

## Question 1: Monte-Carlo Control in Easy21

- a. To implement Monte-Carlo control in Easy21, an iterative on-policy  $\epsilon$ -greedy first-visit MC control algorithm is used. This iterative algorithm was chosen over a batch approach as it has lower time complexity and memory requirements, both of which are important as the number of episodes used will be very large and the process will run on a system with moderate computing power. The first-visit method is preferred over every-visit as it has been shown that for a large enough number of episodes it will result in lower average MSE in approximating the optimal value function<sup>1</sup>. Finally, this algorithm does not require exploring starts, making it applicable to scenarios where starting from any state is not feasible.

The State-Action Value function is updated after each new trace  $\tau$  has been generated. For all unique state-action pairs appearing in  $\tau$ :  $\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \alpha[R(s, a) - \hat{Q}(s, a)]$  where  $R(s, a)$  is the return following the first visit to the state-action pair  $(s, a)$  (note that since all non-zeros rewards occur at the end of the episode and  $\gamma = 1$ ,  $R(s, a) = r_{\text{episode}}$  for all  $(s, a)$  in  $\tau$ ). The parameter  $\alpha$  is chosen to be equal to the reciprocal of the cumulative number of first visits to the state-action pair ( $\alpha = 1/N(s, a)$ ), which satisfies the Robbins-Monroe conditions). Initializing  $\hat{Q}(s, a)$  as zero for all state-action pairs and setting  $\alpha = 1/N(s, a)$  makes the update rule an online average of the returns following the first visit to the state-action pair. The policy is improved by setting  $\hat{\pi}(s, a) = 1 - \epsilon/2$  if  $a$  maximizes  $\hat{Q}(s, a)$ ,  $\hat{\pi}(s, a) = 0.5$  if  $\hat{Q}(s, \text{'hit'})$  is equal to  $\hat{Q}(s, \text{'stick'})$  and  $\hat{\pi}(s, a) = \epsilon/2$  otherwise. Here,  $\epsilon = \frac{N_0}{N_0 + \sum_{a \in A} N(s, a)}$  which results in an  $\epsilon$ -greedy policy that is GLIE. The parameter  $N_0$  is used to balance exploration vs exploitation for the number of episodes to be run (given time and computational constraints). As  $N_0$  increases, the exploration increases but so does the time required for convergence. For smaller values of  $N_0$ , convergence is quicker but the results likely inaccurate. The value of  $N_0$  will be determined by inspecting learning curves for different  $N_0$  (see Q1b).

- b. The algorithm is run for  $10^6$  episodes, as it produces good results without being overly time consuming. Training stops every few episodes and  $N$  simulations of the game are run using the current policy. The mean and standard deviation of the rewards are then calculated. Since the standard deviation is approximately 1 and the mean in the order of 0.01,  $N$  must be in the order of  $10^4$  for a learning curve with reasonable SNR (the standard deviation of the mean decreases approximately with  $1/\sqrt{N}$ ). Hence, due to limitations in computational power, testing is performed every  $5 \cdot 10^4$  episodes. Empirically it was found that for  $N=16 \cdot 10^4$ , the plot becomes sufficiently clear. Nine learning curves are plotted, three separate instances for each of  $N_0 = 10, 100$  and  $1000$ .

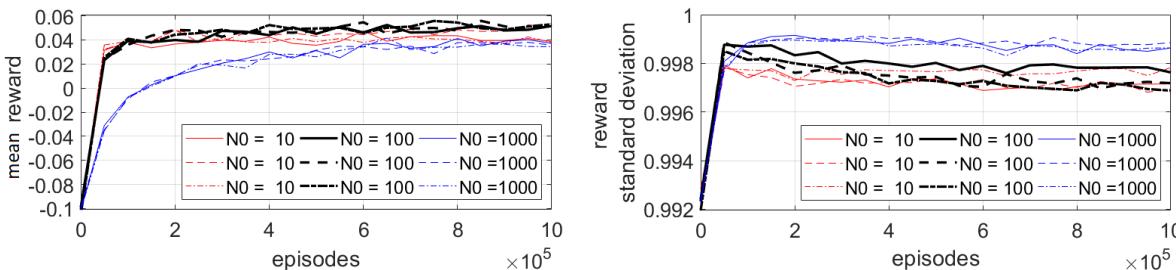


Figure 1: Learning Curves for  $N_0 = \{10, 100, 1000\}$ .  $N_0 = 100$  offers the best overall performance.  $N_0 = 10$  converges faster but 2 out of 3 times converges to a worse policy than  $N_0 = 100$ . When  $N_0 = 1000$ , convergence has not occurred yet after  $10^6$  episodes.

<sup>1</sup>Singh, S.P. & Sutton, R.S. Mach Learn (1996) 22: 123. <https://doi.org/10.1007/BF00114726>

- c. The estimate for the optimal value function after  $10^6$  episodes is shown in the figure below. It is notable that the value of states where the player sum is close to 21 have value of approximately 1. Additionally, as the value of the dealer card increases, the value of the states decreases. While the specific trend is covered by noise and hard to establish, it seems that there is a local maximum when player sum is equal to 11. This can be explained by noting that when the player sum is 11, the player cannot go bust in the next draw. This is further supported by the estimated optimal policy (also shown below).

