AN OVERVIEW OF CAUSAL DIRECTED ACYCLIC GRAPHS FOR SOCIAL WORK RESEARCHERS[1]

Michael A. LEWIS[2,3]
Hunter College School of Social Work, New York, USA

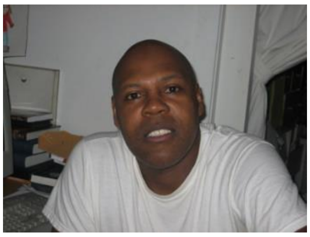

E-mail: mle0050@hunter.cuny.edu

Abstract:
Given the mission of social work to improve people's lives by intervening in ways to enhance their well-being, social work researchers are interested in the causal effects of various types of social interventions (de Anda, 2007; Ohmer, and Jorr, 2006; Hawkins, 2006). Statisticians, econometricians, and other experts in quantitative methods tend to view randomized controlled trials (RCTs) as the "gold standard" when it comes to estimating causal effects, and there's a key reason for this--RCTs are based on randomly assigning units to at least two different intervention groups. Random assignment tends to result in the groups being balanced on variables that may have a causal relationship with the outcome of interest other than the intervention which is the focus of a given study. At least one of these "balanced variables" may also be causally related to the intervention itself. "Balanced" means that the average values of these other variables are equal across the different groups.

Key words: directed acyclic graphs, DAG, social work researchers, randomized controlled trials, RCT

Introduction

Variables that affect both the outcome and the intervention under study are called *confounding variables*. The fact that randomization tends to generate balance on potentially confounding variables means that if a researcher observes a difference (or differences) across intervention groups on the outcome of interest, they can be confident that this difference is due to the intervention being evaluated.

The problem many social work researchers face is that RCTs are often neither ethical nor feasible, yet these researchers are still interested in estimating causal effects. Perhaps unbeknownst to many social work researchers, computer scientists have been concerned with causality as well. This concern has come largely out of the area of *artificial intelligence* (AI). Some of the leaders in this field are UCLA electrical engineer/computer scientist Judea Pearl and Carnegie Mellon faculty members Peter Sprites, Clark Glymour, and Richard Scheines (SGS).

AI has focused on trying to program thinking in computers. One feature of thinking is thinking about causal relationships. Over the years Pearl (2000) and SGS (2000) have been working on the problem of how to represent thinking about such relationships in computers, and they've developed mathematical tools for modeling such thinking. These mathematical tools are called causal Directed Acyclic Graphs (DAGS) and are the subject of this paper.

Causal DAGs aren't only useful to researchers in AI but are to quantitative researchers more generally who're interested in estimating causal effects using non-experimental or observational data. This is so for two reasons. First, causal DAGs provide a way of precisely, yet intuitively, specifying a researcher's causal assumptions. That is, they provide a language researchers can use to clearly state their assumptions about what is causing what. Such clarity allows other researchers to critically review those assumptions. Precise statements of assumptions and constructive critiques of them is a big part of the enterprise of science. Thus, causal DAGs can serve as an additional resource in a scientific approach to quantitative social work research.

Second, causal DAGs provide rules for addressing the problem of confounding when a researcher is faced with non-experimental or observational data. These rules can be drawn upon to guide specification of statistical models for use in social work research. For example, causal DAGs can provide guidance regarding which variables ought and ought not to be included in a *regression or propensity score matching model*, assuming that the causal assumptions encoded in a given DAG are true.

In order to make the fairly abstract ideas discussed in this paper more concrete, I'll refer to an example related to social work, more specifically social policy. I chose a social policy example because this is my substantive area of expertise. Hopefully, it'll be clear how causal DAGs can be used in other areas of social work research as well.

The Basic Income and Economic Security

As any social worker will know from either professional experience or coursework in social policy, the U.S., like all "advanced industrialized societies," is considered a welfare state. One of the hallmarks of welfare states is that, under certain conditions, they provide income to some of their residents that these residents don't have to work for. By "work" I mean sell one's labor in the formal labor market in return for a wage. For the purposes of this paper, I'll refer to such support as government provided non-wage (GNW) income.

Examples of GNW income programs in the U.S. are Temporary Assistance for Needy Families (TANF), Supplemental Security Income (SSI), Social Security, and Unemployment Insurance (UI). In social work we tend to focus on the degree to which these programs promote social justice or meet people's needs. Following this concern, I'll focus in this paper on what I'll call Economic Security, a variable meant to capture how economically or financially secure one feels.

A particular type of GNW a number of social scientists have written about (Widerquist, Lewis, and Pressman, 2005) is called a Basic Income (BI). BI is essentially an **unconditional** minimum income provided by government. That is, some level of government would grant all citizens or residents a minimum income without requiring recipients to engage in certain types of behavior in return for the benefit (e.g., enrolling in a workfare program, forming "proper" family structures, voting in elections, etc.). The policy appears to be based on the idea that people have a right to at least a subsistence level income, regardless of whatever else they might or might not be doing.

In this paper, I'll use a hypothetical study of the causal effect of BI on Economic Security to provide an overview of causal DAGs. Imagine that some political jurisdictions in the U.S. chose to enact a law granting their residents a BI while other jurisdictions did not. Those who ended up receiving or not receiving a BI weren't randomly assigned to their "treatment" groups. Instead, normal processes operating in day to day politics led to these differences in treatment. Imagine that we end up with non-experimental data for a sample of residents from various jurisdictions. For each resident we have data on whether they received BI, their level on the Economic Security variable, and their levels on other variables to be defined below. We think BI is a cause of Economic Security and we'd like to estimate this causal effect.

What is meant by Cause?

Perhaps the most appropriate place to begin is with the concept of *cause*. In a paper co-authored with one of his students (Chen and Pearl, 2013), Pearl discusses cause in terms of probability distributions as well as expected values of outcomes. As those familiar with mathematical statistics know, the term "expected value" can be thought of as synonymous with "average". Since social work researchers often use regression models in their work and since such models typically focus on expected values of outcomes, I'll rely on the expected value based conception of cause.

