

Homogeneity, selection, and the faithfulness condition

Daniel Steel

Received: 13 July 2006 / Accepted: 13 July 2006 /
Published online: 17 August 2006
© Springer Science+Business Media B.V. 2006

Abstract The faithfulness condition (FC) is a useful principle for inferring causal structure from statistical data. The usual motivation for the FC appeals to theorems showing that exceptions to it have probability zero, provided that some apparently reasonable assumptions obtain. However, some have objected that, the theorems notwithstanding, exceptions to the FC are probable in commonly occurring circumstances. I argue that exceptions to the FC are probable in the circumstances specified by this objection only given the presence of a condition that I label homogeneity, and furthermore that this condition typically does not obtain in the FC's intended domain of application.

Keywords Causal inference · Directed acyclic graphs · Faithfulness condition

Introduction

Directed acyclic graphs (DAGs) have proven to be an extremely fruitful system of representation in work on causal inference from statistical data (cf. Glymour and Cooper 1999; Neapolitan 2004; Pearl 2000; Spirtes, Glymour, & Scheines 2000). DAGs consist of arrows linking nodes: the arrows represent the relationship of unmediated causation, while the nodes correspond to random variables. In this framework, causal inference relies on principles that specify which probabilistic dependence and independence relationships would obtain among these random variables if the DAG were an accurate representation of the causal structure. The most fundamental rule for deriving predictions about probabilistic independence from DAGs is the Causal Markov Condition (CMC). This states that every variable in the DAG is probabilistically independent of its non-descendants (intuitively, non-effects) conditional on its parents (intuitively, direct causes). The CMC, however,

D. Steel (✉)
Department of Philosophy, Michigan State University, 503 S. Kedzie Hall, East Lansing
MI 48824, USA
e-mail: steel@msu.edu

only provides information concerning what probabilistic *independencies* obtain—it would be trivially satisfied in any distribution in which there were no probabilistic dependencies whatever. Clearly, causal inference from statistical data requires that causal hypotheses make predictions not only about probabilistic independence but about *dependence* as well. The faithfulness condition (FC) is an important and frequently used rule for deriving predictions about probabilistic dependence from DAGs: the FC says that the *only* probabilistic independencies generated by acyclic causal structures are those entailed by the CMC. For example, the FC requires that causes and effects are probabilistically dependent, as are joint effects of a common cause. The FC also makes claims about conditional independence, for instance, in the chain $X \rightarrow Y \rightarrow Z$, Y and Z are probabilistically dependent conditional on X .

The motivation for the FC rests on theorems showing that, given certain apparently plausible assumptions, exceptions to the principle require very finely-tuned, extremely improbable parameterizations (cf. Meek 1995; Spirtes, Glymour, & Scheines 1993, 2000, pp. 41–42). However, some have objected that, such theorems notwithstanding, there are commonly occurring circumstances in which exceptions to the FC are likely (Cartwright 1999a, pp. 117–118; 1999b, pp. 16–17; Hoover 2001, pp. 168–170). This objection could be interpreted as saying either that, in the specified circumstances, *strict* exceptions to the FC are probable or only that *very near* exceptions are. A strict exception is a case in which a probabilistic dependence required by the FC is absent. A very near exception is a case in which a probabilistic dependence required by the FC is present, but the dependence is so slight that it would be almost impossible to detect in any practical sample size. In any situation in which strict exceptions to the FC are probable, one or more of the assumptions of the aforementioned theorems must fail to obtain. In a situation in which very near exceptions are probable, one would naturally expect that one or more assumptions of the theorem are nearly false.

I suggest that the pertinent assumption is that subsets of the parameter space of Lebesgue measure zero also have probability zero, a proposition I shall refer to as **L**. I argue that *strict* exceptions to **L** are extremely unlikely in the intended domain of application of the FC but that the situation with regard to *near* exceptions is more complex. In particular, I articulate a pair of conditions, which I call *selection* and *homogeneity*, that are jointly but not individually sufficient for failures of the FC. Although it is very rare that these two conditions are *perfectly* satisfied, very near exceptions to the FC would be probable were they closely approximated. But the objection shows at most only that selection is commonly present, while homogeneity usually does not hold even approximately in the intended domain of the FC, such as social science and biology.

The theorem

There are several ways in which the FC might be false. Exceptions to the FC can arise through a lack of variation in the features under investigation. For example, even if X is a cause of Y , X and Y would be independent if all members of the population have the same value of X . The FC can also fail if some variables are deterministic causes of others. For instance, suppose that X is a deterministic cause of Y and that Y is a cause of Z . Then contrary to the FC, Y and Z are independent conditional on X . However, neither of these types of counterexample is deeply

troubling, at least as far as strict exceptions to the FC are concerned. An utter lack of variance in a measured variable will be obvious from the data, and hence serve to indicate that the sample is an inadequate basis for causal inference. Moreover, applications of the FC are generally limited to circumstances, such as those typical in the biological and social sciences, in which causal relations are noisy and probabilistic rather than deterministic. But there is a further sort of counterexample that will chiefly concern us here. Consider the DAG in Fig. 1.

