UvA-DARE (Digital Academic Repository)

Toward a causal interpretation of the common factor model

Van Bork, R.; Wijsen, L.D.; Rhemtulla, M.

DOI

[10.1515/disp-2017-0019](https://doi.org/10.1515/disp-2017-0019)

Publication date

2017

Document Version

Final published version

Published in

Disputatio

License

CC BY-NC-ND

[Link to publication](#)

Citation for published version (APA):

Van Bork, R., Wijsen, L. D., & Rhemtulla, M. (2017). Toward a causal interpretation of the common factor model. *Disputatio*, 9(47), 581–601. <https://doi.org/10.1515/disp-2017-0019>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Toward a Causal Interpretation of the Common Factor Model

Riet Van Bork¹

University of Amsterdam

Lisa D. Wijsen¹

University of Amsterdam

Mijke Rhemtulla

University of California, Davis

BIBLID [0873-626X (2017) 47; pp. 581–601]

DOI: 10.1515/disp-2017-0019

Abstract

Psychological constructs such as personality dimensions or cognitive traits are typically unobserved and are therefore measured by observing so-called indicators of the latent construct (e.g., responses to questionnaire items or observed behavior). The Common Factor Model (CFM) models the relations between the observed indicators and the latent variable. In this article we argue in favor of interpreting the CFM as a causal model rather than merely a statistical model, in which common factors are only descriptions of the indicators. When there is sufficient reason to hypothesize that the underlying causal structure of the data is a common cause structure, a causal interpretation of the CFM has several benefits over a merely statistical interpretation of the model. We argue that (1) a causal interpretation conforms with most research questions in which the goal is to *explain* the correlations between indicators rather than merely summarizing them; (2) a causal interpretation of the factor model legitimizes the focus on *shared*, rather than unique variance of the indicators; and (3) a causal interpretation of the factor model legitimizes the assumption of local independence.

¹ These two authors contributed equally.

Keywords

Causality, reflective model, statistical model.

1 Toward a causal interpretation of the common factor model

One of the many pursuits of psychology is to establish causal relations between properties of the mind and human behavior. Psychologists are interested in the motivations, cognitive abilities and personality traits that explain why people behave in a certain way. This endeavor has proved to be anything but easy. Part of the problem comes down to the very nature of psychological constructs. In psychology, the attributes that are used to explain human behavior (e.g., personality characteristics that explain individual differences in behavior) are typically *latent variables*, that is, theoretical constructs that are unobserved (Borsboom, Mellenbergh and van Heerden 2003). Because latent variables are by definition unobserved, psychological tests are constructed to measure them. The observed responses on such tests are believed to reflect the latent variable that underlies them. For example, *intelligence* cannot be directly observed. Yet, it is assumed that intelligence can be measured by administering a set of IQ items to which the responses can be observed. The measurement model on which this mechanism is based is the *reflective model* (Figure 1). The IQ test is supposed to measure intelligence because it is built on the assumption that the responses to the items are a direct *effect* of one unobserved entity, intelligence. Consequently, the variance shared among the test scores is assumed to reflect this theoretical entity. The variance that is unique to each item is assumed to reflect measurement error as well as unique causes of the responses to that item (represented by ϵ_1 to ϵ_4 in Figure 1). The reflective model is an example of a common cause model because the latent variable functions as a cause of all of the item responses. In the remainder of the paper, we refer to the observed variables (such as the test scores in this example) as the *indicators* of the latent variable.

Figure 1. The reflective model. Squares represent indicators (e.g., X_1 to X_4 could be the responses on four different IQ items) and η represents the latent variable (e.g., intelligence). ε_1 to ε_4 refer to measurement error and unique causes of the indicators. Eliminating these unique factors would imply that the indicators are fully determined by the latent variable, which is typically not assumed for causal processes in psychology.

Latent variable modeling originated with the construction of the *common factor model* (CFM; Spearman 1904). Spearman observed what is now known as the positive manifold: item responses on a variety of mental ability tests were all positively correlated with one another. Using the principle of the common cause in which a correlation between two variables is explained by a third variable, Spearman abstracted the general factor of intelligence (*g*-factor) from these test scores, arguing that one underlying entity explains the shared variance in all branches of intellectual activity: namely, general intelligence. In the CFM all indicators are a linear function of the common factor and the indicators are statistically independent conditional on this latent variable. After all, if a common cause explains the covariation between two indicators, the indicators no longer correlate when conditioning on this common cause. So, a CFM of general intelligence implies that all branches of intellectual activity are rendered independent, conditional on general intelligence. This principle is fundamental to Spearman's CFM and to reflective models in general, and is called the *principle of local independence*. The process of fitting a CFM to data is called *factor analysis*.

