Distributed range adaptation in human parietal encoding of numbers
Supplementary Information

Arthur Prat-Carrabin[1,⋆], Gilles de Hollander[2,3⋆], Saurabh Bedi[2,3], Samuel J. Gershman[1], and Christian C. Ruff[2,3]

[1]Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA, USA
[2]Zurich Center for Neuroeconomics, Department of Economics, University of Zurich, Switzerland
[3]University Research Priority Program (URPP), Adaptive Brain Circuits in Development and Learning, University of Zurich, Zurich, Switzerland

⋆equal contribution
⋆*arthurpc@fas.harvard.edu, gilles.de.hollander@gmail.com*

1 Qualitative assessment of nPRF fits

Numerosity-selective population receptive fields (nPRFs) are expected to be considerably noisier than standard visuospatial PRFs, for two reasons: numerosity tuning is much less pronounced in terms of signal-to-noise ratio than retinotopic tuning, and we used an event-related rather than a block design (as is common in visuospatial PRF studies). To nonetheless provide a qualitative sense of model fit quality, we randomly selected 12 voxels from the top 60 voxels ranked by normalised percent signal change (NPCr) with non-monotonic tuning in three example participants (S12, S20, and S26). For each voxel, we plot the mean BOLD response to each numerosity (averaged over ∼10 trials) alongside the best-fitting model prediction (efficient shift and fixed increase in dispersion).

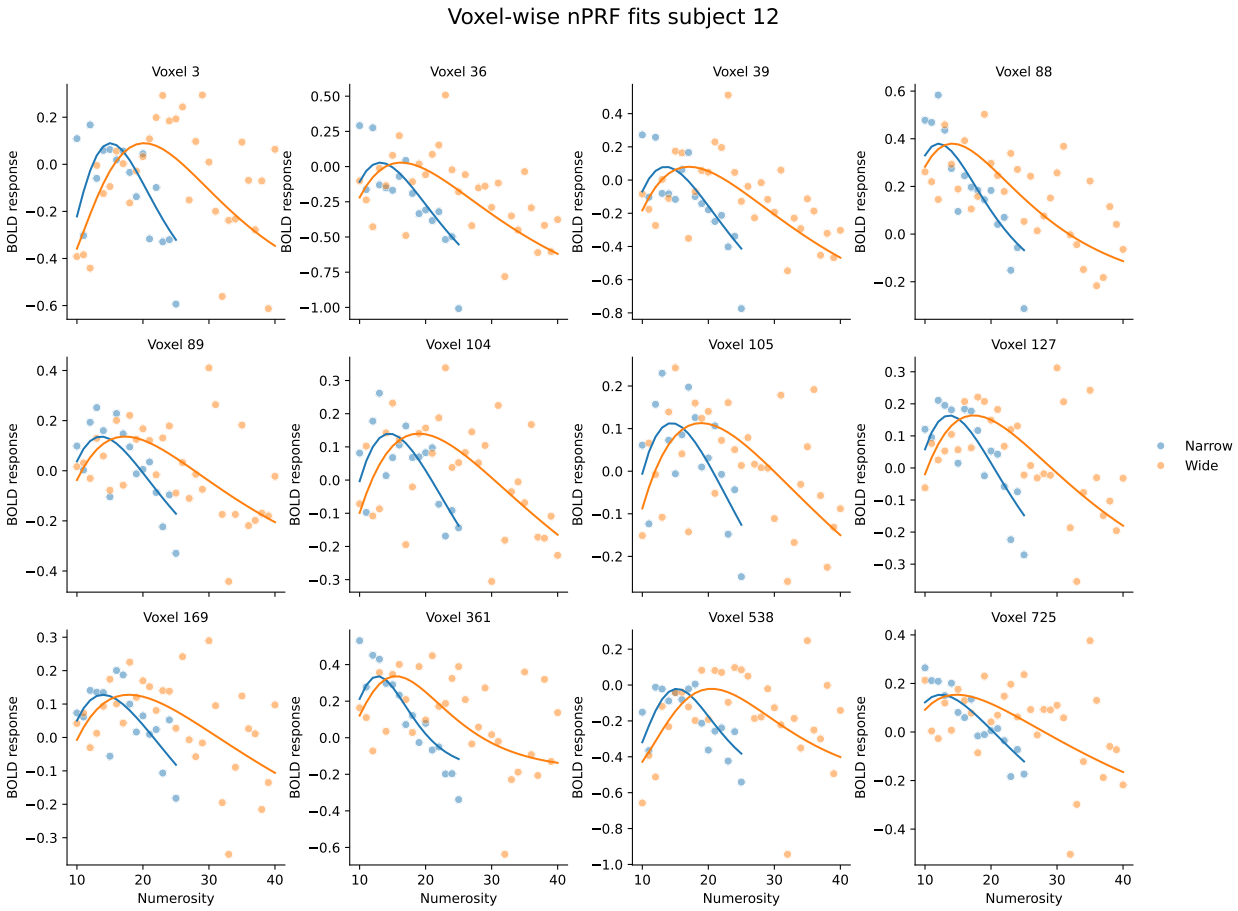

Fig. S1: Qualitative nPRF model fits for participant 12. Each panel shows the mean BOLD response (filled circles) as a function of presented numerosity, separately for the Narrow (blue) and Wide (orange) conditions, alongside the fitted nPRF model prediction (solid lines). Voxels were selected from the top 60 NPCr voxels with non-monotonic tuning.

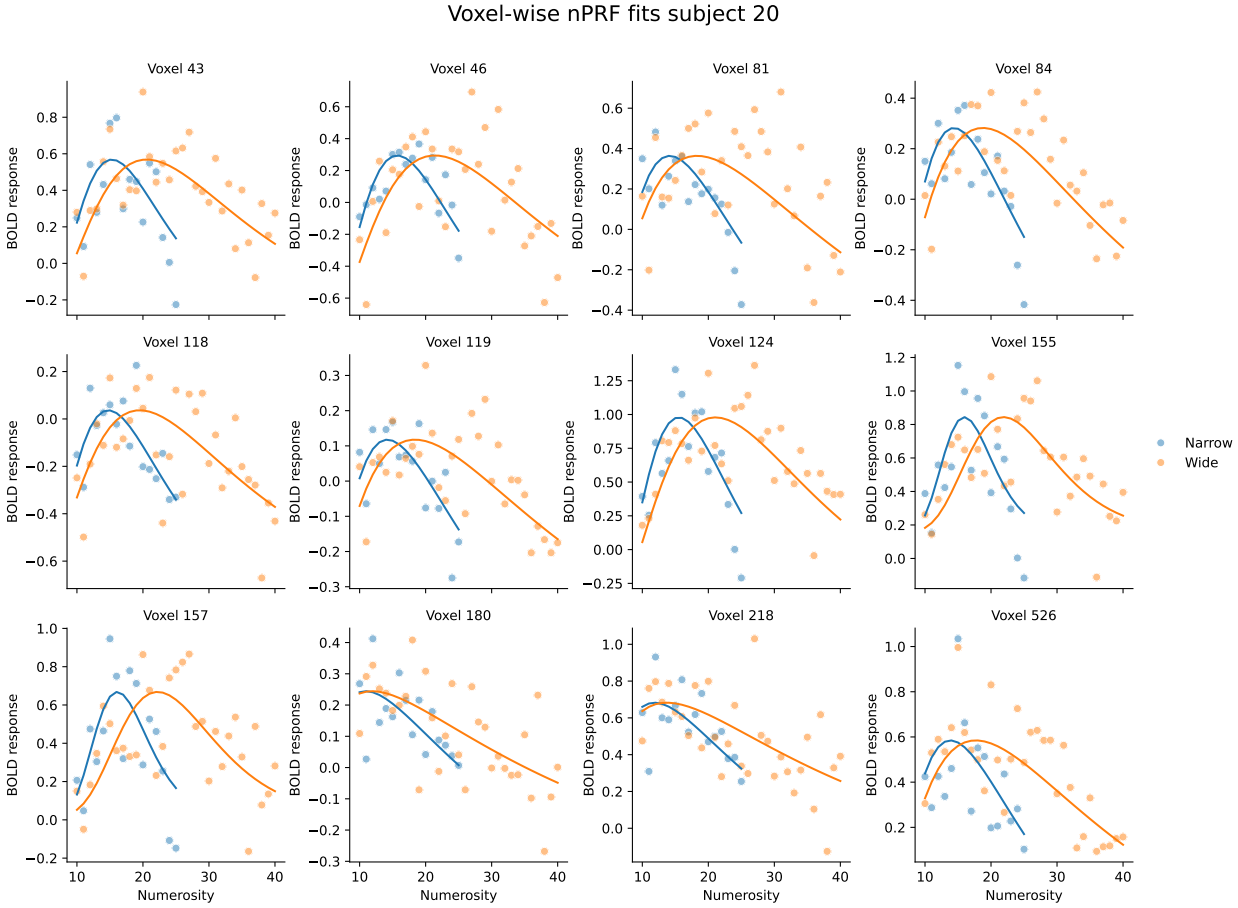

Fig. S2: Qualitative nPRF model fits for participant 20. As Fig. S1, for participant 20.

