Folk Epistemology and Epistemic Closure*

Tim Kraft (Regensburg, tim.kraft@ur.de)

Alex Wiegmann (Göttingen, alex.wiegmann@psych.uni-goettingen.de)

Abstract

According to the epistemic closure principle, if someone knows some proposition *P* and also knows that *P* entails *Q*, she knows *Q* as well. This principle is often defended by appealing to its intuitiveness. But only recently was epistemic closure put to the empirical test: Turri ran experiments in which closure is (allegedly) violated in folk knowledge ascriptions surprisingly often. We disagree with this diagnosis. It is by no means obvious which experimentally testable hypothesis proponents of epistemic closure should accept. In this paper we formulate and test a hypothesis different from Turri's. We argue that our hypothesis is more apt for empirically testing epistemic closure. In a series of experiments we manipulated the strength of entailment between two propositions and found that the stronger the entailment was, the lower was the proportion of participants who violated closure, indicating that folk knowledge ascriptions are sensitive to entailment. We conclude that closure is a principle of folk epistemology after all.

*Both authors contributed equally to this paper and are listed in alphabetical order. We are grateful for the help and comments we received from Christina Botros, David Deml-Müller, Tobias Klauk (and the students of his experimental philosophy course), Hans Rott (and the other members of his research seminar), three anonymous reviewers for and the editors of *Oxford Studies in Experimental Philosophy* and an audience at the 2015 conference of the *Experimental Philosophy Group Germany* at Bochum.

Keywords

experimental epistemology, folk epistemology, epistemic closure, Harman-Vogel cases, Turri

1 Introduction

The epistemic closure principle – if someone knows some proposition *P* and also knows that *P* entails *Q*, she knows *Q* as well – is a principle about knowledge that is easily stated and grasped, is a crucial ingredient to sceptical arguments and other knowledge paradoxes and is accepted by a vast number of epistemologists independently of the side on which they stand on the many dividing lines within epistemology – i.e., no matter whether they are internalists or externalists, infallibilists or fallibilists, invarantists or contextualists and so on. Since closure plays a central role in many epistemological debates, it has given rise to several controversies: How should it be formulated precisely? Does closure lead to scepticism and must scepticism rely on it? How should proponents of closure respond to (alleged) counterexamples of various sorts? Despite these controversies proponents of closure often limit their case for closure to pointing out that closure is a highly intuitive and natural principle about knowledge. Only recently, however, has this argument for epistemic closure been put to the test. The methods of experimental philosophy are especially well suited to test such claims, but the results so far do not speak in favour of closure's intuitiveness (Turri, 2015a; 2015b). To the contrary, folk ascriptions of knowledge violate closure in rates that are at odds with the claim that closure is a highly intuitive principle. If robust, the consequences of the experimental results are far-reaching. On the one hand, a major argument for closure would be blocked and a re-evaluation of closure would be called for. On the other hand, the

sceptical argument (or at least one prominent version of it) and other epistemic paradoxes could no longer be claimed to arise solely from natural assumptions about knowledge and would thus lose much of its intuitive grip.

Before accepting these consequences, however, we should look closely at the experimental studies against the intuitiveness of closure. In this paper we both offer a theoretical critique of Turri's study and present our own experimental study. The main conclusion of our paper is that closure *is* a principle of folk epistemology: Folk knowledge ascriptions by and large adhere to the closure principle. But another important lesson of our paper is a methodological one that has wider application: Epistemological theses are difficult to test experimentally because it is by no means clear how to 'translate' them into experimentally testable hypotheses. Instead of polling participants about single cases we should identify ways in which acceptance of an epistemic principle would influence knowledge ascriptions and test whether this influence can actually be observed. Closure is, therefore, an interesting test case for illustrating both the prospects and perils of experimental epistemology.

We proceed as follows: In section 2 we survey the relevant literature on closure and offer a theoretical critique of Turri's experimental results. In sections 3 to 6 we report a series of four experiments that taken together defend closure as a principle of folk epistemology. In section 7 we summarise and discuss our results.

2 Survey of the literature

2.1 Closure of knowledge

Although the precise formulation of the principle of closure of knowledge under known entailment – also called 'epistemic closure', hereafter simply 'closure' – is controversial[1], the core idea behind closure is easy to state: If someone knows some proposition *P* and also knows that *P* entails *Q*, she knows *Q* as well. For example, if you know that Socrates is in this room and is relaxed (and you know the relevant entailment), you also know that at least one person in this room is relaxed. Or, if you know that your daughter is in the kindergarten (and you know the relevant entailment), you also know that you have a daughter. Or, if you know that the bird in the garden is a goldfinch (and you know the relevant entailment), you also know that it is not a canary. To express this idea in more general terms, whenever we recognise that a given claim is entailed by something we already know, we can and do gain knowledge of that further claim.

1 For recent surveys see Kvanvig (2006), Baumann (2011), Luper (2016). *First*, 'single-premise' closure and 'multi-premise' closure need to be distinguished. According to multi-premise closure, if someone knows a *set* of propositions *S* and also knows that *S* entails proposition *Q*, then she knows *Q* as well. Because of its role in the lottery and the preface paradox multi-premise closure is rejected more often than single-premise closure (cf. Williamson, 2000, pp. 116–117; Hawthorne, 2004, pp. 46–50). *Second*, the formulation in the main text suffers from the 'not putting one and one together' problem and the 'losing belief' problem. Someone may know *P* and the relevant entailment, but may simply not form the belief that *Q*. Alternatively, she may know *P*, but lose her belief that *P* as soon as she becomes aware of the entailment. To avoid counter-examples like these either the antecedent has to be strengthened ("and competently deduces *Q* from *P*") or the consequent has to be weakened ("is in a position to know *Q*") (cf. Hawthorne, 2004, ch. 1).

We do not offer these examples merely for the sake of illustrating the closure principle, but also because that is more or less the central argument in favour of closure: Reflecting on examples like these, or so it is often argued on behalf of closure, suffices to see how obvious, self-evident, natural or intuitive closure is.[2] Although other arguments in favour of closure are rare, the argument from intuitiveness is not the only one. One argument is based on closure's explanatory power: It can explain other features of knowledge in a systematic fashion. For example, closure can explain how inferential knowledge is possible or how knowledge claims can be defeated by error-possibilities that the subject is unable to rule out. However, since theoretical arguments in favour of closure are rarely given, the debate is understood best as crediting closure with a default status: Closure is a highly intuitive principle and epistemologists should accept it even in the absence of additional arguments in its favour *unless* there are very good reasons to give it up. In the current state of debate the burden of proof is put onto the shoulders of opponents of closure. This state of the debate opens a door for experimental research. Since the intuitiveness of closure is so crucial for proponents of closure, it would be instructive to know whether this claim can be vindicated experimentally. Of course, if folk knowledge ascriptions do not agree with closure, that is no conclusive

2 Two claims need to be distinguished here: To say that closure is an intuitive (or natural and so on) principle can mean either that the principle itself is intuitive or that its instances are intuitive. In the main text we assume that the latter claim is at issue. To say that closure is an intuitive principle is shorthand for saying that closure is supported by intuitions about a wide variety of cases. This seems to be a tacit assumption in debates about closure and is also made by Turri (2015a) who concludes that closure is not intuitive because intuitions about particular cases violate closure. Unless the tacit assumption is made, the conclusion does not follow. For intuitions are not closed under known entailment. It is possible to find a general statement intuitive, but not all of its instances and *vice versa*. For example, intuitions about free speech or moral luck principles can come apart from intuitions about particular cases (Shoah denial, drunk drivers who by luck do or do not hit someone).

argument against closure. Nevertheless it would put the burden of proof back on the shoulders of proponents of closure. If, on the other hand, it turns out that folk knowledge ascriptions *do* agree with closure, that is not a conclusive argument for closure either, but would again shift back the burden of proof.

2.2 Objections and counterexamples to closure

Before turning to experimental studies we give an overview of the objections and counterexamples to closure already to be found in the literature. Opponents of closure have argued against it both with counterexamples to closure (thereby challenging the intuitiveness of closure) and with theoretical objections. The theoretical objections to closure are of three kinds. *First*, there are analyses of knowledge that do not guarantee closure, e.g. Nozick's truth-tracking account (Nozick, 1981). *Second*, it has been argued that giving up closure is the best response to scepticism (Dretske, 1970; Nozick, 1981). *Third*, it has been argued that since no other propositional attitude besides knowledge is closed under known entailment, not even other factive propositional attitudes or the modes of knowledge, it would be surprising if knowledge was an exception (Dretske, 2005, p. 13; Turri, 2015a, pp. 2f.).[3]

In this paper we leave these theoretical objections against closure aside. No matter how convincing one find these challenges, it would useful to know more about potential counterexamples and their intuitiveness. This is interesting both for its own sake and for

3 To clarify the objection it is useful to distinguish between a propositional attitude, *f*, being closed under known entailment and its being closed under *f*-ed entailment. Knowledge might be the only propositional attitude closed under known entailment, but it is arguably not the only *f* that is closed under *f*-ed entailment. For example, surprise is not closed under known entailment, but seems to be closed under *surprising* entailment. Several other propositional attitudes, *f*, are *vacuously* closed under *f*-ed entailment, for example, seeing and regret. Presumably, it is impossible to see or regret an entailment.

estimating how revisionary it would be to give up closure. There are at least four kinds of alleged counterexamples to closure (cf. Table 1).[4] Our list of potential counterexamples to closure serves two purposes. On the one hand, it paves the way for experimentally investigating closure by identifying cases that can be used in experimental studies. On the other hand, it shows that it is an exaggeration that epistemologists universally accept closure as highly intuitive.[5]

4 Other classifications have been proposed as well: The first and second (and occasionally even the third) kinds of cases are sometimes grouped together when discussing the "problem of easy knowledge" (Cohen, 2002). Nagel (2011) argues that the second and fourth kinds of cases should be grouped together, whereas Vogel (1990) argues that they are distinct kinds of cases.

