske's semi-penetrating knowledge satisfies: (i) $K(\varphi \wedge \psi) \rightarrow K\varphi$ and (ii) $K\varphi \rightarrow K(\varphi \vee \psi)$; but fails to satisfy: (iii) $K(\varphi \rightarrow \psi) \rightarrow (K\varphi \rightarrow K\psi)$.

RAT (Cont'):

Dretske's RAT was supposed to explain the failure of closure.

- **Alternative:** Q is an alternative to P only if Q is incompatible with P;

RAT (Cont'):

Dretske's RAT was supposed to explain the failure of closure.

- **Alternative:** Q is an alternative to P only if Q is incompatible with P;
- **Relevant Alternatives:** "A relevant alternative is an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized" [ibid, p.1021].

RAT (Cont'):

- For any P, define the *contrasting set* ($CS(P)$) as the set of all alternatives to P; then the set of relevant alternatives (*relevancy set*, $RS(P)$) is always a *proper* subset of $CS(P)$, and may not be the same "from situation to situation even though what is known remain the same" [1981, p.371].

RAT (Cont'):

- For any P, define the *contrasting set* ($CS(P)$) as the set of all alternatives to P; then the set of relevant alternatives (*relevancy set*, $RS(P)$) is always a *proper* subset of $CS(P)$, and may not be the same "from situation to situation even though what is known remain the same" [1981, p.371].
- This can be depicted as follows:

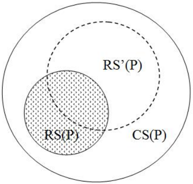

Figure: Dretske's RAT

RAT (Cont'):

- **Dretske-Knowledge:** "I propose to think of knowledge as an evidential state in which all relevant alternatives (to what is known) are eliminated" [1981, p.367].

RAT (Cont'):

- **Dretske-Knowledge:** "I propose to think of knowledge as an evidential state in which all relevant alternatives (to what is known) are eliminated" [1981, p.367].
- To see how closure fails, let P and Q be propositions such that $CS(Q) \subseteq CS(P)$ and $RS(P) \subset RS(Q)$, and E be an evidential state in which exactly all alternatives in $RS(P)$ are eliminated; then the subject knows that P and knows P implies Q, but fails to know that Q.

RAT (Cont'):

- **Dretske-Knowledge:** "I propose to think of knowledge as an evidential state in which all relevant alternatives (to what is known) are eliminated" [1981, p.367].
- To see how closure fails, let P and Q be propositions such that $CS(Q) \subseteq CS(P)$ and $RS(P) \subset RS(Q)$, and E be an evidential state in which exactly all alternatives in $RS(P)$ are eliminated; then the subject knows that P and knows P implies Q, but fails to know that Q.
- However, if we understand above example in standard semantics, it may also falsify (i) $K(\varphi \wedge \psi) \rightarrow K\varphi$ and (ii) $K\varphi \rightarrow K(\varphi \vee \psi)$; to see this, it suffices to notice that, when $[\![P]\!] \subseteq [\![Q]\!]$, $[\![P \wedge Q]\!] = [\![P]\!]$ for case (i), and $[\![Q]\!] = [\![P \vee Q]\!]$ for case(ii).

Lewis's Elusive Knowledge [1996]:

- Lewis's analysis is also known as a version of relevant alternatives theory, but it differs from Dretske's RAT at one significant aspect: what Lewis calls "alternatives" are possibilities, which obviously cannot be incompatible propositions. Since Lewis said:
- "The possibility actually obtains is never *properly ignored*; actuality is always a *relevant alternative;* nothing false may *properly be presupposed*". [1996, p.554]
- Notice that, Lewis uses "not properly ignored", " relevant alternatives", and "properly presupposed" interchangeably.

Elusive Knowledge (Cont'):

Definition (1, 1996, p.554)

S knows that P iff S's evidence eliminates every possibility in which not-P - Psst!-except for those possibilities that we are properly ignoring.

Definition (2, ibid, p.554)

S knows that P iff S's evidence eliminates every possibility in which not-P - Psst!-except for those possibilities that conflict with our proper presuppositions.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Prohibitive):
   - **Actuality:** The possibility actually obtain is never properly ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Prohibitive):
   - **Actuality:** The possibility actually obtain is never properly ignored.
   - **Belief:** A possibility that the subject believes or ought to believe to obtain is not properly ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Prohibitive):
   - **Actuality:** The possibility actually obtain is never properly ignored.
   - **Belief:** A possibility that the subject believes or ought to believe to obtain is not properly ignored.
   - **Attention:** A possibility not ignored at all is *ipso facto* not properly ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Prohibitive):
    - **Actuality:** The possibility actually obtain is never properly ignored.
    - **Belief:** A possibility that the subject believes or ought to believe to obtain is not properly ignored.
    - **Attention:** A possibility not ignored at all is *ipso facto* not properly ignored.
    - **Resemblance:** If one possibility is not properly ignored according to above three rules, then any possibility saliently resembles it is not properly ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Cont', Permissive):
    - **Reliability:** Possibilities concerning failures of reliable processes (such as perception, memory, and testimony) may properly be ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Cont', Permissive):
   - **Reliability:** Possibilities concerning failures of reliable processes (such as perception, memory, and testimony) may properly be ignored.
   - **Method:** Possibilities in which (a) the sample is not representative or (b) the best explanation is not the true explanation may properly be ignored.