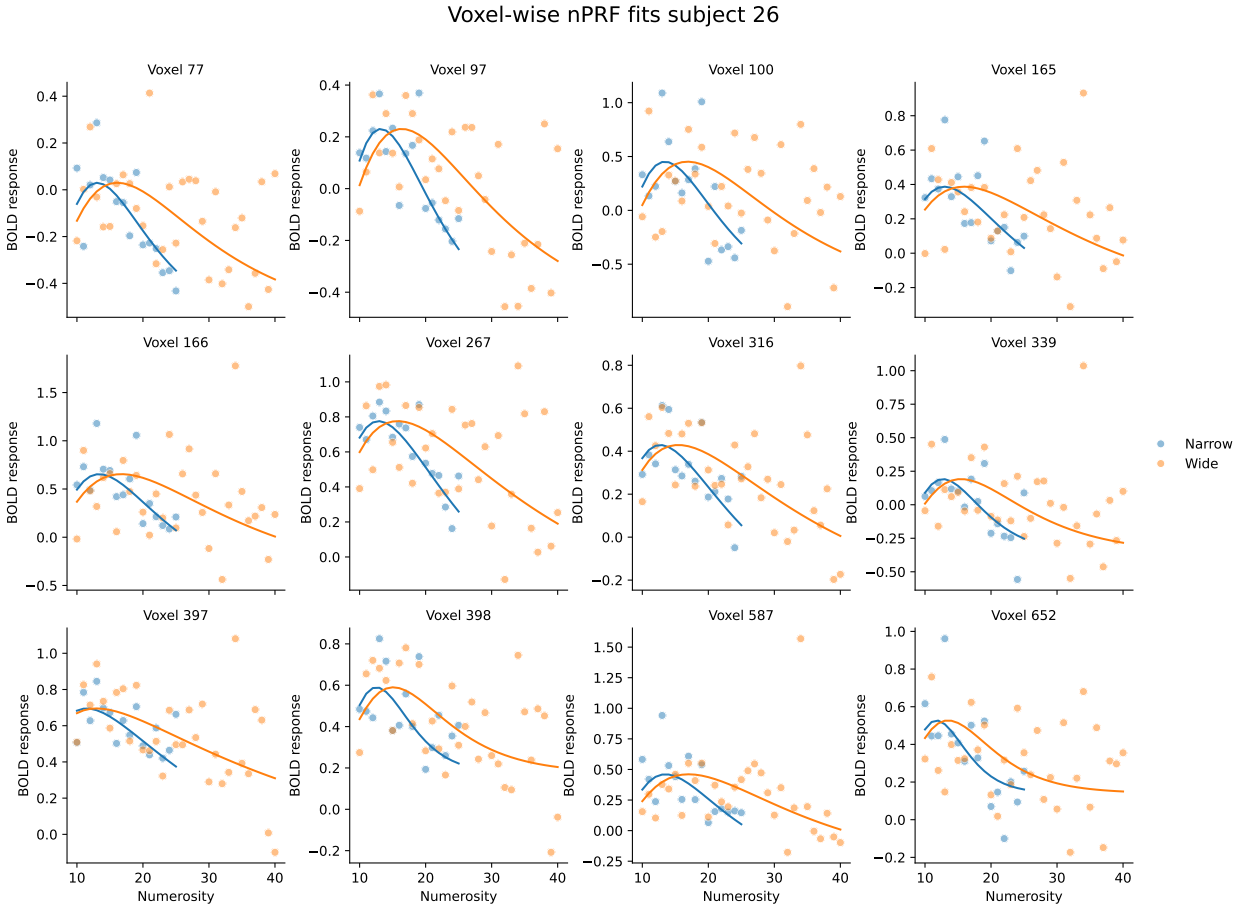

Fig. S3: Qualitative nPRF model fits for participant 26. As Fig. S1, for participant 26.

2 Model comparisons

Figure S4 shows the proportion of voxels with cv$R^2 > 0$ with the different nPRF models that we fit. Models fit on the participants' responses (orange line) consistently outperform the models fit on the ground-truth, presented number (teal line).

Models with amplitude changes

The precision of the neural encoding depends on the amplitude of the populations' receptive fields, in addition to their preferred numerosities and their widths. Here we test for changes in the amplitudes. Thus we fit models in which the amplitude of each voxel may vary across conditions. Specifically, first we consider a model that features efficient shifts of the preferred numerosities, fixed scaling of the widths (as in the main text), and in which for each voxel we let the amplitudes in the two conditions be free parameters, a_n and a_w. This model yields a significantly lower fit than our best-fitting model (with estimates: $t(38) = 6.12$, $P = 3.9 \times 10^{-7}$; with correct numbers: $t(38) = 6.41$, $P = 1.6 \times 10^{-7}$; Fig. S4, model labeled 'Eff. shifts, fixed width-scaling, free amplitudes'). Second, we enforce a relation of proportionality between each participant's voxels amplitudes in the Narrow condition and in the Wide condition, with a fixed ratio for each participant (i.e., $a_w = r_a a_n$). This model yields a better fit than the previous one, but a significantly worse one than our best-fitting model (with estimates: $t(38) = 2.36$, $P = 0.024$; with correct numbers: $t(38) = 2.71$, $P = 0.0099$; Fig. S4, model labeled 'Eff. shifts, fixed width-scaling, p.-sp. amplitude-scaling').

We also consider a 'Participant-specific amplitude-scaling' model, which allows for scaled changes in the amplitudes, while keeping the other parameters constant (rightmost model in Fig. S4). Inspection of all models performance shows that this model does not provide a better account of the data than the efficient-shift model; rather, the evidence points in the opposite direction. With the response-based fit, the 'Participant-specific shifts' model ($\mu_w = 10 + r_\mu(\mu_n - 10)$, with other parameters constant; fifth model) yields a numerically higher proportion of voxels with positive cvR^2 than the 'Participant-specific amplitude-scaling' model, and it fits significantly better than the 'No change' model (in which all the parameters are kept identical across conditions; $t(38) = 2.33$, $P = 0.025$). The same result obtains when considering the median cvR^2 (see below). By contrast, the 'Participant-specific amplitude-scaling' model does not fit better than the 'No change' model ($t(38) = 1.51$, $P = 0.14$), and it performs significantly worse than the 'Efficient shifts' model (with constant width and amplitude; $t(38) = 2.3$, $P = 0.027$). Thus, allowing for shifts in preferred numerosities offers a better account of the data than allowing for an increase in amplitudes.

Models with alternative shift hypotheses

We consider alternatives to our efficient-shift hypothesis (which posits that $\mu_w = 10 + 2(\mu_n - 10)$ if $\mu_n > 10$, otherwise $\mu_w = \mu_n$). First we consider a model in which for all voxels the preferred numerosity in the Wide condition is twice the preferred numerosity in the Narrow condition (i.e., $\mu_w = 2\mu_n$). We compare this model to the model with efficient shifts

