

Causal Interpretation of Graphical Models

Glenn Shafer

Rutgers Business School - Newark and New Brunswick

180 University Avenue

Newark, New Jersey 07102 USA

gshafer@andromeda.rutgers.edu

How do we combine probability with causality? What does it mean for a graphical probability model to be causal? I can hardly review here the many ways statisticians (even those participating in this session) have addressed these questions over the past decade. I will, however, hazard one generalization: In order to answer these questions clearly, you must first adopt some interpretation of probability.

Those who see probability as objective typically combine probability with causality by thinking of a human decision as an exogenous intervention that suspends the stochastic action represented by a probability model. Those who see probability as subjective typically take an equally decision-theoretic point of view, but they see causality in the rule that determines a consequence from a decision and a state of nature, while they see subjectivity in the probabilities for the states of nature (Heckerman and Shachter 1965, Pearl 2000).

Today I propose to bring something new into this discussion: the game-theoretic framework for probability advanced in my recent book with Vladimir Vovk (Shafer and Vovk 2001). In this framework, probability is interpreted by adopting some form of Cournot's principle, which says that reality will not allow a gambler to get rich betting at given odds. What can it mean, under this interpretation, for a graphical probability model to be causal? I shall offer a few brief answers to this question, but my main purpose is to pose it clearly.

The game-theoretic framework for probability

It is hardly novel to say that mathematical probability derives from intuitions about gambling. But it was only in Vovk's work in the mid 1990s that this insight was first made into a rigorous mathematical foundation for probability theory. Inspired by Lindeberg (1922) and Ville (1939), Vovk realized that the classical limit theorems can be formulated and proven in a purely game-theoretic framework, with no measure-theoretic assumptions.

To see how this goes, consider the simplest and most central example in probability—coin tossing. This involves the following protocol:

$\mathcal{K}_0 := 1$.

For $n = 1, 2, \dots$:

House announces $p_n \in [0, 1]$.

Gambler announces $\alpha_n \in \mathbb{R}$.

Reality announces $x_n \in \{0, 1\}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + \alpha_n(x_n - p_n)$.

The players move in the order indicated (not simultaneously), and each player sees the other players' moves as they are made. Reality's move on the n th round is either heads ($x_n = 1$) or tails ($x_n = 0$). House's move p_n is the probability for heads. Gambler's move α_n represents the size of his bet (on heads if $\alpha_n > 0$, on tails if $\alpha_n < 0$). Gambler starts with \$1, and \mathcal{K}_n is his capital after n steps.

To define a game, we need a rule for who wins. Here it is: Gambler wins if both (i) his capital never becomes negative and (ii) either $\lim_{n \rightarrow \infty} \mathcal{K}_n = \infty$ or

$$(1) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (x_i - p_i) = 0.$$

Otherwise, House and Reality win.

Here is Vovk's formulation of the strong law of large numbers for coin tossing: *Gambler has a winning strategy in this game*. Say it this way: Gambler can become infinitely rich unless Reality chooses his x_n to make (1) hold. If we assume that Reality will not allow Gambler to become infinitely rich, we may conclude that (1) will hold. Roughly speaking, the proportion of heads $\frac{1}{n} \sum_{i=1}^n x_i$ will converge to the average probability $\frac{1}{n} \sum_{i=1}^n p_i$.

At first glance, most readers will find the game-theoretic formulation more amusing than important. But as they reflect, they may come to acknowledge it as fundamental. The established measure-theoretic formulation is prized for its mathematical power rather than for its intuitive basis. The game-theoretic formulation is more intuitive (it does not rely on unmotivated assumptions like the axiom of countable additivity), and in some respects it is even more powerful mathematically. For example:

- The game-theoretic formulation does not require that the odds offered by House define a full probability distribution for x_n . We may restrict Gambler by requiring $\alpha_n \geq 0$, making House's odds one-sided, and this leads to an entirely new variant on the central-limit theorem (§6.3 of Shafer and Vovk 2001).
- Even when Gambler is allowed to choose any value—positive, negative, or zero—for α_n , House is not providing a joint probability distribution for Reality's moves x_1, x_2, \dots . House's move p_n plays the role of a conditional probability for x_n given x_1, \dots, x_{n-1} , but this says nothing about how House might have moved had Reality chosen earlier x_i differently. In this respect, the game-theoretic formulation is prequential (Dawid 1984).
- Because of its prequential nature, the game-theoretic framework is straightforwardly compatible with elementary quantum mechanics (§8.4 of Shafer and Vovk 2001).
- The measure-theoretic versions of the classical limit theorems can be deduced easily from the game-theoretic versions, but it is not so easy to go in the other direction (§8.1 of Shafer and Vovk 2001).

Cournot's principle

Vovk and I call the assumption that Reality will not allow Gambler to become infinitely rich *Cournot's principle*. This is not a mathematical assumption. It is not a rule of the game. It is an interpretative assumption—a way of relating the abstract mathematical game to the world. (I hasten to acknowledge the finitary nature of real experience, but for lack of space I leave aside the finitary version of Cournot's principle.)

In any particular application of Cournot's principle, we must fill in many details. Who, if anyone, plays House and Gambler? Who believes that Cournot's principle holds? The many different answers to these questions lead to many different interpretations of probability, including some that are decidedly subjective and others that are more objective. We would call given probabilities objective only if they appear to satisfy Cournot's principle in practice. But any individual who believes that given probabilities do not permit a bankruptcy-free strategy to make him very rich—whether he made up the probabilities or obtained them from a theory or another person—may reasonably call them his personal subjective probabilities. Bear in mind, though, that he is not merely saying that he is uncertain how to get very rich betting at these probabilities. We require him to be convinced that no bankruptcy-free strategy that uses only information available to him will make him very rich. He thinks he will at best more or less break even.