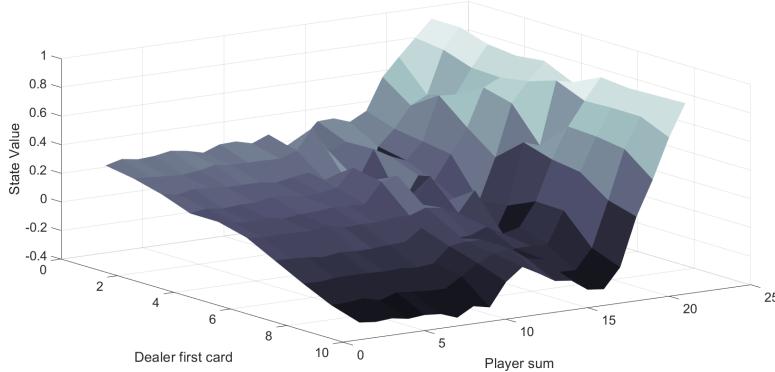
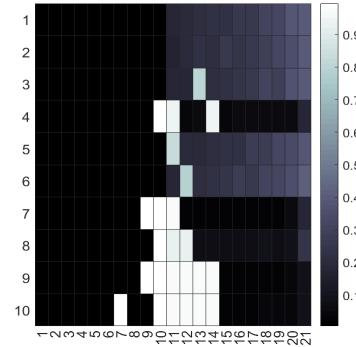
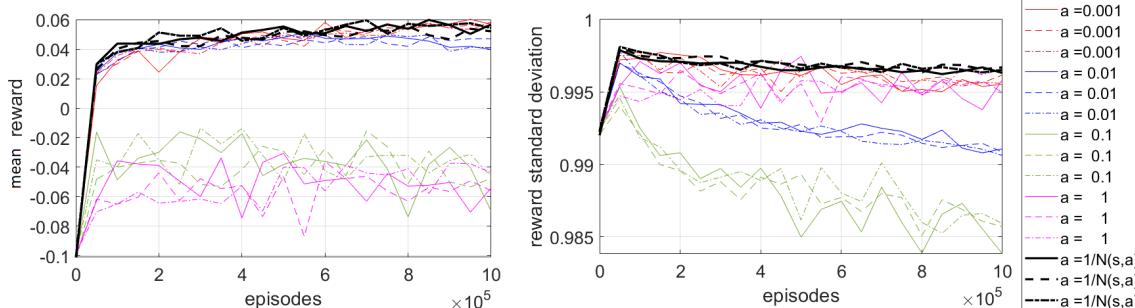
Figure 2: Optimal Value function estimation ( $N_0=100$ )

Figure 3: Optimal Policy for 'hit'

## Question 2: TD Learning in Easy21

- a. In this section, the SARSA on-policy learning TD control algorithm is used. This method offers lower variance than MC control. The State-Action Value function is updated online using the rule:  $\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \alpha[R(s, a) + \gamma\hat{Q}(s', a') - \hat{Q}(s, a)]$ , where  $(s', a')$  represent the next state-action pair. Again,  $\alpha = 1/N(s, a)$  which satisfies Robbins-Monroe. It is expected that this decreasing implementation of  $\alpha$  should perform better than constant step sizes.  $\epsilon = \frac{N_0}{N_0 + \sum_{a \in \mathcal{A}} N(s, a)}$ , which results in an  $\epsilon$ -greedy policy that is GLIE.  $N_0=100$  balances accuracy vs convergence speed (chosen empirically by plotting learning curves, Appendix A).
- b. Learning curves are plotted using the same testing parameters as in Question 1b. To investigate the effects of  $\alpha$ , plots were produced for different constant values of parameter  $\alpha=0.001, 0.01, 0.1, 1$  as well as  $\alpha = 1/N(s, a)$  (three separate instances were run for each parameter value resulting in 15 plots). The plots below suggest that for small values of  $\alpha$ , the process takes longer to converge and is relatively accurate ( $\alpha = 0.001$  has very similar performance to  $\alpha = 1/N(s, a)$ ), whereas higher values of  $\alpha$  significantly under-perform. This suggests that while for constant  $\alpha$  SARSA does not converge to the optimal state-value function (Robbins-Monroe is not satisfied), when  $\alpha$  is sufficiently small the approximation can be very good.

Figure 4: Learning Curves for  $\alpha=\{0.001, 0.01, 0.1, 1, 1/N(s,a)\}$

- c. The estimated optimal value function for SARSA with  $N_0=100$  after  $10^6$  episodes is plotted. The plot's characteristics are identical to those in MC control (Question 1c) apart from a slight increase in smoothness (can be attributed to the lower variance of SARSA vs MC).

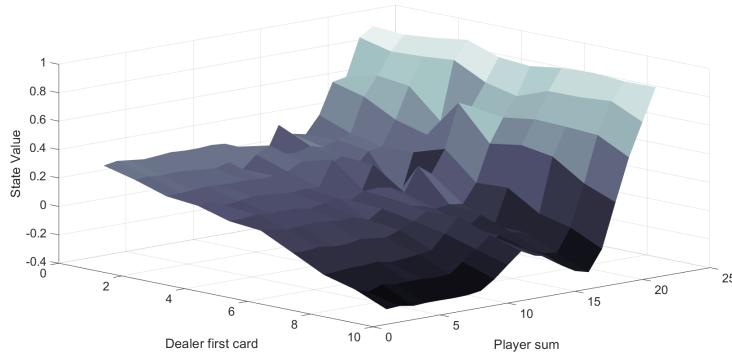
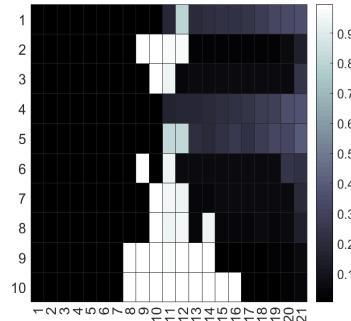
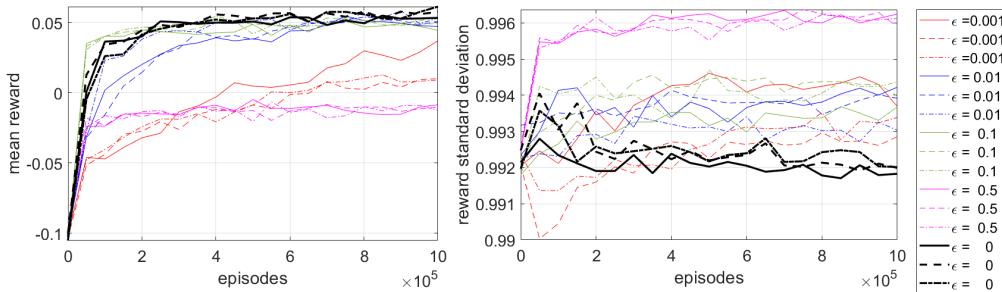
Figure 5: Optimal Value function estimation ( $N_0=100$ )

Figure 6: Optimal Policy for 'hit'

