Published in final edited form as:

J Am Stat Assoc. 2012 December 21; 107(500): 1297–1309. doi:10.1080/01621459.2012.695640.

Functional Causal Mediation Analysis With an Application to Brain Connectivity

Martin A. Lindquist

Associate Professor, Department of Statistics, Columbia University, New York, NY 10027

Martin A. Lindquist: martin@stat.columbia.edu

Abstract

Mediation analysis is often used in the behavioral sciences to investigate the role of intermediate variables that lie on the causal path between a randomized treatment and an outcome variable. Typically, mediation is assessed using structural equation models (SEMs), with model coefficients interpreted as causal effects. In this article, we present an extension of SEMs to the functional data analysis (FDA) setting that allows the mediating variable to be a continuous function rather than a single scalar measure, thus providing the opportunity to study the functional effects of the mediator on the outcome. We provide sufficient conditions for identifying the average causal effects of the functional mediators using the extended SEM, as well as weaker conditions under which an instrumental variable estimand may be interpreted as an effect. The method is applied to data from a functional magnetic resonance imaging (fMRI) study of thermal pain that sought to determine whether activation in certain brain regions mediated the effect of applied temperature on self-reported pain. Our approach provides valuable information about the timing of the mediating effect that is not readily available when using the standard nonfunctional approach. To the best of our knowledge, this work provides the first application of causal inference to the FDA framework.

Keywords

Brain connectivity; Causal inference; fMRI; Functional data analysis; Instrumental variable; Mediation; Structural equation models

1. INTRODUCTION

To date, human brain mapping has been used to primarily construct maps indicating regions of the brain that are activated by certain tasks. Recently, there has been an increased interest in augmenting this type of analysis with connectivity studies that describe how different brain regions interact and how these interactions depend on experimental conditions and behavioral measures. Pathways are considered fundamental properties of brain organization and the ability to understand them is critically important for determining how psychological processes map onto brain function. They provide a means for studying the mechanisms by which experimental manipulations, brain activity, and psychological/physiological outcomes

affect one another. In this work, we discuss a set of tools for modeling functional pathways using functional magnetic resonance imaging (fMRI) time series data.

Functional MRI measures changes in blood flow and oxygenation in the brain in response to neural stimuli, thereby providing a means to noninvasively study changes in mental activity in response to a certain task (Lindquist 2008). The technique offers the potential to measure brain activity while experimentally manipulating treatments, thereby providing information beyond the psychological/physiological outcomes of the treatments typically obtained from psychological experiments. A recent example involves the use of brain imaging to study the relationship between a painful thermal stimulus and self-reported pain (Wager et al. 2008). In this experiment, a noxious heat stimuli was applied at one of the two different levels (high and low) to each of the 20 subjects. In response, subjects gave subjective pain ratings at a specific time point following the offset of the stimulus. While the experiment was being performed, brain activity was measured using fMRI. The goal of the study was to find brain regions whose activity acted as potential mediator of the relationship between temperature and pain rating.

When the effect of the treatment variable Z on the outcome variable Y is at least partially directed through the intervening variable M , then M is said to be a mediator. The three-variable path diagram shown in Figure 1 is often used to represent such relationships. The influence of the intermediate variable on the outcome is then frequently ascertained using structural equation models (SEMs), with the model coefficients interpreted as causal effects. Though the idea of mediation was originally developed in the psychometric and behavioral sciences literature (e.g., Baron and Kenny 1986; MacKinnon 2008), the topic has also received attention in the statistics literature (e.g., Holland 1988; Robins and Greenland 1992; Angrist, Imbens, and Rubin 1996; Ten Have et al. 2007; Albert 2008; Jo 2008; Sobel 2008; VanderWeele 2009; Imai, Keele, and Yamamoto 2010). The brain imaging experiment can be placed into the three-variable path model by letting the variable Z represent the applied pain level, the variable Y the reported pain, and the variable M the brain response. Here, both Z and Y are univariate, while the brain response consists of time series data. In this setting, standard mediation techniques are only applicable if the mediating time course is summarized as a single univariate response, such as peak amplitude or area under the curve. However, the brain response is not necessarily well described by a single summary measure and, in addition, such a measure provides no temporal information about the relationship between Z and Y . Therefore, an important extension of the mediation framework would be to allow the variable M to use information across the entire response—something currently not possible in standard mediation models.

In this work, we consider the same simple three-variable path model described above, with the novel feature that the intermediate variable M is treated as a continuous function (see Figure 2 for the analogous path diagram). This provides an extension of current mediation analysis techniques to the functional data analysis (FDA) setting (Ramsay and Silverman 2006). Conceptually, functional data are thought of as sample paths of a continuous-time stochastic process. Although the observed trajectories are often rough and fluctuating, in many applications of FDA, there is scientific reason to believe that the true underlying trajectory is a smooth function observed with random error. Extending the mediation

framework to the FDA setting allows for the decomposition of the effects of the mediating variable across the support of the function, providing an opportunity to study *functional mediation*.

To assess functional mediation, we propose a linear functional structural equation model (lfSEM), which can be used to both test whether an intermediate variable mediates the relationship between treatment and outcome variables and to provide information about the timing of these relationships. We applied the approach to data from the thermal pain experiment and found that we could discriminate between brain regions that mediate the relationship between temperature and pain at the time of the offset of the stimulus and at the time immediately preceding the later pain report. The identification of these separate mechanisms, not possible using standard SEMs, is important as it promises to significantly increase our knowledge about the underlying brain networks associated with different stages of pain processing.

In statistics, randomized experiments are considered the gold standard for estimating causal effects. A common criticism of SEMs and mediation analysis is that only the assignment to treatment group (i.e., Z) is randomized, while the mediators are self-selected treatments. For this reason, the effect of the mediator (the coefficient for the mediator in the regression of the outcome on the mediator and treatment assignment) does not typically warrant a causal interpretation (Holland 1988). In this work, we adopt the potential outcomes framework for causal inference (Rubin 1974) to show that under certain assumptions, it is possible to obtain a valid estimate of the average causal effects of the mediator from the parameters of the lfSEM. While a number of researchers have considered different strategies for the identification of direct and indirect effects between treatment and outcome (e.g., Robins and Greenland 1992; Pearl 2001; Rubin 2004; VanderWeele 2009), we are primarily concerned with determining sufficient conditions required to equate the parameters of lfSEMs with the effects of mediators. A literature has been recently developed that deals with the identification of direct and indirect effects for situations when interactions and nonlinearities are present (e.g., Petersen, Sinisi, and van der Laan 2006; VanderWeele 2009). However, we leave these issues for later work and focus solely on the type of additive models commonly used in the social sciences. We further extend results from Sobel (2008), which deals with the standard nonfunctional case described in Figure 1, and provide conditions under which the instrumental variable (IV) estimand may be interpreted as an effect when the intermediate variable is a function.

This article is organized as follows. In Section 2, we use potential outcomes notation to construct a causal linear functional model (CLFM) that allows us to specify the causal parameters of interest used in mediation analysis. These are causal analogs of the parameters shown in the path diagram in Figure 2. In general, because we are not able to observe all potential outcomes, we cannot directly estimate the parameters of the CLFM. We therefore proceed by discussing an lfSEM whose parameters can be estimated from the observable data. Under certain assumptions, the parameters of the lfSEM are equivalent to those of the CLFM, thus allowing for the identification of the causal parameters of interest. In Section 3, we compare the IV estimand for the effect of the mediator on the outcome with the equivalent effect in the CLFM. This provides alternate assumptions for allowing a causal

interpretation of the effect of the mediator on the outcome. In Section 4, we describe a technique for estimating the parameters of the lfSEM and the IV estimand, as well as simple procedures for performing inference on the model coefficients using resampling methods. Finally, in Sections 5 and 6, we illustrate the utility of the method in a series of simulations and an application to the fMRI study of thermal pain.

2. IDENTIFICATION OF MEDIATED AND UNMEDIATED EFFECTS

Consider the path diagram in Figure 2 where the mediating variable M_t is assumed to be a continuous function. Here, the subscript t indicates that a variable/coefficient is a function of t , and we write $\mathbf{M} = \{M_t | t \in [0, 1]\}$ to represent the value of the function over its entire range. In the brain imaging setting, Z is the treatment assignment, Y is reported pain, and M_t is a time series of brain data following each stimulation (that will be treated as samples from a continuous underlying function). Further, the path coefficients α_t and β_t are functions that describe the time-varying relationship between the variables.

