Supplementary Information

Here we provide further analysis of the data in Figures 3, 4, and 5 shown in the **Results** in the main text. In each case, the learning curves were fitted with an asymptotic function and the curve-fit parameters are reported.

However, the analysis should be treated with caution. A curve fit is not a simple quantification of the data. It is an assumption about the data, and that assumption may be incorrect or at best approximate. In the present data, the asymptotic curve applied better to conditions in which the learning was comparatively better. In conditions in which the agent learned comparatively less well, when the learning curve was flatter or more variable, the asymptotic function fit the data poorly, and in some conditions failed to fit altogether.

For example, in Figure 3A, the blue line shows a rise to an asymptotic level. An asymptote curve can be fit to the data, and the parameters of the function fit are given in Table S1. In contrast, the orange line shows a rise and then a subsequent fall in performance. The assumption that the line fits an asymptote is not right, except in a most approximate way. One can perform the curve fit and compute the parameters (as shown in Table S1), but the results are not as meaningful for the orange line as for the blue line. The meaningful finding in Figure 3A is not dependant on whether or how the data may fit to functions or reach approximate asymptotes. Rather, the meaningful finding is that learning, over the course of the 1500 training iterations (150,000 training games), is substantially better in the blue line (with functioning attention) than in the orange line (without functioning attention).

With this caveat in mind – that the curves in which learning was relatively poorer do not fit as well to an asymptotic function – here we report the results of fitting the data. Tables S1, S2, and S3 shows the results for experiments 1, 2, and 3 respectively. For each learning curve shown in the main text, we fit the data to the curve: $y = a-b*\exp(-x/c)$. In this equation, a = the estimated final asymptotic level of performance; b = the change in performance between initial and asymptotic level; and c = the number of iterations required to reach two thirds of the asymptotic level.

The curve fit failed to converge on an answer in the first two conditions of experiment 3 (running the fit multiple times resulted in widely divergent, nonsensical answers). This failure of the curve fitting occurred for the two conditions in which the attention tracking score was measured. Thus, if one wishes to use these approximate curve fits to evaluate the data, then, in experiment 3, one must rely on the catch scores only and not on the attention tracking scores. The result is the same, however one evaluates it. The data match the predicted pattern of better performance with an attention schema and worse performance without.

Condition	Parameter a	Parameter b	Parameter c
Attention score, normal agent	3.59	4.41	430
Attention score, attention disabled	0.31	0.94	116
Catch score, normal agent	1.66	2.69	1081
Catch score, attention disabled	-0.38	0.46	507

Table S1: Curve fits for Experiment 1. For an explanation of the four different conditions, see Figure 3 in the main text. The data were fitted to the equation y = a-b*exp(-x/c), where parameter a = the estimated final asymptotic level of performance; b = the change in performance between initial and asymptotic level; and c = the number of iterations before 66.66% of learning occurred.

Condition	Parameter a	Parameter b	Parameter c
Attention score, normal agent	3.58	4.44	423
Attention score, attention schema disabled	-0.34	0.77	34
Catch score, normal agent	1.08	2.12	722
Catch score, attention schema disabled	0.19	1.03	1403

Table S2: Curve fits for Experiment 2. For an explanation of the four different conditions, see Figure 4 in the main text. The data were fitted to the equation y = a-b*exp(-x/c), where parameter a = the estimated final asymptotic level of performance; b = the change in performance between initial and asymptotic level; and c = the number of iterations before 66.66% of learning occurred.

Condition	Parameter a	Parameter b	Parameter c
Attention score, attention normal	--	--	--
Attention score, attention disabled	--	--	--
Catch score, attention normal	-0.12	0.69	748
Catch score, attention disabled	-0.38	0.43	575

Table S3: Curve fits for Experiment 3. For an explanation of the four different conditions, see Figure 5 in the main text. The data were fitted to the equation y = a-b*exp(-x/c), where parameter a = the estimated final asymptotic level of performance; b = the change in performance between initial and asymptotic level; and c = the number of iterations before 66.66% of learning occurred. For the first two conditions, the asymptotic curve fit failed to converge on a stable answer, presumably because the curves were too flat to constrain the curve fit. For the second two conditions, the curve fit was successful.