HIGH SCHOOL EDITION
Journal of Student Research

Volume 14 Issue 1 (2025)

Bayesian Game Theory: Limitations of Incomplete Information in Determining Probabilities and Payoffs

Ashley Lattyak[1] and Patricia Jacobsen[#]

[1]Sacramento Country Day School, USA
[#]Advisor

ABSTRACT

Game theory is a complex area of study based on principles of mathematics and statistics. Simple, two-player games such as the Prisoner's Dilemma demonstrate the basics of probabilities, outcomes, and payoff matrices. Often, game theory involves games of complete information, where all information is known to all players. Bayesian games explore a specific type of multiplayer game, called games of incomplete information, where some information about the players or the game is unknown. This study analyzes Bayesian games with unknown player identities. The study uses the Prisoner's Dilemma as a basis for game theory payoff matrices. The Prisoner's Dilemma is then used as a simple Bayesian game example, followed by a more detailed analysis of the Sheriff's Dilemma. The study shows that probabilities are much more complex in Bayesian games than in games of complete information. The results demonstrate that the ideal strategy of a player could depend on multiple factors in a Bayesian game. The study concludes that ethical considerations are a significant factor in Bayesian games, and this type of analysis can be applied to a variety of fields. Broader applications include studies in economics, sociology, law, and many other fields.

Introduction

People make many choices that can be affected by various factors, such as other people, internal or external biases, etc. We use game theory to mathematically analyze a person's choices and outcomes in these situations.

Game Theory

Game theory is a topic of mathematics involving payoff matrices for players' decisions in certain situations. These payoffs can be interpreted as benefits—positive results—or costs—negative results. Game theory applies to multiplayer situations where each player's outcome depends not only on their actions but also on the actions of other players. Game theory is relevant to a plethora of fields: biology, business, psychology, politics, finance—the list goes on. Situations range from business decisions about lowering prices to moral decisions about whether to confess a crime. Game theory can be useful for evaluating the advantages and disadvantages of certain choices and for analyzing the ways in which the human brain makes decisions (Romanowski 2014).

Invented by Hungarian-American mathematician John von Neumann, game theory began as an extension of poker strategies. The concept then extended to warfare, representing conflicts between the United States and the Soviet Union in the Cold War. As game theory gained popularity, mathematicians began experimenting with different types of games and their characteristics (Chen et al.).

Games defined by game theory payoff matrices adhere to certain rules. Each player has a set number of choices to make, independent of the other players, but they must consider the choices and outcomes of other players. There is a winner or a loser, determined after a certain number of moves have been made, and players either have a benefit or a cost. In some games, players are aware of all the information in the game, and in others, players lack some knowledge during play. Strategies vary for each type of game (Romanowski 2014).

The simplest type of game is a zero-sum game with just two players. In a zero-sum game, the benefits and costs perfectly balance each other, and the sum of the players' outcomes for any given decision is zero. Zero-sum games with two players are the least complex because an ideal approach for each player can be easily determined by simply comparing the payoffs—this is commonly referred to as the Minimax Theorem. There are also more complex games called nonzero-sum games, which are much more common in realistic life situations. The benefits and costs are not necessarily equal in magnitude, and the Minimax Theorem cannot be applied because the probabilities are more complex (Romanowski 2014).

The Prisoner's Dilemma

One popular example of a nonzero-sum game with two players is the Prisoner's Dilemma. In this scenario, two people are being held in prison for a crime that the police know they have committed, but they don't have substantial evidence to keep them in prison for a long time. Each prisoner, or player, has the choice to confess and turn the other in or keep quiet (Ross 2023). In the payoff matrix below, the payoffs are on a scale from 0 to 10 of how good or bad the situation would be for each player, 0 being very bad and 10 being very good. If neither player confesses, both experience a good situation. If both confess, they experience a mild situation. If only one player confesses and the other doesn't, the one who confesses experiences a very good situation, and the one who does not confess experiences a very bad situation. The best strategy for each player in this dilemma, underlined in the matrix, is to confess because the situation is better in each scenario, no matter what choice the other player makes—5 is better than 0, and 10 is better than 8. This simple example not only demonstrates the basics of logic behind game theory but also highlights the connection between game theory and moral reasoning: in some situations, players might consider whether a choice is not just beneficial to them but also whether it is morally right.

Table 1. Payoff matrix for the Prisoner's Dilemma.