This section defines the causal parameters of interest for studying functional mediation, as well as an lfSEM that can be used to identify these parameters under conditions to be described. Our goal is to make explicit the assumptions required to equate the parameters of the lfSEM with the causal effects of interest, thus allowing us to determine when the estimates of the former can be given a causal interpretation. In Section 2.1, we define mediated and unmediated effects at both the subject and the population level. A causal model is constructed using potential outcomes for comparison with the lfSEM introduced in Section 2.2. Unlike the parameters of the causal model, the parameters of the lfSEM are identifiable from the observed data. In general, the causal and lfSEM parameters are not equal, implying that the latter should not be interpreted as causal effects. However, we show in Section 2.3 that under certain conditions, the causal parameters of interest are equal to the corresponding parameters of the lfSEM and thus identifiable.

2.1 Defining the Causal Parameters of Interest

In this section, we construct a causal model, analogous to the path diagram shown in Figure 2, using potential outcomes. Consider a randomized experiment consisting of n subjects receiving either the treatment or the control condition, where $Z_i = 1$ if subject i is assigned to treatment and 0 otherwise. For each subject i in the population, let $M_{it}(0)$ denote the value of the mediator at time t in the absence of treatment and $M_{it}(1)$ the value at time t under treatment. Similarly, let $Y_i(0, \mathbf{M}_i(0))$ denote i 's value on the outcome when assigned to the control group and $Y_i(1, \mathbf{M}_i(1))$ the value if assigned to the treatment group. Throughout, we make the stable unit treatment value assumption (SUTVA; Rubin 1980) that there is no interference between subjects. This implies that potential outcomes for each subject are unrelated to the treatment given to other subjects, which is reasonable in the brain imaging setting discussed in this article.

Under SUTVA, the unit-level causal effects of Z on $\mathbf{M}$ and Y , respectively, for subject i are given by the expressions

$$M_{it}(1) - M_{it}(0) \quad \forall t \quad (1)$$

$$Y_i(1, \mathbf{M}_i(1)) - Y_i(0, \mathbf{M}_i(0)). \quad (2)$$

The latter term can be decomposed as follows:

$$\begin{aligned} Y_i(1, \mathbf{M}_i(1)) - Y_i(0, \mathbf{M}_i(0)) &= \{Y_i(1, \mathbf{M}_i(0)) - Y_i(0, \mathbf{M}_i(0))\} + \{Y_i(1, \mathbf{M}_i(1)) - Y_i(1, \mathbf{M}_i(0))\} \\ &= \{Y_i(1, \mathbf{M}_i(1)) - Y_i(0, \mathbf{M}_i(1))\} + \{Y_i(0, \mathbf{M}_i(1)) - Y_i(0, \mathbf{M}_i(0))\}. \end{aligned} \quad (3)$$

In both decompositions, the first term represents an unmediated (or direct) effect of Z on Y , while the second represents a mediated (or indirect) effect. The direct effects of Z on Y represent the difference in the outcome if one were to change the treatment of subject i from $Z_i = 0$ to $Z_i = 1$, while holding the value of the mediator fixed at $\mathbf{M}_i(0)$ (or $\mathbf{M}_i(1)$). In contrast, the indirect effects represent the difference in the outcomes if the value of the observed mediator were changed from $\mathbf{M}_i(0)$ to $\mathbf{M}_i(1)$, while keeping the actual treatment fixed at 0 (or 1). If the treatment has no effect on the mediator, then $\mathbf{M}_i(1) = \mathbf{M}_i(0)$ and the indirect effects would be zero.

In practice, since only one of the potential outcomes is observable, it is not possible to identify unit-level effects. Instead, we average over subjects and seek to estimate population-level effects such as the average causal effect of Z on Y :

$$\tau^{(c)} \equiv E(Y(1, \mathbf{M}(1)) - Y(0, \mathbf{M}(0))). \quad (4)$$

To define other causal effects of interest, we construct a CLFM in which the relationship between the outcome and the mediator follows a functional linear model (e.g., Ramsay and Silverman 2006). Potential outcomes are used to express the causal relationship between Z , $\mathbf{M}$, and Y using the system of equations

$$M_{it}(z) = \delta_{1t}^{(c)} + \alpha_t^{(c)} z + \varepsilon_{it}(z) \quad (5)$$

$$Y_i(z, \mathbf{m}) = \delta_2^{(c)} + \gamma^{(c)} z + \int_0^1 \beta_t^{(c)} m_t dt + \eta_i(z, \mathbf{m}), \quad (6)$$

where $E(\varepsilon_t(z)) = 0 \quad \forall t$ and $z = 0, 1$, and $E(\eta(z, \mathbf{m})) = 0$ for all values of the pair $(z, \mathbf{m})$. Here, the superscript (c) is used to identify the parameters as being causal, in contrast to the analogous parameters $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\gamma^{(s)}$ in the IfSEM presented below.

Using this model, we can write the average direct effect of Z on Y , sometimes referred to as the controlled direct effect (VanderWeele 2009), as

$$E(Y(1, \mathbf{m}) - Y(0, \mathbf{m})) = \gamma^{(c)} \quad (7)$$

for all values of $\mathbf{m}$. This effect differs from the pure direct effect (Robins and Greenland 1992) that fixes the intermediate variable for each individual to the level it would have received under the absence of treatment. For the CLFM, these effects are equivalent under the additional condition $E(\tau(1, \mathbf{M}(0)) - \tau(0, \mathbf{M}(0))) = 0$.

Averaging (3) over all subjects, we obtain

$$\begin{aligned}\tau^{(c)} &= \gamma^{(c)} + E(Y(1, \mathbf{M}(1)) - Y(1, \mathbf{M}(0))) \\ &= \gamma^{(c)} + E(Y(0, \mathbf{M}(1)) - Y(0, \mathbf{M}(0))),\end{aligned}\quad (8)$$

which allows us to express the average total effect as the sum of the average direct and indirect effect of Z on Y . Note that according to (7), the average of both formulations of the direct effect stated in (3) can be expressed using $\gamma^{(c)}$.

Finally, the average causal effect of Z on $\mathbf{M}$, and $\mathbf{M}$ on Y at level $\mathbf{m}$ versus $\mathbf{m}^*$ can be written, respectively, as

$$E(M_t(1) - M_t(0)) = \alpha_t^{(c)} \quad \forall t \in [0, 1] \quad (9)$$

and

$$E(Y(z, \mathbf{m}) - Y(z, \mathbf{m}^*)) = \int_0^1 \beta_t^{(c)} (m_t - m_t^*) dt. \quad (10)$$

Both α_t and β_t depend on t and their value can vary across the range of $[0, 1]$, thus allowing the treatment to have a functional effect on the intermediate variable. In our application, this provides a temporal decomposition of the effect that applied pain has on the brain response. This is important as the response can potentially vary with regard to its onset, width, or amplitude for different treatments.

The fundamental problem of causal inference (Holland 1988) is that for any given subject, one cannot observe the potential outcomes under both the treatment and the control conditions at the same time. In addition, terms such as $Y_i(1, \mathbf{M}(0))$ can never be observed, regardless of treatment. Hence, randomization of the treatment alone cannot help identify the terms $\gamma^{(c)}$ and $\beta^{(c)}$ and additional assumptions are required to proceed.

2.2 Linear Functional Structural Equation Models

Due to the fundamental problem of causal inference, none of the causal effects defined in the previous section can generally be estimated from the observed data. Therefore, in this section, we discuss an lfSEM, corresponding to the path diagram in Figure 2, whose parameters can be directly estimated from the data.

A great deal of research has focused on extending linear models to the functional setting (e.g., Hastie and Tibshirani 1993; Fan and Zhang 1999; James 2002; Cardot et al. 1999; Müller and Stadtmüller 2005). Currently, techniques exist for performing regression where both/either the response and/or explanatory variables are functional rather than scalar. Using

these techniques, under conditions described in Section 2.3, mediation can be assessed using a *functional* analog to a linear SEM of the form:

$$M_{it}(Z_i) = \delta_{1t}^{(s)} + \alpha_t^{(s)} Z_i + \varepsilon_{it} \quad \forall t \quad (11)$$

$$Y_i(Z_i, \mathbf{M}_i(Z_i)) = \delta_2^{(s)} + \gamma^{(s)} Z_i + \int_0^1 \beta_t^{(s)} M_{it}(Z_i) dt + \eta_i. \quad (12)$$

where $t \in [0, 1]$. Here, the superscript (s) denotes that the parameters are defined in an SEM that is estimable using the observed data. The model parameters are identified through the definitions $E(\varepsilon_t | Z = z) = 0 \quad \forall t$ and $E(\eta | Z = z, \mathbf{M}(Z) = \mathbf{m}) = 0$. While not necessary for identification, for estimation purposes, we also assume that ε_{it} and η_i are independent. Using (11) and (12), we can write

$$\alpha_t^{(s)} = E(M_t(1) | Z=1) - E(M_t(0) | Z=0) \quad \forall t, \quad (13)$$

$$\gamma^{(s)} = E(Y(Z, \mathbf{M}(Z)) | \mathbf{M}(Z) = \mathbf{m}, Z=1) - E(Y(Z, \mathbf{M}(Z)) | \mathbf{M}(Z) = \mathbf{m}, Z=0), \quad (14)$$

$$\int \beta_t^{(s)} (m_t - m_t^*) dt = E(Y(Z, \mathbf{M}(Z)) | \mathbf{M}(Z) = \mathbf{m}, Z=z) - E(Y(Z, \mathbf{M}(Z)) | \mathbf{M}(Z) = \mathbf{m}^*, Z=z). \quad (15)$$

