

**Original citation:**

Spencer, Simon E. F., Hill, Steven M. and Mukherjee, Sach. (2015) Inferring network structure from interventional time-course experiments. *Annals of Applied Statistics*, Volume 9 (Number 1). pp. 507-524. ISSN 1932-6157

**Permanent WRAP url:**

<http://wrap.warwick.ac.uk/66169>

**Copyright and reuse:**

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.


Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

**Publisher's statement:**

Link to published version: <http://dx.doi.org/10.1214/15-AOAS806>

**A note on versions:**

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here. For more information, please contact the WRAP Team at: [publications@warwick.ac.uk](mailto:publications@warwick.ac.uk)



<http://wrap.warwick.ac.uk/>

## INFERRING NETWORK STRUCTURE FROM INTERVENTIONAL TIME-COURSE EXPERIMENTS

BY SIMON E. F. SPENCER\*, STEVEN M. HILL<sup>†</sup> AND SACH MUKHERJEE<sup>†,‡,1</sup>

*University of Warwick,\* MRC Biostatistics Unit, Cambridge<sup>†</sup>  
and University of Cambridge<sup>‡</sup>*

Graphical models are widely used to study biological networks. Interventions on network nodes are an important feature of many experimental designs for the study of biological networks. In this paper we put forward a causal variant of dynamic Bayesian networks (DBNs) for the purpose of modeling time-course data with interventions. The models inherit the simplicity and computational efficiency of DBNs but allow interventional data to be integrated into network inference. We show empirical results, on both simulated and experimental data, that demonstrate the need to appropriately handle interventions when interventions form part of the design.

**1. Introduction.** Network inference approaches are widely used to study biological networks, including gene regulatory and signaling networks. Since processes underlying such networks are dynamical in nature, time-course data can help to elucidate regulatory interplay. Network inference methods for time-course data have been investigated in the literature, with contributions including (among many others) Hill et al. (2012), Husmeier (2003), Bansal, Gatta and di Bernardo (2006). Scalable assays spanning multiple molecular variables continue to advance and network inference applied to such data offers the potential to provide biological insights over many variables at once. Inferred networks can be used to generate testable hypotheses that are context specific in the sense of reflecting regulatory events in the specific cells under study [Hill et al. (2012), Maher (2012)]. In disease biology, such context-specific networks can be used to shed light on disease-specific processes and thereby inform drug targeting and personalized medicine approaches [Akbari et al. (2014), Ideker and Krogan (2012)].

Interventions, for example, gene knockouts, RNA-interference (RNAi), gene editing or inhibition of kinases, play an important role in experimental designs aimed at elucidating network structure. This is due to the fact that association does not imply causation: interventions can reveal whether a given node has a causal influence on another as opposed to merely being co-expressed. As data acquisition costs fall, interventional time-course designs are becoming more common. It

---

Received July 2013; revised December 2014.

<sup>1</sup>Supported in part by NCI U54 CA112970 and the Cancer Systems Biology Center grant from the Netherlands Organisation for Scientific Research. S. Mukherjee is a recipient of a Royal Society Wolfson Research Merit Award.

*Key words and phrases.* Bayesian inference, network inference, structure learning, causal inference, dynamic Bayesian network, causal Bayesian network.

is important to note that in interventional designs the number of interventions is often much smaller than the number of molecular variables (leave alone the number of possible interventions); this may be due to lack of suitable experimental interventions or cost or both. This means that causal edges cannot simply be directly identified from corresponding interventional experiments; however, causal inference may still be possible using a small subset of all possible interventions [Hytinen, Eberhardt and Hoyer (2013)].

In this paper, we put forward an approach to network inference from time-course data with interventions. To fix ideas, we briefly introduce a data set that we study (and describe in detail) below and that motivated the work described here. The data comprise time-course assays of  $p = 48$  signaling proteins in human cancer cell lines. Experiments were carried out under four conditions: no interventions; intervention on the AKT protein nodes; intervention on the EGFR nodes; intervention on both AKT and EGFR (all interventions were carried out using drugs that inhibit the enzymatic activity of the target, as we describe in detail below). Intuitively, the interventional data are valuable because they give information not only on the causal influences of the target nodes (AKT and EGFR), but also on the wider graph structure, since causal descendants of the target nodes are expected to change under intervention. On the other hand, since the number of interventions carried out is small, a causal graph cannot be estimated by modeling of interventions alone. Rather, a network inference approach is needed that can model the time-course data itself as well as the changes seen under intervention and that is the goal of the present paper.