Conceivably, the strengths the two paths from X and Z could precisely counterbalance one another and make X and Z probabilistically independent, thereby contradicting the FC. In fact, some object to the FC on the grounds that it is relatively common that causal paths exactly cancel out.

In contrast, the motivation for FC rests on the observation that exceptions to it require very special—and arguably, improbable—parameterizations. To see what this means, consider the linear causal model in Fig. 2 corresponds to the DAG in Fig. 1.

The subscripted, lowercase e 's are called “error terms,” and represent any source of variation in the dependent variable not accounted for by its direct causes. Since it is assumed that error terms are normally distributed with zero means, a *parameterization* of a linear causal model consists in specifying numerical values for the coefficients and for the variances of the error terms. In the example in Fig. 2, there are six parameters: the coefficients a , b , and c and the variances of the error terms for X , Y , and Z . The CMC characterizes those probabilistic independence relationships that follow from acyclic causal structure and the assumption that error terms are independent.¹ In an acyclic linear causal model, the independence relationships required by the CMC would hold no matter the values of the coefficients and variances of the error terms. In contrast, each of the exceptions to the FC described above depended on special assumptions about the parameters. If there is no variation in X , then the variance of its error term must be zero. Likewise, if X is a deterministic cause of Y , then the variance of e_y must be zero. Finally, if the two paths from X to Z precisely cancel out, then $b + ac = 0$.

Spirtes, Glymour and Scheines (hereafter, SGS) prove a theorem that specifies conditions under which parameterizations that violate the FC have probability zero.² The theorem is stated in terms of possible parameterizations of linear causal models, like that in Fig. 2, though with the additional assumption that the variances of error terms are positive. Notice that this assumption rules out the first two types of counterexample to the FC described above, but leaves exceptions resulting from canceling out paths. Given this set up, the theorem proceeds as follows. Consider a linear causal model with n parameters. For example, in the linear causal model in Fig. 2, n equals six since there are three coefficients and three variances of error terms. Let the set of all parameterizations of the model be represented by an n -dimensional real space, call it \mathfrak{R}^n . Now consider a subset of \mathfrak{R}^n in which every parameterization violates the FC. With regard to the linear model in Fig. 2, an example of such a subset would be one in which each parameterization makes $b + ac = 0$. SGS prove that any subset of \mathfrak{R}^n containing *only* parameterizations that violate the FC has Lebesgue measure zero. This means that any such subset of \mathfrak{R}^n is

¹ See Steel (2005).

² See their theorem 3.2 (2000, p. 42) and its proof (2000, pp. 383–384).




Fig. 1 Counteracting causal paths

$$\begin{aligned}
 X &= e_x \\
 Y &= aX + e_y \\
 Z &= bX + cY + e_z
 \end{aligned}$$

Fig. 2 A linear causal model

of $n - 1$ dimensionality or less.³ For instance, it is easy to see that making $b + ac = 0$ has this effect since it collapses three dimensions to just two. SGS then assume **L**, the proposition that subsets of \mathfrak{R}^n of Lebesgue measure zero also have probability zero, thereby yielding the conclusion that any subset of \mathfrak{R}^n containing only parameterizations that violate the FC has probability zero.

The following analogy may help to motivate the intuition underlying SGS’s theorem. Imagine two children, Sue and Mary. Sue has just finished her bath and is about to get out of the tub, while Mary will climb into the bathtub as soon as Sue leaves. Consider these two hypotheses about the water level of the tub during the period of time in which Sue is getting out and Mary getting in. T_1 : the water level in the tub does not remain constant throughout this time period; T_2 : the water level in the tub does remain constant. It seems obvious that T_1 is immensely more probable than T_2 . Although it is conceivable that Sue and Mary displace identical quantities of water and that their exit and entry are perfectly coordinated so as to maintain a constant water level, such perfect coordination seems monstrously improbable in the real world. Even if Sue and Mary were identical twins, it is unlikely that they could pull it off if they tried. For the slightest quiver, the least motion not coordinated with the other would alter the water level in the tub. Thus, given the information that the water level in the tub had remained constant throughout a particular period of time, one would infer that there had been no change in occupants of the tub during that time. The bathtub example parallels the symmetrical canceling required for exceptions to the FC: the entry of Mary into the water must exactly compensate for the exiting of Sue, and vice versa.⁴

The objection

The upshot of SGS’s theorem, then, is that exceptions to the FC involve probabilistic independences that depend upon precise, unstable parameterizations and that such

³ See Halmos (1950, p. 152) and Port (1994, pp. 54–55) for definitions of an n -dimensional Lebesgue measure.

⁴ Pearl’s chair analogy (2000, pp. 48–49), in which one considers the possibility that a second chair is hidden behind a visible one in a photograph, lacks this symmetry, since it is impossible for each chair to conceal the other. Thus, the chair analogy is problematic, since it would be easy to hide a small chair behind a big one or to hide a second chair of the same size by placing it at a sufficient distance behind the first.

exact balances are unlikely to occur. However, not all find this a convincing argument in favor of the FC, and some argue that, the theorem notwithstanding, there are circumstances in which violations of the FC are not rare. For example, Nancy Cartwright writes:

It is not uncommon for advocates of DAG-techniques to argue that cases of cancelation will be extremely rare, rare enough to count as non-existent. That seems to me unlikely, both in the engineered devices that are sometimes used to illustrate the techniques and in the socioeconomic and medical cases to which we hope to apply the techniques. For these are cases where means are adjusted to ends and where unwanted side effects tend to be eliminated wherever possible, either by following an explicit plan or by less systematic fiddling. ... The bad effects of a feature we want—or are stuck with—are offset by enhancing and encouraging its good effects. Whether we do it consciously or unconsciously, violating the FC is one of the ways we minimize damage in our social systems and in our mechanical regimens. (1999a, p. 118)

A similar argument is made by Kevin Hoover.