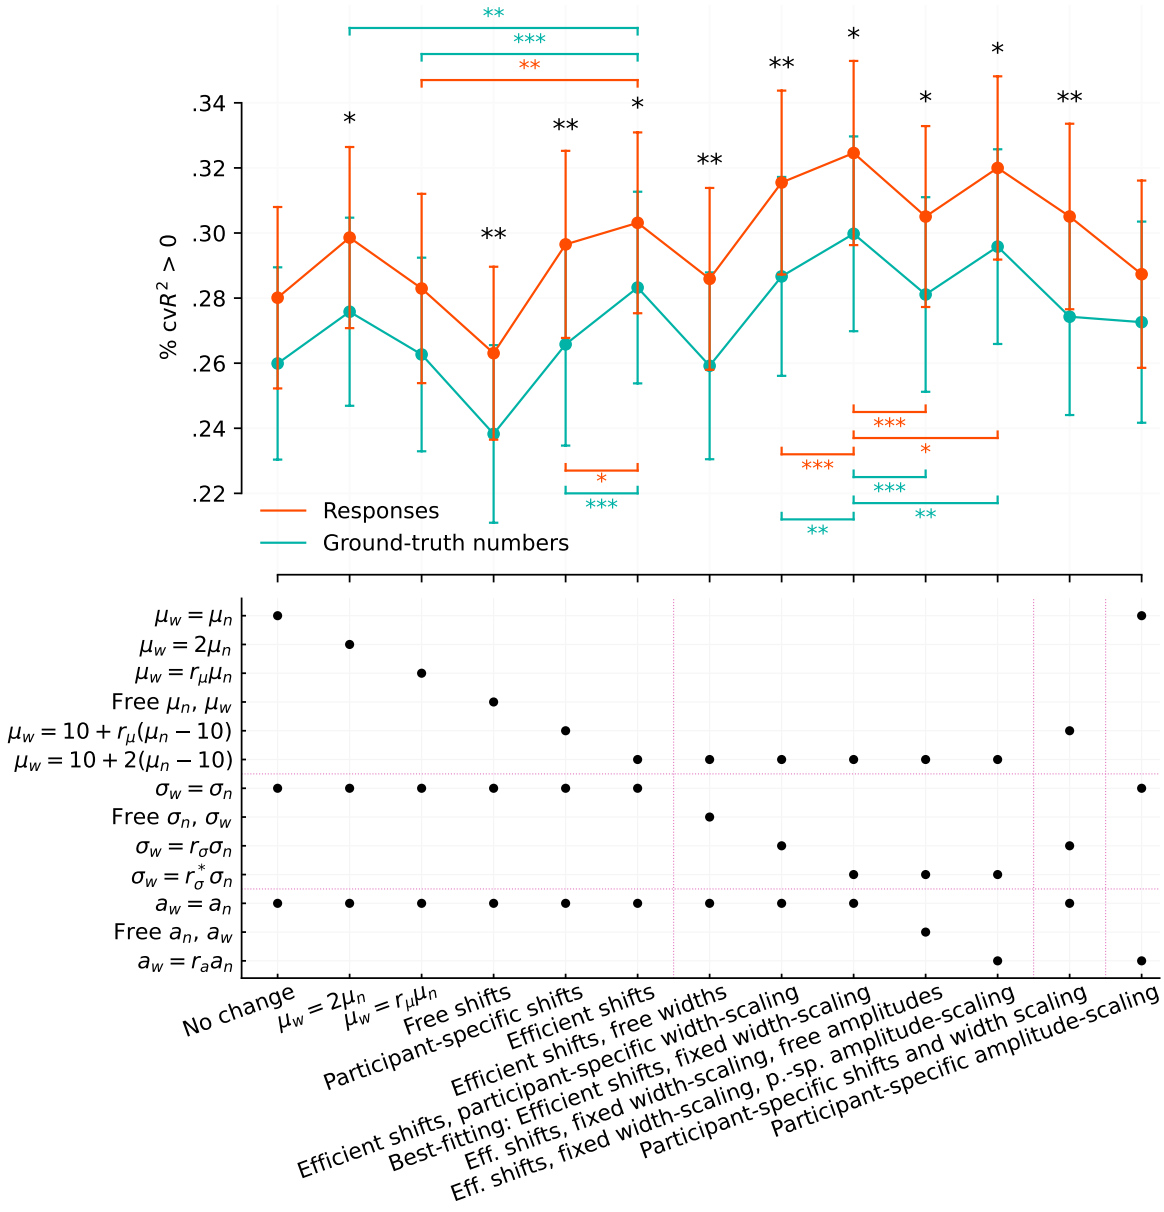

Fig. S4: Proportion of voxels with positive cross-validated variance explained (cvR^2>0) for the different models. Error bars show ± 1 standard error of the mean. The bottom panel indicates the specifications of the models' parameters across priors. Each model (vertical line) has three dots corresponding to the model's specifications regarding the preferred-numerosity parameters (μ_n, μ_w), the width parameters (σ_n, σ_w), and the amplitude parameters (a_n, a_w), indicated on the left. The ratios r_μ, r_σ, and r_a are participant-level parameters, while r^*_σ is fixed at 1.3 (and shared by all participants). "p.-sp.": participant-specific. ***: $P < 0.001$, **: $P < 0.01$, *: $P < 0.05$.

(and unchanging widths and amplitudes): it yields a lower proportion of voxels with cvR^2 positive, with a significant difference when fitting on the correct numbers ($t(38) = 2.63$, $P = 0.012$; Fig. S4, second model). Then we consider a model in which the preferred numerosity in the Wide condition is proportional to that in the Narrow condition, with a fixed parameters for each participant (i.e., $\mu_w = r_\mu \mu_n$; Fig. S4, third model). This model also fits significantly worse than the model featuring the efficient-shift hypothesis (fitting on estimates: $t(38) = 3.19$, $P = 0.003$, fitting on correct numbers: $t(38) = 3.85$, $P = 0.0004$). Thus we conclude that the efficient-shift hypothesis ($\mu_w = 10 + 2(\mu_n - 10)$ if $\mu_n > 10$) better accounts for the data than the hypothesis that μ_w is proportional to μ_n.

Robustness check

To investigate whether our results are robust to the metric used for model comparison, we consider another measure of fit, the median cvR^2. The majority of voxels in the ROI do not contain information about the presented numbers and yield very large negative cvR^2s. We thus compute for each model the median cvR^2 over all the voxels for which at least one model yields a positive cvR^2. Figure S5 shows the results obtained with this metric. The figure is very similar to that obtained when looking at the proportion of voxels with positive cvR^2 (Fig. S4), and calls for the same conclusions. In particular, the ordering of models is identical, and thus the best-fitting model with this analysis is also the 'Efficient-shifts, fixed width-scaling' model.

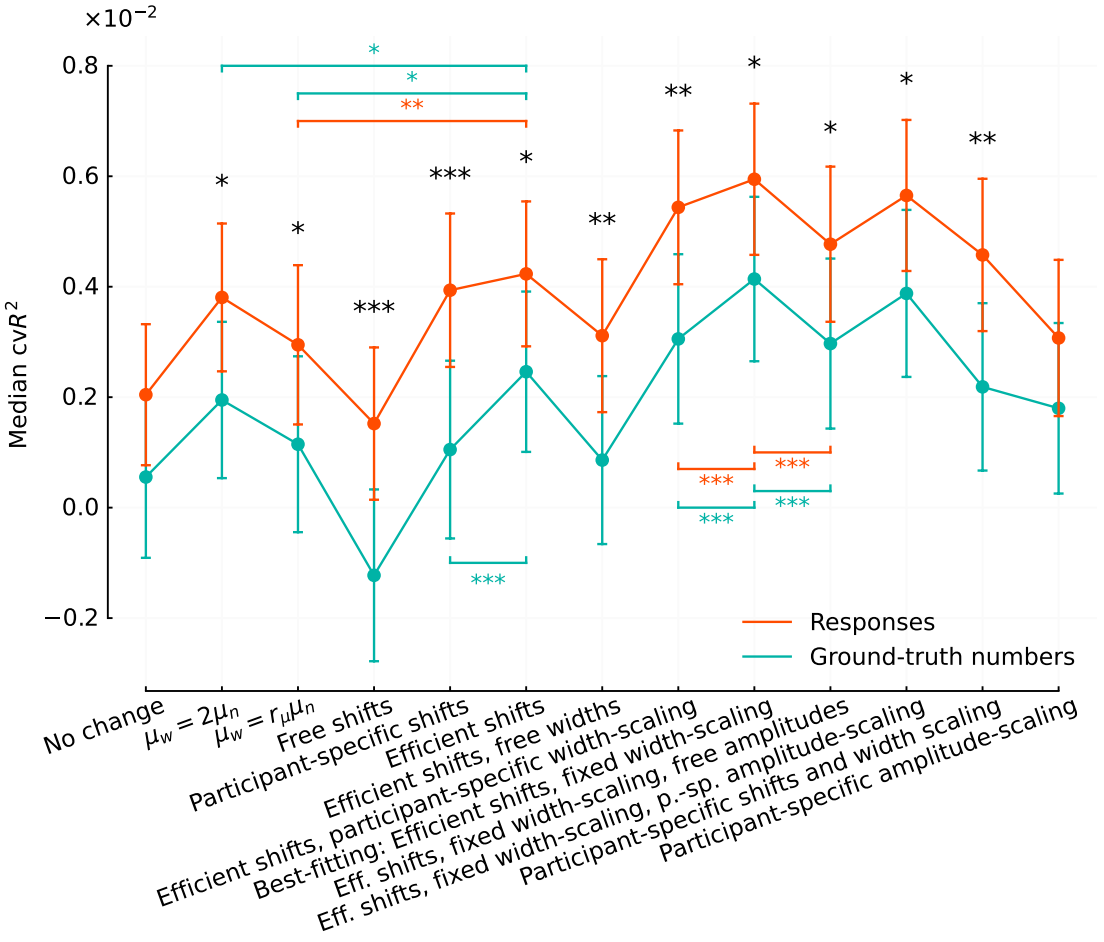

Fig. S5: Median cvR^2 for the different models. Error bars show ± 1 standard error of the mean. ***: $P < 0.001$, **: $P < 0.01$, *: $P < 0.05$.

3 Results when fitting nPRFs with ground-truth numbers

We report the results obtained when fitting the models of population receptive-field with the ground-truth, presented numbers rather than with the participants' responses. Figure S6 is the counterpart of Figure 3 in the main text, and shows the shifts in preferred numerosities across priors. Figure S7 is the counterpart of Figure 4 in the main text, and shows the changes in receptive fields widths across priors. Below, we report for the models fit with the ground-truth numbers the same statistics reported in the main text for the models fit with the responses.

Preferred numerosities

- Correlation between preferred numerosities μ_n and μ_w in the two conditions: participants pooled: Spearman's $\rho = 0.74$, $P < 10^{-320}, N = 7264$; across participants: average $\rho = 0.64$, interquartile range (IQR): 0.57-0.82.
- Across-participant t-test that the mean shift $\mu_w - \mu_n$ is zero, for $\mu_n \geq 10$: $t(38) = 3.45$, $P = 0.0014$.
- For 32 out of 39 participants, the across-voxels median preferred numerosity is larger in the Wide condition than in the Narrow condition (Fig. S6c; across-participants paired t-test of equality of the medians in the two conditions: $t(38) = 4.31$, $P = 1.1 \times 10^{-4}$). Overall, the preferred numerosities of 78% of voxels increase in the Wide condition as compared to the Narrow condition (Fig. S6d).
- μ_w for voxels with $\mu_n \in [14.5, 15.5]$ is on average 20.45; sem: 0.50.
- 2.4% of voxels have preferred numerosities greater than 25 in the Narrow condition.
- 'Efficient-shift' model compared to 'free-shift' model: across-participants paired t-test of equality of the proportions of voxels with cv$R^2 > 0$: $t(38) = 8.34$, $P = 4 \times 10^{-10}$. Compared to 'no-shift' model: $t(38) = 4.21$, $P = 1.5 \times 10^{-4}$. Compared to 'participant-specific slopes' model (lower-right inset in Fig. S6a): $t(38) = 3.86$, $P = 4.3 \times 10^{-4}$. See Figure S6f.
- Slope parameter in the 'participant-specific slopes' model: across-participant median: 1.91, mean: 2.32, sem: 0.24 (Fig. S6e). The mean best-fitting value is significantly different from 1 (t-test $t(38) = 5.59$, $P = 2 \times 10^{-6}$), but not from 2 ($t(38) = 1.36$, $P = 0.18$).