According to Chen and Pearl (2013), x is a cause of y if intervening to change the value of x results in a change in the expected value of y. For example, suppose we take BI as x and Economic Security as y. Then BI would be a cause of Economic Security if intervening to change the value of BI resulted in a change in the expected value of Economic Security.

To be a bit more concrete, let BI take two possible values—$15,000 per year and $0 per year. Since no political jurisdiction in the U.S. currently has a BI program, let's take this to mean that the BI variable currently takes on the value $0 per year for everyone in the country.[4] Suppose currently the average level of Economic Security is 5 units. Now let's say some jurisdictions enact a BI grant of $15,000 per year to all their residents and, as a result, these residents' average Economic Security increases from 5 to 25 units. This would indicate that BI is a cause of Economic Security.

Note that the idea of intervening or doing something to change the value of some variable is crucial to the notion of causality. Also note that an intervention isn't necessarily tied up with an RCT, although it can accommodate RCTs. We could intervene by randomly assigning some people to receive a BI of $15,000 per year and others to receive one of $0 per year. Or, as in our hypothetical "natural experiment," laws could be passed stipulating that residents receive a BI of $15,000 per year. As long as the BI intervention changes the expected value of Economic Security, **by definition** we have a causal effect.

Identification versus Estimation

In order to understand the role causal DAGs might play in social work research, one must understand the distinction between *identification* and *estimation*. In Elwert (2013, p. 247 http://www.ssc.wisc.edu/soc/faculty/pages/docs/elwert/Elwert%202013.pdf) we find the following passage: "identification…determines whether and under what conditions, it is possible to strip an observed association of all its spurious components."

Social work researchers are no doubt familiar with the concept of correlation and its relationship to causality. Two variables, let's call them x and y, are correlated if x causes y, y causes x, or they share a common cause. What Elwert is getting at in the passage above is the idea that identification is related to conditions under which one can isolate "causal correlation" from "non-causal correlation".

Estimation has to do with statistical methods, such as OLS regression, Two-Stage Least Squares regression, and others used to obtain estimates of causal effects. Identification "comes before" estimation in the sense that one needs to determine if causal association can be isolated from non-causal association before trying to estimate the magnitude of causal association.

As will become clear later in this paper, part of the utility of causal DAGs to social work researchers is that they provide clear rules for isolating causal from non-causal correlation, **assuming the causal relationships encoded in a given DAG actually hold true**. That is, causal DAGs tell us how to identify the causal effects encoded in a given DAG and, in that way, they provide guidance for the process of model specification and estimation.

Elements of Causal DAGs

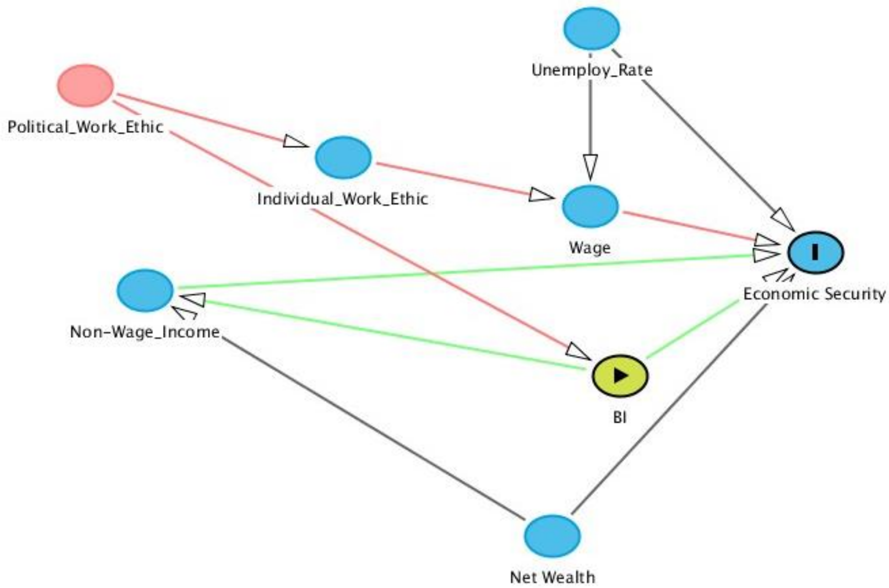

Figure 1. Causal DAG Representing Causes of Economic Security

Figure 1 will be used to explain the basic elements of causal DAGs. It was created using an open source online program for drawing causal graphs called Dagitty (Textor, Hardt, and Knüppel, 2011). Dagitty was explicitly written to implement the ideas of Pearl and others regarding DAGs. The color scheme in the diagram is done automatically in Dagitty as a way of labelling the role of certain variables in the diagram. Since knowing the details of that color scheme isn't necessary for our purposes, I'll have nothing to say about them.

I used Dagitty to draw Figure 1 on the basis of assumptions I made about the causal relationship between BI and Economic Security as well as how those two variables are causally related to a set of other variables. To do so I drew largely on economic theory and intuition. For example, on the basis of economic theory I'm assuming that the unemployment rate in an area where a person resides can influence the wage they receive for two reasons: 1) the unemployment rate can influence whether or not they receive a wage at all and 2) if they do receive a wage, their wage level can be affected by the unemployment rate because the more people there are searching for jobs the less employers might be able to pay current workers (because higher unemployment might give employers an advantage over employees when it comes to wage bargaining).

Drawing a causal DAG is the first step involved in using such tools to aid social work research. That is, a researcher must draw a DAG which encodes the assumptions they're[5] making about what causes what. However, causal DAGs must be drawn according to certain rules.

The diagram in Figure 1 is an example of what mathematicians call a *graph*. In fact, the "G" in DAG stands for the word "graph". This is a different use of the term from what many readers might be accustomed to. That is, the term "graph" doesn't refer to a coordinate system along with a curve in that system. Instead it is a set of vertices or nodes along with edges or arrows connecting those vertices. The set of vertices in Figure 1 is {Unemploy Rate, Political Work Ethic, Individual Work Ethic, Wage, Economic Security, Non-Wage Income, Net Wealth, BI}. The vertices in a causal DAG represent variables. The edges or arrows in a causal DAG represent causal relationships.

