

MULTIVARIATE DYNAMIC MEDIATION ANALYSIS UNDER A REINFORCEMENT LEARNING FRAMEWORK

BY LAN LUO^{1,a*}, CHENGCHUN SHI^{2,b*}, JITAO WANG^{3,c*},
ZHENKE WU^{3,d} AND LEXIN LI^{4,e}

¹*Department of Biostatistics and Epidemiology, Rutgers University, l.luo@rutgers.edu*

²*Department of Statistics, London School of Economics and Political Science, C.Shi7@lse.ac.uk*

³*Department of Biostatistics, University of Michigan, jitwang@umich.edu; zhenkewu@umich.edu*

⁴*Division of Biostatistics, University of California at Berkeley, lexinli@berkeley.edu*

Mediation analysis is an important analytic tool commonly used in a broad range of scientific applications. In this article, we study the problem of mediation analysis when there are multivariate and conditionally dependent mediators, and when the variables are observed over multiple time points. The problem is challenging, because the effect of a mediator involves not only the path from the treatment to this mediator itself at the current time point, but also all possible paths pointed to this mediator from its upstream mediators, as well as the carryover effects from all previous time points. We propose a novel multivariate dynamic mediation analysis approach. Drawing inspiration from the Markov decision process model that is frequently employed in reinforcement learning, we introduce a Markov mediation process paired with a system of time-varying linear structural equation models to formulate the problem. We then formally define the individual mediation effect, built upon the idea of simultaneous interventions and intervention calculus. We next derive the closed-form expression, propose an iterative estimation procedure under the Markov mediation process model, and develop a bootstrap method to infer the individual mediation effect. We study both the asymptotic property and the empirical performance of the proposed methodology, and further illustrate its usefulness with a mobile health application.

1. Introduction. Mediation analysis is an important analytic tool, which seeks to explain the mechanism or pathway that underlies an observed relationship between a treatment and an outcome variable, through the inclusion of an intermediary variable known as a mediator. It decomposes the effect of the treatment on the outcome into a direct effect and an indirect effect, the latter of which indicates whether the mediator is on a pathway from the treatment to the outcome (Baron and Kenny, 1986). Mediation analysis is widely employed in a range of scientific applications, including psychology (MacKinnon, 2008; Rucker et al., 2011), genomics (Huang and Pan, 2016; Bi et al., 2017), economics (Celli, 2022), social science (Kaufman and Kaufman, 2001), neuroscience (Zhao and Luo, 2022), among many others. See VanderWeele (2016) for a comprehensive review and the references therein.

Mediation analysis has seen considerable progress in recent years. There are two lines of research of particular interest, mediation analysis with multivariate mediators, and mediation analysis with time-varying variables. The first line targets the scenario where there are multiple mediators, and one central goal is to evaluate and quantify the contribution of the treatment on the outcome attributed to each individual mediator. There are three main

*Equal contribution.

MSC2020 subject classifications: Primary 00X00, 00X00; secondary 00X00.

Keywords and phrases: Longitudinal data, Markov process, Mediation analysis, Mobile health, Reinforcement learning.

categories of solutions. One category explicitly imposes that the multivariate mediators are conditionally independent given the treatment, which substantially simplifies the analysis (Boca et al., 2014; Huang and Pan, 2016; Zhang et al., 2016; Guo et al., 2023; Shuai et al., 2023; Yuan and Qu, 2023). Another category does not impose such a condition, but instead marginalizes each individual mediator, which in effect neglects any potential interaction and dependency among the mediators (Sampson et al., 2018; Djordjilović, Hemerik and Thoresen, 2022; Zhao, Li and Initiative, 2022; Zhao and Luo, 2022). The last category does allow correlated mediators, characterizes their dependency through some unknown directed acyclic graph (DAG), then carries out mediation analysis based on the estimated DAG (Maathuis, Kalisch and Bühlmann, 2009; Chakraborty, Nandy and Li, 2018; Cai, Song and Lu, 2021; Shi and Li, 2022; Wei et al., 2024a). The second line targets the scenario where the treatment, mediator and outcome are observed over multiple time points, or stages. There have been some pioneering works along this line (Selig and Preacher, 2009; Preacher, 2015; VanderWeele and Tchetgen Tchetgen, 2017; Lin et al., 2017; Huang and Yuan, 2017; Zheng and van der Laan, 2017; Zhao et al., 2018; Hejazi et al., 2020; Cai et al., 2022; Díaz, Williams and Rudolph, 2022; Ge et al., 2023). Nevertheless, they all focus on the case where there is only a single mediator, and most only consider the case where the number of time points, or the time horizon, is finite.

In this article, we study the problem of mediation analysis when there are multivariate and conditionally dependent mediators, and when the variables are observed over multiple time points. Our motivation is a mobile health application from the Intern Health Study (IHS, NeCamp et al., 2020). It is a 26-week prospective longitudinal randomized trial targeting first-year training physicians in the United States (NeCamp et al., 2020). A key objective of this study is to investigate the effectiveness of in-the-moment mobile prompts in improving an intern’s mood score while minimizing user burden and expense. The prompts, delivered through a customized mobile app, consist of practical tips and life insights, such as reminders to have a break, take a walk, or prioritize sleep, that aim at promoting healthy behaviors among interns. These prompts not only affect an intern’s mood, but also influence other physiological measurements, including physical activity, sleep duration, and heart rate variability, which are closely associated with an individual’s mental state. We can naturally formulate this problem in the framework of mediation analysis, where the treatment is the binary variable that encodes receiving a prompt or not, regardless of the prompt type, the mediators consist of measurements of physical activities, sleep duration, and heart rate variability, and the outcome is the mood score. Given the presence of multiple mediators and the longitudinal nature of the data, it calls for a mediation analysis approach that tackles both multivariate mediators and multiple time points.

However, the problem is challenging, for several reasons. First, multivariate and conditionally correlated mediators introduce considerable complexity. Unlike the single mediator or conditionally independent mediator case, the indirect effect of each individual mediator involves not only the path from the treatment to this mediator itself, but also all possible paths pointed to this mediator from its upstream mediators. The total number of potential paths that go through any mediator is super-exponential in the number of mediators. It is thus crucial to carefully disentangle the individual mediation effect when the multivariate mediators are interrelated with unknown dependency. Second, time-varying variables introduce another layer of complexity, as there are carryover effects along time. Unlike the single-time-point case, the mediator at a given time point could affect both the current and future outcomes, and thus it is crucial to learn both the immediate effect and the delayed effect. Infinite time horizon further complicates the analysis when the mediator effect accumulates over infinite time.

To address those challenges, we propose a novel multivariate dynamic mediation analysis approach. Our proposal consists of four key components. First, drawing inspiration from

the Markov decision process (MDP, [Puterman, 2014](#)) model that is frequently employed in reinforcement learning (RL, [Sutton and Barto, 2018](#)) to handle the carryover effects over time, we introduce a Markov mediation process (MMP) framework to formulate the dynamic mediation analysis problem. We then introduce a system of time-varying linear structural equation models (SEMs) to specifically characterize the relations among the treatment, mediator and outcome variables. Second, within the framework of MMP, we formally define the *individual mediation effect*, which is built upon the idea of simultaneous interventions and intervention calculus ([Pearl, 2000](#)). This individual mediation effect can be further decomposed into a sum of the *immediate effect* and the *delayed effect*, the latter of which quantifies the carryover effects of the past treatments and mediators. Third, under the proposed MMP and the linear SEMs, we derive the closed-form expression for the individual mediation effect; see Theorems 1 and 2. These theorems form the basis of our proposed procedure, allow us to express the individual mediation effect using a set of within-stage and cross-stage intermediate quantities that can be estimated through some recursive formulations, and consist of the main contributions of our proposal. We next propose to estimate those within-stage quantities through linear regressions with backdoor covariate adjustment, and estimate those cross-stage quantities through the transition equations under SEMs. Finally, we study the asymptotic property of our proposed estimator of the individual mediation effect, develop a bootstrap method to construct its confidence interval (CI), and conduct extensive numerical studies to demonstrate the effectiveness of our method. In summary, our approach allows us to separately evaluate the contribution of each individual mediator, where multiple mediators are interrelated with unknown dependency and are observed over multiple time points.

The rest of the article is organized as follows. We present the Markov mediation process, the structural equation models, and the definition of the individual mediation effect in Section 2. We derive the intermediate quantities, the recursive formulation of the individual mediation effect, and the estimation algorithms in Section 3. We establish the asymptotic properties in Section 3.5. We carry out the simulations in Section 4, and revisit the IHS example in Section 5. We relegate all technical proofs to the Supplementary Appendix.

2. Model and Definition. In this section, we first present our model setup, then formally define the mediation effect of interest in the multivariate dynamic mediation analysis setting.

2.1. Markov mediation process and structural equation models. We first introduce a Markov mediation process framework to formulate the problem we target; see Figure 1(b) for a graphical illustration. Consider the treatment-mediator-outcome triplets $\{(A_t, \mathbf{M}_t, R_t) : t = 1, \dots, T\}$ over time, where T denotes the number of time points or stages. At each time point or stage t , a random treatment $A_t \in \mathbb{R}$ is administered, which subsequently affects a d -dimensional vector of potential mediators $\mathbf{M}_t \in \mathbb{R}^d$, and an outcome variable $R_t \in \mathbb{R}$. We assume they satisfy the Markov assumption, in that

$$(1) \quad (R_t, \mathbf{M}_t) \perp\!\!\!\perp (R_s, \mathbf{M}_s, A_{s+1})_{s < t-1} \mid (R_{t-1}, \mathbf{M}_{t-1}, A_t), \text{ for any } t \geq 1,$$

where $\perp\!\!\!\perp$ denotes statistical independence. We remark the Markov condition like (1) is widely imposed in sequential data problems ([Sutton and Barto, 2018](#)). Suppose the observed data consists of n independent and identically distributed (i.i.d.) realizations of the triplets $\{(A_t, \mathbf{M}_t, R_t) : t = 1, \dots, T\}$. We allow the data to be either densely or sparsely observed. Moreover, for simplicity, we assume that all subjects have the same T . Nevertheless, our proposed method can be adapted to accommodate the setting when T varies among subjects, and when the time lags between two time points differ. See Section 6.4 for more details.

In addition to the Markov condition, we further assume that the random treatment assignment is independent of the prior information, in that

$$(2) \quad A_t \perp\!\!\!\perp (A_s, \mathbf{M}_s, R_s), \text{ for any } s < t.$$

Condition (2) holds in the sequentially randomized trials naturally, including our IHS example, which is the main setting we target in this article. Meanwhile, our framework can also be extended to the setting that involves certain static baseline confounders that may influence the treatment selection, thus accommodating the data from observational studies. Specifically, letting X denote the baseline confounders, Condition (2) can be relaxed to $A_t \perp\!\!\!\perp (A_s, \mathbf{M}_s, R_s) | X$, for any $s < t$. The individual mediation effect that we define later can be similarly derived by incorporating X into the model. For presentation simplicity, however, we choose not to include those confounders in this article.

We next introduce a system of time-varying structural equation models,

$$(3) \quad (\mathbf{M}_t - \boldsymbol{\mu}_t) = \mathbf{W}_t^\top (\mathbf{M}_{t-1} - \boldsymbol{\mu}_{t-1}) + \boldsymbol{\epsilon}_{wt},$$

where $\boldsymbol{\mu}_t \equiv \mathbb{E}(\mathbf{M}_t | A_t, \mathbf{M}_{t-1}, R_{t-1})$ is the condition mean function, $\mathbf{W}_t \in \mathbb{R}^{d \times d}$ is the weight matrix, such that $(\mathbf{W}_t)_{ij} \neq 0$ if and only if mediator M_{ti} is a parent of M_{tj} , i.e., M_{ti} is in the parent set, $M_{ti} \in \text{Pa}(M_{tj}) \equiv \{M_{ti} : (\mathbf{W}_t)_{ij} \neq 0, i \in \{1, \dots, d\} \setminus \{j\}\}$, and $\boldsymbol{\epsilon}_{wt} = (\epsilon_{w1t}, \dots, \epsilon_{wdt})^\top \in \mathbb{R}^d$ is a vector of mean zero random errors. In model (3), the weight matrix $\mathbf{W}_t$ models the interactions among the mediators $\mathbf{M}_t$ at each time point, and the conditional mean $\boldsymbol{\mu}_t$ characterizes the dynamic dependence over time.

We then consider the linear models for $\boldsymbol{\mu}_t$ and R_t , in that

$$(4) \quad \begin{aligned} \boldsymbol{\mu}_t &= \boldsymbol{\alpha}_{1t} + \delta_{1t} A_t + \boldsymbol{\Gamma}_{1t} \mathbf{M}_{t-1} + \boldsymbol{\zeta}_{1t} R_{t-1}, \\ R_t &= \alpha_{2t} + \delta_{2t} A_t + \boldsymbol{\gamma}_{2t}^\top \mathbf{M}_{t-1} + \zeta_{2t} R_{t-1} + \boldsymbol{\kappa}_t^\top \mathbf{M}_t + \epsilon_{rt}, \end{aligned}$$

for some $\boldsymbol{\alpha}_{1t}, \delta_{1t}, \boldsymbol{\zeta}_{1t} \in \mathbb{R}^d$, $\boldsymbol{\Gamma}_{1t} \in \mathbb{R}^{d \times d}$, and for $\boldsymbol{\gamma}_{2t}, \boldsymbol{\kappa}_t \in \mathbb{R}^d$, $\alpha_{2t}, \delta_{2t}, \zeta_{2t} \in \mathbb{R}$, and some mean zero errors $\{\epsilon_{rt}\}_t$ independent over time, respectively. In model (4), $\boldsymbol{\Gamma}_{1t}$ characterizes the effect of $\mathbf{M}_{t-1}$ on $\mathbf{M}_t$. All random variables in (4) are assumed to have finite second moments. Let $\boldsymbol{\Theta}_{1t} \equiv (\boldsymbol{\alpha}_{1t}, \delta_{1t}, \boldsymbol{\Gamma}_{1t}, \boldsymbol{\zeta}_{1t}) \in \mathbb{R}^{d \times (d+3)}$ collect all the parameters for $\boldsymbol{\mu}_t$, and $\boldsymbol{\Theta}_{2t} \equiv (\alpha_{2t}, \delta_{2t}, \boldsymbol{\gamma}_{2t}^\top, \zeta_{2t}, \boldsymbol{\kappa}_t^\top) \in \mathbb{R}^{(2d+3) \times 1}$ collect all the parameters for R_t .

We make a few remarks. First, we draw a connection between the proposed MMP and MDP commonly studied in RL – a powerful machine learning technique for optimal sequential decision making (Murphy, 2003; Mnih et al., 2015; Silver et al., 2017; Qin, Zhu and Ye, 2021; Zhang, Chen and Yang, 2023). Both capture the carryover effects over time. MDP achieves this by introducing a sequence of time-varying feature variables, referred to as the states. It then models the carryover effects through state transitions, allowing past treatments to affect future outcomes through their impact on future states; see Figure 1(a) for an illustration. This approach has gained substantial attention for policy evaluation, serving to model both immediate and long-term effects of a target policy (Luckett et al., 2020; Hao et al., 2021; Liao, Klasnja and Murphy, 2021; Kallus and Uehara, 2022; Hu and Wager, 2022; Liao et al., 2022; Ramprasad et al., 2022; Shi et al., 2022; Liu et al., 2023; Wang, Qi and Wong, 2023). In a similar vein, our MMP operates by modeling the indirect influence of a preceding mediator via the transitions of mediator-outcome pairs. In essence, a past mediator influences future outcomes by exerting its effects on both the prior outcome and the subsequent mediator; see Figure 1(b) for an illustration. Second, the relation in (3) should be understood as a data generating mechanism, rather than as a mere association. It corresponds to a directed acyclic graph (DAG). Third, to keep the presentation simple, we do not include any time-varying confounders, which may be incorporated into our solution in a relatively straightforward fashion. We do not consider any unobserved confounders either, because we consider random treatment assignments. We leave the case with unobserved confounders as future research. Finally, we consider linear type models in both (3) and (4). The analysis of mediation has been dominated by linear regression paradigms and such linearity assumption is commonly adopted in existing work, see for example, Nandy, Maathuis and Richardson




Fig 1: (a) Diagram of Markov decision process (MDP), where treatments depend on current states only, and (S_t, A_t, R_t) represents the state-treatment-reward triplet; (b) Diagram of the proposed Markov mediation process (MMP), where current mediators depend on previous mediator-reward pairs, and (A_t, M_t, R_t) represents the treatment-mediator-outcome triplet.

(2017); Chakraborty, Nandy and Li (2018); Shi and Li (2022). It is possible to extend to nonlinear type models, under which the definition of individual mediation effect we give later still holds, but is more difficult to evaluate.

In this article, we consider two different settings: the finite-horizon and the infinite-horizon, borrowing the concepts from the RL literature (Sutton and Barto, 2018). For the finite-horizon setting, the number of time points T is fixed and finite. It usually corresponds to the traditional longitudinal studies, where subjects receive a finite number of treatments over a predetermined time period. For the infinite-horizon setting, T approaches infinity theoretically. It often applies to the applications involving continuous monitoring systems such as wearable devices, where the data is densely collected at high frequencies, e.g., every second or minute. For the finite-horizon setting, we allow the process $\{A_t, M_t, R_t\}$ to be non-stationary. However, for the infinite-horizon setting, we require the process to be stationary over time; that is, the parameters $\mathbf{W}_t$, Θ_{1t} and Θ_{2t} are constant over time. Without stationarity, we need to extrapolate to infer system dynamics beyond the observed time.

In summary, we believe that our model framework provides a reasonable starting point for multivariate dynamic mediation analysis. The methodology under our setting is already complex enough, and deserves a careful investigation.

2.2. Total, direct and indirect effects. Before we formally define the individual mediation effect, which is the main target of interest in this article, we first introduce the notions of *total effect*, *natural direct effect*, and *natural indirect effect* in the finite-horizon setting. This is to facilitate the understanding of our target, and also to establish the connection with the classical mediation analysis literature.

