$K \not\subseteq E$*

Elia Zardini

LanCog, Language, Mind and Cognition Research Group

Philosophy Centre

University of Lisbon

elia.zardini@campus.ul.pt

October 11, 2016

1 Knowledge and Evidence

In a series of very influential works, Tim Williamson has advanced and defended the much discussed thesis that *knowledge and evidence coincide* in the sense that:

*This paper grew out of my response to the late Tony Brueckner's paper '∼K∼SK' at the 2009 Arché Basic Knowledge Conference *Contemporary Perspectives on Scepticism* (University of St Andrews). Earlier versions of the material in the paper have been presented in 2013 at the session *Scepticism and Epistemic Circularity* of the 23rd World Congress of Philosophy *Philosophy as Inquiry and Way of Life* (University of Athens), at the 6th NIP Basic Knowledge Workshop *Themes in Basic Knowledge* (University of Aberdeen) and at the LOGOS Seminar (University of Barcelona); in 2015, at a research seminar at Yonsei University and at the LanCog Metaphysics, Epistemology, Logic and Language Seminar (University of Lisbon); in 2016, at the IIF-SADAF Seminar (University of Buenos Aires). I'd like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Eduardo Barrio, the late Tony Brueckner, José Díez, Dylan Dodd, Philip Ebert, Domingos Faria, Miguel Ángel Fernández, Filippo Ferrari, Dan López de Sa, Aidan McGlynn, Alberto Moretti, Eleonora Orlando, Nikolaj Pedersen, Sven Rosenkranz, Diogo Santos, Ricardo Santos, Martin Smith, Ernie Sosa, Crispin Wright, Dave Yates, José Zalabardo and an anonymous referee (who provided me with two rounds of comments that went far beyond the call of duty). At different stages during the writing of the paper, I've benefitted from an AHRC Postdoctoral Research Fellowship, from the FP7 Marie Curie Intra-European Research Fellowship 301493 on *A Non-Contractive Theory of Naive Semantic Properties: Logical Developments and Metaphysical Foundations* (NTNSP) and from the FCT Research Fellowship IF/01202/2013 on *Tolerance and Instability: The Substructure of Cognitions, Transitions and Collections* (TI), as well as from partial funds from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish Ministry of Science and Innovation on *Philosophy of Perspectival Thoughts and Facts* (PERSP), from the FP7 Marie Curie Initial Training Network 238128 on *Perspectival Thoughts and Facts* (PETAF), from the project FFI2011-25626 of the Spanish Ministry of Science and Innovation on *Reference, Self-Reference and Empirical Data* and from the project FFI2012-35026 of the Spanish Ministry of Economy and Competition on *The Makings of Truth: Nature, Extent, and Applications of Truthmaking.*

$K = E$ One knows P iff P is part of one's evidence[1]

(see e.g. Williamson [1997]; [2000], pp. 184–208). As a straightforward consequence, Williamson's theory of evidence contains in particular the thesis that *all knowledge is evidence* in the sense that:

$K \subseteq E$ If one knows P, P is part of one's evidence.

$K \subseteq E$ is *in itself* an interesting (and controversial) thesis. Moreover, as Williamson extensively articulates, if true $K \subseteq E$ would have *substantial repercussions* on several topics of crucial epistemological interest.[2] However, I shall argue that $K \subseteq E$ is false, and indeed that it is so for a reason that Williamson himself essentially provides.

2 $JTB \not\subseteq E$

It will prove helpful first to examine one specific part of Williamson's overall argumentation in favour of $K = E$. That part belongs to Williamson's defence of the *converse* of $K \subseteq E$, to the effect *all evidence is knowledge* ($E \subseteq K$). During that defence, Williamson critically considers the issue whether it is only, say, *justified true belief* (or some other similar good epistemic status falling short of knowledge) rather than knowledge that is *necessary* for a proposition to be part of one's evidence. Williamson offers several arguments against that kind of alternative, of which one is of particular interest for our purposes (I'll briefly discuss another one in fn 14). The argument naturally assumes that the thesis that the *only* good epistemic status *necessary* for a proposition to be part of one's evidence is justified true belief implies the thesis that all justified true belief is *sufficient* for a proposition to be part of one's evidence, for then effectively targeting the latter thesis. It proceeds as follows:

> If evidence required only justified true belief, or some other good cognitive status short of knowledge, then a critical mass of evidence could set off a kind of chain reaction. Our known evidence justifies belief in various true hypotheses; they would count as evidence too, so this larger evidence set would justify belief in still more true hypotheses, which would in turn count as further evidence.... The result would be very different from our present conception of evidence. (Williamson [2000], p. 201; *cf* Williamson [1997], p. 732)

[1]Throughout, I follow Williamson (and many others) and assume a *propositionalist* conception of evidence. As far as I can tell, insofar as $K = E$ and related theses can be so reformulated as to suit other conceptions of evidence, so can the point of this paper.

[2]One such topic is the treatment of a certain prominent kind of *scepticism*. In arguing against $K \subseteq E$, however, by no means do I intend to pave the way to that kind of scepticism, which I've argued in Zardini [2014a]; [2014b] is actually flawed in deeper ways than are made out by $K \subseteq E$.

While this argument is certainly already telling as it stands, I think it'll benefit from being supplemented with some *clarification* and *illustration*. Starting with the clarification, and focussing (as Williamson does) on justification, notice that the conception of justification presupposed by the argument would seem to be one according to which one can have a *justified true belief that is not epistemically good enough to amount to knowledge even in the absence of any "funny business"* of the "external" sorts variously exemplified by Gettier cases:[3] it is *the grounds themselves on which the belief is based* that, while *strong enough to make the belief justified*, are *too weak to make it knowledgeable*. Instances of this kind of conception are: the view according to which, given the usual probabilistic grounds, one can have a justified true belief that one's lottery ticket is a loser without knowing that it is; the view according to which, given the usual inductive grounds, one can have a justified true belief that one will be alive next year without knowing that one will; the view according to which, given the currently available historiographical grounds, one can have a justified true belief that Mary Stuart was involved in the murder of Lord Darnley without knowing that she was.[4] Of course, this kind of conception is far from being uncontroversial, as witnessed, for example, by the popularity in contemporary epistemology of the alternative kind of conception according to which, very roughly, one has a justified belief in P (if and) *only if one has a possible "internal" duplicate who knows* P (an early presentation and defence of this conception is offered by Bird [2007]). But, for better or worse, the former kind of conception would seem to be presupposed by the argument as stated, given that the argument clearly presupposes that the justifiedly believed evidence would be "larger" than the known evidence while apparently also presupposing that the "belief in various true hypotheses" need not be Gettiered (for not only no *mention* is made that the justified beliefs in question have any special feature such as being Gettiered, but also, without the second presupposition, the envisaged "chain reaction" would arguably be neither as *common* across bodies of known evidence nor as *protracted* for each body of known evidence as the argument would seem to assume).[5] I'll henceforth assume that conception as providing the most natural back-

[3]Throughout, I'll adopt a broad understanding of 'Gettier case' and its relatives, covering every case where some such "funny business" is present.

[4]It's tempting to put the idea by saying that, on the conception in question, a belief in P is justified if P's *epistemic probability* $\mathsf{Pr}(P)$ (i.e., roughly, the degree to which one's epistemic position supports P) is such that $\mathsf{Pr}(P)$ *is high enough even if* $\mathsf{Pr}(P) < 1$. But, setting aside the fact that the conception of justification presupposed by the argument need not endorse any such general or semi-technical sufficient condition for justification, we really shouldn't put the idea that way for the purposes of the view that it is only justified true belief rather than knowledge that is necessary for a proposition to be part of one's evidence. For, on the vast majority of conceptions of evidence and epistemic probability, that view implies that, if P is true and justifiedly believed, $\mathsf{Pr}(P) = 1$, which is at odds with saying that, as long as $\mathsf{Pr}(P)$ is high enough, a belief in P is justified even if $\mathsf{Pr}(P) < 1$. Thanks to an anonymous referee for bringing up this issue.