In this formulation, it is easy to show that the “total effect” of Z on Y ¹

$$\tau^{(s)} \equiv E(Y(1, \mathbf{M}(1) | Z=1)) - E(Y(0, \mathbf{M}(0) | Z=0)) \quad (16)$$

can be expressed as

$$\tau^{(s)} = \gamma^{(s)} + \int_0^1 \alpha_t^{(s)} \beta_t^{(s)} dt. \quad (17)$$

Here, $\gamma^{(s)}$ represents the “direct effect” of Z on Y , while the term $\int \alpha_t^{(s)} \beta_t^{(s)} dt$ represents the “indirect effect.” Mediation can then be assessed under suitable conditions by determining whether the integral is significantly different from zero. In addition, the product $\alpha_t^{(s)} \beta_t^{(s)}$ provides a functional decomposition of the “indirect effect,” allowing the specific intervals driving the mediation to be determined. Hence, the proposed framework provides the opportunity under suitable conditions to assess the effects of *functional mediation*.

2.3 Comparison of Causal Parameters and IfSEM Parameters

In general, the parameters of the IfSEM ($\tau^{(s)}$, $\gamma^{(s)}$, $\alpha_t^{(s)}$, and $\beta_t^{(s)}$) are not equal to their counterparts in the CLFM ($\tau^{(s)}$, $\gamma^{(s)}$, $\alpha_t^{(c)}$, and $\beta_t^{(c)}$). To equate the parameters, we need to

¹Quotation marks are used in this subsection to differentiate the effects in the IfSEM from their counterparts in the CLFM.

make a number of assumptions. We begin by assuming that treatment assignment is ignorable (Rosenbaum and Rubin 1983).

Assumption 1—The treatments are assigned independently of the potential outcomes, that is,

$$Y(0, \mathbf{M}(0)), Y(1, \mathbf{M}(1)), \{M_t(0), M_t(1)\}_{t \in [0,1]} \perp\!\!\!\perp Z. \quad (18)$$

Under this assumption, it is easy to show that $\tau^{(c)} = \tau^{(s)}$ and $\alpha_t^{(c)} = \alpha_t^{(s)}$. However, in general, $\gamma^{(c)} \neq \gamma^{(s)}$ and $\beta^{(c)} \neq \beta^{(s)}$. Under Assumption 1, we can express (15) as

$$\int \beta_t^{(s)} (m_t - m_t^*) dt = E(Y(z, \mathbf{m}) | \mathbf{M}(Z) = \mathbf{m}) - E(Y(z, \mathbf{m}^*) | \mathbf{M}(Z) = \mathbf{m}^*), \quad (19)$$

which is not directly comparable with the expression in (10). The term $\beta_t^{(s)}$ compares the values of the outcome in two different subpopulations of units, while $\beta_t^{(c)}$ considers the same subjects under different conditions. For similar reasons, the term $\gamma^{(s)}$ is not directly comparable with (7). A sufficient condition for equality of the remaining parameters is to assume that the mediator is ignorable with respect to the potential outcomes.

Assumption 2—The potential outcomes are independent of the mediator, that is,

$$Y(z, \mathbf{m}) \perp\!\!\!\perp \{M_t(z)\}_{t \in [0,1]} \quad (20)$$

for $z = 0, 1$, and all $\mathbf{m}$.

Under Assumptions 1 and 2, it is easy to show the equality of the remaining parameters, as they imply that

$$\begin{aligned} E(Y(Z, \mathbf{M}(Z)) | Z=z, \mathbf{M}(Z)=\mathbf{m}) &= E(Y(z, \mathbf{m}) | Z=z, \mathbf{M}(z)=\mathbf{m}) \\ &= E(Y(z, \mathbf{m}) | \mathbf{M}(z)=\mathbf{m}) \\ &= E(Y(z, \mathbf{m})). \end{aligned} \quad (21)$$

Hence, under Assumptions 1 and 2, $\gamma^{(s)} = \gamma^{(c)}$ and $\beta_t^{(s)} = \beta_t^{(c)}$. Note that these are the same assumptions required for equating the parameters of SEMs with the parameters of an analogous causal model (Sobel 2008).

In addition, it should also be noted that under Assumption 2, we can write the indirect effect of Z on Y as follows:

$$\begin{aligned} E(Y(0, \mathbf{M}(1)) - Y(0, \mathbf{M}(0))) &= E(Y(1, \mathbf{M}(1)) - Y(1, \mathbf{M}(0))) \\ &= \int_0^1 \alpha_t^{(c)} \beta_t^{(c)} dt. \end{aligned} \quad (22)$$

This, in turn, allows us to express both decompositions of the total effect in (8) as

$$\tau^{(c)} = \gamma^{(c)} + \int_0^1 \alpha_t^{(c)} \beta_t^{(c)} dt, \quad (23)$$

providing expressions for the direct and indirect effects in terms of the parameters of the CLFM.

If Assumptions 1 and 2 are valid, it is appropriate to make causal claims using the results from the lfSEM. Assumption 1 states that treatment assignment is ignorable, as would be the case in a randomized fMRI experiment. Assumption 2 states that the potential outcomes $Y(z, \mathbf{m})$ are ignorable with respect to the intermediate outcomes $M_t(z)$, given Z , as would be the case if subjects were randomly assigned to $\mathbf{M}$ at both levels of Z . Obviously, this is a strong assumption that is untestable in practice, and it is easy to construct examples where it is violated (e.g., Lindquist and Sobel 2011, 2012). For example, unobserved variables may exist that confound the relationship between the outcome and the mediator variable even after conditioning on the treatment status, for example, a latent variable, such as pain resilience, that simultaneously results in decreased brain response and lesser reported pain. In the next section, we look at ways of relaxing Assumption 2.

3. IDENTIFICATION OF CAUSAL EFFECTS USING AN INSTRUMENTAL VARIABLE

In this section, we give conditions under which the IV estimand (the effect of Z on Y divided by the effect of Z on $\mathbf{M}$) can be used to identify the causal effect $\beta^{(c)}$. The results follow from similar work by Holland (1988) and Sobel (2008) on SEMs and provide an alternate set of assumptions for identifying the causal effect of $\mathbf{M}$ on Y . The IV estimand has also been considered in connection with causal effects defined in potential outcomes in work by Angrist and colleagues (Angrist and Imbens 1995; Angrist, Imbens, and Rubin 1996).

Under (5) and (6), it is possible to write

$$Y_i(1, \mathbf{M}_i(1)) - Y_i(0, \mathbf{M}_i(0)) = \gamma^{(c)} + \int_0^1 \beta_t^{(c)} (M_{it}(1) - M_{it}(0)) dt + \eta_i(1, \mathbf{M}_i(1)) - \eta_i(0, \mathbf{M}_i(0)). \quad (24)$$

To proceed, we assume that the difference in potential errors is 0, on average.

Assumption 3— $E(\eta(1, \mathbf{M}(1)) - \eta(0, \mathbf{M}(0))) = 0$.

Because $\tau^{(c)} = \tau^{(s)}$ and $\alpha_t^{(c)} = \alpha_t^{(s)}$ when Assumption 1 holds, using Assumption 3 and averaging over both sides of (24) gives

$$\tau^{(s)} = \gamma^{(c)} + \int_0^1 \alpha_t^{(s)} \beta_t^{(c)} dt, \quad (25)$$

which yields one equation in the unknown scalar $\gamma^{(c)}$ and the unknown function $\beta_t^{(c)}$. In some applications, it may be reasonable to assume that the treatment effect is transmitted solely through the mediator, leading to the following assumption.