### Question 3: Q-Learning in Easy21

- a. In this section, Q-Learning, which is an Off-Policy TD Control algorithm, is implemented. Q-learning aims to improve two policies: the target policy whose value function is the target of the learning process, and the behaviour policy controlling the agent's actions. To achieve this the following update rule is used:  $\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \alpha[R(s, a) + \gamma \max_a \hat{Q}(s', a) - \hat{Q}(s, a)]$ . Taking  $\max_a \hat{Q}(s', a)$ , means that the target policy is greedy w.r.t.  $\hat{Q}(s, a)$  while, as in both previous methods, the behaviour policy is  $\epsilon$ -greedy w.r.t.  $\hat{Q}(s, a)$ .  $\alpha = 1/N(s, a)$  as in previous sections and satisfies Robbins-Monroe.  $\epsilon = \frac{N_0}{N_0 + \sum_{a \in \mathcal{A}} N(s, a)}$ , which results in an  $\epsilon$ -greedy policy that is GLIE.  $N_0=100$  balances accuracy vs convergence speed (chosen empirically by plotting learning curves, Appendix A).
- b. The learning curves are plotted using the same testing parameters as in Questions 1b, 2b. To investigate the effect of different strategies for handling  $\epsilon$ -greediness, plots were produced for different constant values of parameter  $\epsilon=0.001, 0.01, 0.1, 1$  as well as  $\epsilon = \frac{100}{100 + \sum_{a \in \mathcal{A}} N(s, a)}$  (three separate instances were run for each parameter value resulting in 15 plots). The plots in the figure below suggest that when  $\epsilon$  is equal to a very small constant value, not enough exploration will be performed resulting in poor learning ( $\epsilon = 0.001$ ). Conversely, when  $\epsilon$  is constant and high, while exploration is achieved, the optimal policy estimate is far enough from convergence that the results are poor ( $\epsilon = 0.5$ ). Finally, for  $\epsilon$  within a range of appropriate values, the learning curves are close (but slightly worse) than those for the time-varying  $\epsilon$ . It is also important to note that any policy with constant  $\epsilon$  is not GLIE.

Figure 7: Learning Curves for  $\epsilon = \{0.001, 0.01, 0.1, 0.5, \frac{100}{100 + \sum_{a \in \mathcal{A}} N(s, a)}\}$

- c. The estimated optimal value function for Q-learning with  $\epsilon = \frac{N_0}{N_0 + \sum_{a \in \mathcal{A}} N(s, a)}$  and  $N_0=100$  after  $10^6$  episodes is plotted. While the plot has the same shape as those in MC control and SARSA (Question 1c, 2c), it is by far the smoothest. This implies that the variance is significantly lower than in both other methods.

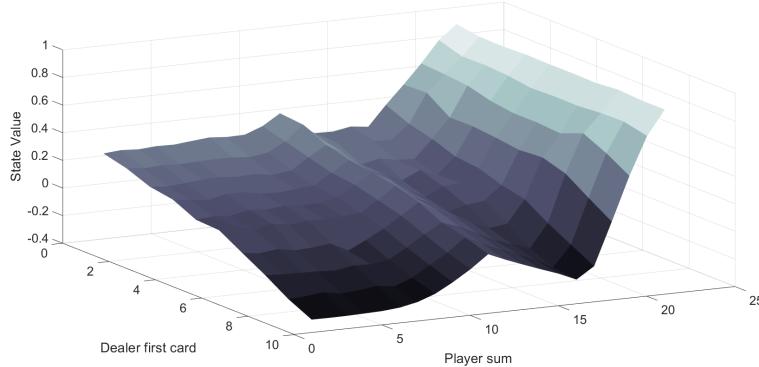


Figure 8: Optimal Value function

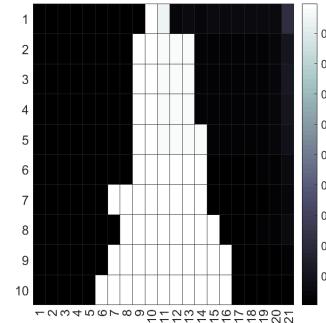


Figure 9: Optimal Policy for 'hit'

## Question 4: Compare the algorithms

To compare the algorithms, learning curves are computed for all three methods. The number of testing simulations has been increased to  $25 \cdot 10^4$  and testing occurs every  $10^4$  episodes. All algorithms use time varying  $\alpha = 1/N(s, a)$  and  $\epsilon = \frac{N_0}{N_0 + \sum_{a \in \mathcal{A}} N(s, a)}$  with  $N_0=100$ . Note that each algorithm has been run three times, producing nine plots.

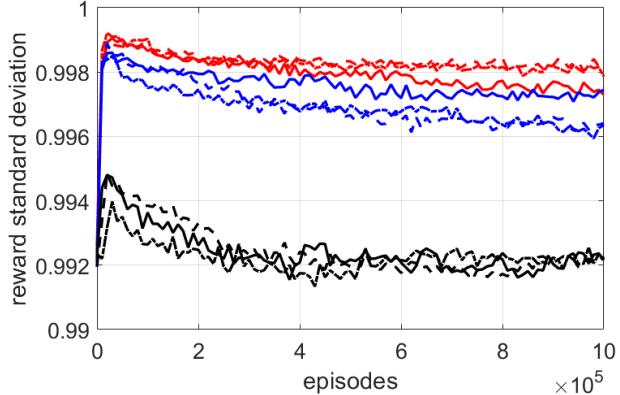
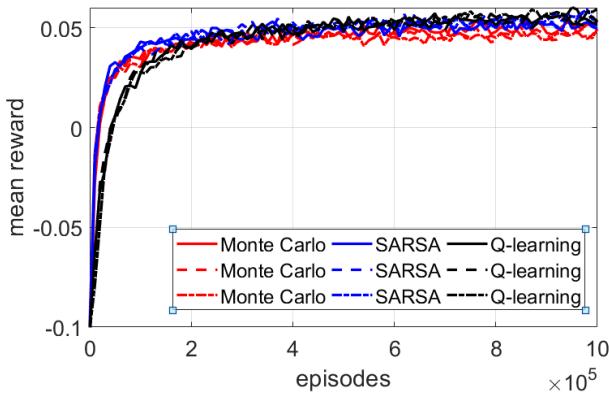


Figure 10: Learning Curves the three methods used

From the above figures, it can be seen that after  $10^6$  episodes, the mean reward for Q-learning is larger than that for SARSA which is larger than that for Monte Carlo. Q-learning producing better results than SARSA can be explained by noting that Q-learning converges closer to the optimal policy, while SARSA converges to a safer policy. Since the penalties aren't large, the reward from following a riskier but more accurate policy will be higher than the reward from following a safe policy. Monte Carlo has high variance and because it is implemented iteratively, it also accumulates error. Thus, it falls below SARSA. It is expected that batch MC would perform better but would be very computationally intensive. The standard deviation of rewards is linked to the variance of the estimator of the optimal value function, thus MC which has the highest estimator variance will also have the highest standard deviation of rewards. Finally, it is worth noting that SARSA and Q-learning exploit the Markov property whereas MC does not, which could explain why the former perform better.