From the perspective of causal inference [Pearl (2000, 2009)], interventional data require special treatment because the intervention modifies the causal graph and thereby the likelihood. We proceed within a graphical models framework, combining ideas from Dynamic Bayesian Networks (DBNs) and Causal Bayesian Networks [CBNs, see Definition 1.3.1 in Pearl (2000)]. We focus on continuous data, as obtained in conventional biological time-course experiments. Interventions are accommodated by modifying the statistical formulation for those experimental samples in which interventions were carried out, following ideas discussed in Eaton and Murphy (2007), Pearl and Bareinboim (2014) and Pearl (2000). Specifically, for experiments in which interventions are carried out, we modify the structure of the directed acyclic graph (DAG) that underlies the DBN and explore various parameterizations of the effect of the intervention. Our modeling of interventions constitutes a pragmatic extension of DBNs to include causal operations that allow analysis of mainstream experimental data. The approaches we propose can be described in terms of CBNs and the “do” operator of Pearl (2000) applied to the DAG underlying a DBN, as we discuss further below. We therefore refer to them as “Causal Dynamic Bayesian Networks” (CDBNs). However, DBNs themselves are not causal models and full formal justification of the approaches we propose requires additional assumptions, including assumptions on the extent and form of the effect of interventions and, for observational or partially interventional data,

on the absence of hidden common causes. Full discussion of causal semantics is beyond the scope of this paper, but we refer the interested reader to Pearl (2000) for further discussion.

The remainder of the paper is structured as follows: First, a Bayesian framework for network inference using DBNs is outlined in Sections 2.1.1 to 2.1.4. Next, the interventional models that constitute our main focus are described in Section 2.2. We illustrate some key points of the approach using examples in Section 2.3. Empirical results appear in Section 3. We apply the methods on both simulated and experimental protein signaling data, exploring the behavior of a number of approaches by which to model interventions, and comparing their performance with respect to network reconstruction. We find that in the context of interventional data, analyses that do not account for interventions do not perform well. We close with discussion of open questions and future prospects.

An R package “*interventionalDBN*” for network inference using interventional data is available on CRAN and on the author’s website: [www2.warwick.ac.uk/fac/sci/statistics/staff/academicresearch/spencer](http://www2.warwick.ac.uk/fac/sci/statistics/staff/academicresearch/spencer).

**2. Methods.** We fix ideas and notation by first reviewing the “classical” DBN formulation (without interventions). We then go on to discuss in detail how the likelihood can be modified to account for interventions. Taken together, this gives an overall approach by which to perform structural network inference from time-course data that includes interventions acting upon a subset of the nodes.

*2.1. Dynamic Bayesian network model.* A DBN uses a graph to describe probabilistic relationships between variables through time, with associated parameters specifying the temporal relationships. Following Friedman, Murphy and Russell (1998), Husmeier (2003), Murphy (2002), we consider DBNs with edges forward in time only (i.e., first-order Markov with no within-time-slice-edges) and assume stationarity in the sense that neither network topology nor parameters change through time (in what follows, we use “DBN” to refer to this specific class of DBN). Then, each variable at time  $t$  depends only on its regulators at the previous time step. Further, since the graph structure does not change with time and edges are always forward in time, the topology can be described by a graph  $G$  with exactly one vertex for each protein under study and edges understood to mean that the child depends on the state of the parent at the previous time-step. Note that the DBN model is in fact a DAG; the graph  $G$  introduced above can be used to construct a DAG with one vertex for each variable at each time point; this is known as the “unrolled” graph [see, e.g., Hill et al. (2012) for further details]. Operations on DBNs can be described in terms of this underlying DAG, but the  $p$ -vertex graph  $G$  offers a convenient summary.

2.1.1. *Statistical formulation.* Let  $x_{j,c,t}$  denote log-expression of variable  $j \in \{1, \dots, p\}$ , at time  $t \in \{0, \dots, T - 1\}$  in the time course obtained under experimental conditions  $c \in \mathcal{C}$ . We use  $\mathbf{X} = \{x_{j,c,t}\}$  to denote the complete data set. The edge set of the graph  $G$  is  $E(G)$ . Let  $\gamma^{(j)} = \{i : (i, j) \in E(G)\}$  denote the set of parents for node  $j$ . Then, for conditions  $c$  without intervention, the DBN model we consider (for node  $j$ ) is

$$(1) \quad x_{j,c,t} = \begin{cases} \alpha_1^{(j)} + \sum_{i \in \gamma^{(j)}} x_{i,c,t-1} \beta_i^{(j)} + \varepsilon_{j,c,t}, & t > 0, \\ \alpha_2^{(j)} + \varepsilon_{j,c,0}, & t = 0, \end{cases}$$

where  $\beta_i^{(j)}$  denotes parameters that govern the dependence on parent nodes in the previous time step,  $\alpha_1^{(j)}, \alpha_2^{(j)}$  are intercept parameters that do not depend on the parent set  $\gamma^{(j)}$  and  $\varepsilon_{j,c,t} \sim N(0, \sigma_j^2)$  is a noise term. The use of two intercept parameters, one for the initial time point, allows the model more flexibility to incorporate the effects of the parents acting on the first observation. Modeling the initial observation also provides extra degrees of freedom, unless the experimental design has only one unreplicated experimental condition.

