

Causal interpretation of stochastic differential equations

Alexander Sokol* Niels Richard Hansen†

Abstract

We give a causal interpretation of stochastic differential equations (SDEs) by defining the postintervention SDE resulting from an intervention in an SDE. We show that under Lipschitz conditions, the solution to the postintervention SDE is equal to a uniform limit in probability of postintervention structural equation models based on the Euler scheme of the original SDE, thus relating our definition to mainstream causal concepts. We prove that when the driving noise in the SDE is a Lévy process, the postintervention distribution is identifiable from the generator of the SDE.

Keywords: Stochastic differential equation; Causality; Structural equation model; Identifiability; Lévy process; Weak conditional local independence.

AMS MSC 2010: Primary Primary 60H10, Secondary Secondary 62A01.

Submitted to EJP on June 25, 2013, final version accepted on October 25, 2014.

Supersedes arXiv:1304.0217.

1 Introduction

The notion of causality has long been of interest to both statisticians and scientists working in fields applying statistics. In general, causal models are models containing families of possible distributions of the variables observed as well as appropriate mathematical descriptions of causal structures in the data. Thus, claiming that a causal model is true amounts to claiming more than statements about the distribution of the variables observed. Causal modeling has several goals, prominent among them are:

1. Estimation of intervention effects from fully or partially observed systems with a given causal structure.
2. Identification of the causal structure from observational data.

One of the most developed theories of causal modeling is the approach based on directed acyclic graphs (DAGs) and finitely many variables with no explicit time component, described in [35, 26]. In recent years, there have been efforts to develop similar notions of causality for stochastic processes, both in discrete time and in continuous

*University of Copenhagen, Denmark. E-mail: alexander@math.ku.dk

†University of Copenhagen, Denmark. E-mail: niels.r.hansen@math.ku.dk

time. For discrete-time results, see for example [9, 10, 11]. As discrete-time models often are defined through explicit functional relationships between variables, as in for example autoregressive processes, such models fit directly into the DAG-based framework. In the continuous-time framework, the uncountable number of variables complicates the question of how to describe causal relationships.

Early discussions of causality in a continuous-time framework can be found in [17, 15, 6]. One of the most recent frameworks for causality in continuous time is based on the concept of weak conditional local independence. For results related to this, see [8, 5, 16, 32, 33]. An alternative notion of causality defined solely through filtrations is developed in [29, 28], and a notion of causality in continuous time for ordinary differential equations is introduced in [25].

In Section 4.1 of [1] it is noted that both ordinary differential equations and stochastic differential equations (SDEs) allow for a natural interpretation in terms of “influence”, and that interventions may be defined by substitutions in the differential equations. In this paper, we make these ideas precise. Our main contributions are:

1. For a given SDE, we give a precise definition of the postintervention SDE resulting from an intervention.
2. We show that under certain regularity assumptions, the solution of the postintervention SDE is the limit of a sequence of interventions in structural equation models based on the Euler scheme of the observational SDE.
3. We prove using (2) that for SDEs with a Lévy process as the driving semimartingale, the postintervention distribution is identifiable from the generator associated with the SDE.

The definition (1) yields a generic notion of intervention effects for SDEs applicable to causal inference in the case where an understanding of the mechanisms of the system under consideration is absent. The results of point (2) clarifies when we may expect this generic notion to be applicable.

The result (3) is stated as Theorem 5.3 and is the main theorem of this paper. Its importance is as follows. In classical DAG-based models of causality such as developed in [26], neither the DAG nor the effect of interventions can be uniquely identified from the observational distribution. This is one of the main difficulties of such causal models, and leads to a rich and challenging theory for partially identifying intervention effects, see for example [23] and the references therein. Theorem 5.3 essentially shows that for Lévy driven SDE models, the effect of interventions can be uniquely identified from the observational distribution, meaning that the intervention effect identification problem present in classical DAG-based models vanishes for these SDE models.

We expect that this result will have considerable applicability for causal inference for time-dependent observations. As argued in the series of examples comprised by Example 2.2, Example 2.5 and Example 5.6, our results for example lead to a dynamic modeling framework where gene knockout effects can be derived from observational data – a difficult problem which has previously only been dealt with, [22, 23], using non-dynamic methods.

Of further particular note is that the identifiability result (3) in the list above corresponds to a case where the error variables are not all independent, as is otherwise often assumed to be the case when calculating intervention effects in the DAG-based framework. For the DAG-based framework, in the case of independent errors, parts of the causal structure may be learned from the observational distribution, as seen in [36], and intervention distributions may be calculated by a truncated factorization formula as in (3.10) of [26]. For dependent errors, such results are harder to come by. In our case, we essentially take advantage of the Markov nature of the solutions to SDEs with

Lévy noise in order to obtain our identifiability result for SDE models, and we are also able to obtain explicit descriptions of the resulting postintervention distributions.

In matters of causality, it is important to distinguish clearly between definitions, theorems and interpretations. Our definition of postintervention SDEs will be a purely mathematical construct. It will, however, have a natural interpretation in terms of causality. Given an SDE model, in order to use the definition of postintervention SDEs given here to predict the effects of real-world interventions, it is necessary that the SDE can be sensibly interpreted as a data-generating mechanism with certain properties: Specifically, as we will argue in Section 4, it is essentially sufficient that the driving semimartingales are autonomous in the sense that they may be assumed not to be directly affected by interventions. This is an assumption which is not testable from a statistical viewpoint. It is, nonetheless, an assumption which may be justified by other means in concrete cases.

The remainder of the paper is organized as follows. In Section 2, we motivate and introduce our notion of intervention for SDEs. In Section 3, we review the terminology of causal inference as developed in [26] and [35], based on structural equation models and directed acyclic graphs. Section 4 shows that under certain conditions, our notion of intervention is equivalent to taking a limit of interventions in the context of structural equation models based on the Euler scheme of the SDE. In Section 5, we give conditions for postintervention distributions to be identifiable from the generator of the SDE. Finally, in Section 6, we discuss our results. Appendix A contains proofs.

2 Interventions for stochastic differential equations

In this section, given an SDE, we define the notion of a postintervention SDE, interpreted as the result of an intervention in a system described by an SDE. This notion yields a causal interpretation of stochastic differential equations.

We begin by considering three examples. Example 2.1 is a classical stochastic control problem. The control over a stochastic process is achieved via a control variable, whose effect on the stochastic system is a part of the model assumptions. Such an assumption is an (implicit) assumption about a causal relationship, or at least about how interventions in the system affect the system. Though the assumption is plausible in the specific example, we want to bring attention to its existence. Example 2.2 discusses a case where our stochastic model, due to the current state of knowledge in the subject matter field, cannot be derived completely from background mechanisms of the system under consideration. It is, however, highly desirable to be able to model and discuss causality and the effect of interventions in this situation. Finally, Example 2.3 provides an example where an understanding of the background mechanisms of a system provides an SDE model and also provides a candidate for how to describe the effects of interventions in the system.

Example 2.1. Consider the following simplified variant of Merton's portfolio selection problem, first formulated in [24]. In this problem, we consider the Black-Scholes model for a financial market in continuous time, consisting of a risk-free asset with price process B and a risky asset with price process S , following the SDEs

$$dB_t = rB_t dt, \quad (2.1)$$

$$dS_t = \mu S_t dt + \sigma S_t dW_t. \quad (2.2)$$

Here, r denotes the risk-free interest rate, μ is the expected return of the risky asset, and σ is the volatility of the risky asset. Now consider an investor endowed with initial wealth V , who invests a constant fraction α of his wealth at time t in the risky asset S and holds the remaining fraction $1 - \alpha$ of his wealth in the risk-free asset B .

Now, as the investor at time t invests $(1 - \alpha)V_t$ in the risk-free asset, yielding ownership of $(1 - \alpha)V_t/B_t$ units of this asset, and invests αV_t in the risky asset, yielding ownership of $\alpha V_t/S_t$ units of this asset, the arguments in Chapter 6 of [4] yield that V satisfies

$$\begin{aligned} dV_t &= (1 - \alpha)(V_t/B_t) dB_t + \alpha(V_t/S_t) dS_t \\ &= (1 - \alpha)V_t r dt + \alpha V_t \mu dt + \alpha V_t \sigma dW_t \\ &= ((r + \alpha(\mu - r))V_t) dt + \alpha V_t \sigma dW_t. \end{aligned} \quad (2.3)$$

In [24], Merton endows the investor with a utility function u , meaning that the utility for the investor of having wealth v is $u(v)$ and proceeds to solve the problem of identifying the portfolio (how α should be dynamically chosen), which optimizes the lifetime value of the portfolio over $[0, T]$, given by

$$Ee^{-rT}u(V_T), \quad (2.4)$$

subject to the constraint that $V_t > 0$. The optimal (Markov) control $\alpha(t, V_t)$, which is a function of time and wealth, can generally be characterized as a solution to the Hamilton-Jacobi-Bellman equation, and for some special choices of utility functions an explicit analytic solution can be found.

Now notice the following subtle point. In the above, we have successfully formulated an optimal control problem, seeking an optimal portfolio for the investor. At no point did it become necessary to consider what the “causal effect” of a particular choice of portfolio on the wealth process is, as the general financial arguments of [4] provides for this: A change of portfolio causes a change in the wealth process, while the opposite is a somewhat insensible statement without a specified control process. This is an example of how, when we have background knowledge of the effects of real-world choices (such as the choice of portfolio) on terms of interest (the wealth of the investor), the causal effects of choices, or interventions, are determined by our background knowledge. In all these arguments there is a hidden assumption, namely that the choice of portfolio doesn’t affect the Brownian motion that drives the price process. For small investors this may be a reasonable assumptions, but it is well known that large investors can affect the price process by their investments. Thus in this classical control problem there are assumptions about how the control variable affects the system, and this includes the assumption that the process driving the SDE is unaffected by the control variable – a notion we later refer to as autonomy of the driving process. ◦

Example 2.2. In this example we discuss the modeling of gene expression in the yeast microorganism *Saccharomyces Cerevisiae*. The genome of this organism was the first eukaryotic genome to be completely sequenced, see [12]. In general, genes of an organism are not active at all times, nor are they simply active or not active. Instead, a gene has a level of expression, indicating the production rate of the protein corresponding to the gene. An important question in connection with genomic research is the understanding of how the expression level of one gene influences the expression level of other genes. An understanding of such causal networks would allow analysis of what interventions to make on gene expression, for example what genes to knock out (that is, turn permanently off) in order to achieve some particular aim, such as optimal growth of an organism or optimal production rate of a particular compound of interest.

