

Causal Games and Causal Nash Equilibrium

Mauricio Gonzalez Soto¹, Luis E. Sucar¹, Hugo Jair Escalante^{1,2}

¹ Instituto Nacional de Astrofísica Óptica y Electrónica,
Coordinación de Ciencias Computacionales,
Mexico

² Centro de Investigaciones y Estudios Avanzados del IPN,
Departamento de Computación,
Mexico

{mauricio, esucar, hugojair}@inaoep.mx

Abstract. Classical results of Decision Theory, and its extension to a multi-agent setting: Game Theory, operate only at the *associative* level of information; this is, classical decision makers only take into account probabilities of events; we go one step further and consider *causal* information: in this work, we define Causal Decision Problems and extend them to a multi-agent decision problem, which we call a causal game. For such games, we study belief updating in a class of strategic games in which any player's action *causes* some consequence via a causal model, which is unknown by all players; for this reason, the most suitable model is Harsanyi's Bayesian Game. We propose a probability updating for the Bayesian Game in such a way that the knowledge of any player in terms of probabilistic beliefs about the causal model, as well as what is caused by her actions as well as the actions of every other player are taken into account. Based on such probability updating we define a Nash equilibria for Causal Games.

Keywords: causal games, causal Nash equilibrium.

1 Introduction

Causal reasoning is a constant element in our lives as it is human nature to constantly ask *why*. Looking for causes is an everyday task and, in fact, causal reasoning is to be found at the very core of our minds [40, 8]. It has been argued that the brain itself is a causal inference machine which uses *effects* to figure out *causes* in order to actively engage with the world [10, 7, 25].

An important aspect of acting in the world is being able to make decisions under uncertain conditions [8, 25]. In their seminal work [39], von Neumann and Morgenstern answered how to make choices if *rational* preferences are assumed and the decision maker knows the stochastic relation between actions and outcomes: maximize expected utility. If such relation is unknown, then J.L. Savage showed in [33] that a rational decision maker must choose *as if* is maximizing her expected utility with respect to a *subjective* probability distribution.

The previous results of von Neumann-Morgenstern and Savage provide formal criteria for decision making if rationality is assumed and information about the environment is considered only at the *asociative* (i.e., probabilistic) level.

These criteria are the basis for many of the techniques used in Artificial Intelligence; for example, Reinforcement Learning algorithms learn *optimal policies* that satisfy the Bellman Equations [37, 31]; therefore, any action prescribed by an optimal policy achieves the maximum expected utility as shown in [41]. Several studies have considered how human beings use causal information when making decisions with uncertain outcomes. It is known that humans tend to prefer causal information over purely probabilistic data [38]; and, in fact, it is shown in [17] that acting in the world is conceived by human beings as *intervening* on it; Therefore, it does not come up as a surprise that humans are able to learn and use causal relations while making single choices as well as in sequential decisions as shown in [8, 35, 11, 28, 24, 15, 27, 16, 32, 14].

Decision problems faced by a rational agent usually involve the decisions made by other agents as well as other, possibly unknown, factors. As seen in several applications, interactive reasoning is a fundamental aspect of human every-day reasoning and it should be addressed by any intelligent agent as argued in [25]. We consider to be of interest the multi-agent setting for causal decision making; for this reason, we consider the interaction of several rational and causal-aware decision makers whose decisions affect each other.

Game Theory [29] deals with situations in which several *rational* decision makers, or players, *interact* while pursuing some well-defined objective; the case in which decision makers make a choice simultaneously without knowing the choice made by the other players is called a *strategic game*; a well-known strategic game is the famous *prisoners' dilemma* in which two detainees must choose between confessing or remaining silent and both know the consequences of any combination of actions, what is ignored by each player is the decision made by the other.

When players ignore both the actions made by other players as well as the knowledge that made them choose a certain action, is called a game with incomplete information, or a *Bayesian Game* which was introduced by Harsanyi in [18, 19, 20]. In this work we will use the Bayesian Game model in order to study what happens when several decision makers have certain knowledge about an environment which is controlled by some, unknown but fixed, causal mechanism. We will first study one-player games, or decision problems, in which the player's actions *cause* some consequence according to some unknown causal model; for this case, we will provide a rational choice criterion which will serve us to define a Nash equilibrium in Causal Games.