I recently studied this understanding of subjective probability in collaboration with two colleagues (Shafer, Gillett, and Scherl 2003). One thing we learned is that Cournot's principle

leads to a new justification for the rule of conditional probability—the rule that one’s new probability for A after learning B and nothing else should be the ratio $P(A \cap B)/P(B)$. The established argument for this rule, which shows that House must obey it in order to avoid sure loss, depends on the assumption that House is required to state in advance what his new probability for A will be, as a function of what is learned about B . The argument based on Cournot’s principle takes the viewpoint of Gambler and does not require House to state any such rule in advance. We observe merely that Gambler can use House’s initial gambling offers ($P(A \cap B)$ and $P(B)$) to accomplish in advance anything that he can accomplish later using the offer $P(A \cap B)/P(B)$ and knowledge that B has happened. This implies that if the initial offers satisfy Cournot’s principle (they do not permit Gambler to get rich), then the later offer of the conditional probability will satisfy it as well. This generalizes to the case where Gambler’s offers do not define full probability distributions for Reality’s moves; we obtain Walley’s updating principle (Walley 1991) when B is our only new information and a more general updating principle when this is not guaranteed.

Using Cournot’s principle to talk about causality

For probabilities to be causal, they must first of all be objective: experience should confirm that they obey Cournot’s principle. But does this mean only that Reality will not allow a particular Gambler, who observes particular things, to beat the probabilities? Or does it mean that there is no information available to anyone that would enable them to beat the probabilities? In the spirit of my 1996 book, I favor the second formulation.¹

We consider probabilities calculated using quantum mechanics causal in this sense. But we can also conjecture the existence of causal probabilities without being able to calculate them. House, Gambler, and Reality are at play, but we see only a few of their moves. In this spirit, we can hypothesize graphical models without specifying the probabilities in them—a procedure that seems incoherent if we rely on a purely subjective interpretation of probability.


To get a quick sense of what more one might say, let us consider the three figures on the next page, which I have borrowed from Dawid (2002).

- First the familiar probabilistic DAG. The nodes represent variables, and probabilities are provided for each node given its parents. The simplest game-theoretic interpretation, (tracking in Shafer 1996), says that Gambler cannot get rich by gambling on each variable at the point where its parents are all determined, even with all information available to science. The more subtle interpretations studied in Shafer 1996, such as linear sign, can also be elaborated game-theoretically. If we know all the probabilities (even though we do not see Reality’s moves), then the subjective use of probabilities conditional on observations can be justified as in the preceding section. But we can think about testing the causal model empirically even if we are not given the probabilities.
- In an influence diagram, square nodes represent decisions, for which probabilities are not provided. Because the game-theoretic approach does not require full probability distributions, the absence of probabilities for these nodes does not affect the causal interpretation. Nor does it matter whether they are absent because the variables are determined by human decision or for some other reason; human intervention is not so central in this approach. This suggests some directions for new research: (1) How are subjective beliefs constrained in this case for those of us who do not see all of Reality’s moves? (2) Is


¹This must be understood within reason, relative to the science of one’s era. What is unpredictable now may be predictable at some time in the future (§1.2 of Shafer 1996). We may also change our minds about what is causal. But when we talk about causality, we are typically talking about some idealized limit of prediction, which usually goes beyond our actual ability to predict with the information we have at hand.

it useful to consider game-theoretical causal graphical models in which Gambler tries to predict the square nodes, subject to specified penalties for the magnitude of his error, even though he is not given probabilities? Here Cournot's principle would say that he cannot do better, asymptotically, using information other than the parent variables.


- The third figure augments the DAG by adding decision nodes to indicate the possibility of an exogenous intervention fixing the value of its child regardless of the values of the other parents. Again, these decision nodes do not hamper a game-theoretic interpretation, but the interpretation does not require us to pretend that such interventions are possible for every variable.



DAG describing probabilistic structure



Influence diagram



Augmented DAG

Three figures from Dawid (2002)

REFERENCES

- Alexander Philip Dawid. Statistical theory: The prequential approach. *Journal of the Royal Statistical Society, Series A* 147:278–292, 1984.
- Alexander Philip Dawid. Influence diagrams for causal modelling and inference. *International Statistical Review* 70:161–189 (corrigenda p. 437), 2002.
- David Heckerman and Ross Shachter. Decision-theoretic foundations for causal reasoning. *Journal of Artificial Intelligence Research* 3:405–430, 1995.
- Jarl Waldemar Lindeberg. Eine neue Herleitung des Exponentialgesetzes in der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift* 15:211–225, 1922.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.
- Glenn Shafer. *The Art of Causal Conjecture*. MIT Press, Cambridge, Massachusetts, 1996. See www.glennshafer.com for excerpts.
- Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001. See www.probabilityandfinance.com for excerpts.
- Glenn Shafer, Peter R. Gillett and Richard B. Scherl. A new understanding of subjective probability and its generalization to lower and upper prevision. *International Journal of Approximate Reasoning* 33:1–49, 2003.
- Jean Ville. *Étude critique de la notion de collectif*. Gauthier-Villars, Paris, 1939.
- Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.

RÉSUMÉ

Shafer et Vovk ont fondé la probabilité sur la théorie des jeux. Dans ce cadre, on donne une signification empirique aux probabilités en privilégiant une forme du principe de Cournot, qui assure que la réalité ne permet pas qu'un joueur gagne outre mesure. On se demande comment ce principe de Cournot se précise quand il s'agit d'une interprétation causale des modèles graphiques.