Notice that there is an arrow "coming out of" BI and "going into" Economic Security. This encodes the assumption that BI causes or has a causal effect on Economic Security. In general, an arrow coming out of x and going into y means that x is assumed to cause y. The use of arrows to represent causal relationships is what makes these kinds of graphs causal graphs.

"Directed Acyclic" means that there can be no cycles or loops in a DAG. Perhaps the best way to understand this is to see a graph where there is a cycle:

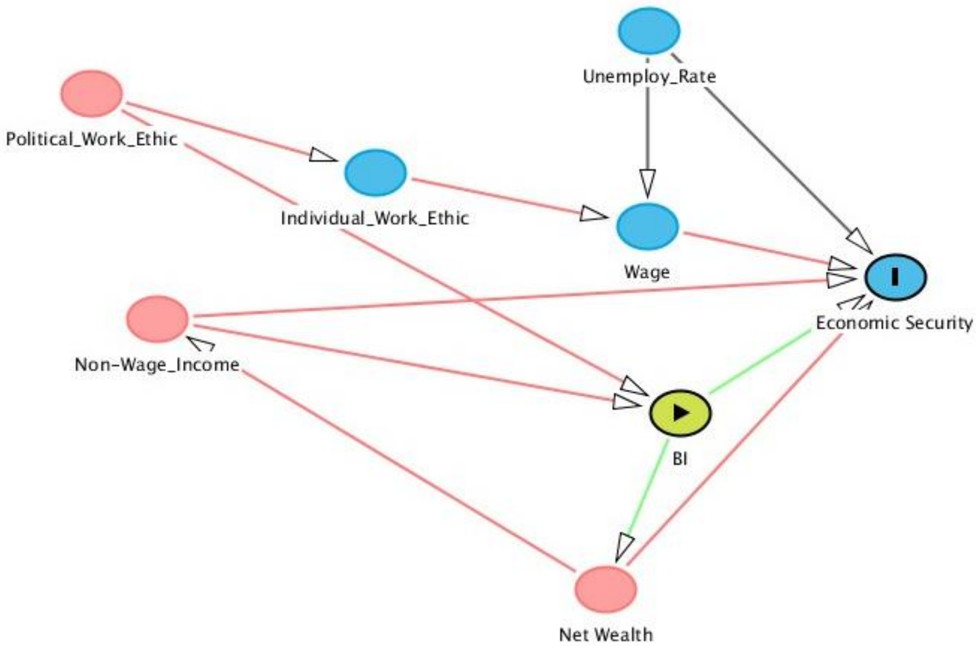

Figure 2. Graph Representing Causes of Economic Security with a Cycle

Notice that there is an arrow going from Non-Wage Income to BI from BI to Net Wealth and from Net Wealth back to Non-Wage Income. This is an example of a cycle because a "train" of causal connections circles back to where it started. Another example of a cycle would be if there were an arrow going from BI to Economic Security and one going from Economic Security back to BI. Again, causal DAGs don't allow these cycles.

Readers familiar with structural equation models might be thinking that causal DAGs are nothing but structural equation models without latent variables. This is true. Structural equation models (as well as path models) can be thought of as special cases of causal DAGs.[6] However, structural equation models are based on specific assumptions regarding functional form, such as linearity, while causal DAGs are not (Greenland, Pearl, and Robins, 1999). For example, the assumption that BI causes Economic Security encoded in Figure 1 makes no assumption at all about the functional form of that causal relationship.

Here are some of the details regarding the variables in Figure 1:

VARIABLE	**VALUES**
BI	1 = $15,000/year 0 = 0$/year
Economic Security	Scale which measures how economically/financially secure one feels
Non-Wage Income	Weekly income in U.S. dollars from sources other than work
Individual Work Ethic	scale which measures the degree of effort expended at work[7]
Political Work Ethic	scale which measures elected officials'(within the study participant's jurisdiction) commitment to idea that able bodied should work for their subsistence[8]
Net Wealth	Total Assets – Total Liabilities measured in U.S. dollars
Unemployment Rate	Jurisdictional unemployment rate measured as a percentage
Wage	Weekly wage measured in U.S. dollars

Table 1. Values of Variables from Figure 1.

As I stated earlier, whenever an arrow comes out of one variable and goes into another, the variable the arrow comes out of is assumed to be a cause and the variable it goes into is assumed to be an effect. More specifically, this is a case of a *direct causal effect*.

Now let's take three variables x, y, and z, and suppose that an arrow comes out of x and goes into y. Also suppose that an arrow comes out of y and goes into z. In this case the causal DAG would be encoding the assumption that x is a direct cause of y, y is a direct cause of z, and x is an *indirect cause* of z. Readers familiar with the use of path analysis and structural equation modeling might recognize this as y being a "mediator" of the causal relationship between x and z. I should also add that the indirect causal effect of one variable on another can "work through" more than one intervening or mediating variable. The total causal effect of one variable on another is the sum of its direct and indirect causal effects on that variable.

Figure 1 also includes examples of indirect causal effects. Notice that Political Work Ethic is assumed to be an indirect cause of Economic Security. The idea is that if we could intervene to change the degree to which a set of elected officials in a given jurisdiction are committed to the idea that all able bodied people should work for their subsistence this would result in a change in the expected value of work effort expended by residents in that jurisdiction. This change in the average level of work effort would result in a change in the average wage, and this, in turn, would result in a change in the expected value of Economic Security.

Readers may not think much of the assumptions I've spelled out regarding the indirect causal effect of Political Work Ethic on Economic Security. But that's beside the point. The point is that a causal DAG should encode whatever assumptions a researcher is making about causal relationships. These types of graphs provide a precise yet intuitive way of doing so with the benefit that others can easily see what a researcher is thinking regarding what causes what.