[5]Another clue pointing to the same conclusion about the conception of justification presupposed by the argument is given by a passage closely preceding the argument: "I have seen draws 1 to n; each was red (produced a red ball). I have not yet seen draw $n + 1$. I reason probabilistically, and form a justified belief that draw $n + 1$ was red too. My belief is in fact true [and the situation has clearly been described so that the belief need not be Gettiered, EZ]. But I do not know that draw $n + 1$ was red" (Williamson [2000], p. 200; *cf* Williamson [1997], p. 731). (I'll briefly discuss the argument where this passage occurs

ground for understanding the argument, although I hasten to add that also those who favour a different conception should be able to make sense of the rest of my discussion by making suitable adjustments (see fn 8 for an indication).[6]

Having clarified the conception of justification presupposed by the argument, consider now, by way of illustrating the kind of situation that the argument envisages, an expert gambler playing 1,000,000 rounds on a huge roulette with pockets from $\mathfrak{p}_0$ to $\mathfrak{p}_{1,000,000}$ and with the ball initially in $\mathfrak{p}_0$. We may suppose that the roulette is no random device: it is controlled by a croupier who, if, just before round $\mathfrak{r}_i$, the ball is in $\mathfrak{p}_j$, on round $\mathfrak{r}_i$ will try to make the ball land in $\mathfrak{p}_{j+1}$ ($1 \leq i \leq 1,000,000$, $0 \leq j < 1,000,000$). The croupier is fairly good but by no means infallible at this job. The gambler is well aware of all this; indeed, we may suppose that all this is part of the gambler's evidence, and that the gambler has no other insight into how the rounds will unfold than what she can work out by reflecting on these initial data.

Now, if any *non-deductive inference based on evidence* was ever epistemically good enough to provide *a justification for believing its conclusion*, albeit not epistemically good enough to provide *a suitable basis for knowing it*,[7] we may suppose that:

(J_1) Given that the set-up of the gambler example is part of one's evidence, the piece of evidence that, just before $\mathfrak{r}_i$, the ball is in $\mathfrak{p}_{i-1}$ provides a justification for believing that, on $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$ ($1 \leq i \leq 1,000,000$), albeit it does not provide a suitable basis for knowing it.[8]

---

in fn 14.)

[6]Thanks to an anonymous referee for comments that prompted this clarification on the conception of justification presupposed by the argument.

[7]Throughout, by 'justification' and its relatives, I mean *propositional* justification. Also, I use 'suitable basis for knowing' and its relatives so that, very roughly, a *suitable basis for knowing* stands to *knowledge* as *(propositional) justification* stands to *justified belief.* Only slightly less roughly, I understand "a suitable basis for knowing P" to be something such that, if P, one believes P on that basis and one's belief is not Gettiered, one knows P. Thus, on the one hand, the existence of "a suitable basis for knowing P" *is not definitionally guaranteed to entail* P, which, anticipating a bit, will be crucial for the formulation of principles governing non-deductive inferential knowledge like (K_1) in section 3 (fn 12); on the other hand, the existence of "a suitable basis for knowing P" *does definitionally guarantee that, if one forms a belief in* P *under certain specifiable circumstances, one knows* P, which, again anticipating a bit, will be crucial for some of the steps in the argument against $K \subseteq E$ in section 3 (something that would not be guaranteed if "a suitable basis for knowing P" were merely something such that it is *possible* that one bases a knowledgeable belief in P on it). Thanks to an anonymous referee for suggestions that led to some of these clarifications.

[8]While (J_1) is indeed plausible given the conception of justification presupposed by the argument, it becomes less so given other conceptions. But the gambler example can be so modified as to be made to fit those too. For instance, given the conception mentioned in the second last paragraph, we should drop the second embedded conjunct of (J_1) but could still secure that the gambler's beliefs that, on $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$ are not knowledgeable by modifying the example so that those beliefs are all Gettiered (say, by setting the example in some sort of "roulette-facade county"; modify the example further if, as described, you don't think that, given that conception, the truth of the first embedded conjunct of (J_1) is already secured). If, under these assumptions, you're now wondering whether it's really relevant for the acceptability of (J_2) (below in the text) that there be all these Gettier factors, you're all set for section 3. (Indeed, it is the assumption of the conception of justification presupposed by the argument, on which,

And, if any non-deductive inference based on evidence was ever epistemically not good enough to provide a justification for believing its conclusion, we may suppose that:

(J_2) The gambler's evidence does not provide a justification for believing that, on $\mathfrak{r}_{1{,}000{,}000}$, the ball will land in $\mathfrak{p}_{1{,}000{,}000}$.

Finally, add to all this that it just so happens that the world cooperates for JTB:

(JWC) On $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$ $(1 \leq i \leq 1,000,000)$.

The resulting situation is fatal for $JTB \subseteq E$, as, in such situation, that principle leads to a sorites-like clash between (J_1) and (J_2). By (J_1), the gambler has a justification for believing that, on $\mathfrak{r}_1$, the ball will land in $\mathfrak{p}_1$. Assuming that the gambler does perform the inference, the gambler has a justified belief that, on $\mathfrak{r}_1$, the ball will land in $\mathfrak{p}_1$, and, by (JWC), such belief is true. By $JTB \subseteq E$, it is then part of the gambler's evidence that, on $\mathfrak{r}_1$, the ball will land in $\mathfrak{p}_1$. But then, by (J_1), the gambler has a justification for believing that, on $\mathfrak{r}_2$, the ball will land in $\mathfrak{p}_2$. Assuming that the gambler does perform the inference, the gambler has a justified belief that, on $\mathfrak{r}_2$, the ball will land in $\mathfrak{p}_2$, and, by (JWC), such belief is true. By $JTB \subseteq E$, it is then part of the gambler's evidence that, on $\mathfrak{r}_2$, the ball will land in $\mathfrak{p}_2$. Another 999,997 analogous arguments yield then that, by (J_1), the gambler has a justification for believing that, on $\mathfrak{r}_{1{,}000{,}000}$, the ball will land in $\mathfrak{p}_{1{,}000{,}000}$, which contradicts ($J_2$). Therefore, keeping fixed (J_1) and (J_2) (and classical logic),[9] $JTB \subseteq E$ is false.

roughly, justification is *essentially easier to come by than knowledge*, that might give one the impression that the "chain reaction" only affects $JTB \subseteq E$ and not $K \subseteq E$, whereas I'll argue in section 3 that, unless knowledge is *so hard to come by as to give rise to a wide-ranging scepticism about non-deductive inference*, such "chain reaction" makes $K \subseteq E$ go kaboom just as well.)

[9]In Zardini [2016b], I've developed a general approach that manages to preserve a certain broad kind of closure principle (prominently including central closure principles for knowledge) from soritification-style arguments while granting surrounding intuitive claims usually thought to be inconsistent with such principles. The approach manages to do so by adopting a certain kind of *non-classical* (more precisely, *non-transitive*) logic on the grounds that, in the relevant examples, *it is vague in which cases the relevant putatively closed property is exemplified*, so that the *true theory of vagueness* should be used (which, as per e.g. Zardini [2008b], implies that the relevant closure principle is a *tolerance principle correctly expressing the vagueness of the relevant property*, and involves adoption of a non-transitive logic where *tolerance principles are consistent with the existence of positive and negative cases*, and so where soritical reasoning is *invalid*). *In principle*, such kind of move is available also for the soritification-style arguments of this section and of section 3 (that is, available if one is willing to buy into that theory of vagueness!): for instance, in the case of the gambler example, one could say that it is vague for which i it is part of the gambler's evidence that, on $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$ $(1 \leq i \leq 1,000,000)$. *In fact*, however, the soritification-style arguments of this section and of section 3 suggest directions of inquiry about the functions of, conditions for and nature of evidence which, as developed in the considerations of section 4, foreclose the kind of move just envisaged, as those considerations indicate that *evidence is not closed under the kinds of non-deductive inferences at play in those arguments*: for instance, in the case of the gambler example, those considerations indicate that, while it is definitely part of the gambler's evidence that the ball is in $\mathfrak{p}_0$, it is definitely not part of the gambler's evidence that, on $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$

3 $K \nsubseteq E$

I shall now argue that $K \subseteq E$ fails *for essentially the same reason* for which we've seen in section 2 that $JTB \subseteq E$ fails. Consider an expert geneticist studying a certain female rabbit $\mathfrak{f}_0$ with a certain designated property F_0 (say, fur colour) plus a series of 1,000,000 male rabbits going from $\mathfrak{m}_0$ to $\mathfrak{m}_{999,999}$ and such that $\mathfrak{m}_i$ has fur colour M_i $(0 \leq i < 1,000,000)$, where it is determined that $\mathfrak{f}_i$ will mate with $\mathfrak{m}_i$ generating another female rabbit $\mathfrak{f}_{i+1}$ $(0 \leq i < 1,000,000)$. We may suppose that fur-colour inheritance is no unfathomable matter: it is a prediction of genetics that, if a female rabbit with fur colour F_i mates with a male rabbit with fur colour M_i and generates another female rabbit, that rabbit will have fur colour F_{i+1} $(0 \leq i < 1,000,000)$.[10] The geneticist is well aware of all this; indeed, we may suppose that all this is part of the geneticist's evidence, and that the geneticist has no other insight into how the generations will unfold than what she can work out by reflecting on these initial data.

