Figure 1. Recognition Memory Procedure used by Faraut et al. (2018)

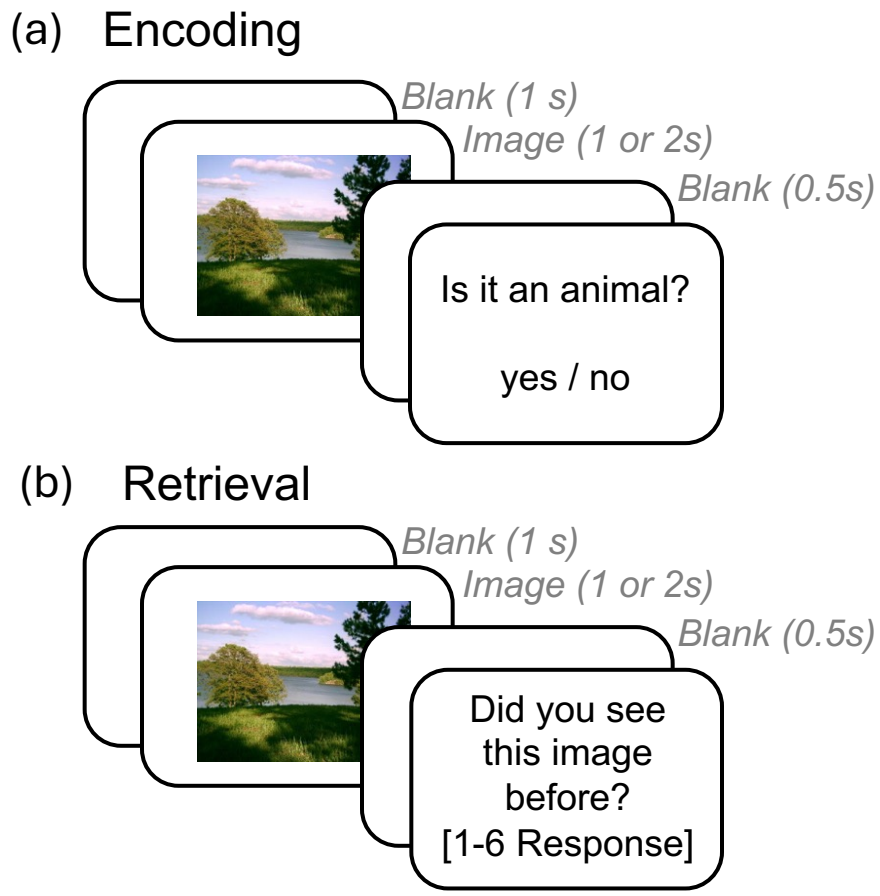

Note. Recognition memory task experimental design (reproduced and adapted from Faraut et al., 2018 under Creative Commons License 4.0). The published dataset consists of single-unit recordings from the hippocampus and amygdala from 59 epilepsy patients while completing a recognition memory test, for a total of 89 experimental sessions (Faraut et al., 2018, Chandravadia et al., 2020). Each session consisted of an encoding and retrieval phase. During the encoding phase (a), participants studied images (n=100) from five different visual categories: animals, people, cars/vehicles, outdoor scenes/houses, and flowers/food items. Each trial consisted of a delay (1s blank screen), the presentation of an image (1- or 2- seconds), a second delay (0.5s blank screen), followed by a yes/no question to promote encoding (i.e., "Is this an animal?), with unlimited time to respond. (b) Approximately 15 minutes after the encoding phase, participants completed an old/new recognition memory test during the retrieval phase. Each trial consisted of a delay (1s blank screen), the presentation of an image (1- or 2-seconds), a second delay (0.5s blank screen), followed by an old/new recognition judgement with confidence ratings (unlimited time to respond). A total of 100 images were shown; 50 were previously studied during the encoding phase ("old"; targets) and 50 were not previously presented during the encoding phase ("new"; foils).

Figure 2. Empirical QQ Plots of Target- vs. Foil-by-Neuron Normalized Spike-Count Distributions at Retrieval