The next key concept relevant to drawing causal DAGs is *path*. A path in a causal DAG is a sequence of variables connected to each other by arrows.[9] A *directed path* between two variables is one where "travel" along arrows between the variables takes place so that travel is always from the tails to the heads of these arrows. Consider Figure 1 again. The path from Political Work Ethic to Individual Work Ethic to Wage to Economic Security is a directed path between Political Work Ethic and Economic Security.

An *undirected path* between two variables is one where travel along arrows takes place ignoring the direction of the arrows along the path. In Figure 1 the path from BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security is an example of an undirected path. The path from Political Work Ethic to Individual Work Ethic to Wage to Economic Security is an undirected path between Political Work Ethic and Economic Security. That is, "ignoring the direction of arrows" means that it doesn't matter which way arrows along the path are pointing. So a directed path, such as Political Work Ethic to Individual Work Ethic to Wage to Economic Security, is a special case of an undirected path.

Having defined directed and undirected path, I can state the no cycles constraint, discussed earlier, a little differently. Causal DAGs aren't allowed to have a directed path that ends with the variable it started with. The path from Non-Wage Income to BI to Net Wealth back to Non-Wage Income is an example of such a directed path. So it isn't allowed in a causal DAG.

Collider and *descendant* are two other concepts crucial to understanding causal DAGs. A collider is a variable which has two arrows coming from two different variables going into it. In Figure 1 Wage is a collider because arrows from both Individual Work Ethic and Unemploy Rate go into it. To understand the notion of descendant, assume that there is a directed path from x to y. There may or may not be variables between x and y. Since the path starts at x and ends at y, y is called a descendant of x. In Figure 1 Individual Work Ethic is a descendant of Political Work ethic and so is Wage.

Next I need to define the notion of *conditioning*. Other terms for conditioning are *subgroup analysis* and *stratification* (Morgan and Winship, 2007). Within the context of causal DAGs, conditioning occurs when an analyst examines the causal relationship between two variables **for given values** of at least one other variable, the variable which is being conditioned on.

For example, suppose in Figure 1 that Individual Work Ethic can only assume 10 possible values, the values 1-10 inclusive. We could estimate the causal relationship between Wage and Economic Security **only** for those with a level of 1 on the Individual Work Ethic variable, only for those with a level of 2 on it, only for those with a level of three on this variable, etc. This would be an example of conditioning on Individual Work Ethic.

It's also possible for two or more variables to be simultaneously conditioned on. For example, in the present case we might examine the causal relationship between BI and Economic Security for given values of both Individual Work Ethic and Political Work Ethic simultaneously.

Confounding, Regression Models, Propensity Score Matching, and Causal DAGS

In this section I'll spend a good deal of time discussing the identification of causal effects. I'll assume that we have a random sample of some population of interest and that each member of the sample has complete data on all the variables in Figure 1. This is to assume away all issues having to do with lack of a random or probability sample as well as issues having to do with missing data. I assume away these matters not because they are unimportant but because the purpose of this paper is to introduce readers to causal DAGs. Bringing in these other issues would needlessly complicate matters, given this purpose.

Regression models are ubiquitous in social work research. A typical example of the use of such models can be found in Studts, Stone, and Barber's paper *Predictors of Access to Health Care Services among Groups of TANF Recipients in Kentucky* (2014). They report findings from a set of logistic regression models involving predictors of access to health care. This use of the term "predictor" is ambiguous because it can be used when one is concerned with correlations between variables or causal relationships between them.

From about the end of November to the beginning/mid-December in the U.S., I suspect there's a strong correlation between the number of Christmas trees bought per day and the number of times the song Rudolf the Red Nosed Reindeer is played per day. So number of Christmas trees bought per day would be a good predictor of number of times Rudolf is played per day, but that wouldn't necessarily mean that the buying of Christmas trees causes the playing of Rudolf the Red Nosed Reindeer.

Social work researchers, including Studts, Stone, and Barber, are no doubt aware of this issue. That, in fact, is why they so often turn to regression models. Such models are thought to address the "correlation does not imply causation problem" by including independent variables other than the one of primary interest as "control variables." These control variables are thought to potentially confound the causal relationship between the independent variable and outcome of interest, and including them in a regression model is believed to take care of this problem. What causal DAGs can provide is guidance regarding the conditions under which including such control variables does take care of it, assuming the causal relationships encoded in a given causal DAG actually hold true.

In addition to regression models, social work researchers have made increasing use of a set of tools from statistics called propensity score matching (PSM) (Guo, Barth, and Gibbons, 2006). In fact, two of the field's most highly regarded quantitative methodologists, Shenyang Guo and Mark W. Fraser recently wrote a book (2010) on this topic. PSM is often used to deal with what econometricians call the *selection problem*.

One of the problems researchers have when it comes to drawing causal conclusions on the basis of non-experimental data is that when the research participants are human beings they often self-select into treatment. This means that there could be at least one factor which causes both selection into treatment and a certain outcome. The result is that treatment and outcome variables will be correlated, and a researcher may mistake this correlation for a causal relationship.

As an example of the selection problem, suppose we're interested in the causal effect of Diet (magnitude of calories one consumes over a period) on Exercise (magnitude of exercise one engages in over a period). We might run into a selection problem because Depression (how depressed one is) may influence both Diet and Exercise.

PSM addresses the selection problem by using observed variables to model selection. That is, values on observed variables are used to estimate the probability that a case ends up "receiving" the intervention or treatment of interest. These estimated probabilities are called *propensity scores*. Cases are then matched on propensity scores and an analysis, often a type of regression model, is run in an effort to estimate the causal effect of some variable of interest. There're a number of ways propensity scores can be used for matching, and Guo and Fraser (2010) have a nice discussion of some of the different methods involved. A detailed discussion of these methods is beyond the scope of this paper. What is relevant is the fact that causal DAGs can provide guidance regarding what variables ought to be included in a propensity score model. In order to understand how causal DAGs can aid in the specification of regression and propensity score models, one first needs to understand the concepts *backdoor path*, *unblocked* or *open backdoor path*, *blocked* or *closed backdoor path*, *intercepting a path*, *confounding path*, and *confounders*.

