We thank all the reviewers for their detailed and insightful feedback.

Reviewer 1: We thank the reviewer for the detailed feedback and pointers to missing literature.

1. We would like to clarify that the buyer in our scenario is not just "strategic up to a certain accuracy factor". Our buyer is an exact utility maximizer whose objective is to maximize her long-term cumulative utility. This is different from the ϵ-strategic buyer proposed in [Mohri and Medina, 2015], which allows for ϵ-suboptimal responses. In our work, the notion of η-DIC is a property of our mechanism, not an assumption on the buyer's behavior. The property ensures that a utility maximizer will always report bids close to her true valuations (i.e., within a bounded additive error).

2. An important distinction with [Drutsa, 2017, 2018] is that our work focuses on stochastic settings in which the buyers' valuations are redrawn independently at each round, whereas their work considers a fixed valuation model. In the stochastic setting, the revenue gap between the optimal dynamic mechanism and the optimal static mechanism could be arbitrarily large [Papadimitriou et al., 2016], while there is no difference for the fixed valuation model. As for [Cohen et al., 2016], they consider robust pricing in repeated contextual auctions for myopic buyers, while we consider a buyer who aims to maximize her long-term utility. We agree, however, that we should improve the coverage of the fixed-valuation model and the existing works on robust pricing in the related work. We will include the citations you brought up and add a discussion of optimal regret bounds under fixed valuations in the revision.

3. Our model of contextual auctions is both stochastic and parametric: The buyer's valuation distribution is characterized by her fixed but private preference vector, while the features of the items and the distributions of market noise at each stage are stochastic. Therefore the buyer's realized value at each stage is stochastic. This directly follows and extends the contextual auction model proposed in the literature [Amin et al., 2014, Golrezaei et al., 2018].

4. We would have loved to discuss more in the main body. But given the technical nature of our proofs and the 8-page limit, we had to move most proofs to Appendix. We will make sure to improve our presentation by providing more high-level ideas and proof sketches, and make space for conclusions in the revision.

5. Our robust non-clairvoyant mechanism only relies on Assumption 2 and does not depend on the parametric valuation model. Our no-regret policy for contextual auctions does depend on the valuation model. It would be an interesting avenue for future work to consider nonlinear valuation models or even non-parametric models. We would like to emphasize though that we view the robust non-clairvoyant mechanism as the most important contribution, as it opens the possibility to take an existing robust pricing mechanism for contextual auctions from [Golrezaei et al., 2018] (which is no-regret against the optimal static benchmark) to obtain a mechanism with no-regret against the optimal dynamic benchmark (albeit with several technical updates), rather than just the optimal static benchmark.

6. We do not have a matching lower bound, and we think it is an important open question to improve the regret guarantee or to provide a matching lower bound. We consider the upper bound a good advancement since there was no previous no-regret bound against the optimal dynamic benchmark in our setting.

Given that these were Reviewer 1's main concerns, we sincerely hope Reviewer 1 would consider revising their score.

Reviewer 2: We thank the reviewer for the positive and encouraging review.

1. Why these four mechanisms: The give-for-free mechanism, the posted-price auction with fee, and the Myerson's auction are used to guarantee a $\frac{1}{3}$ approximation against the dynamic benchmark (when the distributional information is perfect), while the random posted price auction is added to obtain robustness.

2. Intuition behind the hybrid mechanism: The hybrid mechanism uses static stage mechanisms for stages with small a_t and dynamic stage mechanisms for stages with large a_t. Intuitively, for the static mechanisms, we first use give-for-free mechanisms for the first few items with small a_t in each phase to accrue a large enough bank account balance for the buyer. Later, this allows us to almost always implement a give-for-free mechanism with extra fee $\mathbb{E}[v_t]$ for later stages with small a_t to extract full welfare. The key observation is that the dynamics of the bank account balance in give-for-free mechanisms with extra fee is a martingale and therefore, starting with a large enough balance, the probability that the balance becomes close to 0 is small. For stages with large a_t, we apply the same dynamic stage mechanisms as the non-clairvoyant mechanism proposed in Mirrokni et al. [2018]. Finally, both the static and dynamic stage mechanisms are made robust using our framework, by mixing in the random posted-price auction.

3. Assumption 1 is only used in the analysis of our no-regret policy and our robust non-clairvoyant mechanism only depends on Assumption 2. Without Assumption 1, simply consider a worst case scenario in which $a_1 = T$ while $a_t = 1$ for all $t > 1$. In this case, the revenue loss could be $\Omega(T)$ from the first stage since the seller has no information to estimate the buyer's private preference at the first stage.

Reviewer 3: We thank the reviewer for the positive and encouraging review. We will make sure to improve our presentation by providing more high-level ideas and proof sketches, and make space for conclusions in the revision.