

# Epistemic closure, assumptions and topics of inquiry

Marcello Di Bello

Received: 13 April 2013 / Accepted: 14 June 2014  
© Springer Science+Business Media Dordrecht 2014

**Abstract** According to the principle of epistemic closure, knowledge is closed under known implication. The principle is intuitive but it is problematic in some cases. Suppose you know you have hands and you know that ‘I have hands’ implies ‘I am not a brain-in-a-vat’. Does it follow that you know you are not a brain-in-a-vat? It seems not; it should not be so easy to refute skepticism. In this and similar cases, we are confronted with a puzzle: epistemic closure is an intuitive principle, but at times, it does not seem that we know by implication. In response to this puzzle, the literature has been mostly polarized between those who are willing to do away with epistemic closure and those who think we cannot live without it. But there is a third way. Here I formulate a restricted version of the principle of epistemic closure. In the standard version, the principle can range over any proposition; in the restricted version, it can only range over those propositions that are within the limits of a given epistemic inquiry and that do not constitute the underlying assumptions of the inquiry. If we adopt the restricted version, I argue, we can preserve the advantages associated with closure, while at the same time avoiding the puzzle I’ve described. My discussion also yields an insight into the nature of knowledge. I argue that knowledge is best understood as a topic-restricted notion, and that such a conception is a natural one given our limited cognitive resources.

**Keywords** Assumption · Bounded rationality · Contextualism · Epistemic closure · Knowledge · Topic of inquiry

---

M. Di Bello (✉)  
Philosophy Department, Stanford University, 450 Serra Mall, Stanford, CA 94305, USA  
e-mail: mdibello@stanford.edu

## 1 A puzzle about epistemic closure

The starting point of this paper is a puzzle that arises if we concurrently maintain epistemic fallibilism and the principle of epistemic closure. Epistemic fallibilism is the view that we can know a proposition *even though* we have not ruled out every possibility of error or deception with respect to that proposition.<sup>1</sup> The principle of epistemic closure, in one of its formulations, runs as follows:

If  $S$  knows that  $p$  and  $S$  knows that  $p$  implies  $q$ , then  $S$  knows that  $q$ .<sup>2</sup>

When we concurrently maintain epistemic closure and fallibilism, a puzzle arises. Let  $A$  and  $B$  be placeholders for certain propositions (examples are to follow shortly). The general form of the puzzle I have in mind is as follows:

- (1)  $S$  knows that  $A$ ;
- (not-2)  $S$  does *not* know that  $B$ ;
- (3)  $S$  knows that  $A$  implies  $B$ ;
- (2)  $S$  knows that  $B$ .

Items (1), (not-2), and (3) hold on the basis of certain intuitions we have about what we know, what we do not know, and what follows from what. Item (2), instead, holds as a consequence of the principle of epistemic closure. We have a puzzle here because each item is *prima facie* plausible, yet (2) contradicts (not-2).

In this paper, I consider two well-known versions of the puzzle. The first version was formulated by Fred Dretske and Robert Nozick in the seventies and eighties.<sup>3</sup> The second, more recent version of the puzzle became prominent thanks to John Hawthorne.<sup>4</sup> Let us begin with the first version. Suppose Mark is an epistemic agent who is endowed with regular sensory and intellectual abilities. Nozick and Dretske observed that Mark knows *ordinary propositions* such as *Mark has hands* provided Mark can perceive his own hands (or *this is a zebra* provided Mark can see what looks like a zebra in front of him). Further, Nozick and Dretske observed that Mark does not know propositions such as *Mark is not a brain-in-a-vat* by merely inspecting his own

<sup>1</sup> The expression “ruling out” is somewhat metaphorical. One way to make it explicit is to say that a possibility is “ruled out” when there is sufficiently strong evidence that the possibility in question does not obtain.

<sup>2</sup> The principle of epistemic closure can be spelled out in different ways; see Kvanvig (2006). In particular, a number of authors have adopted a formulation of the principle in terms of competent inference; see, among others, Williamson (2000), Hawthorne (2004) and Barker and Adams (2010). Here is an example of such a formulation:

*Inference-based closure.* Suppose  $S$  knows that  $p$ ;  $S$  knows that  $p$  implies  $q$ ; and  $S$  competently infers  $q$  from  $p$ , thereby coming to believe  $q$ . Then,  $S$  knows  $q$ .

The inference-based version is more explicit about the underlying psychological process that allows the epistemic agent to know a proposition through known implication. To ease exposition, I will work with the implication-based formulation, although the claim of this paper should also apply *mutatis mutandis* to inference-based closure.

<sup>3</sup> See Nozick (1981), Dretske (1970) and Dretske (1971).

<sup>4</sup> See Hawthorne (2004).

hands (nor does Mark know *this is not a cleverly disguised mule* by merely seeing what looks like a zebra); call the latter *heavyweight propositions*. So, we have:

(ND-1) Mark knows that he has hands.

(not-ND-2) Mark does *not* know that he is not a brain-in-a-vat.

But now notice that *Mark has hands* implies *Mark is not a brain-in-a-vat* because brains-in-a-vat do not have hands. (By the same token, *this is a zebra* implies *this is not a cleverly disguised mule* because zebras are not mules). Mark should realize, on reflection, that the implication holds good. So, the following is true:

(ND-3) Mark knows that *Mark has hands* implies *Mark is not a brain-in-a-vat*.

By the principle of epistemic closure, from (ND-1) and (ND-3), we have:

(ND-2) Mark knows that he is not a brain-in-a-vat.

The problem is that (ND-2) and (not-ND-2) form a contradiction. To avoid the contradiction, we can deny one of the premises—(ND-1), (not-ND-2), or (ND-3)—or we can deny epistemic closure, or formulate a suitable variation thereof. There is a puzzle here because each of the three premises is *prima facie* plausible, and so is epistemic closure. I shall refer to this as *Nozick–Dretske’s version of the puzzle*, or more simply, as *Nozick–Dretske’s puzzle*.

More recently, John Hawthorne has formulated a different version which follows closely the general form of the puzzle given above. Suppose Jack is of modest means and a wealthy friend of his asks Jack whether he wants to come on a luxury vacation with him. In this situation, the following seems true:

(H-1) Jack knows that he will be unable to afford a luxury vacation.

Of course, Jack could win the lottery; he could receive a gift from a relative; etc. Yet, if fallibilism holds, Jack need not rule out every alternative situation in order to know that he will be unable to afford a luxury vacation. In support of (H-1), and of fallibilism more generally, one can point out that if ruling out every alternative possibility of error or deception was a requirement for knowing a proposition, we would know very little if anything.

Suppose now that Jack bought a lottery ticket. Despite the small chance of winning, many of us have the intuition that the following is true:

(not-H-2) Jack does *not* know that he will not win the lottery.<sup>5</sup>

(Some might disagree here. It is difficult to articulate why we do not know a so-called *lottery proposition* such as *Jack will not win the lottery*, and why we know an *ordinary proposition* such as *Jack will be unable to afford a luxury vacation*. After all, both

<sup>5</sup> As Vogel (1990) has noted, the phenomenon generalizes to situations beyond lottery scenarios. Imagine John parked his car in a relatively crime-free neighborhood and went to eat in a nearby restaurant; unless we knew John returned to his car and found it still parked, we would be disinclined to take ourselves to know the proposition *John’s car has not been stolen*, despite its high probability. Further, consider this scenario: Mike is young and in good health; the proposition *Mike will not suffer a sudden heart attack* is highly probable, but again we would be disinclined to take ourselves to know it, unless we had more precise information about Mike.

propositions are highly probable, and one could argue that no epistemically relevant difference can distinguish between the two. This paper will not be concerned with this issue, and for the purpose of my argument, I ask the reader to concede that (H-1) and (not-H-2) are both true).<sup>6</sup>

Next, Jack would have no problem in accepting that if he does win the lottery, he will gain a considerable amount of money and be able to afford a luxury vacation. Thus, the following holds:

(H-3) Jack knows that *Jack will be unable to afford a luxury vacation* implies *Jack will not win the lottery*.

It is worth noting that the implication in (H-3) does not hold as a matter of logic alone, nor does it hold in virtue of the meaning of the words used. Rather, the implication holds in virtue of our common knowledge of what the world is like.<sup>7</sup>

But now a problem arises. If the principle of epistemic closure holds, from the conjunction of (H-1) and (H-3), it follows that:

(H-2) Jack knows that he will not win the lottery.

The problem is that (H-2) contradicts (not-H-2). I shall refer to this as *Hawthorne's version of the puzzle*, or more simply, as *Hawthorne's puzzle*.

(Some might be unimpressed by Hawthorne's puzzle. They might argue that so long as Jack is participating in a fair lottery, he does not know that he will be unable to afford a luxury vacation, and if so, (H-1) would be false and no puzzle would arise.<sup>8</sup> Other readers might again resist the truth of (H-1) by arguing that the true content of what Jack knows in (H-1) is that Jack will be unable to afford a luxury vacation *unless* he wins the lottery or *assuming* he does not win the lottery.<sup>9</sup> Admittedly, it is difficult to state Hawthorne's puzzle in a completely uncontroversial way. But even

<sup>6</sup> Philosophers have advanced different accounts for why (H-1) and (not-H-2) are both true. Some have argued that our evidence is not causally or explanatorily connected with a lottery proposition, whereas it is connected with an ordinary proposition; see [Achinstein \(1978\)](#) and [Nelkin \(2000\)](#). Others have invoked modal notions such as "safety" and "tracking", and argued that the belief in a lottery proposition is not safe nor tracking, whereas the belief in an ordinary proposition is safe or tracking; see [Pritchard \(2005\)](#), [Williamson \(2000\)](#), [DeRose \(1996\)](#) and [Roush \(2006\)](#). Still others have argued that we are not in a position to rule out the scenario in which a lottery proposition is false, but we are in such a position while considering an ordinary proposition; see [Dretske \(1971\)](#), [Vogel \(1999\)](#) and [Lewis \(1996\)](#).

