

ULTIMATE BLACKJACKTM

USERS GUIDE

KT Enterprises

Copyright September 13, 1996
Revised January 27, 1997

For Version 1.3

Clemmons, NC 27012

PO Box 473

TABLE OF CONTENTS

	page	
Introduction		1
The Blackjack Table	3	
File Menu		5
New Game		5
Open		5
Close		6
Save		6
Save As		7
Keyboard Shortcuts	7	
Quit		8
Practice Menu		9
Play Blackjack		9
Hand Value		9
Train With		9
Check Play		9
Select Strategy		10
Define Card Count		11
Show Card Count		11
Simulation Menu		13
About Test Strategy	13	
Test Strategy	14	
Graph Of Table		16
Show Results	17	
Show Strategy		17
About Genetic Algorithm Learning		18
Genetic Algorithm Learning		20
Show Best Player's DNA		21
About Learning Agent		22
Learning Agent		23
Options Menu		26
Pause		26
Change Players		26

Casion Rules

27

Betting Strategies	29
Show Basic Strategy	30
Show Statistics	30
Show Tally	32
Reset Tally	32

Tally Statistics

33

Navigating The Tally Window	35
Interpretating The Tally Statistics	36
Tally Statistics and The Three Simulations	41
Glossary Of Terms	42

ULTIMATE BLACKJACKTM

USERS GUIDE

INTRODUCTION

Ultimate Blackjack is a commercial quality program for playing Blackjack on a computer which also has three simulation packages: test strategy, genetic algorithm learning, and learning agent.

Playing Blackjack:

The user can gamble at a realistic casino table with up to 6 computerized players. There are several features for learning to play Blackjack better and also card counting.

Test Strategy:

To find out how much money a particular strategy will make, the user can tell the computer how many players are to try the strategy and what their betting style is. The simulation will deal cards to these electronic players, and graph the holdings as the simulation is running. The range of financial outcomes is shown in a histogram.

Genetic Algorithm Learning:

This simulation looks at the artificial intelligence routine modeled after the evolutionary process. The population starts out with only a randomly generated strategy and as they play Blackjack the 6 best players are used to generate each subsequent generation. Over time the playing strategy will improve and converge on the best strategy.

Learning Agent:

This simulation uses the players acting as information gathering agents for a shared statistical memory. The result of every action is recorded and tabulated. In a systematic manner the agents will try standing, doubling, drawing, splitting, and insurance until a winning outcome is reached. Thus starting with no

knowledge, the computer can learn the best action for each situation.

Because of the complexity of the above modules and the processing time required, only one of these four modes can be active at a time. Also, switching out of one of the simulation modes, will lose the results obtained so far. So the program will present a warning dialog box in these cases to make sure the user is aware that useful information may be lost.

The active mode is shown in the menu list with a check mark next to either Play Blackjack, Test Strategy, Genetic Algorithm Learning, or Learning Agent. Also, the menu selections for commands related only to a specific mode will be dimmed until that mode is the current activity. For instance if Test Strategy is not active, then the three menu items listed just below Test Strategy will be dimmed.

Fig. 1 Blackjack Table With Wagering Controls

THE BLACKJACK TABLE

In Ultimate Blackjack you can gamble at a realistic casino table with six positions. These are signified by the player information box which includes the player's name, their current holdings, and their wager. Human players have a highlighted player information box to make them easier to identify. The computer players have either a C, M, or A in parentheses in the player information box, depending on whether their betting style is conservative, moderate, or aggressive. The betting styles are randomly chosen for the computer players on start up but may be changed using the Change Players menu option. The betting style definitions can be observed or changed in the Betting Strategies menu item.

A round starts with the dealer's hand icon indicating to each player to make their wagers. The program will make the wagers automatically for the computer players per their assigned betting style, but the dealer's hand will stop and wait for each human player to make their wagers. Fig. 1 shows the wagering controls as they appear in the center of the screen. Think of the buttons as chips that can be added to your wager. Clicking on the 10 and 25 buttons will add 35 to the current wager. To subtract from your wager, click on the Minus button and then the denomination to be subtracted. The Clear button will reset the wager to zero, and the All button will bet everything. The Half button will reduce the wager by half, while the Double button will double it. The half and double features can be useful when doing a geometric progression type betting strategy, where you keep doubling your wager till you win. There are table minimums and maximums that can be set to match your favorite casino in the Casino Rules menu option. Press on the ok button when the desired wager has been obtained.

After wagering is complete, the first two cards are dealt to each player. Then the dealer's hand icon will point to each player for them to decide how to play their cards. The program will play the hands of the computer player per the active strategy. Use Select Strategy to change the active strategy. For the human players, they must choose whether to double down, draw, stand, or split by clicking on the appropriate command button displayed in the upper left hand corner of the window and shown in Fig. 2. If the command button is dimmed, then that option is not currently available. For instance, the command button for splitting is dimmed unless you have a pair. Also, if after drawing a card the hand value goes over 21 then the dealer automatically goes to the next player.

After all hands have been played, the program will show the result of each hand by writing "won", "lost", "tied", or "BJ" for blackjack on the top card. Then the

dealer will collect the lost wagers and give chips to the hands that won. The motion of these stacks of chips is animated on the screen. After celebrating or crying, press the Deal button to start the next hand, where wagering will begin anew.

The shoe is shown in the lower left hand corner of the table. The white line marks the reshuffle point. The reshuffle point is set when the shoe is shuffled and is a random point between 70 and 90 percent of the shoe. As cards are dealt from the shoe, the graphic can be used to determine how much of the shoe has been played, or is left to play. The number of decks which make up the shoe can be changed in the Casino Rules dialog box.

Fig. 2 Blackjack Table With Command Controls

FILE MENU

NEW GAME

The New Game option will allow the user to start over by returning the program to the default starting conditions. Play at the Blackjack table or any currently running simulations are immediately discontinued and a new table is created with new players. The user will see the initial dialog box saying which table position is available. The funds for all the players are set at \$500 and a new shoe is shuffled. The statistics for each player, as well as the tally statistics, are all reset to zero.

A number of settings are returned to their default values when a New Game is selected. These are the definitions of aggressive, moderate, and conservative betting, the casino rules, the card counting definitions, and the custom strategy.

