Peer Review File

Manuscript Title: Cortical ensembles orchestrate social competition via hypothalamic outputs

Reviewer Comments & Author Rebuttals

Reviewer Reports on the Initial Version:

Referee #1 (Remarks to the Author):

I am a behavioral neurophysiologist, and I reviewed this paper in conjunction with a computational neuroscientist/neurotheorist. Comments reflect our pooled suggestions to primarily aid the authors in improving the rigor and reproducibility of their work and enhancing communication of their findings for a broad scientific audience, and to secondarily aid in editorial consideration of this work for publication. The richness of our suggestions should not be misconstrued as a lack of support for this work, as noted in the summary paragraph.

This study by Padilla-Coreano et al seeks to investigate the role of the PFC, and more specifically, a PFC->LH circuit in encoding social rank and dominance behavior. To achieve this, the authors begin by determining the social rank using a classic test assay in pairs of co-housed mice. Since the primary objective of this study is to uncover the neural representations of social rank, the authors suggest that the tube test has several clear limitations. First, the behavioral data that can be extracted from the tube test is less rich than other behavioral assays broadly deployed in mice. Secondly, I think the authors seem to suggest that the task does not enable the measurement of the electrical activity from the same cells across many trials (which is necessary to facilitate statistical analysis). Here I think its worth largely emphasizing the first point. As previously demonstrated in Zhou et al, Science 2017 (cited by the authors), mice can be repeatedly subjected to the tube test. Since this is a widely accepted and stable (across trials and sessions) measure of social dominance in male mice, and PFC cellular responses to competition have already been described in Zhou et al, Science 2017, a major contribution here is showing that these responses are presented in a more ethologically relevant context.

Thus, the authors develop a novel competition assay. Using a set of mice in which social hierarchy is previously established using the tube test, the authors successfully show that social dominance predicts behavioral outcomes in their competition assay. Nevertheless, here I would raise my next suggestion. With the tube test assay, five animals can be ranked 1-5 (to social rank). Mouse #2 has a high social rank, though it is submissive to Mouse 1. With competition assay and the analysis strategy taken by the authors, it is imperative to note that they are not probing social rank (i.e., Mouse #2 within the colony). Rather, they are probing relative social rank (Mouse 1 vs. Mouse 2)… the neural processes that encode being dominant to your partner. At this point, its unclear if mouse 2 would look any different than mouse 4 if all five animals were performing the competition assay at the same time. At some places in the paper the authors appropriately use the term relative social rank, and in other places, they simply use the term social rank.

After the authors establish their novel behavioral assay as a surrogate for relative social rank that also facilitates multiple trials, they strive to extract a rich set of behavioral measures using a novel ML coding algorithm. Importantly, their approach can monitor the behavior of non-labeled individualmice, and it can extract joint behavioral features from pairs of mice. Here, however, I would suggest toning down the novelty/utility claims of this approach a tiny bit. Its nice that their approach can track non labeled mice, but its not that difficult to label mice. Furthermore, while tracking non- labeled mice within a large colony might be exceptionally useful, they only show its effectiveness in tracking two mice. Thus, I would suggest that they revise the description of AlphaTracker to state that it tracks two mice, not multiple (in the abstract, text, and figure legend). Its a more precise claimthat doesn't take anything away from the utility of their scientific findings

Once the authors identify behavioral clusters using AlphaTracker, they set out to determine whether PFC firing represents social rank and the behaviors observed during their task used to quantify relative social rank. Specifically, the authors implant animals with 32 wire electrodes, and perform wireless while mice engage in their task. Here, some missing details/confusing text made it more difficult for me to evaluate the rigor of their findings. First, I had a hard time determining the number of implanted mice subjected to their assay. I searched the text, figure legends, and methods, and I am sorry if I missed it. Second (and also related), it was hard for me the interpret the following statement (line 99) without that data, "when recording during the reward competition task, we did not detect statistically significant differences in the number of rewards earned by dominant and subordinate mice, allowing us to make comparisons about dominance behavior and competitive success without being confounded by the volume of reward consumption." I was totally lost by this sentence. In Fig 1c, the authors validate their assay as a surrogate for social rank by showing that commutative rewards were related to outcomes in the tube task. Here, the authors seem to be suggesting that the outcomes of the tube task do not relate to successful competition for limited resources. There are several reasonable explanations for this. For example, the electrophysiological experiments may not be powered to detect this behavioral difference (but I can't assess that without N). Did the partner mice wear the dummy headstage during the ephys competition (seems so in the video)? For these mice, there appears to be less cumulative rewards in the unimplanted mice.

There are different behavioral measures shown in figures for the initial version of the competition task in un-implanted mice (cumulative rewards, ratio of port occupation, pushing success and time displaced), than for the implanted mice (cumulative rewards, pushing success, latency pushing success during tone). Here the only common measure that shows the same result is the pushing success. For the latencies, the trials appear to be treated as independent for this latter analysis. The authors should treat these behavioral observations as repeated observations of the same phenomenon for Fig S5c (specifically, within trial for a given subject). Also, the legend for Fig S5c says 'slower latencies.' It would be better to describe the data as 'longer latencies.'

Since the nature/validity of the behavioral task is important for exploring the link between relative social-dominance and neural activity, my further commentary assumes that the suggestions I've provided up until this point can be easily addressed with additional experimental details and minor revisions to the analysis. Thus, I am providing additional suggestion based on the assumption that the authors will clearly establish behavioral outcomes of their competition task as measures of social dominance for the same animals used for their neural analysis in a revised version of the manuscript.

The language around the authors population analysis is a little unclear. As I understand it, the authors used AlphaTracker to discover nine behavioral conditions defined by two mice in an arena. They then trained a SWM using PFC multiunit activity to predict these behavior conditions. Next, they hypothesized that PFC may better encode hidden states. Here, rather than PFC directly encoding the 9 specific behavioral conditions (or better stated, rather than discriminating the behavioral states from each other), the authors tested a model in which PFC encoded a number of states (six in this case) in which distinct profiles of their behavioral conditions were likely to emerge. They found that PFC decoding was more accurate using this latter model in comparison to their models where PFC encoded the behavioral conditions.

The results on the modeling are unusual, where gains between the autoregressive model compared to the SVM and the GLM are drastic. It is rare that you see such dramatic gains when moving to an HMM based model, and the manuscript would be much improved by explaining this outcome. In fact, the authors show that there are drastic gains by using 2 clusters in the HMM based model (noting that the 1 cluster case is equivalent to the GLM model in your formulation). In this case, they are now distinguishing between 9 classes based on a mixture of 2 GLMs, where the GLM is chosen through the historical information. Compared to the classes, this is really a minor gain in modeling complexity, and I'm unclear from a mathematical perspective as to how such gains can be achieved. Additional justification and exploration would be warranted to ensure robustness of these findings.

In any case, if my understanding of their objective is correct, the authors should revise their text for clarity. The term 'behavioral state' is used under too many different contexts. For example, the authors state "the proportion of time spent in each hidden state did not different by competitive success or by social rank, and the model performed equally well across ranks, suggesting that mPFC encoding of social competition behavior states is common across ranks." It would be more appropriate to state that the encoding of social competition hidden states is common across ranks. Also, again, the authors should revise most statements regarding rank to read 'relative rank'. Assuming that my understanding is correct and that the authors clarify the text accordingly, I believe the data argues that PFC activity encodes behavioral states in which distinct profiles of behavioral are likely to emerge. There is plenty of evidence to support this interpretation including studies from the senior author and many of the works literature cited in this manuscript. Nevertheless, the relevance for this particular study remains unclear since these behavioral states (hidden states) do not show a relationship to rank.

Next, the authors set out to determine if PFC unit activity predicted social rank and task relevant behaviors in their competition assay. My concern with this approach, in particularly analyzing winning and losing, is that the authors did not find that winning vs. losing was related to social rank in the cohort of mice used for their electrophysiological study. Specifically, the authors found that the dominant and submissive mice consumed the same amount of reward. Thus, I'm not entirely convinced that decoding winning and losing is optimal. Imagine a situation where basal difference in PFC firing activity are related to relative social rank (a fair hypothesis). One would expect that a model using PFC activity could differentiate relative social rank irrespective of the behavioral condition as long as two mice were in the same arena. This appears the be the case, since the AUCs differentiate relative rank even prior to cue presentation. How then to interpret the trial by trial decoding of winning a losing? Would the PFC show this response in any area, even one in which resources were freely available? What if there were two poke holes, and the animals didn't have to

compete? Here, again, N (number of mice) would be useful. Is winning vs. losing unrelated to social dominance due to a lower number of mice (and thus pairs). Does the assay not work in implanted mice? Is the assay simply not reproducible? The strategy for addressing each of these options is quite different.

The authors could certainly strengthen their claims by comparing data within subject across trials. For example, mice ranked 2-4 could be tested in conditions where they were both the dominant and the submissive. Then they can ask whether decoding of the task relevant variables change within mouse based on the rank of the partner, or does the neural activity across reflect fixed rank (and not relevant rank). This analysis should require no additional data and it would really add a lot to this study.

Next, the authors analyzed whether single unit data represented task relevant variables in a manner that was different across ranks. I'm a huge fan of modeling, but the goal of models should be to make things explainable. Here, I worry that the authors models make their interpretations more complicated than they need to be. I have two suggestions. The authors apply an approach to cluster cells based on 6 response properties. However, the response properties appear to be different that the task variables used in other parts of the study. This analysis now includes responses to port entryduring the inter-trial-interval, and I'm unclear why this response property was included as an input to their joint clustering model. 1) Its important to address this issue since two cells which show identical response properties in a given condition (winning) and different responses to another condition could be clustered differently. To this point, the authors go on to compare how the distribution of cells within the clusters differ between submissive and dominant trials. I found the text describing the interpretation of this analysis to be confusing and I spent some significant time with Extended Figure 9 trying to sort it out. Ultimately, I could only assume that reporting on the clustering analysis findings 182-185 was suboptimal and did not clearly describe that the model structure (and thus segregation of cells) was based on joint cellular properties. For example, in principle a cluster is defined by responses to condition A, B, and E, BUT NOT C, D and F. A cell that shows response to A, B, and F only could fall in a different cluster. Thus, 2) it is necessary to fully describe what is represented by a cluster, then one could perform post-hoc analysis.

I found the subsequent analysis of individual response properties to be clearer. In fact, since the goal of the authors was to determine if single unit data represented task relevant variables in a manner that was different across ranks, I wondered why the complicated clustering model was even useful. As a general rule of thumb, a model should only be as complex as is necessary to explain the phenomenon of interest. Thus, in my opinion, the analysis in 4a-c should be removed. The analysis shown in 4f-I should expanded to include winning and losing (Extended Fig 9c). Discussion of these results should be simple and clear. If the authors choose to keep the analysis shown in 4a, a clearer description of the meaning of joint response properties and its importance in social rank should be discussed. For example, Clusters 2, 3, and 4 show strong excitation to winning. Though sub and dom show differences in % cluster 4 (Fig. 4e), no differences are observed in % of cells showing winning or losing (Fig. E9c). Thus, the authors would need to explain all of this more clearly if they intended to keep the complicated model (which I, again, argue is unnecessary for the goal of their study).

Next the authors explore whether PFC->LH neurons contributed to the code for social rank. The authors chosen analysis pathway could be simplified substantially. In 5c-d they show neural

responses to the 6 conditions outlined in their prior analysis. Then they show that a greater portion of PFC->LH neurons show responses to port entry during the inter trial interval. Did the authors test neural responses for all 6 behavioral conditions? Did they compare the three cell types? I'm assuming the answer is yes, but the authors didn't correct for multiple comparisons (where is the statistics. N for PFC->BLA seems to be missing for legend 5e). If not, why did they only choose a limited number of comparisons. Differences were only observed in the self-entry for the ITI. How should this be interpreted?

If the goal is to show that PFC->LH cells encode rank, why does it matter whether they respond to behavioral variables in the task. If the goal is to show that they carry unique information, why is the information in Fig 5e even necessary. Overall, the analysis in 5f-g seems more appropriate. Here, the strategy is to subtract the cells from the model, and determine if the model prediction improves. Butseveral details are missing here. First of all, what is the classifier being trained on (i.e., what behaviorcondition)? Next, the comparison between the BLA and LH projecting neurons is not entirely warranted: a different number of cells are recorded for the two pathways, and it's unclear how many mice that cellular activity was acquired from (important information because only 10 BLA neurons were recorded).

Minor comments:

There is some unusual language in the methods description of the machine learning methods. In the Support Vector Machines, it is noted that the SVM generates a likelihood. It does not, as it is not a probabilistic model. Additionally, the terminology of a "one-hot encoding" is not the standard usage of that term (which usually means that each of L classes is encoded as a 1 on the lth entry of an L length vector), but the one-vs-all approach seems fine. However, details are lacking here: what parameters were tuned? How was the model selection run? What method was used to deal with class imbalance? The definition of the AUC metric should also be given, since there are differing definitions in the multi-class setup. I assume that this is the average of the one-vs-all AUCs, but that needs to be stated. Also, an SVM looks at all data points while choosing the support vectors, so that language is unclear.

The updates on the EM steps are complex, so the communication of this would be enhanced by the authors highlighting the differences in their EM algorithm in comparison to the Escola et al paper.

The communication behind the figures can be improved by focusing on terminology. It was super hard to track at time. Winning/loosing seems to be event locked to the tone. Reward is event locked to port entry (but these trials are referred to as tone). ITI is event locked to port entry (but is referred to as ITI). Perhaps something simple like competition start (tone), competition end (port entry), and ITI for the blocks within trial.

In general, the authors show adjust their training and testing strategy such that testing is performed using a trials from a hold-out with N-1 fold cross validation.

In summary: The study by Padilla-Coreano explored the important question of the role of PFC in social hierarchy. A primary issue that needs to be addressed is whether the cohort of animals used for electrophysiological analysis actually show their primary competitive outcome measure based on rank. This may just be an issue of simply reporting experimental details (N, and needing a few more mice). Many suggestions for improving communication/clarity are provided to the authors; nevertheless, it is important to clearly state that much of the presented results and undergirding ML analyses are unnecessary to support the main theme of the manuscript. My suggestion would be to simplify the manuscript, and to focus the presentation of the findings towards the unit encoding of social preference in PFC, and the role of the PFC-LH circuit in this behavior. Especially considering theprior discoveries outlined by Zhou et al, projection specific findings is the knowledge advance I foundmost exciting. If these broader contributions (AlphaTracker, HiddenStates, MUA modeling, single cellclustering analysis approach) remain in the manuscript, substantial revisions to the communication of these findings and the analyses methodology are warranted as noted above to enable a clearer evaluation of rigor and robustness.

Referee #2 (Remarks to the Author):

In this study from Padilla-Coreano et al., in the Tye laboratory, the team presents research and computational efforts to examine the neurobiological basis for social dominance - a key feature of aggression behavior. They examine the role of the media prefrontal cortex (mPFC) because some prior evidence has suggested that in mammalian species this structure plays a critical role in determining social rank. They go on to develop a trial-based social competition assay, alongside using a custom machine learning approach they developed (Alpha-tracker). Using these two former approaches alongside electrophysiological measures (using a newer wireless method to better facilitate social interaction), they identify a unique behavioral states (9 in total, they claim), which can be decoded from mPFC ensemble activity - predicting social rank and competition winners and losers. Furthermore, they reveal that mPFC to LH projections are better predictions of behavior than an alternative circuit. Overall, this is a very thorough, elegant, and exciting development for the field of behavioral neuroscience, ethnologically relevant social behavior, and computational efforts in the field. It provides a potentially new (or perhaps just a bit better) method for identifying critical behavioral states in interacting/"socializing" mice , while also methods for decoding unique behavioral and neuronal ensembles in tandem within a specific circuit.

However, in spite of this high enthusiasm for the provocative nature of the work, the potential for extensions of this work into other domains, behaviors, and neural circuits, there were some concerns. These primarily centered around controls and/or additional data analysis that would better substantiate their conclusions; along with some stronger rationale provided for particular experiments chosen over others, including the focus on mPFC to LH, vs a whole host of other

possible PFC-related circuits which could have been examined. Below is a list of major concerns (those requiring new analysis, experimental data, or major text discussion/rationale), and then minor comments or concerns follow, that are of less consequence, and more suggestive in nature.

Major:

1) The biggest control that I see which is missing here, is to show data and the entire processing pipeline described (behavior, ensembles, and circuit specificity) in animals that engage in the task alone, and/or with the scent of more dominant animals nearby (the later not being as critical, just an interesting experiment). The "alone control" is important in particular for the trial-based task, that would allow for cross comparisons of activity while engaging in a social dominance bout from the activity and behavior which is independent and is reward - related or orthogonal to the social dominance behavior itself.

1a) This is a sub-point of question 1, which can be more tightly controlled in the analysis, because the alone mouse can also account for starting positions and general posture of the animals in that regard. (which relates a bit to the point 2). The authors might just use their algorithms to look at other behaviors that would perhaps provide some clues to positioning (some ideas, approach, avoidance, orientation relative to mouse vs target reward, etc)

2) The rationale for the tube-test isn't well defined as currently presented. Is social dominance encoded at the neuronal level or is it just dictated by which mouse is bigger in the tube, and who's stronger? I know it is well established, but since it is used a priori here, it is a more germane to the conclusions drawn. Conceptually both make sense , but the tube model of identification in this case leaves one wondering if the data are preselected and therefore biased in some way because of it. Can they shuffle the data in a way that accounts for differences in rank that the tube test doesn't account for? There are multiple social defeat behaviors and postures that can be easily quantified in social aggression and dominance interactions. Why were none of those examined here? They are very granular in nature and would provide rich confirmation and validation in their behavior tracker machine learning algorithms.

3) The authors may already have these data, but as presented their data sets are binarized into subordinates and dominants. Do they see any correlations or decoding properties in mice that are ranked in the middle of a cages social hierarchy? (If a group of 5 has 1 dominant and 4 submissive, which one is the subordinate in their analysis, and would there be differences in accuracy of the decoder based on that social rank within the group - that would be an exciting finding, and it would lend additional biological credence to their predictive algorithms).

4) Did the authors try to present any reward omissions in the task? Does it impact the neural coding and behavior? It might be worth controlling for especially if position and posturing are important.

5) The paper ends with an experiment which shows that the PFC to LH circuit is most critical for the encoding the dominance behavior and predicting it therein. Yet, it wasn't clear to me from the data presented and/or the discussion why one would predict this circuit in the first place, when many other circuits could be better candidates, particularly within the context of social interaction/dominance. Do the photo-stimulation parameters used for the circuit really align with the decoder predictions? How is it so selective? Panel K of figure 5 doesn't appear to show a difference between the control eYFP and the CHR2 group? Is this just underpowered or due to other pathways or circuits being more necessary and sufficient?

6) The authors use the AlphaTracker approach throughout the manuscript and make claims that this approach is superior to other methods. It very well may be, but as presented the authors should either tone down that claim, or substantiate it more with more stringent comparisons with benchmarked alternatives. Considering one of the novelty selling points of the paper is the computational technology it brings to the field. More evidence, data they may already have, and/or discussion for why this is a superior algorithm to other methods would help solidify the novelty of the method and impact.

Minor-

Figure 1 A. The cartoons are somewhat helpful, but it would almost be better presented if the authors took actual photos of the mice in these positions. Presumably they used a high frame rate camera to capture the behavior for Alphatracker, so they should be available. It would be more powerful to see the behavior live in snap shots

This is more stylistics, but just a suggestion. Most of us aren't used to reading / or understanding the plots presented in Fig 3 panel a or c. The manuscript might benefit in this case from a cartoon of what these plots might look like (mock results) in various predictions followed by the plots presented here with the real data.

Figure 3 panel f target plots are very hard to see. They should be increased in size, and/or de-pixelated (smoothened) to make a better case.

Figure 4. The data in A and B are nice, and thorough, but difficult to see unless you are at 400X, it might make sense to zoom in on a few and simplify the figure and put the rest in supplemental. No

one wants to part with all their data on main figures, but it'll help the reader digest the work in this dense paper.

Referee #3 (Remarks to the Author):

This is a very impressive tour-de-force study involving diverse expertise from multiple laboratories combining neurobiological, behavioral, and computational approaches, as well as a novel behavioral tracking methodology in order to establish a causal function of the mPFC-to-LH pathways in social rank and related dominance behaviors. The uses of both the AlphaTracker as well as the HMM-GLM using neural activity for detecting unique behavioral and neural states are innovative for linking across complex social behavioral patterns and neural activity. The results of this study demonstrate the role of the mPFC in processing social rank and provide new knowledge on the mPFC-to-LH pathways in generating dominance behaviors. The authors also have used a wide array of behavioral assays in mice to better examine the selectivity of the findings. Overall, this new knowledge provides an important causal “bridge” between social rank representation in the mPFC and behavioral regulation by the LH. The use of the BLA as a control projection area (mPFC-->BLA) was very helpful in supporting the specificity of the mPFC-->LH pathways. Additionally, these findings involving the LH nicely support the principle behind social homeostasis previously proposed by the authors, which I find it to be very exciting. For these reasons summarized above, I am very enthusiastic about this work.

The reported findings and the ways that the authors have elegantly applied technical innovations should be highly interesting to both basic sciences and social sciences fields interested in social behaviors broadly, as well as clinical fields interested in regulating aggressive behaviors in social settings. All the statistical and data analytic approaches are sound, with appropriate displays of error bars and other useful information in the Extended Data for better understanding the data.

I have the following specific comments for the authors to consider.

1) The authors of the paper have developed a novel paradigm, combining wireless recording, automated tracking of behavior through computer vision, and a competitive reward gathering task aimed at investigating social ranks and dominance behavior. However, from the data presented, it requires more clarification with respect to how much of this task is measuring dominance and
competitive behaviors on *every trial. On line 76, the authors state that “Importantly, differences in winning were not driven by overall location in the arena or distance to port prior to tone onset” and reference Extended Data Figure 1. However, data shown in Extended Data Figure 1B and in Figure

predictive of which animal wins the trial, i.e., in Extended Data Figure 1B, Figure 3F and 3G it is clear that even at –5 seconds before the cue, the winning animals (regardless of "dom win" and "sub win") are both closer to the port. The 2-way ANOVA in the legend of Extended Data Figure 1B found a main effect of trial type (presumably dominant win/subordinate lose vs. dominant lose/subordinate win), and no effect of rank (p=0.071) or interaction. It is unclear from the text what time period of the fifteen seconds shown this ANOVA is examining.

If distance from the port prior to the cue onset is strongly predictive of trial outcome, then not *all trials may be actually "competitive" or "measuring dominance". An alternative explanation would be that during some trials the nearest mouse to the port can more easily gather the reward, and this trial is either weakly or not contested by the other mouse (could be due to dominance or motivation level, for example). If this is the case, then it would change the interpretation of mPFC population dynamics that track "competitive success". To address this, we would suggest the authors do one of the following:

a) Specifically examine if the winner of each trial and the neural activity predictive of winning were not simply determined by distance from the port during the baseline periods.

b) Exclude any trials from the analyses that do not include pushing, resistance, or displacement behaviors to provide additional neural insights when specifically examining only those trials with clearly observable competition (pushing, resistance, displacement).

c) Examine if prior behavioral dynamics and underlying neural activity between the mice prior to the cue determined their relative positioning during the baseline period (i.e., the competitive behaviors were occurring in advance of the cue presentation) to better understand what behavioral dynamics and neural processes occur that led to the better positioning to begin with that ultimately resulted in winning.

For future directions (note: I am definitely not suggesting the authors to collect more data), I would recommend including a condition where pairs of mice in an experimental chamber with multiple reward ports, where the location of the upcoming cue cannot be predicted and where presumably the positioning during baseline between the dominant and subordinate would average out to being equivalent.

2) Similarly, further clarifications of the behavior would be useful in interpreting some of the data. For example, on line 74 the authors state that “Dominant animals, as defined by the tube test, obtained more rewards, spent more time at the reward port, and were more successful at displacing the competitor from the port (Fig. 1c).” This seems to be in contrast to line 99 of the manuscript which states “When recording during the reward competition task, we did not detect a statistically significant difference in the number of rewards earned by dominant and subordinate mice, allowing us to make comparisons about dominance behavior and competitive success without being confounded by the volume of reward consumption.”

Does this indicate two separate behavioral datasets were used, one during neural recording and one without neural recording? Related: please add the test and the number of observations used for “we did not detect a statistically significant difference”.

3) In all population level neural analyses, the authors have pooled together activity from multiple subjects and use the total number of recorded units as the sample size. This approach is not necessarily unorthodox in the field of behavioral neurophysiology where collection of spiking activity during interactive social behaviors is extremely challenging and large pseudo-populations of units are sometimes requisite for advanced analyses such as hidden markov models. In some areas, the authors made specific efforts to address the potential shortcomings of pooling together activity from multiple subjects. For example, in (Extended Data Fig. 7k-n) the pooled data was split between two different randomly selected subsets over 50 bootstrapping iteration, or in Figure 3B/3D where the authors employed the “leave one out” method excluding neurons from a single animal in each iteration to control for an individual mouse from strongly biasing the results. In other analyses, however, there is no control for potential between-subject variations. Specifically, if it is not possible to construct a HMM-GLM for each subject, or to split individual mPFC cells by rank-dependent responses for each subject, then a good alternative would be to also show inter-subject correlation coefficient that assesses the homogeneity (or potential / interesting differences) of neural responses between mice.

4) In the comparison between PCA vector length, it is stated that subordinates had longer neural trajectories, speculated to be due to either higher or faster firing rate changes in the mPFC population activity. This result is quite intriguing, and although this is later addressed tangentially (Line 187; “subordinates had phasic responses of greater amplitude in response to events…”), the interpretation of this result would greatly improve if that question could be examined quantitatively with the existing data.

5) In Figure 5J, the difference between cumulative rewards obtained by light OFF session vs light ON session appears to only emerges late in the experiment (past 15 trials). This is interesting. Do the authors have any interpretation for this time course? For example, the non-immediate effect of light delivery may suggest that the modulation of social dominance by the cortico-hypothalamic circuit is mediated through slower mechanisms, such as learning.

6) Different analyses required using various sample sizes for good reasons. For example, both Figure 3B and 3D use a sample size of thirteen (resulting from 'leave one out' method). This is only explained in the supplemental text, and the accompanying figures (3A and 3C) use a different samplesize (507 and 490 neurons). Some other figures lack descriptions of sample size (e.g., Extended Data Figure 7). Please add N information and also add how many total mice were used in these studies and which mice (if any) overlapped between different aspects of the manuscript (which can even be in the method section).

7) In several figures (Figure 1G, Extended Data Figure 3A and 3B) output from UMAP (Uniform Manifold Approximation and Projection) clustering is presented without any explanation in the main,or supplemental, text. Although this visualization is informative and appropriate, this technique is still relatively novel, and some readers may not be able to interpret the output (i.e., non-labelled axes) without some explanation or reference to McInnes Et al. 2018.

