Maximizing Winnings on Final Jeopardy!

Jessica Abramson, Natalie Collina, and William Gasarch

August 2017

1 Abstract

Alice and Betty are going into the final round of Jeopardy. Alice knows how much money both she and Betty have and her and Betty's probabilities of getting the question right. Betty knows how much money both she and Alice have and her and Alice's probabilities of getting the question right. How much should the two players wager? This depend on their goals. Prior work defined the player's payoffs as their probability of winning the game. We find strategies using payoffs of (a) the expected amount of money that the player takes home that day, and (b) the expected about of money that the player takes home overall (since if a player wins they play again). We employ a maximin strategy in order to determine the optimal wager. We give mathematical results to question (a) and use assumptions based on historical Jeopardy data to determine optimal wagers for (b).

2 Introduction

Consider a final round of Jeopardy! with players Alice and Betty [1]. We assume that going into this round Alice has the highest monetary total. The procedure is as follows:

1. The players hear the category of the final question.
2. The players place wagers.
3. The host gives them the clue in the form of an answer and the players attempt to come up with the correct response in the form of a question.
4. The correct response is revealed.
5. If a player responds correctly, she adds the amount she wagers to her total. If she responds incorrectly, she loses that amount from her total.
6. The player with the most money takes home her total and is invited back to play another game the next day. In Jeopardy! the second-place player ends the

[1] While Jeopardy! always begins with three participants, in some cases one of these players has a non-positive sum of money and therefore cannot participate in Final Jeopardy! Furthermore, using a two-player model is a useful and common simplification when considering Final Jeopardy!, as we discuss in the Previous Work section. Rarely does this simplification impact results.

game with 2000 dollars. However, for the purposes of this paper we will treat the losing player's payoff as 0 dollars, since this payoff is usually tiny compared to the first place players' winnings. As the average Jeopardy! winner leaves the game with about 20,000 dollars, winning overwhelmingly yields more money [1].
7. If there is a tie, the players answer tie-breaking questions (without wagering) until exactly one of the players responds correctly. This player wins the game and takes home her money. The other player finishes in second place and leaves with 2000 dollars. But for the reasons stated above we will treat this player's winnings as 0 dollars.

For the purposes of this paper, Alice will always be defined as the player with the most money. Betty will always trail Alice going into the final round. Throughout the entire paper we assume that Alice and Betty each know (a) how much money she and the other has, and (b) her probability of responding to the final question correctly.

The breakdown of our paper is as follows: in Section 3 we review previous work. In Sections 4-12, this payoff is strictly defined as the expected amount of money that the player will bring home on that day. Here we do not consider that she may win money when she comes back the next day. We will refer to our exploration of this payoff as our *single-day analysis*. In Sections 4-7 we establish notation and a useful diagram. In Section 8 we present Alice's optimal strategy, assuming she knows Betty's probability of responding correctly. In Section 9 we do the same for Betty. In Sections 10 and 11 we provide cheat sheets that Alice and Betty can actually use while playing Jeopardy!. In Section 12 we discuss how to apply the cheat sheets if a player does not know the other player's probability of responding correctly, using historical Jeopardy! data.

Our maximin strategies and solutions for this single-day scenario may be generalized to other simultaneous single-action wagering games. However, in Jeopardy!, the winning player goes on to play in later games. In Section 13 we consider a scenario in which each player's payoff includes the expected amount of money she would win on following days, given that she triumphs in the current game.

3 Previous Work

Gilbert and Hatcher [2] defined the structure of the two-player payoff matrix which we utilize in this paper. In their paper, they define the game's payoff as the players' probabilities of winning and identify equilibrium points and mixed strategies based on this payoff. Ferguson and Melolidakis [3] utilized this same payoff matrix in their 1997 paper to identify minimax points, which are wagers that one player can make which minimizes the maximum possible playoff for the other player regardless of their wager. In a zero-sum game this is equivalent to maximizing the minimum possible payoff for oneself regardless of the other player's wager.

Both these papers define the payoff as a player's probability of winning the game. In this paper, payoff will be defined as the expected amount of money the player will win. As a result, our payoff matrix is continuous and discontinuous in the same areas as previous matrices, but the payoffs themselves are different. Our paper will more closely follow Ferguson and Melolidakis's minimax strategy than Gilbert and Hatcher's equalibrium strategy. However, unlike in Ferguson and Melolidakis's paper, our payoffs are not zero-sum, since an increase in one player's payoff does not always correlate with a decrease of equal magnitude in the other player's payoff. Therefore we utilize a maximin strategy instead of a minimax strategy.

