We would like to thank you for your time and valuable feedback. Thank you for helping us to improve our manuscript! We hope to have correctly understood your questions, and will try to exhaustively address all your comments.

Reviewer 4 questions the applicability of our theory for hard or atomic interventions. It is important to note that in our framework (Setting 1), interventions are represented as additional variables in the graph, not as edge deletions (also see section 3.2.2 of *Causality*[1]). This captures a wide range of interventions, including atomic interventions. The d-separation statement which defines intervention stable sets (Definition 2.1) is a statement made on this extended graph, instead of on the mutilated graph resulting from removing incoming edges to variables which receive atomic interventions. Loosely put, the idea is that stable sets block the flow of information between the response Y and the exogenous variables whose distributions change with the environment, resulting in invariance (c.f. proof of Proposition 2). We agree to be more specific as to what we mean by "other types of interventions" in footnote 1, p. 3, and will change this to echo line 129. We thank reviewer 4 for the additional comments on the manuscript.

The proposed policies consider single-variable interventions (line 195): at each iteration, a sample is collected from an experiment where only one variable is intervened on. The theoretical results presented in section 2 also hold for multiple-variable interventions (line 129). Since the policies shrink the pool of possible targets and then pick at random, we could simply pick k targets instead of one. We considered single-variable interventions for simplicity and to compare with ABCD, which also only considers single-variable interventions.

We agree that the base assumptions of ICP are strong; regarding their validity in practice, the method has shown competitive performance when applied to real gene expression data[2]. We agree with reviewer 4 that it is difficult to claim that there are no confounders between the target and its parents. Section 5 in ICP [27] outlines ideas for generalizations of the method in the presence of hidden confounders; these were materialized in the causal Dantzig[3] method. Combining this method with A-ICP is interesting future work. In any case, section 6.3 of the ICP paper considers the behaviour of the method under model violations (such as hidden confounders), with the result that these often do not have severe consequences for the coverage guarantees of ICP (albeit result in a loss of power).

Reviewer 2 asks if it is possible to directly intervene on the response variable to identify the parents. We are not entirely sure what the reviewer has in mind. It is true that if the Markov equivalence class is known, one could intervene on the response variable and perform conditional independence testing to orient all its surrounding edges. However, this would limit the interventions to be do (hard, atomic) interventions. Furthermore, estimating the Markov equivalence class is difficult with a finite sample (see related work), even more so when only a small observational sample is available, as we allow for our method. Another approach would be to intervene on the response and check which variables undergo a change in their marginal distributions, i.e. directly testing for a causal effect. However, this would only partition the predictors into variables which are downstream of the response and not. We hope this answers the reviewer's question.

Increasing the number of predictors (i.e. graph size) had the effect of lowering the relative performance of the random strategy, as its chances of picking a parent were reduced. Otherwise, results were very similar and we decided to not present them separately. An interesting question is how different signal-to-noise ratios affect the performance of the policies. To partially answer this question, Figure C.11 illustrates their performance for different intervention strengths.

We will provide a detailed analysis of the computational complexity in the final version. In any case, it is exponential in the number of predictors, which limits the applicability of our method to "large p" settings. However, we envision settings where the time needed to carry out an experiment far outweighs the computation time to select the next experiment, which is common in empirical sciences like biology. Thus, A-ICP could still be useful in settings with a moderate number of variables. We would like to note that in our experiments, the time required to run a single iteration of ABCD far exceeded that of A-ICP, due to the approximation of the posterior.

We believe that ABCD and A-ICP are complimentary in many ways, not only regarding their input requirements. We agree with reviewer 1 that, in terms of Jaccard similarity, ABCD has a better performance in the first few iterations, and would be the preferred method if this was the goal. On the other hand, if false positives incur a high cost, the more conservative A-ICP would be preferable, as it offers a strict control over the family-wise error rate.

Finally, we would like to stress that A-ICP is directly applicable to the nonlinear extension of ICP [14], as the theoretical results used to construct the policies make no assumption of linearity (lines 119, 322-325). We did not show results for nonlinear models for two reasons. First, the performance of nonlinear ICP is not great even for simple graphs; we feared this would confound the performance of the policies, whose evaluation was our main goal. Second, and along the same line, the Markov blanket estimation (for the policies that use it) is not as straightforward for nonlinear systems as it is in the linear case.

[1] Pearl, J., 2009. Causality. Cambridge university press. [2] Meinshausen, N., Hauser, A., Mooij, J.M., Peters, J., Versteeg, P. and Bühlmann, P., 2016. Methods for causal inference from gene perturbation experiments and validation. PNAS, 113(27), pp.7361-7368.

[3] Rothenhäusler, D., Bühlmann, P. and Meinshausen, N., 2019. Causal Dantzig: Fast inference in linear structural equation models with hidden variables under additive interventions. The Annals of Statistics, 47(3), pp.1688-1722.