(McInnes, Leland & Healy, John & Saul, Nathaniel & Grossberger, Lukas. (2018). UMAP: Uniform Manifold Approximation and Projection. Journal of Open Source Software. 3. 861. 10.21105/joss.00861.)

Signed below to opt in for Nature's new transparent peer review scheme:

Reviewed by Steve W. C. Chang with assistance from Philip T. Putnam, a postdoctoral associate in his lab.

Author Rebuttals to Initial Comments:

Point-by-Point Response:

Overall, we were delighted to receive the highly constructive and overwhelmingly positive comments on the first round of review from all three of the excellent reviewers who each provided insightful and thoughtful critical feedback on our manuscript, and collectively referred to our manuscript as "***a very impressive tour-de-force study involving diverse expertise from multiple laboratories combining neurobiological, behavioral, and computational approaches, as well as a novel behavioral tracking methodology in order to establish a causal function of the mPFC-to-LH pathways in social rank and related dominance behaviors***" that was "***Overall, a very thorough, elegant, and exciting development for the field of behavioral neuroscience, ethologically relevant social behavior, and computational efforts in the field***" and "***a major contribution***" comprised of "***elegantly applied technical innovations***." We are also grateful that the Reviewers were such experts in the field they could appreciate that, "***the uses of both the AlphaTracker as well as the HMM-GLM using neural activity for detecting unique behavioral and neural states are innovative for linking across complex social behavioral patterns and neural activity***."

All three reviewers described our work as "***exciting***." Reviewer 1 prefaced their deep, comprehensive, and constructive feedback by emphasizing that "***the richness of our suggestions should not be misconstrued as a lack of support for this work***," while Reviewer 2 held "***high enthusiasm***" and Reviewer 3 was also "***very enthusiastic***" about our manuscript. We are pleased that Reviewer 3 so concisely articulated our biological advance "***Overall, this new knowledge provides an important causal "bridge" between social rank representation in the mPFC and behavioral regulation by the LH."***

We thank the reviewers for their constructive comments and suggestions on our manuscript, and for their patience, as fully addressing their comments required several new, technically-challenging experiments. We thank the editor and reviewers for granting us the opportunity to respond to the critical comments from their thoughtful reading of our manuscript and feel that the incorporation of their very helpful comments have substantially expanded the scope of claims that we can make and thus strengthened the study.

In this new version of the manuscript we have incorporated both ***new experiments*** and ***additional analyses***:

1. We have validated our AlphaTracker tool for tracking of four mice (Extended Data Fig. 2).
2. We have compared our HMM-GLM model with additional models for increased rigor (Fig. 2 and Extended Data Fig. 5).
3. We have now analyzed how mPFC neural dynamics change in intermediates when they are relative subordinates vs dominants (Fig. 3c, l; Extended Data Fig. 9).
4. We have included the important control of recording from mPFC when mice perform the reward task alone as well as during competition. ***This experiment reflects new data collected using wirelessly-recording devices from the mPFC of 24 mice while performing the reward task alone***. We now report that the mPFC dynamics we observe at the population and single cell level are not driven by baseline social rank differences (Fig. 3j-k; Fig. 4b-d).
5. In a new experiment, we directly compared encoding of a reward received alone vs in social competition. We recorded the same neurons by using a continuous recording sessions that include alone and competition epochs. We demonstrate that encoding of receiving the reward alone vs winning is distinct and decodable (Fig. 3l-m).

We believe we have responded to each Reviewer‟s concerns completely and comprehensively, and in so doing, have included a table of contents for your convenience (comments are color-coded by Reviewer, Authors‟ responses are in black font).

Referees' comments:

Referee #1 (Remarks to the Author):

I am a behavioral neurophysiologist, and I reviewed this paper in conjunction with a computational neuroscientist/neurotheorist. Comments reflect our pooled suggestions to primarily aid the authors in improvingthe rigor and reproducibility of their work and enhancing communication of their findings for a broad scientific audience, and to secondarily aid in editorial consideration of this work for publication. The richness of our suggestions should not be misconstrued as a lack of support for this work, as noted in the summary

This study by Padilla-Coreano et al seeks to investigate the role of the PFC, and more specifically, a PFC->LH circuit in encoding social rank and dominance behavior. To achieve this, the authors begin by determining the social rank using a classic test assay in pairs of co-housed mice. Since the primary objective of this study is to uncover the neural representations of social rank, the authors suggest that the tube test has several clear limitations. First, the behavioral data that can be extracted from the tube test is less rich than other behavioral assays broadly deployed in mice. Secondly, I think the authors seem to suggest that the task does not enable the measurement of the electrical activity from the same cells across many trials (which is necessary to facilitate statistical analysis). Here I think its worth largely emphasizing the first point. As previously demonstrated in Zhou et al, Science 2017 (cited by the authors), mice can be repeatedly subjected to the tube test. Since this is a widely accepted and stable (across trials and sessions) measure of social dominance in male mice, and PFC cellular responses to competition have already been described in Zhou et al, Science 2017, a major contribution here is showing that these responses are presented in a more ethologically relevant context. Thus, the authors develop a novel competition assay. Using a set of mice in which social hierarchy is previously established using the tube test, the authors successfully show that social dominance predicts behavioral outcomes in their competition assay.

We thank the reviewer for acknowledging that there is a major contribution of studying the PFC in a more ethologically relevant context. First let us say that the groundbreaking study of Zhou et al., Science 2017 was a landmark study that we are building off of and going beyond – **our results are completely aligned with this published work, and we do not intend to replace the tube test, instead, we propose to contribute to the existing toolbox for probing the neural representation for social rank and provide a new paradigm optimized for neural recordings/systems-level investigations**. Indeed, the tube test allows for repeated tests that can be used as trials, as previously published and it was an oversight on our part to not be clear with our the more clearly defining what we meant by “trial structure.” We have clarified in line 70 that we mean structured trials that are not initiated by the animal, but initiated by the experimenter. With the tube test, trial duration is determined by when the animals initiate and terminate movement in each trial and is therefore non-uniform, which represents a challenge for averaging across trials (which is necessary for certain types of population dynamics visualization). Our task provides a trial structure that is defined by the experimenter rather than initiated by the animal‟s behavior, that contains the identical conditioned stimulus to help determine which responses are to the value of the predicted outcome rather than the features of the unconditioned stimulus.

Nevertheless, here I would raise my next suggestion. With the tube test assay, five animals can be ranked 1-5 (to social rank). Mouse #2 has a high social rank, though it is submissive to Mouse 1. With competition assay and the analysis strategy taken by the authors, it is imperative to note that they are not probing social rank (i.e.,Mouse #2 within the colony). Rather, they are probing relative social rank (Mouse 1 vs. Mouse 2)… the neural processes that encode being dominant to your partner. At this point, its unclear if mouse 2 would look any different than mouse 4 if all five animals were performing the competition assay at the same time. At some places in the paper the authors appropriately use the term relative social rank, and in other places, they simply use the term social rank.

Indeed the reviewer is correct that in our previously-submitted manuscript we were probing relative rank, and should specify such. Now in addition to stating that comparisons made will be relative rank in the manuscript (lines 78-79), we have now clarified in each figure where we mean relative rank. In this new version of the manuscript, we have introduced some analyses of absolute social rank, which we refer to as absolute rank where necessary.

We are also now reporting the outcome of the competition by absolute social rank, showing that rank 1 animals win the highest percent of competitions, in Figures 1c and **Extended Data Fig. 5b**, shown here in **Rebuttal Figure 1**. Finally, to clarify that in all cases in our reward competition is done in pairs of mice, not all cage mates simultaneously, we have added extra details to legends for figures 1c-d and on the ***methods section 6*** we have the following details" During all competition sessions, two mice were placed in one chamber. Mice competed against one cagemate per session."

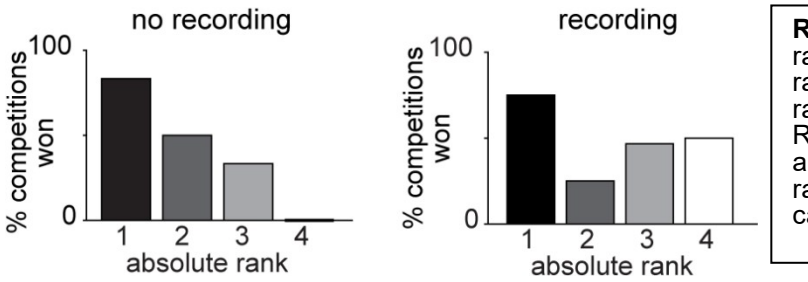

Rebuttal Figure 1: Competitions won by absolute rank. Left, percent competitions won by absolute rank without recordings (number of competitions: rank 1=5/6, rank 2=3/6, rank 3=2/6, rank 4=0/6). Right, percent competitions won by absolute rank in animals during recording sessions (rank 1 n=9/12; rank 2 n=3/12; rank 3 n=7/15; rank 4 7/14). In both cases rank 1 animals won more competitions.

After the authors establish their novel behavioral assay as a surrogate for relative social rank that also facilitates multiple trials, they strive to extract a rich set of behavioral measures using a novel ML coding algorithm. Importantly, their approach can monitor the behavior of non-labeled individual mice, and it can extract joint behavioral features from pairs of mice. Here, however, I would suggest toning down the novelty/utility claims of this approach a tiny bit. Its nice that their approach can track non labeled mice, but its not that difficult to label mice. Furthermore, while tracking non-labeled mice within a large colony might be exceptionally useful, they only show its effectiveness in tracking two mice. Thus, I would suggest that they revise the description of AlphaTracker to state that it tracks two mice, not multiple (in the abstract, text, and figure legend). Its a more precise claim that doesn"t take anything away from the utility of their scientific findings using the tool.

We thank the reviewer for this suggestion. We have now included new data and analyses in **Extended Data Fig. 2a**, also included here in **Rebuttal Figure 2**, where we show that when tracking 2 or 4 unmarked mice AlphaTracker can track mice with near or higher accuracy than humans with an identity error rate of less than 1%. We have also toned down our claim by explicitly stating that existing tools can track multiple marked animals in line 83.

However, while the reviewer trivializes a feature of our tool with the statement "it"s not that difficult to label mice" we would like to highlight the utility of this feature:

1) computer vision tools for pose estimation/animal tracking in labeled and unlabeled subjects are distinct classes of tools;

2) using markers can be trivial in some tasks, but in other tasks (for example, in social identity experiments where a marker could confound social identity recognition or when multiple mice are frequently occluding each other), markerless pose estimation is a significant advantage in any case when the marker could confound an experimental variable or where the marker could be occluded; and

3) AlphaTracker certainly performs even better when animals are marked, but importantly performs well tracking in multiple unmarked mice (which may be a useful feature for certain experiments).

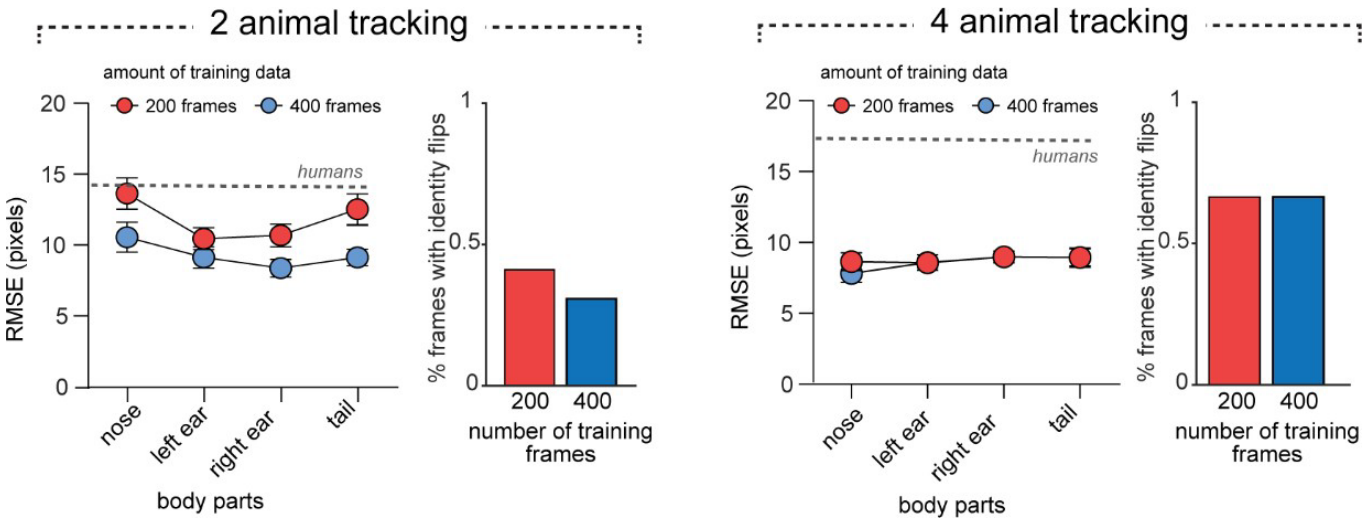

Rebuttal Figure 2: **AlphaTracker performance for multiple animal tracking.** Left, the root mean squared error (RMSE) in pixels for two animal tracking. Right, the RMSE in pixels for four animal tracking. In both cases identity tracking is more than 99% correct.

Once the authors identify behavioral clusters using AlphaTracker, they set out to determine whether PFC firing represents social rank and the behaviors observed during their task used to quantify relative social rank. Specifically, the authors implant animals with 32 wire electrodes, and perform wireless while mice engage in their task. Here, some missing details/confusing text made it more difficult for me to evaluate the rigor of their findings. First, I had a hard time determining the number of implanted mice subjected to their assay. I searched the text, figure legends, and methods, and I am sorry if I missed it.

We apologize for the confusion and thank the reviewer for pointing out this omission. We have now indicated the number of mice utilized for recordings in **Figures 2-5** legends and in the ***methods section 6***. For **Figure 2** we used 13 mice with similar video settings (resolution and camera angle) to allow for the automated behavioral analysis. For **Figure 3** we used 20 mice for the reward competition recordings, 24 mice for the alone recordings and 12 mice for the experiment in which mice received alone and competition trials in the same recording session. For **Figure 4** we used 10 mice for the alone recordings (5 rank 1 and 5 rank 4) and 20 mice for the reward competition recordings. For **Figure 5** we recorded 20 mice of which 9 had viral injections for mPFC-LH phototagging and 8 had viral injections for mPFC-BLA phototagging.

Second (and also related), it was hard for me the interpret the following statement (line 99) without that data, “when recording during the reward competition task, we did not detect statistically significant differences in the number of rewards earned by dominant and subordinate mice, allowing us to make comparisons about dominance behavior and competitive success without being confounded by the volume of reward consumption.” I was totally lost by this sentence. In Fig 1c, the authors validate their assay as a surrogate for social rank by showing that commutative rewards were related to outcomes in the tube task. Here, the authors seem to be suggesting that the outcomes of the tube task do not relate to successful competition for limited resources. There are several reasonable explanations for this. For example, the electrophysiological experiments may not be powered to detect this behavioral difference (but I can"t assess that without N). Did the partner mice wear the dummy headstage during the ephys competition (seems so in the video)? For these mice, there appears to be less cumulative rewards in the unimplanted mice.

We have now removed this sentence from the manuscript and clarified in the ***methods section 6*** that during recording sessions both animals always wore the wireless devices and that to maximize the number of trials obtained sessions consisted of 30 trials for recording sessions (to enhance statistical comparison) and 20 trials

for unimplanted mice. Furthermore, we now report the number of animals used in the figure legends and ***methods section 6***.

There are different behavioral measures shown in figures for the initial version of the competition task in un-implanted mice (cumulative rewards, ratio of port occupation, pushing success and time displaced), than for the implanted mice (cumulative rewards, pushing success, latency pushing success during tone). Here the onlycommon measure that shows the same result is the pushing success. For the latencies, the trials appear to be treated as independent for this latter analysis. The authors should treat these behavioral observations as repeated observations of the same phenomenon for Fig S5c (specifically, within trial for a given subject). Also, the legend for Fig S5c says „slower latencies." It would be better to describe the data as „longer latencies."

The statistical test we used for the latencies is a two sample Kolmogorov-Smirnov test which is adequate for this data because it does not require normality. Yet, it doesn"t have a special case for repeated measures within group. Now we also report a repeated measures one way ANOVA showing a consistent effect of relative rank (One way RM-ANOVA $F_{(1,24)}$=2.06, p=0.002). And yes, “longer latencies” is more accurate, we have made this change.

Since the nature/validity of the behavioral task is important for exploring the link between relative social-dominance and neural activity, my further commentary assumes that the suggestions I"ve provided up until this point can be easily addressed with additional experimental details and minor revisions to the analysis. Thus, I am providing additional suggestion based on the assumption that the authors will clearly establish behavioral outcomes of their competition task as measures of social dominance for the same animals used for their neural analysis in a revised version of the manuscript.

During the recordings, food-restricted mice (weighing an average of 25g) carry on their heads a recording device that weighs 6-7g including cement of headstage (~25% of their body weight); it is inevitable that there will be some changes in behavior (**Rebuttal Figure 3**), even if they are minor in comparison to tethered mice. We have also made sure to report the same metrics across cohorts in the manuscript where relevant (**Figure 1** and **Extended Data Fig. 1** for no recording cohort and **Extended Data Fig. 5** for recording cohort).

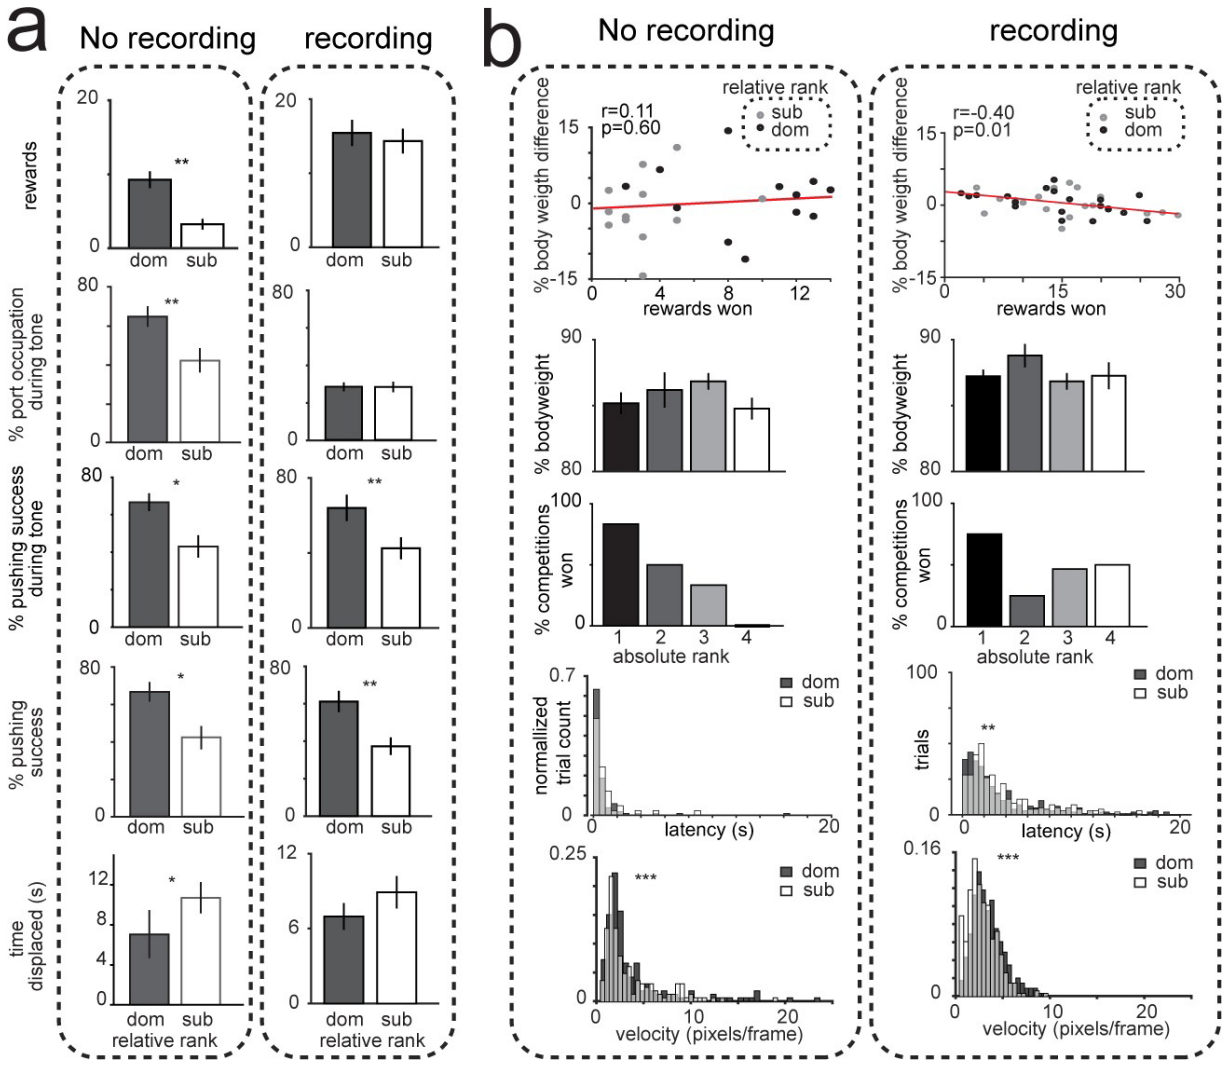

Rebuttal Figure 3: **Reward competition behavior summary for recording vs no recording cohorts**. **a**, Number of rewards, % port occupation during tone time, % pushing success during tone and total % pushing success, and time displaced from reward port across relative dominant and subordinate mice. Across cohorts % pushing success during tone and total % pushing success are higher for relative dominant mice. **b**, Rewards won as a function of bodyweight difference, % bodyweight and % competitions won by absolute rank, distribution of latency to pick up reward across trials and distribution of velocity across trials. Across cohorts rank 1 mice win the majority of competitions and relative dominant mice have higher velocities during the social competition. *p<.05,**p<.01,***p<.001

In both cohorts, pushing success, which represents periods of high competition for both animals, is higher for relative dominant mice than for relative subordinate mice. In addition, our physiology experiments show that in the mPFC there are unique rank-related patterns of activity that emerge during the social competition but are not present at baseline when the animals are performing the task alone (**Fig. 3j-m** and **Figure 4**).

The language around the authors population analysis is a little unclear. As I understand it, the authors used AlphaTracker to discover nine behavioral conditions defined by two mice in an arena. They then trained a SWM using PFC multiunit activity to predict these behavior conditions. Next, they hypothesized that PFC may better encode hidden states. Here, rather than PFC directly encoding the 9 specific behavioral conditions (or better stated, rather than discriminating the behavioral states from each other), the authors tested a model in which PFC encoded a number of states (six in this case) in which distinct profiles of their behavioral conditions

were likely to emerge. They found that PFC decoding was more accurate using this latter model in comparison to their models where PFC encoded the behavioral conditions.

First, we would like to clarify that the HMM-GLM model does indeed show that PFC distinctly encodes 9 specific behavioral states, and it does so by assuming there is a hierarchical model wherein the animals hidden state contains probabilities of the behavioral states identified that are distinct from each other hidden state.

Second, we would also like to clarify that the 9 behavioral conditions (or labels) analyzed in Figure 2 are indeed analyzed using AlphaTracker tracking data, but they are unrelated to the clusters of behavior shown for the arena behavior shown in Figure 1 which is a proof of principle use of the tool we have created and open sourced. We apologize for the confusion and have now added additional clarifying language in the ***methods section 17*** "Decoding of behavior" under the first subsection titled "Dataset."

The results on the modeling are unusual, where gains between the autoregressive model compared to the SVM and the GLM are drastic. It is rare that you see such dramatic gains when moving to an HMM based model, and the manuscript would be much improved by explaining this outcome. In fact, the authors show that there are drastic gains by using 2 clusters in the HMM based model (noting that the 1 cluster case is equivalentto the GLM model in your formulation). In this case, they are now distinguishing between 9 classes based on a mixture of 2 GLMs, where the GLM is chosen through the historical information. Compared to the classes, this is really a minor gain in modeling complexity, and I"m unclear from a mathematical perspective as to how such gains can be achieved. Additional justification and exploration would be warranted to ensure robustness of these findings.

The purpose of the analyses in Figure 2 is to determine if this novel social competition task is prefrontal dependent. The literature of the tube test suggests that mPFC would be involved in this social competition assay, but since this is a novel assay we needed to confirm the role of mPFC. To do that, we proceeded to ask if mPFC population activity could predict the behavior of mice in this social competition assay. We compared multiple models and report the model that worked best. We see the best performance with a model that combines hidden Markov models with generalized linear models. This has interesting implications given that this model uses hidden states that may reflect different internal states of the animals during the social competition. In the revised version of the manuscript we have added some additional controls for increased rigor as the reviewer suggested and we now explain the model complexity better in lines 112-126 and the control comparisons we performed in the results lines 132-141, and the details on these new model comparisons are in the ***methods section 17*** Decoding of behavior under subsection Dataset in pages 7-

8.We have now amended our main text in Lines 112-141 such that it now reads:

"*Once AlphaTracker facilitated the identification of 9 different behavioral states (Fig. 2a), we then wanted to determine how the mPFC may predict behavior during the reward competition assay. The mPFC is known to be important for a number of higher cognitive functions(Sawaguchi and Goldman-Rakic, 1991; Miller and Cohen, 2001; Wallis et al., 2001; Hornak et al., 2004; Ridderinkhof et al., 2004), and has been shown to used "mixed selectivity" to maximize computational power – which refers to the ability of mPFC neurons to be selective for different stimulus features under different contexts(Rigotti et al., 2013; Fusi et al., 2016).*

We posited that mPFC neural activity could be dynamic, representations may be hierarchical and may be influenced by internal hidden states. Therefore, we turned to a recently-developed unsupervised method to identify hidden states by combining a hidden Markov model (HMM) with generalized linear models (GLMs)(Escola et al., 2011a; Calhoun et al., 2019a) and adapted it to use mPFC neural activity to predict each of nine behavioral labels. We trained a set of multinomial GLMs to predict the transition probabilities between

hidden states. In addition, each hidden state is paired with another multinomial GLM that describes the relationship between neural activity and the behavior of that particular hidden state (Fig. 2a-b).

To create a dynamical model of the temporal relationship between neural activity and behavior, each component of our model followed the first-order Markovian property, to help preserve information about past events when predicting the future (Fig. 2d). An HMM-GLM model with 6 hidden states decoded behavioral labels from neural activity with superior performance to static models (Fig. 2c-e, Extended Data Fig. 5j, Extended Data Fig. 6, and Supplementary Movie 2). Interestingly, the model performed equally well when training for one relative rank and testing on the other (Extended Data Fig. 5k-l), suggesting that mPFC encoding of social competition behavior is generalizable across relative ranks.

