

Closure, Counter-Closure, and Inferential Knowledge

Branden Fitelson¹

April 29, 2015

1 Closure and counter-closure: Some stage-setting

The traditional (or “received”) theory of deductive inferential knowledge includes the following two fundamental epistemic principles.²

Closure (C). If S knows that P and S competently deduces Q from P (while maintaining her knowledge that P), then S (thereby) comes to know that Q (*via* deductive inference).

Counter-Closure (CC). If S competently deduces Q from her belief that P , (thereby) coming to know Q (*via* deductive inference), then S knew that P (and she maintained her knowledge of P throughout the inference).

The first thing to note about these principles is that (as I will be understanding them), they are not merely classificatory principles for demarcating/separating cases of knowledge from non-knowledge. As I will be understanding (and using) them, (C) and (CC) are intended to be featured in *epistemological explanations* of *why* some agent S *knows* that Q (as opposed, *e.g.*, to *merely truly believing* that Q). This is why I have added the locutions “thereby” and “*via* deductive inference,” which are not typically included in the statements of (C) and (CC). The explanatory (and dialectical) role(s) of these locutions will become clearer in due course.

The second thing to note about these two principles is that, although there is an obvious sense in which they are epistemologically symmetrical, there are also (perhaps less obvious) senses in which they are epistemologically *asymmetrical*. Specifically, consider the following central epistemological explanandum.

¹Much of the dialectic in this paper has recently been covered in a similar way by Federico Luzzi [12]. I was unaware of Luzzi’s work when I (independently, and often in very different ways) arrived at many of the same general conclusions he does [8]. However, my present emphasis on epistemological *vs.* psychological *explanation* (as well as my emphasis on the relationship between closure and counter-closure, and my discussions regarding generalized counter-closure and ampliative inference) is rather different than his. Having said that, I have learned a great deal from Luzzi’s work (and from many fruitful conversations with him about these issues). My thinking about these issues has been informed by useful discussions with many people (in addition to Federico Luzzi) over the past several years. I cannot list them all here, but I am especially indebted to the following people (in alphabetical order): Brian Ball, Michael Blome-Tillmann, Tim Button, Cian Dorr, Jane Friedman, John Hawthorne, Allen Hazlett, Peter Klein, Matthew McGrath, Martin Montminy, Susanna Rinard, Miriam Schoenfield, Ian Schnee, Jonathan Vogel and Fritz Warfield.

²For instance, in his excellent survey of analytic epistemology, Audi [2, Ch. 8] emphasizes the traditional importance of both of these principles. Interestingly, though, he does voice some worries about Counter-Closure. The name “Counter-Closure” seems to have been coined by Federico Luzzi [13].

- (1) *S* came to *know* (in contrast to coming to *merely truly believe*) that *Q* (*via* a deductive inference from her belief that *P*).

The epistemological symmetry between (C) and (CC), with respect to (1), involves the following central epistemological explanans.

- (2) *S* knew that *P* (and she maintained this knowledge of *P* throughout a competent deduction of *Q* from her belief that *P*).

(CC) and (C) are epistemologically symmetric here, in the sense that (CC) identifies (2) a necessary condition for (1), while (C) identifies (2) as a sufficient condition for (1).

In addition to this obvious epistemological symmetry between (C) and (CC), there are also some interesting epistemological *asymmetries* between them. Here, I'll focus on two of these asymmetries. First, those who accept (CC) are also inclined (or even mandated, given their other epistemological commitments) to accept the following generalization of (CC).

Generalized Counter-Closure (GCC). If *S* infers *Q* from her belief that *P*, (thereby coming to know *Q* (*via* said inference), then *S* knew that *P* (and she maintained her knowledge of *P* throughout the inference).

On the other hand, the corresponding generalization of (C) is (of course) not something that anyone should accept. That is to say, the following principle is clearly unacceptable.

Generalized Closure (GC). If *S* knows that *P* and *S* infers *Q* from *P* (while maintaining her knowledge that *P*), then *S* (thereby) comes to know that *Q* (*via* said inference).

(GC) is unacceptable for various reasons. For instance, (GC) doesn't even require that the inference in question was competently performed. This reveals an important epistemological asymmetry between (C) and (CC). The motivations/reasons for accepting (CC) — which we will be calling into question below — seem to support (GCC), whereas the motivations/reasons for accepting (C) do not support its (symmetrical) generalization (GC).