Spirtes et al. (1993, p. 95) acknowledge the possibility that particular parameter values might result in violations of faithfulness, but they dismiss their importance as having “measure zero.” But this will not do for macroeconomics. It fails to account for the fact that in macroeconomic and other control contexts, the policymaker aims to set parameter values in just such a way to make this supposedly measure-zero situation occur. To the degree that policy is successful, such situations are common, not infinitely rare. (2001, p. 171)

Cartwright and Hoover are both quite clear that they do not intend to claim that the FC is *never* a reasonable assumption, but only that it is inappropriate in certain commonly occurring circumstances, namely, when there is some process that selects for canceling out causal paths.⁵

And indeed, whether SGS’s theorem provides a motivation for the FC in a given context clearly depends on whether its assumptions are reasonable there. The most apparent limitation of the theorem is its restriction to linear causal models. But SGS conjecture that their theorem holds in the non-linear cases as well (2000, p. 42), and the theorem has in fact been extended to causal models with discrete variables (Meek 1995). Moreover, the point of Cartwright and Hoover’s objection is not that the causal relationships might be non-linear but that selection for FC-violating parameterizations could occur.

The other significant assumption in the theorem is **L**. In the linear model in Fig. 2, **L** entails that any subset of \mathfrak{R}^n consisting solely of parameterizations in which $b + ac = 0$ must have probability zero, since any such subset must be of $n - 1$ dimensionality or less. But it is precisely this assumption that is challenged in Cartwright and Hoover’s objection. If their objection is correct, then it is likely that the actual parameterization falls within Lebesgue measure zero subsets of \mathfrak{R}^n when there is selection for counteracting causal paths. And **L** does seem to be a promising

⁵ Glymour (1999, p. 161) responds to Cartwright by claiming that reliable causal inference is impossible without the FC. Even if this claim were true, it would not follow that the objection is mistaken, since reliable causal inference might be more narrowly restricted than one would have hoped.

target of criticism. SGS do not provide any motivation for **L**, yet the assumption not obviously true, since it is not plausible that subsets of Lebesgue measure zero must *always* have probability zero. For such a claim would entail that we must be certain a priori that no quantity is equal to any other quantity. This point can be appreciated by reference to the diagram in Fig. 2.

In the diagram, the subset of pairs of values in which a equals b is represented by the diagonal line in the square, and in a two dimensional space lines have Lebesgue measure zero. Consequently, SGS's theorem can serve as a motivation for the FC only provided some explication of the conditions under which **L** is true and of why we should think that those conditions hold in the domain of application of the FC.⁶

Deciding whether **L** is a reasonable assumption is difficult in part because it is unclear how the probability distribution over the parameter space should be interpreted. For example, **L** might be regarded as a constraint on rational degrees of belief. But given the well-known difficulties facing objective Bayesianism in motivating such constraints above and beyond the axioms of probability,⁷ it seems likely that any argument that **L** deserves such status would be difficult to make. A more promising suggestion is that the probability distribution over \mathfrak{R}^n represents a physical chance process. Given this interpretation, it is impossible to know whether **L** is a reasonable assumption unless some information concerning this chance process is provided. But if the relevant features of the physical process were described, it might be possible to decide whether **L** is appropriate. Pursuing this line of thought requires some exploration of what the relevant features of such a chance process are. In the remainder this essay I examine this issue and its connection to Cartwright and Hoover's objection to the FC.

As was noted in the introduction, the objection that violations of the FC are common in some types of circumstance can be interpreted as making a claim about either *strict* or *very near* exceptions. Let us consider these two cases in turn.

When **L** is strictly true

Judea Pearl (1998, p. 121) and Jim Woodward (1998, pp. 142–147) make similar proposals about when **L** and, hence the FC, is true. Consider again a linear causal model, such as that in Fig. 2. Pearl and Woodward's thought is that in such fields as biology or social science, causal structures represented in such models are often stable across a range of populations, while the specific values of the parameters of the model are extremely sensitive to the conditions of particular populations at particular times. For example, the structure specified by the laws of supply and demand is very stable, while the parameters representing quantitative aspects of the causal relationships vary from one economy to the next. This is in effect a built in assumption of SGS's theorem, wherein one supposes that a given causal structure is fixed while the parameter values vary. Pearl and Woodward's suggestion is that **L** (and thereby the FC) is appropriate when the causal structure is constant but

⁶ That it would be unreasonable to insist that sets of Lebesgue measure zero must always have probability zero is noted by Pearl (1998, p. 121). SGS also acknowledge the point (2000, p. 66) but do not explain why sets of Lebesgue measure zero should have zero probability the in the sorts of cases relevant to their theorem.