Factor analysis enabled psychologists to discover and test theories

about plausible explanations for human behavior. It is therefore not surprising that CFMs have become very popular, not only in intelligence research, but also in the fields of psychopathology (e.g., Asmundson 2000, Caspi et al. 2014) and personality (e.g., McCrae and Costa 1987, Muecke 2007).

We reserve the name *reflective model* for the theoretical causal interpretation of the CFM. In contrast to the reflective model, the CFM is not defined as a causal model but is typically defined as the set of equations that equate each indicator X_i to a function of the latent variable η and a unique component ε_i that is typically called *the residual* in factor analysis. For example, consider four indicators X_1 to X_4 of the same latent variable η . Each indicator is a linear function of the same latent variable and a unique residual component:

$$X_1 = \lambda_1 \eta + \varepsilon_1$$

$$X_2 = \lambda_2 \eta + \varepsilon_2$$

$$X_3 = \lambda_3 \eta + \varepsilon_3$$

$$X_4 = \lambda_4 \eta + \varepsilon_4$$

These equations are typically graphically represented in the same way as the reflective model (see Figure 1). The covariance matrix of the indicators Σ is a function of the vector of factor loadings λ , the variance of the latent variable ψ and the covariance matrix of the residuals Θ :

$$\Sigma = \lambda \psi \lambda + \Theta$$

Although the equations above are agnostic with respect to causality, we argue that when factor analysis is used to measure a psychological construct, the CFM benefits from being interpreted as a reflective model, that is, interpreting the latent variable as the common cause of the indicators. The reflective model in Figure 1 can be seen as a specific case of the CFM: it includes the CFM equations above, and adds to them a causal interpretation (Bollen and Lennox 1991, Edwards and Bagozzi 2000).

Even though latent variable modeling has had a profound influence on psychological research, methodologists and modelers have not reached a general consensus about how to understand the nature

of latent variables (Bollen 2002, Borsboom et al. 2003, Jonas and Markon 2016). As latent variables are not directly observed, it is typically impossible to perform a manipulation on them to determine (a) whether they exist at all, and (b) whether they really *cause* the observed responses that they are supposed to cause. For example, it is impossible for us to manipulate a person's intelligence to test whether the answers on an IQ test change as a result of this intervention².

It is not only the unobservable character of psychological constructs that makes causal modeling in psychology a complicated endeavour. Factor analysis is most frequently performed on cross-sectional data, gathered at a single time point. Such data cannot be used to distinguish between models that have different causal structures but are statistically equivalent. Whenever a CFM is fit to a dataset, there are alternative causal structures that may equally well have generated the data (van der Maas et al. 2006). Because of this ambiguity, many psychometricians advocate sticking to a strict statistical interpretation of the CFM, to avoid inferring causality without direct evidence for it. They argue instead for a descriptivist approach in which the latent variable merely represents the shared variance among a set of indicators. The descriptivist approach understands latent variables as a parsimonious summary of the data, rather than an underlying cause of the indicators (Jonas and Markon 2016).

Thus, although the CFM was developed as a model in which the common factor is hypothesized to represent an existing causal entity (i.e., general intelligence) that explains patterns in different branches of intellectual activity (Spearman 1904), the descriptivist approach views the common factor that is obtained with factor analysis merely as “just a convenient way of summarizing patterns of observed relationships” (Jonas and Markon 2016: 91). Jonas and Markon (2016) argue:

Reflective latent variable models are agnostic with regard to the nature of the etiological process: this is the heart of the descriptivist paradigm. Reflective latent variable models can describe data generated

² Note that it is possible to hypothesize an intervention on intelligence, e.g., drinking a lot of alcohol (Borsboom, Mellenbergh and van Heerden 2003). However, it is problematic to test whether this intervention actually affects intelligence or whether it instead affects the indicators directly, because intelligence is only visible via its reflection in the indicators.

from any number of etiologic processes; consequently, the form of the latent variable model cannot arbitrate questions of causality. (Jonas and Markon 2016: 91–2)

Factor analysis does not test any causal relations but rather absorbs shared variance in a common factor. For any arbitrary set of indicators that share variance, the shared variance will constitute a common factor, no matter what caused this shared variance. However, in order to measure a psychological construct of interest, one does not consider any randomly chosen set of variables, but rather a specific set of variables that are hypothesized to be affected by the construct. The hypothesis that the indicators reflect the psychological construct of interest enables the researcher to interpret the common factor as this construct, by the logic that, *if* the indicators share a common cause, the shared variance among these variables reflects this cause (Edwards and Bagozzi 2000).