Another interesting epistemological asymmetry between (C) and (CC) can be brought out *via* an analogy³ between inference and entailment. Entailment (or whatever your favorite explication of entailment is) involves the (necessary) preservation of certain alethic features of premises. Classically, entailment involves truth-preservation (fancier theories now exist, but they typically also involve the necessary preservation of some “good” alethic property

³Note: this is *merely* an analogy (and a superficial one at that!). I, like Harman [9], am quite skeptical about the claim that “logic is normative for thought.” That is, I share Harman's worries about articulating probative bridge principles [14] between inference and entailment. Steinberger's excellent recent paper [18] does a great job of explaining (in detail) why this is so difficult. So, this analogy is not intended to be a deep one. It is only meant to highlight a specific alethic (and epistemic) asymmetry between (CC) and (C).

of premises). To be sure, entailment does *not* involve *falsity*-preservation. In other words, in logic, the “rules of inference” are meant to identify argument forms which must have true conclusions, if their premises are true. From this perspective, it would be very odd to try to describe argument forms which must have false conclusions, if their premises are false. This is simply not what an explication of entailment is supposed to do.⁴ Analogously, it is illuminating to recognize that (C) and (CC) can be re-stated in the following way.

Closure (C). The epistemic good-making feature of premises: *being known* is necessarily preserved by (single-premise) competent deductions (provided that the premise retains the property *being known* throughout the competent deduction).

Counter-Closure (CC). The epistemic bad-making feature of premises: *being unknown* is necessarily preserved by (single-premise) competent deductions.⁵

The idea that epistemically *good*-making features of premises should be preserved by virtuous inferences is a very natural one (and this is analogous to the idea that truth should necessarily be preserved by entailment). After all, I take it that one of the key functions of virtuous inferences is to expand our knowledge. However, the idea that epistemically *bad*-making features of premises should be preserved by virtuous inferences does not seem as natural (and this is analogous to the unnaturalness of the idea that falsity should necessarily be preserved by entailment). Why should it be incumbent upon a theory of virtuous inferences to explain (in any systematic way) what happens when we make virtuous inferences from *bad* premises? I’ll return to this crucial question at the end of this essay.

Finally, I close this opening section with one last (preliminary) asymmetry between (C) and (CC) — the significant epistemological differences between the alleged counterexamples to the two principles. Counterexamples to (C) — if there be such — tend to involve some sort of skeptical (or, at least, “heavyweight” [6]) conclusions (and non-skeptical/“lightweight” premises). As we’ll see throughout the rest of the paper, alleged counterexamples to (CC) tend to be quite mundane (*i.e.*, they tend to involve “lightweight” *P*’s and *Q*’s). I think this is another (and related) crucial epistemological asymmetry between the two principles. In any case, this opening section was merely intended to set the stage. Now, onto the main event.

⁴Indeed, any attempt to explicate both truth and falsity preservation simultaneously (as a matter of logical form) will inevitably lead to a trivial entailment relation, according to which the only valid form is $p \vDash p$.

⁵Indeed, as I explained above, I suspect that most advocates of (CC) would also be inclined to accept (GCC), which means that we actually have the stronger claim that

Generalized Counter-Closure (GCC). The epistemic bad-making feature of premises: *being unknown* is necessarily preserved by *all* (single-premise) inferences.

2 Alleged counterexamples to (CC) and The Standard Response

There is a burgeoning recent literature involving alleged counterexamples to counter-closure. Interestingly, this literature traces back to an alleged counterexample to generalized counter-closure (GCC) — an example which involves ampliative inference.⁶ Here, I present my own variation on this example (which is a bit simpler than the original, but having the same gist).

Urn. An urn contains 2 balls of unknown (to Sam) color distribution (each ball is either red or blue). Sam samples one ball (randomly, with replacement) from the urn many, many times. He is a very reliable counter and observer (and Sam knows all of the above facts). Sam then (competently, ampliatively) performs the following inference: “(*P*) I have (randomly, with replacement) sampled a red ball from the urn exactly n times in a row (where n is sufficiently large). \therefore (*Q*) Both of the balls in the urn are red.” And, in fact, both of the balls in the urn *are* red.