Our distinct payoff definition leads to significant strategical differences between our work and both previous papers. For example, when Alice has more than twice the amount of money that Betty has, she can always wager 0 dollars to ensure herself a win. A strategy focused on win probability will always recommend this wager, possibly along with other wagers small enough to guarantee Alice a win. This recommendation holds in the equilibrium and minimax strategies when defining payoffs only according to chances of winning the current game. However, according to our payoff definition, in which Alice wants to maximize her expected monetary winnings, there are cases in which she should wager all her money instead.

The Gilbert and Hatcher paper and the Ferguson and Melolidakis paper have both excluded a third player from their calculations, and we follow this model. While Jeopardy! is always played with three players, Ferguson and Melolidakis note that given their primary analysis of a three-player strategy, adding a third player with the least money is unlikely to significantly alter the strategy. We similarly make this simplification. However, some of our maximin proofs can be extended to a three-player game, as will be noted.

4 Variables

Table 1: Variables

M_A	Alice's money coming into the final question
M_B	Betty's money coming into the final question
P_A	Alice's probability of responding to the final question correctly
P_B	Betty's probability of responding to the final question correctly
W_A	The amount of money Alice wagers
W_B	The amount of money Betty wagers
P_{A_T}	Alice's probability of winning a tiebreaker
P_{B_T}	Betty's probability of winning a tiebreaker

5 The Maximin Strategy

In order to find optimal strategies for players, we consider what a player can wager to maximize her minimum payoff. The wager which accomplishes this is her maximin wager, and the smallest possible payoff given that wager is her maximin.

The maximin strategy is a common game theoretical approach for simultaneous single-action games, as it provides a worst-case payoff guarantee regardless of other players' actions. This strategy simulates a scenario in which players have no knowledge of others' wagers and do not have multiple rounds of wagering to reach an equilibrium point.

In the first part of our single-day analysis, we will examine when a dominant strategy wager exists for the leading player Alice. A wager that is a dominant strategy must be a maximin wager, because by definition a dominant strategy is a strategy that is optimal in every case no matter what the other player wagers. Therefore the scenario resulting in the dominant wager's lowest possible payoff results in at least as low payoffs for all other wagers. If a dominant strategy exists, then, this is the maximin wager.

When no dominant strategy exists, we find maximins by utilizing our structural knowledge of the two-player payoff matrix.

6 The Payoff Matrix

The optimal strategies for the players can be determined using a payoff matrix. The matrix represents all possible combinations of wagers for Alice and Betty, where the row represents Alice's wager and the column represents Betty's wager. This general payoff matrix has a predictable structure.

This matrix has distinct regions defined by their own payoff equations and are discontinuous from region to region. These different regions relate to the four possible outcomes of the players responding to the final question: both players respond correctly (CC), neither player responds correctly (II), only Alice responds correctly (CI), or only Betty responds correctly (IC). Whether the outcome is a win, loss, or a tie for a player in each case depends on how much each player wagers. Each region of the matrix represents combinations of wagers that have the same winner for each outcome.

The general shape of the payoff matrix, based off the payoff matrix developed by Gilbert and Hatcher [2], can be seen in Figure 1 with the winner in each of the four Final Jeopardy! outcomes (CC, II, IC, and CI, respectively) listed along with each region. 'A' indicates that Alice wins, 'B' indicates that Betty wins, 'T' that there is a tie, and 'L' that, as they both end with 0 dollars, both players lose. For example, region 6 is labeled ABBA because Alice ends the game with the most money in the CC and CI case and Betty does otherwise. The diagonal regions 1, 3 and 5 are "tying lines," places in which some outcome results in a tie. Regions 2 and 4 are intersections of tying lines where more than one outcome results in a tie. These regions each exist only for a single wager combination. The shape of the payoff matrix changes depending on the ratio between M_A and M_B, but the regions below remain relatively consistent when $M_B < M_A < 2M_B$.

Figure 1: General Structure of the Payoff Matrix

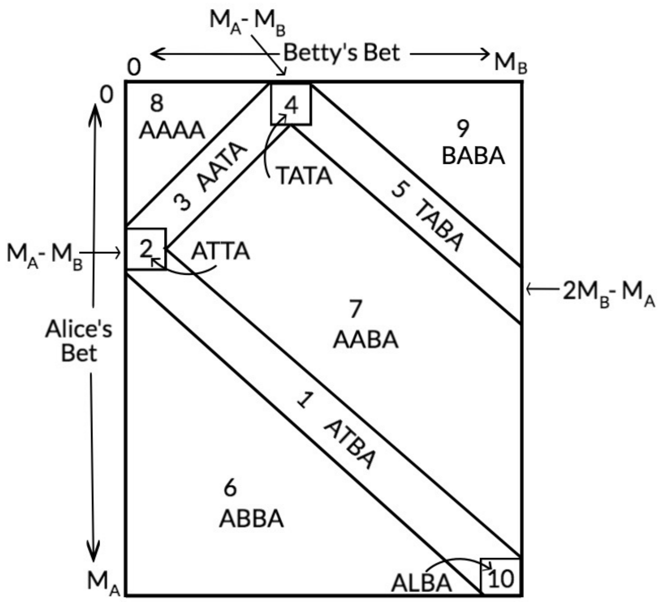

The following table identifies the payoff equations for Alice within each region.