In this paper, we argue against the view that the common factor is merely a convenient summary of the data and argue for a causal interpretation of the CFM when used to measure constructs. When a causal interpretation is not justified, for example because it is unlikely that the construct of interest causes the indicators, other statistical models might be more appropriate (see for example Bollen 2011, Diamantopoulos, Riefler and Roth 2008, Edwards and Bagozzi 2000). Although a causal interpretation of the CFM is neither always justified nor necessary, it offers several benefits over a purely statistical reading of the model in cases where it is justified. To make this argument we first distinguish between statistical models and causal models, and explain why researchers might be reluctant to interpret the CFM causally. Subsequently, we outline the descriptivist approach in more detail. Finally, we present three arguments that highlight the benefits of a causal interpretation over a purely statistical interpretation of the factor model.

2 Statistical vs. causal models

We distinguish between statistical models on the one hand, and statistical models with a causal interpretation on the other. For brevity, we use *causal model* to refer to the latter, though we acknowledge that

causal models need not be statistical. A statistical model is a set of probability distributions on some sample space (McCullagh 2002). This means that no conceptual interpretation is yet included; a statistical model only includes the statistical dependencies among indicators (Moneta and Russo 2014). In contrast, a causal model has an additional causal interpretation because one or more of the parameters in the model reflect causal relations. A causal model could be seen as a representation of a real-world phenomenon, rather than a model that only plots the relations in the data. We are aware that there exists a large literature on the definition of a model and the distinction between different types of models in philosophy of science (e.g., van Fraassen 2008, Frigg and Hartmann 2012), but a discussion of this literature is beyond the scope of this paper.

Two models that represent different causal relations can be statistically equivalent. For example, many structural equation models are statistically equivalent to a model in which one of their structural relations is reversed (MacCallum, Wegener, Uchino and Fabrigar 1993). When two models are statistically equivalent, any observational dataset provides equal support for both of them; however, because they have different causal implications, it is possible to distinguish them via experimental intervention. Consider two statistically equivalent models: (I) $A \rightarrow B \rightarrow C$ and (II) $A \leftarrow B \rightarrow C$. Any possible observed correlation matrix between the variables A, B, and C will provide equal support for both models, though the arrow connecting A and B is reversed. The causal relation between A and B is therefore different for each model. Now suppose that A, B and C are measured at two points in time, and that between these time points, variable B is experimentally manipulated. Both the correlational structure at time point 1 and time point 2 will provide equal support for both models, but model (I) and (II) imply different expectations with respect to the increase or decrease in variable A as a result of the intervention on variable B. If intervening on B changes A to the extent predicted by model (II), that is evidence for model (II) over model (I).

Because causally different models can be statistically equivalent, some researchers and psychometricians argue that it is unjustified to make any causal interpretations of statistical models, and prefer the descriptivist approach. Against this view, we argue that it is often best to enrich a given statistical model with a causal interpretation

and thus, we argue for the interpretation and use of the CFM as a causal model. In the next section, we take issue with one of the main arguments for interpreting the CFM merely as a statistical model rather than a causal model. According to this view, a causal interpretation is not justified because factor models are typically estimated from the covariance structure of cross-sectional data without considering interventions on any of the variables over time.

3 Correlation does not entail causation

Factor analysis is typically applied to the correlation structure of cross-sectional data. However, a causal interpretation of the CFM implies that factor analysis infers causal relations between the latent variable and the indicators from observational data without having manipulated or intervened on anything. So how can factor analysis be used to hypothesize causal relations if it is based on correlational data?

Although correlations do not entail causality, causality does entail covariation. In addition, particular causal structures imply particular patterns of conditional independence among variables. For example, when $A \leftarrow B \rightarrow C$ is hypothesized, in which B is a common cause of A and C , then one expects covariation between A , B and C , and one expects that A and C are conditionally independent given B (Reichenbach 1956). Also, one expects that the correlation between A and C is smaller than the correlations between A and B and between B and C . These are all implications of a causal structure that can be verified from a correlation matrix. Similar expectations can be laid out for the reflective model (Figure 1), in which the latent variable is a common cause of its indicators: one expects the indicators to covary in a certain way, e.g., two indicators that are strongly correlated with each other should also share much variance with the other indicators. After all, indicators that consist of relatively more shared variance and less unique variance are more reliable indicators than those that consist of relatively little shared variance and more unique variance. Reliable indicators will also correlate more strongly with each other. Unlike in the 3-variable example cited earlier, however, in latent variable modeling it is impossible to test conditional independence given the common cause, because the common cause is unobserved. All in all, causation does lay out some expectations

for observational data, though one should always consider the possibility of alternative causal structures that imply similar expectations for the data (e.g., the expectations summarized for $A \leftarrow B \rightarrow C$ also correspond to the structure $A \rightarrow B \rightarrow C$).

