Comments on the optimal policy

- The die with the most sides is generally favored as it has the highest expected outcome
- The dice with less sides can be used to increase the probability of hitting a ladder square or missing a snake square
  - However, in some cases, missing ladders and hitting snakes might be beneficial
- When the set of possible dice is expanded, the player has more possible choices and the expected amount of turns needed should decrease or stay the same

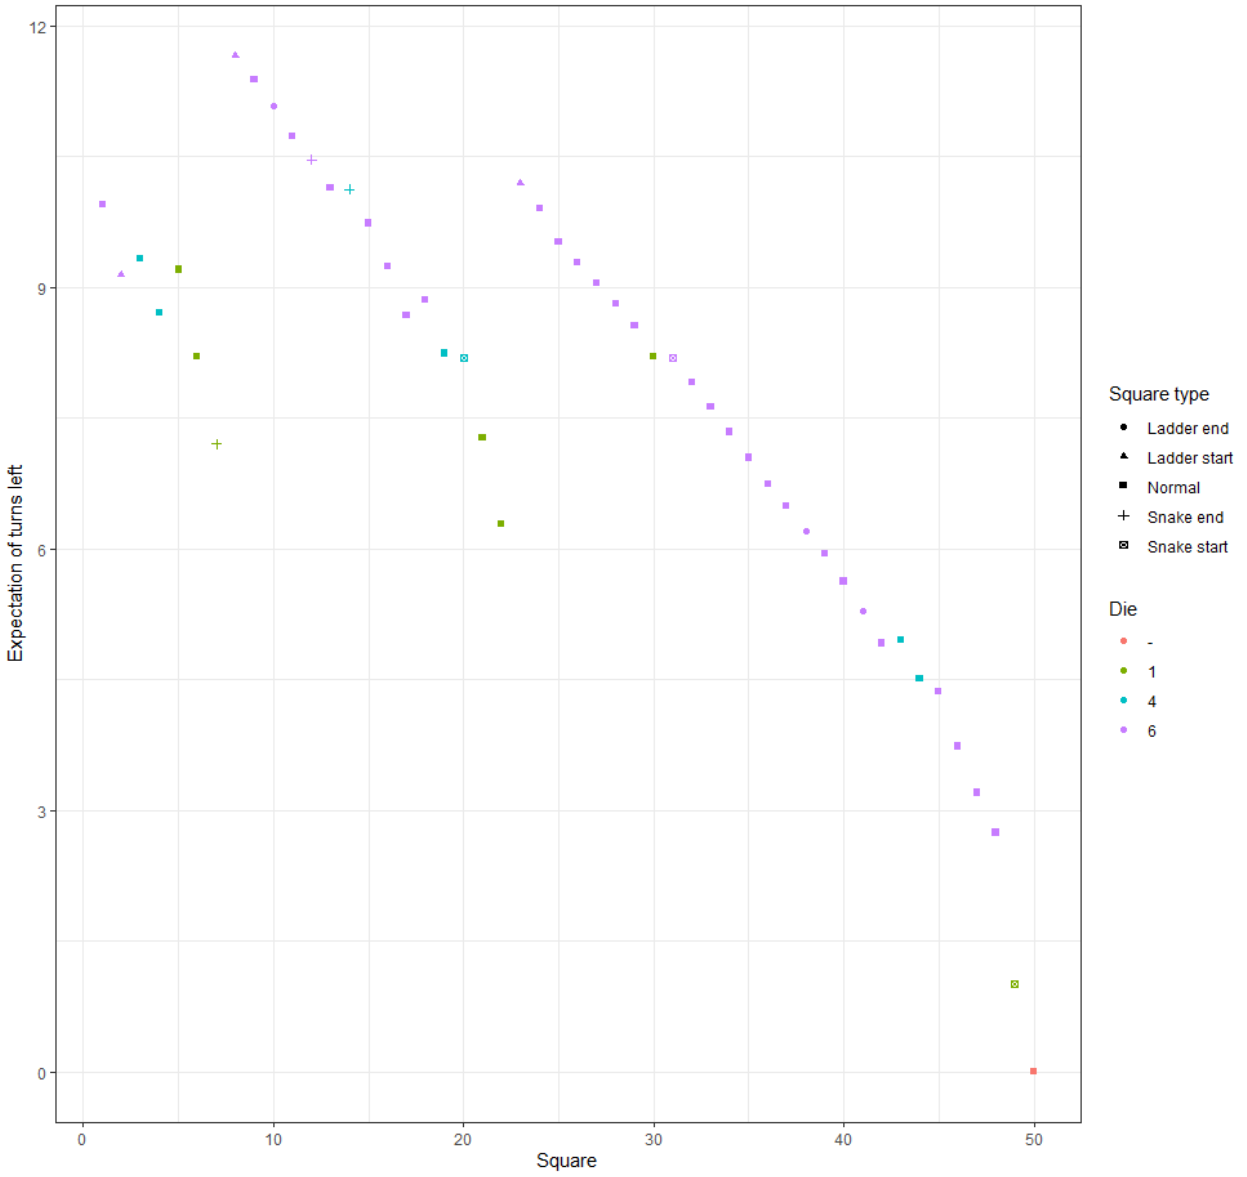

Figure 1: The solution of the problem with $U = \{1, 4, 6\}$ is visualized in the graph.