<sup>7</sup> [Roush \(2006\)](#) disagrees with (H-3); see note 13.

<sup>8</sup> Here is a response. Let's suppose Jack is not participating in any fair lottery. There is still a non-zero chance that he might inherit one million dollars from a distant relative. If one denies that Jack knows he will be unable to afford an expensive vacation because he is participating in a fair lottery, by parity of reasoning, one will also have to deny that Jack knows he will be unable to afford an expensive vacation because he might inherit one million dollars. But if so, there would be hardly any situation in which Jack can know he will be unable to afford an expensive vacation. This is a dangerously skeptical conclusion.

<sup>9</sup> I am sympathetic with this point. I say something along these lines in Sect. 4. However, instead of construing the knowledge claim as 'Jack knows [Jack will be unable to afford a luxury vacation *assuming* Jack does not win the lottery]', I suggest that we construe the knowledge claim as 'Jack knows [Jack will be unable to afford a luxury vacation] *assuming* Jack does not win the lottery'. The two construals differ in the scope of the knowledge operator; on the significance of this difference, see note 27.

if we dismiss it, we can find in the literature possibly less controversial formulations, such as Nozick–Dretske’s puzzle.<sup>10</sup>

In this paper, I propose a solution to both versions of the puzzle. The solution is based on a restricted version of the principle of epistemic closure. In the standard version, the principle can range over any proposition; in the restricted version, it can only range over those propositions that are within the limits of a given epistemic inquiry. If we adopt the restricted version, I argue, we can preserve some of the advantages associated with closure, while at the same time avoiding the puzzle. I should say that the idea of restricting closure is not new.<sup>11</sup> The novelty of my approach consists in restricting closure to what I call the *topic of inquiry* and in motivating the restriction on the basis of considerations that appeal to our limited cognitive resources.

## 2 Giving up unrestricted closure

We can tackle our puzzle by adopting a number of revisionary approaches. We can go the “skeptical way” and argue that we do not know ordinary propositions, thereby denying premises (ND-1) and (H-1). We can go the “dogmatic way” and argue that we do know lottery and heavyweight propositions, thereby denying premises (not-ND-2) and (not-H-2). We can deny that a heavyweight or a lottery proposition follows from an ordinary one, thereby denying premises (ND-3) and (H-3). We can take issue with the notion of knowledge. We can deny epistemic closure. And there may very well be other options.

One of the most prominent revisionary strategies in the literature is the denial of epistemic closure.<sup>12</sup> I shall discuss this strategy here and leave the discussion of the other revisionary options to a footnote.<sup>13</sup> The puzzle we’ve seen at the beginning

<sup>10</sup> Another version of the puzzle goes as follows:

(T-1) I know that it is 3 PM (e.g. by reading my watch).

(not-T-2) I do not know that my watch is not mistakenly reporting that it is 3 PM (or at least, I do not know that by simply reading my watch).

(T-3) I know that ‘it is 3 PM’ *implies* ‘my watch is not mistakenly reporting that it is 3 PM’. (The implication holds because if it is 3 PM and my watch says that it is 3 PM, it follows that my watch is not mistakenly reporting that it is 3 PM.)

(T-2) I know that my watch is not mistakenly reporting that it is 3 PM (by closure).

For reasons of brevity, I do not discuss this version of the puzzle here, but I believe that the machinery offered in the paper should be able to handle it, as well.

<sup>11</sup> Authors working in the framework of “relevant alternative theories of knowledge” have recently proposed restrictions for epistemic closure, e.g. Lawlor (2013) and Barker and Adams (2010).

<sup>12</sup> Nozick (1981) proposed to characterize knowledge in terms of the modal notion of “tracking”. Dretske (1970) and Dretske (1971) held that knowledge that  $p$  consists in having a conclusive reason for  $p$  or in having ruled out all relevant alternatives in which  $p$  is false. On both accounts of knowledge, epistemic closure fails. Epistemologists who recently questioned epistemic closure include Sharon and Spectre (2013), Adams et al. (2012), Sherman and Harman (2011) and Lawlor (2005).

<sup>13</sup> One way to solve our puzzle is by endorsing a form of skepticism or dogmatism. These solutions, however, are usually considered unattractive because they entail that we would know too little or too much. In other words, they contradict epistemic fallibilism which many epistemologists endorse. Another option consists in denying that lottery/heavyweight propositions follow from ordinary propositions. I find this difficult to accept, especially if the implication holds between ordinary and heavyweight propositions.

of this paper suggests that epistemic closure, at least in its standard formulation, is problematic. If we were to give up epistemic closure, the puzzle would disappear. Shall we give up closure, then? Many authors think that closure is one of the most intuitive epistemological principles and that giving it up would be theoretically too costly. For instance, Richard Feldman emphatically writes that “the idea that no version of this principle is true strikes me, and many other philosophers, as of one of the least plausible ideas to come down the philosophical pike in recent years”.<sup>14</sup>

I will now rehearse three arguments that are sometimes given in support of epistemic closure. As will soon become apparent, these arguments can be plausibly construed as arguments in support of a restricted principle of epistemic closure. To begin with, the idea that knowledge is closed under known implication is intuitively appealing because logic is closed under implication. Just as *modus ponens* tells us that from  $p$  and  $p \rightarrow q$ , it follows that  $q$ , epistemic closure tells us that from knowing  $p$  and knowing  $p \rightarrow q$ , it follows that one knows  $q$ . So a first argument in support of epistemic closure goes as follows:

MODUS PONENS ARGUMENT. *Modus ponens* is a hardly objectionable logical rule. Epistemic closure mirrors *modus ponens* in the epistemic context.

Countless examples can be given. You know that the trash is collected once a week on Fridays. You know, on reflection, that if the trash is collected once a week on Fridays, it won't be collected on Mondays. So, you know that the trash won't be collected on Mondays. Or you know that there are ten people in the room. You know that if there are ten people in the room, there are at least five. So, you know there are at least five people in the room. Or again, you know it is three o'clock. You know that if it is three o'clock, it is not evening yet. So, you know it is not evening yet. And so on. These are all straightforward examples. But those who deny epistemic closure do not deny that

---

Footnote 13 continued

Now, *Mark has hands* implies *Mark is not a brain-in-a-vat*, in virtue of the meaning of the words used, because a handed creature such as Mark cannot be a handless creature such as a brain. So long as we master the words 'hands' and 'brain', the implication is unobjectionable. More plausible is to think that an ordinary proposition need not imply a lottery proposition. For instance, if Jack will be unable to afford a luxury vacation, it need not follow that his lottery ticket is a loser, for even if he were to win the lottery and collect the money, he might suddenly lose the money and thus be unable to afford the vacation. All we can say, it would seem, is that if Jack wins, it is highly likely that he will be able to afford the vacation, or that if Jack wins, he will be able to afford the vacation *on the assumption that* he does not suddenly lose the money Roush (2006). This is a fair point, but all it shows is that sometimes a proposition follows from another only tentatively, not definitely. This conclusion fits well with the claim of this paper, i.e. that we should take into account the role that assumptions play in our epistemic inquiries (see Sect. 5). Finally, another way to tackle our puzzle is to take issue with the notion of knowledge. Some philosophers endorse probabilism. This is the view that a rational epistemic agent should assign degrees of belief to propositions in accordance with the axioms of probability theory; see Easwaran (2011). The relation between degrees of belief and knowledge is unclear, but the probabilist can argue that if an epistemic agent assigns the same degree of belief to two propositions, she cannot exhibit any form of epistemic asymmetry relative to these propositions. The probabilist, then, can argue that since ordinary, lottery, and heavyweight propositions are (roughly) equally probable, an epistemic agent cannot be in a position to know some and not others. The probabilist can try to solve our puzzle by resisting our intuitions about knowledge. Instead, I try to provide a solution that is more faithful to our intuitions.

<sup>14</sup> See Feldman (1995).

knowing an implication and knowing the antecedent often puts the epistemic agent in a position to know the consequent. They simply deny that that's always the case. After all, we have seen some problematic cases, such as Nozick–Dretske's puzzle and Hawthorne's puzzle. This suggests that epistemic closure might need to be suitably restricted.

Let us try a second argument in defense of closure:

**WARRANT TRANSMISSION ARGUMENT.** Warrant transmits across a known implication. If  $S$  has warrant for  $p$  and  $S$  knows  $p \rightarrow q$ , then  $S$  must have warrant for  $q$  as well.

The argument here is that knowledge is *closed* under a known implication because warrant *transmits* across a known implication.<sup>15</sup> However we think of it, the Warrant Transmission Argument is a seemingly general argument. For it applies to any  $p$  or  $q$  provided the epistemic agent knows that  $p$  implies  $q$ . Yet, despite the seeming generality, some have argued—I think rightly—that warrant fails to transmit in some cases. Suppose Mark has warrant for believing that he has hands, because, say, he is looking at his hands. Mark also knows, on reflection, that having hands implies not being a brain-in-a-vat. Still, it seems that Mark cannot have warrant for believing that he is not a brain-in-a-vat *just in virtue of* looking at his hands.<sup>16</sup> Warrant often transmits, but not always. The Warrant Transmission Argument makes a good case for some version of closure, but it is not enough to make a case for unrestricted closure.

