SUPPLEMENTAL FIGURES

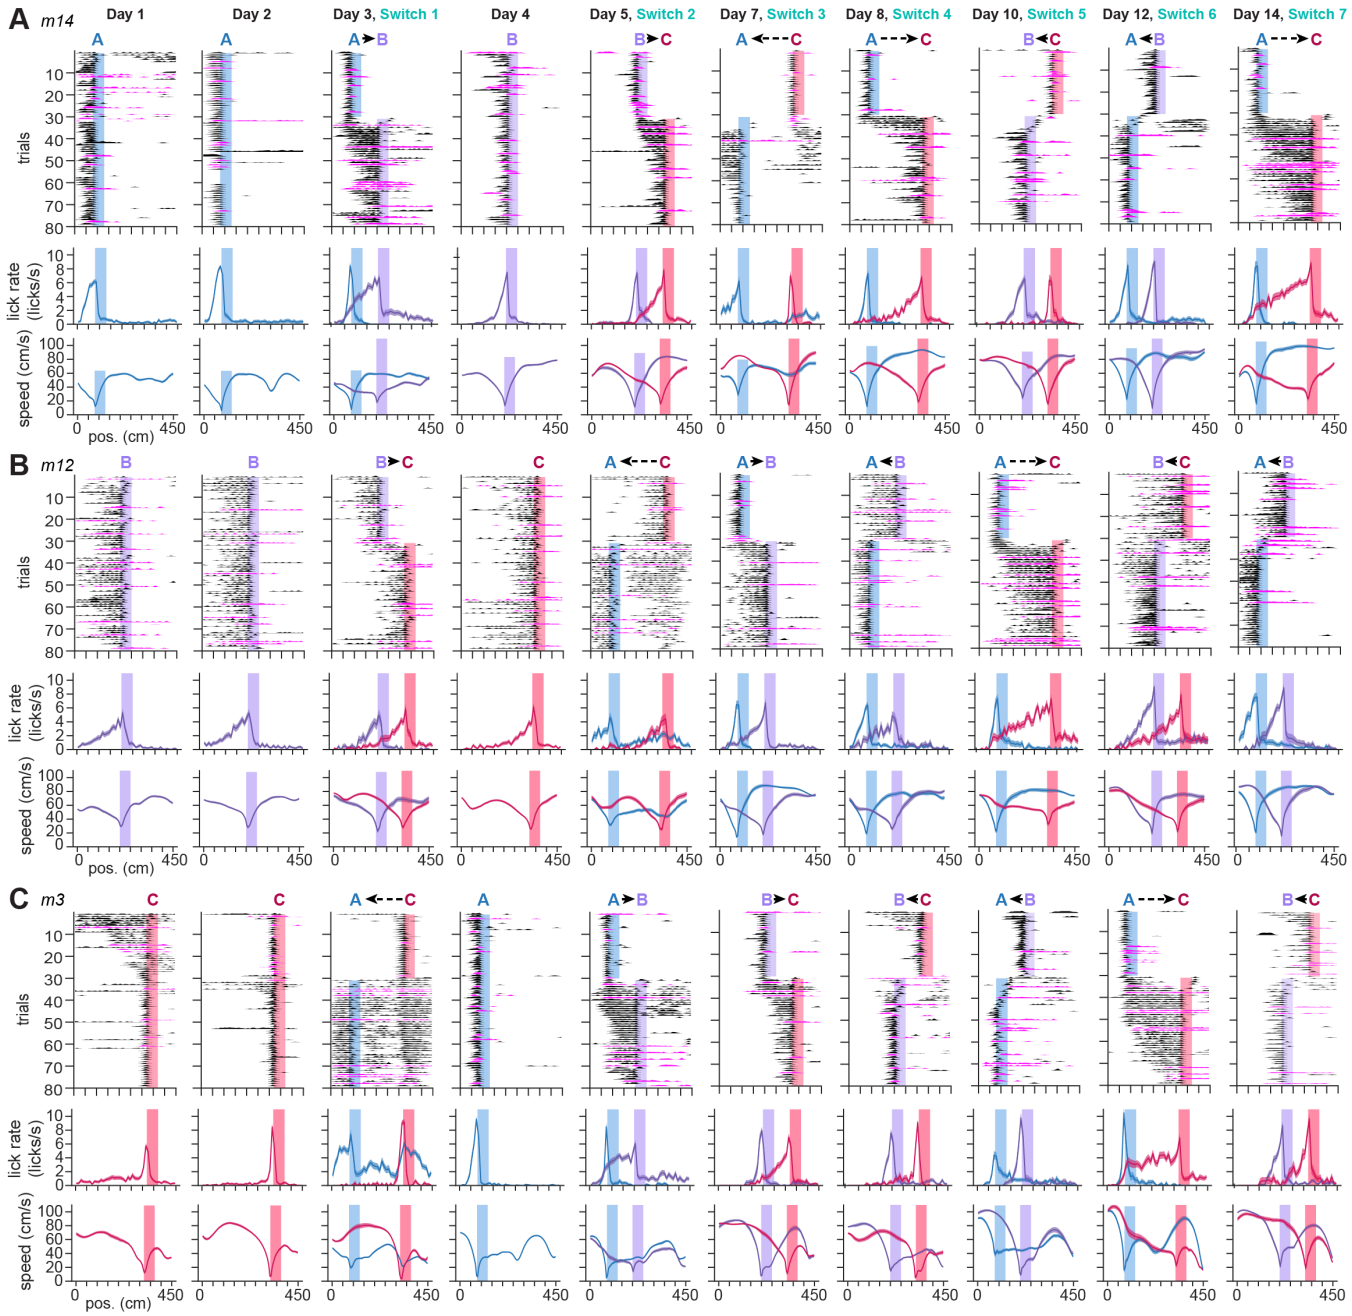

D

mouse	sex	start zone	Env 1 seq.	Env 2 seq.
m3	M	C	(C, A, B), C	(B, A, C), B
m4	F	B	(B, A, C), B	(C, A, B), C
m7	F	A	(A, C, B), A	(B, C, A), B
m11	M	B	(B, A, C), B	(C, A, B), C
m12	M	B	(B, C, A), B	(A, C, B), A
m13	F	C	(C, B, A), C	(A, B, C), A
m14	F	A	(A, B, C), A	(C, B, A), C
m2	M	A	(A, A, A ...)	n/a
m6	F	B	(B, B, B ...)	n/a
m10	M	C	(C, C, C ...)	n/a

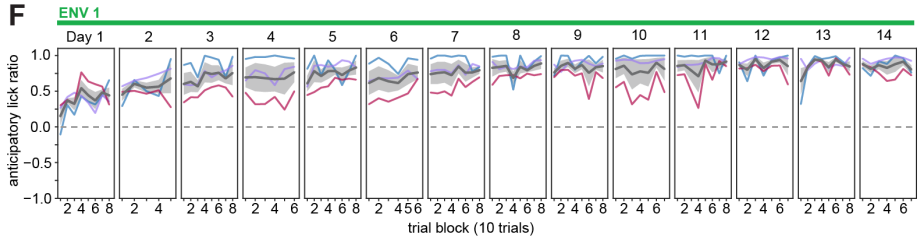

E

Figure S1. Behavior across multiple reward switch sequences, related to Figure 1.

A) Expanded behavior over the 14-day task for the example animal in Fig. 1F which started the task with reward zone "A" (m14), showing days 1-5 and then each subsequent switch (day 7, 8, 10, 12, 14). Top: smoothed lick raster over 80 trials per session; black indicates rewarded trials, and magenta indicates omission trials. Shaded regions indicate the reward zone active on each set of trials. Middle: The mean ± s.e.m. lick rate for trials at each reward location (blue=A, purple=B, red=C). Bottom: Mean ± s.e.m. speed for trials at each reward location.

B) Behavior for an example animal that started the task with reward zone "B" (m12), organized as in (A).

C) Expanded behavior for the example animal in Fig. 1G which started the task with reward zone "C" (m3), organized as in (A).