## Appendix A

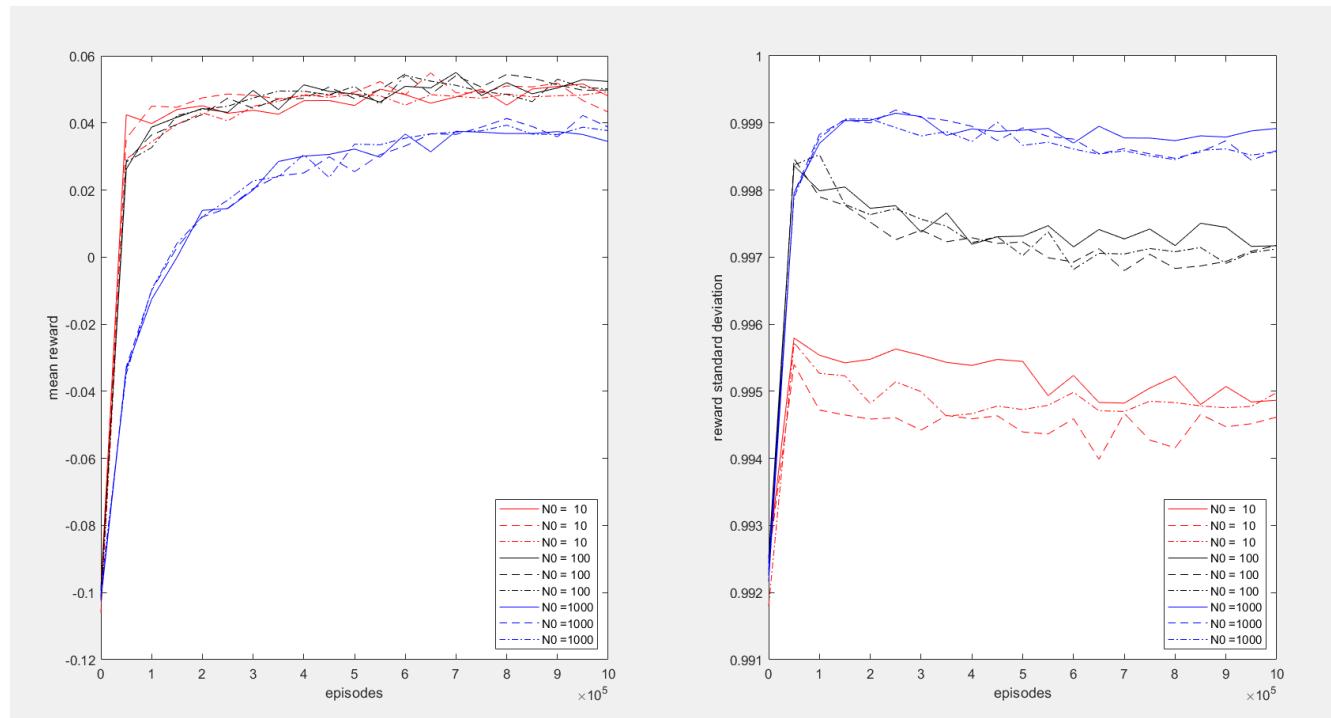


Figure 11: Learning Curves for SARSA with varying  $N_0$ .  $N_0=100$  offers best performance

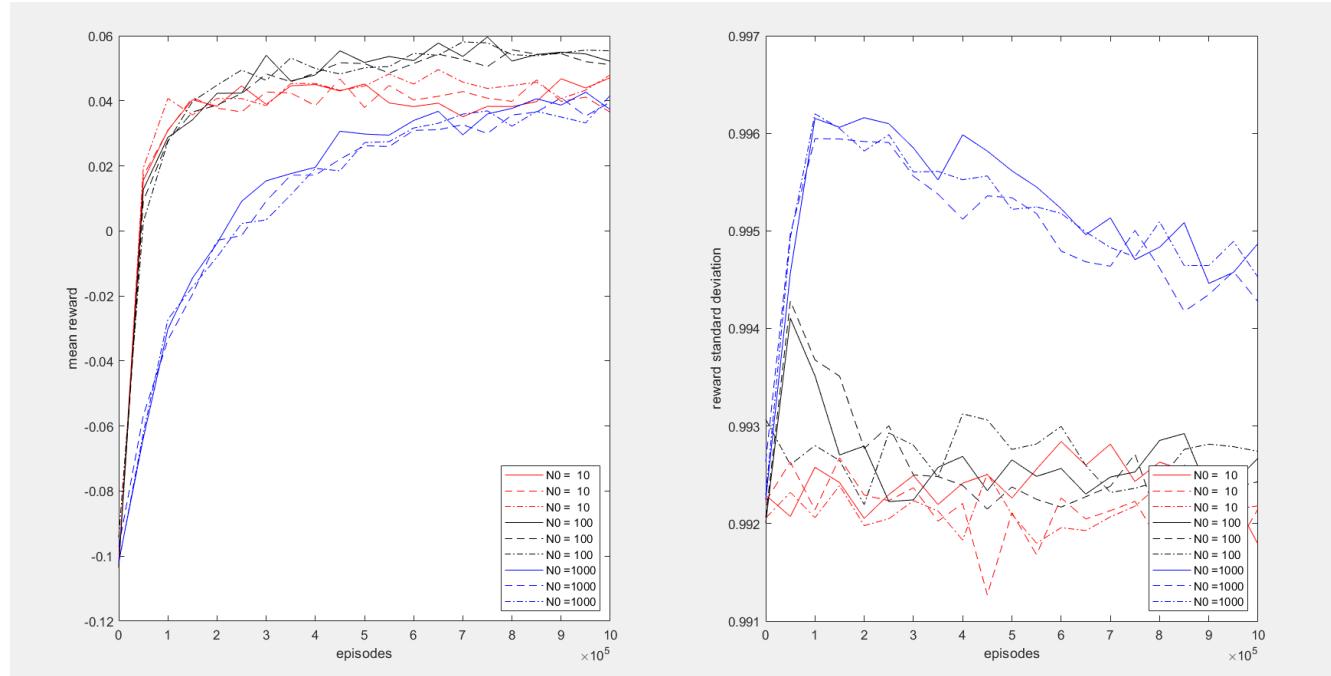


Figure 12: Learning Curves for Q-learning with varying  $N_0$ .  $N_0=100$  offers best performance