Now, if any *non-deductive inference based on evidence* was ever epistemically good enough to provide *a suitable basis for knowing its conclusion*, we may suppose that:

(K_1) Given that the set-up of the geneticist example is part of one's evidence, the piece of evidence that $\mathfrak{f}_{i-1}$ will have fur colour F_{i-1} provides a suitable basis for knowing that $\mathfrak{f}_i$ will have fur colour F_i $(1 \leq i \leq 1,000,000)$.[11,12]

$(1 \leq i \leq 1,000,000)$. (Zardini [2013b] considers another approach on which considerations of vagueness might block soritification-style arguments against a principle, which consists in *modifying the principle rather than the logic*: for instance, in the case of the gambler example and $JTB \subseteq E$, one could say that it is vague for which i the gambler has a justified true belief that, on $\mathfrak{r}_i$, the ball will land in $\mathfrak{p}_i$ $(1 \leq i \leq 1,000,000)$, and maintain that it is only all *determinately* justified true belief that is evidence. However, for the reasons developed in Zardini [2013b], pp. 779–782, such kind of move would also seem foreclosed.)

[10]To make it completely explicit, 'F_i' is not supposed to be *non-canonically* defined as 'the fur colour of $\mathfrak{f}_i$'; it is rather supposed to be *canonically* defined as, say, 'the colour with hex triplet #9F8170'. It is not supposed to be a *trivial truth*, but a *substantial prediction of genetics* (under the assumption that $\mathfrak{f}_{i-1}$ has fur colour F_{i-1} and $\mathfrak{m}_{i-1}$ has fur colour M_{i-1}) that $\mathfrak{f}_i$ will have fur colour F_i $(1 \leq i \leq 1,000,000)$.

[11]I hope it's clear that the role of the *specific* example I'm constructing isn't to be completely realistic, but to bring out the plausibility of an abstract point by illustrating it with some concreteness (which is why I've gone into some detail about the rabbits). While, even in this role, the example might certainly benefit from further elaboration and defence, I think that *some* principle along the broad lines of (K_1) must be true lest the science of genetics be epistemically bankrupt and incapable of providing knowledgeable predictions. Even if, for some reason, you think that genetics does fare so badly, I trust you'll be able to substitute *some* example that still does the trick and that conforms to the strictures of your theory of knowledge (if you're not so able, you might want to consider whether something's wrong with your theory of knowledge). Relatedly, I should stress that I'm focussing on *scientific* inference only because of (the names of some of my sponsors and) its especially paradigmatic status as a source of non-deductive inferential knowledge; clearly, on pain of a wide-ranging scepticism about *ordinary* inference, the argument I'm developing could be run by focussing instead on ordinary non-deductive inference. The problem with $K \subseteq E$ is thus not limited to science: more generally, $K \subseteq E$ would seem not to pay due attention to *knowledge yielded by non-deductive inferences* (the kinds of knowledge typically taken as examples by defenders of $K \subseteq E$ are in this sense revealing: for instance, Williamson [2000], pp. 184–208 does not contain a single clear-cut example of knowledge yielded by non-deductive inferences).

[12]Notice the crucial role played by the notion of a suitable basis for knowing (as defined in fn 7) in the

And, if any non-deductive inference based on evidence was ever epistemically not good enough to provide a suitable basis for knowing its conclusion, we may suppose that:

(K_2) The evidence of the geneticist does not provide a suitable basis for knowing that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$.

Finally, add to all this that it just so happens that the world cooperates for K:

(KWC) $\mathfrak{f}_i$ will have fur colour F_i and no Gettier factor is there ($0 \leq i \leq 1,000,000$).

The resulting situation is fatal for $K \subseteq E$, as, in such situation, that principle leads to a sorites-like clash between (K_1) and (K_2). By (K_1), the geneticist has a suitable basis for knowing that $\mathfrak{f}_1$ will have fur colour F_1. Assuming that the geneticist does perform the inference, the geneticist believes on a suitable basis that $\mathfrak{f}_1$ will have fur colour F_1, and, by (KWC), such belief is knowledgeable (recall fn 7). By $K \subseteq E$, it is then part of the geneticist's evidence that $\mathfrak{f}_1$ will have fur colour F_1. But then, by (K_1), the geneticist has a suitable basis for knowing that $\mathfrak{f}_2$ will have fur colour F_2. Assuming that the geneticist does perform the inference, the geneticist believes on a suitable basis that $\mathfrak{f}_2$ will have fur colour F_2, and, by (KWC), such belief is knowledgeable. By $K \subseteq E$, it is then part of the geneticist's evidence that $\mathfrak{f}_2$ will have fur colour F_2. Another 999,997 analogous arguments yield then that, by (K_1), the geneticist has a suitable basis for knowing that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$, which contradicts (K_2). Therefore, keeping fixed (K_1)[13] and (K_2) (and classical logic), $K \subseteq E$ is false.[14]

formulation of a principle governing non-deductive inferential knowledge like (K_1): that notion allows us to formulate *a completely general principle while leaving room for the possibility that some of the relevant non-deductive inferences just so happen to lead from truth to falsity* (in particular, we leave room for the possibility that, while the proposition that $\mathfrak{f}_{i-1}$ will have fur colour F_{i-1} is true and part of the geneticist's evidence, thus providing the geneticist with a suitable basis for knowing that $\mathfrak{f}_i$ will have fur colour F_i, it just so happens that $\mathfrak{f}_i$ will not have fur colour F_i).

[13]Maybe anti-scepticism about scientific inference can *by and large* be preserved by maintaining that (K_1) only fails for *one* i ($1 \leq i \leq 1,000,000$)? That would seem not to be a *stable* position: given that, assuming that the relevant premises are indeed part of the geneticist's evidence, all her inferences would seem *equally epistemically good*, it would rather seem that, if (K_1) fails for *one* i, it must fail for *every* i. Or maybe anti-scepticism about scientific inference can *to some extent* be preserved by maintaining that scientific inference is only epistemically good enough to provide a suitable basis for knowing *general laws* rather than *particular facts*? That would seem not to address the *essence* of the problem: presumably, an example similar to the geneticist example can be devised so that, instead of reasoning *induction-like* from particular facts to other particular facts, the relevant subject reasons *abduction-like* from general laws to other general laws (in this paper, I won't try though to give an even minimally realistic example thereof). Thanks to Sven Rosenkranz and Dave Yates for mention of these options.

[14]Williamson offers another argument against $JTB \subseteq E$:

> Suppose that balls are drawn from a bag, with replacement. In order to avoid issues about the present truth-values of statements about the future, assume that someone else has already made the draws; I watch them on film. For a suitable number n, the following situation can arise. I have seen draws 1 to n; each was red (produced a red ball). I have not yet seen draw $n + 1$. I reason probabilistically, and form a justified belief that draw $n + 1$

4 From Evidence to Knowledge

If $K \subseteq E$ is false, what does that mean for evidence? While a systematic account of evidence lies beyond the (mainly negative) scope of this paper, it is in order to mention some of the (positive) directions of inquiry suggested by the argument in section 3.[15] Let's start with some considerations that that argument suggests concerning the *functions* of evidence. In a dialectical context related to the argument discussed in section 2, Williamson himself appeals to *theoretically central functions of evidence* in order to argue for the thesis that evidence is propositional (which is in turn a presupposition of $K = E$; Williamson [1997], pp. 725–728; [2000], pp. 194–197). For example, he appeals to the function of evidence of *ruling out hypotheses inconsistent with it* (Williamson [1997], p. 727; [2000], p. 196). But, just as evidence provides a suitable basis for knowledgeably *rejecting* hypotheses *in conflict with* it, so evidence arguably provides a suitable basis for knowledgeably *accepting* hypotheses *adequately supported by* it. And such knowledgeability arguably covers cases where the relations of "being in conflict with" and "being

