Supplemental Information: Figures S1-S5

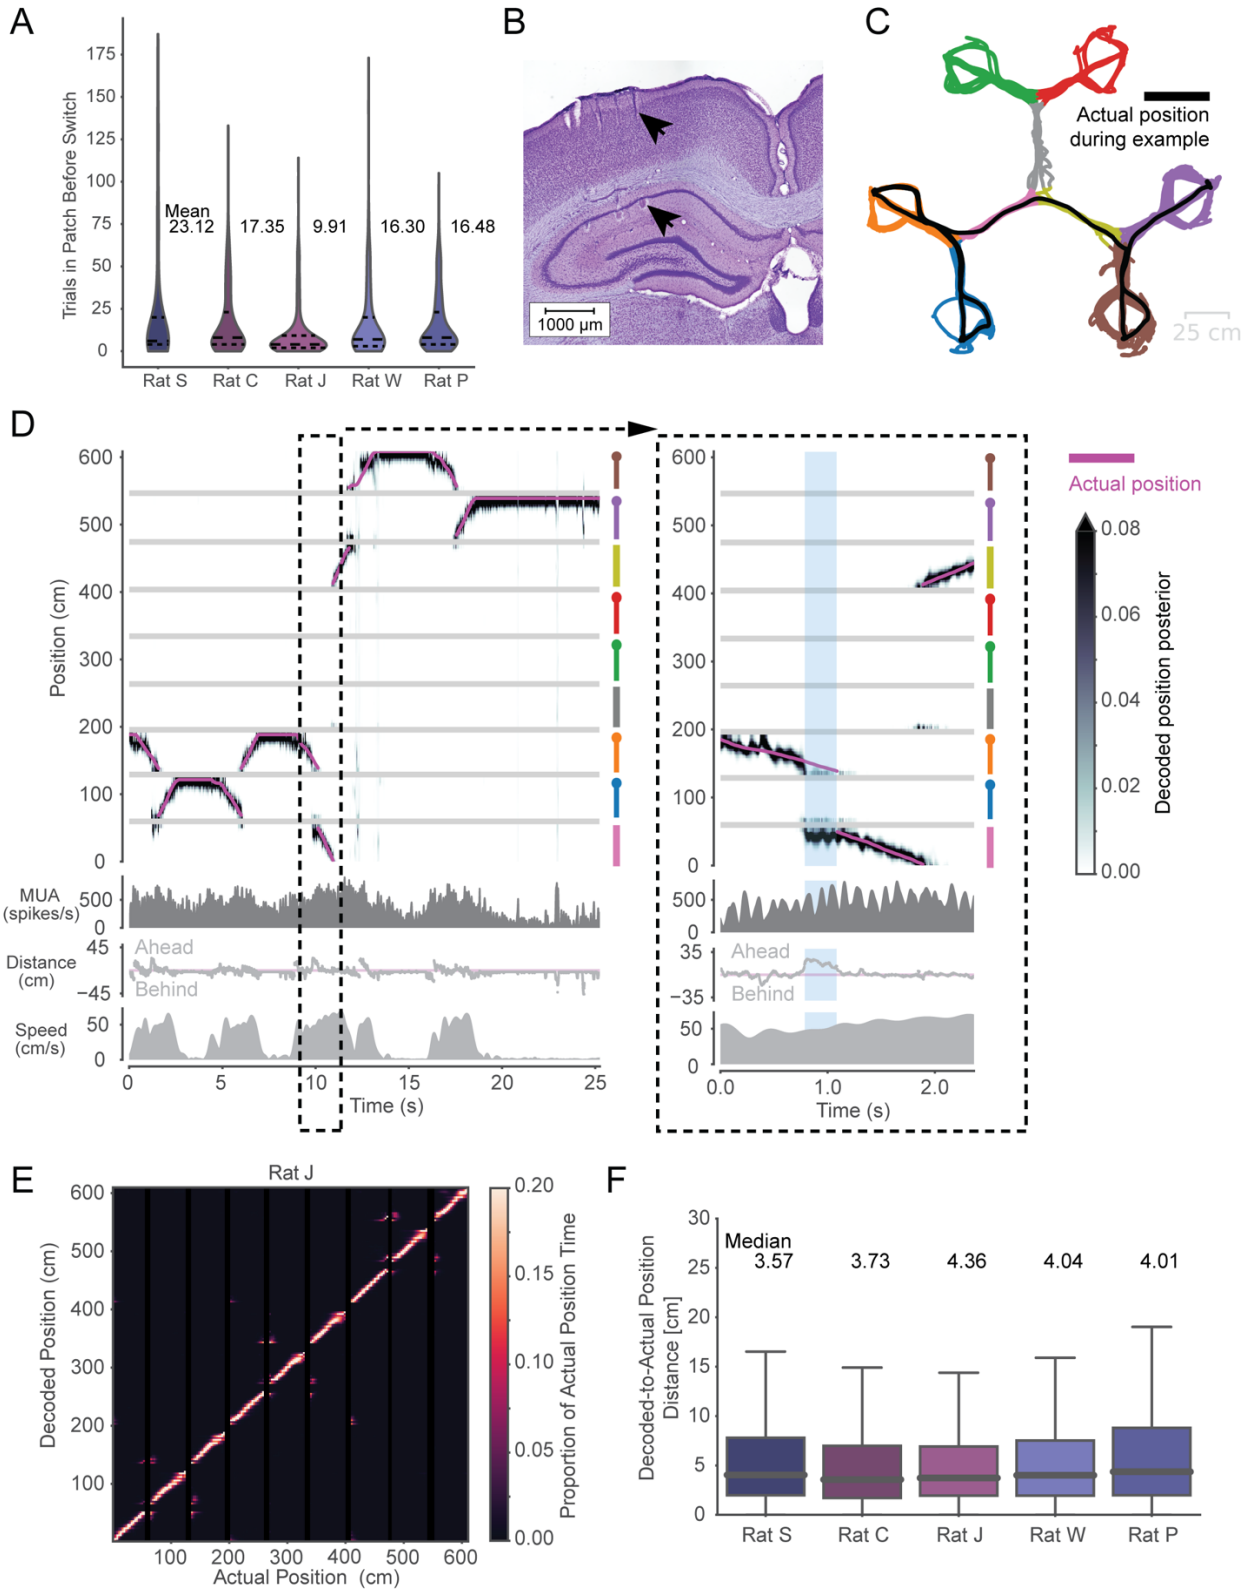

Figure S1. Decoding animal position from Hippocampal spiking during patch foraging.
(A) Violin plots showing distribution of number of trials spent in a patch before Switching, with horizontal dashed line at median, and upper and lower dashed lines indicating quartiles. Mean trial bout durations are labeled per animal.
(B) Nissl-labeled coronal brain tissue section representative example showing targeting of tetrodes to dCA1 of the hippocampus. Black arrows point to a tetrode track in dorsal cortex above and a lesion from the tetrode tip in the pyramidal cell layer below.
(C) Animal head position over the course of one example session, with position data colored by the currently-occupied track segment. Black line highlights the animal's trajectory throughout the same ~25s period shown in **D**. Colored track segments correspond to the colored track segments in **D**.
(D) Decoded position from hippocampal population spiking tracks the animal's actual position at a behavioral timescale and can sweep ahead or behind the animal to represent non-local positions at a sub-second timescale. Dashed lines highlight a small period that is enlarged at right. This period shows an example non-local representation (blue vertical bars) where the decoded position is in a segment distinct from the animal's actual position. Top: Actual head position in 1D (linearized) is shown in pink and decoded position posterior is shown in greyscale. Track segments are aligned on the right y-axis, and segment colors correspond to segments in **C**. Circles at the end of colored segment lines represent reward port locations. Note that actual and decoded position are stationary at reward port positions, as animal pokes into reward port. Horizontal grey lines correspond to 15 cm gaps introduced between track segments in linear position space. Top-middle: Multiunit spike rate across hippocampal tetrodes. Note fluctuations at roughly 8 Hz theta frequency, as expected during running. Bottom-middle: Distance of decoded position from animal's actual position, which can be either ahead (positive values), at (zero cm), or behind (negative values) the actual position. Bottom: Animal head speed.
(E) Confusion matrix showing actual position and decoded position for one rat. Decoded position largely tracks animal's actual position across all track segments. Small amounts of off-diagonal density tend to occur at intersections of track segments, corresponding to adjacent positions in 2D space, as expected.
(F) Boxplots of distribution of distance between decoded and actual position across valid decoded run times for each animal. Boxes show quartiles, whiskers correspond to the data range, and horizontal lines indicate medians, which are also labeled above.