Elusive Knowledge (Cont'):

1. Rules of Relevance (Cont', Permissive):
    - **Reliability:** Possibilities concerning failures of reliable processes (such as perception, memory, and testimony) may properly be ignored.
    - **Method:** Possibilities in which (a) the sample is not representative or (b) the best explanation is not the true explanation may properly be ignored.
    - **Conservatism:** Possibilities normally and commonly known to be ignored by people around us may properly be ignored.

Elusive Knowledge (Cont'):

- Back to skeptical argument (skeptic's *modus tollens*):
  - "The premise that 'I know I have hands' was true in its everyday context, where the possibility of deceiving demons was properly ignored. The mention of that very possibility switched the context midway. The conclusion 'I know that I am not handless and deceived' was false in *its* context, because that was a context in which the possibility of deceiving demons was being mentioned, hence was not being ignored, hence was not being properly ignored. ... Closure, rightly understood, survives the test. If we evaluate the conclusion for truth not with respect to the context in which it is uttered, but instead wit respect to the different context in which the premise was uttered, then truth is preserved."

"Maybe we do know a lot in daily life; but maybe when we look hard at our knowledge, it goes away."

— David Lewis, *Elusive Knowledge*, p.550.

Formalization:

1. To formalize Lewis's definition, we have at least three options:
   - (1) The first option is to represent uneliminated and relevant possibilities *wholly*, and define knowledge in terms of one accessible relation;

Formalization:

1. To formalize Lewis's definition, we have at least three options:
   - (1) The first option is to represent uneliminated and relevant possibilities *wholly*, and define knowledge in terms of one accessible relation;
   - (2) The second is to represent uneliminated and relevant possibilities *separately*, and define knowledge in terms of intersection of two accessible relations.

Formalization:

1. To formalize Lewis's definition, we have at least three options:
   - (1) The first option is to represent uneliminated and relevant possibilities *wholly*, and define knowledge in terms of one accessible relation;
   - (2) The second is to represent uneliminated and relevant possibilities *separately*, and define knowledge in terms of intersection of two accessible relations.
   - (3) The third option is similar to the second one, but it represents all relevant possibilities by a fixed set [Holliday, 2010, basic *RA* model; Rebuschi & Lihoreau, 2008].

Formalization:

1. To formalize Lewis's definition, we have at least three options:
   - (1) The first option is to represent uneliminated and relevant possibilities *wholly*, and define knowledge in terms of one accessible relation;
   - (2) The second is to represent uneliminated and relevant possibilities *separately*, and define knowledge in terms of intersection of two accessible relations.
   - (3) The third option is similar to the second one, but it represents all relevant possibilities by a fixed set [Holliday, 2010, basic *RA* model; Rebuschi & Lihoreau, 2008].
2. Since option (1) is too coarse and option (3) has its own problem (which we'll see latter), I prefer option (2).

Formalization (Cont'):

1. Properties for relevance(or properly presupposing) relation:
   - From Rule of Actuality, we can easily see that the relevance relation is *reflexive*.

Formalization (Cont'):

1. Properties for relevance(or properly presupposing) relation:
   - From Rule of Actuality, we can easily see that the relevance relation is *reflexive*.
   - Can we get more?
     No and Yes.
     No because other prohibitive rules said nothing on the that.
     Yes because from an analogy between "evidence" and "proper presuppositions", we can even assume it is an equivalence relation, i.e., the relevant possibilities are which satisfying our proper presuppositions.

Formalization (Cont'):

1. Properties for relevance(or properly presupposing) relation:
   - From Rule of Actuality, we can easily see that the relevance relation is *reflexive*.
   - Can we get more?
     No and Yes.
     No because other prohibitive rules said nothing on the that.
     Yes because from an analogy between "evidence" and "proper presuppositions", we can even assume it is an equivalence relation, i.e., the relevant possibilities are which satisfying our proper presuppositions.
   - For the underlying intuition, suppose you are not sure whether *w* or *v* is the actual possibility, then no matter which one turns out to the truth, the relevant possibilities should be the same.