2.1.2. *Variable selection.* Under the stationarity and Markov assumptions above, there is a close relationship between inference concerning the DBN graph  $G$  and variable selection for the above regression formulation. As discussed in detail in Hill et al. (2012), exploiting this connection allows efficient inference regarding the graph  $G$ . Specifically, if  $P(i \in \gamma^{(j)} | \mathbf{X})$  is the posterior probability that variable  $i$  appears in the regression model for variable  $j$  above (i.e., the posterior inclusion probability in the variable selection sense) and assuming a modular graph prior  $P(G)$  (i.e., a prior that can be factorized over nodes), we have

$$(2) \quad \begin{aligned} P((i, j) \in E(G) | \mathbf{X}) &= P(i \in \gamma^{(j)} | \mathbf{X}) \\ &= \sum_{\gamma \in \mathcal{M}} I(i \in \gamma) P(\gamma^{(j)} = \gamma | \mathbf{X}), \end{aligned}$$

where  $\mathcal{M}$  denotes the set of all possible variable subsets and  $I$  is the indicator function (for simplicity we assume that  $\mathcal{M}$  does not depend on  $j$ , but it could do so).

Thus, due to the structure of the DBN model, for estimation of posterior probabilities of edges in graph  $G$ , it suffices to perform variable selection for each node in turn, with variables at the previous time point considered as potential predictors. To ease notational burden, we leave dependence on node  $j$  implicit in the following sections. Let  $\mathbf{x} = \{x_{j,c,t}\}$  denote all observations for protein  $j$ ; let  $n (= T \times |\mathcal{C}|)$  be the total number of such observations. Then, model (1) can be written as

$$(3) \quad \mathbf{x} = \mathbf{X}_0 \boldsymbol{\alpha} + \mathbf{X}_\gamma \boldsymbol{\beta} + \boldsymbol{\varepsilon},$$

where  $\mathbf{X}_\gamma$  denotes the matrix formed by selection of columns of  $\mathbf{X}$  corresponding to indices  $\gamma$ ,  $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$ ,  $N_n$  denotes the  $n$ -dimensional multivariate Normal distribution,  $\mathbf{0}_n$  is the  $n$ -dimensional vector of zeros and  $\mathbf{I}_n$  is the  $n \times n$  identity matrix. The design matrix is split into two parts:  $\mathbf{X}_0$ , which is the same for every model and has parameter vector  $\boldsymbol{\alpha}$ ; and  $\mathbf{X}_\gamma$ , which depends on the choice of parents given by  $\gamma$  and has parameter vector  $\boldsymbol{\beta}$ . Let  $a$  be the length of  $\boldsymbol{\alpha}$  and  $b$  be the length of  $\boldsymbol{\beta}$ , then  $\mathbf{X}_0$  has dimension  $n \times a$  and  $\mathbf{X}_\gamma$  has dimension  $n \times b$ . Following equation (1), we see that here  $a = 2$  and  $\mathbf{X}_0 = [\mathbf{1}_{\{t>0\}} \mathbf{1}_{\{t=0\}}]_{n \times 2}$ . In the absence of interventions, observations of the parent proteins from the previous time point form the columns of  $\mathbf{X}_\gamma$  (we discuss interventions below). For the first observation, where there are no previous observations, zeros are inserted into  $\mathbf{X}_\gamma$  in the place of the parent observations.

We can assume without loss of generality that the two parts of the design matrix ( $\mathbf{X}_0$  and  $\mathbf{X}_\gamma$ ) are orthogonal, that is,  $\mathbf{X}_0^T \mathbf{X}_\gamma = \mathbf{0}_{a \times b}$ . This reparameterization ensures the predictors have mean zero; for details see supplementary material [Spencer, Hill and Mukherjee (2015)].