For this particular microorganism, gene expression data are available, both for non-mutated specimens and for mutations corresponding to deletion of particular genes, see [13]. Inference of the effect of interventions based on gene expression levels of non-mutated specimens has been carried out in [22] using IDA (an acronym for “Intervention calculus when the DAG is absent”), see [23], and compared to intervention data resulting from deletion mutants with favorable results.

The method investigated in [22] is not based on a dynamic model of gene expression, but rather on a multivariate Gaussian model of cross sectional data. It suffers, for instance, from the inability to include feedback loops. As a simple alternative suppose that the $p = 5361$ genes of a non-mutant specimen of *S. Cerevisiae* evolves according to an Ornstein-Uhlenbeck process solving the SDE

$$dX_t = B(X_t - A) dt + \sigma dW_t, \tag{2.5}$$

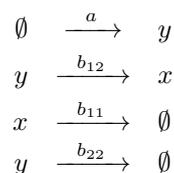
where B is a $p \times p$ matrix, A is a p -dimensional vector, σ is a $p \times d$ matrix and W is a d -dimensional Brownian motion. One benefit of such a model is its mean reversion properties, corresponding to gene expression levels fluctuating over time, but generally remaining stable over periods of the life of the specimen. Depending on the data available, standard statistical methods may then be applied to obtain estimates of some or all of the parameters of the model, yielding a description of the distribution of our data.

As discussed above, the effect of knocking out gene m (corresponding to setting X^m to zero for some m) in the model (2.5) is of central importance. However, as we in this case do not have a sufficiently detailed biochemical understanding of how genes influence each other over time, it is less obvious than in Example 2.1 how the knockout intervention of gene m affects the system.

In other words, our lack of a generic concept for causality for SDEs, applicable in the absence of knowledge of particular mechanisms of causality, in this case prevents us from considering intervention effects in our model. \circ

Example 2.3. Chemical kinetics is concerned with the evolution of the concentrations of chemicals over time, given in terms of a number of coupled chemical reactions, see [37]. In this example, we consider two chemicals and derive a simple system of SDEs from the fundamental mechanisms of the chemical reactions. If the concentration of one chemical is fixed (as an alternative to letting it evolve according to the chemical reactions) the fundamental mechanisms allow us to obtain an SDE for the concentrations of the remaining chemicals. This SDE then describes the system after the intervention, and can be obtained from the original system by a purely mechanical deletion and substitution process.

The chemicals are denoted x and y and the corresponding concentrations are denoted X and Y , respectively. We assume that four reactions are possible, namely:



Here, the first reaction denotes the creation or influx of chemical y with constant rate a , the second reaction denotes the change of y into x at rate $b_{12}Y$, and the third and fourth reactions denote degradation or outflux of x and y with rates $b_{11}X$ and $b_{22}Y$, respectively. We collect the rates into the vector

$$\lambda(X, Y) = \begin{pmatrix} a \\ b_{12}Y \\ b_{11}X \\ b_{22}Y \end{pmatrix}. \tag{2.6}$$

The so-called stoichiometric matrix

$$S = \begin{pmatrix} 0 & 1 & -1 & 0 \\ 1 & -1 & 0 & -1 \end{pmatrix} \tag{2.7}$$

collects the information about the number of molecules, for each of the two chemicals (rows), which are created or destroyed by each of the four reactions (columns). The rates $\lambda(X, Y)$ and the stoichiometric matrix S form the fundamental parameters of the system. We are interested in using $\lambda(X, Y)$ and S to construct a model for the evolution of X and Y over time.

Several different stochastic and deterministic models are available. One stochastic model is obtained by considering a Markov jump process on \mathbb{N}_0^2 , where each coordinate denotes the total number of molecules of each chemical x and y , and the transition rates are given in terms of S and $\lambda(X, Y)$. A system of SDEs approximating the Markov jump process, see [2], is given by

$$\begin{pmatrix} X_t \\ Y_t \end{pmatrix} = \begin{pmatrix} X_0 \\ Y_0 + at \end{pmatrix} + \int_0^t B \begin{pmatrix} X_s \\ Y_s \end{pmatrix} ds + \int_0^t \Sigma(X_s, Y_s) dW_s \quad (2.8)$$

where W_s denotes a four-dimensional Wiener process, and the matrices $\Sigma(x, y)$ and B are given by

$$\begin{aligned} \Sigma(x, y) &= S \text{diag} \sqrt{\lambda(x, y)} \\ &= \begin{pmatrix} 0 & \sqrt{b_{12}y} & -\sqrt{b_{11}x} & 0 \\ \sqrt{a} & -\sqrt{b_{12}y} & 0 & -\sqrt{b_{22}y} \end{pmatrix} \end{aligned} \quad (2.9)$$

and

$$B = \begin{pmatrix} -b_{11} & b_{12} \\ -b_{12} & -b_{22} \end{pmatrix}. \quad (2.10)$$

If we are able to fix the concentration Y_t at a level ζ , we effectively remove the first and last of the reactions and the second will have the constant rate $b_{12}\zeta$. By arguments as above we then derive the SDE

$$X_t = X_0 + tb_{12}\zeta - \int_0^t b_{11}X_s ds + \int_0^t \sigma(X_s) d\widetilde{W}_s, \quad (2.11)$$

with \widetilde{W}_s a two-dimensional Wiener process and $\sigma(x) = (\sqrt{b_{12}\zeta}, -\sqrt{b_{11}x})$. We observe that this SDE, describing the intervened system, can be obtained from (2.8) by deleting the equation for Y_t and substituting ζ for Y_t in the remaining equation. \circ

We now proceed to our main definition. Recall that in Example 2.2, we were stopped short in our discussion of the effect of interventions in our model due to the lack of a generic notion of interventions for SDEs. We will now use the conclusions from Example 2.3 to introduce such a generic notion of interventions.

In the DAG-based framework, the DAG is a direct representation of the causal structure of the system. We do not directly provide such a representation of causality for SDEs. In general, the precise meaning of ‘‘causation’’ is a point of contemporary debate, see for example [7]. For our purposes, it suffices to take a practical standpoint: The causal structure of a system is sufficiently elucidated for the purposes of our discussion if we know the effects of making interventions in the system. For this reason, we restrict ourselves in Definition 2.4 to defining the effect of making interventions.

In Example 2.3, we obtained results on the effects of intervention in a system from a model for the entire system. In this particular example, the resulting model for the intervention was justified by reference to the fundamental mechanisms (the chemical reactions) driving the system, and interventions resulted in SDEs modified by substitution and deletion. While noting that this correspondence between interventions and substitution and deletion in the original equations may not always be justified, we will use this principle as a general, purely mathematical definition of interventions in SDEs.

Consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ satisfying the usual conditions, see [30] for the definition of this and other notions related to continuous-time stochastic processes. In order to formalize our definition in a general framework, let Z be a d -dimensional semimartingale and assume that $a : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$ is a continuous mapping, where $\mathbb{M}(p, d)$ denotes the space of real $p \times d$ matrices. We consider the stochastic differential equation

$$X_t^i = X_0^i + \sum_{j=1}^d \int_0^t a_{ij}(X_{s-}) dZ_s^j, \quad i \leq p. \tag{2.12}$$

This SDE is written in integral form. Using differential and matrix notation, (2.12) corresponds to the SDE $dX_t = a(X_{t-}) dZ_t$ with initial condition X_0 . In the following, x^{-m} denotes the $(p - 1)$ -dimensional vector where the m 'th coordinate of $x \in \mathbb{R}^p$ has been removed.

Definition 2.4. Consider some $m \leq p$ and $\zeta : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$. The stochastic differential equation arising from (2.12) under the intervention $X_t^m := \zeta(X_t^{-m})$ is the $(p - 1)$ -dimensional equation

$$(Y^{-m})_t^i = X_0^i + \sum_{j=1}^d \int_0^t b_{ij}(Y_{s-}^{-m}) dZ_s^j, \quad i \neq m, \tag{2.13}$$

where $b : \mathbb{R}^{p-1} \rightarrow \mathbb{M}(p - 1, d)$ is defined by $b_{ij}(y) = a_{ij}(y_1, \dots, \zeta(y), \dots, y_p)$ for $i \neq m$ and $j \leq d$ and the $\zeta(y)$ is on the m 'th coordinate.

By Definition 2.4, intervening takes a p -dimensional SDE as its argument and yields a $(p - 1)$ -dimensional SDE as its result. Note that existence and uniqueness of solutions are not required for Definition 2.4 to make sense, although we will mainly take interest in cases where both (2.12) and (2.13) have unique solutions. By Theorem V.7 of [30], this is for example the case whenever the mappings a and ζ are Lipschitz.

We stress that while Definition 2.4 is motivated by actual results from Example 2.3, we do not claim that it universally describes the effects of actual interventions in a system. The discussion in Section 4 gives indications for *whether* Definition 2.4 properly describes causality for a particular SDE system. Our other results, such as those of Section 5, are devoted to analyze the consequences *if* Definition 2.4 is a valid description of the effect of interventions (and thus also a valid description of the causal structure of the system, since knowing the effects of interventions yields causal information about the system).

As discussed in Example 2.2, an intervention with a constant function ζ is of some interest, and in the context of gene expression a knockout intervention, corresponding to $\zeta(y) = 0$, is one of the only control mechanisms currently possible. If ζ is a constant we identify the function with this constant, and we write $X_t^m := \zeta$ for the intervention that puts the m 'th coordinate constantly equal to ζ .

Also note that the process Y^{-m} above for which the SDE is formulated is a $(p - 1)$ -dimensional process indexed by $\{1, \dots, p\} \setminus \{m\}$. When Y^{-m} is a solution to (2.13), we also define $Y_t^m = \zeta(Y_t^{-m})$, and the p -dimensional process Y is then the full result of making the intervention $X^m := \zeta(X_t^{-m})$. The process Y^{-m} is simply Y with its m 'th coordinate removed. In general, the p -dimensional process Y will not satisfy any p -dimensional SDE except in special cases. One such special case is when ζ is constant. In this case Y will satisfy the p -dimensional SDE

$$Y_t^i = Y_0^i + \sum_{j=1}^d \int_0^t c_{ij}(Y_{s-}) dZ_s^j, \quad i \leq p, \tag{2.14}$$

where $Y_0^i = X_0^i$ for $i \neq m$ and $Y_0^m = \zeta$, and $c : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$ is given by letting $c_{ij}(x) = a_{ij}(x)$ for $i \neq m$ and $c_{mj}(x) = 0$ for all $x \in \mathbb{R}^p$ and $j \leq d$.