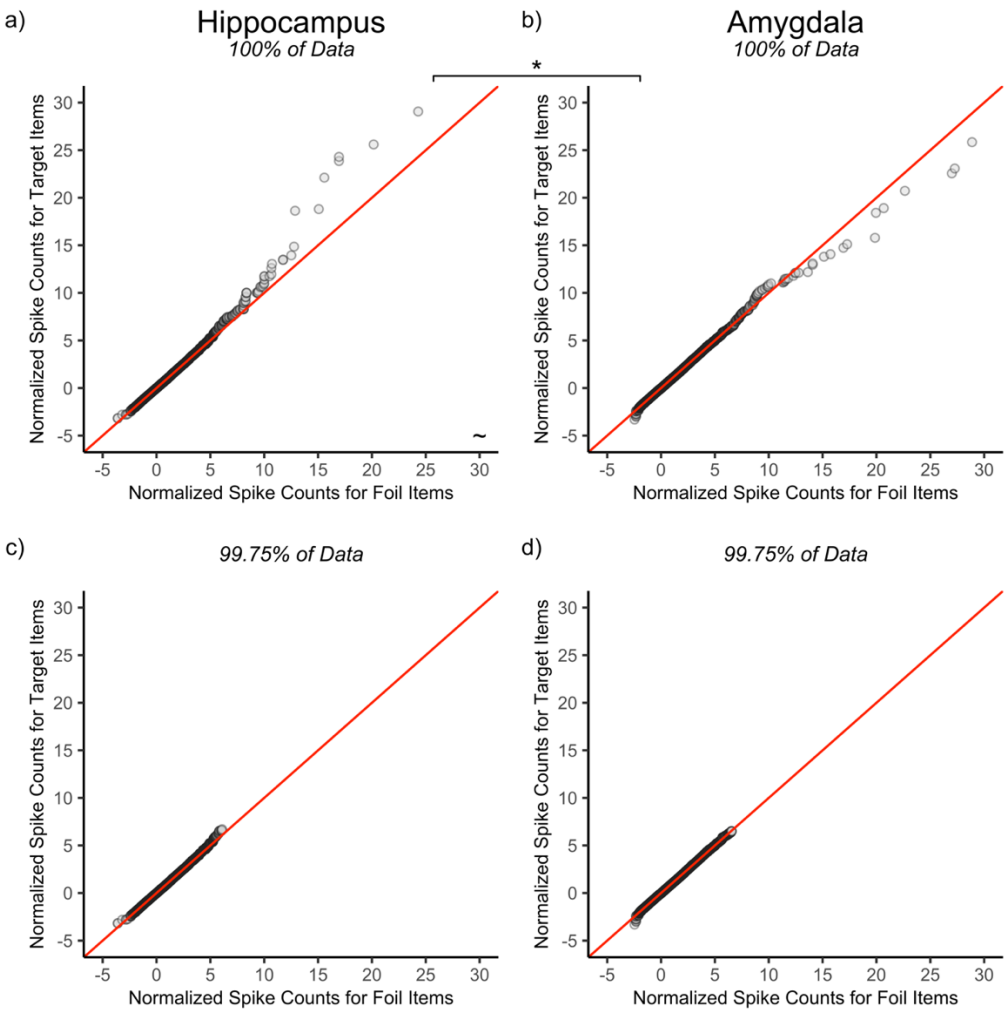

Note. Empirical quantile-quantile (QQ) plots of target-by-neuron normalized spike count distributions (y-axis) vs. foil-by-neuron normalized spike-count distributions (x-axis) in the hippocampus (a and b) and the amygdala (c and d). The top panels plot the distributions for 100% of the recordings from the hippocampus (panel a, 72,538 recordings) and the amygdala (panel b, 99,944 recordings). Points that are densely grouped (dark grey) represent thousands of recordings whereas less densely grouped points (light grey) represent relatively few recordings. If the points all fell on the red line of equivalence, it would indicate that the two distributions are identical. If the points deviate from the red line but still yield a linear relationship, it would indicate that the two distributions have the same form (e.g., both exponential or both lognormal) but differ in their location and scale parameters (e.g., different mean and/or standard deviation). The observed deflection of points away from the red line of equivalence towards one axis in (a) indicates that the target-by-neuron distribution has a different form than the foil-by-neuron distribution. More specifically, the deflection indication that the target distribution is more skewed to the right relative to the foil distribution. A less pronounced effect in the opposite direction is evident in the amygdala (b). The deflection in

(a) is predicted by a sparse coding account, which further predicts that the deflection reflects the strong responses of a small percentage of neurons in response to target items. After removing the top 0.25% of both the target-by-neuron recordings and foil-by-neuron recordings (c and d), the deflections in (a) and (b) were no longer apparent. The skewness values of the two distributions (targets vs. foils) were compared to determine if the theoretically predicted deflection in panel (a) was statistically reliable. The difference between the target vs. foil distributions was significantly different between the brain regions (hippocampus vs. amygdala) and was marginally significant only within the hippocampus (refer to Methods and Table 1 for detailed statistical reporting, * $p < .05$, ~ $p < .10$).

Figure 3. Statistical Moments of Target- and Foil-by-Neuron Normalized Spike-Count Distributions