Convergence of Value Iteration

- The default number of iterations is quite conservative. To solve the default problem, approximately 30 iterations are needed.
- Small updates are still done after many iterations
- If the size of the state space is increased, there are more states to loop over and updating the costs is more complicated
- If the size of the action space is increased, there are more actions to be considered

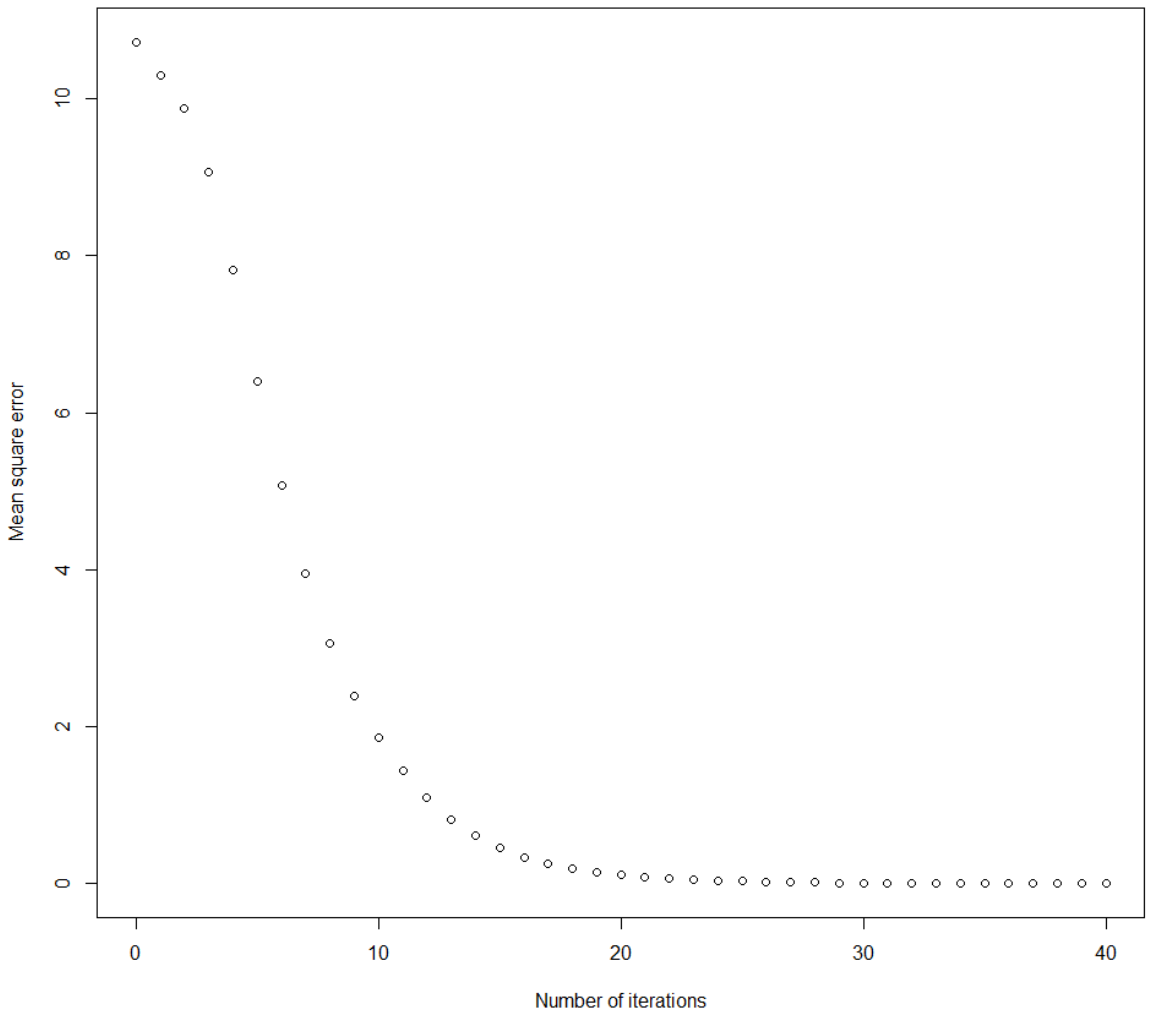

Figure 2: With the default settings, 30 iterations should be enough to get a good estimate of the cost vector.

