We thank the reviewers for their detailed and constructive comments. Please find our answers below.

Algorithm	**Lazy TS**	**RAGE**	**Oracle**
dimension	Mean (Std)	Mean (Std)	Mean (Std)
$(d = 4)$	**3346.29** (125.3)	8033.51 (464.0)	3968.72 (163.9)
$(d = 7)$	**4405.74** (143.5)	9675.00 (537.8)	4107.09 (160.6)
$(d = 10)$	**5602.55** (180.9)	9780.54 (360.0)	4321.84 (167.8)

Algorithm	**Lazy TS**	**RAGE**	**$\mathcal{XY}$-Adaptive**	**Oracle**
confidence level	Mean (Std)	Mean (Std)	Mean (Std)	Mean (Std)
$(\delta = 0.5)$	**3080.56** (119.1)	5840.21 (373.6)	6192.34 (373.8)	3016.91 (133.1)
$(\delta = 0.05)$	**3699.84** (130.1)	7751.79 (434.2)	8167.51 (368.3)	3610.33 (146.1)
$(\delta = 0.005)$	**4297.23** (131.8)	9810.19 (543.8)	9278.82 (315.8)	4219.38 (165.3)

Reviewer 1. 1) Experiments. An important part of our work is the design of the first optimal algorithm whose implementation and performance guarantees are completely independent of the number K of arms. That is why our experiments focused on scenarios with large set of arms. We will include more experiments (essentially all scenarios considered in [16]) – e.g. the table above corresponds to the benchmark used in Soare et al. [12] and all other papers on the topic: here the angle $\omega = 0.1$, the left part of the table is for $\delta = 0.01$, and the right part for $d = 6$. In all experiments we have done, the results suggest that our algorithm is very competitive.
2) Explore-and-Verify (EV) framework. Thanks for suggesting that combining RAGE and the EV framework of [13] may lead to an efficient algorithm. As you mention, an idea of the same flavour is used by Chen et al. [0] for combinatorial pure exploration. But to get guarantees in expectation, one needs parallel simulations (Lemma 4.8 in [0]) or repetitive calls of RAGE (see Appendix 2 of [13]). Thus, as already noticed by Fiez et al., the authors of RAGE, the EV framework may be impractical (see [16] Page 7). In fact, neither the authors of [0] nor those of [13] implemented their framework in practice. In our paper, we devise a very simple and asymptotically optimal algorithm using a track-and-stop framework in the spirit of [18]. To this aim, we needed to tackle the following challenges: (i) the optimal allocation is not unique, which poses tracking issues; (ii) computing an optimal allocation in each round might be computationally demanding (we solve this issue with the laziness of our tracking rule); (iii) most of the components of the algorithm, including the stopping rule, must be independent of the number of arms K.
3) Continuous sets of arms. Our main contribution for such sets of arms is to derive the first (problem-specific) sample complexity lower bound at least for the sphere. This has not been done earlier, and proved to be very challenging. It might be the case that discretizing the set of arms using an $\epsilon/2$-net of the sphere and applying our TS algorithm would yield a sample complexity with the right scaling in d, ϵ, δ. The algorithm we propose is even simpler, and in particular does not need to work with a set of arms with cardinality growing exponentially with the dimension.

Reviewer 3. 1) Novelty of the problem and of our solution. The problem of best arm identification in linear bandits is not new, but it is fundamental. To the best of our knowledge, we propose the first asymptotically optimal algorithm whose implementation and performance guarantees are completely independent of the number K of arms. To this aim and to use the track-and-stop framework [18], we needed to tackle the following challenges: (i) the optimal allocation is not unique, which poses tracking issues; (ii) computing an optimal allocation in each round might be computationally demanding (we solve this issue with the laziness of our tracking rule); (iii) most of the components of the algorithm, including the stopping rule, must be independent of the number of arms K. We agree that a very interesting research direction is to derive non-asymptotic performance guarantees for our algorithm.

Reviewer 4. Thanks for pointing us to the recent paper of Degenne et al. (it was not available when we submitted our paper). We will cite and discuss it. We also appreciate your suggestions for improving the clarity of the paper. We will elaborate further on the broader impact section as suggested.
1) Experiments. Please refer to answer 1) to Reviewer 1.
2) The constant $c_{\mathcal{A}_0}$. This constant controls the rate at which forced exploration is performed. Having a too large constant could indeed lead to worse performance. Theoretically, however, it does not affect the asymptotic performance of the algorithm. In practice, we note that forced exploration rarely occurs for the considered level of laziness in our experiments. We conjecture that it may not even be needed. We will provide further discussion on this in the paper.
3) The goal of $\mathcal{A}_0$. We confirm that the goal of $\mathcal{A}_0$ is to ensure that the growth rate of $\lambda_{\min}(\sum_{s=1}^{t} a_s a_s^\top)$ is not too small. We also note that the chosen growth rate $f(t) = O(\sqrt{t})$ is not too large because, ideally, if we knew an optimal allocation, sampling according to it would yield a growth rate of order exactly t.
4) Estimation of μ along problem-dependent directions. These directions also appear implicitly in the sampling rule. More precisely, the tracked weights w_t in the sampling rule are solutions of the optimization problem $\max_{w \in \Lambda} \psi(\hat{\mu}_t, w)$.
5) The stopping rule interpretation. Intuitively, the stopping rule can be viewed as an empirical version of the problem-specific sample complexity lower bound. Roughly speaking, the stopping rule would correspond to the lower bound where μ is replaced by $\hat{\mu}_t$, and w by $(N_a(t)/t)_{a \in \mathcal{A}}$.
6) Continuous set of arms. Regarding the continuous set of arms, our main result is the lower bound. Deriving such an explicit bound was challenging: in the change-of-measure argument, we needed to propose an appropriate reduction of the set of *confusing* parameters, see Appendix G, Steps 2 and 3 of the proof. We are only able to derive it for the sphere for now. For a generic continuous set of arms, the problem is even harder because in the change-of-measure argument, the aforementioned reduction of the set of confusing parameters becomes even more challenging.