Let us consider a third argument in defense of closure:

**SPREADING ARGUMENT.** Denying epistemic closure typically requires the denial of other plausible epistemic principles, such as the equivalence principle (i.e. if  $S$  knows a priori that  $p$  and  $q$  are equivalent,  $S$  knows  $p$  iff  $S$  knows  $q$ ) or the addition principle (i.e. if  $S$  knows  $p$ , then  $S$  knows  $p$ -or- $q$ ).

I will not discuss the details of the Spreading Argument here.<sup>17</sup> Let me simply note that this argument is compatible with a restricted version of epistemic closure. If we endorse a restricted version of closure, it is plausible to endorse a restricted version of its satellite principles, such as the addition or the equivalence principle. On one reading of it, the Spreading Argument suggests that closure is an unrestricted principle provided both the equivalence principle and the addition principle are unrestricted. But do we have independent reasons to think that both these principles must be unrestricted? I think not. Arguably, the addition principle can be left unrestricted, but it is clear that

<sup>15</sup> Tucker (2010) makes a roughly similar point. The Warrant Transmission argument can also be seen with probabilities. Let  $Pr(p) = k$  and let  $Pr(p \rightarrow q) = 1$ . It follows that  $Pr(q) \geq k$ , because  $Pr(p) \leq Pr(q)$ . This suggests that the (probabilistic) strength of the warrant does not decrease across an implication. Importantly, this holds provided  $Pr(p \rightarrow q) = 1$ .

<sup>16</sup> See Wright (2004) and Pryor (2012). Similarly, Jack has warrant for believing that he will not be able to afford an expensive vacation, because, say, he has a modest salary and his bank account contains little money. Yet, Jack does not have warrant for believing that he will not win the lottery *just because* of his modest salary and the little money in his bank account.

<sup>17</sup> For a more detailed discussion, see Hawthorne (2004) and Holliday (2014).

the equivalence principle is prone to a number of counterexamples which are similar to those typically leveled against epistemic closure.<sup>18</sup>

If I am right, the above three arguments in defense of closure fail to support unrestricted closure, although they make a good case for some version of closure. In fact, these arguments can be taken to support a restricted principle of closure, or at least they are not downright incompatible with it. It is now time to explain the criteria that should govern the workings of a restricted closure principle. I will proceed in two steps. First, I will defend the claim that knowledge rests on tacit, background assumptions, which are not themselves items of knowledge (Sects. 3, 4). Second, I will argue that knowledge and epistemic closure are restricted to the ongoing topic of inquiry (Sects. 5, 6, 7). These two steps are connected: the epistemic assumptions lie outside the topic of inquiry, and since closure is restricted to the topic, assumptions will be excluded from the scope of closure. I will also compare my account with epistemic contextualism, i.e. the view according to which knowledge attributions and epistemic closure are relativized to context-dependent epistemic standards. The comparison is instructive because there are some similarities between my approach and epistemic contextualism (Sect. 8).

### 3 Evidence, assumptions and knowledge

As David Lewis reminds us, we know a lot.<sup>19</sup> We know we have hands; we know we will be unable to afford a luxury vacation; we know the earth revolves around the sun; we know our car is parked around the block; and so on. How do we know these things? Our sense perceptions tell us we have hands; our modest monthly income tells us what we can and cannot afford; satellite observations tell us the earth revolves around the sun; our memory, together with information about the low rate of car thefts in the neighborhood, tells us where our car is parked; and so on. In short, knowledge of a proposition requires *evidence* for that proposition.

Yet, the evidence we have for what we know cannot rule out every possibility of error or deception with respect to what we know. We know we have hands, but the deliverances of our senses are not enough to rule out the possibility of widespread deception; we know we will be unable to afford an expensive vacation, but our modest income cannot rule out the possibility that we might win the lottery; we know the earth revolves around the sun, but satellite observations cannot rule out the possibility that we might be under a systematic illusion or that an evil genius might deceive us; we know our car is parked around the block, but our evidence cannot rule out the possibility that our car might have been stolen; and so on. It is hard to rule out every possibility of error or deception. Even though (presumably) we know many things, we might have failed to know them. Our knowledge, after all, is fallible.

<sup>18</sup> The proposition *Mark has hands* and *Mark has hands and Mark is not a handless brain-in-a-vat* are equivalent, and Mark should realize, on reflection, that they are equivalent. Yet, it is rather problematic to say that if Mark knows that he has hands, then Mark knows that he has hands and that he is not a brain-in-a-vat. This is at least as problematic as saying that Mark knows he is not a brain-in-a-vat. So, the equivalence principle (in its unrestricted form) seems suspect.

<sup>19</sup> See Lewis (1996).



To maintain the possibility of fallible knowledge and avoid skepticism, I suggest that we think of fallible knowledge as knowledge that rests on evidence *as well as* on assumptions. To know that we have hands on the basis of our sense perceptions, we must rely on the assumption that we are not systematically mistaken; to know that we will be unable to afford a luxury vacation given our modest income, we must rely on the assumption that we will neither win the lottery nor inherit one million dollars; to know that the earth revolves around the sun on the basis of satellite observations, we must rely on the assumption that the observations are not misleading; to know that our car is parked where we left it, we must rely on the assumption that it has not been stolen; and so on. To have knowledge of a proposition, then, we need to rely on our evidence as well as on certain background assumptions.<sup>20</sup>

I should say that the idea that knowledge rests on assumptions has been defended by a numbers of authors, among whom are Ludwig Wittgenstein in *On Certainty*, Crispin Wright, and more recently, Gilbert Harman and Brett Sherman.<sup>21</sup> In this section, I will distill and systematize some ideas on the role of assumptions which are inspired by these authors. This will pave the way for my discussion of epistemic closure in the sections to come.

Let me now be more explicit about what it means for an epistemic agent to assume a proposition and what types of propositions can serve as assumptions. Simply put, an epistemic agent assumes a proposition whenever she takes it for granted in the sense that she implicitly relies on it as though it were true and even though she lacks compelling or conclusive evidence for its truth. Depending on the situation at hand, an epistemic agent can take for granted different types of propositions. For the purpose of my argument, I will focus on two types only. First, consider propositions such as: our satellite observations are not systemically wrong; no evil genius is deceiving us all the time; we are not brains-in-a-vat; scientific practice does not rest on a colossal mistake; etc. These propositions are denials of various skeptical scenarios. Call them *anti-skeptical propositions*. Second, consider propositions such as: my car has not been stolen; Mark will not win the lottery; Joe will not have a sudden heart attack in young age; etc. These are not anti-skeptical propositions; rather, they are “instantiations” of certain statistical generalizations in the world, such as: cars *rarely* get stolen (at least, in safe neighborhoods); people *rarely* win lotteries; young people *rarely* have heart attacks; etc. I shall call *regularity propositions* those which are instantiations of statistical generalizations.

With the distinction between anti-skeptical and regularity propositions in hand, let us look at the role that assumptions play in our epistemic lives. I will now discuss four core features of assumptions.

- (a) *An epistemic agent’s putative knowledge of certain propositions can be undermined if what the agent assumes is false.* An epistemic agent sometimes needs to take for granted certain propositions in order to have knowledge, because the

<sup>20</sup> The view expounded here is that it is epistemic agents (inquirers or potential knowers) who rely on assumptions. This view should be contrasted with a possible alternative view, i.e. one on which knowledge *attributors* rely on assumptions when they utter sentences of the form ‘*S* knows that *p*’. I do not develop this alternative view here. See Sect. 8 for some remarks on this score.

<sup>21</sup> See e.g. Harman and Sherman (2004), Sherman and Harman (2011) and Wright (2004).

falsity of these propositions can undermine the agent's knowledge. Consider first anti-skeptical propositions. A plausible picture is one on which we assume that we are not brains-in-a-vat or that scientific research is not systematically mistaken. If the assumptions in question were false, we would not know that we have hands, and all presumed scientific knowledge would be unwarranted, thereby failing to count as knowledge. Next, consider regularity propositions. Suppose Jack is mistakenly assuming that he is going to lose the lottery, because he is, in fact, going to win. If so, it would be false that Jack cannot afford a luxury vacation and thus it would be false that he knows he cannot afford a luxury vacation. All in all, the falsity of one's assumptions can undermine one's knowledge of a proposition, because it can undermine one's evidential basis for believing the proposition or it can make the proposition in question downright false.

- (b) *An epistemic agent relies on an assumption even though she does not know it.* It is important to understand what it means to claim that an epistemic agent, at least in some circumstances, does not know what she assumes. Consider anti-skeptical propositions. Our lack of knowledge of anti-skeptical propositions is quite widespread, even systemic and unavoidable. Anti-skeptical propositions, after all, are propositions which we hardly ever know because we cannot—even if we wanted to—gather any decisive evidence about them. In most if not all circumstances, we do not know that we are not brains-in-a-vat; we do not know that our observations are not systematically misleading; we do not know that laboratory analyses are not manipulated by an evil genius; etc. If we could know such things, refuting skepticism would be easy, but it isn't.<sup>22</sup> When it comes to regularity propositions as opposed to anti-skeptical ones, we are often in a position to know them. For instance, Jack can certainly come to know whether or not he lost the lottery by reading a trustworthy newspaper. The important point here is that unless Jack gathered additional evidence, he would not know that he lost (or won) the lottery by simply buying a ticket. Regularity and anti-skeptical propositions, then, are different. We lack knowledge of regularity propositions only momentarily, and we can come to know them by gathering additional evidence; our lack of knowledge of anti-skeptical propositions, instead, is more widespread and systemic. Both types of propositions, however, are similar. As epistemic agents, we sometimes rely on these two types of propositions even though we do not know them in the sense that we momentarily or systemically lack sufficient supporting evidence.
- (c) *The lack of sufficient supporting evidence does not mean that epistemic agents have no reason to rely on what they assume; on the contrary, such a reliance is justified.* First, consider anti-skeptical propositions. We have a reason to rely on them, despite the likely impossibility of refuting skepticism. It is hard to conceive of a world where we are handless brains-in-a-vat; where observations

<sup>22</sup> Some might think that, in particular circumstances, it is possible to know anti-skeptical propositions such as 'I am not dreaming right now'. I do not have a compelling response here. I maintain that, in most if not all circumstances, we do not know anti-skeptical propositions and that this is a starting point for the puzzle raised by Nozick and Dretske. If we thought we knew anti-skeptical propositions, this would presumably solve the puzzle, but it would solve it in a way that dismisses it at the outset.

are always wrong; or where an evil genius confuses us all the time. We would need to radically rethink our lives to make sense of a world in which our anti-skeptical assumptions do not hold. This means that these propositions are deeply entrenched in our practical and intellectual lives; this is what gives us a reason to hold on to them, despite lacking knowledge of them.<sup>23</sup> Next, consider regularity propositions. When we need to, we can rely on regularity propositions because they are instances of well-justified statistical generalizations, although we don't know whether a given instance is true in the specific case. After all, our car might be one of the few cars which got stolen; our ticket might be one of the few winning tickets; etc.