2 Causation and Classical Decision Problems

2.1 Causation

The notion of causation deals with regularities found in a given environment which are stronger than probabilistic (or associative) relations in the sense that

a causal relation allows for evaluating a change in the *consequence* given that a change in the *cause* is performed, while probabilistic relations only capture patterns that appear on observed data. For example, when training only on observed samples (x, y) , a Bayesian Network can be equally trained as $X \rightarrow Y$ or $Y \rightarrow X$, see [2] Appendix A for a theoretical argument.

We adopt here the *manipulationist* interpretation of Causality as expressed by Woodward in [42]. The main paradigm is clearly expressed in [6] as *manipulation of a cause will result in a manipulation of the effect*. Consider the following example from [42]: manually forcing a barometer to go down won't cause a storm, whereas the occurrence of a storm will cause the barometer to go down.

We restrict ourselves to probabilistic causation and adopt the formal definition of Causality given in [36]; i.e., a stochastic relation between events which is *irreflexive*, *antisymmetric* and *transitive*; such formal definition is encompassed by the manipulationist interpretation. Similar descriptions of the manipulationist approach can be found in [21] and [9]. Causal inference tools, such as Pearl's do-calculus, stated in [30], allows to find the effect of an intervention in terms of probabilistic information when certain conditions are met. For what remains, we assume the *causal axioms* found in [36] with the condition known as *causal sufficiency*.

2.2 Classical Decision Theory

Classical decision making consist of a set \mathcal{A} of available options to a rational decision maker, and a family \mathcal{E} of uncertain events which will affect the consequence of the action made by the decision maker; any knowledge by the decision maker of such uncertainties is available only at the associative, or probabilistic, level of information. We now state the formal framework for classical decision making as we will use it in order to build upon the causal version of it:

Definition 1. *An uncertain environment is the tuple $(\Omega, \mathcal{A}, \mathcal{C}, \mathcal{E})$. Where \mathcal{A} is a non-empty set of available actions, \mathcal{C} a set of consequences and \mathcal{E} an algebra of events over Ω .*

When we consider the preferences of some decision maker over the set of consequences of some uncertain environment we have a Decision Problem under Uncertainty:

Definition 2. *A Decision Problem under Uncertainty is an uncertain environment $(\Omega, \mathcal{A}, \mathcal{C}, \mathcal{E})$ plus a preference relation \succeq defined over \mathcal{C} .*

2.3 Causal Environments and Causal Decision Problems

In this section we define a Causal Environment to be an *uncertain environment* for which there exists a Causal Graphical Model (CGM) \mathcal{G} which controls the environment. Details on CGMs can be found in [22].

Definition 3. *A Causal Environment is a tuple $(\Omega, \mathcal{A}, \mathcal{G}, \mathcal{C}, \mathcal{E})$ where $(\Omega, \mathcal{A}, \mathcal{C}, \mathcal{E})$ is an uncertain environment and \mathcal{G} is a CGM such that the set of variables of \mathcal{G} correspond to the uncertain events in \mathcal{E} .*

2.4 Rational Choice in Causal Environments

Consider a decision maker who knows that any action she takes will *cause* a certain action, but she does not explicitly know the form of such causal relation, she only has probabilistic beliefs about such relation. We define in this section a formal framework for studying such situations.

Definition 4. We define a Causal Decision Problem (CDP) as $(\mathcal{A}, \mathcal{G}, \mathcal{E}, \mathcal{C}, \succeq)$ where $(\mathcal{A}, \mathcal{G}, \mathcal{E}, \mathcal{C})$ is a Causal Environment and \succeq is a preference relation.

For the CGM in a given CDP we will distinguish two particular variables: one corresponding to the available actions, and one corresponding to the caused outcome. We are considering that only one variable can be intervened upon and that the values of such variable represent the actions available to the decision maker; i.e., the value forced upon such variable under an intervention represents the action taken by the decision maker. The intuition behind the definition of a Causal Decision Problem is this: a decision maker chooses an action $a \in \mathcal{A}$, which is automatically inputted into the model \mathcal{G} , which outputs the *causal outcome* $c \in \mathcal{C}$. We say a CDP is *finite* if the set \mathcal{A} is finite. We now provide a decision criterion for rationally choosing in a Causal Decision Problem.