2.1.3. *Marginal likelihood.* The marginal likelihood  $p(\mathbf{x}|\gamma)$  for node  $j$  is obtained by marginalizing over all model parameters, that is,

$$(4) \quad p(\mathbf{x}|\gamma) = \int p(\mathbf{x}|\boldsymbol{\theta}, \gamma) p(\boldsymbol{\theta}|\gamma) d\boldsymbol{\theta},$$

where  $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \sigma)$  is the full set of model parameters. We make use of widely used parameter priors from the Bayesian literature [Denison et al. (2002)]. First, we use improper priors for  $\boldsymbol{\alpha}$  and  $\sigma$ , namely, that  $p(\boldsymbol{\alpha}, \sigma|\gamma) \propto \frac{1}{\sigma}$  for  $\sigma > 0$ . Note that as this prior is improper, for meaningful comparisons to be made between models in  $\mathcal{M}$ , this prior must be the same for all of the models. Second, we use Zellner’s  $g$ -prior for the regression coefficients so that  $\boldsymbol{\beta} | (\boldsymbol{\alpha}, \sigma, \gamma) \sim N_b(\mathbf{0}_b, g\sigma^2(\mathbf{X}_\gamma^T \mathbf{X}_\gamma)^{-1})$ . Following Kohn, Smith and Chan (2001), Smith and Kohn (1996), we set  $g = n$ . With this prior the covariance matrix for  $\boldsymbol{\beta}$  is proportional to  $(\mathbf{X}_\gamma^T \mathbf{X}_\gamma)^{-1}$ , which has some nice properties, for example, invariance to rescaling of the columns of  $\mathbf{X}_\gamma$  [Smith and Kohn (1996)]. Using standard results [Denison et al. (2002)], the marginal likelihood is then given in closed form as

$$(5) \quad p(\mathbf{x}|\gamma) = \frac{K}{(n+1)^{b/2}} \left( \mathbf{x}^T \left( \mathbf{I}_n - \mathbf{P}_0 - \frac{n}{n+1} \mathbf{P}_\gamma \right) \mathbf{x} \right)^{-(n-a)/2},$$

where  $\mathbf{P}_0 = \mathbf{X}_0(\mathbf{X}_0^T \mathbf{X}_0)^{-1} \mathbf{X}_0^T$ ,  $\mathbf{P}_\gamma = \mathbf{X}_\gamma(\mathbf{X}_\gamma^T \mathbf{X}_\gamma)^{-1} \mathbf{X}_\gamma^T$  and the normalizing constant  $K = \frac{1}{2} \Gamma(\frac{n-a}{2}) \pi^{-(n-a)/2} |\mathbf{X}_0^T \mathbf{X}_0|^{-1}$ .

We also wish to consider the model  $\gamma = \emptyset$  (in which  $b = 0$ ). The regression equation is simply  $\mathbf{x} = \mathbf{X}_0 \boldsymbol{\alpha} + \boldsymbol{\varepsilon}$  and the marginal likelihood is given by  $p(\mathbf{x}|\gamma = \emptyset) = K (\mathbf{x}^T (\mathbf{I}_n - \mathbf{P}_0) \mathbf{x})^{-(n-a)/2}$ .

2.1.4. *Model prior.* Following [Scott and Berger \(2010\)](#), we include a multiplicity correction to properly weight models in light of the number of possible parent sets. Since there are  $\binom{p}{k}$  possible models for node  $j$  with  $k$  parents, the prior probability is chosen so that for absent prior information on specific edges we have  $P(\gamma^{(j)} = \gamma) \propto \binom{p}{|\gamma|}^{-1}$ .

We may also wish to include existing biological knowledge in the model prior, which we do by specifying a prior network  $G_0$ , following [Hill et al. \(2012\)](#), [Mukherjee and Speed \(2008\)](#), [Werhli and Husmeier \(2007\)](#). Such a prior network is based on causal biochemistry and should be regarded as a prior on causal structures. A penalty is applied to each candidate graph  $G$  based on the number of edge differences with the prior graph  $G_0$ . That is,

$$(6) \quad P(\gamma^{(j)} = \gamma | \gamma_0^{(j)}) \propto \binom{p}{|\gamma|}^{-1} \exp(-\lambda(|\gamma \setminus \gamma_0^{(j)}| + |\gamma_0^{(j)} \setminus \gamma|)),$$

where  $\gamma_0^{(j)}$  is the parent set of node  $j$  in the prior graph  $G_0$  and  $\lambda$  is a scalar hyperparameter that controls the strength of the prior. Detailed discussion of informative priors for networks is beyond the scope of this paper; we refer interested readers to the references above for further discussion. The prior graph used for results reported in Section 3.2 is shown in [Spencer, Hill and Mukherjee \(2015\)](#). The prior strength parameter was chosen subjectively to be  $\lambda = 4$ .