- (d) *The truth of an assumption cannot be enough to guarantee the truth of the agent's putative knowledge.* An assumption cannot be as strong as to guarantee the truth of what one knows, because if that were the case, the contribution of evidence to knowledge would become pointless.<sup>24</sup> Although I have suggested that knowledge rests on propositions which epistemic agents justifiably assume without knowing them, I am not denying that we need appropriate evidence as a basis for what we know. Assumptions should be enriched with evidence to yield knowledge, and conversely, the evidence should be properly supplemented by a number of assumptions to yield knowledge. I will not go into the details here of how assumptions and evidence should be combined to yield knowledge, but I should emphasize that we cannot gain knowledge by cleverly assuming whatever we please. As is apparent from (c) above, not any proposition can serve as an assumption. What we assume needs to be highly probable on the basis of our evidence or deeply entrenched in our conceptual and practical lives.

To recapitulate, then, a proposition  $\alpha$  serves as an assumption which underlies  $S$ 's knowledge that  $p$  whenever the following conditions hold:

- (a) If  $\alpha$  were false,  $S$ 's knowledge that  $p$  would be undermined;
- (b)  $S$  relies on  $\alpha$  as though  $\alpha$  were true, yet  $S$  does not know  $\alpha$ ;
- (c)  $S$  justifiably relies on  $\alpha$ , i.e.  $\alpha$  is deeply entrenched or highly probable; and
- (d) the truth of  $\alpha$  alone does not guarantee the truth of  $p$ .

It is useful to distinguish the *attitude* of assuming from the *proposition* that is being assumed. The attitude of assuming amounts to an implicit and justified reliance; the proposition being assumed can be of different types, e.g. regularity and anti-skeptical propositions. The latter are deeply entrenched in our practical and intellectual lives; the former are highly probable given certain statistical data.

At this point, some might object that what we know should rest on our evidence, not on assumptions. They might say that, given our evidence, certain propositions are more or less probable, or more or less evidentially supported, so that an appeal to assumptions would be unnecessary. The matter is difficult, and I can only offer a couple

<sup>23</sup> Propositions that are deeply entrenched are similar to Wittgenstein's *bedrock propositions* from *On Certainty*.

<sup>24</sup> According to condition (d), a proposition such as *Mark has hands AND Mark is not a brain-in-a-vat* would not count as an assumption that underlines Mark's knowledge that he has hands, because *Mark has hands AND Mark is not a brain-in-a-vat* alone guarantees the truth of *Mark has hands*.

of tentative remarks here. First, note that for a proposition to count as evidence for another proposition, the evidence needs to be assessed and interpreted. This requires an appeal to a set of assumptions. To illustrate, suppose we look through a telescope and notice that the moon has an irregular surface. This observation, presumably, is evidence that the moon's surface is irregular. But why? The observation is evidence because we rely on the assumption that the telescope can tell us something about the moon. Thus, it seems difficult to escape the role that assumptions play in our epistemic enterprises.<sup>25</sup>

The second remark I want to make concerns how my assumption-based account of knowledge compares with the “relevant alternatives” framework. The suggestion that what we know rests on a combination of evidence and assumptions is not very different from what Fred Dretske and David Lewis had in mind.<sup>26</sup> They believed that in order to know  $p$ , we should rule out certain relevant alternatives to  $p$ , where the “ruling out” is an evidence-based process. For instance, in order to know that we have hands, we should rule out the alternative that our hands have been cut off, and we can do so by gathering sensory evidence, e.g. by looking at our arms and noticing that our hands are still attached to them. Dretske and Lewis, at the same time, believed that in order to know  $p$ , we need *not* rule out every alternative to  $p$ , because certain alternatives can be “properly ignored” insofar as they are irrelevant. For Dretske and Lewis, we know as a result of “ruling out” relevant alternatives and of “properly ignoring” irrelevant alternatives. Importantly, while the process of “ruling out” is based on evidence, the process of “properly ignoring” is not (entirely) based on evidence, and we might very well say that it rests on certain assumptions regarding what counts as an irrelevant alternative. So, even if the reader is unsympathetic with the prominent role that assumptions play in my account, I do not think I have made a drastic departure from some of the commitments of a theory of knowledge based on relevant alternatives.

#### 4 Assumptions and known implication

If knowledge rests on assumptions, a statement such as ‘ $S$  knows that  $p$  on the basis of evidence  $e$ ’ can be made explicit as follows:

$S$  knows [ $p$ ] assuming  $\alpha_1, \alpha_2$ , etc. on the basis of evidence  $e$ .<sup>27</sup>

<sup>25</sup> On the role of assumptions in scientific inquiries, see chapter 3 of Longino (1990).

<sup>26</sup> See Dretske (1971) and Lewis (1996).

<sup>27</sup> I have in mind a *narrow scope* construal of the knowledge operator. Alternatively, ‘ $S$  knows that  $p$  on the basis of evidence  $e$ ’ can be made explicit by using a *wide scope* construal, namely:

$S$  knows [ $p$  assuming  $\alpha_1, \alpha_2$ , etc.] on the basis of evidence  $e$ .

I suspect that the wide scope analysis will turn out to be problematic. For suppose Jack claims that he will not be able to afford a luxury vacation. In response, suppose his interlocutor points out that Jack will receive one million dollars from a distant relative, as shown in a notarized letter Jack just received. At this point, Jack will have to correct himself in some way; he will have to admit he misspoke, or at least, he will have to clarify his earlier claim. It seems that Jack's knowledge claim has been contradicted. But if the true content of Jack's knowledge claim were simply that he will be unable to afford a luxury vacation *assuming* e.g. he

For ease of exposition, I will drop the specification “on the basis of evidence  $e$ ”. The question I address in this section is the following: on an assumption-based conception of knowledge, do we know—through known implication—what follows from what we know? My answer will be that we do *provided what follows is NOT an assumption*. To motivate this answer, I will work with the propositions *Jack will be unable to afford a luxury vacation* and *Mark has hands*. If you recall, these propositions were part of the puzzle I presented at the beginning of this paper.

Suppose Jack knows he will be unable to afford a luxury vacation (abbreviated *not-afford*). On my account, Jack knows *not-afford* tacitly assuming that he will not win the lottery; that he will not inherit a million dollars; etc. Further, suppose Mark knows he has hands (abbreviated, *hands*). Again, Mark knows *hands*, tacitly assuming that he is not a brain-in-a-vat; that he is not systematically deceived; etc. We therefore have:

(H-1\*) Jack knows [*not-afford*], assuming *not-win*, *not-inheritance*, etc.

(ND-1\*) Mark knows [*hands*], assuming *not-deception*, *not-brain-in-vat*, etc.

The question here is whether, given (H-1\*), it is the case that Jack knows he will not win the lottery (abbreviated *not-win*); and whether, given (ND-1\*), it is the case that Mark knows he is not a brain-in-a-vat (abbreviated *not-brain-in-vat*). Indeed, Jack and Mark should know that *not-afford* implies *not-win* and that *hands* implies *not-brain-in-vat*. But is a known implication a path to knowledge?

Keeping in mind that knowledge rests on assumptions, it is plausible to say that a known implication—e.g. *hands* implying *not-brain-in-vat*, and *not-afford* implying *not-win*—cannot interfere with the assumptions underlying one’s knowledge. Here is a principle that captures this idea:

ASSUMPTION-INERTIA. Other things being equal, a known implication does not cancel one’s assumptions. More precisely, suppose  $S$  knows that  $p$  by assuming  $\alpha_1, \alpha_2$ , etc. and suppose  $S$  knows that  $p$  implies  $q$ . Then, other things being equal,  $S$  knows that  $q$  (on the basis of knowing  $p$  and knowing  $p \rightarrow q$ ) only by (still) assuming  $\alpha_1, \alpha_2$ , etc.<sup>28</sup>

If we were to apply the standard principle of epistemic closure, it would follow that Jack knows *not-win* and that Mark knows *not-brain-in-vat*. Yet, in accordance with the principle of assumption-inertia, Jack and Mark can be in a position to know *not-afford* and *not-brain-in-vat*, only by tacitly assuming a number of things. In other words, Jack’s and Mark’s epistemic situations would be:

(H-2\*) Jack knows [*not-win*], assuming *not-win*, *not-inheritance*, etc.

---

Footnote 27 continued

will not win the lottery or *assuming* e.g. he will not inherit one million dollars, then Jack would not need to correct himself. It would be as though his interlocutor did not really contradict him. This seems odd, and that’s why I prefer the narrow scope construal.

