2.2 The Causal Modeling Framework

We view the task of causal modeling as an induction game that scientists play against Nature. Nature possesses stable causal mechanisms that, on a detailed level of descriptions, are deterministic functional relationships between variables, some of which are unobservable. These mechanisms are organized in the form of an acyclic structure, which the scientist attempts to identify from the available observations.

Definition 2.2.1 (Causal Structure)
A causal structure of a set of variables V is a directed acyclic graph (DAG) *in which each node corresponds to a distinct element of V, and each link represents direct functional relationship among the corresponding variables.*

A causal structure serves as a blueprint for forming a "causal model"—a precise specification of *how* each variable is influenced by its parents in the DAG, as in the structural equation model of (1.40). Here we assume that Nature is at liberty to impose arbitrary functional relationships between each effect and its causes and then to perturb these relationships by introducing arbitrary (yet mutually independent) disturbances. These disturbances reflect "hidden" or unmeasurable conditions and exceptions that Nature chooses to govern by some undisclosed probability function.

Definition 2.2.2 (Causal Model)
A causal model is a pair $M = <D, \Theta_D>$ consisting of a causal structure D and a set of parameters Θ_D compatible with D. The parameters Θ_D assign a function $x_i = f_i(pa_i, u_i)$ to each $X_i \in V$ and a probability measure $P(u_i)$ to each u_i, where PA_i are the parents of X_i in D and where each U_i is a random disturbance distributed according to $P(u_i)$, independently of all other u.

As we have seen in Chapter 1 (Theorem 1.4.1), the assumption of independent disturbances renders the model *Markovian* in the sense that each variable is independent of all its nondescendants, conditional on its parents. This Markov assumption is more a convention than an assumption; for it merely defines the granularity of the models we wish

to consider as candidates before we begin the search. We can start in the deterministic extreme, where all variables are explicated in microscopic details and where the Markov condition certainly holds. As we move up to macroscopic abstractions by aggregating variables and introducing probabilities to summarize omitted variables, we need to decide at what stage the abstraction has gone too far and where useful properties of causation are lost. Evidently, the Markov condition has been recognized by our ancestors (the authors of our causal thoughts) as a property worth protecting in this abstraction; correlations that are not explained by common causes are considered *spurious*, and models containing such correlations are considered *incomplete*. The Markov condition guides us in deciding when a set of parents PA_i is considered complete in the sense that it include *all* the relevant immediate causes of variable X_i. It permits us to leave some of these causes out of PA_i (and be summarized by probabilities), but not if they also affect other variables modeled in the system. If a set PA_i in a model is too narrow, there will be disturbance terms that influence several variables simultaneously and the Markov property will be lost. Such disturbances will be treated explicitly as "latent" variables (see Definition 2.3.2). Once we acknowledge the existence of latent variables and represent their existence explicitly as nodes in a graph, the Markov property is restored.

Once a causal model M is formed, it defines a joint probability distribution $P(M)$ over the variables in the system. This distribution reflects some features of the causal structure (e.g., each variable must be independent of its grandparents, given the values of its parents). Nature then permits the scientist to inspect a select subset $O \subseteq V$ of "observed" variables and to ask questions about $P_{[O]}$, the probability distribution over the observables, but it hides the underlying causal model as well as the causal structure. We investigate the feasibility of recovering the topology D of the DAG, from features of the probability distribution $P_{[O]}$.[2]

[2]This formulation invokes several idealizations of the actual task of scientific discovery. It assumes, for example, that the scientist obtains the distribution directly, rather than events sampled from the distribution. Additionally, we assume that the observed variables actually appear in the original causal model and are not some aggregate thereof. Aggregation might result in feedback loops, which we do not