As it happens (and unbeknownst to Sam), the streak of red balls observed by Sam actually had length $n + 1$. So, *P* is false (hence unknown by Sam), but (intuitively) Sam (still) knows that *Q*. If this is right, then we *seem* to have a counterexample to (GCC). Of course, whether we *do* have a counterexample to (GCC) here will depend on whether Sam came to know that *Q* *via* his (competent, ampliative) inference from his belief that *P*. We’ll return to that key question shortly. Meanwhile, we’ll continue with a quick tour of the recent history of the dialectic concerning alleged counterexamples to (CC).

Since that original alleged counterexample to (GCC) appeared, many similar examples of apparent deductive inferential knowledge from a false premise — *i.e.*, many alleged counterexamples to (CC) — have been discussed. Here is a representative case.⁷

Handouts. Counting with some care the number of people present at my talk, I reason: “(*P*) There are 53 people at my talk; therefore (*Q*) 100 handout copies are sufficient.”

As it happens (and unbeknownst to me), my belief that *P* is incorrect. There are, in fact, only 52 people in attendance — I double counted one person who changed seats during

⁶The **Urn** example (or one very similar to it) was (as far as I know), the first published alleged counterexample to (GCC). And, it appeared in the context of Saunders and Champawat’s [16] reply to Clark’s [3] “no false lemmas” (NFL) response to the Gettier problem. It is interesting to note that many people seem to have rejected the (NFL) response to the Gettier problem. But, this is usually because they think it is too weak to rule out all Gettier cases. Interestingly, Saunders and Champawat argued both that (NFL) is too weak and that it is too strong. It is their latter claim that pioneered the contemporary discussions regarding (GCC) and (CC). Of course, this latter claim of theirs has been more controversial. See, for instance, Schnee’s [17] for a useful recent discussion of the relationship between (GCC)/(CC) and the Gettier problem.

⁷**Handouts** is reported by Luzzi [12], and it is very similar to many other examples discussed in the recent literature [11, 19, 7, 1]. As far as I know, the first person to discuss examples like **Handouts** was Risto Hilpinen [10]. Although, apparently, it wasn’t until Ted Warfield’s paper [19] (17 years later) that such examples were taken up as serious challenges to (CC).

the count. Again, intuitively, I (still) know that Q . If this is right, then we *seem* to have a counterexample to (CC).⁸ Of course, whether this *really is* a counterexample to (CC) will depend on whether I came to know that Q *via* my (competent, deductive) inference from my belief that P . At this point in the dialectic, the most popular defensive maneuver made by defenders of (CC) or (GCC) is to deny that I came to know Q *via* my deduction of Q from my belief that P . More precisely, the standard response is as follows.⁹

The Standard Response. In alleged counterexamples to (CC), the best epistemological explanation of why S knows that Q (supposing, *arguendo*, that S does know that Q in these cases¹⁰) does not make essential reference (*qua explanans*) to the fact that S competently deduced Q from her belief that P . Rather, the best epistemological explanation of why S knows that Q involves appeal to some other proposition P' — which I will call *the epistemic proxy* (or *the proxy*, for short) — such that (a) S knows that P' (or S is in a position to know that P'), and (b) P' serves to epistemicize Q for S (*viz.*, it is P' , and not P , that epistemically undergirds S 's knowledge that Q).

Here's the basic idea behind The Standard Response. In these alleged counterexamples to (CC), there are two distinct processes involved: a *psychological* process and an *epistemic* process. The psychological process involves S competently deducing Q from her belief that P , and (thereby) coming to *believe* that Q . But, the epistemic process does not (essentially) involve S 's deduction of Q from her belief that P . To be more precise, The Standard Response urges us to distinguish the following two explananda.

The Psychological Explanandum. S *believes* that Q (in contrast to *not believing* Q).

The Epistemological Explanandum. S *knows* that Q (in contrast to *merely truly believing* that Q).

According to The Standard Response, the psychological explanandum is best explained by appealing to the fact (*i.e.*, the psychological explanans) that S competently deduced Q from

⁸I'm being a little sloppy here about P and Q , since we need P to *entail* Q here. Just interpret " n is a sufficient number of handouts" as " $n \geq m$," where m is the number of people present at my talk.

