Dear reviewers and chairs,

We thank all reviewers for their careful reading of the manuscript.

Reviewer # 1:

We will try to add more examples on when it helps to be comparator-adaptive. Note that in the paper we already have a nice but subtle application of the comparator-adaptive property in line 256, where we use that the regret against $\mathbf{0}$ is $O(1)$. Without using this property the analysis of the unconstrained smooth case actually becomes troublesome as the surrogate losses for the unconstrained smooth case are not Lipschitz on $[0, \infty)$.

Reviewer # 2:

First, thank you for the references. Especially the the second reference appears to be related as it seems to extend the ideas of MetaGrad to the Bandit Convex Optimization setting.

As for our technical contribution, understanding how to utilize the ideas of Cutkosky and Orabona [10] in the Bandit Convex Optimization setting is one of the main contributions of this paper. There are subtle but important differences between the full-information setting and the Bandit Convex optimization setting that require us to make (not so) subtle adaptations to be able to apply the ideas of Cutkosky and Orabona [10]. For example, as we point out in section 3.2, the projections of Cutkosky and Orabona [10] do not directly work in the bandit setting. Another example is the gradient estimator. Even though it appears as a natural estimator some care is required to be able to use it as it is tightly connected to the surrogate losses we feed to Interface 3. However, in the final version we will reduce the emphasis on the gradient estimator as a stand-alone novelty. Also note that there is an entire line of work that expands the ideas of Cutkosky and Orabona [10] to other settings, see for example [8], [20], and [26].

Reviewer # 3:

Even though the extensions indeed appear to be natural, as reviewer # 1 does we would like to argue that this is a strength rather than a weakness. One of the main contributions is that our techniques represent a way to utilize the ideas of Cutkosky and Orabona [10] in the Bandit Convex Optimization setting and we hope that our insights will allow future researchers to design further extensions. As reviewer # 1 points out, the way we extend the ideas of Cutkosky and Orabona [10] to the Bandit Convex Optimization setting may be of interest to other applications. With simple and clean proofs the ideas will be easier to use by other researchers, otherwise our hope could be in vain as the ideas might be too obfuscated to be of use.

Regardless, there are other subtle modifications that one may miss. For example, as we mentioned in our comments to reviewer # 1, in the unconstrained smooth case the surrogate losses are not Lipschitz on $[0, \infty)$, which means that a direct application of Interface 3 would not work. Instead, we restrict the algorithm to play on a smaller domain where the surrogate losses are Lipschitz. This leaves us with a problem if the comparator is outside of this smaller domain, but, inspired by [8], we manage to solve this by utilizing the fact that our comparator-adaptive algorithm has $O(1)$ regret against $\mathbf{0}$.

We would also like to address our assumption of zero loss at zero, which is helpful for overcoming one of the key challenges mentioned in Line 199. While it's obviously best to have as few assumptions as possible, this particular assumption actually holds in many practical scenarios, as we discuss in the paper. Moreover, finding improved performance on special problem classes can open up interesting new research directions - for example, it may even be that an assumption similar to ours is actually *necessary* in order to obtain the comparator-adaptive bounds!

As for the dependencies on c, if the domain is a unit ball (in any norm) then the dependency on c disappears as it is 1 in that case. For more general domains it is not clear if the dependency on c is tight. To our knowledge, the only place in literature where a similar constant appears in the regret bounds is in Flaxman et. al. (2005) [13], which was removed in subsequent work. Because of this we suspect that, as in the case where the domain is a unit ball, the optimal regret bound should not depend on c. As for the dependency on d, it is tight in the linear case (see Dani et. al. (2008) [11]).

Finally, to clarify, the assumption $\ell_t(0) = 0$ is used in Theorems 3, 4, and 5. Thank you for pointing out the typo in Table 1.