

Commentary for Symposium on “Normative Inference Tickets” by Jennifer Foster and Jonathan Ichikawa

What is a normative inference ticket?

Foster and Ichikawa provide a theoretical framework to help us understand common, “knee-jerk,” inferences to “normatively significant conclusions” that are underwritten by stereotypes (p. 1). Like the inference tickets discussed in Prior (1960), these normative inference tickets have input and output conditions that license inferences of various forms.¹ But the reasoning involved need not be deductive.

The authors’ examples of inference tickets include HE-SAID-SHE-SAID

Input: “There are conflicting reports”

Output: “No one can know what happened”

and CAT LADY

Input: “having five cats”

Output: “being unmarried and undesirable”

Foster and Ichikawa emphasize that inference tickets can be epistemically unreliable (they may yield false beliefs) and morally pernicious (they can contribute to concrete harms and the perpetuation of oppression). They tend to be deeply entrenched, making them hard to avoid or unlearn even in the face of apparent counterexamples.

Yet Foster and Ichikawa also emphasize that these same features can make good normative inference tickets very helpful; they may constitute hermeneutical advances that remedy hermeneutical injustice in Miranda Fricker’s (2007) sense.²

Four General Questions about the Framework

1. What is the precise nature of the relationship between normative inference tickets and the concepts associated with the stereotypes they exploit?

The authors describe some examples, like HE-SAID-SHE-SAID, as being both “normative inference tickets” and “concepts” (pp. 1, 2).

But the authors don’t quite want to equate normative inference tickets with their closely linked concepts. Their response to Fricker is that in some key cases of hermeneutical lacunae, normative inference tickets (e.g. to the wrongness of sexual harassment) were unavailable even while the concept (SEXUAL HARASSMENT) was indeed possessed.

So what is the relationship between normative inference tickets and the relevant concepts and stereotypes with which they are closely associated?

¹ Prior, A. N. (1960). The Runabout Inference-Ticket. *Analysis*, 21(2), 38-9.

² Fricker, Miranda. (2007). *Epistemic Injustice*. Oxford University Press.

The authors also sometimes describe concepts as “supply[ing]” the mental shortcuts involve in normative inference tickets. (pp. 1, 2). How does this work, exactly?

Some of the authors’ remarks seem to suggest that each normative inference ticket will be associated with one cluster concept, where that cluster includes (among additional stereotypes) both the input and output conditions for the inference ticket. But it’s not totally clear to me that the authors intend for all normative inference tickets to work like this. According to the authors:

Normative inference tickets are closely connected to stereotype activation. (p. 2)

Stereotypes, as we will understand them, are *prototypical features* (like *x is athletic* or *having four legs*) which are associated *more or less tightly* with certain concepts or words (like BOY and ‘chair’), and are made more or less *salient* by occasions of those concepts or words. (p. 2)

Jenkins and Ichikawa use a dense core model of the concept DOG for general illustration (pp. 2-3). For the concept DOG, we find the properties *barks* and *four legs* close to the core, while *isn’t allowed in apartments* and *mammal* are further out.

How do inference tickets work, then? According to the authors:

[I]nference tickets exploit stereotypes that are close to the core of concepts. They can also bring stereotypes that might not be close to the core otherwise, closer. (p. 3)

The inference ticket for chick flicks then works like this:

The stereotype *is a movie with female leads* underwrites an inference ticket to *is a chick flick*; and *is a chick flick*, for a very significant population of speakers, is an inference ticket to *is not worth seeing*. (pp. 3-4)

The authors are naturally read as suggesting here that the chick flick inference ticket operates via activation of the cluster concept CHICK FLICK, where close to the core of the concept are the properties *is a movie with female leads* and *is not worth seeing*, which serve as input and output conditions for a non-deductive inference. This would be analogous, then, to an inference that proceeds from mention of Fido’s *ability to bark* (input) to activation of the whole concept DOG, yielding the non-deductive conclusion that Fido *has a good sense of smell* (output).

However, the authors’ later discussion of *he-said-she-said* emphasizes the “close to analytic” feel of an inference pattern from one stereotype to another. They do not emphasize, here, any role played by a large cluster of properties associated with a single concept. They do not provide a diagram of a dense core model of the concept HE-SAID-SHE-SAID. Instead, the authors offer simply the following input and output rules (p. 8):

X said p
Y said *not-p*

it’s a he-said-she-said situation
with respect to p

it’s a he-said-she-said situation
with respect to p

there is no way to know whether p

Yet it looks like the authors could have chosen to emphasize again here, by analogy once again to the DOG example, that we have a cluster concept of HE-SAID-SHE-SAID, where the input *X said p and Y said not-p* and the output *there is no way to know whether p* are close to the core of the concept. This would enable the easy, non-deductive inference via activation of the central concept. Further out from the core, maybe we would find properties like *scandalous* and *a source of gossip*, as well as *happened in private* (with no other witnesses).

- How important is it, in order to make use of a normative inference ticket, to have a single, central concept associated with a cluster of more and less tightly linked properties working in the background in order to supply the input and output conditions?

I suspect that the authors may have been hesitant to claim that all normative inference tickets are supplied by one concept, where both the input and output conditions are always properties associated with that single concept, because they want to deny that the inferences we make when we use normative inference tickets are reports of analytic truths. The claim that the input and output conditions are both “contained” (or anything like it) within the relevant concept may sound too close to analyticity for comfort.

The authors also want to distance their normative inference tickets from the kind of “runabout” inference tickets that Prior (1960) discussed with the infamous example of the problematic *tonk* logical connective, which is defined (solely) in terms of its introduction and elimination rules.