Assuming that (2.12) and (2.13) have unique solutions for all interventions, we refer to (2.12) as the observational SDE, to the solution of (2.12) as the observational process, and to the distribution of the solution of (2.12) as the observational distribution. We refer to (2.13) as the postintervention SDE, to the solution of (2.13) as the postintervention process and to the distribution of the solution to (2.13) as the postintervention distribution. Note how our definition of the postintervention SDE has the same structure as the SDE obtained in Example 2.3 by reference to fundamental mechanisms.

As a first application of Definition 2.4, we show in Example 2.5 that by Definition 2.4, intervention with constant functions in an Ornstein-Uhlenbeck process yields another Ornstein-Uhlenbeck process. Recalling Example 2.2, we thus find that if Definition 2.4 is applicable in the SDE model of Example 2.2, and if we can identify the correct parameters of the SDE, then we can reason about the effects of interventions.

Example 2.5. Let $x_0 \in \mathbb{R}^p$, $A \in \mathbb{R}^p$, $B \in \mathbb{M}(p, p)$ and $\sigma \in \mathbb{M}(p, d)$. The Ornstein-Uhlenbeck SDE with initial value X_0 , mean reversion level A , mean reversion speed B , diffusion matrix σ and d -dimensional driving noise is

$$X_t = X_0 + \int_0^t B(X_s - A) ds + \sigma W_t, \tag{2.15}$$

where W is a d -dimensional (\mathcal{F}_t) Brownian motion, see Section II.72 of [31]. Fix $m \leq p$ and $\zeta \in \mathbb{R}$. Under the intervention $X^m := \zeta$, we obtain that the postintervention process satisfies

$$Y_t^i = X_0^i + \int_0^t \sum_{j \neq m}^p B_{ij}(Y_s^j - A_j) + B_{im}(\zeta - A_m) ds + \sum_{j=1}^d \sigma_{ij} W_t^j. \tag{2.16}$$

for $i \neq m$. Now let \tilde{B} be the submatrix of B obtained by removing the m 'th row and column of B , and assume that \tilde{B} is invertible. We then obtain

$$Y_t^{-m} = X_0^{-m} + \int_0^t \tilde{B}(Y_s^{-m} - \tilde{A}) ds + \tilde{\sigma} W_t, \tag{2.17}$$

where $\tilde{\sigma}$ is obtained by removing the m 'th row of σ and $\tilde{A} = \alpha - \tilde{B}^{-1}\beta$, where α and β are obtained by removing the m 'th coordinate from A and from the vector whose i 'th component is $B_{im}(\zeta - A_m)$, respectively. Thus, Y^{-m} solves an $(p - 1)$ -dimensional Ornstein-Uhlenbeck SDE with initial value X_0^{-m} , mean reversion level \tilde{A} , mean reversion speed \tilde{B} and diffusion matrix $\tilde{\sigma}$. ◦

Now note that for the SDE

$$dX_t = B(X_t - A) dt + \sigma dW_t, \tag{2.18}$$

considered in Example 2.2, the solution distribution depends only on σ through $\sigma\sigma^t$. Therefore, the parameters of the SDE are not uniquely identifiable from the observational distribution. As we thus cannot identify the parameters of the SDE, it appears that we cannot identify the postintervention SDE in Definition 2.4. In Example 5.6, we show how to use our main theorem, Theorem 5.3, on identifiability of postintervention distributions to circumvent this problem. Though we cannot identify the postintervention SDE, we can in fact identify the postintervention distribution.

Note also that in Example 2.3, the matrix

$$\Sigma(X, Y)\Sigma(X, Y)^t = \begin{pmatrix} b_{12}Y + b_{11}X & -b_{12}Y \\ -b_{12}Y & a + b_{12}Y + b_{22}Y \end{pmatrix} \tag{2.19}$$

is not diagonal, implying that the martingale parts of the semimartingale (X, Y) are not orthogonal. This shows that there are naturally occurring situations where it is necessary to consider models with non-orthogonal martingale parts. This is a situation excluded in the WCLI framework of [16] and is a motivating factor for the level of generality in our definition.

3 Terminology of SEMs, DAGs and interventions

In this section, we review the basic notions related to intervention calculus for structural equation models (SEMs). For a detailed overview, see [26, 35]. We will use these notions in Section 4 to interpret our definition of intervention for SDEs in terms of intervention calculus for structural equation models.

As remarked above, Definition 2.4 takes an SDE as an argument and yields another SDE, in contrast to, for example, taking the distribution of an SDE, and yielding another distribution. This corresponds to how an intervention in the framework of SEMs, see [26], takes a SEM and returns another SEM, instead of taking a distribution and yielding another distribution. This is a key point, and allows us in Section 4 to use SEMs and DAGs to interpret Definition 2.4, and view SDEs as a natural extension of SEMs to continuous time models.

Let V be a finite set, and let E be a subset of $V \times V$. A directed graph G on V is a pair (V, E) . We refer to V as the vertex set, and refer to E as the edge set. Note that by this definition, there can be at most one edge between any pair of vertices. A path is an unbroken series of vertices and edges such that no vertices are repeated except possibly the initial and terminal vertices. A directed cycle is a path with the same initial and terminal vertices and all arrows pointing in the same direction. We say that G is an acyclic directed graph (DAG) if G contains no directed cycles. Note that this in particular excludes that the graph contains an edge with the same initial and terminal vertex. For any graph G and $i \in V$, we write $\text{pa}(i) = \{j \in V \mid (j, i) \in E\}$, and refer to $\text{pa}(i)$ as the parents of the vertex i . If we wish to emphasise the graph G , we also write $\text{pa}_G(i)$.

A structural equation model (SEM) consists of three components:

1. Two families $(X_i)_{i \in V}$ and $(U_i)_{i \in V}$ of random variables.
2. A directed acyclic graph G on V .
3. A set of functional relationships $X_i = f_i(X_{\text{pa}_G(i)}, U_i)$.

We refer to $(X_i)_{i \in V}$ as the primary variables and $(U_i)_{i \in V}$ as the noise variables. Note that we do not a priori assume that the noise variables are independent. The idea behind a SEM is that the DAG provides the sequence in which the functional relationships are evaluated, thus yielding an algorithm for obtaining the values of $(X_i)_{i \in V}$ from $(U_i)_{i \in V}$. A SEM does not only yield the distribution of the variables $(X_i)_{i \in V}$, but also a description of a data generating mechanism. This is made precise by the notion of an intervention, see Definition 3.2.1 of [26]. Chapter 3 of [26] discusses interventions where a subset of variables are set to a constant value. We will need to consider a more general type of interventions where variables are set to values depending on other variables. Therefore, our definition below extends Definition 3.2.1 of [26]. See Chapter 4 of [26] for more on this type of interventions.

Definition 3.1. Consider given a SEM with primary variables $(X_i)_{i \in V}$, noise variables $(U_i)_{i \in V}$, DAG G and functional relationships $X_i = f_i(X_{\text{pa}_G(i)}, U_i)$. Let A be a subset of V , and for $i \in A$ let $I(i) \subseteq V \setminus A$ and $\zeta_i(X_{I(i)})$ be a function of the primary variables with indices in $I(i)$. We form a new graph G' by replacing $\text{pa}_G(i)$ with $I(i)$ for $i \in A$. We assume that G' is a DAG. The postintervention SEM obtained by doing $X_i := \zeta_i(X_{I(i)})$

for $i \in A$ is a SEM with primary variables $(X_i)_{i \in V}$, noise variables $(U_i)_{i \in V}$, DAG G' and functional relationships obtained by substituting all occurrences of X_i by $\zeta_i(X_{I(i)})$ for $i \in A$.

In short, Definition 3.1 describes the effect of intervening and setting X_i for $i \in A$ to be a function of certain variables in $V \setminus A$. In the case where the ζ_i are constant, this reduces to Definition 3.2.1 of [26].

4 Interpretation of postintervention SDEs

In this section, we show that under Lipschitz conditions on the coefficients in (2.12) and the intervention mapping, the solution to the postintervention SDE described in Definition 2.4 essentially is the limit of a sequence of postintervention SEMs as described in Definition 3.1 based on the Euler scheme of (2.12). We use this to clarify the role of the driving semimartingales Z^1, \dots, Z^d . Also, we will use this result to prove the main theorem on identifiability in Section 5.

Definition 4.1. *The signature of the SDE (2.12) is the graph S with vertex set $\{1, \dots, n\}$ and an edge from i to j if it holds that there is k such that the mapping a_{jk} is not independent of the i 'th coordinate.*

Letting $a_{j \cdot} = (a_{j1}, \dots, a_{jd})$, another way of describing the signature S in Definition 4.1 is that there is an edge from i to j if $x_i \mapsto a_{j \cdot}(x)$ is not constant, or equivalently, there is no edge from i to j if it holds for all k that a_{jk} does not depend on the i 'th coordinate.

Definition 4.2. *We say that X^j is locally unaffected by X^i in the SDE (2.12) if there is no edge from i to j in the signature of (2.12).*

Being locally unaffected is a property of two coordinates of an SDE. If there is no risk of ambiguity, we leave out the SDE and simply state that X^j is locally unaffected by X^i .

The signature is used in the following definition to define a SEM corresponding to the Euler scheme for (2.12). With a slight abuse of notation, we choose in Definition 4.3 for convenience to consider the initial variables X_0^1, \dots, X_0^p as primary variables, even though these variables have no associated noise variables in the SEM. This is not a problem as it is nonetheless clear how interventions for the SEM given in Definition 4.3 should be understood.

Definition 4.3. *Fix $T > 0$ and consider $\Delta > 0$ such that T/Δ is a natural number. Let $N = T/\Delta$ and $t_k = k\Delta$. The Euler SEM over $[0, T]$ with step size Δ for (2.12) consists of the following:*

1. *The primary variables are the $p(N + 1)$ variables in the set $(X_{t_k}^\Delta)_{0 \leq k \leq N}$, indexed by $\{0, \dots, N\} \times \{1, \dots, p\}$.*
2. *For $1 \leq k \leq N$, the noise variable for the i 'th coordinate of $X_{t_k}^\Delta$ is the d -dimensional variable $Z_{t_k} - Z_{t_{k-1}}$.*
3. *The DAG is the graph $G = (V, E)$ with vertex set $\{0, \dots, N\} \times \{1, \dots, p\}$ defined by having $((i_1, j_1), (i_2, j_2))$ be an edge of D if and only if $i_2 = i_1 + 1$ and either $j_2 = j_1$ or (j_1, j_2) is an edge in the signature of (2.12).*
4. *The functional relationships are given by:*

$$(X_{t_k}^\Delta)^i = (X_{t_{k-1}}^\Delta)^i + \sum_{j=1}^d a_{ij}(X_{t_{k-1}}^\Delta)(Z_{t_k}^j - Z_{t_{k-1}}^j). \tag{4.1}$$

A visualization of the DAG for the SEM of Definition 4.3 is shown in Figure 1. The figure shows how the signature S determines the DAG describing the algorithm for

calculating the variables in the Euler SEMs. Making the constant intervention $X_{t_k}^1 := \zeta$ for all k corresponds to removing the top row in Figure 1.