5 Philosophers' claims about what the majority of philosophers believes seem to be rather untrustworthy. For example, reading the literature on fake barns one is easily left with the impression that philosophers almost universally deny knowledge in this case. There is a growing body of evidence that this is false, cf. Horvath & Wiegmann (2016).

Type of case	Entailing item	Entailed item
Limiting propositions (Moore)		
External world	this is a hand	there are external objects
Past	I have slept for several hours	the universe is older than five minutes
Numbers	7 is prime	numbers exist
Perceptually indistinguishable propositions (Dretske)		
Wall	this wall is red	this isn't a cleverly illuminated white wall
Zebra	the animal in the pen is a zebra	the animal in the pen isn't a cleverly disguised mule
Dogmatism (Kripke)		
Stove	I turned off the stove	evidence I didn't is misleading
Ravens	all ravens are black	all alleged sightings of non-black ravens have been non-veridical
Lottery-type propositions (Harman and Vogel)		
Safari (Harman-Hawthorne)	I can't afford a safari this year	I won't win the lottery
Car theft (Vogel)	my car is where I parked it an hour ago	my car hasn't been stolen
President (Vogel)	NN is the current president	NN hasn't suffered a lethal heart attack within the last five minutes

Table 1: Counterexamples to closure

The *first* group of potential counterexamples consists of limiting propositions entailed by ordinary knowledge claims (cf. Dretske, 2003; Luper, 2016, also called 'cornerstones', 'heavyweight propositions' etc.). Several authors argue, or at least suspect, that knowing that 7 is prime, that this is a hand or that one has slept for several hours does not entail knowing that there are numbers, physical objects or an extended past. However, intuitions with respect to these cases vary. While many agree that *arguments* like "this is a hand; therefore, external objects exist" are question-begging (or defective in some other way), this is not by itself an objection to closure. For closure only requires that one cannot know the entailing proposition without knowing the entailed proposition, but not that knowledge of the entailed proposition can be *based solely* on knowing the entailing proposition.

A *second* group of potential counterexamples is based on cases of ordinary perceptual knowledge. Several authors argue, or at least suspect, that knowing that the animal is a zebra or that the wall is red does not entail knowing that the animal is not a cleverly disguised mule or that the wall is not a cleverly illuminated white wall (Dretske, 1970, 1015f.; Cohen, 2002). What these cases have in common is that they introduce an error-possibility that cannot be ruled out due to its being perceptually indistinguishable. Again, intuitions with respect to these cases vary. Several authors have argued that although one's perception *alone* does not rule out that the animal is a disguised mule or that the wall is illuminated in some unexpected way, perception together with background knowledge (about zoos and walls in one's vicinity) suffices to know the entailed proposition (cf. Vogel, 1990, p. 14; Pritchard, 2012, pp. 77–81).

A *third* group of potential counterexamples derives from cases that seem to imply that closure entails dogmatism (Kripke, 2011, pp. 43–45; Harman, 1973, pp. 147–149, Nozick, 1981, pp. 236–239). Several authors argue, or at least suspect, that knowing that all ravens are black or that one has turned off the stove does not entail knowing that all evidence to the contrary must be misleading. That is supposed to be true in general: In almost all cases of knowledge one simply cannot know *in advance*, i.e. without looking at it, listening to it or reflecting on it, that evidence to the contrary must be misleading. As puzzling and interesting as these cases are philosophically, they are difficult to test experimentally because of the technical notions ('misleading evidence', 'non-veridical') needed to formulate them. An experimental investigation of these cases must wait for another occasion.

A *fourth* and final group of potential counterexamples consists of lottery-type propositions entailed by ordinary knowledge claims, also known as Harman-Vogel cases. These cases are inspired by Kyburg's (1961) lottery paradox, but add a new twist to it. Is it possible to know

(or believe rationally) of a particular ticket in a fair lottery with a very high number of tickets that it will lose? While Kyburg's paradox derives a contradiction from answering this question positively, the Harman-Vogel cases arise from answering it negatively. Harman (1973, pp. 160f.) noted that if one cannot know that one's lottery ticket will lose, one cannot know various propositions entailing it either. For example, nobody knows that they will not be able to afford a safari this year, unless they know that they will not win the lottery. More examples of this kind are due to Vogel (1990, pp. 15f., pp. 20f.): Nobody knows that Obama is the current president of the United States unless they know that he did not die from a lethal heart attack within the last five minutes. Nobody knows that their car is where they left it unless they know that it has not been stolen. What these examples have in common is that an ordinary proposition entails a lottery-type proposition, i.e. a proposition that is very likely true, but also member of a set of equiprobable propositions one of which must be (or very likely is) false. For example, that the car of an arbitrary person has not been stolen is as likely as that another arbitrary person's car has not been stolen, but at the same time we are quite certain that recently someone's car has been stolen. Or, an arbitrary person's not having died from a heart attack the last five minutes is as likely as another arbitrary person's not having died that way, but at the same time we are quite certain that unfortunately someone or other has died that way within the last five minutes. Faced with lottery-type cases all available options are *prima facie* unwelcome: *Either* we can know after all that, for example, a particular car has not been stolen (counterintuitive, how could one possibly know this?) *or* we must give up our claims to know where the car is (counterintuitive, leads to scepticism) *or* we must give up closure (counterintuitive, sounds absurd that we can know the location of a

car without knowing whether it has been taken away). Yet, given how unpalatable the other two options are, giving up closure may not be as unattractive as it initially seemed.

To sum up, although its intuitiveness is usually taken to be the central argument in favour of closure, the literature on closure contains a wealth of potential counterexamples. Arguably the hardest challenge to closure are the fourth kind of counterexamples, lottery-type cases. A common response to these potential counterexamples is to acknowledge that they give rise to recalcitrant intuitions while also offering an account of those examples and our intuitions about them that is consistent with closure (Vogel, 1990; Hawthorne, 2004; Douven, 2007; Roush, 2010; Nagel, 2011). In this paper we do not discuss the merits of these proposals, but investigate folk intuitions about knowledge and whether they do or do not violate closure in the first place.

2.3 Closure and experimental epistemology: Turri's study

Fortunately, we need not start from scratch. It has already been argued that closure is not a principle of folk epistemology on the basis of lottery-type cases (Turri, 2015a; see also Turri, 2015b). Putting Vogel's car theft case (Vogel, 1990, pp. 15f.)[6] to the test, Turri found that the folk quite often ascribe knowledge in a proposition without ascribing knowledge in a

6 In one of his studies (2015b) Turri also uses a case from the second group of potential counterexamples that is loosely inspired by Dretske's zebra case (called "specimen"). Here participants are asked whether the protagonist, Michelle, knows that the animal in the cage is a jaguar and whether she knows that the animal in the cage is not a leopard. However, we doubt that this case is sufficiently similar to Dretske's zebra case. Jaguars and leopards may be perceptually indistinguishable to an average zoo visitor, but nothing in the vignette suggests that they are indistinguishable to Michelle who is introduced as a regular zoo patron ("has visited the city zoo every day for the past ten years", 2015b, p. 313).

proposition entailed by it. Here is one of the vignettes used by Turri (square brackets and slashes indicate the differences between the two conditions used):

> "When Mr. Maxwell arrives at work in the morning, he always parks in one of two spots: C8 or D8. Half the time he parks in C8, and half the time he parks in D8. Today Maxwell parked in C8.
>
> It's lunchtime at work. Maxwell and his assistant are up in the archives room searching for a particular document. Maxwell says, 'I might have left the document in my car.' The assistant asks, 'Mr. Maxwell, is your car parked in space C8? It's not unheard of for cars to be stolen.'
>
> Maxwell [looks carefully out the window / thinks carefully for a moment] and then responds, 'No, my car has not been stolen. It is parked in C8.'" (Turri, 2015a, p. 6)

In the perception condition Maxwell bases his belief that his car is parked in C8 on seeing it, in the inference condition on remembering where he left it (and "careful thinking"). In all conditions participants were presented with knowledge ascriptions and asked to tick those they agreed with:[7]

> A: Maxwell knows that his car is parked in C8.
>
> B: Maxwell knows that his car has not been stolen.

Turri found that participants tend to ascribe both knowledge of the A- and the B-item in the perception condition, but tend to ascribe knowledge of the A-item but not of the B-item in the inference condition. Since similar patterns are observed in a between-subjects design, when knowledge ascriptions are in conjunctive form, when participants are additionally asked about knowledge of the entailment and when clarifying additions to the vignette are made, this basic

7 The labels our ours. Here and below, when reporting our own study, we call the *entailing* proposition the "A-item" and the entailed proposition the "B-item", Turri calls them "positive" and "negative" respectively.

result is very robust. For the sake of brevity, we report only some data from a within-subjects design (experiment 2 in Turri, 2015a, p. 8, fig. 2): In this experiment participants saw both items and their responses can be sorted into four groups, namely (a) the knowledge-friendly view (yes/yes), (b) the anti-closure view (yes/no), (c) scepticism (no/no) and (d) other (no/yes, not chosen by any participant).[8] In the perception condition most participants responded with the knowledge-friendly view (80%, most of the other participants chose the anti-closure view and only a small minority scepticism), whereas in the inference condition close to half of the participants responded with the anti-closure view (45%, other participants chose the knowledge-friendly view more often than scepticism, Turri, 2015a, p. 8, esp. fig. 2). Running binomial tests shows that in the inference condition the anti-closure view was chosen significantly ($p<.001$) *above* chance, whereas in the perception condition the anti-closure view was chosen non-significantly (but trending towards significance, $p<.09$) *below* chance (Turri, 2015a, p. 8).