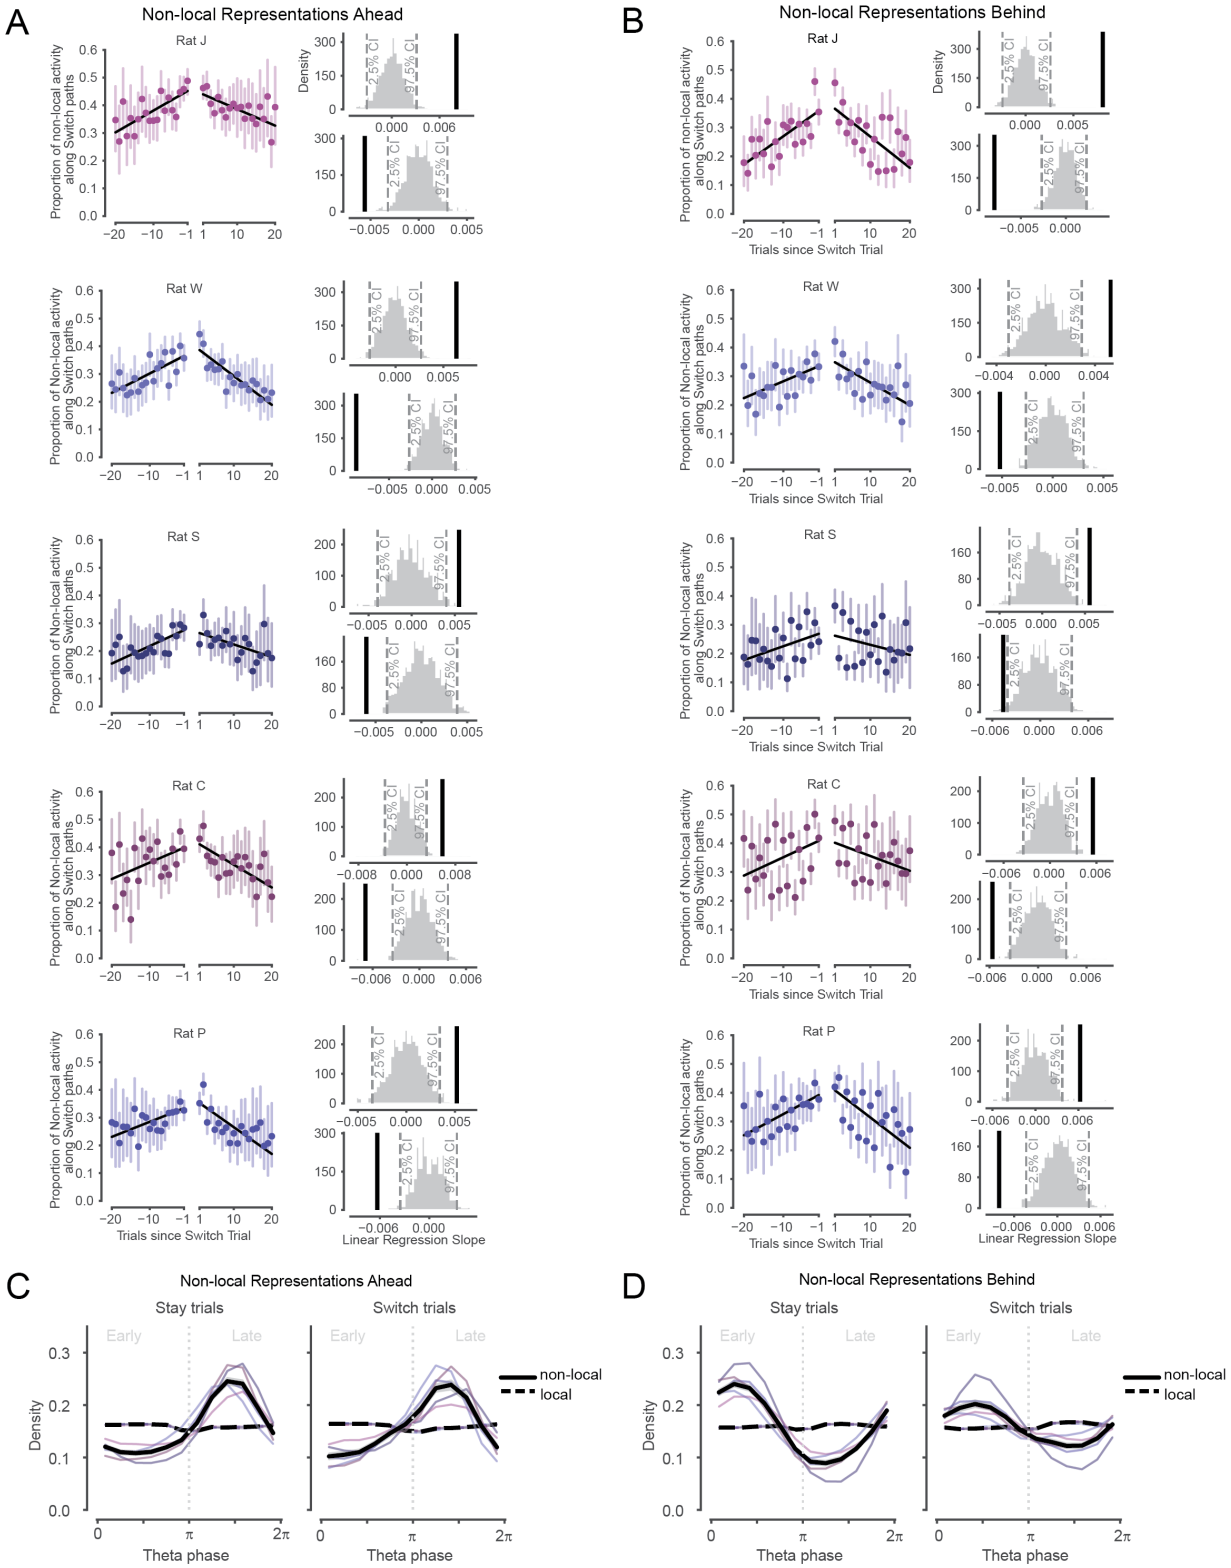

Figure S2. Non-local representations of alternative paths Ahead and Behind are enriched across trials before and after patch Switching.

(A) Proportion of all non-local activity that represents paths consistent with Switching on Stay trials before and after Switch trial for each animal. Data correspond to Stay trials when animals were located in the

first track segment and approaching the choice point, as in Fig. 2B-D. Error bars are 95% CIs on the mean. Pre- and post-Switch linear regressions overlaid in black. All slopes are significantly different than 0 (p_{pre}=0.002, 0.002, .0.014, 0.002, 0.004, p_{post}=0.002, 0.002, 0.008, 0.006, 0.004). Upper right plot: slope of pre-switch linear regression (black) is greater than the 97.5% CI on the slopes from 1000 shuffles of the underlying data (grey). Lower right plot: slope of post-switch linear regression (black) is less than the 2.5% CI on the slopes from 1000 shuffles of the underlying data (grey).

(B) Same as **A**, but for non-local activity that represents paths consistent with Switching as the animal traverses the final segment of each trial and approaches the reward port, as in Fig. 3B-D. All slopes are significantly different than 0 (p_{pre}=0.002, 0.002, 0.004, 0.01, 0.008, p_{post}=0.002, 0.032, 0.004, 0.002, 0.002).