These implied constraints of the common cause structure are tested in the fit of a CFM. Violation of local independence will result in poor model fit. As noted earlier, although the hypothesis of a common cause can be rejected because of the testable constraints it puts on the data (e.g., on tetrad rank constraints see Bollen and Ting 1993), the common cause structure cannot be verified from observational data alone because there will always exist alternative causal structures that can explain the observations in the data. But even though factor analysis cannot verify a common cause structure, hypothesizing such a structure justifies the use of factor analysis to measure the latent variable in the first place. After all, the constraints that are tested when fitting a CFM to the data are all implied by the common cause structure.

In the next section we elaborate on the descriptivist approach and then we proceed to developing three arguments in favor of a causal interpretation of the CFM.

4 What is the common factor in the descriptivist approach?

As described by Jonas and Markon (2016), the descriptivist approach entails that a common factor should not be interpreted as a real-world entity, but rather as mere shared variance that is nothing but a parsimonious summary of the data. From this perspective, latent variables are not postulated as concrete entities that have direct causal effects on their indicators, but rather as *summaries* of the covariance among indicators. When the latent variable is defined as the shared variance of a set of indicators, it cannot also be a common cause of the indicators. After all, the shared variance among indicators does not cause the shared variance among indicators; something cannot have caused itself. Thus, when the latent variable is the shared variance rather than being a cause of the shared variance, the arrows in Figure 1 pointing from the latent variable to the indicators should be interpreted as “is part of”. The latent variable *is part of* the indicators, by the logic that the variance in the indicators *consists of* shared

variance (arrow pointing from a common component to the indicator) and unique variance (arrow pointing from a unique component to the indicator). That is, when the latent variable is merely shared variance, the relation *intelligence* $\rightarrow$ *IQ test* should be interpreted in the same manner as *boys* $\rightarrow$ *people*: just as the set of boys is part of the set of people, intelligence is part of IQ test scores. Whether a variable is part of another variable or causes another variable, statistically speaking there is no difference between these two relations. More concretely, whether the latent variable is a part of the variance of the indicators (i.e., the variance that is shared among indicators) or is a cause of the indicators, both result in the same statistical model.

5 Toward a causal interpretation

Whereas descriptivists argue for a statistical reading of the factor model, we believe that a causal interpretation of the model is in many cases more appropriate. It is important to note here that our objections to the descriptivist approach are limited to cases in which the CFM is used as a measurement model rather than as a method for data reduction. When factor analysis is used as a data reduction method, the common factor can be used to predict other variables, but the common factor does not refer to anything outside the model itself. In data reduction, researchers are only interested in bringing a large set of variables down to a smaller set of dimensions. In other words, they are interested in a concise summary of the data, rather than in causal connections between latent variables and indicators. Again, in this case the obtained factor *is part of* the original set of indicators, rather than referring to a common cause that generated the data, and a causal reading is not sensible.

In the next subsections we present three arguments for a causal interpretation of a CFM over a merely statistical interpretation when the goal is to measure psychological constructs. We argue that (1) establishing causal relations conforms with most research questions in which the goal is to *explain* the correlations between indicators rather than merely summarizing them; (2) a causal interpretation of the CFM legitimizes *why* we are interested in the shared variance rather than in the unique variance of the indicators; and (3) a causal interpretation of the CFM legitimizes the assumption of local independence.

5.1 “Look, we found shared variance!”

Social scientists are typically interested in the best possible explanation for their observations. They want to know how and why observations occur the way they do. Borsboom et al. (2003) argued that latent variable models require a realist ontology in order to use them for establishing causal connections between latent variables and indicators. The choice for latent variable models, in which a set of indicators is assumed to covary *because* of this latent variable, already implies that one assumes this latent variable exists. Choosing a latent variable model therefore naturally implies a realist ontology. In other words, when the aim is to detect and theorize about a putative causal relation, it does not make sense to avoid a realist, causal interpretation all together.

To illustrate this point, we will use a study performed by Caspi et al. (2014) as a special case that in our view highlights why a causal reading is sensible. In their study, Caspi et al. (2014) “evaluate alternative hypotheses about the latent structure underlying 10 common mental disorders” (2014: 120). The authors conclude that a bifactor model, with three group factors and one general factor, best explains the structure of psychopathology, and they name this general psychopathology factor *the p factor*. The descriptivist view would only allow the researchers to conclude that their study provides support for shared variance among a set of disorders. With a causal interpretation though, their findings would provide support for a meaningful hypothesis, namely that the p factor causes mental disorders to covary, and explains comorbidity. Here, a descriptivist approach is simply unsatisfying and uninteresting.