However, we also wanted to consider the alternative hypothesis that mPFC neural representations are static or simple, so we next tested simple models to see if they too could predict social behavior. The simplest model that could predict behavior is to utilize the fixed probabilities of the behaviors occurring (frequency table) to try to predict behavior labels, but this model without neural information failed, suggesting the need of incorporating neural activity for the prediction (Fig. 2c-d). Next, we trained a multinomial support vector machine (SVM) classifier and a multinomial generalized linear model (GLM) to decode these behavioral labels using mPFC multi-unit activity (Fig. 2c-d). The SVM failed and the GLM performed above chance (Fig. 2c-d). Next, we tested if the performance increased by utilizing an ensemble of GLMs, one per behavioral label, but there was no improvement in performance (Fig. 2c-d). The lack of improvement suggested that the static nature of the GLM did not fully capture the relationship between mPFC neural activity and behavior."

The dramatic gains of the HMM-GLM models are due to having multiple GLMs and incorporating a hidden Markov model that uses temporal information. This model incorporates one GLM per type of hidden transition state (hidden state 1 to 2, hidden state 1 to 3, etc.) and one GLM per type of behavior observation (hidden state 1 to behavior 1, hidden state 1 to behavior 2, etc.). For example, the 2-state GLM HMM consists of 22 GLMs (4 for the hidden state transitions and 18 for the behavior emissions), leading to a significant increase in complexity. In addition, the HMM-GLM model has an autoregressive component of the model, i.e. the HMM, that improves the performance by incorporating temporal information of the previous timepoint using a Markov chain. These hidden states are meant to represent the internal states of the brain and each of them define distinct relationships between neural activity and behavior via the emission GLMs.

For example, at time 't', we have a probability distribution for the 6 states. Using neural activity at this time, we can now define a probability of transition for each of the states at time 't+1' using the transition GLMs. Thus, we have 36 possible configurations. We then identify the most probable hidden state for time 't' as the one which produces the highest probability for the true behavior label at time 't. Thus, using the Markovian property, we reduce the number of possible configurations to 6. This is how the HMM uses the information of past events to narrow down probabilities for the future. Next, each hidden state (6 of them) can lead to each behavior (9 of them) using the corresponding emission GLM per behavior/hidden state. Thus, providing 54 probabilities of behavioral outcomes coming from hidden state and behavior combinations. We concatenate the probabilities across these for each behavior to obtain one final value per behavior. For example, say the probability distribution for the hidden states (H) at time 't+1' is: P(H1) = 0.1, P(H2) = 0.3, P(H3) = 0.07, P(H4) = 0.2, P(H5)
= 0.3 & P(H6) = 0.03, and the conditional probabilities for behavior 1 (B1) given each hidden state are: P(B1|H1) = 0.01, P(B1|H2) = 0.05, P(B1|H3) = 0.002, P(B1|H4) = 0.12, P(B1|H5) = 0.17, P(B1|H6) = 0.2, then the probability of observing B1 (using the law of total probability) becomes: P(B1) = 0.1x0.01 + 0.3x0.05 + 0.07x0.002 + 0.2x0.12 + 0.3x0.17 + 0.03x0.2 = 0.09714. Thus, we use this complex set of interconnected GLMs driven by neural activity to describe a diverse behavior distribution in conjunction with the Markovian property. The combination of GLMs and HMMs with the Markovian property is what gives such a good performance. Importantly, given the complexity of this model preventing overfitting was key. To control for overfitting we did two methods in which we always tested with held-out data: 10-fold cross validation in which we trained with 90% of the trials and tested on 10% held-out trials and leave one out method in which we trained with all animals except one, and tested the model on the held out animal. These details about how the

HMM-GLM works are detailed in the ***methods section 17*** subsections HMM-GLM and Model Architecture & Assumptions in pages 8-16.

To be conservative, we refrain from making explicit claims in our manuscript, but the idea that a model with a hidden layer would perform better is consistent with the notion that the PFC uses a hierarchical representation, and is consistent with the notion that the PFC relies on mixed selectivity to exert its computational power. PFC neurons are known to encode different stimuli depending on context (Rigotti et al., 2013), and hidden states may be the neural representation of a given context. We can now add this speculative interpretation of the additional significance to the manuscript.

In any case, if my understanding of their objective is correct, the authors should revise their text for clarity. The term „behavioral state" is used under too many different contexts. For example, the authors state "the proportion of time spent in each hidden state did not different by competitive success or by social rank, and the model performed equally well across ranks, suggesting that mPFC encoding of social competition behavior states is common across ranks." It would be more appropriate to state that the encoding of social competition hidden states is common across ranks. Also, again, the authors should revise most statements regarding rank to read „relative rank". Assuming that my understanding is correct and that the authors clarify the text accordingly, I believe the data argues that PFC activity encodes behavioral states in which distinct profiles of behavioral are likely to emerge. There is plenty of evidence to support this interpretation including studies from the senior author and many of the works literature cited in this manuscript. Nevertheless, the relevance for this particular study remains unclear since these behavioral states (hidden states) do not show a relationship to rank.

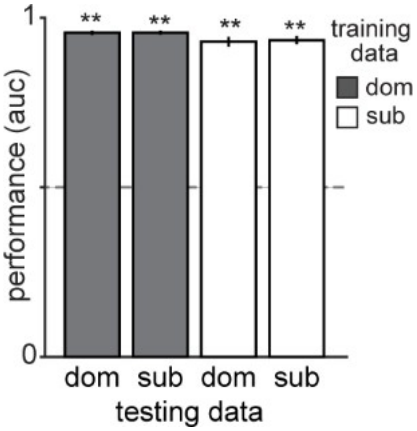

Rebuttal Figure 4: The 6-state HMM GLM was predictive of behavior with high performance regardless of which dataset was used for training or testing. Training on relative dominant (dom) or relative subordinate (sub) data and testing on the opposing dataset still resulted in high performance (n=9 behavior labels using 482 trials for *dom* vs 478 trials for *sub*; Sign test performance vs 0.5 (chance) p=0.004 for all tests).

We thank the reviewer for pointing out that behavioral state is a confusing term given that we also have hidden states. We want to clarify that we were using behavioral states to mean the behaviors themselves, therefore now we are referring to them as behaviors or behavioral labels to avoid confusion with the hidden states, which although are determined based on the behavior and neural activity they are not the same.

Regarding the point on hidden states not showing relationship to relative rank: That is indeed the case, we have now more extensively explored to what extent the HMM-GLM model predicts behavior differently across relative ranks and we show that if you train the HMM-GLM model on exclusively one group and testing on opposite group we can still predict behavior labels with high accuracy (**Rebuttal Figure 4**). Based on these analyses we conclude that mPFC has a common model of encoding social competition behaviors across relative ranks, even though the proportion of time performing different behaviors may differ across ranks.

Next, the authors set out to determine if PFC unit activity predicted social rank and task relevant behaviors in their competition assay. My concern with this approach, in particularly analyzing winning and losing, is that the authors did not find that winning vs. losing was related to social rank in the cohort of mice used for their electrophysiological study. Specifically, the authors found that the dominant and submissive mice consumed the same amount of reward. Thus, I"m not entirely convinced that decoding winning and losing is optimal. Imagine a situation where basal difference in PFC firing activity are related to relative social rank (a fair hypothesis). One would expect that a model using PFC activity could differentiate relative social rank irrespective of the behavioral condition as long as two mice were in the same arena. This appears the be the

case, since the AUCs differentiate relative rank even prior to cue presentation. How then to interpret the trial by trial decoding of winning a losing? Would the PFC show this response in any area, even one in which resources were freely available? What if there were two poke holes, and the animals didn"t have to compete? Here, again, N (number of mice) would be useful. Is winning vs. losing unrelated to social dominance due to a lower number of mice (and thus pairs). Does the assay not work in implanted mice? Is the assay simply not reproducible? The strategy for addressing each of these options is quite different.

We thank Reviewer #1 for this deep and thoughtful line of questioning. Indeed, the reviewer helps to point out that we needed to more clearly articulate the advantages of this new task and these conditions. One of the most powerful advantages of our task is that, by having many trials, we can parse related information across different time scales.

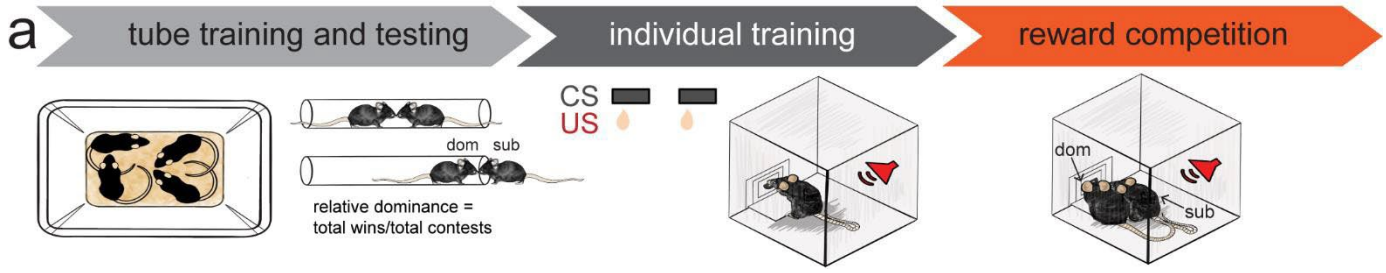

Main Fig. 1a: Schematic of reward competition behavioral paradigm. First, social rank was assayed using the tube test. Once ranks were stable, mice were trained individually to associate a tone with reward delivery of Ensure 2 s post tone onset. Finally, after training, mice competed for rewards during competition with cagemates of known social ranks.

Specifically, wins and losses occur on a short time scale, to contribute to the longer time scale representation of rank. The Reviewer is concerned that our task does not have all the advantages of the Tube Test, yet we want to emphasize that our paradigm combines use of the Tube Test (as depicted in **Figure 1a**) with the Reward Competition Assay, and thereby combines the different strengths of these assays to allow us to make new insights. The Tube Test is more all-or-nothing (qualitative) and is designed to be robust, while the Reward Competition Assay is designed to reveal subtle gradations that can be quantified. <u>We are not proposing that the Reward Competition Assay replace the Tube Test</u>, but that the use of both paradigms in combination is more powerful.

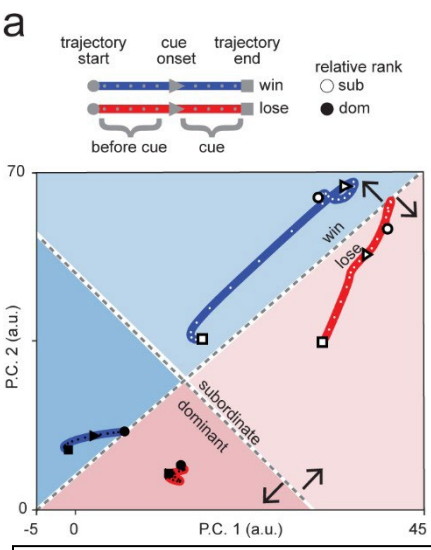

One exciting set of results is that although wins and losses are highly related to social rank (as rank would likely influence win/loss probability, and win/loss probability would also influence and define social rank), they are represented orthogonally and occupy distinct subspaces within the greater activity space (**Figure 3a and b**). We speculate that this could represent a strategy for ensembles to keep related information on different timescales separate. This central discovery (that wins and losses would be represented orthogonally and in segregated subspaces) would not have been possible without the trial structure and features of this task.

Main Fig. 3a: Neural trajectories of mPFC population firing rate differ for relative dominants (*dom*) and subordinates (*sub*) during the tone presentation for both *win* and *lose* trials in a lower dimensional common principal component (PC) sub-space (trajectories are the average across leave one out iterations leaving out one mouse at a time, total neurons recorded from dominants: n=507 and subordinates: n=490 units from 20 mice). *Win* and *lose* trials are aligned to the cue

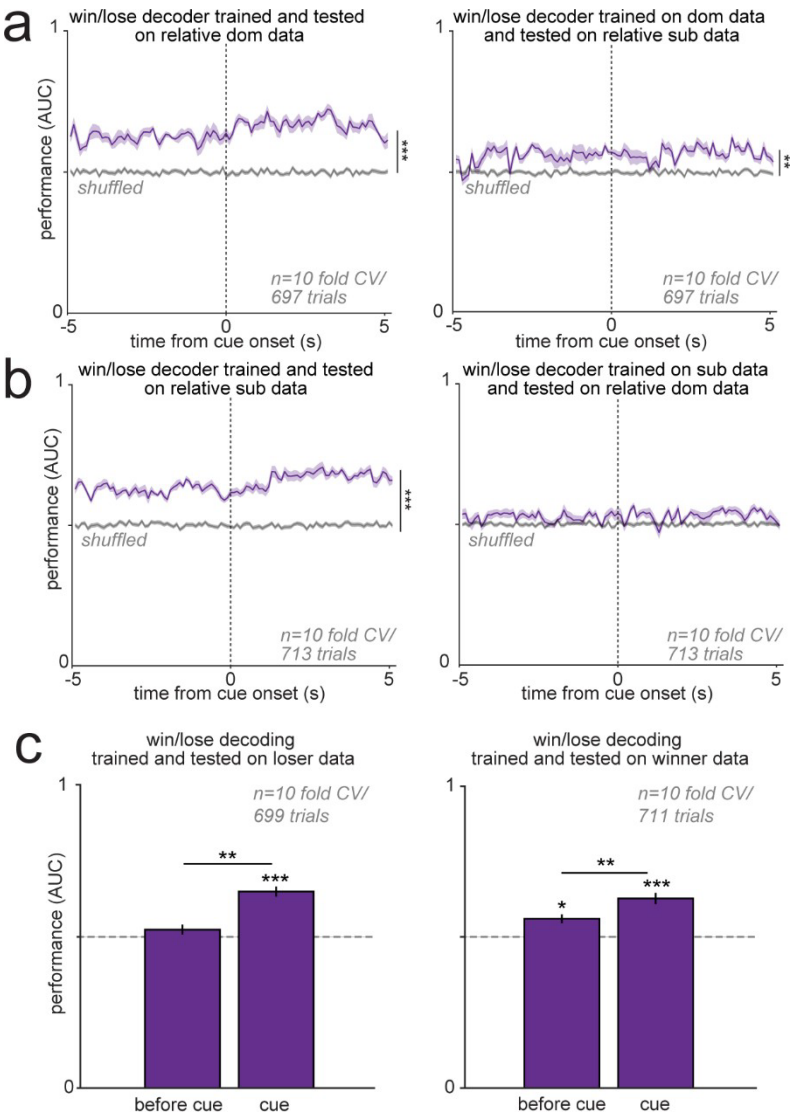

Rebuttal Figure 5: Decoding performance for competitive success using different datasets. **a**, Decoding performance (area under the receiving operating curve; AUC) when training and testing on relative dominant data (left) or training on dominant and tested on relative subordinate data (right) was higher than chance (shuffled dataset indicated in gray). (Wilcoxon rank sum, *dom/dom* p=0.0002, *dom/sub* p=0.003). **b**, Decoding performance (area under the receiving operating curve; AUC) when training and testing on relative subordinate data (left) was higher than chance but not when testing on relative dominant data (shuffled dataset indicated in gray). (Wilcoxon rank sum, *sub/sub* p=0.0002, *sub/dom* p=0.14). **c**, Decoder performance for classifying competition outcome using training and testing data from loser data (e.g. mouse lost majority of trials) and using training and testing data from winner data (e.g. mouse won majority of trials) (Wilcoxon rank sum: loser baseline vs shuffle p=0.10, loser cue vs shuffle p=0.0002, winner baseline vs shuffle p=0.02, winner cue vs shuffle p=0.0002; Wilcoxon sign rank: loser base vs cue p=0.002, winner base vs cue p=0.004).

We hope that emphasizing the following points will provide clarity. First, we rely on a ground truth measure to validate ranks (tube test). Second, we provide a task wherein the degree of hierarchical despotism can be quantified (the degree to which there is a disparity in resources by rank), and ranks can still be represented and expressed when social hierarchies are more egalitarian (as opposed to being more despotic, at the opposite end of the spectrum) – while we do not explore that in this manuscript, this task would lend itself well to investigation of hierarchy structure, and we aim to establish paradigms that can propel the study of social rank forward. Third, a major confound of other dominance tasks is that subordinates do not have enough win trials and therefore it is difficult to compare the neural activity across ranks and wins/losses – so we view the similar number of wins and losses as a strength of our paradigm because it allows us to draw our conclusions with confidence (without concern of a sampling bias). Here, we continue to rely on the robust, well-validated Tube Test while also providing a strategy to examine competitive success on a shorter time scale that is optimized for rigorous statistical comparison of parametric potentially non-stationary) (an

To ensure that our decoding of competitive success using mPFC population data was optimal we took two approaches: We explored how relative social rank and overall competition outcome (who won the majority of rewards) affected the decoding and we also added a new experiment in which we compared mPFC encoding of reward while alone vs during competition, a context that allows the dissociation between receiving the reward and winning.

First, mPFC population activity was able to predict competitive success per trial (win/lose) even when training and testing within relative subordinates or dominants (**Rebuttal Figure 5a-b**). These results suggest that mPFC is not exclusively using baseline rank differences decode competitive success. Next, to understand if competitive success decoding generalized across relative ranks, we asked if a decoder trained on relative dominant (dominant model) data was sufficient to decode competitive success for subordinates,

and vice versa. When testing on data from the opposite relative rank the model trained on dominants was able to generalize and perform above chance, but the model trained on subordinates failed to decode outcome for dominant data sets (**Rebuttal Figure 5a-b**), suggesting an asymmetry associated with rank. It is of note that in both cross-group decoding tests the performance level decreased, suggesting that relative rank influences how competitive success is encoded in mPFC. Finally, to understand how much the outcome of the competition affects the trial by trial competitive success decoding we restricted our data to overall winner data (sessions during which the animal obtained the majority of rewards) or overall loser data (sessions during which the animal obtained a minority of rewards). In these extreme scenarios the mPFC could still decode trial by trial competitive success during the cue (**Rebuttal Figure 5c**). We have now included all these informative iterations of the decoder of competitive success in the manuscript **Extended Data Fig. 8** and brief interpretation in these results lines 174-192.

To rule out the possibility that the competitive success decoding was exclusively a result of responses to the reward and therefore was unrelated to social competition, we performed a new experiment in which the same neurons were recorded during trials alone followed by trials in social competition. To understand how similar or different the mPFC population response was to the reward alone vs winning we asked if an SVM could decode between these two conditions. Indeed, mPFC population activity could decode alone vs win trials with high

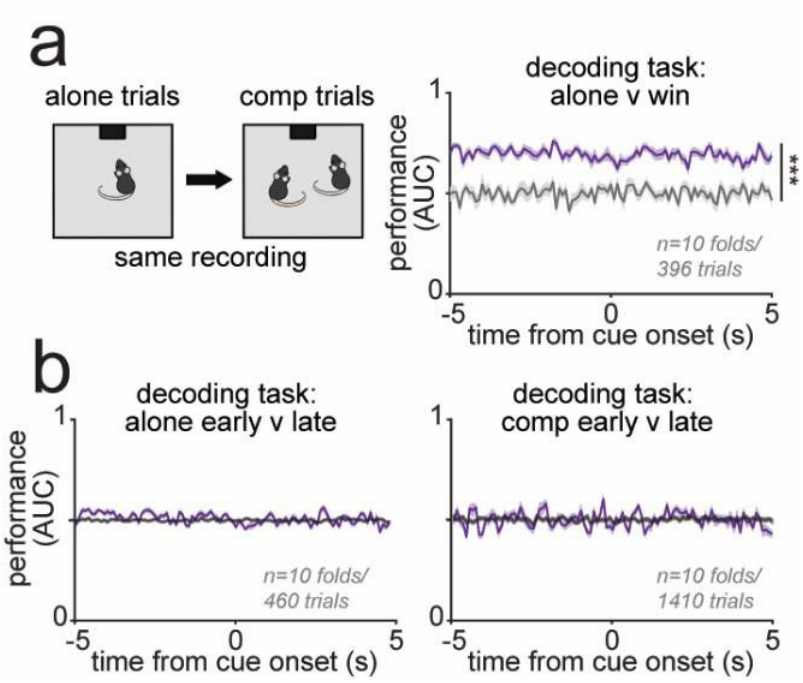

Rebuttal Figure 6: Decoding between winning and receiving reward alone. **a,** Left, the mice were left alone in the behavior cage for the first half with the auditory CS followed by introducing a competitor for the second half. Right, decoding performance for classifying alone trials vs competition win trials (shuffle performance indicated by gray line; mean AUC vs shuffled AUC Wilcoxon rank sum $p=1.8x10^{-4}$). **b,** Decoding performance for classifying whether a trial came from the first half of the session vs the second half for (left) alone trials and (right) competition trials.

accuracy, but early and late trials of the same condition were not decodable (**Rebuttal Figure 6**). These data are now included in **Figure 3** and strengthen our claims that mPFC population dynamics encode competitive success during social competition. Altogether these new analyses and experiments support the robustness of our claim that mPFC population activity decodes competitive success.

The authors could certainly strengthen their claims by comparing data within subject across trials. For example, mice ranked 2-4 could be tested in conditions where they were both the dominant and the submissive. Then they can ask whether decoding of the task relevant variables change within mouse based on

the rank of the partner, or does the neural activity across reflect fixed rank (and not relevant rank). This analysis should require no additional data and it would really add a lot to this study.

We thank the reviewer for this excellent suggestion, which provides a within-subject comparison for interactions with subordinates and dominants – though we did need to perform new experiments and collect substantially more new data to be able to make these comparisons rigorously and with confidence. We have now expanded our manuscript to include several new experiments and analyses to distinguish neural dynamics that are predictive of absolute rank and relative rank. Below is a list of the new analyses and

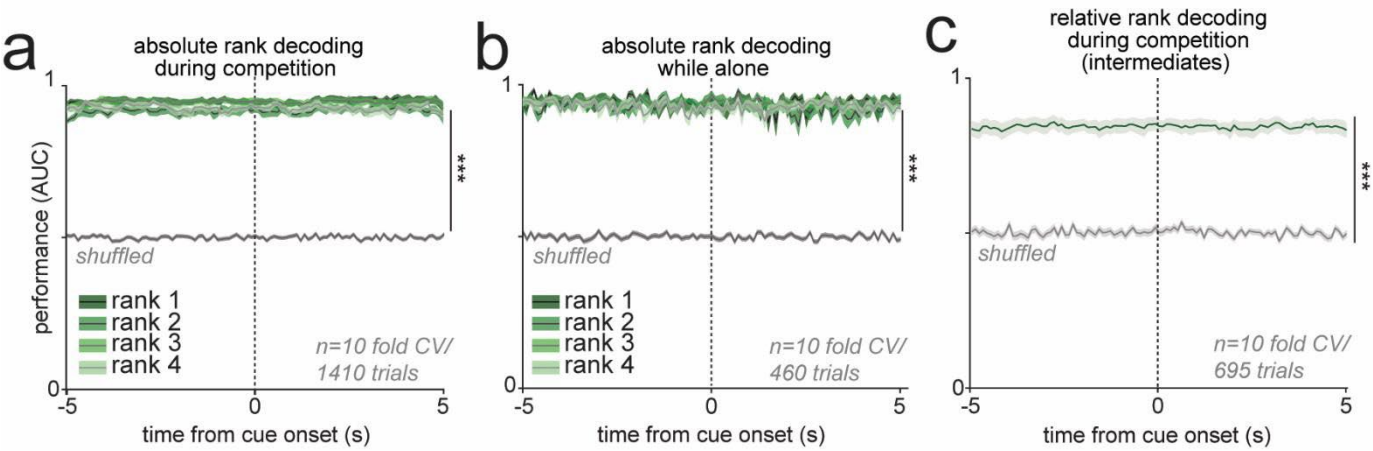

Rebuttal Figure 7: mPFC population dynamics are predictive of absolute rank and relative rank. **a,** Absolute rank can be decoded from mPFC population activity during social competition (performance vs shuffled data; Wilcoxon rank-sum p=0.0002) and **b,** when performing reward task alone (performance vs shuffled data; Wilcoxon rank sum p=0.0002). **c,** Relative rank in intermediate mice (ranks 2 and 3) can be decoded from mPFC population activity during social competition (performance vs shuffled data; Wilcoxon rank sum $p=1.8x10^{-4}$).

conclusions contained in our revised manuscript.

a) Given our previous reports that mPFC activity was predictive of relative rank (**Figure 3** of manuscript) during the reward competition, we asked if this was also true for absolute rank (1,2,3,4 in our cages of 4 mice). We see that mPFC population activity is highly predictive of absolute rank during the social competition (**Rebuttal Figure 7a)**.
b) To understand if the absolute rank code is also there at baseline, in a new experiment we recorded from mPFC single cells while socially ranked mice performed the reward task alone. We see that again mPFC population activity is highly predictive of absolute rank during the reward task alone **(Rebuttal Figure 7b)**.
c) Given that mPFC population activity is predictive of social rank at baseline, we considered how much the relative rank differences observed during social competition are driven by baseline rank differences independent of the competition.
    i. If all population dynamic differences seen during competition were driven by absolute rank differences, then decoding relative rank in intermediates would be impossible. However, we see that when training and testing an SVM decoder using exclusively intermediate rank animals (ranks 2 and 3) as relative dominants or subordinates the prediction of relative rank is highly accurate **(Rebuttal Figure 7c**).
    ii. We explored the possibility that rank differences observed during the reward competition were due to baseline differences. For this, we recorded more than 400 mPFC neurons from 24 mice performing the task in the *alone* condition and compared single cell responses during the reward task alone in rank 1 vs rank 4 animals. We did not observed any statistical differences in the number of responsive cells nor magnitude to the tone nor port entries (**Rebuttal Figure 8**). This suggests that the relative rank differences that emerge during the social competition in

response to winning and port entries are not due to baseline differences in how the mice respond to the reward task. These data are now included in the manuscript in figure 4 and strengthen our claims that the single cell differences observed are due to relative rank differences during the social competition.