Receptive fields widths

- Correlation between the widths parameters σ_n and σ_w in the two conditions: participants pooled: $r = 0.27$, $P = 8 \times 10^{-131}$, $N = 7977$; across participant: average $r = 0.45$, IQR: 0.25-0.60.
- For 31 out of 39 participants, the median width (across voxels) in the Wide condition is larger than in the Narrow condition (paired t-test of equality: $t(38) = 4.9$, $P = 2 \times 10^{-5}$; Fig. S7c). Overall, the widths of 81% of voxels increase in the Wide condition as compared to the Narrow condition (Fig. S7d).
- Scale parameter of the 'participant-specific scaling' model: across-participants median: 1.3, median: 1.3, sd: 0.25 (Fig. S7e). The mean best-fitting value is significantly greater than 1 (t-test $t(38) = 7.23$, $P = 1.2 \times 10^{-8}$), and significantly lower than 2 ($t(38) = 17.9$, $P = 4 \times 10^{-20}$.

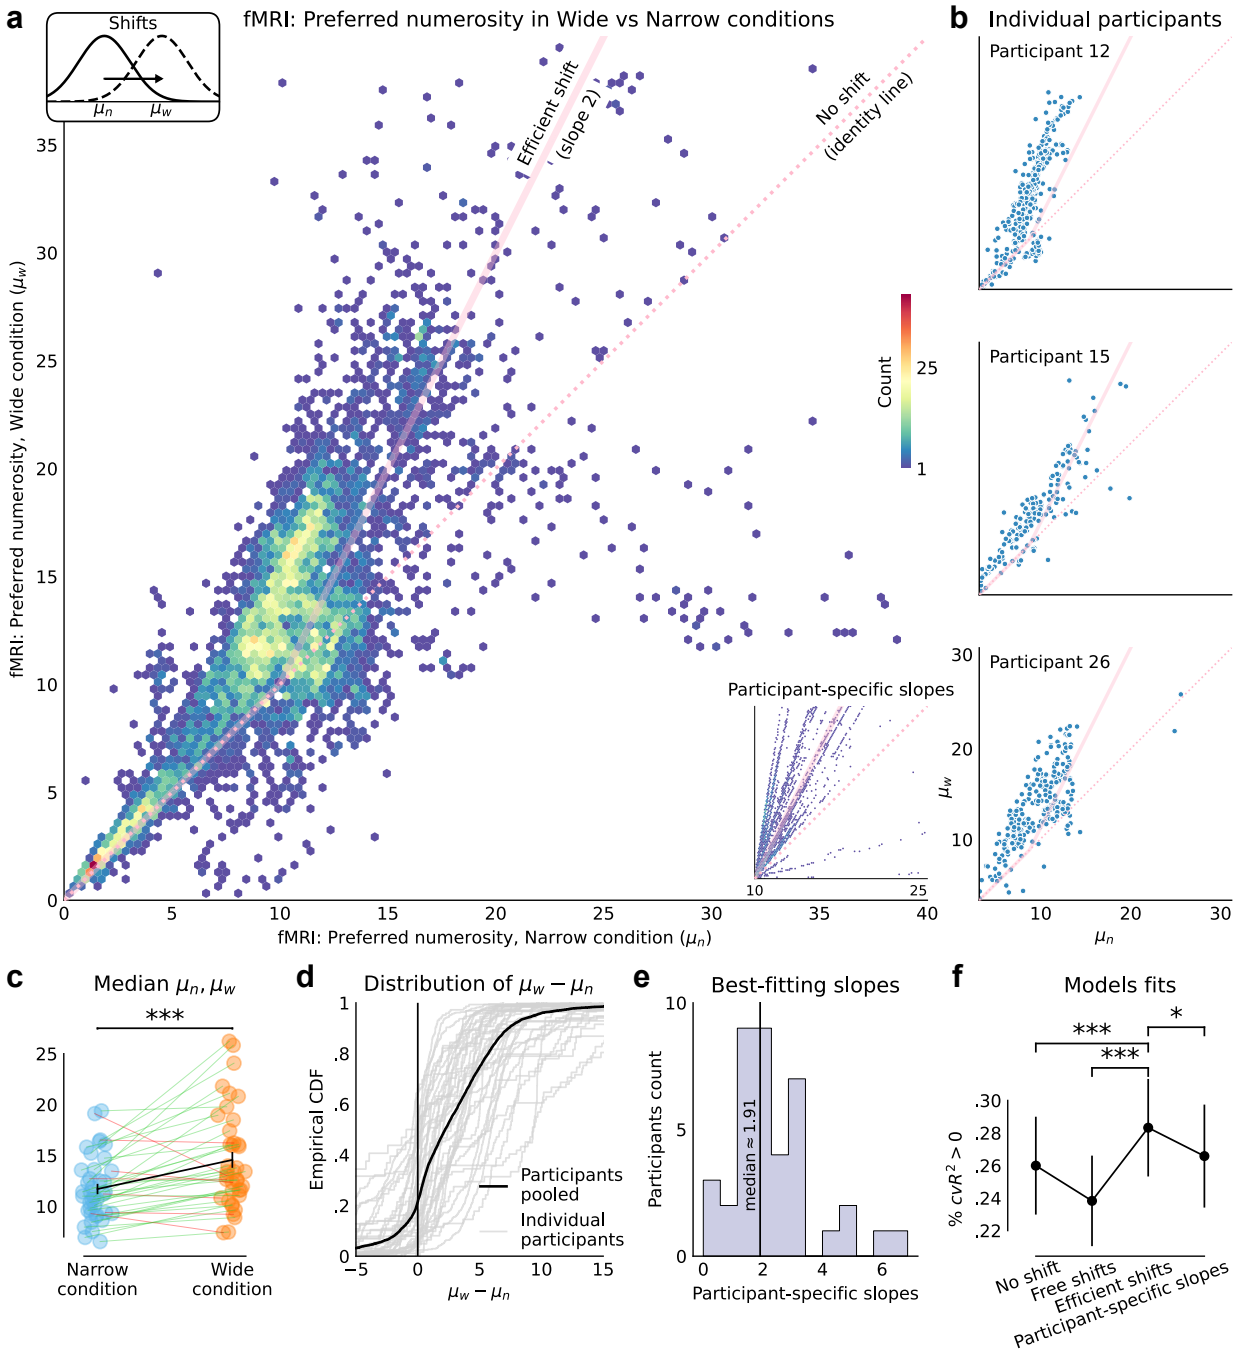

Fig. S6: Ground-truth fit: Efficient shifts of neural receptive fields across priors. As Figure 3 in the main text, but with receptive-field models fit using the ground-truth numbers rather than the responses.

- 'Fixed scaling' model ($r_\sigma = 1.3$) compared to 'participant-specific scaling' model: across-participants paired t-test of equality of the proportions of voxels with cv$R^2 > 0$: $t(38) = 3.5$, $P = 1.3 \times 10^{-3}$. Compared to 'free widths' model: $t(38) = 9.0$, $P = 7 \times 10^{-11}$. Compared to 'no width change' model: $t(38) = 3.4$, $P = 1.6 \times 10^{-3}$. See Figure S7f.