(C) Non-local representations of paths ahead (solid lines) are concentrated in late phases of the theta rhythm, compared to local representations corresponding to the current track segment (dashed lines), on both Stay trials (left) and Switch trials (right). All animal data in black with 95% CI on the mean in grey band, and individual animal data colored per **A.** Early and late phases are separated with a vertical dotted grey line, and labeled in grey text above. Local and non-local distributions are significantly different in each animal for Stay trials (p = 4.4e-64, 1.0e-16, 2.5e-175,8.2e-129, 2.9e-18, Kuiper test) and Switch trials (p = 1.2e-7, 1.4e-11, 3.3e-12, 3.8-19, 9.3e-14).

(D) Same as **C**, but for non-local representations of paths behind (solid lines) and local representations corresponding to the current track segment (dashed lines). Non-local paths behind are concentrated in early phases of theta. Local and non-local distributions are significantly different in each animal for Stay trials (p = 2.2e-7, 1.3e-52, 2.9e-202, 3.5e-23, 1.0e-14, Kuiper test) and Switch trials (p = 8.7e-4, 0.040, 1.3e-16, 0.015, 0.016).

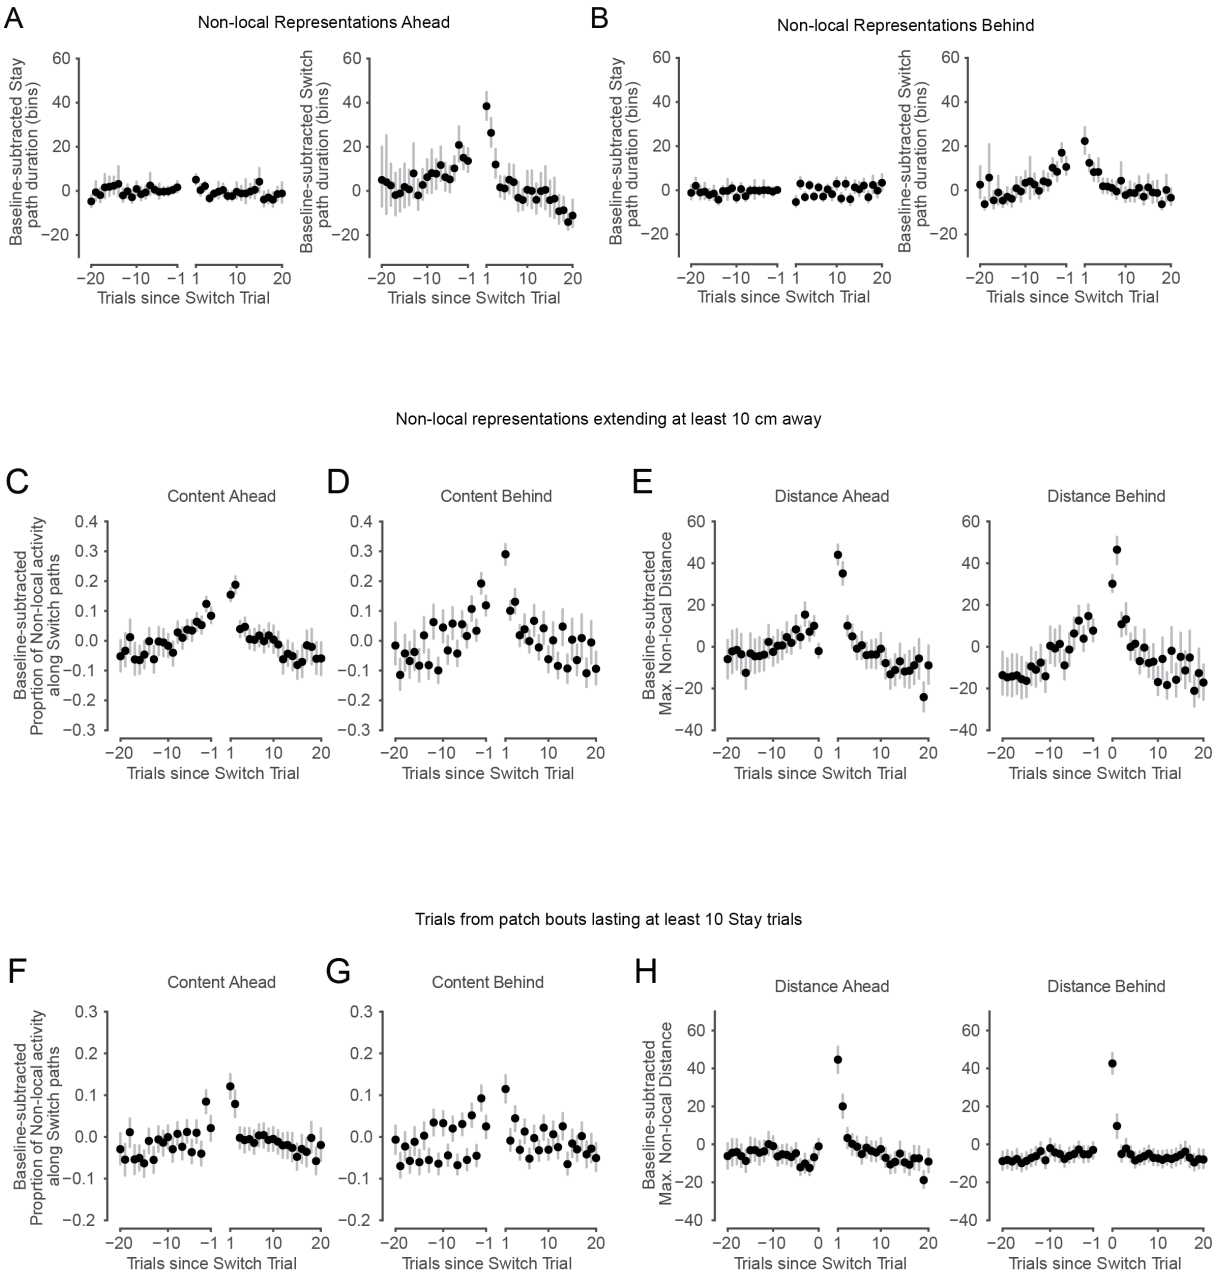

Figure S3. Non-local representations Ahead and Behind are flexibly engaged around Switch trials.

(A) Baseline-subtracted Stay path duration (left) and Switch path duration (right) across all animals. Durations are in 2 ms bins and quantify the number of bins where non-local representations were present during the approach of the first choice point across trials before and after Switch trials. Error bars are 95% CIs on the mean. Switch path durations were modulated more than were Stay path durations leading up to and following Switch trials. Related to Fig. 2D. Baseline durations per animal ranged 21.12-25.68 bins for Stay path representations and 14.61-28.03 bins for Switch path representations.

(B) Same as **A**, but for non-local representations occurring during traversal of the final track segment and approach of the reward port across trials before and after Switch trials. Related to Fig. 3D. Baseline durations per animal ranged 13.99 – 26.65 bins for Stay path representations and 5.7-14.17 bins for Switch path representations.

(C) Baseline-subtracted proportion of non-local representations along Switch paths, across all animals. These representations were expressed during the approach of the first choice point across trials before

and after Switch trials. Only non-local representations at least 10 cm from the animal are included. Error bars are 95% CIs on the mean. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Fig. S2A ($slope_{pre}$=0.0073, p=0.002; $slope_{post}$=-0.0094, p=0.002). The similarity to Fig. 2D demonstrates that our results are not driven by representations very close to the animal.