For the total effect, consider a hypothetical intervention applied to the system, where we set all treatments to some value a uniformly over the entire population. This can be realized through Pearl's do-operator, $do(A_t = a)$, which generates an interventional distribution by removing the edges leading into $\{A_s : s \leq t\}$ in the corresponding DAG (Pearl, 2000). We denote the post-interventional expectation of R_t by $\mathbb{E}[R_t | do(A_s = a, \forall s \leq t)]$. We then define the total effect over T time points as

$$(5) \quad \text{TE} = \sum_{t=1}^T \frac{\partial}{\partial a} \mathbb{E}[R_t | do(A_s = a, \forall s \leq t)].$$

We make two remarks. First, each summand $\partial \mathbb{E}[R_t | do(A_s = a, \forall s \leq t)] / \partial a$ on the right-hand-side of (5) measures the total effect of treatment $\{A_s\}_{s \leq t}$ on R_t at time t . Second, when

the treatment takes a discrete value, the derivative in those terms can be substituted by the difference operator. For instance, if the treatment is binary, we obtain that

$$\text{TE} = \sum_{t=1}^T \left\{ \mathbb{E}[R_t \mid \text{do}(A_s = 1, \forall s \leq t)] - \mathbb{E}[R_t \mid \text{do}(A_s = 0, \forall s \leq t)] \right\}.$$

Next, we decompose the total effect into the sum of the natural direct effect and the natural indirect effect. Specifically, at each time point $t = 1, \dots, T$, we define the direct effect to be the portion of the total effect of a sequence of treatment variables $\{A_s : s \leq t\}$ on the outcome R_t that does not go through $\{M_s : s \leq t\}$. To measure such an effect, we consider the joint intervention on (A_s, M_s) through $\text{do}(A_s = a, M_s = \mathbf{m}_s)$, and denote the post-interventional expectation of R_t by $\mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}_s, \forall s \leq t)]$. We then define the natural direct effect over T time points as

$$\begin{aligned} \text{DE} = & \sum_{t=1}^T \int_{\mathbf{m}_1, \dots, \mathbf{m}_T} \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}_s, \forall s \leq t)] \\ & \times f(\mathbf{m}_1, \dots, \mathbf{m}_T \mid \text{do}(A_t = a, \forall 1 \leq t \leq T)) d\mathbf{m}_1 \cdots d\mathbf{m}_T, \end{aligned}$$

where f denotes the interventional probability density function of $(M_1, \dots, M_T)$ under the assumption that all treatments are set to a , and the term $\partial \mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}_s, \forall s \leq t)] / \partial a$ corresponds to the interventional effect of $\{A_s\}_{s \leq t}$ on R_t when setting the interventional values of $\{M_s\}_{s \leq t}$ to constants.

We then define the natural indirect effect as the difference between TE and DE, i.e., $\text{IE} = \text{TE} - \text{DE}$, which quantifies the portion of the effect of the treatment sequence on the outcome that goes through the mediator sequence.

We remark that, the definitions of TE, DE and IE do *not* require linear structural equation models like (4). However, imposing such models greatly simplify the analysis. More specifically, in a general Markov mediation process, both $\partial \mathbb{E}[R_t \mid \text{do}(A_s = a, \forall s \leq t)] / \partial a$ and $\partial \mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}_s, \forall s \leq t)] / \partial a$ can depend on $(a, \mathbf{m}_1, \dots, \mathbf{m}_t)$ in a very complex manner, as also mentioned in Chakraborty, Nandy and Li (2018). However, (4) simplifies these expressions to be functions of the regression coefficients that are independent of specific values of $\mathbf{m}_s$ or a ; see Section 3 for more details. As a result, DE is equivalent to

$$(6) \quad \sum_{t=1}^T \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}, \forall s \leq t)],$$

for any $\mathbf{m}$, where (6) is the *controlled direct effect*, computed by setting all values of mediators to $\mathbf{m}$. Correspondingly, by (5) and (6), the natural indirect effect becomes

$$(7) \quad \sum_{t=1}^T \left\{ \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, \forall s \leq t)] - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, M_s = \mathbf{m}, \forall s \leq t)] \right\}.$$

In other words, the linear structural equation model avoids the need to estimate the interventional probability density function f , thus substantially simplifying the forms of DE, IE, and their subsequent estimation and inference procedures.

2.3. Individual mediation effect. We now formally define the individual mediation effect, the main target in our dynamic mediation analysis. We begin with the finite-horizon setting.

DEFINITION 1 (Individual mediation effect in finite-horizon). *In a finite-horizon setting, the individual mediation effect of the j th mediator over T time points is defined as*

$$\eta_j^{(T)} \equiv \sum_{t=1}^T \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, \forall s \leq t)] - \sum_{t=1}^T \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a, M_{sj} = m, \forall s \leq t)].$$

When the treatment is binary, $\eta_j^{(T)}$ can be defined as

$$(8) \quad \eta_j^{(T)} = \sum_{a=0}^1 (-1)^{a+1} \left\{ \sum_{t=1}^T \mathbb{E}[R_t \mid do(A_s = a, \forall s \leq t)] - \sum_{t=1}^T \mathbb{E}[R_t \mid do(A_s = a, M_{sj} = m, \forall s \leq t)] \right\}.$$

We again make a few remarks. First, our definition of the individual mediation effect is related to the definition of IE in (7), and is defined as the difference between the total effect and the interventional effect $\partial \mathbb{E}[R_t \mid do(A_s = a, M_{sj} = m, \forall s \leq t)] / \partial a$ over time. However, the key difference is that, whereas IE measures the mediation effect of *all* mediators, $\eta_j^{(T)}$ focuses on quantifying the portion of the effect that specifically passes through the *individual* j th mediator. As such, we only intervene the j th mediator as opposed to all mediators in the intervention, and we refer to it as an *individual* mediation effect.

Second, our definition is consistent with the existing literature. Specifically, when there is only one mediator M and a single time point, our definition aligns with the classical definition of [Baron and Kenny \(1986\)](#), $\partial \mathbb{E}[R \mid do(A = a)] / \partial a - \partial \mathbb{E}[R \mid do(A = a, M = m)] / \partial a$, where the first term is the regression coefficient by regressing R on A , and the second term is the partial regression coefficient by further including M in the regression. This difference measuring the reduction in the total effect due to controlling for M is widely used to quantify the effect mediated through M ; see also [Pearl \(2012\)](#). When there is a single time point, our definition is consistent with the one proposed by [Chakraborty, Nandy and Li \(2018\)](#). When there is a single mediator, (8) is similar to the definitions proposed by [VanderWeele and Tchetgen Tchetgen \(2017\)](#) and [Ge et al. \(2023\)](#) under a potential outcome framework.

Third, under the linear structural equation model (4), $\eta_j^{(T)}$ is a constant function with respect to m and a . It can depend on (a, m) in a more general Markov mediation process, e.g., when either the transition model from a given time point to the next or the DAG structural equation model at a given time point is nonlinear. In addition, we observe that, under (4), the second post-interventional expectation term, $\mathbb{E}[R_t \mid do(A_s = a, M_{sj} = m, \forall s \leq t)]$, may be analyzed by the path method ([Wright, 1921](#)). That is, it can be calculated by summing up the effects along all directed paths from $\{A_s\}_{s \leq t}$ to R_t that do not pass through $\{M_{sj}\}_{s \leq t}$. However, this approach can be computationally expensive, since the number of paths grows exponentially fast as the number of time points T increases.

Finally, $\eta_j^{(T)}$ in Definition 1 measures the *cumulative* effect mediated through the j th mediator M_j across all T stages. Alternatively, one may be interested in the *incremental* effect,

$$\Delta_j^{(t)} \equiv \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, \forall s \leq t)] - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, M_{sj} = m, \forall s \leq t)],$$

which measures the individual mediation effect attributed to M_j at time t . By definition, we see that this incremental effect is related to the cumulative effect, in that $\eta_j^{(T)} = \sum_{t=1}^T \Delta_j^{(t)}$.

To better understand our definition of the individual mediation effect, we further decompose $\Delta_j^{(t)}$ into the sum of the *immediate* individual mediation effect (IIME) and the *delayed* individual mediation effect (DIME), as defined next.

DEFINITION 2 (Immediate and delayed individual mediation effects in finite-horizon). *Define the immediate individual mediation effect (IIME) and the delayed individual mediation effect (DIME) as,*

$$\text{IIME}_j^{(t)} \equiv \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_t = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_t = a, M_{tj} = m)],$$

$$\text{DIME}_j^{(t)} \equiv \Delta_j^{(t)} - \text{IIME}_j^{(t)}.$$

At a given time point t , $\text{IIME}_j^{(t)}$ can be interpreted as the change in the total effect of A_t on R_t when M_{tj} is knocked out, whereas $\text{DIME}_j^{(t)}$ captures the individual mediation effect of the j th mediator that is carried over from all previous stages $s < t$ up to time t .

We next turn to the infinite-horizon setting. In this setting, the individual mediation effect is defined as the *average* effect over time, so to prevent it from being unbounded.

DEFINITION 3 (Individual mediation effect in infinite-horizon). *In an infinite-horizon setting, the individual mediation effect of the j th mediator is defined as $\eta_j^{(\infty)} \equiv \lim_{T \rightarrow \infty} \eta_j^{(T)} / T$, provided that the limit exists.*

We comment that all above definitions are natural and agree with the intuitions. We next discuss how to evaluate and estimate the individual mediation effect given the data.

3. Evaluation of Individual Mediation Effect. In this section, we propose approaches to estimate and infer the individual mediation effects given in Definitions 1 and 3. The problem, however, is very challenging, as we have to deal with both the multiple mediators with unknown correlation structure, as well as the carryover effects from the upstream treatments and mediators. Our proposed solution involves deriving a closed-form expression for the individual mediation effect, as detailed in Theorems 1 and 2. This expression is dependent on several intermediate quantities that are computable through a set of recursive relations. Leveraging these theoretical findings, we develop an efficient recursive computation algorithm to estimate these intermediate quantities, which subsequently facilitate both the estimation and inference of the individual mediation effect.

3.1. An illustrative example. To illustrate our theory, we first consider a simple example as shown in Figure 2, in which there are only $d = 2$ mediators and $T = 2$ time stages. We then extend our observations to more general cases with d mediators and T stages.

Suppose we focus on the individual mediation effect for the second mediator, i.e., $j = 2$. We begin with the first stage $t = 1$, as shown in Figure 2, left panel. By Definition 1,

$$(9) \quad \eta_2^{(1)} = \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid \text{do}(A_1 = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid \text{do}(A_1 = a, M_{12} = m)] \equiv \theta_{A_1 \rightarrow R_1} - \theta_{A_1 \rightarrow R_1}^{M_{12}},$$

where we use the notation $\theta_{A_1 \rightarrow R_1}$ to denote the total effect from A_1 to R_1 , and use $\theta_{A_1 \rightarrow R_1}^{M_{12}}$ to denote the total effect from A_1 to R_1 when the second mediator is intervened in the first stage.

We compute $\theta_{A_1 \rightarrow R_1}$ by summing up the effects along all directed paths from A_1 to R_1 , namely, $A_1 \rightarrow R_1$, $A_1 \rightarrow M_{11} \rightarrow R_1$, $A_1 \rightarrow M_{12} \rightarrow R_1$, and $A_1 \rightarrow M_{11} \rightarrow M_{12} \rightarrow R_1$. Meanwhile, we compute $\theta_{A_1 \rightarrow R_1}^{M_{12}}$ by eliminating all paths that go from A_1 to R_1 through M_{12} . This corresponds to subtracting the effects along $A_1 \rightarrow M_{12} \rightarrow R_1$ and $A_1 \rightarrow M_{11} \rightarrow M_{12} \rightarrow R_1$, which leads to

$$(10) \quad \theta_{A_1 \rightarrow R_1}^{M_{12}} = \theta_{A_1 \rightarrow R_1} - \theta_{A_1 \rightarrow M_{12}} \theta_{M_{12} \rightarrow R_1},$$

where we use the notation $\theta_{A_1 \rightarrow M_{12}}$ and $\theta_{M_{12} \rightarrow R_1}$ to denote the effect from A_1 to M_{12} , and from M_{12} to R_1 , respectively. The relation in (10) has an intuitive interpretation: to evaluate the total effect from A_1 to R_1 when M_{12} is intervened, we subtract from $\theta_{A_1 \rightarrow R_1}$ the effects along the paths that go through M_{12} . Plugging (10) into (9), we obtain that




Fig 2: An illustrative example with two mediators and two time stages. The left panel shows the DAG in the first stage, and the right panel the first two stages. The dashed lines highlight the paths that go through the intervened mediator, i.e., the second mediator.

$\eta_2^{(1)} = \theta_{A_1 \rightarrow M_{12}} \theta_{M_{12} \rightarrow R_1}$, which coincides with the classical product type representation of the individual mediation effect in a single-stage analysis (Nandy, Maathuis and Richardson, 2017).

We next move to the second stage $t = 2$, as in Figure 2, right panel. By Definition 1,

$$(11) \quad \eta_2^{(2)} = \eta_2^{(1)} + \left\{ \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a, M_{12} = M_{22} = m)] \right\} \equiv \eta_2^{(1)} + \theta_{A_1, A_2 \rightarrow R_2} - \theta_{A_1, A_2 \rightarrow R_2}^{M_{12}, M_{22}}$$

where we use the notation $\theta_{A_1, A_2 \rightarrow R_2}$ to denote the cumulative total effect of (A_1, A_2) on R_2 , and use $\theta_{A_1, A_2 \rightarrow R_2}^{M_{12}, M_{22}}$ to denote the cumulative total effect of (A_1, A_2) on R_2 when the second mediator is intervened in both the first and second stages.

Because the treatments are randomly assigned and their parent sets are empty, we have,

$$(12) \quad \theta_{A_1, A_2 \rightarrow R_2} = \theta_{A_1 \rightarrow R_2} + \theta_{A_2 \rightarrow R_2}, \quad \theta_{A_1, A_2 \rightarrow R_2}^{M_{12}, M_{22}} = \theta_{A_1 \rightarrow R_2}^{M_{12}, M_{22}} + \theta_{A_2 \rightarrow R_2}^{M_{22}},$$

where we use the notation $\theta_{A_1 \rightarrow R_2}$ and $\theta_{A_2 \rightarrow R_2}$ to denote the total effects from A_1 to R_2 , and from A_2 to R_2 , respectively, use $\theta_{A_2 \rightarrow R_2}^{M_{22}}$ to denote the effect from A_2 to R_2 when the second mediator is intervened in the second stage, and use $\theta_{A_1 \rightarrow R_2}^{M_{12}, M_{22}}$ to denote the effect from A_1 to R_2 when the second mediator is intervened in both stages.

We compute $\theta_{A_1 \rightarrow R_2}$ and $\theta_{A_2 \rightarrow R_2}$ similarly as that for $\theta_{A_1 \rightarrow R_1}$. We compute $\theta_{A_2 \rightarrow R_2}^{M_{22}}$ similarly as in (10), i.e., $\theta_{A_2 \rightarrow R_2}^{M_{22}} = \theta_{A_2 \rightarrow R_2} - \theta_{A_2 \rightarrow M_{22}} \theta_{M_{22} \rightarrow R_2}$. Also, similar to (10), we have,

$$(13) \quad \theta_{A_1 \rightarrow R_2}^{M_{12}, M_{22}} = \theta_{A_1 \rightarrow R_2}^{M_{22}} - \theta_{A_1 \rightarrow M_{12}} \theta_{M_{12} \rightarrow R_2}^{M_{22}},$$

where we use $\theta_{A_1 \rightarrow R_2}^{M_{22}}$ and $\theta_{M_{12} \rightarrow R_2}^{M_{22}}$ to denote the effects from A_1 to R_2 , and from M_{12} to R_2 , when the second mediator is intervened in the second stage. Intuitively, $\theta_{M_{12} \rightarrow R_2}^{M_{22}}$ can be interpreted as the carryover effect of the second mediator on the outcome from the first stage to the second stage when it is intervened. We again compute $\theta_{A_1 \rightarrow R_2}^{M_{22}}$ similarly to (10), i.e., $\theta_{A_1 \rightarrow R_2}^{M_{22}} = \theta_{A_1 \rightarrow R_2} - \theta_{A_1 \rightarrow M_{22}} \theta_{M_{22} \rightarrow R_2}$. In addition, we compute $\theta_{M_{12} \rightarrow R_2}^{M_{22}}$ by eliminating all paths that go from M_{12} to R_2 through M_{22} , which leads to

$$(14) \quad \theta_{M_{12} \rightarrow R_2}^{M_{22}} = \theta_{M_{12} \rightarrow R_2} - \theta_{M_{12} \rightarrow M_{22}} \theta_{M_{22} \rightarrow R_2},$$

where we use $\theta_{M_{12} \rightarrow R_2}$, $\theta_{M_{12} \rightarrow M_{22}}$, and $\theta_{M_{22} \rightarrow R_2}$ to denote the effects from M_{12} to R_2 , from M_{12} to M_{22} , and from M_{22} to R_2 , respectively.

Based on the derivations so far, we see that, to evaluate the individual mediation effect, it is crucial to calculate those intermediate quantities, such as $\theta_{A_1 \rightarrow R_1}$, $\theta_{A_1 \rightarrow M_{22}}$, $\theta_{M_{12} \rightarrow M_{22}}$, $\theta_{M_{22} \rightarrow R_2}$, among others. Next, we briefly discuss how to evaluate those intermediate quantities after imposing the linear structural equation models (3) and (4). We also discuss some recursive relation that facilitates both the computational and statistical efficiencies.

First, we note that, under (3) and (4), we can estimate those intermediate quantities through linear regressions. For instance, we can estimate $\theta_{A_1 \rightarrow R_1}$ as the coefficient of A_1 by linearly regressing R_1 onto A_1 . Meanwhile, we can estimate $\theta_{M_{22} \rightarrow R_2}$ as the coefficient of M_{22} by linearly regressing R_2 onto M_{22} , however, with some additional covariate adjustment. This is because, unless the two mediators (M_{21}, M_{22}) are conditionally independent given $(M_{11}, M_{12}, R_1, A_2)$, we need to adjust for a set of covariates that satisfy Pearl's backdoor criterion (Pearl, 2000). In other words, we should block the effects flowing from $(A_2, M_{11}, M_{12}, M_{21}, R_1)$ to M_{22} . Therefore, we need to adjust for the covariate set $\text{Pa}(M_{22}) \cup M_{11} \cup M_{12} \cup R_1 \cup A_2$ in this regression.

Second, we note that, the set of intermediate quantities involve both within-stage quantities such as $\theta_{A_1 \rightarrow R_1}$, $\theta_{M_{22} \rightarrow R_2}$, as well as cross-stage quantities such as $\theta_{A_1 \rightarrow M_{22}}$, $\theta_{M_{12} \rightarrow M_{22}}$. There is some useful recursive relation between the within-stage and cross-stage quantities under the linear structural equation models (3) and (4). For instance,

$$(15) \quad \theta_{A_1 \rightarrow M_2} = \Gamma_{12} \theta_{A_1 \rightarrow M_1} + \zeta_{12} \theta_{A_1 \rightarrow R_1},$$

where $\theta_{A_1 \rightarrow M_2} = (\theta_{A_1 \rightarrow M_{21}}, \theta_{A_1 \rightarrow M_{22}})^\top$, and $\theta_{A_1 \rightarrow M_1} = (\theta_{A_1 \rightarrow M_{11}}, \theta_{A_1 \rightarrow M_{12}})^\top$. This suggests that the cross-stage carryover effect from A_1 to $M_2 = (M_{21}, M_{22})^\top$ is a combination of its within-stage effect on $M_1 = (M_{11}, M_{12})^\top$ and on R_1 , respectively, whereas the coefficients $(\Gamma_{12}, \zeta_{12})$ can be viewed as the weights for such a cross-stage transition. In our estimation algorithm, we first estimate the within-stage quantities via linear regressions, then update the cross-stage quantities following (15). As we show later, the relation such as (15) not only expedites the computation, but also improves the estimation efficiency by leveraging more information from the conditional model (4).