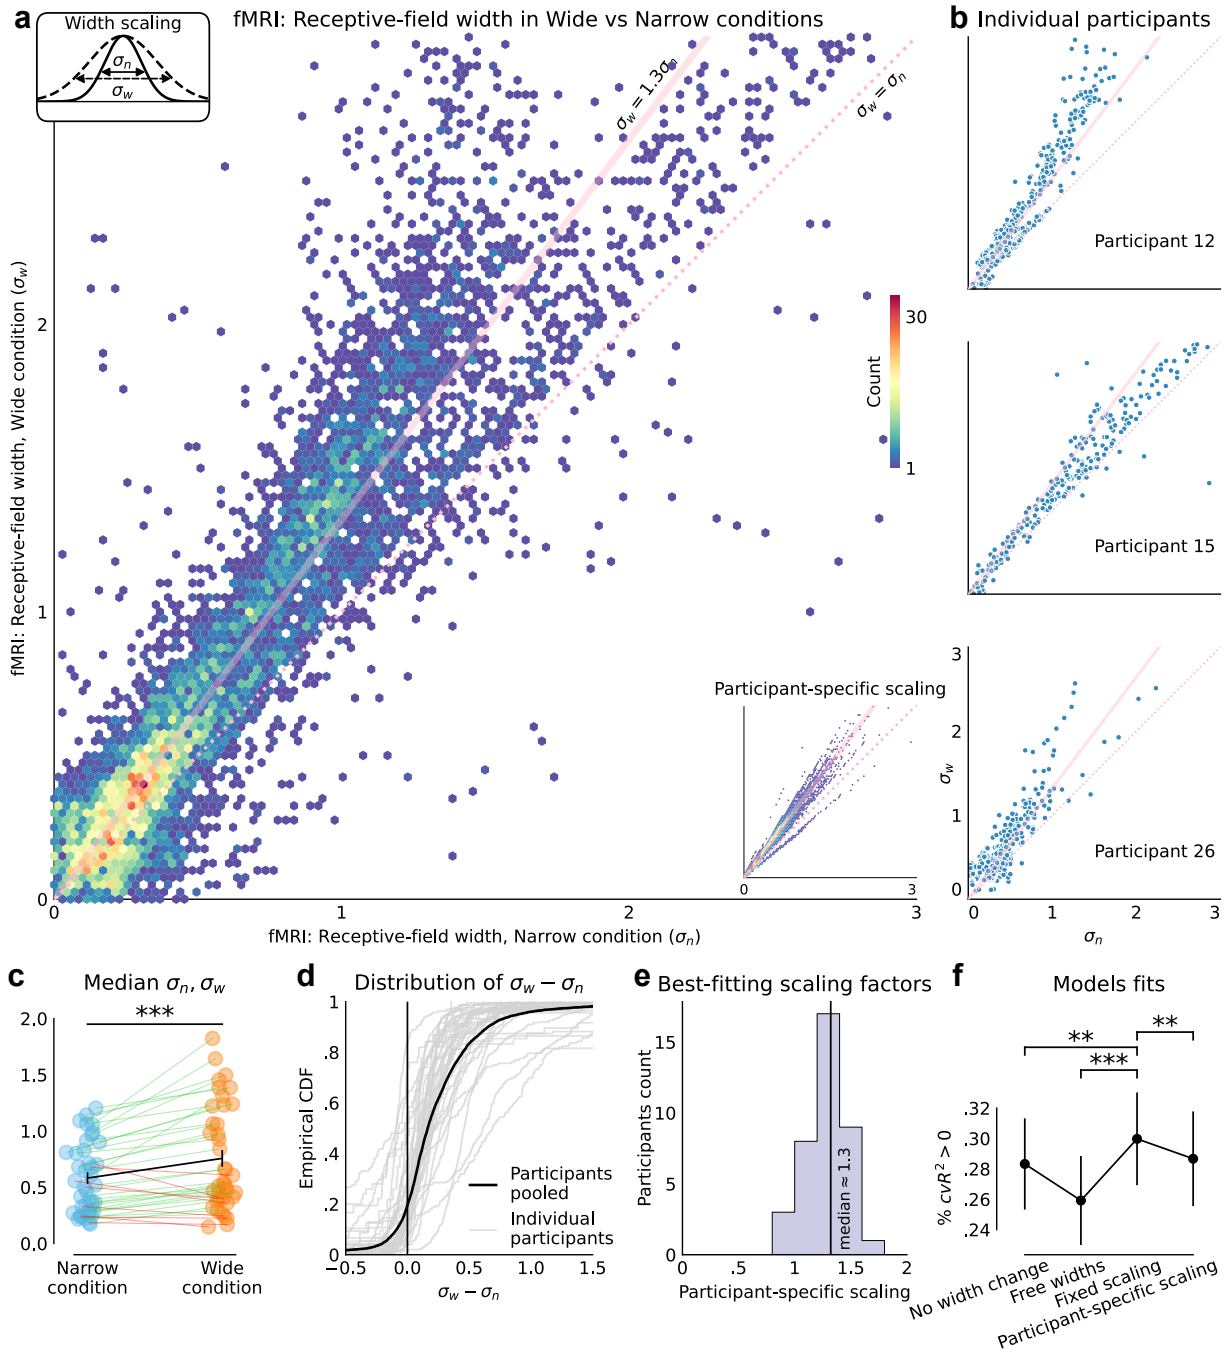

Fig. S7: Ground-truth fit: The receptive fields broaden under the wider range. As Figure 4 in the main text, but with receptive-field models fit using the ground-truth numbers rather than the responses.

In short, the results obtained with the receptive fields models fit with the correct numbers are qualitatively identical, and quantitatively very similar, to those obtained when fitting with the participants' estimates.

4 nPRF estimation checks

Model recovery confirms unbiased nPRF estimation across stimulus ranges

A potential concern with comparing nPRF parameters across conditions with different stimulus ranges is that the model fitting procedure itself might be biased by the range of numerosities presented. To address this, we performed a model recovery analysis (see Methods in main text for details). Briefly, we simulated 250 voxels under a 2×2 factorial design, crossing two generating slopes ($\delta = 1$, no shift; $\delta = 2$, full range adaptation) with two stimulus designs (full, matching our actual experiment; censored, restricted to the shared range $[10, 25]$). We repeated each condition 50 times and applied the same nPRF fitting procedure used in the main analyses.

As shown in Fig. S8, the recovered slope parameter closely matched the generating value in all four conditions. Critically, restricting the stimulus range to the shared interval (censored design) did not introduce any systematic bias relative to the full design, confirming that differences in stimulus sampling between the Narrow and Wide conditions cannot explain the tuning shifts reported in the main text. Code for this simulation is available at https://github.com/ruffgroup/neural_priors/blob/main/neural_priors/revision%20feb%202026/simulate_data.py.

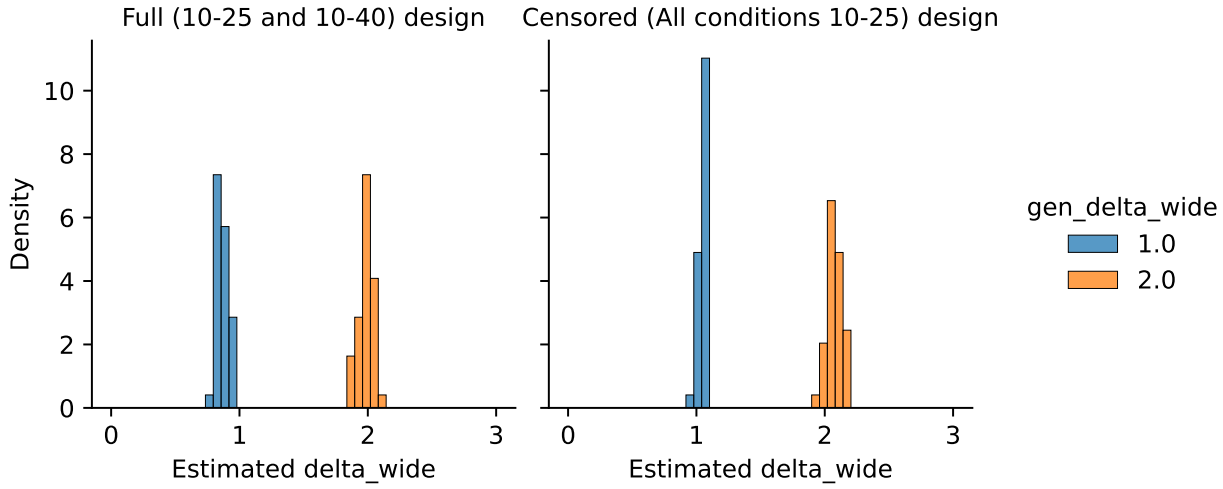

Fig. S8: Model recovery analysis confirms unbiased nPRF parameter estimation across stimulus ranges. Results of the 2×2 factorial simulation study. Each panel shows the distribution of recovered slope parameters (relating μ_{wide} to μ_{narrow}) across 50 simulation repetitions, for generating slopes of 1 (no shift, left) and 2 (full range adaptation, right), and for the full (top) and censored (bottom) stimulus designs. Dashed vertical lines indicate the generating slope value. The recovered slope parameter is close to the generating value in all four conditions, with negligible bias, confirming that nPRF model fitting is not systematically distorted by differences in stimulus range between conditions.

Preferred numerosities and widths with censored fits

One might argue that the shift in preferred numerosities might (partly) arise due to the fact that in the Wide condition, the encoding model is fit to a broader range of stimuli, potentially biasing preferred numerosities toward more extreme values. To rule out this possibility, we re-fit the 'Free shifts' model and the 'Participant-specific shifts' model, restricting the input numerosities to the range shared between the Narrow and Wide conditions (henceforth *censored* fits). Note that these censored fits are necessarily noisier, as the Wide condition then contains only half the amount of data. Reassuringly, the shift in median preferred numerosity remains highly significant under the censored fits (paired t-test: $t(38) = 5.60$, $P = 2 \times 10^{-6}$), and the slope parameter remains distributed (across participants) around 2 (median: 1.71, mean: 2.04, sem: 0.23), with a mean significantly different from 1 ($t(38) = 4.55$, $P = 5 \times 10^{-5}$), but not from 2 ($t(38) = 0.18$, $P = 0.86$), closely mirroring the results from the uncensored fits. Figure S9 shows the obtained preferred numerosities, which exhibit the same patterns as with the uncensored fits. These results confirm that the reported tuning shifts are not an artifact of differential stimulus sampling between conditions.

We conduct the same censored-fits analysis for the width parameter. Specifically, we fit the 'Efficient shifts, free widths' model and the 'Efficient shifts, participant-specific width-scaling' model on data restricted to the range shared between the Narrow and Wide conditions. Figure S10 shows the obtained widths, which exhibit the same patterns as with the uncensored fits. With the free-widths model, the two widths are significantly correlated (participants pooled: Spearman's $\rho = 0.80$, $P < 10^{-320}$, $N = 6586$; across participant: average $\rho = 0.66$, IQR: 0.59-0.85). For a majority of participants (30 out of 39), the median width (across voxels) in the Wide condition is larger than in the Narrow condition (paired t-test of equality: $t(38) = 3.74$, $P = 6 \times 10^{-4}$; Fig. S10c). Overall, the widths of 68% of voxels increase in the Wide condition as compared to the Narrow condition (Fig. S10d).