To sum up, there is at least one case, the inferential car theft case, in which a large number of the folk accepts that someone knows a proposition, but does not know an entailed proposition. Of course, not everyone gives responses that violate closure, but the number of responses that violate closure is apparently high enough to show that closure is not a universally accepted principle of folk epistemology. Therefore, as Turri concludes,

> "the results from our five experiments paint a very different image of the status of epistemic closure in folk epistemology. The overall pattern of results is definitely not what closure's proponents had led us to expect. We repeatedly observed

8 Again, the labels are ours. Alternative labels are, for example, (a) anti-scepticism, (b) semi-scepticism and (c) scepticism.

patterns that are hard to reconcile with the claim that (unqualified) epistemic closure is a defining feature of folk epistemology." (Turri, 2015a:, p. 14)

2.4 Discussion of Turri's study

Turri's study makes an interesting and strong case against the intuitiveness of closure and the assumption that folk knowledge ascriptions abide by closure. Yet, it also suffers from several weaknesses. Some simple worries about details of the vignette can be dealt with easily by small alterations of the vignette used (see below, section 3.2), but there are also principled methodological and philosophical objections that cannot be overcome that easily and have wider implications. In this section we raise and discuss three objections to Turri's study.

First, the role of the experimental manipulation in Turri's study requires clarification. Although there are conditions in his study, namely a perception and an inference condition, they play no argumentative role in his case against closure. Turri manipulates the epistemic source (of the entailing item) and shows that folk knowledge ascriptions are sensitive to this difference. However, the epistemic source, especially of the entailing item, is irrelevant to closure. For, according to closure, if someone (who knows the relevant entailment) knows something – no matter whether perceptually, inferentially or in some other way –, she must also know the entailed proposition – again, no matter whether perceptually, inferentially or in some other way.[9] Turri's case rests solely on the pattern of knowledge ascriptions observed in one condition, the inference condition. For example, the important part of his discussion of the second experiment (2015a, p. 8) is not that knowledge ascriptions differ across the two

9 Some versions of closure (cf. fn. 1) require that knowledge of the *entailed* item (= B-item) is or at least can be inferential. But even these versions of closure are not affected by Turri's manipulation because his manipulation targets the epistemic source of the *entailing* item (= A-item).

conditions, but that knowledge ascriptions in the inference condition are significantly different from chance (as shown by running binomial tests). The argument against closure is not based on an experimental manipulation, but just on polling the folk.[10]

Second, strictly speaking, that my car is parked in C8 does not *entail* that it has not been stolen. To begin with, the conditional *if my car is parked in C8, it has not been stolen* is only plausible given the background assumption that it was indeed parked in C8 this morning, that Maxwell has not lent his car keys to someone else and so on. Once alerted to such background assumptions, even more remote possibilities spring to mind. For example, the car may have been stolen twice. The second thief parked the car in C8 intending to hide it there. Alternatively, the car may have been stolen only once, but the absentminded thief drove it back to C8 intending to hide it there. Again, the conditional *if the car is parked in C8, it has not been stolen* is true only if a lot of background assumptions are taken for granted.

The distinction between material and logical entailments is important because it lends itself to an explanation of why participants let the knowledge ascriptions of the two items come apart: Since *the car is parked in C8* does not entail *it has not been stolen*, it is all too easy to judge the two items independently of each other. Unless the intended entailment is made salient to participants, they can direct attention at the location of the car when considering the first item and at the possibility of a car theft when considering the second item. The intended entailment is also slightly obscured by the formulation that the car "is parked in C8" instead of "is in C8 now". The first formulation may misdirect attention to the

10 It may be objected that at least in the between-subjects design a statistical comparison between conditions is central to the argument against closure (Turri, 2015a, p. 7). But the statistical comparison is needed here only because in a between-subjects design participants are only asked about one of the two items and closure violations are not observed directly.

act of parking instead of directing it at the current location of the car. This is different in the case of obvious logical entailments: One cannot ignore the logical entailment between, for example, *the car is parked in C8* and *the car is not somewhere else than in C8 (because it was stolen)*, even momentarily. To put the same objection slightly differently, "the car hasn't been stolen" can be interpreted in two ways, namely literally and in an enriched way. When considered in isolation the sentence is solely about the possibility of a car theft, but when considered as being entailed by "my car is in C8" the content of the sentences must be enriched to be also about the car's not being in C8 (i.e. the sentence must be interpreted as saying that the car has not been taken away from C8 by thieves). It is unfortunate that the argument against closure rests on the enriched interpretation since it cannot be taken for granted that folk knowledge ascriptions are indeed based on that interpretation.[11]

Third, Turri assumes that proponents of closure expect that folk knowledge ascriptions do not violate closure in the car theft case. However, in discussions of lottery-type cases it is often predicted that violating closure *is* intuitive to some degree. For example, Vogel writes:

11 Note that this is not a *post hoc* objection. Roush (2010) argues that even *philosophers'* intuitions about the brain in a vat (= BIV) argument suffer from such difficulties. When considering whether knowing that one has hands requires knowing that one is not a BIV, we automatically assume the enriched (entailed, easier to know) interpretation *I'm not a handless BIV*. But when considering whether we can rule out being a BIV, we shift back to the literal (not entailed, harder to know) interpretation *I'm not a BIV*. According to Roush, this shift partly explains the (misleading) appeal of the sceptical argument. Douven (2007) offers a somewhat complementary explanation. According to his "pragmatic dissolution", it is the entailing item that can either be taken literally or as pragmatically enriched. In its pragmatically enriched interpretation an assertion of "my car is in C8 now" is modified by a *ceteris paribus* or *ceteris absentibus* clause. What is asserted is that the car is in C8 *unless* it has been stolen, was hit by a meteorite and so on. Accordingly, *my car is in C8* only entails that it has not been stolen when taken literally, but not in its pragmatically enriched interpretation.

> "Many people have the intuition that you would not know that [your car hasn't been stolen]. If this intuition is combined with the previous one [that you know where your car is], then it seems that the closure principle is violated." (Vogel, 1990, pp. 15f.)

Thus, contrary to what Turri assumes, Vogel concedes that "many people" have the intuitions Turri found them to have. A similar stance is taken by Nagel:

> "Notwithstanding its intuitive appeal, closure might seem to be threatened by certain patterns of judgment that also seem intuitive to us." (Nagel, 2011, p. 1)

Like Vogel, Nagel concedes that the pattern of folk ascriptions found by Turri is to be expected, but argues at length why these (apparent) intuitions do not tell against closure. Thus, the contested issue is not so much what intuitions the folk have about the car theft case, but whether their intuitions do or do not indicate a rejection of closure. Proponents of closure can reconcile intuitions about particular cases with a general acceptance of closure in one of two ways. On the one hand, closure violations may be explained away as a performance error. Although closure is a principle of folk epistemology, there may be features of particular cases that induce us to commit mistakes. On the other hand, we may need to accept that intuitive responses can be inconsistent. Counterexamples to closure are not called paradoxes for no reason and it is a hallmark of paradoxes that they can induce inconsistent responses.

The first strategy is exemplified by Vogel's and Nagel's account of our intuitive responses. Vogel's basic idea (1990, p. 19, p. 22) is that we are mistaken about which propositions are lottery-type propositions. Whereas we recognise easily the lottery character of *my car has not been stolen*, *NN had a lethal heart attack within the last five minutes* and so on, we do not recognise the lottery character of *my car is where I left it* or *NN is the current president* right

away. Their lottery character is only revealed when one thinks about the entailment carefully. Thus, intuitions do not indicate a rejection of closure, but only that we are unaware of the lottery character of some propositions.[12] Nagel's basic idea (2011) is that we need to distinguish between system 1 and system 2 responses to knowledge ascriptions. Several features of *my car has not been stolen* are features that usually activate a system 2 response: Its content is negative, it is concerned with probabilities and the possibility is not something we are ordinarily confronted with. System 2 responses, however, tend to be more cautious and knowledge is denied more readily. Thus, intuitions do not indicate a rejection of closure, but only that some contents activate system 2 which tends to induce a denial of knowledge.[13]

The second strategy takes seriously that lottery-type cases including the car theft case can be regarded as paradoxes (e.g. Douven, 2007; Nagel, 2011). Paradoxes consist of claims that are individually intuitive, but jointly inconsistent. In the car theft case the paradox is made up of three statements:

1. Maxwell knows that his car is parked in C8.
2. Maxwell does not know that his car has not been stolen.
3. If Maxwell knows that his car is in C8, he also knows that it has not been stolen.

12 Vogel also mentions that our assessment of the probabilities involved may be unstable. Thus, the tendency to deny knowledge of the lottery propositions may be due to a shift in probability estimates, cf. 1990, p. 19. Turri's study, however, found not evidence for such shifts. In the within-subjects experiments the two responses that amount to violating closure are given *at the same time*.