Yet I think this way of thinking (in terms of clusters of closely associated properties) helps the authors to explain why most or all normative inference tickets have a “feel” that is “close” to analyticity (when both the input and output conditions are properties near the core of one relevant concept) while still not involving inferences of strictly analytic truths. Using a dense core model of the central concept, the authors can deny that the associated properties are necessary and jointly sufficient components of the concept. It need not be an analytic truth that an example of a dog, or of a he-said-she-said situation, possesses any one of the properties associated with the central concept.

And as another friendly point, I think the authors need not be worried at all about an alleged similarity between their normative inference tickets and *tonk*. This is not just because the authors deny that normative inference tickets involve recognition of analytic truths. The authors can acknowledge that the normative inference tickets we use may be supplied by relevant concepts. Most importantly, the authors never claim that an inference is a *good* or *legitimate* one so long as it is supplied by concepts we already possess. This is an important component of their view, since they want to emphasize both the harm of pernicious stereotypes and the social benefit of hermeneutical advances through the introduction of (appropriate) new inference tickets.

It remains an important, substantive question for the authors which normative inference tickets are helpful ones and which are pernicious. This will presumably depend in part upon which concepts we should employ in the first place. Ethicists will have much to say about which normative properties we should associate with which non-normative ones (and how tightly). Discussion of *tonk*, on the other hand, was intended to illustrate problems with the claim that certain inferences are valid solely in virtue of the meaning of the expressions they involve.

2. What is involved in a hermeneutical advance, theoretically speaking?

I've mentioned that in response to Fricker's (2007) work, Foster and Ichikawa suggest that what was missing in many cases of hermeneutical injustice, and specifically hermeneutical lacunae, was not the whole concept itself (e.g. SEXUAL HARASSMENT) but a communal normative inference ticket, e.g. from "this is happening" to "this is wrong" (p. 1, 12).

When a hermeneutical advance occurs, what do the authors think must, or usually does, occur?

- Do we add properties or other features to the relevant concept (or relevant cluster of concepts), which serve as new input and/or output conditions for a new inference?

In this example, do we add *this is wrong* to the concept of SEXUAL HARASSMENT, so that it now serves as an output condition for a new normative inference?

- Or does the original concept remain totally unchanged, while the disposition to make a particular inference (e.g. from "is occurring" to "needs to be reported"), becomes much more common and natural among members of our language community?

Or as a middle ground between these first two options, perhaps the latter occurs when the property *is wrong* moves closer to the *core* of the concept SEXUAL HARASSMENT?

- Do these advances ever involve the adoption of entirely new concepts, which would support new inferences? (That is, even if Fricker's sexual harassment case is not supposed to be one of them?)

I might suggest STRUCTURAL INJUSTICE or HEALTH DISPARITY as potential examples of relatively new concepts, which now, in our community of language users, license non-deductive inferences from descriptive input conditions to normative output conditions.

I'm also curious what Foster and Ichikawa think of the process of "reclaiming" pejorative terms. I presume this mostly involves changing the output conditions for associated inferences.

e.g. from *deserves to be humiliated* to *is a source of pride*

3. What is involved in achieving a hermeneutical advance, practically speaking?

To pick up on the authors' concluding thoughts about the broadest aims of their project, this question could be rephrased as a request for, among other things, examples of input and output conditions for the inference ticket NORMATIVE INFERENCE TICKET. How do we employ the relevant concept to help us understand and improve our normative reasoning?

But I presume that hermeneutical advances will require practical achievements as well as merely intellectual ones; we cannot, unfortunately, just reason ourselves into better habits.

How might we bring ourselves, and others, to make fewer problematic inferences (e.g. from “He said S and She said Y” to “We can’t be sure what happened”) and more productive ones (e.g. from “This student has claimed harassment” to “This needs to be reported to the university”)?

- Should we object to others when we hear them make problematic inferences?
- Encourage ourselves, and others, to replace problematic inferences with better ones?
- Refrain from employing certain concepts altogether? (e.g. CAT LADY, CHICK FLICK)
- Something else?

4. What, if anything, makes normative inference tickets uniquely normative?

Foster and Ichikawa are especially interested in inference tickets where the output conditions are evaluative judgments (e.g. “is undesirable” for CAT LADY) or normative claims (e.g. “y should call x on Mother’s Day”). This is because the authors are especially interested in the ethical and political ramifications of our inference patterns.

It seems to me that there’s nothing about this general framework for describing normative inference tickets that requires that the relevant output conditions be normatively “loaded.” Presumably we also have many non-deductive inference tickets that take us quickly and almost automatically in a very similar way from one description (e.g. has pine needles) to another description (e.g. grows in a place where it snows) via the activation of concepts like PINE TREE. Perhaps the authors think most of our non-deductive everyday inferences work like this? This isn’t necessarily a problem; it may be a feature of the view that it employs a widely applicable general framework.

I’m reminded by the authors’ work of the meta-ethical discussion of “thick” ethical concepts that involve both descriptive and evaluative elements (like COURAGEOUS as involving traits like *willing to subject oneself to danger* and *able to continue working under intense stress* as well as *praiseworthy*).

- On the authors’ intended picture, do most of our concepts (like DOG) include some evaluative and normative elements in addition to a wide variety of non-normative properties?

Since the authors emphasize the ubiquitousness of normative inference tickets, I wonder whether they’re attracted to the interesting and provocative (but highly plausible, I think!) conclusion that most of our concepts are, in a sense, thick evaluative and/or normative concepts. That is, most of our concepts involve evaluative and normative content that may serve as output conditions for inference tickets.