With the model with participant-specific width-scaling, the scaling parameter r_σ is significantly greater than 1 (t-test $t(38) = 4.16$, $P = 1.8 \times 10^{-4}$), and significantly lower than 2 ($t(38) = 8.6$, $P = 1.8 \times 10^{-10}$). Its across-participant average is 1.3 (median: 1.2, sem: 0.08; Fig. S10e). These results, consistent with those obtained with the uncensored fits, attest to the robustness of our parameter estimation, and substantiate our conclusions.

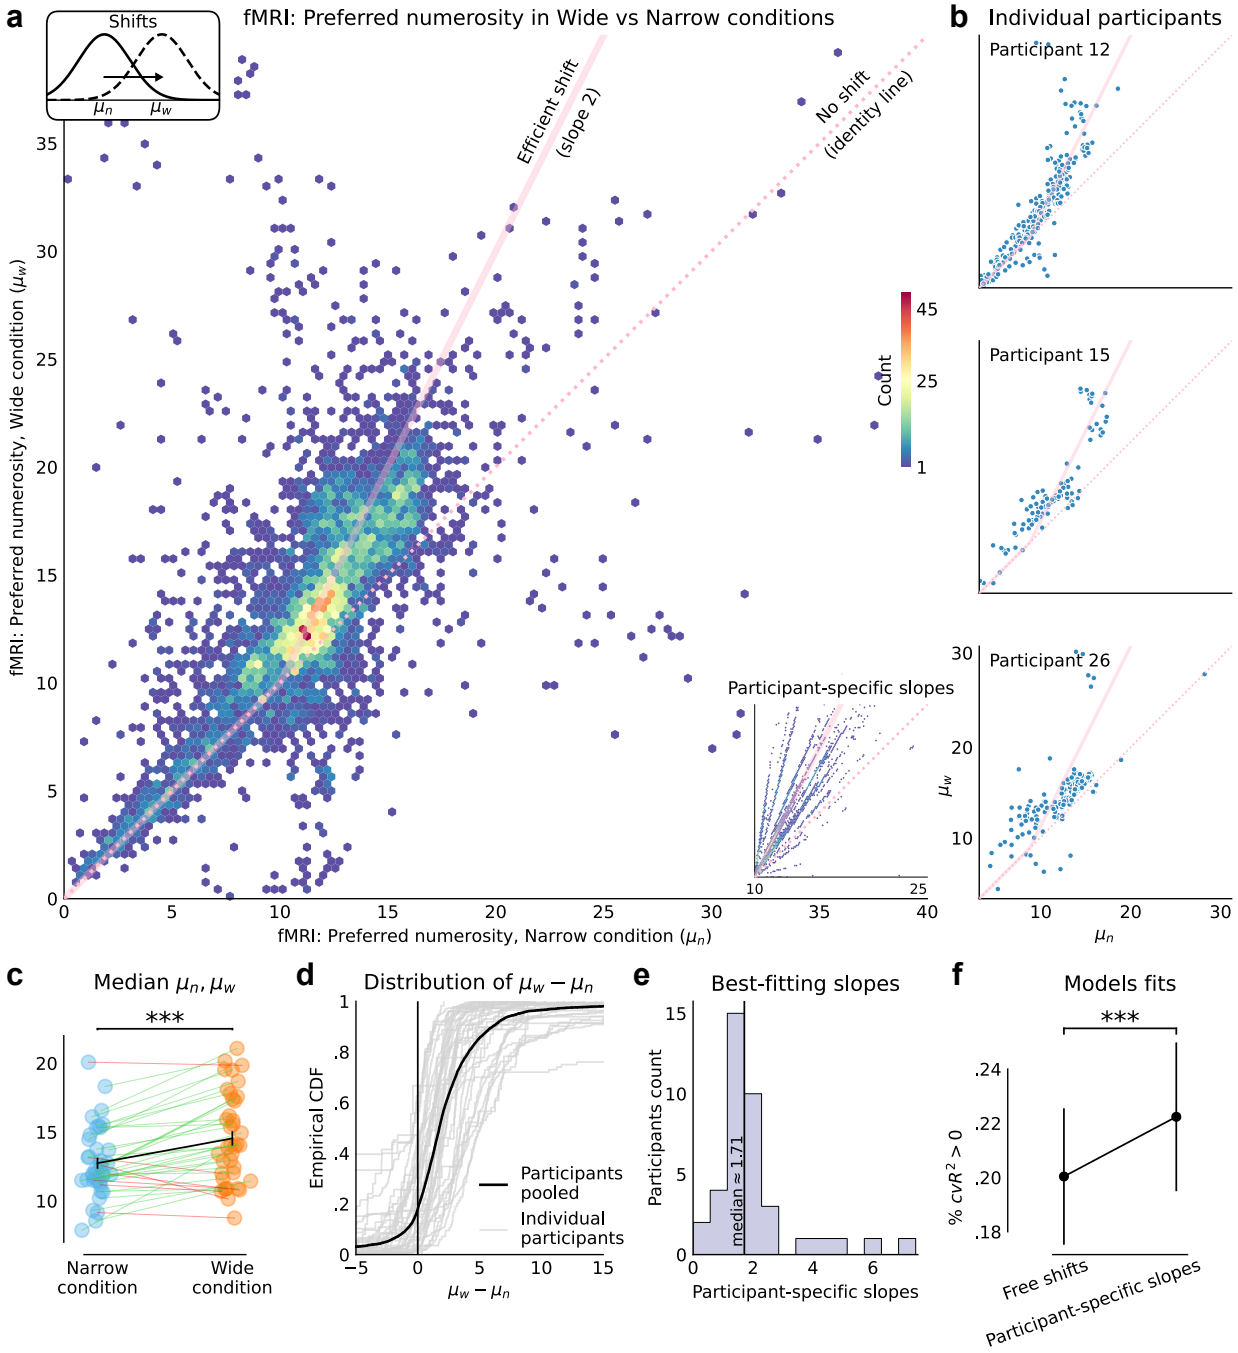

Fig. S9: Censored fits: preferred numerosities. As Figure 3 in the main text, but with fits restricted to the range shared between the Narrow and Wide conditions.

