

DRAFT

Epistemic Closure in Context

Yves Bouchard

1. Epistemic closure

The general principle of epistemic closure stipulates that epistemic properties are transmissible through logical means. The principle of *epistemic closure under known entailment* (ECKE), a particular instance of epistemic closure (EC), has received a good deal of attention since the last thirty years or so. ECKE states that: if one *knows that p*, and she *knows that p entails q*, then she *knows that q*. It is generally accepted that ECKE constitutes an important piece of the skeptical argument, but the acceptance of an unrestricted version of ECKE is still a matter of debate. I agree with Richard Feldman when he writes that “an unrestricted closure principle is false” (1995, 487). The question remains: *under what conditions does EC apply?*

Since Dretske’s (1970) seminal paper, several strategies have been explored by epistemologists to evaluate the EC principle, most notably the strategy based upon the analysis of the necessary conditions of knowledge, or k-conditions, (e.g. Nozick, 1981; Warfield, 2004; Brueckner, 2004; Murphy, 2006), and the strategy based upon the relevant alternatives view (e.g. Dretske, 1970; Stine, 1976; Heller, 1999). Several solutions provided by RA theorists have been spelled out in externalist terms, essentially by means of a counterfactual analysis. Dretske (1970; 2006) and Nozick (1981, 204 ff.) have notoriously argued against EC from two different perspectives, but they nonetheless share a counterfactual interpretation of relevant alternatives. As for me, I will rather follow Williams (1996) regarding what he takes to be a *non sequitur* between the Dretske-Nozick strategy and nonclosure. I want to preserve both EC and the notion of relevant alternatives, which is, as Goldman underlines it, a major issue:

The qualifier ‘relevant’ plays an important role in my view. If knowledge required the elimination of all logically possible alternatives, there would be no knowledge (at least of contingent truths). If only *relevant* alternatives need to be precluded, however, the scope of knowledge could be substantial. This depends, of course, on which alternatives are relevant. (1976, 775)

Dretske’s defines a relevant alternative in the following way: “A relevant alternative is an alternative that might have been realized in the existing circumstances if the actual state of affairs had not materialized.” (1970, 1021) Dretske was well aware of the limits of his own definition, which demands that a knowledge claim be evaluated in function of possible worlds. But the possible worlds semantics is not without difficulties. For instance, selecting an epistemic criterion for ordering the set of possible worlds in order to make possible a measure of world proximity and to fix the boundaries of relevance is not an uncontroversial task. Ontological relevance is an obscure notion and world proximity alone does not seem to help much here. What I would like to submit is a definition of RA that is strictly epistemological rather than ontological, and that will enable me to preserve EC. It is worth noting at this point that the gist of my suggestion lies in moving the starting point from an externalist perspective to an internalist one.

In a very clarifying paper, which presents the several epistemological motivations behind the *relevant alternatives* strategy (RA), Rysiew concludes that “dressing disagreements about

closure or contextualism up in the language of disputes about the RA approach itself is, and has been, counter-productive.” (2006, 276) I am afraid (*alas!*) that I will be ‘counter-productive’, since what I want to do is precisely to link tightly the notion of relevant alternatives to the notion of context in order to restrict the domain of application of the EC principle, so that the domain of the closure principle is fixed by the domain of the relevant alternatives. Therefore, in the end, what I am promoting is a restricted version of the EC principle. This should not sound more awkward than defending the idea that the principle of elimination of double negation is, in some way, *restricted* to classical logic (and stronger systems).

2. Epistemic context

From a conversational point of view, a context is what is *taken for granted*. From an epistemological point of view, I will define an epistemic context as a set of basic (or contextual) beliefs. That set includes beliefs about our environment but also beliefs about our own epistemic attitudes, epistemic standards, and so on. For instance, I may entertain beliefs like: my sensory experience is generally reliable, or, in normal conditions, if I see an object *o* with property *P* then I know that $P(o)$.¹ These beliefs are nothing but *pragmatic presuppositions*, as Stalnaker has correctly pointed out:

The distinction between presupposition and assertion should be drawn, not in terms of the content of the propositions expressed, but in terms of the situations in which the statement is made – the attitudes and intentions of the speaker and his audience. Presuppositions, on this account [pragmatic account], are something like the background beliefs of the speaker – propositions whose truth he takes for granted, or seems to take for granted, in making his statement. (1999, 48)²

The set of contextual beliefs exhibits two noticeable features: (1) the contextual beliefs are simply taken as true (and, of course, they *may* be true); (2) the contextual beliefs are not justified, i.e. the contextual beliefs are not part of the *justification space*. In order to convert a contextual belief *p* into knowledge, *p* has to meet the epistemic standards in use. The first consequence of this analysis is that an epistemic agent *cannot know a presupposition*. Dretske has already underlined that point: “These presuppositions, although their truth is entailed by the truth of the statement, are not part of what is *operated on* when we operate on the statement with one of our epistemic operators.” (1970, 1014) Wittgenstein (1969) also insisted particularly on that aspect too, by means of his analysis of the difference between certainty and knowledge (contra Moore).

Now, in virtue of this characterization of an epistemic context, I can define a relevant alternative as *an alternative that does not affect the given epistemic context*. An alternative may affect an epistemic context in many ways. For instance, it may require an additional presupposition, or several presuppositions, or it may require that a given presupposition becomes justified. But as soon as such a change has taken place, a new epistemic context has been set. How can one determine if a new alternative involves a context shift? A formal answer to this question calls for a complete theory of presupposition, a theory linguists and philosophers of language are still trying to develop. But, fortunately, from an informal standpoint, the following condition will do the work: *if the presuppositions of the conclusion*

¹ Unlike Goldman’s (1986) transition rules, these beliefs are all spelled out in epistemic terms.

² See also Lewis (1979).

of KMP do not affect the epistemic context, then EC applies.³ What is determinative here is what I do take for granted in the *actual* epistemic context, not what would happen in a *counterfactual* situation, a possible world close or not to the one I am in.

This characterization of the notion of epistemic context can serve as a basis for further definitions, that will prove useful when addressing other epistemological puzzles (e.g. KK thesis). For example, a *permissible positive context change* can be defined in this way:

A permissible positive context change with respect to an epistemic context c is a change that involves the addition of a (compatible) presupposition r to c , such that if $c = \{p, q\}$ then $r \rightarrow \neg p \vee \neg q$ and the agents must satisfy $Believe(p \wedge q \wedge r)$.

In the case where a tacit epistemic standard (a qualification method) is at issue, the context can be envisaged as an *epistemological context*, otherwise it is simply a *normal epistemic context*, i.e. a context in which knowledge simply operates.

3. Context shift

Before considering the EC principle in the light of my proposal, there is an aspect of the closure principle that is in need of clarification because there is a threat of equivocation that lurks in the debate. This aspect seems to have been overlooked in the literature. An important obstacle to a proper evaluation of the problem is the *apparent homogeneity* of the K-operator. In the general formulation of the EC principle, we have to pay close attention to the characterization of the three K-operators involved: $K_1\phi$, $K_2(\phi \supset \psi)$, $K_3\psi$. The reasons why an agent s knows that ϕ and knows that $\phi \supset \psi$ might be empirical and different from each other, but the reason why s knows that ψ has to be logical. S has to believe in the validity of the modus ponens to get $K_3\psi$. Consequently, s knows that ϕ and s knows that ψ do not convey exactly the same meaning. There is a significant difference for s with regard to her epistemic commitment to ϕ (or $\phi \supset \psi$) and to ψ . That means if s is challenged about her knowledge of ψ , s would rely upon logical reasons (classical logic, for instance) and if in some circumstances modus ponens cannot be applied, then neutralizing the logical reasons would have no effect on the non-logical reasons for the knowledge of ϕ and the knowledge of $\phi \supset \psi$ — as far as the logical and non-logical reasons are independent. So, there is a shift in the epistemic standards used to qualify the knowledge of ϕ and the knowledge of ψ .⁴

Now, let's consider Dretske's (1970, 1015 ff.) paradigmatic case:

K(p): I know that this animal is a zebra.


K(p \supset \neg q): I know that if this animal is a zebra, then this animal is not a cleverly disguised mule.