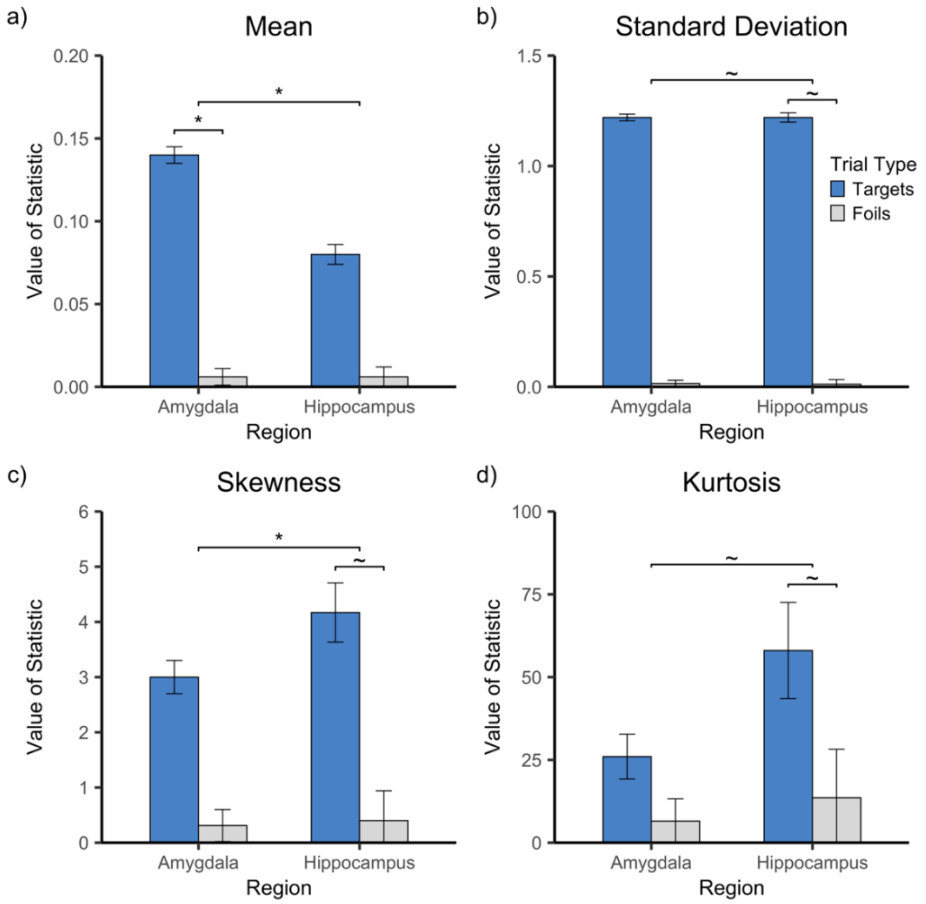

Note. Statistical moments of the target-by-neuron and foil-by-neuron normalized spike count distributions that yielded the QQ plots shown in Figures 1a and 1b: (a) mean, (b) standard deviation, (c) skewness, and (d) kurtosis. Within each panel, data from the amygdala are shown on the left, and data from the hippocampus are shown on the right. Significant differences between the target (blue) and foil (grey) distributions within a region, and significant interactions between regions (hippocampus vs. amygdala), are indicated by horizontal lines above the relevant bars (* = $p < .05$, ~ = $p < .10$). As detailed in Table 1, a significant interaction was observed for mean values (a), with greater average spiking in the amygdala in response to foil items compared to targets, with no significant difference observed for the hippocampus. A marginally significant interaction was present for standard deviation (b) and kurtosis (d), with greater values for the target distribution relative to the foil distribution in the hippocampus, but not the amygdala. Critically, a significant interaction between the hippocampus and amygdala was present for skewness (c), the statistical moment most relevant to the item-specific memory signal under investigation here.

Table 1. Statistical Results for the Statistical Moments of the Target vs. Foil Distributions at Retrieval

	Interaction	Hippocampus			Amygdala		
	p	Targets	Foils	*p*	Targets	Foils	*p*
Mean	.046*	0.08	0.08	.334	0.14	0.16	.019*
SD	.065~	1.22	1.18	.054~	1.22	1.23	.344
Skewness	.048*	4.17	3.05	.070~	3.00	3.28	.254
Kurtosis	.054~	58.0	29.9	.086~	26.0	34.9	.178

Note. A significant interaction was defined as a greater difference in a given statistical moment (e.g., skewness) between the target-vs.-foil normalized spike count distributions in the hippocampus compared to the corresponding difference in the amygdala according to bootstrap tests (k=10,000, $p < .05$, two-tailed). The interactions for mean and skewness were significant. Differences in each statistical moment between the two distributions (target-vs.-foil) were also considered separately within each structure (hippocampus and amygdala; bootstrapped k=10,000, $p < .05$, one-tailed). In the hippocampus,, differences in the upper moments of the target and foil spike count distributions were not significant but trended in that direction (standard deviation, $p = .054$; skewness, $p = .07$; kurtosis, $p = .086$; one-tailed). In the amygdala, the differences were not significant (standard deviation, $p = .344$; skewness, $p = .254$; kurtosis, $p = .178$). Finally, mean firing in the amygdala distinguished between target and foil items ($p=0.019$) but not in the hippocampus ($p = .334$). * = $p < .05$, ~ = $p < 0.10$.

Table 2. Statistical Results for Skewness of the Target vs. Foil Distributions at Retrieval Partitioned by the Pattern of Target Firing at Encoding

		Skewness Value					
	Interaction	Hippocampus			Amygdala		
	p	Targets	Foils	*p*	Targets	Foils	*p*
Low-High	.007**	6.79	3.05	.004**	3.27	3.28	.989
High-High	.549	2.61	3.05	.716	3.18	3.28	.878
High-Low	.504	2.21	3.05	.527	2.64	3.28	.388
Low-Low	.965	3.10	3.05	.924	2.42	3.28	.173

Note. As in Table 1, a significant interaction was defined as a greater difference in skewness between the target-vs.-foil normalized spike count distributions in one region (e.g., the hippocampus) compared to the corresponding difference the other region (e.g., the amygdala) according to bootstrap tests (k=10,000, $p < .05$, Bonferroni corrected). The interaction was significant for the test involving targets associated with excitable neurons at encoding (Low-High, $p < .001$) but not for any other category of targets at encoding (High-Low, $p = .549$, High-High, $p = .504$, or Low-Low, $p = .965$). Within the hippocampus, a significant difference between target vs. foil skewness (theoretically indicative of the item-specific memory signal) was observed only for targets that were associated with excitability at encoding (Low-High; $p < .001$). Within the amygdala, no significant difference in skewness between the target and foil distributions was observed (Low-High, $p = .989$; High-High, $p = .878$; High-Low, $p = .388$; Low-Low, $p = .173$). Corresponding analyses for the other statistical moments are reported in Supplemental Table 1. ** = $p < 0.01$