13 Nagel also argues that the examples fail to be counterexamples to closure because due to the system 1 vs. system 2 distinction we cannot actually infer the entailed proposition from the entailing one. This objection, however, applies only to those versions of closure that require not only knowing the entailment, but actually drawing the inference. Since our discussion is not restricted to particular versions of closure, we leave this part of her account aside.

The paradox account allows that intuitions may be inconsistent and predicts that all three claims are accepted by the folk. In the light of this observation Turri's data can be interpreted in two ways: That many participants accept the first two claims can indicate *either* that they reject the third claim *or* that they accept all three claims and have inconsistent intuitions. Since Turri only asks participants about the first two claims, his data do not rule out that all three claims are intuitive and are, therefore, in principle consistent with closure being intuitive. Of course, if folk intuitions actually turned out to be inconsistent[14], that would not decide the issue in favour of closure either. That is why a different experimental approach would be helpful.

To sum up the third objection, proponents of closure need not, and as a matter of fact at least some do not, predict that folk knowledge ascriptions never violate closure in the car theft case. To the contrary, some take the car theft case and similar cases to push us exactly towards the intuitions reported by Turri. Accordingly, his data do not refute a prediction unanimously made by proponents of closure.

2.5 A different experimental approach

Taken together the three objection point in the same direction: Closure is difficult to test experimentally because its proponents do not formulate empirically testable predictions. For they predict both that closure is highly intuitive as a general principle and that there are puzzling, maybe even paradoxical, cases in which intuitions stray off course. Instead of

14 Paradoxical cases are tricky to test experimentally. Philosophers may be used to the possibility of having inconsistent intuitions, but participants in experimental studies may tend to report consistent intuitions. Such a tendency to minimise consistency is predicted by cognitive dissonance theory and is mentioned within the experimental philosophy literature, for example, by Schwitzgebel & Cushman (2012, p. 148, cf. Lombrozo, 2009) as a potential explanation of order effects.

refuting a prediction made by proponents of closure, Turri's study rather confirms that lottery-type cases are indeed puzzling cases. But what should their prediction be instead? They should not rest content with the unspecific and vague expectation that in most, but not all cases folk knowledge ascriptions agree with closure. Fortunately, our objections are not merely destructive, but also point towards a different experimental approach to testing closure. In response to the first objection, we should not manipulate epistemic source, but something else. In response to the second objection, we should look at the kind of entailment as a candidate for experimental manipulation. In response to the third objection, we should not focus on the mere number of closure violations, but on whether and how this number changes across conditions. Our proposal, therefore, is to understand the assumption that the folk endorse closure as predicting that folk knowledge ascriptions are sensitive to entailments, i.e. when comparing folk knowledge ascriptions of two logically independent items with knowledge ascriptions of two logically dependent, but otherwise very similar items, the number of split knowledge ascriptions (yes/no) must go down considerably. This prediction has two parts: On the one hand, split knowledge ascription must be low when the entailment is a logical entailment, on the other hand, they must be lower when the strength of entailment increases. Both predictions are important: Suppose we found that almost all participants choose the split view when the items are independent and most, but not all when the items are logically dependent (i.e. lower, but not low). Although this result would call for an explanation, it would not indicate that the folk endorse closure. Suppose we found that the split view is chosen very rarely no matter how the items are related (i.e. low, but not lower). This result would not indicate that the folk endorse closure because it could be explained by a wide variety of reasons (features of the vignette that are not changed between conditions,

pragmatic effects, knowledge-friendly view or scepticism being independently appealing and so on).

To explain our perspective on closure further, if knowledge ascriptions are influenced by closure, the mere fact that a proposition is entailed by another proposition should make a difference to knowledge ascriptions even if everything else is kept constant. Thus, when considering two logically independent propositions, A and B, with A being an ordinarily known proposition and B being a lottery-type proposition, we expect knowledge ascriptions of A to be much higher than of B. (For example, *my car is in C8* and *my car hasn't been damaged*). When replacing B by a slightly different proposition which is entailed by A, the epistemic standing of A and B does not change or, at least, if the propositions are well chosen, the kind of evidence does not change. (For example, *my car is in C8* and *my car hasn't been stolen*). However, if folk knowledge ascriptions adhere to closure, the pattern of knowledge ascriptions must change. Importantly, since the epistemic standing is kept constant and only the logical dependence is changed, the difference of knowledge ascriptions, if it can indeed be observed, can only be explained by knowledge ascriptions being sensitive to entailments.

Two further predictions are suggested by the objections above. We objected that the car theft case is, strictly speaking, not a case of logical entailment. This suggests to test whether the distinction between material and logical entailment has an effect on knowledge ascriptions. Moreover, material entailments can be more or less salient to readers of a vignette. This suggests picking up the conjecture from above that "is parked in" and "is now in" may be interpreted differently due to a pragmatic difference. If knowledge ascriptions are sensitive to this distinction as well, this suggests an explanation of the number of seeming closure violations found by Turri that is consistent with closure being a principle of folk

epistemology: If the number of closure violations goes down when replacing “is parked in” by “is now in”, this indicates that the folk interpret the first pair of items as independent, but the second pair as dependent. To the very least, it shows that they do not treat “is parked in” and “is now in” synonymously.

3 Experiment 1

In the first experiment we put our criticism to the test and changed the methodology accordingly. Instead of manipulating epistemic source we manipulated the kind of entailment. All participants read one vignette, a modified version of Turri’s car theft case, and were asked about two items. The first item, the A-item, is something we commonly think of as knowledge (in this case, where one’s car is), while the second item, the B-item, is a lottery-type proposition that is *either* independent of the A-item *or* materially or logically entailed by it (in this case, whether one’s car has been damaged or stolen). Responses were classified as follows: The sceptical view (no/no) denies that any of the two items are known, whereas the knowledge-friendly view (yes/yes) consists in the opposite stance that both items are known. The split view (yes/no) is what we are especially interested in. If the A-item entails the B-item, taking the split view amounts to violating closure. The view labeled “other” (no/yes) is only included for the sake of comprehensiveness. It is difficult to make sense of this view and it was almost never chosen by participants. As described above, our hypothesis was that the stronger the inferential dependence, the lower the number of splitters (i.e. closure violations) will be.

3.1 Participants

The experiment was run online in the U.K. 308 participants started the experiment, 297 were included in the analysis (we excluded participants who did not finish the survey or went through it in less than 40 seconds).[15] 54% identified as male, 46% as female. Mean age was 31 years and participants earned £0.25 for their participation which took on average 92 seconds.

3.2 Material and Procedure

The experiment was conducted on the internet. Upon clicking on a link on the Prolific Academics webpage[16] they were redirected to a website containing the experiment. Participants first read general instructions and were then randomly assigned to one of four conditions in a between-subjects design. They were presented with one of the following stories and two randomly ordered items. We used a modified version of Turri's vignette: Instead of a dialogical setting we used a monological setting thereby ruling out that participants' responses are influenced by their judgement about the adequacy of making an assertion. Instead of introducing the protagonist by describing his parking habits we introduced him as someone who memorises where he parked his car. Instead of leaving it

15 In all four experiments reported in this paper we also asked a comprehension question placed on a separate screen before the demographic questions. Unfortunately, participants in one condition (in the third experiment) mistakenly received the comprehension question belonging to a different condition. To keep the methodology consistent we relied only on the time criterion in all experiments. Excluding participants that failed the comprehension test does not affect the results in any of the four experiments reported below.

16 Prolific Academics was established as a European alternative to the most popular database to recruit participants online, namely MTurk, which originally could only be accessed by providing a US social security number. Most members are Caucasian and have an undergraduate degree. For more details, see https://www.prolific.ac/demographics.

implicit that the protagonists' beliefs are true we explicitly stated so in the final sentence of the vignette. Finally, we kept Turri's formulation of the A-item ("parked in C8") in one condition, but replaced it with a formulation explicitly about the car's current location ("is in C8 now") in the other three conditions. The four conditions then resulted from manipulating the kind of entailment between the A- and the B-item. The differences between conditions are explained below and are also summarised in Table 2 (text in bold was not seen by participants).

Condition	A-item	B-item
Independent	car is in C8 now	car has not been damaged
Weak Material	car is parked in C8	car has not been stolen
Strong Material	car is in C8 now	car has not been stolen
Logical	car is in C8 now	car has not been taken away from C8 by thieves

Table 2: Experiment 1: Conditions

In the first condition, Independent, participants were asked about two logically independent items (*car is in C8 now*, *car hasn't been damaged*). The second item is a lottery-type proposition comparable to *my car hasn't been stolen*: That a particular car has been damaged or stolen is both unlikely, but not unheard of. Thus, we assume that one's justification or evidence for both lottery-type propositions is of the same quality.

Independent

When Maxwell arrived at work this morning, he parked his car in C8. He kept in mind where he parked it because he wants to get home quickly after work. Now it's 5 p.m. and Maxwell leaves his office. On the way to the parking lot it comes to his mind that it's not unheard of for cars to get damaged. After considering this

carefully for a moment, he thinks: “No, my car hasn’t been damaged. It’s in C8 now.” And he is right: His car hasn’t been damaged and is in C8 now.