⁷ For an overview of this issue, see Howson and Urbach (1993, chapter 4).

parameters “vary independently” (cf. Pearl 1998, p. 121; Woodward 1998, p. 145). However, it turns out that, given what seems the most natural interpretation, independently varying parameters is neither necessary nor sufficient to ensure that all sets of Lebesgue measure zero receive probability zero. In this section, I explain how this is so, identify a necessary and sufficient condition for **L**, and explain why that condition is quite reasonable in the usual domain of application of the FC.

Addressing this issue requires associating each parameter with a random variable whose values correspond to possible values of the parameter. Let the set of all these random variables be $\mathbf{V} = \{V_1, \dots, V_n\}$ where n is the number of parameters. It is important to keep the distinction between the members of \mathbf{V} and the variables in a causal model (e.g., X , Y , and Z in Fig. 2) firmly in mind. Variables in a causal model vary *within* populations, while parameters and hence the members of \mathbf{V} are fixed within particular populations at particular times but vary *across* populations. As before, \mathfrak{R}^n is the n -dimensional real space of all combinations of values of the members of \mathbf{V} . A *joint distribution function* of \mathbf{V} , then, is a function that assigns probabilities to subsets of \mathfrak{R}^n . For example, if $\mathbf{V} = \{V_1, V_2\}$, then the joint distribution function, $F(v_1, v_2)$, specifies $P(V_1 \leq v_1, V_2 \leq v_2)$ for any v_1 and v_2 . The question at issue can then be rephrased as follows: what conditions must the joint distribution function of \mathbf{V} satisfy to ensure that subsets of \mathfrak{R}^n of Lebesgue measure zero also have probability zero?

One way to proceed is by considering ways in which **L** might be false, that is, to consider ways in which a joint distribution function might assign positive probability to an $n - 1$ dimensional subset of \mathfrak{R}^n . The concept of a *support* of a probability measure is useful in this regard. A set A is said to be a *support* of the probability distribution P exactly if $P(A) = 1$.⁸ For example, suppose that \mathfrak{R}^n is a plane and that A is a line in \mathfrak{R}^n . Then if A is a support of the probability distribution, there is a subset of Lebesgue measure zero that receives probability greater than zero (e.g., A itself). Thus, a set of Lebesgue measure zero receives positive probability whenever a subset of $n - 1$ dimensionality or less is a support of the probability distribution over the parameter space. Let us consider how such a thing might occur.

The simplest case occurs when one or more of the parameters are entirely fixed and invariable. It is easy to see that some subsets of Lebesgue measure zero will have positive probability when this is the case. For example, in the two dimensional case, if one of the members of \mathbf{V} has zero variance, then one of the two variables is constant and a straight horizontal or vertical line is a support of the probability distribution. But this is not the only way for a Lebesgue measure zero subset of \mathfrak{R}^n to have a positive probability. Returning to the two-dimensional case again, suppose that although the variance of each variable is positive, it is completely certain that the values of the two variables are equal. In this case, a 45° diagonal line, and hence to a set of Lebesgue measure zero, would be a support of the probability distribution (as in Fig. 3).

Pearl and Woodward’s requirement that the parameters are unstable and vary independently of one another is presumably intended to rule out such cases. Although neither Pearl nor Woodward gives a precise characterization of “vary independently,” it is quite natural to interpret this phrase as referring to the technical concept of *variation independence*. Variation independence means that no

⁸ See Billingsley (1995, p. 23).




Fig. 3 A set of Lebesgue measure zero

group of parameter values restricts the possible range of values of any other parameter.⁹ In the present context, for any V_i in \mathbf{V} and any a and b , if the marginal probability that $a \leq V_i \leq b$ is greater than zero, then the probability of $a \leq V_i \leq b$ conditional on any specification of values of any combination of the other members of \mathbf{V} is also greater than zero. The example described above in which \mathbf{V} has just two members, each with positive variance but whose values are certainly equal, is ruled out by variation independence. For instance, suppose that the marginal probability that V_1 is between 2 and 3 is greater than zero. But then the probability that V_1 falls within this interval is zero if V_2 equals, say, 4.

So, is **L** true whenever no parameter has a fixed, constant value and variation independence obtains? That the answer to this question is “no” can be seen from the example in Fig. 4. Continuing with the two-dimensional case, suppose that the two parallel lines are a support of the probability distribution.

For instance, let the joint distribution function $F(v_1, v_2)$ equal $0.5v_2$ if v_1 is either 0.25 or 0.75, and zero otherwise. Then the marginal distribution of V_1 is discrete with $P(V_1 = 0.25) = P(V_1 = 0.75) = 0.5$, and V_2 is continuous with a probability density function of 1, so that the marginal distribution function of V_2 is $F_2(v_2) = \int_0^{v_2} dv_2 = v_2$. Since V_1 and V_1 are probabilistically independent, variation independence is clearly satisfied, yet a subset of Lebesgue measure zero is a support of the probability distribution.

