We thank the reviewers for the constructive feedback and are happy to provide clarifications.

To reviewer 1:

We would like to stress the benefit of this work. The advantage of dynamic models such as Markov Process models is the ability to predict the effect of interventions on complex systems (e.g., drug interventions). Applying policy evaluation methods from reinforcement learning to the policy of selecting optimal interventions is useful, as it increases the efficiency of the process. This manuscript considers policy evaluation for this class of models, and extends it to counterfactual policy evaluation.

The counterfactual policy evaluation is useful in two ways. First, it saves resources. For example, it gives a drug company a way to predict the effectiveness of a perturbation using past data collected under a different policy, without expending resources in collecting data under the proposed policy. Second, predictions anchored on past data are more robust to model misspecifications than the pure intervention predictions of Markov process models. The revised manuscript will emphasize these benefits of this approach, will provide a detailed discussion of the counterfactual offline policy evaluation use case in Supplementary Materials, and will comparing intervention predictions from the derived SCM with predictions from the original Markov process model under model misspecification.

Reviewer 1 comments that our focus on the Binomial distribution is restrictive, and suggests for coverage of more general cases. We would like to clarify that the proposed approach works with any closed-form equilibrium conditional probability distribution. The challenge is in whether it is possible to derive this distribution from a given Markov Process model. Generally, with enough simplifying assumptions, it possible to derive any conditional probability distribution. We previously avoided this discussion because the validity of the assumptions depends heavily on the application domain. However, we agree with the reviewers that we undershot the mark, and will add this discussion to the manuscript, as well as a more extensive discussion to Supplementary Materials. To be specific, we will include in Supplementary Materials a demonstration of the approach with a Poisson-distributed conditional probability distribution, derived from a slight adjustment in the assumptions in our examples.

To reviewer 2:

Reviewer 2 advocates for proofs of soundness and completeness. Indeed, we would like to clarify that for any probability model there exists a class of SCM models that are equivalent to that probability model in distribution [1]. We will emphasize this point in the revised manuscript by including a lemma that translates this general result to our case. The lemma will state that if there exists a closed-form equilibrium probability distribution of the Markov Process (MP), then there exists a class of Structural Causal Models (SCM) that are equivalent to the MP model's equilibrium probability model both in observations and in interventions. Our monotonicity constraint then selects an SCM model from that class, in a way that enables the identification of key counterfactual quantities necessary for policy evaluation. As mentioned in response to Reviewer 1, modelers can apply simplifying assumptions that attain an equilibrium probability model, subject to the validity of these assumptions within a problem domain.

Reviewer 2 suggests that the audience would benefit from a clearer and less technically dense description of the key research findings. We take this suggestion to heart and will adjust the text to make the key findings more succinct, and move biological exposition and mathematical details that are not essential to communicating those findings to Supplementary Materials.

To reviewer 3:

We are grateful to Reviewer 3 for the suggested corrections and clarifications and will fix or clarify each of these points in the final manuscript. Due to the space constraint of this letter, we answer three of Reviewer 3's broader questions.

Reviewer 3 asks about extensions to cyclic graphs. In this work, we limited our modeling assumptions to the acyclic graphs. However, we agree that since MP models and SCMs accommodate cycles, this is a logical avenue for future work. We will discuss this limitation and extension to future research in the discussion section.

Reviewer 3 asks about the equivalence of distributions for X_i^* and X_i under the do-operator. This equivalence is implicit in the proof in Section 2.4.1 of Supplementary Materials. In response, will provide more explicit answers. We will also include in Supplementary Materials a discussion of the recent work connecting interventions on ODE models and SCMs, and its relationship to the proposed method.

Regarding the reviewer's question about assuming the "noise" to be the same across these distributions in lines 207-209, the answer is yes.

[1] Peters, Janzing, Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*, MIT Press, 2017