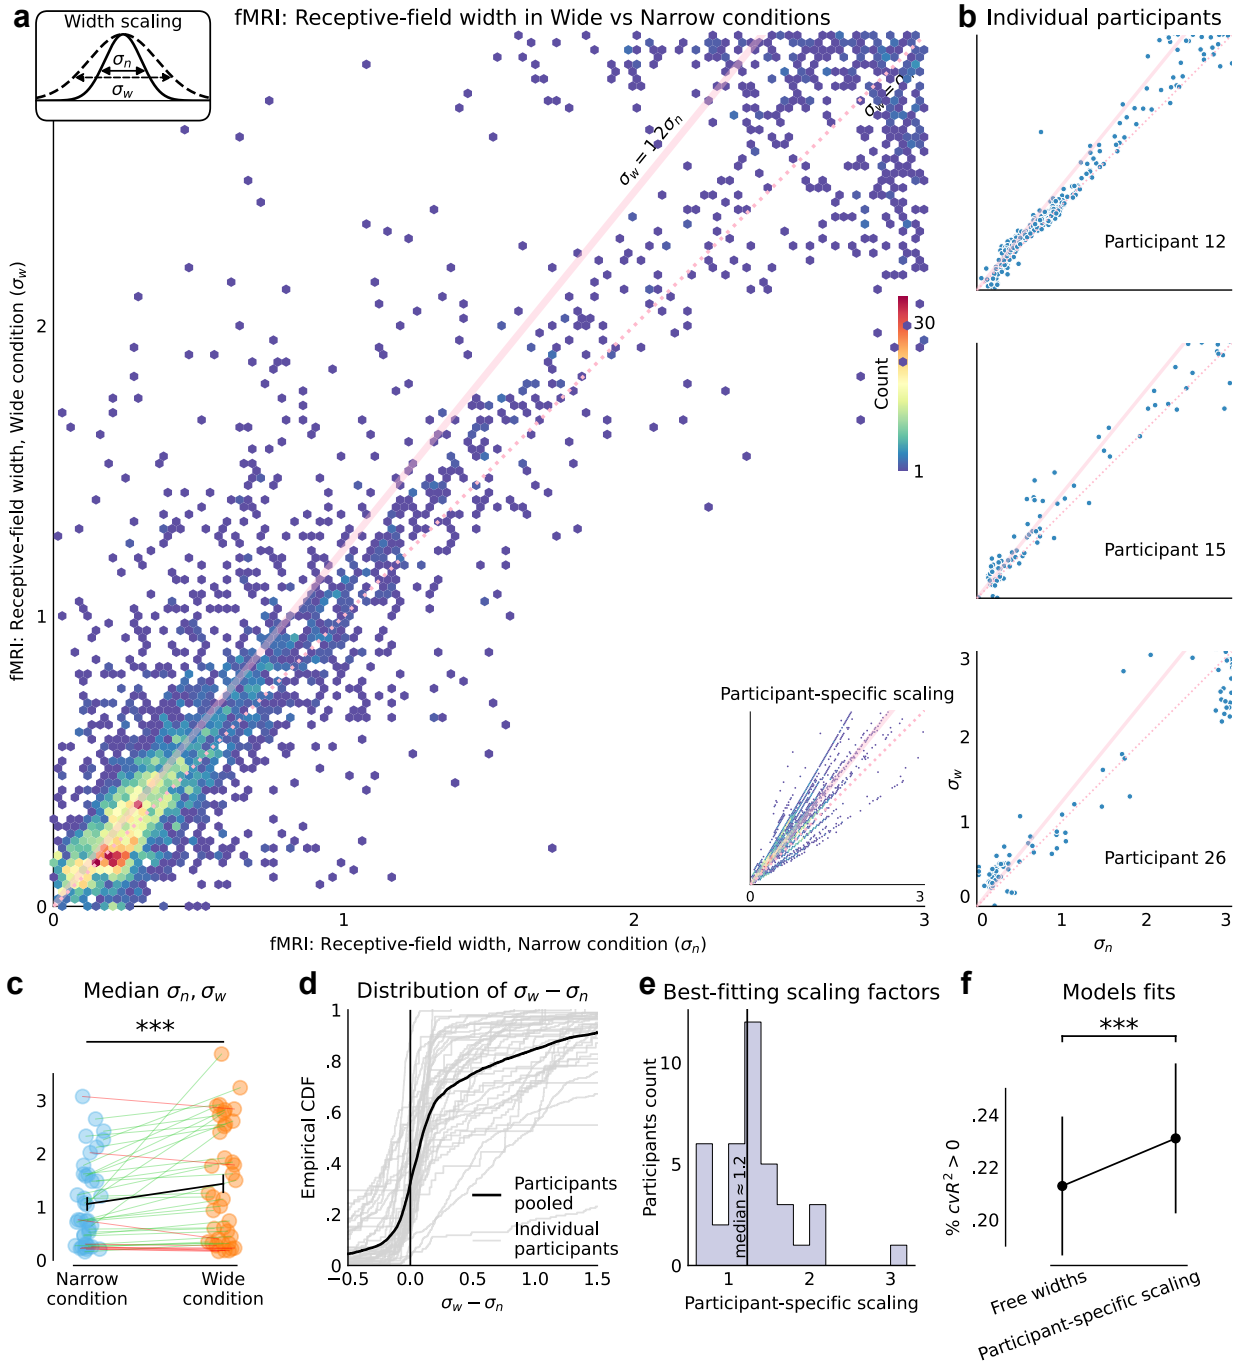

Fig. S10: Censored fits: widths. As Figure 4 in the main text, but with fits restricted to the range shared between the Narrow and Wide conditions.

5 Participants' responses and bias

Figure S11 shows the average responses provided by the participant in response to each numerosity, as well as their bias. The behavior of participants exhibit all the typical patterns found in magnitude estimation tasks (including with numerosity)[1–4]: overestimation of small magnitudes, underestimation of large magnitudes, a 'central tendency of judgment' wherein the crossover point (where the bias vanishes) depends on the width of the prior, and a negative bias for large numerosities that is larger than the positive bias for small numerosities. We also show the biases and standard deviations of individual participants in Figures S12 and S13.

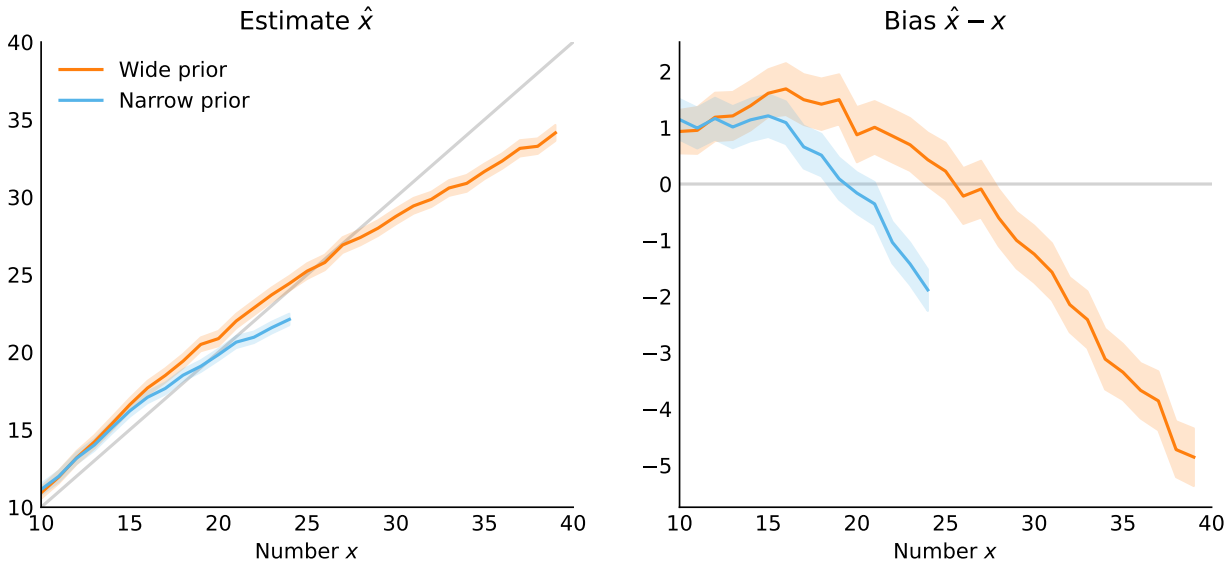

Fig. S11: Participants' responses and bias. Mean participants' response $\hat{x}$ (*left*), and bias $\hat{x} - x$ (*right*), as a function of the presented number x. Shaded areas show the 5%-95% credible intervals (see Methods for details on the statistical estimation procedure).