The counterexample in Fig. 4 relies on having one of the variables vary discretely rather than continuously. Thus, we could rule out the counterexample by requiring that each member of \mathbf{V} be marginally continuous, that is, that each varies continuously considered separately from the others. Indeed, perhaps Pearl and Woodward intended the instability of parameters to imply that the variation of each parameter is marginally continuous. It turns out that variation independence together with the marginal continuity of each member of \mathbf{V} entails that any Lebesgue measure zero subset of \mathfrak{R}^n receives probability zero. However, variation independence and marginal continuity are stronger than necessary, as can be shown by identifying a weaker sufficient condition for **L**.

A sufficient condition for **L** is that the variables in \mathbf{V} are *jointly continuous*.¹⁰ Joint continuity is simply an extension of the concept of a continuous distribution for a single variable to joint distributions of arbitrarily many variables. In a continuous

⁹ See Lindsey (1996, p. 29). Thanks to an anonymous referee for suggesting this interpretation.

¹⁰ This seems to be what Meek (1995, p. 411) has in mind when he writes that all Lebesgue measure zero subsets of \mathfrak{R}^n have probability zero when the distribution function of \mathbf{V} is “smooth.” For a more formal definition of jointly continuous random variables than that provided here, see Stirzaker (2003, chapter 8).




Fig. 4 Why variation independence is not sufficient

distribution of a single variable, the distribution function specifies areas over intervals of a line. Similarly, a jointly continuous distribution for two variables specifies volumes over areas of a plane. When $n = 1$, the probability of any interval of \mathfrak{R}^n corresponds to the area above it. Likewise, when $n = 2$, the probability of any area of \mathfrak{R}^n equals the volume above. The idea generalizes to examples in which n is greater than two, though such cases are not easy to visualize.

The joint continuity of the members of \mathbf{V} ensures that all subsets of Lebesgue measure zero receive probability zero because it entails that no subset of \mathfrak{R}^n of $n - 1$ dimensionality or less receives a positive probability. For concreteness and without loss of generality, consider the two dimensional case. Then it can easily be seen that if the members of \mathbf{V} are jointly continuous, any line receives probability zero, since the volume over a line is zero. By similar reasoning, any point must also receive probability zero. In short, any one-dimensional or less subset of \mathfrak{R}^n must receive probability zero if \mathbf{V} is jointly continuous and $n = 2$. But in a two-dimensional space, any set of Lebesgue measure zero is of one or less dimension. So, in general, if the members of \mathbf{V} are jointly continuous, then any subset of \mathfrak{R}^n of $n - 1$ dimensionality or less, and hence any set of Lebesgue measure zero, has probability zero.

It is obvious that marginal continuity and variation independence together entail joint continuity, and hence that the conjunction of these two conditions is also sufficient for \mathbf{L} . However, variation independence is not necessary, since it is not entailed by joint continuity. If the members of \mathbf{V} are jointly continuous, then each member of \mathbf{V} is continuously distributed not only marginally but also conditional on any of the other members of \mathbf{V} . But joint continuity does not require that variation independence be true, since the range of possible values of one variable may be restricted by the value of another even if each variable is continuously distributed conditional on any combination of other variables. For instance, suppose that, marginally, the probability V_1 is uniformly distributed over a particular interval $[a, b]$, but that conditional on certain values of V_2 it is uniformly distributed over a proper subset of $[a, b]$. In such a case, variation independence is violated yet V_1 and V_2 may nevertheless be jointly continuous.

Let us return to the question of whether strict exceptions to \mathbf{L} are ever probable within the FC's intended domain of use. The answer to this question is *no* as long as each parameter varies continuously considered separately and conditional on any subset of the others. And this will be case so long as the parameter values depend upon continuously variable factors that are not themselves perfectly coordinated. Yet it is quite plausible that this is indeed the case in biology and social science, and indeed, in any field that studies complex systems in which the strength of causal relationships depend on a plethora of variable factors. In such systems, the distribution function of \mathbf{V} might be tightly focused on a Lebesgue measure zero subset of \mathfrak{R}^n , but it nevertheless is extremely probable that there will be some continuous

variation in the parameter values. However, even if *strict* failures of **L** (and hence the FC) are indeed highly unlikely, it does not follow that the same is true of *near* exceptions, which for practical purposes may be just as problematic as strict counterexamples. Let us turn, then, to a consideration of near failures of the FC.

Homogeneity, selection, and near exceptions to the FC

Near exceptions to the FC raise several complications not found in the case of strict exceptions. Near exceptions might occur as a result of counteracting causal paths that very nearly, though not exactly, cancel out. But they can also occur in other ways. For instance, suppose that *X* is a cause of *Y* and that these two variables are otherwise unrelated. Then if the influence of *X* upon *Y* is extremely minute, so too will be the probabilistic dependence between them. Such near failures of the FC are not particularly worrisome, since it is not much of a mistake to conclude that an extremely weak causal link is no causal connection at all. Consequently, I shall not be concerned with near exceptions of this sort in what follows. A second complication arises in connection to the variation in the variables. As explained in Sect. 2, strict exceptions to the FC can result if the variances of some error terms are zero. Yet it seems reasonable to assume that such circumstances are unlikely in the intended domain of application of the FC. Presumably, this is what motivates the assumption in SGS's theorem that the variances of the error terms are positive. However, near exceptions to the FC can result from positive yet very minimal variances of error terms that make some variables almost constant or some connections almost deterministic. I will be primarily concerned with cases involving almost canceling out paths, but I will also briefly discuss near exceptions arising from minimal variation in variables.