K(\neg q): I know that this animal is not a cleverly disguised mule.

Let's represent the situation this way:

³ For a similar view, see Williams (1996, 329).

⁴ This is more obvious in the case of Gettier's type II problems, where the principle at stake is epistemic closure for disjunction: $K\phi \supset K(\phi \vee \psi)$. This shift in justification is generally not allowed in non-strictly logical contexts.



$\{a, b, K_1, K_2\}$ is the set of contextual beliefs, call it the *presuppositional belt*, in which $\{K_1, K_2\}$ are epistemic standards, the solid arrows give the meaning of the K-operator, and the dotted arrows mean ‘is presupposed by’. So, the difference in the epistemic standard used to qualify p and $\neg q$ is made explicit. Knowledge of p satisfies the epistemic standard K_1 , whereas knowledge of $\neg q$ satisfies the epistemic standard K_2 (in this case, classical logic). Taking advantage of my previous definitions, I can say that if q is a relevant alternative, then the presuppositional belt remains unaffected and even though q is justified only in virtue of logic it *can* be (and eventually it *should* be) evaluated against the epistemic standard K_1 . On the other hand, if q is not a relevant alternative, then to make it relevant would require to alter the epistemic context. These changes, more or less pervasive, would be echoed in the entire belief network. In Dretske’s example, q presupposes a number of things that are not part of the context, and some of the required presuppositions by q are clearly incompatible with the actual context. For instance, q presupposes that someone has cleverly disguised the mules and that someone had the intention to fool the zoo visitors for some reason, call this presupposition b . This presupposition is incompatible with the one, call it a , according to which the visitor believes she is visiting the zoo ‘under normal circumstances’, viz. if the zoo authorities have marked the pen with ‘zebras’ then the animals in the pen *are* zebras.⁵ So, in order to accommodate the context for q , presupposition a has to be denied (or simply rejected) — exactly in the same manner Lewis’ rules of accommodation would require it. If the agents agree to do so (tacitly or explicitly), and by the same token agree to raise the epistemic standards, then K_1 becomes inappropriate as an epistemic standard, a stronger one is needed, one that will permit the discrimination between p (being a zebra) and q (being a cleverly disguised mule).⁶ As a result, $K(p)$ is henceforth false (it is false that the visitor knows it is a zebra), $K(q)$ cannot be justified logically, and the whole case is not a case for epistemic closure anymore. Williams (1996, 329) has brought this last point into clear focus.

My point here is not only a matter of accent, but most importantly a matter of what comes first. For instance, Stine, who wants to preserve deductive closure while accepting the RA view, suggests the following analysis:

“to say that John knows that p does normally presuppose that not- p is a relevant alternative. This is, however, a pragmatic, not a semantic presupposition. That is, it is the speaker, not the sentence

⁵ The difference between ‘normal’ and ‘bizarre’ worlds does not rely upon a distance relation among a set of possible worlds.

⁶ See Goldman’s (1976), minus the counterfactual analysis, and Lewis’ (1979), minus the automatic character of the rise of standards.

(or proposition) itself, who does the presupposing. Thus, the presupposition falls in the category of those which Grice labels ‘cancellable’.” (1976, 255)

The problem with such an account is that *relevance itself* is presupposed and therefore taken as a kind of primitive notion. On my view, the notion of pragmatic presupposition comes first and can serve as a *definiens* for the notion of relevant alternative. From that perspective, relevance is a property of a relation between an alternative and an epistemic context (not an absolute property), and the principle of deductive closure is valid insofar the pragmatic context remains the same, in other words, insofar the presuppositions remain the same. For Stine, it is rather the set of relevant alternatives itself that has to remain fixed: “my account holds the set of relevant alternatives constant from beginning to end of the deductive closure argument. This is as it should be; to do otherwise would be to commit some logical sin akin to equivocation.” (1976, 256; also Lewis, 1996) Moreover, Stine’s view may allow an agent to refuse explicitly to presuppose an alternative (imagine a crucial one, for instance), since relevant alternatives are alternatives taken for granted. But, in order for an agent to refuse a relevant alternative in my sense, it would not suffice to simply rule out one particular alternative by *fiat* or collective agreement. The agent would have to refuse the whole logic that justifies it, as well as what it presupposes. This is one reason why deductive closure is so important as an instrument in our scientific endeavors: it can force us, as long as we agree on (one) logic, to explore hidden possibilities within the limits of a given and fix epistemic context.