2.1.5. *Computation.* Combining the marginal likelihood (5) and model prior  $P(\gamma^{(j)})$  gives the posterior  $P(\gamma^{(j)} | \mathbf{x}) \propto p(\mathbf{x} | \gamma^{(j)}) P(\gamma^{(j)})$  over parent sets (so far, without interventions). Posterior probabilities for individual edges in the graph are obtained directly from the posterior over parent sets by (2). As discussed in detail in [Hill et al. \(2012\)](#), placing a bound  $m$  on graph in-degree, following common practice in structural inference for graphical models [e.g., [Husmeier \(2003\)](#)], allows exact computation of the posterior scores.

2.2. *Modeling interventions.* In statistical terms, interventions may alter the edge structure of the graphical model, or model parameters, or both. Here, we discuss the modeling of interventional data as a causal extension of the DBN model outlined above, which we call a Causal Dynamic Bayesian Network (CDBN). For experimental conditions  $c$  that involve an intervention, Section 2.2.1 below outlines different approaches by which to form the likelihood  $p_c(\mathbf{x} | G, \theta_c)$  for an interventional experiment  $c$ . We first give a general typology of interventions following [Eaton and Murphy \(2007\)](#), with extensions to accommodate the wide range of interventions seen in biological experiments, and then go on to discuss kinase inhibition in more detail.

It is important to note that throughout we assume that the nodes targeted by the interventions are known and so the intervention has no additional unmodeled effects elsewhere in the network, an assumption which is integral to the definition



of a CBN [Pearl (2000)]. We also assume that interventions are in effect during the entire period of the experiment (e.g., a gene that is knocked out remains knocked out throughout). This is a reasonable assumption for mainstream interventional designs, but for some interventions that are mediated by reversible biochemistry this may require that the time course is of appropriate total length. Although we do not pursue this direction in this paper, we note that since the approaches described here provide a likelihood that incorporates interventions, they could in principle be used to estimate the targets of interventions.

*2.2.1. Approaches for modeling interventions.* In a *perfect intervention* certain edges that the target node participates in are removed. We call an intervention that corresponds to removal of edges leading out of the target node a *perfect-out* intervention and one that corresponds to removal of edges leading into the target node a *perfect-in* intervention. For example, a knockout with known target gene  $j$  can be thought of as externally setting the transcription level of node  $j$  to zero. This removes the causal influence of other nodes on  $j$  and therefore constitutes a perfect-in intervention. However, since the change to  $j$  may have causal influences on other nodes, outgoing links are allowed to remain.

When applied to a DBN, such an intervention corresponds to a compound “do” [Pearl (2000)] that operates on multiple nodes in the underlying unrolled DAG. For example, the knockout of gene  $j$  mentioned above would correspond to  $\text{do}(X_{j,0} = 0, \dots, X_{j,T-1} = 0)$ , where  $X_{j,t}$  is the vertex (and associated random variable) in the unrolled graph corresponding to gene  $j$  at time  $t$ .

In a *mechanism change intervention* the structure of the graph remains unchanged, but parameters associated with edges that the target participates in are allowed to change. In a *mechanism-change-out* intervention, parameters are re-estimated for the case where the target is a parent; in a *mechanism-change-in* intervention parameters are re-estimated when the target is the child.

In a *fixed-effect intervention*, the effect of the inhibitor is modeled by an additional, additive parameter in the regression equation. In a *fixed-effect-in* intervention the effect appears in the equation for the target itself, while in a *fixed-effect-out* intervention the effect appears in the equations for the children of the target. These formulations can be useful in settings where the intervention results in a change in the average level of the target or its causal descendants. All of the intervention models can be described within the framework of CBNs using the “do” operator of Pearl [Pearl (2000), Pearl and Bareinboim (2014)]; for more details see Spencer, Hill and Mukherjee (2015).

In our empirical results, we focus on a specific type of intervention, namely, drug inhibition of kinases, as used in studies of protein signaling. This application illustrates the need to consider the biological mechanism of the intervention in selecting from the interventional formulations outlined above.

Kinase inhibition blocks the kinase domain of the target, removing the ability of the target to enzymatically influence other nodes. However, such inhibitors may



No intervention  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \xrightarrow{\beta X} Y$	Perfect  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \quad Y$	Mechanism change  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \xrightarrow{\beta^* X} Y$
Fixed effect  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \xrightarrow{\beta X + \delta} Y$	Perfect and fixed effect  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \quad +\delta \quad Y$	Mechanism change and fixed effect  $X \xrightarrow{\beta X} Y$ $X_i \dashv X \xrightarrow{\beta^* X + \delta} Y$

FIG. 1. Diagrammatic representation of intervention models in their “-out” forms, for variables  $X$ ,  $Y$  with directed acyclic graph  $X \rightarrow Y$ ; “ $X_i \dashv X$ ” denotes inhibition of variable  $X$  by an inhibitor  $X_i$  (see text for details).

not prevent phosphorylation of the target itself. Therefore, we focus on “-out” interventions for modeling kinase inhibitors. These intervention models can be used in combination (see Figure 1) to reflect understanding of the biological action of the interventions. The perfect and mechanism change intervention models cannot be used together, as this would introduce a column of zeros into the design matrix. Perfect interventions in combination with fixed-effect interventions are well suited to modeling kinase inhibition using log-transformed data, since they capture the blocking of enzymatic ability and also allow estimation of the quantitative effect of inhibition on child nodes.