Assumption 4 $Y_i(0, \mathbf{m}) = Y_i(1, \mathbf{m}) \quad \forall i, \mathbf{m}$.—This assumption, which implies $\gamma^{(c)} = 0$ and is often referred to as an exclusion restriction, states that treatment assignment is unrelated to the potential outcomes once the mediator has been taken into account. Under this assumption, we have

$$\tau^{(s)} = \int_0^1 \alpha_t^{(s)} \beta_t^{(c)} dt, \quad (26)$$

which can be viewed as an ill-posed homogeneous integral equation of the first kind. In this case, the standard IV estimand $\tau^{(s)} / \alpha_t^{(s)}$ is a solution for $\beta_t^{(c)}$, albeit a nonunique one. This particular solution becomes unique if one assumes that the indirect effect is constant across the range of the mediating variable. As the assumption of a constant mediation effect may not be reasonable in fMRI studies, we instead seek a different solution, namely the least-squares solution of minimum norm. We begin by expressing (26) in the form of an operator equation

$$\mathcal{A} \beta^{(c)} = \tau^{(s)}, \quad (27)$$

where $\mathcal{A}$ is a linear operator such that $\mathcal{A} \beta^{(c)} = \int \alpha_t^{(s)} \beta_t^{(c)} dt$. This, in turn, can be replaced by the approximating matrix problem

$$\mathbf{A} \tilde{\beta}^{(c)} = \tau^{(s)}, \quad (28)$$

whose solution is given by

$$\tilde{\beta}^{(c)} = \mathbf{A}^+ \tau^{(s)}, \quad (29)$$

where $\mathbf{A}^+$ is the Moore–Penrose pseudo-inverse of $\mathbf{A}$ and $\tilde{\beta}^{(c)}$ is the discretized version of $\beta_t^{(c)}$. This solution has many attractive features, including that if $\mathbf{A}$ has real entries, then so does its pseudo-inverse $\mathbf{A}^+$, and if $\mathbf{A}$ is invertible, then $\mathbf{A}^+$ and $\mathbf{A}^{-1}$ coincide. The choice of the minimum norm solution constitutes an additional assumption. However, we differentiate it from the other assumptions due to our belief that they must be verified each time the methodology is applied to a new dataset, while the same is not true for the use of the minimum norm solution.

In sum, under Assumptions 1 and 2, we can equate the parameters of the CLFM with the equivalent parameters of the lfSEM. In contrast, under Assumptions 1, 3, and 4, we can use the IV estimand to identify the causal effect of $\mathbf{M}$ on Y . It should be noted that Assumption 2, which states that there is no selection on the mediators with respect to the potential outcomes $Y(z, \mathbf{m})$, is more restrictive than Assumption 3, which states that there is no selection on the difference between potential errors. To see this, note that Assumption 2 is equivalent to

$$\eta(z, \mathbf{m}) \perp\!\!\!\perp \{M_t(z)\}_{t \in [0,1]} \quad (30)$$

for $z = 0, 1$, and all $\mathbf{m}$. Hence,

$$\begin{aligned} E(\eta(z, \mathbf{M}(z))) &= E(E(\eta(z, \mathbf{M}(z))) | \mathbf{M}(z) = \mathbf{m}) \\ &= E(E(\eta(z, \mathbf{m})) | \mathbf{M}(z) = \mathbf{m}) \\ &= 0, \end{aligned} \quad (31)$$

which implies Assumption 3. However, Assumption 3 does not imply Assumption 2.

Assumption 4 (exclusion restriction) is not always plausible, and the IV estimand will be biased by a factor $\mathbf{A}^+ \gamma^{(c)}$ when it does not hold. In this situation, assumptions on the sign and/or magnitude of $\gamma^{(c)}$ might be used to bound the causal effect $\beta_t^{(c)}$.

4. ESTIMATION AND INFERENCE

In this section, we describe an approach for estimating the parameters of the lfSEM defined in (11)–(12) and the IV estimand defined in (29). We also discuss inferential procedures based on the use of resampling techniques.

4.1 Estimation

We begin by describing techniques for estimating the lfSEM. Equation (11) consists of a scalar predictor and a functional response, while (12) consists of a functional predictor and a scalar response. For both types of models, there are a variety of estimation techniques; see Ramsay and Silverman (2006) for a comprehensive overview. In the following, we will for simplicity suppress the intercept terms when illustrating the estimation procedure and instead consider the models:

$$M_{it} = \alpha_t^{(s)} Z_i + \varepsilon_{it} \quad (32)$$

$$Y_i = \gamma^{(s)} Z_i + \int \beta_t^{(s)} M_{it} dt + \eta_i. \quad (33)$$

We approach (32) by expressing $\alpha_t^{(s)}$ as a linear combination of K known basis functions ϕ_{kt} , that is,

$$\alpha_t^{(s)} = \sum_{k=1}^K \phi_{kt} a_k = \Phi_t \mathbf{a}. \quad (34)$$

where $\Phi_t = (\phi_{1t}, \dots, \phi_{Kt})$ and $\mathbf{a} = (a_1, \dots, a_K)^T$. The coefficients of the expansion, $\mathbf{a}$, are determined by minimizing the least-squares criterion

$$\text{LMSSE}(\alpha_t^{(s)}) = \sum \int [M_{it} - \alpha_t^{(s)} Z_i]^2 dt. \quad (35)$$

If we reexpress $M_{it} = \alpha_t^{(s)} Z_i$ as $M_{it} - Z_i \Phi_t \mathbf{a}$ and let $\Psi_{it} = Z_i \Phi_t$, then the solution is given by

$$\hat{\mathbf{a}} = \left[\sum \int \Psi_{it}^T \Psi_{it} dt \right]^{-1} \left[\sum \int \Psi_{it}^T M_{it} dt \right], \quad (36)$$

which provides us with an estimate of $\alpha_t^{(s)}$.

The space of functions $\beta_t^{(s)}$ satisfying (33) is infinite-dimensional and thus simply minimizing the sum of squares will not provide a meaningful estimate. Instead, we minimize the penalized square error

$$\text{PENSSE}_\lambda \left(\gamma^{(s)}, \beta_t^{(s)} \right) = \sum \left[Y_i - \gamma^{(s)} Z_i - \int \beta_t^{(s)} M_{it} dt \right]^2 + \lambda \int D^2 \beta_t^{(s)} dt, \quad (37)$$

where the integral portion of the second term represents the integrated squared second derivative and λ represents a constant smoothing parameter. We proceed by expressing the regression function as

$$\beta_t^{(s)} = \sum_{k=1}^K \phi_{kt} b_k = \Phi_t \mathbf{b}, \quad (38)$$

where $\mathbf{b} = (b_1, \dots, b_K)^T$. Using this expression, we can write (33) as

$$\begin{aligned} Y_i &= \gamma^{(s)} Z_i + \int \beta_t^{(s)} M_{it} dt + \eta_i \\ &= \gamma^{(s)} Z_i + \left[\int \Phi_t M_{it} dt \right] \mathbf{b} + \eta_i \\ &= \gamma^{(s)} Z_i + \Lambda_i \mathbf{b} + \eta_i, \end{aligned} \quad (39)$$

where $\Lambda_i = \int \Phi_t M_{it} dt$. Concatenating the results across all subjects, we obtain

$$Y = G\xi + \eta, \quad (40)$$

where

$$Y = \begin{pmatrix} Y_1 \\ \vdots \\ Y_n \end{pmatrix}, G = \begin{pmatrix} Z_1 & \Lambda_1 \\ \vdots & \vdots \\ Z_n & \Lambda_n \end{pmatrix} \text{ and } \xi = \begin{pmatrix} \gamma^{(s)} \\ \mathbf{b} \end{pmatrix}. \quad (41)$$

The solution of (37) is given by

$$\hat{\xi} = (G^T G + \lambda R)^{-1} G^T Y, \quad (42)$$

where R is a matrix with elements R_{ij} equal to zero if j or $k = 1$ and $\int D^2 \varphi_{js} D^2 \varphi_{ks} ds$ otherwise.

To compute the IV estimand, we first obtain the path coefficient $\alpha_t^{(s)}$ as outlined above. The discretized version of the linear operator $\mathcal{A}$, denoted $\mathbf{A}$, can then be obtained by evaluating $\Phi_t \hat{\mathbf{a}}, t \in [0, 1]$, at N equidistant points. Next, the path coefficient $\tau^{(s)}$ is computed using the equation

$$Y_i(Z_i, \mathbf{M}_i(Z_i)) = \delta_3^{(s)} + \tau^{(s)} Z_i + \varepsilon_i, \quad (43)$$

which can be solved using standard linear regression methods. Finally, the IV estimand is obtained by taking the pseudo-inverse of $\mathbf{A}$ and computing $\mathbf{A}^+ \tau^{(s)}$.

4.2 Inference

Once the parameters of the lfSEM and/or the IV estimand have been estimated, we can perform inference on the indirect effect. In brain imaging, hypothesis testing is considered to be of primary importance. In mediation analysis, we are particularly interested in determining whether the total effect is stronger than the direct relationship controlling for $\mathbf{M}$. Analogous to the fully univariate setting (Sobel 1982; Baron and Kenny 1986), this can be done by testing the significance of the integral of the $\alpha_t^{(s)} \beta_t^{(s)}$ product. As the exact distribution of the function $\alpha_t^{(s)} \beta_t^{(s)}$ is unknown, we use bootstrap methods to perform inference, providing a functional equivalent to an approach often used in the univariate setting (Shrout and Bolger 2002).