OPEN

The Open command will let the user return to a game that was saved earlier. The screen will be restored to the earlier session described in the file after aborting any currently running simulation or Blackjack practice session. The program will return to what was happening at the point the file was saved. If the user was playing Blackjack, then the casino table will appear on the screen and the player's names, funds, and other attributes will be the same as before. Even the cards left in the shoe will be the same. If the user was conducting a simulation run when the file was saved, then the simulation window will be opened and the other windows closed. The data pertaining to the simulation will be read from the file and the simulation will continue running where it left off.

The file is able to hold information on a number of program settings so that these do not have to be specified every time the program is run. The values of the following settings are read from the saved session during the Open command.

1. The names, funds and other attributes of the players at the Blackjack table.
2. The definitions of aggressive, moderate, and conservative betting.
3. The casino rule settings.

4. The player statistics.
5. The tally statistics.
6. The custom strategy is read from the file if being used when file was saved.
7. The definitions of the card values for card counting.
8. Simulation data and variables describing current status of active simulation
9. The cards in the shoe and hands of the players at the table.

The Open command will close any already open file without updating and the new file will be left open so that the save command can be used to update the file periodically.

CLOSE

The Close option will close the currently open file. If play has progressed since the file was last saved then a warning dialog will ask the user if they wish to update the file before closing it.

SAVE

This option allows a long simulation run or a good Blackjack session to be continued at a later time. This is accomplished by writing the current status of the Blackjack table or simulation into a file. If the user is playing Blackjack then the Save command will record the player's names, holdings and other attributes into the file. It will also write the cards in the shoe and the cards in the players' hands. If the user is performing a simulation, then the data collected and simulation variables are stored in the file. Also stored in the file, are current values held in player statistics and the tally statistics.

Also put in the file are several parameters which will eliminate the user from having to configure these options every time the program is run. These parameters are: the definitions of aggressive, moderate, and conservative betting, the casino rules, the card values used for card counting, and the custom strategy if being used. For instance, you can set the casino rules for your favorite casino and then save the file with the casino's name as the file name. That way this file can be double clicked to start the game, and the casino rules will be set

automatically.

The program parameters that are recorded during the save option are listed below:

1. The names, funds and other attributes of the players at the Blackjack table.
2. The definitions of aggressive, moderate, and conservative betting.
3. The casino rule settings.
4. The player statistics and the tally statistics.
5. The custom strategy is written to the file if being used when the file was saved.
6. The definitions of the card values for card counting.
7. Simulation data and variables describing current status of active simulation
8. The cards in the shoe and hands of the players at the table.

SAVE AS

The Save As option is used to specify the name of the file that the game session should be saved under. After the new file name is chosen, the information described above will be written to the file. If a file is currently open, then the current file is closed without updating, and the new file becomes the active file.

KEYBOARD SHORTCUTS

Some people may wish to use the keyboard instead of the mouse to choose between hitting a hand, standing, doubling down, or splitting a pair. Or even to tell the computer to deal the next hand. Each of these commands has a keyboard equivalent which will do the same as using the mouse. These keyboard shortcuts are shown in Fig. 3. The user can also use this dialog box to change which keys represent which command in order to create the most comfortable layout on the keyboard for themselves. Only a single character may be assigned for each command. For instance 's' may be used to represent the split command but not 'sp'. Also, a command may be defined as the space bar or the return key by typing these in the dialog text entry box. Some people may prefer this for the Deal command.

Once the keyboard shortcuts have been changed, the program will use these settings from then on. Even if the program is turned off and restarted, it will

remember the current settings. This way you don't have to redefine your key layout every time the program is run.

QUIT

The Quit command allows the user to exit the program. A dialog box will ask the user if they wish to save before quitting if the current game has not been saved.

Fig. 3 Keyboard Shortcuts Dialog

PRACTICE MENU

PLAY BLACKJACK

This menu command will bring up the casino table and allow the user to play Blackjack. This option will end the other simulation modes of Test Strategy, Genetic Algorithm Learning, and Learning Agent and also close any windows associated with them. An alert dialog box will warn the user that the current simulation will be aborted.

HAND VALUE

Some users might find playing Blackjack more entertaining if the value of the hands were displayed automatically on the screen rather than having to add up the cards every time. Select Show Hand Value to toggle the displaying of the hand values on the screen. The value of the hand will appear in a small box at the bottom of the topmost card in the hand.

TRAIN WITH

This menu option allows you to learn a strategy by concentrating on different portions of the basic strategy. For instance by selecting pairs, the first two cards dealt to each player will be the same. For hard hands, the hands do not include an Ace while for soft hands, the first card dealt to each player will be an ace. With insurance, the dealer's up-card will always be an ace. Multiple selections can be made too. You can practice insurance and pairs for instance. Remember though while this is an excellent way to obtain concentrated training, the cards dealt are not from a statistically valid shoe. Extra cards are being created to make the hand required. To return to normal play, select All Types.

CHECK PLAY

To help you learn to play better, the program can monitor your play and give varying degrees of guidance depending on the level of involvement you specify. If Beep Mistakes is selected then the program will beep when a human player tries to do a play that differs from the active strategy. If Suggest Better Play is selected then a hint box is presented in the table window which recommends a better alternative. With both of these options, the correction is only made the first time the wrong action is tried. Once the beep alerts the player or the hint window shows the proper action, the player will be able to pick any action without further criticism. The play checking feature can be turned off by selecting the No Check option.

The active strategy can be chosen using the Select Strategy menu option.

SELECT STRATEGY

The select strategy option allows for the Basic strategy or for a custom strategy to be designated. The strategy defined in this manner is used by the program two ways. First when a check play mode is employed, this is the strategy that will be used to determine if a human player has made a correct play and to suggest a better one. Second, for the computer players at the table, all of their decisions are based on the strategy selected here.

This is handy if you want to change to a particular strategy, and then monitor how well the computer players do by watching them play every hand and watching their holdings go up or down.

Or if you want to learn a strategy different than the basic one, then select the Custom... item. This will bring up a window with the current playing strategy which can then be modified as necessary. Playing then with one of the check play options, will test how well you know this new strategy.