> was red too. My belief is in fact true. But I do not know that draw $n+1$ was red. Consider two false hypotheses:
>
> h: Draws 1 to n were red; draw $n+1$ was black.
>
> h^*: Draw 1 was black; draws 2 to n+1 were red.
>
> It is natural to say that h is consistent with my evidence and that h^* is not. In particular, it is consistent with my evidence that draw $n+1$ was black; it is not consistent with my evidence that draw 1 was black. Thus my evidence does not include the proposition that draw $n+1$ was red. Why not? After all, by hypothesis I have a justified true belief that it was red. The obvious answer is that I do not *know* that draw $n+1$ was red; the unsatisfied necessary condition for evidence is knowledge. (Williamson [2000], pp. 200–201; *cf* Williamson [1997], p. 731)

It seems to me that, if this kind of consideration can be taken to offer a good argument against $JTB \subseteq E$, it can equally be taken to offer a good argument against $K \subseteq E$. To wit, in the geneticist example, consider two false hypotheses:
h': $\mathfrak{f}_0$ has fur colour F_0; $\mathfrak{f}_1$ will not have fur colour F_1.
h'^*: $\mathfrak{f}_0$ does not have fur colour F_0; $\mathfrak{f}_i$ will have fur colour F_i $(1 \leq i \leq 1,000,000)$.
It is natural to say that h' is consistent with the geneticist's evidence and that h'^* is not. In particular, it is consistent with the geneticist's evidence that $\mathfrak{f}_1$ will not have fur colour F_1; it is not consistent with the geneticist's evidence that $\mathfrak{f}_0$ does not have fur colour F_0. Thus the geneticist's evidence does not include the proposition that $\mathfrak{f}_1$ will have fur colour F_1, even though, by (K_1) and (KWC) (and by performing the relevant inference), the geneticist *knows* that $\mathfrak{f}_1$ will have fur colour F_1! For what it's worth, I note that I find this kind of consideration—which relies on untutored "natural things to say" about examples where one has been primed to certain contrasts (notice that such priming seems to be the sole function of all the stuff about h^* and h'^* not being consistent with the evidence)—much less convincing and illuminating than the kind of consideration discussed in the text—which, more in keeping with the methodology prevailing in natural and formal sciences, relies instead on the rigorous articulation of the consequences of the interaction of theoretically interesting general principles when these are tested under extreme conditions. Thanks to an anonymous referee for discussion of the argument treated in this fn.

[15]Thanks to an anonymous referee for recommending to include more considerations of a constructive nature than were present in a previous draft of this paper, and for making a number of suggestions in this direction.

adequately supported by" are *non-deductive*, lest a wide-ranging scepticism about non-deductive inference ensue. It is the function of evidence of providing a suitable basis for knowledgeably accepting hypotheses adequately supported by it that offers a very natural *general framework* able to *ground* a *specific claim* such as (K_1), and so it is this theoretically central function of evidence that the argument in section 3, granted its other assumptions, shows to ground $K \nsubseteq E$.[16]

Such view of the functions of evidence leads to a "*stratified*" conception of the functions played in *epistemic structure* by evidence and knowledge respectively. Williamson [2000], p. 186 writes: "$E = K$ suggests a very modest kind of foundationalism, on which all one's knowledge serves as the foundation for all one's justified beliefs" (*cf* Williamson [1997], p. 719). But the main thrust of the last paragraph shows that epistemic structure is richer than is dreamt of in $E = K$'s suggestions: by grounding (K_1) and so $K \nsubseteq E$, the function of evidence of providing a suitable basis for knowledgeably accepting hypotheses adequately supported by it requires that, deeper down, at the foundation of "all one's knowledge", there be a set of propositions that are *epistemically privileged vis-à-vis knowledge*: propositions whose *epistemic status* is such that, on the one hand, and contrary to knowledge, propositions enjoying it are guaranteed to be able to serve as *premises* for certain knowledge-conferring (deductive or non-deductive) inferences, but also such that, on the other hand, and again contrary to knowledge, it is not guaranteed to be conferred on the *conclusions* of those evidence-based inferences. It's because evidence is *epistemically so good* as to be guaranteed to be able to occur as premise of certain knowledge-conferring inferences that evidence is *epistemically too good* to be guaranteed to occur as conclusion of such inferences. The foundations provided by evidence for knowledge unfailingly sustain without indefinitely extending.[17,18]

[16] Suppose that (KWC) fails, and, in particular, that, for some i, $\mathfrak{f}_i$ will not have fur colour F_i ($1 < i \leq 1,000,000$). Does the geneticist's false belief that $\mathfrak{f}_i$ will have fur colour F_i then not reflect badly on the *reliability* of the geneticist's belief that $\mathfrak{f}_1$ will have fur colour F_1? No, because, from the point of view developed in this paper, the two beliefs *are not formed in a relevantly similar way*: the former belief is formed by inference *from something that is not part of the geneticist's evidence*, whereas the latter belief is formed by inference *from the geneticist's evidence*. Thanks to an anonymous referee for bringing this matter to my attention.

[17] Does this broadly chime with the *philosophical account* given e.g. in Zardini [2015] of the *non-transitive logics* developed e.g. by Zardini [2008a]? Indeed it does, but this paper is not the place to explore this connection further.

[18] In arguing that there is, as it were, a one-way road from evidence to knowledge, by no means do I intend to suggest an identification of evidence with *foundations* (i.e., roughly, with those propositions which enjoy a good epistemic status not in virtue of other propositions enjoying some such status, and typically in virtue of which other propositions enjoy some such status). In one direction, *not all evidence is foundational* (for example, it is part of my evidence that the cube is coloured but that proposition is not plausibly foundational, as it is plausible that its good epistemic status depends on the good epistemic status of the more specific proposition that the cube is blue; see fns 27 and 32 for further kinds of examples). In the other direction, *not all foundations are evidence* (for example, some sort of inductive principle is plausibly foundational, as it is plausible that its good epistemic status does not depend—deductively or inductively—on the good epistemic status of some propositions, but it is not part of my evidence; see the sixth to eighth next paragraphs for a related discussion of the relationships between evidence and non-inferential knowledge).

Let's briefly explore further this view of the functions of evidence by sketching how it might be fruitfully applied to a vexed question in epistemology that this time concerns the knowledge-yielding power of *deductive* rather than *non-deductive* inference. Consider the well-known *preface-paradox-style* argument purportedly refuting *multi-premise* closure of knowledge (see Makinson [1965] for the original preface paradox).[19] One natural and (to me at least) appealing take on the matter is that the argument does succeed, but that it does not threaten a *weaker* principle of *multi-premise derivation of knowledge from evidence* according to which, if P_1, P_2, $P_3 \ldots$ are *part of one's evidence* (rather than merely *known by one*), and one knows that P_1, P_2, $P_3 \ldots$ entail Q, one knows Q (*such principle in its full generality* is entailed at least by a strong understanding of the view that one of the functions of evidence is to provide a suitable basis for knowledgeably accepting hypotheses adequately supported by it; *its single-premise cases* are entailed by that view anyhow).[20] The preface-paradox-style argument would seem not to threaten this weaker principle because, in the best preface-paradox-style examples, there is little plausibility to the thought—required to make trouble for the principle—that the premises are *not only* each known, but are *also* each part of one's evidence.[21] More strongly, the multi-premise derivation principle would seem to take care of the root of the problem of the multi-premise closure principle in that, for a variety of reasons, *evidence collection* (if P_1 is part of one's evidence and P_2 is part of one's evidence, $P_1 \wedge P_2$ is part of one's evidence) is in the relevant respects more plausible than *knowledge collection* (if one knows

[19]Here is a sketch of the preface-paradox-style argument that will do well enough for our purposes. Take all the propositions P_1, P_2, $P_3 \ldots$ intuitively known by Dave, a normal, fallible human being. Dave intuitively does not know that all of P_1, P_2, $P_3 \ldots$ are true (for one of many things, Dave intuitively does not know that, with respect to P_1, P_2, $P_3 \ldots$, he's always got it right; see Zardini [2016b] for the other things). But P_1, P_2, $P_3 \ldots$ jointly entail that all of P_1, P_2, $P_3 \ldots$ are true, and so, since Dave intuitively knows each of those premises, multi-premise closure seems to fail.

