Better Rationality Attributions and Requirements: Contextualising the Local and the Global

Sara intends to send off her paper by the end of the day. She believes that skipping her lunch break is a necessary means for achieving this. But Sara's favourite Italian restaurant runs a pizza lunch special today and she goes for a 2h-lunch break. What should we say about Sara's rationality?

There are at least two different available answers. We could say that Sara is clearly irrational. She fails to intend the means she believes to be necessary for her ends, thus violating a meansend coherence requirement of rationality. Alternatively, we might think that in the interest of selfcare, Sara has good reason to stop obsessing over her paper and take a lunch break. She might have complied with a different requirement of rationality – to intend what you have most reason to intend.

These two different judgments of Sara's attitudes can be explained as the difference between a *local* and a *global* perspective of evaluation. Proponents of the *local* perspective argue that Sara is irrational because of the specific conflict between her intention to send off her paper by the end of the day, her belief that skipping lunch is necessary for this, and her absence of intending to skip lunch. They hold that "our ordinary judgments about rationality [...] are focused on specific conflicts among one's attitudes" (Kolodny 2005: 515-516). Proponents of the *global* perspective leave it open that Sara might be rational if taking a lunch break is rational in view of all of her attitudes. They argue that rational requirements are "global constraint[s] on all of one's attitudes" (Kolodny 2005: 515-516) and "specify what's rational in light of the totality of an agent's attitudes" (Brunero 2012: 129).

This distinction between the *local* and the *global* perspective is at the heart of this paper. Note that the local/global distinction extends to both rationality requirements and attributions. Choosing between these perspectives amounts to taking a stand on whether rationality is a matter of coherence amongst *all* of an agent's propositional attitudes or a *smaller subset* of these attitudes, with rationality requirements applying to and rationality attributions referring to either the entire set or a subset.

I think that it is a mistake to treat *local* and *global* as mutually exclusive, thereby depriving oneself of the benefits and insights of the other perspective. In this paper, I develop a new contextual account of rationality requirements and attributions that has two advantages: it attributes an important role to an agent's propositional attitudes and is flexible enough to provide a nuanced and accurate answer for cases like Sara's.

Before turning to my account, I consider two plausible contenders and identify their

advantages and disadvantages. These are the desiderata that my account will have to meet. The options are the strictly local view (LV) – which only allows for local requirements and attributions – and the mixed view (MV) – which allows for both local and global requirements and attributions.

Whilst (LV) is compatible with Kolodny's claim that our intuitive rationality attributions are local, it is unable to account for process requirements without losing its status of a *strictly* local requirement. Endorsing (MV) instead could be a natural reaction. This has the additional advantage of providing a reply to the symmetry objection against the wide-scope version of a means-end coherence requirement of rationality. The objection criticises that we cannot differentiate between the different options of complying with the requirement, despite their differing rationality. But by allowing for the local and global perspective, we can see that "there is no inconsistency in claiming both that (1) two ways of proceeding are on a par as far as instrumental rationality goes, but (2) considering the other attitudes of an agent, one way of proceeding is more rational than another" (Brunero 2012: 129). But (MV) has its shortcomings. In particular, it lacks flexibility. What is relevant is often neither the very small subset, nor the entire set of attitudes but rather a subset in-between. (MV) however only operates on the extremes.

The benefits and shortcomings of (LV) and (MV) suggest four desiderata: my account should (i) respect intuitive rationality judgments, (ii) account for process requirements, (iii) explain asymmetry and (iv) be flexible.

On my account, rationality attributions of the form "*S* is rational" should be understood as short-hand for "*S* is rational with regard to subset a_i of her attitudes and degree of robustness *r*". As a result, rationality attributions are always indexed to a particular subset of the agent's attitudes. This makes the account flexible: we can maintain the local and global perspective (corresponding to the small subsets a_1 , a_2 , ..., and the entire set a_n) and everything in-between, since we can always take more attitudes into consideration.

The robustness measure guarantees accuracy and intuitiveness of the account. *r* tracks the agent's rational performance, were we to assess a larger attitude subset. This is to avoid distorted pictures, in cases of agents that do well with regard to small subsets but immediately violate rational requirements when we take into consideration bigger subsets.

Finally, rationality attributions are also sensitive to context. Depending on the deliberative question, the set of attitudes to be evaluated shifts. In some contexts, it is necessary to take into account many or all of an agent's attitudes in order to assess her rationality. In other contexts, plausibly only a small amount of attitudes is relevant. Importantly, we can also imagine intermediary contexts where we do not have to consider all attitudes but still more than just a few. Note that a shift in evaluated sets of attitudes leads to different requirements becoming relevant,

namely the ones triggered by the attitudes in the set. Regardless of the size of the subset at issue, robustness still features in the resulting rationality attribution. This way, we have an idea of the agent's overall rationality even if our primary concern is a specific local conflict. In my conclusion, I apply this to Sara's case and show how my account meets (i)–(iv).