Consider a social planner attempting to do what Hoover describes in his objection, that is, create a compensating mechanism to precisely counteract an undesired side effect of some policy. For example, imagine a road improvement program that involves resurfacing and widening a number of large thoroughfares and some smaller side streets. Although improved, safer roads contribute to fewer traffic accidents, they also have the unfortunate side effect of increasing speeding, which is a significant cause of traffic fatalities. Letting *R*, *S*, and *T* be variables denoting road improvement, rates of speeding and traffic fatalities, respectively, the causal relationships can be represented in the graph in Fig. 5.

Suppose that, initially, the net effect of the road improvement is to increase the rate of traffic fatalities. To offset this problem, more police are hired to patrol the newly improved roads and the fines for speeding are increased. However, given a tight budgetary situation, the social planners do not want to spend more money on




Fig. 5 Road improvement and traffic fatalities

speeding prevention than necessary. They want to do just enough to make the two causal paths cancel out, and no more.

The strategy of the social planners in this case is to implement changes in the situation that will weaken the positive influence of R upon S so as to even the balance between the two paths. In principle, if the strength of influence of R upon S can be fine-tuned independently of the other parameters, this would be possible. But the relevant question is whether the social planners really can make the exact canceling out occur, or at least be sufficiently approximated for practical purposes. Their ability to do so requires the following two things:

Selection of parameters: a process that tends to concentrate the weight of the distribution of parameterizations on a subset in which the FC is violated.

Homogeneity of parameters: the absence of factors that perturb parameter values and thereby alter their distribution in uncontrolled ways.

It may be helpful to consider what these two conditions mean with respect to an example of a particular causal model, for instance, the linear causal model in Fig. 2. Recall that parameters are assumed to be fixed for particular populations at particular times, although they may vary from one population to the next. Imagine a collection of populations in which the causal structure represented in the linear causal model in Fig. 2 is constant, but across which the values of the parameters (the coefficients and variances of error terms) are variable. Selection, then, is a process that if applied to such a collection of populations, would over time focus the distribution of parameters on a subset of \mathfrak{P}^n in which the FC is violated. Homogeneity refers to a lack of other perturbations of the parameters besides those involved in the selection process, so that parameters set to particular values by the selection process tend to stay put. For convenience, I will generally abbreviate “selection of parameters” and “homogeneity of parameters” to “selection” and “homogeneity.” However, it is important to keep homogeneity of *parameters* of a model distinct from a lack of variance in the *variables* of that model. For instance, the fixity of the parameters in the linear causal model in Fig. 2 entails nothing about the variances of the variables in any population accurately characterized by that model.

If selection and homogeneity were perfectly achieved, then the probability distribution of parameterizations would be restricted to an $n - 1$ dimensional subset of the parameter space. Thus, when selection and homogeneity are perfectly satisfied, assumption **L** of SGS’s theorem (that all $n - 1$ dimensional subsets of the parameter space receive probability zero) is false. Of course, perfect homogeneity and selection rarely if ever occur in real life. But as Christopher Meek points out, “the interesting questions about reliably inferring Bayesian networks from data (rather than a population distribution) have to do with near violations of faithfulness” (1995, p. 411). In other words, near exceptions to the FC would not be a concern for an omniscient being with perfect knowledge of the probability distribution, since such a being would always be able to distinguish zero correlations from non-zero ones, no matter how miniscule the difference. However, for humans who must estimate probability distributions from finite samples of data, very near exceptions to the FC can be just as bad as strict ones.

I claim, then, that selection and homogeneity are jointly but not individually sufficient for (near) failures of the FC to be probable. That they are jointly sufficient is straightforward, since when perfect selection and homogeneity are closely

approximated, the distribution over the space of parameterizations will be tightly focused on an FC-violating subset. And in this case it is very probable that the parameterization will be one in which the FC is nearly false. Hence, it only remains to show that neither condition is sufficient in the absence of the other. The case in which homogeneity but not selection is present is simple. Although lack of variation in the *variables* of a model can produce exceptions to the FC, the same is not true with respect to homogeneity of *parameters*. Of course, the *values* of the parameters—particularly, the variances of the error terms—matter to the variation of the variables. But that the parameters are constant across distinct populations entails nothing about what the values of those parameters are. And since strict or near exceptions depend on the specific parameter values, nothing can be inferred about violations of the FC from homogeneity of parameters alone.

