Philos Stud (2019) 176:1789–1806
https://doi.org/10.1007/s11098-018-1096-5

Higher order ignorance inside the margins

Sam Carter[1]

Published online: 26 May 2018

Abstract According to the KK-principle, knowledge iterates freely. It has been argued, notably in Greco (J Philos 111:169–197, 2014a), that accounts of knowledge which involve essential appeal to normality are particularly conducive to defence of the KK-principle. The present article evaluates the prospects for employing normality in this role. First, it is argued that the defence of the KK-principle depends upon an implausible assumption about the logical principles governing iterated normality claims. Once this assumption is dropped, counter-instances to the principle can be expected to arise. Second, it is argued that even if the assumption is maintained, there are other logical properties of normality which can be expected to lead to failures of KK. Such failures are noteworthy, since they do not depend on either a margins-for-error principle or safety condition of the kinds Williamson (Mind 101:217–242, 1992; Knowledge and its limits, OUP, Oxford, 2000) appeals to in motivating rejection KK. "Introduction: KK and Being in a Position to Know" Section formulates two versions of the KK-Principle; "Inexact Knowledge and Margins for Error" Section presents a version of Williamson's margins-for-error argument against it; "Knowledge and Normality" and "Iterated Normality" Sections discuss the defence of the KK-Principle due to Greco (J Philos 111:169–197, 2014a) and show that it is dependent upon the implausible assumption that the logic of normality ascriptions is at least as strong as K4; finally, "Knowledge in Abnormal Conditions" and "Higher-Order Ignorance Inside the Margins" Sections argue that a weakened version of Greco's constraint on knowledge is plausible and demonstrate that this weakened constraint will, given uncontentious assumptions, systematically generate counter-instances to the KK-principle of a novel kind.

✉ Sam Carter
samjbcarter@gmail.com

[1] Rutgers University, 106 Somerset St, New Brunswick, NJ 08901, USA

Keywords Higher-order knowledge · KK · Positive introspection · Iterated knowledge · Margins for error · Normality

1 Introduction: KK and being in a position to know

Call the thesis, imprecisely stated, that knowledge iterates freely the KK-Principle. (KK) is the simplest formulation of the KK-Principle.

(KK) (S knows that *p*)⊃(S knows that (S knows that *p*)).

(KK) is widely taken to be untenable. It is standardly assumed that knowing that *p* entails believing that *p* (though see Radford (1966) and Myers-Schulz and Schwitzgebel (2013)). If so, then since, plausibly, S can know that *p* without believing that S knows that *p* (though see Greco (2014a, 174, fn17, 2014b)), (KK) will admit counter-instances. Similarly, if S must possess the concept of *F* to know that $F(a_i, a_j \ldots)$, then S may fail to know that S knows that *p* despite knowing that *p* due to failing to possess the concept of knowledge (Feldman (1981), Das and Salow (2016); though see Stalnaker (1999)). As a result, much discussion of the KK-Principle has instead focused on weaker variants of (KK), such as (KK^{-}).[1]

(KK^{-}) (S knows that *p*)⊃(S is in a position to know that (S knows that *p*)).

(KK^{-}) employs the notion of being in a position to know. No full analysis of this state is necessary for present purposes. However, we can assume, minimally, that being in a position to know that *p* is factive (i.e., that S is in a position to know that *p* implies that *p*) and strictly weaker than knowledge that *p* (i.e., that S knows that *p* implies that S is in a position to know that *p* but that S is in a position to know that *p* does not imply that S knows that *p*).

It has been argued, notably in Greco (2014a), that accounts of knowledge which involve essential appeal to normality are particularly conducive to defence of the KK-principle. This paper investigates the plausibility of that claim. It is argued, first, that Greco's defence of the KK-principle depends upon an implausible assumption about the logical principles governing iterated normality claims. Once this assumption is dropped, counter-instances to the principle can be expected to arise. Second, it is argued that even if the assumption is maintained, there are other logical properties of normality which can be expected to lead to failures of KK.

Section 2 presents a version of Williamson's margins-for-error argument against KK; Sects. 3, 4 discuss the defence of the KK-Principle due to Greco (2014a) and show that it is dependent upon the implausible assumption that the logic of normality ascriptions is at least as strong as K4; finally, Sects. 5, 6 argue that a weakened version of Greco's constraint on knowledge is plausible and demonstrate that this weakened constraint will, given uncontentious assumptions, systematically generate counter-instances to the KK-principle of a novel kind.

[1] For the purposes of the present paper, the KK-Principle can be treated as the disjunction of (KK) and (KK^{-}).

2 Inexact knowledge and margins for error

Sometimes our knowledge of the world is inexact. Say that S has inexact knowledge of the value of a parameter x iff, for any value taken by x, there is some non-trivial constant size of interval centred on that value such that S knows that the value of x falls within the interval, but there is no subset of that interval such that S is in a position to know that x falls in that subset.

Inexact knowledge can arise in a variety of ways. Our information about some physical magnitude may depend upon a measuring device which issues reports with a non-trivial degree of inaccuracy. Alternatively, our best available judgement about a state of affairs may involve estimation, as when evaluating the size of crowd, or the number of coffee beans in a jar. Williamson (1992, 2000) has argued that the existence of known inexact knowledge is incompatible with (KK) (and, more generally, any acceptable precifisication of the KK-Principle). Call this the MARGINS-FOR-ERROR ARGUMENT.

Suppose that S has inexact knowledge of the temperature in a room. Whatever the value of the temperature, there is a non-trivial range such that S knows that the temperature falls within that range of its actual value. However, she is not in a position to know, of any value within that range, that it is not the value of the temperature in the room. Suppose that this range is ±5 °F. Then for all n, if the temperature in the room is n °F, the strongest proposition S knows is that the temperature is between $(n-5)$ °F and $(n+5)$ °F. It follows, assuming closure, that if the temperature is n °F, then, for any $n' \leq 5$, S does not know that the temperature is not $(n+n')$ °F.

Assume that, by reflecting on the imperfections of her means of coming to know the temperature, S knows, for some $0 < i \leq 5$, that if the temperature is n °F, then she does not know that the temperature is not $(n+i)$ °F. That is, she knows that any knowledge she has of the temperature is inexact.[2] Likewise assume that, due to her diligence, her knowledge is closed under entailment. While these assumptions may be idealisations for normal agents under normal conditions, they are harmless ones in the present context.