⁹For two nice recent articulations of The Standard Response, see Martin Montminy's [15] and Ian Schnee's [17]. Peter Klein [11] seems to have been the first to articulate and defend The Standard Response.

¹⁰One could try to maintain that the agents in these cases do *not* know that Q . But, this strikes me as an implausible position to take. In any event, this won't really matter (for present purposes), since our aim here is to investigate the various maneuvers a defender of (CC) might make in her attempt to salvage (CC). We have chosen to focus on cases in which S does know that Q . And, I think this will allow us to cover the terrain of possible defensive maneuvers in a nearly exhaustive way. We could have chosen to examine the landscape in a different way (*e.g.*, by focusing on cases in which S does not know that Q and then examining the various explanations defenders of (CC) might offer as to why that is the case [17]), but it seems to me that this choice is a conventional one, which doesn't undermine the probative value of our discussion.

her belief that P . However, according to The Standard Response, the epistemological explanandum is *not* best explained by appealing to the fact that S competently deduced Q from her belief that P (which, of course, S was not in a position to know). Rather, there is some *other* proposition P' — which S is in a position to know — that features in the best explanation of the epistemological explanandum. In other words, P does *not* feature in the best epistemological explanans, but there is some proxy proposition P' that does. Two natural questions arise at this point: What *is* this proxy; and, *how* does it feature in the best explanation of why S knows that Q (as opposed to merely truly believing that Q)?

Initially, a natural conjecture about the proxy P' is that P' is something like the following (as applied to **Handouts**, but similar proxies will exist for similar examples, like **Urn**).

(P') There are *approximately* 53 people present at my talk.

It seems that P' is a plausible candidate proxy in **Handouts**. After all, (a) I am in a position to know P' , and (b) P' seems equally capable of doing whatever epistemicizing P was supposed to be doing (*vis-à-vis* my belief that Q). Unfortunately, this initial conjecture about the content of P' cannot be correct (in general). Consider the following case.¹¹

Marbles. As they swiftly roll by on the wooden track I have assembled for them, I count a series of marbles. The procedure yields 53 as a result. With some confidence, I come to believe that (P) there are 53 marbles on the wooden track. Recalling that my logic professor told me earlier that day that precision entails approximation, I competently deduce that (Q) there are approximately 53 marbles (without any loss of confidence in my belief that there are 53).

In **Marbles**, the corresponding “approximately proxy” would be:

(P') There are *approximately* 53 marbles.

Unfortunately, P' does not seem capable of serving as a proxy for P in **Marbles**. While (a) I am in a position to know P' , it seems incorrect to claim that (b) P' is equally capable of doing whatever epistemicizing P was supposed to be doing (*vis-à-vis* my belief that Q). After all, P' just is Q in **Marbles**, and (presumably) whatever epistemicizing P' needs to do here cannot be done by the conclusion Q itself.¹²

¹¹The **Marbles** example is attributed (by Federico Luzzi [12]) to Crispin Wright. Several years ago, I [8] independently came up with a class of examples that is very similar to **Marbles**. But, because Wright's example is simpler and more direct (for my present purposes), I have chosen to use it instead.

¹²There is an implicit assumption here that the epistemicizing P' needs to do here is *analogous* to the (alleged) epistemicizing that those who reject (CC) attribute to P in these cases. That is, in some sense, P'

In light of these considerations, it seems that this initial “approximately proxy” version of The Standard Response is inadequate. However, the defender of (CC) need not give up on The Standard Response (and, in fact, defenders of (CC) have tended to stick with The Standard Response here). There are other possible proxies that seem more suitable (in general). While there are many specific proposals out there (see, *e.g.*, [11, 15, 17, 19]), I think all of the most plausible precisifications of the content of P' fall under one of the following two (generic) proposals. Let E_P be S 's total evidence for her belief that P .

(P'_α) E_P .

(P'_ξ) E_P and if E_P , then Q .

Here, I am using the notation P'_α because (except perhaps in some rare cases, unlike the ones we've been discussing) P'_α will not entail Q . Hence, in general, the (propositional, evidential support) relation between P'_α and Q will be *ampliative*. On the other hand, P'_ξ will (always, *a fortiori*) entail Q . That is, the relation between P'_ξ and Q is (by design) *deductive*.