A: Maxwell knows that his car is in C8 now. (yes/no)

B: Maxwell knows that his car has not been damaged. (yes/no)

In the second and third condition, Weak and Strong Material, participants were asked about two materially dependent items, i.e. although the A-item does not, strictly speaking, entail the B-item, the conditional *if A, then B* is a natural background assumption. While the B-item is in both conditions *my car has not been stolen*, the two conditions differ with respect to the formulation of the A-item. For Weak Material we used Turri’s formulation “parked in C8”, for Strong Material we used “is in C8 now” (as we did in the other conditions as well). Our hypothesis was that the latter formulation makes the entailment more salient than Turri’s and that, therefore, reduced numbers of closure violation are to be expected.

Weak Material

When Maxwell arrived at work this morning, he parked his car in C8. He kept in mind where he parked it because he wants to get home quickly after work. Now it’s 5 p.m. and Maxwell leaves his office. On the way to the parking lot it comes to his mind that it’s not unheard of for cars to be stolen. After considering this carefully for a moment, he thinks: “No, my car hasn’t been stolen. It’s parked in C8.” And he is right: His car hasn’t been stolen and is parked in C8.

A: Maxwell knows that his car is parked in C8. (yes/no)

B: Maxwell knows that his car has not been stolen. (yes/no)

Strong Material

When Maxwell arrived at work this morning, he parked his car in C8. He kept in mind where he parked it because he wants to get home quickly after work. Now it's 5 p.m. and Maxwell leaves his office. On the way to the parking lot it comes to his mind that it's not unheard of for cars to be stolen. After considering this carefully for a moment, he thinks: "No, my car hasn't been stolen. It's in C8 now." And he is right: His car hasn't been stolen and is in C8 now.

A: Maxwell knows that his car is in C8 now. (yes/no)

B: Maxwell knows that his car has not been stolen. (yes/no)

Finally, in the fourth condition, Logical, we changed the formulation of the B-item such that it is logically entailed by A: "my car hasn't been taken away from C8 by thieves". Our hypothesis was that in the case of logical entailment the dependence is even more salient than in the case of material dependence. Note that our assumption is not that participants understand and apply the distinction between material and logical dependence or entailment, but that logical entailments are usually more salient, especially so in this case because the location of the car, C8, is explicitly mentioned in both items.

Logical

When Maxwell arrived at work this morning, he parked his car in C8. He kept in mind where he parked it because he wants to get home quickly after work. Now it's 5 p.m. and Maxwell leaves his office. On the way to the parking lot it comes to his mind that it's not unheard of for cars to be stolen. After considering this carefully for a moment, he thinks: "No, my car hasn't been taken away from C8 by thieves. It's in C8 now." And he is right: His car hasn't been taken away from C8 by thieves and is in C8 now.

A: Maxwell knows that his car is in C8 now. (yes/no)

B: Maxwell knows that his car has not been taken away from C8 by thieves. (yes/no)

Afterwards participants were asked some demographic questions.

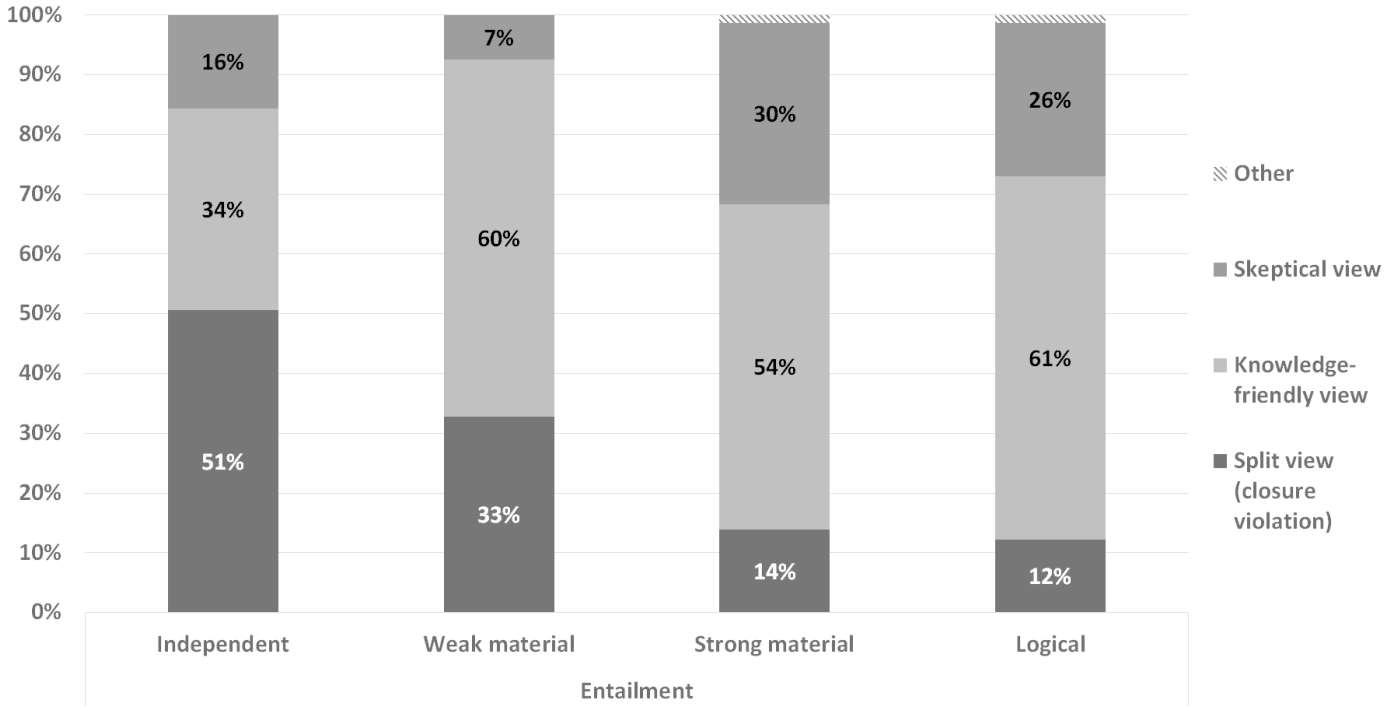

Figure 1: Percentage of participants' responses as a function of the strength of entailment in experiment 1

3.3 Results and Discussion

The results of the experiment are summarised in Figure 1. Overall, the stronger the entailment was, the lower was the proportion of participants who violated closure, i.e. who agreed with the A-item but disagreed with the B-item. A trend-test for frequencies (Pfanzagl, 1974, p. 193), =5.88, p<0.001, confirmed this impression. While there were 51% splitters in

Independent, only 12% of participants violated the closure principle in Logical (Weak Material: 33%, Strong Material: 14%). Comparing Independent with Weak Material revealed a significant difference, χ^2 (df = 1, N = 144) = 4.66, p < 0.05, V=0.18, with closure violations being higher in Independent. The comparison of Weak Material with Strong Material also revealed a significant effect, χ^2 (df = 1, N = 146) = 7.41, p < 0.01, V=0.23, with closure violations being higher in Weak Material. There was no significant difference between Strong Material and Logical, χ^2 (df = 1, N = 153) = 0.10, p = 0.75, V=0.03, which might be due to both the subtlety of our manipulation and the fact that the proportion of participants violating closure in Strong Material was already quite low. For Logical the percentage of participants ascribing knowledge was significantly below chance (binominal test, *N*=74, test value =0.25, two-tailed, $p<0.001$). Together with the reported trend this result indicate that closure is endorsed by the majority of participants. Interestingly, in Weak Material our numbers are only slightly different from Turri's (experiment 2, inference condition; 2015a, p. 8). While in our study 33% violated closure, 45% did so in his study.

4 Experiment 2

Our first experiment may raise the worry that the vignette in car may have been misunderstood by the participants: Since Maxwell is described as being on his way to the parking lot and the final sentence mentions that his car is where he left it, participants may interpret the vignette as Maxwell having already arrived at his car in the final sentence. To address this worry we reran the experiment with the changes explained below.

4.1 Participants

The experiment was again run online in the U.K. 320 participants started the experiment, 313 were included in the analysis (we excluded participants who did not finish the survey or went through it in less than 40 seconds). 42% identified as male, 57% as female, and 1% as neither. Mean age was 32 years and participants earned £0.25 for their participation which took on average 91 seconds.

4.2 Material and Procedure

We made two changes to our original car vignette. Questions (two items in randomised order) and design (four conditions: Independent, Weak Material, Strong Material, Logical) remained the same. The two differences consisted, *first*, in leaving out the last sentence of the vignette ("And he is right: His car hasn't been damaged and is in C8 now.") and, *second*, in replacing "Now it's 5 p.m. and Maxwell leaves his office. On the way to the parking lot" with "Now it's 5 p.m. and Maxwell leaves his office for the parking lot. While waiting for the elevator […]". Both changes were made to clarify that Maxwell has not already arrived at his car and cannot see it from his current location.

4.3 Results and Discussion

The results of the experiment are summarised in Figure 2. The overall trend was similar to our previous experiment: The stronger the entailment was, the lower was the proportion of participants who violated closure, i.e. who agreed with the A-item but disagreed with the B-item. This trend was statistical significant, T=8.68, p<0.001. While there were 66% splitters in Independent, only 10% of participants violated the closure principle in Logical (Weak Material: 70%, Strong Material: 23%). Comparing Independent with Weak Material revealed

no significant difference this time, χ^2 (df = 1; N = 155) = 0.26; p = 0.61, V= 0.04.[17] The comparison of Weak Material with Strong Material revealed a significant effect, χ^2 (df = 1; N = 159) = 33.61, p < 0.001, V= 0.46, with closure violations being higher in Weak Material. There was also a significant difference between Strong Material and Logical, χ^2 (df = 1; N = 158) = 5.08; p < 0.05, V= 0.18. For Logical the percentage of participants ascribing knowledge was significantly below chance (binominal test, N=78, test value =0.25, two-tailed, p<0.001). Together with the reported trend this result indicates that closure is endorsed by the majority of participants.