(KK) is inconsistent with S's knowledge of her inexact knowledge of the temperature. Suppose that the temperature is 70 °F. Then S knows that the temperature is between 65 °F and 75 °F. Accordingly, by closure, she knows that the temperature is not 80 °F. By (KK), she knows that she knows that the temperature is not 80 °F. Furthermore, by contraposition of the above, S knows, for some $0 < i \leq 5$, that if she knows that the temperature is not $(n+i)$ °F, then the temperature is not n °F. Hence, by closure, S knows that the temperature is not $(80-i)$ °F. By a second application of (KK), S knows that S knows that the temperature is not $(80-i)$ °F. Thus, by a second application of closure, S knows that the temperature is not $(80-2i)$ °F… Repetition of this procedure will yield the conclusion that S knows that the temperature is not 70 °F. Yet by factivity, this is inconsistent with the original

[2] Even if she is not in a position to know precisely how inexact her knowledge is, it is plausible that there is some $i > 0$ such that she is in a position to know that she is not in a position to know the temperature to within i °F.

supposition that S does not have exact knowledge of the temperature. Accordingly, Williamson concludes that we must reject (KK) in its full generality.

The margins-for-error argument only directly supports rejection of (KK). However, Williamson suggests (2000, 115), it can be extended into an argument against (KK^{-}) with the additional assumption that if S is in a position to know that *p* then S knows that *p*. That is, S acquires all knowledge she is in a position to acquire (making (KK^{-}) and (KK) equivalent). Hence, if the margins-for-error argument succeeds in demonstrating the existence of counter-examples to (KK), it also succeeds in demonstrating the existence of counter-examples to (KK^{-}) in those cases in which this further idealisation holds.

3 Knowledge and normality

The conclusion of the margins-for-error argument has been claimed (by e.g., Hawthorne and Magidor (2009), Stalnaker (2009) and Greco (2014a, b)) to be incompatible with the success of putative explanations which make essential appeal to the KK-principle in pragmatics, computer science, game theory and other areas. Insofar as one takes these concerns to be well-founded and yet seeks to preserve such explanations, there is reason to investigate alternative accounts of knowledge which validate some version of the KK-principle.

A number of defences of the KK-principle do so by appealing to an account of knowledge stated in terms of normality. Versions of this strategy can be found in Dretske (1981), Greco (2014a), Stalnaker (2015) and Goodman and Salow (2018). As the most explicit such defence of the KK-principle, I will focus on Greco's version of the normality theory (though the concerns generalise to the other accounts). Where C is a schematic variable ranging over agents' total cognitive states, Greco's normality theory can be formulated as (NTK):[3]

(NTK) S knows, in C, that *p* iff (*i.*) Normally ((S is in C)$\supset$*p*); and
(*ii.*) Conditions are normal.

[3] In fact, Greco includes a third conjunct in the RH-clause of the biconditional:

(NTK.*iii*) Being in C causes or constitutes S believing that *p*.

However, he proposes that this conjunct can be dropped under the idealising assumption that if normally, being in C entails *p* then normally being in C causes S to believe that *p* (2014a, 184) (similar qualifications can be found in Dretske (1981) and Stalnaker (2015)). This idealisation can be treated as analogous to Williamson's idealising assumption in the extension of the margins-for-error argument against (KK^{-}), that S knows that *p* if S is in a position to know that *p*. Satisfying (NTK.*i-ii*) can reasonably be treated as necessary and sufficient for being in a position to know under Greco's theory. Thus, dropping the idealising assumptions, Greco's argument in fact only establishes the following version of the KK-principle (which is neither stronger nor weaker than (KK)):

$(KK^{\sim})$ (S is in a position to know that *p*)$\supset$(S is in a position to know that (S is in a position to know that *p*)).

For present purposes, an agent's cognitive state can be thought of as the minimal state of affairs comprising all and only those states of the agent to which her acquisition of knowledge is sensitive. For example, it may include, but need not be limited to, facts about her evidence, the concepts in her possession, the methods of belief formation available to her, &c. Construed in this way, (NTK) will not offer a non-circular analysis of knowledge, but might nevertheless be hoped to generate informative predictions regarding, e.g., structural properties of knowledge.

Like Dretske (1981) and Stalnaker (2015), Greco motivates (NTK) by appeal to the notion of a state 'carrying' information. Under a popular account of this notion, a system carries the information that p in a given state iff normally, if the system is in that state, p is true (see, e.g., Stampe (1977), Dretske (1981), Millikan (1984) and Stalnaker (1999), a.o.)[4] On such accounts, (NTK.*i*) (as a necessary condition on knowledge) follows directly from the bridge principle that S knows that p only if her cognitive state carries the information that p. (NTK) can then be restated as the theory that S knows that p iff S is in a cognitive state which carries the information that p and conditions are normal.

Greco argues that (NTK) entails (KK). We can consider an informal instance of this argument, before stating it precisely in more general terms. Suppose that S knows that the temperature is between 65 and 75 °F. Then, by the L-to-R direction of (NTK): (a) conditions are normal and (b) S is in a state C such that normally, if S is in C then the temperature is between 65 and 75 °F. Next, by the R-to-L direction of (NTK), it is sufficient for S to know that S knows that the temperature is between 65 and 75 °F that: (a′) conditions are normal; and (b′) S is in some state C* such that normally, when S is in C*, (a) and (b) obtain. Yet, Greco claims, C itself is such a state. That is, he claims it follows from (b) that normally, if S is in C, (b) obtains. Thus, (b) entails (b′), under the assumption that, normally, conditions are normal. Since (a) = (a′), (a) trivially entails (a′). If this is correct, then S's satisfaction of the necessary conditions for knowledge that p is sufficient for S's satisfaction of the sufficient conditions for knowledge that S knows that p. If S knows that p then S knows that S knows that p. Accordingly, Greco concludes, (NTK) entails (KK).

Formulating the argument in greater generality, let $\blacksquare\varphi$ be the proposition that normally φ and $\square\varphi$ the proposition that necessarily φ. We assume that $\blacksquare$ is weaker than $\square$; that is, $\vDash\square\varphi\supset\blacksquare\varphi$. Let C_S be the proposition that S is in C.[5] Let $K_S p$ be the proposition that S knows that p. $(\forall p\ (p\supset\neg\blacksquare\neg p))$ is the proposition that conditions are normal. In a Kripke semantics, if $N(w)$ is the set of worlds normal at w, $w\vDash(\forall p\ (p\supset\neg\blacksquare\neg p))$ iff $w\in N(w)$. Greco's argument assumes the principle (SR)—that is, the principle that normally, conditions are normal.

(SR) $\blacksquare(\forall p\ (p\supset\neg\blacksquare\neg p))$.[6]

[4] Note that, on certain formulations, the information carried by a system in a given state is said to be the *strongest* such proposition. On this variant, the relevant bridge principle will be that S knows, when in C, that p only if the information carried by C *entails* that p.