We propose using a wild bootstrap procedure as it is known to perform well for regression models (Flachaire 2005; Zhu et al. 2007). To produce a bootstrap sample $(Z_i, Y_i^*, \mathbf{M}_i^*), i = 1, \dots, n$, we use the following data-generating process (DGP):

$$M_{it}^*(Z_i) = \hat{\delta}_{1t}^{(s)} + \hat{\alpha}_t^{(s)} Z_i + v_i^* \hat{\varepsilon}_{it} \quad \forall t \quad (44)$$

$$Y_i^*(Z_i, \mathbf{M}_i^*(Z_i)) = \hat{\delta}_2^{(s)} + \hat{\gamma}^{(s)} Z_i + \int_0^1 \hat{\beta}_t^{(s)} M_{it}^*(Z_i) dt + v_i^* \hat{\eta}_i, \quad (45)$$

where $\hat{\delta}_{1t}^{(s)}, \hat{\delta}_2^{(s)}, \hat{\alpha}_t^{(s)}, \hat{\beta}_t^{(s)}$, and $\hat{\gamma}^{(s)}$ are estimates of the parameters in (11) and (12), $\hat{\varepsilon}_{it}$ and $\hat{\eta}_i$ are the i th residuals, and v_i^* are independent and identically distributed as

$$v_i^* = \begin{cases} 1 & \text{with probability } 0.5 \\ -1 & \text{with probability } 0.5 \end{cases}.$$

This particular DGP has been studied extensively in the context of IV regression (Davidson and MacKinnon 2010) and shown to outperform other comparable bootstrap methods.

The bootstrap procedure is performed as follows:

1. Independently generate a bootstrap sample $(Z_i, Y_i^*, \mathbf{M}_i^*)$, $i = 1, \dots, n$, using the wild boot-strap DGP described in (44) and (45).
2. Using the resampled data, refit (11) and (12) and record the estimate of $\alpha_t^{(s)} \beta_t^{(s)}$ for $t \in [0, 1]$. When working with the IV estimand, refit (29) rather than (12).
3. Repeat the procedure outlined in steps (1) and (2) B (e.g., 1000) times.

The B replications are used to compute the bootstrap distribution of $\alpha_t^{(s)} \beta_t^{(s)}$ for all values of t , which can be used to test whether $\alpha_t^{(s)} \beta_t^{(s)}$ differs significantly from 0. Because the suggested DGP does not impose the null hypothesis on the resampled data, it is necessary to instead test whether $\alpha_t^{(s)} \beta_t^{(s)}$ differs from $\hat{\alpha}_t^{(s)} \hat{\beta}_t^{(s)}$ to obtain a valid test. It should be noted that the bootstrap distribution can also be used to compute percentile confidence intervals (Efron and Tibshirani 1998) for $\alpha_t^{(s)} \beta_t^{(s)}$ for all t .

Because we are interested in testing for significance across the range of t , it is necessary to correct for multiple comparisons. Throughout the article, we use the Benjamini–Hochberg procedure (Benjamini and Hochberg 1995) to control the false detection rate (FDR). We use FDR-controlling methods, rather than methods that control the familywise error rate (FWER), because choosing the appropriate threshold to control the FWER is difficult due to dependencies in the data.

5. SIMULATION STUDY

In this section, we report the results of a series of simulation studies where data were generated using different combinations of the assumptions described in Sections 2 and 3. In the first five simulations (summarized in the schematic shown in Figure 3), data were generated according to the model described in (11) and (12). This is meant to represent a situation under which sequential ignorability (Assumptions 1 and 2) holds. The first three simulations were designed to determine the empirical false positive rate of the method when no mediation effect is present, while the next two were designed to study the power in the presence of a significant mediation effect. The sixth and final simulation, generated using the causal models (5) and (6), was designed to mimic a situation where Assumption 2 is violated but the assumptions for identifying $\beta_t^{(c)}$ using the IV estimand hold (Assumptions 3 and 4).

In five of the simulations (1–4 and 6), data were generated for 20 “subjects,” 10 randomly assigned to a treatment group ($Z = 1$) and 10 to a control group ($Z = 0$). In the fifth simulation, data were generated for n “subjects,” with n allowed to vary between 10 and 50, split equally between $Z = 0$ and $Z = 1$. The values of M_t and Y differed between the six simulations, as outlined below, and were chosen to mimic plausible brain imaging settings.

Throughout, we restricted M_t to a uniform grid of 60 points in the interval $[0, 1]$. In five of the simulations, the shape of M_t is based on a common time-varying function, denoted h_t (see bottom-left-hand side of Figure 3), which equals 0 in the ranges $[0, 1/3]$ and $(2/3, T]$, and varies smoothly in the range $[1/3, 2/3]$. The function corresponds to the so-called canonical hemodynamic response function (HRF; Lindquist 2008) typically used to model changes in fMRI signal in response to a stimulus. This function consists of an early rise in activation level, followed by a decline and a subsequent postactivation undershoot below baseline level.

Each of the six simulations was repeated 500 times. For each repetition of Simulations 1–5, the models (11) and (12) were fitted using the approach outlined in Section 4. When fitting the lfSEM model, we restrict M_t to a uniform grid of 60 points in the interval $[0, 1]$. The path coefficients $\alpha_t^{(s)}$ and $\beta_t^{(s)}$ are modeled using a B-spline basis set of order 6 with 30 knots. The path equations are computed using custom Matlab software, incorporating functions from the Matlab toolbox (<http://www.psych.mcgill.ca/misc/fda/index.html>) for FDA. The roughness penalty required to fit (12) is taken as the L_2 -norm of its second derivative, with the smoothing parameter selected by leave-one-out cross-validation; no smoothing is used in the initial step of representing M_t in terms of the B-spline basis. A bootstrap test was performed to determine time points where $\alpha_t^{(s)}\beta_t^{(s)}$ was significantly different from 0. The results were controlled for multiple comparisons using the Benjamini–Hochberg procedure ($q = 0.05$).

For each repetition of Simulation 6, $\alpha_t^{(c)}\beta_t^{(c)}$ is obtained by equating $\alpha_t^{(c)} = \alpha_t^{(s)}$ and computing the IV estimand $\beta_t^{(c)}$ using (29). A bootstrap test was performed to determine whether $\alpha_t^{(c)}\beta_t^{(c)}$ was significantly different from 0, controlling for multiple comparisons using the Benjamini–Hochberg procedure ($q = 0.05$).

Simulation 1—The data are generated assuming $M_{it} = \varepsilon_{it}$ and $Y_i = Z_i + \eta_i$ for all $i = 1, \dots, 20$, where both ε_{it} and η_i follow a standard normal distribution. In this simulation, there should be a significant relationship between Z and Y , but it should not be mediated by M_t . This holds because there is no effect of Z on M_t and similarly no effect of M_t on Y , that is, $\alpha^{(s)} = \beta^{(s)} = 0$. The first row of Figure 4 shows estimates of $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\alpha_t^{(s)}\beta_t^{(s)}$, together with 95% bootstrap percentile confidence intervals. None of the effects appear to deviate significantly from 0 across the time interval. Figure 5(A) shows the proportion of times $\alpha_t^{(s)}\beta_t^{(s)}$ was deemed significant as a function of t in the 500 replications. Clearly, all time points fall well below 0.05.

Simulation 2—The data are generated in an analogous manner as in Simulation 1, except $M_{it} = (Z_i + \xi_i)h_t + \varepsilon_{it}$ for all $i = 1, \dots, 20$, where h_t is defined as above, ε_{it} follows a standard normal distribution, and $\xi_i \sim N(1, 0.1)$. Again, there should be a significant relationship between Z and Y , but it should not be mediated by M_t as there is no effect of M_t on Y since $\beta_t^{(s)} = 0$ for all t . Note that in this case, $\alpha_t^{(s)} \neq 0$, differentiating it from the setting of Simulation 1. The second row of Figure 4 shows estimates of $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\alpha_t^{(s)}\beta_t^{(s)}$,

together with 95% bootstrap percentile confidence intervals. The estimates coincide with the simulated values, with only $\alpha_t^{(s)}$ being significantly different from 0 in the time interval corresponding to where h_t is nonzero.

Figure 5(B) shows the proportion of times $\alpha_t^{(s)}\beta_t^{(s)}$ was deemed significant as a function of t . Again, all time points fall below 0.05. However, the proportion of false positives are increased and approaches 0.05 in the interval where h_t is nonzero.