Despite its internalist commitment, my proposal remains in line with some contentions of Goldman and Heller, who both defend a counterfactual interpretation of the RA view. My account only makes explicit something that is kept more or less implicit in their views. For instance, Goldman acknowledges the role of the context in the determination of some alternatives:

“It is not only the circumstances of the putative knower’s situation, however, that influence the choice of alternatives. The speaker’s own linguistic and psychological contexts are also important. If the speaker is in a class where Descartes’ evil demon has just been discussed, or Russell’s five-minute-old-world hypothesis, he may think of an alternative he would not otherwise think of and will perhaps treat them seriously.” (1976, 776)

Heller is more explicit regarding the relation between the context and the set of possible worlds. According to him, the pragmatic context plays two roles: (1) “the context in which the utterance is made determines which respects of similarity are to be assigned the most weight when ordering worlds”, and (2) it determines “how similar enough a world has to be to the actual world to be similar enough to be relevant.” (1999, 203)

I believe this shows, in last analysis, that what is primarily determinative in the selection of relevant alternatives is not the counterfactual situations at stake, but rather structural elements of the epistemic context at play. In qualifying alternatives, we do not need a semantic theory that will account for a difference between ‘normal’ and ‘bizarre’ worlds in terms of a proximity relation for a set of possible worlds, we only need to track context shifts. In the suggested perspective, a relevant alternative is context preserving, and that means it is a function of epistemic (and pragmatic) presuppositions.

4. Conceptual gains

If the above analysis is correct, then we get three conceptual gains out of it:

1. EC is vindicated within the limits of relevant alternatives, which are understood in terms of pragmatic presuppositions rather than in terms of counterfactual situations.
2. EC can be used as an epistemic contextual marker. A failure of EC clearly indicates a context shift, i.e. an alteration of the presuppositional belt.
3. This general framework provides interesting clues on the kinematics of epistemic normativity. As it turns out, epistemic normativity normally operates in the background, as the epistemic standards are presupposed by the agents, and provides the agents with a justification space where all the epistemic practices are well regimented by epistemic standards — this is the place where you plug your favorite epistemological option in the theory.

REFERENCES

- Brueckner, Anthony. 2004. "Strategies for Refuting Closure for Knowledge." *Analysis* 64: 333-335.
- Dretske, Fred. 1970. "Epistemic Operators." *The Journal of Philosophy* 67: 1007-1023.
- Dretske, Fred. 2006. "Information and Closure." *Erkenntnis* 64: 409-413.
- Feldman, Richard. 1995. "In Defence of Closure." *Philosophical Quarterly* 45: 487-494.
- Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73: 771-791.
- Goldman, Alvin I. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Heller, Mark. 1999. "Relevant Alternatives and Closure." *Australasian Journal of Philosophy* 77: 196 - 208.
- Lewis, David. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339-359.
- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74: 549-567.
- Murphy, Peter. 2006. "A Strategy for Assessing Closure." *Erkenntnis* 65: 365-383.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge: Belknap Press.
- Rysiew, Patrick. 2006. "Motivating the Relevant Alternatives Approach." *Canadian Journal of Philosophy* 36: 259-280.
- Stalnaker, Robert C. 1999. "Pragmatic Presuppositions." In *Context and Content*. Oxford: Oxford University Press.
- Stine, Gail C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies* 29: 249.
- Warfield, Ted A. 2004. "When Epistemic Closure Does and Does Not Fail: A Lesson From the History of Epistemology." *Analysis* 64: 35-41.
- Williams, Michael. 1996. *Unnatural Doubts. Epistemological Realism and the Basis of Scepticism*. Princeton: Princeton University Press.
- Wittgenstein, Ludwig. 1969. *On Certainty*. Translated by D. Paul, and G. E. M. Anscombe. Edited by G. E. M. Anscombe, and G. H. von Wright: Basil Blackwell.