<sup>28</sup> I am here assuming that  $S$ ’s knowledge of  $p$  rests on assumptions  $\alpha_1, \alpha_2, \dots$ , while  $S$ ’s knowledge of  $p \rightarrow q$  rest on no assumptions. This is a simplification. If the knowledge of  $p \rightarrow q$  also rests on assumptions, say,  $\beta_1, \beta_2, \dots$ , the resulting knowledge of  $q$  will rest on assumptions  $\alpha_1, \alpha_2, \dots$  as well as  $\beta_1, \beta_2, \dots$ .

(ND-2\*) Mark knows [*not-brain-in-vat*], assuming *not-deception*, *not-brain-in-vat*, etc.

Notice that the assumptions in (H-1\*) and (ND-1\*) are the same as the assumptions in (H-2\*) and (ND-2\*), in accordance with the principle of assumption-inertia. Can we say that—by known implication—Jack knows *not-win* or that Mark knows *not-brain-in-vat*? I think not. Jack’s putative knowledge of *not-win* and Mark’s putative knowledge of *not-brain-in-vat* are made explicit in (H-2\*) and (ND-2\*). The latter describe situations of the form ‘*assuming p, S knows p*’. But one cannot know a proposition by simply assuming it.<sup>29</sup> So the following restriction applies:

NO EASY KNOWLEDGE RESTRICTION. If *S* knows that *p*, then *p* is not among (or entailed by) the assumptions that underlie *S*’s knowledge that *p*.

Consequently, if Jack’s and Mark’s epistemic situations are adequately described by (H-2\*) and (ND-2\*) and if the above restriction holds, then Jack does not know *not-win* nor does Mark know *not-brain-in-vat*. This accords with our intuitions.<sup>30</sup>

We can draw two morals here. First, we should always be wary that knowledge rests on certain tacit epistemic assumptions. As we consider what is implied by what we know, we should not forget the tacit assumptions on which our knowledge rests. Tacit assumptions cannot be canceled by appraising what follows from what we know. The second moral is that we do not acquire knowledge without some effort. We can be in a position to know *p* provided we justifiably rely on a set of suitable assumptions *and* we have suitable evidence for *p*. To acquire knowledge, we have to supplement our assumptions with evidence. We do not get knowledge for free, and thus, we do not know that *p* if *p* itself is among (or entailed by) our tacit assumptions.<sup>31</sup>

This is all well and good. But what about epistemic closure? Unfortunately, the “no easy knowledge restriction” conflicts with the standard principle of closure. The conflict between the two arises because standard closure can range over all propositions, even those that constitute the epistemic assumptions that underlie one’s knowledge. What we need, it seems, is a way to shield our epistemic assumptions from the reach of closure, so to say. This can be done by imposing a restriction on the principle of closure so that the principle does not apply to assumptions. But why should we adopt this restriction? Why is this a plausible way to conceive of epistemic closure, and more generally, of fallible human knowledge? In what follows, I shall address these questions by doing two things: first, I will introduce a new notion, namely topics of inquiry; second, I will relate the new notion to assumptions, fallible knowledge, and our resource-bounded nature.

<sup>29</sup> This would violate, for example, clauses (b) and clause (d) from the previous section.

<sup>30</sup> Some authors deny just that. For instance, [Stine \(1976\)](#) writes that ‘one does not know what one takes for granted in normal circumstances’ (or assume, in my terminology). Although I am unable to offer an argument against this view, I can say that, on this view, we would know, by implication, that we are not brains-in-a-vat or we would know other anti-skeptical propositions insofar as they are part of what we take for granted. But claiming that we know anti-skeptical propositions strikes me as *prima facie* implausible; see also footnote 22.

<sup>31</sup> On this score, some authors believe that knowledge is a form of cognitive achievement; it is the result of an epistemic effort. See, among others, [Heller \(1999\)](#), [Vogel \(1999\)](#), [Greco \(2010\)](#) and [Holliday \(Forthcoming\)](#).

## 5 Topics and assumptions

So far I have been discussing what we might call the *statics of knowledge*. I have been discussing *S-knows-that-p* as the state in which *S* stands in the appropriate epistemic relation to *p*. But the statics of knowledge cannot be entirely divorced from the *dynamics of knowledge*—from the way epistemic agents come to acquire knowledge. That we know a proposition must have something to do with how we came to know that proposition. And this is especially true in the case of many empirical propositions whose knowledge we must have acquired at some point. Although my focus remains the statics of knowledge, in this section I will offer a few remarks that pertain to the dynamics of knowledge. These remarks will serve to introduce the notion of a topic of inquiry.

As epistemic agents seek to gain knowledge, they partake in epistemic inquiries. Through the collection of evidence, agents may acquire knowledge of a number of propositions, although this knowledge will always be determined by the inquiry. I suggest that we think of epistemic inquiries as restricted to a *topic of inquiry*, i.e. a set of propositions whose truth value the inquirers seek to determine; I call this the *topic set*. Academic research programs are ways for officially drawing the boundaries of a topic of inquiry. Similarly, in daily affairs, we want to find out the truth about certain propositions because we have a purpose in mind; our purpose will delimit the topic of our inquiry, although this might happen less officially. Either officially or unofficially, explicitly or implicitly, epistemic inquiries will be limited by their topic.<sup>32</sup> How the boundaries of an epistemic inquiry are drawn is a difficult question—one which I will not address here.<sup>33</sup> What matters is that when a boundary is set, implicitly or explicitly, the resulting knowledge will be topic-restricted knowledge. This is in line with talk of scientific knowledge, common-sense knowledge, practical knowledge, mathematical knowledge, knowledge from the physical sciences, etc. These expressions remind us that knowledge might very well be a topic-restricted notion.

In Sect. 7, I will examine more in depth the rationale behind a notion of topic-restricted knowledge, but the above general remarks should suffice for the time being. The notion of a topic of inquiry makes it possible to formulate a topic-restricted version of epistemic closure which can address the puzzle I am examining in this paper, or so I shall argue in the next section. But before I offer a topic-restricted formulation of epistemic closure, I should spell out the relationship between ‘topic of inquiry’ and ‘epistemic assumption’.

<sup>32</sup> Along similar lines, some recent work in linguistics and formal semantics attempts to offer an account of meaning in terms of what is at issue. For instance, Groenendijk (1999) distinguishes between indicative sentences, which provide the *data*, and interrogative sentences, which raise the *issues* in a conversation. The meaning of a sentence is thus understood as its *context change potential*, where a context encodes the data and what is at issue in a conversation. In other words, the meaning of a sentence amounts to how the sentence can change the data given what is at issue in the conversation.

<sup>33</sup> The notion of a topic is related to other notions such as *subject matter*, *aboutness*, *domain*, *framework*. On the notion of aboutness and subject matter, see Yablo (2014). In order to define a topic (or a domain, subject matter, etc.), we can begin by defining a vocabulary for all the objects, entities, concepts we want to refer to (or use) in the course of our inquiry. The topic set, then, will contain at least the *atomic propositions* which can be constructed from our vocabulary. An interesting question is whether all *complex propositions* which can be constructed from the atomic ones will be in the topic set. To answer this, we should not forget that we, as epistemic inquirers, have bounded resources (see Sect. 7 for more clarifications).

On a topic-restricted view of epistemic inquiries, the knowable propositions are within a topic of inquiry. The propositions that lie outside are either entirely disregarded or kept in the background. And here is where assumptions come into play again. For it is plausible to think that an epistemic inquiry, which focuses on a certain topic, takes for granted a number of propositions. These will be the assumptions underlying the inquiry. If you recall the content of Sect. 3, an assumption  $\alpha$  underlines one's knowledge that  $p$  provided (a)  $\alpha$  is needed for one's knowledge that  $p$ ; (b)  $\alpha$  is relied upon without being an item of knowledge; (c)  $\alpha$  is either deeply entrenched or highly probable; and finally, (d)  $\alpha$  alone does not guarantee the truth of  $p$ .

Now, since one's knowledge is often the result of an epistemic inquiry, conditions (a), (b), (c), and (d) can be interpreted *mutatis mutandis* as conditions that describe how an assumption underlies an epistemic inquiry (and not simply one's knowledge). To these conditions, we can add a further one, along the following lines:

(e) *An assumption is part of the background of an epistemic inquiry and lies outside the topic of inquiry.*

More carefully, I should say that an assumption lies outside the topic of inquiry *at a certain stage of the inquiry*, although it may become part of the topic at a later stage. For consider the following conversation:

- Joe, you will be unable to afford a luxury vacation.
- But what if I won the lottery?
- Come on!
- In fact, I won the lottery. [*said by Joe showing his winning ticket*]

At the initial stage, Joe's interlocutor might have very well been in a position to know of Joe's inability *at that time* to afford a luxury vacation. As a ground for her putative knowledge, Joe's interlocutor relied—not without a reason—on the assumption that Joe did not, or will not, win the lottery. In response, however, Joe introduced evidence showing that the assumption in question was false. Between the initial and the final stage, a significant change occurred: though it was initially kept in the background, the assumption that Joe did not win the lottery became part of the topic of inquiry, and it was eventually discarded as false.

It is natural to ask when an assumption can transition from the background to the foreground and become part of the topic set. Merely mentioning it will not do, for consider this conversation:

- Jack, we're going to a luxury resort this weekend. Wanna come?
- I'd like to, but I am short of money. I won't be able to afford it.
- But Jack, you bought a lottery ticket. What if you win?
- So what!?

It seems that Jack is in a position to know that he will be unable to afford a luxury vacation, and I have insisted that this putative knowledge rests on a number of tacit assumptions: Jack will not win the lottery; Jack will not inherit one million dollars; etc. The question 'What if you win?' is an attempt to bring one of these assumptions to the foreground; it is an attempt to make an assumption part of the topic of inquiry.