Of course it is legitimate to claim that a single dimension is able to account for most of the covariation among disorders, resulting in a parsimonious description of the data. But if the goal is to explain the structure of psychopathology, this claim is greatly unsatisfying: shared variance as such has no explanatory value. In contrast, when the shared variance reflects a common cause, this common cause does explain the correlations between indicators, resulting in a better understanding of the observations and an opportunity to further the research programme by searching for the identity of the common cause. A common cause not only renders the indicators independent when accounted for, the indicators correlate *because* they share a common

cause. Explanation therefore has additional value to mere description, and in the end, establishing causal relations and finding explanations is essential to science. Thus, the p factor model has much greater theoretical import as a causal model than as a merely statistical model. Researchers should, however, try to verify whether this causal model for the structure of psychopathology is justified in each case.

5.2 Why shared variance at all?

The CFM distinguishes between shared variance (that which is shared by all indicators) and unique variance (variance which is unique to each indicator). Shared variance is attributed to the latent variable, while unique variance is attributed to residual influences, such as measurement error. The distinction between shared and unique variance, of which only the former is of interest in latent variable modeling, makes sense when a common cause underlies the indicators but is not sensible under alternative underlying causal structures. Put differently, the belief that a common cause underlies the indicators legitimizes that only shared variance is of interest rather than unique variance of the indicators.

Consider an example in which the indicators reflect some variable that is not a common cause of the indicators. For example, *healthy eating* is a variable that can be indicated by responses to questions like “do you eat a lot of junk food?”, “do you often eat fruit?”, and “how much sugar do you eat?”. These three items are all indicators of healthy eating, however, *eating healthy* is probably not a cause of these indicators. In this case it is not sensible to only take into account the variance that is *shared* by these items: the unique variance of the indicators is equally important to the construct *healthy eating*, even though these indicators are correlated and would likely result in a well-fitting reflective model.

The items *eating vegetables* and *avoiding sugar* are both relevant to the construct *healthy eating*, regardless of what causes them or how much of that variance is shared. That is, it does not matter *what* causes someone to eat a lot of vegetables (e.g., having a big vegetable garden) and whether that is the same cause as for not eating too much sugar (e.g., having strict parents). Both of these are indicators of healthy eating, including their unique variance. When a factor model is fit

to these data, variance due to these unique causes is relegated to the uninteresting unique variance component of the model. But arguably, the unique variance due to vegetable gardens and strict parents is no less relevant to *healthy eating* than the shared variance that may be due to causes such as a person's motivation to be healthy. As such, when the construct of interest is not a common cause, there is no theoretical reason to disentangle the shared and unique variance of the indicators. Doing so risks a biased representation of the construct of interest.

The above argument can just as well be applied to psychological constructs of which the exact nature is unclear. For example, *extraversion* is typically measured with a reflective model. Suppose that *extraversion* is just a summary of the scores on a set of extraversion items. In that case, why use the shared variance as a summary, rather than the unique variance or all variance? What is the justification for discarding unique variance if the shared variance holds no special status? There is none. The shared variance is just one way of reducing high dimensional data to fewer dimensions, and there are alternative techniques for data reduction that do not discard unique variance, e.g., Principal Component Analysis (PCA; James, Witten, Hastie and Tibshirani 2013). Thus, when the status of the construct is either known to be not a common cause (as in the *healthy eating* example) or is unknown (as in the *extraversion* example), the practice of interpreting only the shared variance among a set of items is not well supported.

In contrast, when one's theory explicitly states that the construct is a common cause of a set of items, it is immediately legitimate to interpret only the shared variance of the indicators because it is precisely this shared variance that must be due to the common cause. Returning to the Caspi et al. (2014) study, we can draw a similar conclusion. The reason why the authors are at all interested in the shared variance rather than the total variance, or the unique variance, is because they must have reasons to believe that the shared variance has certain explanatory power. Thus, the difference between a latent variable that is interpreted as merely shared variance and a latent variable that is believed to be a common cause of the indicators is that the former *results* from the distinction between unique and shared variance in a CFM, while the latter *legitimizes* this distinction.