Altering other parameters

- If the cost of throwing a die is set to be equal as the expected outcome of a throw, the results will be more complicated
- The main observations will remain the same when changes to the snakes or ladders are made or the board size is changed
- In a competitive game, the optimal strategy should be dependent on the positions of other players

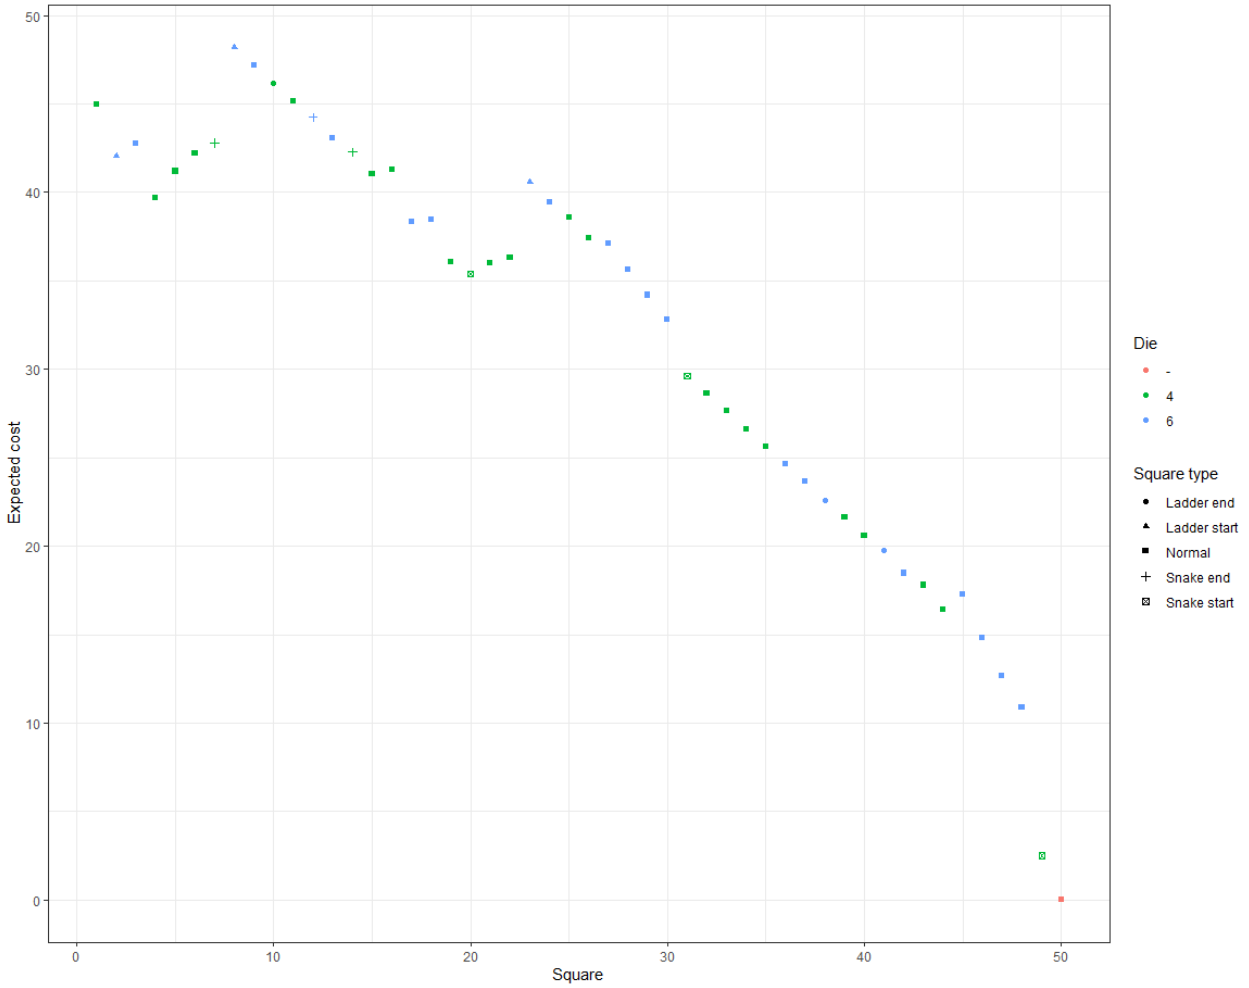

Figure 3: The plot shows the optimal strategy, when the usable dice are $\{4, 6\}$ and the cost of using a die is equal to the expected outcome of a throw.