(D) Same as **C**, but for non-local representations occurring during traversal of the final track segment and approach of the reward port. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Fig. S2B ($slope_{pre}$=0.0099, p=0.002; $slope_{post}$=-0.011, p=0.002). The similarity to Fig. 3D demonstrates that our results are not driven by representations very close to the animal.
(E) Baseline-subtracted maximum non-local distance occurring during approach of the first choice point (left) and traversal of the final track segment and approach of the reward port (right), across all animals. Only non-local representations at least 10 cm from the animal are included. Error bars are 95% CIs on the mean. Distances are especially elevated upon Switching patches and for a couple trials thereafter, and then decrease toward baseline levels. The similarity to Figs. 4D and G demonstrates that our results are not driven by representations very close to the animal.
(F) Baseline-subtracted proportion of non-local representations along Switch paths, across all animals. Data correspond to the approach of the first choice point across trials before and after Switch trials. Here, non-local representations are only included from trials in bouts within a patch lasting at least 10 Stay trials. Error bars are 95% CIs on the mean. Proportion increases before Switch trials and decreases following Switch trials. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Fig. S2A ($slope_{pre}$=0.0033, p=0.002; $slope_{post}$=-0.0047, p=0.002). The approximately symmetrical pattern around the Switch is similar to that seen in Fig. 2D, indicating that the results are not only driven by periods where animals Stayed in a patch for a small number of trials.
(G) Same as **F**, but for non-local representations occurring during traversal of the final track segment and approach of the reward port across trials before and after Switch trials. Error bars are 95% CIs on the mean. Pre- and post-switch linear regression slopes are significantly different than 0 compared to shuffles, as in Fig. S2B ($slope_{pre}$=0.0034, p=0.002; $slope_{post}$=-0.0036, p=0.002). The approximately symmetrical pattern around the Switch is similar to that seen in Fig. 3D, indicating that the results are not only driven by periods where animals Stayed in a patch for a small number of trials.
(H) Baseline-subtracted maximum non-local distance during the approach of the first choice point (left) and traversal of the final track segment and approach of the reward port (right) across all animals. Here, non-local representations are only included from trials in bouts within a patch lasting at least 10 Stay trials. Distances are asymmetric around Switch trials, and are especially enhanced upon Switching patches, and then decrease across trials after the Switch. Related to Figs. 4D and G, respectively.