17 That is, folk responses to a pair of independent items – *my car is in C8 now* and *my car has not been damaged* – and to Turri's pair of items – *my car is parked in C8* and *my car has not been stolen* – are similar. In our view the best explanation is that in some contexts they understand Turri's pair as an independent pair, not that they violate closure. Not mentioning that the car actually is in C8 may trigger interpreting "is parked in" as being about where the car has been parked.

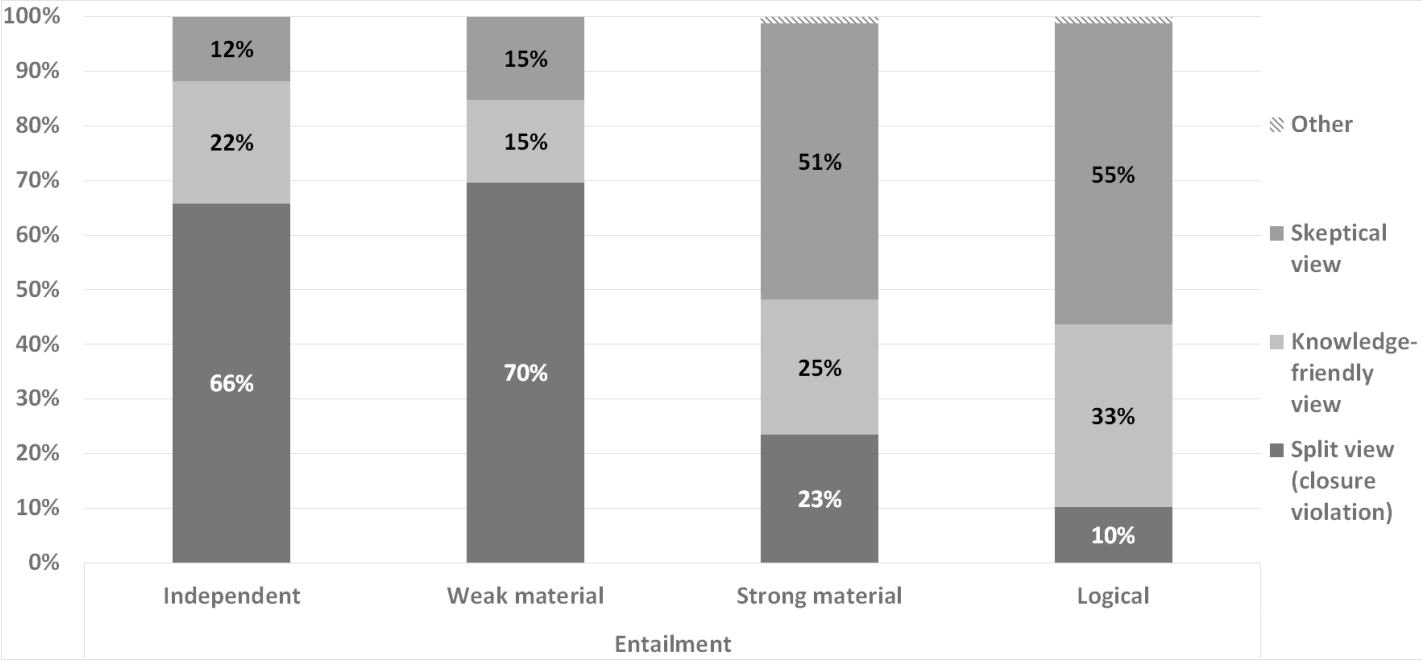

Figure 2: Percentage of participants' responses as a function of the strength of entailment in experiment 2

5 Experiment 3

Contrary to the manipulation used by Turri, the manipulation used in experiments 1 and 2 suggests that folk ascriptions of knowledge adhere to closure after all. Unlike his study, our study reveals a driving factor behind the observed knowledge ascriptions. The large number of spilt view responses in material entailment conditions found by Turri is consistent with responses being sensitive to entailment. However, the two experiments also have limitations. *First*, we did not find a significant difference between strong material dependence and logical dependence in the first experiment and no significant difference between independence and weak material dependence. *Second*, the distinction between weak and strong material entailment was suggested by the specifics of the car theft case, namely the presumably pragmatic difference between "parked in" and "is now in". It is not to be expected that a

similar distinction can be drawn in all other cases as well. *Finally*, the two experiments do not tell whether our result is limited to this particular case or can be generalised. For these reasons we ran a similarly designed experiment with two new vignettes. Apart from using different vignettes, we changed the number of conditions from four to three: Independent, Material and Logical.

5.1 Participants

The experiment was run online in the U.K. 468 participants started the experiment, 447 were included in the analysis (we excluded participants who did not finish the survey or went through it in less than 40 seconds). 54% identified as male, 46% as female. Mean age was 29 years and participants earned £0.25 for their participation which took on average 129 seconds.

5.2 Material and Procedure

The experiment was conducted on the internet. Upon clicking on a link on the Prolific Academics webpage participants were redirected to a website containing the experiment. They first read general instructions and were then randomly assigned to one of the six conditions in a between-subjects design. Participants were either presented the Supermarket or the Party vignette in one of the three versions differing in the strength of entailment (Independent, Material, Logical) and they were again asked two randomly ordered items. The vignettes read as follows, differences between conditions are marked by square brackets and are separated by a slash:

Supermarket

Susan likes to drink her morning coffee with a lot of milk. After drinking her coffee this morning, she notices that there is no milk left in the fridge. Since she is

working the late afternoon shift today, she plans to go to the local supermarket after work at about 9 p.m. She has bought her groceries there almost daily during the past year and it has always been open 6 a.m. to 12 a.m. While commuting to work, it comes to her mind that it's not unheard of for supermarkets [to be seriously understaffed, for example due to several cases of illness, and that therefore she may have to wait longer than usual in the queue / to change their opening hours at very short notice, either by extending or by cutting them / to change their opening hours at very short notice]. After considering this carefully for a moment, she thinks "No, [the supermarket isn't seriously understaffed today / the supermarket hasn't changed its opening hours / the supermarket hasn't changed its opening hours to 6 a.m. to 8 p.m.]. It will be open at 9 p.m. tonight." And she is right: [The supermarket isn't seriously understaffed / the supermarket hasn't changed its opening hours / the supermarket hasn't changed its opening hours to 6 a.m. to 8 p.m.] and will be open at 9 p.m. tonight.

A: Susan knows that the supermarket will be open at 9 p.m. (yes/no)

B: Susan knows [that the supermarket is not seriously understaffed today / that the supermarket has not changed its opening hours / that the supermarket has not changed its opening hours to 6 a.m. to 8 p.m]. (yes/no)

Party

Arthur and his best friend David are both invited to Betty's flat for her birthday party. Yesterday David told Arthur that he would be coming. David lives within walking distance of Betty's flat and promised Betty to come early to help her with moving furniture. On his way to Betty Arthur is stuck in traffic. Running late it

comes to his mind [that David might be wearing his awful purple jumper today he hasn't worn for at least a year / that David might have hurt his back again while working in his garden / that David might be in the hospital because he has hurt his back again while working in his garden]. After considering this carefully for a moment, he thinks, "No, [David is already at Betty's flat, but he isn't wearing his purple jumper. / David hasn't hurt his back. He's already at Betty's flat. / David isn't in the hospital right now, but is already at Betty's flat.". And he is right: [David is already at Betty's flat and isn't wearing his purple jumper. / David hasn't hurt his back again and is already at Betty's flat. / David isn't in the hospital now, but already at Betty's flat.]

A: Arthur knows that David is already at Betty's flat. (yes/no)

B: Arthur knows that [David is not wearing his purple jumper / that David has not hurt his back / that David is not in the hospital right now]. (yes/no)

Afterwards participants were asked some demographic questions.

5.3 Results and Discussion

The results of the experiment are summarised in Figure 3. Overall, we found again that the stronger the entailment was, the lower was the proportion of participants who violated closure, i.e. who agreed with the A-item but disagreed with the B-item. Since this trend was evident in both vignettes we collapsed the data for the analysis.[18] Again, the trend-test for frequencies (Pfanzagl, 1974, p. 193) was significant, T=7.09,p<0.001. While there were 41% splitters in Independent, only 7% of participants violated the closure principle in Logical

18 Using logistic regression, adding the interaction term of vignette * entailment to the model did not significantly improve it, .

(Material: 16%). A χ^2-test revealed a significant difference between Independent and Material, χ^2 (df = 1, N = 297) = 21.87, p < 0.001, V=0.27, with closure violations being higher in Independent. The comparison of Material with Logical also revealed a significant effect, with closure violations being higher in Material. For Logical the percentage of participants ascribing knowledge was significantly below chance (binominal test, *N*=150, test value =0.25, two-tailed, p<0.001). Together with the reported trend this result indicate that closure is endorsed by the majority of participants. To sum up, our results indicate again that folk knowledge ascriptions are sensitive to entailment.