In summary, this simple example reveals a number of important relations. First, we see that those intermediate quantities form the building blocks for our evaluation of the individual mediation effect. Second, (10), (13), and (14) suggest some useful recursive representations that in effect reduce the number of mediators in the superscript that are intervened upon. Third, under models (3) and (4), we can estimate those intermediate quantities through linear regressions, with possibly backdoor adjustment. Finally, we can further improve both the computational and statistical efficiencies through another set of recursive relations between the within-stage and cross-stage quantities. All these observations are crucial for our derivation of the expression for the individual mediation effect.

3.2. Intermediate quantities. We now formally define the set of all intermediate quantities that are needed for the evaluation of the individual mediation effect. We then discuss how to estimate those quantities.

For any $1 \leq s \leq t \leq T$, and $j = 1, \dots, d$, define

$$(16) \quad \begin{aligned} \theta_{A_s \rightarrow R_t} &\equiv \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_s = a)] \in \mathbb{R}, & \theta_{M_{sj} \rightarrow M_t} &\equiv \frac{\partial}{\partial m} \mathbb{E}[M_t \mid \text{do}(M_{sj} = m)] \in \mathbb{R}^d, \\ \theta_{A_s \rightarrow M_t} &\equiv \frac{\partial}{\partial a} \mathbb{E}[M_t \mid \text{do}(A_s = a)] \in \mathbb{R}^d, & \theta_{M_{sj} \rightarrow R_t} &\equiv \frac{\partial}{\partial m} \mathbb{E}[R_t \mid \text{do}(M_{sj} = m)] \in \mathbb{R}. \end{aligned}$$

These intermediate quantities in (16) can be obtained through linear regressions, either directly, or by some backdoor covariate adjustment. In particular, $\theta_{A_s \rightarrow R_t}$ can be obtained as the coefficient of A_s by linearly regressing R_t onto A_s with an intercept, and we write this coefficient as β_{A_s, R_t} . Similarly, $\theta_{A_s \rightarrow M_t}$ can be obtained as the coefficient of A_s by linearly regressing M_t onto A_s , denoted as β_{A_s, M_t} . Meanwhile, $\theta_{M_{sj} \rightarrow R_t}$ can be obtained as the coefficient of M_{sj} by linearly regressing R_t onto M_{sj} , along with the adjusted covariate set, $\text{Pa}(M_{sj}) \cup M_{s-1} \cup R_{s-1} \cup A_s$. Similarly, $\theta_{M_{sj} \rightarrow M_t}$ can be obtained as the coefficient of M_{sj} by linearly regressing M_t onto M_{sj} , along with the adjusted covariate set, $\text{Pa}(M_{sj}) \cup M_{s-1} \cup R_{s-1} \cup A_s$. Putting together, we have,

$$(17) \quad \begin{aligned} \theta_{A_s \rightarrow R_t} &= \beta_{A_s, R_t}, & \theta_{M_{sj} \rightarrow M_t} &= \beta_{M_{sj}, M_t | \text{Pa}(M_{sj}) \cup M_{s-1} \cup R_{s-1} \cup A_s}, \\ \theta_{A_s \rightarrow M_t} &= \beta_{A_s, M_t}, & \theta_{M_{sj} \rightarrow R_t} &= \beta_{M_{sj}, R_t | \text{Pa}(M_{sj}) \cup M_{s-1} \cup R_{s-1} \cup A_s}. \end{aligned}$$

Moreover, similar to (15), we obtain the following recursive relations between the within-stage quantities and the cross-stage quantities. That is, under models (3) and (4),

$$(18) \quad \begin{aligned} \theta_{A_s \rightarrow R_t} &= \zeta_{2t} \theta_{A_s \rightarrow R_{t-1}} + \kappa_t^\top \theta_{A_s \rightarrow M_t} + \gamma_{2t}^\top \theta_{A_s \rightarrow M_{t-1}}, \\ \theta_{A_s \rightarrow M_t} &= \Gamma_{1t} \theta_{A_s \rightarrow M_{t-1}} + \zeta_{1t} \theta_{A_s \rightarrow R_{t-1}}, \\ \theta_{M_{sj} \rightarrow R_t} &= \zeta_{2t} \theta_{M_{sj} \rightarrow R_{t-1}} + \kappa_t^\top \theta_{M_{sj} \rightarrow M_t} + \gamma_{2t}^\top \theta_{M_{sj} \rightarrow M_{t-1}}, \\ \theta_{M_{sj} \rightarrow M_t} &= \Gamma_{1t} \theta_{M_{sj} \rightarrow M_{t-1}} + \zeta_{1t} \theta_{M_{sj} \rightarrow R_{t-1}}. \end{aligned}$$

In our implementation, we first estimate the within-stage quantities, $\{\theta_{A_t \rightarrow R_t}, \theta_{M_{tj} \rightarrow R_t}, \theta_{A_t \rightarrow M_t}, \theta_{M_{tj} \rightarrow M_t}\}$, using (17), then estimate the cross-stage quantities, $\{\theta_{A_s \rightarrow R_t}, \theta_{M_{sj} \rightarrow R_t}, \theta_{A_s \rightarrow M_t}, \theta_{M_{sj} \rightarrow M_t}\}$, for $s = 1, \dots, t-1$, using (18). This helps to improve both the estimation efficiency and the computation efficiency.

3.3. Individual mediation effect in a finite-horizon setting. We next derive the expression of the individual mediation effect through the intermediate quantities in (16) under the finite-horizon setting.

We first sketch the key ideas of the derivation here, and relegate more details to the Supplementary Appendix. By Definition 1, the individual mediation effect of the j th mediator across t stages, $j = 1, \dots, d$, $t = 1, \dots, T$, can be expressed as

$$(19) \quad \eta_j^{(t)} = \eta_j^{(t-1)} + \theta_{A_1, \dots, A_t \rightarrow R_t} - \theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}}.$$

Because all the treatments $(A_1, \dots, A_t)$ are randomly assigned, and thus are independent of each other and all other covariates, similar to (12), we have,

$$(20) \quad \begin{aligned} \theta_{A_1, \dots, A_t \rightarrow R_t} &= \theta_{A_1 \rightarrow R_t} + \theta_{A_2 \rightarrow R_t} + \dots + \theta_{A_t \rightarrow R_t}, \\ \theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} &= \theta_{A_1 \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} + \theta_{A_2 \rightarrow R_t}^{M_{2j}, \dots, M_{tj}} + \dots + \theta_{A_t \rightarrow R_t}^{M_{tj}}. \end{aligned}$$

Then, similar to (10), (13), and (14), we obtain the following recursive relations that help reduce the number of mediators being intervened upon. That is, for any $s = 1, \dots, t-1$,

$$(21) \quad \begin{aligned} \theta_{A_s \rightarrow R_t}^{M_{sj}, \dots, M_{tj}} &= \theta_{A_s \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} - \theta_{A_s \rightarrow M_{sj}} \theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}} = \dots \\ &= \theta_{A_s \rightarrow R_t} - \sum_{i=s}^t \theta_{A_s \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}, \end{aligned}$$

$$(22) \quad \begin{aligned} \theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}} &= \theta_{M_{sj} \rightarrow R_t}^{M_{(s+2)j}, \dots, M_{tj}} - \theta_{M_{sj} \rightarrow M_{(s+1)j}} \theta_{M_{(s+1)j} \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} = \dots \\ &= \theta_{M_{sj} \rightarrow R_t} - \sum_{i=s+1}^t \theta_{M_{sj} \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}. \end{aligned}$$

Algorithm 1 Estimation and inference for the finite-horizon setting.

Input: Observed data $\{(A_{it}, M_{it}, R_{it}) : i = 1, \dots, n, t = 1, \dots, T\}$, the number of bootstrap samples B , and the significance level α .

Output: Individual mediation effect $\hat{\eta}_j^{(t)}$, for $t = 1, \dots, T, j = 1, \dots, d$, and their CIs.

- 1: Estimate within-stage quantities $(\hat{\theta}_{A_1 \rightarrow R_1}, \hat{\theta}_{M_{1j} \rightarrow R_1}, \hat{\theta}_{A_1 \rightarrow M_1}, \hat{\theta}_{M_{1j} \rightarrow M_1})$, and compute $\hat{\eta}_j^{(1)} = \hat{\theta}_{A_1 \rightarrow M_{1j}} \hat{\theta}_{M_{1j} \rightarrow R_1}$.
 - 2: **for** $t = 2, \dots, T$ **do**
 - 3: Estimate parameters $\hat{\mathbf{W}}_t, \hat{\Theta}_{1t}, \hat{\Theta}_{2t}$ in models (3) and (4).
 - 4: Estimate within-stage quantities $(\hat{\theta}_{A_t \rightarrow R_t}, \hat{\theta}_{M_{tj} \rightarrow R_t}, \hat{\theta}_{A_t \rightarrow M_t}, \hat{\theta}_{M_{tj} \rightarrow M_t})$,
 - 5: **for** $s = 1, \dots, t - 1$ **do**
 - 6: Estimate cross-stage quantities $(\hat{\theta}_{A_s \rightarrow R_t}, \hat{\theta}_{M_{sj} \rightarrow R_t}, \hat{\theta}_{A_s \rightarrow M_t}, \hat{\theta}_{M_{sj} \rightarrow M_t})$ using (18).
 - 7: **end for**
 - 8: Compute $\hat{\eta}_j^{(t)}$ using (23).
 - 9: **end for**
 - 10: **for** $b = 1, \dots, B$ **do**
 - 11: Sample n trajectories from the observed data with replacement.
 - 12: Repeat Lines 1 to 9 to compute $\hat{\eta}_j^{(t,b)}$ using the bootstrap samples.
 - 13: **end for**
 - 14: **for** $t = 1, \dots, T$ **do**
 - 15: Construct the CI $[\hat{\eta}_j^{(t,L)}, \hat{\eta}_j^{(t,U)}]$ for $\eta_j^{(t)}$, where $\hat{\eta}_j^{(t,L)}$ and $\hat{\eta}_j^{(t,U)}$ are the empirical lower and upper
 - 16: $\alpha/2$ quantiles of $\{\hat{\eta}_j^{(t,b)}\}_{b=1}^B$, respectively.
 - 17: **end for**
-

Combining (19), (20), (21), and (22), we obtain the following theorem with respect to the identification of individual mediation effect in finite-horizon settings. We relegate a more detailed derivation to the Supplementary Appendix.

THEOREM 1 (Individual mediation effect for finite-horizon). *Suppose $\{A_t\}_{t \geq 1}$ are randomly assigned treatments and satisfy (2), and $\{M_t, R_t\}_{t \geq 1}$ follow (3) and (4). Then,*

$$(23) \quad \eta_j^{(t)} = \eta_j^{(t-1)} + \sum_{s=1}^t \sum_{i=s}^t \theta_{A_s \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}, \quad t = 1, \dots, T,$$

where $\theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}}$ is computed following (22), and we set $\eta_j^{(0)} = 0$.

As for estimation, based on Theorem 1, we estimate the individual mediation effect in a recursive manner. That is, for stage $t - 1$, we first estimate the weight matrix $\mathbf{W}_t$ in (3), and the parameters Θ_{1t}, Θ_{2t} in (4) for stage t . Estimation of $\mathbf{W}_t$ is needed for determining the parent set of each mediator, which in turn is used for backdoor covariate adjustment. There are multiple algorithms available to estimate $\mathbf{W}_t$, including those developed in single-time-point settings (Zheng et al., 2018; Yuan et al., 2019; Bello, Aragam and Ravikumar, 2022), and those developed in time-varying settings (Hyvärinen et al., 2010; Malinsky and Spirtes, 2018; Pamfil et al., 2020). We employ the DAGMA algorithm recently proposed by Bello, Aragam and Ravikumar (2022) for its simplicity and effectiveness. Recall that, in the finite-horizon setting, we allow the DAG to be nonstationary over time. We thus estimate the time-specific DAGs using DAGMA, with data pooled over multiple subjects but restricted to that time point only. We then estimate the within-stage quantities $\{\theta_{A_t \rightarrow R_t}, \theta_{M_{tj} \rightarrow R_t}, \theta_{A_t \rightarrow M_t}, \theta_{M_{tj} \rightarrow M_t}\}$ using (17), and estimate the cross-stage quantities, $\{\theta_{A_s \rightarrow R_t}, \theta_{M_{sj} \rightarrow R_t}, \theta_{A_s \rightarrow M_t}, \theta_{M_{sj} \rightarrow M_t}\}$, for $s = 1, \dots, t - 1$, using (18). Finally, we plug-in these estimators into (23) to estimate the individual mediation effect.

As for inference, based on Theorem 3 in Section 3.5, our plug-in estimator is asymptotically normal. This motivates us to employ the nonparametric bootstrap approach (Efron, 1979) to conduct statistical inference. Specifically, we first sample n trajectories from the observed data with replacement. We next refit the weight matrix $\mathbf{W}_t$ and re-estimate the within-stage and cross-stage quantities using the bootstrap samples. During the refitting, there are two options. One is to use the bootstrap samples to refit the coefficients of the nonzero entries in the weight matrix only, without re-estimating the structure of $\mathbf{W}_t$. The other is to refit both the structure of $\mathbf{W}_t$ and its nonzero entries. When the structure learning algorithm is consistent, the two solutions are asymptotically equivalent. For finite samples, the former is computationally more efficient, while the latter is expected to be more accurate. In our implementation, we adopt the first option due to its simplicity, and our numerical experiments suggest that it works well empirically. We next construct the plug-in estimators using these refitted parameters. Finally, we use the empirical lower and upper $\alpha/2$ quantiles of the plug-in estimators to construct the confidence interval under a given significance level α .

Algorithm 1 summarizes our estimation and inference procedure.

3.4. Individual mediation effect in an infinite-horizon setting. We next derive the expression of the individual mediation effect under the infinite-horizon setting. Unlike the finite-horizon setting, we now require the Markov mediation process $\{A_t, \mathbf{M}_t, R_t\}_{t \geq 1}$ to be stationary. This is to ensure the existence of the limit in Definition 3. Correspondingly, the parameters $\mathbf{W}, \Theta_1, \Theta_2$ in (3) and (4) remain the same across different time stages, which leads to the following relations. For any $1 \leq s \leq t \leq T$, and $j = 1, \dots, d$,

$$\begin{aligned} \theta_{A_s \rightarrow R_t} &= \theta_{A_{s+1} \rightarrow R_{t+1}} = \dots = \theta_{A_{s+T-t} \rightarrow R_T}, \quad \theta_{M_{sj} \rightarrow \mathbf{M}_t} = \theta_{M_{(s+1)j} \rightarrow \mathbf{M}_{t+1}} = \dots = \theta_{M_{(s+T-t)j} \rightarrow \mathbf{M}_T}, \\ \theta_{A_s \rightarrow \mathbf{M}_t} &= \theta_{A_{s+1} \rightarrow \mathbf{M}_{t+1}} = \dots = \theta_{A_{s+T-t} \rightarrow \mathbf{M}_T}, \quad \theta_{M_{sj} \rightarrow R_t} = \theta_{M_{(s+1)j} \rightarrow R_{t+1}} = \dots = \theta_{M_{(s+T-t)j} \rightarrow R_T}. \end{aligned}$$

Based on this observation, we obtain a simplified representation for $\eta_j^{(T)}$ as,

$$\eta_j^{(T)} = \theta_{M_{1j} \rightarrow R_1} \sum_{t=1}^T (T-t+1) \theta_{A_1 \rightarrow M_{tj}} + \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{\mathbf{M}_{1j}, \dots, \mathbf{M}_{(t+1)j}} \left(\sum_{s=1}^{T-t} (T-t-s+1) \theta_{A_1 \rightarrow M_{sj}} \right).$$

By dividing the right-hand-side by T and taking the limit $T \rightarrow \infty$, we obtain the following theorem for identifying the individual mediation effect in infinite-horizon settings. We relegate a more detailed derivation to the Supplementary Appendix.

THEOREM 2 (Individual mediation effect for infinite-horizon). *Suppose $\{A_t\}_{t \geq 1}$ are randomly assigned treatments and satisfy (2), $\{\mathbf{M}_t, R_t\}_{t \geq 1}$ follow (3) and (4), the Markov process $\{A_t, \mathbf{M}_t, R_t\}_{t \geq 1}$ is stationary, and the parameters $\{\mathbf{W}, \Theta_1, \Theta_2\}$ in (3) and (4) are time-invariant. Then,*

$$(24) \quad \eta_j^{(\infty)} \equiv \lim_{T \rightarrow \infty} \frac{\eta_j^{(T)}}{T} = \theta_{M_{1j} \rightarrow R_1} B_{1j} + (1 + B_{2j})^{-1} (B_{5j} - \theta_{M_{1j} \rightarrow R_1} B_{4j}),$$

where B_{1j}, B_{2j}, B_{4j} are the j th element of $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_4$, respectively,

$$\begin{aligned} \mathbf{B}_1 &= \{(1 - \zeta_2)(\mathbf{I} - \Gamma_1) - \zeta_1(\boldsymbol{\kappa} + \gamma_2)^\top\}^{-1} \{(1 - \zeta_2)\boldsymbol{\delta}_1 + \delta_2 \zeta_1\} \in \mathbb{R}^d, \\ (25) \quad \begin{pmatrix} \mathbf{B}_2 \\ \mathbf{B}_{3j} \end{pmatrix} &= (\mathbf{I} - \mathbf{B}_6 - \mathbf{B}_7)^{-1} \mathbf{B}_7 \begin{pmatrix} \theta_{M_{1j} \rightarrow \mathbf{M}_1} \\ \theta_{M_{1j} \rightarrow R_1} \end{pmatrix} \in \mathbb{R}^{d+1}, \\ \begin{pmatrix} \mathbf{B}_4 \\ \mathbf{B}_{5j} \end{pmatrix} &= (\mathbf{I} - \mathbf{B}_6 - \mathbf{B}_7)^{-1} \mathbf{B}_7 \begin{pmatrix} \theta_{M_{1j} \rightarrow \mathbf{M}_1} \\ \theta_{M_{1j} \rightarrow R_1} \end{pmatrix} B_{1j} \in \mathbb{R}^{d+1}, \\ \mathbf{B}_6 &= \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_d \\ \boldsymbol{\kappa}^\top & 0 \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}, \text{ and } \mathbf{B}_7 = \begin{pmatrix} \Gamma_1 & \zeta_1 \\ \gamma_2^\top & \zeta_2 \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}. \end{aligned}$$

Algorithm 2 Estimation for the infinite-horizon setting.

Input: Observed data $\{(A_{it}, M_{it}, R_{it}) : i = 1, \dots, n, t = 1, \dots, T\}$, the number of bootstrap samples B , and the significance level α .

Output: Estimated individual mediation effect $\hat{\eta}_j^{(\infty)}$, for $j = 1, \dots, d$.

- 1: Pool data across all t stages, $t = 1, \dots, T$.
- 2: Estimate parameters $\bar{\mathbf{W}}, \bar{\Theta}_1, \bar{\Theta}_2$ using the pooled data.
- 3: Estimate $\theta_{M_{1j} \rightarrow M_1}$ and $\theta_{M_{1j} \rightarrow R_1}$ using the pooled data.
- 4: Compute B_1 to B_7 using (25).
- 5: Compute $\hat{\eta}_j^{(\infty)}$ using (24).
- 6: **for** $b = 1, \dots, B$ **do**
- 7: Sample n trajectories from the observed data with replacement.
- 8: Repeat Lines 1 to 5 to compute $\hat{\eta}_j^{(\infty, b)}$ using the bootstrap samples.
- 9: **end for**
- 10: Construct the CI $[\hat{\eta}_j^{(\infty, L)}, \hat{\eta}_j^{(\infty, U)}]$ for $\eta_j^{(\infty)}$, where $\hat{\eta}_j^{(\infty, L)}$ and $\hat{\eta}_j^{(\infty, U)}$ are the empirical lower and upper $\alpha/2$ quantiles of $\{\hat{\eta}_j^{(\infty, b)}\}_{b=1}^B$, respectively.

