General Response. Since some comments overlapped, we first provide a general response before individual questions.

• Some reviewers mentioned the exact prediction assumption in the paper. Definitely the inexact prediction case is very interesting and important, and we are currently working on that. However, we would like to emphasize:

a. Exact predictions are very common in online learning literature (e.g., [23], [24], [16]), where the focus is on how future information improves algorithmic optimality (the metric is typically regret or competitive ratio).

b. Note that the disturbance w_t in our model is much more general than just "noise" — in Section 4 we consider the most general jointly stochastic w_t (not necessarily i.i.d.) and in Section 5 we consider adversarial w_t. Our setting supports a lot of real-world applications. For example, in the LQ tracking example in Section 2, predictions of w_t are from desired trajectories, which are usually predefined and exact in robotics (as Reviewer 2 mentioned, *"An online high-level trajectory planning algorithm often outputs some states in the future for the low-level controller to track."*). In other words, the exact prediction assumption is actually practical in many scenarios, since our w_t is beyond unstructured "noise".

c. If the predictions are inexact, the exponentially decaying properties will be kept and we will have a residual term for prediction errors (i.e., regret becomes $O((\lambda^k + \sum_{i=0}^{k-1} \epsilon_i \lambda^i)^2 T)$, where ϵ_i is the prediction error), according to our current research results. We would like to add the inexact case since it is a straightforward add-on.

• Some reviewers asked about the relationship between our results and more general MPC settings, especially robust MPC, learning unknown dynamics in MPC, and stochastic MPC. Some of these settings are indeed relevant to our work or can be integrated with our work, so we will definitely cite more papers in this field and reorganize them (if accepted we will have one more page to do that). However, we would like to emphasize:

a. In this paper, we adopt a conventional definition of MPC as an online deterministic optimal control problem with a finite-time horizon with dynamics constraints. This type of MPC is the most natural and simple policy given deterministic and finite future predictions — it is why we call it a "greedy and naive" policy in our paper. Definitely the general MPC family is more sophisticated (e.g., considering constraints, robustness and safety), but our goal is to show future predictions allowing simple algorithmic ideas to be effective and to connect online learning with control.

b. Completely different from robust MPC, learning unknown dynamics in MPC and stochastic MPC settings, this paper focuses on **optimality** and uses standard metrics in the online learning community (dynamics regret and performance ratio). We focus on the optimality gap between MPC, the optimal online policy and the optimal offline policy (with the knowledge of all future) in this paper. On the other hand, other settings in MPC mentioned by R1 and R4 focus on robustness, stability and safety. As mentioned by R2 (*"As far as the reviewer is aware, rigorous bounds for this setting are unavailable in prior art."*), this paper is the first result that guarantees non-asymptotic optimality of the most standard MPC policy. We will make this comparison clearer and cite these papers.

c. This paper not only analyses the performance of MPC, but also studies the **fundamental limit** given k future exact predictions (i.e., we characterized the optimal policies given k predictions, and derived the regret lower bounds of any online policy), which is not covered in the prior arts.

Reviewer 1: Thanks for your constructive feedback and pointing out related work about learning to MPC. We hope the general response helps, and we would definitely reorganize related literature and add more systematical discussions.

Reviewer 2: Thanks for your valuable comments and suggestions. We hope the general response addresses your concern about exact predictions.

Reviewer 3: Thanks for your constructive reviews. We hope the general response addresses your concern about the exact predictions. Regarding empirical studies, we would like to add some motivating numerical examples in the main body since we will have one more page if accepted. For technical questions:

Line 168-172: I suggest the authors to give more intuition. The key idea is that $\max(f-g) \geq \max f - \max g$.

Line 219: Why $\tilde{Q}_f = P$? The terminal cost Q_f is given, but the virtual terminal cost $\tilde{Q}_f$ used in MPC is to be designed.

Line 594: Where is "Ω" used in the proof? When reducing the adversarial case to the stochastic case, we require that the support of the distribution of each w_t is contained in Ω.

The difference between $O(), \Omega()$ and $\Theta()$. We say $f(k) = O(g(k))$ if $\exists C > 0, \forall k \geq 1, |f(k)| \leq C\, g(k)$; $\Omega()$ is similar except that the last "$\leq$" is replaced by "$\geq$"; $\Theta()$ means both $O()$ and $\Omega()$. This is stronger than the standard definition where $f(k) = O(g(k))$ if $\exists C > 0, k^* > 0, \forall k \geq k^*$, $|f(k)| \leq C\, g(k)$. We will make it clearer in the paper.

Reviewer 4: Thanks for your constructive reviews and we hope the general response addresses your concern about the relationship with other MPC settings and the practicality of predicted disturbances.

The papers you mentioned have studied MPC in various settings, but they demonstrate their effectiveness mainly through empirical ways. Some papers have theoretical results for stability, but optimality guarantees are lacking. As far as we know, our paper is the first one to provide theoretical guarantees for the dynamic regret of MPC. We will definitely make this comparison clearer, and add more references and discussions.