Yet, Jack dismissed this attempt, and I think he did so appropriately. Why? When can an assumption become part of the topic set? Answer: when an inquirer has sufficiently strong evidence showing that the assumption is false or more likely to be false than previously thought. In the first conversation, Joe had good evidence that his interlocutor's assumption was false; in the second conversation, Jack's interlocutor merely mentioned the possibility that an underlying assumption could be false, yet she had no evidence to that effect. Jack, of course, was well aware that he could win the lottery; he was simply assuming—and with a reason—that he was not going to win. Merely mentioning that an assumption could be false is otiose and pointless, and thus, it cannot suffice to bring the assumption to the foreground.

The above remarks suggest that the topic of inquiry—defined as the set of propositions under inquiry, under focus, or under epistemic assessment—can change dynamically over time as new pieces of evidence are made available to the epistemic inquirers, whether or not the inquirers were actively seeking these pieces of evidence. For instance, if I am looking for a parking spot, I will be concerned with where I am more likely to find a parking spot; with whether it is rush hours or not; with where a less busy street is located; etc. I won't be directly concerned with airplanes flying on top of my head; asteroids which might collide with the earth; etc. But suppose that, as I am driving down a street, a plane suddenly crashes in the middle of the street. Well, I would instantly come to know that there is a plane which has just crashed in the middle of the street. This was not in the topic of inquiry until a few seconds ago, but it now certainly is. Just as an assumption can transition from the background to the foreground (and thus, it can enter the topic of inquiry) because new evidence proves the assumption to be true or false, in a similar manner a proposition which was previously wholly disregarded could suddenly become part of the topic of inquiry because new evidence (almost effortlessly) becomes available to the epistemic inquirer.<sup>34</sup>

More needs to be said about epistemic inquiries and their dynamics. But giving a theory of epistemic inquiries is not my task here. More modestly, I have suggested that the statics of knowledge—our theory of *S-knows-that-p*—can make

<sup>34</sup> An anonymous reviewer rightly points out that the notion of a topic of inquiry might be too restrictive, in the sense that we sometimes acquire new knowledge even without having a topic of inquiry explicitly in focus. One example of this is the airplane crash example mentioned in the text. Another example—which I owe to the reviewer—is this. Suppose I look outside and see the neighbor's orange cat. Indeed, after looking outside I come to know that there's an orange cat outside. This scenario suggests that we can gain new items of knowledge with no inquiry whatsoever, and if there is no inquiry, it's difficult to say what one's topic of inquiry should be. I wish to say two things in response to this objection. First, when we are looking outside the window, we are probably looking for something *that is outside the window*. Thus, the presence or absence of a cat of a certain color seems to be part of a (suitably defined) topic of inquiry, namely the topic of inquiry that answers the question 'what is outside the window?' If we were not looking for anything whatsoever outside the window, we would hardly notice the cat, and thus we would fail to learn that there's an orange cat outside. The second point I want to make is that the notion of a topic of inquiry is not as rigid as it might appear at first. As emphasized in the text, the available evidence—whether the inquirers actively seek it or not—does affect what the topic of inquiry is going to be. At times, the evidence in support of a proposition is so compelling, forceful, and unambiguous that we do automatically come to know the proposition, almost with no effort whatsoever. What the orange cat example suggests is that, at times, it is the evidence we encounter which selects (parts of) our topic of inquiry. When this occurs, we are not fully in control of our topic of inquiry. This is unavoidable: our topic of inquiry is shaped by the interactions between the world and us as epistemic inquirers.

progress by locating it in the larger context of the dynamics of knowledge. One feature of the dynamics of knowledge I have emphasized is that our knowledge is the result of epistemic inquiries whose scope is limited by a topic. Consequently, a theory of *S-knows-that-p* can profit by taking into account the suggestion that our knowledge is topic-restricted knowledge. The next section will apply this suggestion to the principle of epistemic closure itself, yielding a topic-restricted principle of closure.

## 6 Topic-restricted epistemic closure

I am now ready to offer a new formulation of the principle of epistemic closure. If you recall, in order to avoid our puzzle, we needed a restricted principle of closure which did not apply to assumptions. This can be done by restricting closure to the topic of inquiry, given that the inquirers' assumptions, by condition (e) from the previous section, must lie outside the topic. So, we can adopt the following topic-restricted version of epistemic closure:

TOPIC-RESTRICTED EPISTEMIC CLOSURE. Let  $p$  and  $q$  be part of the topic set. If  $S$  knows that  $p$  and  $S$  knows that  $p$  implies  $q$ , then  $S$  knows that  $q$ .<sup>35</sup>

The above principle of epistemic closure provides us with a way out of Hawthorne's puzzle. Suppose Jack knows he will be unable to afford a luxury vacation. On a fallibilist conception of knowledge, Jack knows that *even though* he has not ruled out certain far-fetched possibilities, such as winning the lottery, inheriting one million dollars, etc. In my terminology, Jack knows that he will be unable to afford a luxury vacation, *tacitly assuming* that certain far-fetched possibilities do not obtain. Now, consider the following implication (which Jack accepts and knows): If Jack is unable to afford a luxury vacation, he will not win the lottery. With standard epistemic closure, it follows that Jack knows he will not win the lottery. With a topic-restricted epistemic closure, however, the same conclusion does not follow. The lottery drawing is not part of the topic of inquiry, for the proposition *Jack will not win the lottery* is one of the tacit assumptions. Consequently, topic-restricted epistemic closure does not apply to the proposition *Jack will not win the lottery*.

We can tackle Nozick–Dretske's puzzle in a similar fashion. Suppose Mark knows that he has hands, given some tacit assumptions. Now, by topic-restricted epistemic closure, it does not follow that Mark knows he is not a brain-in-a-vat, because *Mark is not a brain-in-a-vat* is among the epistemic assumptions and is not part of the topic of inquiry.

Let me now turn to a complication. The standard principle of epistemic closure has a contrapositive formulation, as follows:

<sup>35</sup> This formulation echoes Stephen Yablo's proposal that epistemic closure applies so long as there is no change in subject matter (or topic of inquiry, in my terminology) (Yablo 2014). But Yablo's conception of subject matter, I think, is not the same as my notion of topic of inquiry. Further, Yablo does not explicitly draw the connection between epistemic assumptions, a topic-restricted principle of closure, and our resource-bounded nature (see, in particular, Sects. 5 and 7 of this paper).

EPISTEMIC CLOSURE [CONTRAPOSITIVE]. If  $S$  does *not* know that  $q$  and  $S$  knows that  $p$  implies  $q$ , then  $S$  does *not* know that  $p$  either.

Suppose Jack does not know that he will lose the lottery. Now, since *Jack will be unable to afford a luxury vacation* implies *Jack will not win the lottery*, the principle above mandates that Jack does not know he will be unable to afford a luxury vacation either. This seems implausible. Can topic-restricted closure avoid this implausible result when it is used in its contrapositive version? The contrapositive of topic-restricted epistemic closure is as follows:

TOPIC-RESTRICTED EPISTEMIC CLOSURE [CONTRAPOSITIVE]. Let  $p$  and  $q$  be part of the topic set. If  $S$  does *not* know that  $q$  and  $S$  knows that  $p$  implies  $q$ , then  $S$  does *not* know that  $p$  either.

In order to see what follows from a topic-restricted formulation, in its contrapositive version, we should not forget that any epistemic inquiry has a topic. Let us work with the propositions *Jack will not win the lottery* and *Jack will be unable to afford a luxury vacation*. We should distinguish two possibilities here. The first possibility is that at least one of the two propositions is not in the topic set. If so, the principle of closure would not apply to them, and no problem would arise. The other possibility is that both propositions are part of the topic set. If so, our new principle of closure would yield a result similar to that of standard epistemic closure, namely that Jack does not know he will be unable to afford a luxury vacation. Yet, this result is now perfectly plausible. How so? Since both propositions are in the topic set, for one reason or another, the proposition *Jack will not win the lottery* won't be among the tacit assumptions. So it should be no surprise that Jack does not know he will be unable to afford a luxury vacation. This phenomenon shows that the fewer propositions are among the tacit assumptions, the fewer propositions we know. The less we assume, the less we know. This is not implausible; that's how knowledge should behave.

I hope the reader is now convinced that a topic-restricted formulation of epistemic closure can help us avoid the Hawthorne and the Nozick–Dretske puzzle. But now a question arises. A topic-restricted principle of epistemic closure can avoid trouble because it only applies to propositions that are in the topic set. Yet, some might think that closure should apply *beyond the boundaries of an epistemic inquiry*. Unfortunately, this is not possible. But why? And isn't restricting epistemic closure too costly a price to pay for our theory of knowledge? In what follows, I will argue that even though giving up unrestricted closure appears painful and costly to many epistemologists, a restricted principle of closure is more in tune with the type of epistemic agents we are; and what we want is a theory of knowledge that is about *us*. To substantiate this claim, I will show that topic-restricted closure is an adequate principle for resource-bounded and fallible agents such as ourselves.

## 7 Why a topic-restricted principle of closure?

I shall begin by defending a view of knowledge which I'd like to call *topic-restricted knowledge*. Take one of the most powerful sources of knowledge: science. Scientific research began to yield powerful results when it compartmentalized itself; when it

restricted itself to investigating a limited number of propositions. Arguably, the success of scientific knowledge is a consequence of abandoning the dream of universal knowledge. This is no more than a conjecture, and I leave it to historians of science to scrutinize it further. The fact is that science, and academic research more generally, is structured in a topic-restricted way (or domain-restricted, discipline-specific—pick the terminology you prefer).