[5] We assume that S is in exactly one cognitive state at each world; that is, $\lambda C.C_S$ is a partition of the subset of modal space in which S exists.

[6] (SR) corresponds to the assumption that the accessibility relation for $\blacksquare$ is *shift reflexive*: if $w'\in N(w)$, then $w'\in N(w')$.

Finally, we assume the logic for normality is a normal modal logic. We can restate (NTK) as follows:

(NTK) $\Box(C_s \supset (K_S p \equiv (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))))$.

Assuming C_S and $K_S p$, Greco argues $K_S(K_S p)$ can be derived as follows:

(P_1)	$K_S p$	(*premise*)
(P_2)	$(K_S p \equiv (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p)))$	(*from* C_S, (NTK), $\vDash \Box\varphi \supset \varphi$, *by* MP)
(P_3)	$\blacksquare(C_S \supset p)$	(*from* (P_1),(P_2), *by* MP, $\wedge$E)
(P_4)	$\forall p\ (p \supset \neg\blacksquare\neg p)$	(*from* (P_1),(P_2), *by* MP, $\wedge$E)
(P_5)	$K_S(K_S p) \equiv \blacksquare(C_S \supset (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))) \wedge \forall p\ (p \supset \neg\blacksquare\neg p)$	(*from* (P_2), (NTK))[7]
(P_6)	$K_S(K_S p) \equiv \blacksquare(C_S \supset \blacksquare(C_S \supset p))$	(*from* (P_4),(P_5),(SR))
(P_7)	$\blacksquare(C_S \supset \blacksquare(C_S \supset p))$	(*from* (P_3), *by* $\blacksquare(\varphi \supset \psi) \vDash \blacksquare(\varphi \supset \blacksquare(\varphi \supset \psi))$)
(C)	$K_S(K_S p)$	(*from* (P_6),(P_7), *by* MP)

Substituting $K_S p/p$ in (P_2) yields (P_5). From the assumption that $K_S p$ and (SR), we know that $\blacksquare\forall p\ (p \supset \neg\blacksquare\neg p)$ and $\forall p\ (p \supset \neg\blacksquare\neg p)$. Thus, ($P_5$) simplifies to ($P_6$) by $\wedge$-elimination and: $\blacksquare(\varphi \supset (\psi \wedge \chi))$, $\blacksquare\chi \vDash \blacksquare(\varphi \supset \psi)$. Yet, Greco claims, $\blacksquare(\varphi \supset \psi) \vDash \blacksquare(\varphi \supset \blacksquare(\varphi \supset \psi))$. Thus, from ($P_3$), it follows that $\blacksquare(C_S \supset \blacksquare(C_S \supset p))$. So, by the R-to-L direction of (P_6), $K_S(K_S p)$. Having derived $K_S(K_S p)$ from $K_S p$ (and auxiliary assumptions), Greco concludes (NTK) entails the KK-principle.

4 Iterated normality

The crucial step in Greco's argument is the derivation of (P_7) from (P_3). This inference depends upon the validity of the schema (CI) (for 'Crucial Inference'):

(CI) $\blacksquare(\varphi \supset \psi) \vDash \blacksquare(\varphi \supset \blacksquare(\varphi \supset \psi))$

In the informal statement of the argument in Sect. 3, (CI) was required infer from (b) (i.e., the claim that normally, if S is in C then the temperature is between 65 and

7 Proof:

($P^\dagger_1$)	$K_S(K_S p) \equiv (\blacksquare(C_S \supset K_S p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))$	(*from* (P_2), $p/K_S p$)
($P^\dagger_2$)	$\Box(C_s \supset (K_S p \equiv (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))))$	(NTK)
($P^\dagger_3$)	$\blacksquare(C_s \supset (K_S p \equiv (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))))$	(*from* $\Box\varphi \vDash \blacksquare\varphi$)
($P^\dagger_4$)	$K_S(K_S p) \equiv \blacksquare(C_S \supset (\blacksquare(C_S \supset p) \wedge \forall p\ (p \supset \neg\blacksquare\neg p))) \wedge \forall p\ (p \supset \neg\blacksquare\neg p)$	(*from* ($P^\dagger_1$), ($P^\dagger_3$))

75 °F), that normally, if S is in C, then (b) obtains. (CI), in turn, entails the axiom schema 4 of modal logic.[8]

(4) $\blacksquare\varphi \vDash \blacksquare\blacksquare\varphi$.

(4) says that what is normal is normally normal—normality iterates freely. This corresponds to the constraint that N be transitive; for any w, if $w' \in N(w)$, then $N(w') \subseteq N(w)$. In order to evaluate the acceptability of Greco's argument then, we need to assess the plausibility of (4) for $\blacksquare$.

Most extant logics of normality either validate the schema (Smith (2008) and Boutilier (1994)), do not permit iterated normality statements (Veltman (1996)) or have trivial logics of iterated normality (Delgrande (1987) and Asher and Morreau (1995)). Nevertheless, cases such as the following suggest there is good reason to reject it.

Suppose that Jaime has a relatively regular schedule; normally, she arrives home from work between 5.00 pm and 6.00 pm, though it would not be abnormal for her to arrive home at any time between 5.00 pm and 6.00 pm. Furthermore, she normally does not come close to arriving home abnormally early or abnormally late. For any time, normally, if she arrived home at that time, it would not have been abnormal for her to have arrived up to 5 min earlier or up to 5 min later.

Jaime's situation appears perfectly coherent. Yet it involves a violation of the principle that normality iterates freely. Since it would not be abnormal for her to have arrived home at any time between 5.00 pm and 6.00 pm, there are normal situations in which she arrives home at 1 min before 6 pm. Yet, normally, whenever she arrives home, it would not have been abnormal for her to arrive up to 5 min later. Thus, in a normal situation in which she arrives home 1 min before 6 pm, it would be normal for her to have arrived home 1 min after 6.00 pm. Thus, despite the fact that she does not normally arrive 6.01 pm, it is not the case that normally, she does not normally arrive at 6.01 pm.

We can model Jaime's situation as follows: Let $w_{n.nn}$ be the world at which Jaime arrives at $n.nn$pm. Then, suppose that $w_{5.30}$ is a world at which (a)–(c) obtain.