[20]Of course, so unguardedly stated, the principle is open to "boring" counterexamples just as unguardedly stated closure principles are. It is a matter of open debate in contemporary epistemology how best to refine closure principles so that they are proof against at least such counterexamples (see e.g. David and Warfield [2008] for an extended discussion), but I think that we can safely assume that whatever refinement turns out to be the right one for closure principles will also be essentially the right one for the derivation principle in the text. Similar qualifications should be understood to apply implicitly to a few other principles below.

[21]This is obviously so for the preface-paradox-style example I've given in fn 19, and it is also so for the original preface-paradox-style example where the premises are all the statements made in a book. But wait—presumably, *perception* and *memory* are sources of evidence, and can't we set up a *good enough*, even if not *best*, preface-paradox-style example by taking as premises all one perceives and all one remembers (of past evidence)? Well, presumably, as far as perception and memory are concerned, what more precisely are sources of evidence are only *manifest* perceptions and memories (rather than *dim* perceptions and memories and the like), so that the target preface-paradox-style example should take as premises only all one perceives in a manifest way and all one remembers in a manifest way. While acknowledging that these matters require more investigation, I note for the record that it does *not* seem clear to me that, in the case we've just ended up envisaging, one would not know the conclusion of the relevant multi-premise argument. Thanks to an anonymous referee for raising this issue.

P_1 and one knows P_2, one knows $P_1 \wedge P_2$).[22,23] Given the rich stratification of epistemic structure, epistemology should focus not only on *intra-level closure* principles, but also on *inter-level derivation* principles.

Of course, this natural and (to me at least) appealing take on the matter becomes non-sense in the presence of $K \subseteq E$, but this circumstance only *amplifies the problematicity of* $K \subseteq E$: keeping fixed minimal positive claims about what evidence one has and what one knows, that principle is not only incompatible with the conjunction of a very plausible principle governing *non-deductive* inferential knowledge (i.e. (K_1)) and a very plausible claim of ignorance (i.e. (K_2)), it is also incompatible with the conjunction of a very plausible principle governing *deductive* inferential knowledge (i.e. the multi-premise derivation principle)—a principle which *arguably preserves what is right about multi-premise closure*—and a very plausible claim of ignorance (i.e. the claim that, in a preface-paradox-style example, one does not know the conclusion of the relevant multi-premise argument)—a claim which *arguably preserves what is right about the preface-paradox-style argument.*

The previous considerations that the argument for $K \nsubseteq E$ suggests concerning the *functions* of evidence put us in a position to determine better what that argument suggests concerning the *conditions* for (and, possibly, even the *nature* of) evidence. Since the argument for $K \nsubseteq E$ mainly puts pressure on formulating an *interestingly weak sufficient*

[22]Let me sketch just but two prominent such reasons. Firstly, if P_1 is part of one's evidence and P_2 is part of one's evidence, on the vast majority of conceptions of evidence and epistemic probability it must be the case that $\mathsf{Pr}(P_1) = \mathsf{Pr}(P_2) = 1$, from which it follows, on the vast majority of conceptions of epistemic probability, that $\mathsf{Pr}(P_1 \wedge P_2) = \mathsf{Pr}(P_1) = \mathsf{Pr}(P_2) = 1$, whereas, if one knows P_1 and one knows P_2, on quite a few conceptions of knowledge and epistemic probability, it might be the case that $\mathsf{Pr}(P_1) < 1$ and $\mathsf{Pr}(P_2) < 1$, from which it would follow, on the vast majority of conceptions of epistemic probability, that, at least if P_1 and P_2 are independent, $\mathsf{Pr}(P_1 \wedge P_2) < \mathsf{Pr}(P_1)$ and $\mathsf{Pr}(P_1 \wedge P_2) < \mathsf{Pr}(P_2)$. Secondly, two main putative sources of evidence are, as per fn 21, perception and memory. But *perception collection* (if one perceives P_1 and one perceives P_2, one perceives $P_1 \wedge P_2$) and *memory collection* (if one remembers P_1 and one remembers P_2, one remembers $P_1 \wedge P_2$) are both in the relevant respects more plausible than knowledge collection.

[23]Obviously, this only "takes care of the root of the problem" if $E \subseteq K$ holds (which I've been implicitly assuming at a few places anyways). But, under our assumptions (fn 1), $E \subseteq K$ does strike me as much more plausible than $K \subseteq E$, although its defence is mainly a battle for another day. (By way of a demonstrative foray, consider that *reliability constraints* applying to knowledge would also seem to apply to evidence: for example, in barn-facade county (Goldman [1976]), it would seem not to be part of one's evidence that one is in front of a barn (although one is indeed in front of the only real barn in the county).) For those who remain sceptical of evidence collection, let me sketch a somewhat more direct argument for *evidence collection into knowledge* (if P_1 is part of one's evidence and P_2 is part of one's evidence, one knows $P_1 \wedge P_2$), the principle that is really needed to "take care of the root of the problem" (for what it's worth, setting $P_1 = P_2$, such principle also virtually entails as a special case $E \subseteq K$). Given what should be an unproblematic application of multi-premise closure, evidence collection into knowledge follows from the principle of *known evidence* (if P_1 and P_2 are part of one's evidence, one knows that P_1 and P_2 are part of one's evidence) and from knowledge of the principle of *true evidence* (if P_1 and P_2 are part of one's evidence, $P_1 \wedge P_2$ is true). Under our assumptions (fn 1), both principles strike me as very plausible (Williamson [1995]; [1996]; [2000], pp. 93–113 would famously disagree about the principle of known evidence, but Zardini [2012] would in turn disagree with him). Thanks to an anonymous referee for waking me up from a dogmatic slumber concerning $E \subseteq K$.

condition for being evidence, I'll focus on proposals in this sense. An initial proposal obtusely goes for one of the most straightforward strengthenings of K and is to the effect that *all knowledge of knowledge is evidence*. The tenability of such proposal depends of course on the behaviour of iterations of knowledge, a notoriously difficult issue on which there are wide and deep disagreements (see e.g. Williamson [1992] *vs* Zardini [2016c]). But I think that there are at least two very plausible sets of assumptions about iterations of knowledge each of which suffices for the argument for $K \nsubseteq E$ to sink $KK \subseteq E$ too.

Firstly, if any non-deductive inference based on evidence was ever epistemically good enough to provide a suitable basis for knowing that one knows its conclusion we may suppose that:

(KK_1) Given that the set-up of the geneticist example is part of one's evidence, the piece of evidence that $\mathfrak{f}_{i-1}$ will have fur colour F_{i-1}, provides a suitable basis for knowing that one knows that $\mathfrak{f}_i$ will have fur colour F_i ($1 \leq i \leq 1,000,000$).[24]

I think that some principle along the broad lines of (KK_1) must be true lest it be the case that, for all we know, the science of genetics is epistemically bankrupt and incapable of providing knowledgeable predictions. If it is on pain of an unacceptable wide-ranging scepticism about *scientific inference* that some principle along the broad lines of (K_1) must be true, it is on pain of a similarly unacceptable wide-ranging scepticism about *the falsity of that scepticism* that some principle along the broad lines of (KK_1) must be true. Often, scientific inference yields knowledge and, often enough, we know that. But then the argument for $K \nsubseteq E$ can be modified in the obvious way to the effect that, keeping fixed (KK_1) and (K_2) (and classical logic), $KK \subseteq E$ is false.

Secondly, the *kind* of principle governing non-deductive inferential knowledge underlying (K_1) would seem to admit of *iteration*, in the sense that it would seem that, if there is an instance of that kind saying that, given background evidence P, a piece of evidence Q provides a suitable basis for knowing R (as, for example, (K_1) in effect does), there is another instance of that kind saying that, given the background evidence that P is part of one's evidence, the piece of evidence that Q is part of one's evidence provides a suitable basis for knowing that one has a suitable basis for knowing R. Let 'P is part of one's evidence^i' and its relatives be short for 'It is part of one's evidence that it is part of one's evidence that it is part of one's evidence... ($i - 1$ times) that P is part of one's evidence' and its relatives ($i \geq 1$), and make analogous conventions for 'One has a suitable basis^i for knowing P' and its relatives as well as for 'One knows^i P' and its relatives. Then the kind of principle in question can be formulated as including, for every $i \geq 1$:

[24] It might not be completely obvious that the kind of evidence mentioned in (KK_1) is what typically supports claims about *knowledge* (as opposed to claims about, say, *rabbits*), but I've argued in Zardini [2016a]; [2016c] that things must indeed be so. Also, notice that, just as (J_1) and (K_1) implicitly assume that the relevant subject is versed in *ludic* and *genetic* matters respectively, so does (KK_1) (as well as (K_1^i) below in the text) assume that the relevant subject is also versed in *epistemic* matters. Thanks to an anonymous referee for observations that gave rise to these clarifications.