Figure 4. Empirical QQ Plots of the Normalized Spike-Count Distributions at Retrieval Partitioned by Neuronal Excitability at Encoding

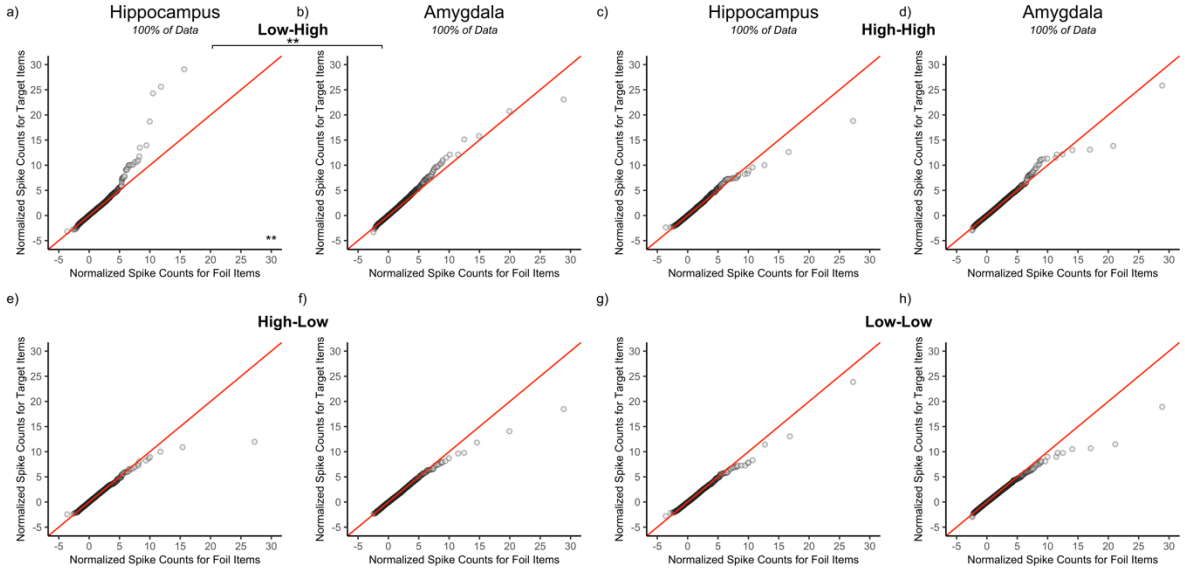

Note. As in Figure 2, normalized spike count distributions were plotted to visualize differences in skewness between the target-by-neuron distributions (y-axis) and foil-by-neuron (x-axis) distributions. Here, separate plots were created for each of the four spiking patterns at encoding (Low-High, High-High, High-Low, and Low-Low; see Methods). Thus, four panels show retrieval data from the hippocampus partitioned by the encoding pattern observed in the hippocampus (panels a, c, e, and g), and four panels show retrieval data from the amygdala partitioned by the encoding pattern observed in the amygdala (panels b, d, f, and h). The same foil-by-neuron (x-axis) distribution based on recordings from the hippocampus is used in all four plots for the hippocampus, and the same foil-by-neuron (x-axis) distribution based on recordings from the amygdala is used in all four plots for the amygdala. Visually, a sharp deflection towards the y-axis (indicative of greater skewness for the target distribution relative to the foil distribution) is evident for targets associated with heightened excitability at encoding (Low-High) in the (a) hippocampus but not in the (b) amygdala. As detailed in Table 2, the interaction between the skewness of the two distributions (targets vs. foils) and brain region (hippocampus vs. amygdala) was significant, as denoted by the horizontal line above panels a & b with double asterisks. Within the hippocampus, the apparent visual difference in skewness was also significant (as denoted by ** in panel (a)), but no significant difference was detected for the amygdala. No visual or statistical evidence of a skewness difference between the target and foil distributions was observed for targets associated with the other spiking patterns at encoding, for either the hippocampus or amygdala Thus, the difference in skewness was statistically reliable only in the hippocampus, and only for targets that were associated with excitable hippocampal neurons at encoding. Supplemental Figure 1 reports graphs with the top 0.25% of both the target-by-neuron and foils-by-neuron distributions were removed. The deflection observed in panel (a) disappeared after removing the top 0.25% of data, indicating that relatively few neurons fired strongly in response to targets mostly associated with excitable neurons at encoding compared to foil items, and only within the hippocampus. ** = $p < 0.01$

Figure 5. Empirical QQ Plots for Remembered Targets vs. Foils and Forgotten Targets vs. Foils