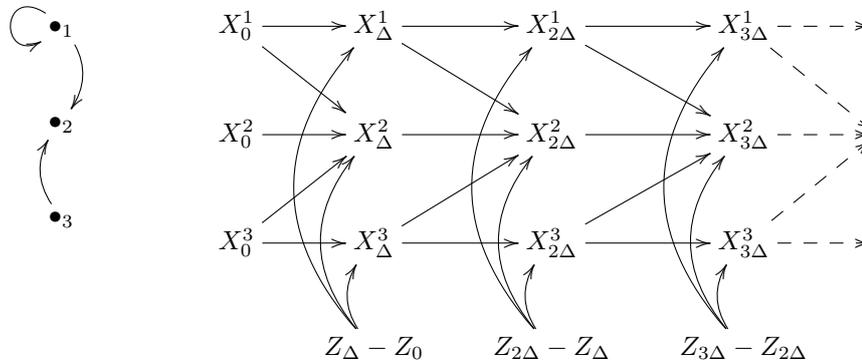


Figure 1: The signature for a three-dimensional SDE (left) and the DAG for the corresponding Euler SEM (right).

Combining the following two lemmas yields the main result of this section.

Lemma 4.4. Assume that $a : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$ is Lipschitz. Fix $T > 0$ and let $(\Delta_n)_{n \geq 1}$ be a sequence of positive numbers converging to zero such that T/Δ_n is natural for all $n \geq 1$. For each n , there exists a pathwisely unique solution to the equation

$$(X_t^n)^i = X_0^i + \sum_{j=1}^d \int_0^t a_{ij}(X_{\eta_n(s-)}^n) dZ_s^j, \quad i \leq p, \quad (4.2)$$

where $\eta_n(t) = k\Delta_n$ for $k\Delta_n \leq t < (k+1)\Delta_n$, satisfying that $((X^n)_{t_k})_{0 \leq k \leq T/\Delta_n}$ are the primary variables in the Euler SEM for (2.12), and $\sup_{0 \leq t \leq T} |X_t - X_t^n|$ converges in probability to zero, where X is the solution to (2.12).

Proof. By inspection, (4.2) has a unique solution, and $((X^n)_{t_k})_{k \leq T/\Delta_n}$ is the primary variables in the Euler SEM for (2.12). That $\sup_{0 \leq t \leq T} |X_t - X_t^n|$ converges in probability to zero is the corollary to Theorem V.16 of [30]. \square

Lemma 4.5. Fix $T > 0$ and consider $\Delta > 0$ such that T/Δ is a natural number. Fix $m \leq p$ and $\zeta : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$. The Euler SEM for the stochastic differential equation (2.13) is equal to the result of removing the m 'th coordinate of the postintervention SEM obtained by the intervention $(X_{t_k}^\Delta)^m := \zeta((X_{t_{k-1}}^\Delta)^{-m})$ for $0 \leq k \leq T/\Delta$ in the Euler SEM for (2.12).

Proof. The functional relationships in the Euler SEM for (2.12) are

$$(X_{t_k}^\Delta)^i = (X_{t_{k-1}}^\Delta)^i + \sum_{j=1}^d a_{ij}(X_{t_{k-1}}^\Delta)(Z_{t_k}^j - Z_{t_{k-1}}^j), \quad (4.3)$$

while for (2.13) and $i \neq m$, they are

$$\begin{aligned} (Y_{t_k}^\Delta)^i &= (Y_{t_{k-1}}^\Delta)^i + \sum_{j=1}^d b_{ij}((Y_{t_{k-1}}^\Delta)^{-m})(Z_{t_k}^j - Z_{t_{k-1}}^j) \\ &= (Y_{t_{k-1}}^\Delta)^i + \sum_{j=1}^d a_{ij}(Y_{t_{k-1}}^\Delta)(Z_{t_k}^j - Z_{t_{k-1}}^j), \end{aligned} \quad (4.4)$$

where $(Y_{t_k}^\Delta)^m = \zeta((Y_{t_{k-1}}^\Delta)^{-m})$. By inspection, (4.4) is the result of the stated intervention in the Euler SEM according to Definition 3.1. \square

Together, Lemma 4.4 and Lemma 4.5 states that the diagram in Figure 2 commutes: Defining interventions directly in terms of changing the terms in the stochastic differential equation has the same effect as intervening in the Euler SEM and taking the limit.

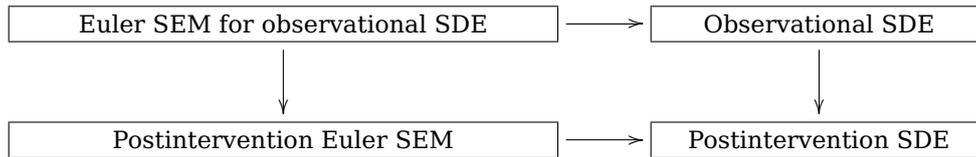


Figure 2: The interpretation of intervention in a stochastic differential equation understood as the limit of interventions in the Euler SEMs.

These results clarify what Definition 2.4 means, and in particular, when this generic definition of intervention is applicable when background mechanisms are unknown, such as Example 2.2. The intuition behind Definition 2.4 is that interventions are assumed not to influence the semimartingale Z directly. This is made concrete by assuming that the family $(Z_{t_k} - Z_{t_{k-1}})_{k \leq N}$ are the noise variables in the Euler SEM, such that there are no arrows in the DAG for the SEM with terminal vertices in $(Z_{t_k} - Z_{t_{k-1}})_{k \leq N}$. The lemmas show that when this condition holds true, the notion of intervention given in Definition 2.4 is consistent with the result of intervention in the Euler SEM. Note that this does not constitute a proof of causality. Rather, it gives guidelines as to when it is reasonable to expect that our notion of intervention will reflect real-world interventions: namely, when none of the coordinates X^i have a direct effect on the driving semimartingale Z . Whether this is the case or not is in general not a testable assumption.

Furthermore, note that the arrows across columns in the Euler SEM is determined by the signature of the SDE. Therefore, if we accept the hypothesis that the DAG of the Euler SEM describes the causal links between the coordinates of the SDE, then the signature S describes which coordinates of the SDE (2.12) are causally dependent on each other in an infinitesimal sense. Also note that as we are not using the Euler SEMs to draw any conclusions about the distribution of the variables, we do not require independence of the noise variables $(Z_{t_k} - Z_{t_{k-1}})_{k \leq N}$. In particular, the variables in the Euler SEM do not need to be Markov with respect to the DAG in the sense of [26].

Concluding this section, we give an example to illustrate that the notion of intervention given in Definition 2.4, and the corresponding causal interpretation outlined above, may not always be applicable.

Example 4.6. Let $X^1 = W$ be a one-dimensional Wiener process, consider a twice continuously differentiable function $f : \mathbb{R} \rightarrow \mathbb{R}$ and assume that for all $t \geq 0$, it holds that

$$X_t^2 = f(X_t^1). \tag{4.5}$$

We now make the following assumption: Assume that (4.5) represents the actual causal relationship between X^1 and X^2 , in the sense that the result on X^2 of the intervention $X^1 := \zeta$ is the process

$$X_t^2 = f(\zeta). \tag{4.6}$$

Now, by Itô's lemma, it holds that

$$\begin{aligned} X_t^2 &= f(X_0^1) + \frac{1}{2} \int_0^t f''(X_s^1) d[X^1]_s + \int_0^t f'(X_s^1) dX_s^1 \\ &= f(0) + \frac{1}{2} \int_0^t f''(X_s^1) ds + \int_0^t f'(X_s^1) dW_s, \end{aligned} \tag{4.7}$$

such that (X^1, X^2) satisfies

$$X_t^1 = \int_0^t dW_s \tag{4.8}$$

$$X_t^2 = f(0) + \frac{1}{2} \int_0^t f''(X_s^1) ds + \int_0^t f'(X_s^1) dW_s, \tag{4.9}$$

which together yields a two-dimensional SDE of the form given in (2.12). Therefore, we may apply Definition 2.4 to this SDE. The resulting postintervention SDE for X^2 under the intervention $X^1 := \zeta$ is

$$X_t^2 = f(0) + \frac{1}{2} \int_0^t f''(\zeta) ds + \int_0^t f'(\zeta) dW_s, \tag{4.10}$$

which yields the result $X_t^2 = f(0) + \frac{1}{2} f''(\zeta)t + f'(\zeta)W_t$, in contradiction with our assumed result in (4.6), $X_t^2 = f(\zeta)$. This shows that we may conceptualize ideas about the effects of interventions which are rather natural, but which are not captured by Definition 2.4. This illustrates the importance of the conclusions made above: We can only argue under certain circumstances that Definition 2.4 is a reasonable description of the effects of intervention. \circ

We note that in Example 4.6, it is not the use of Itô's lemma which yields the discrepancy between the results of Definition 2.4 and the assumed result, (4.6). Rather, it is the subsequent substitution of X^1 by W . In fact, if we intervene directly in (4.7) by replacing X^1 by the constant ζ , the result would be that $X_t^2 = f(\zeta)$, in accordance with (4.6). However, Definition 2.4 does not allow for such interventions on the integrators. To do so generally would complicate matters, and we will not pursue this any further.

5 Identifiability of postintervention distributions

In this section we formulate a result, Theorem 5.3, giving conditions for the postintervention distributions to be determined by uniquely identifiable aspects of the SDE. We show that if the SDE is driven by a Lévy process, the postintervention distribution is determined by the generator.

To introduce the generator associated with the SDE (2.12), when it is driven by a Lévy process, we need to introduce Lévy triplets. A Lévy measure on \mathbb{R}^d is a measure ν assigning zero measure to $\{0\}$ such that $x \mapsto \min\{1, \|x\|^2\}$ is integrable with respect to ν . A d -dimensional Lévy triplet is a triplet (α, C, ν) , where α is an element of \mathbb{R}^d , C is a positive semidefinite $d \times d$ matrix and ν is a Lévy measure on \mathbb{R}^d . Recall that for any bounded neighborhood D of zero in \mathbb{R}^d and any d -dimensional Lévy process X , there is a Lévy triplet (α, C, ν) such that

$$Ee^{iu^t X_1} = \exp\left(iu^t \alpha - \frac{1}{2} u^t C u - \int_{\mathbb{R}^d} e^{iu^t x} - 1 - iu^t x 1_D(x) d\nu(x)\right), \tag{5.1}$$

and this triplet uniquely determines the distribution of X , see Theorem 1.2.14 of [3]. We refer to (α, C, ν) as the characteristics of X with respect to D , or as the D -characteristic

triplet of X . Conversely, for any bounded neighborhood D of zero in \mathbb{R}^d and any Lévy triplet (α, C, ν) , there exists a Lévy process having (α, C, ν) as its D -characteristic triplet.