(K_1^i) Given that the set-up of the geneticist example is part of one's evidence^i, the piece of evidence^i that $\mathfrak{f}_{j-1}$ will have fur colour F_{j-1} provides a suitable basis^i for knowing that $\mathfrak{f}_j$ will have fur colour F_j $(1 \leq j \leq 1,000,000)$

(obviously, our original (K_1) is tantamount to (K_1^1)). Assume now that the geneticist performs all the relevant inferences and that she knows$^{2^{999,999}-1}$ that, for every relevant proposition P, if she has a suitable basis for knowing P, then (the world cooperates for knowledge of P in the sense exemplified in (KWC) so that) she knows P. Given a modicum of multi-premise closure of knowledge, these assumptions imply that the geneticist has a suitable basis^i for knowing P iff she knows^i P $(1 \leq i \leq 2^{999,999}$; I'll henceforth leave appeal to these equivalences tacit). Then the argument for $K \nsubseteq E$ can be modified as follows (assuming a modicum of multi-premise closure of evidence and that, if $KK \subseteq E$ holds, it is part of the geneticist's evidence$^{2^{999,999}-1}$ that it does). Given that the geneticist knows^2 about the set-up of the geneticist example, if she knows^2 that $\mathfrak{f}_{999,999}$ will have fur colour $F_{999,999}$, by $KK \subseteq E$ and (K_1^1) she knows that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$. But, given that she knows^4 about the set-up of the example, if she knows^4 that $\mathfrak{f}_{999,998}$ will have fur colour $F_{999,998}$, by its being part of her evidence that $KK \subseteq E$ holds and by (K_1^2) she does know^2 that $\mathfrak{f}_{999,999}$ will have fur colour $F_{999,999}$. But, given that she knows^8 about the set-up of the example, if she knows^8 that $\mathfrak{f}_{999,997}$ will have fur colour $F_{999,997}$, by its being part of her evidence^3 that $KK \subseteq E$ holds and by (K_1^4) she does know^4 that $\mathfrak{f}_{999,998}$ will have fur colour $F_{999,998}$. Putting these observations together with the other relevant 999,997 ones, we can conclude that, given that the geneticist knows$^{2^{1,000,000}}$ about the set-up of the example, if she knows$^{2^{1,000,000}}$ that $\mathfrak{f}_0$ has fur colour F_0, by its being part of her evidence$^{2^{999,999}-1}$ that $KK \subseteq E$ holds and by (K_1^i) she does know that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$. But the example can be described so that it very plausibly verifies the further assumption that the geneticist knows$^{2^{1,000,000}}$ about the set-up of the example and (thus) knows$^{2^{1,000,000}}$ that $\mathfrak{f}_0$ has fur colour F_0, while (K_2) remains compelling even under those circumstances. Therefore, keeping fixed all the relevant assumptions, (K_1^i) and (K_2) (and classical logic), $KK \subseteq E$ is false. Notice that, while strengthening the proposal to the effect that, for some i, all knowledge^i is evidence $(i \geq 3)$ might avoid the first objection, it would not do much to avoid the second objection (since, for every $i \geq 1$, the geneticist example can be described so that it very plausibly verifies the further assumption that the geneticist knows$^{i^{1,000,000}}$ about the set-up of the example as well as the initial assumptions that she knows$^{i^{999,999}-1}$ that, for every relevant proposition P, if she has a suitable basis for knowing P, she knows P and that, if $K^i \subseteq E$ holds, it is part of her evidence$^{i^{999,999}-1}$ that it does, while (K_2) remains compelling even under those circumstances). Evidence is inaccessible even to arbitrary iterations of knowledge.[25]

A better proposal, more perceptive to the workings of the argument for $K \nsubseteq E$, has it that *all non-inferential knowledge is evidence*. The tenability of such proposal depends of course on what counts as non-inferential knowledge, another notoriously difficult issue

[25]Essentially, the second objection *gets around restrictions on iteration* by *considering higher-order transition principles* and by *relying on iteration only on platitudes*. In this respect, the *dialectical move* and the *argumentative strategy* of the objection are the same as those of certain *paradoxes of higher-order vagueness* (see e.g. Zardini [2013a]).

on which there are wide and deep disagreements (see e.g. Hume [1748], Section X *vs* Reid [1764], Chapter VI, Section XXIV). But I think that there are at least two very plausible assumptions about non-inferential knowledge of which one suffices for the argument for $K \nsubseteq E$ to sink $K_{ni} \subseteq E$ too and the other one suffices for making $K_{ni} \subseteq E$ inconsistent with the multi-premise derivation principle (which, as I've observed in the fifth last paragraph, is entailed at least by a strong understanding of the view that one of the functions of evidence is to provide a suitable basis for knowledgeably accepting hypotheses adequately supported by it, a view that in turn, as I've observed in the seventh last paragraph, naturally accounts for a crucial premise—(K_1)—of the argument for $K \nsubseteq E$).

Firstly, it is very plausible that *testimony* typically provides non-inferential knowledge. But then the argument for $K \nsubseteq E$ can be modified as follows (substituting in (K_1) and (K_2) 'geneticists' for 'geneticist'). By (K_1), our original geneticist $\mathfrak{g}_1$ has a suitable basis for knowing that $\mathfrak{f}_1$ will have fur colour F_1. Assuming that $\mathfrak{g}_1$ does perform the inference, $\mathfrak{g}_1$ believes on a suitable basis that $\mathfrak{f}_1$ will have fur colour F_1 and, by (KWC), such belief is knowledgeable. Suppose then that $\mathfrak{g}_1$ communicates the conclusion of this inference to another geneticist $\mathfrak{g}_2$, who thereby comes to have non-inferential knowledge that $\mathfrak{f}_1$ will have fur colour F_1. By $K_{ni} \subseteq E$, it is then part of $\mathfrak{g}_2$'s evidence that $\mathfrak{f}_1$ will have fur colour F_1. But then, by (K_1), $\mathfrak{g}_2$ has a suitable basis for knowing that $\mathfrak{f}_2$ will have fur colour F_2. Assuming that $\mathfrak{g}_2$ does perform the inference, $\mathfrak{g}_2$ believes on a suitable basis that $\mathfrak{f}_2$ will have fur colour F_2 and, by (KWC), such belief is knowledgeable. Suppose then that $\mathfrak{g}_2$ communicates the conclusion of this inference to another geneticist $\mathfrak{g}_3$, who thereby comes to have non-inferential knowledge that $\mathfrak{f}_2$ will have fur colour F_2. By $K_{ni} \subseteq E$, it is then part of $\mathfrak{g}_3$'s evidence that $\mathfrak{f}_2$ will have fur colour F_2. Another 999,997 analogous arguments yield then that, by (K_1), geneticist $\mathfrak{g}_{1,000,000}$ has a suitable basis for knowing that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$, which contradicts (K_2). Therefore, keeping fixed (K_1) and (K_2) (and classical logic), $K_{ni} \subseteq E$ is false.

Secondly, assuming the take on multi-premise closure that I've recommended in the seventh last paragraph, it is very plausible that there are *good enough preface-paradox-style examples taking as premises only pieces of non-inferential knowledge* (suppose, for example, that *perception*, even when not manifest, typically provides non-inferential knowledge, and that one perceives that a far-away object in the twilight is a blue cube, that another far-away object in the twilight is a red sphere, that another far-away object in the twilight is a yellow pyramid etc.; *cf* fn 21). If the multi-premise derivation principle is true, $K_{ni} \subseteq E$ is false.[26,27]

An even better proposal, even more perceptive to the workings of the argument for $K \nsubseteq E$, abandons the idea of providing a sufficient condition for evidence in terms of some

[26] Beyond the issues raised by the argument for $K \nsubseteq E$, notice that $K_{ni} \subseteq E$ might be false also because of the last example presented in fn 18.