D) Mouse identities, sexes, and reward switch sequence allocation. Top 7 mice are "switch" mice, bottom 3 mice are "fixed-condition" mice that maintained the same reward zone and environment throughout the experiment. Sequences are shown as "(0, 1, 2), 0" where 0 is the first reward zone experienced in a given environment, 1 is the first switch (2nd zone), 2 is the second switch (3rd zone), and the third switch which is a return to zone 0.

E) Schematic depicting anticipatory lick ratio calculation around reward zone "B" as an example. The anticipatory lick ratio compares lick rate in an anticipatory zone 50 cm before the start of the reward zone vs. everywhere outside the reward zone (the reward zone is excluded to eliminate consummatory licks). A value of 1 indicates licking exclusively in the anticipatory zone, 0 is chance everywhere outside the anticipatory zone and the reward zone.

F) Anticipatory lick ratio across blocks of 10 trials for fixed-condition animals (n=3 mice) which maintained the same reward zone and environment throughout the entire experiment. Colored lines represent individual mice and the reward zone active for each mouse (blue=A, purple=B, red=C); gray lines and shading show mean ± s.e.m across mice.

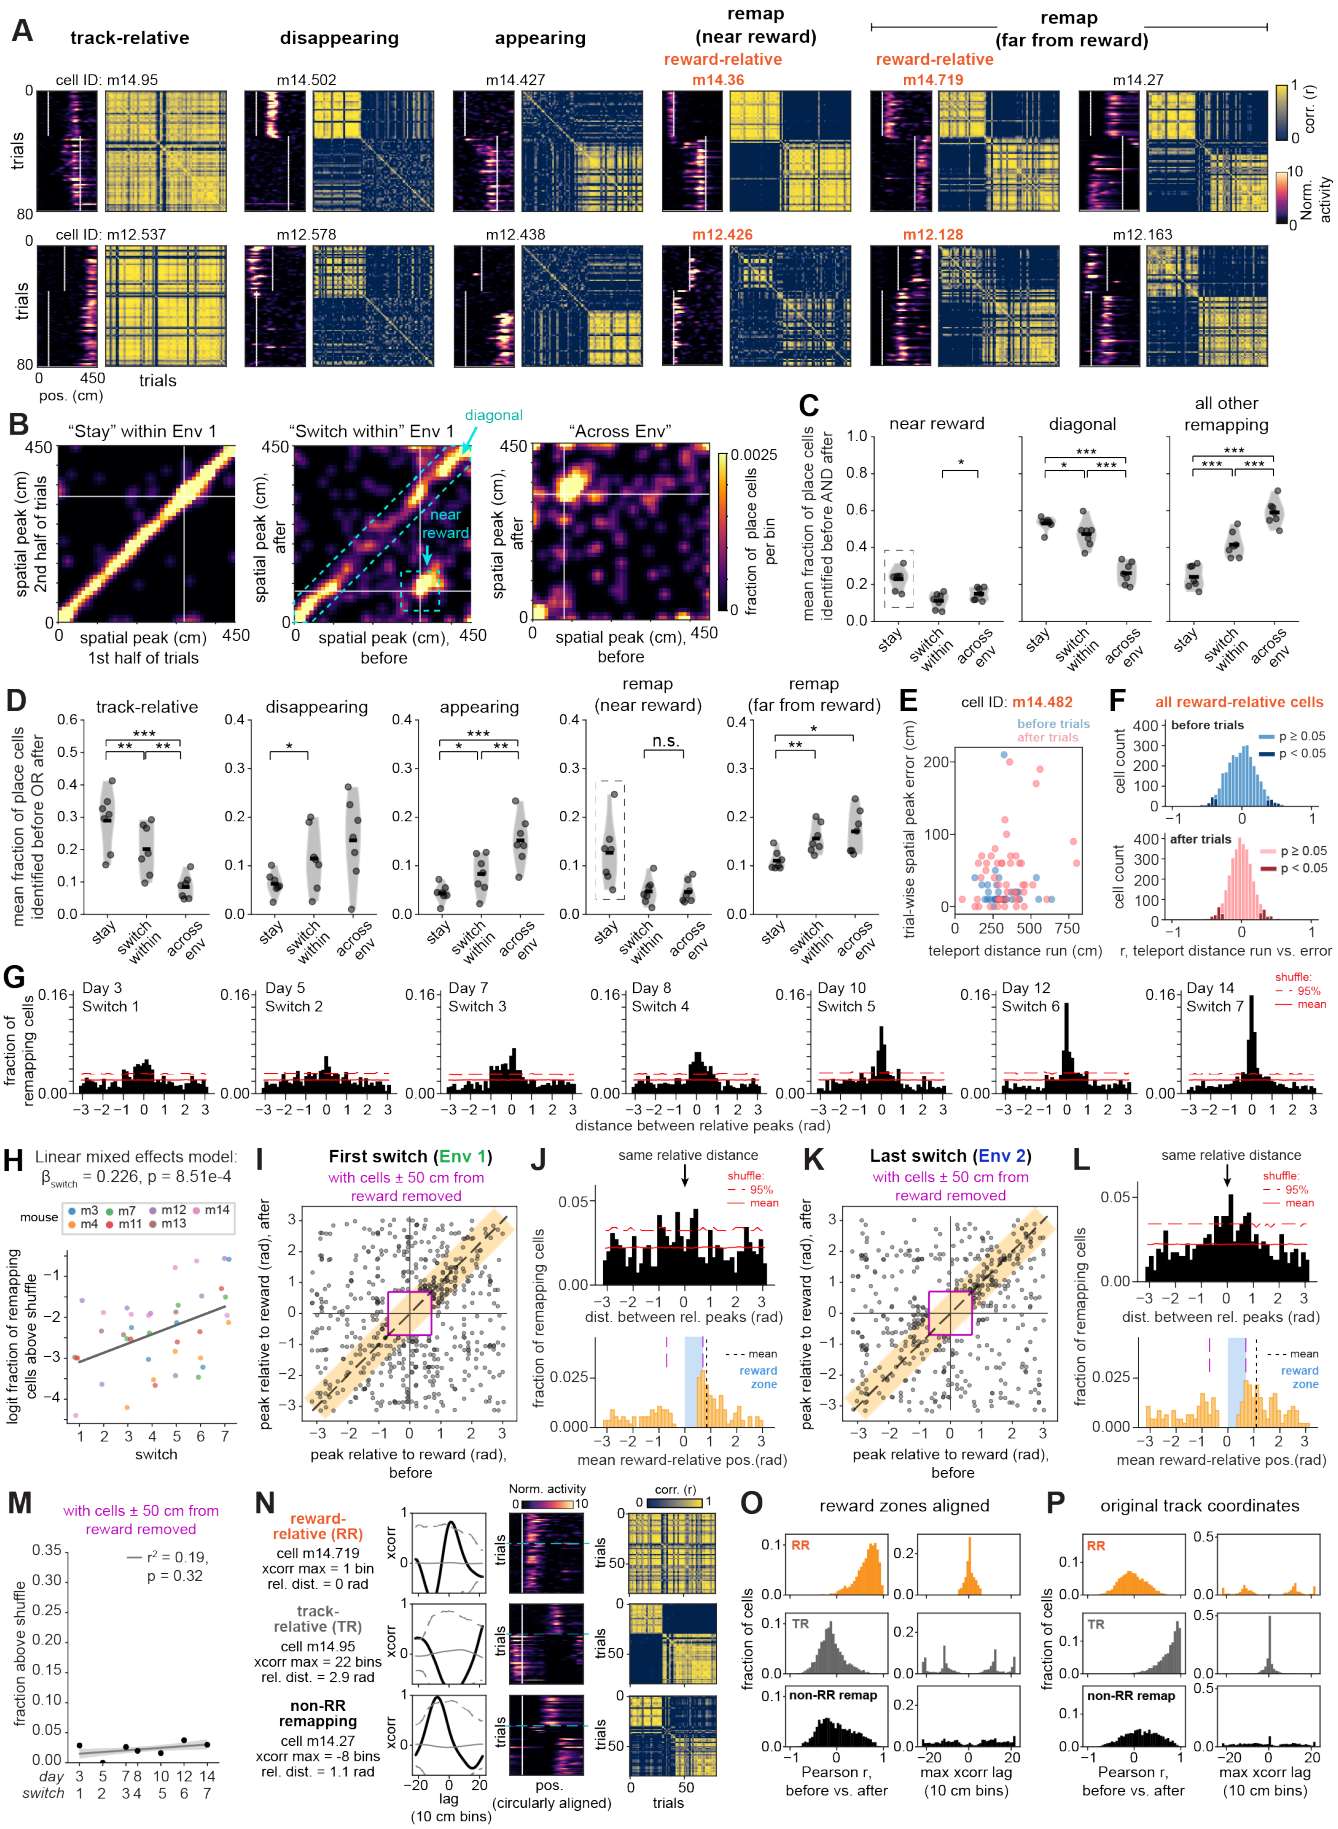

Figure S2. Quantification of remapping types and controls for reward-relative remapping, related to Figure 2.