Theorem 1. In a finite Causal Decision Problem $(\mathcal{A}, \mathcal{G}, \mathcal{E}, \mathcal{C}, \succeq)$, where \mathcal{G} is a Causal Graphical Model, we have that the preferences \succeq of a decision maker are Savage-rational if and only if there exists a probability distribution P_C over a family \mathcal{F} of causal models such that for $a, b \in \mathcal{A}$:

$$a \succeq b \text{ if and only if:}$$

$$\sum_{c \in \mathcal{C}} u(c) \left(\sum_{g \in \mathcal{F}} P_g(c|do(a)) P_C(g) \right),$$

$$\geq$$

$$\sum_{c \in \mathcal{C}} u(c) \left(\sum_{g \in \mathcal{F}} P_g(c|do(b)) P_C(g) \right),$$

where P_g is the probability distribution associated with the causal model g .

Proof. The decision maker is facing an environment in which any action she takes will stochastically cause an outcome $c \in \mathcal{C}$. For this reason, the decision making is facing a very particular case of decision making under uncertainty. Assuming rationality, we invoke Savage's Theorem [33, 23, 12] to obtain a utility function u^S and a probability measure P^S which satisfy that the preference relation is represented by the expectation of u^S with respect to P^S .

In such a causal environment, the CGM \mathcal{G} contains all of the information which connects actions, uncertain events and outcomes, and noting that we can identify any action a with $\{c_j|E_j : j \in J\}$ where J a countable set of indexes [3] we have that:

$$\mathbb{E}_{P^S}[u(c)] = \sum_{j \in J} u(c_j) P^S(E_j).$$

For each action $a = \{c_j | E_j : j \in J\}$, $P^S(E_j)$ is the probability of *causing* consequence c_j by choosing action a . In order for the decision maker to find the probability of a certain consequence c_j given that an action a is performed then she must have in mind a single causal model g and a way to assign probabilities over a family of causal models; i.e., the uncertainty component $P^S(E_j)$ is formed by two parts: a distribution P_C which represents the degree of belief of the decision maker about a specific model g being the true one, and within g , a distribution P_g used to calculate the probability of causing some consequence c_j given that action a is chosen. Using the Caratheodory Extension Theorem [1] a probability measure P_C whose support is a sufficiently general family of causal models \mathcal{F} can be shown to exist. For $g \in \mathcal{F}$, the decision maker considers g to be the true causal model with probability $P_C(g)$, and within g , we use the classical von Neumann-Morgenstern Theorem in order to obtain the best action (see section 4.1 of [30] for details). Let P_g the probability distribution associated with the causal model g .

Then:

$$\mathbb{E}_{P^S}[u(c)] = \sum_{j \in J} u(c_j) P^S(E_j), \quad (1)$$

$$= \sum_{j \in J} u(c_j) \left(\sum_{g \in \mathcal{F}} P_g(c_j | do(a)) P_C(g) \right). \quad (2)$$

We have shown what is the expected utility for some action $a \in \mathcal{A}$, and by Savage's Theorem the result follows.

2.5 Interpretation

Theorem 1 says that a decision maker who faces a Causal Decision Problem is considering a probability distribution P_C over a family \mathcal{F} and, within each structure, using the term $P_g(c | do(a))$ in order to find the probability of obtaining a certain consequence given that the intervention $do(a)$ is performed; in this way, the optimal action a^* is given by:

$$a^* = \operatorname{argmax}_{a \in \mathcal{A}} \sum_{c \in \mathcal{C}} u(c) \left(\sum_{g \in \mathcal{F}} P_g(c | do(a)) P_C(g) \right). \quad (3)$$

We note that a^* is obtained by taking into account the utility obtained by every possible consequences weighted using both the probability of causing such action within a specific causal model g and the probability that the decision maker assign to such $g \in \mathcal{F}$.

We are considering a *normative* interpretation for Theorem 1 according to which a decision maker must use any causal information in order to obtain the best possible action.

Such action must be obtained by considering the *beliefs* of the decision maker about the causal relations that hold in her environment (the distribution P_C),

how such relations could produce the best action when considered *as if* they were true (distribution P_g), and the satisfaction (utility u) produced by the consequences of actions [22].

3 Classical Strategic Games

A *strategic game* is a model of a situation in which several players must take an action and afterwards they will be affected both by the outcome of their own action as well as the actions of the other players. In a strategic game it is assumed that no player knows the action taken by any other players; this is:

Definition 5. *a strategic game ([29]) consists of:*

- A finite set N of n players.
- For each player, a nonempty set A_i of available actions.
- For each player, a preference relation \succeq_i defined over $A = A_1 \times \dots \times A_n$.

Definition 6. *A Nash equilibrium of a strategic game $G = (N, (A)_{i \in N}, (\succeq_i)_{i \in N})$ is a vector of strategies $a^* = (a_1, a_2, \dots, a_n)$ such that:*

$$(a_{-i}^*, a_i^*) \succeq_i (a_{-i}^*, b_i) \text{ for all } b_i \in A_i,$$

where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$.