Any extra parameters introduced by the intervention models are handled in exactly the same way as the existing regression coefficients denoted by  $\beta$  in equation (3). The design matrix  $\mathbf{X}_Y$  is augmented to include the effects of the interventions in the conditions when they are active and so once the parameters have been integrated out, the marginal likelihood takes the same form as before [equation (5)]. Regressions that include causal components can be described within the framework of Structural Equation Models [for a comprehensive discussion see Chapter 5 of Pearl (2000)]. Full technical details of how to apply these interventions in practice are given in Spencer, Hill and Mukherjee (2015), along with a toy example illustrating their application.

2.3. Protein data example. We now illustrate the foregoing approaches using a simple, real data example (Figure 2) in which a known three-node network is interrogated by inhibition (data courtesy Gray Lab, OHSU Knight Cancer Institute, Portland, OR, USA). Three phospho-proteins—the receptor EGFR, phosphorylated on tyrosine residue #1173 (“EGFRpY1173”), and two nodes downstream of EGFR, namely, AKTpS473 and MAPKpT202—were observed through time under several experimental conditions. The conditions included the following: no inhibitors (green in Figure 2), with an AKT inhibitor (blue), with an EGFR inhibitor

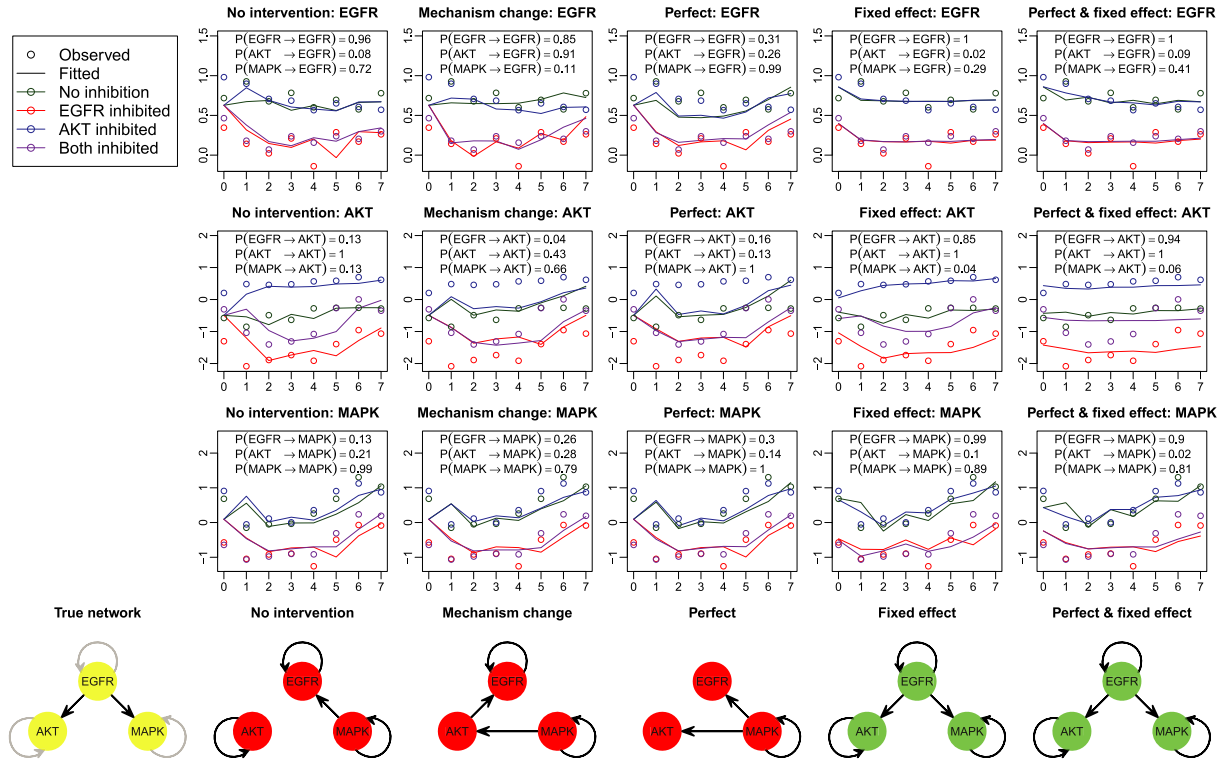


FIG. 2. Real data illustration of the behavior of interventional and noninterventional approaches. Data (circles; expression level vs time index) are from a breast cancer cell line AU565 for three proteins EGFRpY1173, AKTpS473 and MAPKpT202. The data were modeled using DBN (no intervention) and CDBN with mechanism change, perfect, fixed effect, and perfect with fixed effect interventions (the latter all in their “out” form). Fitted values are shown as lines; inferred networks are shown in the last row, with marginal posterior edge probabilities shown in the legends. “True network” indicates what we believe to be the correct causal graph (AKT and MAPK are known to be downstream of EGFR, these edges are verified by the interventional data shown here).

(red) and with both EGFR and AKT inhibition (purple). In line with the known network, the data show a clear reduction in the observed level of AKTpS473 and MAPKpT202 under EGFR inhibition.

