

1 We thank the reviewers for their constructive and fair reviews. We do our best to address the points brought up by the  
2 reviewers in the text below.

3 We agree that including the ethics and impact sections in the main body would have been preferable, given unlimited  
4 space. Given the space constraints, we chose to discuss our bias analysis and other broader impacts work in the separate  
5 Broader Impacts section, as we felt it was important to give this work a more thorough, contextual treatment, which  
6 would have been difficult in the main body. We hope that future NeurIPS formats will be more supportive of a tighter  
7 integration between the main body and Broader Impacts section.

8 One reviewer expressed curiosity about whether increasing scale will continue to yield results. As noted in the paper,  
9 we do not see significant deviation from the power law trends across eight orders of magnitude. We expect an eventual  
10 plateau, but we do not yet have evidence as to where this plateau will take place.

11 Finally, we appreciate the suggestions of adding additional studies of (1) fine-tuning, (2) the effect of context length, (3)  
12 causal reasoning, (4) numerical reasoning, (5) stylistic variations and (6) long-term coherence. We agree that all of  
13 these studies would be valuable, but felt that they could each merit their own papers to fully explore, so we leave them  
14 as fertile avenues for future work.