Table 2: Payoff Equations by Matrix Region

Region	Alice's Payoff
1	$P_A(M_A+W_A)+P_{A_T}*(1-P_A)(1-P_B)(M_A-W_A)$
2	$P_{A_T}*(1-P_A)(M_A-W_A)+P_A(M_A+W_A)$
3	$P_A(M_A+W_A)+(1-P_A)((1-P_B)(M_A-W_A)+P_{A_T}(P_B)(M_A-x))$
4	$(1-P_B)(P_A(M_A+W_A)+(1-P_A)(M_A-W_A))+$ $P_{A_T}P_B((P_A)(M_A+W_A)+(1-P_A)(M_A-W_A))$
5	$P_{A_T}P_AP_B(M_A+W_A)+(1-P_A)(1-P_B)(M_A-W_A)+P_A(1-P_B)(M_A+W_A)$
6	$P_A(M_A+W_A)$
7	$P_A(M_A+W_A)+(1-P_A)(1-P_B)(M_A-W_A)$
8	$P_A(M_A+W_A)+(1-P_A)(M_A-W_A)$
9	$(1-P_B)(P_A)(M_A+W_A)+(1-P_B)(1-P_A)(M_A-W_A)$
10	$P_A(M_A+W_A)$

7 Payoff Matrix Maximin Strategy

The final payoff from the payoff matrix is determined by both Alice and Betty's wagers. Alice can only impact her own wager, and has no control over Betty's choice. Thus, she can choose which row of probabilities from the matrix to accept, but cannot control which value within that row will represent her actual payoff. Similarly, Betty can choose the column of the wager combination, but not the row. The maximin strategy accounts for this uncertainty, maximizing the lowest possible payoff for a player.

As the expected amount of money a player gains depends not only on the cases in which the player wins but how much money the players wagers, payoffs are not constant within each region. However, the payoffs are still defined by continuous equations in each region. Furthermore the payoff equations for a player never include the other player's wager as a variable. Therefore given a particular wager by Alice, her payoff is impacted only by which region Betty makes the wager combination fall in, not what specific value Betty wagers. To illustrate this concept, let us refer to Table 2. If Alice wagers 3 dollars, Betty could cause the wager combination to fall into region 2 or 7. When Betty wagers more than 0 dollars, Alice's payoff will be in region 7. Regardless of Betty's specific wager, Alice's payoff will be 4.4 dollars. Similarly, if Betty wagers 0 dollars, Alice can cause the wager combination to fall into region 8, 2, or 6. If

Alice wagers more than 3 dollars, the wager combination will fall into region 6 and Betty's payoff will be 2.8 dollars regardless of Alice's specific wager.

The number of payoff possibilities for a player that could result given a single wager is therefore equal to the number of distinct regions which exist in the row or column that she wagers.

Given a player's single wager, W, P_A, P_B, M_A, and M_B are all constant. Therefore, the payoff equations that this wager could fall into can be compared directly. In order to find the lowest possible payoff for a player's single wager, we can determine the least favorable payoff equation. As the players win, lose or tie in different combinations in each region, it is often clear which payoff equation is inferior. For example, region 6's winner results are ABBA, while region 7's are AABA. Alice wins in a case in region 7 where she loses in region 6. Therefore, given a certain wager by Alice, if these are the two potential payoff equations, region 6 would certainly offer her the lower payoff.

Once lowest payoff(s) for each of the players' potential wagers are found, the next step is to maximize them. Within the wagers which share the same min region of 6, for example, the best choice is the wager where the payoff in region 6 is the highest. Region 6's payoff equation is optimized when Alice wagers M_A. Therefore if the maximin is in region 6, it will be when Alice goes all in.

In this way, the number of potential maximins can be reduced down to at most a handful. From there, it may be necessary to directly compare payoff equations to determine the maximin and therefore the maximin wagering strategy.

Overall, this technique for finding the maximin for a player involves finding the region with the worst payoff equation for each of her wagers, maximizing the payoffs within those least desirable regions, then comparing those payoffs to find the one that affords the largest value, which is the maximin. We use this payoff matrix technique when using dominant strategies is not a possibility.

8 Alice's Strategy

Before utilizing the payoff matrix maximin strategy, we will examine when general dominant strategies can be found.