(a) Normally, Jaime arrives between 5.00 pm and 6.00 pm;
(b) For any time between 5.00 pm and 6.00 pm, it would not be abnormal for Jaime to arrive at that time;

[8] Proof:

$(P^{\ddagger}_1)$	$\blacksquare(\top\supset\varphi)\supset\blacksquare(\top\supset\blacksquare(\top\supset\varphi))$	(*from* (CI))
$(P^{\ddagger}_2)$	$\blacksquare(\varphi)\supset\blacksquare(\top\supset\blacksquare(\varphi))$	(*from* $(P^{\ddagger}_1)$, $\varphi \dashv\vDash \top\supset\varphi$)
$(P^{\ddagger}_3)$	$\blacksquare(\varphi)\supset(\blacksquare\top\supset\blacksquare\blacksquare(\varphi))$	(*from* $(P^{\ddagger}_2)$, K)
$(P^{\ddagger}_4)$	$(\blacksquare(\varphi)\wedge\blacksquare\top)\supset\blacksquare\blacksquare(\varphi))$	(*from* $(P^{\ddagger}_3)$, $((\varphi\wedge\psi)\supset\chi)\dashv\vDash(\varphi\supset(\psi\supset\chi)))$
$(P^{\ddagger}_5)$	$\blacksquare\top$	(*from* $\blacksquare$Nec)
$(P^{\ddagger}_6)$	$\blacksquare(\varphi)\supset\blacksquare\blacksquare(\varphi)$	(*from* $(P^{\ddagger}_4)$,$(P^{\ddagger}_5)$, $\{(\varphi\wedge\psi)\supset\chi), \varphi\}\vDash \psi\supset\chi$)

(c) Normally, Jaime does not come within 5 min of arriving abnormally early or abnormally late.[9]

By (b), $w_{5.59}$ is accessible from $w_{5.30}$ (since otherwise, it would be abnormal for Jaime to arrive at 5.59 pm). Yet, since $w_{5.59}$ is normal relative to $w_{5.30}$, by (c), $w_{6.01}$ is accessible from $w_{5.59}$ (since otherwise, Jaime would not normally not come within 5 min of arriving abnormally late). Yet, by (a), $w_{6.01}$ is inaccessible from $w_{5.30}$. Thus, in order for (a)–(c) to be satisfied at $w_{5.30}$, the relevant accessibility relation must be non-transitive.

On the assumption that what is normal is a contingent matter, failures of normality to iterate appear unsurprising. Plausibly, facts about what is normal depend on other, non-modal facts (e.g., statistical facts, facts about evolutionary history, &c.). Furthermore, it is plausible that the facts on which facts about normality depend can vary, both across normal worlds and between the actual world and worlds which are actually normal. For example, facts about when Jaime normally arrives home can be expected to depend (in part) on facts about when she in fact arrives home. If, over a sufficiently long period of time (and holding all else fixed), Jaime arrives home later in w' than she does in w, then the latest she would normally arrive home in w' can be expected to be (marginally) later than the latest she would normally arrive home in w. A little lateness can make lateness the new normal. Yet, if Jaime's arrivals in w' would not have be abnormal in w, then, in w, there will be some times at which it would be abnormal for Jaime to arrive home, but which would be normal in normal conditions.

As long as failures of (4) are permitted, we can construct counterexamples to (KK) while assuming (NTK). Suppose that, at the actual world, S is in a cognitive state C such that normally, if S is in C then the temperature is between 65 and 75 °F and that normally, if S is in C then it would not be abnormal for the temperature to be 1 degree cooler or hotter. Suppose, additionally, that conditions are normal (meaning that the temperature must actually be between 66 and 74 °F). By (NTK), S knows that the temperature is between 65 and 75 °F. Nevertheless, there is some normal world at which S is in C, the temperature is 65 °F and at which it is not abnormal for S to be in C and for the temperature to be 64 °F. Hence, by (NTK), at that world, S does not know that the temperature is between 65 and 75 °F. Thus, it is not the case that in the actual world, at all normal worlds in which S is in C, S knows that the temperature is between 65 °F and 75 °F. So, by (NTK), at the actual world, S knows that the temperature is between 65 °F and 75 °F but does not know that she knows that the temperature is between 65 °F and 75 °F.[10]

[9] It may be that the iterated reading of (c) is dispreferred, in favor of the non-iterated reading (on which both normality claims are evaluated with respect to the world of utterance). This is, it should be clear, not the reading on which the case is probative regarding the plausibility of 4 for $\blacksquare$.

[10] Goodman and Salow (2018) propose an alternative defense of (KK) via appeal to a normality-based account of knowledge. Their models include a single ordering of worlds for comparative normality, which results in KD45 for the logic of $\blacksquare$. If their models are expanded to allow for contingency in facts about comparative normality (i.e., by introducing distinct orderings for each world), then if the resulting logic of maximal normality (i.e. $\blacksquare$) is weaker than K4, counter-instances to (KK) are predicted.

5 Knowledge in abnormal conditions

Section 4 argued that (4) was too strong a constraint on the logic of normality. Some abnormal worlds are not normally abnormal. As a result, it was argued, counter-instances to (KK) should be expected under (NTK). Yet, even in the absence of failures of (4), there is reason to think that normality-based accounts of knowledge are less conducive to a defence of the KK-principle than has generally been thought. In the remainder of the paper, we will turn our attention to a less contentious logical property of normality and show how it also creates problems for the free iteration of knowledge.

Conditions are sometimes abnormal. The logic of normality must, as such, be weaker than KT; some worlds do not count as normal by their own standards. This corresponds to the constraint that N be non-reflexive; for some w, $w \notin N(w)$. Sections 5, 6 argue that regardless of whether (4) is accepted, failures of the T axiom schema[11] make (NTK) implausibly strong, and, furthermore, that such failures can be expected to yield counter-instances to (KK) on the theory's most natural weakening.

(NTK.*ii*) entails that if conditions are abnormal, then, no matter what cognitive state S is in, there is no *p* such that she knows that *p*. Any actual abnormality implies actual global scepticism for Greco (though not necessary global scepticism). Since the actual world is in fact admirably abnormal in numerous ways, this is an especially serious concern. A number of revisions to (NTK) are available which would allow the theory to avoid this problem. However, each revision faces serious problems of its own.

One response would be to index the relevant normality claims in (NTK) to the cognitive state of the individual. Let $\blacksquare^{\Psi}\varphi$ be the proposition that φ is true in all worlds normal with respect to ψ. The current proposal is that each instance of $\blacksquare$ is replaced with $\blacksquare^{C_S}$.[12] That is, S knows that *p* when in C iff in all worlds normal with respect to S being in C, if S is in C, then *p*, and conditions are normal with respect to S being in C. Clearly, under this revision, S can know that *p* when in C despite conditions being abnormal, as long as conditions are nevertheless normal with respect to S being in C. A version of this strategy is considered in passing by Greco (2014a, 181, fn35), possibly motivated by this concern. Stalnaker's proposal (2006, 2015) that normality be indexed to particular sources of information can also be seen as a variant of the same strategy.