A) Example cells from 2 mice not shown in Fig. 2A on the last switch (day 14), replicating the co-existence of remapping types within animal shown in Fig. 2A. Top row: Data from mouse m14. Bottom row: Data from mouse m12. The "reward-relative" designation (orange text) encompasses cells both near and far from reward that maintain their firing distance with respect to the reward zone (see Fig. 2B and Methods).

B) Examples of population maps from mouse m14 illustrating a day that is "stay" (day 6, same reward location, same environment), "switch within" (day 7, change reward location, same environment), or "across environment" (day 8, change reward location, change environment). Each plot shows the smoothed 2D histogram of the fraction of place cells with peak firing in each 10 cm bin before vs. after the reward zone switch (or in the 1st vs. 2nd half of trials on "stay" days). White lines indicate the start of each reward zone. For this analysis, only cells with significant spatial information both before and after are included. On stay days, this produces a strong density along the diagonal, corresponding to place cells that remain track-relative throughout the session. On switch days, this diagonal degrades as more cells remap, and the off-diagonal band around the reward intersection appears (see also Fig. 2C). Upon the across-environment switch, the diagonal degrades further indicating more global remapping, but the reward-relative band is maintained. Cyan boxes mark bins used for the "diagonal" and "near reward" categories shown in (C).

C) Quantification of population-level remapping using bins of the histograms illustrated in (B) (unsmoothed for quantification). Each dot is the mean fraction of cells across days of each category and environment per animal (n=7 mice); violins show the distribution across mice, horizontal lines show the mean across mice. In (C) and (D), *p<0.05, **p<0.01, ***p<0.001 from pairwise paired t-tests between categories ("stay", "switch within" environment, or "across env"), performed on the logit-transformed fractions. P-values are Holm corrected for multiple comparisons within each plot. Left: Near reward; cells with peaks ≤50 cm from both reward zone starts; gray dashed box signifies that on stay days, this value reflects the fraction of total place cells near the reward across the 1st to 2nd half of the session, thus no t-tests are shown for this category as the reward is not moved. Switch within vs. across env: t = -3.0, p = 0.024. Middle: Diagonal; peaks within ≤50 cm of the same linear position before vs. after. Stay vs. switch within: t = 2.6, p = 0.042; stay vs. across env: t = 11.82, p = 6.6e-5; switch within vs. across env: t = 8.2, p = 3.6e-4. Right: All other remapping; all non-near-reward and non-diagonal bins. Stay vs. switch within: t = -7.7, p = 2.5e-4; stay vs. across env: t = -29.4, p = 3.1e-7; switch within vs. across env: t = -9.1, p = 2.0e-4.

D) Mean fractions of total place cells defined as track-relative (peaks ≤50 cm apart), disappearing, appearing, remap near reward (≤50 cm from both reward zone starts), or remap far from reward (>50 cm from at least one reward zone start), averaged across days of each type per mouse as in (C). See (C) for description of statistical tests. Track-relative: Stay vs. switch within: t = 5.7, p = 0.0012; stay vs. across env: t = 18.4, p = 5.0e-6; switch within vs. across env: t=6.6, p = 0.0011. Disappearing: Stay vs. switch within: t = -3.6, p = 0.033; stay vs. across env: t = -2.5, p = 0.093; switch within vs. across env: t = -0.96, p = 0.37. Appearing: Stay vs. switch within: t = -3.2, p = 0.018; stay vs. across env: t = -8.6, p = 4.1e-4; switch within vs. across env: t = -5.9, p = 0.002. Remap near reward: switch within vs. across env: t = -0.23, p = 0.83. As in (C), the gray dashed box indicates a fraction near reward on stay days that cannot be identified as "reward-specific", as there is no reward switch; thus no statistical comparisons are shown for this category. Remap far from reward: Stay vs. switch within: t = -4.8, p = 0.0093; stay vs. across env: t = -3.8, p= 0.018; switch within vs. across env: t = -1.6, p = 0.16.

E) To test the possibility that the putative reward-relative neurons encode distance run since the last reward, we leveraged the variable length teleport zone to ask whether distance run in the teleport zone predicted an offset (error) in the cell's spatial firing on the subsequent trial (see Methods). Here, an example reward-relative cell is shown (cell m14.482 on day 8, switch 4 across environments; also shown in Fig. 2B). Each dot is a trial (blue = trials before the reward switch, pink = trials after the switch). This cell shows no significant Spearman correlation between the teleport distance run and spatial peak error for either set of trials (before: r = -0.10, p = 0.62; after: r = 0.19, p = 0.18).

F) Quantification of Spearman correlation coefficients between teleport distance run and spatial peak error on the subsequent trial, for all reward-relative cells across all animals and switch days (n = 3104 cells, 7 mice, 7 switch

days). Histograms are stacked; dark shades highlight cells with significant correlation coefficients ($p < 0.05$): 7.6% of reward-relative cells on before trials, 8.9% on after trials. Light shades highlight non-significant cells: 92.4% of reward-relative cells on before trials, 91.1% on after trials.

G) Histograms of the circular distance between relative peaks after minus before for remapping place cells as in Fig. 2E, G, shown for all switch days (days 3 and 14 are repeated from Fig. 2). n=7 mice; day 3 n=677 cells, day 5 n=583 cells, day 7 n=726 cells, day 8 n=706 cells, day 10 n=508 cells, day 12 n=506 cells, day 14 n=707 cells.

H) Increase in the fraction of cells remapping relative to reward above the shuffle over experience, at the level of individual animals. Each point is color-coded by mouse (n=7 mice) and represents the fraction of cells (logit-transformed) exceeding the "random-remapping" shuffle calculated within each individual mouse and switch day. Regression coefficient β and p-value are from a linear mixed effects model with switch index as the fixed effect and mice as random effects (random intercepts were allowed, and using random slopes did not change the outcome of the model; data not shown). Gray line shows the model prediction.

I-M) Control for Fig. 2D-H, excluding cells with peaks within ± 50 cm (~0.698 radians) from both reward zone starts (indicated by the magenta lines), intended to exclude putative reward-dedicated cells reported previously[54]. (I) n=557 cells, 7 mice. (K) n=466 cells, 7 mice. All points are jittered slightly for visualization (see Methods). The fraction of cells remapping relative to reward at distances greater than 50 cm from reward still exceeds the shuffle (J, top; L, top). (J, bottom) and (L, bottom): the distribution of mean reward-relative positions for the remaining cells shown in the orange shaded region around the unity line in (I) and (K), respectively. The means of these distributions are significantly non-zero: (J) non-zero mean = 0.848, 95% confidence interval [lower, upper] = [0.443, 1.253], circular mean test, n = 163 cells. (L) non-zero mean = 1.092, 95% confidence interval [lower, upper] = [0.573, 1.612], circular mean test, n = 167 cells. Black dotted lines mark the means, blue shaded region marks the extent of the reward zone. Note that cells that fire within 50 cm of one reward but not the other may have a mean relative distance of <0.698 radians (<50 cm); this is visible as the small fractions between the magenta lines marking ±50 cm from the start of the reward zone in (J, bottom) and (L, bottom). In (M), note that the fraction of cells above the shuffle shows a non-significant increase across task experience.