In all of the cases we have seen so far, this P'_α may *seem* suitable as a proxy for P . After all, it is plausible that (a) S is in a position to know that P'_α .¹³ And, it may *seem* that (b) P'_α is capable of doing whatever epistemicizing P was supposed to be doing (*vis-à-vis* S 's belief that Q). For instance, in **Marbles**, P'_α is distinct from Q , and it *seems* that P'_α should support Q strongly enough for S , so as to epistemicize S 's belief that Q . However, there is an important wrinkle here, when it comes to the *way* in which P'_α is able to satisfy (b). In the alleged counterexamples to (CC), it is important that the premise P entails the conclusion Q . After all, these are all (*prima facie*) instances of deductive inferential knowledge (and, in any event, Q surely is deduced from P in all of these cases). So, one might worry that, although P'_α seems to be able to epistemicize Q for S , it can only do so *ampliatively*; whereas, P was (if these alleged counterexamples are *bona fide*) doing so deductively.¹⁴ This explains why I have also introduced P'_ξ as a candidate proxy for P . For P'_ξ will also, generally, be such that

has to epistemicize Q in something like an “inferential” way. This is why if P' just is Q , then P' cannot properly epistemicize Q . This leaves open the conceptual possibility that P' is epistemicizing Q , but in a way that is not even analogous to the way premises epistemicize conclusions of virtuous inferences (in good cases). While I grant that this is a conceptual possibility, it doesn't strike me as a plausible (or helpful) one. For this reason, I will not discuss this line of defense of the initial version of The Standard Response.

¹³One might worry that (some of) S 's “evidence” for P could itself be false, or (more generally) that S might not be in a position to know (some of) S 's “evidence” for P [1, 5]. As it happens, in all the contemporary alleged counterexamples to (CC)/(GCC), this possibility is (as far as anyone seems to be concerned) not actualized. In the end, I will endorse a package of principles which has the consequence that some “evidence/reasons” (*i.e.*, some essential parts of S 's epistemic basis for believing that Q) may be false (*fn.* 16). But, in the present examples, I do think that true/known proxies can be found. In any case, I am now simply trying to charitably reconstruct a response to the alleged counterexamples, on behalf of (CC) defenders.

¹⁴Luzzi [12] voices a similar worry about P'_α -proxy strategies (although, he states the worry somewhat less generally). Luzzi does not, however, consider the possibility of moving to a “deductive proxy” such as P'_ξ .

(a) S is in a position to know that P'_ξ (*modulo* the worries in *fn.* 13), and (b) P'_ξ is capable of epistemicizing S 's belief that Q . Moreover, unlike P'_α , P'_ξ can satisfy (b) in the same way that P does — if these alleged counterexamples to (CC) are *bona fide* — *viz.*, by entailing Q .

So much for the alleged counterexamples to (CC), and The Standard Response to them. Next, I will explain my own worries about The Standard Response (in any of its forms).

3 My worries about The Standard Response

I think it is useful to visualize both the objections to (CC) that we've been discussing and The Standard Response to them (in a generic way) using diagrams that depict both the psychological process leading to S 's belief that Q , and the epistemological process leading to S 's knowledge that Q . Figure 1 provides just such a pair of diagrams.




Figure 1: Two psychological vs. epistemological pictures of alleged (CC) counterexamples

In Figure 1, the double-arrows represent the epistemological process leading to S 's knowing that Q , and the single-arrows represent the psychological process leading to S 's believing that Q . Both parties agree that the psychological process leading to S 's believing that Q involves S 's having performed a competent deduction of Q from her belief that P (this is why there is a single arrow — depicting this psychological inferential process — from P to Q in both diagrams). The diagram on the right of Figure 1 — depicting a view according to which the alleged counterexamples to (CC) are *bona fide* — collapses these two processes. On this *non*-bifurcated picture, (CC) must come out false, since P (which is not known by S) epistemicizes S 's belief that Q *via* S 's competent deduction of Q from her belief that P (which also explains why S believes that Q). The diagram on the left of Figure 1 — depicting The Standard Response — *bifurcates* these two processes. On this bifurcated picture, it is P' (which is known by S) that epistemicizes Q for S — despite the fact that S believes Q because S deduced Q from her belief that P (and not from P').