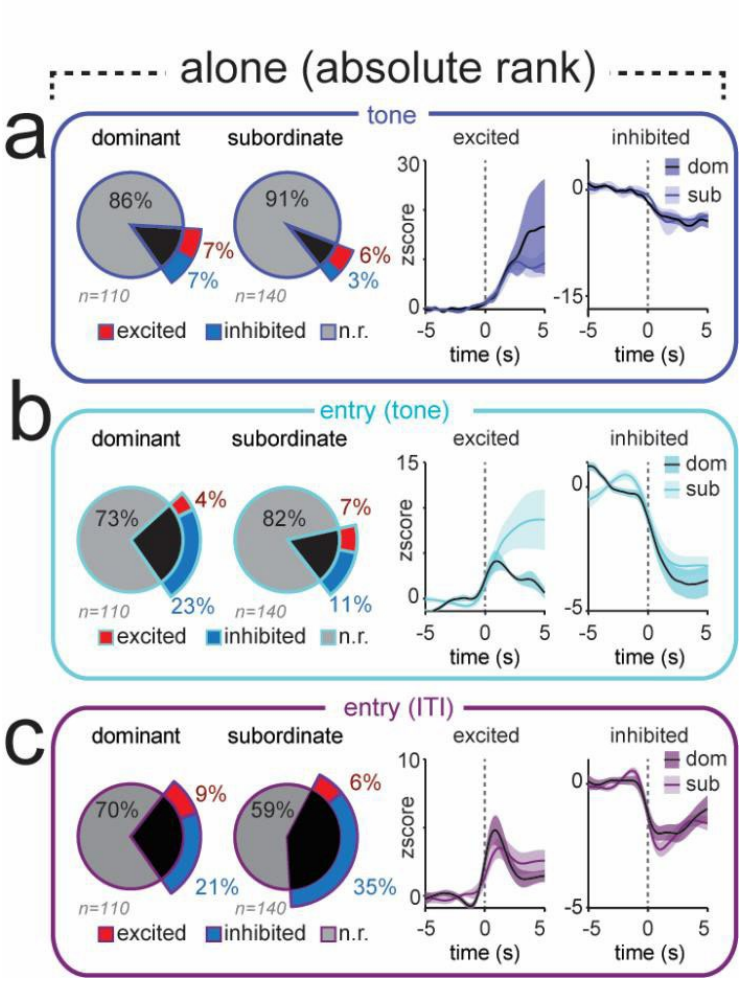

Rebuttal Figure 8: mPFC neurons do not show rank-related differences in general population response profiles to task events when a mouse is performing the reward task alone. a, Number of responsive cells and response magnitude to tone does not differ across absolute *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=8 *dom* inh=8 *sub* exc= 8 *sub* inh=4; Fisher"s exact test, total responsive per group p=0.16; Wilcoxon rank sum for firing rate across groups: exc p=0.87, inh p=1.0) **b**, Number of responsive cells and response magnitude to *port entries* during tone does not differ across *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=5 *dom* inh=25 *sub* exc=9 *sub* inh=16; Fisher"s exact test, total responsive per group p=0.09; Wilcoxon rank sum for firing rate across groups: exc p=0.23, inh p=0.62). **c**, Number of responsive cells and response magnitude to *port entries* during ITI does not differ across *dom* (rank 1) vs *sub* (rank 4) mice(*dom* exc=10, *dom* inh=23 *sub* exc=9 *sub* inh=49; Fisher"s exact test, total responsive per group p=0.06; Wilcoxon rank sum for firing rate across groups: exc p=0.84, inh p=0.17).

Next, the authors analyzed whether single unit data represented task relevant variables in a manner that was different across ranks. I"m a huge fan of modeling, but the goal of models should be to make things explainable. Here, I worry that the authors models make their interpretations more complicated than they need to be. I have two suggestions. The authors apply an approach to cluster cells based on 6 response properties. However, the response properties appear to be different that the task variables used in other parts of the study.This analysis now includes responses to port entry during the inter-trial-interval, and I"m unclear why this response property was included as an input to their joint clustering model.

Our interest in reporting and quantifying single cell responses to these task events, specifically cue onset (separated for win vs lose trials) and port entries (during tone and during ITI, separated for self vs other), comes from observing that at the population level mPFC has significant relative rank differences during these 6 task events (shown in **Figure 3a-d** and **Extended Data Fig. 7d**).

1) Its important to address this issue since two cells which show identical response properties in a given condition (winning) and different responses to another condition could be clustered differently. To this point, the authors go on to compare how the distribution of cells within the clusters differ between submissive and dominant trials. I found the text describing the interpretation of this analysis to be confusing and I spent some significant time with **Extended Data Fig.** 9 trying to sort it out. Ultimately, I could only assume that reporting on the clustering analysis findings 182-185 was suboptimal and did not clearly describe that the model structure (and thus segregation of cells) was based on joint cellular properties. For example, in principle a cluster is defined by responses to condition A, B, and E, BUT NOT C, D and F. A cell that shows response to A, B, and F only could fall in a different cluster. Thus, 2) it is necessary to fully describe what is represented by a cluster, then one could perform post-hoc analysis. I found the subsequent analysis of individual response properties to be clearer. In fact, since the goal of the authors was to determine if single unit data represented task relevant variables in a manner that was different across ranks, I wondered why the complicated clustering model was even useful. As a general rule of thumb, a model should only be as complex as is necessary to explain the phenomenon of interest. Thus, in my opinion, the analysis in 4a-c should be removed. The analysis shown in 4f-I should expanded to include winning and losing (Extended Fig 9c). Discussion of these results should be simple and clear. If the authors choose to keep the analysis shown in 4a, a clearer description of the meaning of joint response properties and its importance in social rank should be discussed. For example, Clusters 2, 3, and 4 show strong excitation to winning. Though sub and dom show differences in % cluster 4 (Fig. 4e), no differences are observed in % of cells showing winning or losing (Fig. E9c). Thus, the authors would need to explain all of this more clearly if they intended to keep the complicated model (which I, again, argue is unnecessary for the goal of their study).

We appreciate the reviewer pointing out the unnecessary perceived complexity of doing the unsupervised clustering with the combined responses to these 6 task events. In response to that we have removed our claims regarding the comparison of clusters obtained with this unsupervised method and use the approach exclusively for visualization. As the reviewer suggested, we now report the number of significant cells and magnitude differences for these 6 events in the manuscript in main figure 4 together with new data of mPFC recordings during the reward task while mice are alone.

However, we feel that the way we are performing our analysis is such that "two cells which show identical response properties in a given condition (winning) and different responses to another condition could be clustered differently", which we would consider a strength, because it provides an additional layer of granularity (**Figure 4a**), as we also provide analyses to examine each task event in isolation (**Figure 4b-g**), which directly addresses the reviewer"s concern.

Next the authors explore whether PFC->LH neurons contributed to the code for social rank. The authors chosen analysis pathway could be simplified substantially. In 5c-d they show neural responses to the 6 conditions outlined in their prior analysis. Then they show that a greater portion of PFC->LH neurons show responses to port entry during the inter trial interval. Did the authors test neural responses for all 6 behavioral

conditions? Did they compare the three cell types? I"m assuming the answer is yes, but the authors didn"t correct for multiple comparisons (where is the statistics. N for PFC->BLA seems to be missing for legend 5e). Ifnot, why did they only choose a limited number of comparisons. Differences were only observed in the self-entry for the ITI. How should this be interpreted?

The reviewer makes an excellent point about the difficulty of interpreting the specificity of this difference. We have removed our claims given the difficulty in interpreting the specificity. Originally, we tested only the proportion of cells responding to port entry behavior because those events are the only ones that showed rank-dependent differences in proportion of cells (manuscript **Figure 4**). Given that it is difficult to interpret the specificity of why only ITI self-entries are affected we removed the claims of that data reported. Now given that we have the new *alone condition* controls we focus the main analysis on tone responses vs win trials during competition – a comparison which provides social context specificity (**Rebuttal Figure 9**).

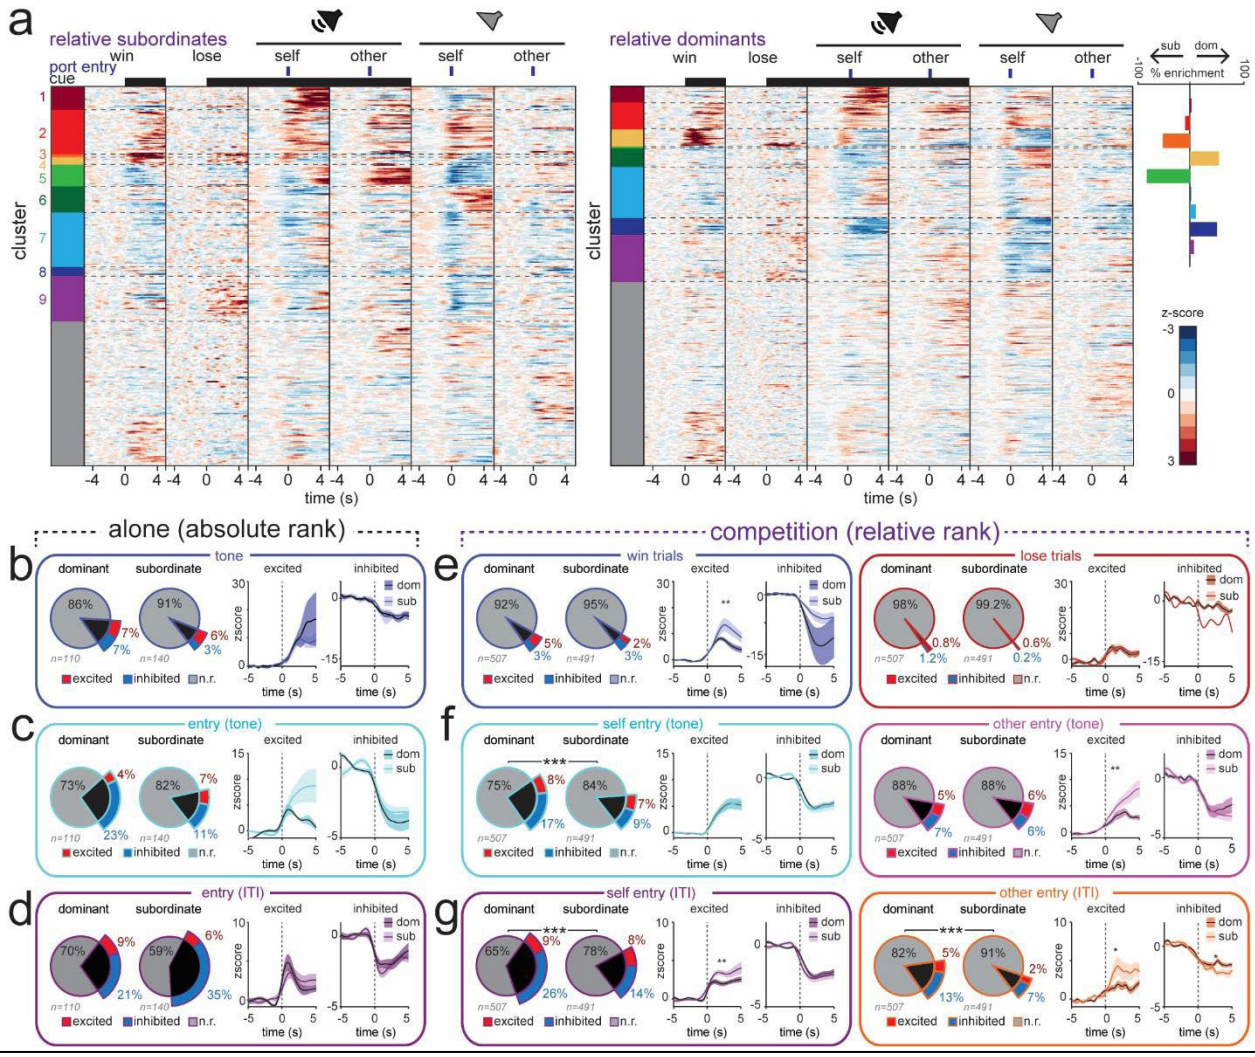

Manuscript Figure 4: Relative dominants have more reward seeking behavior cells while subordinates have larger responses to competitor behavior. **a,** Heatmaps (left=subordinate; right=dominant) for firing rate responses to the six task-relevant events during the reward competition. Color bar indicates functional clusters obtained by hierarchical agglomerative clustering. Clustering was performed with all cells, but only responsive cells are shown in the heatmap (mean firing rate was larger than 2 or smaller than -1 z-scores to any behavioral event). On the right, Percent difference between relative dominant (*dom)* and subordinate (*sub)* cells (% enrichment) in functional cluster membership showed differences across ranks. **b,** mPFC tone responsive cells when mice perform the reward task alone. Number of responsive cells and response magnitude to the tone does not differ across absolute *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=8 *dom* inh=8 *sub* exc= 8 *sub* inh=4; Fisher"s exact test, total responsive per group p=0.16; Wilcoxon rank sum across groups: exc p=0.87, inh p=1.0). **c,** mPFC tone *port entries* responsive cells when mice perform the reward task alone. Number of responsive cells and response magnitude to *port entries* during tone does not differ across *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=5 *dom* inh=25 *sub* exc=9 *sub* inh=16; Fisher"s exact test, total responsive per group p=0.09; Wilcoxon rank sum across groups: exc p=0.23, inh p=0.62). **d,** mPFC inter trial interval (ITI) *port entries* responsive cells when mice perform the reward task alone. Number of responsive cells and response magnitude to *port entries* during ITI does not differ across *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=10, *dom* inh=23 *sub* exc=9 *sub* inh=49; Fisher"s exact test, total responsive per group p=0.06; Wilcoxon rank sum across groups: exc p=0.84, inh p=0.17). **e,** mPFC tone responsive cells during *win* vs *lose* trials in the social competition task. Left, number of responsive cells did not differ by relative rank while response magnitude for *win* trials differed (*dom* exc=24 *dom* inh=14 *sub* exc=12 *sub* inh=16; Fisher"s exact test, total responsive cells per group p=0.30; Wilcoxon rank-sum across groups: excited p=0.01; inhibited p=0.06). Right, number of responsive cells and response magnitude for *lose* trials did not differ by relative rank (*dom* exc=4, *dom* inh=6, *sub* exc=3, *sub* inh=1; Fisher"s exact test, total responsive per group p=0.17; Wilcoxon rank-sum: excited p=0.62; inhibited p=0.28). **f,** mPFC *self* vs *other port entry* responsive cells during tones in the social competition task. Left, number of responsive cells was higher for relative dominants while there was no firing rate magnitude difference (*dom* exc=40 *dom* inh=87, *sub* exc=33 *sub* inh 44; Fisher"s exact test, total responsive cells per group p=0.0002; Wilcoxon rank-sum: excitation p=0.28 and inhibition p=0.99). Right, there was no relative rank difference in number of responsive cells to *other port entries* during the tone, while the excitation magnitude was higher for relative subordinates (*dom* exc=25 *dom* inh=37 *sub* exc=30 *sub* inh=29; Fisher"s exact test, total responsive cells per group p=0.84; Wilcoxon rank-sum: excitation p=0.006, inhibition p=0.11). **g,** mPFC self vs other port entry responsive cells during ITI periods in the social competition task. Left, relative dominants had more responsive cells to *self port entries* during the ITI while relative subordinates had larger excitation magnitude (*dom* exc=45 *dom* inh=132 *sub* exc=38 *sub* inh=71, Fisher"s exact test, $p=9.8x10^{-6}$; Wilcoxon rank-sum: excitation p=0.006 and inhibition p=0.28). Right, relative dominants had more responsive cells to *other port entries* during ITI while subordinates had larger magnitudes (*dom* exc=25 *dom* inh=64 *sub* exc=13 *sub* inh=33; Fisher"s exact test, p=0.00019; Wilcoxon rank-sum: excitation p=0.015 and inhibition p=0.04). Data presented comes from recordings from 10 mice (rank 1; n=5 and rank 4; n=5) for the alone condition and 20 mice for the competition condition.

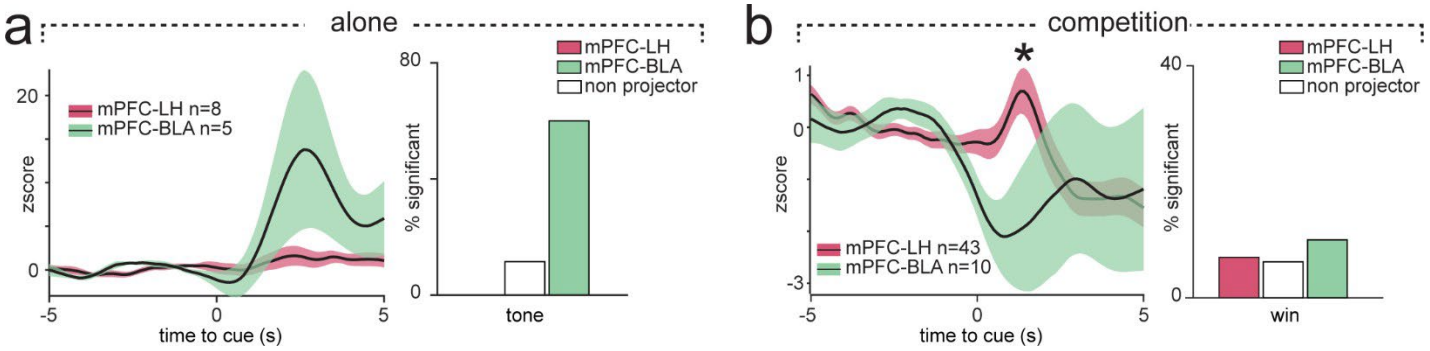

Rebuttal Figure 9: mPFC-LH neurons are more responsive to the winning trials compared to mPFC-BLA neurons during reward competition. **a,** Right, firing rate of mPFC-LH is higher than mPFC-BLA during the reward delivery period in *win* trials in the competition (mPFC-BLA n=10 from 3 mice, mPFC-LH n=43 from 3 mice, Wilcoxon rank sum p=0.045). Left, percent cells responding to tone during *win* trials in the social competition (mPFC-LH n= 3/43, mPFC-BLA n=1/10 and non-phototagged n=57/620). **b,** Left, firing rate for projector populations during tones while animals performed the reward task alone (alone vs comp mean zscore during tone; mPFC-LH p= 0.043, mPFC-BLA p=0.049). Right, percent cells responding to the reward predictive tone when animal is alone (LH=0/8 from 3 mice, BLA 4/5 from 3 mice, non phototagged=54/470).

If the goal is to show that PFC->LH cells encode rank, why does it matter whether they respond to behavioral variables in the task. If the goal is to show that they carry unique information, why is the information in Fig 5e even necessary. Overall, the analysis in 5f-g seems more appropriate. Here, the strategy is to subtract the cells from the model, and determine if the model prediction improves. But several details are missing here. First of all, what is the classifier being trained on (i.e., what behavior condition)? Next, the comparison between the BLA and LH projecting neurons is not entirely warranted: a different number of cells are recorded for the two pathways, and it"s unclear how many mice that cellular activity was acquired from (important information because only 10 BLA neurons were recorded).

The reviewer makes a reasonable point that the number of BLA projector neurons is very low. In 8 new mice, we injected two viruses, implanted electrodes in mPFC, tube tested mice, trained them for the reward competition assay, recorded during competitions, spike sorted and analyzed the data, a process that took 3 months, in an attempt to record from additional mPFC-BLA projectors. Unfortunately, despite our best efforts to perform this experiment (we added several new authors, working alongside original authors, for several months), and despite good histological expression of the virus (**Rebuttal Figure 10**) and high neuronal yield across mice, we failed to collect recordings from any neurons that qualified as mPFC-BLA projector cells given our rigorous criteria for classifying a neuron as “phototagged” or “photoidentified.” Given the remaining disparity in number of projector populations we agree that the comparison of removing 10 cells vs 43 cells is not warranted, so we removed this analysis from the manuscript and instead opted to report the firing rate and responsiveness of these subpopulations of cells in alone vs competition conditions (**Rebuttal Figure 9**).

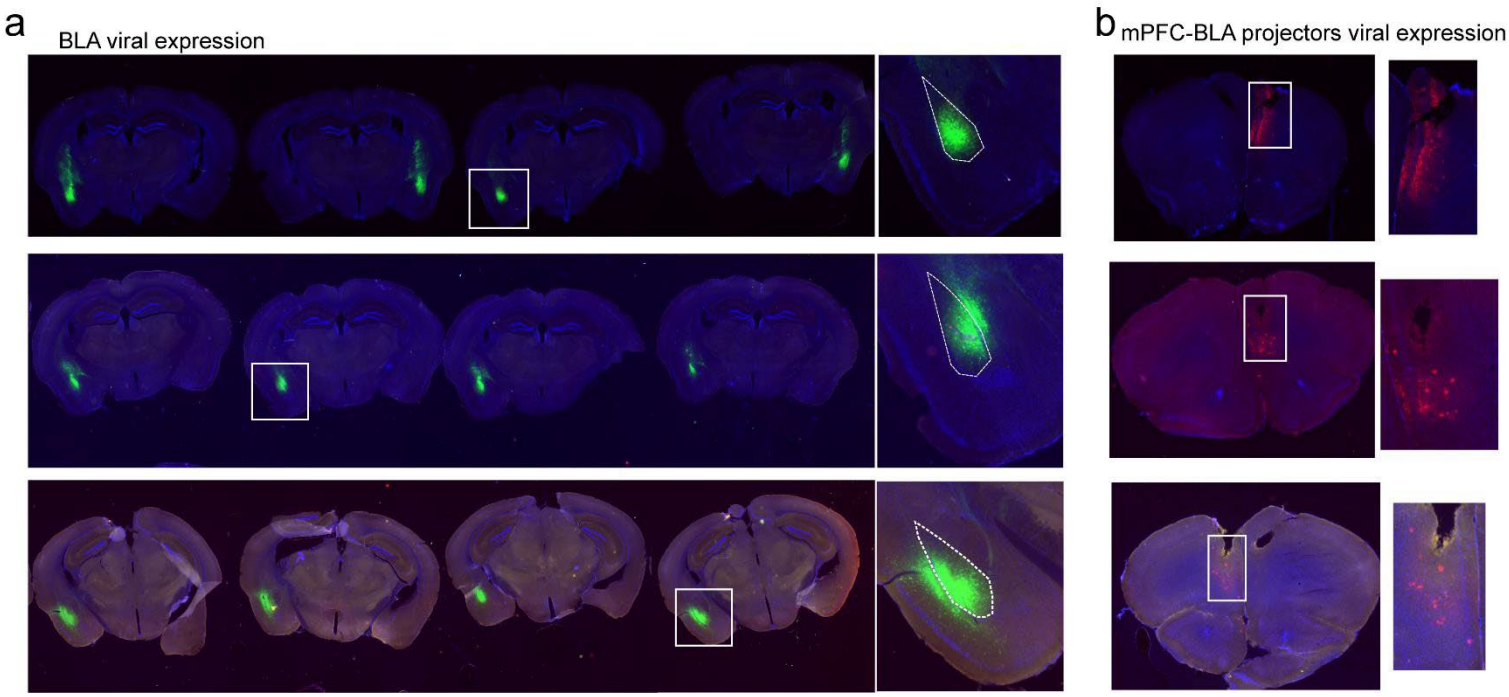

Rebuttal Figure 10: Example successful viral expression in BLA and mPFC for recording additional mPFC-BLA neurons. a, Three example animals showing successful expression of CAV-Cre mixed with eYFP in the BLA (images taken at 4x). Insets show closer look of the selected BLA image indicated with white square. **b,** Example images for mPFC for the corresponding animals shown in b. showing expression of DIO-Chrimson-tdTomato in mPFC neurons that project to BLA.

Minor comments:

There is some unusual language in the methods description of the machine learning methods. In the Support Vector Machines, it is noted that the SVM generates a likelihood. It does not, as it is not a probabilistic model. Additionally, the terminology of a “one-hot encoding” is not the standard usage of that term (which usually means that each of L classes is encoded as a 1 on the lth entry of an L length vector), but the one-vs-all approach seems fine. However, details are lacking here: what parameters were tuned? How was the model selection run? What method was used to deal with class imbalance? The definition of the AUC metric should also be given, since there are differing definitions in the multi-class setup. I assume that this is the average of the one-vs-all AUCs, but that needs to be stated. Also, an SVM looks at all data points while choosing the support vectors, so that language is unclear.

The reviewer is right to point out that an SVM is inherently not a probabilistic model. In order to make the SVM probabilistic we used an appropriate score-to-posterior-probability transformation function, as proposed by Platt in 1999, which is a widely adopted approach (Platt, 1999). If the classes are inseparable, the transformation function is a sigmoid function, and if they are perfectly separable, it"s a step function. We thank the reviewer for catching the one-hot typo which we have now corrected. We utilized a one-vs-all approach to decode the 9 types of behavioral classes. For the SVM, we used a gaussian kernel and default model hyperparameters as defined in MATLAB 2019B. Since we"re reporting the area under the receiver operating characteristic curve (AUC-ROC), the performance metric is already sensitive to class imbalance. We have also tried to balance class sizes wherever possible. This information is now reflected in the ***methods section 14*** Electrophysiology data analysis, under the subsection titled SVM classifier.

The updates on the EM steps are complex, so the communication of this would be enhanced by the authors highlighting the differences in their EM algorithm in comparison to the Escola et al paper.

The differences in the Expectation-Maximization (EM) algorithm when compared to the Escola paper (Escola et al., 2011b) are now reported in the ***methods section 17*** "Decoding of behavior" under the subsection for HMM-GLM and added here for convenience:

The model formulation and EM training algorithm are similar to Escola et al. The primary differences are: 1) the input to the transition and state GLMs for the current time point is neural activity from the immediately previous time point only and not further back, 2) the GLM link function being a sigmoid/softmax has different gradients and dynamics as compared to the Poisson, 3) The outputs are categorical behaviors in our case instead of the binned spike counts and 4) we utilize an iterative optimization technique since the second order derivatives in our case are highly complex and computationally inefficient to use.

We also added the following language: Escola et al. first introduced the idea of using a Hidden Markov Model (HMM) for modeling the change in the internal state of the organism as transitions of „hidden states" and the neural spikes as emissions obtained from the hidden states. Escola and colleagues used Poisson GLMs for computing probabilities of each transition between hidden states as well as for determining emissions of neural spikes (Escola et al., 2011b). This approach provided a mechanism where both the state transition probabilitiesas well as neural spike emission probabilities varied with time. Calhoun and colleagues proposed a modification of the HMM-GLM in which both the input and output were behavioral states (Calhoun et al., 2019b), using a logit GLM for the transition and emission probabilities instead of Poisson GLMs used by Escola. In our study, we use a similar model as Calhoun, where the input to the GLM layers is the mPFC neural activity and the emissions are the behaviors of the mice.

The communication behind the figures can be improved by focusing on terminology. It was super hard to track at time. Winning/loosing seems to be event locked to the tone. Reward is event locked to port entry (but these trials are referred to as tone). ITI is event locked to port entry (but is referred to as ITI). Perhaps something simple like competition start (tone), competition end (port entry), and ITI for the blocks within trial.