N) Illustration of cross-correlation criterion to identify reward-relative remapping cells (top row) vs. track-relative cells (middle row) and non-reward-relative (non-RR) remapping cells (bottom row). Each example cell is also shown in (A), top row. Left of each cell: cross-correlation (xcorr) between trial-averaged spatial firing with reward zones circularly aligned, before vs. after the switch. Gray lines mark the mean (solid) and 95% confidence interval (dashed) of the shuffle for that cell. The offset of the xcorr maximum above the shuffle and the distance between relative spatial peaks in radians is listed next to each cell. Note that the reward-relative cell is maximally cross-correlated very close to zero lag, while the track-relative cell is anticorrelated at 0 lag in reward-aligned coordinates. Middle of each cell: spatial activity with reward zones aligned as in Fig. 2B; blue line marks the switch between reward zones. Right of each cell: trial-by-trial correlation matrix computed on the reward-aligned spatial activity. Note the more uniform structure of the correlation matrix for the reward-relative cell vs. the block-like matrices for the track-relative and non-RR remapping cell, which is the opposite of how these correlation matrices appear in original linear track coordinates (A, top row).

O) Left: Distributions of Pearson correlation coefficients for each subpopulation between the cell's trial-averaged activity maps pre- vs. post-switch in reward-aligned coordinates. The reward-relative (RR) population exhibits significantly higher correlations in these coordinates than the track-relative (TR) and non-RR remapping populations: RR vs. TR U = 67.7, $p << 1e\text{-}12$, Wilcoxon rank-sum test (n = 3104 RR cells, 3901 TR cells across 7 mice, 7 switch days); RR vs. non-RR remapping U = 53.8, $p << 1e\text{-}12$, Wilcoxon rank-sum test (n = 3104 RR cells, 2518 TR cells across 7 mice, 7 switch days). Right: Distributions of maximum xcorr offsets above the shuffle for each subpopulation, with the RR distribution thresholded at ±5 bins.

P) Same as (O), but computed between trial-averaged activity of each subpopulation in original linear track coordinates. The TR population exhibits significantly higher correlations in these coordinates than the RR and non-RR remapping populations: TR vs. RR U = 69.1, $p << 1e\text{-}12$, Wilcoxon rank-sum test; TR vs. non-RR remapping U = 21.0, $p << 1e\text{-}12$, Wilcoxon rank-sum test.