8.1 When $P_A \geq \frac{1}{2}$

When Alice's chance of answering the question correctly is greater than 50%, there is always a dominant strategy for her regardless of how much money each player has or the other players' probabilities of winning. We will find her maximin wager by finding this dominant strategy.

We can upper bound Alice's payoffs for all wagers with the assumption that she wins the game in all Final Jeopardy! outcomes. This is expressed by the equation:

$$E_{W_A} \leq P_A(M_A + W_A) + (1 - P_A)(M_A - W_A)$$

$$= M_A + W_A(2P_A - 1)$$

When Alice wagers M_A dollars, and therefore $W_A = M_A$, she will always lose if she responds to the question incorrectly and always win when she responds correctly, so her payoff is

$$P_A(M_A + Z_A) + (1 - P_A) * 0$$

$$= P_A(M_A + M_A)$$

We will choose to rewrite this as

$$M_A + 2P_A M_A - M_A$$

$$= M_A + M_A(2P_A - 1)$$

As $W_A = M_A$, this equation is exactly equal to the equation

$$M_A + W_A(2P_A - 1)$$

We would like to prove that a wager of M_A for Alice yields a greater payoff than any other wager does. We have determined an upper bound for all Z_A: $M_A + W_A(2P_A - 1)$. In order to prove that M_A is optimal, we will prove that the payoff is always larger than this upper bound.

There is a linear relationship between W_A and E_{W_A}, expressed by the equation $E_{W_A} = M_A + W_A(2P_A - 1)$.

When $P_A \geq \frac{1}{2}$, $2P_A - 1 \geq 0$, so we can maximize the payoff by maximizing W_A to M_A. As wager M_A yields a higher payoff than the upper bounds of all other wagers when $P_A \geq \frac{1}{2}$, this wager is a dominant strategy. Therefore

Alice's maximin wager is M_A.

Note that as long as Alice remains the leading player, this proof extends to the three-player case without loss of generality.

8.2 When $P_A < \frac{1}{2}$

When $P_A < \frac{1}{2}$, Alice's optimal wager has multiple cases:

8.2.1 When $M_A > 2M_B$

When the leading player Alice has more than twice Betty's money, we can again find a dominant strategy. Alice will always win the game when wagering 0 dollars. Therefore when $W_A = 0$ her payoff is M_A. This value can again be expressed by the equation

$$M_A + W_A(2P_A - 1)$$

In this case, therefore, a wager of 0 is exactly equal to the upper bound equation of all wagers.
We would like to prove that a wager of 0 for Alice yields a greater payoff than any other wager does. Again, we know an upper bound for all Z_A: $M_A + W_A(2P_A - 1)$. In order to prove that a wager of 0 is optimal, we will prove that the payoff is always larger than this upper bound.

There is a linear relationship between W_A and E_{W_A}, expressed by the equation $E_{W_A} = M_A + W_A(2P_A - 1)$. When $P_A \leq \frac{1}{2}$, $2P_A - 1 \leq 0$, so we can maximize E_{W_A} by minimizing W_A to 0. As wagering 0 yields a higher payoff than the upper bounds of all other wagers when $P_A \leq \frac{1}{2}$, this wager is a dominant strategy. Therefore

$$\textbf{Alice's maximin wager is } 0$$

.

Note that as long as Alice has at least twice as much money as all other players, this proof extends to the three-player case without loss of generality.

8.2.2 When $\frac{3}{2}M_B < M_A < 2M_B$

For the remaining cases, there are no clear dominant strategies for Alice in a three-player game. Therefore we will find Alice's maximins utilizing the payoff matrix maximin strategy. The matrix's structure changes depending on the relationship between M_A and M_B, and in the case where $\frac{3}{2}M_B < M_A < 2M_B$, the bottom of region 9 is at least one row above the top or region 6.

Given this relationship between M_A and M_B, Alice has three potential maximin wagers.

If she wagers less than $2M_B - M_A$, Alice's worst region is region 9. As $P_A < \frac{1}{2}$, this equation is maximized when $W_A = 0$. This will yield a minimum payoff of $(1 - P_B)M_A$ when wagering 0.

If Alice wagers exactly $2M_B - M_A$ dollars, her worst region is region 5.

If Alice wagers exactly $M_A - M_B$ dollars, her worst region may be either region 2 or region 7.

If Alice wagers more than $M_A - M_B$, her worst region is region 6. This is maximized when Alice wagers M_A dollars, yielding a min payoff of $2P_AM_A$.

Finally, if Alice wagers such that $2M_B - M_A < W_A < M_A - M_B$, her worst region is region 7. We refer to these rows in which region 7 is the minimum payoff for Alice as min-7. As seen in Table 3, region 7's payoff equation is

$$P_A(M_A + W_A) + (1 - P_A)(1 - P_B)(M_A - W_A)$$

When $P_A > (1 - P_A)(1 - P_B)$, this equation is optimized by maximizing W_A within min-7, and otherwise it is optimized by minimizing W_A within min-7. We must determine which of the potential maximin wagers discussed above yield the highest payoff. We will define two different cases given the two optimizations of region 7.

