We would like to thank all the reviewers for their detailed feedback, for appreciating our writing style (R1,R2, R3), and for recognizing that we provide an *interesting answer to a natural question* (R2), also in a *practical manner* (R1).

REVIEWER 1. Q: Regret bound in terms of the variability of the sequence can be obtained for FTRL style algorithms.
A: This is indeed true, thank you for pointing it out! However, to obtain it you would have to run FTRL with full losses, rather than the common **linearized** version. Unfortunately, this would entail solving a constrained convex optimization problem whose size (in terms of number of functions) grows each step, that would be slow even when the implicit updates have closed form expressions, e.g., linear classification with hinge loss.

Q: Unless there is some special structure, the implicit update requires an "inner loop" of optimization to compute.
A: It is true that implicit updates in general require a heavier computational burden. However, as illustrated in Appendix C, there are some very important and practical cases where the updates can be computed in closed form. Also, the "inner loop" does not seem a serious problem in practice, people do use these kind of algorithms. See for example the nice posts on Alex Shtof's blog on proximal (aka implicit) updates.

REVIEWER 2. Q: Does the lower bound extends to an arbitrary algorithm?
A: The lower bound does hold for any deterministic algorithm on constrained domains. For randomized ones, the proof should be changed and we suspect the lower bound to hold.

Q: Please confirm there is no absolute value in (1). Is this correct?
A: That's correct, no absolute value in Eq. 1, contrarily to the usual definition of temporal variability in dynamic regret papers. So, our definition is stronger, since it could lead to a negative term. This is not so surprising: as R1 points out, FTRL with full losses would depend on a similar quantity at the expense of a growing computational complexity.

Q: What does the result of the paper imply for the changing dependency framework in prediction with expert advice?
A: We also think it would be interesting to extend our results to the dynamic case. Even if we are not dealing specifically with the setting of online regression, we believe that it should be possible to get bounds which retain the minimum between two quantities, one involving V_T and the other in terms of the variation of the comparator sequence (aka path length), as done in Moroshko et al. [2015] and Kalnishkan [2016].

REVIEWER 3. **Q**: This paper considers a new setting that lies between the static setting and the dynamic setting.
A: **This is not a new setting**, we are dealing with the **static** setting, as specified in the abstract and in the introduction (see lines 5; 19–25). The reviewer seems to imply that the dynamic setting is the only one where losses vary slowly, but this is incorrect: it only differs from the static setting for the choice of a set of competitors rather than a single one. The confusion of the reviewer seems to stem from the fact that no prior analysis of implicit updates had a term that depends on the variability of the losses. Shedding light on this novel aspect of implicit updates is our main contribution.

Q: This paper does not discuss its relation with prior works in the dynamic online setting.
A: This seems factually incorrect: We do have a discussion with prior work for dynamic environments: please see lines 60–64, 71–74, and related references.

Q: The authors claim that this work is more suitable to "small" V_T, but without further explanation or justification.
A: This seems incorrect: we **never** claim that our algorithm is suitable for small V_T. Instead, our bound is a minimum between two quantities. In the worst case our algorithm recovers the $\sqrt{T}$ bound, but in other situations where V_T is small it adapts to it (contrarily to standard algorithms like linearized MD or FTRL) and can have a better bound. In other words, we have a classic "best of both worlds" bound.

Q: This work has a strong overlap FIOL [31], thus the contribution and novelty is very limited.
A: Our paper has not only overlaps with FIOL, but also with [18] and [20]. All of them provide an intuition that it's possible to get a gain from the analysis, *but they fail to quantify it in the final bound*. We also **clearly** point out this overlap, see lines 125–127. In particular, from the analysis in FIOL we can see a potential gain in the bound in their Eqs. (6–7). On the other hand, it is not clear how their negative term in the final bound could lead to something which is less than $\sqrt{T}$. Also, any potential gain is entirely destroyed by their learning rate $\eta_t \propto 1/\sqrt{t}$ which leads to a $\sqrt{T}$ bound.

Q: The main novelty is to use V_T to bound the gain, which is quite straightforward to deduce.
A: We respectfully disagree: almost anything in Science is straightforward after somebody points it out. Yet, none of the previous work pointed out the connection between implicit updates and V_T. Unless we missed other related work, this is the first work where implicit updates could have a quantifiable advantage (i.e., possibly O(1) regret) over OMD ones.

Q: I then wonder what is "small" V_T (this even contradict the experiment setup where [...])
A: We do show in our synthetic experiment a situation when V_T is small and our algorithm is much better compared to the other baselines. On standard real-world datasets, we say that there is no reason to believe that V_T is small, but we want to show that our algorithm is still competitive. On a side note, we are not aware of any paper on dynamic regret with experiments on real world datasets with small V_T.