This is, in a Nash equilibrium no player can find a better action given the actions taken by the rest of the players. We adopt here the *deductive* interpretation of an equilibrium, according to which an equilibrium results from rationality principles [4, 5].

4 Causal Games

In this section, we define a *causal strategic game* as a strategic game within a causal environment; this is, consider a *strategic game* between N rational players who are situated in a causal environment. We assume that it is *common knowledge* the causal nature of the environment as well as the rationality assumption for each player. We also assume that the causal mechanism, which represented by a Causal Graphical Model \mathcal{G} , remains fixed and it is unknown for each player. In this game, players ignore the actions taken by any other player, and since the causal model which controls the environment is unknown by every player, then players also ignore the information that players will use in order to take their respective actions: for this reason, we will work within the framework of *bayesian games*.

4.1 Bayesian Games

Strategic games are games in which no player knows the action taken by the other players; we now consider a type of game in which no player knows both the actions taken by any other player, nor the *private information* that made each player to take any action. Such model is called a *Bayesian Game*, introduced in [18, 19, 20]

Definition 7. A Bayesian strategic game ([29]), consists of:

- A finite set N of players.
- A finite set Ω of states of nature.
- For each player, a nonempty set A_i of actions.
- For each player, a finite set T_i and a function $\tau_i : \Omega \mapsto T_i$ the signal function of the player.
- For each player, a probability measure p_i over Ω such that $p_i(\tau_i^{-1}(t_i)) > 0$ for all $t_i \in T_i$.
- A preference relation \succeq_i defined over the set of probability measures over $A \times \Omega$ where $A = A_1 \times \dots \times A_n$.

4.2 Bayesian Causal Games

In this section, we consider a *strategic game* between N rational players who are situated in a causal environment. A game is a model of a situation in which several players must take an action and afterwards they will be affected both by the outcome of their own action as well as the actions of the other players. In a strategic game it is assumed that no player knows the action taken by any other players; we also assume that the causal mechanism, which represented by a Causal Graphical Model \mathcal{G} , remains fixed and it is unknown for each player.

In this game, players ignore the actions taken by any other player, and since the causal model which controls the environment is unknown by every player, then players also ignore the information that players will use in order to take their respective actions: strategic games of this type are called *Bayesian Game*, introduced in [18, 19, 20]. In the games we will consider, the uncertainty of every player consists of two levels: on a first level, the true causal model \mathcal{G} ; on a second level, what an action $do(a)$ causes if a certain CGM ω is considered to be the causal model.

We will consider the set Ω to be a family of possible causal models; in this way, $\omega \in \Omega$ being the true state of nature fixes a causal model which controls the environment in which the players make their choices. In classical Bayesian games, once $\omega \in \Omega$ is realized as the true state, then each player receives a signal $t_i = \tau_i(\omega)$ and the posterior belief $p_i(\omega | \tau_i^{-1}(t_i))$ given by $p_i(\omega) / p_i(\tau_i^{-1}(t_i))$ if $\omega \in \tau_i^{-1}(t_i)$. In the case for causal bayesian games, we must consider both the probability p_i of ω being the true state as well as the probability p_i^a of observing a certain consequence when doing some action a_i if ω is the true model.

Following [29], we define a new game G^* in which its players are all of the possible combinations $(i, t_i) \in N \times T_i$, where the possible actions for (i, T_i) is A_i .

We see that fixing a player $i \in N$, the posterior probability $p(\omega|\tau_i^{-1}(t_i))$ induces a lottery over the pairs $(a^*(j, \tau_j(\omega)), \omega)$ for some other $j \in N$. This lottery assigns to $(a^*(j, \tau_j(\omega)), \omega)$ the probability $p_i(\omega)/p_i(\tau_i^{-1}(t_i))$ if $\omega \in \tau_i^{-1}(t_i)$. The classical Bayesian game will simply call a Nash equilibrium for the game G^* a Nash equilibrium of the original game; but we have the second level of uncertainty: the consequences caused by some action a through a causal model ω . We notice that the posterior probability itself induces a probability distribution defined over *actions* for each player once a *desired consequence* is fixed, this distribution, according to Theorem 1 is given by $p_i^\omega(c|do(a_i^*), a_{-i}^*)p_i(\omega|\tau_i^{-1}(t_i))$. This motivates the following definition of a *Causal Nash equilibrium*.