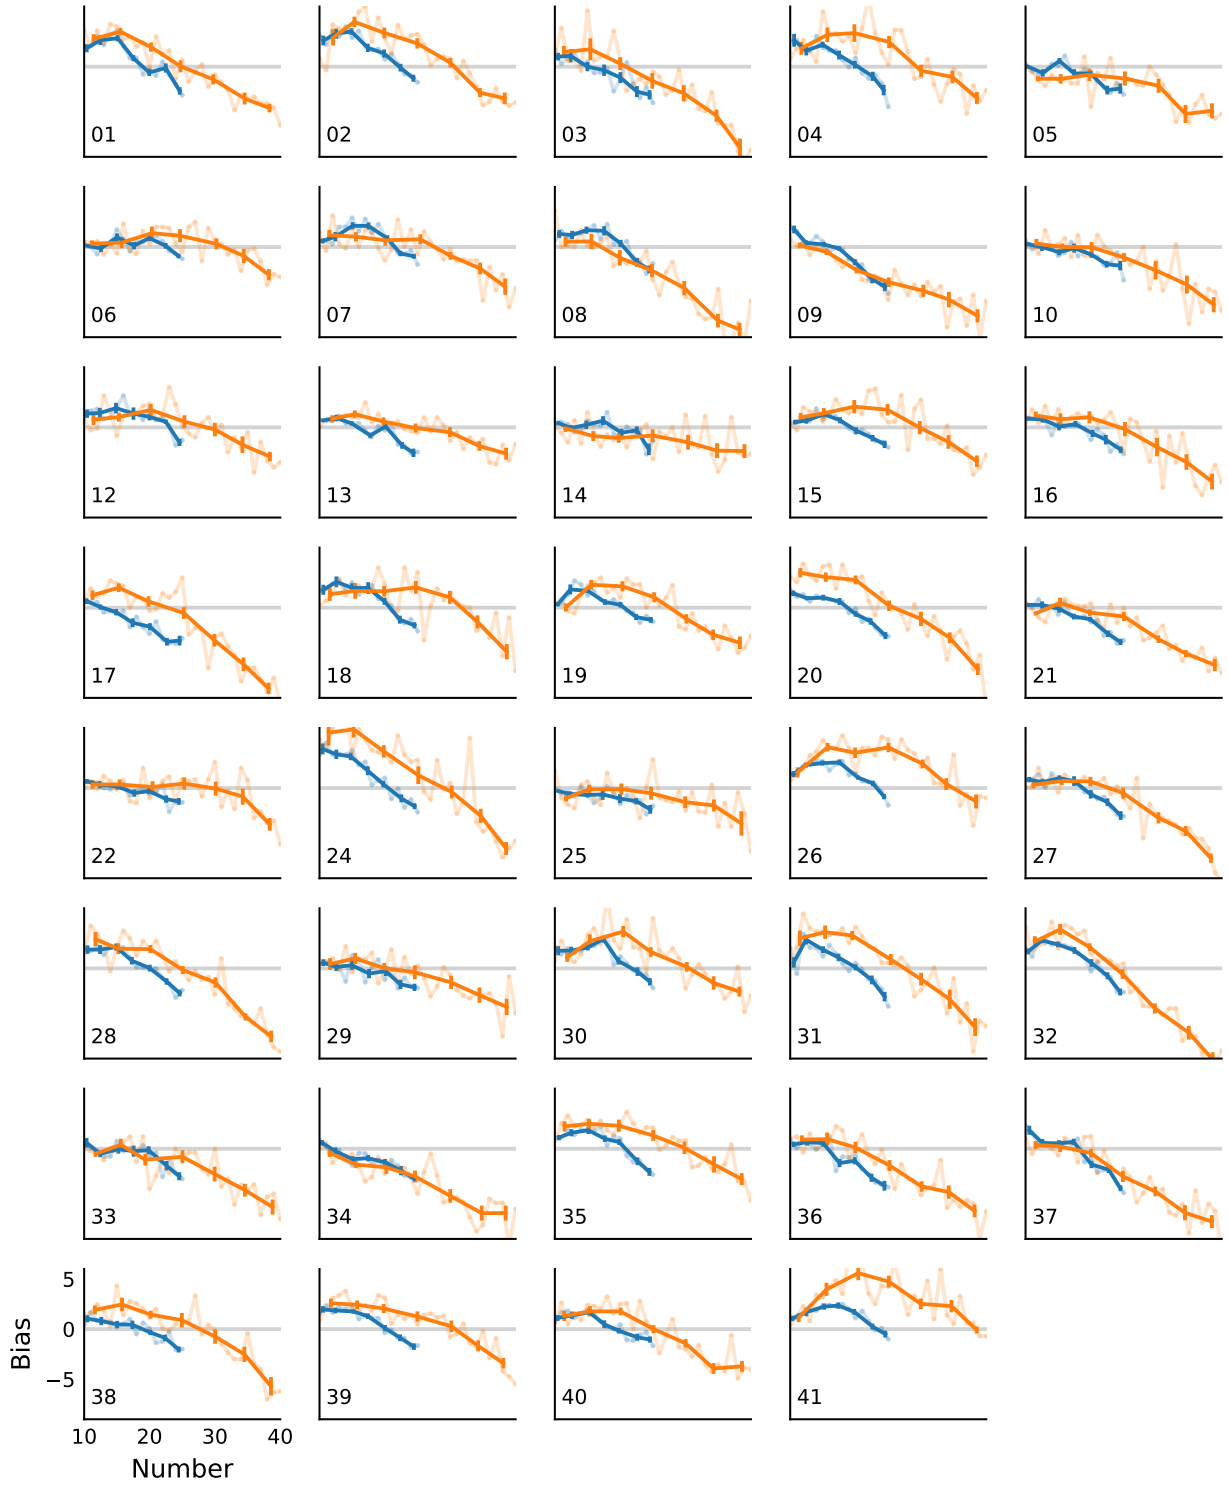

Fig. S12: Responses of individual participants. Empirical average of each participant's bias (difference between the response and the presented number) as a function of the presented number; with numbers pooled in bins (thick lines) or not (thin lines). Error bars: sem.

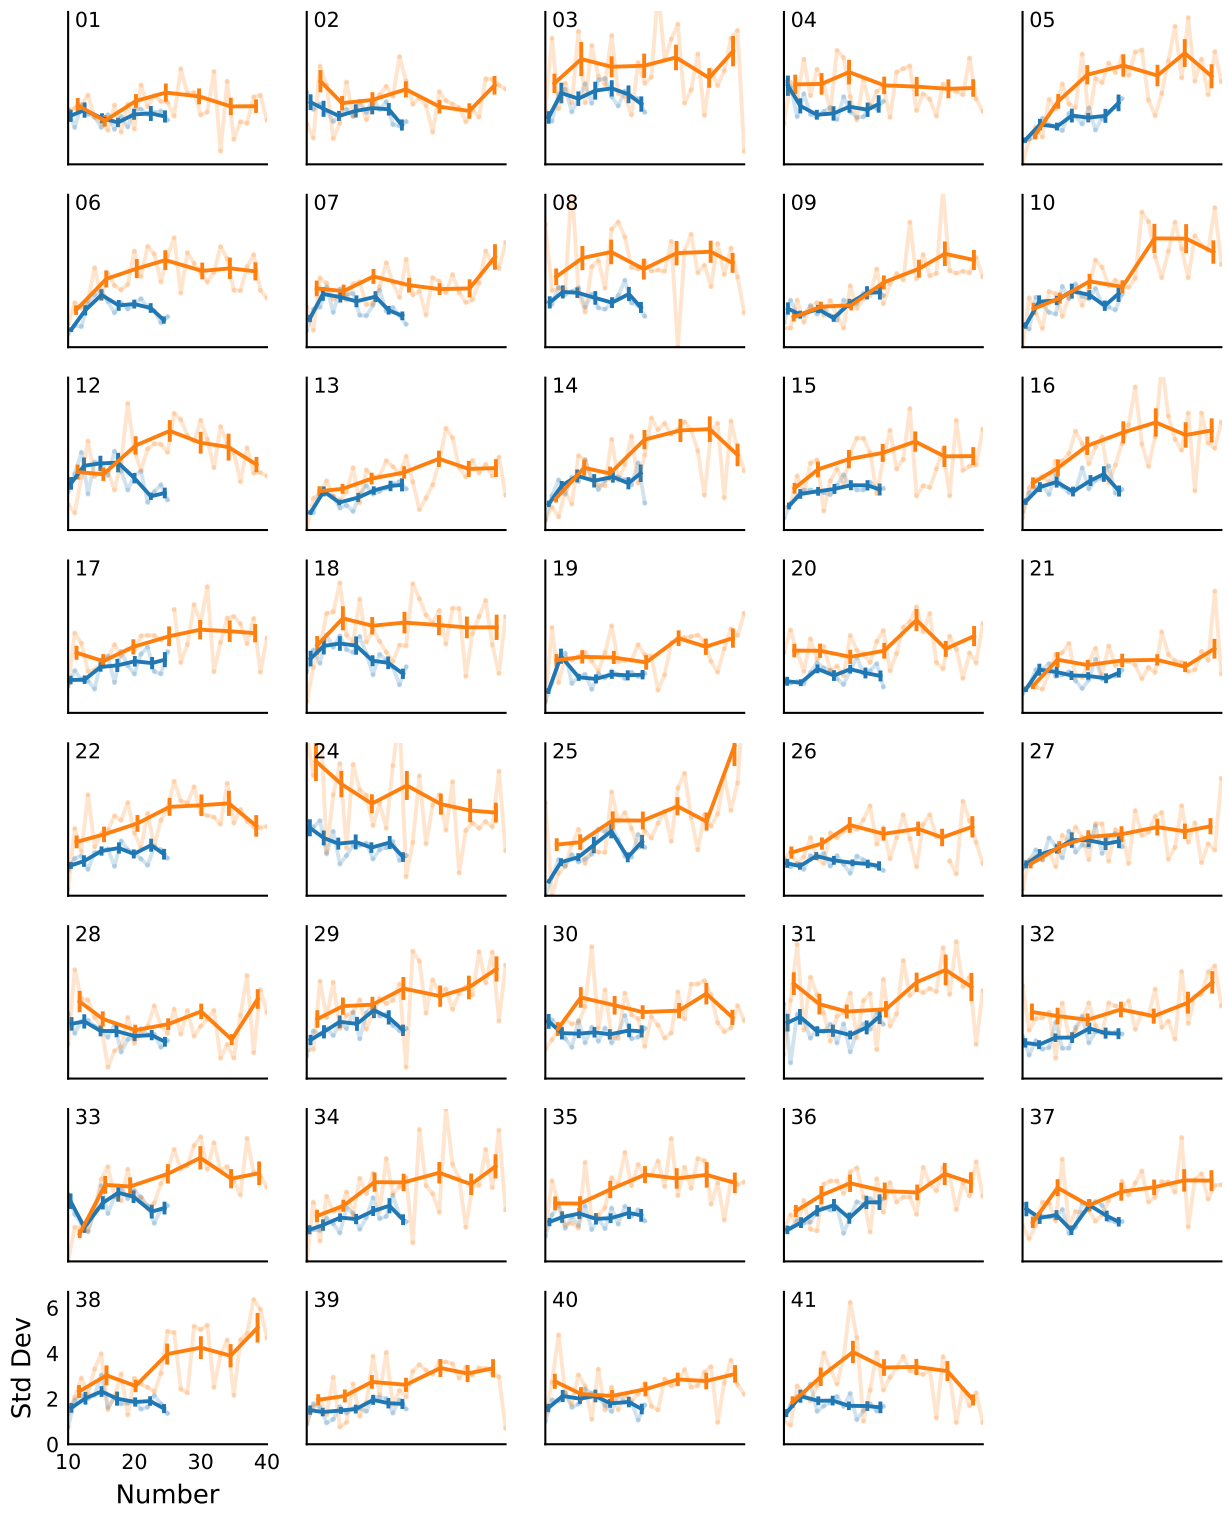

Fig. S13: Standard deviations of individual participants. Empirical standard deviation of each participant's responses as a function of the presented number; with numbers pooled in bins (thick lines) or not (thin lines). Error bars: sem.

References

1. Anobile, G., Cicchini, G. M. & Burr, D. C. Linear Mapping of Numbers onto Space Requires Attention. *Cognition* **122,** 454–459. ISSN: 0010-0277 (Mar. 2012).
2. Petzschner, F. H., Glasauer, S. & Stephan, K. E. A Bayesian Perspective on Magnitude Estimation. *Trends in Cognitive Sciences* **19,** 285–293. ISSN: 13646613 (May 2015).
3. Xiang, Y., Graeber, T., Enke, B. & Gershman, S. J. Confidence and Central Tendency in Perceptual Judgment. *Attention, Perception, & Psychophysics* **83,** 3024–3034. ISSN: 1943-393X (Oct. 2021).
4. Prat-Carrabin, A. & Gershman, S. J. Bayesian Estimation Yields Anti-Weber Variability. *PNAS Nexus* (Aug. 2025).