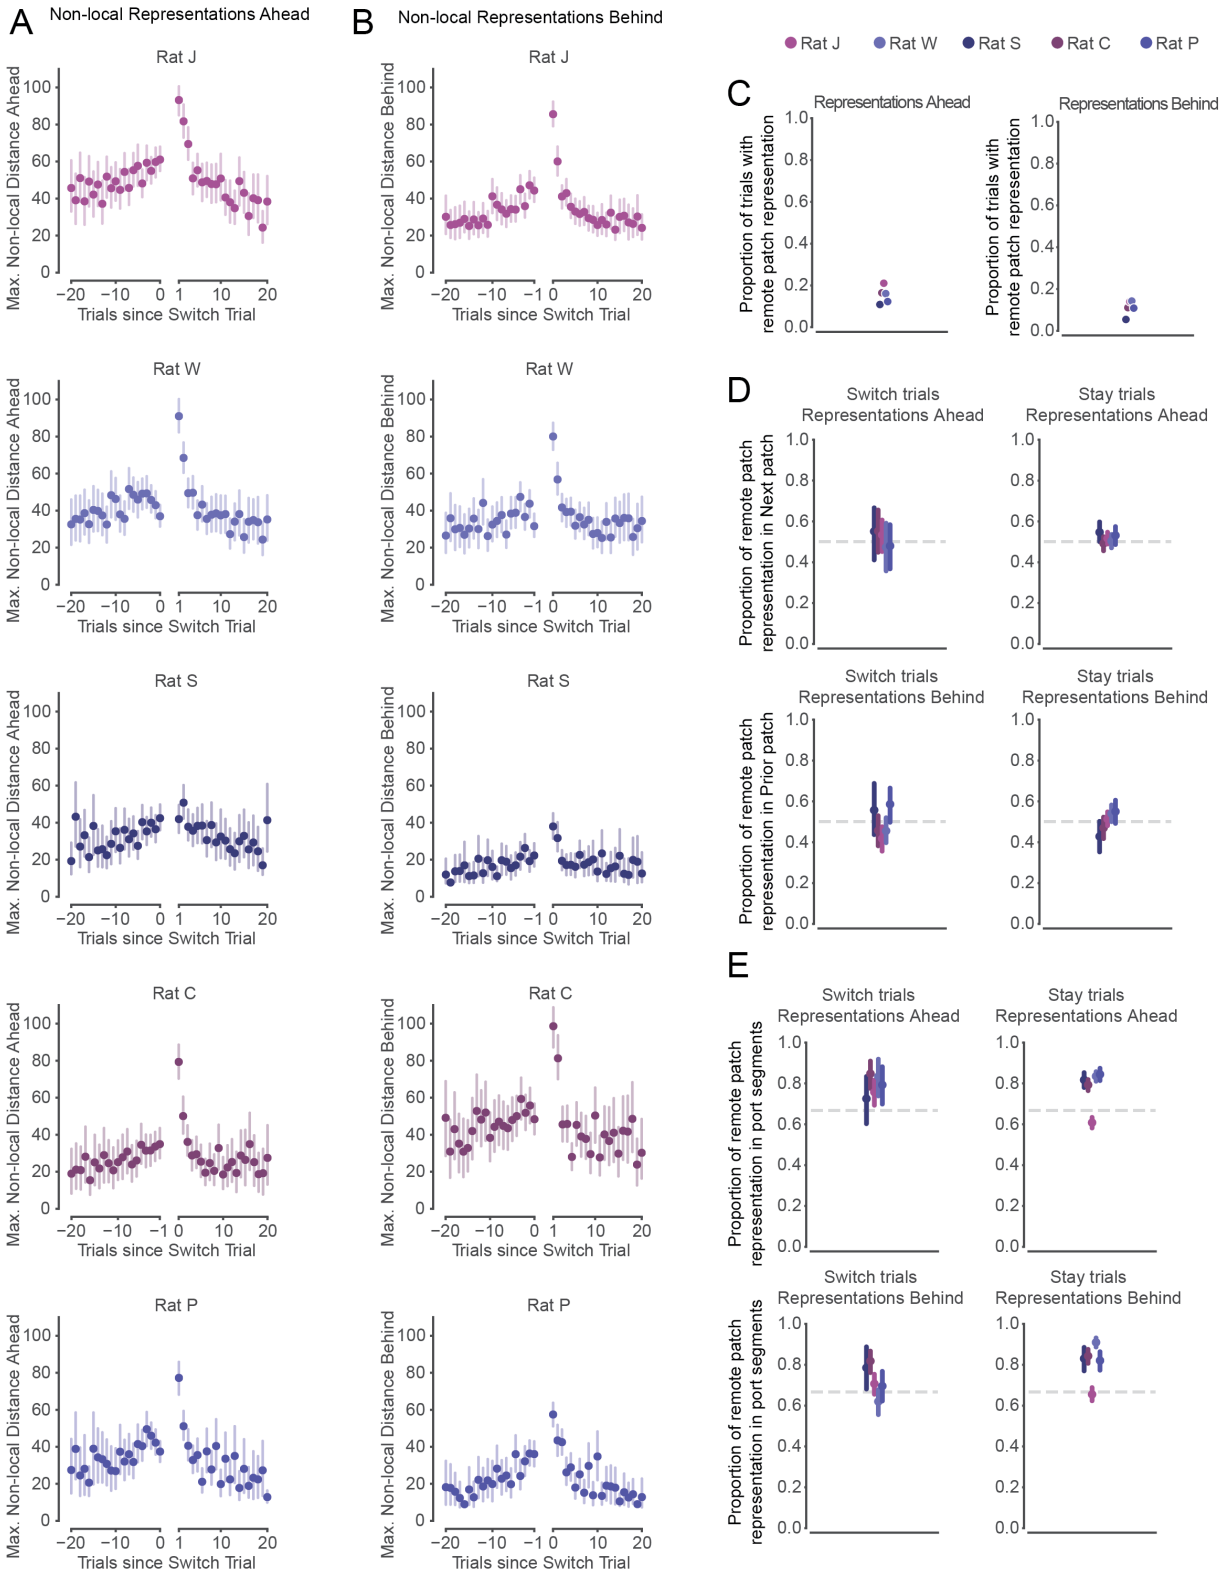

Figure S4. Non-local representations of distant locations.