Figure 2: Structure of the Payoff Matrix when $\frac{3}{2}M_B < M_A < 2M_B$

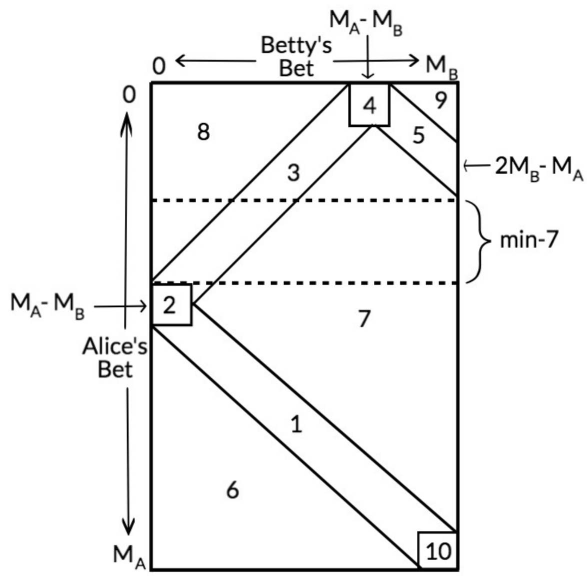

Case 1: $P_A > (1 - P_A)(1 - P_B)$
Here, the higher that W_A is in region 7's payoff equation, the greater the payoff is. However, to stay in min-7 Alice can only wager as much as $M_A - M_B - 1$. This means one potential maximin wager for Alice is $M_A - M_B - 1$.

However, this payoff is always lower than that of going all in. As mentioned previously, the minimum payoff equation for Alice when going all in is

$$P_A(M_A + W_A)$$

When $W_A = M_A$, this can be written as

$$P_A(M_A + W_A) + (1 - P_A)(1 - P_B)(M_A - W_A)$$

as $M_A - W_A = M_A - M_A = 0$. This is the payoff equation in region 7.

Therefore we can consider Alice wagering M_A as a continuation of region 7. As Alice wants to maximize Z_A given region 7's payoff equation, she would rather wager M_A than any other value. Therefore the potential maximin when wagering M_A always has a higher payoff than wagering the most in min-7.

Wagering $M_A - M_B$ is also inferior to going all in for Alice. When region 2 has a higher payoff than region 7, it is not even the minimum in its row. When it has a lower payoff, it is inferior to some payoff equation in region 7, which we have shown is itself inferior to going all in.

Wagering M_A is also superior to wagering $2M_B - M_A$ to have a minimum in region 5, using similar logic.

Finally, wagering M_A is also a better strategy than wagering nothing. As in this case $P_A > (1 - P_B)(1 - P_A)$ and $1 - P_A > \frac{1}{2}$:

$$P_A > (1 - P_B) * \frac{1}{2}$$

Which by algebra, is equivalent to:

$$2P_A M_A > (1 - P_B)M_A$$

The value on the left is Alice's minimum payoff when wagering M_B and the value on the right is her minimum payoff when wagering 0.

Therefore when $P_A > (1-P_A)(1-P_B)$ and $P_A < \frac{1}{2}$, Alice's minimum payoff given a wager of M_A is higher than the minimum payoffs for all her other wagers.

Alice's maximin wager is M_A

.

Case 2: $P_A < (1 - P_A)(1 - P_B)$

In this case, we would like to minimize W_A in min-7 to $2M_B - M_A + 1$. We can again consider the maximin when going all in as a continuation of region 7's payoff equation, and therefore when region 7 should be minimized, going all in is clearly inferior to any wager in region 7 itself.

Region 2 is never a maximin in this case, for similar reasons as in the former case.

Therefore we must compare three payoffs for Alice: wagering $2M_B - M_A + 1$ in region 7, $2M_B - M_A$ in region 5, or 0 in region 9. Without making further assumptions,

Alice's maximin wager could be $2M_B - M_A + 1$, $2M_B - M_A$ or 0

See section 8 for the specific equation comparisons.

Determining the optimal wager is therefore trivial, as Alice can calculate the values of these three equations using her specific values of M_A, M_B, P_A and P_B and wager according to the equation with the highest value.

We note that only in cases with both extremely high values of P_{A_T} and extremely low values of P_A, P_B and M_B is it possible for region 5 to be the maximin. For example, when $P_A = .16$, $P_B = .1$, $M_B = 1700$, and $P_{A_T} = .99$, it is best to wager in region 5. For practical purposes players may be better off disregarding this option.