4.3 Causal Nash Equilibrium

For each player $i \in N$ in the strategic game, we define the following probability distribution over consequences:

$$p_i^a(c) = p_i^\omega(c|do(a_i), a_{-i})p_i(\omega) \text{ for } a \in A = A_1 \times \dots \times A_N, \quad (4)$$

where p_i^ω is the probability of causing a certain consequence within a causal model ω and p_i are the player's *posterior beliefs* about the causal model that controls the environment, and $do()$ is the well known intervention operator from [30]. We now define:

$$u_i^C(a) = \sum_{c \in C} u_i(c)p_i^a(c) \text{ for } a \in A = A_1 \times \dots \times A_N. \quad (5)$$

Notice that u_i^C evaluates an action profile $a \in A$ in terms of: The knowledge about the causal model of each player represented by p_i , which allows each player to evaluate the probability of causing outcomes in terms of actions by using the *do* operator as well as the other actions taken by the other players, given by a_{-i} and the preferences of each player u_i . Using this new function, we define the equilibrium for a strategic game with causal information and Bayesian players as:

Definition 8. *An action profile $a^* \in A$ is a Nash equilibrium for this causal strategic game if and only if:*

$$u_i^C(a^*) \geq u_i^C(a_i, a_{-i}^*) \text{ for any other } a_i \in A_i. \quad (6)$$

This is, an action profile is a Nash equilibrium if and only if each player uses her current knowledge about the causal model of the environment in order to (causally) produce the best possible outcome given the actions taken by the other players. The existence of the Causal Nash Equilibrium is guaranteed if every A_i is a nonempty compact convex set in some \mathbb{R}^n and if the preference relation induced by u_i^C is continuous and quasi-concave.

5 Conclusion

We have studied Decision Making under uncertainty in the case where a Causal Graphical Model is responsible for producing an outcome given an action (intervention) of the decision maker. We have provided a rational decision making criterion for the case in which the decision maker does not know the causal model, but has probabilistic beliefs about possible models.

Using our decision making result, and taking as a basis Harsanyi's model of a Bayesian Game in which every player has incomplete information about both the actions taken by other players as well as the information that made each player take his action we have been able to provide a definition of a Causal Nash Equilibrium in which every player is aware that there exists a Causal Mechanism that will produce some consequence once he takes an action.

Our decision making result (i.e., Theorem 1), besides motivating the Causal Nash Equilibrium, also provides an optimality criterion for learning algorithms in causal settings such as those presented in [26, 34, 13]. Our definition of Causal Equilibrium takes into account classical game theory through the incorporation of the classical von Neumann-Morgenstern utility function as well as the fundamental notion in Causation of Pearl's *do* operator.

We hope this work contributes to recent efforts of giving Causation its well deserved place in Artificial Intelligence as well as motivating further research in computational aspects of Causal Decision Theory.

References

1. Ash, R.B., Doleans-Dade, C.A.: Probability and Measure Theory. Academic Press, 2nd edn. (2000)
2. Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., Pal, C.: A meta-transfer objective for learning to disentangle causal mechanisms (2019)
3. Bernardo, J.M., Smith, A.F.M.: Bayesian theory. Wiley Series in Probability and Statistics. (2000)
4. Binmore, K.: Modeling rational players: Part i. *Economics & Philosophy* 3(2), 179–214 (1987)
5. Binmore, K.: Modeling rational players: Part ii. *Economics & Philosophy* 4(1), 9–55 (1988)
6. Campbell, D.T., Cook, T.D.: Quasi-experimentation: Design & analysis issues for field settings. Rand McNally College Publishing Company Chicago (1979)
7. Clark, A.: Surfing uncertainty: Prediction, action, and the embodied mind. Oxford University Press (2015)
8. Danks, D.: Unifying the mind: Cognitive representations as graphical models. MIT Press (2014)
9. Freedman, D.: From association to causation via regression. *Advances in applied mathematics* 18(1), 59–110 (1997)
10. Friston, K.: The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience* 11(2), 127–138 (2010)
11. Garcia-Retamero, R., Hoffrage, U.: How causal knowledge simplifies decision-making. *Minds and Machines* 16(3), 365–380 (2006)

