We are grateful to hear that the reviewers find our manuscript worth of being accepted. All reviewers make good summaries of our contributions and provide helpful suggestions for us to further improve the paper. Here we provide some specific discussions to address the reviewers' questions. We will also incorporate them into our final version.

==== *To reviewer 1* ====

Prediction model. We agree with the reviewer that our current prediction model is idealized and should be generalized. Indeed, our ongoing work is considering noisy predictions and studying the role of bias and variance of the prediction. One issue we face is finding a meaningful yet simple enough noisy prediction model for insightful theoretical analysis. We really appreciate the suggestion of the doubly robust estimators and will try it in our ongoing work.

Competitive ratio. This is a good point and in fact we had investigated whether a constant competitive ratio could be obtained before our initial submission. However, we have not obtained many meaningful results without adding strong assumptions on the objectives. We will continue to work on this as future work and discuss it in the final version.

Parameterization in terms of control inputs. We agree that the more direct parameterization in terms of controls does not yield the result. [Richter, Jones, Morari, 2012] shows that when the system matrix A is unstable, the condition number of the W-step optimization on the control inputs goes to infinity as $W \to +\infty$, causing numerical issues. But for our parameterization, the condition number is independent of W and is bounded even when A is unstable.

Effect of linear transform to canonical form on the strong convexity and smoothness. This is a good point. The strong convexity and smoothness are preserved under the linear transformation to the canonical form, but the strong convexity number and the smoothness number would change with the transformation. Taking $\hat{f}_t(\hat{x}_t) = f_t(S_x^{-1}\hat{x}_t)$ as an example, the smoothness of $\hat{f}_t$ can be upper bounded by $\|S_x^{-1}\|_2^2 l_f$, and the strong convexity is lower bounded by $\mu_f/\|S_x\|^2$. Roughly speaking, not only f_t, g_t but also A, B affect the condition number ζ. Getting more elegant and insightful formulas are worth exploring and left as future work. More discussion will be added to the paper.

On the $x_{t+1} = x_t + u_t$. Yes, this is a typo and should be $x_{t+1} = Ax_t + Bu_t$.

==== *To reviewer 2*====

Infinite regret when $\zeta \to +\infty$? The regret bound will not go to infinity as $\zeta \to +\infty$ because $((\sqrt{\zeta}-1)/\sqrt{\zeta})^K \leq 1$ $((\sqrt{\zeta}-1)/(\sqrt{\zeta}+1))^K \leq 1, \forall\, \zeta \geq 1$. When ζ is large, the regret decays slowly when prediction W increases.

The strong convexity of g_t is not used. This is true and is a good observation. We will clarify this in the final version.

The gap between the upper and lower bounds. We think it would be extremely exciting, but not an easy task, to close this gap, because this gap is closely related to a long lasting gap between the upper and lower bounds of the convergence rate of the first-order optimization algorithms for strongly convex and smooth functions. In particular, our regret upper bound is based on the convergence rate of triple momentum (the best convergence rate to our knowledge) [39], and our regret lower bound, perhaps surprisingly, matches the convergence rate lower bound by Nesterov [38]. This interesting connection is worth further exploring. Research advances for either gap will help close the other gap. Here, we leave this gap as future work since our major focus, as summarized by the reviewer, is to quantify the effect of prediction in online control instead of solving the long lasting open question in first-order optimization.

Comparing with learning-based control and noisy unknown dynamics. In some sense, our current results are orthogonal to that of learning-based control, because learning-based control usually considers a time-invariant environment and aims to learn system parameters or optimal controllers by data; while our current paper considers a time-varying scenario with known dynamics but changing objectives and studies decision making with limited predictions. We agree that assuming noiseless known dynamics is a main drawback of our setup. It is our ongoing work to relax the assumptions to generalize the application of our results. Firstly, we note that our current setup can handle deterministic but time-varying disturbances in the same way of handling the time-varying objectives. However, for general randomness and unknown dynamics, it requires much more work. One way is to draw ideas from learning-based control, trying both model-based and model-free learning. To get meaningful results, we conjecture that the environment's volatility and nonasymptotic learning guarantees will be crucial. We also hope certainty equivalence will facilitate the regret analysis to some degree. More references and discussions will be added to the paper.

Why nonlinear f_t, g_t? Our algorithm naturally applies to general f_t, g_t and the proof only uses strong convexity and smoothness. But we agree LQR is a nice candidate and allows more customized results as shown in section 5.

==== *To reviewer 3* ====

Comparison to [Agarwal et al. 2019]. Thank you very much for mentioning this. We also found this paper, but shortly after our submission. It is very relevant and we will cite it and compare with it in the final version.

Algorithm presentation. We thank the suggestion and will present the vanilla gradient method first in the final paper.