We appreciate the reviewer expressing the nomenclature and terminology is confusing. We have added some language in the results and figure legends to clarify the terminology. Indeed win and lose trials are aligned to the cue/tone onset. On the other hand self and other entry events are aligned to the port entry timepoint. Port entries occur either during the tone or during the ITI period, even in the absent of the tone as this is a strategic behavior, so we distinguish between tone port entries and ITI port entries. In lines 151-154 we now specify the alignment to cue vs port entry: "Neural trajectories during the cue for *win* or *lose* trials occupied segregated PCA subspaces – even before the cue onset (**Fig. 3a; Extended Data Fig. 7c**). Similarly, neural trajectories aligned to port entries of the self or other (i.e. competitor) were segregated in the PCA subspaces…"

In the legends for **Fig. 3a-b** and **Extended Data Fig. 7** we have added the following language respectively:
"*Win* and *lose* trials are aligned to the cue onset"
"Self entry events are aligned to port entries of the subject mouse while other entry events are aligned to the competitor"s port entries"
"Self entry events are aligned to port entries of the subject mouse while other entry events are aligned to the competitor"s port entries. ITI port entries refer to port entries that occurred outside of the tone period."

In general, the authors show adjust their training and testing strategy such that testing is performed using a trials from a hold-out with N-1 fold cross validation.

All of our models have a 10-fold cross-validation such that they are tested on data that was held out from the training set. We report this information in the corresponding figure subpanels and legends in **Figures 2-3** and ***methods sections 14 and 17***.

In summary: The study by Padilla-Coreano explored the important question of the role of PFC in social hierarchy. A primary issue that needs to be addressed is whether the cohort of animals used for electrophysiological analysis actually show their primary competitive outcome measure based on rank. This may just be an issue of simply reporting experimental details (N, and needing a few more mice). Many

suggestions for improving communication/clarity are provided to the authors; nevertheless, it is important to clearly state that much of the presented results and undergirding ML analyses are unnecessary to support the main theme of the manuscript. My suggestion would be to simplify the manuscript, and to focus the presentation of the findings towards the unit encoding of social preference in PFC, and the role of the PFC-LH circuit in this behavior. Especially considering the prior discoveries outlined by Zhou et al, projection specific findings is the knowledge advance I found most exciting. If these broader contributions (AlphaTracker, HiddenStates, MUA modeling, single cell clustering analysis approach) remain in the manuscript, substantial revisions to the communication of these findings and the analyses methodology are warranted as noted above to enable a clearer evaluation of rigor and robustness.

Referee #2 (Remarks to the Author):

In this study from Padilla-Coreano et al., in the Tye laboratory, the team presents research and computational efforts to examine the neurobiological basis for social dominance - a key feature of aggression behavior. They examine the role of the media prefrontal cortex (mPFC) because some prior evidence has suggested that in mammalian species this structure plays a critical role in determining social rank. They go on to develop a trial-based social competition assay, alongside using a custom machine learning approach they developed (Alpha-tracker). Using these two former approaches alongside electrophysiological measures (using a newer wireless method to better facilitate social interaction), they identify a unique behavioral states (9 in total, they claim), which can be decoded from mPFC ensemble activity - predicting social rank and competition winners and losers. Furthermore, they reveal that mPFC to LH projections are better predictions of behavior than an alternative circuit. Overall, this is a very thorough, elegant, and exciting development for the field of behavioral neuroscience, ethnologically relevant social behavior, and computational efforts in the field. It provides a potentially new (or perhaps just a bit better) method for identifying critical behavioral states in interacting/"socializing" mice , while also methods for decoding unique behavioral and neuronal ensembles in tandem within a specific circuit.

We are gratified that this reviewer appreciates a number of the innovations and features of this study that we have efforted to establish. However, we do wish to clarify that this line of investigation is focused on social rank, rather than aggression (and in our hands, aggression and dominance do not always correlate). Here, we are examining how different ranks within a stable, established hierarchy are representing competitive success and both absolute and relative rank.

However, in spite of this high enthusiasm for the provocative nature of the work, the potential for extensions of this work into other domains, behaviors, and neural circuits, there were some concerns. These primarily centered around controls and/or additional data analysis that would better substantiate their conclusions; along with some stronger rationale provided for particular experiments chosen over others, including the focus on mPFC to LH, vs a whole host of other possible PFC-related circuits which could have been examined. Below is a list of major concerns (those requiring new analysis, experimental data, or major text discussion/rationale), and then minor comments or concerns follow, that are of less consequence, and more suggestive in nature.

Major:

1) The biggest control that I see which is missing here, is to show data and the entire processing pipeline described (behavior, ensembles, and circuit specificity) in animals that engage in the task alone, and/or with the scent of more dominant animals nearby (the later not being as critical, just an interesting experiment). The "alone control" is important in particular for the trial-based task, that would allow for cross comparisons of activity while engaging in a social dominance bout from the activity and behavior which is independent and is reward - related or orthogonal to the social dominance behavior itself.

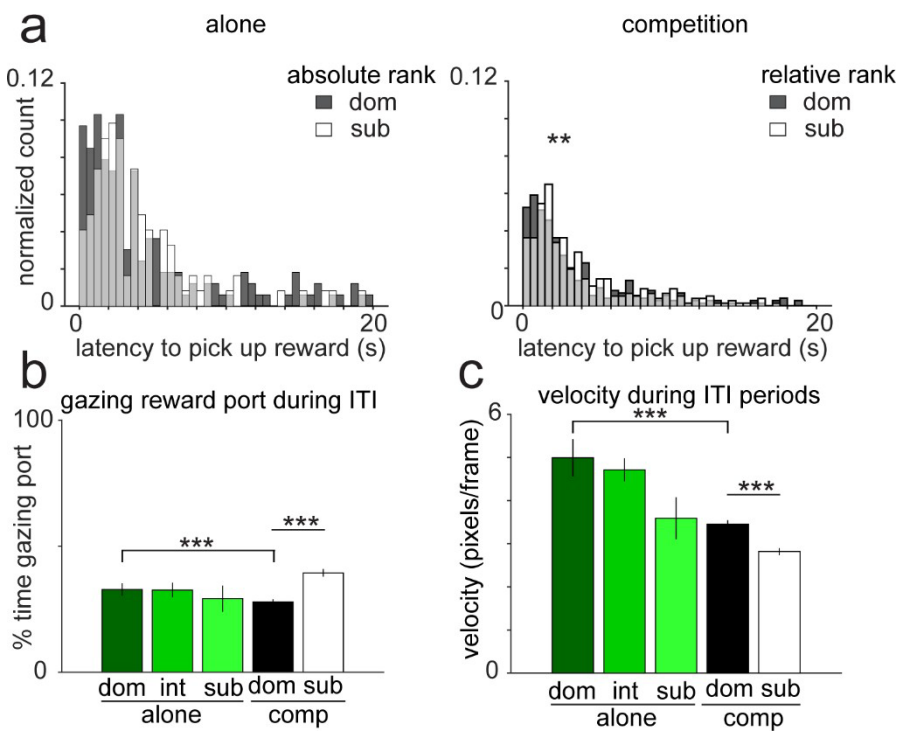

Rebuttal Figure 11: Behavioral differences between alone and competition condition. a, Left, while performing the reward task alone there was no difference in latency to pick up reward between dominant (rank 1) and subordinate (rank 4) mice across trials (rank 1 n=165, rank 4 n=122, Kolmogorov-Smirnov 2 sample test p=0.46). Right, during the social competition relative subordinate animals increased the latency to pick up the reward (*dom trials* n=326, *sub trials* n=358, Two-sample Kolmogorov-Smirnov test, *dom* vs *sub trials* p=0.015; One way RM-ANOVA $F_{(1,24)}$=2.06, p=0.002). **b,** In a subset of videos we corrected identity errors to quantify behavioral metrics through the session. Gazing to the reward during intertrial interval (ITI) periods increased during social competition most prominently in relative subordinate mice (alone trials: dom/rank 1= 27, int/ranks 2&3=55, sub/rank 4=17; competition: relative dom=706 and relative sub=706; ***p<.001 Wilcoxon rank sum comparison). **c,** Velocity during ITI periods decreased during social competition most prominently in relative subordinate mice (alone trials: dom/rank 1= 27, int/ranks 2&3=55, sub/rank 4=17; competition: relative dom=706 and relative sub=706;

The reviewer makes an excellent point. To address these missing controls, we have now included data from several new experiments where animals engage in the reward task alone (Pavlovian approach conditioning). We see behavioral differences in the animals performing the task alone vs in social competition. Interestingly, our results suggest that the subordinate mice are more sensitive to the competition condition, as they show larger changes in latency to pick up reward, velocity and gazing towards the reward port (**Rebuttal Figure 11).** These behavioral metrics provide additional evidence for the social-context dependency of the animal‟s behavior during this novel task.

Next, we explored the differences in mPFC population dynamics in the alone condition vs the social competition. We found that subordinates had the largest changes in population dynamics comparing tone responses alone vs in competition (**Rebuttal Figure 12**). These data are now included in figure 3 and strengthen our claims that the mPFC population changes we observe relate to relative rank differences during the social competition.

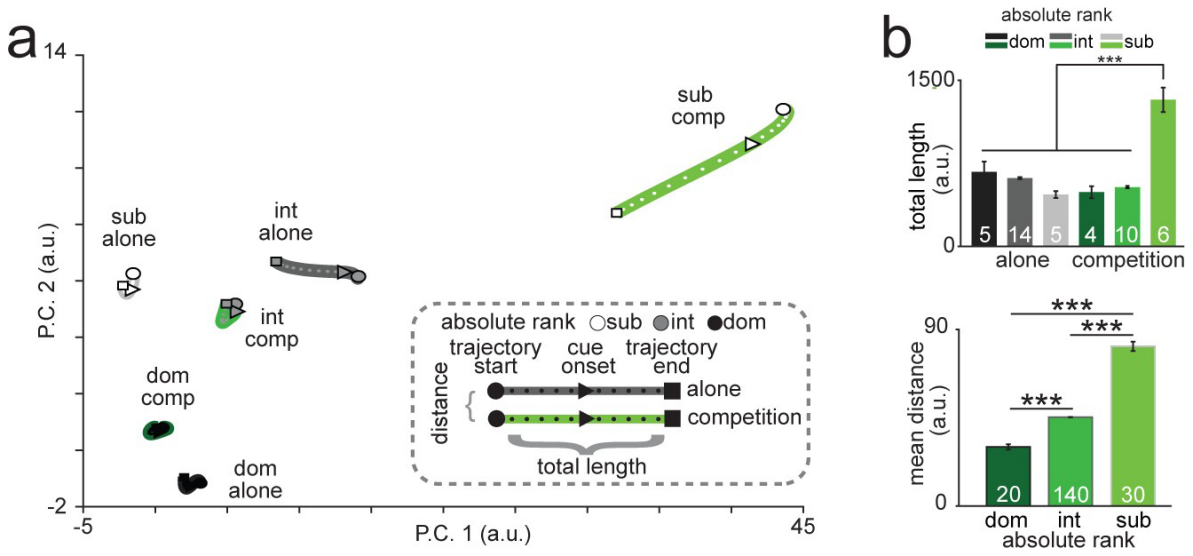

Rebuttal Figure 12: Subordinates mPFC population dynamics are most affect by the social competition. a, Neural trajectories of mPFC population firing rate by absolute ranks (*dom*=rank 1; *int*=ranks 2 & 3; *sub*=rank 4) when performing the reward task alone vs in competition in a lower dimensional common principal component (PC) *sub*-space (include neurons alone: *dom*=111, *int*=259, *sub*=140; competition *dom*=309, *int*=359, *sub*=330). **b**, Top, trajectory lengths (using PCs that captured 90% of variance) during the tone is higher for subordinates only during competition trials (2-way ANOVA main effect of rank $F_{(2,38)}$=30.4, p=1x10^{-8}, task $F_{(1,38)}$=26.1, p=9x10^{-6} and interaction $F_{(2,38)}$=70.1, p=1x10^{-13}). Bottom, distances between alone and competition trajectories during the tone increases with rank (n reflects all possible combinations of trajectories across iterations;1-way ANOVA main effect of rank $F_{(2,187)}$=536, p=3x10^{-78}).

To rule out the possibility that the competitive success decoding was exclusively a result of responses to the reward and not related to social competition, we performed a new experiment in which the same neurons were recorded during trials alone followed by trials in social competition. To understand how similar or different the mPFC population response was to the reward alone vs winning we asked if an SVM could decode between these two conditions. Indeed, mPFC population activity could decode alone vs win trials with high accuracy, but early and late trials within the same condition (early vs late tones during alone or early vs late win trials) were not decodable (**Rebuttal Figure 6**). These data are now included in figure 3 and strengthen our claims that mPFC population dynamics encode competitive success during social competition.

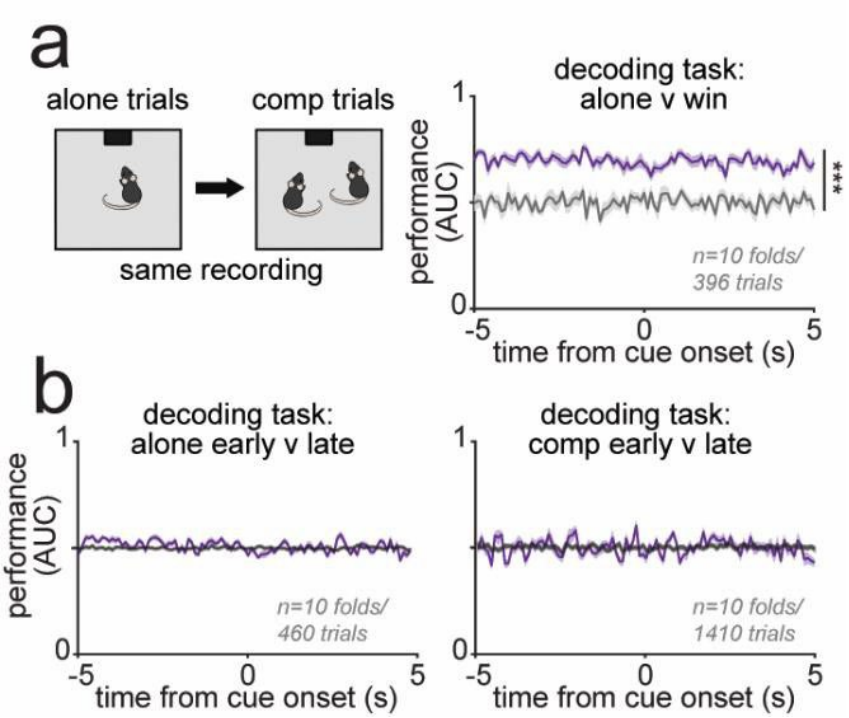

Rebuttal Figure 6: Decoding winning vs receiving the reward alone. **a**, Left, the mice were left alone in the behavior cage for the first half with the auditory CS followed by introducing a competitor for the second half. Right, decoding performance for classifying alone trials vs competition win trials (shuffle performance indicatedby gray line; mean AUC vs shuffled AUC Wilcoxon rank sum $p=1.8x10^{-4}$). **b**, Decoding performance for classifying whether a trial came from the first half of the session vs the second half for (left) alone trials and (right)competition trials.

Next, by comparing the mPFC ensemble dynamics when the mice perform the task alone vs in competition we addressed if any ensemble differences were due to baseline differences. For this, we recorded more than 400 mPFC neurons from 23 mice performing the task alone and compared single cell responses during the reward task alone in rank 1 vs rank 4 animals. We did not observe any statistical differences in the number of responsive cells nor magnitude to the tone nor port entries (**Rebuttal Figure 8**). This suggests that the relative rank differences that emerge during the social competition in response to winning and port entries are not due to baseline differences in how the mice respond to the reward task. These data are now included in the manuscript in figure 4 and strengthen our claims that the single cell differences observed are due to relative rank differences during the social competition.

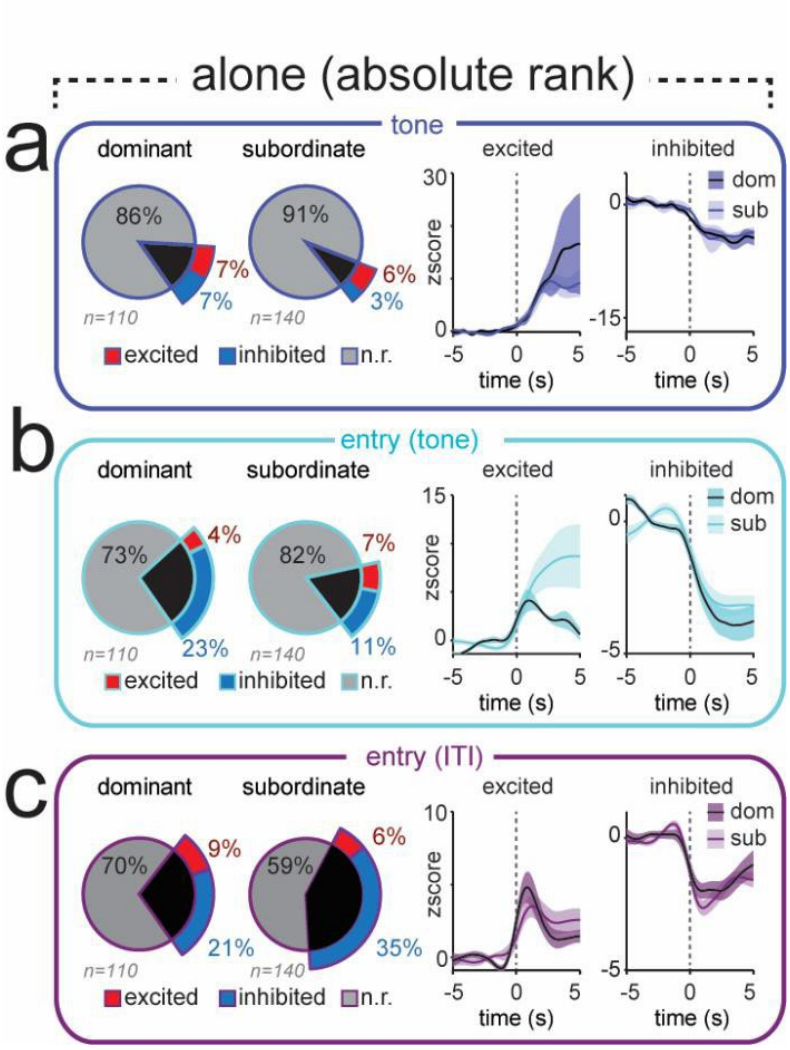

Rebuttal Figure 8: mPFC neurons do not show social rank differences when the mouse is performing the reward task alone. **a**, Number of responsive cells and response magnitude to tone does not differ across absolute *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=8 *dom* inh=8 *sub* exc= 8 *sub* inh=4; Fisher"s exact test, total responsive per group p=0.16; Wilcoxon rank sum for firing rate across groups: exc p=0.87, inh p=1.0) **b**, Number of responsive cells and response magnitude to *port entries* during tone does not differ across *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=5 *dom* inh=25 *sub* exc=9 *sub* inh=16; Fisher"s exact test, total responsive per group p=0.09; Wilcoxon rank sum for firing rate across groups: exc p=0.23, inh p=0.62). **c**, Number of responsive cells and response magnitude to *port entries* during ITI does not differ across *dom* (rank 1) vs *sub* (rank 4) mice (*dom* exc=10, *dom* inh=23 *sub* exc=9 *sub* inh=49; Fisher"s exact test, total responsive per group p=0.06; Wilcoxon rank sum for firing rate across groups: exc p=0.84, inh p=0.17).

Finally, we have also added new phototagging data from mice performing the reward task alone. In a subset of animals we phototagged mPFC-LH or mPFC-BLA projector neurons after mice performed the reward task alone. We see that none of the mPFC-LH neurons significantly responded to the reward-signaling tone when mice are alone while mPFC-BLA neurons were highly responsive to the reward-signaling tone in the alone condition. On the other hand, during competition mPFC-BLA neurons were less responsive to the tone during win trials and mPFC-LH neurons firing rate was significantly higher that mPFC-BLA firing (**Rebuttal Figure 9**).

Altogether these new experiments and analyses control for the reward processes that do not relate to social competition and strengthen the claims of our paper.

1a) This is a sub-point of question 1, which can be more tightly controlled in the analysis, because the alone mouse can also account for starting positions and general posture of the animals in that regard. (which relates a bit to the point 2). The authors might just use their algorithms to look at other behaviors that would perhaps provide some clues to positioning (some ideas, approach, avoidance, orientation relative to mouse vs target reward, etc).

As suggested by the reviewer, we have quantified position-associated behaviors using the AlphaTracker output. Doing so was laborious as it necessitated correcting social identity errors which albeit are very low (<1% of frames) require close and careful correction. We indeed observe several see rank related differences, with relative subordinates having lower velocities and gazing the reward port more than the dominant competitors (**Rebuttal Figure 13**).

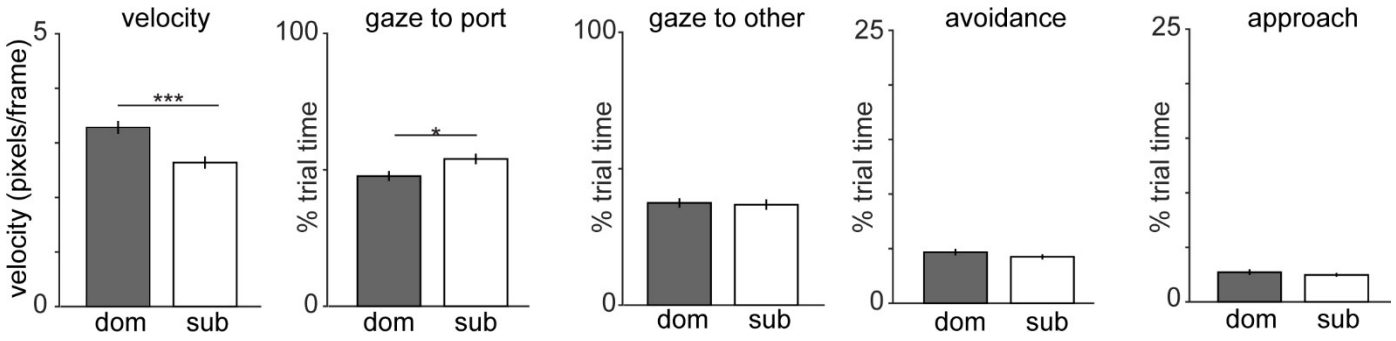

Rebuttal Figure 13: Behavioral metrics during the social competition. Average velocity, time gazing to the port, time gazing to the other (competitor), time avoiding competitor (as defined by distance increasing and angles between mice being higher than 120 degrees) and time approaching competitor as defined by distance decreasing and angles between mice being lower than 60 degrees). Wilcoxon rank sum test ***p<.001, *p<0.05 n=359 trials.

2) The rationale for the tube-test isn't well defined as currently presented. Is social dominance encoded at the neuronal level or is it just dictated by which mouse is bigger in the tube, and who's stronger? I know it is well established, but since it is used a priori here, it is a more germane to the conclusions drawn. Conceptually both make sense, but the tube model of identification in this case leaves one wondering if the data are preselected

and therefore biased in some way because of it. Can they shuffle the data in a way that accounts for differences in rank that the tube test doesn't account for? There are multiple social defeat behaviors and postures that can be easily quantified in social aggression and dominance interactions. Why were none of those examined here? They are very granular in nature and would provide rich confirmation and validation in their behavior tracker machine learning algorithms.

We now provide shuffled controls for our decoder models for competitive success and relative rank shown in **Fig. 3i** (also shown in **Rebuttal Figure 14**). In addition, in **Extended Data Fig. 7h-k**, to control for social identity we separated our data into two groups and quantified the differences in mPFC population dynamics and found that the trajectory lengths, which were longer for subordinates, did not differ across groups (**Rebuttal Figure 15**).

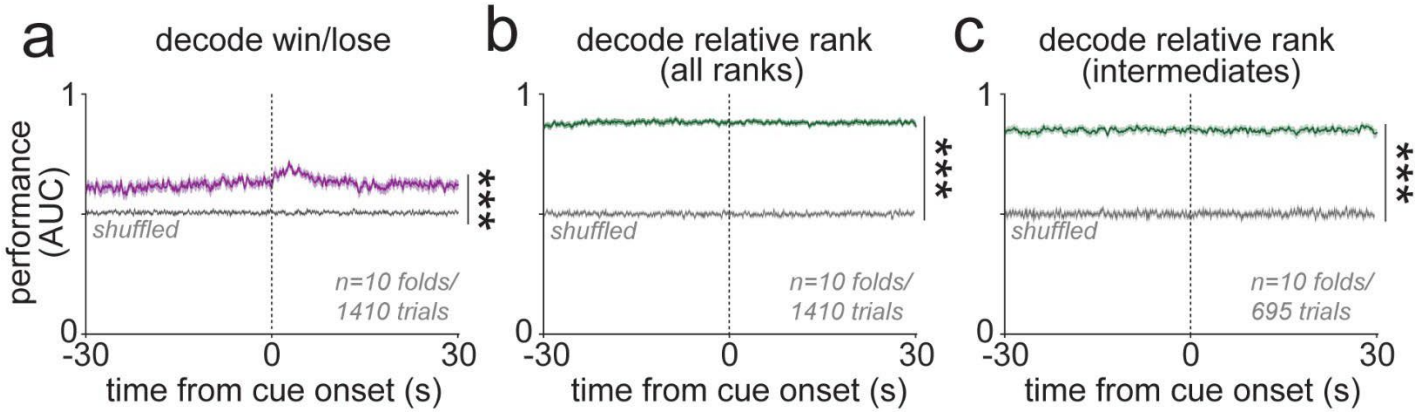

Rebuttal Figure 14: SVM performance is higher than chance for decoding (**a**) competitive success and (**b**) relative rank for all mice and specifically for (**c)** intermediate mice (area under the receiving operating curve: AUC, gray indicates performance when shuffling data; mean AUC vs shuffled AUC Wilcoxon rank sum: competitive success $p=1.8x10^{-4}$, relative rank all mice $p=1.8x10^{-4}$ and relative rank for intermediates $p=1.8x10^{-4}$).

Rebuttal Figure 15: mPFC population dynamics are predictive of relative rank in intermediate mice. a, Neural trajectory lengths (using principal components that captured 90% of variance) for *win* (left) and *lose* trials (right) for intermediate mice (ranks 2&3) are higher for relative subordinates (n indicated on plots; *win* 2-way RM-ANOVA main effects of relative rank $F_{(1,14)}=165$, $p=2x10^{-6}$; *lose* 2-way RM-ANOVA effect of relative rank $F_{(1,14)}=262$, $p=6x10^{-7}$). **b**, SVM performance is higher than chance for decoding relative rank, specifically for intermediate mice (area under the29 receiving operating curve: AUC, gray indicates performance when shuffling data; mean AUC vs shuffled AUC Wilcoxon rank sum: competitive success $p=1.8x10^{-4}$ and relative rank for intermediates $p=1.8x10^{-4}$).