A backdoor path from x to y is a path which has an arrow going into x. (Morgan and Winship, 2007; Greenland, Pearl, and Robins, 1999). A path is intercepted by a variable if that variable is on the path but isn't one of the variables on either end of the path (Greenland and Pearl, 2009). For example, the BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security path is intercepted by Political Work Ethic because Political Work Ethic is on the path but isn't on either end of the path. A path is intercepted by a set of variables if the members of that set intercept the path. So the BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security path is intercepted by the set {Political Work Ethic, Individual Work Ethic, Wage}.

Now suppose we have a backdoor path between x and y. This path would be blocked or closed if it has at least one collider and unblocked or open if it doesn't (Greenland, Pearl, and Robins, 2000). Take a look at Figure 1 again:

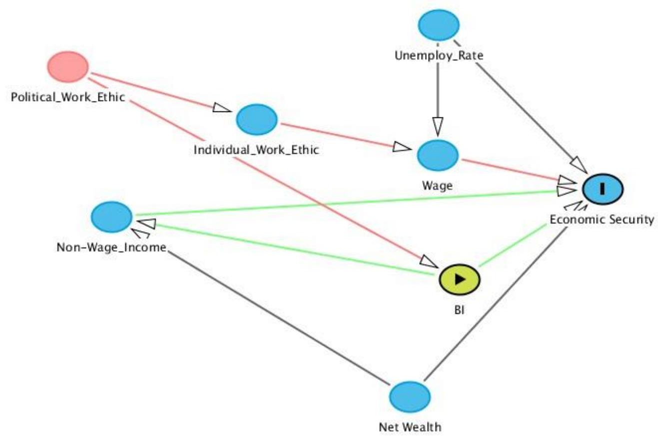

The path BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security is a backdoor path because BI and Economic Security are connected by a path that begins with an arrow going into BI. The path BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate to Economic Security is also a backdoor path for the same reason.

The BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security path is unblocked because it doesn't contain a collider along it. The BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate to Economic Security path is blocked because it contains a collider, namely Wage.

A confounding path is an unblocked or open backdoor path, and the variables which intercept that path are confounders (Greenland and Pearl, 2009). In the BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security path, members of the set {Political Work Ethic, Individual Work Ethic, Wage} are confounders.

Given the causal assumptions encoded in Figure 1 and the rules of causal DAGs, the existence of the confounding path referred to in the previous paragraph means that the relationship between BI and Economic Security is confounded by Political Work Ethic, Individual Work Ethic and Wage. So in order to identify the causal effect of BI on Economic Security, we need to do something about these confounders. This brings us to the notion of conditioning on a set of variables to block a confounding path.

Suppose **Z** is a set of confounders along a confounding path. Conditioning on the variables in Z blocks this path if:

1. any node along the path which has an arrow coming out of it is a member of Z or
2. the path has at least one collider which is not a member of Z and no descendant of any collider is a member of Z (Greenland and Pearl, 2009).

Conditioning on a set of variables to block an otherwise confounding path is the causal DAG version of controlling for confounding variables (or confounders) to identify a causal effect if interest. Thus, if the assumptions encoded in a particular causal DAG hold true, there's a set of variables Z which intercepts **all** confounding paths between x and y, and conditioning on Z blocks all backdoor paths between x and y, then the causal effect of x on y can be identified by conditioning on the elements of **Z.** (Pearl, 2000 and Greenland and Pearl, 2009).

Going back to Figure 1, notice again that the backdoor path BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security is intercepted by Political Work Ethic Individual Work Ethic and Wage. The backdoor path BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate, to Economic Security is intercepted by Political Work Ethic, Individual Work Ethic, Wage, and Unemploy Rate. These are the only backdoor paths in the DAG in Figure 1.

The second backdoor path in Figure 1 can be blocked without conditioning on anything since Wage is a variable along that path and Wage is a collider. In this case we'd say that we're conditioning on the empty set Z = {} because the empty set is the one with no members. The second backdoor path can also be blocked by conditioning on Z = {Political Work Ethic, Individual Work Ethic, Unemploy Rate} since all the variables in this set have arrows coming out of them along that path.

The first backdoor path in Figure 1 can be blocked by conditioning on Z = {Political Work Ethic, Individual Work Ethic} since all the variables in this set have arrows coming out of them along that path. Notice that Political Work Ethic and Individual Work Ethic block both (all) paths between BI and Economic Security. Hence, the causal effect of BI on Economic Security can be identified by conditioning on Political Work Ethic and Individual Work Ethic.

To see how conditioning on a collider variable enters into all this, take a look at Figure 1 again. The path BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate to Economic Security contains a collider variable, namely Wage. Suppose in an effort to identify the causal effect of BI on Economic Security, we conditioned **only** on Wage. That is suppose Z = {Wage}. Now take a look again at the second criterion for blocking a confounding path by conditioning on a set of variables[10]: the path has at least one collider which is **not** a member of Z and no descendant of any collider is a member of Z. Since the only member of Z is Wage and Wage is a collider, the BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate to Economic Security is **unblocked** or **opened** by conditioning on Wage. That is, by conditioning on Wage we create a confounding path which is exactly what we don't want to do. Hence, if a researcher is interested in identifying a causal effect while addressing the problem of confounding, one shouldn't condition only on a collider variable.