The more interesting case is that in which selection is present and homogeneity is absent. Even if there is a process at work that tends to focus the probability distribution of parameterizations around a FC-violating subset, it does not follow that exceptions or near exceptions the FC are probable, since the distribution of parameters might also be influenced by other trends that undo the work of the selection process. Suppose that there are a wide variety of difficult to predict or control factors at play that are capable of altering the values of the parameters, i.e., that homogeneity does not hold even approximately. Clearly, these disturbing factors would be expected to increase the variance of the distribution of parameterizations, increasing the chance that the actual parameterization would fall in a region distant from a subset in which the FC is false. In addition to enlarging the variance of the distribution, factors that alter the values of parameters can also change its mean if not all parameters are uniformly susceptible to disturbance. For instance, if some parameters are more susceptible than others to factors that alter their values in a particular direction, then the mean of the distribution may be driven away from an FC-violating subset. In the above example, if the effect of *R* upon *S* is sensitive to factors that tend to increase its value while the other parameters are relatively stable, then the mean of the parameterizations will move towards a positive net effect of *R* upon *T*.

The simple moral, then, is that the existence of a selection process can fail to make exceptions or near exceptions to the FC probable when a variety of uncontrollable factors that perturb parameter values are present. Consequently, noting the presence of a selection process does not suffice to show that (near) violations of the FC are likely to occur. Yet Cartwright and Hoover's objection points out that it is common that selection processes are present or at least that some effort is made to create them and thence concludes that exceptions or near exceptions to the FC are likewise commonplace. This argument is invalid on two grounds. First, effectively designing and implementing a selection process may be very difficult, so the fact there is some effort afoot to create one provides little assurance that one exists. Secondly, even if selection processes were common, Cartwright and Hoover's conclusion would follow only if homogeneity generally obtained in the intended domain of application of the FC, and there is reason to suspect that the opposite is true.

For simple systems whose parts are easily manipulated and inspected, causal knowledge can often be readily attained without the need for sophisticated analysis of large samples of statistical data. The intended domain of principles that facilitate causal inference from statistical data, therefore, consists of more complex systems whose workings are not so easily ascertained, such as social groups, biological

organisms, ecosystems, and so forth. But such complex systems are ones in which parameter values depend upon variable factors that are difficult to predict or control. Hence, the intended domain of the FC consists of causal systems of which it is quite doubtful that homogeneity is typically true or approximately true.¹¹

In short, Cartwright and Hoover's objection has failed to show that exceptions or near exceptions to the FC are common in its intended domain of use. Nevertheless, it would be a mistake to conclude that the FC is *always* an entirely unproblematic assumption with regard to complex systems. Since near exceptions to the FC are to be expected when selection and homogeneity are well approximated, a natural question is whether realistic examples of such cases occur within the intended domain of application of the FC. In fact, cases of this sort sometimes arise in gene knockout experiments.¹²

For example, consider a gene that serves as a template for the transcription of a protein that normally performs a specific set of functions in a cell, but when that protein is not present in sufficient quantities, the transcription of a distinct yet functionally similar protein from a second gene is increased. It is plausible that there would be an adaptive benefit in having the quantitative strengths of the two paths counterbalance one another. For example, maintaining the function may require that the sum quantity of two products be kept within certain bounds. Hence, it would not be optimal for the genes for both to normally be transcribed together, while it is beneficial that the function be maintained at the normal rate when the usual product is not present in adequate quantities. Thus, natural selection would constitute a selection process that favors parameterizations in which the counteracting paths exactly or very nearly cancel out. As explained above, the presence of a selection process alone is not sufficient to make exceptions or near exceptions to the FC probable; homogeneity is also required. But this latter condition is much more likely to be approximated in the context of a gene knockout experiment than in a wild population. Organisms in knockout experiments are typically generated from extremely genetically homogenous strains that have been reared for numerous generations under standard laboratory conditions.

Given the above analysis, one would expect near exceptions to the FC to be a concern in gene knockout experiments, and this is indeed the case (cf. Pearson 2002, p. 8). For example, there are examples of nearly canceling out causal paths (cf. Scarff et al. 2004). Other types of near exceptions to the FC can also be found in the gene knockout literature, particularly, cases resulting from a lack of variance. For example, there are cases in which two genes perform the same function, and nearly all units of the experimental population possess both genes, which has the result that disabling one of the genes makes no detectable difference to the function (cf. Liljegren et al. 2000). Lack of variation in measured variables and homogeneity of parameters likely arise from a common source in gene knockout experiments, namely, the genetic and environmental homogeneity of the experimental organisms.

¹¹ A consequence of this point is that examples of relatively simple technological devices in which near violations of the FC can be made probable do not show that near exceptions to the FC are likely in its intended domain of application. In regard to this, see Cartwright's "solition" example, which she uses to motivate her objection to the FC (1999a, pp. 30–31, 118).

¹² Chu, Glymour, Scheines and Spirtes (2003) demonstrate an obstacle to the Causal Markov Condition in studies that aim to infer gene regulatory networks from microarray data. However, as they observe (Chu et al. 2003, p. 1147), this difficulty is not relevant to gene knockout experiments.

In the context of gene knockout experiments, then, it would be unwise to presume the FC in anything but a very tentative manner, a point which is clearly reflected in practice.¹³ This conclusion shows that Cartwright and Hoover were correct to assert that exceptions, or at least near exceptions, to the FC may be common in some circumstances, SGS's theorem notwithstanding. However, the range of circumstances in which the FC is a doubtful premise for causal inference is much more narrowly restricted than they suggest. Neither attempts to devise a process that selects for FC violating parameterizations nor the presence of such a process are sufficient to make near exceptions to the FC probable. It is also necessary that factors capable of perturbing parameter values in uncontrollable and unpredictable ways be significantly diminished. But such homogeneity is unlikely to occur most circumstances in which one wishes to utilize the FC, although it may occur in some tightly controlled laboratory settings.