Formalization (Cont'):

Definition (Language $\mathcal{L}$)

The formal language $\mathcal{L}$ consists of all formulas generated by the following BNF, where Δ is a countable set of propositional atoms, s is the subject of knowledge, a is the attributor, and $p \in \Delta$:

$$\mathcal{L} := p|\neg\varphi|\varphi \wedge \psi|E_s\varphi|P_a\varphi|K_s^a\varphi.$$

Remark

Here E_s, P_a, and P_s^a are all box operators. We can define their diamond dual as usual. Intuitively, $E_s\varphi$ means "s has evidence for φ" or "φ is true in very possibility conforming to s's evidence"; "$P_a\varphi$ means "a properly presuppose φ" or "φ is true in all relevant possibilities according to attributor a"; and $K_s^a\varphi$ means "s knows that φ according to attributor a" or simply "a thinks s knows that φ".

Formalization (Cont'):

Definition (Lewisian Epistemic Model)

A Lewisian epistemic model $\mathfrak{M}$ *is a tuple* $\langle W, \approx, R, V\rangle$, *where* W *is a non-empty set,* $\approx$ *and* R *are equivalence relations on* W, *and* $V : \Delta \rightarrow \wp(W)$ *is a valuation function.*

Formalization (Cont'):

Definition (Satisfaction)

For any pointed epistemic model $\mathfrak{M}, w$ and any formula $\varphi \in \mathcal{L}$, the satisfaction relation $\vDash$ is defined recursively (with Boolean cases as usual):

- $\mathfrak{M}, w \vDash E_s\varphi$ iff $\forall v \in W : w \approx v \Rightarrow \mathfrak{M}, v \vDash \varphi$;
- $\mathfrak{M}, w \vDash P_a\varphi$ iff $\forall v \in W : wRv \Rightarrow \mathfrak{M}, v \vDash \varphi$;
- $\mathfrak{M}, w \vDash K_s^a\varphi$ iff $\forall v \in W : (w \approx v \& wRv) \Rightarrow \mathfrak{M}, v \vDash \varphi$.

Formalization (Cont'):

Definition (Axiomatization)

The logic LEL can be axiomatized by the axioms and rules below, where $\Box \in \{E_s, P_a, K_s^a\}$:

1. *Taut: All substitutional instances of propositional tautologies;*
2. *S5:*
   - *K:* $\Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi)$;
   - *T:* $\Box\varphi \to \varphi$;
   - *4:* $\Box\varphi \to \Box\Box\varphi$;
   - *E:* $\neg\Box\varphi \to \Box\neg\Box\varphi$;
3. *IA:* $E_s\varphi \vee P_a\varphi \to K_s^a\varphi$;
4. *MP: From* $\vdash \varphi$ *and* $\vdash \varphi \to \psi$, *derive* $\vdash \psi$;
5. *Generalization: From* $\vdash \varphi$, *derive* $\vdash \Box\varphi$.

Formalization (Cont'):

Theorem (Completeness)

The logic LEL is sound and complete with respect to the class of all Lewisian epistemic models.

Remark

- *First, our formalization conforms to Timothy Williamson's remark that Lewis accepts S5 for epistemic logic; it also answers partly to Holliday's question how a relevant alternatives theorist should handle higher-order knowledge.*

Formalization (Cont'):

Theorem (Completeness)

The logic LEL is sound and complete with respect to the class of all Lewisian epistemic models.

Remark

- *First, our formalization conforms to Timothy Williamson's remark that Lewis accepts S5 for epistemic logic; it also answers partly to Holliday's question how a relevant alternatives theorist should handle higher-order knowledge.*
- *Second, our formalization has axiom T, thus avoids the factivity problem of knowledge; and it also avoids the problem of evaluating epistemic formula at irrelevant possibilities (which we'll see later).*

Interaction with Belief:

- **Belief:** A possibility that the subject believes or ought to believe to obtain is not properly ignored.
- Here we have "believe" and "ought to believe" (which Lewis interprets as "evidence and arguments justify him in believing"). So, we introduce two operators for them, B_s and B_s^o, respectively.

Interaction with Belief (Cont'):

Definition (Semantics)

For the semantics, we expand Lewisian epistemic model with two relations R_s *and* R_o*, which are serial, transitive and Euclidean, and satisfying* $R_s \subseteq R$, $R_o \subseteq R$ *and* $R_o \subseteq \approx$ *(?). Then the truth for* $B_s\varphi$ *and* $B_s^o\varphi$ *are defined as follows:*

- $\mathfrak{M}, w \vDash B_s\varphi$ *iff* $\forall v \in W : wR_sv \Rightarrow \mathfrak{M}, v \vDash \varphi$;
- $\mathfrak{M}, w \vDash B_s^o\varphi$ *iff* $\forall v \in W : wR_ov \Rightarrow \mathfrak{M}, v \vDash \varphi$.