Fig. 4 Define Card Count Dialog Box

DEFINE CARD COUNT...

This menu option is used to display or change the value assigned to various card types during card counting. A dialog box, depicted in Fig. 4, will appear which shows the point value set for each type of card. Typically, the ten's and Ace's get a negative number, like -1, while the 2's, 3's, 4's, 5's, and 6's are given a positive value such as 1. The 7's, 8's, and 9's are neutral with 0. With this point system as the cards are dealt the value for each card type can be counted and if the number is positive then the shoe has extra ten's and if it is negative then the shoe is light on ten's. You can customize you own card counting point system if you wish by changing the default values. When the OK button on the dialog box is pressed the current card count and true count values are automatically recalculated for the current point in the shoe.

SHOW CARD COUNT

Select this option to make a continuous display of the card count appear in the casino table window. In card counting, each type of card is given a point value as described above. Sometimes during a shoe the count may go high in the positive direction just through normal variation, for instance at the beginning of the deck if several small cards are dealt without any tens. One way to counteract this false impression is to factor in the position in the shoe. The true count is the card count divided by the number decks still in the shoe. Your play can be adjusted based on the current value of the true count and card count to take advantage of shoes which becomes heavily loaded with tens or to minimize exposure when the shoe runs out of tens.

You can practice card counting by keeping a running count in your head and then display the computers value at regular intervals to check yourself.

When the Train With option is in use, the card count and true count will be correct for the cards displayed, but will not be representative of a true shoe. For instance, when practicing soft hands, the first card dealt to each player is an ace that is created automatically. The rest of the cards are dealt from the shoe. Thus, the count will accurately portray the cards shown on the screen, but these cards do not portray a valid shoe.

SIMULATION MENU

ABOUT TEST STRATEGY...

The Test Strategy mode allows you to set different Blackjack strategies and see how much money can be made or lost. You tell the program how many computer players are to use the strategy and how many hands they are to play each before stopping. The number of hands can be used to approximate an amount of playing time, say 8 hours, or a weekend. There are two components to the Blackjack strategy being simulated: the betting style and the playing strategy.

Several betting styles can also be simulated. The simplest is to bet the same amount every time. But, for more realistic play, an aggressive, moderate, and conservative style are available. In these modes, the players will have a minimum base wager, but will double it a certain percentage of the time. Also, if they obtain a balance over a set threshold, the base wager is calculated as a percentage of the current balance. If the player's balance goes below zero then a loan of \$500 is extended. The particular base wagers and balance thresholds for each style can be seen or adjusted by selecting "Betting Strategies" under the "Options" menu.

The playing strategy determines what action will be taken for each possible hand value and dealer card. It says whether to draw, double down, stand, split a pair, or take insurance in each situation. After the Test Strategy dialog box is closed, the playing strategy for the players to use is presented. The default Basic strategy is given to be modified as necessary. Then the simulation will start. Since only six players can sit at the table at one time, it may take several tables to complete all the players.

While the simulation is running a graph of the players' holdings is displayed. The "Show Results" menu will present summary results in a separate window and will be updated after each table finishes. The statistics for the biggest ending balance, smallest, and average ending amount are shown. Also, a histogram gives a visual representation of the likely outcomes.

The "Show Tally" item under the "Options" menu will bring up statistical data about the Blackjack strategy being used. It is set to zero at the beginning of the

run, and can be examined to learn whether a player's hand and dealer's up-card combination makes money or loses money. It calculates the winning percentage for the actions tried, and suggests the best action.

To refer back to the strategy that is being tested, then select "Show Strategy" under the "Simulation" menu.

Fig. 5 Test Strategy Dialog Box

TEST STRATEGY

This menu option will present the test strategy dialog box shown in Fig. 5. This allows the proper information to be gathered before starting the test strategy simulation run. Make the following requested entries to customize the run.

Betting Strategy: Choose between aggressive, moderate, conservative, or the same bet every time. The definitions for aggressive, moderate, and conservative betting are available in the Betting Strategies menu item. If choosing the same bet every time, then also enter a wager amount. The wager amount is not used by the aggressive, moderate, or conservative selection.

Starting Balance: Enter the initial amount of money for each player.

Number of times: Enter the number of players that will try the strategy.

Number of hands: Enter the number of hands that each player should play the strategy for. This can be used to approximate playing time at a casino. 360 hands is roughly 8 hours.

After the values have been entered for the simulation, click on OK. The test strategy mode will then become active and end any other currently running simulation. A warning dialog box will alert the user about any aborted simulation runs before proceeding. The final dialog box will then appear (Fig. 6) so the playing strategy can be adjusted. Modify the Basic strategy shown and then click on OK. The program will automatically close any windows for any nonactive simulations and open the player holdings window.

The tally statistics will be automatically reset to zero at the start of this simulation, so that they can be inspected at regular intervals to check on progress.

Fig. 6 Test Strategy Information Dialog

GRAPH OF TABLE

While the Test Strategy simulation is running, it will continuously update the graph showing how each player is doing. The graph (see Fig. 7) shows the amount of money possessed by each player after each hand played. From this you can gain insight into the ups and downs of gambling. Since only six players can sit at a table at one time, the status bar at the bottom of the window shows how far along in completing all the players the simulation is. It also states how many tables have been competed out of the total to be done.

This window is the driver for the Test Strategy simulation. It needs to be the front most window for the simulation to operate. The simulation will wait, if other windows are in front of this one, so the user to interact with those windows unabated from a CPU time standpoint which might slow performance. Closing the Graph of Table window will not abort the current simulation, but will just suspend it from continuing. Re-displaying the window with the Graph of Table menu option will put the simulation back in operation.

EMBED Word.Picture.6

Fig. 7 Graph of Table Window

SHOW RESULTS

Use this menu option to display the results of the test strategy simulation. As each player finishes playing the required number of hands, their ending balance is recorded and included in a histogram of outcomes. The histogram is shown in Fig. 8 and has the ending amounts along its horizontal axis. On the vertical axis is the number of players with that amount. The vertical axis is located on the horizontal axis at the starting amount for the players. The vertical dashed line represents the average ending amount for the players. The histogram illustrates how with the same strategy there can be a wide dispersion in the amount of money won or lost between many players. Experiment with various betting styles to see the influence of aggressive to conservative wagering, on the histogram.