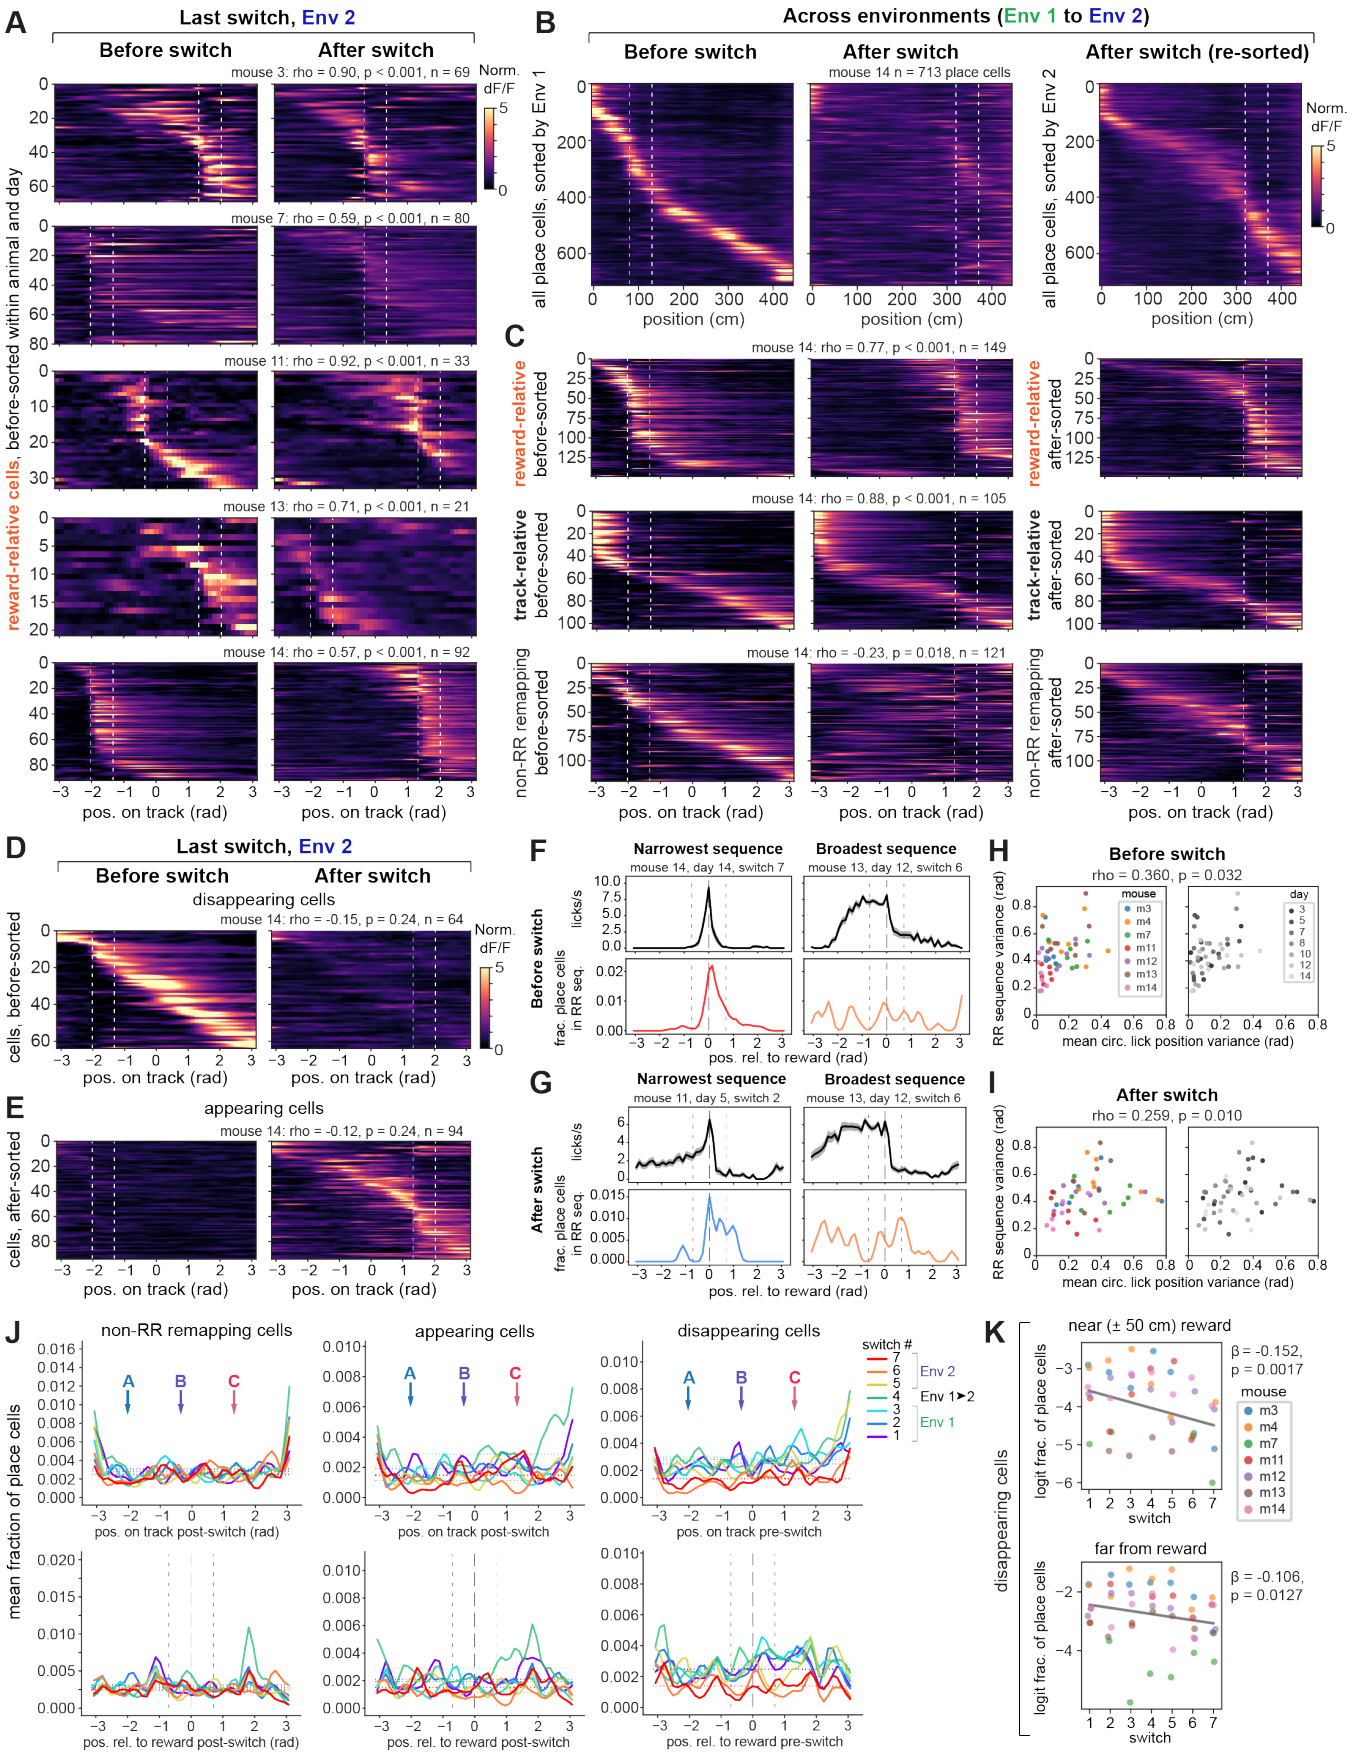

Figure S3. Sequences across animals, remapping types, and environments, related to Figure 3.

A) Reward-relative sequences on the last switch (day 14) for the 5 mice not shown in Fig. 3A (organized as in Fig. 3A).

B) Population level remapping upon the switch from Env 1 to novel Env 2 (day 8). All place cells identified in mouse m14 are shown, sorted in the left two panels by their order in Env 1 before the switch, and by their order in Env 2 in the right-most panel. Note population-level overrepresentation in both sorts as the bend away from the diagonal around the start of the reward zone.

C) Remapping (or lack thereof) of specific subpopulations on the switch from Env 1 to Env 2 (day 8), shown for mouse m14. Top: reward-relative (RR) sequences are preserved across environments; Middle: a subset of track-relative place cells also retain their firing positions across environments, especially at the beginning of the track; Bottom: non-reward-relative, non-track-relative cells ("non-RR remapping") with significant spatial information in both Env 1 and Env 2 instead remap globally, mirroring the overall population. The cells in the left two columns in (C) are sorted by their position in Env 1; the right-most column is sorted (and cross-validated) by position in Env 2, which improves the sort for the non-RR remapping cells but not the reward-relative or track-relative cells.

D) Example sequences of disappearing cells from mouse m14. In (D) and (E), cross-validated sort was performed on the trials in which cells had significant spatial information (before for disappearing, after for appearing).

E) Example sequences of appearing cells from mouse m14.

F) Example licking (top) and distribution of reward-relative (RR) sequence peak firing positions during trials before the switch, aligned to the start of the reward zone. Left column: animal and day with the narrowest RR sequence (lowest circular variance); Right column: animal and day with the broadest sequence (highest circular variance). Sequence colors correspond to switch index as shown in (J). Note more distributed, less precise licking coinciding with a more distributed sequence that is less dense around the reward zone (right). Vertical gray dashed lines indicate the reward zone start and surrounding ±50 cm.

G) Same as (F), but showing the narrowest and broadest sequences and corresponding licking for trials after the switch.

H-I) Circular variance in licking positions is significantly correlated with RR sequence variance both before (H) and after the reward switch (I). Each dot is a session, colored by mouse (left; n=7 mice) or by switch day (right; n=7 days). Note no apparent trend over days.

J) Mean distributions of peak firing positions across days (n=7 mice, s.e.m. omitted for clarity) for non-RR remapping cells (left), appearing cells (middle), and disappearing cells (right). Top row: sequences according to original linear track position, converted to radians. Bottom row: sequences according to position relative to the starts of reward zones. Horizontal dotted lines indicate the expected uniform distribution for each day. Vertical gray dashed lines indicate the reward zone start and surrounding ±50 cm.

K) Both near (top) and far from (bottom) the start of the reward zone, the fraction of disappearing place cells out of total place cells decreases across experience. Fractions are shown as a logit transform here to account for values close to zero, and sessions with zero disappearing cells on any given day have been removed. Regression coefficients β and p-values are from linear mixed effects models with switch index as the fixed effect and mice as random effects. Non-RR remapping and appearing fractions did not show significant changes over experience (data not shown).

