Many thanks to the reviewers for the insightful and constructive feedback, which has significantly improved the manuscript. We were pleased to see quite a few remarks collectively from the reviewers highlighting the novelty and strength of our paper, including our *state-of-the-art theoretical result* and *strongly supportive simulation*. Due to space constraints, instead of responding point-by-point, we address points in common with multiple reviews. All minor comments have been addressed and incorporated into a revised version of the paper.

Assumption on i.i.d. Gaussian Design. In general, AMP theory sheds light mostly on i.i.d. Gaussian data, and thus to quantify the same diagram for *general* covariance matrix, one may need to develop stronger AMP tools. As sharply pointed out by reviewer #3, though Gaussianity with identity covariance is restrictive, it is already unclear what will happen to the Donoho-Tanner phase transition without this assumption, so it is desirable yet extremely difficult to achieve this generality for our more refined result. Though it is hard to theoretically quantify the exact complete diagram for general design (especially the lower boundary), empirically we find our diagram is *still* correct up to small differences on the lower boundary for a wide range of designs. For example, in Figure 1, we illustrate the Lasso diagram for various designs: namely, Gaussian design with AR(0.05) covariance matrix (Top), Bernoulli design with each entry being i.i.d. Bern(0.5) (Middle), and Cauchy design with each entry being Cauchy$(0, 1/n)$ (Bottom). In the Gaussian and Bernoulli case, our claimed region (enclosed by the black lines) is still almost exact. When the design comes from Cauchy distribution, where its mean or variance is not even well-defined, the simulation result has a higher lower boundary. This is easy to understand: the difficulty of the Cauchy design complicates the model selection problem, and the Lasso generally cannot achieve the best case as in the i.i.d. Gaussian case.

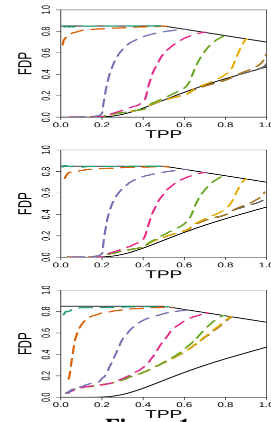

Figure 1

Statistical Implication from our Complete Diagram. We want to re-emphasize the motivation and the statistical implication for studying the complete tradeoff diagram. First of all, our finding is a novel result that quantifies the exact complete achievable region of Lasso, and it is of great theoretical interest. Also, the usage of homotopy methods is rare in the statistics and the machine learning community, and our framework can be used to establish similar results for other methods like SLOPE, SCAD, group Lasso, etc. Secondly, the complete Lasso diagram allows us insight to the Lasso's performance. As illustrated in Figure 2, we can have a very narrow estimate of the false discoveries when the Lasso has large power. According to the lower bounds of FDP in the three cases ($21\%, 36\%$, and 16%) are the best possible value achievable when the TPP is close to its maximum. However, our complete Lasso tradeoff diagram also guarantees that it is impossible to have a much worse FDP than the best possible ones when the TPP is large.

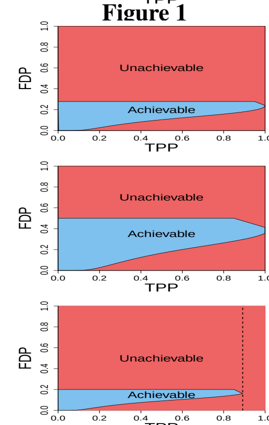

Figure 2

Level Plot of the Lasso Tradeoff Diagram. To better illustrate our result, we present Figure 3 suggested by reviewer. In each diagram, we plot $\delta = n/p\,(x-\text{axis})$ versus $\epsilon = k/p\,(y-\text{axis})$, and fix FDP to be 0.2 (Top), 0.4 (Middle), and 0.6 (Bottom). The color of each point represents the largest TPP (since trivially, minimum TPP is 0) achievable (red for 0 and white for 1). We see that for large FDP, the TPP is always decrease with the sparsity ratio ϵ, no matter beyond or below the DT phase transition. However, for small FDP, the maximum power first decreases with the increase of sparsity, and then increase with sparsity when above the DT phase transition. Our more refined result exactly characterizes this complication beyond DT transition. These plots, though being mathematically equivalent, complement to our tradeoff diagrams from a different perspective.

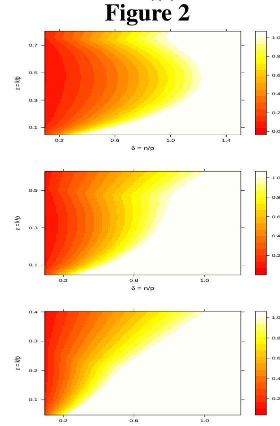
Figure 3

Other Minor Details and Comments. We have addressed all the corrections suggested by the reviewers and updated a revised version of the paper to define more clearly all notations and terminologies. We remark on some confusion as follows: 1. The q^* is well defined when we fixed ϵ and δ, however for notation simplicity we omit its dependence on ϵ and δ when it is clear from the context. 2. To be clear, as stated in the first assumption on page 3, we consider n_l, p_l, k_l for some $l \geq 0$, and the asymptotic regime is when $n_l/p_l \to \delta$ and $k_l/p_l \to \epsilon$. Specifically, the k_l here is not a random variable for any l. The prior Π (where k is random) is only used in (3.3) (3.4) to define $(\text{tpp}^\infty, \text{fdp}^\infty)$. 3. The finite second moment assumption of the prior is an assumption needed by AMP. We believe this is an artifact—in practice, a large second moment can be desirable, since it often results in large "effect-size heterogeneity" (a new notion proposed recently), where the Lasso's performance would be very close to the lower boundary q^*, which is also enclosed in our Lasso diagram.