At the bottom of the window is a list of the parameters used for the simulation and the summary statistics which can be used to describe the results. The average ending balance is a primary measure of how much as a group this strategy would win or lose. The maximum and minimum ending balances are also shown along with the range (which is the minimum balance subtracted from the maximum balance).

The results window is updated automatically after each table is completed. The results window will not show a histogram until the first table is completed. So don't let the temporarily blank screen concern you.

SHOW STRATEGY

If after starting the test strategy, you have a question about the strategy that is being simulated, the Show Strategy option will create a window for examining the strategy. This is the same information that the user input at the start of the test strategy run. Since this is just for reference, the information can not be modified.

Fig. 8 Results Histogram

ABOUT GENETIC ALGORITHM LEARNING

This simulation mode explores the capability of a computer to learn how to play Blackjack by using a popular theory called the Genetic Algorithm which is based on the principals of evolution and natural selection. In the animal kingdom, the ability of each animal to survive in their environment and prosper is influenced significantly by the animal's genetic makeup. The animals that are the strongest will earn the right to mate and produce offspring. In this way the better genes are passed on to the next generation. And the offspring of two parents that are superior in their ability to survive, should also be superior. An important wildcard aspect of evolution is genetic mutation which allows children to have abilities their parents did not. Mutations in nature are genetic changes that occur from incorrect copying of the DNA code or from environmental factors like radiation. Some mutations can be an advantage if they help the animal survive while others can be a disadvantage that will die out.

Back to Blackjack. The program uses a population of 36 players, whose Blackjack DNA is the strategy information which tells them whether to take a card, stand, double down, split a pair, or take insurance. At the start this information is just randomly initialized. Or the user may choose to set this information himself for the six players that produce the first generation.

To give each player an opportunity to prove themselves, six tables are played for the specified number of hands, with six players sitting at each. The best player at each table is selected to produce the next generation. By mating the six players with each other, a new generation of 36 players is obtained. The children's playing strategy, or DNA, is created by each card combination being copied from one parent or the other, with a 50%, 50% chance for each. During this offspring production, mutations which are just random changes to the DNA code (or playing strategies) will occur. The user specifies a mutation rate in the dialog box as the number of changes per 1000 genes. You may wish to select a larger mutation rate at the beginning and then reduce it as the players are getting better.

Normally, the offspring will be a result of just the two parents and mutations, but if the user selects the Filter With Tally Statistics option, the results being tabulated in the tally statistics will be examined. If the tally statistics for a hand and dealer combination shows an advantage to a specific action after a sufficient

number of entries (100), then the offspring DNA will be adjusted to this action. This will allow a mix of learning due to genetics and by trying various actions.

Like real evolution the genetic algorithm takes time. After each table is completed, the statistics are shown on the screen for the best player. You can keep tabs on progress by watching the ending balance to see if it is improving. Also the DNA of the best players can be monitored as they converge to the best strategy. Select "Show Best Player's DNA" under the "Simulation" menu to examine the DNA of the best player at each table. The tables that have not yet been completed will not show any values. The program will run until stopped by the user selecting another simulation mode or to play Blackjack.

Fig. 9 Genetic Algorithm Learning Dialog Box

GENETIC ALGORITHM LEARNING

This menu option will present the dialog box shown in Fig. 9, which will allow the user to enter the parameters needed to start the genetic algorithm learning simulation.

Player Strategy: The playing strategy for the starting 6 players will be randomly generated, or the user can select to set these manually. The simulation starts by mating the 6 starting players to generate the first generation of 36 players. If the user selects to set these manually, the dialog box to adjust the playing strategy will appear after this one. This option is only available at the start or restart of the simulation.

Reproduction Method:

Player Alone: If selected, the 6 best players alone will be used to determine the next generation.

Filter With Tally Statistics: Select this to incorporate information in the tally statistics. If a card combination has been tried at least 100 times, and shows that one action is better than the others, then this action will be set for the next

generation.

No. Hands: Each table will play these many hands.

Initial Balance: \$500. This is the starting amount for each player.

All Bets \$5. This is the betting strategy for the simulation.

Mutation Rate: Enter a value here which will be the number of genes that mutate per 1000. A higher mutation rate may help find better plays at the start, but may create too many backward steps as the solution converges. So you may want to start with a slightly higher value at the start and then lower it as the players get better.

Continue: Use this to change the simulation parameters and then continue the simulation already in progress.

Restart: Select this to change the simulation parameters and to start over.

After selecting the desired options and entering the requested information click on the OK button. The active mode will switch to the genetic algorithm learning, and a notice to the user will warn about aborting any active simulation. Then if applicable, the window to manually set the strategies for the initial 6 players will be opened. Click on OK when done.

Upon starting or restarting, the tally statistics will be automatically reset to zero, so that they can be inspected at regular intervals to check on progress.

The Genetic Algorithm Learning window has to be front most for the simulation to continue. The simulation will pause when other windows are in front.

SHOW BEST PLAYER'S DNA

The genetic algorithm selects the best player at each table to produce the next generation. The strategy information, or the player's DNA, can be observed with this menu item for the 6 best players. A window will be created which allows you to access the strategy information. There is a radio button to display the DNA for each of the best players. The information can also be modified to play what if scenarios, or to introduce your own mutations. The changes are being made to the players which will generate the upcoming generation.

Each generation is comprised of six tables being played in turn. The DNA

information for the best player shown in the Best Player's DNA window will be blank until the table for that player is complete. For instance, information for best player number 5 will not be observable until table 5 is finished.

ABOUT LEARNING AGENT

The Learning Agent is another method of computer learning used in artificial intelligence. It uses an entity, called an agent which can interact with the environment and learn from the outcome. In our Blackjack situation, six players act as agents with a shared memory for their actions and responses. The tally statistics serve this purpose.

For each hand and dealer card combination, the actions are each tried in turn for the specified number of times (set by the user). First standing is tried, after which doubling down is tried, followed by drawing. If one of these actions is found to be profitable then the remaining action(s) are not evaluated. If none make money then after the performance data for all three actions has been compiled, the action that mitigates the losses is taken from then on.