5.3 Local independence

The principle of local independence is a fundamental assumption of the reflective model and comprises the idea that when a common cause underlies a set of indicators, conditioning on this common cause renders the indicators statistically independent (Borsboom et al. 2003). When factor analysis is performed on a dataset, it searches for a factor solution that meets this criterion. The assumption of local independence does not apply to all data reduction models. As stated before, PCA also results in a parsimonious summary of the data but in contrast to factor analysis, PCA does not use local independence as a criterion. PCA rather composes a variable that accounts for most of the variability in the indicators (James et al. 2013). In this process, PCA takes all variance into account, not only shared variance. Thus in the case of PCA, the indicators are not rendered independent, yet PCA provides a parsimonious summary of the data. A purely statistical interpretation of the reflective model does not legitimize the principle of local independence and therefore the use of the reflective model. In contrast, a causal interpretation of the reflective model implies that local independence should hold, so a latent variable model should be preferred when theory holds that a common cause is responsible for covariation among items. Getting back to the example of the p factor, the reason why the p factor is constructed in a way that it renders the disorders independent is because it is believed to reflect a psychopathology factor that forms a common cause to these disorders. A mere summary of the data is not bound to such a constraint on the data. So not only does a causal theory justify a focus on the shared variance, it also justifies the assumption of local independence. Instead of local independence merely *resulting from* absorbing shared variance, a common cause *explains why* local independence is assumed.

6 How to assess whether a common cause structure is correct

Using a CFM in the descriptivist framework does not require that a specific causal structure is hypothesized. It simply does not matter exactly what causal structure underlies the data for the latent variable

to be a useful summary of the data. For the causal interpretation, in contrast, it is of central importance that the model accurately represents the underlying causal structure. Attaching causal meaning to a CFM that does not accurately represent the actual causal system will result in incorrect predictions.

Consider four indicators that all influence each other, resulting in shared variance among these variables. A CFM can be applied to the correlation matrix of these variables, and, depending on the strengths of the causal relations between the indicators, this may result in a well-fitting model. In such a situation, the factor model may be useful as a parsimonious prediction model, but as a causal model this factor model makes incorrect predictions with respect to intervention on any of the indicators. For example, a causal interpretation of the factor model implies that intervention on any of the indicators would not affect the other indicators, because all covariation is purported to arise due to a single latent common cause. In contrast, the true underlying causal structure implies that intervention on any of the indicators leads to changes in the other indicators, because the shared variance truly reflects direct causal relations among indicators. So a causal interpretation of an estimated model can lead to false conclusions when the causal structure of the data generating model is not represented accurately.

This begs the question of how to assess whether the true data generating model has a common cause structure. How do we figure out whether the principle of the common cause applies in cases in which we do not have easy access to the real underlying causal structure? As stated before, ideally one would manipulate the common cause, and see what happens to the variation in the indicators (Borsboom et al. 2003). If this were possible, one could directly observe the effects of the manipulation and consequently give causal meaning to the model. But when the common cause is latent, such direct manipulation can be impossible. In some cases, however, an alternative causal model might offer different predictions for manipulations on the indicators. For example, if an alternative model states that the indicators cause each other, such that the shared variance is explained by causal influences between the variables rather than by a common cause, intervening on the indicators would differentiate the two models. Whereas a model with causal influences between indicators

implies that certain indicators change as a result of interventions on other indicators, a common cause model implies that indicators are not affected by other indicators. These diverging predictions enable the researcher to differentiate between such alternative models that explain the shared variance among a set of variables.

When intervention is not possible, empirical tests of the goodness of fit of the reflective model may shed light on whether the common cause model is a likely generating model for a particular data set. Although a researcher cannot determine what causal structure underlies cross-sectional data, the common cause model does put testable constraints on the covariance structure of the data. Widely used test statistics and fit indices for confirmatory factor analysis test the hypothesis that the covariance structure of the data matches the covariance structure implied by a common cause model. Tests based on the pattern of partial correlations in the data may also allow researchers to determine whether a common cause model is more likely to underlie the data than an alternative model that posits direct effects among indicators (van Bork, Rhemtulla, and Borsboom 2015, van Bork, Rhemtulla, Waldorp and Borsboom 2016).

We believe that it is especially important to think about the data generating process in the phase of test construction. As we mentioned before, researchers do not pick an arbitrary set of indicators but select those indicators that are hypothesized to be affected by the construct of interest. This is an important point that concerns the stage of test construction rather than test-analysis. Borsboom, Mellenbergh and van Heerden write:

a century of experience with test construction and analysis clearly shows that it is very hard to find out where the scores are coming from if tests are not constructed on the basis of a theory of item response processes in the first place (Borsboom, Mellenbergh and van Heerden 2004: 1067)

They continue with the conclusion:

Thus, it is suggested here that the issue may not be first to measure and then to find out what it is that is being measured but rather that the process must run the other way. It does seem that if one knows exactly what one intends to measure, then one will probably know how to measure it, and little if any validation research will be necessary. (Borsboom, Mellenbergh and van Heerden 2004: 1067)

Another possibility to give causal meaning to the factor model, is that a factor model can be posited as a plausible hypothesis for a certain phenomenon (Haig 2005, 2014). When theoretical considerations imply that the principle of the common cause describes a certain mechanism accurately, the factor model can be used as a method for theory generation. Specifically, exploratory factor analysis can be used for the generation of plausible theories, and through confirmatory factor analysis, these theories can then be evaluated. Through methods such as cross-validation and replication, the theory can gain additional support. This of course does not result in absolute certainty that the causal reading is in fact the true reading, but by choosing the models that have more explanatory value over those that are less explanatory, the best model is left standing. This way, it is possible to gather evidence for the hypothesis that a common cause structure is indeed applicable to the phenomenon in question.