However, while this will avoid the immediate problem that *any* actual abnormality implies global scepticism, it remains too strong. Abnormal conditions can be beneficial to the acquisition of knowledge. Suppose that, while in C, S finds herself in an environment which is more epistemically hospitable than is normal (with-respect-to-C_S). Plausibly, this will be to her epistemic advantage—acquisition of knowledge should, if anything, be easier. Yet, on the proposed revision, any

[11] i.e., $\blacksquare\varphi \vDash \varphi$.

[12] The argument against (4) for $\blacksquare$ can be given, *mutatis mutandis*, for $\blacksquare^{C_S}$.

abnormality (with-respect-to-C_S) is incompatible with S possessing knowledge while in C.

For example, imagine an animal which has evolved an acute sense of smell in order to navigate the subterranean environment it normally inhabits. Under normal conditions, below ground, its olfactory system is subject to a non-negligible degree of noise. When the same system is employed above ground (i.e., in abnormal conditions with respect to cognitive states produced by the system) it is subject to a lesser degree of noise (though the states it produces are of the same kind). Intuitively, if the animal can acquire knowledge via its olfactory system below ground, it can acquire knowledge via the same system above ground at least as easily. Yet, since conditions above ground are abnormal with respect to the cognitive state it is in (since normally, when in that state, it is below ground), any acquisition of knowledge in such conditions is predicted to be impossible.

Or, consider a measuring device designed for use in environments with a high level of interference (e.g., a radio receiver designed for use in areas of high electromagnetic interference). Under normal conditions for its use, it generates inexact information about the parameter it measures; reports on the value of the parameter are subject to a risk of error; however, there will be (abnormal) conditions of lower interference in which its reports are subject to a less substantial risk of error. Intuitively, it should be easier to acquire knowledge of the value of the relevant physical parameter in the latter case than in the former. Yet this is not what is predicted under the proposed revision.

Note that it need not be claimed that abnormally hospitable epistemic conditions are always knowledge-conducive. If a system is normally wholly unreliable, it may never be capable of generating knowledge, even if its present behaviour in fact abnormally reliable. That is, there may be a difference in kind between a normally accurate system functioning with abnormally high accuracy, and a normally inaccurate system functioning with the same level of abnormally high accuracy.

Rather than indexing, the proponent of (NTK) might alternatively appeal to the apparent context-sensitivity of normality ascriptions.[13] As Dretske (1981) and Stalnaker (2015) observe, standards of normality appear capable of variation across contexts. For example, the range within which the acidity of a cup of coffee must fall to qualify as normal for Colombian medium roast will plausibly vary depending on whether the assessor is a distributor selecting beans for a high-end blend or a non-connoisseur evaluating whether her stock of beans has expired.

Dretske appeals to the context-sensitivity of normality ascriptions to explain how knowledge can be attained by agents whose cognitive states are compatible with sceptical alternatives. S may know that *p* when in C, despite being unable to rule out *p* scenarios on the basis of C, as long as, given the contextually determined standards of normality, all $C_S \wedge p$-worlds are abnormal. For S to know that *p* when in C, the standards of normality must be sufficiently high that they do not categorise any $C_S \wedge p$-world as normal.

[13] Greco (p.c.) is sympathetic to a version of this response.

In contrast, the present sceptical concern does not arise from an inability to rule out sceptical alternatives. According to (NTK), knowledge that *p* requires the world of evaluation to be normal. Thus, for S to know that *p*, the standards of normality must be sufficiently low that they do not categorise the world at which S is located as abnormal.

The problem is that these two anti-sceptical requirements impose conflicting demands on context. The elimination of sceptical alternatives requires the standards determined by context to be sufficiently high for sceptical scenarios to qualify as abnormal. In contrast, the satisfaction of (NTK.*ii*) requires the standards determined by context to be sufficiently low for the world of evaluation to qualify as normal.

In any world which is itself less normal than the most normal sceptical alternative worlds, these two conditions will be incapable of being jointly satisfied. Yet, plausibly, where the kind of abnormality exhibited by the world of evaluation is conducive to true belief formation (as discussed above), a degree of abnormality at that world which is (arbitrarily) higher than the abnormality of the most normal sceptical alternative world is consistent with knowledge.

Eliminating (NTK.*ii*) would leave (NTK.*i*) as a sufficient condition on knowledge that *p* (conditional on C_S). (NTK.*i*) is clearly too weak for this role, since it does not even entail factivity. Since conditions are sometimes abnormal, (NTK.*i*) can be satisfied despite *p* being false.[14] In summary, (NTK.*ii*) is too strong as a necessary condition. Yet without (NTK.*ii*), (NTK.*i*) is too weak as a sufficient condition by itself.

However, the prospects of identifying an interesting principle relating normality and knowledge should not be dismissed outright. In particular, adopting (NTK.*i*) as a merely necessary condition on knowledge avoids each of these issues. Call the resulting thesis the Weak Normality Condition on knowledge:

(WNC) S knows, in C, that *p* only if Normally($C_S \supset p$).

(WNC) states that S knows, in C, that *p* only if in all normal worlds in which S is in C, *p* is true. That is, S does not know that *p* if it would not be abnormal for S to be in the state she is in and *p* be false. (WNC) is asymmetrically entailed by (NTK). However, it avoids each of the above problems. For all that has been said thus far, it is a potential necessary condition on knowledge. Indeed, there are at least some theoretical and abductive reasons in favour of adopting (WNC).

Most significantly, (WNC) is the weakest condition on knowledge compatible with the requirement that S knows that *p* only if her cognitive state carries the information that *p*. If a system carries the information that *p* iff it is in a state which normally implies *p*, and S knows that *p* only if she is in state which carries the

[14] The obvious, if ad hoc, solution would be to adopt the conjunction of (NTK.*i*) and *p* as a sufficient condition instead. Yet, while accounting for factivity, this response still cannot accommodate the apparent incompatibility of knowledge with gettierization. In many cases in which S has a gettierized belief that *p*, S will be in a cognitive state C such that, normally, if S is in C, then *p*. The proponent of the unrevised version of (NTK) can at least suggest that in such cases, conditions are abnormal. Hence, S's failure to know that *p* is to be explained as a result of the failure of (NTK.*ii*). Yet, if (NTK.*ii*) is replaced with *p*, then we appear committed to ascribing knowledge to S in such cases, contrary to the gettier intuition (since *ex hypothesi*, (NTK.*i*) is satisfied and if S has a gettierized belief that *p* then *p*).

information that *p*, then S knows that *p* only if she is in a state which normally implies *p*. To the extent this account of carrying information supports (NTK), it also supports (WNC). Indeed, (WNC) follows on the weaker condition that a system in given state carries the information that *p* only if it is in a state which normally implies *p*. This appears *prima facie* reasonable. If it wouldn't be abnormal for the system to be in its present state and *p* be false, then the information its present state carries cannot include *p*.