In a manner similar to the other actions, insurance and splitting pairs are “learned” by first declining for the specified number of tries. If this alternative is not a money maker then taking insurance or splitting for the same number of times is tried. Afterwards the alternative that makes the most money is selected from then on. But, insurance and splitting pairs are dependent on the hands being played correctly in order to obtain a valid result. For this reason the Learning Agent dialog box has a checkbox for you to tell the program to start learning insurance and splitting. This way you can wait until the drawing, standing, and doubling is learned reasonably well first.

Also in the dialog box are checkboxes to tell the program to re-learn various actions. To re-learn an action, the tally statistics for that action are reset to zero and the action is tried the set number of times again. This feature can be used, for instance, to see if drawing were tested again would the tally statistics show the same conclusion. If the number of tries is small there is a greater possibility that the results may not be representative of the true statistical outcome. If the number of tries is large, the results are more statistically valid, but the simulation will take longer to run. Individual hand and dealer combinations can also be chosen to be re-learned in the tally window. After clicking the Erase checkbox, select which squares are to be set to zero and thus re-learned.

The Learning Agent mode will evaluate alternatives till it finds those that make money or minimize losses. Use the Tally Window to examine the statistical outcomes obtained. See how well these compare to the Basic strategy. In general,

a table of six players play till the desired number of hands is reached. At this point the statistics for the best player is shown on the screen which allows the ending balance to be monitored to see how it improves over time. The balances are then reset to \$500 and the deck is shuffled as the next table starts. The simulation will continue till stopped by the user.

Fig. 10 Learning Agent Dialog Box

LEARNING AGENT

The dialog box, Fig. 10, for starting the learning agent simulation will be presented by this menu option. Enter the items described below.

No. of Tries: This value (N) tells the program how many times to try an action like standing, before going on to try other actions. If this number is small it is possible that the results may not be representative of the true statistical outcome. But, alas the larger the number of tries the longer the simulation will take to run, so choose a number depending on the time you want to wait for an answer and the quality of the answer you want. For a short run 5, and 10, will finish quicker, but the results will be more of an approximation. A middle run of 15, 20, to 40 will take longer to run but will be a better approximation. Large values of 50, 100, and 200 will provide the most accurate results but the program will require the most time to try each alternative this many times.

No. Hands: Each table of six players will play this many hands.

Initial Balance: This is the starting amount for each player.

All Bets \$5: This is the betting style used in this simulation run.

Start/Restart: By checking this box, the simulation will abort and restart with the new parameters.

Continue: If the user desires, the learning agent parameters can be adjusted and the simulation continue on. This is useful when selecting re-learning of various options.

Learning:

Since some actions like splitting pairs and taking insurance, depend on the hands being played properly to obtain the best results, the dialog box has two checkboxes for the user to tell the program when the strategy has been learned reasonably well enough that learning about splitting and insurance should be started.

Splitting: Click on this checkbox to have the program start to learn splitting pairs. After checking this box, the tally information for pairs will be reset to zero. Then not splitting will be tried for N times. If that does not make money then splitting will be tried N times. After which the best financial move will be played.

Insurance: Click on this checkbox to have the program start to learn insurance. After checking this box, the tally information for insurance will be reset to zero. Then declining insurance will be tried for N times. If that does not make money then buying insurance will be tried N times. After which the best financial move will be played.

Re-Learning: Since the learning is slightly dependent on the number of tries that are set. There may be occasion that the user would like the program to re-learn a particular action without having to erase everything and start over. So the re-learn checkboxes allow the tally statistics for specific actions (drawing, standing, etc.) to be reset to zero and then the program will start collecting data about them again.

There is another way to re-learn a specific card and hand combination. If the tally information is looking fine except for a few squares that are not matching the expected results, then these can re-learned if desired. With the tally window displayed, click on the Erase checkbox and then click on all the squares that need to be re-learned. The tally information for each square selected will be reset to zero for all actions. Continuing to run the simulation now will allow new data to be collected. Remember to click on the Erase checkbox again when you are finished erasing specific squares, so that the erase mode will be turned off.

After entering all the designated information into the agent learning dialog box, click on the OK button. This will make the learning agent the active mode and if leaving another simulation an alert dialog will warn the user that the other simulation will be aborted.

If start or restart is selected, then the tally statistics will be automatically reset to zero at the start of this simulation, so that they can be inspected at regular

intervals to check on progress.

The Learning Agent window has to be the front most window on the screen for the simulation to run. Otherwise it pauses to allow interacting with the other windows.

OPTIONS MENU

PAUSE

This menu command will halt any of the simulations that may be currently running: test strategy, genetic algorithm, learning agent. This will also suspend activity for the play Blackjack mode. Once the program is suspended the Pause menu changes to “Continue” to signify that if it is selected again normal play will be continued.

CHANGE PLAYERS

The user can change several characteristics about the players to create various effects. You can play Blackjack with a friend at one of the other positions, or you can play multiple hands. Or you can practice playing from each position. A table position can be vacant, filled by a computer player, or a human player. Vacant positions are skipped over during play at the table. The human players are signified at the table by the highlighted player information box. While the human players have to make all their own wagers and play commands, the computer players will perform these automatically. The wagering of the computer players is determined by the setting in the Change Players dialog box (Fig. 11). The computer players can be conservative, moderate, or aggressive. A “(C)”, “(M)”, or “(A)” is shown in the player information box to indicate the style of wagering the computer players are making. The definition of the three betting styles can be studied in the Betting Strategies dialog box. For the human players the betting style setting is ignored.

Also in the Change Players dialog box is shown a current financial statement of each player’s holdings. The current amount the player has and how much they have borrowed are both shown. Whether the player is in the black or the red can be determined by mentally subtracting the borrowed amount from the current balance. These financial numbers can be changed by the user.

Lastly, the player’s name can be changed in the dialog box. Be creative.

If the user wants to play multiple hands then they can use the Change Players dialog to make some of the computer players be human. This also allows up to six friends to play Blackjack at the same table. An auto mode results if you designate all the players as "computer". You can watch the screen as all the hands and wagers are handled automatically. This can be used as a short run simulation too. This automatic play can be stopped with the pause menu option.