Furthermore, our ordinary, everyday knowledge can also be viewed as topic-restricted. When we claim to know we have hands, we do not worry about the micro-physical world consisting of atoms and subatomic particles. Rather, we focus on what appears to us through the senses. Our everyday knowledge is topic-restricted in the sense that we do not take into consideration all the facets of reality; we bracket some and we focus on others.

(Incidentally, topic-restriction is not confined to scientific and ordinary knowledge. Consider our actions. As we walk down a street, we are aware of other individuals, cars, obstacles, pickpockets, etc., but we tend to bracket certain things, such as birds flying around, the shapes of the clouds, chemical reactions at the molecular level, etc. As we walk, we take some aspects into consideration, and leave out some others. If, as some believe, there is a close connection between knowledge and action, the topic-restriction of our action-oriented considerations would be a further reason to think that knowledge is topic-restricted).

Let us now ask: Why are scientific and ordinary knowledge topic-restricted in the sense I've just illustrated? I want to suggest that human knowledge tends to be topic-restricted *because* human epistemic resources are bounded. That our cognitive resources are limited, after all, must affect the scope of our epistemic inquiries. As we investigate the truth of certain propositions, we should strive to collect evidence about them, but since our resources are limited, we are unable to collect evidence about *any* proposition which could possibly be related to the propositions under investigation. The result is that we will focus on some propositions and leave others in the background. So, what we come to know will be topic-restricted. If this is correct, a topic-restricted conception of knowledge would be perfectly adequate for resource-bounded agents such as ourselves; a topic-restricted principle of epistemic closure, then, would be the right principle for us.

At this point, some might object that an appeal to our limited cognitive resources undermines my proposed solution to our puzzle, because the puzzle concerns knowledge as such, irrespective of us having limited or unlimited resources. For, as Fred Dretske has noted, although the ideally astute logician is able to appraise all the logical consequences of what he knows, he is not in a position to know that he is not a brain-in-vat by knowing that he has hands. And if the Nozick–Dretske puzzle arises even for the ideally astute logician, it would seem that invoking our resource-bounded nature can do nothing for solving the puzzle.

I think the objection is misplaced, because the ideally astute logician should not be the paradigm of the ideal, unbounded agent. In fact, our resources are limited in at least two senses: for one, we have limited intellectual and inferential capabilities, and for another, we have limited empirical and information-processing capabilities. We are intellectually bounded in the sense that we cannot appraise all logical and inferential connections among all possible propositions. In Dretske's terminology, we are not

ideally astute logicians and reasoners. But another side of our resource-bounded nature is that our capabilities to acquire and process information about the environment are limited. For example, our sensory apparatus can only select and process some inputs and has to neglect others. And further, we can only and solely acquire and process information within a perspective: we are unable to perceive and understand reality from the point of view of the absolute, as it were. The ideally astute logician, then, is only partly an ideal agent. The truly ideal, unbounded agent, in contrast, is one who is endowed with both infinite intellectual capabilities and infinite empirical capabilities. And if the ideally astute logician is not the paradigm of the unbounded agent, the fact that the Nozick–Dretske puzzle arises even for the ideally astute logician is no objection to invoking our resource-bounded nature as a way to address the puzzle.<sup>36</sup>

All in all, I have suggested that we can solve our puzzle provided we formulate a principle of epistemic closure that is in line with our resource-bounded nature. Since we are resource-bounded agents, the topic of our epistemic inquiries cannot include every proposition because our limited intellectual and empirical capabilities force us to focus on some propositions and neglect others. A topic-restricted view of knowledge and epistemic closure, then, is well-suited for resource-bounded agents such as ourselves.

## 8 Contexts versus topics of inquiry

My argument is now complete. In this final section, I compare my topic-restricted account of knowledge with epistemic contextualism. Epistemic contextualism is the view that knowledge attributions are relative to context. This means that sentences like ‘*S* knows that *p*’ can be true or false depending on the context in which they are uttered. What is a context? Very roughly, a context encodes features of the attributor’s psychology, values, and practical interests.<sup>37</sup> What makes epistemic contextualism a controversial and interesting thesis is that a context typically does not encode differences in the epistemic agent’s available evidence. That the truth-values of our knowledge attributions vary depending on the agent’s available evidence, after all, should be uncontroversial.<sup>38</sup> But it is certainly more controversial to say that the truth-values of our knowledge attributions vary depending on non-evidential factors, such as the attributor’s interests and practical stakes in a situation. So, an important ingredient

<sup>36</sup> Here one might object that we have no guarantee that the Nozick–Dretske puzzle would not arise even for the truly ideal, unbounded agent who is endowed with infinite intellectual and empirical capabilities. Here is a brief, tentative response. The Nozick–Dretske puzzle arises provided we adopt a fallibilist epistemology according to which an epistemic agent can know things without ruling out all the possibilities of error and deception. Now, I do not think that a fallibilist epistemology is adequate to theorize about the ideal agent who is endowed with infinite intellectual and empirical capabilities. The agent with infinite empirical and intellectual capabilities should be able to rule out all possibilities in which a given proposition is false. But if we admit that the ideal agent can rule out all possibilities of error, the Nozick–Dretske puzzle would no longer arise.

<sup>37</sup> This should make clear that I am referring to so-called *attributor contextualism*, according to which the truth-values of knowledge attributions depend on the attributor’s purposes, goals, expectations, practical interests, etc. See, among others, [Cohen \(1986\)](#) and [DeRose \(2002\)](#).

<sup>38</sup> The agent’s evidence, or lack thereof, affects the agent’s justification, which, in turn, affects whether the agent is in a position to know or not.

of contextualism is that, without any change in  $S$ 's evidence, it can be true to say that  $S$  knows that  $p$ , and also true to say that  $S$  does *not* know that  $p$ , because each knowledge attribution is associated with a different epistemic context. What accounts for the variability is that, in the two contexts, different epistemic standards are at work, one standard being more evidentially demanding than the other.

Epistemic contextualism is a view about knowledge attributions, not a view about knowledge as such.<sup>39</sup> In what follows, I shall speak somewhat sloppily as though epistemic contextualism were simply a view about the contextuality of knowledge. I will do so for ease of exposition. (If I were to treat epistemic contextualism as only a view about knowledge attributions, and not at all as a view about knowledge, no comparison would be possible with my account, which is primarily a view about knowledge).

What interests me here is that epistemic contextualism offers a putative solution to the puzzle I've discussed in this paper. If you recall, Hawthorne's version of the puzzle had the following form:

- (H-1) Jack knows that he will be unable to afford a luxury vacation.
- (not-H-2) Jack does *not* know that he will not win the lottery.
- (H-3) Jack knows that *Jack will be unable to afford a luxury vacation* implies *Jack will not win the lottery*.
- (H-2) Jack knows that he will not win the lottery.

And Nozick–Dretske's version had the following form:

- (ND-1) Mark knows that he has hands.
- (not-ND-2) Mark does *not* know that he is not a brain-in-a-vat.
- (ND-3) Mark knows that *Mark has hands* implies *Mark is not a brain-in-a-vat*.
- (ND-2) Mark knows that he is not a brain-in-a-vat.

According to the contextualist, Jack can certainly know he will be unable to afford a luxury vacation and Mark can know he has hands. The proviso here is that we are dealing with an ordinary epistemic context, one that is not too demanding in terms of the evidence required for knowledge. Call this a *low standard context*. If (H-1) and (DN-1) both hold, by applying unrestricted epistemic closure, Jack would know he lost the lottery (even before learning about the outcome of the lottery drawing) and Mark would know he is not a brain-in-a-vat. These are no longer counterintuitive results—the contextualist will argue—because Jack and Mark know such things relative to a context that is associated with a low or ordinary epistemic standard.

At the same time, the contextualist can accommodate the intuition that Jack does not know he lost the lottery or that Mark does not know he is not a brain-in-a-vat. This intuition can be vindicated by appealing to a context that is associated with a more demanding standard in terms of what can count as knowledge. Call this a *high standard context*. Now, if Jack does not know he lost the lottery and if Mark does not know he is not a brain-in-a-vat, then Jack would not know that he cannot afford a luxury vacation, and Mark would not know that he has hands. How so? This is a

<sup>39</sup> One could say that insofar as knowledge attributions can be true or false, contextualism is also, though indirectly, a view about knowledge. I do not enter into this difficult question here.

consequence resulting from the application of the closure principle. But Jack's and Mark's failure to know things which they would ordinarily know is not problematic—so the contextualist argues—because we are now dealing with a high standard context, and in a high standard context, we would expect Mark not to know that he has hands and Jack not know that he will be unable to afford a luxury vacation.

We can now see a first difference between my account and epistemic contextualism. The contextualist has to say that, in a low standard context, Jack knows he will be unable to afford a luxury vacation *and* that Jack knows he will not win the lottery, or alternatively, the contextualist has to say that, in a high standard context, Jack fails to know both things. (Similarly, the contextualist has to say that, in a low standard context, Mark knows that he has hands *and* that Mark knows he is not a brain-in-a-vat, or alternatively, the contextualist has to say that, in a high standard context, Mark fails to know both things.) In order to maintain the closure principle, the contextualist cannot say that, in the same context, Jack knows he will be unable to afford a luxury vacation and does not know he will not win the lottery. (Similarly, in order to maintain the closure principle, the contextualist cannot say that, in the same context, Mark knows he has hands and does not know he is not a brain-in-a-vat.) Instead, my account does exactly what the contextualist cannot (or does not want to) do. The trick is to replace 'epistemic contexts' with 'topics of inquiry'. Let me explain.