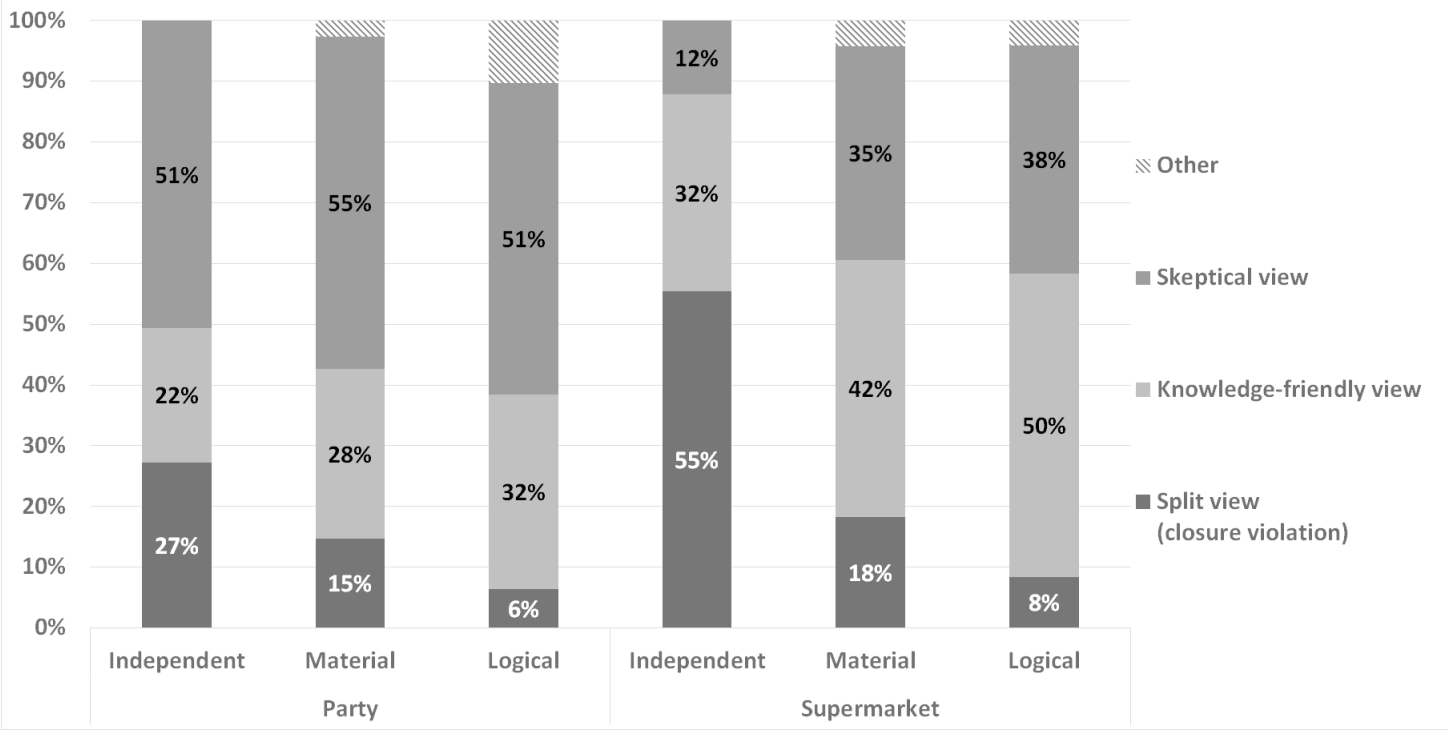

Figure 3: Percentage of participants' responses as a function of the strength of entailment in experiment 3

6 Experiment 4

In this section we address a general worry about the design of our experiments: Since the B-items have to be different for the kind of entailment to be different, there are bound to be more differences between conditions than merely the kind of entailment. Our thesis is that the different response patterns are driven by the difference in the kind of entailment, but one may object that they may well be driven by other differences instead. For example, participiants may think that it is more likely that a particular car has been damaged than that it has been stolen. Accordingly, their responses could be explained by its being easier to know that a car has not been stolen than that it has not been damaged. Since participants who (tend to) ascribe knowledge in the B-item independently of the A-item cannot violate closure, such a difference in independent knowability could explain why the number of closure violations decreases in Material and Logical. Can we rule out that our data are explained by differences in independent knowability?

To address this worry we ran an additional experiment. But before turning to our new experiment, let us look at what the data we already have tell us. If the explanation of our data is that the B-items in Logical and Material are independently easier to know than the B-item in Independent, we should expect only the number of participants that choose the knowledge-friendly response to go up. But what we found is an increase in both the knowledge-friendly and the sceptical response (see Figure 1, Figure 3). Thus, it is rather unlikely that the differences between conditions are driven by the differences in independent knowability of the B-items. In the fourth experiment we collected additional data to confirm this conjecture. In order to estimate the independent knowability of the respective B-items we added modified versions of the vignettes used in previous experiments. Participants in One Item read vignettes

in which the protagonists consider only the B-item and were asked only about that item. Participants in Two Items read the orginal vignettes and were asked about both items.

6.1 Participants

The experiment was again run online in the U.K. 737 participants started the experiment, 707 were included in the analysis (we excluded participants who did not finish the survey or went through it in less than 90 seconds). 48% identified as male and 52% as female. Mean age was 35 years and participants earned £0.30 for their participation which took on average 202 seconds.

6.2 Material and Procedure

Participants were randomly assigned to one of six conditions in a 2 (between: One Item vs. Two Items) * 3 (between: Independent vs. Material vs. Logical) * 3 (within: Car vs. Supermarket vs. Party) mixed design, i.e. each participant read three vignettes (in random order) and either rated only their B-items or both items. For the Two Items conditions we used the vignettes from the second and third experiment and for the One Item conditions we altered these vignettes by leaving out all mentions of the protagonist thinking about the A-item.[19] We compared the rate of knowledge ascriptions when both items were presented with the rate of knowledge ascriptions when only the B-item was presented. While we did not predict an effect for the independent condition (if the B-item is logically independent of the A-item,

19 Since Weak Material and Strong Material differ only in their A-item, in this experiment Car has only three conditions. Here is a sample full vignette. **Car, Material:** When Maxwell arrived at work this morning, he parked his car in the nearby parking lot, spot C8. Now it's 5 p.m. and Maxwell leaves his office. On the way to the parking lot it comes to his mind that it's not unheard of for cars to be stolen. After considering this carefully for a moment, he thinks: "No, my car hasn't been stolen."

knowledge ascriptions should not depend on whether the A-item is mentioned or not), we expected an effect for the material and the logical condition (if the B-item is entailed by the A-item, knowledge ascription should be lower, as compared to only the B-item being present, according to our theory). Moreover, we expected to find only a small effect since it seems plausible to assume that participants in the One Item condition thought about the second item as well. For example, in Car it is a natural interpretation of the scenario that the protagonist does not think exclusively about the possibility of his car being stolen, but also about its current whereabouts. However, we refrained from altering the vignettes even more because a comparison would not tell us anything meaningful if the One Item and Two Item vignettes differed in too many respects.

6.3 Results and Discussion

The results are summarised in Figure 4. To achieve sufficient test power to find a statistical difference we collapsed the data for the three cover stories (all data can be found in Table 3).