However, social defeat models are used for inducing depressive like behavior in C57 mice (Golden et al.,

2011; Challis et al., 2013) so they are not a good alternative for what we want to study, which is the neural representation of stable, established ranks wherein cagemates are performing more affiliative behaviors than aggressive behaviors. We are not presently investigating social rank learning/establishment, but aim to do so in future studies. Our goal is to study the neural dynamics underlying social ranks of naturally forming hierarchies in groups living together. In C57 mice the tube test has been validated to reflect social dominance of co-housed mice (Wang et al., 2011). Other rodents, such as CD1 mice or Syrian Hamsters, are more aggressive at baseline and social aggression serves to quantify social dominance (CD1 mice: So et al., 2015; Williamson et al., 2016) (Syrian hamsters: Morrison et al., 2011, 2014).

3) The authors may already have these data, but as presented their data sets are binarized into subordinates and dominants. Do they see any correlations or decoding properties in mice that are ranked in the middle of a cages social hierarchy? (If a group of 5 has 1 dominant and 4 submissive, which one is the subordinate in their analysis, and would there be differences in accuracy of the decoder based on that social rank within the group - that would be an exciting finding, and it would lend additional biological credence to their predictive algorithms).

We thank the reviewer for this suggestion. We have now included several analyses that show that the neural dynamics we report reflect relative social rank, as we see changes in intermediate animals when they become relative subordinates vs dominants. First, when only analyzing the mPFC neural trajectories for intermediate mice, we see that when the mice are relative subordinates the trajectories are longer (**Rebuttal Figure 15a**), consistent with our previous findings in all mice. Furthermore, an SVM decoder predicted relative rank in intermediate animals with high accuracy (**Rebuttal Figure 15b**). This data is now included in **Figure 3** and we agree with the reviewer that it is exciting and strengthens our findings.

In addition, we now report the single cell responses to task events (tone onset for win or lose trials; port entries of self vs other) for intermediate animals in **Extended Data Fig.** 9. In intermediates, we also see that more neurons in relative dominant mice encode reward-seeking behavior (**Rebuttal Figure 16**), consistent with the conclusions from our main findings reported in **Figure 4**.

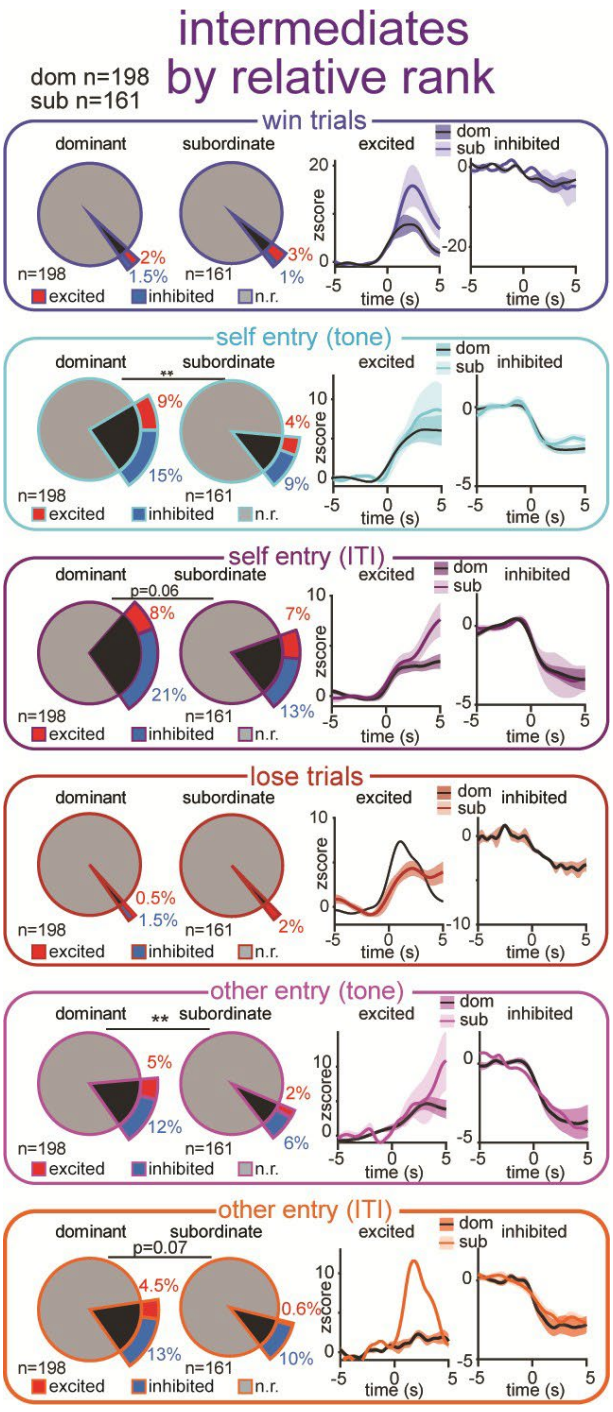

Rebuttal Figure 16: Single cell responses to task events during social competition in intermediates change with relative social rank. Total responsive cells and response magnitude to task-relevant event during social competition for intermediate rank mice (ranks 2 and 3) by relative rank (win trials: *dom* exc=4, *dom* inh=3, *sub* exc=5, *sub* inh=2, Fisher"s exact test p=0.76, Wilcoxon rank sum exc p=0.11, inh p=0.80; lose trials: *dom* exc=1 *dom* inh=3, *sub* exc=3, *sub* inh=0, Fisher"s exact test p=1, Wilcoxon rank sum exc p=1; self entries tone: *dom* exc=17, *dom* inh=30, *sub* exc=7, *sub* inh=14, Fisher"s exact test p=0.01, Wilcoxon rank sum exc p=0.89, inh p=0.57; other entries tone: *dom* exc=10, *dom* inh 23, *sub* exc=3, *sub* inh=10, Fisher"s exact test p=0.01, Wilcoxon rank sum exc p=0.46, inh p=0.79; self entries ITI *dom* exc=15, *dom* inh=42, *sub* exc=11, *sub* inh=21, Fisher"s exact test p=0.06, Wilcoxon rank sum exc p=0.11, inh p=0.44; other entries ITI *dom* exc=9, *dom* inh=26, *sub* exc=1, *sub* inh=16, Fisher"s exact test p=0.07, Wilcoxon rank sum exc p=0.20, inh p=0.91).

4) Did the authors try to present any reward omissions in the task? Does it impact the neural coding and behavior? It might be worth controlling for especially if position and posturing are important.

We did not present any reward omissions during this task, as each animal is already only collecting ~50% of the rewards with a Fixed Ratio 1 Reinforcement Schedule, and we wanted to prioritize examination of motivation as modulated by rank, rather than reinforcement probability. For example, how would a recent reward omission following an effortful win impact the decision to actively compete (as opposed to passively take turns), given its non-stationary impact on cost/benefit ratio? While these questions are interesting future directions, this study first sought to reveal the neural representations of social rank and competitive success,

and future investigations may broaden this task to probe the neural representations of reward contingency and its impact on competitive success.

However, to control for the effects of position in neural activity we took several approaches. We verified that the distance to the reward port did not differ by relative rank in our main analysis window (-5 to 5 sec surrounding cue onset; **Rebuttal Figure 17a**). We quantified how correlated mPFC single units are with distance to the reward port across relative ranks and see no difference (**Rebuttal Figure 17b**). Finally, we also analyzed our data by dividing per location: trials in which mice are close to the reward port vs far from the reward port. In both cases we see that neural trajectories are longer for relative subordinates (**Rebuttal Figure 17c**), demonstrating that the mPFC population dynamics we report are not driven by position differences in the reward chamber.

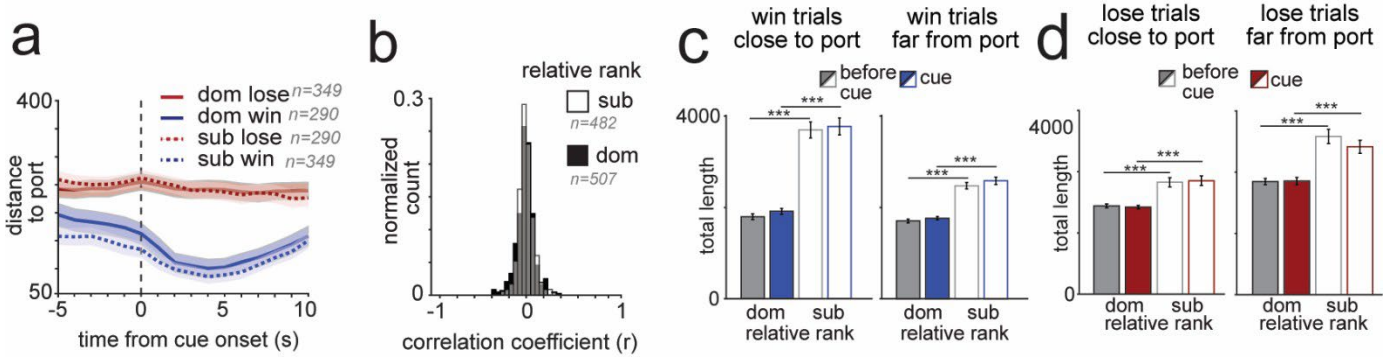

Rebuttal Figure 17: Position related controls. a, Distance to reward port differed by trial-type but not by rank (trials: *dom win*=290, *dom lose*=349, *sub win*=349, *sub lose*=290; 2-way ANOVA, main effect of trial-type $F_{(1,1274)}$=353, p=8.8x10^{-70}, rank p=0.098 and interaction p=0.066). **b**, Distribution of the correlation coefficients for firing rate and distance to port for the population of mPFC single units did not differ by rank (*dom*=321, *sub*=479; KS test, p=0.48). **c**, To determine if distance to reward port affected the population dynamics during *win* and *lose* trials a subset of data with matched video conditions was split by distance to reward port. Neural trajectory lengths were higher for relative subordinates during *win* trials in which mice were close or far to the reward port during tone onset (*dom* n=19 sessions, *sub* n=18 sessions; *win close to port*: 2-way RM-ANOVA main effect of rank $F_{(1,35)}$=738, p=5x10^{-21}; *win far from port*: 2-way RM-ANOVA main effect of rank $F_{(1,35)}$=588, p=3x10^{-20}). **d**, Neural trajectory lengths were higher for relative subordinates during *lose* trials in which mice were *close* or far from reward port during tone onset (*lose close to port*: 2-way RM-ANOVA main effect of rank $F_{(1,35)}$=588, p=3x10^{-20}; *lose far from port*: 2-way RM-ANOVA main effect of rank $F_{(1,35)}$=46.7, p=5x10^{-11}).

5) The paper ends with an experiment which shows that the PFC to LH circuit is most critical for the encoding the dominance behavior and predicting it therein. Yet, it wasn't clear to me from the data presented and/or the discussion why one would predict this circuit in the first place, when many other circuits could be better candidates, particularly within the context of social interaction/dominance. Do the photo-stimulation parameters used for the circuit really align with the decoder predictions? How is it so selective? Panel K of figure 5 doesn't appear to show a difference between the control eYFP and the CHR2 group? Is this just underpowered or due to other pathways or circuits being more necessary and sufficient?

We thank the reviewer for pointing out that our selection of PFC-LH was not clear. We have now added the following information to the manuscript which provides are rationale of why we selected to investigate this pathway:

Lines 223-233 "*Several mPFC subcortical pathways have been implicated in social behaviors(Ko, 2017), including the mPFC projection to the lateral hypothalamus (LH) in social aggression(Biro et al., 2018).*

The LH is comprised of a diversity of cell types and has been shown to drive hypersocial behavior and social investigation(Nieh et al., 2016), modulate social defensive behaviors(Rangel et al., 2016; Li et al., 2018), and promote reward/aversion(Aston-Jones et al., 2009; Jennings et al., 2015; Nieh et al., 2015, 2016). Further, it has been shown to play a critical role that in energy balance homeostasis(Burton et al., 1976) that suggests it is capable of performing the computations of a homeostatic control center(Cannon, 1929). Based on the conceptual framework for social homeostasis, after social information is detected and evaluated in a rank-dependent manner, it would be sent to a control center for comparison to a social homeostatic set point(Matthews and Tye, 2019; Lee et al., 2021). Given the multiplicity of roles for the LH in modulating homeostatic functions as well as both reward and social behaviors, we hypothesized that the LH would be a prime target to integrate social rank information from the mPFC with other homeostatic needs of the animal, during social competition."

In addition to these motivations from the literature of the role of LH and this prefrontal-LH pathway, we also see that mPFC-LH neurons are more responsive to winning trials compared to mPFC-BLA neurons, but this is not the case when mice are alone in the reward task (**Rebuttal Figure 9**). This data is now included in the manuscript in **Figure 5**.

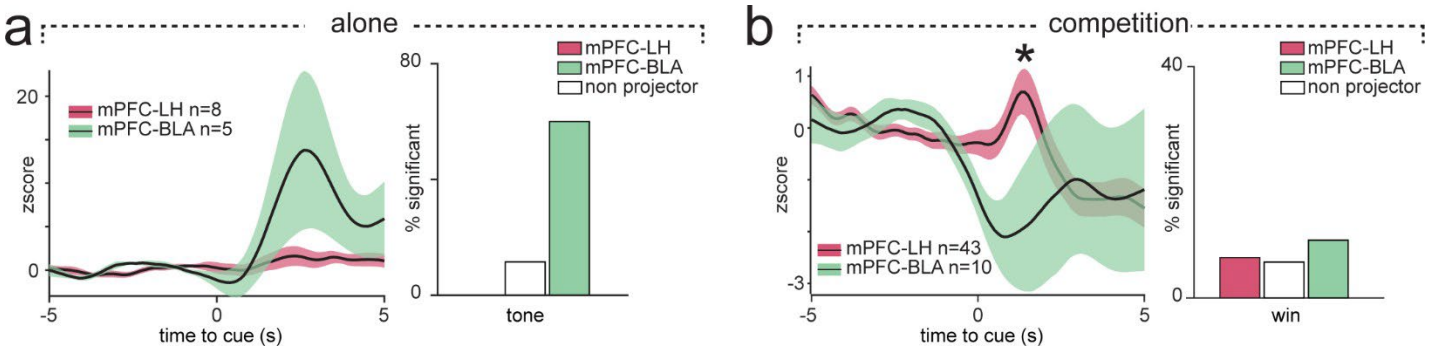

Rebuttal Figure 9: mPFC-LH neurons are more responsive to the wnning trials compared to mPFC-BLA neurons during reward competition. **a**, Right, firing rate of mPFC-LH is higher than mPFC-BLA during the reward delivery period in *win* trials in the competition (mPFC-BLA n=10 from 3 mice, mPFC-LH n=43 from 3 mice, Wilcoxon rank sum p=0.045). Left, percent cells responding to tone during *win* trials in the social competition (mPFC-LH n= 3/43, mPFC-BLA n=1/10 and non-phototagged n=57/620). **b**, Left, firing rate for projector populations during tones while animals performed the reward task alone (alone vs comp mean zscore during tone; mPFC-LH p= 0.043, mPFC-BLA p=0.049). Right, percent cells responding to the reward predictive tone when animal is alone (LH=0/8 from 3 mice, BLA 4/5 from 3 mice, non phototagged=54/470).

Finally, we report the statistics for **Fig. 5k** (now **Fig. 5i**) in the legends showing that mPFC-LH cell stimulation significantly increased the number of trials won (ChR2 n=9, eYFP n=6; 2-way RM ANOVA interaction of virus and light $F_{(1, 14)}$=5.82, p=0.03; Bonferroni corrected t-test ChR2 p=0.01).

6) The authors use the AlphaTracker approach throughout the manuscript and make claims that this approach is superior to other methods. It very well may be, but as presented the authors should either tone down that claim, or substantiate it more with more stringent comparisons with benchmarked alternatives. Considering one

of the novelty selling points of the paper is the computational technology it brings to the field. More evidence, data they may already have, and/or discussion for why this is a superior algorithm to other methods would help solidify the novelty of the method and impact.

We would like to clarify that we don"t make claims that AlphaTracker is better than other methods. Here is whatwe say about AlphaTracker in the manuscript:

Abstract: “With the development of a deep learning computer vision tool (AlphaTracker) and wireless electrophysiology recording devices, we have established a novel platform to facilitate quantitative examination of how the brain gives rise to social behaviors.”

Results: “The development and application of AlphaTracker represents a new platform for the emerging field of computational neuroethology of social behaviors (Nilsson et al., 2020; Pereira et al., 2020)”

Given that all tools have different architectures and require different training times they are difficult to compare. Instead we make our comparisons to human error levels calculated in the same datasets we are evaluating. We have now added to the manuscript in **Extended Data Fig. 2** AlphaTracker"s performance when tracking 4 mice (**Rebuttal Figure 2**).

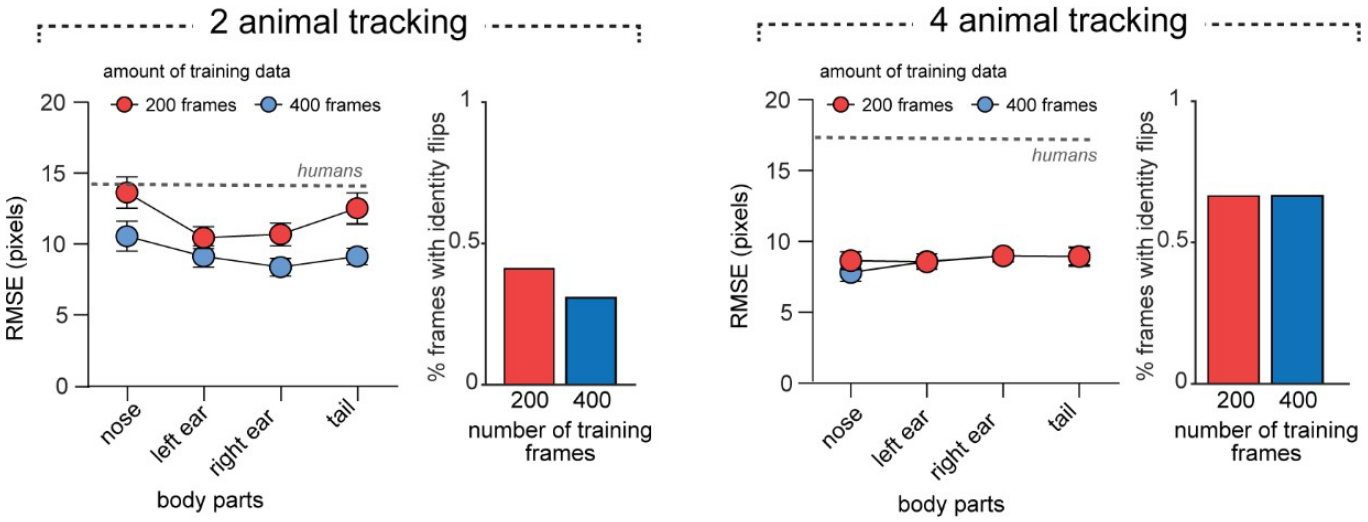

Rebuttal Figure 2: AlphaTracker performance for multiple animal tracking. Left, the root mean squared error (RMSE) in pixels for two animal tracking. Right, the RMSE in pixels for four animal tracking. In both cases identity tracking is more than 99% correct. Average RMSE error between two humans is indicated with a dashed line.

We have also compared AlphaTracker to other tools that have been recently released as preprints or open source code, and see that when tracking two unmarked mice. Despite having a lower median, AlphaTracker has a similar error rate to the other tools, and there was no detectable difference (**Rebuttal Figure 18**). Given the vast differences in network architecture and that these tools are not yet peer reviewed, we feel uncomfortable including these comparisons in the manuscript and instead believe that the correct comparison is human tracking.

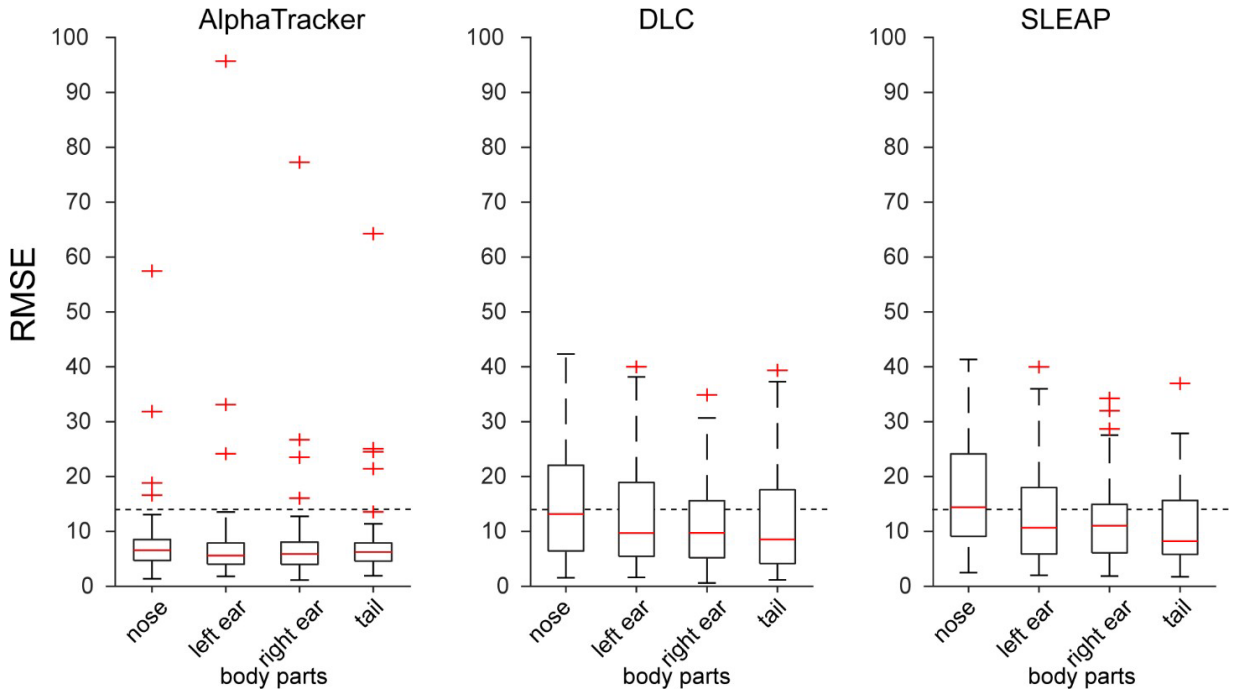

Rebuttal Figure 18: Error rate of tracking 2 mice across tools. Root mean square of the error (RMSE) is reported by body part for tracking 2 mice in the same frames across tools. Dash line indicates human level error, red lines in the box plot indicate median, and outliers are plotted. While AlphaTracker has a lower median, it has more outliers than DeepLabCut (DLC) and SLEAP. There was no difference in RMSE across tools. 56 frames were subsampled from a video and tested. AT produced all body points for 56/56, DLC 50/56 and SLEAP 48/56. For each tool, only frames with all tracked points were included in the analysis (2way ANOVA; no effect of body part $F(3,604) = 1.84$ $p=0.14$; no effect of tool $F(2,604) = 0.68$ $p=0.51$).

Minor-

Figure 1 A. The cartoons are somewhat helpful, but it would almost be better presented if the authors took actual photos of the mice in these positions. Presumably they used a high frame rate camera to capture the behavior for Alphatracker, so they should be available. It would be more powerful to see the behavior live in snap shots.

This is more stylistics, but just a suggestion. Most of us aren't used to reading / or understanding the plots presented in Fig 3 panel a or c. The manuscript might benefit in this case from a cartoon of what these plots might look like (mock results) in various predictions followed by the plots presented here with the real data.

Figure 3 panel f target plots are very hard to see. They should be increased in size, and/or de-pixelated (smoothened) to make a better case.

Figure 4. The data in A and B are nice, and thorough, but difficult to see unless you are at 400X, it might make sense to zoom in on a few and simplify the figure and put the rest in supplemental. No one wants to part with all their data on main figures, but it'll help the reader digest the work in this dense paper.

We appreciate all these suggestions and have made the following changes based on them:

a) Now we have included example frames of the behavior in **Extended Data Fig. 1**. In addition, we include example video clips in supplementary video 3.
b) We have superimposed a diagram of the PC space summary findings on the neural trajectory plots to facilitate the visual understanding of the data.
c) We have moved panel **Fig. 3f** to **Extended Data Fig. 7** and increased it in size 3x.
d) We have simplified figure 4 which has allowed us to make the data larger.

Referee #3 (Remarks to the Author):

This is a very impressive tour-de-force study involving diverse expertise from multiple laboratories combining neurobiological, behavioral, and computational approaches, as well as a novel behavioral tracking methodology in order to establish a causal function of the mPFC-to-LH pathways in social rank and related dominance behaviors. The uses of both the AlphaTracker as well as the HMM-GLM using neural activity for detecting unique behavioral and neural states are innovative for linking across complex social behavioral patterns and neural activity. The results of this study demonstrate the role of the mPFC in processing social rank and provide new knowledge on the mPFC-to-LH pathways in generating dominance behaviors. The authors also have used a wide array of behavioral assays in mice to better examine the selectivity of the findings. Overall, this new knowledge provides an important causal "bridge" between social rank representation in the mPFC and behavioral regulation by the LH. The use of the BLA as a control projection area (mPFC-->BLA) was very helpful in supporting the specificity of the mPFC-->LH pathways. Additionally, these findings involving the LH nicely support the principle behind social homeostasis previously proposed by the authors, which I find it to be very exciting. For these reasons summarized above, I am very enthusiastic about this work.

The reported findings and the ways that the authors have elegantly applied technical innovations should be highly interesting to both basic sciences and social sciences fields interested in social behaviors broadly, as well as clinical fields interested in regulating aggressive behaviors in social settings. All the statistical and data analytic approaches are sound, with appropriate displays of error bars and other useful information in the Extended Data for better understanding the data.

I have the following specific comments for the authors to consider.

1) The authors of the paper have developed a novel paradigm, combining wireless recording, automated tracking of behavior through computer vision, and a competitive reward gathering task aimed at investigating social ranks and dominance behavior. However, from the data presented, it requires more clarification with respect to how much of this task is measuring dominance and competitive behaviors on *every trial. On line 76,the authors state that "Importantly, differences in winning were not driven by overall location in the arena or distance to port prior to tone onset" and reference Extended Data Fig. 1. However, data shown in Extended Data Fig. 1B and in Figure 3F/3G seem to suggest that the distance from the port during the baseline period is clearly predictive of which animal wins the trial, i.e., in Extended Data Fig. 1B, Figure 3F and 3G it is clear that even at –5 seconds before the cue, the winning animals (regardless of "dom win" and "sub win") are both closer to the port. The 2-way ANOVA in the legend of Extended Data Fig. 1B found a main effect of trial type (presumably dominant win/subordinate lose vs. dominant lose/subordinate win), and no effect of rank ($p=0.071$) or interaction. It is unclear from the text what time period of the fifteen seconds shown this ANOVA isexamining.