Unlike the contextualist account, my account manages to satisfy three desiderata while holding fixed the same topic set. The three desiderata are, first, to accommodate the intuition that Jack knows he will be unable to afford a luxury vacation (and that Mark knows he has hands); second, to accommodate the intuition that Jack does *not* know he will not win the lottery (and that Mark does *not* know he is not a brain-in-a-vat); and third, to secure that some version of closure still holds. Epistemic contextualism can satisfy the first two desiderata across different contexts, while my account can do so within the same topic of inquiry. The difference is that epistemic contexts can apply to any proposition whatsoever, whilst topics of inquiry are demarcated "from the inside" by those propositions which are at issue and they are demarcated "from the outside" by those propositions which are simply assumed. Consequently, if some propositions lie outside the ongoing topic of inquiry, the epistemic agent will fail to know—relative to that topic—the propositions that are off topic.<sup>40</sup> In contrast, an epistemic agent can in principle know any proposition relative to a fixed epistemic context. In short, one difference between my account and epistemic contextualism is that epistemic contexts are proposition-neutral (in the sense that any proposition can be at issue within a given epistemic context), while topics of inquiry are proposition-sensitive (in the sense that some proposition are part of the topic of inquiry and some aren't).<sup>41</sup>

Another noteworthy difference is that, on the contextualist picture, it is the knowledge *attributor* (or group of attributors) who fixes the epistemic context against which statements such 'S knows that *p*' are evaluated. On my account, instead, it is the epistemic *inquirer* (or group of inquirers) who fixes the topic of inquiry. At least, this

<sup>40</sup> This does not mean that the agent will fail to know *simpliciter* the propositions that are off topic. Consider two different topics, *T*1 and *T*2. It might well be that the epistemic agent can know certain propositions relative to *T*1, while she fails to know them relative to *T*2 because they are off topic.

<sup>41</sup> For a similar, proposition-sensitive view of topics of inquiry, see Chapter 4 of Lawlor (2013).

is how I've been formulating my account in this paper. What determines the topic of inquiry are the propositions that are at issue, and it is the epistemic inquirer who seeks to determine the truth-value of the propositions at issue while taking others for granted. I am, however, open to a version of my account in which the topic of inquiry is affected not only by the epistemic inquirer, but also by the knowledge attributor. I leave it for future work to explore what such an account might look like.

Let me now focus on a similarity. A feature of the contextualist picture is that the principle of epistemic closure is relativized to contexts, and it does not apply across contexts. If it did apply across contexts, we would have our puzzle back. This is a commonality with my account. The contextualist holds that the principle of closure applies within a given *context*, and similarly, I hold that the principle applies within the *topic* of a given epistemic inquiry. Recall that a topic of inquiry is partly determined by the background assumptions of the inquiry, in the sense that the assumptions lie outside the topic of inquiry. So, on my account, *S* might be in a position to know that *p*, given a certain topic and certain epistemic assumptions, and also, *S* might *not* be in a position to know that *p*, given a certain *other* topic and epistemic assumptions. For instance, if *Mark is not a brain-in-a-vat* were not among the tacit assumptions, Mark would *not* be in a position to know that he has hands. But if *Mark is not a brain-in-a-vat* were added to the set of tacit assumptions, Mark would presumably be in a position to know that he has hands. This “assumption variability” has a seemingly contextualist flavour. But let us look more closely at how assumptions can vary.

Assumptions can vary in ways that are *evidentially driven*. For example, an assumption should be dropped if new evidence comes in showing that the assumption in question is false or not very likely (see Sect. 5). Now, if assumptions were to change only in ways that are evidentially driven, my proposed account would not be contextualist at all. For, recall that a contextualist is committed to the idea that, without any change in the available evidence for the subject *S*, epistemic standards can still vary.

My account would become more similar to contextualism if assumptions were allowed to vary in ways that were not evidentially driven—and I am open to this possibility. This may occur when an assumption is added or dropped depending on the topic of inquiry, and regardless of the available evidence. Suppose a group of people decided to investigate whether we are brains-in-a-vat, and whether we have hands, arms, legs, and other body parts. Now, *relative to that topic of inquiry*, they would not know whether we are brains-in-a-vat, and since topic-restricted closure applies, they would not know whether we have hands either. Instead, in normal circumstances, we know that we have hands because our topic of inquiry does not include far-fetched skeptical scenarios, or so I claim. This is a somewhat contextualist conclusion: knowledge can come and go, without any change in the evidence, depending on the topic of inquiry and on the underlying assumptions we decide to live with.<sup>42</sup>

<sup>42</sup> To be sure, there is a way to resist this conclusion, namely by saying that there cannot be a topic of inquiry that includes whether we are brains-in-a-vat. A topic of inquiry—some might argue—is largely determined by the evidence that the inquirers reasonably expect to obtain. If so, it would be impractical for a group of epistemic inquirers to investigate propositions about which they have no reasonable hope to gather suitable evidence. I do not know if this line of argument is a defensible one, and I leave it as an open question whether it is or not.



## 9 Conclusion

I have put forward a topic-restricted formulation of the principle of epistemic closure that avoids the puzzle I've described at the beginning of this paper. What does the trick is that the principle of closure thus formulated does not range over any proposition, but only over those that belong to the topic set. In particular, my argument weaves together two insights about the nature of human knowledge:

1. Knowledge rests on assumptions, and not simply on evidence; and
2. Epistemic inquiries are limited by a topic, so knowledge must be topic-restricted.

I have argued that if knowledge is topic-restricted, epistemic closure must be topic-restricted, as well. This means that epistemic closure does not apply to the propositions that lie outside the topic of inquiry. And since among the propositions that lie outside the topic of inquiry are the underlying assumptions of the inquiry, epistemic closure cannot apply to the inquiry's assumptions. (This allows us to address both the Nozick–Dretske puzzle and the Hawthorne puzzle insofar as propositions such as *Mark is not a brain-in-a-vat* or *Jack will not win the lottery* count as assumptions which are not within closure's scope of application.) Finally, I have argued that our epistemic inquiries, our knowledge, and the principle of epistemic closure are topic-restricted *because* our cognitive resources are limited.

**Acknowledgments** I would like to thank Krista Lawlor, Rahul Chaudhri, Carlos Nunez, Mark Crimmins, Samuel Asarnow, Peter Hawke, and two anonymous reviewers for helpful comments.

## References

- Achinstein, P. (1978). Concepts of evidence. *Mind*, 87(345), 22–45.
- Adams, F., Barker, J. A., & Figurelli, J. (2012). Towards closure on closure. *Synthese*, 188(2), 179–186.
- Barker, J. A., & Adams, F. (2010). Epistemic closure and skepticism. *Logos and Episteme*, 1(2), 221–246.
- Cohen, S. (1986). Knowledge and context. *Journal of Philosophy*, 83(10), 574–583.
- DeRose, K. (1996). Knowledge, assertion and lotteries. *Australian Journal of Philosophy*, 76, 568–580.
- DeRose, K. (2002). Assertion, knowledge and context. *The Philosophical Review*, 111(2), 167–203.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (1971). Conclusive reasons. *Australian Journal of Philosophy*, 49, 1–22.
- Easwaran, K. (2011). Bayesianism I: Introduction and arguments in favor. *Philosophy Compass*, 6(5), 312–320.
- Feldman, R. (1995). In defense of closure. *Philosophical Quarterly*, 45, 487–494.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge: Cambridge University Press.
- Groenendijk, J. (1999). The logic of interrogation: Classical version. In T. Matthews & D. Strolovitich (Eds.), *Proceedings from semantics and linguistic theory IX*. Ithaca, NY: Cornell University Press.
- Harman, G., & Sherman, B. (2004). Knowledge, assumptions, lotteries. *Philosophical Issues*, 14, 492–500.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Heller, M. (1999). Relevant alternatives and closure. *Australian Journal of Philosophy*, 77(2), 196–208.
- Holliday, W. H. (2014). Epistemic closure and epistemic logic I: Relevant alternatives and subjunctivism. *Journal of Philosophical Logic*. doi:10.1007/s10992-013-9306-2.
- Holliday, W. H. (Forthcoming). Fallibilism and multiple paths to knowledge. *Oxford Studies in Epistemology* 5.
- Kvanvig, J. L. (2006). Epistemic closure principles. *Philosophy Compass*, 1(3), 256–267.
- Lawlor, K. (2005). Living without closure. *Grazer Philosophische Studien*, 69(1), 25–50.

- Lawlor, K. (2013). *Assurance: An Austinian view of knowledge and knowledge claims*. Oxford: Oxford University Press.
- Lewis, D. (1996). Elusive knowledge. *Australian Journal of Philosophy*, 74, 549–567.
- Longino, H. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton, NY: Princeton University Press.
- Nelkin, D. N. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review*, 109, 373–409.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge: Harvard University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pryor, J. (2012). When warrant transmits. In A. Coliva (Ed.), *Wittgenstein, epistemology, and mind: Themes from the philosophy of Crispin Wright*. Oxford: Oxford University Press.
- Roush, S. (2006). *Tracking truth: knowledge, evidence, and science*. Oxford: Oxford University Press.
- Sharon, A., & Spectre, L. (2013). Epistemic closure under deductive inference: What is it and can we afford it? *Synthese*, 190(14), 2731–2748.
- Sherman, B., & Harman, G. (2011). Knowledge and assumptions. *Philosophical Studies*, 156, 131–140.
- Stine, G. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29, 249–261.
- Tucker, C. (2010). When transmission fails. *Philosophical Review*, 119, 497–529.
- Vogel, J. (1990). Are there counterexamples to the closure principle. In M. Roth & G. Ross (Eds.), *Doubling: Contemporary perspectives on skepticism* (pp. 13–27). Dordrecht: Kluwer.
- Vogel, J. (1999). The new relevant alternatives theory. *Nous*, 33, 155–180.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (2004). Warrant for nothing (and foundation for free). *Aristotelian Society Supplementary*, 78, 167–212.
- Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.