## Appendix B: MATLAB Code

Note that while the algorithms have been coded quite efficiently, the plotting section is not implemented elegantly and some cosmetic editing of the plots (line-widths, colours etc) has been done through the MATLAB GUI and not through coding.

```

1 %STEP FUNCTION
2 function [s_n , r , fin]=step(s , a)
3 ds=s(1);
4 ps=s(2);
5 fin=0;
6 if a==1
7     val=randi(10);
8     mult=2*(rand()>(1/3))-1; % 2/3 chance 1, 1/3 chance -1
9     ps=ps+mult*val;
10    if ps<1 || ps>21
11        r=-1;
12        fin=1;
13    else
14        r=0;
15    end
16 else
17     fin=1;
18     while ds<17 && ds>=1
19         val=randi(10);
20         mult=2*(rand()>(1/3))-1;
21         ds=ds+mult*val;
22     end
23     if ds <1||ds >21||ds<ps
24         r=1;
25     elseif ds==ps
26         r=0;
27     else
28         r=-1;
29     end
30 end
31 s_n=[ds ps];
32 end
33
34 %MONTE CARLO fUNCTION
35 function [R, R_test ,Q, pol , count_state_action]=mc(n_test ,n_epi ,d_test ,N0 ,
36 ev ,dshow)
37 ind_show=1;
38 ind_test_vec=[1 d_test :d_test :n_epi ];
39 ind_test=1;
40 %
41 t0=tic;%timer
42 pol=0.5*ones(10 ,21 ,2); % pi('hit' ,s)
```

```

42 count_state_action=zeros(10,21,2);
43 Q=zeros(10,21,2);
44 R=zeros(n_test,length(ind_test_vec));
45 R_test=zeros(1,n_epi);
46 for epi=1:n_epi
47     fin=0;
48     s=randi(10,[1 2]);
49     trace=[];
50     while fin==0
51         temp=rand();
52         a=2-(temp<pol(s(1),s(2)));
53         [sn,r,fin]=step(s,a);
54         trace=[trace;[s a]];
55         s=sn;
56     end
57     R_test(epi)=r;
58     checked=zeros(10,21,2);
59     for ind=1:size(trace,1)
60         if ev || ~checked(trace(ind,1),trace(ind,2),trace(ind,3))
61             count_state_action(trace(ind,1),trace(ind,2),trace(ind,3))
62             =count_state_action(trace(ind,1),trace(ind,2),trace(ind,3))+1;
63             Q_hat=Q(trace(ind,1),trace(ind,2),trace(ind,3));
64             Q_hat=Q_hat+(r-Q_hat)/count_state_action(trace(ind,1),
65                 trace(ind,2),trace(ind,3));
66             Q(trace(ind,1),trace(ind,2),trace(ind,3))=Q_hat;
67             checked(trace(ind,1),trace(ind,2),trace(ind,3))=1;
68         end
69     end
70     count_state=count_state_action(:,:,1)+count_state_action(:,:,2);
71     checked=zeros(10,21);
72     for ind=1:size(trace,1)
73         if ev || ~checked(trace(ind,1),trace(ind,2))
74             eps=N0/(N0+count_state(trace(ind,1),trace(ind,2)));
75             if Q(trace(ind,1),trace(ind,2),1)>Q(trace(ind,1),trace(ind,
76                 2),2)
77                 pol(trace(ind,1),trace(ind,2),1)=1-eps/2;
78                 pol(trace(ind,1),trace(ind,2),2)=eps/2;
79             end
80             if Q(trace(ind,1),trace(ind,2),1)<Q(trace(ind,1),trace(ind,
81                 2),2)
82                 pol(trace(ind,1),trace(ind,2),2)=1-eps/2;
83                 pol(trace(ind,1),trace(ind,2),1)=eps/2;
84             end
85             if Q(trace(ind,1),trace(ind,2),1)==Q(trace(ind,1),trace(
86                 ind,2),2)
87                 pol(trace(ind,1),trace(ind,2),1)=0.5;

```

```

83         pol(trace(ind,1),trace(ind,2),2)=0.5;
84     end
85     checked(trace(ind,1),trace(ind,2))=1;
86   end
87 end
88 if epi==ind_test_vec(ind_test)
89   for train_epi=1:n_test
90     fin=0;
91     s=randi(10,[1 2]);
92     trace=[];
93     while fin==0
94       temp=rand();
95       a=2-(temp<pol(s(1),s(2))); % is 1 for 'hit' 2 for 'stick'
96       [sn,r,fin]=step(s,a);
97       s=sn;
98     end
99     R(train_epi,ind_test)=r;
100   end
101   ind_test=ind_test+1;
102 end
103 if (epi/n_epi)/dshow>ind_show
104   ind_show=ind_show+1;
105   [round(100*(epi/n_epi)) toc(t0)]
106   t0=tic;
107 end
108 end
109 end
110 %SARSA FUNCTION
111 function [R,R_train,Q,pol,count_state_action]=fsarsa(n_train,n_epi,
112   d_test,N0,dshow,a_choice)
113 % counter stuff for visualizing progress
114 ind_show=1;
115 ind_test_vec=[1:d_test:d_test:n_epi];
116 ind_test=1;
117 %
118 t0=tic;%timer
119 pol=0.5*ones(10,21,2); % pi('hit',s)
120 count_state_action=zeros(10,21,2);
121 Q=zeros(10,21,2); %dim 1 is 'hit' dim 2 'stick'
122 R_train=zeros(1,n_epi);
123 R=zeros(n_train,length(ind_test_vec));
124 for epi=1:n_epi
125   fin=0;
126   s=randi(10,[1 2]);
127   temp=rand();
   A=2-(temp<pol(s(1),s(2))); % is 1 for 'hit' 2 for 'stick',

```

```

128 while fin==0
129     [sn ,r ,fin]=step (s ,A) ;
130     temp=rand () ;
131     count_state_action (s (1) ,s (2) ,A)=count_state_action (s (1) ,s (2) ,A)
132         )+1;
133     if a_choice==0
134         a=1/count_state_action (s (1) ,s (2) ,A) ;
135     else
136         a=a_choice ;
137     end
138     if fin==0
139         An=2-(temp<pol (sn (1) ,sn (2) )); % is 1 for 'hit' 2 for '
140             stick ,
141         Q(s (1) ,s (2) ,A)=Q(s (1) ,s (2) ,A)+a*( r+Q(sn (1) ,sn (2) ,An)-Q(s
142             (1) ,s (2) ,A)) ;
143     else
144         Q(s (1) ,s (2) ,A)=Q(s (1) ,s (2) ,A)+a*( r-Q(s (1) ,s (2) ,A)) ;
145     end
146     eps=N0/(N0+count_state_action (s (1) ,s (2) ,1)+count_state_action (
147         s (1) ,s (2) ,2));
148     if Q(s (1) ,s (2) ,1)>Q(s (1) ,s (2) ,2)
149         pol(s (1) ,s (2) ,1)=1-eps /2;
150         pol(s (1) ,s (2) ,2)=eps /2;
151     end
152     if Q(s (1) ,s (2) ,1)<Q(s (1) ,s (2) ,2)
153         pol(s (1) ,s (2) ,2)=1-eps /2;
154         pol(s (1) ,s (2) ,1)=eps /2;
155     end
156     if fin==0
157         A=An;
158         s=sn ;
159     end
160 end
161 R_train(epi)=r ;
162 if epi==ind_test_vec(ind_test)
163     for train_epi=1:n_train
164         fin=0;
165         s=randi(10 ,[1 2]) ;
166         while fin==0
167             temp=rand () ;
168             a=2-(temp<pol (s (1) ,s (2) )); % is 1 for 'hit' 2 for '
169                 stick ,
[sn ,r ,fin]=step (s ,a) ;