As for estimation, based on Theorem 2, we pool the data across all T stages to estimate the model parameters $\mathbf{W}, \Theta_1, \Theta_2$. Recall that, in the infinite-horizon setting, we require the DAG to be stationary over time. We thus apply DAGMA to the data pooled across all time points to learn this time-invariant DAG. This is different from the finite-horizon setting where the model parameters can differ from one stage to another, while in the infinite-horizon setting, they remain the same. We then estimate the quantities B_1 to B_7 in (25) by plugging in the estimates of the corresponding model parameters, and we estimate the individual mediation effect using (24). As for inference, we adopt a similar nonparametric bootstrap procedure as in the finite-horizon setting. Algorithm 2 summarizes our estimation and inference procedure.

3.5. Asymptotic theory. We establish the asymptotic properties of our estimator of the individual mediation effect. We first present a set of regularity conditions.

ASSUMPTION 1 (Invertibility). *Let $\mathbf{X}_t \equiv (1, A_t, \mathbf{M}_{t-1}^\top, R_{t-1}, \mathbf{M}_t^\top)^\top \in \mathbb{R}^{2d+3}$. There exists some constant $C > 0$ such that $\lambda_{\min}(\mathbb{E} \mathbf{X}_t \mathbf{X}_t^\top) \geq C$ for any t , where $\lambda_{\min}(\bullet)$ denotes the minimum eigenvalue of a given matrix.*

ASSUMPTION 2 (Error residuals). *(i) The error terms $\epsilon_{w1t}, \epsilon_{w2t}, \dots, \epsilon_{wdt}$ in (3) are jointly normally distributed and independent, for $t = 1, \dots, T$. In addition, their variances are constant, in that $\text{Var}(\epsilon_{w1t}) = \text{Var}(\epsilon_{w2t}) = \dots = \text{Var}(\epsilon_{wdt})$. (ii) The error terms ϵ_{rt} in (4) satisfy that $\mathbb{E}(\epsilon_{rt}^4) < \infty$, for $t = 1, \dots, T$.*

ASSUMPTION 3 (Structure learning consistency). *The estimated DAG is a consistent estimator of the true underlying DAG; i.e., (i) for the finite-horizon setting, for $t = 1, \dots, T$, $\mathbb{P}(\hat{\mathcal{G}}_t \neq \mathcal{G}_t) \rightarrow 0$ as $n \rightarrow \infty$, where $\mathcal{G}_t$ is the true DAG in stage t ; (ii) for the infinite-horizon setting, $\mathbb{P}(\hat{\mathcal{G}} \neq \mathcal{G}) \rightarrow 0$ as $T \rightarrow \infty$, where $\mathcal{G}$ is the true time-invariant DAG.*

ASSUMPTION 4 (Stationarity). *For the infinite-horizon setting, the process $\{A_t, \mathbf{M}_t, R_t\}_{t \geq 1}$ is a strictly stationary β -mixing process.*

We give some remarks about these conditions. Assumption 1 guarantees the identifiability of model parameters in (4). Assumption 2(i) is commonly imposed for identifying DAG structures; see, e.g., Peters and Bühlmann (2014); Nandy, Maathuis and Richardson (2017);

Yuan et al. (2019); Shi and Li (2022). Assumption 2(ii) requires the fourth moments of the errors to be finite, which is mild. Assumption 3 holds for numerous DAG estimation algorithms, including the one by Bello, Aragam and Ravikumar (2022) that we use in our implementation. Assumption 4 is required for the infinite-horizon setting only, and the β -mixing condition is again commonly imposed; see, e.g., Bradley (2005). Overall, these regularity conditions are reasonable.

We next establish the asymptotic normality of our individual mediation effect estimator for both the finite-horizon and infinite-horizon settings.

THEOREM 3 (Asymptotic distribution for finite-horizon). *Suppose Assumptions 2 and 3(i) hold. Then, for $t = 1, \dots, T$,*

$$\sqrt{n}(\hat{\eta}_j^{(t)} - \eta_j^{(t)}) \xrightarrow{d} \mathcal{N}(0, \sigma_{tj}^2), \text{ as } n \rightarrow \infty,$$

where σ_{tj}^2 denotes the asymptotic variance of $\hat{\eta}_j^{(t)}$. Its form is given in the Supplementary Material, and it can be consistently estimated via the bootstrap method.

THEOREM 4 (Asymptotic distribution for infinite-horizon). *Suppose Assumptions 2, 3(ii) and 4 hold. Then,*

$$\sqrt{nT}(\hat{\eta}_j^{(\infty)} - \eta_j^{(\infty)}) \xrightarrow{d} \mathcal{N}(0, \sigma_j^2), \text{ as } nT \rightarrow \infty,$$

where σ_j^2 denotes the asymptotic variance of $\hat{\eta}_j^{(\infty)}$. Its definition is given in the Supplementary Material, and it can be consistently estimated using the bootstrap method.

We briefly remark that Theorem 3 requires the number of realizations n of $(A_t, \mathbf{M}_t, R_t)$ to diverge to infinity for every $t = 1, \dots, T$, with T being finite. Meanwhile, Theorem 4 requires either n or the number of time points T or both to diverge to infinity.

4. Simulations. In this section, we investigate the empirical performance of our proposed method through intensive simulations. We also compare with some baseline methods.

4.1. Simulation setup. We generate n copies of random samples $\{A_t, \mathbf{M}_t, R_t\}_{t=1}^T$ following models (3) and (4), i.e.,

$$\begin{aligned} \mathbf{M}_t &= (\mathbf{I} - \mathbf{W}^\top)^{-1} \boldsymbol{\epsilon}_{wt} + \boldsymbol{\alpha}_{1t} + \boldsymbol{\delta}_{1t} A_t + \boldsymbol{\Gamma}_{1t} \mathbf{M}_{t-1} + \boldsymbol{\zeta}_{1t} R_{t-1}, \\ R_t &= \alpha_{2t} + \delta_{2t} A_t + \boldsymbol{\gamma}_{2t}^\top \mathbf{M}_{t-1} + \zeta_{2t} R_{t-1} + \boldsymbol{\kappa}_t^\top \mathbf{M}_t + \epsilon_{rt}. \end{aligned}$$

We generate the sequence of treatments $A_t \stackrel{iid}{\sim} \text{Bernoulli}(0.5)$, $t = 1, \dots, T$, and the error terms $\boldsymbol{\epsilon}_{wt}$ and ϵ_{rt} from a standard normal distribution. Let $\boldsymbol{\Theta}_{1t} = (\boldsymbol{\alpha}_{1t}, \boldsymbol{\delta}_{1t}, \boldsymbol{\Gamma}_{1t}, \boldsymbol{\zeta}_{1t})^\top \in \mathbb{R}^{(d+3) \times d}$ and $\boldsymbol{\Theta}_{2t} = (\alpha_{2t}, \delta_{2t}, \zeta_{2t}, \boldsymbol{\gamma}_{2t}^\top, \boldsymbol{\kappa}_t^\top)^\top \in \mathbb{R}^{(2d+3)}$ collect the model parameters, and we generate the entries of $\boldsymbol{\Theta}_{1t}, \boldsymbol{\Theta}_{2t}$ from a uniform distribution on $(-0.5, 0.5)$. For the finite-horizon setting, we generate different $\boldsymbol{\Theta}_{1t}, \boldsymbol{\Theta}_{2t}$ for different time points t , whereas for the infinite-horizon setting, we only generate one copy of $\boldsymbol{\Theta}_{1t}, \boldsymbol{\Theta}_{2t}$, and keep them fixed across all time points. We fix $d = 3$, and generate the matrix $\mathbf{W} \in \mathbb{R}^{3 \times 3}$ in two steps. We first begin with a zero matrix, then replace every entry $(\mathbf{W})_{ij}, i < j$, by the product of two random variables $E_{ij}^{(1)} \times E_{ij}^{(2)}$, where $E_{ij}^{(1)}$ is a Bernoulli variable with probability 0.9, indicating a random directed edge is added from mediator M_i to M_j , and $E_{ij}^{(2)}$ is the edge weight, which is randomly drawn from a uniform distribution on $[-0.9, -0.5] \cup [0.5, 0.9]$. Following this

generation process, we obtain $\mathbf{W} = \begin{pmatrix} 0 & -0.80 & 0.61 \\ 0 & 0 & -0.82 \\ 0 & 0 & 0 \end{pmatrix}$. We set the initial values of mediators and outcome all equal to 0.

We recognize that there is no existing solution in the literature for this problem. Instead, we compare with two baseline solutions. One method is termed “independent time points”, which ignores all temporal dependence across different time points. That is, it estimates the individual mediation effect at every single time stage with stage-specific data, without taking into account the dependence to prior stages nor the carryover effects. The other method is termed “independent mediators”, which ignores all dependence among the multivariate mediators, and essentially treats $\mathbf{W}$ as a zero matrix. We evaluate all estimation methods using three criteria: the estimation bias, the empirical standard error (SE), and the coverage probability (CP). For inference, we set the number of bootstrapped samples to be 500 and the significance level α to be 0.05. We employ DAGMA with linear models for structural learning, and the corresponding hyperparameters, `w_threshold` and `lambda1`, are set to 0.1 and 0.0, respectively.

4.2. Finite-horizon setting. For the finite-horizon setting, we consider the number of time points $T \in \{10, 20, 30\}$ and the sample size $n \in \{100, 250, 500\}$. Note that the true individual mediation effect cannot be derived analytically in our setting, and is computed numerically based on 10 million Monte Carlo samples. Table 1 reports the results based on 500 replications. We observe that, when the number of subjects n increases, the estimation bias and standard error both decrease for our proposed method, which agrees with our theory. Moreover, our method achieves the desired coverage probability across all scenarios. The same is not true for the two baseline methods. For instance, for the independent mediators method, the estimation bias of $\eta_2^{(T)}$ and $\eta_3^{(T)}$ does not decrease. This is because M_{t1} is in the parent set of M_{t2} and M_{t3} for all $t = 1, \dots, T$, and ignoring the effects along the paths $M_{t1} \rightarrow M_{t2}$ and $M_{t1} \rightarrow M_{t3}$ lead to a larger bias and invalid coverage probability in estimating $\eta_2^{(T)}$ and $\eta_3^{(T)}$. On the contrary, the estimation bias of $\eta_1^{(T)}$ is much smaller, since the parent set of M_{t1} is empty given $(\mathbf{M}_{t-1}, R_t, A_t)$ in our simulation example.

4.3. Infinite-horizon setting. For the infinite-horizon setting, we consider the number of time points $T \in \{100, 250, 500\}$ and the sample size $n \in \{20, 50, 100\}$. Again, the true individual mediation effect cannot be derived analytically in our setting, and is computed numerically based on 10 million Monte Carlo samples with 3,000 time points. We also drop the first five steps as a warm up to mitigate the potential influence of the initial conditions. Table 2 reports the results based on 500 data replications. We again observe that the estimation bias and standard error both decrease for our method as n increases, and the desired coverage probability is achieved, across all scenarios. On the other hand, the bias increases sharply for the independent time points method, because ignoring the carryover effects becomes exaggerated with a large T . Moreover, the estimation bias of the independent mediators method is much larger than the proposed method for $\eta_2^{(\infty)}$ and $\eta_3^{(\infty)}$. Both baseline methods fail to achieve the desired coverage probability in most cases, except for the independent mediators method with $\eta_1^{(\infty)}$. This is due to the specific simulation setting with an empty parent set given the history. In general, these results demonstrate the importance of taking into account both the dependence among multiple time points and among multiple mediators.

5. Data Application. In this section, we revisit the motivating Intern Health Study (NeCamp et al., 2020). The IHS is a 26-week sequentially randomized trial with the objective of understanding the biological mechanisms of depression, and the ultimate goal of

TABLE 1

Simulations for the finite-horizon setting: the bias, the empirical standard error (SE) and the coverage probability (CP) between the estimated and true individual mediator effects with varying number of time points T and number of subjects n .

Method	n		100			250			500		
	T	Param	Bias	SE	CP	Bias	SE	CP	Bias	SE	CP
Proposed method	10	$\eta_1^{(T)}$	0.000	0.594	0.972	0.003	0.351	0.966	0.008	0.249	0.960
		$\eta_2^{(T)}$	-0.016	0.533	0.960	-0.036	0.283	0.970	-0.006	0.209	0.952
		$\eta_3^{(T)}$	0.007	0.696	0.964	0.023	0.420	0.962	0.011	0.312	0.950
	20	$\eta_1^{(T)}$	-0.046	1.960	0.950	0.008	1.215	0.946	-0.023	0.840	0.950
		$\eta_2^{(T)}$	-0.037	1.134	0.970	-0.018	0.717	0.944	-0.021	0.490	0.948
		$\eta_3^{(T)}$	0.027	1.236	0.956	0.012	0.721	0.954	-0.003	0.491	0.972
	30	$\eta_1^{(T)}$	0.044	2.655	0.948	-0.078	1.691	0.938	0.055	1.159	0.940
		$\eta_2^{(T)}$	-0.067	1.610	0.962	-0.002	0.915	0.964	-0.020	0.651	0.954
		$\eta_3^{(T)}$	-0.067	1.838	0.958	-0.055	1.122	0.960	0.051	0.783	0.946
Independent time points	10	$\eta_1^{(T)}$	-0.577	0.373	0.712	-0.558	0.204	0.370	-0.569	0.131	0.132
		$\eta_2^{(T)}$	0.199	0.759	0.962	0.171	0.447	0.926	0.175	0.308	0.888
		$\eta_3^{(T)}$	-0.768	1.472	0.896	-0.762	0.861	0.866	-0.716	0.656	0.806
	20	$\eta_1^{(T)}$	0.823	1.301	0.882	0.844	0.665	0.724	0.852	0.441	0.580
		$\eta_2^{(T)}$	6.158	4.075	0.636	6.071	2.479	0.286	6.075	1.797	0.066
		$\eta_3^{(T)}$	0.694	3.095	0.952	0.665	1.804	0.936	0.683	1.315	0.914
	30	$\eta_1^{(T)}$	1.662	1.491	0.854	1.626	0.831	0.598	1.607	0.530	0.370
		$\eta_2^{(T)}$	-10.111	6.992	0.700	-9.898	4.801	0.410	-9.989	3.207	0.126
		$\eta_3^{(T)}$	-2.352	1.880	0.860	-2.446	1.081	0.374	-2.392	0.715	0.144
Independent mediators	10	$\eta_1^{(T)}$	0.000	0.594	0.972	0.003	0.351	0.954	0.008	0.249	0.964
		$\eta_2^{(T)}$	-0.139	0.640	0.944	-0.175	0.342	0.944	-0.142	0.263	0.904
		$\eta_3^{(T)}$	0.164	0.603	0.942	0.178	0.368	0.938	0.172	0.276	0.892
	20	$\eta_1^{(T)}$	-0.045	1.959	0.948	0.008	1.215	0.950	-0.023	0.840	0.952
		$\eta_2^{(T)}$	-0.159	1.370	0.956	-0.135	0.814	0.950	-0.157	0.576	0.930
		$\eta_3^{(T)}$	-0.764	1.390	0.926	-0.728	0.830	0.882	-0.779	0.602	0.772
	30	$\eta_1^{(T)}$	0.044	2.656	0.946	-0.078	1.691	0.932	0.055	1.159	0.948
		$\eta_2^{(T)}$	-1.286	2.087	0.922	-1.154	1.114	0.872	-1.227	0.864	0.692
		$\eta_3^{(T)}$	-0.741	2.128	0.934	-0.724	1.279	0.892	-0.632	0.920	0.884

improving the mental health outcomes of medical interns in the United States. The study developed and deployed a mobile app to deliver prompt notifications, such as reminders to have a break, take a walk, or prioritize sleep, which aims to improve the well-being of interns who often work under stressful environments. In each week, each intern was randomized into receiving the notifications or not. Meanwhile, wearable devices (Fitbit) recorded daily measurements of step count (Step), sleep duration (Sleep, minutes), resting heart rate (RHR, beats per minute), and heart rate variability (HRV, milliseconds). Interns also self-reported a daily mood score in the app (Shaffer and Ginsberg, 2017). We formulate the problem in the framework of multivariate dynamic mediation analysis, with the binary status of receiving the notification or not as the treatment, the transformed measurements, i.e., the cubic-root of step count, the square-root of sleep duration, RHR and HRV as $d = 4$ potential mediators, and the mood score as the outcome. We average all the measurements within each week, resulting in $T = 26$ weeks of data, for $n = 1196$ interns undergoing the sequential randomization.

TABLE 2

Simulations for the infinite-horizon setting: the bias, the empirical standard error (SE) and the coverage probability (CP) between the estimated and true individual mediator effects with varying number of time points T and number of subjects n .

Method	n		20			50			100		
	T	Param	Bias	SE	CP	Bias	SE	CP	Bias	SE	CP
Proposed method	100	$\eta_1^{(\infty)}$	.002	.070	.918	.002	.046	.928	.001	.029	.934
		$\eta_2^{(\infty)}$	.003	.045	.924	.001	.029	.932	.001	.020	.940
		$\eta_3^{(\infty)}$	.003	.036	.924	.002	.021	.946	.000	.015	.946
	250	$\eta_1^{(\infty)}$	-.001	.042	.932	-.002	.026	.946	-.001	.019	.930
		$\eta_2^{(\infty)}$	-.003	.029	.914	-.001	.019	.938	.001	.012	.958
		$\eta_3^{(\infty)}$	.001	.020	.938	.000	.014	.938	-.001	.010	.950
	500	$\eta_1^{(\infty)}$	.000	.029	.944	.000	.018	.956	-.001	.014	.932
		$\eta_2^{(\infty)}$	.000	.020	.928	-.001	.012	.954	.000	.009	.942
		$\eta_3^{(\infty)}$	.000	.016	.924	-.001	.010	.932	.000	.007	.928
Independent time points	100	$\eta_1^{(\infty)}$	.371	.047	.000	.370	.030	.000	.371	.021	.000
		$\eta_2^{(\infty)}$	-.065	.076	.818	-.070	.049	.658	-.071	.034	.404
		$\eta_3^{(\infty)}$	.132	.086	.628	.133	.060	.314	.132	.040	.076
	250	$\eta_1^{(\infty)}$	.372	.027	.000	.371	.018	.000	.371	.013	.000
		$\eta_2^{(\infty)}$	-.075	.048	.616	-.073	.030	.300	-.073	.020	.052
		$\eta_3^{(\infty)}$	.134	.055	.286	.135	.034	.028	.133	.024	.000
	500	$\eta_1^{(\infty)}$	.370	.021	.000	.371	.013	.000	.371	.010	.000
		$\eta_2^{(\infty)}$	-.073	.035	.368	-.074	.021	.058	-.073	.015	.002
		$\eta_3^{(\infty)}$	.131	.039	.074	.131	.024	.000	.134	.018	.000
Independent mediators	100	$\eta_1^{(\infty)}$	.002	.070	.918	.002	.046	.934	.001	.029	.936
		$\eta_2^{(\infty)}$	.322	.031	.000	.324	.019	.000	.325	.015	.000
		$\eta_3^{(\infty)}$	-.264	.038	.000	-.266	.024	.000	-.265	.016	.000
	250	$\eta_1^{(\infty)}$	-.001	.042	.934	-.002	.026	.934	-.001	.019	.942
		$\eta_2^{(\infty)}$	.324	.019	.000	.324	.012	.000	.325	.009	.000
		$\eta_3^{(\infty)}$	-.265	.024	.000	-.265	.014	.000	-.265	.010	.000
	500	$\eta_1^{(\infty)}$	.000	.029	.936	.000	.018	.952	-.001	.014	.936
		$\eta_2^{(\infty)}$	.324	.014	.000	.325	.009	.000	.324	.006	.000
		$\eta_3^{(\infty)}$	-.265	.017	.000	-.265	.011	.000	-.265	.007	.000

We apply our method to this data adopting the finite-horizon setting. This is partly because we have a relatively limited number of time points $T = 26$ in this study, and partly because the data is likely to be non-stationary over time (NeCamp et al., 2020; Li et al., 2022; Wang, Shi and Wu, 2023). Our analysis, as shown in Figure 4, also confirms that the individual mediation effect varies over time. Figure 3(a) shows the estimated DAG structure among the four mediators, while Figures 3(b) and (c) show the estimated individual mediation effects over time, along with the associated 95% confidence intervals, that are smoothed with a natural cubic spline and further decomposed as the immediate and delayed effects.