Fig. 11 Change Players Dialog Box

CASINO RULES

You can customize the casino rules to match your favorite casino with this menu option. Or use this option to change the casino rules and see how the results of the test strategy simulation are affected. The following items can be changed using the dialog box shown in Fig. 12. These items are in effect while playing at the casino table, and while any simulation is being conducted. In the simulations and with the computer players if the desired wager is outside the table minimum or maximum, it is set equal to the closest limit. For human players, a warning dialog box appears describing the situation.

The number of decks can be changed from 1 to 8, with the default being 8.

The table minimum can be set as one of the following: 1, 2, 5, 10, 25, 100, or 500. 5 is the default.

The table maximum can also be set at 200, 500, 1000, 5000, or to have no limit. The no limit setting can be used to try with various geometric progression betting strategies.

Typically, the dealer has to draw on 16 and stand on 17 or above, but some casinos offer a variation on this where the dealer will draw if he has a soft 17. The user can choose between these two settings.

Casinos offer a variation on when a player can double down. The least limiting allows doubling on any two cards. The next limiting allows doubling on hands with two cards that total 9, 10, or 11. The most limiting only allows doubling on hands that total 10 or 11 with 2 cards. Use the popup menu to select between these three.

Some casinos do not allow doubling after a hand is split and so the Yes/No radio buttons can be used to set this parameter.

Fig. 12 Casino Rules Dialog Box

BETTING STRATEGIES

With this option you can refer to the betting style definitions or change them to produce your own effects. The algorithm is versatile and can be used to have a computer player place the same wager every time, or to double the base wager a specific percentage of the time, or to even have the amount of the base wager increase to keep up with increased holdings.

First consider the minimum wager parameter as shown in Fig. 13. The amount of the wager will be at least this amount, say \$100. If the player's balance goes below zero a loan is made automatically to give the player \$500 gambling money. But sometimes doing the same wager all the time gets boring and so the base wager can be doubled (to \$200 in our example) a certain percentage of the time. This percentage is entered as a number between 0 and 100 and is the first parameter in the dialog box.

If the player wins several hands in a row, the winnings may become so large that making a \$100 wager every time may not move the earnings further or reduce them for that matter. Thus, there is a provision that if the balance goes above a specified threshold, then the wager is calculated as a percentage of the holdings. Both the threshold level and the percentage are the next two parameters in the dialog box. If the wager calculated this way does not exceed the minimum wager, then the minimum wager is made.

Essentially, the difference between the default conservative, moderate, and aggressive betting style is just a difference in degrees. The minimum wager is higher and the frequency that the base wager is doubled is greater as the betting style becomes more aggressive. Also after the threshold is reached, the subsequent wagers are a larger percentage of the holdings for the more aggressive styles.

Through manipulation of the betting style parameters, several variations can be obtained. For instance:

A. To make the same wager every time.

1. Set the percentage of the time for doubling to 0
2. Set minimum to the desired wager
3. Set the threshold balance to any number

4. Set the percent of holdings to 0

B. To make the same wager most of the time but double it every now and then.

1. Set the percentage of the time for doubling to the desired frequency
 2. Set minimum to the desired wager
 3. Set the threshold balance to any number
 4. Set the percent of holdings to 0
- C. To make every wager a percentage of the current holdings
1. Set the percentage of the time for doubling to 0
 2. Set minimum to 1
 3. Set the threshold balance to 0
4. Set the percent of holdings to desired proportion

Fig. 13 Betting Strategies Dialog Box

SHOW BASIC STRATEGY

This option is included for reference. There may be times you wish to refresh your memory about the Basic Blackjack strategy. This option will bring up a window with the Basic strategy drawn in it. The information in this window can be examined fully, but it can not be changed by the user.

SHOW STATISTICS

Show Statistics offers another way to gauge how well a player is doing besides how much money they have made. Detailed statistical information (see Fig. 14) is available on how many hands were won, lost, and tied for all hands together, as well as a break down by straight hands and double down hands. The number of blackjacks received and number of hands that were split is also included. These statistics can help understand why a player is winning or losing.

This statistical information is reset when ever a simulation starts a new table, or when switching from one simulation to another or to playing Blackjack. In the simulations a table is completed when all the players have played the specified number of hands. At the start of the next table these statistics will be reset. If the casino table window is visible then the statistics represent the performance of the human and computer players.

The statistics are taken at a snapshot in time. They can be examined in detail but if left open for many hands, the window will need to be closed and re-selected from the Options menu to display more recent information.

Fig. 14 Statistics Window

SHOW TALLY

The tally window is displayed through this command. The tally window comprises very detailed statistical information on the times that a particular dealer and player hand combination won, lost, or tied for each possible action that can be taken. The tally window and the tally statistical package are discussed in detail in the Tally Statistics section. The tally statistics are stored every time a card is played. But the tally window shows the statistics at a particular point in time. If the window is left open for many hands, then close the tally window and re-select Show Tally to display the latest information.

RESET TALLY

This option is a convenient way to reset all the tally information described above to zero. This can be used to start collecting data a new and to erase the previous information.

When conducting the Learning Agent simulation, resetting the tally information will cause all the alternatives of standing, doubling, drawing, splitting, and insurance to be retried and the learning process to be started over again.

When the genetic algorithm learning is active and the filter with tally option is in use, resetting the tally information will suppress the effect of the filter until a positive result has been obtained and confirmed by 100 tries.

TALLY STATISTICS

Ultimate Blackjack includes a detailed and comprehensive statistical package for analyzing Blackjack card situations, called tally statistics. The tally statistics were developed to determine the best action in each Blackjack card situation and also to find out which hand and dealer card combinations win more often than lose. The tally statistics got the name because it refers to keeping a tally on which actions produce positive results and which ones have undesirable outcomes. The tally statistics can be broken down into four main portions for: hard hands, soft hands, pairs, and insurance. In each of these four areas, every possible value of a player's hand is matched with every possible dealer's up-card to form a matrix which describes every possible Blackjack situation.

The tally statistics try to show whether each action is the right play or not and the most effective way to judge this is based on the outcome of the hand, because it all comes down to did the hand win or lose. If the hand won then all the decisions used to play the hand (draw, stand, split, take insurance, etc.) all get credited with a win. If the hand loses, then all the decisions or actions taken during the hand are credited with a loss. Over time as more hands are played the information about each card combination will yield how many times an action, like drawing, produced a winning hand versus a losing hand or a tie.