The generator of (2.12) is defined as a linear operator on the set $C_0^2(\mathbb{R}^p)$ of twice continuously differentiable functions such that the function itself together with all its first and second partial derivatives vanish at infinity.

Definition 5.1. Let D be a bounded neighborhood of zero in \mathbb{R}^d . Consider the SDE (2.12), where Z is a d -dimensional Lévy process with D -characteristic triplet (α, C, ν) and $a : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$. We define the generator A of (2.12) on $C_0^2(\mathbb{R}^p)$ by

$$\begin{aligned}
 Af(x) = & \sum_{i=1}^p \sum_{j=1}^d a_{ij}(x) \alpha_j \frac{\partial f}{\partial x_i}(x) + \frac{1}{2} \sum_{i=1}^p \sum_{j=1}^p (a(x)Ca(x)^t)_{ij} \frac{\partial^2 f}{\partial x_i \partial x_j}(x) \\
 & + \int_{\mathbb{R}^p} f(x + a(x)y) - f(x) - 1_D(y) \sum_{i=1}^p \frac{\partial f}{\partial x_i}(x) \sum_{j=1}^d a_{ij}(x)y_j \, d\nu(y) \quad (5.2)
 \end{aligned}$$

for $f \in C_0^2(\mathbb{R}^p)$ and $x \in \mathbb{R}^p$.

It holds that for any choice of a that the generator A is well defined on $C_0^2(\mathbb{R}^p)$ with values in the set of functions on \mathbb{R}^p . If we are willing to put restrictions on a , the range of the generator can be restricted as well.

The interest in the generator stems from the fact that when Z is a Lévy process, the generator of (2.12) can usually be determined by the semigroup of transition probabilities for the Markov process that solves (2.12). We state one such result here. Lemma 5.2 is a folklore result, and follows from the results in Chapter 6 of [3].

Lemma 5.2. If Z is a Lévy process and $a : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$ is Lipschitz and bounded then there exists a unique Feller semigroup (P_t) with the property that all solutions of (2.12) are Feller processes with semigroup (P_t) . Moreover, the generator A of (2.12) satisfies that

$$Af = \lim_{t \rightarrow 0} t^{-1}(P_t f - P_0 f) \quad (5.3)$$

for $f \in C_0^2(\mathbb{R}^p)$, where convergence is in the uniform norm on $C_0^2(\mathbb{R}^p)$.

For a treatment of the theory of Markov processes and Lévy processes, and in particular for notions such as Feller processes, Feller semigroups, generators, Lévy processes and so forth, see [14, 3, 34]. We are now ready to state our main result on identifiability.

Theorem 5.3. Consider the SDEs

$$X_t^i = X_0^i + \sum_{j=1}^d \int_0^t a_{ij}(X_s) \, dZ_s^j, \quad i \leq p, \quad (5.4)$$

and

$$\tilde{X}_t^i = \tilde{X}_0^i + \sum_{j=1}^d \int_0^t \tilde{a}_{ij}(\tilde{X}_s) \, d\tilde{Z}_s^j, \quad i \leq p, \quad (5.5)$$

where Z is a d -dimensional Lévy process and \tilde{Z} is a \tilde{d} -dimensional Lévy process. Assume that (5.4) and (5.5) have the same generator, that $a : \mathbb{R}^p \rightarrow \mathbb{M}(p, d)$ and $\zeta : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$ are Lipschitz and that the initial values have the same distribution. Then the postintervention distributions of doing $X^m := \zeta(X^{-m})$ in (5.4) and doing $\tilde{X}^m := \zeta(\tilde{X}^{-m})$ in (5.5) are equal for any choice of ζ and m .

Theorem 5.3 is proven in Appendix A. Theorem 5.3 states that for SDEs with a Lévy process as the driving semimartingale, postintervention distributions are identifiable from the generator. In the remainder of this section, we discuss the content of Theorem 5.3.

First, recall that a main theme of the DAG-based framework for causal inference as in [35, 26] is to identify conditions for when postintervention distributions are identifiable from the observational distribution. Theorem 5.3 gives a criterion for when postintervention distributions are identifiable from the generator of the SDE, which is not exactly the same. Nonetheless, in a large family of naturally occurring cases, the semigroup is identifiable from the observational distribution. This is for example the case if the solutions to (5.4) and (5.5) are irreducible, as the family of transition probabilities in this case will be identifiable from the observational distribution, allowing us to obtain the generator through Lemma 5.2.

Next, we comment on the relationship between the result of Theorem 5.3 and identifiability results of DAG-based causal inference. Consider the Euler SEM of Definition 4.3, illustrated in Figure 1. In the DAG of this SEM, the orientation of all arrows is assumed known: All orientations for arrows from primary variables point forward in time. If the error variables for each primary variable were independent, it would hold that the distribution of the variables would be Markov with respect to the DAG in the sense of [26]. In this case, by the results of [36], we would be able to identify the skeleton of the graph (that is, its undirected edges) from the observational distribution. As all orientations are given, this leads to identifiability of the entire graph. Using the truncated factorization (3.10) of [26], this leads to identifiability of intervention distributions from the observational distribution. Thus, in this case, identifiability would not be a surprising result.

However, when the driving semimartingale Z is a Lévy process, the error variables are independent across time, but are not independent across coordinates: For each k , the variables $X_{\Delta k}^1, \dots, X_{\Delta k}^p$ have the same d -dimensional error variable, namely $Z_{\Delta k} - Z_{\Delta(k-1)}$, and so the Euler SEM illustrated in Figure 1 is not Markov with respect to its DAG. Therefore, our scenario differs from the conventional causal modeling scenario of [26] in two ways: Both by considering a continuous-time model with uncountably many variables and by considering a particular type of dependent errors.

We end the section with three examples. Example 5.4 considers a particularly simple scenario where identifiability of postintervention distributions can be seen explicitly from the transition probabilities. In Example 5.5, we show that it is possible for two SDEs with the same distribution to have different signatures. Remarkably, this shows that while postintervention distributions are identifiable by Theorem 5.3, the signature of the true SDE is not generally identifiable. However, we expect that the behaviour observed in Example 5.5 is atypical, similarly to the absence of faithfulness in Gaussian SEMs, see Theorem 3.2 of [35]. Finally, in Example 5.6, we show how Theorem 5.3 allows us to infer intervention effects of knocking out genes in our previous example on *S. Cerevisiae*, Example 2.2.

Example 5.4. Let W and \tilde{W} be d -dimensional and \tilde{d} -dimensional Brownian motions, let B and \tilde{B} be $p \times p$ matrices, and let σ and $\tilde{\sigma}$ be $p \times d$ and $p \times \tilde{d}$ matrices. Consider two processes X and Y being the unique solutions to the Ornstein-Uhlenbeck SDEs

$$X_t = X_0 + \int_0^t BX_t dt + \sigma W_t \tag{5.6}$$

and

$$Y_t = Y_0 + \int_0^t \tilde{B}X_t dt + \tilde{\sigma} W_t. \tag{5.7}$$

We will show by a direct analysis that if the generators of the SDEs are equal and the initial distributions are the same, then the postintervention distributions are equal as well. For notational simplicity, we consider intervening on the first coordinate, making the interventions $X^1 := \zeta$ and $Y^1 := \zeta$. It will suffice to show equality of distributions for the non-intervened coordinates in the postintervention distributions. Consider block decompositions of the form

$$B = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \quad \text{and} \quad \sigma = \begin{pmatrix} \sigma_1 \\ \sigma_2 \end{pmatrix}, \tag{5.8}$$

where B_{11} is a 1×1 matrix and B_{22} is a $(p - 1) \times (p - 1)$ matrix and σ_1 is a $1 \times d$ matrix and σ_2 is a $(p - 1) \times d$ matrix. Also consider corresponding decompositions of \tilde{B} and $\tilde{\sigma}$.

Assume that the generators of the SDEs are equal, and assume that X_0 and Y_0 have the same distribution. The transition probabilities for X and Y are then the same. With $P_t(x, \cdot)$ denoting the transition probability of moving from state x in time t for X , the results of [20] show that

$$P_t(x, \cdot) = \mathcal{N} \left(\exp(tB)x, \int_0^t \exp(sB)\sigma\sigma^t \exp(sB^t) ds \right), \tag{5.9}$$

where the right-hand side denotes a Gaussian distribution, and similarly for the transition probabilities of Y . As these are equal for all $x \in \mathbb{R}^p$ and $t \geq 0$, we obtain $\exp(tB) = \exp(t\tilde{B})$ for all $t \geq 0$, so by differentiating, $B = \tilde{B}$ as well. Likewise, as $\int_0^t \exp(sB)\sigma\sigma^t \exp(sB^t) ds = \int_0^t \exp(s\tilde{B})\tilde{\sigma}\tilde{\sigma}^t \exp(s\tilde{B}^t) ds$ for all $t \geq 0$, we obtain $\sigma\sigma^t = \tilde{\sigma}\tilde{\sigma}^t$. Note that

$$\sigma\sigma^t = \begin{pmatrix} \sigma_1\sigma_1^t & \sigma_1\sigma_2^t \\ \sigma_2\sigma_1^t & \sigma_2\sigma_2^t \end{pmatrix}, \tag{5.10}$$

and similarly for $\tilde{\sigma}\tilde{\sigma}^t$. Therefore, we obtain in particular that $\sigma_2\sigma_2^t = \tilde{\sigma}_2\tilde{\sigma}_2^t$.