My main worry about the bifurcated picture (*i.e.*, The Standard Response) is that, although it allows us to salvage the *truth* of (CC), it does so at the expense of the (simultane-

ous) *explanatoriness* of (C) and (CC). And, ultimately, I think this incurs greater theoretical (epistemological) costs than any benefits that might accrue from salvaging the truth of (CC).

To see my worry, let's forget about the alleged *bad* cases of counter-closure for a moment, and instead let's think about apparent *good* cases of *closure*. These are (supposed to be) cases in which an agent *S* knows that *P*, competently deduces *Q* from *P*, and (thereby) comes to know *Q* *via* her competent deduction from *P*. Presumably, we can all think of many *prima facie* good cases of closure. Here's a toy case, to fix ideas. Suppose Miriam knows that (*P*) the glass before her contains water, and water is H₂O. Miriam then competently deduces that (*Q*) the glass before her contains H₂O (which we may assume she did not believe before performing the inference), thereby coming to believe that *Q* *via* her competent deduction from *P* (while maintaining her knowledge that *P*). According to the advocate of non-bifurcated picture in Figure 1 (right), Miriam's competent deduction of *Q* from *P* will feature essentially in any adequate (psychological) explanation of why she believes that *Q* and in any adequate (epistemological) explanation of why she knows that *Q*. Moreover, the advocate of the non-bifurcated picture will be able to offer this unified explanation for all such apparent good cases of (C). What about the advocate of The Standard Response, *i.e.*, the bifurcated picture in Figure 1 (left)? Are they entitled to the same unified explanation regarding all apparent good cases of (C)? This is unclear to me, for reasons I will now explain.

Let's call The Standard Response (*a.k.a.*, the bifurcated picture) *the disunified view*, and let's call the non-bifurcated picture *the unified view*. According to the unified view:

Actual Reasons are Epistemic Reasons (or, **Reasons**, for short). Whenever an agent *S* competently deduces *Q* from her belief that *P* (while maintaining her belief that *P*), thereby coming to believe *Q*, *P* is (an epistemologically explanatorily essential) part of *S*'s epistemic basis for her belief that *Q*.¹⁵

It is because the unified view accepts **Reasons** that it can offer a unified explanation of both the psychological and the epistemological explananda — in all good cases of (C). In contrast, according to the disunified view (*viz.*, The Standard Response), an agent *S* can competently deduce *Q* from her belief that *P* (while maintaining her belief that *P*), thereby coming to believe *Q*, *without P* being (an epistemologically explanatorily essential) part of *S*'s epistemic basis for *Q*. This is what allows the disunified view to salvage the truth of (CC), despite the apparent bad cases of (CC). Unfortunately, this same feature of the disunified view seems to prevent it from being able to undergird the desired kind of unified explanation of all apparent good cases of (C). For instance, let's return to our example involving Miriam. She competently deduced *Q* from her belief that *P* (while maintaining her belief that *P*),

¹⁵A similar principle has been proposed and defended by Arnold [1]. My emphasis on the psychological vs. epistemological explanatory roles of *S*'s belief that *P* is what makes my **Reasons** distinctive.

thereby coming to believe that Q . As such, the unified view (*viz.*, **Reasons**) implies that P is (an epistemologically explanatorily essential) part of her epistemic basis for Q , which is consonant with our usual intuitions about such cases. However, the disunified view does not seem (automatically) entitled to this same sort of explanation. It seems possible (as far as the disunified view is concerned) that Miriam’s epistemic position regarding Q is best explained *not* by the fact that she competently deduced Q from her belief that P , but instead by the fact that some proxy proposition P' has epistemicized her belief that Q . Specifically, consider (P'_α) Miriam’s (total) evidence for her belief that P (*i.e.*, her P -relevant evidence). Presumably, P'_α will have the requisite (availability and) power to epistemicize Miriam’s belief that Q — just as it would in an alleged bad case of (CC). Without some principled way of deciding when proxies (actually) play their (potential) explanatory roles, there would seem to be no principled way to block this alternative epistemological explanation of the fact that Miriam knows that Q (as opposed to merely truly believing that Q).