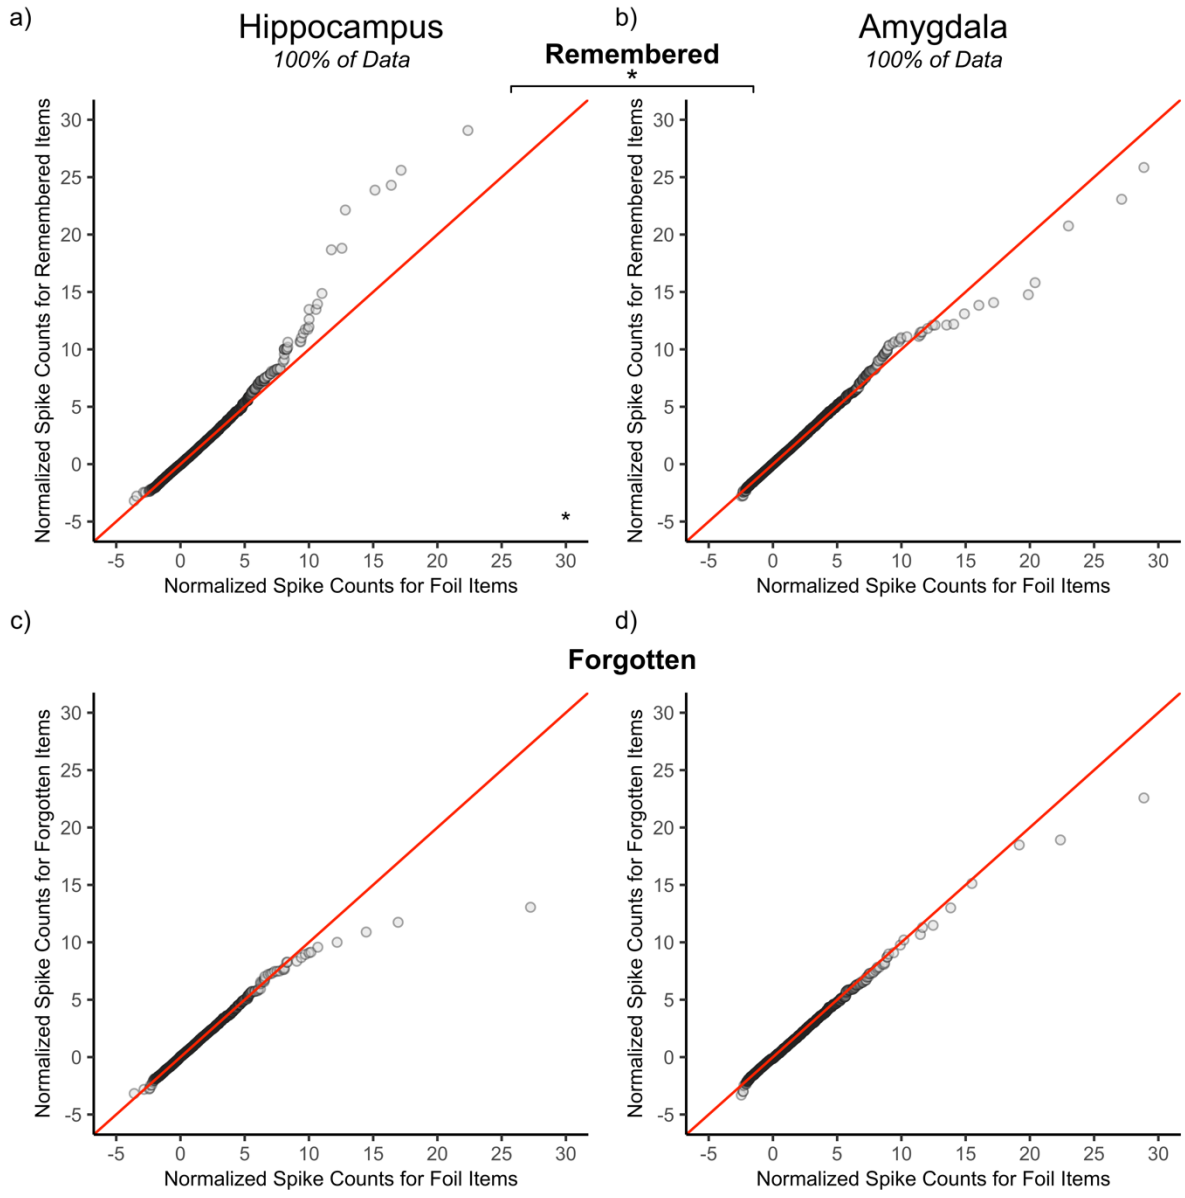

Note. As in Figures 2 & 3, normalized spike count distributions were plotted to visualize differences in skewness between the target-by-neuron (y-axis) and foil-by-neuron (x-axis) distributions, separately for targets that were remembered (top row) or forgotten (bottom row), and for both the hippocampus (a & c) and amygdala (b & d). Visually, in the hippocampus, a sharp deflection towards the y-axis (indicative of greater skewness of the target distribution relative to the foil distribution) is visually evident for remembered targets (a). A smaller trend in the opposite direction was observed in the amygdala (b). As detailed in Table 3, the interaction between the skewness of the two distributions (remembered targets vs. foils) and brain regions (hippocampus vs. amygdala) was significant, as denoted by the line with an asterisk above panels (a) and (b). Within each region considered separately, the difference in

skewness for the remembered target-by-neuron distribution was significantly greater than the foil-by-neuron distribution (denoted by * in the lower right of panel a) in the hippocampus, but the slight difference in the opposite direction in the amygdala was not significant. No visual or statistical evidence of a difference in skewness was evident for forgotten targets in either the hippocampus (c) or amygdala (d). Thus, the difference in skewness associated with the item-specific memory signal (the theoretical signature of a sparse episodic memory code) was statistically reliable only in the hippocampus, and only for targets that were remembered. Supplemental Figure 2 reports graphs with the top 0.25% of both the target-by-neuron and foils-by-neuron distributions were removed. The deflection observed in panel (a) disappeared after removing the top 0.25% of data, indicating that relatively few neurons fired strongly in response to remembered targets compared to foil items, and only within the hippocampus.