(WNC) also derives abductive support from its ability to explain judgements about certain minimal pairs. Nelkin (2000) imagines a computer screen whose background color is selected at random each time it is switched on. On any given occasion, there is a .99999999 chance that the color selected is blue, and a .00000001 chance that the color selected is red. Without having seen the color of the screen, Nelkin suggests, it is not possible for an agent to come to know that it is blue on the basis of probability alone, even if the screen is in fact blue (imagine that the agent and screen are in separate rooms). However, if an agent enters the room and sees the red screen, it appears possible for her to come to know it is red (assuming that it is, in fact, red). In particular, this is possible in spite of the fact that the likelihood having a (hallucinatory) appearance as of a red screen when the screen is in fact blue will be higher than the likelihood of the screen being red.[15]

As has been noted by, e.g., Smith (2010), facts about normality can easily explain the contrast between these cases.[16] Nothing abnormal would have to happen for the screen to be red. As such, (WNC) predicts that the agent without direct evidence of the color of the screen is not in a position to come to know that it is blue on the basis of the probabilities alone (even if it is, in fact, blue). It would be unlikely, but not abnormal, for her to be in the cognitive state she is in and for the screen be red. However, it would be abnormal for an agent to have an appearance as of the screen being red when it is in fact blue. (WNC) precludes attribution of knowledge of the screen color in former, but not the latter case (regardless of the low probability of a red screen). As such, the condition appears well-suited to explain otherwise problematic judgements. In line with the theoretical motivation for (WNC), it is plausible to suppose that the relevant difference between the agents in the two cases is that the latter, but not the former, is in a state which carries (non-probabilistic) information about the colour of the screen.

[15] Suppose that the probability of having a hallucinatory appearance as of a red screen is .00000001 (i.e., one-in-ten-million). Assume it is independent of the colour of the screen. Then the probability of having an appearance as of a red screen while the screen is in fact blue is $.00000001 \times .99999999 = .000000099999999$ – or almost 10 times as likely as the screen being red.

[16] Smith defends a normality condition on justification, rather than knowledge. Assuming (propositional) justification for *p* is a necessary condition on knowledge that *p*, Smith's condition will entail (WNC).

6 Higher-order ignorance inside the margins

(WNC) constitutes a more plausible implementation of the normality-based picture of knowledge underlying (NTK). (WNC) is the weakest principle entailed by Greco's primary argument for (NTK) and avoids the latter's sceptical implications. However, failures of the KK-principle can be expected to arise from (WNC) in cases in which agents are subject to *systematic gettierization*.

Suppose that S comes to truly believe that the time is 8am as a result of consulting a clock which stopped precisely 12 hours prior. Under these conditions, S's belief that the time is 8am is *gettierized*; it is merely a matter of luck that S's belief is true. Gettierized belief that *p* is (i.) factive and (ii.) incompatible with knowledge that *p*.

Say that S is *systematically gettierized* when in C with respect to *p* iff normally, if S is in C, then S believes that *p* and S's belief that *p* is gettierized. Systematic gettierization is unusual, but does not appear to be in any way incoherent. However, given (WNC), counter-examples to (KK) and (KK^{-}) arise naturally from cases of systematic gettierization. I will first present the abstract structure of such cases before considering concrete examples instantiating this structure.

Suppose that S is systematically gettierized when in C with respect to *p*. Since gettierization is factive, S satisfies (WNC) with respect to *p*. Normally, if she is in C, *p* is true. It is consistent with S being systematically gettierized that she is in C, believes that *p* and her belief is not gettierized—sometimes, conditions are abnormal. In such cases, since she is not gettierized and satisfies (WNC), under the assumption that any further conditions on knowledge are satisfiable, there appears *prima facie* reason to think that she can know that *p*. Nevertheless, she cannot know that she knows that *p*. Since gettierization is incompatible with knowledge, normally, if she is in C, she does not know that *p*. Hence, by (WNC) she is not in a position to know that she knows that *p*. (KK^{-}) fails. Assuming that knowing that *p* entails being in a position to know that *p*, (KK) also fails.

Such counter-examples are compatible with the assumption that normality iterates freely. Instead, they depend on the fact that a world may fail to be normal by its own standards. An agent may be systematically gettierized without being in fact gettierized. We can consider a pair of concrete cases exemplifying this structure:

> INDUSTRIAL SABOTEURS.
> Suppose that it is Xander's job to check the pressure in the water cooler of a nuclear power plant. Every morning, Xander checks a gauge at his desk which is connected to the cooler and fixes the conditions in the plant relative to its pressure. Unbeknownst to Xander, a saboteur, Yasmin, arrives before him in the morning and tampers with the gauge so that it reports the wrong pressure. Luckily, though also unbeknownst to Xander, a counter-saboteur, Zadie, manually changes the pressure in the cooler so that it matches the report Yasmin has set on the gauge.
> However, suppose that one day, both Yasmin and Zadie oversleep and the gauge remains untampered with, reliably reporting the pressure in the cooler

as it is designed to do. Xander checks the gauge and forms a true belief about the pressure inside the water cooler.

Building Facades.

Suppose that Ava has recently moved to a new city. Unbeknownst to her, every building in the city is currently undergoing renovation, and is hidden behind a façade. However, each façade is designed to perfectly match the exterior of the building behind it. Whenever Ava sees a building, her beliefs about its appearance are based on its façade, which she falsely believes to be its exterior.

However, suppose that one day, the façade is temporarily removed from the building opposite Ava's house. She sees its exterior directly and forms a true belief about the building's appearance.

Xander's beliefs about the pressure in the water cooler are systematically gettierized. Normally, when he forms a belief about the pressure, his belief is gettierized. However, on the day on which Yasmin and Zadie oversleep, it appears reasonable to judge that, upon consulting the report on the gauge, Xander knows the pressure in the cooler. In particular, Xander's belief about the temperature satisfies (WNC) and is not gettierized. However, by (WNC), he is not in a position to know that he knows, since normally, his belief would be gettierized and hence fail to constitute knowledge. Hence, if Xander possesses first-order knowledge about the pressure, he is a counter-instance to the KK-Principle.