(A) Maximum non-local distance represented on trials leading up to and following patch Switches for each animal. Data are from the period of each trial in which the animal approached the first choice point. Error bars are 95% CIs on the mean. Related to Fig. 4D.

(B) Maximum non-local distance represented on trials leading up to and following patch Switches, during the period of each trial in which the animal traversed the final track segment towards the reward port, for each animal. Error bars are 95% CIs on the mean. Related to Fig. 4G.
(C) Proportion of trials with non-local representations corresponding to remote, unoccupied patches, out of all trials with any non-local representations, for each animal. Left: calculated for non-local representations occurring as animals approached first choice point. Right: calculated for non-local representations occurring as animals traversed final track segment and approached reward port. In both cases, ~10-20% of trials contain representations corresponding to locations in remote patches.
(D) Non-local representations corresponding to locations in remote patches are not biased to represent subsequently chosen or immediately previous patches. Top row: proportion of remote patch representations occurring during approach of first choice point on Switch trials (left) and Stay trials (right) that correspond to subsequently chosen (next) patch. Bottom row: proportion of remote patch representations occurring during traversal of final track segment and approach of reward port on Switch trials (left) and Stay trials (right) that correspond to most recently chosen (prior) patch. Error bars are 95% CIs on the mean. Grey dashed lines correspond to an equal (50%) representation of the chosen and non-chosen patches.
(E) Non-local representations in remote patches are approximately equally distributed across track segments. Top row: proportion of remote patch representations occurring during approach of first choice point on Switch trials (left) and Stay trials (right) that correspond to track segments containing reward ports. Bottom row: proportion of remote patch representations occurring during traversal of final track segment and approach of reward port on Switch trials (left) and Stay trials (right) that correspond to track segments containing reward ports. Error bars are 95% CIs on the mean. Grey dashed lines correspond to a chance (66.67%) representation of the 4 segments containing reward ports out of 6 total track segments in remote patches.