Simulation 3—The data are generated assuming $M_{it} = \xi_i h_t + \varepsilon_{it}$ and $Y_i = \int_{1/3}^{2/3} M_{it} dt + \eta_i$ for all $i = 1, \dots, 20$, where h_t is defined as above, both ε_{it} and η_i follow standard normal distributions, and ξ_i is $N(1, 0.1)$. In this simulation, there should be no relationship between Z and Y or Z and M as $\tau^{(s)} = 0$ and $\alpha_t^{(s)} = 0$ for all t . The third row of Figure 4 shows estimates of $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\alpha_t^{(s)}\beta_t^{(s)}$, together with 95% bootstrap percentile confidence intervals. Again, the estimates seem to coincide with the simulated values, with only $\beta_t^{(s)}$ being significantly different from 0 in the time interval $[1/3, 2/3]$. Figure 5(C) shows the proportion of times $\alpha_t^{(s)}\beta_t^{(s)}$ was deemed significant as a function of t in the 500 repetitions. Again, all time points fall well below 0.05.

Simulation 4—The data are generated in an analogous manner as described in Simulation 3, except $M_{it} = (Z_i + \xi_i)h_t + \varepsilon_{it}$ for all $i = 1, \dots, 20$, where ε_i follows a standard normal distribution and ξ_i is $N(1, 0.1)$. In this simulation, the relationship between Z and Y is significant and should be mediated by M_t in the range $[1/3, 2/3]$ as both $\alpha_t^{(s)}$ and $\beta_t^{(s)}$ are nonzero in that interval. The fourth row of Figure 4 shows estimates of $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\alpha_t^{(s)}\beta_t^{(s)}$, together with 95% bootstrap percentile confidence intervals. Again, all three estimates coincide with the simulated values, as they are significantly different from 0 in the appropriate intervals of time. Figure 6(A) shows the proportion of times $\alpha_t^{(s)}\beta_t^{(s)}$ was deemed significant as a function of t . In the intervals $[0, 1/3]$ and $(2/3, 1]$, all time points fall below 0.05. In addition, we see a significant relationship in the range $[1/3, 2/3]$, as expected. The proportion of significant results peaks at 45% around the point where h_t takes its maximum.

Simulation 5—The data are generated in an analogous manner as described in Simulation 4, except the sample size was allowed to vary between 10 and 50 “subjects.” Here, the relationship between Z and Y should be mediated by M_t in the range $[1/3, 2/3]$, as before.

Figure 6(B) shows the proportion of times $\alpha_t^{(s)}\beta_t^{(s)}$ was deemed significant as a function of t for a number of sample sizes ranging from 10 to 50. As expected, the higher the sample size, the better the sensitivity and specificity.

Simulation 6—The data are generated according to the causal model in (5) and (6) with $\delta_2^{(c)} = \gamma^{(c)} = 0$, $\delta_{1t}^{(c)} = 0 \forall t$, $\alpha_t^{(c)} = h_t$ and $\beta_t^{(c)} = 1$ in the range $[2/6, 3/6]$ and 0 otherwise. We ensured that $\varepsilon_{it}(0), \varepsilon_{it}(1) \perp Z_i$ (Assumption 1) by generating both potential outcomes from a

standard normal model and randomly assigning the observed outcome to each subject.

Further, we set $\eta_i(z, \mathbf{m}) = \int_{2/6}^{3/6} \varepsilon_{it}(z) dt + \eta_i^*(z, \mathbf{m})$, where $\eta_i^*(z, m)$ is standard normal. This ensures that Assumption 2 is violated since $\eta_i(z, m)$ is not independent of $\varepsilon_i(z)$ for $z = 0, 1$.

However, since Assumptions 3 and 4 hold, we can estimate $\beta_t^{(c)}$ using the IV estimand. In this simulation, the relationship between Z and Y is significant and should be mediated by M_t in the range $[2/6, 3/6]$ as both $\alpha_t^{(c)}$ and $\beta_t^{(c)}$ are nonzero in that interval. Figure 6(C) shows the proportion of times $\alpha_t^{(c)} \beta_t^{(c)}$ was deemed significant as a function of t . In the intervals $[0, 2/6]$ and $(3/6, 1]$, almost all time points fall below 0.05. In addition, we see a significant relationship in the range $[2/6, 3/6]$, as expected. The proportion of significant results almost reaches 100% around the point where h_t takes its maximum.

6. EXPERIMENTAL DATA

In this section, we study data from the fMRI study of thermal pain ($n = 20$) described in the Introduction. Functional MRI data were extracted from 21 different classic pain-responsive brain regions. Each time course consisted of 23 equidistant temporal measurements made every 2 sec, providing a total of 46 sec of brain activation, ranging from the time of the application of the heat stimuli to the pain report. The stimuli consisted of thermal stimulations delivered to the left volar forearm that participants judged to be nonpainful versus near the limit of pain tolerance. The temperature of these painful (i.e., hot) and nonpainful (i.e., warm) stimulations was determined via a pain calibration task that took place prior to the experiment on the day of scanning. Following an 18-sec interval of thermal stimulation, a fixation cross was presented for a 14-sec interval until the words “How painful?” appeared on the screen. After a few seconds of contemplation, the participants rated the overall pain intensity on a 10-point numerically anchored visual analog scale (VAS), similar to those commonly used in clinical practice. Participants respond by indicating a position along a continuous line between two endpoints. The continuous aspect of the scale differentiates it from similar discrete scales (e.g., the Likert scale) and allows us to use the suggested model, instead of a variant specifically designed for ordinal responses.

Each time course was placed into the three-variable path model shown in Figure 2, where the variable Z represents the applied pain level, the variable Y the reported pain, and the variable M the brain response. We began by estimating the parameters using the lfSEM framework. We restrict M_t to a uniform grid of 23 points in the range $[0, 46]$. The path coefficients $\alpha_t^{(s)}$ and $\beta_t^{(s)}$ were modeled using a B-spline basis set of order 6, with 10 equidistantly spaced knots. Inference was performed using a bootstrap test to determine whether $\alpha_t^{(s)} \beta_t^{(s)}$ was significantly different from 0 in any time interval between pain application and pain report. The Benjamini–Hochberg procedure was used to control the false detection rate ($q = 0.05$). Next, we estimate the IV estimand using the technique outlined in Section 4.

6.1 Results of Functional Mediation Analysis

Figure 7 shows results from the right anterior insula, which has been shown to be related to negative emotional experience. This region is believed to mediate the relationship between temperature and pain rating. Estimates of the $\alpha_t^{(s)}$ and $\beta_t^{(s)}$ pathway functions are shown on either side of the path model. These functions suggest a sustained period for which activity is modulated by temperature, and a more phasic response predicting perceived pain, controlling for temperature. A 95% bootstrap percentile confidence interval shows that the $\alpha_t^{(s)}\beta_t^{(s)}$ effect is significantly nonzero between 20 and 24 sec after the start of the trial, indicating the key time interval driving the mediation. This corresponds to the first 4 sec following the end of the application of heat. This somewhat delayed effect is due to the sluggish nature of brain hemodynamics, which peaks roughly 6 sec after peak neuronal activation, and is consistent with timings of other fMRI experiments (Lindquist 2008).

Figure 8 shows time points with significant nonzero $\alpha_t^{(s)}\beta_t^{(s)}$ effects for each of the 21 brain regions obtained using a bootstrap test. The results show that data in many classic “pain-responsive regions” such as the anterior insula (AINS) show significant mediation of the temperature–report relationship, particularly around the end of the treatment, that is, heat application (20–24 sec). The subjective pain grows during the stimulation period, and is often maximal around the end of stimulation, which typically drives pain ratings made after the trial. Notably, some other regions show significant mediation effects around the time immediately preceding pain reporting (38–44 sec), perhaps signaling a contribution of activity during “pain recall.” In particular, regions in the insular cortex appear to be active during pain judgment, which is reasonable due to their link to decision making (Moulton et al. 2005; Grinband, Hirsh, and Ferrera 2006). In addition to responding to physical pain, the anterior insula has been shown to be activated by viewing others in pain, receiving an unfair offer in an economic game, or viewing aversive pictures (Amodio and Frith 2006; Wager et al. 2008).

The ability to identify these disjoint time intervals as being involved in mediating the effect between temperature and reported pain is a unique property of the functional mediation framework and is not possible using standard mediation techniques. Hence, we believe that the application of IfSEM provides invaluable information about different components of pain processing in the human brain. This, in turn, furthers our understanding of the components involved in the processing and reporting of applied pain.