8.2.3 When $\frac{3}{2}M_B > M_A$

In this case, any wager from Alice could result in at least one of region 6 or region 9, as the regions overlap vertically. These are the two potential maximin regions in the matrix. As we have previously shown, the payoff equation for region 9 is maximized to $(1 - P_B)M_A$ when $W_A = 0$. The payoff equation for region 6 is maximized when $W_A = M_A$, so the payoff is $2P_A M_A$.

Figure 3: Structure of the Payoff Matrix when $\frac{3}{2}M_B > M_A$

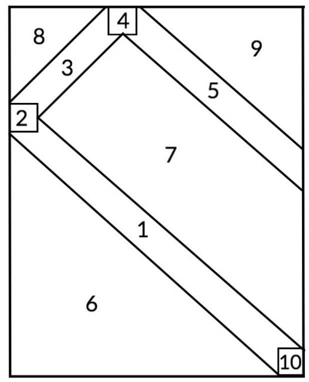

These two payoffs represent potential maximin payoffs for Alice.

$$\textbf{Alice's maximin wager is } 0 \textbf{ when } (1 - P_B)/2 > P_A$$

And

$$\textbf{Alice's maximin wager is } M_A \textbf{ when } (1 - P_B)/2 < P_A$$

. If the two equations are equal, these wagers are equally optimal payoffs for Alice.

9 Betty's Strategy

Let us consider the optimal strategies for Betty, the trailing player.

9.1 When $M_A > 2M_B$

In this scenario, when Alice wagers 0 dollars, she can guarantee that the result is in region 8. Therefore all wagers for Betty lead to the same maximin of 0 dollars, so

Betty's wager does not matter.

9.2 When $M_A < 2M_B$

If Betty wagers less than $2M_B - M_A$, then her maximin is certainly in region 8, which leads to a payoff of 0. For any wagers greater than this, her worst payoff equation is region 7. Therefore Betty's maximin is in region 7. Thus, Betty would like to maximize her payoff in region 7. The payoff equation in this region is

$$(1 - P_A)(P_B)(M_B + W_B)$$

Maximizing W_B to M_B will maximize this equation to $(1-P_A)(2P_BM_B)$.

Betty's maximin wager is M_B.

10 Cheat Sheet for Alice

We will now restate the theorem for Alice's optimal wager.

10.1 $P_A \geq \frac{1}{2}$

Alice's maximin wager is M_A.

10.2 $P_A < \frac{1}{2}$

10.2.1 $M_A > 2M_B$

Alice's maximin wager is 0.

10.2.2 $\frac{3}{2}M_B < M_A < 2M_B$

Case 1: $P_A > (1-P_A)(1-P_B)$
Alice's maximin wager is M_A.

Case 2: $P_A < (1-P_A)(1-P_B)$
Consider three equations:
Equation 1:

$$M_A(1-P_B)$$

Equation 2:

$$P_A(2M_B+1)+(1-P_A)(1-P_B)(2M_A+2M_B+1)$$

Equation 3:

$$P_{A_T}(2M_B)+(1-P_A)(1-P_B)(2M_A+2M_B)+P_A(1-P_B)(2M_B)$$

If Equation 1 yields the largest value, Alice's maximin wager is 0. If Equation 2 yields the largest value, Alice's maximin wager is $2M_B-M_A+1$. If Equation 3 yeilds the largest value, Alice's maximin wager is $2M_B-M_A$.

10.2.3 $\frac{3}{2}M_B > M_A$

Alice's maximin wager is 0 when $\frac{(1-P_B)}{2} > P_A$ and M_A otherwise.

11 Cheat Sheet for Betty

11.1 $M_A > 2M_B$

Betty's wager does not matter, as she always has a maximin of zero dollars.

11.2 $M_A < 2M_B$

Betty's maximin wager is M_B.

12 Taking a Realistic Perspective

What if Alice Does Not Know P_B? What if Betty Does Not Know P_A? In Sections 8-11 we assumed that each player knew the other's probability of responding to the question correctly. We assume Alice knows how likely she is to respond to the final question correctly based on her knowledge of the category of the final question. However, she likely does not know how knowledgeable Betty is on the topic of the final question, so this secondary assumption may not be realistic.

Historical Jeopardy data shows that in the last 10 seasons the average correct response rate on the final question is about 49.6% with a standard deviation of about 3.6% [1]. Without knowing anything about Betty's knowledge of the topic of the final question, Alice's best assumption is that Betty has the same chance as the average player of responding correctly. Hence, it would be reasonable for Alice (Betty) to use the cheat sheet assuming $P_B = \frac{1}{2}$ ($P_A = \frac{1}{2}$). Note that this does not reduce the number of possible cases.