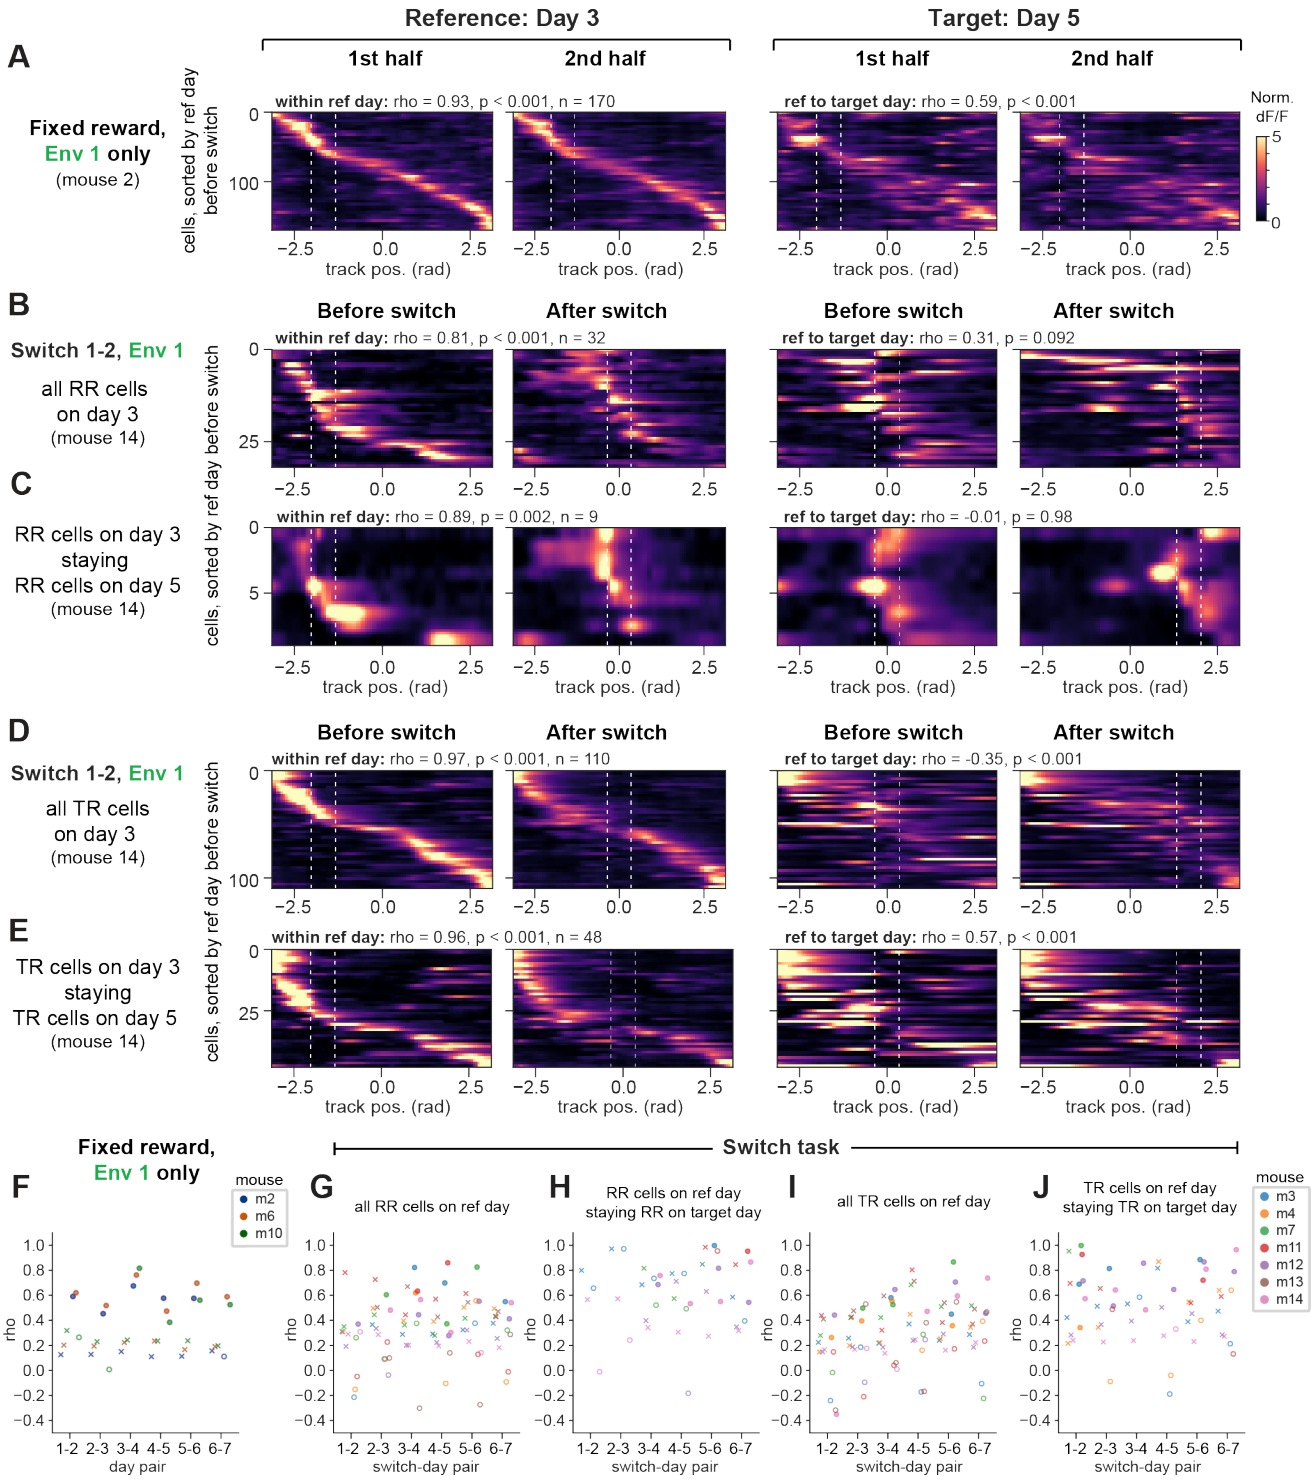

Figure S4. Drift in sequences tracked across days, related to Figure 4.

A) An example place cell sequence tracked from days 3 to 5 in an animal (m2) from the fixed-condition cohort, which experienced the same reward zone and environment for all 14 days. Neurons in this animal were identified using the same criteria as reward-relative cells (i.e. their field position relative to reward had to be stable in terms of its peak position and spatial cross-correlation, from the 1st to the 2nd half of trials), but without a reward switch to identify reward-relative cells, this likely includes other cell categories. In (A-E), cells within an animal are sorted by their cross-validated peak activity in the 1st trial set (1st half, or before the switch) on the reference (ref) day, which is defined as the first day in each day pair. This sort then is applied to activity in the 2nd trial set on the reference day, as well as the 1st and 2nd trial sets on the target day (the second day in the pair). Also in (A-E), the

sequence correlation coefficient, p-value relative to shuffle (two-sided permutation test), and n of cells tracked across the day pair are shown above each plot (within ref day correlation on the left, across-day correlation on the right; the across-day correlation is computed between the 2nd trial set on the ref day and the 1st trial set on the target day). In (A), note that drift is apparent across days in this fixed-condition animal as denoted by the drop in correlation coefficient, despite the cell ordering remaining significantly correlated across days.

B) A reward-relative sequence identified in animal m14 on day 3 (switch 1) and tracked to day 5 (switch 2), agnostic of the cell categories for the cells on day 5.

C) Reward-relative cells from day 3, m14 that were tracked to day 5 and remained reward-relative on day 5.

D) A track-relative sequence identified in animal m14 on day 3 (switch 1) and tracked to day 5 (switch 2), agnostic of the cell categories for the cells on day 5.

E) Track-relative cells from day 3, m14 that were tracked to day 5 and remained track-relative on day 5.

F) Across-day circular-circular correlation coefficients for place cell sequences tracked across pairs of days in the fixed-condition cohort (n = 3 mice; 95 ± 54 tracked cells per day pair, mean ± s.d. across mice and day pairs). Place cells included for these sequences were identified as described in (A). In (F-J), coefficients are plotted with the upper 95th percentile of 1000 shuffles of cell IDs (marked as x, jittered to the left of each set of coefficients, and color-coded by animal). Closed circles indicate coefficients with $p < 0.05$ compared to shuffle using a two-sided permutation test; open circles indicate $p \geq 0.05$.

G) Across-day circular-circular correlation coefficients for all cells in reward-relative sequences on each reference day tracked to each target day, for the animals performing the hidden reward zone switch task (n = 7 mice; 26 ± 20 cells, mean ± s.d. across mice and day pairs).

H) Across-day circular-circular correlation coefficients for reward-relative cells on each reference day that were tracked and remained reward-relative on each target day (n = 7 mice; 8 ± 10 cells, mean ± s.d. across mice and day pairs). Note slightly higher correlation coefficients for this subset compared to (G), though the small numbers of cells that were both successfully tracked and remained reward-relative preclude a robust comparison.

I) Across-day circular-circular correlation coefficients for all cells in track-relative sequences on each reference day tracked to each target day (n = 7 mice; 45 ± 37 cells, mean ± s.d. across mice and day pairs).

J) Across-day circular-circular correlation coefficients for track-relative cells on each reference day that were tracked and remained track-relative on each target day (n = 7 mice; 15 ± 17 cells, mean ± s.d. across mice and day pairs).