Player 1	Player 2		
		Confess	*Don't Confess*
	Confess	5, 5	10, 0
	Don't Confess	0, 10	8, 8

One important detail of The Prisoner's Dilemma and other similar games is that these are games of "perfect information," where the players know all possible moves, outcomes, and characteristics of the other players (Romanowski 2014). What would happen if there were multiple types of players with different payoff matrices? What if each player wasn't aware of the other players' types, and therefore couldn't make a fully informed decision? These questions and complexities are explored with the concept of Bayesian games—games of incomplete information in which players must consider the probabilities of other players being certain types. Bayesian games more accurately model real-life situations because, in reality, most people don't have complete knowledge of every complexity in a given situation (Zamir 2009). This paper will explore these complexities, analyze examples of Bayesian games, show calculations of probabilities, and discuss moral and ethical implications.

HIGH SCHOOL EDITION
Journal of Student Research
Volume 14 Issue 1 (2025)

The Prisoner's Dilemma as a Bayesian Game

The Prisoner's Dilemma can be represented as a Bayesian game if we introduce incomplete information, such as an unknown player type. In this example, we will assume that one player isn't sure if the other player is altruistic. Let's say that Player 2 is not altruistic, so their payoffs are the same in every scenario. Their ideal strategy is still to confess. However, Player 1 could be altruistic or could be not altruistic. If Player 1 is altruistic, they will always prefer to not confess because, for moral reasons, they choose to be selfless and do not want to convict the other player—4 is better than 2, and 9 is better than 6. If Player 1 is not altruistic, the payoff matrix is the same as in the previous example, when both players were assumed to be not altruistic, and Player 1 would prefer to confess.

Table 2. Payoff matrix for when Player 1 is altruistic.

	Player 2		
Player 1		*Confess*	*Don't Confess*
	Confess	2, 5	6, 0
	Don't Confess	4, 7	9, 9

Table 3. Payoff matrix for when Player 1 is not altruistic.

	Player 2		
Player 1		*Confess*	*Don't Confess*
	Confess	5, 5	10, 0
	~~*Don't Confess*~~	~~0, 10~~	~~8, 8~~

The ideal strategy for Player 2 is more complex in this game. In the second payoff matrix, the ideal strategy for both players is still to confess. However, in the first payoff matrix, Player 2's strategy can only be determined by remembering that Player 1 will always choose to not confess. Then, we can modify the payoff matrix.

Table 4. Payoff matrix for when Player 1 is altruistic, some options crossed out.

	Player 2		
		Confess	*Don't Confess*

Player 1	~~*Confess*~~	~~2, 5~~	~~6, 0~~
	Don't Confess	4, 7	9, 9

If Player 2 knows that Player 1 will not confess, then they will also choose to not confess—9 is better than 7, and the ideal choice for each player is to not confess, as shown in underlined in the matrix above. However, if Player 2's best choice could either be to confess or to not confess depending on Player 1's type, their ideal strategy is not obvious, so we can use probabilities. Assume that the probability of Player 1 being altruistic is p, and the probability of Player 1 being not altruistic is 1-p. The expected payoff of Player 2 confessing or not confessing can be calculated in terms of p. For confessing, we can multiply the payoff of confessing when Player 1 is altruistic, 7, by probability p, and add that to the product of the payoff of confessing when Player 1 is not altruistic, 5, and probability 1-p, as shown below.

Player 2's expected payoff for confessing:

$$7p + 5(1 - p) = 2p + 5$$

This method can be repeated to calculate Player 2's expected payoff for not confessing:

$$9p + 0(1 - p) = 9p$$

Finally, we can set these payoffs equal to each other to determine Player 2's ideal strategy:

$$2p + 5 = 9p$$
$$5 = 7p$$
$$p = 5/7$$

If $p > 5/7$, Player 2 should confess. If $p < 5/7$, Player 2 should not confess. If $p = 5/7$, Player 2 could make either choice.

By turning the Prisoner's Dilemma into a Bayesian game, we can see the complexities of multiple types of players with varying strategies. This example may more accurately reflect reality because a person's reasoning in this situation could greatly vary depending on their moral compass or other factors.

Case Study: Sheriff's Dilemma

We can further analyze the ethics of Bayesian games through the Sheriff's Dilemma. In this scenario, the two players are a sheriff and a suspect, both choosing whether to shoot the other player. There are some underlying assumptions about each player's preferences. The sheriff prefers to make the same choice as the suspect; if the suspect chooses to shoot, the cop would like to as well, but if the suspect chooses not to shoot, the cop would not like to shoot. On the contrary, the suspect's preferences depend on which type of player they are: guilty or innocent. The guilty suspect prefers to shoot no matter what, and the innocent suspect prefers to not shoot no matter what. The primary complexity of this example is that the sheriff does not know whether the suspect is innocent or guilty. To start, we can see the payoff matrices for each scenario (innocent suspect or guilty suspect).

Table 5. Payoff matrix for the guilty suspect.

Suspect	Sheriff		
		Shoot	*Don't Shoot*
	Shoot	0, 0	5, -3

	Don't Shoot	-6, 2	-3, 4

Table 6. Payoff matrix for the innocent suspect.