Table 3. Analyses of the Statistical Moments of Remembered Target vs. Foil Distributions at Retrieval

	Interaction	Hippocampus			Amygdala		
	p	Remembered	Foils	*p*	Remembered	Foils	*p*
Mean	.424	0.09	0.08	.316	0.16	0.16	.395
SD	.042*	1.26	1.18	.015*	1.24	1.23	.634
Skewness	.015*	4.82	3.05	.036*	2.94	3.28	.456
Kurtosis	.022*	72.3	29.9	.064~	24.9	34.9	.345

Note. The interaction between a given statistical moment (e.g., skewness) of the two distributions (remembered targets vs. foils) and brain region (hippocampus vs. amygdala) was significant for SD, skewness, and kurtosis (bootstrapped, k=10,000, p=.042, p = .015, p = .022, respectively). Within each region considered separately, significant differences in the upper moments of the remembered vs. foil item distributions were observed in the hippocampus (bootstrapped, k=10,000, standard deviation, p = .015; skewness, p = .036; kurtosis, p = .064, respectively) but not the amygdala (bootstrapped, k=10,000, standard deviation, p = .634; skewness, p = .456; kurtosis, p = .345). Additionally, mean firing in the amygdala distinguished between remembered and foil items (p = .019) but not in the hippocampus (p = .33). *p < .05 and ~p < .10.

Table 4. Analyses of the Statistical Moments of Forgotten Target vs. Foil Distributions at Retrieval

	Interaction	Hippocampus			Amygdala		
	p	Forgotten	Foils	*p*	Forgotten	Foils	*p*
Mean	<.001***	0.07	0.08	.470	0.09	0.16	<.001***
SD	.304	1.15	1.18	.542	1.18	1.23	.085~
Skewness	.676	2.30	3.05	.478	3.15	3.28	.834
Kurtosis	.602	12.8	29.9	.545	28.78	34.9	.660

Note. The interaction between a given statistical moment of the two distributions (forgotten targets vs. foils) and brain region (hippocampus vs. amygdala) was significant for the mean only ($p < .0001$). Within each region considered separately, the difference between the mean of the forgotten vs. foils item distribution was significant in the amygdala ($p < .001$) but not the hippocampus ($p = .470$). ***$p < .001$ and ~$p < 0.10$.

Figure 6. Empirical QQ Plots for Remembered vs. Forgotten Targets

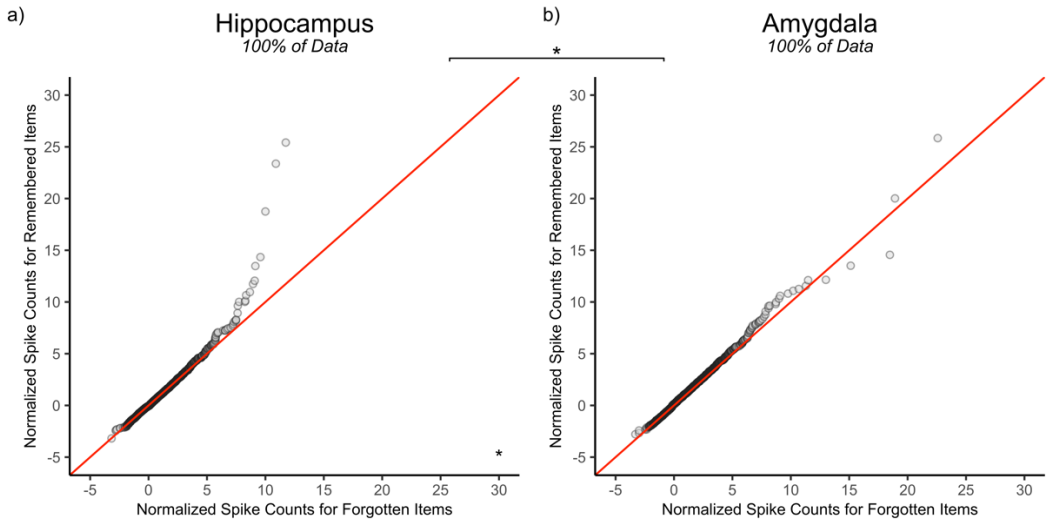

Note. Normalized spike count distributions were again plotted to visualize differences in skewness between the remembered target-by-neuron (y-axis) and forgotten item-by-neuron (x-axis) distributions for the (a) hippocampus and (b) amygdala. Visually, a sharp deflection towards the y-axis was evident for the remembered vs. forgotten item spike count distributions in the (a) hippocampus but not in the (b) amygdala recordings. As detailed in Table 5, the interaction between the skewness of the two distributions (remembered targets vs. forgotten targets) and brain region (hippocampus vs. amygdala) was significant, as denoted by the * above the horizontal line connecting panels (a) and (b). Within each region considered separately, the difference in skewness for the remembered target distribution was significantly greater than the forgotten target distribution in the hippocampus (denoted by * in lower right of panel a) but not in the amygdala. The deflection observed in panel (a) disappeared after removing the top 0.25% of data (Supplemental Figure 3), indicating that relatively few neurons fired strongly in response to remembered targets compared to forgotten targets, and only within the hippocampus.