We make a number of observations. First of all, we see from Figure 3(a) that, the four mediators have complex relationships between each other, indicating the importance of accounting for the dependence structure among the mediators. Second, we see from Figure 3(b) and (c) that, the magnitude of the delayed effects is generally greater than that of the immediate effects, indicating the importance of accounting for the carryover effects in our




Fig 3: Analysis of IHS mobile health data: (a) the estimated DAG among the four mediators, dashed/solid arrows represent negative/positive effects; (b) the estimated immediate individual mediation effects over time; (c) the estimated delayed individual mediation effects over time. The shaded area indicates the 95% confidence interval.




Fig 4: Analysis of IHS mobile health data: the estimated effects from the treatment to the four individual mediators, decomposed as the immediate (first row) and delayed (second row) effects, respectively.

understanding of the effects of push notification to intern's mood. Third, most individual mediation effects are insignificant. This could potentially be attributed to the relatively weak treatment effect, as also noted by [NeCamp et al. \(2020\)](#). Fourth, there is a significantly negative carryover effect from the push notification to mood score mediated by the RHR around week 10. Meanwhile, this carryover effect is predominantly negative across most weeks. To illustrate this effect, Figure 4 further plots the smoothed estimated effects from the treatment to the four individual mediators, decomposed as the immediate and delayed effects, respectively. We see that the push notification leads to an increased RHR after week five, suggesting that the push notification can lead to a lower mood score by increasing the RHR. This agrees with our prior knowledge that an increased RHR usually indicates a more tired and stressful state, thus a worse mood. Finally, the carryover effect from the push notification to sleep duration is mostly negative, and is nearly significant between weeks 15 and 20. As seen in Figure 4, the push notification generally results in reduced sleep hours, aligning with our knowledge that fewer sleep hours lead to poorer mood too.

6. Conclusion and Discussion. This paper proposes a reinforcement learning framework for multivariate dynamic mediation analysis where there are multivariate and conditionally dependent mediators observed over multiple time points. While most existing studies focus on individual mediation effects among multivariate mediators or the effects of a single set of mediators over time, they typically do not consider both. A recent study by [Wei et al. \(2024b\)](#) explores individual mediation effects across multiple time points. However, it does not account for delayed mediation effects, nor does it analyze the indirect effects of each mediator from its upstream mediators.




Fig 5: An illustration of the difference between the individual mediation effect and the Granger causality.

Below, we discuss the relation between our individual mediation effect and the Granger causality, numerous other mediation effects that may be of potential interest, and some additional extensions.

6.1. Connection to Granger causality. The Granger causality is a concept that leverages the temporal ordering inherent to time series so to draw causal statements restricted to the “past” causing the “future” (Granger, 1969). A variable is said to Granger-cause the other, if including this variable’s past values provides valuable information for forecasting the other variable’s future behavior. The proposed models in (3) and (4) are mathematically similar to those under the Granger causality framework, since R_t may be predicted by its past value R_{t-1} , M_t and M_{t-1} , and M_t also depends on its past value M_{t-1} and R_{t-1} .

However, our notion of individual mediation effect differs from the concept of Granger causality, in that a mediator’s individual mediation effect and its ability to Granger-cause the outcome is *not* equivalent. On the one hand, a mediator can have a zero individual mediation effect, but still Granger-causes the outcome, if it is not influenced by the treatment. Figure 5, Case 1, gives an illustrative example. In this case, we have a single mediator that impacts the outcome at each time point, thereby Granger-causing the outcome. However, since the mediator is unaffected by the treatment, there is no direct causal pathway from the treatment through the mediator to the outcome, resulting in a zero individual mediation effect. On the other hand, a mediator with a nonzero individual mediation effect might not necessarily Granger-cause the outcome, if it does not improve the forecast accuracy of the outcome beyond other mediators. Figure 5, Case 2, gives another example. In this case, the first set of mediators M_{11} and M_{21} directly influence the outcome, thus Granger-causing the outcome. The second mediator M_{12} , however, affects the outcome indirectly through its influence on the first mediator M_{21} , resulting in a nonzero individual mediation effect. Despite this, M_{12} does not Granger-cause the outcome R_1 or R_2 , as its influence is mediated indirectly.

6.2. Sequential mediation effect. As discussed in Section 2.3, the individual mediation effect $\eta_j^{(T)}$ measures the cumulative effect mediated through the j th mediator M_j across all the upstream mediators and all T time points. Meanwhile, it may be of separate interest to analyze the portion of the individual mediation effect of M_j that is attributed to its upstream mediator M_k . We denote this effect as $\eta_{k \rightarrow j}$, and call it the *sequential* mediation effect, as it quantifies the effect that is sequentially transmitted from the treatment, through the k th mediator, to the j th mediator, then to the outcome.

We next formally define the sequential mediation effect. Toward that end, we require the DAG structure to remain stationary over time. That is, if the k th mediator is an ancestor

of the j th mediator at any time t , i.e., there exists a direct path from M_{tk} to M_{tj} , then the same relation holds at all time points. For the finite-horizon setting, we define the sequential mediation effect over T time points as

$$\eta_{k \rightarrow j}^{(T)} = \sum_{t=1}^T \left\{ \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, \forall s \leq t)] - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, M_{sj} = m, \forall s \leq t)] \right. \\ \left. - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, M_{sk} = m, \forall s \leq t)] + \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_s = a, M_{sj} = M_{kj} = m, \forall s \leq t)] \right\},$$

if the k th mediator is an ancestor of the j th mediator, and set $\eta_{k \rightarrow j}^{(T)} = 0$ otherwise. In this definition, the first term corresponds to the total treatment effect, the second and third terms measure the effects that do not go through the j th or the k th mediator, respectively, and the last term quantifies the effect that does not simultaneously pass both the j th and k th mediators. Then, by the principle of inclusion-exclusion, $\eta_{k \rightarrow j}^{(T)}$ measures the desired sequential mediation effect. For the infinite-horizon setting, we define $\eta_{k \rightarrow j}^{(\infty)} = \lim_{T \rightarrow \infty} \eta_{k \rightarrow j}^{(T)}/T$, provided the limit exists.

To further illustrate this definition, we revisit the examples in Figure 2. For the single-stage example in Figure 2, left panel, the mediator M_{12} is a child of M_{11} and is affected by M_{11} . Its individual mediation effect that passes through M_{11} is

$$\begin{aligned} \eta_{1 \rightarrow 2}^{(1)} &= \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid do(A_1 = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid do(A_1 = a, M_{12} = m)] \\ &\quad - \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid do(A_1 = a, M_{11} = m)] + \frac{\partial}{\partial a} \mathbb{E}[R_1 \mid do(A_1 = a, M_{11} = M_{12} = m)] \\ &= \theta_{A_1 \rightarrow R_1} - \theta_{A_1 \rightarrow R_1}^{M_{12}} - \theta_{A_1 \rightarrow R_1}^{M_{11}} + \theta_{A_1 \rightarrow R_1}^{M_{11}, M_{12}} \\ &= \theta_{A_1 \rightarrow M_{11}} \theta_{M_{11} \rightarrow M_{12}} \theta_{M_{12} \rightarrow R_1}. \end{aligned}$$

It is clear to see that the term $\theta_{A_1 \rightarrow M_{11}} \theta_{M_{11} \rightarrow M_{12}} \theta_{M_{12} \rightarrow R_1}$ corresponds to the effect along the path $A_1 \rightarrow M_{11} \rightarrow M_{12} \rightarrow R_1$, which coincides with the results from the path analysis. Similarly, for the two-stage example in Figure 2, right panel, we have that,

$$\begin{aligned} \eta_{1 \rightarrow 2}^{(2)} &= \eta_{1 \rightarrow 2}^{(1)} + \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a, M_{11} = M_{21} = m)] \\ &\quad - \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a, M_{12} = M_{22} = m)] \\ &\quad + \frac{\partial}{\partial a} \mathbb{E}[R_2 \mid do(A_1 = A_2 = a, M_{11} = M_{21} = M_{12} = M_{22} = m)] \\ &= \eta_{1 \rightarrow 2}^{(1)} + \theta_{A_1, A_2 \rightarrow R_2} - \theta_{A_1, A_2 \rightarrow R_2}^{M_{11}, M_{21}} - \theta_{A_1, A_2 \rightarrow R_2}^{M_{12}, M_{22}} + \theta_{A_1, A_2 \rightarrow R_2}^{M_{11}, M_{12}, M_{21}, M_{22}} \\ &= \eta_{1 \rightarrow 2}^{(1)} + \theta_{A_1 \rightarrow R_2} - \theta_{A_1 \rightarrow R_2}^{M_{11}, M_{21}} - \theta_{A_1 \rightarrow R_2}^{M_{12}, M_{22}} + \theta_{A_1 \rightarrow R_2}^{M_{11}, M_{21}, M_{12}, M_{22}} \\ &\quad + \theta_{A_2 \rightarrow R_2} - \theta_{A_2 \rightarrow R_2}^{M_{21}} - \theta_{A_2 \rightarrow R_2}^{M_{22}} + \theta_{A_2 \rightarrow R_2}^{M_{21}, M_{22}}, \end{aligned}$$

where the terms in the last two lines of the above equation can be calculated recursively with the formulas in equations (21) and (22), allowing for the estimation and inference of the sequential mediation effect.

Finally, we extend beyond the above illustrative examples to outline the estimation and inference procedure for the sequential mediation effect under a general setting involving d mediators across T time points. As for estimation, we begin by applying an existing structural

learning algorithm to the observed data to learn the DAG structure. From the estimated DAG, if the k th mediator is not an ancestor of the j th mediator, we set the estimated sequential effect to zero. Otherwise, we obtain the estimated $\hat{\eta}_{k \rightarrow j}$ through the quantities $\theta_{A_1, \dots, A_t \rightarrow R_t}$, $\theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}}$, and $\theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1k}, M_{1j}, \dots, M_{kj}, M_{tj}}$, which in turn can be recursively updated following the method developed in Section 3. As for inference, we again apply the bootstrap approach in Section 3 to quantify the uncertainty of this plug-in estimator $\hat{\eta}_{k \rightarrow j}$.

6.3. Direct, indirect and conditional effects. We discuss the direct effect DE and indirect effect IE in Section 2.2. We now briefly discuss their estimation and inference, which may be of separate interest. Recall DE in (6) under linear SEMs. It satisfies a recursive equation,

$$\text{DE}^{(t)} = \text{DE}^{(t-1)} \zeta_{2t} + \delta_{2t},$$

for $t = 2, \dots, T$, with $\text{DE}^{(1)} = \delta_{21}$. Here we add the superscript (t) to DE to indicate the direct effect across t time points. This recursive formulation allows us to employ a similar plug-in method as in Section 3 to estimate DE, as well as a similar bootstrap approach for inference. To estimate IE, we first compute the total effect as a byproduct of the proposed procedure in Section 3. We then estimate IE as the difference between the estimated total effect and DE. Again, we can employ the bootstrap approach to infer IE.

Our definition of the individual mediation effect is concerned with the *marginal* effect,

$$\frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a)] - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a, M_{tj} = m)],$$

which marginalizes over the history $\{A_s, \mathbf{M}_s, R_s\}$ for all $s < t$. This marginalization is appropriate because our interventions on the treatment and mediator at time t do not causally influence those historical variables. Henceforth, the distribution of those historical variables remains unchanged before and after the intervention. Meanwhile, it may be of separate interest to study the *conditional* effect,

$$\begin{aligned} & \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a), \{A_s, \mathbf{M}_s, R_s\}, \forall s < t] \\ & - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a, M_{tj} = m), \{A_s, \mathbf{M}_s, R_s\}, \forall s < t], \end{aligned}$$

given the history. Nevertheless, we note that, under the linear structural model, the conditional effect is equivalent to the marginal effect that we target in this article. We next sketch a few lines to show the equivalence. Specifically, the first term is $\frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a)] = \theta_{A_t \rightarrow R_t} = \beta_{A_t, R_t \mid \{A_s, \mathbf{M}_s, R_s\}, \forall s < t}$ in our current definition. It is the partial coefficient of A_t by linearly regressing R_t on A_t given the historical variables that could affect the distribution of A_t . Under the linear model structures (3) and (4), the effects of A_t and the historical variables on R_t are additive. Consequently, the resulting conditional effect $\frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a), \{A_s, \mathbf{M}_s, R_s\}, \forall s < t] = \beta_{A_t, R_t \mid \{A_s, \mathbf{M}_s, R_s\}, \forall s < t}$ remains the same as the marginal effect, regardless of the values of the historical variables. The second term is $\frac{\partial}{\partial a} \mathbb{E}[R_t \mid \text{do}(A_t = a, M_{tj} = m)] = \theta_{A_t \rightarrow R_t}^{M_{tj}} = \theta_{A_t \rightarrow R_t} - \theta_{A_t \rightarrow M_{tj}} \theta_{M_{tj} \rightarrow R_t}$. Again, the value $\theta_{A_t \rightarrow R_t}$ remains the same whether or not conditioning on the historical variables. Similarly, under (3) and (4), the effects of M_{tj} and the historical variables on R_t are additive. Consequently, the value of the product $\theta_{A_t \rightarrow M_{tj}} \theta_{M_{tj} \rightarrow R_t}$ remains the same, regardless of whether the history is included in the conditioning set or not.

6.4. *Extension to varying T and time lags.* Finally, we remark that our method can be adapted to accommodate the setting when the number of time points T varies among subjects, and when the time lags between two time points differ.

For the finite-horizon setting, T can be different for different subjects, because the model parameters and intermediate quantities at any given time t are estimated by pooling the data across subjects who have observations available at that time point. Besides, Theorem 3 remains valid when T varies across subjects, provided that the sample size at each time point is sufficiently large. Meanwhile, for the time lag between two time points, we require it to follow the same distribution across all subjects, but it does not need to follow the same distribution across all time points, since we allow the data generating process to be non-stationary over time.

For the infinite-horizon setting, T again can be different, because the data are pooled across both subjects and time points. Similarly, Theorem 4 remains valid when T varies across subjects, provided that the total sample size is sufficiently large. Meanwhile, the time lag between two time points needs to follow the same distribution, due to the requirement for stationarity for the infinite-horizon setting. However, this requirement does not extend to between subjects, and the distributions can vary among different subjects. In that case, we can use each subject’s own data to estimate their subject-specific mediation effects.

Acknowledgments. The authors are grateful for the contributions of the researchers, administrators, and participants involved in the Intern Health Study (<https://clinicaltrials.gov/study/NCT03972293>). The authors also thank the editor, associate editor and reviewers for their comments, which have led to substantial improvements of the manuscript.

Funding. Luo’s research was partly supported by NIH grant R21AG083364. Shi’s research was supported in part by the EPSRC grant EP/W014971/1. Wu’s work was partly supported by NIH grants R01 MH101459 and R01 NR013658. Li’s research was partially supported by NSF grant CIF-2102227, NIH grants R01AG062542 and R01AG080043.

SUPPLEMENTARY MATERIAL

The supplement provides the proofs of the main theorems in the paper.

REFERENCES

- BARON, R. M. and KENNY, D. A. (1986). The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology* **51** 1173–1182.
- BELLO, K., ARAGAM, B. and RAVIKUMAR, P. (2022). DAGMA: Learning DAGs via M-matrices and a Log-Determinant Acyclicity Characterization. *arXiv preprint arXiv:2209.08037*.
- BI, X., YANG, L., LI, T., WANG, B., ZHU, H. and ZHANG, H. (2017). Genome-wide mediation analysis of psychiatric and cognitive traits through imaging phenotypes. *Human brain mapping* **38** 4088–4097.
- BOCA, S. M., SINHA, R., CROSS, A. J., MOORE, S. C. and SAMPSON, J. N. (2014). Testing multiple biological mediators simultaneously. *Bioinformatics* **30** 214–220.
- BRADLEY, R. C. (2005). Basic properties of strong mixing conditions. A survey and some open questions. *Probability Survey* **2** 107–144. Update of, and a supplement to, the 1986 original.
- CAI, H., SONG, R. and LU, W. (2021). ANOCE: Analysis of Causal Effects with Multiple Mediators via Constrained Structural Learning. In *International Conference on Learning Representations*.
- CAI, X., COFFMAN, D. L., PIPER, M. E. and LI, R. (2022). Estimation and inference for the mediation effect in a time-varying mediation model. *BMC Medical Research Methodology* **22** 113.
- CELLI, V. (2022). Causal mediation analysis in economics: Objectives, assumptions, models. *Journal of Economic Surveys* **36** 214–234.