Some examples might clarify the tally procedure:

1. Hard and soft hands. Say a player has an ace and a 3 for a total of 14. The dealer shows an 8. So the player takes a hit and gets a 10 for a total of 14, but now the hand is a hard hand. The player takes another card and is dealt a 6 for 20, upon which he stands. The dealer turns over a 10 for 18. The player wins. In the tally statistics for soft hands and the dealer having an 8, a win will be recorded for the player drawing with 14. In the tally statistics for hard hands and the dealer having an 8, a win will be recorded for drawing on 14, and for standing with 20.

2. Pairs. Another player is dealt a pair of 8's for a total of 16 and the dealer has a 5 showing. The player chooses to split the pair and is dealt a 10 on the first hand and a 9 on the second hand. The player is happy with these cards and stands on both. The dealer turns over a hole card that is a 6 and then draws a 10 for a total of 21. Both hands lose. So in the tally statistics for hard hands when the dealer

shows a 5, a loss is recorded for standing with 18 and for standing with 17. In the tally statistics for pairs and the dealer showing a 5, two losses are recorded for a pair of 8's.

3. Insurance. Imagine a player has 20 with the first two cards, but the dealer has an ace. So the player purchases the insurance by putting up an additional amount equal to half his wager. The dealer peeks and does not have a blackjack. Later the player wins with his 20. The tally statistics for insurance situations will record a win for when the player takes insurance, and the dealer does not have blackjack, and the player's hand wins.

Fig. 15 Tally Window -- View All

NAVIGATING THE TALLY WINDOW

The tally window is shown in Fig. 15. In the upper right hand corner are the view radio buttons. The default is All Of The Above which shows the 4 main sections on the screen at once. Each of these 4 areas can be selected by clicking on them. A read box will highlight the last one selected. In addition, to help identify each square, you can click on a square to select it and information about the square will appear in the square status area. The information includes the hand type, the player's hand value, the up-card, and the recommended action based on the data collected so far.

In the bottom right hand corner are two radio buttons which can change what the colored tally squares are showing. They can be set to show the best action at each card combination or to show the winning percentage. Both of these are based on the data accumulated up to now. The color shade of each square is set to portray the range of best actions and the winning percentages. The color legend is shown just above the "SHOW" radio buttons and the meaning changes depending on whether the winning percentage or the best play is being displayed.

The winning percentage looks at the recommended best play for which data has been recorded and calculates the financial outcome with that action. Four colors are shown to portray the following degrees of profitability. The first is if the winning percentage is greater than 20%. The next is if the winning percentage is greater than 0 and less than or equal to 20%. The next one, is if the action breaks even or loses less than or equal to 20%. The last group is if the action loses more than 20%. You might ask why would the best action have an outcome that loses money. The answer is that all the other possible actions lose more money or have not been evaluated. There is also a color that represents no data being available.

If the best move is being displayed then the color codes squares refer to the move recommended by the collected statistics. For hard and soft hands, they comprise four options: 1) to double down when available otherwise draw a card; 2) to double down if possible or else stand pat; 3) to draw a card; 4) or to stand. A color is also used to indicate that no data is available to make a recommendation. The color legend for pairs and insurance is simplified since there are only two alternatives for each. Do split the pair/take insurance, or do not.

Both the winning percentage, and the best play, results are based on the data collected. If only three hands have been played and all were won, then the data

would show a 100% likelihood. In actuality, the next 7 hands may be losers, and this will give a different perspective. Just remember if the number of entries is small, there is the possibility that the data collected may suggest something that does not match the conclusion of more data. Learn to inspect the data of various squares that look unusual to see what the data is based upon.

By clicking on the first 4 radio buttons, the desired area to be examined more closely is brought to the screen, whether it be hard hands, soft hands, pairs, or insurance. These views are used to display the detailed statistical data, when a square is selected with the mouse. To show the statistics accumulated for a hand value of 11 and a dealer's up-card of 3, just click on the colored square with 11 on the left and 3 on the top. A box will highlight the square selected and the statistics for this card combination will be presented at the bottom of the screen. If all the data is zero then the occasion for this particular card combination has not come up yet during play. Each of the four areas for hard hands, soft hands, pairs, and insurance can be examined in this manner to see the data that has been collected.

INTERPRETING THE TALLY STATISTICS

Hard Hands and Soft Hands

For hard hands and soft hands, there are three actions that can be tried: doubling, drawing, and standing. The performance of each is segregated in Fig. 16 for comparison. For each, the number of hands that were won, lost, and tied was tallied. From these the total is summed and the percent won, percent lost, and percent tied is calculated. These percentages are used to determine the gain percentage or the whether an action will win or lose if played. The gain percentage is calculated by taking the percentage of the time that the action won and subtracting the percentage of the time that the action lost. For example if drawing won 20 % of the time, lost 10 %, and tied 70 %. Then the gain percentage would be $20 - 10 = 10\%$. Your first thought might be to say that the percent won is less than 50 and so is not a money maker, when in actuality you will win more than you lose and the profit will be 1 hand for every 10 played.

Fig. 16 Tally Statistics For Hard Hands

After calculating the gain possible with each action, this information can be used to compare the actions to determine which one makes the most money. One thing to consider is that the double down wager is double the normal hand and so the same percentage gain will return more money than for standing or drawing. For instance, a 10% return on \$10 is better than a 10% return on \$5. So to determine if doubling is advantageous, the gain with doubling needs to be multiplied by 2 and then compared to the gain with both drawing and standing. This can be expressed as

do doubling: if $2 * (\text{gain with doubling}) > (\text{gain with drawing})$
 and $2 * (\text{gain with doubling}) > (\text{gain with standing})$

do drawing if $(\text{gain with drawing}) > (\text{gain with standing})$

do standing

if $(\text{gain with standing}) > (\text{gain with drawing})$

The frequency shown in the tally window is the way the program communicates the recommended best play. The frequency indicates the percentage of the time that each action should be tried. If it says 100 percent for one action, then that action is the recommended best play and the other two are worse financially. When doubling is advantageous, then the frequency says what percentage of the time you will be doubling and what percentage of the time, the other action will be executed. To determine the frequency of doubling, start with the number of times doubling was allowable which is the “doubles available” number tallied and displayed in the tally window. Then divide the total number of hands played for all actions into the number of available doubles to yield the percentage of the time that doubling is possible.

