We thank all the reviewers for insightful comments and suggestions.

Reviewer 1: Thanks for the spot-on comments and for being our champion! We address the two remarks below.

Remark 1: In the batch setting, optimality in TV class indeed implies optimality in enclosed Sobolev and Holder classes as the reviewer pointed out. However it is not true for forecasting due to the dependence of $[C'_n]^2$ in the optimal regret rate as in theorem 8. While bounding the regret of ARROWS, we get a ground truth dependent L2 norm term $\|D\theta\|_2^2$ in equation (20). This enables the adaptive minimaxity for Sobolev and Holder classes. However, a minimax strategy whose regret bound contains the term $\|D\theta\|_1^2$ in the place of $\|D\theta\|_2^2$ in (20) will be optimal for TV class but fails to get the correct dependence on $[C'_n]^2$ for the Sobolev class.

Remark 2: Achieving minimax forecasting in the TV-constrained comparator setting with a polynomial time algorithm is an intriguing open question. Our results do not directly apply to that stronger setting. Although some of our techniques might be reusable but we believe nontrivial new algorithmic ideas/proof techniques are probably needed. Our work is better viewed under the lens of non-stationary sequential stochastic optimization as in Besbes et al [1] with squared error loss and noisy gradient feedback.

Reviewer 2: Thanks for the detailed and insightful review. Please see Remark 2 above on the comparison to the TV-constrained comparator setting and detailed responses to other questions.

Re More general loss functions: Generalization to other convex costs is regarded as a future work. Thanks for the suggestion of self-concordant losses. It is a good direction to explore.

Re Relation to Gaillard and Gerchinovitz[2015]: The regret bound of $O(n^{1/3})$ in [2] attained by an $O(n^{7/3})$ runtime policy holds for Holder class which features more regular functions than TV class. Their regret bound in theorem 11 fails to capture the optimal dependence on the Lipschitz constant and hence cannot be used to construct the correct lowerbound with precise dependence on all of the problem parameters in our setting.

Re Boundedness of theta and C_n^2 term in the lowerbound: If we assume all theta to be bounded by B then we would be able to get a better $\Omega(BC_n)$ bound. For instance we can consider packing functions that alternates C_n/B times between 0 and B. This also points to the fact that forecasting is harder than smoothing. However, this boundedness constraint implies that we will be focusing only on a smaller subset of all sequences whose TV is bounded by C_n. Of course this B in worst case is at most $U + C_n$ where U is the bound on first data point.

Re Adaptivity to C_n: Adaptivity to unknown variational functionals are usually nontrivial. Contrary to the reviewer's comment, the uniform restarting proposed in [1] is in fact not adaptive. It requires knowing C_n to set the optimal restarting intervals. To the best of our knowledge, Zhang et al. 2018 [3] was the first paper that made it adaptive — albeit suboptimally in our setting — as $\sqrt{C_n}$ to the total variation. Even there, they achieve adaptivity with a very nice new idea of connecting to strongly adaptive regret minimizing algorithms.

That said, the reviewer's question challenged us to look into the problem further. We are now convinced that with a simple tweak in the restart rule, it is possible to transform ARROWS to **an anytime algorithm that optimally adapts to** C_n — the TV of ground truth. Let the expression in LHS of the restart rule be $\hat{C}$. The idea is to replace n and C_n in the RHS of restart rule by k and $\hat{C}$ respectively. So we restart when $\hat{C} > \sigma k^{-1/2}$. All the results can be proved to be true with this almost fully adaptive restart scheme. We do not have space for a proof in this short rebuttal, instead we present below but the regret plot with the new restart rule as an empirical validation. We will include this update in main paper if accepted. σ if unknown, can be robustly estimated (thanks to sparsity of the wavelet coeffs of Bounded Variation functions) using first few observations as mentioned in line 69.

Reviewer 3: Thanks for appreciating our contributions!

References:

[1] Besbes et al. *Non-stationary stochastic optimization*, In Operations Research 2015.

[2] Gaillard et al. *A chaining algorithm for online nonparametric regression*. In COLT 2015.

[3] Zhang et al. *Dynamic Regret of Strongly Adaptive Methods*, In ICML 2018.

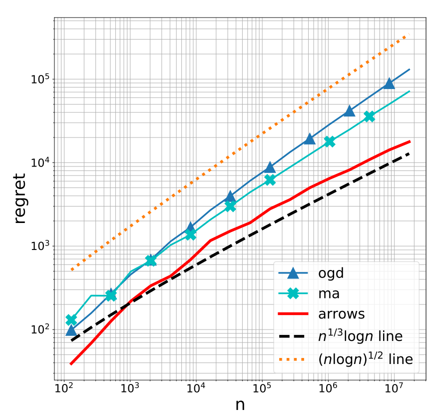

Figure 1: *Regret plot for function in Fig.2 of main paper with the new restart scheme that makes ARROWS optimally adaptive to both n and C_n.*