I said earlier that conditioning on the variables Political Work Ethic and Individual Work Ethic would allow us to identify the causal effect of BI on Economic Security because these variables block all backdoor paths between BI and Economic Security. The set of two variables Z = {Political Work Ethic, Individual Work Ethic} is an example of what's called a *sufficient set*. This set of variables is sufficient in the sense that conditioning on them would be enough to identify the causal effect of BI on Economic Security. In general a set of variables is sufficient for identifying the causal effect of x on y if conditioning on those variables blocks all backdoor paths between x and y. A set of variables is *minimally sufficient* for identifying the causal effect of x on y if no proper subset of the set is sufficient (Greenland and Pearl, 2009 and Pearl, 2000).[11]

The set of variables {Political Work Ethic, Individual Work Ethic} is a sufficient set but it isn't minimally sufficient. To see this look again at the first criterion for blocking a backdoor path by conditioning on a set of variables along that path: any node along the path which has an arrow coming out of it. In Figure 1 the variable in set Z = {Political Work Ethic} has an arrow coming out of it along both backdoor paths. So conditioning on it identifies the causal effect of BI on Economic Security. The same would be true by conditioning on the variable in Z = {Individual Work Ethic}. The sets Z = {Political Work Ethic} and Z = {Individual Work Ethic} are both subsets of {Political Work Ethic, Individual Work Ethic} so {Political Work Ethic, Individual Work Ethic} isn't a minimally sufficient set.

The relevance of all this talk about conditioning on confounders to block backdoor paths is this: if we were willing to make the causal assumptions encoded in Figure 1 and willing to assume a linear (in parameters) functional form, then conditioning on a set of variables would amount to including them in an OLS regression model, along with the causal variable of interest (BI in this case), or including them in a model of propensity scores and subsequently running an appropriate statistical model on the basis of those scores. Let's take the OLS regression example first. Given the causal assumptions encoded in Figure 1, the effect of BI on Economic Security can be estimated by the following Ordinary Least Squares (OLS) regression model:

$$\text{Economic Security} = a + b_1\text{BI} + b_2\text{Political Work Ethic} + b_3\text{Individual work ethic}$$

The coefficient b_1 would be an estimate of the total causal effect of BI on Economic Security, controlling for the effects of Political Work Ethic and Individual Work Ethic on Economic Security. That is, b_1 would estimate how much the average level of Economic Security would change **as a result of** an intervention to change BI by one unit (recall the "units" of BI are 1 = $15,000 per year and 0 = $0 per year). This is the sense in which causal DAGs can provide guidance regarding how to take care of the "correlation does not imply causation problem" through regression modeling.

Given what I said about sufficiency and minimal sufficiency, the following models could also be run to estimate the effect of BI on Economic Security:

$$\text{Economic Security} = a + b_1\text{BI} + b_2\text{Political Work Ethic}$$

$$\text{Economic Security} = a + b_1\text{BI} + b_2\text{Individual Work Ethic}$$

Earlier I discussed propensity score matching as a way of addressing the selection problem. What causal DAGs add to propensity score methodology is guidance on what variables should be included in a propensity score model. Those variables are the ones which appear in a given set **Z**—the ones which block all backdoor paths between a cause and effect variable of interest (Morgan and Winship, 2007). Thus, if one were interested in using propensity score matching to estimate the causal effect of BI on Economic Security, one could do the following.

First, use Political Work Ethic and/or Individual Work Ethic to model selection into treatment, BI in the present case. Second, use the saved propensity scores along with an appropriate method of matching, such as *nearest-neighbor*, *kernel*, etc., in a model to estimate the effect of BI on Economic Security. Both Guo and Fraser (2010) and Morgan and Winship (2007) discuss various matching methods for use with propensity scores.

OLS regression models and propensity score matching are often discussed as two different methodologies and in some respects they are. The insight that causal DAGs provide, however, is that in a sense they are equivalent ways of adjusting for or controlling for confounding variables.

Much of this section of the paper has been fairly theoretical and abstract. What I want to do now is discuss two things researchers can do if they're interested in applying these ideas. One relates to determining if confounding is present in a given causal DAG. The other has to do with finding a sufficient set of variables to condition on in order to control for confounding.

It's relatively easy to visually determine if confounding is present in a causal DAG if one recognizes the following fact: backdoor paths between x (cause of interest) and y (effect of interest) which contain a variable along them which is a common cause of both x and y are the "candidates" for confounding paths. Thus, one can use the following procedure to determine if confounding is present (Greenland, Pearl, and Robbins, 1999):

1. delete all arrows coming out of x
2. check to see whether the remaining graph contains variables which cause both x and y whether directly or indirectly
3. if there are common causes of x and y, then the backdoor paths going through those common causes are confounding paths (unless such a backdoor path goes through a collider or descendant of one) and so confounding is present; if there aren't such common causes of x and y, then confounding is absent

If this algorithm were applied to the DAG of Figure 1, we'd have to delete the arrow going from BI to Economic Security (step 1). We'd then find that Political Work Ethic is the only common cause of BI and Economic Security (step 2). Thus, all back door paths going through Political Work Ethic would be candidates for confounding paths (step 3). The BI to Political Work Ethic to Individual Work Ethic to Wage to Economic Security is unblocked so it's a confounding path. The BI to Political Work Ethic to Individual Work Ethic to Wage to Unemploy Rate, to Economic Security is blocked by Wage since Wage is a collider. In order to address the problem of confounding we have to block all backdoor paths between BI and Economic Security by conditioning on a sufficient set.

To find a sufficient set of variables to condition on we could use the following procedure:

1. For each backdoor path collect the variables which intercept that path into a set.
2. Any such set which contains a collider or descendant of a collider blocks the relevant backdoor path.
3. Any variables in any such set which are non-colliders or non-descendants of colliders can be conditioned on the block the relevant path.
4. The sufficient set of variables is that set which blocks all backdoor paths between the variables of interest. A minimally sufficient set can be obtained by deleting variables from a sufficient set one at a time, thereby ending up with a series of subsets of the sufficient set until no other variables can be dropped without ending up with a set of variables which no longer blocks all backdoor paths between the cause and effect variables of interest.

If we followed the two procedures or algorithms spelled out above in regard to Figure 1, we end up with Z = {Political Work Ethic, Individual Work Ethic} as a sufficient set and with Z = {Political Work Ethic} and Z = {Individual Work Ethic} as minimally sufficient sets. These algorithms can be tedious to implement, however, without a computer. Dagitty, the program I used to draw the DAGs in this paper, can also be used to implement these two procedures. Thus, social work researchers interested in applying the ideas of causal DAGS would do well to acquaint themselves with Dagitty or some similar program.