Table 5. Analyses of the Statistical Moments of Remembered Target vs. Forgotten Target Distributions at Retrieval

	Interaction	Hippocampus			Amygdala		
	p	Remembered	Forgotten	*p*	Remembered	Forgotten	*p*
Mean	<.001***	0.09	0.07	.156	0.16	0.09	<.001***
SD	.213	1.26	1.15	.020*	1.24	1.18	.055~
Skewness	.012*	4.82	2.30	.022*	2.94	3.15	.732
Kurtosis	.022*	72.3	12.8	.050*	24.9	28.8	.789

Note. A significant interaction reflects a greater difference in a statistical moment (e.g., skewness) of the spike count distributions for the remembered vs. forgotten targets in one region compared to the other region (hippocampus vs. amygdala) according to bootstrap tests (k=10,000, $p < 0.05$)). The interactions were significant for the mean, skewness, and kurtosis statistics ($p<.001$, $p = .012$, $p = .02$, respectively). Within each region considered separately, differences in the upper moments of the distributions were detected in the hippocampus (skewness, $p = .022$ kurtosis, $p = .05$) but not the amygdala (skewness, $p = .732$; kurtosis, $p = .789$). Conversely, mean firing was greater for remembered compared to forgotten items ($p < .001$) in the amygdala but not in the hippocampus ($p = .156$). Although the interaction was not significant for standard deviation, within each region, both the hippocampus and amygdala exhibited a significant difference for that statistical moment ($p = .022$, $p = .055$, respectively), with greater SD values for remembered compared to forgotten items.

Figure 7. Empirical QQ Plots for Excitable Neurons as a Function of Subsequent Memory