To investigate the behavior of the interventional schemes described above, we carried out network inference for these data using a CDBN with the respective intervention scheme (in their “out” forms). We show the data itself, posterior expected fitted values obtained via model averaging (hereafter abbreviated to fitted values) from the various models and the corresponding inferred networks. Strikingly, although several of the methods fit the data reasonably well, only fixed effect and perfect fixed effect are able to both fit the data and estimate what we believe to be the correct network.

It is noteworthy that even in this simple example it is possible to fit the data well while estimating a plainly incorrect network. For example, the no intervention model fits the data (including the inhibitor time courses) reasonably well, but does not estimate the known edges from EGFR to MAPK and AKT, despite the fact that both MAPK and AKT change dramatically under EGFR inhibition in the very data being analyzed. This is an example of statistical confounding that arises due to the fact that the data are analyzed “blind”: the analysis does not know which time course was obtained under EGFR inhibition, rendering the easily seen causal effect of EGFR on AKT and MAPK invisible to network inference. In contrast, the fixed-effect intervention approaches can directly incorporate this information in the overall network inference. Note also that the inhibitors can be seen to affect the concentration of their target proteins, most likely due to feedback mechanisms that are represented by self-edges in the estimated network. For more discussion about the role of the self-edge in the network, see Section 4.

### 3. Results.

3.1. *Simulation study.* We performed a simulation study to compare the network inference methods with different intervention models. Data for 15 nodes were simulated from a CDBN using a data-generating graph  $G^*$  [see Figure S3 in Spencer, Hill and Mukherjee (2015)]. Mimicking the design of typical real proteomic experiments, for each protein we simulated a small number of time points (8) in four experimental conditions (no inhibitor; inhibition of node “A”; inhibition of node “B”; inhibition of both node “A” and node “B”), giving  $n = 32$  multivariate datapoints. We sampled coefficients for the node-specific linear models uniformly at random from the set  $(-1, -0.5) \cup (0.5, 1)$  in order to create associations of various strengths, while avoiding associations close to zero. To simulate data under *in silico* inhibition requires an intervention model: since each interventional scheme also corresponds to such a data-generating model, to avoid bias, we simulated data based on all four intervention models that were considered, namely, perfect, fixed effect, perfect with fixed effect and mechanism change (all in their “out” forms). Network inference was then carried out as described above using