Causal DAGs and RCTs

For most of this paper, I've used a policy example to illustrate the core ideas involved in causal DAGs. I've also focused on the use of DAGs in research with observational or non-experimental data. It should be said, however, that causal DAGs can be used by researchers working at a more "micro" level, and they can be used by those conducting RCTs. Here's an example to give readers a feel for what might be involved in such uses.

Suppose a social work researcher is interested in the effectiveness of a treatment designed to decrease trauma. More specifically, the researcher believes that traumatized persons exposed to this treatment will see a bigger reduction in their symptoms than such persons not so exposed. This researcher also believes that level of violent crime in a traumatized person's neighborhood and level of stress in a traumatized person's life also causally affect magnitude of traumatic symptoms. Further assumptions this researcher makes are 1) violent crime in a traumatized person's neighborhood causes treatment and 2) violent crime causes stress. By violent crime causing treatment I have in mind the idea that if we intervened to increase the level of violent crime in traumatized persons' neighborhoods this would cause an increase in the average level of treatment for trauma that they chose to undergo.[12]

The causal assumptions spelled out in the previous paragraph are encoded in the following graph.

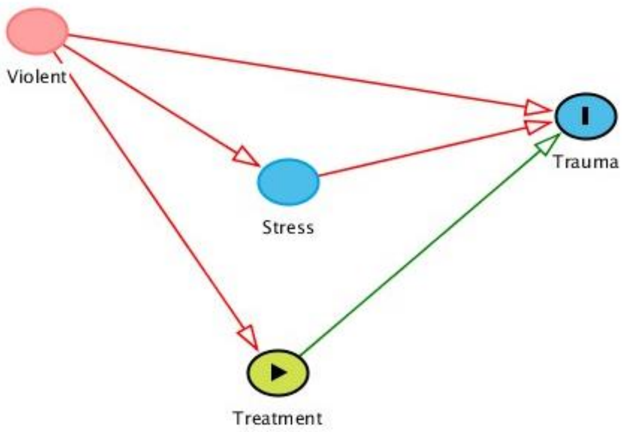

Figure 3 Causal Effect of Treatment on Trauma

Notice that there are two backdoor paths in this graph. One is Treatment to Violent to Trauma and the other is Treatment to Violent to Stress to Trauma.

We can see if there's confounding present in the graph by applying the algorithm discussed earlier and which is repeated here for convenience:

1. delete all arrows coming out of x
2. check to see whether the remaining graph contains variables which cause both x and y whether directly or indirectly
3. if there are common causes of x and y, then the backdoor paths going through those common causes are confounding paths (unless such a backdoor path goes through a collider or descendant of one) and so confounding is present; if there aren't such common causes of x and y, then confounding is absent

Here x is Treatment and y is Trauma, and the only arrow coming out of Treatment is the one going into Trauma. So, applying step 1, that arrow would be deleted. Applying step 2, we'd find that Violent is the only common cause of Treatment and Trauma. Since there are no colliders or descendants of colliders along the two backdoor paths and Violent is a variable on each of these paths, applying step 3 we find that both of those paths are confounding paths—the confounder, of course, is Violent.

If the researcher conducted an RCT and things went as planned, the connection between Violent and Treatment would be broken by randomization. That is, by randomly assigning people to either receiving the treatment or not, the researcher would end up with a graph like that in Figure 4:

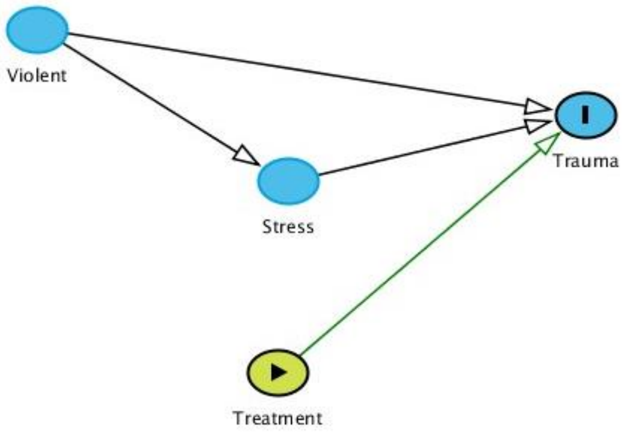

Figure 4

That is by way of random assignment Violent is no longer a cause of treatment. What this example shows is that when it comes to conducting RCTs, causal DAGs can help researchers think through the confounders they need to "watch out for". This, of course, is only true if the causal assumptions encoded in the DAG hold.

Another use of causal DAGs for RCT researchers can be seen by considering what can go wrong. Humans have a way of "not behaving themselves" as RCT participants. This misbehavior can occur in different ways but for the purposes of this paper let's just consider attrition. Attrition is when participants drop out of a study or are loss to follow up.

Suppose in an RCT to test the effectiveness of the treatment for trauma, "drop outs" from the group who received the treatment tend to be those who live in relatively safer neighborhoods while drop outs from the control group tend to be those from less safe neighborhoods. Thus, the group who received the treatment would tend to come from less safe neighborhoods while those in the control group would tend to come from more safe ones. This is another example of the selection problem, which was discussed earlier in the paper and we've already discussed how causal DAGs can help address it. Since Violent is a confounder, the treatment effect could be estimated by either a propensity score model which represented the probability of receiving treatment as a function of Violent before estimating this effect or a regression model of the effect of Treatment on Trauma, controlling for the effect of Violent.

RCTs are often talked about as if there's a sharp divide between experimental and non-experimental or observational data. In theory there is, but once humans start doing all they can to "mess up" pristinely designed RCTs, the line between experimental and observational data starts to look finer. Causal DAGs can help researchers think through how to address confounding in such situations.

Conclusion

I've said nothing in this paper so far about the relationship between hypothesis testing and causal DAGs. This is because causal DAGs have entered the social sciences more as a way of guiding how statistical models are set up in an effort to estimate causal effects than how to test hypotheses regarding such effects. At this point that is, causal DAGs have entered the social sciences more as guides to model specification. But in this closing section I'll make some brief comments regarding what I see as the relationship between causal DAGs and hypothesis testing.