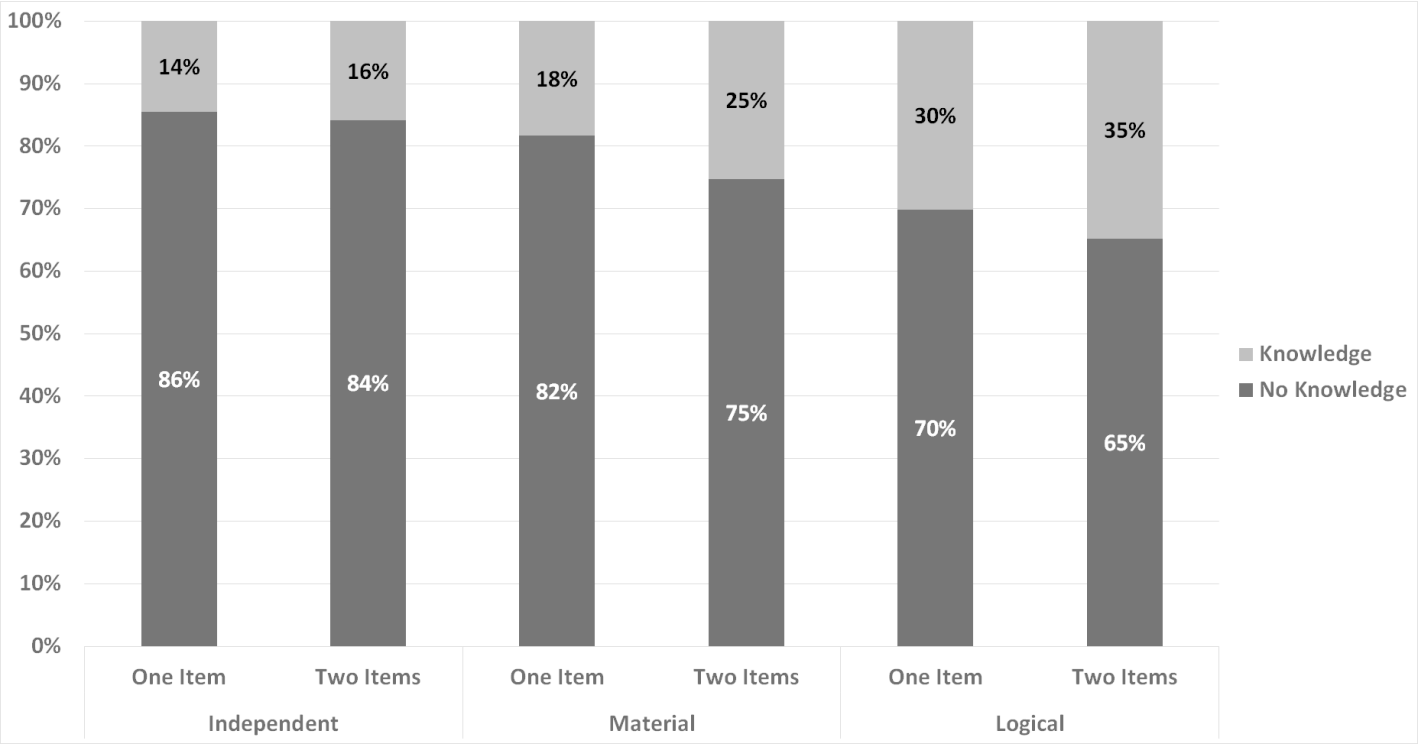

Figure 4: Percentage of participants ascribing knowledge of the B-item

Vignette	Entailment	1 vs. 2 items	No Knowledge	Knowledge
Car	Independent	One	99 (86%)	16 (14%)
Car	Independent	Two	97 (81%)	23 (19%)
Car	Material	One	99 (84%)	19 (16%)
Car	Material	Two	94 (78%)	26 (22%)
Car	Logical	One	95 (81%)	22 (19%)
Car	Logical	Two	86 (74%)	31 (27%)
Supermarket	Independent	One	96 (83%)	19 (17%)
Supermarket	Independent	Two	100 (83%)	20 (17%)
Supermarket	Material	One	90 (76%)	28 (24%)
Supermarket	Material	Two	84 (70%)	36 (30%)
Supermarket	Logical	One	66 (56%)	51 (44%)
Supermarket	Logical	Two	68 (58%)	49 (42%)
Party	Independent	One	100 (87%)	15 (13%)
Party	Independent	Two	106 (88%)	14 (12%)
Party	Material	One	100 (85%)	18 (15%)
Party	Material	Two	91 (76%)	29 (24%)
Party	Logical	One	84 (72%)	33 (28%)
Party	Logical	Two	75 (64%)	42 (36%)

Table 3: Data for all vignettes

Independent: The percentage of participants ascribing knowledge was slightly higher when both items were present (16%), as compared to only the B-item being present (14%), with this difference not being statistically significant, χ^2 (df = 1; N = 705) = 0.25; p = 0.62, V= 0.019. This finding was predicted by our theory.

Material: The percentage of participants ascribing knowledge was higher when both items were present (25%), as compared to only the B-item being present (18%), with this difference being statistically significant, χ^2 (df = 1; N = 714) = 5.00; p = 0.025, V= 0.084. This finding was predicted by our theory.

Logical: The percentage of participants ascribing knowledge was slightly higher when both items were present (35%), as compared to only the B-item being present (30%), with this difference not being statistically significant, χ^2 (df = 1; N = 714) = 1.66; p = 0.197, V= 0.049. While the non-significance was not predicted by our theory, the direction of the found difference was.

To sum up, the findings of this experiment do not suggest that closure violations in Material and Logical decrease just because the B-items in these conditions are independently rated as easy to know. To the contrary, they suggest that knowledge ascriptions of the B-item when considered in isolation are rather different from knowledge ascriptions of the B-item when considered in combination with the respective A-item (unless both items are independent of each other).

7 Discussion

In this paper we discussed two approaches to testing epistemic closure experimentally. After criticising an approach that was meant to show that folk knowledge ascription do not agree

with closure, we offered a different 'translation' of closure into an empirically testable hypothesis. The assumption that closure is a principle of folk epistemology predicts not just that the number of closure violations is low in cases of clear logical entailment, but also that folk knowledge ascriptions are sensitive to the kind of entailment. By keeping the epistemic standing of the pairs of items constant and manipulating their logical dependence we tested whether folk knowledge ascriptions are affected by our manipulation. Our data paint a uniform, but surprisingly different picture than Turri's data do: In all vignettes we observed that the number of split responses, i.e. violations of closure, decreases when the strength of the entailment is increased. Hence, folk knowledge ascriptions actually are sensitive to the kind of entailment: If the B-item, a lottery-type proposition, is not entailed by something we ordinarily take ourselves to know, participants very often do not think that it is known. Yet, if the B-item is not only a lottery-type proposition, but also entailed by something we ordinarily take ourselves to know, folk knowledge ascriptions respond by *either* taking the B-item to be known as well (knowledge-friendly response) *or* by rejecting that the ordinary proposition is known after all (sceptical response). We observed not merely a low number of closure violations when the A-item logically entails the B-item. That alone would not suffice as an argument in favour of closure being a principle of folk epistemology. For a low number of such responses can be explained in various ways, for example, not violating closure need not indicate acceptance of closure, but may indicate that the sceptical or the knowledge-friendly view are independently appealing in the cases at hand. Therefore, our argument depends not merely on the low number of closure violations (when the items are logically dependent), but to the difference in closure violations across conditions. Since the number of split responses, or closure violations, is not constant across conditions, responses cannot be explained by the

sceptical or the knowledge-friendly view being independently appealing. At the same time none of our data contradict the data found by Turri. Like Turri, we found that in vignettes with (weak) material entailment closure violations occur in rates exceeding chance. But if we put these data in perspective, we see something not visible in his studies: Only when comparing responses to his case (a case of weak material entailment) with the other conditions can we assess what the number of seeming closure violations he found actually indicates.

Our study also has some limitations which are worth pointing out. A first limitation is that our study does not allow to decide on the pros and cons of the various philosophical explanations of why it is that closure violations in lottery-type cases seem intuitive. Such an explanation would complement our case for closure as a principle of folk epistemology. Our data are consistent with Vogel's explanation (we are easily unaware that the A-item is a lottery-type proposition as well), with Nagel's explanation (intuition can be explained by distinguishing between system 1 vs. system 2 responses) and with Roush's and Douven's explanation (our intuitions result from not clearly distinguishing between literal and (pragmatically) enriched interpretations of either the A- or the B-item). Yet, our data contributes something to that debate by showing that there is indeed something that stands in need of an explanation, namely why folk knowledge ascriptions both agree by and large with closure as well as (apparently) violate it sometimes. A second limitation is that, although we have argued that acceptance of closure is a factor driving folk knowledge ascriptions, we cannot rule out that there are further cases in which the number of closure violations does not go down when the entailment is made stricter. Since philosophers are usually not interested in whether closure holds for many or most cases, but in whether it holds without exception, opponents of closure need not and should not give up. What we do think, however, is that the

burden of proof is back on the shoulders of closure-deniers. That may not be much news to proponents of closure, but it is not old hat either.

Bibliography

Baumann, P. (2011). Epistemic closure. In S. Bernecker & D. Pritchard (Eds.), *The Routledge Companion to Epistemology* (597–608). London: Routledge.

Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research, 65,* 309–329.

Douven, I. (2007). A pragmatic dissolution of Harman's paradox. *Philosophy and Phenomenological Research, 74,* 326–345.

Dretske, F. (1970). Epistemic operators. *Journal of Philosophy, 67,* 1007–1023.

Dretske, F. (2003). Skepticism: What perception teaches. In S. Luper (Ed.), *The Skeptics* (105–118). Hampshire: Ashgate.

Dretske, F. (2005). Is knowledge closed under known entailment? The case against closure. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (13–26). Malden, Blackwell.

Harman, G. (1973). *Thought.* Princeton: Princeton University Press.

Hawthorne, J. (2004). *Knowledge and Lotteries.* Oxford: Oxford University Press.

Horvath, J. & Wiegmann, A. (2016). Intuitive expertise and intuitions about knowledge. *Philosophical Studies, 173,* 2701–2726.

Kripke, S. (2011). On two paradoxes of knowledge. In S. Kripke, *Philosophical Troubles. Collected Papers, Volume I* (27–51). Oxford: Oxford University Press.

Kvanvig, J. (2006). Closure principles. *Philosophy Compass, 1,* 256–267.

Kyburg, H. (1961). *Probability and the Logic of Rational Belief.* Middletown, CT: Wesleyan University Press.

Lombrozo, T. (2009). The role of moral commitments in moral judgment. *Cognitive Science, 33*, 273–286.

Luper, S. (2016). The epistemic closure principle. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), http://plato.stanford.edu/archives/spr2016/entries/closure-epistemic/

Nagel, J. (2011). The psychological basis of the Harman-Vogel paradox. *Philosopher's Imprint, 11,* 1–28.

Nozick, R. (1981). *Philosophical Explanations.* Cambridge, Ms.: Harvard University Press.

Pfanzagl, J. (1974). *Allgemeine Methodenlehre der Statistik. Band II.* Berlin: de Gruyter.

Pritchard, D. (2012). *Epistemological Disjunctivism.* Oxford: Oxford University Press.

Roush, S. (2010). Closure on skepticism. *Journal of Philosophy, 107*, 243–256.

Schwitzgebel, E. & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language, 27*, 135–153.

Turri, J. (2015a). An open and shut case: Epistemic closure in the manifest image. *Philosopher's Imprint, 15,* 1–18.

Turri, J. (2015b). Skeptical appeal: The source-content bias. *Cognitive Science, 39,* 307–324.

Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth, & G. Ross (Eds.), *Doubting* (13–27). Dordrecht: Kluwer.

Williamson, T. (2000). *Knowledge and Its Limits.* Oxford: Oxford University Press.