Explaining Miriam’s knowledge that Q *via* P'_α (and *not* her deduction of Q from P) seems to have two undesirable consequences. First, it would seem undermine the intuition that Miriam’s knowledge that Q was deductive, inferential knowledge (since P'_α does not entail Q). This problem can (apparently) be fixed by using P'_β , rather than P'_α , as the proxy for P , since P'_β does entail Q . However, this “fix” introduces a new explanatory problem, which can be brought out by the following chain of two single-premises inferences (the first deductive, and the second ampliative), which combines the key features of the **Urn** and **Marbles** cases.

Wright’s Urn. An urn contains 2 balls of unknown (to Wright) color distribution (each ball is either red or blue). Wright samples one ball (with replacement) from the urn many, many times. Wright is a very reliable counter and observer (and Wright knows all of the above facts). Wright then reasons as follows: “(P) I have sampled a red ball from the urn *exactly* n times in a row. \therefore (Q_1) I have sampled a red ball from the urn *approximately* n times in a row. \therefore (Q_2) Both of the balls in the urn are red.” And, in fact, both of the balls in the urn *are* red.

As in the **Urn** and **Marbles** cases, Wright seems to know both Q_2 and Q_1 in this case. Moreover, it seems that Wright knows Q_2 *via* ampliative inference and he knows Q_1 *via* deductive inference. But, his pair of single-premise inferences trace back to a false initial premise P . At this point, the defender of (CC) and (GCC) will choose some proxy for P . And, there are various proxy strategies that could be adopted. For instance, the defender of (CC) could choose one of the following proxies for P (in order to explain why Wright knows that Q_1).

(P'_α) E_P .

(P'_β) E_P and if E_P , then Q_1 .

If they choose P'_α as the proxy for P , then the inference from P to Q_1 — which seemed deductive — becomes *non*-deductive, since P'_α does not entail Q_1 . This problem can be fixed, by choosing $P^{1'}_\delta$ as the proxy for P , since $P^{1'}_\delta$ does entail Q_1 . However, in light of this maneuver, there seems to be nothing preventing the defender of (GCC) from providing a more direct response to this example, which involves the following alternative proxy for P .

$(P^{2'}_\delta)$ E_P and if E_P , then Q_2 .

If $P^{2'}_\delta$ is chosen as the proxy for P , then (a) there is no reason to appeal the first inference at all in the explanation of why Wright knows that Q_2 , and (b) the (now, direct) inference to Q_2 — which seemed ampliative — would be rendered deductive (or conclusive), since $P^{2'}_\delta$ entails Q_2 . What all of these examples illustrate is an important difficulty with any view of inferential knowledge that rejects **Reasons**. Once we reject **Reasons**, it becomes very difficult to provide a principled account of which propositions constitute the epistemic basis for conclusions reached *via (prima facie)* knowledge-yielding inferences.¹⁶

There is a second (and more patent) problem with appealing to the proxy strategy for explaining why S knows that Q , in good cases of closure. In good cases of closure, we are inclined to say that the deduction S performs is both psychologically and epistemologically explanatorily essential. That is, in good cases of closure we are inclined to accept (the implications of) **Reasons**. Indeed, I suspect that closure can be explanatory (in all the ways we want it to be) *only if* we accept **Reasons**. This is why (ultimately) I prefer the unified picture, and also why I'm inclined to reject (CC), while accepting (C). As I see it, there are (basically) the following two competing “packages” of views regarding inferential knowledge.

Package #1 (unified). Accepts **Reasons**, the truth of (C), and the explanatoriness of (C). Rejects the truth (and with it the explanatoriness) of (CC)/(GCC).

Package #2 (disunified). Accepts the truth of both (C) and (CC)/(GCC). Rejects **Reasons** and with it the explanatoriness of (C) and/or the explanatoriness of (CC)/(GCC).

That is to say, in the end, I think one has to decide whether the epistemological explanatory costs (as outlined above) of rejecting **Reasons** in order to salvage the truth of (CC)/(GCC) are outweighed by the benefits of salvaging the truth of (CC)/(GCC). I, for one, do not view salvaging the truth of (CC)/(GCC) as beneficial in the first place. This brings us back to my (then, mainly rhetorical) question from the opening section, to which I now return.