We thank the reviewer for finding this mistake and apologize about this error. We have corrected our statement (in lines 83-84 of the manuscript) to say that differences in relative rank were not driven by differences in overall location and distance to port. Indeed, as the reviewer points out there is a very large effect of trial type (win vs lose) in distance to port.

In the legend of **Extended Data Fig. 1f** (corresponding to **Extended Data Fig. 1b** in the previous version of the manuscript) we have now clarified the time period used for the statistics. For thoroughness, we now

report the statistics for 3 different time periods and in all cases, there is no effect of relative rank. For early in the baseline period (from 30 sec to 20 sec prior to cue) there was no difference between trial type nor rank, while closer to the cue onset (from -5 to 0) and during tone time there was an effect of trial type **(Rebuttal Figure 19).**

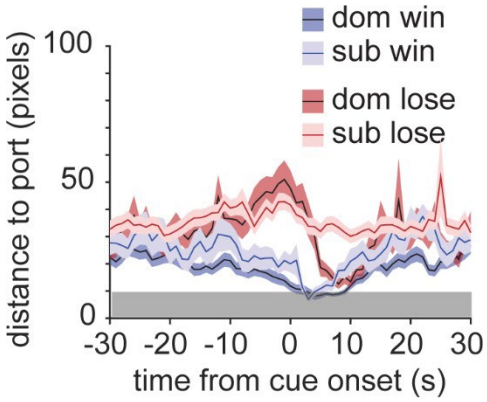

Rebuttal Figure 19: Distance to reward port differed by trial type but not by relative rank (trials n: *dom win*=68, *dom lose*=24, *sub win*=24, *sub lose*=68 from 12 dyads; early baseline -30 to -20 s prior to cue there is no effect of trial nor relative rank; 2-way ANOVA using the mean distance from -5 s to cue onset: main effect of trial type $F_{(1,180)}$=44.4, p=3x10^{-10}, rank p=0.94, interaction p=0.09; 2-way ANOVA using the mean distance from
5 seconds prior to tone until 10 seconds post tone: main effect of trial type $F_{(1,180)}$=68, p=2.5x10^{-14}, rank p=0.071, interaction p=0.79). Gray line indicates contact range for the reward port.

If distance from the port prior to the cue onset is strongly predictive of trial outcome, then not *all trials may be actually “competitive” or “measuring dominance”. An alternative explanation would be that during some trials the nearest mouse to the port can more easily gather the reward, and this trial is either weakly or not contested by the other mouse (could be due to dominance or motivation level, for example). If this is the case, then it would change the interpretation of mPFC population dynamics that track “competitive success”. To address this, we would suggest the authors do one of the following:

a) Specifically examine if the winner of each trial and the neural activity predictive of winning were not simply determined by distance from the port during the baseline periods.

b) Exclude any trials from the analyses that do not include pushing, resistance, or displacement behaviors to provide additional neural insights when specifically examining only those trials with clearly observable competition (pushing, resistance, displacement).

c) Examine if prior behavioral dynamics and underlying neural activity between the mice prior to the cue determined their relative positioning during the baseline period (i.e., the competitive behaviors were occurring in advance of the cue presentation) to better understand what behavioral dynamics and neural processes occur that led to the better positioning to begin with that ultimately resulted in winning.

For future directions (note: I am definitely not suggesting the authors to collect more data), I would recommend including a condition where pairs of mice in an experimental chamber with multiple reward ports, where the location of the upcoming cue cannot be predicted and where presumably the positioning during baseline between the dominant and subordinate would average out to being equivalent.

We thank the reviewer for these thoughtful suggestions. Indeed, understanding the contribution of behavioral differences to the decoding accuracy is important. Considering that earlier periods prior to the tone did not have a difference in distance to port for winning vs losing trials (**Rebuttal Figure 19**), we decided to modify our decoding analyses to include a much larger time window from -30 sec to 30 sec post cue onset. We see that decoding of competitive success is highly stable, including in times -30 to -20 and 20 to 30 s post cue onset (**Rebuttal Figure 14**). On the other hand, if we use distance to reward port to predict competitive success the decoder can only predict above chance for periods closer to the tone onset (**Rebuttal Figure 20**). Despite distance to reward port being predictive of competitive success, this is not as stable and, importantly, distance to reward port explains a small portion of the variance of mPFC firing rate, with most mPFC cells having and very low correlation coefficient between firing rate and distance to reward port (mean correlation coefficient across cells -0.004; **Rebuttal Figure 21**). Furthermore, when removing mPFC reward place cells (cells with significant correlation between distance to reward and firing rate) we still observed the same population dynamics of subordinates having longer trajectories (**Rebuttal Figure 22**).

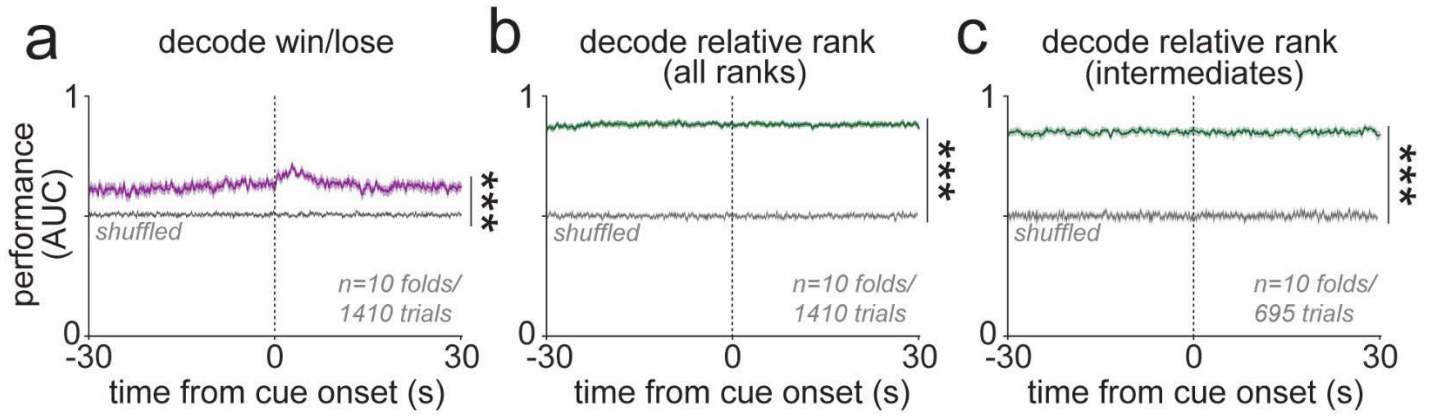

Rebuttal Figure 14: SVM performance is higher than chance for decoding (**a**) competitive success and (**b**) relative rank for all mice and specifically for (**c)** intermediate mice (area under the receiving operating curve: AUC, gray indicates performance when shuffling data; mean AUC vs shuffled AUC Wilcoxon rank sum: competitive success $p=1.8x10^{-4}$, relative rank all mice $p=1.8x10^{-4}$ and relative rank for intermediates $p=1.8x10^{-4}$).

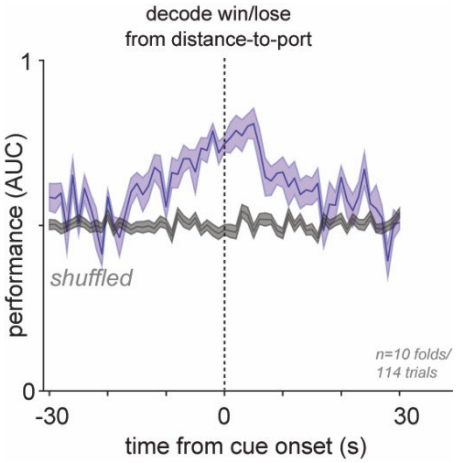

Rebuttal Figure 20: Decoding competitive success using distance from reward port. Performance of decoder when predicting win/lose from distance to reward port degrades as time to cue increases.

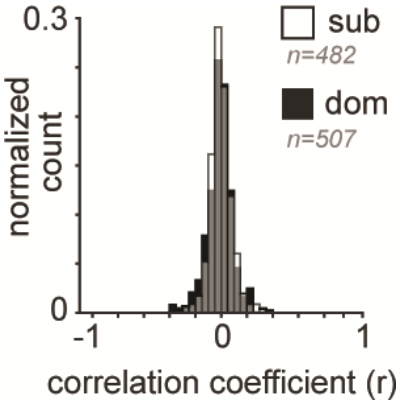

Rebuttal Figure 21: Distribution of the correlation coefficients for mPFC firing rate and distance to port did not differ by relative rank (number of neurons: *dom*=321, *sub*=479; KS test, p=0.48).

2) Similarly, further clarifications of the behavior would be useful in interpreting some of the data. For example, on line 74 the authors state that "Dominant animals, as defined by the tube test, obtained more rewards, spent more time at the reward port, and were more successful at displacing the competitor from the port (Fig. 1c)." This seems to be in contrast to line 99 of the manuscript which states "When recording during the reward competition task, we did not detect a statistically significant difference in the number of rewards earned by dominant and subordinate mice, us to make comparisons about dominance behavior and competitive success without being confounded by the volume of reward consumption." Does this indicate two separate behavioral datasets were used, one during neural recording and one without neural recording? Related: please add the test and the number of observations used for "we did not detect a statistically significant difference".

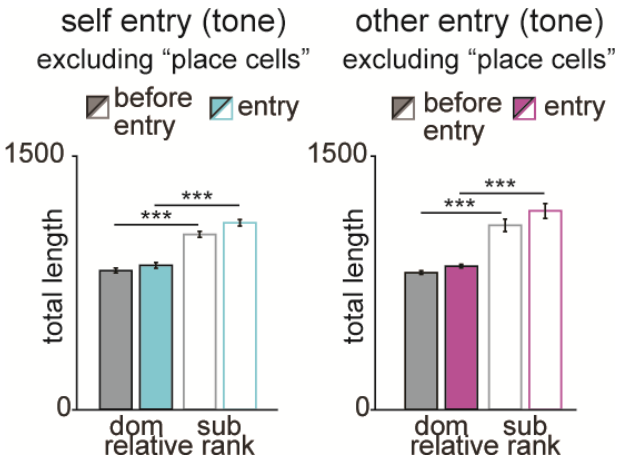

Rebuttal Figure 22: To determine if reward "place cells" contributed to neural trajectory rank differences we calculated the neural trajectory lengths without cells that were correlated to distance to port in a subset of data with equivalent video settings (see methods). Left, neural trajectories for *self entry* during the tone are highest for relative subordinates without the distance correlated cells (*dom* n=19 sessions, *sub* n=18 sessions; 2-way RM-ANOVA main effect of rank $F_{(1,35)}=94.4$, $p=1x10^{-13}$). Right, neural trajectories are highest for relative subordinates without the distance correlated cells (excluding correlated cells: 2-way RM-ANOVA main effect of rank $F_{(1,35)}=100$, $p=1x10^{-13}$).

We apologize for the lack of clarity in our previous version of the manuscript. Indeed, there were separate behavioral cohorts for the no recording vs recording experiments. Not surprisingly, when wearing 6 g headstages (25% body weight for many of the mice) there were changes in the behavior of the mice (**Rebuttal Figure 3**). We have now included the sample size for the animals used for mPFC recordings and behavioral dataset descriptions in the ***methods section 6*** "Reward competition recording experiments" and in the relevant figure legends. For **Figure 2** we used 13 mice with similar video settings (resolution and camera angle) to allow for the automated behavioral analysis. For **Figure 3** we used 20 mice for the reward competition recordings, 24 mice for the alone recordings and 15 mice for the experiment in which mice received alone and competition trials in the same recording session. For **Figure 4** we used 10 mice for the alone recordings (5 rank 1 and 5 rank 4) and 20 mice for the reward competition recordings. For **Figure 5** we recorded 20 mice of which 9 had viral injections for mPFC-LH phototagging and 8 had viral injections for mPFC-BLA phototagging.

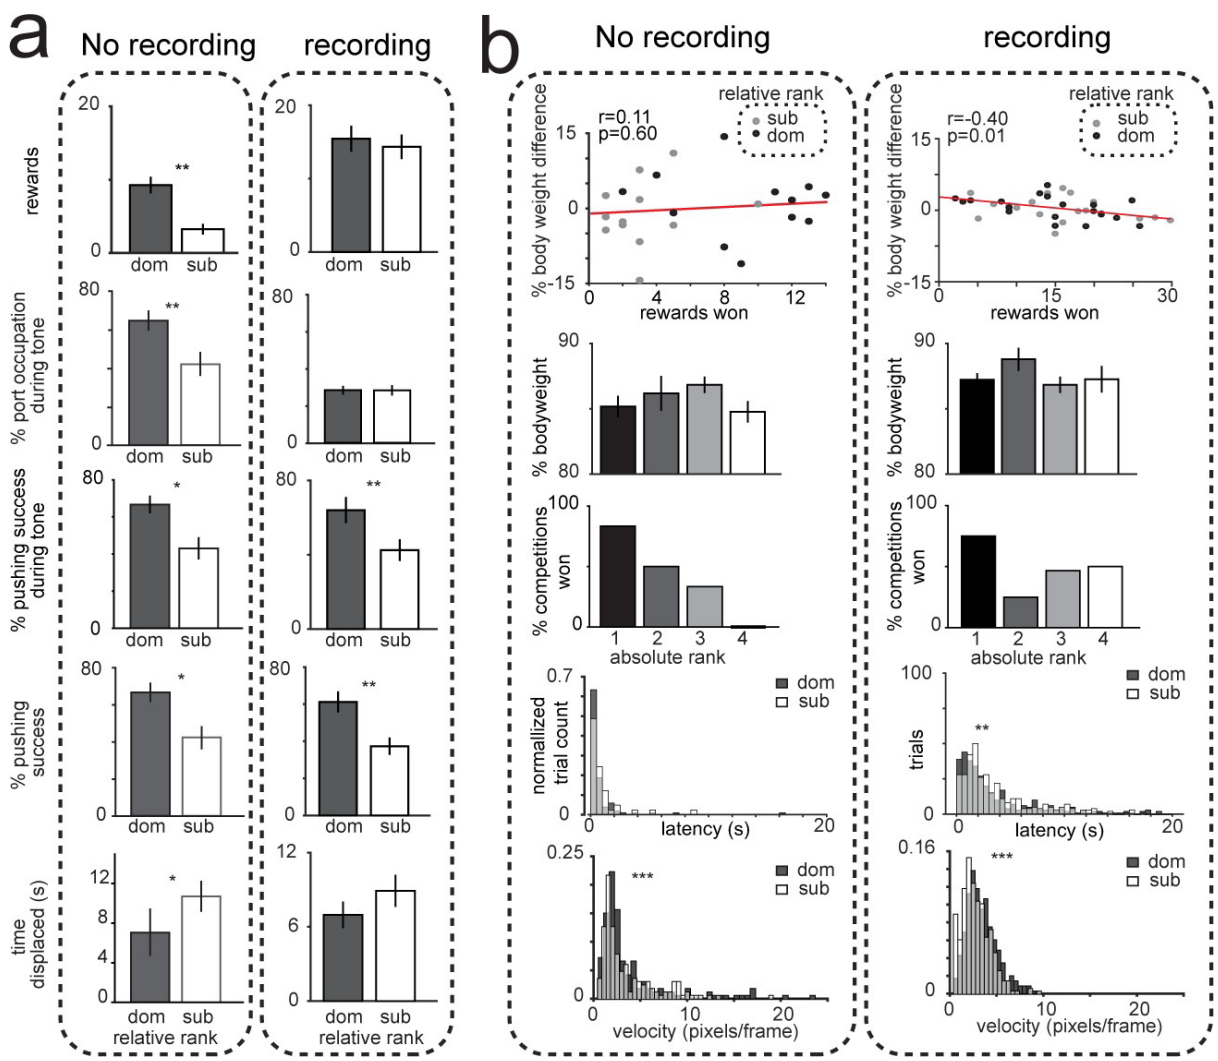

Rebuttal Figure 3: **Reward competition behavior summary for recording vs no recording cohorts**. **a**, Number of rewards, % port occupation during tone time, % pushing success during tone and total % pushing success, and time displaced from reward port across relative dominant and subordinate mice. Across cohorts % pushing success during tone and total % pushing success are higher for relative dominant mice. **b**, Rewards won as a function of bodyweight difference, % bodyweight and % competitions won by absolute rank, distribution of latency to pick up reward across trials and distribution of velocity across trials. Across cohorts rank 1 mice win the majority of competitions and relative dominant mice have higher velocities during the social competition. *p<.05,**p<.01,***p<.001

3) In all population level neural analyses, the authors have pooled together activity from multiple subjects and use the total number of recorded units as the sample size. This approach is not necessarily unorthodox in the field of behavioral neurophysiology where collection of spiking activity during interactive social behaviors is extremely challenging and large pseudo-populations of units are sometimes requisite for advanced analyses such as hidden markov models. In some areas, the authors made specific efforts to address the potential shortcomings of pooling together activity from multiple subjects. For example, in (Extended Data Fig. 7k-n) the pooled data was split between two different randomly selected subsets over 50 bootstrapping iteration, or in Figure 3B/3D where the authors employed the “leave one out” method excluding neurons from a single animal in each iteration to control for an individual mouse from strongly biasing the results. In other analyses, however, there is no control for potential between-subject variations. Specifically, if it is not possible to construct a HMM-GLM for each subject, or to split individual mPFC cells by rank-dependent responses for

each subject, then a good alternative would be to also show inter-subject correlation coefficient that assesses the homogeneity (or potential / interesting differences) of neural responses between mice.

We have now validated that we have the same results with the GLM-HMM when using a leave-one-animal-out cross-validation method vs the originally used 10-fold CV (**Rebuttal Figure 23a**). In addition, to test the robustness of our HMM-GLM model we trained and tested on the relative dominant vs subordinate dataset and saw equivalent performances when trained on one dataset and tested in the other (**Rebuttal Figure 23b**). We believe that altogether these additional validations strengthen our manuscript and have now included these findings in **Extended Data Fig. 5.**

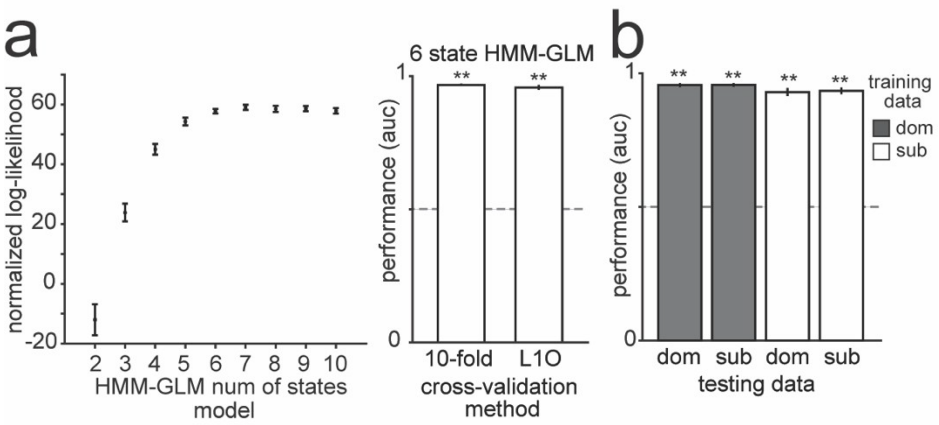

Rebuttal Figure 23: Additional controls for HMM-GLM decoding of behavior. a, Left, model selection forHMM-GLM state number using the leave one out method (LOO) results in a 6 states model being optimal. Right, HMM-GLM 6 states model performance predicts behavioral label regardless of training method utilized (AUC n=9, one per each behavior label; Sign test of model performance vs chance p=0.004 for both methods) **b,** HMM-GLM 6 states model predicted behavioral label regardless of which dataset was used for training or testing (n=9 behavior labels using 482 trials for *dom* vs 478 trials for *sub*; Sign test performance vs 0.5 (chance) p=0.004 for all tests).

4) In the comparison between PCA vector length, it is stated that subordinates had longer neural trajectories, speculated to be due to either higher or faster firing rate changes in the mPFC population activity. This result is quite intriguing, and although this is later addressed tangentially (Line 187; "subordinates had phasic responses of greater amplitude in response to events…"), the interpretation of this result would greatly improve if that question could be examined quantitatively with the existing data.

We thank the reviewer for this great suggestion! We quantified single cell mPFC data to determine if either higher (larger amplitude) or faster firing rate changes underlie the group differences observed in trajectory lengths. To test whether the longer lengths are due to larger amplitude changes, we computed the variance of each neuron after mean-centering the trial-averaged activity. Subtracting the mean makes the variance a good approximation of the amplitude. To test whether the longer lengths are due to faster firing rate changes, we computed the mean rate of change of the z-scores of the firing rates of each neuron. We found that the longer trajectory lengths reflect higher firing rate changes as the variance was much higher for relative subordinates vs dominants (**Rebuttal Figure 24**). We have now included these analyses in **Fig. 3e**. These additional findings point to a biological explanation for the trajectory length difference, and also better connect the population dynamics observed with the individual cell analyses reported in **Figure 4**.

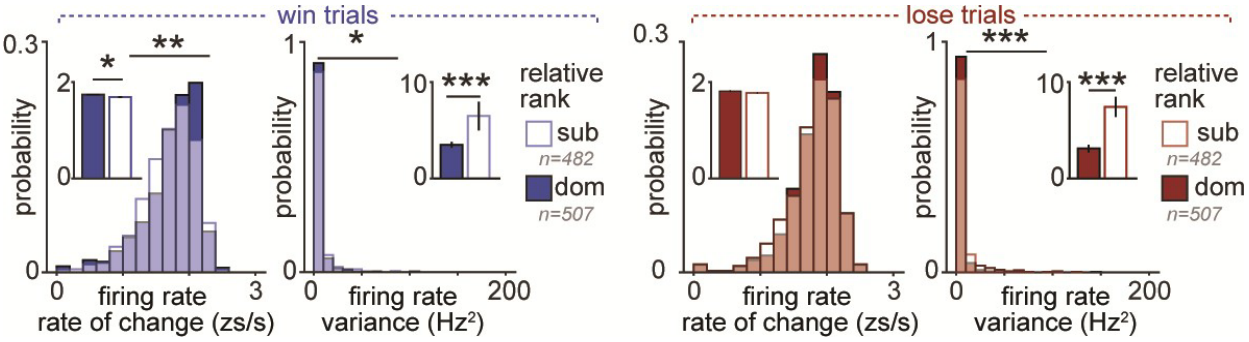

Rebuttal Figure 24: Firing rate variance is higher for relative subordinates in both *win* and *lose* trials, while rate of change is higher for relative dominants only in *win* trials (number of neurons indicated in plots, inset plot has average across groups; *win* trials rate of change: Kolmogorov-Smirnov (KS) test p=0.009, Wilcoxon rank sum p=0.01; *win* trials variance: KS test p=0.01, Wilcoxon rank sum p=0.19; *lose* trials rate of change: KS test p=0.40, Wilcoxon rank sum p=0.19; *lose* trials variance KS test $p=5x10^{-7}$, Wilcoxon rank sum $p=2x10^{-9}$).

5) In Figure 5J, the difference between cumulative rewards obtained by light OFF session vs light ON session appears to only emerges late in the experiment (past 15 trials). This is interesting. Do the authors have any interpretation for this time course? For example, the non-immediate effect of light delivery may suggest that the modulation of social dominance by the cortico-hypothalamic circuit is mediated through slower mechanisms, such as learning.

We agree that this delayed effect is interesting, and we believe that it may reflect that there is plasticity in this pathway with the repeated stimulation that leads to the effect, or that the behavioral effect is mediated by slow acting neuropeptides released by LH subpopulations, such as orexin. In addition, as the reviewer mentioned, the late emergence could indicate learning being necessary for the behavioral effect to emerge. All of these possibilities are not mutually exclusive and could co-exist. We added the following line in the results to comment on this slow effect in lines 260-262: “Interestingly, the increases in rewards obtained were most notable towards the end of the session suggesting that experience and stimulation-dependent plasticity and/or neuromodulation in the mPFC-LH pathway may mediate this change in behavior.”

6) Different analyses required using various sample sizes for good reasons. For example, both Figure 3B and 3D use a sample size of thirteen (resulting from „leave one out“ method). This is only explained in the supplemental text, and the accompanying figures (3A and 3C) use a different sample size (507 and 490 neurons). Some other figures lack descriptions of sample size (e.g., **Extended Data Fig.** 7). Please add N information and also add how many total mice were used in these studies and which mice (if any) overlapped between different aspects of the manuscript (which can even be in the method section).

We apologize for the complexity, the missing information and confusion. Since we have different methods of analysis the sample size varies per analysis. We have now indicated the number of mice utilized for recordings in **Figures 2-5** legends and in the methods. For **Figure 2** we used 13 mice with similar video settings (resolution and camera angle) to allow for the automated behavioral analysis. For **Figure 3** we used 20 mice for the reward competition recordings, 24 mice for the alone recordings and 15 mice for the experiment in which mice received alone and competition trials in the same recording session. For **Figure 4** we used 10 mice for the alone recordings (5 rank 1 and 5 rank 4) and 20 mice for the reward competition recordings. For **Figure 5** we recorded 20 mice of which 9 had viral injections for mPFC-LH phototagging and 8 had viral injections for mPFC-BLA phototagging. In the ***methods section 6*** **“**Reward competition and “alone” recording experiments” we indicate the overlap as follows:

For the mPFC recordings during the reward competition 24 new mice were implanted with electrodes. Out of those 24 mice only 20 mice had cells but the remaining 4 were still used for behavioral purposes as competitors. Additional animals were implanted (7 mice) and together with a subset of the mice used for competition (17 mice), had recording sessions after reaching training criteria while performing the task alone before ever having a social competition session (24 mice in total). Finally, 12 out of the 20 mice used for competition recording we performed an additional experiment consisting in alone trials followed by competition trials (**Fig. 3l**) in the same recording session. For those recording sessions, the mouse started the reward task alone and the experimenter opened the chamber and added a competitor midway through the session. For all experiments, several times recording sessions were excluded because of battery failures. Competition sessions during recording experiments consistent of 30 trials to maximize the amount of data obtained, both mice always wore loggers to match conditions across competitors. However, in 6 mice we utilized first generation loggers that had shorter battery life and therefore sessions for those animals consisted in 20 trials. Electrode-implanted mice were trained to wear a weight matched dummy headstage (mimicking the wireless logger weight) during their individual reward training. Individual reward training continued until mice reached the mean 5 sec latency criteria.