- CHAKRABORTTY, A., NANDY, P. and LI, H. (2018). Inference for individual mediation effects and interventional effects in sparse high-dimensional causal graphical models. *arXiv preprint arXiv:1809.10652*.
- DÍAZ, I., WILLIAMS, N. and RUDOLPH, K. E. (2022). Efficient and flexible causal mediation with time-varying mediators, treatments, and confounders. *arXiv preprint arXiv:2203.15085*.
- DJORDJLOVIĆ, V., HEMERIK, J. and THORESEN, M. (2022). On optimal two-stage testing of multiple mediators. *Biometrical Journal* **64** 1090–1108.
- EFRON, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* **7** 1–26.
- GE, L., WANG, J., SHI, C., WU, Z. and SONG, R. (2023). A Reinforcement Learning Framework for Dynamic Mediation Analysis. *arXiv preprint arXiv:2301.13348*.
- GRANGER, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society* 424–438.
- GUO, X., LI, R., LIU, J. and ZENG, M. (2023). Statistical inference for linear mediation models with high-dimensional mediators and application to studying stock reaction to COVID-19 pandemic. *Journal of Econometrics* **235** 166–179.
- HAO, B., JI, X., DUAN, Y., LU, H., SZEPESVARI, C. and WANG, M. (2021). Bootstrapping fitted q-evaluation for off-policy inference. In *International Conference on Machine Learning* 4074–4084. PMLR.
- HEJAZI, N. S., RUDOLPH, K. E., VAN DER LAAN, M. J. and DÍAZ, I. (2020). Nonparametric causal mediation analysis for stochastic interventional (in)direct effects. *arXiv preprint arXiv:2009.06203*.
- HU, Y. and WAGER, S. (2022). Switchback experiments under geometric mixing. *arXiv preprint arXiv:2209.00197*.
- HUANG, Y.-T. and PAN, W.-C. (2016). Hypothesis test of mediation effect in causal mediation model with high-dimensional continuous mediators. *Biometrics* **72** 402–413.
- HUANG, J. and YUAN, Y. (2017). Bayesian dynamic mediation analysis. *Psychol. Methods* **22** 667–686.
- HYVÄRINEN, A., ZHANG, K., SHIMIZU, S. and HOYER, P. O. (2010). Estimation of a structural vector autoregression model using non-gaussianity. *Journal of Machine Learning Research* **11**.
- KALLUS, N. and UEHARA, M. (2022). Efficiently breaking the curse of horizon in off-policy evaluation with double reinforcement learning. *Operations Research* **70** 3282–3302.
- KAUFMAN, J. S. and KAUFMAN, S. (2001). Assessment of structured socioeconomic effects on health. *Epidemiology* 157–167.
- LI, M., SHI, C., WU, Z. and FRYZLEWICZ, P. (2022). Testing stationarity and change point detection in reinforcement learning. *arXiv preprint arXiv:2203.01707*.
- LIAO, P., KLASNJA, P. and MURPHY, S. (2021). Off-policy estimation of long-term average outcomes with applications to mobile health. *Journal of the American Statistical Association* **116** 382–391.
- LIAO, P., QI, Z., WAN, R., KLASNJA, P. and MURPHY, S. A. (2022). Batch policy learning in average reward markov decision processes. *Annals of statistics* **50** 3364.
- LIN, S.-H., YOUNG, J. G., LOGAN, R. and VANDERWEELE, T. J. (2017). Mediation analysis for a survival outcome with time-varying exposures, mediators, and confounders. *Statistics in medicine* **36** 4153–4166.
- LIU, W., TU, J., ZHANG, Y. and CHEN, X. (2023). Online Estimation and Inference for Robust Policy Evaluation in Reinforcement Learning. *arXiv preprint arXiv:2310.02581*.
- LUCKETT, D. J., LABER, E. B., KAHKOSKA, A. R., MAAHS, D. M., MAYER-DAVIS, E. and KOSOROK, M. R. (2020). Estimating Dynamic Treatment Regimes in Mobile Health Using V-Learning. *Journal of the American Statistical Association* **115** 692–706.
- MAATHUIS, M. H., KALISCH, M. and BÜHLMANN, P. (2009). Estimating high-dimensional intervention effects from observational data. *The Annals of Statistics* **37** 3133 – 3164.
- MACKINNON, D. P. (2008). *Introduction to statistical mediation analysis*. Taylor and Francis.
- MALINSKY, D. and SPIRITES, P. (2018). Causal structure learning from multivariate time series in settings with unmeasured confounding. In *Proceedings of 2018 ACM SIGKDD workshop on causal discovery* 23–47. PMLR.
- MNIH, V., KAVUKCUOĞLU, K., SILVER, D., RUSU, A. A., VENESS, J., BELLEMARE, M. G., GRAVES, A., RIEDMILLER, M., FIDJELAND, A. K., OSTROVSKI, G. et al. (2015). Human-level control through deep reinforcement learning. *nature* **518** 529–533.
- MURPHY, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **65** 331–355.
- NANDY, P., MAATHUIS, M. H. and RICHARDSON, T. S. (2017). Estimating the effect of joint interventions from observational data in sparse high-dimensional settings. *Annals of Statistics* **45** 647–674.
- NECAMP, T., SEN, S., FRANK, E., WALTON, M. A., IONIDES, E. L., FANG, Y., TEWARI, A. and WU, Z. (2020). Assessing Real-Time Moderation for Developing Adaptive Mobile Health Interventions for Medical Interns: Micro-Randomized Trial. *Journal of Medical Internet Research* **22** e15033.

- PAMFIL, R., SRIWATTANAWORACHAI, N., DESAI, S., PILGERSTORFER, P., GEORGATZIS, K., BEAUMONT, P. and ARAGAM, B. (2020). Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics* 1595–1605. Pmlr.
- PEARL, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge University Press.
- PEARL, J. (2012). The causal mediation formula—a guide to the assessment of pathways and mechanisms. *Prevention science* **13** 426–436.
- PETERS, J. and BÜHLMANN, P. (2014). Identifiability of Gaussian structural equation models with equal error variances. *Biometrika* **101** 219–228.
- PREACHER, K. J. (2015). Advances in mediation analysis: A survey and synthesis of new developments. *Annual Review of Psychology* **66** 825–852.
- PUTERMAN, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- QIN, Z. T., ZHU, H. and YE, J. (2021). Reinforcement learning for ridesharing: A survey. In *2021 IEEE international intelligent transportation systems conference (ITSC)* 2447–2454. IEEE.
- RAMPRASAD, P., LI, Y., YANG, Z., WANG, Z., SUN, W. W. and CHENG, G. (2022). Online bootstrap inference for policy evaluation in reinforcement learning. *Journal of the American Statistical Association* 1–14.
- RUCKER, D. D., PREACHER, K. J., TORMALA, Z. L. and PETTY, R. E. (2011). Mediation analysis in social psychology: Current practices and new recommendations. *Social and Personality Psychology Compass* **5** 359–371.
- SAMPSON, J. N., BOCA, S. M., MOORE, S. C. and HELLER, R. (2018). FWER and FDR control when testing multiple mediators. *Bioinformatics* **34** 2418–2424.
- SELIG, J. P. and PREACHER, K. J. (2009). Mediation models for longitudinal data in developmental research. *Research in Human Development* **6** 144–164.
- SHAFFER, F. and GINSBERG, J. P. (2017). An overview of heart rate variability metrics and norms. *Frontiers in Public Health* 258.
- SHI, C. and LI, L. (2022). Testing mediation effects using logic of Boolean matrices. *Journal of the American Statistical Association* **117** 2014–2027.
- SHI, C., ZHANG, S., LU, W. and SONG, R. (2022). Statistical inference of the value function for reinforcement learning in infinite-horizon settings. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **84** 765–793.
- SHUAI, K., LIU, L., HE, Y. and LI, W. (2023). Mediation pathway selection with unmeasured mediator-outcome confounding. *arXiv preprint arXiv:2311.16793*.
- SILVER, D., SCHRITTWIESER, J., SIMONYAN, K., ANTONOGLOU, I., HUANG, A., GUEZ, A., HUBERT, T., BAKER, L., LAI, M., BOLTON, A. et al. (2017). Mastering the game of go without human knowledge. *nature* **550** 354–359.
- SUTTON, R. S. and BARTO, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- VANDERWEELE, T. J. (2016). Mediation Analysis: A Practitioner’s Guide. *Annual Review of Public Health* **37** 17–32.
- VANDERWEELE, T. J. and TCHETGEN TCHETGEN, E. J. (2017). Mediation analysis with time varying exposures and mediators. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **79** 917–938.
- WANG, J., QI, Z. and WONG, R. K. (2023). Projected state-action balancing weights for offline reinforcement learning. *Annals of Statistics* **accepted**.
- WANG, J., SHI, C. and WU, Z. (2023). A robust test for the stationarity assumption in sequential decision making. In *International Conference on Machine Learning* 36355–36379. PMLR.
- WEI, H., CAI, H., SHI, C. and SONG, R. (2024a). On Efficient Inference of Causal Effects with Multiple Mediators. *arXiv preprint arXiv:2401.05517*.
- WEI, K., XUE, F., XU, Q., YUAN, Y., ZHANG, Y., QIN, G., WANI, A. H., AIELLO, A. E., WILDMAN, D. E., UDDIN, M. et al. (2024b). Time-varying mediation analysis for incomplete data with application to DNA methylation study for PTSD. *bioRxiv* 2024–02.
- WRIGHT, S. (1921). Correlation and causation. *Journal of Agricultural Research* **20** 557–585.
- YUAN, Y. and QU, A. (2023). De-confounding causal inference using latent multiple-mediator pathways. *arXiv preprint arXiv:2302.05513*.
- YUAN, Y., SHEN, X., PAN, W. and WANG, Z. (2019). Constrained likelihood for reconstructing a directed acyclic Gaussian graph. *Biometrika* **106** 109–125.
- ZHANG, H., CHEN, X. and YANG, L. F. (2023). Adaptive Liquidity Provision in Uniswap V3 with Deep Reinforcement Learning. *arXiv preprint arXiv:2309.10129*.
- ZHANG, H., ZHENG, Y., ZHANG, Z., GAO, T., JOYCE, B., YOON, G., ZHANG, W., SCHWARTZ, J., JUST, A., COLICINO, E. et al. (2016). Estimating and testing high-dimensional mediation effects in epigenetic studies. *Bioinformatics* **32** 3150–3154.

- ZHAO, Y., LI, L. and INITIATIVE, A. D. N. (2022). Multimodal data integration via mediation analysis with high-dimensional exposures and mediators. *Human Brain Mapping* **43** 2519–2533.
- ZHAO, Y. and LUO, X. (2022). Pathway lasso: pathway estimation and selection with high dimensional mediators. *Statistics and Its Interface* **15** 39–50.
- ZHAO, Y., LUO, X., LINDQUIST, M. and CAFFO, B. (2018). Functional mediation analysis with an application to functional magnetic resonance imaging data. *arXiv preprint arXiv:1805.06923*.
- ZHENG, W. and VAN DER LAAN, M. (2017). Longitudinal mediation analysis with time-varying mediators and exposures, with application to survival outcomes. *Journal of causal inference* **5** 20160006.
- ZHENG, X., ARAGAM, B., RAVIKUMAR, P. K. and XING, E. P. (2018). DAGs with NO TEARS: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems* 9472–9483.

1. Proof of Theorem 1. By Definition 1, the individual mediation effect of the j th mediator, $j = 1, \dots, d$, is of the form,

$$\begin{aligned}
 \eta_j^{(t)} &= \sum_{s=1}^t \frac{\partial}{\partial a} \mathbb{E}[R_s \mid do(A_i = a, \forall i \leq s)] - \sum_{s=1}^t \frac{\partial}{\partial a} \mathbb{E}[R_s \mid do(A_i = a, M_{ij} = m, \forall i \leq s)] \\
 &= (t-1)\eta_j^{(t-1)} + \left\{ \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_i = a, \forall i \leq t)] \right. \\
 &\quad \left. - \frac{\partial}{\partial a} \mathbb{E}[R_t \mid do(A_i = a, M_{ij} = m, \forall i \leq t)] \right\} \\
 (26) \quad &= (t-1)\eta_j^{(t-1)} + \left(\theta_{A_1, \dots, A_t \rightarrow R_t} - \theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} \right).
 \end{aligned}$$

The key is to express (26) as a function of the intermediate quantities in Section 3.2 of the paper.

First, because all the treatments $(A_1, \dots, A_t)$ are randomly assigned, and thus are independent of each other and all other covariates, we have

$$\begin{aligned}
 \theta_{A_1, \dots, A_t \rightarrow R_t} &= \theta_{A_1 \rightarrow R_t} + \theta_{A_2 \rightarrow R_t} + \dots + \theta_{A_t \rightarrow R_t}, \\
 \theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} &= \theta_{A_1 \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} + \theta_{A_2 \rightarrow R_t}^{M_{2j}, \dots, M_{tj}} + \dots + \theta_{A_t \rightarrow R_t}^{M_{tj}}.
 \end{aligned}$$

Next, we expand the terms in the summation of the second equation. Note that, when $s = t$, $\theta_{A_t \rightarrow R_t}^{M_{tj}} = \theta_{A_t \rightarrow R_t} - \theta_{A_t \rightarrow M_{tj}} \theta_{M_{tj} \rightarrow R_t}$. For any $s = 1, \dots, t-1$, the effect of A_s to R_t in a joint intervention on $(M_{sj}, M_{(s+1)j}, \dots, M_{tj})$ is

$$\begin{aligned}
 \theta_{A_s \rightarrow R_t}^{M_{sj}, \dots, M_{tj}} &= \theta_{A_s \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} - \theta_{A_s \rightarrow M_{sj}} \theta_{M_{sj} \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} \\
 &= \theta_{A_s \rightarrow R_t}^{M_{(s+2)j}, \dots, M_{tj}} - \theta_{A_s \rightarrow M_{(s+1)j}} \theta_{M_{(s+1)j} \rightarrow R_t}^{M_{(s+2)j}, \dots, M_{tj}} - \theta_{A_s \rightarrow M_{sj}} \theta_{M_{sj} \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} \\
 (27) \quad &= \dots \\
 &= \theta_{A_s \rightarrow R_t} - \sum_{i=s}^t \theta_{A_s \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{i+1j}, \dots, M_{tj}}.
 \end{aligned}$$

Next, we derive the recursive formula for computing $\theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}}$, with $s = 1, \dots, t-1$. The idea is similar to (27), i.e.,

$$\begin{aligned}
 (28) \quad \theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}} &= \theta_{M_{sj} \rightarrow R_t}^{M_{(s+2)j}, \dots, M_{tj}} - \theta_{M_{sj} \rightarrow M_{(s+1)j}} \theta_{M_{(s+1)j} \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} \\
 &= \theta_{M_{sj} \rightarrow R_t}^{M_{(s+3)j}, \dots, M_{tj}} - \theta_{M_{sj} \rightarrow M_{(s+2)j}} \theta_{M_{(s+2)j} \rightarrow R_t}^{M_{(s+2)j}, \dots, M_{tj}} - \theta_{M_{sj} \rightarrow M_{(s+1)j}} \theta_{M_{(s+1)j} \rightarrow R_t}^{M_{(s+1)j}, \dots, M_{tj}} \\
 &= \dots \\
 &= \theta_{M_{sj} \rightarrow R_t} - \sum_{i=s+1}^t \theta_{M_{sj} \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}.
 \end{aligned}$$

Given $\theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}$ for $i = s+1, \dots, t$, the above recursion allows us to evaluate $\theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}}$.

Given the expanded forms in (27) and (28), we compute the total effects from $(A_1, \dots, A_t)$ on R_t in a joint intervention on $(M_{1j}, \dots, M_{tj})$ as,

$$\begin{aligned}
 \theta_{A_1, \dots, A_t \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} &= \theta_{A_1 \rightarrow R_t}^{M_{1j}, \dots, M_{tj}} + \theta_{A_2 \rightarrow R_t}^{M_{2j}, \dots, M_{tj}} + \dots + \theta_{A_t \rightarrow R_t}^{M_{tj}} \\
 &= \left(\theta_{A_1 \rightarrow R_t} - \sum_{i=1}^t \theta_{A_1 \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}} \right) + \left(\theta_{A_2 \rightarrow R_t} - \sum_{i=2}^t \theta_{A_2 \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}} \right) \\
 &\quad + \dots + \left(\theta_{A_t \rightarrow R_t} - \theta_{A_t \rightarrow M_{tj}} \theta_{M_{tj} \rightarrow R_t}^{M_{tj}} \right) \\
 &= \sum_{s=1}^t \theta_{A_s \rightarrow R_t} - \sum_{s=1}^t \sum_{i=s}^t \theta_{A_s \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}},
 \end{aligned}$$

where $\theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}$ can be calculated via (28).

Finally, we plug in the above expansion to (26), and obtain that,

$$\eta_j^{(t)} = \eta_j^{(t-1)} + \sum_{s=1}^t \sum_{i=s}^t \theta_{A_s \rightarrow M_{ij}} \theta_{M_{ij} \rightarrow R_t}^{M_{ij}, \dots, M_{tj}}, \quad t = 1, \dots, T.$$

This completes the proof of Theorem 1. □

2. Proof of Theorem 2. First, by the method of induction, we have,

$$\begin{aligned}
 \eta_j^{(1)} &= \theta_{M_{1j} \rightarrow R_1} \theta_{A_1 \rightarrow M_{1j}}, \\
 \eta_j^{(2)} &= \theta_{M_{1j} \rightarrow R_1} (2\theta_{A_1 \rightarrow M_{1j}} + \theta_{A_1 \rightarrow M_{2j}}) + \theta_{M_{1j} \rightarrow R_2}^{M_{1j}, M_{2j}} \theta_{A_1 \rightarrow M_{1j}}, \\
 \eta_j^{(3)} &= \theta_{M_{1j} \rightarrow R_1} (3\theta_{A_1 \rightarrow M_{1j}} + 2\theta_{A_1 \rightarrow M_{2j}} + \theta_{A_1 \rightarrow M_{3j}}) + \theta_{M_{1j} \rightarrow R_3}^{M_{1j}, M_{2j}, M_{3j}} \theta_{A_1 \rightarrow M_{1j}} \\
 &\quad + \theta_{M_{1j} \rightarrow R_2}^{M_{1j}, M_{2j}} (2\theta_{A_1 \rightarrow M_{1j}} + \theta_{A_1 \rightarrow M_{2j}}), \\
 &\dots \\
 \eta_j^{(T)} &= \theta_{M_{1j} \rightarrow R_1} \sum_{t=1}^T (T-t+1) \theta_{A_1 \rightarrow M_{tj}} \\
 &\quad + \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} \left\{ \sum_{s=1}^{T-t} (T-t-s+1) \theta_{A_1 \rightarrow M_{sj}} \right\}.
 \end{aligned}$$

Dividing the right-hand side of the above equation by T and taking $T \rightarrow \infty$, we obtain the formula of individual mediation effect in an infinite-horizon setting:

$$\begin{aligned}\eta_j^{(\infty)} &= \theta_{M_{1j} \rightarrow R_1} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T (T-t+1) \theta_{A_1 \rightarrow M_{tj}} \\ &\quad + \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} \left\{ \sum_{s=1}^{T-t} (T-t-s+1) \theta_{A_1 \rightarrow M_{sj}} \right\} \\ &\equiv \theta_{M_{1j} \rightarrow R_1} I_{1j} + I_{2j}.\end{aligned}$$

We next derive the two terms I_{1j} and I_{2j} , respectively,

For I_{1j} , we have

$$\begin{aligned}I_{1j} &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T (T-t+1) \theta_{A_1 \rightarrow M_{tj}} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \{ T \theta_{A_1 \rightarrow M_{1j}} + (T-1) \theta_{A_1 \rightarrow M_{2j}} + \dots + \theta_{A_1 \rightarrow M_{Tj}} \} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \sum_{s=1}^t \theta_{A_1 \rightarrow M_{sj}}.\end{aligned}$$