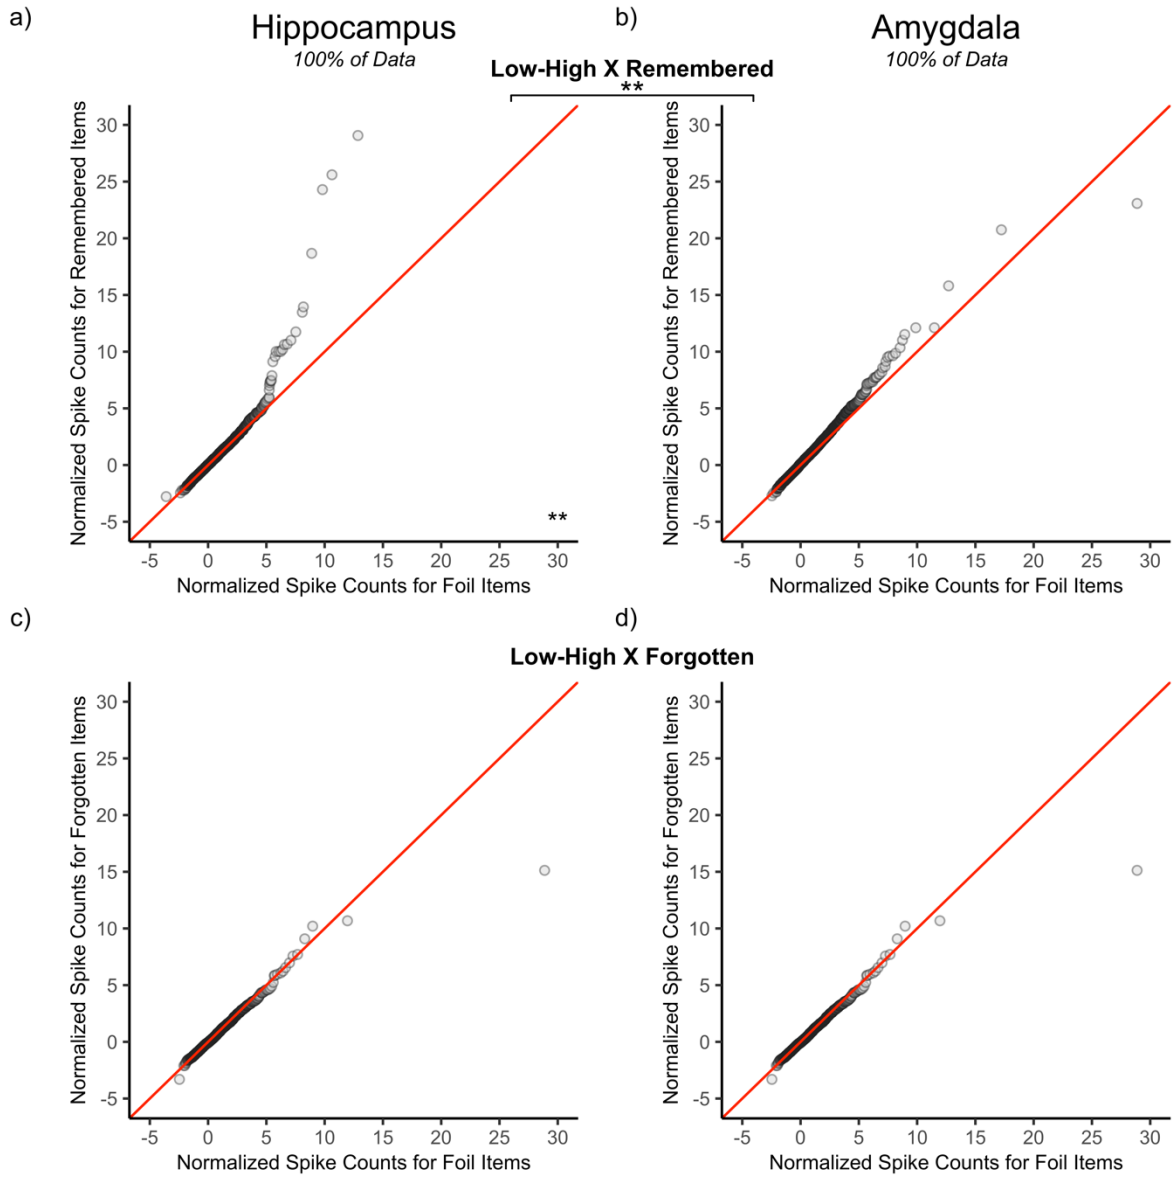

Note. Normalized empirical QQ plots of the excitable (i.e., Low-High pattern at encoding) target-by-neuron response distribution (y-axis) and the foil-by-neuron response distribution (x-axis). Panel (a) plots remembered targets associated with excitable hippocampal neurons at encoding, panel (b) plots remembered targets associated with excitable amygdala neurons at encoding, panel (c) plots forgotten targets associated with excitable hippocampal neurons at encoding, and panel (d) plots forgotten targets associated with excitable amygdala neurons at encoding. In panel (a), a sharp deflection towards the y-axis is evident in the hippocampus for remembered targets that were associated with heightened excitability at encoding. By contrast, in the amygdala (b), most points fell densely on the diagonal line. Additionally, no other subset of targets split by other levels of spiking at encoding (Low-Low, High-High, High-Low) for

remembered or forgotten items showed visual evidence of a difference in skewness (Supplemental Figure 4). The deflection observed in panel a disappeared after removing the top 0.25% of data, indicating that relatively few neurons fired strongly in response to remembered targets compared to foil items, and only within the hippocampus (Supplemental Figure 5). Thus, the difference in skewness associated with the item-specific memory signal was statistically reliable and selective to only the hippocampus, only to targets that were remembered, and only to targets associated with heightened excitability at encoding (bootstrapped, k = 10,000, $p < 0.05$; ** $p < 0.01$).

Figure 8. Statistical Moments of Item-by-Neuron Normalized Spike Distributions by Neuronal Excitability, Subsequent Memory, and Region

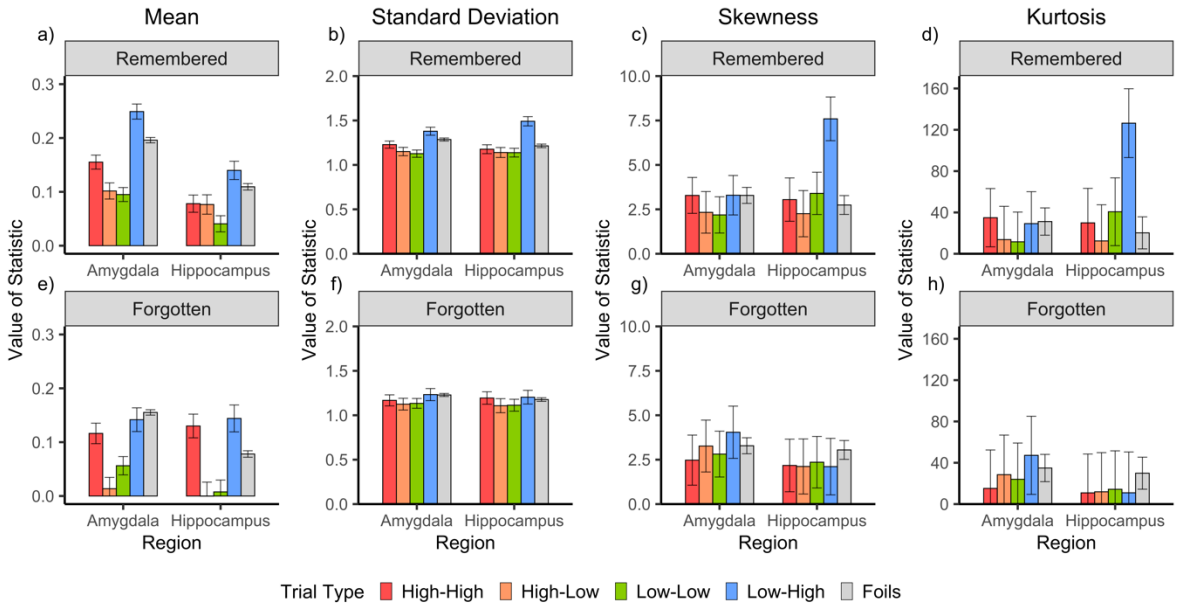

Note. Bar charts representing the statistical moments of the normalized spike counts at retrieval for targets, categorized by their 1) level of spiking at encoding (red: neurons that were consistently high spiking [High-High]; orange: neurons that decreased in spiking [High-Low]; green: neurons that were consistently low spiking [Low-Low]; blue: neurons with heightened excitability [Low-High]) and foils (grey) and 2) subsequent memory response (remembered items: top row, forgotten items: bottom row). The primary statistical moment suggestive of sparse coding, (c) skewness, was conspicuously elevated for targets with heightened excitability at encoding (changed from low pre-stimulus to high post-stimulus firing at encoding, [Low-High: blue]). Additionally, elevations in the (a) mean, (b) standard deviation, and (d) kurtosis values of the spike counts at retrieval were observed for targets that changed from low pre-stimulus to high post-stimulus firing at encoding (Low-High; blue) compared to foils (grey) for the hippocampus. This pattern was selective only to targets that were later remembered (top row) and not forgotten (bottom row).