Conclusion

The FC entails that where there is no probabilistic dependence, there is also no causal connection. This principle significantly facilitates causal inferences from statistical data and is implicit in nearly any study that reports no statistically significant correlation among a certain pair of variables and thereby concludes that neither is a cause of the other. The potential of the FC to aid causal inference makes it worthwhile to consider what basis there is for regarding it as an appropriate assumption.

In this essay, I have examined an objection, made by Cartwright and Hoover, to attempts to justify the FC by reference to theorems showing that exceptions to it have probability zero given some apparently reasonable assumptions. I maintained that selection and homogeneity are jointly but not individually sufficient for exceptions to the FC to be probable. Given this, Cartwright and Hoover's objection fails to show that exceptions to the FC are common in its intended domain, since they only point out that one of these two conditions—selection—not infrequently occurs. In contrast, homogeneity is usually not a plausible assumption with regard to the complex systems that fall into the intended range of application of the FC. Of course, that does not mean that the FC is always an unproblematic assumption. But it does indicate that the status of the FC as a general principle of causal inference is on firmer ground than the objection would suggest.

Acknowledgments I would like to thank Richard Scheines for helpful comments on an earlier draft and Clark Glymour for discussion. I would also like to thank audiences at the 2004 meeting of the British Society for the Philosophy of Science (especially Dan Hausman) and the 2004 Konstanz Summer School on Causality, Ignorance, and Uncertainty, and two anonymous referees.

References

- Billingsley, P. (1995). *Probability and measure* (3rd edn.). New York, NY: Wiley.
 Cartwright, N. (1999a). *The dappled word: a study of the boundaries of science*. Cambridge: Cambridge University Press.

¹³ Robins and Wasserman (1999) also maintain that epidemiologists would be hesitant to assume the FC for a rather different reason that described here.

- Cartwright, N. (1999b). Causal diversity and the Markov condition. *Synthese*, 121, 3–27.
- Chu, T., Glymour, C., Scheines, R., & Spirtes, P. (2003). A statistical problem for inference to regulatory structure from associations of gene expression measurement with microarrays. *Bioinformatics*, 19, 1147–1152.
- Glymour, C. (1999). Rabbit hunting. *Synthese*, 121, 55–78.
- Glymour, C., & Cooper, G. (Eds.) (1999). *Computation, causation, and discovery*. Menlo Park, CA: American Association for Artificial Intelligence.
- Halmos, P. (1950). *Measure theory*. New York: D. Van Nostrand Company.
- Hoover, K. (2001). *Causality in macroeconomics*. Cambridge, UK: Cambridge University Press.
- Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach* (2nd edn.). La Salle, IL: Open Court Press.
- Liljegren, S., Ditta, G., Eshed, Y., Savidge, B., Bowman, J., & Yanofsky, M. (2000). Shatterproof MADS-box genes control seed dispersal in *Arabidopsis*. *Nature*, 404, 766–770.
- Lindsey, J. K. (1996). *Parametric statistical inference*. Oxford: Clarendon Press.
- Meek, C. (1995). Strong completeness and faithfulness in Bayesian networks. In P. Besnard (Ed.), *Uncertainty in artificial intelligence: Proceedings of the 11th conference* (pp. 411–418). San Francisco: Morgan Kaufman Publishers.
- Neopolitan, R. (2004). *Learning Bayesian networks*. Upper Saddle River, NJ: Prentice Hall.
- Pearl, J. (1998). TETRAD and SEM. *Multivariate Behavioral Research*, 33, 119–128.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Pearson, H. (2002). Surviving a knockout blow. *Nature*, 415, 8–9.
- Port, S. C. (1994). *Theoretical probability for applications*. New York: Wiley.
- Robins, J. M., & Wasserman, L. (1999). On the impossibility of inferring causation from association without background knowledge. In: C. Glymour, & G. Cooper (Eds.), *Computation, causation, and discovery* (pp. 305–322). Menlo park, CA, Cambridge, MA: AAAI Press/The MIT Press.
- Scarff, K., Ung, K., Nandurkar, H., Crack, P., Bird, C., & Bird, P. (2004). Targeted disruption of SPI3/Serpib6 does not result in developmental or growth defects, leukocyte dysfunction, or susceptibility to stroke. *Molecular and Cellular Biology*, 24, 4075–4082.
- Spirtes, P., Glymour, C., & Scheines, R. (1993). *Causation, prediction, and search*. Berlin Heidelberg New York: Springer.
- Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction, and search*, (2nd edn.). Cambridge, MA: MIT Press.
- Steel, D. (2005). Indeterminism and the Causal Markov Condition. *British Journal for the Philosophy of Science*, 56, 3–26.
- Stirzaker, D. (2003). *Elementary probability* (2nd edn.). Cambridge, UK: Cambridge University Press.
- Woodward, J. (1998). Causal independence and faithfulness. *Multivariate Behavioral Research*, 33, 129–148.