All in all, an inevitable conclusion of the causal interpretation of the factor model we defend, is that the data generating process matters. Therefore, one should always try to assess whether it is plausible that the data generating mechanism has a common cause structure.

7 Conclusion

Even though the fit of a CFM to cross-sectional data does not test the causal structure that generated the data, if the goal is to measure psychological constructs and make meaningful claims about how such constructs relate to each other, we argue for a causal interpretation of the CFM rather than a merely statistical interpretation proposed by the descriptivist approach. First, a causal interpretation matches the scientific aim of explaining observed patterns in behavior rather than summarizing them. Second, a causal interpretation of the CFM legitimizes *why* the shared variance is of interest rather than the unique variance of the indicators. And lastly, a causal interpretation of the CFM legitimizes the assumption of local independence. A statistical reading of the model is not incorrect but does not provide an explanation for the observed data, it does not explain why shared variance rather than unique variance is of interest and it does not explain why local independence should hold.

We are aware that a causal interpretation of CFMs brings a host

of other problems. Reading a CFM causally does not suddenly turn a non-causal mechanism into one that is in fact causal. CFMs are only tools to posit plausible causal relations, and they need additional input in the form of causal knowledge or assumptions to establish what kind of causal relation it is. These models are only meaningful when users impose meaning onto them, and it is up to these users to then verify whether their hypotheses are confirmed or not. In other words, researchers construct plausible causal hypotheses, fit a model that imposes these hypotheses on the data, and then consider how the output supports their hypothesis. This process has a built-in buffer that balances formulating strong causal statements that refer to real-world mechanisms and being cautious about such statements.

We also acknowledge that there are situations in which a statistical non-causal interpretation of the CFM makes sense. When a common cause structure is unlikely to underlie the data, a purely statistical interpretation of the CFM can still be useful for prediction. For example, although the existence of general intelligence is often disputed, the *g*-factor can be used to merely predict school success, without having to establish its objective existence. But even for data reduction, other methods are available (e.g., PCA) that do not rely on latent variables, local independence and a distinction between shared and unique variance in the indicators. It is not entirely clear why factor analysis should be used rather than other data reduction methods such as PCA, when the first model implies constraints on the data that can not be justified from a descriptivist approach whereas the latter does not.

Ultimately, we plead for an understanding of the CFM as an explicit causal hypothesis that can be falsified. In situations where the data generating mechanism matters for the theory about a psychological construct (e.g., should depression be understood as a cause of its symptoms or as the name for a system of interacting symptoms?), treating the CFM as a summary rather than a hypothesis about this causal structure takes away the need to test the data-generating mechanism. Rather than concealing the urge to infer causal relations by restricting the interpretations of these models to descriptions of the data without reference outside of the data, we argue that one should dare to hypothesize. By making hypotheses explicit (e.g., “we hypothesize that construct A is the common cause of the indicators X_1 to X_j ”), they

are open to falsification. Additionally, interpreting the CFM causally stresses the need for theories about *how* differences in the latent variable result in differences in the responses; theories that we think are crucial for the measurement of psychological constructs.³

Riet van Bork
University of Amsterdam
Nieuwe Achtergracht 129B
1018 WS Amsterdam
The Netherlands
r.vanbork@uva.nl

Lisa Wijsen
University of Amsterdam
Nieuwe Achtergracht 129B
1018 WS Amsterdam
The Netherlands
l.d.wijsen@uva.nl

Mijke Rhemtulla
University of California, Davis
102E Young Hall
1 Shields Avenue
95616 Davis CA
mrhemtulla@ucdavis.edu

References

- Asmundson, Gordon. J.G.; Frombach, Inge; McQuaid, John; Pedrelli, Paulo; Lenox, Rebecca; and Stein, Murray B. 2000. Dimensionality of posttraumatic stress symptoms: a confirmatory factor analysis of DSM-IV symptom clusters and other symptom models. *Behaviour Research and Therapy* 38: 203–14.
- Bollen, Kenneth. A. 2002. Latent variables in psychology and the social sciences. *Annual Review of Psychology* 53: 605–34.