Following the recursive relation in (18) of the paper, we first obtain the following equations with the time-invariant model parameters,

$$\begin{aligned}\theta_{A_1 \rightarrow M_1} &= \delta_1, \\ \theta_{A_1 \rightarrow M_2} &= \Gamma_1 \theta_{A_1 \rightarrow M_1} + \zeta_1 \theta_{A_1 \rightarrow R_1}, \\ &\dots \\ \theta_{A_1 \rightarrow M_t} &= \Gamma_1 \theta_{A_1 \rightarrow M_{t-1}} + \zeta_1 \theta_{A_1 \rightarrow R_{t-1}}.\end{aligned}$$

Taking summation over the above equations leads to

$$\sum_{s=1}^t \theta_{A_1 \rightarrow M_s} = \delta_1 + \Gamma_1 \sum_{s=1}^{t-1} \theta_{A_1 \rightarrow M_s} + \zeta_1 \sum_{s=1}^{t-1} \theta_{A_1 \rightarrow R_s}.$$

Taking another summation of the above equations leads to

$$(29) \quad \frac{1}{T} \sum_{t=2}^T \sum_{s=1}^t \theta_{A_1 \rightarrow M_s} = \frac{T-1}{T} \delta_1 + \Gamma_1 \frac{1}{T} \sum_{t=1}^{T-1} \sum_{s=1}^t \theta_{A_1 \rightarrow M_s} + \zeta_1 \frac{1}{T} \sum_{t=1}^{T-1} \sum_{s=1}^t \theta_{A_1 \rightarrow R_s}.$$

Similarly, to evaluate $\sum_{s=1}^t \theta_{A_1 \rightarrow R_s}$, we first obtain the recursion relations,

$$\begin{aligned}\theta_{A_1 \rightarrow R_1} &= \delta_2 + \kappa^\top \theta_{A_1 \rightarrow M_1}, \\ \theta_{A_1 \rightarrow R_2} &= \zeta_2 \theta_{A_1 \rightarrow R_1} + \kappa^\top \theta_{A_1 \rightarrow M_2} + \gamma_2^\top \theta_{A_1 \rightarrow M_1}, \\ &\dots \\ \theta_{A_1 \rightarrow R_t} &= \zeta_2 \theta_{A_1 \rightarrow R_{t-1}} + \kappa^\top \theta_{A_1 \rightarrow M_t} + \gamma_2^\top \theta_{A_1 \rightarrow M_{t-1}}.\end{aligned}$$

Taking double summation over the above equations leads to

$$\frac{1}{T} \sum_{t=2}^T \sum_{s=1}^t \theta_{A_1 \rightarrow R_s} = \frac{T-1}{T} \delta_2 + \zeta_2 \frac{1}{T} \sum_{t=1}^{T-1} \sum_{s=1}^t \theta_{A_1 \rightarrow R_s}$$

$$(30) \quad +\kappa^\top \frac{1}{T} \sum_{t=1}^{T-1} \sum_{s=1}^t \theta_{A_1 \rightarrow M_s} + \gamma_2^\top \frac{1}{T} \sum_{t=1}^{T-1} \sum_{s=1}^t \theta_{A_1 \rightarrow M_s}$$

Denote $\mathbf{B}_1 \equiv \lim_{T \rightarrow \infty} T^{-1} \sum_{t=2}^T \sum_{s=1}^t \theta_{A_1 \rightarrow M_s} \in \mathbb{R}^d$, and $\tilde{R} \equiv \lim_{T \rightarrow \infty} T^{-1} \sum_{t=2}^T \theta_{A_1 \rightarrow R_s}$. Combining (29) and (30), we obtain that

$$\begin{aligned} \mathbf{B}_1 &= \delta_1 + \Gamma_1 \mathbf{B}_1 + \zeta_1 \tilde{R}, \\ \tilde{R} &= \delta_2 + \zeta_2 \tilde{R} + \kappa^\top \mathbf{B}_1 + \gamma_2^\top \mathbf{B}_1. \end{aligned}$$

Solving the above equations leads to

$$\mathbf{B}_1 = \{(1 - \zeta_2)(\mathbf{I} - \Gamma_1) - \zeta_1(\kappa + \gamma_2)^\top\}^{-1} \{(1 - \zeta_2)\delta_1 + \delta_2\zeta_1\}.$$

Then I_{1j} is the j th element in the vector $\mathbf{B}_1$.

For I_{2j} , we have

$$\begin{aligned} I_{2j} &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} \left\{ \sum_{s=1}^{T-t} (T-t-s+1) \theta_{A_1 \rightarrow M_{sj}} \right\} \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} \left(\sum_{s=1}^{T-t} \sum_{i=1}^s \theta_{A_1 \rightarrow M_{ij}} \right). \end{aligned}$$

For notational simplicity, let $L_{tj} \equiv \sum_{s=1}^t \sum_{i=1}^s \theta_{A_1 \rightarrow M_{ij}}$. Next, we derive the recursive relationship for the summation $T^{-1} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} L_{(T-t)j}$.

Following the recursion for $\theta_{M_{sj} \rightarrow R_t}^{M_{sj}, \dots, M_{tj}}$ in (28), and multiplying both sides by L_{tj} , we obtain that

$$\begin{aligned} \theta_{M_{1j} \rightarrow R_2}^{M_{1j}, M_{2j}} L_{(T-1)j} &= L_{(T-1)j} \theta_{M_{1j} \rightarrow R_2} - \theta_{M_{1j} \rightarrow M_{2j}} \theta_{M_{1j} \rightarrow R_1} L_{(T-1)j}, \\ \theta_{M_{1j} \rightarrow R_3}^{M_{1j}, \dots, M_{3j}} L_{(T-2)j} &= L_{(T-2)j} \theta_{M_{1j} \rightarrow R_3} - \theta_{M_{1j} \rightarrow M_{3j}} \theta_{M_{1j} \rightarrow R_1} L_{(T-2)j} - \theta_{M_{1j} \rightarrow M_{2j}} \theta_{M_{1j} \rightarrow R_2}^{M_{1j}, M_{2j}} L_{(T-2)j}, \\ \theta_{M_{1j} \rightarrow R_4}^{M_{1j}, \dots, M_{4j}} L_{(T-3)j} &= L_{(T-3)j} \theta_{M_{1j} \rightarrow R_4} - \theta_{M_{1j} \rightarrow M_{4j}} \theta_{M_{1j} \rightarrow R_1} L_{(T-3)j} - \sum_{i=2}^3 \theta_{M_{1j} \rightarrow M_{ij}} \theta_{M_{1j} \rightarrow R_{4-i+1}}^{M_{1j}, \dots, M_{(4-i+1)j}} L_{(T-3)j}, \\ &\dots \\ \theta_{M_{1j} \rightarrow R_T}^{M_{1j}, \dots, M_{Tj}} L_{1j} &= L_{1j} \theta_{M_{1j} \rightarrow R_T} - \theta_{M_{1j} \rightarrow M_{Tj}} \theta_{M_{1j} \rightarrow R_1} L_{1j} - \sum_{i=2}^{T-1} \theta_{M_{1j} \rightarrow M_{ij}} \theta_{M_{1j} \rightarrow R_{T-i+1}}^{M_{1j}, \dots, M_{(T-i+1)j}} L_{1j}. \end{aligned}$$

Taking summation over the above equations leads to

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} L_{(T-t)j} &= \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow R_t} L_{(T-t)j} - \theta_{M_{1j} \rightarrow R_1} \frac{1}{T} \sum_{t=1}^{T-1} \theta_{M_{1j} \rightarrow M_{tj}} L_{(T-t)j} \\ &\quad - \sum_{i=2}^{T-1} \theta_{M_{1j} \rightarrow M_{ij}} \frac{1}{T} \sum_{t=1}^{T-i} \theta_{M_{1j} \rightarrow R_{t+1}}^{M_{1j}, \dots, M_{(t+1)j}} L_{(T-i-t+1)j}. \end{aligned}$$

Define $\mathbf{B}_2 \equiv \lim_{T \rightarrow \infty} \sum_{t=2}^T \theta_{M_{1j} \rightarrow M_t} \in \mathbb{R}^d$, $\mathbf{B}_{3j} \equiv \lim_{T \rightarrow \infty} \sum_{t=2}^T \theta_{M_{1j} \rightarrow R_t}$, $\mathbf{B}_4 \equiv \lim_{T \rightarrow \infty} T^{-1} \sum_{t=2}^T \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j}$, and $\mathbf{B}_{5j} \equiv \lim_{T \rightarrow \infty} T^{-1} \sum_{t=2}^T \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j}$. Taking the limit with $T \rightarrow \infty$ on both sides of the above equation leads to

$$(31) \quad I_{2j} = (1 + \mathbf{B}_{2j})^{-1} (\mathbf{B}_{5j} - \theta_{M_{1j} \rightarrow R_1} \mathbf{B}_{4j}).$$

Next, we obtain the recursive relationship for $\sum_{t=2}^T \theta_{M_{1j} \rightarrow M_t}$, $\sum_{t=2}^T \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j}$, and $\sum_{t=2}^T \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j}$. Let $B_6 \equiv \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_d \\ \boldsymbol{\kappa}^\top & 0 \end{pmatrix}$ and $B_7 \equiv \begin{pmatrix} \Gamma_1 & \zeta_1 \\ \gamma_2^\top & \zeta_2 \end{pmatrix}$. The recursion between $(\theta_{M_{1j} \rightarrow M_t}, \theta_{M_{1j} \rightarrow R_t})$ and $(\theta_{M_{1j} \rightarrow M_{t-1}}, \theta_{M_{1j} \rightarrow R_{t-1}})$ takes the following form:

$$(32) \quad \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} \\ \theta_{M_{1j} \rightarrow R_t} \end{bmatrix} = \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_d \\ \boldsymbol{\kappa}^\top & 0 \end{pmatrix} \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} \\ \theta_{M_{1j} \rightarrow R_t} \end{bmatrix} + \begin{pmatrix} \Gamma_1 & \zeta_1 \\ \gamma_2^\top & \zeta_2 \end{pmatrix} \begin{bmatrix} \theta_{M_{1j} \rightarrow M_{t-1}} \\ \theta_{M_{1j} \rightarrow R_{t-1}} \end{bmatrix}$$

Taking summation from $t = 2$ to T of the above equation leads to

$$\begin{aligned} \sum_{t=2}^T (\mathbf{I} - B_6) \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} \\ \theta_{M_{1j} \rightarrow R_t} \end{bmatrix} &= \sum_{t=2}^T B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_{t-1}} \\ \theta_{M_{1j} \rightarrow R_{t-1}} \end{bmatrix} \\ (\mathbf{I} - B_6) \sum_{t=2}^T \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} \\ \theta_{M_{1j} \rightarrow R_t} \end{bmatrix} &= B_7 \sum_{t=2}^{T-1} \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} \\ \theta_{M_{1j} \rightarrow R_t} \end{bmatrix} + B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_1} \\ \theta_{M_{1j} \rightarrow R_1} \end{bmatrix} \end{aligned}$$

Taking the limit with $T \rightarrow \infty$ on both sides leads to

$$(33) \quad \begin{bmatrix} B_2 \\ B_{3j} \end{bmatrix} = (\mathbf{I} - B_6 - B_7)^{-1} B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_1} \\ \theta_{M_{1j} \rightarrow R_1} \end{bmatrix}.$$

Multiply both sides of the recursion in (32) by $L_{(T-t+1)j}$ leads to

$$\begin{aligned} \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j} \end{bmatrix} &= B_6 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j} \end{bmatrix} + B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_{t-1}} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_{t-1}} L_{(T-t+1)j} \end{bmatrix}, \\ \frac{1}{T} \sum_{t=2}^T (\mathbf{I} - B_6) \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j} \end{bmatrix} &= \frac{1}{T} \sum_{t=2}^T B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_{t-1}} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_{t-1}} L_{(T-t+1)j} \end{bmatrix} \\ (\mathbf{I} - B_6 - B_7) \frac{1}{T} \sum_{t=2}^T \begin{bmatrix} \theta_{M_{1j} \rightarrow M_t} L_{(T-t+1)j} \\ \theta_{M_{1j} \rightarrow R_t} L_{(T-t+1)j} \end{bmatrix} &= \frac{1}{T} B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_1} L_{(T-1)j} \\ \theta_{M_{1j} \rightarrow R_1} L_{(T-1)j} \end{bmatrix} \end{aligned}$$

Furthermore, we have $\lim_{T \rightarrow \infty} \frac{1}{T} L_{(T-1)j} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^{T-1} \sum_{i=1}^t \theta_{A_1 \rightarrow M_{ij}} = B_{1j}$. We now take the limit with $T \rightarrow \infty$ on both sides of the summations, we have

$$(34) \quad \begin{bmatrix} B_4 \\ B_{5j} \end{bmatrix} = (\mathbf{I} - B_6 - B_7)^{-1} B_7 \begin{bmatrix} \theta_{M_{1j} \rightarrow M_1} \\ \theta_{M_{1j} \rightarrow R_1} \end{bmatrix} B_{1j}.$$

We then combine (31), (33) and (34) and solve for I_{2j} .

This completes the proof of Theorem 2. $\square$

3. Proof of Theorem 3. We begin with some notations. Let $\text{vech}(\mathbf{B}) \in \mathbb{R}^{q(q+1)/2}$ denote the half-vectorization of a symmetric $q \times q$ matrix $\mathbf{B}$ that vectorizes the lower triangular part of $\mathbf{B}$. Let $\partial \mathbf{y} / \partial \mathbf{x}$ denote the $r \times s$ matrix of the derivative of $\mathbf{y}(\mathbf{x}) = (y_1(\mathbf{x}), \dots, y_r(\mathbf{x}))^\top \in \mathbb{R}^r$ with respect to $\mathbf{x} = (x_1, \dots, x_s)^\top \in \mathbb{R}^s$ whose (i, j) th entry is $\partial y_i / \partial x_j, i = 1, \dots, r, j = 1, \dots, s$.

We first prove the case of $T = 1$, then the case of $T = 2$. The case for any finite $T > 2$ can be proved in a similar fashion.

For $T = 1$, recall $\eta_j^{(1)} = \theta_{A_1 \rightarrow M_{1j}} \theta_{M_{1j} \rightarrow R_1} = \beta_{1j} \beta_{M_{1j}, R_1 | \text{Pa}_j \cup A_1}$, and $\hat{\eta}_j^{(1)} = \hat{\theta}_{A_1 \rightarrow M_{1j}} \hat{\theta}_{M_{1j} \rightarrow R_1} = \hat{\beta}_{1j} \hat{\beta}_{M_{1j}, R_1 | \text{Pa}_j \cup A_1}$ as $n \rightarrow \infty$. To establish the asymptotic distribution of $\hat{\eta}_j^{(1)}$, the main idea is to first show the asymptotic distribution of $(\hat{\theta}_{A_1 \rightarrow M_{1j}}, \hat{\theta}_{M_{1j} \rightarrow R_1})^\top$, then apply the delta method.

Let $\mathbf{U}_{1j} \equiv (U_{1j1}, \dots, U_{1jq_1})^\top = (A_1, \text{Pa}_{1j}^\top, M_{1j}, R_1)^\top \in \mathbb{R}^{q_1}$. Let $\Sigma_{1j} = \text{cov}(\mathbf{U}_{1j}) \in \mathbb{R}^{q_1 \times q_1}$, and let $\widehat{\Sigma}_{1j}$ denote the corresponding sample covariance matrix. Since all fourth moments of the variables in $\mathbf{U}_{1j}$ are finite, the central limit theorem implies that

$$(35) \quad \sqrt{n} \left\{ \text{vech}(\widehat{\Sigma}_{1j}) - \text{vech}(\Sigma_{1j}) \right\} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}_1),$$

where $\mathbf{V}_{1j} = \text{cov} \left(\text{vech}(\mathbf{U}_{1j} \mathbf{U}_{1j}^\top) \right)$.

Let $\boldsymbol{\theta}_{1j} \equiv (\theta_{A_1 \rightarrow M_{1j}}, \theta_{M_{1j} \rightarrow R_1})^\top = (\beta_{1j}, \beta_{M_{1j}, R_1 | \text{Pa}_{1j} \cup A_1})^\top$, and $\widehat{\boldsymbol{\theta}}_{1j} \equiv (\widehat{\theta}_{A_1 \rightarrow M_{1j}}, \widehat{\theta}_{M_{1j} \rightarrow R_1})^\top = (\widehat{\beta}_{1j}, \widehat{\beta}_{M_{1j}, R_1 | \text{Pa}_{1j} \cup A_1})^\top$. There exists a differentiable function $h_{11} : \mathbb{R}^{q_1(q_1+1)/2} \rightarrow \mathbb{R}^2$, such that $h_{11}(\text{vech}(\Sigma_{1j})) = \boldsymbol{\theta}_{1j}$, and $h_{11}(\text{vech}(\widehat{\Sigma}_{1j})) = \widehat{\boldsymbol{\theta}}_{1j}$. Therefore, by (35) and the multivariate delta method, we have

$$\sqrt{n}(\widehat{\boldsymbol{\theta}}_{1j} - \boldsymbol{\theta}_{1j}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Lambda_{1j} \mathbf{V}_{1j} \Lambda_{1j}^\top),$$

where $\Lambda_{1j} = \partial \boldsymbol{\theta}_{1j} / \partial \text{vech}(\Sigma_{1j})$.

Consider another differentiable function $h_{12}(x_1, x_2) = x_1 x_2$, and apply the multivariate delta method again, we have

$$\sqrt{n}(\widehat{\eta}_j^{(1)} - \eta_j^{(1)}) \xrightarrow{d} \mathcal{N}(0, \mathbf{S}_{1j} \Lambda_{1j} \mathbf{V}_{1j} \Lambda_{1j}^\top \mathbf{S}_{1j}^\top),$$

where $\mathbf{V}_{1j} = \text{cov} \left(\text{vech}(\mathbf{U}_{1j} \mathbf{U}_{1j}^\top) \right)$, $\Lambda_{1j} = \partial \boldsymbol{\theta}_{1j} / \partial (\text{vech}(\Sigma_{1j}))$, $\boldsymbol{\theta}_{1j} = (\theta_{A_1 \rightarrow M_{1j}}, \theta_{M_{1j} \rightarrow R_1})^\top$, and $\mathbf{S}_{1j} = \partial \eta_j^{(1)} / \partial \boldsymbol{\theta}_{1j}$.

For $T = 2$, let $\mathbf{U}_{2j} \equiv (U_{2j1}, \dots, U_{2jq_2})^\top = (A_1, \mathbf{M}_1^\top, R_1, A_2, \text{Pa}_{2j}^\top, M_{2j}, R_2)^\top \in \mathbb{R}^{q_2}$, where $\text{Pa}_{2j} = \{M_{2i} : (\mathbf{W}_t)_{ij} \neq 0, i \in \{1, \dots, d\} \setminus \{j\}\}$ is the parent set of M_{2j} in $\mathcal{G}_2$. Let $\Sigma_{2j} = \text{cov}(\mathbf{U}_{2j}) \in \mathbb{R}^{q_2 \times q_2}$ and let $\widehat{\Sigma}_{2j}$ denote the corresponding sample covariance matrix. Similar to (35), we have

$$\sqrt{n} \left\{ \text{vech}(\widehat{\Sigma}_{2j}) - \text{vech}(\Sigma_{2j}) \right\} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}_{2j}),$$

where $\mathbf{V}_{2j} = \text{cov} \left(\text{vech}(\mathbf{U}_{2j} \mathbf{U}_{2j}^\top) \right)$.

