We would like to thank the reviewers for their comments and feedback. We are aware that in a largely conceptual paper like ours there are subtleties, and highly appreciate the time and effort that the reviewers are putting in to digest these.

Reviewer #1: Causal Shapley values (SVs) are defined in Section 2. These do *not* coincide with what [9] and others call the interventional SVs (marginal SVs in our terminology). Janzing et al. [9] write down the same equation, but then choose to ignore any dependencies between the features in the real world (e.g., that in summer it tends to be warmer than in winter). We do choose to incorporate these dependencies and hence cannot simplify to $P(\mathbf{X}_{\bar{S}}|do(\mathbf{X}_S = \mathbf{x}_S)) = P(\mathbf{X}_{\bar{S}})$, but keep $P(\mathbf{X}_{\bar{S}}|do(\mathbf{X}_S = \mathbf{x}_S))$ in our definition of the causal SVs. We will follow the reviewer's suggestion to make this more explicit in Section 2. This distinction then hopefully also resolves the reviewer's issue about the indirect effect: it indeed vanishes for marginal SVs, but need not vanish for causal and conditional SVs. See also the examples in Section 4 (Figure 1). The decomposition for conditional SVs follows by replacing "conditioning by intervention" with "conditioning by observation", i.e., by replacing $do(\mathbf{X} = \mathbf{x})$ with $\mathbf{x}$ on the righthand side of the bar. The decomposition is introduced in Section 3 to assist our illustration of how the different SVs attribute a model's prediction to the features involved in this prediction in Section 4 for different causal models. Here we also discuss in which cases (most notably the fork and the confounder) conditional SVs fail to provide an intuitive causal attribution.

Causal chain graphs are introduced as a means to compute causal SVs (whether symmetric or asymmetric) when users are willing/able to specify a (partial) causal ordering, but not a full-fledged causal model. The asymmetric SVs of [6] indeed rely on the same information. On top of [6] we offer a formalization in terms of causal chain graphs and show that, with "conditioning by intervention" instead of "by observation" as in [6], there is no need for asymmetry in the SVs. Unlike conditional (asymmetric) SVs, causal SVs provide the right intuition in the case of common confounding.

Reviewer #2: W.r.t. the novelty in comparison to [6]: asymmetric (conditional) SVs as defined in [6] in some cases coincide with symmetric or asymmetric (causal) SVs, but are different in general. See also the previous paragraph.

Section 4 aims to illustrate the behavior of the various SVs in simple cases that can be analyzed analytically and then to argue which is the most intuitive, indeed also linking to psychological literature when appropriate. Here one prominent theory, dating back to [15], states that humans sample over different possible scenarios to judge causation. Translating this to a situation in which there are two possible causes, X_1 and X_2, where it is unknown which one is intervened upon first, may suggest that the natural interpretation is to consider both options and average over them.

We fully agree that quantifying causal influence is a difficult topic and any method has its weaknesses, but causal SVs appear to fare better than the reviewer suggests. Discontinuity w.r.t. arrows with zero strength is an issue for the asymmetric SVs, but not for the symmetric SVs that consider all orderings, not just those consistent with the causal DAG. After averaging over all these orderings, the indirect effect already does incorporate all possible paths (so we do not see how or why it needs to be generalized), but of course in the game-specific way inherent to the Shapley value approach. We will add comments and disclaimers to clarify this and adapt our description of Janzing et al. and related work as suggested by the reviewer. Our statement 'not every causal query need be identifiable (see e.g., [24])' did not presume DAGs with all variables observed, but more general causal structures possibly including latent variables.

Reviewer #3: W.r.t. the scope, see our answer to Reviewer #1 (third paragraph) and the beginning of Section 5: causal SVs are generally applicable when a user is willing/able to specify a causal model among the features that are used as input to the model and when all causal queries are indeed identifiable. Specifying when this is the case is a topic on its own: we will add more references (see also the supplement). Causal chain graphs are "just" proposed as a practical approach to handle partial causal knowledge. In causal chain graphs, all causal queries are guaranteed to be identifiable and can be answered based on the available observational data. These graphs allow for handling cycles, confounders, etc (see Figure 2). In fact, all examples in Figure 1 are easily translated to causal chain graphs. An illustrative example for the fork could be predicting hotel occupation (Y), based on season (X_2) and temperature (X_1).

We miss the point the reviewer tries to make w.r.t. counterfactual analysis. As far as we can tell, the counterfactual question posed by the reviewer (assuming all features are known) can be answered simply by reading off the output of the model. Our analysis can be interpreted as counterfactual (third rung) reasoning to analyze what the model prediction would have been had we not known some of the input features (see second paragraph of Section 4). Counterfactual *explanations* as in e.g. [33] may be improved with similar techniques, but are beyond the scope of the current paper.

Causal relationships are indeed asymmetric, but that does not prevent the causal SVs from being symmetric according to the standard symmetry axiom for SVs (see the definition in Section 2 and the elaborate discussion in [9], Section 3 in response to Sundarajan and Najmi, 2019). We chose not to repeat this argumentation, but will add a reference.

Figure 4 is meant to illustrate the difference between the various SVs (asymmetric SVs focus on the root cause, marginal SVs on the direct effect, symmetric causal SVs consider both), not necessarily to claim that one is always better than the other. We will extend the supplement with additional empirical analyses, e.g., on (deep) neural networks.

(7) indeed should have been (6). We will fix the other minor issues, also those rightfully indicated by **Reviewer #4**.