Ava's beliefs about the appearance of buildings in her town are likewise systematically gettierized. Normally, when she forms a belief about the appearance of a building in the city, her belief is gettierized. However, when the façade on the building opposite her house is removed, it appears reasonable to judge that she can come to know what its exterior looks like. Yet by (WNC), she is not in a position to know that she knows, since normally, her beliefs about its appearance (along with the appearance of other buildings in the city) would be gettierized and hence fail to constitute knowledge. Hence, if Ava possesses first-order knowledge of the appearance of the building, she is, likewise a counter-instance to the KK-principle.

Judgements here are by no means conclusive. The cases are primarily intended to function as concrete examples of the theoretical features which allow for failures of higher-order knowledge under (WNC), even when normality is assumed to iterate freely. Accordingly, there is still potential for a proponent of (WNC) to preserve the KK-principle by resisting the attribution of first-order knowledge to any agent who is systematically gettierized. Before concluding, it will be worthwhile to briefly consider the prospects for this form of response.

The KK-principle sympathiser who accepts (WNC) must deny that S can know, in C, that p if in all normal C_S-worlds it is merely a matter of luck that p; that is, they must deny that systematic gettierization is compatible with first-order knowledge. An analogous concern is raised by Stalnaker (2015) regarding safety. Stalnaker suggests that S's belief, in C, that p, is no better if in all nearby C_S-worlds, p could easily have been false, than it would be if in some nearby C_S-worlds, p were in fact false.[17] If it was merely a matter of luck that one's success was not merely a matter of luck, one is no better off (Stalnaker claims) than if one's success was merely a matter of luck. The present objection would correspondingly hold that S's belief, in C, that *p*, is no better if in all normal C_S-worlds, *p* could easily have been false than it would be if in in some normal C_S-worlds *p* were in fact false. If one's success is normally merely a matter of luck, one is no better off than if it would not be abnormal for one to be unsuccessful. However, accepting the latter as a constraint on knowledge has some seemingly undesirable consequences. In particular, it implies that agents are never in a position to know precisely what could easily be the case.

Consider the following simple model of the kind employed in Williamson (2013) and Goodman (2013), and which is well-suited to our current purposes. A world in the model is a pair of real numbers, $\langle e, f \rangle$. Informally, a world $\langle e, f \rangle$ represents two pieces of information: (i.) the real value, *e*, of some physical parameter, and (ii.) the value, *f*, of the parameter as represented by some system. For present purposes, the second element of each point can be identified with the (class of) state(s) of a cognitive system which represent the parameter as having the relevant value.

Let $N(\langle e, f \rangle)$ be the set of normal worlds at $\langle e, f \rangle$ in which the system is in the state *f*. Let $R(\langle e, f \rangle)$ be the set of worlds which could easily obtain at $\langle e, f \rangle$. We assume (potentially idealizing) that what is normal given that the system is in the state *f* is determined by the value of *f*. That is, there is some $i > 0$, such that $N(\langle e, f \rangle) = \{\langle e', f \rangle : |f-e'| \leq |f-i|\}$. The normal worlds at $\langle e, f \rangle$ are worlds at which the value of the parameter does not differ from its apparent value by more than *i*. Likewise we assume that what could easily be the case at $\langle e, f \rangle$ is determined by *e*. That is, there is some $j > 0$, such that $R(\langle e, f \rangle) = \{\langle e', f \rangle : |e-e'| < |e-j|\}$. The nearby worlds at $\langle e, f \rangle$ are the worlds at which the value of the parameter does not differ from *e* by more than *j*. Let Easy(*k*) be the proposition: $\{\langle e, f \rangle : \langle k, f \rangle \in R(\langle e, f \rangle)\}$; that is, Easy(*k*) is the set of worlds at which *k* could easily be the value of the parameter. We can consider two constraints on K_Sp, the set of worlds at which S knows that *p*:

I. If $\langle e, f \rangle \in K_Sp$, then $N(\langle e, f \rangle) \subseteq p$.
II. If $\langle e, f \rangle \in K_Sp$, then $R(\langle e', f' \rangle) \subseteq p$, for each $\langle e', f' \rangle \in N(\langle e, f \rangle)$.

(I) and (II) correspond to (WNC) and the proposed further constraint, respectively. The former says that S can know that *p* only if *p* is true at all the normal worlds where the system is in *f*. The latter says that S can know that *p* only if *p* is true at all of the worlds near to those worlds. Yet, it follows from (II) alone that S is never in a

[17] Stalnaker (2015, 38): "Does it make a belief any safer, in a sense of safety that has epistemic merit, if all the very similar cases are Gettier cases(cases of justified true belief without knowledge) rather than cases of false belief? […] However the relevant nearness relation is spelled out, it does not seem reasonable to think that a belief being true by coincidence in nearby situations should contribute to the robustness and stability of the belief in the actual situation." Thanks to an anonymous *Philosophical Studies* reviewer for raising this point.

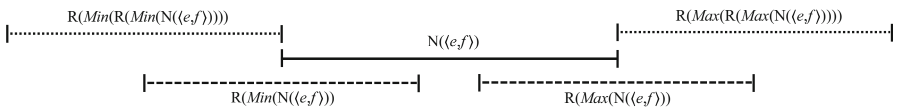

Fig. 1 Diagrammatic representation of proof in footnote 18

position to gain precise knowledge of what the value of the physical parameter could easily be. That is, for all k, $K_S(\text{Easy}(k)) = \varnothing$.[18]

To see why, in intuitive terms, consider the diagram above. Worlds, as pairs of reals, can be identified with points on a plain. N(⟨ e, f ⟩) and R(⟨ e, f ⟩) can be identified with line segments (since their elements differ only with respect to e). The solid line in Fig. 1 corresponds to N(⟨ e, f ⟩). The dashed lines correspond to the sets of worlds which could easily be the case at the minimal and maximal worlds in N(⟨ e, f ⟩), respectively. The dotted lines correspond to the worlds which could have easily been the case at the minimal and maximal worlds in *those* sets, respectively (for comprehensibility, line segments are separated on the vertical axis, though, clearly, they will in fact share the same y-co-ordinate). As can be seen, there is no world which: (i.) could easily have been the case at the minimal world easily the case at the minimal N(⟨ e, f ⟩)-world, and also (ii.) could easily have been the case at the maximal world easily the case in the maximal N(⟨ e, f ⟩)-world. Thus, there will be no k such that, for every N(⟨ e, f ⟩)-world, Easy(k) is true throughout the worlds easily the case at that world. As such, by (II), there is no k such that $K_S(\text{Easy}(k))$.