Figure 9 shows the IV estimands for each of the 21 brain regions. The results show that for most regions, there are three distinct time intervals of significant effects. The first, roughly 4–16 sec following the start of the trial, corresponds to a negative value of the IV estimand. Hence, high brain activation in the early portion of the trial has a negative effect on reported pain. This time period is typically thought of as a baseline period before task-related activation has had time to appear and any activation may be related to anticipation. In contrast, the time periods 20–28 sec and 38–44 sec correspond to a positive value of the IV estimand. In these intervals, high activity has a positive effect on reported pain, and they are related to the timing of task-related activation and pain recall, respectively, as discussed above.

6.2 Assessment of Assumptions

The validity of the causal model defined in Section 2.1 depends upon the appropriateness of SUTVA and models (1) and (2). In our application, neither the brain response nor the reported pain of an individual subject should be affected by the treatment given to any other subject in the study, thus validating SUTVA. In addition, linear additive models are commonly used in the analysis of neuroimaging data, and their usage in this particular application appears reasonable. To partially verify linearity and additivity, we studied a number of independent datasets that used the same paradigm and matched values of the mediator to compute $Y_i(1, \mathbf{m}) - Y_i(0, \mathbf{m})$ for different values of $\mathbf{m}$. The results indicate only minor deviations from linearity or additivity.

To apply a causal interpretation to the parameters of the lfSEM and the IV estimand, we require certain combinations of Assumptions 1–4 to hold. The requirements for equating the parameters of the lfSEM with the effects of the mediators are given by Assumptions 1 and 2. As treatments are randomly assigned to subjects, it is reasonable to assume that the treatment is ignorable. Assumption 2 would be valid if the mediators were randomly assigned to the subjects. However, this is not the case here, and instead, we must assume that they behave as if they were. This assumption is unverifiable in practice and ultimately depends on context. In the neuroimaging setting, its validity may differ across brain regions, making causal claims more difficult to access. For example, assume subjects in the study are either resilient or not to pain. Suppose that resilient subjects tend to have a lower brain response than nonresilient subjects when a painful stimulus is applied. Further suppose that resilient subjects always report lower pain ratings than their nonresilient counterparts at any value of the mediating variable. In this situation, $Y(1, \mathbf{m})$ will not be independent of $\mathbf{M}(1)$ and Assumption 2 will be violated. In practice, this assumption could potentially be weakened by allowing for conditioning on potential confounders. However, no such covariates were available in this particular study.

For a valid causal interpretation of the IV estimand, we require Assumptions 1, 3, and 4. The assumption of no average difference in potential errors appears plausible, in particular in comparison with Assumption 2. Revisiting the example with resilient/nonresilient subjects, here it suffices that the resilient (nonresilient) subjects on average fall above (below) the mean pain rating by the same amount under both conditions. Finally, Assumption 4 is untestable in practice but amounts to attributing the effect of Z on Y to the change in the brain response rather than the change in the treatment. One caveat is that since a single mediator model is assumed and multiple mediators have been targeted in the study, this assumption may potentially be violated as the effects of omitted mediators may become confounded with the direct effect. Ideally, we would wish to use a model that simultaneously considers the effects of multiple mediators on the outcome to circumvent this issue. We leave this for future research and simply note that any violation of Assumption 4 gives rise to bias. As a final note on assumptions, the application considers M_{it} (and the corresponding coefficients) as functions of time. Hence, it is conceivable that values of the mediator at earlier times could affect values at later times. The implicit assumption of our model is that there is no causal connection among the multiple

measurements of the mediator, that is, that they behave as if they were measured contemporaneously.

7. DISCUSSION

This article introduces linear functional structural equation modeling as a means of studying time-varying mediation effects. It also examines conditions under which the lfSEM and IV methods can be used to identify causal effects of a mediating functional variable on an outcome. To the best of our knowledge, this work provides the first application of causal inference to the FDA framework.

The causal interpretation for the parameters of the lfSEM rests on a strong untestable assumption, namely sequential ignorability (Assumptions 1 and 2). Recent advances in imaging may allow us to avoid Assumption 2 as neuroimaging data have been combined with transcranial magnetic stimulation (TMS) to integrate the ability of neuroimaging to observe brain activity with the ability of TMS to manipulate brain function (Bohning et al. 1997). Using this technique, one can simulate temporary “brain lesions” while the subject performs certain tasks. One can then attempt to infer causal relationships by studying differences in a brain network when a region is functioning and when it is not. Using this technique will allow us to experimentally manipulate the mediating variable, as well as the treatment, thus allowing us to circumvent the sequential ignorability assumption.

The proposed lfSEM framework assumes that the treatment and outcome variables are both univariate, while the intermediate variable is continuous. The method can be extended to allow the variables Z , $\mathbf{M}$, and Y to all be functional. There are several ways to formulate an analogous system of equations to (11) and (12) for this case. One example that uses the time-varying, or concurrent, model for functional regression (Hastie and Tibshirani 1993; Ramsay, Hooker, and Graves 2009) can be expressed as follows:

$$M_{it}(Z_{it}) = \delta_{1t} + \alpha_t Z_{it} + \varepsilon_{it} \quad (46)$$

$$Y_{it}(Z_{it}, M_{it}(Z_{it})) = \delta_{2t} + \gamma_t Z_{it} + \beta_t M_{it}(Z_{it}) + \eta_{it} \quad (47)$$

for $t \in [0, 1]$. In this formulation, the total effect of Z on Y can be expressed as

$$\tau_t = \gamma_t + \alpha_t \beta_t \quad (48)$$

for all $t \in [0, 1]$. Note that, here, the total effect is allowed to vary across the entire range of t .

Finally, an interesting aspect of our approach that we have not previously discussed is its ability to handle missing, or nonuniformly sampled, data in the mediator variable. One of the main benefits of FDA is its ability to handle such data. Though our example did not contain such data, it would be an interesting application for future work, as fMRI scans are sometimes dropped due to artifacts in the data collection at certain time points.

Acknowledgments

The author thanks Michael Sobel for all his help with the preparation of this article, Niall Bolger for helpful comments, and Tor Wager for supplying the data. This work was partially funded by the National Institute on Drug Abuse (NIDA) grant 1RC1DA028608.

References

- Albert JM. Mediation Analysis via Potential Outcomes Models. *Statistics in Medicine*. 2008; 27:1282–1304. [PubMed: 17691077]
- Amodio DM, Frith CD. Meeting of Minds: The Medial Frontal Cortex and Social Cognition. *Nature Reviews Neuroscience*. 2006; 7:268–277.
- Angrist JD, Imbens GW. Two-Stage Least Squares Estimation of Average Causal Effects in Models With Variable Treatment Intensity. *Journal of the American Statistical Association*. 1995; 90:431–442.
- Angrist JD, Imbens GW, Rubin DB. Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*. 1996; 91:444–455.
- Baron RM, Kenny DA. The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic and Statistical Considerations. *Journal of Personality and Social Psychology*. 1986; 51:1173–1182. [PubMed: 3806354]
- Benjamini Y, Hochberg Y. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society, Series B*. 1995; 57:289–300.
- Bohning DE, Pecheny AP, Epstein CM, Speer AM, Vincent DJ, Dannels W, George M. Mapping Transcranial Magnetic Stimulation (TMS) Fields in Vivo With MRI. *Neuroreport*. 1997; 8:2535–2538. [PubMed: 9261822]
- Cardot H, Ferraty F, Sarda P. Functional Linear Model. *Statistics & Probability Letters*. 1999; 45:11–22.
- Davidson R, MacKinnon JG. Wild Bootstrap Tests for IV Regression. *Journal of Business & Economic Statistics*. 2010; 28:128–144.
- Efron, B.; Tibshirani, RJ. *An Introduction to the Bootstrap*. London: Chapman & Hall/CRC; 1998.
- Fan J, Zhang J-T. Two-Step Estimation of Functional Linear Models With Applications to Longitudinal Data. *Journal of the Royal Statistical Society, Series B*. 1999; 62:303–322.
- Flachaire E. Bootstrapping Heteroskedastic Regression Models: Wild Bootstrap vs. Pairs Bootstrap. *Computational Statistics & Data Analysis*. 2005; 49:361–376.
- Grinband J, Hirsh J, Ferrera VP. A Neural Representation of Categorization Uncertainty in the Human Brain. *Neuron*. 2006; 49:757–763. [PubMed: 16504950]
- Hastie T, Tibshirani RJ. Varying-Coefficient Models. *Journal of the Royal Statistical Society, Series B*. 1993; 55:757–796.
- Holland PW. Causal Inference, Path Analysis, and Recursive Structural Equations Models. *Sociological Methodology*. 1988; 18:449–484.
- Imai K, Keele L, Yamamoto T. Identification, Inference, and Sensitivity Analysis for Causal Mediation Effects. *Statistical Science*. 2010; 25(1):51–71.
- James GM. Generalized Linear Models With Functional Predictors. *Journal of the Royal Statistical Society, Series B*. 2002; 64:411–432.
- Jo B. Causal Inference in Randomized Experiments With Mediational Processes. *Psychological Methods*. 2008; 13:314–336. [PubMed: 19071997]
- Lindquist MA. The Statistical Analysis of fMRI Data. *Statistical Science*. 2008; 23:439–464.
- Lindquist MA, Sobel ME. Graphical Models, Potential Outcomes and Causal Inference: Comment on Ramsey, Spirtes and Glymour. *NeuroImage*. 2011; 57:334–336. [PubMed: 20970507]
- Lindquist, MA.; Sobel, ME. Cloak and DAG: A Response to the Comments on our Comment. *NeuroImage*. 2012. available at <http://dx.doi.org/10.1016/j.neuroimage.2011.11.027>
- MacKinnon, DP. *Introduction to Statistical Mediation Analysis*. Mahwah, NJ: Erlbaum; 2008.