13 Maximizing Overall Winnings: A Jeopardy!-Specific Scenario

In Jeopardy!, players who win a game come back to play the next day. Thus, if a player aims to maximize her payoff from Jeopardy! as a whole, not just on one specific day, she should consider her expected future winnings. We define "win benefit" to be the amount of money a player expects to win in future Jeopardy! games should she win the game on the current day. In the following section, we consider a realistic perspective of a leading player going into Final Jeopardy! who hopes to take home the most possible money from all the games she participates in.

13.1 Simplifications and Statistical Assumptions

In order to realistically model the perspective of Alice, a leading player about to wager in Final Jeopardy!, we assume $P_B = \frac{1}{2}$ for the remainder of the paper. Refer to section 12 for justification.

In the previous sections, making this assumption could be useful for determining which case and corresponding optimal wager applies, but does not reduce the total number of cases. Therefore it was not necessary to make this assumption in the case calculations themselves. However, when considering the win benefit, assuming $P_B \frac{1}{2}$ does reduce the number of cases to a reasonable number and allows us to give advice to Alice which could be applied to a real game of Jeopardy!

We also use historical Jeopardy! data to quantify the win benefit. The average winnings on Jeopardy! are about 20,000 dollars per game. Furthermore, the returning champion wins Jeopardy! in about 46.26% of games [4]. Thus, using $\frac{a}{1-r}$, where $a = 0.4626$ and $r = 0.4626$, the equation for the sum of an infinite geometric series, a player can expect to win Jeopardy! about 0.86 more times after winning the first game. Thus, we multiply these values together to estimate that a player can expect to win about 17,000 dollars in future Jeopardy! games if she wins a current game. We use 17,000 dollars as the win benefit for the remainder of this paper.

We consider the cases in which Alice has between 0 and 50,000 dollars going into Final Jeopardy! The largest one-day winnings in Jeopardy! history is 77,000 dollars from a game in which the winning player entered Final Jeopardy! with 47,000 dollars [5]. Thus we assume Alice is unlikely to enter Final Jeopardy! with more than 50,000 dollars and provide advice to Alice accordingly.

We will not consider tying lines in this section. As we have shown in the previous section, these ties rarely impact a player's optimal strategy.

13.2 Strategy

When considering the payoff with the win benefit, the first two steps for finding the maximin for a player using the payoff matrix strategy do not change. These steps are finding the region with the worst payoff equation for each of her wagers and maximizing the payoffs within those least desirable regions. This is because the payoff equation for each region when a win benefit is added in can be expressed as:

$$\text{expected winnings that day} + (\text{probability of winning that day}) * (\text{win benefit})$$

We previously compared the payoffs of different regions given a single wager by comparing the outcomes in which the regions afforded each player a win. The expectation of the win benefit also scales by the number of cases in which a player wins.

In all cases where we determined that one region was strictly worse than another in a given row or column, it was because this region afforded the player in question equal or worse win/tie/loss results for every outcome. Therefore these regions also have lower win probabilities and thus their payoffs are still lower given a single wager.

However, adding in the win benefit can impact the final step, which is comparing the potential maximin values. In determining in which cases each minimum payoff is optimal, we utilized a combination of algebra and the Desmos online graphing calculator. As many of these comparisons are trivial but tedious, some of our inequalities have been solved only graphically.

13.3 $M_A > 2M_B$

Figure 4: Maximin Wagers for Alice when $M_A > 2M_B$

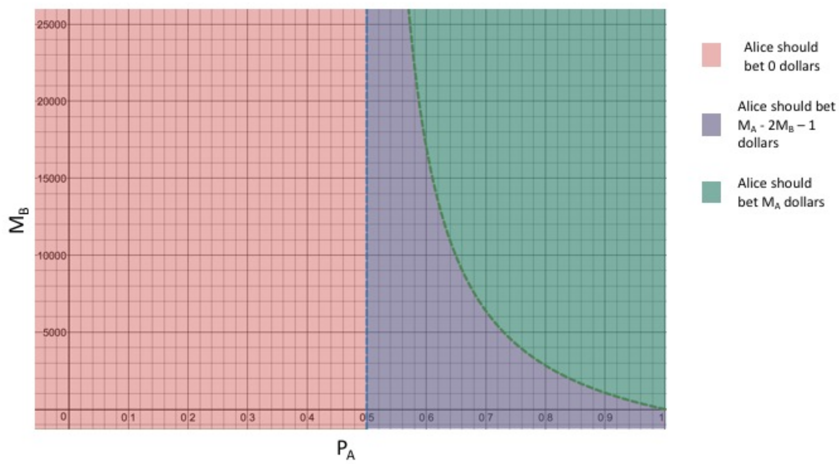

Given this relationship between M_A and M_B, the three potential maximins for Alice are region 8's top row, region 8 in the lowest row which contains only region 8, and region 6 in the bottom row. The shaded areas indicate when each payoff is the greatest. Given our assumptions about P_B and the win benefit, this depends only on M_B and P_A.