```

```

170             s==sn ;
171         end
172     R( train_epi , ind_test )=r ;
173     end
174     ind_test=ind_test+1;
175 end
176 %timer/counter for progress monitoring
177 if ( epi/n_epi )/dshow>ind_show
178     ind_show=ind_show+1;
179     [ round(100*( epi/n_epi )) toc(t0) ]
180     t0=tic ;
181 end
182 end
183 end
184 %Q LEARN FUNCTION
185 function [R,R_train,Q,pol , count_state_action]=Q_learn( n_train , n_epi ,
186     d_test ,N0,dshow , eps_choice )
187 % counter stuff for visualizing progress
188 ind_show=1;
189 ind_test_vec=[1 d_test:d_test:n_epi ];
190 ind_test=1;
191 %
192 t0=tic;%timer
193 pol=0.5*ones(10,21,2); % pi('hit',s)
194 count_state_action=zeros(10,21,2);
195 Q=zeros(10,21,2); %dim 1 is 'hit' dim 2 'stick'
196 R_train=zeros(1,n_epi);
197 R=zeros(n_train , length(ind_test_vec));
198 for epi=1:n_epi
199     fin=0;
200     s=randi(10,[1 2]);
201     temp=rand();
202     A=2-(temp<pol(s(1),s(2))); % is 1 for 'hit' 2 for 'stick'
203     while fin==0
204         [sn,r,fin]=step(s,A);
205         temp=rand();
206         count_state_action(s(1),s(2),A)=count_state_action(s(1),s(2),A)
207             +1;
208         a=1/count_state_action(s(1),s(2),A);
209         if fin==0
210             An=2-(temp<pol(sn(1),sn(2))); % is 1 for 'hit' 2 for 'stick'
211             Q(s(1),s(2),A)=Q(s(1),s(2),A)+a*( r+max(Q(sn(1),sn(2),1),Q(
212                 sn(1),sn(2),2))-Q(s(1),s(2),A));
213         else
214             Q(s(1),s(2),A)=Q(s(1),s(2),A)+a*(r-Q(s(1),s(2),A));
215         end

```

```

213     if eps_choice==0
214         eps=N0/(N0+count_state_action(s(1),s(2),1) +
215             count_state_action(s(1),s(2),2));
216     else
217         eps=eps_choice;
218     end
219     if Q(s(1),s(2),1)>Q(s(1),s(2),2)
220         pol(s(1),s(2),1)=1-eps/2;
221         pol(s(1),s(2),2)=eps/2;
222     end
223     if Q(s(1),s(2),1)<Q(s(1),s(2),2)
224         pol(s(1),s(2),2)=1-eps/2;
225         pol(s(1),s(2),1)=eps/2;
226     end
227     if Q(s(1),s(2),1)==Q(s(1),s(2),2)
228         pol(s(1),s(2),1)=0.5;
229         pol(s(1),s(2),2)=0.5;
230     end
231     if fin==0
232         A=An;
233         s=sn;
234     end
235 end
236 R_train(epi)=r;
237 if epi==ind_test_vec(ind_test)
238     for train_epi=1:n_train
239         fin=0;
240         s=randi(10,[1 2]);
241         trace=[];
242         while fin==0
243             temp=rand();
244             a=2-(temp<pol(s(1),s(2))); % is 1 for 'hit' 2 for 'stick'
245             [sn,r,fin]=step(s,a);
246             s=sn;
247         end
248         R(train_epi,ind_test)=r;
249     end
250     ind_test=ind_test+1;
251 end
252 %timer/counter for progress monitoring
253 if (epi/n_epi)/dshow>ind_show
254     ind_show=ind_show+1;
255     [round(100*(epi/n_epi)) toc(t0)]
256     t0=tic;
257 end

```

```

258 end
259 %PLOTTING
260
261
262 clc ; clear ; close all
263 %%%
264 n_test = 16*10^4;
265 n_epi = 10^6;
266 dshow = 0.1;
267 d_test = 5*10^4;
268 figure(1)
269 N0_vec = [10 100 1000];
270 c_vec = [ 'r' ; 'k' ; 'b' ];
271 for ind=1:length(N0_vec)
272     N0=N0_vec(ind);
273     [R, R_train, Q, pol, count_state_action] = mc(n_test, n_epi, d_test, N0, 0,
274         dshow);
275     subplot(1,2,1)
276     plot([1 d_test : d_test : n_epi], mean(R), 'color', c_vec(ind))
277     hold on
278     subplot(1,2,2)
279     plot([1 d_test : d_test : n_epi], std(R), 'color', c_vec(ind))
280     hold on
281     [R, R_train, Q, pol, count_state_action] = mc(n_test, n_epi, d_test, N0, 0,
282         dshow);
283     subplot(1,2,1)
284     plot([1 d_test : d_test : n_epi], mean(R), 'color', c_vec(ind), 'Linestyle',
285         '---')
286     hold on
287     subplot(1,2,2)
288     plot([1 d_test : d_test : n_epi], std(R), 'color', c_vec(ind), 'Linestyle',
289         '--')
290     hold on
291     subplot(1,2,1)
292     plot([1 d_test : d_test : n_epi], mean(R), 'color', c_vec(ind), 'Linestyle',
293         '-.')
294     hold on
295 end
296 subplot(1,2,1)
297 legend(strcat('N0 = ', num2str([10 10 10 100 100 100 1000 1000 1000])),
298         'Location', 'southeast')