¹⁶It is important at this point to register a potential theoretical cost for the advocate of the unified view who accepts **Reasons** and (because of this) rejects (CC)/(GCC). This version of the unified view ends-up being committed to the possibility of “false evidence” (*i.e.*, false essential parts of S 's epistemic basis for her belief that Q). Because I want to limit the scope of this paper to the dialectic involving (CC)/(GCC) vs. (C) and the theory of inferential knowledge, I will not attempt to assess the potential costs of this commitment in a broader epistemological setting. Moreover, because I think several excellent essays have already been written about this issue, *e.g.*, [1, 5], I do not think I would be able to add much (here) to that debate anyhow.

4 Coda: What is required of a theory of inferential knowledge?

Recall my (then, mainly rhetorical) question from the end of section 1: Why should it be incumbent upon a theory of virtuous inferences to explain (in any systematic way) what happens when we make virtuous inferences from *bad* premises? In light of the considerations above, I am inclined to say that it is *not* incumbent upon a theory of inference to provide any systematic account/explanation of what happens when we make virtuous inferences from bad premises. Sometimes, doing this will saddle us with bad conclusions, and sometimes it won't. But, there is no reason to think that there will be a nice, (epistemologically) lawlike structure to these cases. Contrast this with what happens when we make virtuous inferences from *good* premises. We are inclined to think that there is a systematic, lawlike epistemological structure to these cases (which explains our adherence to principles like closure). As I explained in the opening section, I think this is analogous to our views about the lawlike structure of truth-preservation vs. the *non*-lawlike structure of falsity-preservation in the context of deductive logic. It would be a fool's errand to try to construct a simultaneous "logic of both truth and falsity preservation" (see *fn.* 4). And, similarly, I think it is a fool's errand to try to construct a theory of inferential knowledge which simultaneously accepts both (C) and (CC)/(GCC) as general (epistemologically explanatory) principles. Instead, I propose that we opt for a view which accepts **Reasons** and both the truth and explanatoriness of closure, but which rejects counter-closure (and, hence, generalized counter-closure).

References

- [1] A. Arnold, *Some Evidence is False*, *Australasian J. of Philosophy*, 2011.
- [2] R. Audi, *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (third edition), Routledge, 2010.
- [3] M. Clark, *Knowledge and Grounds: A Comment on Mr. Gettier's paper*, *Analysis*, 1963.
- [4] E.J. Coffman, *Warrant Without Truth?*, *Synthese*, 2008.
- [5] J. Comesaña and M. McGrath, *Having False Reasons*, to appear in *Epistemic Norms*, OUP, 2015.
- [6] F. Dretske, *The Case Against Closure*, in M. Steup and E. Sosa, eds., *Contemporary Debates in Epistemology*, Blackwell, 2004.
- [7] B. Fitelson, *Strengthening the Case for Knowledge From Falsehood*, *Analysis*, 2010.

- [8] ———, *Notes on Warfield's "Knowledge from Falsehood"*, http://fitelson.org/seminar/notes_6.pdf, 2010.
- [9] G. Harman, *Change in View*, MIT Press, 1986.
- [10] Hilpinen, R, *Knowledge & Conditionals, Philosophical Perspectives*, 1988.
- [11] P. Klein, *Useful False Beliefs*, In *Epistemology: New Essays*, 2008.
- [12] F. Luzzi, *What Does Knowledge-Yielding Deduction Require of its Premises?*, *Episteme*, 2014.
- [13] ———, *Counter-Closure*, *Australasian Journal of Philosophy*, 2010.
- [14] J. MacFarlane, *In What Sense (if any) is Logic Normative for Thought?*, manuscript, 2004.
- [15] M. Montminy, *Knowledge Despite Falsehood*, *Canadian Journal of Philosophy*, 2014.
- [16] J.T. Saunders and N. Champawat, *Mr. Clark's Definition of 'Knowledge'*, *Analysis*, 1964.
- [17] I. Schnee, *There is No Knowledge from Falsehood*, *Episteme*, 2014.
- [18] F. Steinberger, *Explosion and the Normativity of Logic*, *Mind*, 2014.
- [19] T. Warfield, *Knowledge from Falsehood, Philosophical Perspectives*, 2005.