Now, applying Definition 2.4 and recalling Example 2.5, the intervened processes minus the first coordinate, \tilde{X}^{-1} and \tilde{Y}^{-1} (note that the superscripts do not denote reciprocals), are Ornstein-Uhlenbeck processes with initial values X_0^{-1} and Y_0^{-1} , mean reversion speeds B_{22} and \tilde{B}_{22} , mean reversion levels $-B_{22}^{-1}B_{21}\zeta$ and $-\tilde{B}_{22}^{-1}\tilde{B}_{21}\zeta$ and diffusion matrices σ_2 and $\tilde{\sigma}_2$. As we above concluded that X_0 and Y_0 have the same distribution, $B = \tilde{B}$ and $\sigma_2\sigma_2^t = \tilde{\sigma}_2\tilde{\sigma}_2^t$, we obtain that the distributions of \tilde{X}^{-1} and \tilde{Y}^{-1} must be equal. Thus, by direct calculation of transition probabilities, we see that for the Ornstein-Uhlenbeck with zero mean reversion level, intervention distributions are identifiable from the observational distribution. \circ

Example 5.5. Consider the mapping $a : \mathbb{R}^2 \rightarrow \mathbb{M}(2, 2)$ defined by

$$a(x) = \begin{bmatrix} \frac{x_1}{x_2^2/\sqrt{x_1^2 + x_2^2}} & 0 \\ -x_1x_2/\sqrt{x_1^2 + x_2^2} & \frac{x_1}{x_2^2/\sqrt{x_1^2 + x_2^2}} \end{bmatrix} \tag{5.11}$$

whenever x is not zero, and $a(0) = 0$. This mapping satisfies

$$\begin{aligned} a(x)a(x)^t &= \begin{bmatrix} \frac{x_1}{x_2^2/\sqrt{x_1^2 + x_2^2}} & 0 \\ -x_1x_2/\sqrt{x_1^2 + x_2^2} & \frac{x_1}{x_2^2/\sqrt{x_1^2 + x_2^2}} \end{bmatrix} \begin{bmatrix} x_1 & x_2^2/\sqrt{x_1^2 + x_2^2} \\ 0 & -x_1x_2/\sqrt{x_1^2 + x_2^2} \end{bmatrix} \\ &= \begin{bmatrix} x_1^2 & x_1x_2^2/\sqrt{x_1^2 + x_2^2} \\ x_1x_2^2/\sqrt{x_1^2 + x_2^2} & x_2^2 \end{bmatrix} \end{aligned} \tag{5.12}$$

whenever $x \neq 0$. We will construct another mapping \tilde{a} which has a different signature from a , but which has the same cross product as a , in the sense of having $\tilde{a}(x)\tilde{a}(x)^t =$

$a(x)a(x)^t$. To do so, define $p : \mathbb{R}^2 \rightarrow \mathbb{M}(2, 2)$ by

$$p(x) = \frac{1}{\sqrt{x_1^2 + x_2^2}} \begin{bmatrix} x_1 & x_2 \\ x_2 & -x_1 \end{bmatrix}, \tag{5.13}$$

for $x \neq 0$ and let $p(0)$ be the identity matrix. Put $\tilde{a}(x) = a(x)p(x)$. We then obtain $\tilde{a}(0) = a(0) = 0$ and

$$\begin{aligned} \tilde{a}(x) &= \frac{1}{\sqrt{x_1^2 + x_2^2}} \begin{bmatrix} x_1 & 0 \\ x_2/\sqrt{x_1^2 + x_2^2} & -x_1x_2/\sqrt{x_1^2 + x_2^2} \end{bmatrix} \begin{bmatrix} x_1 & x_2 \\ x_2 & -x_1 \end{bmatrix} \\ &= \frac{1}{\sqrt{x_1^2 + x_2^2}} \begin{bmatrix} x_1^2 & x_1x_2 \\ 0 & (x_2^3 + x_1^2x_2)/\sqrt{x_1^2 + x_2^2} \end{bmatrix} \\ &= \begin{bmatrix} x_1^2/\sqrt{x_1^2 + x_2^2} & x_1x_2/\sqrt{x_1^2 + x_2^2} \\ 0 & x_2 \end{bmatrix}. \end{aligned} \tag{5.14}$$

Note that the first row of a depends only on the first coordinate, while the second row depends on both coordinates. On the other hand, the first row of \tilde{a} depends on both coordinates, while the second row of \tilde{a} depends only on the second coordinate. This translates into a and \tilde{a} corresponding to different signatures, shown in Figure 3.



Figure 3: Left: The signature corresponding to a . Right: The signature corresponding to \tilde{a} .

As $p(x)$ is orthonormal for all x , it holds that $\tilde{a}(x)\tilde{a}(x)^t = a(x)a(x)^t$ and so the solutions to the two SDEs

$$dX_t = a(X_t) dW_t \tag{5.15}$$

$$dX_t = \tilde{a}(X_t) dW_t \tag{5.16}$$

have the same distribution. Thus, we have explicitly constructed two SDEs with the same solution distributions but with different signatures. Note now that the intervention $X^2 := \zeta$ in (5.15) yields an SDE where the first coordinate satisfies

$$dX_t^1 = X_t^1 dW_t^1 \tag{5.17}$$

while the intervention $X^2 := \zeta$ in (5.16) yields an SDE where the first coordinate satisfies

$$dX_t^1 = \frac{(X_t^1)^2}{\sqrt{(X_t^1)^2 + \zeta^2}} dW_t^1 + \frac{X_t^1 \zeta}{\sqrt{(X_t^1)^2 + \zeta^2}} dW_t^2 \tag{5.18}$$

The distribution of the solution to (5.18) is a Markov process whose generator on $C_0^2(\mathbb{R})$ is given by

$$Af(x) = \frac{x^4 + (x\zeta)^2}{x^2 + \zeta^2} \frac{d^2f}{dx^2}(x) = x^2 \frac{d^2f}{dx^2}(x), \tag{5.19}$$

which is the generator of a geometric Brownian motion with zero drift. This is the same as the generator of the solution to (5.17). Thus, as required in Theorem 5.3, the postintervention distributions are the same, even in this case where the signatures are different. \circ

Example 5.5 illustrates a rather curious fact: For SDE models, the postintervention distributions are identifiable from the observational distribution, even when the signature and thus the resulting DAGs of the Euler SEMs are not identifiable from the observational distribution. One interpretation of this is that for SDEs, the postintervention distributions will be the same for all signatures and thus all resulting DAGs which are compatible with the observational distribution. From this perspective, and in concordance with Theorem 5.3, the agreement of the two postintervention distributions in Example 5.5 is not so much related to the dependence structure of $a(x)$, but rather on the dependence structure of $a(x)a(x)^t$, or equivalently, $\tilde{a}(x)\tilde{a}(x)^t$.

This also indicates that in order to obtain a successful theory of causality for SDEs, the relevant concept to consider is postintervention distributions, and not the signatures, since the latter is identifiable from the observational distribution while the former is not. This contrasts with the classical DAG-based case, where a natural methodology consists of first identifying the DAGs compatible with the observational distribution and then, in order to partially infer intervention effects, consider the intervention effect for each possible DAG, as in [23].

Example 5.6. Consider again the yeast microorganism *S. Cerevisiae*. In Example 2.2, we assumed that we were given observations X_{t_0}, \dots, X_{t_p} over time of all $p = 5361$ genes of a non-mutant specimen of *S. Cerevisiae*, and we modeled these observations using an Ornstein-Uhlenbeck process, given by

$$dX_t = B(X_t - A) dt + \sigma dW_t. \quad (5.20)$$

In Example 2.5, we used Definition 2.4 to calculate postintervention distributions from Ornstein-Uhlenbeck processes. We concluded that if Definition 2.4 is applicable (a non-testable hypothesis, according to the discussion of Section 4) and we could identify the parameters in (5.20), then it would be possible to draw inferences about the SDEs resulting from interventions such as knocking out a single gene by setting $X^m := 0$. We also concluded that the parameter σ is not identifiable in our model. Thus, we cannot identify the true SDE, and thus cannot identify the true postintervention SDE.

Assume now, however, that we are satisfied by only knowing the distributional effects of interventions, corresponding to the postintervention distribution. In this case, Theorem 5.3 states that we are in fact capable of inferring the postintervention distribution from the observational distribution. One way of understanding this is that for all SDEs with the same distribution as the observational distribution, the postintervention distribution will be the same. This can be seen explicitly in Example 5.4.

This conclusion allows us, for example, to infer the effects of knocking out genes only from observational distributions. \circ

6 Discussion

In this section, we will reflect on the results of the preceding sections and discuss opportunities for further work.

The definition of the postintervention SDE, Definition 2.4, is a natural way to define how interventions should affect stochastic dynamic systems. It constitutes a generic notion of intervention effects applicable in cases such as Example 2.2 where the background mechanisms of the system are not known and the statistical model is primarily based on observational data. However, the definition reflects assumptions about the

underlying causal nature of the system being modeled, and it is important to make precise when the definition can be assumed to reflect an actual real-world intervention and when the definition is simply a mathematical construct. This is clarified in Section 4, where we used the DAG-based intervention calculus to show that the postintervention SDE of Definition 2.4 can be assumed to reflect real-world interventions when the following hold:

1. The SDE reflects a data-generating mechanism in which the variables at a given timepoint are obtained as a function of the previous timepoints and the driving semimartingales.
2. The driving semimartingales are not directly affected by interventions, in the sense that they can be taken to be noise variables in the Euler SEMs.

In full generality, causal mechanisms of a model are not identifiable from the observational distribution, see [36]. However, when considering only restricted classes of structural equation models, the underlying causal mechanisms may often be identifiable, see for example [38, 19, 27]. In such cases, linearity of the functional relationships or Gaussianity of the noise variables often determine identifiability. In our case, as shown by Theorem 5.3, identifiability holds whenever the driving semimartingale is a Lévy process. This is a key result for the applicability of our notion of intervention, and yields the line of inference depicted in Figure 4: Under sufficient regularity conditions such as appropriate notions of irreducibility, the generator of a Markov process is identifiable from the observational distribution, and Theorem 5.3 allows for deducing postintervention distributions from the generator.

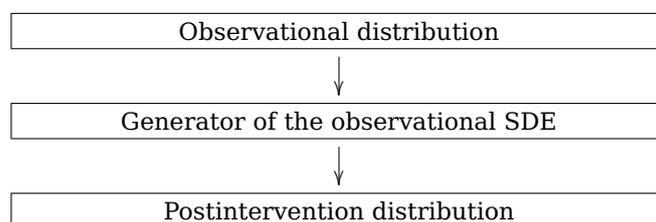


Figure 4: Line of inference for causality in SDEs.

As argued in the series of examples comprised by Example 2.2, Example 2.5 and Example 5.6, in the case of for example time-dependent observations of gene expression data, this allows for identification of knockout effects of genes using only observational data.