```

```

297 xlabel('episodes')
298 ylabel('mean reward')
299 subplot(1,2,2)
300 legend(strcat('N0 = ', num2str([10 10 10 100 100 100 1000 1000 1000])), ...
301 , 'Location', 'southeast')
301 xlabel('episodes')
302 ylabel('reward standard deviation')
303 %%  

304 n_test=1;
305 n_epi=10^6;
306 N0=100;
307 [R,R_train,Q,pol,count_state_action]=mc(n_test,n_epi,d_test,N0,0,dshow);
308 [X,Y] = meshgrid(1:10,1:21);
309 v=max(Q,[],3);
310 figure(2)
311 colormap('bone')
312 surf(X,Y,v,'EdgeColor','none');
313 xlabel('Dealer first card')
314 ylabel('Player sum')
315 zlabel('State Value')
316 figure(3)
317 h=heatmap(pol(:,:,1));
318 h.Colormap=colormap('bone');
319 %%  

320 n_test=16*10^4;
321 n_epi=10^6;
322 dshow=0.1;
323 d_test=5*10^4;
324 figure(4)
325 N0_vec=[10 100 1000];
326 c_vec=['r';'k';'b'];
327 for ind=1:length(N0_vec)
328 N0=N0_vec(ind);
329 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0 ...
330 ,dshow,0);
330 subplot(1,2,1)
331 plot([1:d_test:d_test:n_epi],mean(R),'color',c_vec(ind))
332 hold on
333 subplot(1,2,2)
334 plot([1:d_test:d_test:n_epi],std(R),'color',c_vec(ind))
335 hold on
336 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0 ...
337 ,dshow,0);
337 subplot(1,2,1)
338 plot([1:d_test:d_test:n_epi],mean(R),'color',c_vec(ind),'LineStyle ...
339 ','--')

```

```

339 hold on
340 subplot(1,2,2)
341 plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind),'Linestyle'
342 ,'-')
343 hold on
344 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0
345 ,dshow,0);
346 subplot(1,2,1)
347 plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind),'Linestyle'
348 ,'.')
349 hold on
350 subplot(1,2,2)
351 plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind),'Linestyle'
352 ,'.')
353 hold on
354 end
355 subplot(1,2,1)
356 legend(strcat('N0 = ',num2str([10 10 10 100 100 100 1000 1000 1000]')),
357 'Location','southeast')
358 xlabel('episodes')
359 ylabel('mean reward')
360 subplot(1,2,2)
361 legend(strcat('N0 = ',num2str([10 10 10 100 100 100 1000 1000 1000]')),
362 'Location','southeast')
363 xlabel('episodes')
364 ylabel('reward standard deviation')
365 %%%
366 n_test=16*10^4;
367 n_epi=10^6;
368 dshow=0.1;
369 d_test=5*10^4;
370 figure(5)
371 a_vec=[0.001 0.01 0.1 1 0];
372 c_vec=['r';'b';'g';'m';'k'];
373 N0=100;
374 for ind=1:length(a_vec)
375 a=a_vec(ind);
376 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0
377 ,dshow,a);
378 subplot(1,2,1)
379 plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind))
380 hold on
381 subplot(1,2,2)
382 plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind))
383 hold on
384 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0
385 ,dshow,a);

```

```

378 subplot(1,2,1)
379 plot([1 d_test:d_test:n_epi],mean(R), 'color',c_vec(ind), 'Linestyle
      , '--')
380 hold on
381 subplot(1,2,2)
382 plot([1 d_test:d_test:n_epi],std(R), 'color',c_vec(ind), 'Linestyle',
      , '-.')
383 hold on
384 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0
      ,dshow,a);
385 subplot(1,2,1)
386 plot([1 d_test:d_test:n_epi],mean(R), 'color',c_vec(ind), 'Linestyle
      , '-.')
387 hold on
388 subplot(1,2,2)
389 plot([1 d_test:d_test:n_epi],std(R), 'color',c_vec(ind), 'Linestyle',
      , '-.')
390 hold on
391 end
392 subplot(1,2,1)
393 legend(strcat('a = ',num2str([0.001 0.001 0.001 0.01 0.01 0.01 0.1 0.1
      0.1 1 1 1 0 0 0]')), 'Location', 'southeast')
394 xlabel('episodes')
395 ylabel('mean reward')
396 subplot(1,2,2)
397 legend(strcat('a = ',num2str([0.001 0.001 0.001 0.01 0.01 0.01 0.01 0.1 0.1
      0.1 1 1 1 0 0 0]')), 'Location', 'southeast')
398 xlabel('episodes')
399 ylabel('reward standard deviation')
400 %%
401 n_test=1;
402 n_epi=10^6;
403 N0=100;
404 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0,
      dshow,0);
405 [X,Y] = meshgrid(1:10,1:21);
406 v=max(Q,[],3);
407 figure(6)
408 colormap('bone')
409 surf(X,Y,v, 'EdgeColor', 'none');
410 xlabel('Dealer first card')
411 ylabel('Player sum')
412 zlabel('State Value')
413 figure(7)
414 h=heatmap(pol(:,:,1))
415 h.Colormap=colormap('bone');
416 %%

```

```

417 n_test=16*10^4;
418 n_epi=10^6;
419 dshow=0.1;
420 d_test=5*10^4;
421 figure(8)
422 N0_vec=[10 100 1000];
423 c_vec=['r'; 'k'; 'b'];
424 for ind=1:length(N0_vec)
425     N0=N0_vec(ind);
426     [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,
427         N0,dshow,0);
428     subplot(1,2,1)
429     plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind))
430     hold on
431     subplot(1,2,2)
432     plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind))
433     hold on
434     [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,
435         N0,dshow,0);
436     subplot(1,2,1)
437     plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind),'Linestyle',
438         ,'-')
439     hold on
440     subplot(1,2,2)
441     plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind),'Linestyle',
442         ,'-')
443     hold on
444     subplot(1,2,2)
445     plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind),'Linestyle',
446         ,'-')
447     hold on
448 end
449 subplot(1,2,1)
450 legend(strcat('N0 = ',num2str([10 10 10 100 100 100 1000 1000 1000])),
451     , 'Location', 'southeast')
452 xlabel('episodes')
453 ylabel('mean reward')
454 subplot(1,2,2)
455 legend(strcat('N0 = ',num2str([10 10 10 100 100 100 1000 1000 1000])),
456     , 'Location', 'southeast')
457 xlabel('episodes')

