We thank all reviewers for their valuable comments, which will be taken into account in the revision of the paper.

R3, R5: *Adding empirical results.* We agree that an extensive numerical experiment would be of interest. We did not develop it in this paper as it is already dense in theoretical results. The current format allows maybe for adding a small numerical example, which cannot be seriously considered as a sufficient evidence. We also mention that the whole line of previous papers connected to ours (Agarwal et al. [1], Bach & Perchet [3], Duchi et al [13], Flaxman et al. [16] Jamieson et al [20], Shamir [32]) presents only theoretical results.
R2, R3: *Give guarantees in the case when the function parameters are unknown.* This is an interesting question related to adaptive techniques. To the best of our knowledge, it was not answered in the literature even for simpler settings where $\beta = 1$ or $\beta = 2$ and one only needs to adapt to the Lipschitz constant L and strong convexity parameter α. When the question is to find algorithms optimizing the bounds, all the previous related work assumes explicitly known L (known α if the functions are strongly convex), cf. references above among others.
R2, R4: *Use of the kernel being redundant in the case $\beta = 2$.* Inspection of the proof of Lemma 2.3 shows that for $\beta = 2$ the bias is of the same order even if $K = 1$ and the randomization in r is suppressed (this is not true for bigger β). For the variance term (Lemma 2.4), the dependency on K does not influence the rate for all $\beta > 1$.
R2: *Strong convexity assumption.* This assumption is quite common in the bandit literature, cf. for example, Agarwal et al. [1], Bach & Perchet [3], Hu et al. [17], Shamir [32]. Discussion of motivation for strong convexity assumption can be found there. We did not reproduce this discussion due to space limitations.
R2. *Noise assumption in Theorem 6.1 (lower bound):* Using the second order expansion of the logarithm w.r.t. v, it is not hard to check that this assumption is satisfied when F has a smooth enough density with finite Fisher information.
R2: *Choice of the learning rate for the unconstrained case (Eq. (6)).* Having two regimes for tuning parameters is a technical point in order to cover small enough strong convexity parameters α. If we drop the first regime, i.e. set $T_0 = 0$ (start from $t = 1$) then we cannot guarantee that in the display below Eq. (34) the term with $\bar{L}^2/\alpha t$ on the r.h.s. is small.
R2: *Choice of strong convexity parameter in Theorem 3.2.* It is a natural assumption since the case of very small strong convexity parameter boils down to the setting without strong convexity, where higher order smoothness does not help and the optimal rate is different. This is discussed in Section 7.
R2: *Additional references on bandit literature:* Thanks for pointing out the papers on contextual bandit, which we will cite in the revision. Dealing with the same Holder assumption on the arm' reward functions [Hu et al. 2019] their setting is significantly different from ours as at each step the learner has to choose between finite number of arms, while in our case it is a continuum. In the optimal rate for contextual bandits (see [Hu et al. 2019]) the dimension appears in the exponent of T, whereas in our setting the dimension is only a factor.
R2: *line 95.* One verifies that for this choice of K, it holds that $\kappa_\beta \leq 3\beta^3$. We'll include this in the revision.
R2: *Line 430.* Indeed, in this argument the simplification comes by roughly dropping the quantities bounded by 1.
R2: *Summarized some of results in [3] in Appendix Section C.* In Appendix C we have summarized some results of [3] by indicating the suboptimal rates that they obtain. Concerning the incorrect results, we have indicated where there is a problem in the argument.
R2: *More discussion.* We will add the definition of cumulative regret in the main paper and give more intuition behind the proofs whenever possible in the paper or in the appendix. Concerning the broader impact, as our work is theoretical, we do not expect a broader impact, nor possibly negative effects on society. If the reviewers detected any relevant aspect we missed, we are happy to fill this section.
R3: *I am unsure of the significance of this contribution.* We show what is the optimal choice of parameters for the algorithm by Bach & Perchet and, notably derive near matching lower bounds. Moreover, we provide the first polynomial time method of estimation of the minimum value $\min_x f(x)$ with near oracle behavior.
R3: *Related work section at the beginning of the paper.* We prefer to keep it at the end of the paper since it is easier to compare to the previous literature once the results have been presented. However notice that in the introduction we briefly comment on key related papers and improvements developed by our work.
R4: *Adding a table.* We have also thought about this option but did not choose it since there are not too many cases to compare with.
R4: *Why generalization to "adversarial noise" is plausible?* The reason is that multiplication of y_t, y_t' by ζ_t, which is zero-mean, makes the stochastic error zero-mean independently of the nature of the noise.
R4: *Lower bound - what is new?* The bound extends over the initial proof technique of Polyak and Tsybakov [28] by accumulating multiple (rather than two) probe functions to account for the dependency on the dimension d and α and applying Assouad's Lemma to obtain the final result.
R4: *Why does def 2.1 need to constrain ζ and r?* There is no def 2.1, probably assumption 2.1? Indeed, informally it is clear that ζ and r should be chosen independently. We explicitly state this requirement in the assumption just to make the theorems formally correct.
R5: *line 180 and 181:* by the fact that it cannot be improved we mean that even if x^* is known and we are sampling T times $f(x^*) +$ "noise", all at the same point x^*, the error is of the order $1/\sqrt{T}$. We will modify the discussion in the revision, also referring to Thm. 4.1.