12. Gilboa, I.: Theory of Decision under Uncertainty. Cambridge University Press (2009)
13. Gonzalez-Soto, M., Sucar, L.E., Escalante, H.J.: Playing against nature: causal discovery for decision making under uncertainty. In: Machine Learning for Causal Inference, Counterfactual Prediction and Autonomous Action (CausalML) Workshop at ICML 2018 (2018), <https://arxiv.org/abs/1807.01268>
14. Hagmayer, Y., Fernbach, P.: Causality in decision-making. In: Waldmann, M.R. (ed.) The Oxford Handbook of Causal Reasoning. Oxford University Press (2017)
15. Hagmayer, Y., Meder, B.: Causal learning through repeated decision making. In: Proceedings of the Annual Meeting of the Cognitive Science Society. vol. 30 (2008)
16. Hagmayer, Y., Meder, B.: Repeated causal decision making. Journal of Experimental Psychology: Learning, Memory, and Cognition 39(1), 33 (2013)
17. Hagmayer, Y., Sloman, S.A.: Decision makers conceive of their choices as interventions. Journal of Experimental Psychology: General 138(1), 22 (2009)
18. Harsanyi, J.C.: Games with incomplete information played by “Bayesian” players, I–III Part I. the basic model. Management science 14(3), 159–182 (1967)
19. Harsanyi, J.C.: Games with incomplete information played by “Bayesian” players part II. Bayesian equilibrium points. Management Science 14(5), 320–334 (1968)
20. Harsanyi, J.C.: Games with incomplete information played by ‘Bayesian’ players, part III. the basic probability distribution of the game. Management Science 14(7), 486–502 (1968)
21. Holland, P.W.: Statistics and causal inference. Journal of the American Statistical Association 81(396), 945–960 (1986)
22. Koller, D., Friedman, N.: Probabilistic graphical models: principles and techniques. MIT press (2009)
23. Kreps, D.M.: Notes On The Theory Of Choice. Routledge New York (1988)
24. Lagnado, D.A., Waldmann, M.R., Hagmayer, Y., Sloman, S.A.: Beyond covariation. Causal learning: Psychology, philosophy, and computation pp. 154–172 (2007)
25. Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J.: Building machines that learn and think like people. Behavioral and Brain Sciences 40 (2017)
26. Lattimore, F., Lattimore, T., Reid, M.D.: Causal bandits: Learning good interventions via causal inference. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29, pp. 1181–1189. Curran Associates, Inc. (2016)
27. Meder, B., Gerstenberg, T., Hagmayer, Y., Waldmann, M.R.: Observing and intervening: Rational and heuristic models of causal decision making. The Open Psychology Journal 3, 119–135 (2010)
28. Nichols, W., Danks, D.: Decision making using learned causal structures. In: Proceedings of the Annual Meeting of the Cognitive Science Society. vol. 29 (2007)
29. Osborne, M.J., Rubinstein, A.: A course in game theory. MIT press (1994)
30. Pearl, J.: Causality: Models, Reasoning and Inference. Cambridge University Press, New York, NY, USA, 2nd edn. (2009)
31. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley & Sons, Inc., New York, NY, USA, 1st edn. (1994)
32. Rottman, B.M., Hastie, R.: Reasoning about causal relationships: Inferences on causal networks. Psychological bulletin 140(1), 109 (2014)
33. Savage, L.: The Foundations of Statistics. New York: John Wiley & Sons (1954)
34. Sen, R., Shanmugam, K., Dimakis, A.G., Shakkottai, S.: Identifying best interventions through online importance sampling. In: International Conference on Machine Learning. pp. 3057–3066 (2017)

35. Sloman, S.A., Hagmayer, Y.: The causal psycho-logic of choice. *Trends in Cognitive Sciences* 10(9), 407–412 (2006)
36. Spirtes, P., Glymour, C.N., Scheines, R.: *Causation, prediction and search*. MIT Press (2000)
37. Sutton, R.S., Barto, A.G.: *Reinforcement Learning: An introduction*. MIT Press (1998)
38. Tversky, A., Kahneman, D.: Causal schemas in judgments under uncertainty. *Progress in social psychology* 1, 49–72 (1980)
39. Von Neumann, J., Morgenstern, O.: *Theory of games and economic behavior*. Princeton University Press (1944)
40. Waldmann, M.R., Hagmayer, Y.: Causal reasoning. In: Reisberg, D. (ed.) *The Oxford Handbook of Cognitive Psychology*. Oxford University Press (2013)
41. Webb, J.N.: *Game theory: Decisions, Interaction and Evolution*. Springer Undergraduate Mathematics Series (2007)
42. Woodward, J.: *Making things happen: A theory of causal explanation*. Oxford Studies in Philosophy of Science, Oxford University Press (2003)