[27] In addition to being false, $K_{ni} \subseteq E$ also appeals to a sufficient condition for evidence which is, in a certain respect, *not interestingly weak*—truckloads of evidence is acquired precisely through inference (for example, it is part of my evidence that there is a statistical correlation between smoke and lung cancer but I've inferred that by collecting several pieces of information about particular cases; see fns 18 and 32 for further kinds of examples).

kind of *knowledge* and appeals instead to the informal notion of *data*, maintaining that *all data are evidence*.[28] Arguably, $D \subseteq E$ has more to go for it than just its consonance with the alphabet. *Like* non-inferential knowledge, data *cannot be acquired by the kind of non-deductive inference* at play in the geneticist example, and so the objections against $KK \subseteq E$ do not tell against $D \subseteq E$. But, *unlike* non-inferential knowledge, data exhibit an *intersubjective dimension* which makes it impossible that what is not a datum for the speaker becomes a datum for the hearer, and so the first objection against $K_{ni} \subseteq E$ does not tell against $D \subseteq E$. And, again *unlike* non-inferential knowledge, data exhibit a *freedom from risk* which makes it unlikely that they could be subject to a good enough preface-paradox-style example,[29] and so it is unlikely that the second objection against $K_{ni} \subseteq E$ tells against $D \subseteq E$ (*cf* fn 21).[30]

All these three features of data can be traced back to their more fundamental nature of being "*fresh information*", delivered by the world and still untainted by human epistemic activities. *Untainted* rather than *untouched*: for such information can often be *acquired* as such through e.g. perception, and be *transmitted* as such through e.g. testimony and inference, even if some other times all such activities, while still *properly executed*, nevertheless involve an element of *elaboration* that *degrades* what was a datum into a mere item of knowledge or of justified belief. It is this nature of data as fresh information that accounts not only for the three features mentioned in the last paragraph, but also, in accordance with the discussion of the functions of evidence at the beginning of this section, for the further feature of data of providing a suitable basis for knowledgeably accepting hypotheses adequately supported by them.[31] And, if data can do that while at the same time they cannot be generated by testimony or be transmitted by the kind

[28] Another proposal abandoning that idea would unhelpfully maintain that *all certainty is evidence* (assuming throughout the informal notion of certainty, which I don't take to be fully equivalent with the semi-technical notion of maximum epistemic probability). Such proposal would be unhelpful as it would appeal to a sufficient condition that is, in a certain respect, *not interestingly weak*—evidence is typically not certain. Worse, the proposal would not even succeed in specifying a (not *interestingly weak* but at least *correct*) sufficient condition for evidence. Firstly, to give just but one straightforward counterexample, it is certain for me that the Belenenses won't win the Liga (there's no way that can happen) but it is not part of my evidence that the Belenenses won't win the Liga. Secondly, it is not even clear that variations on the first objection against $KK \subseteq E$ and on the second objection against $K_{ni} \subseteq E$ would not apply against $C \subseteq E$ (see Zardini [2016b] for some discussion of the latter variation and fn 29 for some discussion of the relationships between certainty and data).

[29] Data's freedom from risk might suggest a connection with *certainty*. But, arguably, any such connection will be at best somewhat indirect, since, arguably, *neither are all data certain nor is all certainty a datum*. For example, it is part of my data that there is a house in front but it is not certain for me that there is a house in front (I can't conclusively rule out that what is in front is just a facade); conversely, it is certain for me that there is no pink elephant inside the moon (that would be incredible) but it is not part of my data that there is no pink elephant inside the moon (*cf* fn 28).

[30] Beyond the issues raised by the argument for $K \nsubseteq E$, notice that data do not include the last example presented in fn 18, and so the point in fn 26 against $K_{ni} \subseteq E$ does not tell against $D \subseteq E$. Notice also that, in addition to being arguably true, $D \subseteq E$ also appeals to a sufficient condition for evidence which is, in the respect focussed on in fn 27, interestingly weak—truckloads of data are acquired precisely through inference (as witnessed by the very same kind of example presented in that fn).

[31] Compare the suggestive definition of 'datum' in the *Oxford Dictionary of English*: "An assumption or premise from which inferences may be drawn".

of non-deductive inference at play in the geneticist example, (K_1) and (K_2) can both be true, and so, for all the argument for $K \nsubseteq E$ shows, all data may well be evidence.[32]

5 One Objection and Three Replies

A natural kind of objection to the argument for $K \nsubseteq E$ consists in equating it with certain apparently compelling *anti-closure* arguments, the idea being that, since such arguments *just have to fail in spite of their apparent compellingness*, and since the argument for $K \nsubseteq E$ is *structurally similar* to them, for essentially the same reason that argument also just has to fail in spite of its apparent compellingness.[33] More in detail, the argument for $K \nsubseteq E$ might be submitted to resemble in structure the preface-paradox-style argument discussed in section 4: it might be observed, roughly, that, just as that argument plays on *aggregating further risks with further premises*, so the argument for $K \nsubseteq E$ plays on *aggregating further risks with further inferences*. And, on the basis of this observation, it might be objected that, whatever turns out to be the right explanation of why, in spite of the apparent compellingness of the preface-paradox-style argument, the aggregation of risk induced by further premises does not eventually preclude knowledge in a preface-paradox-style example, essentially the same explanation will also show why, in spite of the apparent compellingness of the argument for $K \nsubseteq E$, the aggregation of risk induced by further inferences does not eventually preclude knowledge in the geneticist example.[34,35]

[32]The question naturally arises whether, conversely, *all evidence is a datum*. I think that that question should be answered in the negative. For one thing, it would seem that *evidence can be acquired through adequately double-checked sustained deductive reasoning that goes beyond the data* (for example, it would seem that it is part of my evidence that there are infinitely many prime numbers but it is not part of my data that there are infinitely many prime numbers). For another thing, it would seem that *adequately well-established theories that go beyond the data can serve as evidence for further theories* (for example, it would seem that it is part of my evidence that the theory of special relativity is true but it is not part of my data that the theory of special relativity is true). Thanks to Diogo Santos for questioning me about this converse inclusion.

[33]Thanks to Philip Ebert for pressing me on this issue.

[34]In its involving a *long chain* of *single-premise* inferences (rather than a *single multi-premise* inference *with many premises*), and, relatedly, in its involving risk generated by *inferences* (rather than risk generated by *premises*), the argument for $K \nsubseteq E$ resembles even more closely the less well-known *long-chain-style* argument purportedly refuting *single-premise* closure which can be extracted from Locke [1690], Book IV, Chapter II, 6; Hume [1739], Book I, Part IV, Section I (see e.g. Lasonen-Aarnio [2008], p. 171 for a contemporary reconstruction). However, as an anonymous referee reminded me, the opposite is the case in other respects: for example, in its involving a final conclusion that is *much riskier* than the initial premise(s) (rather than a final conclusion that is *at worst as risky as* the initial premise(s)), the argument for $K \nsubseteq E$ resembles more closely the preface-paradox-style argument. While I've chosen to conduct my discussion in terms of the preface-paradox-style argument as it is more familiar and better understood, I note that all the three points I'll be making apply *mutatis mutandis* to the relationships between the argument for $K \nsubseteq E$ and the long-chain-style argument.

[35]A related objection is that the argument for $K \nsubseteq E$ does actually *rely on* multi-premise closure. Although, as I've officially presented it, the argument does not mention multi-premise closure, I think that it is indeed naturally understood as envisaging a situation that involves a series of multi-premise inferences on the part of the geneticist (in putting together the intermediate conclusion that $\mathfrak{f}_i$ will have fur colour F_i with the information about e.g. the set-up of the example which is additionally needed to

In reply, firstly (and least importantly), I don't think we should uncritically assume that the preface-paradox-style argument "just has to fail", for, as I've argued in section 4, there is a very plausible weaker principle governing deductive inferential knowledge which allows us not to throw out the baby of the *knowledge-yielding power of deduction* with the bath water of *multi-premise closure*. And, as I've noted there, the bath water turns out crucially to include $K \subseteq E$ itself.