Yet, it is frequently possible for us to know that there is some precise value a parameter could easily take, even if it is not possible for us to know the precise value it in fact takes. For example, consider a set of scales which are subject to an error of ±.5 kg. Suppose that the scales display 100 kg, and the object on the scales in fact weighs 100 kg. It appears plausible that it is possible to know, under these conditions, that the object could easily be 100 kg (even if one is not in a position to know that it is in fact 100 kg). Yet the proponent of the revised principle must deny this. In order to preserve the KK-principle under (WNC) the defender of the normality-based approach is forced to deny the possibility of precise knowledge of what could easily be the case.[19]

[18] Proof: Let *Min*(A)={⟨e,*f*⟩: ¬∃⟨e′,*f*⟩ ∈ A: e′<e}. *Min*(A) is the set of worlds in A in which the parameter takes an A-minimal value. Let *Max*(A)={⟨e, *f*⟩: ¬∃⟨e′,*f*⟩ ∈ A: e<e′}. *Max*(A) is the set of worlds in A in which the parameter takes an A-maximal value. By (II) $K_S p$ only if *Min*(R(*Min*(N(⟨*e*, *f*⟩))) ∈ *p* and *Max*(R(*Max*(N(⟨*e*,*f*⟩))) ∈ *p*. Yet *Max*(R(*Min*(R(*Min*(N(⟨e, f⟩)))))=*Min*(N⟨*e*, *f*⟩) and *Min*(R(*Max*(R(*Max*(N(⟨*e*, *f*⟩)))))=*Max*(N⟨*e*, *f*⟩). Thus, since |N(⟨*e*,*f*⟩)|>1, there is no *k* such that ⟨*k*, *f*⟩ is R-accessible from every world in both *Min*(R(*Min*(N(⟨*e*, *f*⟩))) and *Max*(R(*Max*(N(⟨*e*,*f*⟩))). So there is no *k* such that Easy(*k*) is true at both.

[19] An anonymous reviewer for *Philosophical Studies* suggests the following, related principle:

S knows, in C, that *p* only if normally, ($C_S \supset$ S knows that *p*).

Assuming that knowledge that *p* is incompatible with it easily being the case that *p*, then this principle will have the same consequence regarding knowledge of what could easily have been the case.

7 Conclusion

The structural properties of normality have been widely assumed to make it well-suited to defence of the KK-principle. This paper has questioned this assumption in two regards. First, it was shown to depend upon the tacit assumption that normality iterates freely, an assumption which, it was claimed, is untenable. Second, it was argued that the most plausible way of formulating the relation between knowledge and normality can be expected to give rise to counter-instances to the principle, even on the assumption that normality iterates freely.

The kind of counter-instances generated by (WNC) are independent of assumptions about inexact knowledge or corresponding margin-for-error constraints (although (WNC) remains compatible with such constraints). The latter forms of counter-instance can arise if an agent's knowledge that *p* depends upon *p*'s truth at some set of appropriately related worlds, where the relation in question is *non-transitive*. In contrast, the form of counter-example discussed in Section 5 arises due to the *non-reflexivity* of the relation in question. Where the set of appropriately related worlds need not include the world of evaluation, an agent can know that *p* despite her belief that *p* possessing a property in all accessible worlds which is factive yet incompatible with knowledge that *p* (for example, gettierization).

It is for this reason that conditions on knowledge stated in terms of normality can be expected to generate counter-instances to the KK-principle which do not arise from conditions stated in terms of relations such as modal closeness, such as safety. Every world is close to itself, but many worlds classify themselves as abnormal.

Acknowledgements This paper has benefited from helpful discussion with and feedback from Brian Ball, Andy Egan, Juan Sebastian Piñeros Glasscock, Daniel Greco, Alex Roberts, Ginger Schultheis, Timothy Williamson and audiences at Yale, Oxford, and The Joint Session of Mind and the Aristotelian Society 2017.

References

Asher, N., & Morreau, M. (1995). What some generic sentences mean. In G. P. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 300–339). Chicago, IL.: University of Chicago Press.

Boutilier, C. (1994). Conditional logics of normality: A modal approach. *Artificial Intelligence, 68,* 87–154.

Das, N., & Salow, B. (2016). Transparency and the KK-Principle. *Noûs, 50,* 3–23.

Delgrande, J. (1987). A first-order conditional logic for prototypical properties. *Artificial Intelligence, 33,* 105–130.

Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

Feldman, R. (1981). Fallibilism and knowing that one knows. *Philosophical Review, 90,* 266–282.

Goodman, J. (2013). Inexact knowledge without improbable knowledge. *Inquiry, 56,* 30–53.

Goodman, J., & Salow, B. (2018). Taking a chance on KK. *Philosophical Studies, 175,* 183–196.

Greco, D. (2014a). Could KK be ok? *The Journal of Philosophy, 111,* 169–197.

Greco, D. (2014b). Iteration and fragmentation. *Philosophy and Phenomenological Research, 88,* 656–673.

Hawthorne, J., & Magidor, O. (2009). Assertion, context, and epistemic accessibility. *Mind, 118,* 377–397.

Millikan, R. (1984). *Language, thought and other biological categories: New foundations for realism*. Cambridge, MA: MIT Press.

Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that-P without believing that-P. *Noûs, 47,* 371–384.

Nelkin, D. K. (2000). The lottery paradox, knowledge and rationality. *The Philosophical Review, 109,* 373–409.

Radford, C. (1966). Knowledge—by examples. *Analysis, 27,* 1–11.

Smith, M. (2008). Ceteris paribus conditionals and comparative normality. *Journal of Philosophical Logic, 36,* 97–121.

Smith, M. (2010). What else justification could be. *Nous, 44,* 10–31.

Stalnaker, R. (1999). *Context and content*. Oxford: OUP.

Stalnaker, R. (2006). On logics of normality and belief. *Philosphical Studies, 128,* 169–199.

Stalnaker, R. (2009). On Hawthorne and Magidor on assertion, context, and epistemic accessibility. *Mind, 118,* 399–410.

Stalnaker, R. (2015). Luminosity and the KK-thesis. In Sanford C. Goldberg (Ed.), *Externalism, self-knowledge and skepticism* (pp. 19–40). Cambridge: CUP.

Stampe, D. (1977). Towards a causal theory of linguistic representation. *Midwest Studies in Philosophy, 2,* 42–63.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic, 25,* 221–261.

Williamson, T. (1992). Inexact knowledge. *Mind, 101,* 217–242.

Williamson, T. (2000). *Knowledge and its limits*. Oxford: OUP.

Williamson, T. (2013). Gettier cases in epistemic logic. *Inquiry, 56,* 1–14.