- Moulton EA, Keaser ML, Gullapalli RP, Greenspan JD. Regional Intensive and Temporal Patterns of Functional MRI Activation Distinguishing Noxious and Innocuous Contact Heat. *Journal of Neurophysiology*. 2005; 93:2183–2193. [PubMed: 15601733]
- Müller HG, StadtMüller U. Generalized Functional Linear Models. *The Annals of Statistics*. 2005; 33:774–805.
- Pearl, J. Direct and Indirect Effects. *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*; San Francisco, CA: Morgan Kaufman; 2001. p. 411–420.
- Petersen ML, Sinisi SE, van der Laan MJ. Estimation of Direct Causal Effects. *Epidemiology*. 2006; 17:276–284. [PubMed: 16617276]
- Ramsay, JO.; Hooker, G.; Graves, S. *Functional Data Analysis With R and MATLAB*. New York: Springer; 2009.
- Ramsay, JO.; Silverman, BW. *Functional Data Analysis*. 2. New York: Springer; 2006.
- Rosenbaum PR, Rubin DB. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*. 1983; 70:41–55.
- Robins JM, Greenland S. Identifiability and Exchangeability for Direct and Indirect Effects. *Epidemiology*. 1992; 3:143–155. [PubMed: 1576220]
- Rubin DB. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology*. 1974; 66:688–701.
- Rubin DB. Comment on ‘Randomization Analysis of Experimental Data: The Fisher Randomization Test,’ by D. Basu. *Journal of the American Statistical Association*. 1980; 75:591–593.
- Rubin DB. Direct and Indirect Causal Effects via Potential Outcomes. *Scandinavian Journal of Statistics*. 2004; 31:161–170.
- Shrout PE, Bolger N. Mediation in Experimental and Nonexperimental Studies: New Procedures and Recommendations. *Psychological Methods*. 2002; 7:422–445. [PubMed: 12530702]
- Sobel, ME. Asymptotic Confidence Intervals for Indirect Effects in Structural Equation Models,” in: Leinhardt, S., editor. *Sociological Methodology*. Washington, DC: American Sociological Association; 1982. p. 290–312.
- Sobel ME. Identification of Causal Parameters in Randomized Studies With Mediating Variables. *Journal of Educational and Behavioral Statistics*. 2008; 33:230–251.
- Ten Have TR, Joffe MM, Lynch KG, Brown GK, Maisto SA, Beck AT. Causal Mediation Analyses With Rank Preserving Models. *Biometrics*. 2007; 63:926–934. [PubMed: 17825022]
- VanderWeele TJ. Marginal Structural Models for the Estimation of Direct and Indirect Effects. *Epidemiology*. 2009; 20:18–26. [PubMed: 19234398]
- Wager TD, Davidson M, Hughes B, Lindquist MA, Ochsner K. Prefrontal-Subcortical Pathways Mediating Successful Emotion Regulation. *Neuron*. 2008; 59:1037–1050. [PubMed: 18817740]
- Zhu H, Ibrahim JG, Tang N, Rowe DB, Hao X, Bansal R, Peterson BS. A Statistical Analysis of Brain Morphology Using Wild Bootstrapping. *IEEE Transactions on Medical Imaging*. 2007; 26:954–966. [PubMed: 17649909]

Figure 1.

The three-variable path diagram used to represent the standard mediation framework. The variables corresponding to Z , Y , and M are all scalar, as are the path coefficients α , β , and γ linking them.

Figure 2.

The three-variable path diagram used to represent the functional mediation framework. The variables corresponding to Z and Y are scalars, while the variable corresponding to M is a function. Both the α_t and the β_t pathways are represented by functions, while the γ pathway is a scalar.

	Z	M_t	Y	αβ(t)-effect
Simulation 1	$Z_i = \{0,1\}$	ε_{it}	$Z_i + \eta_i$	None
Simulation 2	$Z_i = \{0,1\}$	$(Z_i + \xi_i)h_t + \varepsilon_{it}$	$Z_i + \eta_i$	None
Simulation 3	$Z_i = \{0,1\}$	$\xi_i h_t + \varepsilon_{it}$	$\int_{T/3}^{2T/3} M_{it} dt + \eta_i$	None
Simulations 4 & 5	$Z_i = \{0,1\}$	$(Z_i + \xi_i)h_t + \varepsilon_{it}$	$\int_{T/3}^{2T/3} M_{it} dt + \eta_i$
Note: $h_t = $ $\varepsilon_{it}, \eta_i \sim N(0,1)$ $\xi_i \sim N(1,0.1)$				Green indicates range of expected significant functional mediation

Figure 3.

A schematic overview of the first five simulation studies, along with expected $\alpha_t^{(s)}\beta_t^{(s)}$ effects (last column). Simulations 1–3 were designed to determine the empirical false positive rate of the method when no mediation effect is present, while Simulations 4 and 5 were designed to study the power in the presence of a mediation effect. In Simulations 1–4, data were generated for 20 “subjects,” where half were assigned to a treatment group ($Z = 1$) and half to a control group ($Z = 0$). In Simulation 5, the sample size was allowed to vary between 10 and 50 “subjects.” The values of M_t and Y differed between the five simulations, as summarized in the schematic, to mimic plausible brain imaging settings.

Figure 4.

Estimates of $\alpha_t^{(s)}$, $\beta_t^{(s)}$, and $\alpha_t^{(s)} \beta_t^{(s)}$, together with uncorrected 95% bootstrap percentile confidence intervals, for data generated according to the settings described in the first four simulations studies.

Figure 5.

Results of the first three simulation studies are shown in (A), (B), and (C), respectively.

Each plot shows the proportion of times $\alpha_t^{(s)} \beta_t^{(s)}$ was deemed significant as a function of time. The results illustrate that the method provides adequate control of the false positive rate in all three simulated scenarios.

Figure 6.

Results of the last three simulation studies: (A) The plot shows the results of the fourth simulation. It shows the proportion of times $\alpha_t^{(s)} \beta_t^{(s)}$ was deemed significant as a function of time and illustrates the power of the method in detecting true positives in the interval $[1/3, 2/3]$, while appropriately controlling for false positives in the intervals $[0, 1/3]$ and $[2/3, 1]$. (B) The plot illustrates the power to detect true positives as a function of sample size, with values ranging from 10 to 50. As expected, the larger the sample size, the better the sensitivity and specificity. (C) The plot shows the results of the sixth simulation and illustrates the power of the IV estimand to detect true positives.

Figure 7.

Results show that activation in the right anterior insula (see the brain map for the anatomical location of the region) mediates the relationship between temperature and pain rating.

Estimates of the $\alpha_t^{(s)}$ and $\beta_t^{(s)}$ pathway functions are shown on either side of the path model. These functions suggest a sustained period for which activity is modulated by temperature, and a more phasic response predicting perceived pain, controlling for temperature. An uncorrected 95% bootstrap percentile confidence interval shows that the $\alpha_t^{(s)}\beta_t^{(s)}$ effect is significantly nonzero between 20 and 24 sec following the start of the trial (shown in red in the bottom plot). This indicates the key time interval driving the mediation.

Figure 8.

Time points with significant nonzero $\alpha_t^{(s)}\beta_t^{(s)}$ effects for each of the 21 brain regions (FDR-corrected p -value < 0.05). The color bar shows the size of the p -value. The results show that classic “pain-responsive regions” show significant mediation of the temperature–report relationship, particularly around the end of the treatment (20–24 sec). Some other pain regions show significant mediation effects around the time of pain reporting, perhaps signaling a contribution of activity during pain recall.

Figure 9.

Time points with significant $\beta_t^{(c)}$ effects for each of the 21 brain regions (FDR-corrected p -value < 0.05). The color bar shows the amplitude of the estimate, with results obtained using the IV estimand.