Let $\boldsymbol{\theta}_{2j} \equiv (\theta_{A_1 \rightarrow M_{1j}}, \theta_{M_{1j} \rightarrow R_1}, \theta_{A_1 \rightarrow M_{2j}}, \theta_{M_{1j} \rightarrow M_{2j}}, \theta_{M_{1j} \rightarrow R_2}, \theta_{A_2 \rightarrow M_{2j}}, \theta_{M_{2j} \rightarrow R_2})^\top$, and $\widehat{\boldsymbol{\theta}}_{2j} \equiv (\widehat{\theta}_{A_1 \rightarrow M_{1j}}, \widehat{\theta}_{M_{1j} \rightarrow R_1}, \widehat{\theta}_{A_1 \rightarrow M_{2j}}, \widehat{\theta}_{M_{1j} \rightarrow M_{2j}}, \widehat{\theta}_{M_{1j} \rightarrow R_2}, \widehat{\theta}_{A_2 \rightarrow M_{2j}}, \widehat{\theta}_{M_{2j} \rightarrow R_2})^\top$. Note that every element in $\boldsymbol{\theta}_{2j}$ represents the total effect of one variable on another, where all variables in this expanded DAG up to the second stage are included in $\mathbf{U}_{2j}$. For example, $\theta_{M_{1j} \rightarrow R_2} = \beta_{M_{1j}, R_2 | \text{Pa}_j \cup A_1}$ where this regression coefficient depends only on the covariance matrix of $\mathbf{U}_{2j}$, which is Σ_{2j} . Then there exists a differentiable function $h_{21} : \mathbb{R}^{q_2(q_2+1)/2} \rightarrow \mathbb{R}^7$, such that $h_{21}(\text{vech}(\Sigma_{2j})) = \boldsymbol{\theta}_{2j}$ and $h_{21}(\text{vech}(\widehat{\Sigma}_{2j})) = \widehat{\boldsymbol{\theta}}_{2j}$. By the multivariate delta method, we have

$$\sqrt{n}(\widehat{\boldsymbol{\theta}}_{2j} - \boldsymbol{\theta}_{2j}) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \Lambda_{2j} \mathbf{V}_{2j} \Lambda_{2j}^\top),$$

where $\Lambda_{2j} = \partial \boldsymbol{\theta}_{2j} / \partial (\text{vech}(\Sigma_{2j}))$.

Consider another differentiable function $h(x_1, x_2, \dots, x_7) = \{x_1 x_2 + (x_3 + x_6) \cdot x_7 + x_1(x_5 - x_4 x_7)\} / 2$, and apply the multivariate delta method again, we have

$$\sqrt{n}(\widehat{\eta}_j^{(2)} - \eta_j^{(2)}) \xrightarrow{d} \mathcal{N}(0, \mathbf{S}_{2j} \Lambda_{2j} \mathbf{V}_{2j} \Lambda_{2j}^\top \mathbf{S}_{2j}^\top),$$

where $\mathbf{V}_{2j} = \text{cov} \left(\text{vech}(\mathbf{U}_{2j} \mathbf{U}_{2j}^\top) \right)$, $\Lambda_{2j} = \partial \boldsymbol{\theta}_{2j} / \partial (\text{vech}(\Sigma_{2j}))$, and $\mathbf{S}_{2j} = \partial \eta_j^{(2)} / \partial \boldsymbol{\theta}_{2j}$.

Generally, for $t = 1, \dots, T$, let $\mathbf{U}_{tj} \equiv (U_{tj1}, \dots, U_{tjq_t})^\top = (\bar{A}_{t-1}, \bar{\mathbf{M}}_{t-1}^\top, \bar{R}_{t-1}, A_t, \text{Pa}_{tj}^\top, M_{tj}, R_t)^\top \in \mathbb{R}^{q_t}$, where $\text{Pa}_{tj} = \{M_{ti} : (\mathbf{W}_t)_{ij} \neq 0, i \in \{1, \dots, d\} \setminus \{j\}\}$ is the parent set of M_{tj} in $\mathcal{G}_t$, and

$\boldsymbol{\theta}_{tj} \equiv (\boldsymbol{\theta}_{(t-1)j}^\top, \theta_{A_1 \rightarrow M_{tj}}, \dots, \theta_{A_t \rightarrow M_{tj}}, \theta_{M_{1j} \rightarrow M_{tj}}, \dots, \theta_{M_{t-1,j} \rightarrow M_{tj}}, \theta_{M_{1j} \rightarrow R_t}, \dots, \theta_{M_{tj} \rightarrow R_t})^\top \in \mathbb{R}^{p_t}$. For $T > 2$, it follows that

$$\sqrt{n}(\hat{\eta}_j^{(t)} - \eta_j^{(t)}) \xrightarrow{d} \mathcal{N}(0, \sigma_{tj}^2), \text{ as } n \rightarrow \infty,$$

where $\sigma_{tj}^2 = \mathbf{S}_{tj} \boldsymbol{\Lambda}_{tj} \mathbf{V}_{tj} \boldsymbol{\Lambda}_{tj}^\top \mathbf{S}_{tj}^\top$, $\mathbf{V}_{tj} = \text{cov}(\text{vech}(\mathbf{U}_{tj} \mathbf{U}_{tj}^\top)) \in \mathbb{R}^{q_t(q_t+1)/2 \times q_t(q_t+1)/2}$, $\boldsymbol{\Lambda}_{tj} = \partial \boldsymbol{\theta}_{tj} / \partial (\text{vech}(\boldsymbol{\Sigma}_{tj})) \in \mathbb{R}^{p_t \times q_t(q_t+1)/2}$, $\boldsymbol{\Sigma}_{tj} = \text{cov}(\mathbf{U}_{tj}) \in \mathbb{R}^{q_t \times q_t}$, and $\mathbf{S}_{tj} = \partial \eta_j^{(t)} / \partial \boldsymbol{\theta}_{tj} \in \mathbb{R}^{1 \times p_t}$.

This completes the proof of Theorem 3. $\square$

.4. Proof of Theorem 4. From the proof of Theorem 2, recall that $I_{1j} = B_{1j}$, and $I_{2j} = (1 + B_{2j})^{-1}(B_{5j} - \theta_{M_{1j} \rightarrow R_1} B_{4j})$. Therefore, I_{1j}, I_{2j} are continuous functions of the intermediate quantities $(\theta_{M_{1j} \rightarrow R_1}, \boldsymbol{\theta}_{M_{1j} \rightarrow M_1})$, and the model parameters $\boldsymbol{\Theta}_1 = (\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \boldsymbol{\Gamma}_1, \boldsymbol{\zeta}_1)^\top \in \mathbb{R}^{(d+3) \times d}$, $\boldsymbol{\Theta}_2 = (\alpha_2, \beta_2, \zeta_2, \boldsymbol{\gamma}_2^\top, \boldsymbol{\kappa}^\top)^\top \in \mathbb{R}^{(2d+3) \times 1}$. To establish the asymptotic distribution of $\hat{\eta}_j^{(\infty)} = \hat{\theta}_{M_{1j} \rightarrow R_1} \hat{I}_{1j} + \hat{I}_{2j}$, the main idea is to first show the asymptotic distribution of the estimators of the model parameters $(\boldsymbol{\Theta}_1, \boldsymbol{\Theta}_2)$ and those intermediate quantities $(B_1, B_2, B_{3j}, B_4, B_{5j})$, then apply the delta method.

We first show the asymptotic normality of $\hat{\boldsymbol{\Theta}}_2$. The asymptotic normality of $\hat{\boldsymbol{\Theta}}_1$ can be shown similarly. Let $\mathbf{X}_{it} \equiv (1, A_{it}, \mathbf{M}_{i(t-1)}^\top, R_{i(t-1)}, \mathbf{M}_{it}^\top)^\top \in \mathbb{R}^{(2d+3) \times 1}$. Since

$$\sqrt{nT}(\hat{\boldsymbol{\Theta}}_2 - \boldsymbol{\Theta}_2) = \left(\frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \mathbf{X}_{it}^\top \right)^{-1} \left(\frac{1}{\sqrt{nT}} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \epsilon_{i,rt} \right),$$

and the stochastic process $\{(A_t, \mathbf{M}_t^\top, R_t) : t \geq 1\}$ is β -mixing, by the weak law of large numbers of a weak mixing process, we have

$$\begin{aligned} \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \mathbf{X}_{it}^\top &\xrightarrow{p} \mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^\top], \text{ where } \mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^\top] = \lim_{nT \rightarrow \infty} \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \mathbb{E}[\mathbf{X}_{it} \mathbf{X}_{it}^\top], \\ \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \epsilon_{i,rt} &\xrightarrow{p} \mathbf{0}, \text{ because } \mathbf{X}_{it} \perp \epsilon_{i,rt}, \text{ and } \mathbb{E}[\epsilon_{i,rt}] = 0, i = 1, \dots, n. \end{aligned}$$

Next, we derive the asymptotic distribution of $\tilde{\mathbf{U}}_T = \sum_{t=1}^T \mathbf{U}_t \equiv \sum_{t=1}^T \mathbf{X}_t \epsilon_{rt} \in \mathbb{R}^{(2d+3) \times 1}$. For notional simplicity, we drop the subscript i in this part. To apply the Cramér-Wold theorem, let $\mathbf{a} = (a_1, \dots, a_{2d+3})^\top \in \mathbb{R}^{2d+3}$ be a nonrandom and nonzero vector, and $\mathbf{U}_t = (U_{t1}, \dots, U_{t,2d+3})^\top$. Then we write

$$\mathbf{a}^\top \tilde{\mathbf{U}}_T = \sum_{t=1}^T \sum_{p=1}^{2d+3} a_p U_{tp} = \sum_{t=1}^T U'_t.$$

Since $\{\mathbf{U}_t\}_{t \geq 1}$ is a β -mixing process with mean zero, it follows that $\{U'_t\}_{t \geq 1}$ is also a β -mixing centered process with the second moment given by

$$\sigma_T^2 = \mathbf{a}^\top \text{var}(\tilde{\mathbf{U}}_T) \mathbf{a} = \text{var}\left(\sum_{t=1}^T U'_t\right) \rightarrow \infty, T \rightarrow \infty.$$

For every $\epsilon > 0$,

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}[(U'_t)^2 \mathbf{1}(|U'_t| > \epsilon \sigma_T)] &= \sum_{t=1}^T \mathbb{E} \left\{ \left(\sum_{p=1}^{2d+3} a_p U_{tp} \right)^2 \mathbf{1} \left(\sum_{p=1}^{2d+3} |a_p U_{tp}| > \epsilon \sigma_T \right) \right\} \\ &\leq \sum_{t=1}^T a_p^2 \mathbb{E} \left\{ U_{tp}^2 \mathbf{1} \left(\sum_{p=1}^{2d+3} |U_{tp}| > \frac{\epsilon \sigma_T}{\max_p |a_p|} \right) \right\}, \end{aligned}$$

where $\mathbf{1}(\cdot)$ is an indicator function. Since $\sigma_T \rightarrow \infty$, and $\max_p |a_p| < \infty$, we have

$$\mathbf{1} \left(\sum_{p=1}^{2d+3} |U_{tp}| > \frac{\epsilon \sigma_T}{\max_p |a_p|} \right) \xrightarrow{a.s.} 0.$$

In addition, because $\mathbb{E}[U_{tp}^2] < \infty$, and $\mathbb{P}(U_{tp} = \infty) = 0$, it follows that

$$\sum_{t=1}^T \mathbb{E}[(U'_t)^2 \mathbf{1}(|U'_t| > \epsilon \sigma_T)] \rightarrow 0, \quad T \rightarrow \infty.$$

By the central limit theorem for martingales, we have

$$\frac{\sum_{t=1}^T U'_t}{\sigma_T} \xrightarrow{d} \mathcal{N}(0, 1), \quad \text{as } T \rightarrow \infty.$$

Applying the Cramér-Wold theorem, we have

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T U_t = \frac{1}{\sqrt{T}} \sum_{t=1}^T \mathbf{X}_{t\epsilon_{rt}} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}_r),$$

where $\mathbf{V}_r = \sigma_r^2 \mathbb{E}(\mathbf{X}_t \mathbf{X}_t^\top) \in \mathbb{R}^{(2d+3) \times (2d+3)}$. Generalizing to n i.i.d. subjects and by Slutsky's theorem,

$$\sqrt{nT}(\hat{\boldsymbol{\Theta}}_2 - \boldsymbol{\Theta}_2) = \left(\frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \mathbf{X}_{it}^\top \right)^{-1} \left(\frac{1}{\sqrt{nT}} \sum_{i=1}^n \sum_{t=1}^T \mathbf{X}_{it} \epsilon_{i,rt} \right) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}_{\theta_2}),$$

where $\mathbf{V}_{\theta_2} = \mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^\top]^{-1} \mathbf{V}_r \mathbb{E}[\mathbf{X}_0 \mathbf{X}_0^\top]^{-1} \in \mathbb{R}^{(2d+3) \times (2d+3)}$.

Since $\mathbf{B}_1 = \{(1 - \zeta_2)(\mathbf{I} - \boldsymbol{\Gamma}_1) - \boldsymbol{\zeta}_1(\boldsymbol{\kappa} + \boldsymbol{\gamma}_2)^\top\}^{-1} \{(1 - \zeta_2)\boldsymbol{\delta}_1 + \delta_2 \boldsymbol{\zeta}_1\}$ is a continuous function of $(\boldsymbol{\Theta}_1, \boldsymbol{\Theta}_2)$, then by multivariate delta method, the estimator of $\mathbf{B}_1$ which is a continuous function of $(\hat{\boldsymbol{\Theta}}_1, \hat{\boldsymbol{\Theta}}_2)$ is also asymptotically normal. Therefore, we establish the asymptotic normality of $\hat{I}_{1j} = \hat{B}_{1j}$.

Next, we show the asymptotic normality of $\hat{\theta}_{M_{1j} \rightarrow R_1}$. Let $\mathbf{Y}_t \equiv (\mathbf{M}_{t-1}^\top, R_{t-1}, A_t, \text{Pa}_j, M_{tj}, R_t)^\top \in \mathbb{R}^{q \times 1}$, and $\mathbb{E}(\mathbf{Y}_t) = \boldsymbol{\mu}_Y$ is independent of t because of the stationarity assumption. Let $\boldsymbol{\Sigma} = \mathbb{E}[(\mathbf{Y}_t - \boldsymbol{\mu}_Y)(\mathbf{Y}_t - \boldsymbol{\mu}_Y)^\top]$, and $\hat{\boldsymbol{\Sigma}} = T^{-1} \sum_{t=1}^T (\mathbf{Y}_t - \bar{\mathbf{Y}})(\mathbf{Y}_t - \bar{\mathbf{Y}})^\top$, where $\bar{\mathbf{Y}} = T^{-1} \sum_{t=1}^T \mathbf{Y}_t$. Since $\{\mathbf{Y}_t - \boldsymbol{\mu}_Y\}_{t \geq 1}$ is a centered stationary process with mean zero, by the finite fourth moments condition on $\mathbf{Y}_t$'s and the Martingale central limit theorem, we have

$$\sqrt{nT} \left\{ \text{vech}(\hat{\boldsymbol{\Sigma}}) - \text{vech}(\boldsymbol{\Sigma}) \right\} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{V}_0),$$

where $\mathbf{V}_0 = \text{cov}[\text{vech}\{(\mathbf{Y}_t - \boldsymbol{\mu}_Y)(\mathbf{Y}_t - \boldsymbol{\mu}_Y)^\top\}] \in \mathbb{R}^{q(q+1)/2 \times q(q+1)/2}$.

Since $\theta_{M_{1j} \rightarrow R_1} = \dots = \theta_{M_{tj} \rightarrow R_t} = \beta_{M_{tj}, R_t | \text{Pa}_j \cup \mathbf{M}_{t-1} \cup R_{t-1} \cup A_t}$ for $t = 2, \dots, T$, there exists a differentiable function $h(\cdot) : \mathbb{R}^{q(q+1)/2} \rightarrow \mathbb{R}$, such that $h(\text{vech}(\mathbf{\Sigma})) = \theta_{M_{1j} \rightarrow R_1}$, and $h(\text{vech}(\hat{\mathbf{\Sigma}})) = \hat{\theta}_{M_{1j} \rightarrow R_1}$. Then by the multivariate delta method, we have

$$\sqrt{nT}(\hat{\theta}_{M_{1j} \rightarrow R_1} - \theta_{M_{1j} \rightarrow R_1}) \xrightarrow{d} \mathcal{N}(0, \mathbf{\Lambda}_0 \mathbf{V}_0 \mathbf{\Lambda}_0^\top),$$

where $\mathbf{\Lambda}_0 = \partial \theta_{M_{1j} \rightarrow R_1} / \partial (\text{vech}(\mathbf{\Sigma})) \in \mathbb{R}^{1 \times q(q+1)/2}$.

Similarly, we can establish the asymptotic normality of $\hat{\theta}_{M_{1j} \rightarrow M_1}$. Furthermore, since (B_2, B_{3j}) and (B_4, B_{5j}) are continuous functions of $(\Theta_1, \Theta_2, \theta_{M_{1j} \rightarrow M_1}, \theta_{M_{1j} \rightarrow R_1})$, by the multivariate delta method again, the estimator of $I_{2j} = (1 + B_{2j})^{-1}(B_{5j} - \theta_{M_{1j} \rightarrow R_1} B_{4j})$ is also asymptotically normal.

Furthermore, since we have established

$$\sqrt{nT} \begin{pmatrix} \hat{\theta}_{M_{1j} \rightarrow R_1} - \theta_{M_{1j} \rightarrow R_1} \\ \hat{I}_{1j} - I_{1j} \\ \hat{I}_{2j} - I_{2j} \end{pmatrix} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_j),$$

where $\mathbf{\Sigma}_j \in \mathbb{R}^{3 \times 3}$ is a covariance matrix. Finally, since $\eta_j^{(\infty)} = \theta_{M_{1j} \rightarrow R_1} I_{1j} + I_{2j}$ and $\hat{\eta}_j^{(\infty)} = \hat{\theta}_{M_{1j} \rightarrow R_1} \hat{I}_{1j} + \hat{I}_{2j}$. Given a differentiable function $h(x_1, x_2, x_3) = x_1 x_2 + x_3$, and by the multivariate delta method again, we have

$$\sqrt{nT}(\hat{\eta}_j^{(\infty)} - \eta_j^{(\infty)}) \xrightarrow{d} \mathcal{N}(0, \sigma_j^2), \text{ as } nT \rightarrow \infty,$$

where $\sigma_j^2 = \mathbf{V}_j^\top \mathbf{\Sigma}_j \mathbf{V}_j \in \mathbb{R}$ where $\mathbf{V}_j = (I_{1j}, \theta_{M_{1j} \rightarrow R_1}, 1)^\top$. This completes the proof of Theorem 4. $\square$