The proofs given in Section 5 use the Markov structure of the solution to the SDE. In the case where the driving semimartingale has independent, but not stationary, increments, the solution to the SDE will be an inhomogeneous Markov process, thus also amenable to operator methods, though requiring more powerful technical results. We expect that Theorem 5.3 extends to this case. It should also be noted that identifiability holds independently of the dimension of the driving Lévy process. This is useful, for instance, in relation to Example 2.3. We do not need to use the specific SDE driven by a four-dimensional Wiener process. We can replace the diffusion term in the SDE by a term involving the positive definite square root of the diffusion matrix and a two-dimensional Wiener process without affecting the postintervention distribution.

It is, however, important to be careful about the interpretation of the identifiability result. The result states that when using Definition 2.4 to model interventions, the

postintervention distributions are identifiable. As discussed above, Definition 2.4 is not always useful as a notion of intervention: This requires that we are willing to interpret the SDE in a particular way. As Example 4.6 shows, not all SDEs are amenable to such an interpretation. This requires separate arguments, such as in Example 2.3.

We also remark on the connection between our notion of intervention and the framework of weak conditional local independence (WCLI) discussed in [5, 16]. Definition 2 of [16] defines WCLI for semimartingales in the class \mathcal{D}' . In Remark 1 of [16], it is explained how WCLI can lose its interpretation if extended to larger classes of semimartingales. However, the definition does make sense for all special semimartingales. Extending it to this class, let X be the solution to an SDE of the type (2.12), driven by a Lévy process and assume that X is a special semimartingale. One relationship between our notion of intervention and WCLI is then this: It holds that if X^i is locally unaffected by X^m in (2.12), then X^i is WCLI of X^m . This follows by considering the semimartingale characteristics of solutions to SDEs, see for example Proposition IX.5.3 of [21].

Our results offer opportunities for further research. One main opportunity concerns latent variables: In the DAG-based framework of [26], the back-door and front-door criteria shows how to calculate intervention effects from the observational distribution in the presence of latent variables. For an SDE, the causal structure is summarized in the signature, see Definition 4.1, which does not need to be acyclic, reflecting the possibility of feedback loops. It is an open question how to obtain similar results in terms of the signature in the case of, for example, a diffusion model with some coordinates being unobserved. Another question concerns criteria for when the signature contains useful high-level information about the effects of interventions. As Example 5.5 shows, this is not always the case. We expect that the behaviour seen in Example 5.5 is atypical, but have not shown any precise results about this.

A Proof of Theorem 5.3

In this appendix, we prove Theorem 5.3. We first state some well known and some simple results. The first result, Lemma A.1, is an elementary yet crucial result about interventions in discrete time Markov chains, allowing us to use the Euler scheme to prove Theorem 5.3. Two additional lemmas are simple facts about generators. We do not give full proofs, but we do briefly state how to use results from the literature to obtain full proofs.

For any $G : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}^p$ and $\zeta : \mathbb{R}^{p-1} \rightarrow \mathbb{R}$ we introduce $H_G : \mathbb{R}^{p-1} \times \mathbb{R}^d \rightarrow \mathbb{R}^{p-1}$ by

$$H_G(y, u)^i = G((y_1, \dots, \zeta(y), \dots, y_p), u)^i \tag{A.1}$$

for $i \neq m$ with $\zeta(y)$ at the m 'th position. Now if U and V are random variables with values in \mathbb{R}^d and $\mathbb{R}^{d'}$, respectively, and if $G : \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}^p$ and $\tilde{G} : \mathbb{R}^p \times \mathbb{R}^{d'} \rightarrow \mathbb{R}^p$ fulfill that for all $x \in \mathbb{R}^p$

$$G(x, U) \stackrel{\mathcal{D}}{=} \tilde{G}(x, V), \tag{A.2}$$

meaning that the variables are equal in distribution, then obviously

$$H_G(y, U) \stackrel{\mathcal{D}}{=} H_{\tilde{G}}(y, V). \tag{A.3}$$

for all $y \in \mathbb{R}^{p-1}$. The important consequence that we can derive from this observation is that postintervention distributions in discrete-time Markov processes can be identified

from their transition distributions. Specifically, consider the Markov processes

$$X_n = G(X_{n-1}, U_n), \tag{A.4}$$

$$\tilde{X}_n = \tilde{G}(\tilde{X}_{n-1}, V_n), \tag{A.5}$$

defined recursively in terms of update functions G and \tilde{G} and sequences (U_n) and (V_n) of independent random variables with values in \mathbb{R}^d and $\mathbb{R}^{d'}$, respectively. We also introduce the corresponding intervened processes

$$Y_n = H_G(Y_{n-1}, U_n), \tag{A.6}$$

$$\tilde{Y}_n = H_{\tilde{G}}(\tilde{Y}_{n-1}, V_n). \tag{A.7}$$

The following lemma is a simple consequence of the considerations above.

Lemma A.1. *If $X_0 \stackrel{D}{=} \tilde{X}_0$ and if*

$$G(x, U_n) \stackrel{D}{=} \tilde{G}(x, V_n) \tag{A.8}$$

for all $n \geq 1$ and $x \in \mathbb{R}^p$, then (Y_n) and (\tilde{Y}_n) have the same distribution.

Proving Theorem 5.3 via the Euler scheme will be done by showing that the update formulas for the Euler schemes for two processes with the same generator satisfy (A.8). The following lemma shows how to write (5.2) of Definition 5.1 in a form which is more suitable for the subsequent proof.

Lemma A.2. *Let E be a neighborhood of zero in \mathbb{R}^p . On $C_0^2(\mathbb{R}^p)$, the generator (5.2) of the SDE (2.12) may be rewritten as*

$$\begin{aligned} Af(x) &= \sum_{i=1}^p \beta_i(x) \frac{\partial f}{\partial x_i}(x) + \frac{1}{2} \sum_{i=1}^p \sum_{j=1}^p (a(x)Ca(x)^t)_{ij} \frac{\partial^2 f}{\partial x_i \partial x_j}(x) \\ &\quad + \int_{\mathbb{R}^p} f(x+y) - f(x) - 1_E(y) \sum_{i=1}^p \frac{\partial f}{\partial x_i}(x) y_i dT^{a(x)}(\nu)(y), \end{aligned} \tag{A.9}$$

where $T^{a_0} : \mathbb{R}^d \rightarrow \mathbb{R}^p$ is defined by $T^{a_0}(y) = a_0 y$ for $a_0 \in \mathbb{M}(p, d)$, and

$$\beta_i(x) = \sum_{j=1}^d a_{ij}(x) \alpha_j + \int_{\mathbb{R}^d} (1_E(T^{a(x)}(y)) - 1_D(y)) \sum_{j=1}^d a_{ij}(x) y_j d\nu(y), \tag{A.10}$$

whenever the integrals are well defined and finite. This finiteness condition is in particular satisfied if E is bounded.

The proof of Lemma A.2 is elementary, as (5.2) and (A.9) are equal for all E such that the integrals in (A.9) and (A.10) are well defined and finite. In the case where E is bounded, this is seen to be the case by the integrability properties of Lévy measures.

As a final preparation, we state a lemma on identity of two functionals on $C_0^2(\mathbb{R}^p)$. The nontrivial implication of the lemma is proving that all coefficients are equal if only the functionals are the same.

Lemma A.3. *Fix $x \in \mathbb{R}^p$ and let D be a bounded neighborhood of zero in \mathbb{R}^p . Let $a, \tilde{a} \in \mathbb{R}^p$ and $b, \tilde{b} \in \mathbb{M}(p, p)$, and let ν and $\tilde{\nu}$ be two measures on \mathbb{R}^p such that $x \mapsto \min\{1, \|x\|^2\}$ is integrable with respect to ν and $\tilde{\nu}$. Consider two linear functionals A and \tilde{A} from $C_0^2(\mathbb{R}^p)$ to \mathbb{R} , where A is given by*

$$\begin{aligned} Af &= \sum_{i=1}^p a_i \frac{\partial f}{\partial x_i}(x) + \frac{1}{2} \sum_{i=1}^p \sum_{j=1}^p b_{ij} \frac{\partial^2 f}{\partial x_i \partial x_j}(x) \\ &\quad + \int_{\mathbb{R}^p} f(x+y) - f(x) - 1_D(y) \sum_{i=1}^p \frac{\partial f}{\partial x_i}(x) y_i d\nu(y), \end{aligned} \tag{A.11}$$

and \tilde{A} is given by the same expression, with \tilde{a} , \tilde{b} and $\tilde{\nu}$ substituted for a , b and ν . It then holds that $A = \tilde{A}$ if and only if $a = \tilde{a}$, $b = \tilde{b}$ and $\nu = \tilde{\nu}$ on $\mathbb{R}^p \setminus \{0\}$.

The proof of Lemma A.3 can be obtained as follows. In the notation of the lemma, assume that $A = \tilde{A}$. Using approximate units such as defined in [18], prove that ν and $\tilde{\nu}$ agree on all sets of the form B^c where B is a bounded neighborhood of zero. This implies $\nu = \tilde{\nu}$ on $\mathbb{R}^p \setminus \{0\}$. From this, $a = \tilde{a}$ and $b = \tilde{b}$ follows.

Using the above, we may now make short work of the proof of Theorem 5.3.

Proof of Theorem 5.3. Fix a bounded neighborhood D of zero in \mathbb{R}^d , a bounded neighborhood \tilde{D} of zero in $\mathbb{R}^{\tilde{d}}$ and a bounded neighborhood E of zero in \mathbb{R}^p . Assume that Z has D -characteristics (α, C, ν) and that \tilde{Z} has \tilde{D} -characteristics $(\tilde{\alpha}, \tilde{C}, \tilde{\nu})$. For $a_0 \in \mathbb{M}(p, d)$ define $T^{a_0} : \mathbb{R}^d \rightarrow \mathbb{R}^p$ by $T^{a_0}(y) = a_0 y$. Also define

$$\beta_i(x) = \sum_{j=1}^d a_{ij}(x)\alpha_j + \int (1_E(a(x)y) - 1_D(y)) \sum_{j=1}^d a_{ij}(x)y_j \, d\nu(y) \tag{A.12}$$

$$\tilde{\beta}_i(x) = \sum_{j=1}^{\tilde{d}} \tilde{a}_{ij}(x)\tilde{\alpha}_j + \int (1_E(\tilde{a}(x)y) - 1_{\tilde{D}}(y)) \sum_{j=1}^{\tilde{d}} \tilde{a}_{ij}(x)y_j \, d\tilde{\nu}(y). \tag{A.13}$$

