We thank the reviewers for their time and effort, and for the helpful feedback!

For those reviewers that we have answers: we hope our comments alleviate your concerns and, in that case, potentially warrant a suitable revision of your reviews, thank you!

Specific remarks:

Reviewer 1:

1. The use of bandit techniques in control comes with technical difficulties, and is not black-box (see more detail in comments for reviewer 2).
2. We in fact compete with disturbance-feedback controllers, which is a very large class that also contains linear-dynamical controllers. This is of course larger than linear controllers, and is SOTA for H2 control. We talk about linear controllers mostly for simplicity of presentation.
3. Actually controllability is a more general condition than stabilizability, and that's what we use. "Strong controllability" just puts a numeric constant for controllability, since the usual control literature is all asymptotic and we have finite-time bounds, this is necessary.
4. You are absolutely right, we think that $T^{\frac{3}{4}}$ is *not* tight. But going all the way to $T^{\frac{1}{2}}$ seems very hard, we leave it to future work.

Reviewer 2:

1. Please note that generalizing FKM to functions with memory is more subtle than appears at first glance and the introduction of delay is not the complete story. The second-order subtlety comes from the fact that Stokes Theorem (which lies at the core of FKM) cannot be immediately applied in the setting of functions with memory. More specifically, the problem arises due to the fact that the boundary of the set defined by a direct sum of balls is *not* a direct sum of spheres. As such, our proof needs to get around this issue by first going through a linear approximation. This issue cannot be solved by ensuring that the introduced perturbations lie on a hypersphere as a group either because this introduces further dependencies that interfere with the application of the FKM approach.
2. Your second comment can also be answered: it is highly non-trivial to use exploration methods such as in Bubeck et al, since for control the parametrization space is different from the cost space. This is due to the use of OCO with memory and the convex relaxation that we use, and was introduced by Agarwal et al.
3. It is true, however, that for strongly convex and smooth losses the techniques of Hazan and Levy can probably be used to give tighter bounds. We didn't explore this since we focused on convex losses, where this is not applicable.

Reviewers 5 and 6: We completely agree with your reviews, thank you!