Interaction with Belief (Cont'):

- After adding belief, we have more constraints on P_a, namely, $P_a\varphi \rightarrow B_s\varphi$ and $P_a\varphi \rightarrow B_s^o\varphi$; therefore, for a presupposition to be proper, it has to be believed and ought to be believed by the subject.

Interaction with Belief (Cont'):

- After adding belief, we have more constraints on P_a, namely, $P_a\varphi \rightarrow B_s\varphi$ and $P_a\varphi \rightarrow B_s^o\varphi$; therefore, for a presupposition to be proper, it has to be believed and ought to be believed by the subject.
- If it is proper to represent "ought to believe" by condition $R_o \subseteq \approx$ (?), then $E_s\varphi \rightarrow B_s^o\varphi$ is also valid; further, we have $R_o \subseteq \approx \cap R$, which validates $K_s^a\varphi \rightarrow B_s^o\varphi$.

Interaction with Belief (Cont'):

- After adding belief, we have more constraints on P_a, namely, $P_a\varphi \rightarrow B_s\varphi$ and $P_a\varphi \rightarrow B_s^o\varphi$; therefore, for a presupposition to be proper, it has to be believed and ought to be believed by the subject.
- If it is proper to represent "ought to believe" by condition $R_o \subseteq\approx$ (?), then $E_s\varphi \rightarrow B_s^o\varphi$ is also valid; further, we have $R_o \subseteq\approx \cap R$, which validates $K_s^a\varphi \rightarrow B_s^o\varphi$.
- **Problem:** However, it seems our definition of B_s^o only takes care of "evidence" part. We still need a modification to account for the "arguments" part.

Retracting Presupposition

Definition (Retracting relation)

Given a formula $\chi \in \mathcal{L}$*, and two pointed Lewisian epistemic models* $(\mathfrak{M}, w)$ *and* $(\mathfrak{M}', w')$*,* $(\mathfrak{M}, w) \xrightarrow{-\chi} (\mathfrak{M}', w')$ *if:*

- $W = W'; w = w'; \approx = \approx'; V = V';$
- $R(w) \nvDash \chi \& R(w) = R'(w)$*, or*
  $R(w) \vDash \chi \& (R(w) \subset R'(w)) \& R'(w) \setminus R(w) \vDash \neg\chi.$

Definition (Retracting Presupposition (Cont'))

we extend $\mathcal{L}$ *with operators* $[-\chi]$ *for each* $\chi \in \mathcal{L}$*, then the truth for* $[-\chi]\varphi$ *is defined as:*

- $\mathfrak{M}, w \vDash [-\chi]\varphi$ *iff*
  $\forall(\mathfrak{M}', w')[(\mathfrak{M}, w) \xrightarrow{-\chi} (\mathfrak{M}', w') \Rightarrow \mathfrak{M}', w' \vDash \varphi].$

Elusive Knowledge

Proposition (Closure)

Lewisian knowledge is closed under known implication w.r.t. a fixed context, but not closed under known implication across context changes:

- $\models K^a_s\varphi \rightarrow (K^a_s(\varphi \rightarrow \psi) \rightarrow K^a_s\psi)$;
- $\not\models K^a_s\varphi \rightarrow [-\psi](K^a_s(\varphi \rightarrow \psi) \rightarrow K^a_s\psi)$.

Problem: how to give a complete set of reduction axioms (esp. for $[-\chi]P_a\varphi$ and $[-\chi]K^a_s\varphi$)?

Conclusion

- **Conclusion:**
- (i) I proposed a formalization of Lewis's "Elusive Knowledge", which is free of the problem of evaluating epistemic formula at irrelevant possibilities and the factivity problem.
- (ii) I proposed a dynamic semantics for retracting presupposition (context shift).

Conclusion

- **Conclusion:**
- (i) I proposed a formalization of Lewis's "Elusive Knowledge", which is free of the problem of evaluating epistemic formula at irrelevant possibilities and the factivity problem.
- (ii) I proposed a dynamic semantics for retracting presupposition (context shift).
- **Problems:**
- (i) How to modify the formal definition of "ought to believe" to take "arguments" into account?
- (ii) How to axiomatize the logic of retracting presupposition?

Conclusion (Cont'):

- **Future Works:**
- (i) Use the formalization to analyze "missed clues" [Schaffer, 2001] and "pure ignorance" [Lihoreau, 2008] challenges to Lewisian knowledge.
- (ii) Use the formalization to analyze the "inconsistency" objection to epistemic contextualism [Baumann, 2008; 2010].
- (iii) Find a proper semantics for Dretske's semi-penetrating knowledge.