Suspect	Sheriff		
		Shoot	*Don't Shoot*
	Shoot	-3, -5	-2, -4
	Don't Shoot	-2, -7	0, 0

Similar to the last example, the sheriff's ideal strategy is not immediately clear based on the payoff matrices, so we have to account for the preferences of the suspect—the innocent suspect will not shoot, but the guilty suspect will. We can therefore modify the payoff matrices.

Table 7. Payoff matrix for the guilty suspect, some options crossed out.

Suspect	Sheriff		
		Shoot	*Don't Shoot*
	Shoot	0, 0	5, -3
	~~*Don't Shoot*~~	~~-6, 2~~	~~-3, 4~~

Table 8. Payoff matrix for the innocent suspect, some options crossed out.

Suspect	Sheriff		
		Shoot	*Don't Shoot*
	~~*Shoot*~~	~~-3, -5~~	~~-2, -4~~
	Don't Shoot	-2, -7	0, 0

Using the same method as the previous example, we can determine the sheriff's ideal strategy using probabilities. Let p represent the probability that the suspect is guilty, and let $1-p$ represent the probability that the suspect is innocent.

Sheriff's expected payoff for shooting:

$$0p - 7(1 - p) = 7p - 7$$

Sheriff's expected payoff for not shooting:

$$-3p + 0(1 - p) = -3p$$

Again, we can set these payoffs equal to each other to determine the sheriff's ideal strategy:

$$7p - 7 = -3p$$
$$10p = 7$$
$$p = 7/10$$

If $p > 7/10$, the sheriff should shoot; if $p < 7/10$, Player 2 should not shoot, and if $p = 7/10$, the sheriff could make either choice. In other words, if there is more than a 70% chance the suspect is guilty, it is in the sheriff's best interest to shoot. If there is less than a 70% chance the suspect is guilty, the sheriff should not shoot. If there is a 70% chance the suspect is guilty, neither choice is better.

Ethical Considerations

As mentioned above, the Sheriff's Dilemma example concludes that logically, the sheriff should choose to shoot the suspect if there is more than a 70% chance the suspect is guilty. However, in real life, the sheriff would have a heavy personal bias guided by their moral compass. Ethically, a sheriff may not want to shoot a suspect if there is even a 1% chance the suspect is innocent. In addition, killing someone without complete certainty that they are guilty could leave lasting guilt and a desire to know the truth. These factors represent an aspect of decision-making not reflected by the conclusions from the payouts above.

From another perspective, a sheriff may feel a compulsive need to protect themselves if there is a chance the suspect could shoot. Rather than take the risk of getting shot, a sheriff may choose to shoot because of their self-defensive instinct.

These unmeasurable biases could dramatically affect the probability of the sheriff making certain decisions. While studying these scenarios, it is important to recognize these ethical considerations.

Conclusion

From this study, we can draw conclusions about the applications of Bayesian games. We can use game theory analysis to determine probabilities based on payoff matrices, adjust the payoffs for given situations or parameters, and use the results to determine the ideal plan of action for players in different games. We can also recognize the ethical considerations of scenarios such as the Sheriff's Dilemma and the impact of internal biases.

This application of analytical thinking can extend beyond games into complexities of business, economics, law, political science, sociology, and more. Statisticians use principles of game theory to make predictions, economists to analyze financial decisions, and sociologists to draw conclusions about human behavioral tendencies.

Limitations

The primary limitation of our study of Bayesian game theory is its oversimplification of reality. To most accurately simulate a real situation, game theorists would need to create extremely complex scenarios with many more players and outcomes. However, our models provide valuable insights into the basics of Bayesian game theory and ways to study its principles.

Acknowledgments

I would like to thank my advisor for the valuable insight provided to me on this topic.

References

Chen, J., Lu, S.-I., & Vekhter, D. (n.d.). Von Neumann and the development of game theory. Stanford University Computer Science. https://cs.stanford.edu/people/eroberts/courses/soco/projects/1998-99/game-theory/neumann.html

Romanowski, P. (2014). Game theory. In K. L. Lerner & B. W. Lerner (Eds.), The Gale Encyclopedia of Science (5th ed.). Gale. https://link.gale.com/apps/doc/CV2644030955/SCIC?u=sacramentoctyds&sid=bookmark-SCIC&xid=38f7bda5

Ross, D. (2023, September 3). Game theory. The Stanford Encyclopedia of Philosophy. https://plato.stanford.edu/entries/game-theory/

Zamir, S. (2009). Bayesian games: Games with incomplete information. Encyclopedia of Complexity and Systems Science, 426-441. https://doi.org/10.1007/978-0-387-30440-3_29