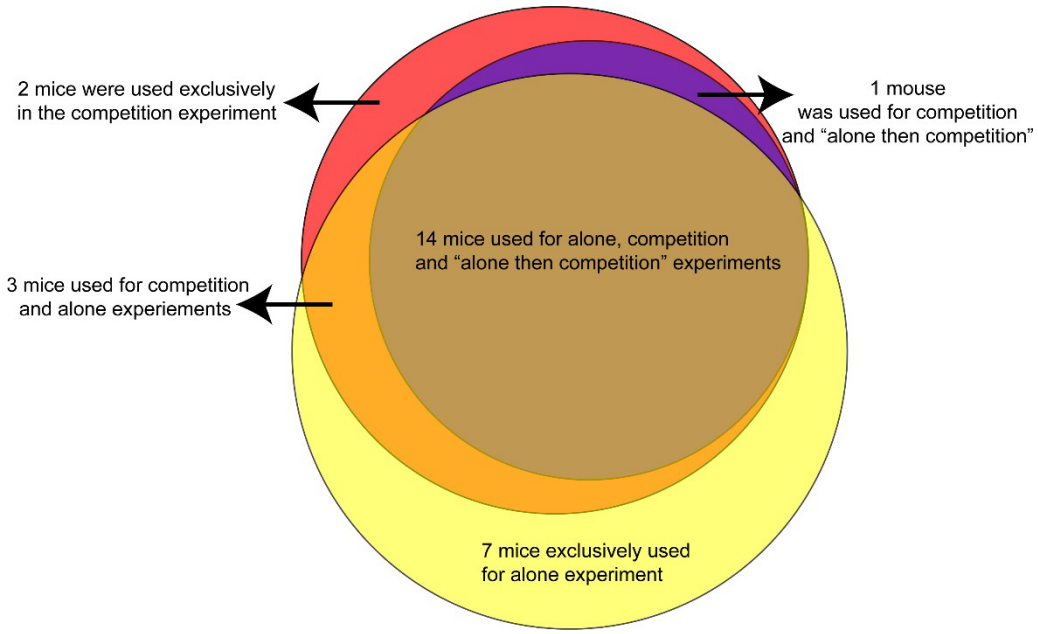

Rebuttal Figure 25: Schematic of workflow for subjects across cohorts. Venn Diagram indicating overlap of mice recorded across experiments. The majority of the mice were used in all experiments, however, to consider absolute rank differences we added additional animals to the alone experiments.

We have also added the following venn diagram in our methods to depict the overlap of mice across experiments.

7) In several figures (Figure 1G, **Extended Data Fig.** 3A and 3B) output from UMAP (Uniform Manifold Approximation and Projection) clustering is presented without any explanation in the main, or supplemental, text. Although this visualization is informative and appropriate, this technique is still relatively novel, and some readers may not be able to interpret the output (i.e., non-labelled axes) without some explanation or reference to McInnes Et al. 2018. (McInnes, Leland & Healy, John & Saul, Nathaniel & Grossberger, Lukas. (2018). UMAP: Uniform Manifold Approximation and Projection. Journal of Open Source Software. 3. 861. 10.21105/joss.00861.)

Thanks so much for noticing we forgot to include an explanation of UMAP in our methods. We have now added a short explanation of this method and the suggested reference in the ***methods section 15*** AlphaTracker under the heading unsupervised clustering for behavioral motifs.

Signed below to opt in for Nature"s new transparent peer review scheme: Reviewed by Steve W. C. Chang with assistance from Philip T. Putnam, a postdoctoral associate in his lab.

References Cited:

Aston-Jones G, Smith RJ, Moorman DE, Richardson KA (2009) Role of lateral hypothalamic orexin neurons in reward processing and addiction. Neuropharmacology 56 Suppl 1:112–121.

Biro L, Sipos E, Bruzsik B, Farkas I, Zelena D, Balazsfi D, Toth M, Haller J (2018) Task Division within the Prefrontal Cortex: Distinct Neuron Populations Selectively Control Different Aspects of Aggressive Behavior via the Hypothalamus. J Neurosci 38:4065–4075.

Burton MJ, Rolls ET, Mora F (1976) Effects of hunger on the responses of neurons in the lateral hypothalamus to the sight and taste of food. Exp Neurol 51:668–677.

Calhoun AJ, Pillow JW, Murthy M (2019a) Unsupervised identification of the internal states that shape natural behavior. Nat Neurosci 22:2040–2049.

Calhoun AJ, Pillow JW, Murthy M (2019b) Unsupervised identification of the internal states that shape natural behavior. Nat Neurosci 22:2040–2049.

Cannon WB (1929) Organization for Physiological Homeostasis. Physiol Rev 9:399–431.

Challis C, Boulden J, Veerakumar A, Espallergues J, Vassoler FM, Pierce RC, Beck SG, Berton O (2013) Raphe GABAergic neurons mediate the acquisition of avoidance after social defeat. J Neurosci Off J Soc Neurosci 33:13978–13988, 13988a.

Escola S, Fontanini A, Katz D, Paninski L (2011a) Hidden Markov models for the stimulus-response relationships of multistate neural systems. Neural Comput 23:1071–1132.

Escola S, Fontanini A, Katz D, Paninski L (2011b) Hidden Markov models for the stimulus-response relationships of multistate neural systems. Neural Comput 23:1071–1132.

Fusi S, Miller EK, Rigotti M (2016) Why neurons mix: high dimensionality for higher cognition. Curr Opin Neurobiol 37:66–74.

Golden SA, Covington HE, Berton O, Russo SJ (2011) A standardized protocol for repeated social defeat stress in mice. Nat Protoc 6:1183–1191.

Hornak J, O"doherty J, Bramham J, Rolls ET, Morris RG, Bullock PR, Polkey CE (2004) Reward-related reversal learning after surgical excisions in orbito-frontal or dorsolateral prefrontal cortex in humans. J Cogn Neurosci 16:463–478.

Jennings JH, Ung RL, Resendez SL, Stamatakis AM, Taylor JG, Huang J, Veleta K, Kantak PA, Aita M, Shilling-Scrivo K, Ramakrishnan C, Deisseroth K, Otte S, Stuber GD (2015) Visualizing hypothalamic network dynamics for appetitive and consummatory behaviors. Cell 160:516–527.

Ko J (2017) Neuroanatomical Substrates of Rodent Social Behavior: The Medial Prefrontal Cortex and Its Projection Patterns. Front Neural Circuits 11 Available at: http://journal.frontiersin.org/article/10.3389/fncir.2017.00041/full [Accessed July 22, 2017].

Lee CR, Chen A, Tye KM (2021) The neural circuitry of social homeostasis: Consequences of acute versus chronic social isolation. Cell 184:1500–1516.

Li Y, Zeng J, Zhang J, Yue C, Zhong W, Liu Z, Feng Q, Luo M (2018) Hypothalamic circuits for predation and evasion. Neuron 97:911–924.

Matthews GA, Tye KM (2019) Neural mechanisms of social homeostasis. Ann N Y Acad Sci 1457:5–25.

Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. Annu Rev Neurosci 24:167–202.

Morrison KE, Bader LR, Clinard CT, Gerhard DM, Gross SE, Cooper MA (2014) Maintenance of dominance status is necessary for resistance to social defeat stress in Syrian hamsters. Behav Brain Res 270:277–286.

Morrison KE, Swallows CL, Cooper MA (2011) Effects of dominance status on conditioned defeat and expression of 5-HT1A and 5-HT2A receptors. Physiol Behav 104:283–290.

Nieh EH, Matthews GA, Allsop SA, Presbrey KN, Leppla CA, Wichmann R, Neve R, Wildes CP, Tye KM (2015) Decoding Neural Circuits that Control Compulsive Sucrose Seeking. Cell 160:528–541.

Nieh EH, Vander Weele CM, Matthews GA, Presbrey KN, Wichmann R, Leppla CA, Izadmehr EM, Tye KM (2016) Inhibitory Input from the Lateral Hypothalamus to the Ventral Tegmental Area Disinhibits Dopamine Neurons and Promotes Behavioral Activation. Neuron 90:1286–1298.

Nilsson SR, Goodwin NL, Choong JJ, Hwang S, Wright HR, Norville ZC, Tong X, Lin D, Bentzley BS, Eshel N, McLaughlin RJ, Golden SA (2020) Simple Behavioral Analysis (SimBA) – an open source toolkit for computer classification of complex social behaviors in experimental animals. bioRxiv:2020.04.19.049452.

Pereira TD, Tabris N, Li J, Ravindranath S, Papadoyannis ES, Wang ZY, Turner DM, McKenzie-Smith G, Kocher SD, Falkner AL, Shaevitz JW, Murthy M (2020) SLEAP: Multi-animal pose tracking. bioRxiv:2020.08.31.276246.

Platt JC (1999) Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods. In: Advances in Large Margin Classifiers, pp 61–74. MIT Press.

Rangel MJ, Baldo MVC, Canteras NS, Hahn JD (2016) Evidence of a Role for the Lateral Hypothalamic Area Juxtadorsomedial Region (LHAjd) in Defensive Behaviors Associated with Social Defeat. Front Syst Neurosci 10:92.

Ridderinkhof KR, Van Den Wildenberg WP, Segalowitz SJ, Carter CS (2004) Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. Brain Cogn 56:129–140.

Rigotti M, Barak O, Warden MR, Wang X-J, Daw ND, Miller EK, Fusi S (2013) The importance of mixed selectivity in complex cognitive tasks. Nature 497:585–590.

Sawaguchi T, Goldman-Rakic PS (1991) D1 dopamine receptors in prefrontal cortex: involvement in working memory. Science 251:947–951.

So N, Franks B, Lim S, Curley JP (2015) A Social Network Approach Reveals Associations between Mouse Social Dominance and Brain Gene Expression Branchi I, ed. PLOS ONE 10:e0134509.

Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. Nature 411:953–956.

Wang F, Zhu J, Zhu H, Zhang Q, Lin Z, Hu H (2011) Bidirectional Control of Social Hierarchy by Synaptic Efficacy in Medial Prefrontal Cortex. Science 334:693–697.

Williamson CM, Lee W, Curley JP (2016) Temporal dynamics of social hierarchy formation and maintenance in male mice. Anim Behav 115:259–272.

Reviewer Reports on the First Revision:

Referee #1 (Remarks to the Author):

I am a behavioral neurophysiologist, and I re-reviewed this paper in conjunction with a computational neuroscientist/neurotheorist. Comments reflect our pooled suggestions to primarily aid the authors in improving the rigor and reproducibility of their work and enhancing communication of their findings for a broad scientific audience, and to secondarily aid in editorial consideration of this work for publication.

I am highly supportive of this work, and I believe that it provides an important advance that is of broad interest to the field. In my original review, I provided a rich feedback that I believed would enhance the reproducibility of this outstanding work. The authors were incredibly responsive. I also suggested that the authors simplify the work, because I believed that even a simplified version would still provide an important advance that is of broad interest to the field. The authors generally chose to take the opposite approach, which was the provide substantially more experimental details and evidence to support their work. It was a great choice on their part. They have addressed nearly all of my concerns. However, a major concern remains unaddressed.

Major Concern: I previously commented that "The results on the modeling are unusual, where gains between the autoregressive model compared to the SVM and the GLM are drastic. It is rare that you see such dramatic gains when moving to an HMM based model, and the manuscript would be much improved by explaining this outcome. In fact, the authors show that there are drastic gains by using 2 clusters in the HMM based model (noting that the 1 cluster case is equivalent to the GLM model in your formulation). In this case, they are now distinguishing between 9 classes based on a mixture of 2 GLMs, where the GLM is chosen through the historical information. Compared to the classes, this is really a minor gain in modeling complexity, and I'm unclear from a mathematical perspective as to how such gains can be achieved. Additional justification and exploration would be warranted to ensure robustness of these findings."

In response the authors provided additional background and reference to supplementary materials. But it is important to clearly articulate why my concern regarding their approach remains. While I appreciate the extra information on the model description, the problem remains in the log-likelihoods of Figure 2(c) persists. First, as they have described it, the authors are modeling discrete behaviors. It is well-known that loglikelihoods of discrete outcomes MUST be negative, but the reported log-likelihoods for the HMM are positive for more than 2 states. Second, it is not explicitly stated what normalization is used in the log-likelihoods, or at least I could not find it. The typical normalization used in machine learning is to normalize by the number of instances. As such, a baseline random guess with probability 1/9 for each outcome would give a loglikelihood of -log(9)~=-2.2. Hence, I would expect all the reported loglikelihoods to be in the range of -2.2 to 0 (since aperfect model would give 0). In contrast, the reported normalized log-likelihoods are -12000 to positive 60. The GLM gives a likelihood that is close to infinitely worse than random

whereas Figure 2(d) shows that the GLM is clearly above chance. Thus, the loglikelihoods must be properly explained in order to evaluate to robustness of the result. I can believe that the temporal history helps to make predictions, so the AUC performances seem believable by themselves; however, given the issues in the loglikelihoods I remain concerned that another group would not be able to effectively reproduce the findings.

As a minor detail, an HMM-GLM with one state is identical to the GLM. I also would like to understand why those two models do not give the same log-likelihoods despite being mathematically the same model.

Finally, the comparison approach on the mixture of GLMs isn't particularly relevant. A typical trick in time series predictions is just to aggregate the last few time steps to form the predictors for a GLM,which would be a much more relevant comparison model. I am not too concerned about this, sincethe major issue is the reported loglikelihoods.

Another minor issue is in the response on counting the GLMs. The authors state that a 2-state GLM HMM consists of 22 GLMs, which is inconsistent with typical statistics and machine learning definitions. The multinomial logistic model is considered a single GLM, so in typical nomenclature there are only 2 behavior emissions GLMS (one for each state).

I do hope that additional description of the methods utilized here will be sufficient to alleviate all of these concerns.

Minor concerns:

1) The authors responded, "We thank the reviewer for pointing out that behavioral state is a confusing term given that we also have hidden states. We want to clarify that we were using behavioral states to mean the behaviors themselves, therefore now we are referring to them as behaviors or behavioral labels" However, the text still refers to 'Behavioral states' in multiple places instead of behaviors (line 112 and throughout Fig 2 legend).

2) Line 257: '5 ms pulses at 100 Hz every 200 ms.' This should say FOUR 5 ms pulses… consistent with the legend. Otherwise, its confusing given the duty cycle.

3) The authors have further clarified the utility of AlphaTracker demonstrating that it 'surpasses human accuracy and has lower error rate.' Note, in the context of this manuscript, it is sufficient to just say that it surpasses human accuracy.

Finally, massaging the term 'YOLO' into a scientific manuscript is pure genius. Congrats to the authors on this brilliant work, and I do hope that they can rapidly address the remaining issues regarding the modeling description so this important work can be rapidly shared with the field.

Referee #2 (Remarks to the Author):

The authors have provided an outstanding, complete, and scholarly rebuttal, with a substantial effort to address each and every concern. The authors have done a commendable job addressing all of my concerns. I only have minor concerns remaining or comments, I do not need to see the paper again, until publication.

Minor-

-The authors claim that the wireless device for recording is 6 grams? Is this correct. Generally a mouse can tolerate and move freely with no more than 10% body weight on their head. Assuming these are about 25g mice, this seems like an incorrect measurement? Perhaps it isn't, but it might be explained how the mice were habituated (if it hasn't been already, and I missed it), and/or they should check the weight to be certain, as it might cause some confusion for readers.

-Extended Fig 6, is difficult to digest as labeled. Perhaps more of a "heat map" approach for the top figures, and/or more labels and/or a key would help, since the legend is sparse. Minor style issue, up to the authors and editors discretion.

-The authors should double check all figures, particularly in the extended for typos and errors. There seem to be some stylistic (font tilts or stretching during conversions, etc), and typos that appear in the supplemental.

Referee #3 (Remarks to the Author):

The authors have thoroughly addressed all of my original comments with direct data analyses, data visualizations, and amending the manuscript's texts and figures. I very much look forward to seeing this work published and seeing new exciting future studies from the broader field encouraged by this work.

Signed below to opt in for Nature's new transparent peer review scheme: Reviewed by Steve W. C. Chang with assistance from Philip T. Putnam, a postdoctoral associate in his lab.

Author Rebuttals to First Revision:

We thank all three reviewers for their very helpful, thoughtful and constructive comments, and want to say a special thank you to Referee #1 who caught a very important error in one of our plots, and we are exceedingly grateful for this catch!

Referees' comments:

Referee #1 (Remarks to the Author):

I am a behavioral neurophysiologist, and I re-reviewed this paper in conjunction with a computational neuroscientist/neurotheorist. Comments reflect our pooled suggestions to primarily aid the authors in improving the rigor and reproducibility of their work and enhancing communication of their findings for a broad scientific audience, and to secondarily aid in editorial consideration of this work for publication.

I am highly supportive of this work, and I believe that it provides an important advance that is of broad interest to the field. In my original review, I provided a rich feedback that I believed would enhance the reproducibility of this outstanding work. The authors were incredibly responsive. I also suggested that the authors simplify the work, because I believed that even a simplified version would still provide an important advance that is of broad interest to the field. The authors generally chose to take the opposite approach, which was the provide substantially more experimental details and evidence to support their work. It was a great choice on their part. They have addressed nearly all of my concerns. However, a major concern remains unaddressed.

Major Concern: I previously commented that "The results on the modeling are unusual, where gains between the autoregressive model compared to the SVM and the GLM are drastic. It is rare that you see such dramatic gains when moving to an HMM based model, and the manuscript would be much improved by explaining this outcome. In fact, the authors show that there are drastic gains by using 2 clusters in the HMM based model (noting that the 1 cluster case is equivalent to the GLM model in your formulation). In this case, they are now distinguishing between 9 classes based on a mixture of 2 GLMs, where the GLM is chosen through the historical information. Compared to the classes, this is really a minor gain in modeling complexity, and I‟m unclear from a mathematical perspective as to how such gains can be achieved. Additional justification and exploration would be warranted to ensure robustness of these findings."
In response the authors provided additional background and reference to supplementary materials. But it is important to clearly articulate why my concern regarding their approach remains. While I appreciate the extra information on the model description, the problem remains in the log-likelihoods of Figure 2(c) persists. First, as they have described it, the authors are modeling discrete behaviors. It is well-known that loglikelihoods of discrete

outcomes MUST be negative, but the reported log-likelihoods for the HMM are positive for more than 2 states. Second, it is not explicitly stated what normalization is used in the log-likelihoods, or at least I could not find it. The typical normalization used in machine learning is to normalize by the number of instances. As such, a baseline random guess with probability 1/9 for each outcome would give a loglikelihood of -log(9)~=-2.2. Hence, I would expect all the reported loglikelihoods to be in the range of -2.2 to 0 (since a perfect model would give 0). In contrast, the reported normalized log-likelihoods are -12000 to positive 60. The GLM gives a likelihood that is close to infinitely worse than random guessing, whereas Figure 2(d) shows that the GLM is clearly above chance. Thus, the loglikelihoods must be properly explained in order to evaluate to robustness of the result. I can believe that the temporal history helps to make predictions, so the AUC performances seem believable by themselves; however, given the issues in the loglikelihoods I remain concerned that another group would not be able to effectively reproduce the findings.

I want to commend Reviewer #1 for their exceptionally careful and thoughtful review, and to express our deep appreciation for the level of scrutiny they applied to this critical issue.

We thank reviewer#1 for asking for further clarification on this important point, and for theirvery careful eye. Going back to the calculations, we realized that the way we plotted the static models and HMM-GLM data together was misleading. For the static models we were reporting the sum of the likelihoods across trials while for the HMM GLM we were reportingthe average across trials because for the static models, we compute it for each timepoint (bin-by-bin) and then combine them at the end, whereas for the HMM-GLM, we use a special dynamic programming algorithm because of the nature of the model that has dependency on time (Markovian property) and hidden state(s). In contrast, the static modelsassume every time bin is independent, making it impossible to calculate the likelihoods in an equivalent way. Erring on the side of caution with these insurmountable differences, we removed the likelihood values for the static models given that the main purpose of the likelihood calculations is to select the optimal number of hidden states.

It is worth mentioning that once we corrected that, the values were above chance, as the reviewer #1 expected. However, given the very different techniques for computing the likelihoods for the static model vs the HMM-GLM we decided not to report that comparison.

Also, a clarification note about the normalization process: it was done with respect to a chance model for which, the likelihood of a trial would be: (-) (9 behaviors and 70 bins) which is equal to and we normalized by subtraction of this number. So, all the previously reported log likelihoods should be in the range of negative infinity to +153.8057. Cases with positive loglikelihoods values meant that the performance was better than chance. However, since the theoretical chance value assumes that each behavior is

equally likely, it does not account for the class imbalance in the data. Therefore, we have decided to forgo the normalization.

We apologize for the confusion, thank you for your careful eye. Now Extended Figure Data 5j only shows the likelihood values for the HMM-GLM model without normalization and without directly comparing it to the static model likelihoods (see plots below for before and after comparison). We thank the reviewer very much for pointing out this crucial detail. This has been a learning opportunity for how to deal with likelihood computations of models that are very different.

Old plot shown in Fig 2c

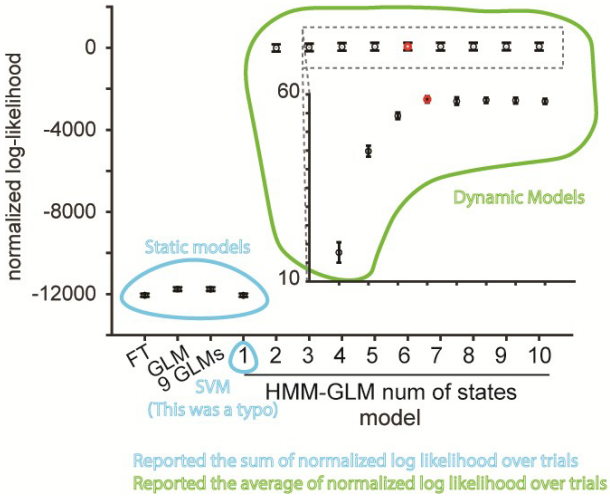

New plot shown now in ED 5j

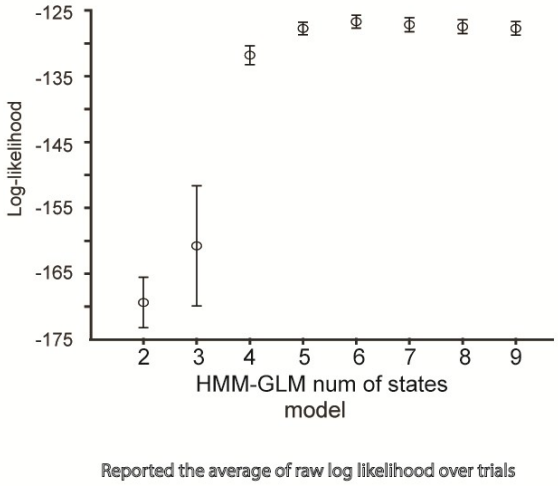

As a minor detail, an HMM-GLM with one state is identical to the GLM. I also would like to understand why those two models do not give the same log-likelihoods despite being mathematically the same model.

We accidentally labeled the SVM likelihood as HMM-GLM state 1. This was a typo. We did not run an HMM-GLM assuming only 1 state because (as Reviewer #1 correctly points out) by definition an HMM must have multiple states. However, in general if you were to run a single GLM multiple times they might not give identical log-likelihoods because of the stochasticity in the individual runs such as random initialization of weights. We thank the reviewer catching this oversight.

Finally, the comparison approach on the mixture of GLMs isn‟t particularly relevant. A typical trick in time series predictions is just to aggregate the last few time steps to form the predictors for a GLM, which would be a much more relevant comparison model. I am not too concerned about this, since the major issue is the reported loglikelihoods.

Another minor issue is in the response on counting the GLMs. The authors state that a 2-state GLM HMM consists of 22 GLMs, which is inconsistent with typical statistics and

machine learning definitions. The multinomial logistic model is considered a single GLM, so in typical nomenclature there are only 2 behavior emissions GLMS (one for each state).

Thank you for this comment! We have now clarified in the methods whether a GLM is binomial or multinomial to prevent any confusions.

I do hope that additional description of the methods utilized here will be sufficient to alleviate all of these concerns.

Minor concerns:
1) The authors responded, "We thank the reviewer for pointing out that behavioral state is a confusing term given that we also have hidden states. We want to clarify that we were usingbehavioral states to mean the behaviors themselves, therefore now we are referring to themas behaviors or behavioral labels" However, the text still refers to „Behavioral states" in multiple places instead of behaviors (line 112 and throughout Fig 2 legend).
2) Line 257: „5 ms pulses at 100 Hz every 200 ms." This should say FOUR 5 ms pulses… consistent with the legend. Otherwise, its confusing given the duty cycle.
3) The authors have further clarified the utility of AlphaTracker demonstrating that it „surpasses human accuracy and has lower error rate." Note, in the context of this manuscript, it is sufficient to just say that it surpasses human accuracy.

Thank you we have fixed these minor details

Finally, massaging the term 'YOLO' into a scientific manuscript is pure genius. Congrats to the authors on this brilliant work, and I do hope that they can rapidly address the remaining issues regarding the modeling description so this important work can be rapidly shared with the field.

Thank you for the compliment, but we cannot claim credit for this bit of genius! Redmon and colleagues came up with YOLO in this original research paper, https://arxiv.org/abs/1506.02640, which is now widely-used as a tool in computer vision.

Referee #2 (Remarks to the Author):

The authors have provided an outstanding, complete, and scholarly rebuttal, with a substantial effort to address each and every concern. The authors have done a commendable job addressing all of my concerns. I only have minor concerns remaining or comments, I do not need to see the paper again, until publication.

Minor-

-The authors claim that the wireless device for recording is 6 grams? Is this correct. Generally a mouse can tolerate and move freely with no more than 10% body weight on their head. Assuming these are about 25g mice, this seems like an incorrect measurement? Perhaps it isn't, but it might be explained how the mice were habituated (if it hasn't been already, and I missed it), and/or they should check the weight to be certain, as it might cause some confusion for readers.

We have added details about the habituation of wearing the device in the methods. Mice were closer to ~30g and as discussed in ample detail in the first round of revision, wearing the recording devices did alter the behavioral patterns of mice.

-Extended Fig 6, is difficult to digest as labeled. Perhaps more of a "heat map" approach for the top figures, and/or more labels and/or a key would help, since the legend is sparse. Minor style issue, up to the authors and editors discretion.

In the current version of this manuscript we chose to remove these plots.

-The authors should double check all figures, particularly in the extended for typos and errors. There seem to be some stylistic (font tilts or stretching during conversions, etc), and typos that appear in the supplemental.

Referee #3 (Remarks to the Author):

The authors have thoroughly addressed all of my original comments with direct data analyses, data visualizations, and amending the manuscript‟s texts and figures. I very much look forward to seeing this work published and seeing new exciting future studies from the broader field encouraged by this work.
Signed below to opt in for Nature‟s new transparent peer review scheme: Reviewed by Steve W. C. Chang with assistance from Philip T. Putnam, a postdoctoral associate in his lab.

Thank you both very much for your very helpful constructive feedback!