Let $A : C_0^2(\mathbb{R}^p) \rightarrow C_0(\mathbb{R}^p)$ be given by (A.9), and let $\tilde{A} : C_0^2(\mathbb{R}^p) \rightarrow C_0(\mathbb{R}^p)$ be given similarly, except with β , a , C , ν , D and α exchanged by $\tilde{\beta}$, \tilde{a} , \tilde{C} , $\tilde{\nu}$, \tilde{D} and $\tilde{\alpha}$. By our assumptions, $A = \tilde{A}$. As a consequence, by Lemma A.2 and the uniqueness result of Lemma A.3, we find that for all $x \in \mathbb{R}^p$ and $i \leq p$, we have

$$\beta_i(x) = \tilde{\beta}_i(x), \tag{A.14}$$

$$a(x)Ca(x)^t = \tilde{a}(x)\tilde{C}\tilde{a}(x)^t, \tag{A.15}$$

$$T^{a(x)}(\nu) = T^{\tilde{a}(x)}(\tilde{\nu}). \tag{A.16}$$

Now let $\Delta > 0$ and $t_k = k\Delta$. The Euler scheme for the process X is a Markov process given by the update function having i 'th coordinate

$$G(x, U_k)^i = x^i + \sum_{j=1}^d a_{ij}(x)U_k^j, \tag{A.17}$$

with $U_k = Z_{t_k} - Z_{t_{k-1}}$. The Euler scheme for the process \tilde{X} is likewise given by the update function having i 'th coordinate

$$\tilde{G}(x, V_k)^i = x^i + \sum_{j=1}^{d'} \tilde{a}_{ij}(x)V_k^j, \tag{A.18}$$

with $V_k = \tilde{Z}_{t_k} - \tilde{Z}_{t_{k-1}}$. The characteristic function of $G(x, U_k)$ is

$$\begin{aligned} E \exp(iu^t(x + a(x)U_k)) &= \exp(iu^t x) E \exp(iu^t a(x)(Z_{t_k} - Z_{t_{k-1}})) \\ &= \exp(iu^t x) E \exp(i(a(x)^t u)^t (Z_\Delta - Z_0)), \end{aligned} \tag{A.19}$$

for $u \in \mathbb{R}^p$. By (5.1) and some algebraic manipulations, we therefore have

$$\begin{aligned} \log E \exp(iu^t(x + a(x)U_k)) &= iu^t(x + a(x)\Delta\alpha) - \frac{1}{2}\Delta u^t a(x)Ca(x)^t u \\ &\quad - \Delta \int_{\mathbb{R}^d} e^{iu^t a(x)y} - 1 - iu^t a(x)y 1_D(y) \, d\nu(y) \\ &= iu^t(x + \Delta\beta(x)) - \frac{1}{2}\Delta u^t a(x)Ca(x)^t u \\ &\quad - \Delta \int_{\mathbb{R}^d} e^{iu^t y} - 1 - iu^t y 1_E(y) \, dT^{a(x)}(\nu)(y). \end{aligned} \tag{A.20}$$

Making the same calculations for the characteristic function of $\tilde{G}(x, V_k)$ and applying (A.14), (A.15) and (A.16), we conclude that (A.8) holds for the two Euler scheme update functions. Lemma A.1 therefore allows us to conclude that the postintervention Euler SEMs have the same distributions. Using Lemma 4.4, we conclude that the postintervention distributions obtained from the two SDEs are equal. \square

References

- [1] O. O. Aalen, K. Røysland, J. M. Gran, and B. Ledergerber, *Causality, mediation and time: A dynamic viewpoint*, J. Roy. Statist. Soc. Ser. A **175** (2012), no. 4, 831–861. MR-2993496
- [2] D. Anderson and T. G. Kurtz, *Continuous time Markov chain models for chemical reaction networks*, Design and Analysis of Biomolecular Circuits, Springer, Heidelberg, 2011, pp. 3–42.
- [3] D. Applebaum, *Lévy processes and stochastic calculus*, second ed., Cambridge Studies in Advanced Mathematics, vol. 116, Cambridge University Press, Cambridge, 2009. MR-2512800
- [4] T. Björk, *Arbitrage theory in continuous time*, Oxford University Press, 2009.
- [5] D. Commenges and A. Gégout-Petit, *A general dynamical statistical model with causal interpretation*, J. R. Stat. Soc. Ser. B Stat. Methodol. **71** (2009), no. 3, 719–736. MR-2749916
- [6] F. Comte and E. Renault, *Noncausality in continuous time models*, Econometric Theory **12** (1996), no. 2, 215–256. MR-1395031
- [7] A. P. Dawid, *Causal inference without counterfactuals*, J. Amer. Statist. Assoc. **95** (2000), no. 450, 407–448, With comments and a rejoinder by the author. MR-1803167
- [8] V. Didelez, *Graphical models for marked point processes based on local independence*, J. R. Stat. Soc. Ser. B Stat. Methodol. **70** (2008), no. 1, 245–264. MR-2412641
- [9] M. Eichler, *Granger causality and path diagrams for multivariate time series*, J. Econometrics **137** (2007), no. 2, 334–353. MR-2354948
- [10] M. Eichler, *Graphical modelling of multivariate time series*, Probab. Theory Related Fields **153** (2012), no. 1-2, 233–268. MR-2925574
- [11] M. Eichler and V. Didelez, *On Granger causality and the effect of interventions in time series*, Lifetime Data Anal. **16** (2010), no. 1, 3–32. MR-2575937
- [12] A. Gofreau et al., *Life with 6000 genes*, Science **274** (1996), no. 5287, 546–567.
- [13] T. R. Hughes et al., *Functional discovery via a compendium of expression profiles*, Cell **102** (2000), 109–126.
- [14] S. N. Ethier and T. G. Kurtz, *Markov processes*, Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics, John Wiley & Sons Inc., New York, 1986, Characterization and convergence. MR-0838085
- [15] J.-P. Florens and D. Fougere, *Noncausality in continuous time*, Econometrica **64** (1996), no. 5, 1195–1212. MR-1403234
- [16] A. Gégout-Petit and D. Commenges, *A general definition of influence between stochastic processes*, Lifetime Data Anal. **16** (2010), no. 1, 33–44. MR-2575938
- [17] J. B. Gill and L. Petrović, *Causality and stochastic dynamic systems*, SIAM J. Appl. Math. **47** (1987), no. 6, 1361–1366. MR-0916245
- [18] G. Grubb, *Distributions and operators*, Graduate Texts in Mathematics, vol. 252, Springer, New York, 2009. MR-2453959
- [19] P. O. Hoyer, D. Janzing, J. M. Mooij, J. Peters, and B. Schölkopf, *Nonlinear causal discovery with additive noise models*, Advances in Neural Information Processing Systems 21 (NIPS), MIT Press, 2009, pp. 689–696.
- [20] M. Jacobsen, *A brief account of the theory of homogeneous Gaussian diffusions in finite dimensions*, Frontiers in Pure and Applied Probability 1, TVP Science Publishers, 1993, pp. 86–94.
- [21] J. Jacod and A. N. Shiryaev, *Limit theorems for stochastic processes*, second ed., Springer-Verlag, Berlin, 2003. MR-1943877

- [22] M. H. Maathuis, D. Colombo, M. Kalisch, and P. Bühlmann, *Predicting causal effects in large-scale systems from observational data*, *Nature Methods* (2010), no. 4, 247–248.
- [23] M. H. Maathuis, M. Kalisch, and P. Bühlmann, *Estimating high-dimensional intervention effects from observational data*, *Ann. Statist.* **37** (2009), no. 6A, 3133–3164. MR-2549555
- [24] R. C. Merton, *Lifetime portfolio selection under uncertainty: the continuous-time case*, *Rev. Econ. Stat.* **51** (1969), no. 3, 247–257.
- [25] J. M. Mooij, D. Janzing, and B. Schölkopf, *From ordinary differential equations to structural causal models: the deterministic case*, *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, 2013.
- [26] J. Pearl, *Causality*, second ed., Cambridge University Press, Cambridge, 2009, Models, reasoning, and inference. MR-2548166
- [27] J. Peters and P. Bühlmann, *Identifiability of Gaussian structural equation models with same error variances*, *Biometrika*, to appear (2014), 1–11. MR-3180667
- [28] L. Petrović and S. Dimitrijević, *Invariance of statistical causality under convergence*, *Statist. Probab. Lett.* **81** (2011), no. 9, 1445–1448. MR-2811860
- [29] L. Petrović and D. Stanojević, *Statistical causality, extremal measures and weak solutions of stochastic differential equations with driving semimartingales*, *J. Math. Model. Algorithms* **9** (2010), no. 1, 113–128. MR-2596243
- [30] P. E. Protter, *Stochastic integration and differential equations*, *Stochastic Modelling and Applied Probability*, vol. 21, Springer-Verlag, Berlin, 2005, Second edition. Version 2.1, Corrected third printing. MR-2273672
- [31] L. C. G. Rogers and D. Williams, *Diffusions, Markov processes, and martingales. Vol. 1*, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2000, Foundations, Reprint of the second (1994) edition. MR-1796539
- [32] K. Røysland, *A martingale approach to continuous-time marginal structural models*, *Bernoulli* **17** (2011), no. 3, 895–915. MR-2817610
- [33] K. Røysland, *Counterfactual analyses with graphical models based on local independence*, *Ann. Statist.* **40** (2012), no. 4, 2162–2194. MR-3059080
- [34] K.-I. Sato, *Lévy processes and infinitely divisible distributions*, *Cambridge Studies in Advanced Mathematics*, vol. 68, Cambridge University Press, Cambridge, 1999, Translated from the 1990 Japanese original, Revised by the author. MR-1739520
- [35] P. Spirtes, C. Glymour, and R. Scheines, *Causation, prediction, and search*, second ed., *Adaptive Computation and Machine Learning*, MIT Press, Cambridge, MA, 2000, With additional material by David Heckerman, Christopher Meek, Gregory F. Cooper and Thomas Richardson, A Bradford Book. MR-1815675
- [36] T. Verma and J. Pearl, *Equivalence and synthesis of causal models*, *Proceedings of the 6th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, Elsevier, 1991, pp. 255–268.
- [37] D. J. Wilkinson, *Stochastic modelling for systems biology*, *Chapman & Hall/CRC Mathematical and Computational Biology Series*, Chapman & Hall/CRC, Boca Raton, FL, 2006. MR-2222876
- [38] K. Zhang and A. Hyvärinen, *On the identifiability of the post-nonlinear causal model*, *Proceedings of the 25th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, Elsevier, 2009, pp. 647–655.

Acknowledgments. The authors would like to thank Marloes Maathuis for fruitful discussions and comments. We also thank two anonymous reviewers and an associate editor whose comments led to considerable improvements of the paper. Finally we thank Steffen Lauritzen, who suggested that a result similar to Lemma A.1 could be used to prove Theorem 5.3 via the Euler scheme.