The program will recommend an untried alternative over one that loses money. But the winning percentage is calculated on the best play for which data is available to calculate a profit which avoids untried alternatives. The winning percentage is determined by multiplying the frequency for the recommended

action by the gain for that action.

Pairs

For dealing with pairs there are two actions: splitting or not splitting the pair. The results with each of these actions are shown in Fig. 17, so they can be compared. Similar to the hard hand and soft hand discussion above, the percent won, percent lost, and percent tied are calculated for both actions. From these the gain with splitting and with not splitting can be calculated again as the percent won - the percent lost.

To determine which action is better, the percent gain for each needs to be compared while taking into account that for splitting hands there is a doubling of the wager. So like with doubling down, splitting that has a 10 % return on \$10 is better than a 10% return on \$5 for an non-split hand. Thus, if two times the gain with splitting is greater than the gain without splitting then splitting is better. This can be written as

do splitting

if $2 * (\text{gain with splitting}) > (\text{gain without splitting})$

Once the best play is chosen the frequency for that action is set to 100.

The winning percentage is just the gain of the best play that data has been collected for.

Fig. 17 Tally Statistics For Pairs

Insurance

Insurance may seem one of the simpler choices because there are only two actions, to insure or not to insure, but as we'll see the possible outcomes when insurance is taken are numerous. Fig. 18 shows the data collected for insurance hands. The statistics for when insurance is declined are calculated in the usual way.

When insurance is taken the outcomes not only deal with whether the hand was won, lost, or tied, but also whether the insurance bet was won, or lost. If the insurance bet was won then the dealer had blackjack. The five outcomes and their unique payoffs are shown in Fig. 18 based on a \$10 wager and a \$5 insurance bet. Two of the situations make money and two lose money. While one does neither. To calculate the gain for when insurance is taken, the percent occurrence for each group is multiplied by its payoff and then dividing by the sum of the wagers, \$15.

To compare the gains of taking insurance to those with declining insurance, remember that the insurance action has a combined wager 1.5 times as large as the other. So similarly to doubling, if the gain with insurance when multiplied by 1.5 is greater than the gain without insurance then insurance should be taken. Or in other terms

take insurance:

if $1.5 * (\text{gain with insurance}) > (\text{gain without insurance})$

The frequency is assigned to 100 percent for the play that makes the most money. The winning percentage is just the gain with the most profitable play for there is data collected.

Fig. 18 Tally Statistics For Insurance

TALLY STATISTICS AND THE THREE SIMULATIONS

Test Strategy

When the test strategy simulation is run the strategy does not change over the course of the simulation, but rather is fixed. In this situation, every time, for instance, a player has a hand total of 7 and the dealer has a 10 it will draw if that's what the strategy dictates. The players will never choose to stand or double down and so the tally statistics will compile data on the effects of drawing, but the results of standing or doubling will be unknown. To use the test strategy method to compare different actions, you need to run the simulation with one strategy, and then change the strategy for the squares that you are interested in and run it again. You may wish to write down the results from the first simulation to compare to the second.

Genetic Algorithm

When the genetic algorithm simulation is running, the strategy used by the players is changing over time as they evolve. The players will be starting out with little ability and will be getting better over time. At the start players may be hitting a hand that has 21 or doubling on 19 because they don't know anything better yet. Successful sequences of drawing cards usually end in a final decision to stand. If the player is hitting high hand totals and the hand loses than all the drawing decisions are tallied with a loss. Even though it may have been the last decision alone that was wrong. In this way the tally statistics for drawing may be skewed for poorer players. Likewise, splitting and insurance are dependent on the hands being played correctly for the most accurate tally results. The tally based conclusions now may be different than when the players' ability improves. As the players' ability improves the tally statistics may become a mix of poor and better play. It may be necessary to reset the tally statistics at intervals to see how the players are doing now. Incidentally, the tally statistics for standing and doubling are independent of the player's abilities since they receive an immediate and direct result of winning or losing.

Learning Agent

The advantage of the learning agent simulation is that the results of all actions are obtained during the same run and so can be compared to each other easier. This information may be used to justify the Basic strategy. The simulation uses a systematic approach to learning each action which will help reduce errors in learning. Although depending on the number of tries used and other factors, there may be tally squares or card situations which show an unusual conclusion

and these can be selectively erased in the tally window in order to be re-learned.

GLOSSARY OF TERMS

The best play is that play which either maximizes the profit or at least minimizes the losses. The best play is determined by comparing the gains of individual actions as discussed in this manual.

The card count is the point value of each card that has been dealt added together.

The doubles available is the number of hands tallied so far for which doubling was permissible even if doubling was not tried.

The frequency is used to refer to the percentage of the time that each action should be pursued. Except when doubling down is involved, the frequency will be 100% for the recommended best play and 0% for the other alternatives. When doubling down is profitable, the frequency for doubling is just the number of times that doubling is available divided by the total number of hands. The second most profitable action (either standing or drawing) has a frequency that is 100 minus the doubling frequency. This secondary action will be done when doubling is not allowed.

The gain is the return obtained by playing an action, like drawing or splitting. A 20% return means that for 10 hands played, the profit will be equal to the wager for 2 hands. In all cases except for insurance the gain is calculated by the percent of hands won minus the percent of hands lost.

A strategy is a collection of rules which describe how to handle each player hand and dealer up-card combination. The strategy can be broken down into four parts dealing with hard hands, soft hands, pairs, and insurance. The actions possible include drawing, standing, doubling down, splitting pairs, and taking insurance.

The true count is the card count divided by the number of decks remaining to be dealt from the shoe.

The winning percentage is the return that will be obtained by pursuing the best play. Except for when doubling is involved, the return is the same as the gain for the best play. When doubling is advantageous, the winning percentage is the

frequency of doubling times the gain for doubling plus the frequency for the second most profitable action times its gain.