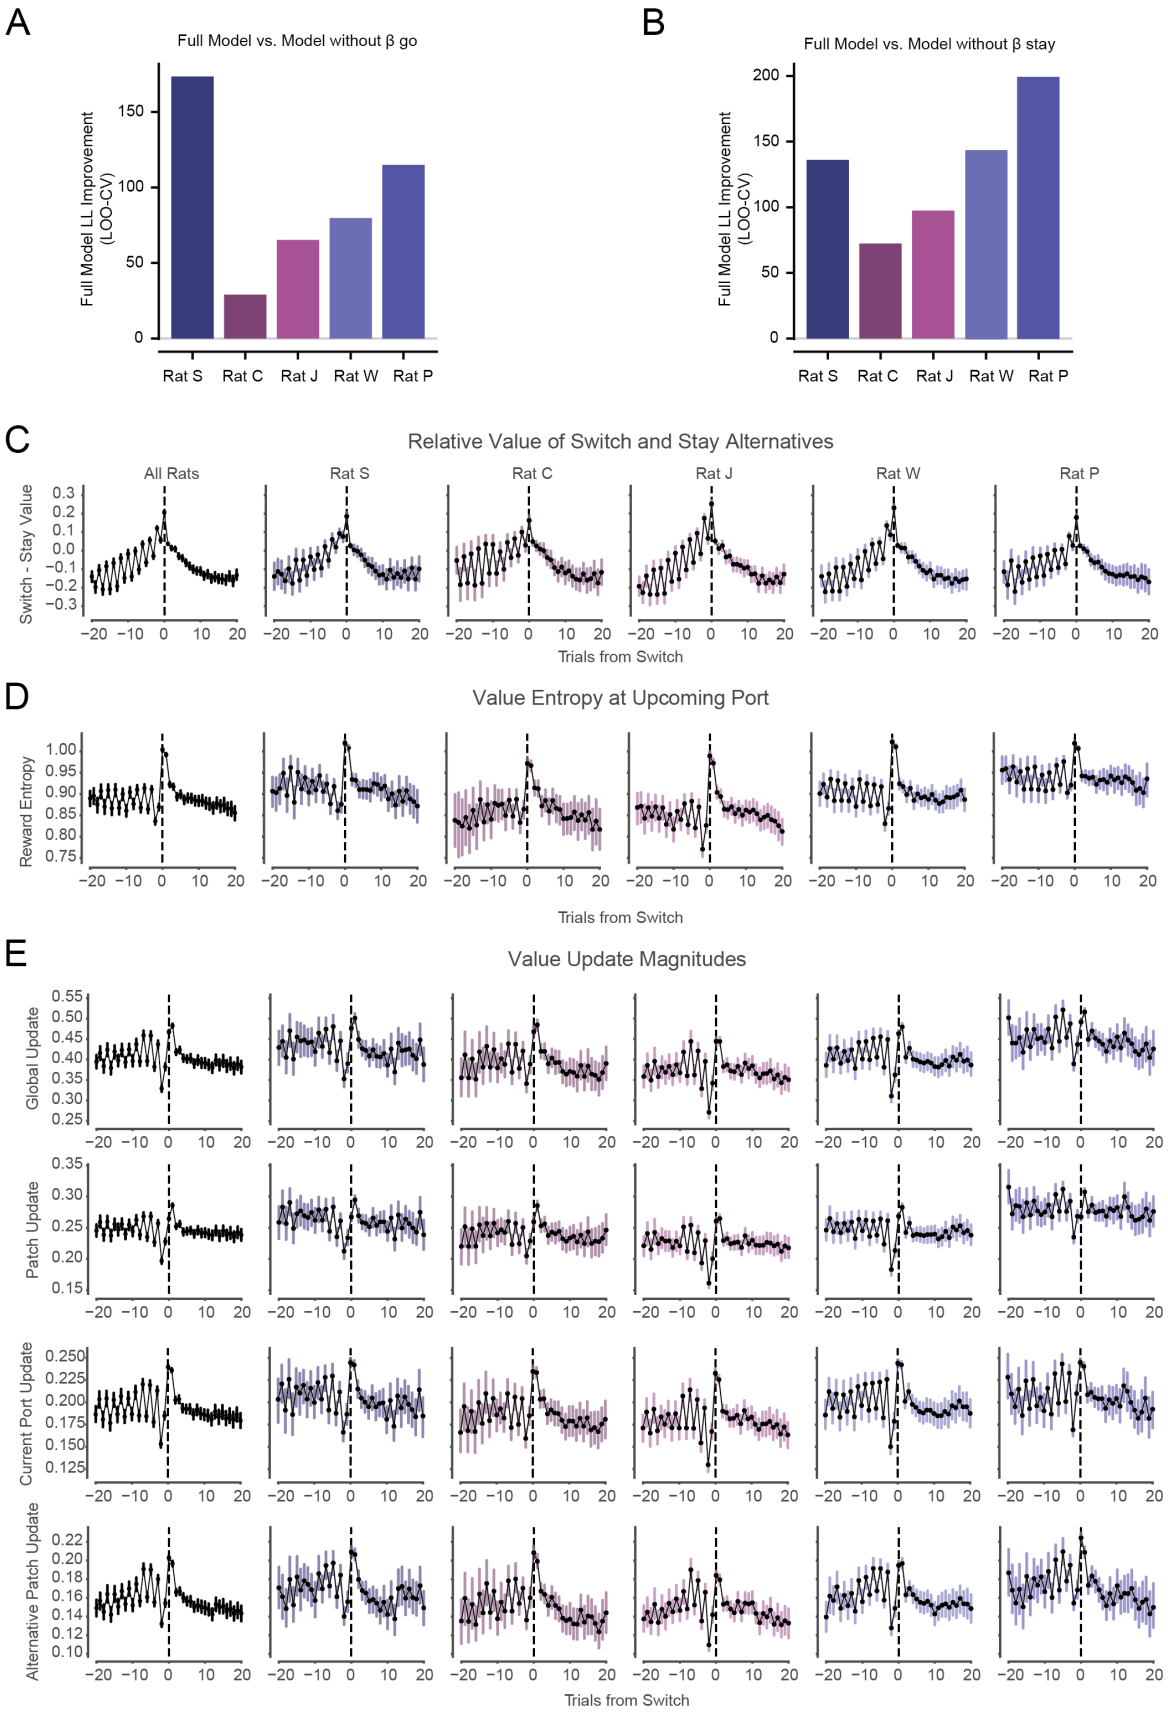

Figure S5. Behavioral model reward sensitivity and variables for individual animals.
(A) Leave-one-out cross-validated log-likelihood improvement (higher is better) of Full Model versus a model where β_{go} was fixed to 0 (such that choice predictions ignored the estimate of the current patch value). Full Model significantly better fit behavior in all animals but one (p=1e-4, 0.113, 0.002, 0.016, 0.001, t-test on cross-validated log-likelihoods per day).
(B) Leave-one-out cross-validated log-likelihood improvement (higher is better) of Full Model versus a model where β_{stay} was fixed to 0 (such that choice predictions ignored the estimate of the alternative patch values). Full Model significantly better fit behavior in all animals (p=0.002, 0.002, 0.002, 0.0002, 0.003, t-test on cross-validated log-likelihoods per day).
(C) Relative value of Switching and Staying increases leading up to Switch trials and decreases afterwards. The value of Staying is the behavioral model-estimated value of the upcoming port within the current patch. The value of Switching is the behavioral model-estimated value of the greater value patch, where patch value is the average of the two ports within. Across-trial dynamics are roughly symmetric around the Switch. Error bars are 95% CIs on the mean. Related to Fig. 1F.
(D) Entropy over behavioral model-estimated value states in the upcoming (chosen) port on each trial, increasing on and after patch Switches. Across-trial dynamics are asymmetric around the Switch. Error bars are 95% CIs on the mean. Related to Fig. 5D.
(E) Value updates on each trial increase leading up to and especially on and after Switch trials, then decay across trials back to baseline. This pattern is observed across ports in the maze, not only at the currently visited port, indicating a period of enhanced learning by value updating upon patch Switching. Value updates are calculated as the absolute magnitude of the change in value across all ports or a subset of ports as follows. Top: Global value updates across all ports in the environment, as in Fig. 5D. Top-middle: Value updates at the current port. Bottom-middle: Value updates at ports in the current patch. Bottom: Value updates at ports in unoccupied, alternative patches. Error bars are 95% CIs on the mean. Related to Fig. 5D.