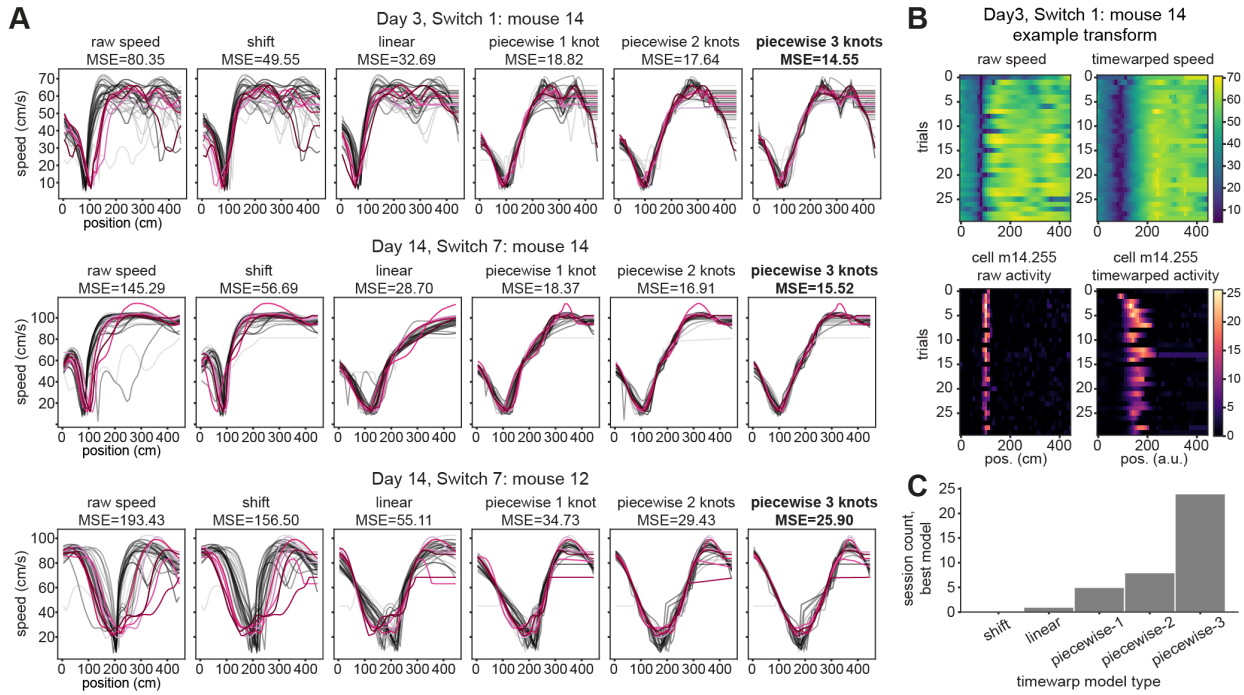

Figure S5. Implementation of time warp models and controls for analysis of rewarded vs. omission trials, related to Figure 5.

A) Illustration of time warp model fitting procedure to running speed profiles. Shades of black indicate speed on rewarded trials, shades of magenta indicate omission trials. Left to right: raw speed, followed by the transformed speed from fitting 5 model types to identify a model for best fit, as assessed by mean squared error (MSE) across the transformed trial-by-trial speed profiles. Bold text indicates the best-performing model. Top to bottom: the same example sessions shown in Fig. 5A-C: first reward switch on day 3, mouse 14; last reward switch on day 14, mouse 14; last reward switch on day 14, mouse 12.

B) Example transformation using the piecewise 3 knots model fit on the speed from day 3, mouse 14 (shown in A, top). Top left: raw trial-by-trial speed. Top right: time warp transformed speed. Bottom left: raw, binned deconvolved activity (normalized to the mean of the session) for example cell m14.255, which is also shown in Fig. 5A (top row of cells). Bottom right: time warp transformed neural activity of cell m14.255.

C) Frequency with which each model type was selected as maximally aligning speed according to MSE. For further analysis, we applied piecewise-3 to all data as this was the most common best model, to avoid variance in results due to model type for individual sessions.

D) Time-warp model fitting significantly reduces MSE across trial-by-trial speed profiles compared to the original speed data (model aligned speed vs. raw speed: W = 0, p = 7.74e-8, Wilcoxon signed-rank test, n = 38 sessions across 7 mice). Each dot is a session, colored by mouse. Black edges indicate sessions with the reward zone at location “A” or “B”, included for the analysis shown in Fig. 5. In (D-F), gray dashed line marks the unity line.

E) Reward vs. omission index (RO index) calculated from the model-transformed neural data, shown for the same sessions in Fig. 5D, is highly correlated with the RO index calculated from the original neural data (Pearson correlations). This indicates that the time warp model scales and aligns the firing rate curves between rewarded and omission trials but does not distort their relationship.

F) MSE between trial-by-trial raw speed profiles before vs. after the reward switch (within each session with at least 3 omission trials before the switch). Each dot is a session, colored by animal as in (D). Black edges indicate sessions included for analysis in Fig. 5. Speed MSE after the switch is significantly higher than before the switch (thus rewards and omissions are more difficult to compare on these trials): W = 130, p = 4.87e-4, Wilcoxon signed-rank test, n = 38 sessions across 7 mice.

G) Raw and model-aligned speed and neural activity of the same session and example cells shown in Fig. 5A, but for the trials after the reward switch (“after” trials). RO indices shown in bold text (top, far right of right column panel) are calculated from the “after” trials. Note preference of cells to fire more following rewards versus omissions, with the caveat that the speed profiles are more different than during the “before” trials.

H) Same session and example cells as Fig. 5B, but for the “after” trials.

I) Same session and example cells as Fig. 5C, but for the “after” trials.

J) A linear mixed effects model shows no significant effect of day (β= 0.01, p = 0.73), original MSE between mean speed on rewarded and omission trials (β= -0.001, p = 0.49), MSE of the time warp model fit (β = 0.002, p = 0.86), or the interactions of these terms (fixed effects) on the median RO index (random effects are mouse identity). Each dot shows the median RO index of the sessions included in Fig. 5 for each mouse, colored by mice as in (D). Gray line shows the model prediction ± 95% confidence interval.

K) Reward vs. omission index across all switch days (sessions) where the reward zone was at “A” or “B”, on the trials after the switch, for the population of reward-relative cells that have spatial firing peaks following the start of the reward zone within each mouse (n = 7 mice). In (K-L), significance level is set at p<0.007 with Bonferroni correction. **p<0.007, ***p<0.001, ****p<0.0001, one-sample Wilcoxon signed-rank test against a 0 median. m3: median = 0.10, W = 5.01e3, p = 3.31e-9, n = 5 days, 197 cells; m4: median = 0.13, W = 2.10e3, p = 3.70e-6, n = 4 days, 126 cells; m7: median = 0.14, W = 4.21e3, p = 7.58e-8, n = 5 days, 177 cells; m11: median = 0.22, W = 106, p = 2.61e-5, n = 3 days, 41 cells; m12: median = 0.13, W = 2.57e4, p = 1.55e-13, n = 5 days, 419 cells; m13: median = 0.21, W = 377, p = 2.12e-4, n = 5 days, 58 cells; m14: median = 0.04, W = 9.84e3, p = 0.085, n = 3 days, 213 cells.

L) Reward vs. omission index across all switch days (sessions) with at least 3 omission trials, at any reward zone before the switch, within each mouse (complementary to Fig. 5E, n = 7 mice). m3: median = 0.13, W = 4.99e3, p = 4.86e-8, n = 5 days, 191 cells; m4: median = 0.15, W = 2.63e3, p = 7.11e-12, n = 5 days, 165 cells; m7: median = 0.09, W = 4.44e3, p = 2.08e-3, n = 4 days, 157 cells; m11: median = 0.20, W = 418, p = 3.96e-6, n = 6 days, 68 cells; m12: median = 0.13, W = 4.66e4, p = 1.96e-13, n = 7 days, 541 cells; m13: median = 0.10, W = 902, p = 0.021, n = 5 days, 72 cells; m14: median = 0.09, W = 3.35e4, p = 9.13e-9, n = 6 days, 442 cells.