Secondly (and more importantly), henceforth assuming multi-premise closure, it is very plausible to expect that, whatever turns out to be the right explanation of why the aggregation of risk induced by further premises does not eventually preclude knowledge in a preface-paradox-style example, it will crucially appeal—to be sure, among other things—to the fact that, in any such example, it is *metaphysically necessary* that, given the truth of what one uncontroversially knows (i.e. the truth of the several premises of the relevant multi-premise argument), the conclusion of the relevant multi-premise argument is true (for one thing, Williamson [2009]'s own explanation does so appeal). But nothing like that is equally a feature of the geneticist example: it is *not even statistically likely*—let alone metaphysically necessary—that, given the truth of what the geneticist uncontroversially knows (e.g. the truth of the propositions describing the set-up of the geneticist example), the final conclusion of her inferences (i.e. that $\mathfrak{f}_{1{,}000{,}000}$ will have fur colour $F_{1{,}000{,}000}$) is true. Thus, since, contrary to a preface-paradox-style example, in the geneticist example *there is no statistically likely*—let alone metaphysically necessary—*truth preservation from the initial premises to the final conclusion*, it is very plausible to expect that, against what the objection presupposes, it is actually *not* the case that the right explanation of why the aggregation of risk induced by further premises does not eventually preclude knowledge in a preface-paradox-style example will essentially also show why the aggregation of risk induced by further inferences does not eventually preclude knowledge in the geneticist example.[36]

infer further that $\mathfrak{f}_{i+1}$ will have fur colour F_{i+1}, $1 \leq i < 1{,}000{,}000$). However, firstly, all such information is very plausibly understood (and can in any event be further stipulated) *not to add any further risk*, thereby pre-empting a preface-paradox-style argument. Secondly, the argument for $K \nsubseteq E$ can (less naturally but) equally compellingly be understood as envisaging a situation that only involves a series of *single-premise inferences* on the part of the geneticist (by inferring as intermediate conclusion the conjunction of the proposition that $\mathfrak{f}_i$ will have fur colour F_i with all the additionally needed information, $1 \leq i < 1{,}000{,}000$).

[36] Notice that the point in the text stands even under the assumption that both the preface-paradox-style argument and the argument for $K \nsubseteq E$ eventually need to appeal to a principle very roughly to the effect that *high objective chance of* $\neg P$ *is incompatible with knowledge of* P (the former argument in justifying the claim that one does not know the conclusion of the relevant multi-premise argument—as it is in fact done e.g. by Hawthorne and Lasonen-Aarnio [2009], to whom Williamson [2009] is replying—the latter argument in justifying (K_2)). For the point in the text is to the effect that, under that assumption, it is very plausible to expect that the right explanation of why high objective chance of $\neg P$ is compatible with knowledge of P in a preface-paradox-style example will crucially appeal to a feature that is completely absent in the geneticist example. Having noted this much, I should hasten to add that, in my view, issues of objective chance are quite a red herring both in the case of the preface-paradox-style argument and in the case of the argument for $K \nsubseteq E$, since, in my view, both arguments can *equally well* be run on examples where there is no issue that the propositions in question *only have extremal objective chances*. Thanks to an anonymous referee for discussion of these issues.

Thirdly (and most importantly), while, in a preface-paradox-style example, it would at least *not* seem ludicrous to think that one knows the conclusion of the relevant multi-premise argument, in the geneticist example it would seem *utterly ludicrous* to think that the geneticist knows that $\mathfrak{f}_{1,000,000}$ will have fur colour $F_{1,000,000}$. Knowledge has its limits, and this is surely one of them. It is thus extremely plausible to expect that *no good explanation of anything*—not even, against what the objection presupposes, the right explanation of why the aggregation of risk induced by further premises does not eventually preclude knowledge in a preface-paradox-style example—will essentially also vindicate such a ludicrous thought. If it did, the result would be very different from our present conception of knowledge.[37]

I thus don't think that its superficial similarity with certain anti-closure arguments detracts from the probative force of the argument for $K \nsubseteq E$. But let me close by summing up my findings in a more neutral, disjunctive fashion: either scepticism about scientific inference prevails, or we can now know what fur colour a rabbit that will be born in a 1,000,000 years will have, or $K \nsubseteq E$.

References

Andrew Bacon. Giving your knowledge half a chance. *Philosophical Studies*, 171:373–397, 2014.

Alexander Bird. Justified judging. *Philosophy and Phenomenological Research*, 74:81–110, 2007.

Marian David and Ted Warfield. Knowledge-closure and skepticism. In Quentin Smith, editor, *Epistemology: New Essays*, pages 137–187. Oxford University Press, Oxford, 2008.

Cian Dorr, Jeremy Goodman, and John Hawthorne. Knowing against the odds. *Philosophical Studies*, 170:277–287, 2014.

[37]Long chains provide extreme conditions under which to test theoretically interesting general principles. In addition to the argument for $K \nsubseteq E$ in this paper, a recent example is offered by Dorr et al. [2014], who use a long chain to show that the extremely plausible principle that, *if one knows that a coin is fair and does not know that it will not be flipped, one does not know that it will not land tails*, the assumption that one knows the set-up of a certain example and the assumption that one knows that 1,000,000 fair coins will not all land tails lead to a sorites-like clash. Somewhat similarly to how the argument for $K \nsubseteq E$ crucially relies on the anti-sceptical principle (K_1), Dorr et al.'s argument crucially relies on the latter assumption, whose rejection, they argue, leads to a wide-ranging scepticism. (I should note though that (K_1) strikes me as being on much firmer grounds than that assumption; I've given my own take on Dorr et al.'s argument in Zardini [2016b] in the context of the general approach mentioned in fn 9.) Something along the lines of that assumption is also taken as a starting point by Bacon [2014] for an investigation that, somewhat similarly to the one in this paper, leads to question the compatibility of $K \subseteq E$ with some theoretically central functions of evidence (not the one though that this paper focusses on), with a somewhat similar conclusion that *the only states that provide evidence should be states that are more recognisably concerned with information gathering than any old state of knowledge is.* Thanks to an anonymous referee for alerting me to the relevance of these works.

Alvin Goldman. Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73:771–791, 1976.

John Hawthorne and Maria Lasonen-Aarnio. Knowledge and objective chance. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*, pages 92–108. Oxford University Press, Oxford, 2009.

David Hume. *A Treatise of Human Nature*. John Noon, London, 1739.

David Hume. *Philosophical Essays Concerning Human Understanding*. Andrew Millar, London, 1748.

Maria Lasonen-Aarnio. Single premise deduction and risk. *Philosophical Studies*, 141: 157–173, 2008.

John Locke. *An Essay Concerning Humane Understanding*. Thomas Basset, London, 1690.

David Makinson. The paradox of the preface. *Analysis*, 25:205–207, 1965.

Thomas Reid. *An Inquiry into the Human Mind*. Andrew Millar, London, 1764.

Timothy Williamson. Inexact knowledge. *Mind*, 101:217–242, 1992.

Timothy Williamson. Does assertibility satisfy the S4 axiom? *Crítica*, 27:3–22, 1995.

Timothy Williamson. Cognitive homelessness. *The Journal of Philosophy*, 93:554–573, 1996.

Timothy Williamson. Knowledge as evidence. *Mind*, 106:717–741, 1997.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000.

Timothy Williamson. Reply to John Hawthorne and Maria Lasonen-Aarnio. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*. Oxford University Press, 2009.

Elia Zardini. A model of tolerance. *Studia Logica*, 90:337–368, 2008a.

Elia Zardini. *Living on the Slippery Slope. The Nature, Sources and Logic of Vagueness.* PhD thesis, Department of Logic and Metaphysics, University of St Andrews, 2008b.

Elia Zardini. Luminosity and vagueness. *Dialectica*, 66:375–410, 2012.

Elia Zardini. Higher-order sorites paradox. *Journal of Philosophical Logic*, 42:25–48, 2013a.

Elia Zardini. Luminosity and determinacy. *Philosophical Studies*, 165:765–786, 2013b.

Elia Zardini. Confirming the less likely, discovering the unknown. Dogmatisms: surd and doubly surd, natural, flat and sharp. In Dylan Dodd and Elia Zardini, editors, *Scepticism and Perceptual Justification*, pages 33–70. Oxford University Press, Oxford, 2014a.

Elia Zardini. Possibility, necessity and probability. A meditation on underdetermination and justification. *Erkenntnis*, 79:639–667, 2014b.

Elia Zardini. Breaking the chains. Following-from and transitivity. In Colin Caret and Ole Hjortland, editors, *Foundations of Logical Consequence*, pages 221–275. Oxford University Press, Oxford, 2015.

Elia Zardini. No state *a priori* known to be factive is mental. *Philosophical Issues*, 2016a. Forthcoming.

Elia Zardini. Closed without boundaries. ms, 2016b.

Elia Zardini. Open knowledge of one's inexact knowledge. ms, 2016c.