```

```

455 ylabel('reward standard deviation')
456 
457 n_test=16*10^4;
458 n_epi=10^6;
459 dshow=0.1;
460 d_test=5*10^4;
461 figure(9); close; figure(9)
462 e_vec=[0.001 0.01 0.1 0.5 0];
463 c_vec=['r'; 'b'; 'g'; 'm'; 'k'];
464 N0=100;
465 for ind=1:length(e_vec)
466     e=e_vec(ind);
467     [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,
468         N0,dshow,e);
469     subplot(1,2,1)
470     plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind))
471     hold on
472     subplot(1,2,2)
473     plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind))
474     hold on
475     [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,
476         N0,dshow,e);
477     subplot(1,2,1)
478     plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind),'Linestyle',
479           ',-')
480     hold on
481     subplot(1,2,2)
482     plot([1 d_test:d_test:n_epi],std(R),'color',c_vec(ind),'Linestyle',
483           ',-')
484     hold on
485     subplot(1,2,1)
486     plot([1 d_test:d_test:n_epi],mean(R),'color',c_vec(ind),'Linestyle',
487           ',-')
488 end
489 subplot(1,2,1)
490 legend(strcat('\epsilon = ',num2str([0.001 0.001 0.001 0.01 0.01 0.01
491     0.1 0.1 0.1 0.5 0.5 0.5 0 0 0])), 'Location', 'southeast')
492 xlabel('episodes')
493 ylabel('mean reward')
494 subplot(1,2,2)

```

```

494 legend(strcat(' \epsilon = ', num2str([0.001 0.001 0.001 0.01 0.01 0.01
        0.1 0.1 0.1 0.5 0.5 0.5 0 0 0])), 'Location', 'southeast')
495 xlabel('episodes')
496 ylabel('reward standard deviation')
497 %%  

498 n_test=1;
499 n_epi=10^6;
500 N0=100;
501 [R, R_train, Q, pol, count_state_action]=Q_learn(n_test, n_epi, d_test, N0,
        dshow, 0);
502 [X, Y] = meshgrid(1:10, 1:21);
503 v=max(Q, [], 3);
504 figure(10)
505 colormap('bone')
506 surf(X, Y, v, 'EdgeColor', 'none');
507 xlabel('Dealer first card')
508 ylabel('Player sum')
509 zlabel('State Value')
510 figure(11)
511 h=heatmap(pol(:,:,1));
512 h.Colormap=colormap('bone');
513 %%  

514 n_test=25*10^4;
515 n_epi=10^6;
516 dshow=0.95;
517 d_test=10^4;
518 figure(12); close; figure(12)
519 N0=100;
520 %MC-RED  

521 [R, R_train, Q, pol, count_state_action]=mc(n_test, n_epi, d_test, N0, 0, dshow
        );
522 subplot(1, 2, 1)
523 plot([1 d_test : d_test : n_epi], mean(R), 'color', 'r', 'Linestyle', '-')
524 hold on
525 subplot(1, 2, 2)
526 plot([1 d_test : d_test : n_epi], std(R), 'color', 'r', 'Linestyle', '-')
527 hold on
528 [R, R_train, Q, pol, count_state_action]=mc(n_test, n_epi, d_test, N0, 0, dshow
        );
529 subplot(1, 2, 1)
530 plot([1 d_test : d_test : n_epi], mean(R), 'color', 'r', 'Linestyle', '--')
531 hold on
532 subplot(1, 2, 2)
533 plot([1 d_test : d_test : n_epi], std(R), 'color', 'r', 'Linestyle', '--')
534 hold on
535 [R, R_train, Q, pol, count_state_action]=mc(n_test, n_epi, d_test, N0, 0, dshow
        );

```

```

536 subplot(1,2,1)
537 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'r', 'Linestyle', '-')
538 hold on
539 subplot(1,2,2)
540 plot([1 d_test:d_test:n_epi],std(R), 'color', 'r', 'Linestyle', '-')
541 hold on
542 %FSARSA-BLUE
543 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0,
      dshow,0);
544 subplot(1,2,1)
545 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'b', 'Linestyle', '-')
546 hold on
547 subplot(1,2,2)
548 plot([1 d_test:d_test:n_epi],std(R), 'color', 'b', 'Linestyle', '-')
549 hold on
550 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0,
      dshow,0);
551 subplot(1,2,1)
552 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'b', 'Linestyle', '--')
553 hold on
554 subplot(1,2,2)
555 plot([1 d_test:d_test:n_epi],std(R), 'color', 'b', 'Linestyle', '--')
556 hold on
557 [R,R_train,Q,pol,count_state_action]=fsarsa(n_test,n_epi,d_test,N0,
      dshow,0);
558 subplot(1,2,1)
559 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'b', 'Linestyle', '-')
560 hold on
561 subplot(1,2,2)
562 plot([1 d_test:d_test:n_epi],std(R), 'color', 'b', 'Linestyle', '-')
563 hold on
564 %Q-BLACK
565 [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,N0,
      dshow,0);
566 subplot(1,2,1)
567 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'k', 'Linestyle', '-')
568 hold on
569 subplot(1,2,2)
570 plot([1 d_test:d_test:n_epi],std(R), 'color', 'k', 'Linestyle', '-')
571 hold on
572 [R,R_train,Q,pol,count_state_action]=Q_learn(n_test,n_epi,d_test,N0,
      dshow,0);
573 subplot(1,2,1)
574 plot([1 d_test:d_test:n_epi],mean(R), 'color', 'k', 'Linestyle', '--')
575 hold on
576 subplot(1,2,2)
577 plot([1 d_test:d_test:n_epi],std(R), 'color', 'k', 'Linestyle', '--')

```

```
578 hold on
579 [R, R_train, Q, pol, count_state_action]=Q_learn(n_test, n_epi, d_test, N0,
      dshow, 0);
580 subplot(1,2,1)
581 plot([1:d_test:d_test:n_epi], mean(R), 'color', 'k', 'Linestyle', '-.')
582 hold on
583 subplot(1,2,2)
584 plot([1:d_test:d_test:n_epi], std(R), 'color', 'k', 'Linestyle', '-.')
585 hold on
586 subplot(1,2,1)
587 xlabel('episodes')
588 ylabel('mean reward')
589 subplot(1,2,2)
590 legend([{ 'Monte Carlo'},{ 'Monte Carlo'},{ 'Monte Carlo'},{ 'SARSA'},{ 'SARSA'},{ 'SARSA'},{ 'Q-learning'},{ 'Q-learning'},{ 'Q-learning' }])
591 xlabel('episodes')
592 ylabel('reward standard deviation')
```