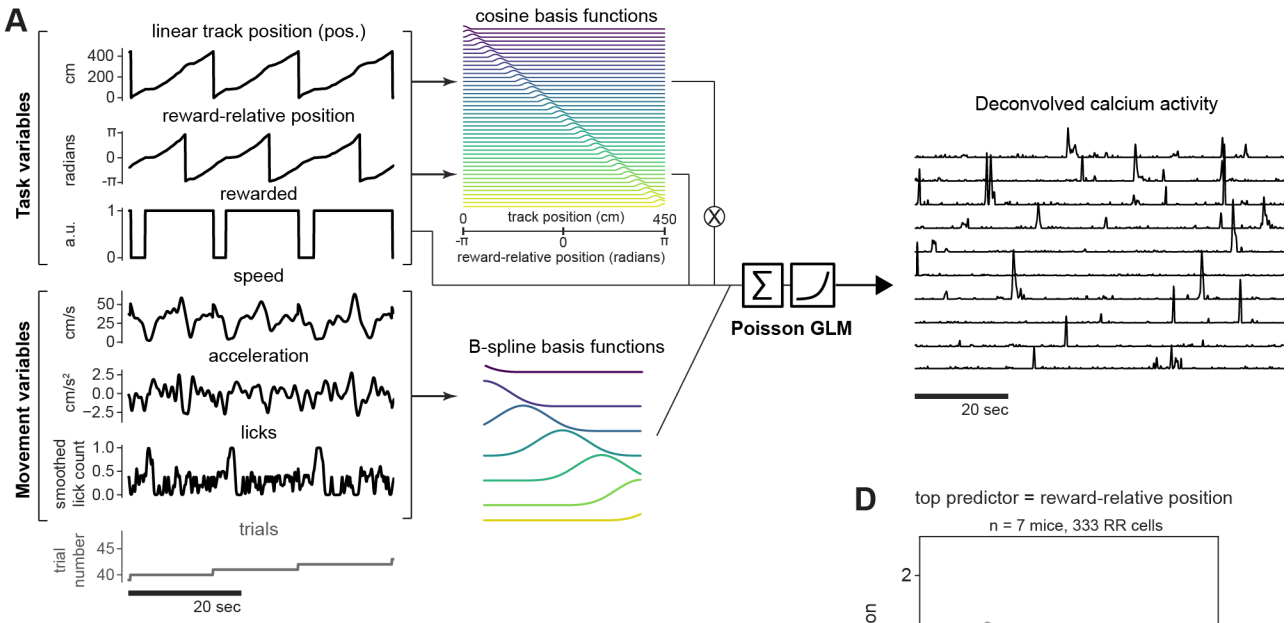

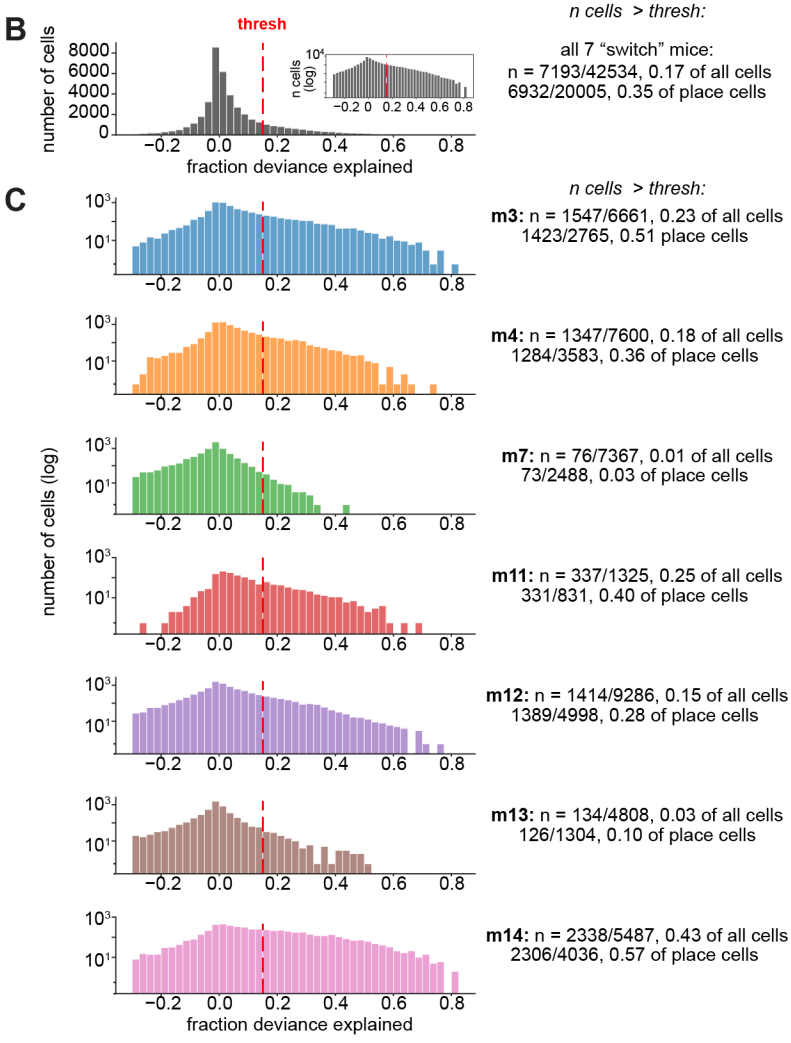

n cells > thresh:

all 7 "switch" mice:
n = 7193/42534, 0.17 of all cells
6932/20005, 0.35 of place cells

n cells > thresh:

m3: n = 1547/6661, 0.23 of all cells
1423/2765, 0.51 place cells

m4: n = 1347/7600, 0.18 of all cells
1284/3583, 0.36 of place cells

m7: n = 76/7367, 0.01 of all cells
73/2488, 0.03 of place cells

m11: n = 337/1325, 0.25 of all cells
331/831, 0.40 of place cells

m12: n = 1414/9286, 0.15 of all cells
1389/4998, 0.28 of place cells

m13: n = 134/4808, 0.03 of all cells
126/1304, 0.10 of place cells

m14: n = 2338/5487, 0.43 of all cells
2306/4036, 0.57 of place cells

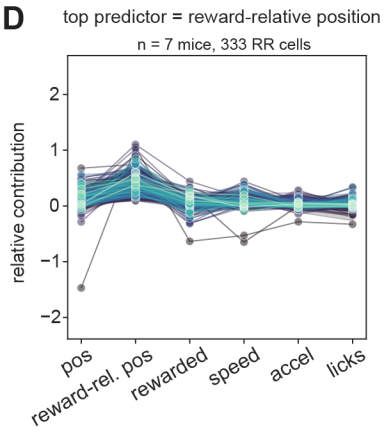

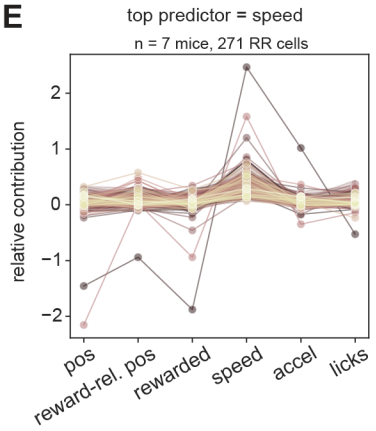

Figure S6. Generalized linear model (GLM) implementation and additional quantification, related to Figure 6.

A) Schematic of the Poisson GLM to predict deconvolved calcium activity from task and movement variables. Linear track position and reward-relative position are both transformed into cosine basis functions tiling the space in each set of coordinates (45 cosine bumps for the 45 10-cm spatial bins used). A binary representing whether reward was received on each trial ("rewarded") is multiplied with the linear track position basis functions to represent the interaction between reward received and position. Speed, acceleration, and smoothed lick count are quantile-transformed into B-splines with 5 knots and 3 polynomial degrees (providing 7 splines total) which smoothly tile the range of speeds, accelerations, and licking within each animal. Trial identities are not provided as an input but are used to group data for the training, cross-validation, and test sets (see Methods).

B) Quantification of model performance (fraction deviance explained on test data) across all animals and cells. Inset shows the histogram on a logarithmic scale. Red dashed line indicates the fit threshold of 0.15 fraction deviance explained used to select neurons for further analysis. Number and fraction of cells that exceeded the threshold is listed relative to all cells and relative to place cells, across all switch days.

C) Quantification of model performance within each animal (n = 7 mice), shown on a logarithmic scale.

D) Relative contributions of each GLM variable within reward-relative cells which were maximally predicted by reward-relative position. Each line and set of dots is an individual cell.

E) Relative contributions of each GLM variable within reward-relative cells which were maximally predicted by running speed, shown as in (D).