³ We hereby declare that all authors contributed in a significant way to the manuscript. R.B. and L.D.W. wrote the draft of this manuscript and prepared the manuscript for submission. M.R. provided critical revisions and wrote parts of the manuscript. All authors have read and approved the final version of the paper for submission. There are no conflicts of interest. This statement applies to all authors. No other people than the authors contributed to this manuscript.

- Bollen, Kenneth. A. 2011. Evaluating effect, composite, and causal indicators in structural equation models. *MIS Quarterly* 35: 359–72.
- Bollen, Kenneth; and Lennox, Richard. 1991. Conventional wisdom on measurement: A structural equation perspective. *Psychological Bulletin* 110: 305–14.
- Bollen, Kenneth A.; and Ting, Kwok-Fai. 1993. Confirmatory tetrad analysis. *Sociological Methodology* 23: 147–75.
- Borsboom, Denny; Mellenbergh, Gideon. J.; and van Heerden, Jaap. 2003. The theoretical status of latent variables. *Psychological Review* 110: 203–19.
- Borsboom, Denny; Mellenbergh, Gideon. J.; and van Heerden, Jaap. 2004. The concept of validity. *Psychological Review* 111: 1061–71.
- Caspi, Avshalom; Houts, Renate. M.; Belsky, Daniel. W.; Goldman-Mellor, Sidra. J.; Harrington, HonaLee; Israel, Salomon; and Moffitt, Terrie. E. 2014. The p factor one general psychopathology factor in the structure of psychiatric disorders? *Clinical Psychological Science* 2: 119–37.
- Diamantopoulos, Adamantios; Riefler, Petra; and Roth, Katharina. P. 2008. Advancing formative measurement models. *Journal of Business Research* 61: 1203–18.
- Edwards, Jeffrey. R.; and Bagozzi, Richard. P. 2000. On the nature and direction of relationship constructs and measurement. *Psychological Methods* 5: 155–74.
- Frigg, Roman; and Hartmann, Stephan. 2006. Models in science. In *The philosophy of science: An encyclopedia*, ed. by S. Sarkar and J. Pfeifer, 740–9. New York, NY: Routledge.
- Haig, Brian. D. 2005. An abductive theory of scientific method. *Psychological methods* 10:371–88.
- Haig, Brian. D. 2014. *Investigating the Psychological World: Scientific Method in the Behavioral Sciences*. Cambridge, MA: MIT Press.
- James, Gareth; Witten, Daniela; Hastie, Trevor; and Tibshirani, Robert. 2013. *An Introduction to Statistical Learning*. Vol. 112. New York, NY: Springer.
- Jonas, Katherine. G.; and Markon, Kristian. E. 2016. A descriptivist approach to trait conceptualization and inference. *Psychological Review* 123: 90–6.
- MacCallum, Robert. C.; Wegener, Duane. T.; Uchino, Bert. N.; and Fabrigar, Leandre. R. 1993. The problem of equivalent models in applications of covariance structure analysis. *Psychological Bulletin* 114: 185–99.
- McCrae, Robert. R.; and Costa, Paul. T. 1987. Validation of the five-factor model of personality across instruments and observers. *Journal of Personality and Social Psychology* 52: 81–90.
- McCullagh, Peter. 2002. What is a statistical model? *The Annals of Statistics* 30: 1225–67.
- Moneta, Alessio; and Russo, Federica. 2014. Causal models and evidential pluralism in econometrics. *Journal of Economic Methodology* 21: 54–76.
- Musek, Janek. 2007. A general factor of personality: evidence for the big one in the five-factor model. *Journal of Research in Personality* 41: 1213–33.
- Reichenbach, Hans. 1956. *The Direction of Time*. Los Angeles: University of California Press.
- van der Maas, Han. L. J.; Dolan, Conor. V.; Grasman, Raoul. P.; Wicherts, Jelte. M.; Huizenga, Hilde. M.; and Raijmakers, Maartje. E. 2006. A dynamical model of general intelligence: the positive manifold of

- intelligence by mutualism. *Psychological Review* 113: 842–61.
- Spearman, Charles. 1904. 'General Intelligence,' objectively determined and measured. *The American Journal of Psychology* 15: 201–92.
- van Bork, Riet; Rhemtulla, Mijke; and Borsboom, Denny. 2015. Latent variable and network model implications for partial correlation structures. Presentation for the 80th Annual Meeting of the Psychometric Society (IMPS), Beijing, July 2015.
- van Bork, Riet; Rhemtulla, Mijke; Waldorp, Lourens. J.; and Borsboom, Denny. 2016. Distinguishing latent variable models and network models. Presentation for the 28th Annual Convention of the Association for Psychological Science (APS), Chicago, May 2016.
- van Fraassen, Bas. C. 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.