Social work researchers will often have ideas about which variables are causally related to others. These may come from theory, previous research findings, "common sense," or some combination of these. What causal DAGs provide is a precise language for visually representing these ideas. Yet they provide more than that—they also provide rules which stipulate the conditions under which causal effects can be identified, assuming the causal relationships encoded in a given graph are true. As we've seen from earlier sections of the paper, these rules about what needs to be done can be "translated" into ways of specifying statistical models to estimate such causal effects.

The role that hypothesis testing plays in all this is that it allows us to bring data to bear to determine if the assumed causal relationships encoded in a DAG mesh with available data. Rejection of the null hypothesis in this context suggests that a causal DAG is consistent with available data because causal relationships encoded in DAGs constitute alternative hypotheses (from the hypothesis testing perspective). Looked at this way what causal DAGs do for us is offer a precise way of determining what hypotheses we ought to test. And in that role they can be a potent addition to the tool set of quantitative social work researchers.

References

1. Chen, Bryant, Pearl, Judea, **"Regression and Causation: a Critical Examination of Six Econometrics Textbooks"** Real World Economics Review (65), 2013, p. 2-20.
2. de Anda, Diane, **"Intervention Research and Program Evaluation in the School Setting",** Children and Schools, vol. 29 (2), 2007, p. 107-121.
3. Elwert, F. **"Graphical Causal Models"** http://www.ssc.wisc.edu/soc/faculty/pages/docs/elwert/Elwert%202013.pdf, 2015
4. Greenland, Sander and Pearl, Judea, **"Causal Diagrams",** International Encyclopedia of Statistical Science, Lovic, M. (editor), 2011, p. 208-216.
5. Greenland, Sander, Pearl, Judea, Robins, James, **"Causal Diagrams for Epidemiologic Research"** http://www.hsph.harvard.edu/wp-content/uploads/sites/343/2013/03/causaldia.pdf, 1999
6. Guo, Shenyang and Fraser, Mark, **Propensity Score Analysis: Statistical Methods and Applications.** California: Sage Publications, Inc, 2010.
7. Guo, S., Barth R. P., Gibbons, C., **"Propensity Score Matching Strategies for Evaluating Substance Abuse Services for Child Welfare Clients",** Children and Youth Services Review, vol. 28 (4), 2006, p. 357-383.
8. Hawkins, J.D, **"Science, Social Work, Prevention: Finding the Intersections",** Social Work Research, vol. 30 (3), 2006, p. 137-152.
9. Morgan, Stephen L., Winship, Christopher. **Counterfactuals and Causal Inference: Methods and Principles for Social Research,** New York: Cambridge University Press, 2007.
10. Ohmer, M.L. and Korr, W.S, **"The Effectiveness of Community Practice Interventions: a Review of the Literature",** Research on Social Work Practice, vol. 16 (2), 2006, p. 132-145.
11. Pearl, Judea, **Causality: Models, Reasoning, and Inference,** New York: Cambridge University Press, 2000.

12. Spirtes, Peter, Glymour, Clark N., Scheines, Richard, **Causation, Prediction, and Search.** Massachusetts: MIT Press, 2000.
13. Studts, Christina, Stone, Romona, Barber, Gerald, **"Predictors of Access to health Care Services among Groups of TANF Recipients in Kentucky",** Social Service Review, 2006, p. 527-548.
14. Textor, J., Hardt, J. Knuppell, S., **"DAGitty: a Graphical Tool for Analyzing Causal Graphs",** Epidemiology, vol. 22 (5), 2011.
15. Widerquist, Karl, Lewis, Michael, Pressman Steve, **The Ethics and Economics of the Basic Income Guarantee**. England: Ashgate Press, 2005.

---

[1] **Acknowledgements**
I'd like to thank Sharon Schwartz, Dan Herman, Kristin Ferguson-Colvin, and Mark Fraser for their helpful comments. Judea Pearl was also helpful at clarifying issues in a series of email exchanges I had with him. I suspect the paper is better than it otherwise would have been as a result of those exchanges. Any errors or omissions are my responsibility alone.

[2] **Correspondence:** Michael A. Lewis, 129 East 79th Street, New York, NY 10075

[3]Michael Lewis holds a masters degree in social work From Columbia University and a doctorate in Sociology from the City University of New York Graduate Center. He teaches courses in social welfare policy and political economy at the Hunter College School of Social Work. Lewis' research interests are in poverty, social welfare policy, and quantitative methods. His work, some co-authored with Eri Noguchi, has appeared in a number of peer-review journals.

[4] Actually, Alaska does have a policy called the Permanent Fund Dividend which is arguably a basic income but, for the sake of discussion, let's ignore this reality.

[5] I'll use "they" and similar constructions, instead of "he," "she," etc., to be gender neutral.

[6] As long as we're only considering recursive path or structural equation models, since these are the ones consistent with the no cycles or loops constraint regarding causal DAGs.

[7] A 0 on the scale simply means that a person doesn't work at all (outside the home that is).

[8] The idea here is that the degree to which elected officials are committed to this view affects the behavior of individuals in those officials' jurisdiction. More concretely, if members of the legislature in a given locality have made public statements indicative of a high level of commitment to the view that able bodied persons should work for their subsistence this is assumed to cause individual residents of that jurisdiction to expend more effort at work.

[9] When I used the term "train" earlier I really had in mind the idea of a path.

[10] Even though I said "variables," the set in question, as is the case here, can have only one member.

[11] Suppose set Z_1 = {a, b, c} and set Z_2 = {a, b}. Then Z_2 is a proper subset of Z_1 because every member of Z_2 is a member of Z_1 but Z_2 and Z_1 are not equal to or the same as one another.

[12] Intervening to increase crime in someone's neighborhood isn't likely to get support from a Human Subjects Review Board but that doesn't matter for present purposes. I'm really just going back to the definition of cause referred to earlier in the paper.