The intuition for why Alice's maximin in this scenario is influenced by M_B and not M_A is as follows: Alice's min payoff when wagering M_A is affected only by the value of M_A, while her payoff when wagering $M_A - 2M_B - 1$ is affected by the relationship between M_A and M_B. As M_A increases, both payoffs increase and the relationship between the mins stays the same. When M_B increases, however, wagering $M_A - 2M_B - 1$ becomes worse and wagering M_A has the same payoff.

13.4 $\frac{3}{2}M_B < M_A < 2M_B$

Given this relationship between M_A and M_B, the three potential maximins for Alice are the top row of region 9, the bottom row of min-7 (the lowest row containing only region 7), and the bottom row of region 6. The shaded areas indicate when each payoff is the greatest. Given our assumptions about P_B and the win benefit, this depends only on M_A and P_A.

Figure 5: Maximin Wagers for Alice when $\frac{3}{2}M_B < M_A < 2M_B$

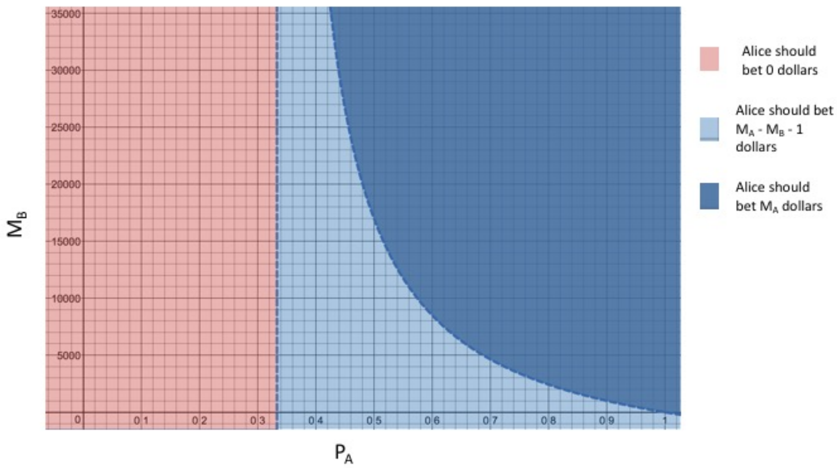

13.5 $M_A < \frac{3}{2}M_B$

Figure 6: Maximin Wagers for Alice when $M_A < \frac{3}{2}M_B$

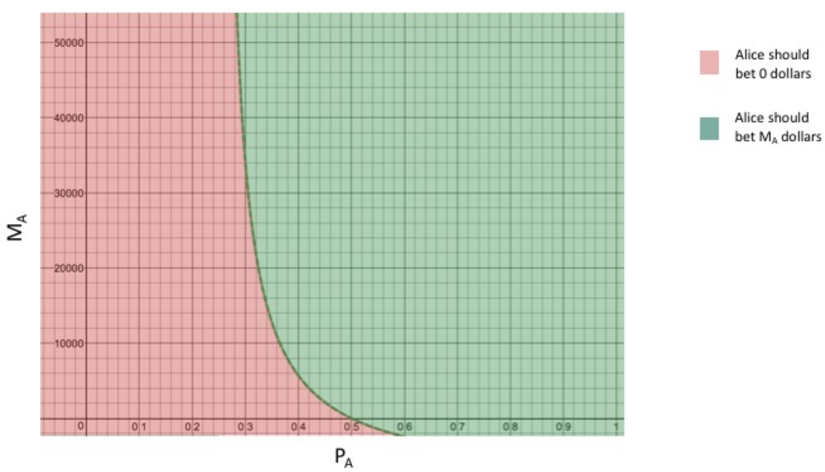

Given this relationship between M_A and M_B, the two potential maximins for Alice are the top row of region 9 and the bottom row of region 6. The shaded areas indicate when each payoff is the greatest. Given our assumptions about P_B and the win benefit, this depends only on M_A and P_A.

References

[1] Final jeopardy! round statistics.

[2] George T. Gilbert and Rhonda L. Hatcher. Wagering in final jeopardy! *Mathematics Magazine*, 67(4):268–277, 1994.

[3] Thomas S. Ferguson and C. Melolidakis. Last round betting. *Journal of Applied Probability*, 34(4):974–987, 1997.

[4] Strategy talk: How much is a win worth on jeopardy! – the jeopardy! fan, Apr 2017.

[5] Ud's roger craig wins unprecedented $77,000 on 'jeopardy', Sep 2010.