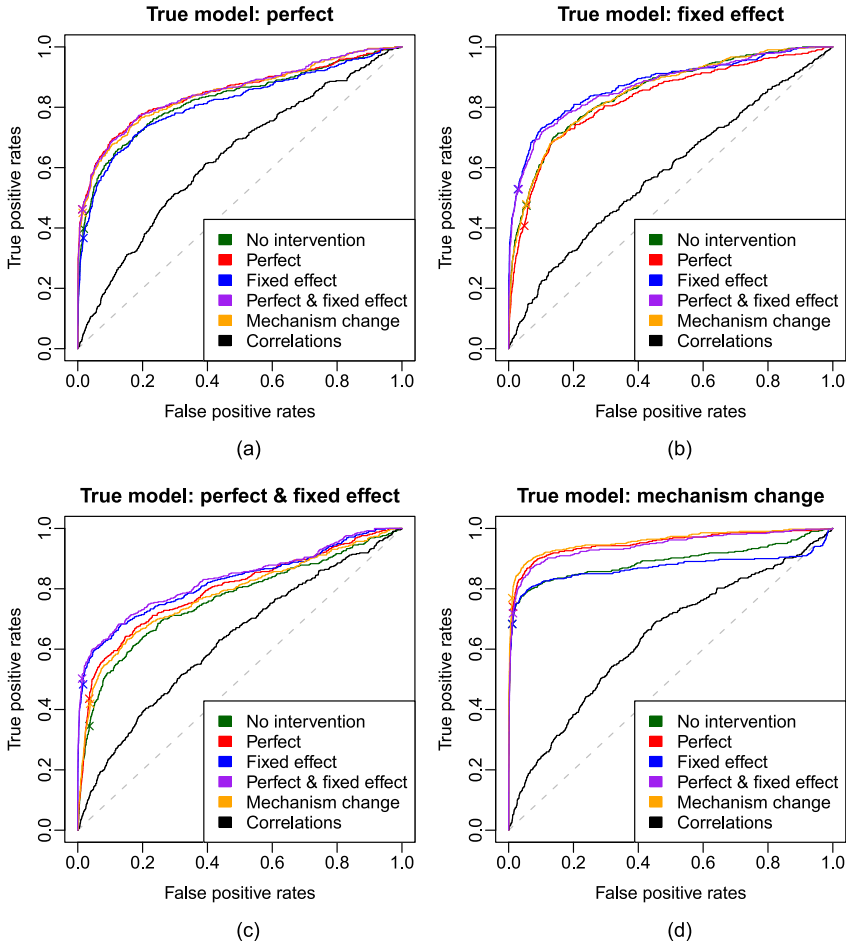


FIG. 3. Simulated data results. Data were generated based on (a) perfect (b) fixed effect, (c) perfect and fixed effect and (d) mechanism change intervention models (all in their “out” form) and analyzed using CDBNs coupled to these four intervention models, plus a classical DBN (“no intervention”) and a baseline co-expression analysis (“correlations”; see text for details). Receiver Operating Characteristic (ROC) curves for each method in each data-generating regime are shown; the crosses correspond to the point estimate of the network obtained by thresholding marginal posterior edge probabilities at 1/2.

these four intervention models plus the model with no intervention and simple, marginal correlations between nodes (i.e., a co-expression network).

Figure 3 shows ROC curves (produced from 20 data sets for each regime) for each combination of intervention method and the underlying model. These curves plot true positive rates (with respect to edges in the data-generating graph) against false positive rates across a range of thresholds on marginal posterior edge probabilities. In each case and as expected, analysis under the data-generating model

gives the best results. However, the perfect and fixed-effect model does consistently well and generally performs almost as well as inference using the data-generating model. The mechanism change model generally appears to perform similarly to the perfect intervention model. Co-expression analysis does much less well than all of the CDBN models. Note that even under the correct data-generating model, in this noisy, small sample example, the area under the ROC curve can be much lower than unity, highlighting the inherent difficulty of the network inference problem and the challenging nature of the simulation.

### 3.2. *Cancer cell line data.*

3.2.1. *Data.* Phospho-protein time courses were obtained from two breast cancer cell lines (AU565 and BT474) for 48 proteins using reverse-phase protein arrays (data courtesy Gray Lab, Knight Cancer Center, OHSU, Portland, OR; these data form part of a larger, ongoing study covering a broad panel of breast cancer cell lines and a larger set of proteins). Data comprised 8 time points (0.5, 1, 2, 4, 8 and 24 hours following Serum stimulation) in 4 experimental conditions: no inhibitor (DMSO); EGFR inhibition by Lapatinib (“EGFRi”, at a dose of 250 nM);<sup>2</sup> AKT inhibition by GSK690693 (“AKTi”, at 250 nM); and inhibition by both EGFRi and AKTi (each at 250 nM). This gave  $n = 32$  datapoints for each protein. For further details concerning experimental protocol see [Spencer, Hill and Mukherjee \(2015\)](#).

3.2.2. *Cell line specific networks.* The two cancer cell lines studied differ in terms of the genetic alterations that they harbor [[Neve et al. \(2006\)](#)] and may differ in terms of underlying signaling network topology. To avoid aggregating potentially heterogeneous data, we analyzed the cell lines separately to obtain cell line-specific networks. Figure 4 shows the inferred networks for the two cell lines. The edges highlighted in green are not inferred with the conventional DBN without interventions [the full networks inferred by the no intervention DBN are shown in [Spencer, Hill and Mukherjee \(2015\)](#)]. We discuss network validity below, but note that full validation of the cell line-specific networks requires further experimental work and is beyond the scope of this paper.

3.2.3. *Network validity.* In this real data example, the true data-generating networks are not known. However, since the experimental design includes interventions, the relevant data (EGFRi and AKTi) can be used to test the causal validity of the estimated networks downstream of the inhibited nodes. For example, suppose a node  $k$  changes under inhibition of AKT. This means that  $k$  is downstream of AKT; since the observation is made under external manipulation of AKT, we can

---

<sup>2</sup>Lapatinib is a dual EGFR/HER2 inhibitor.

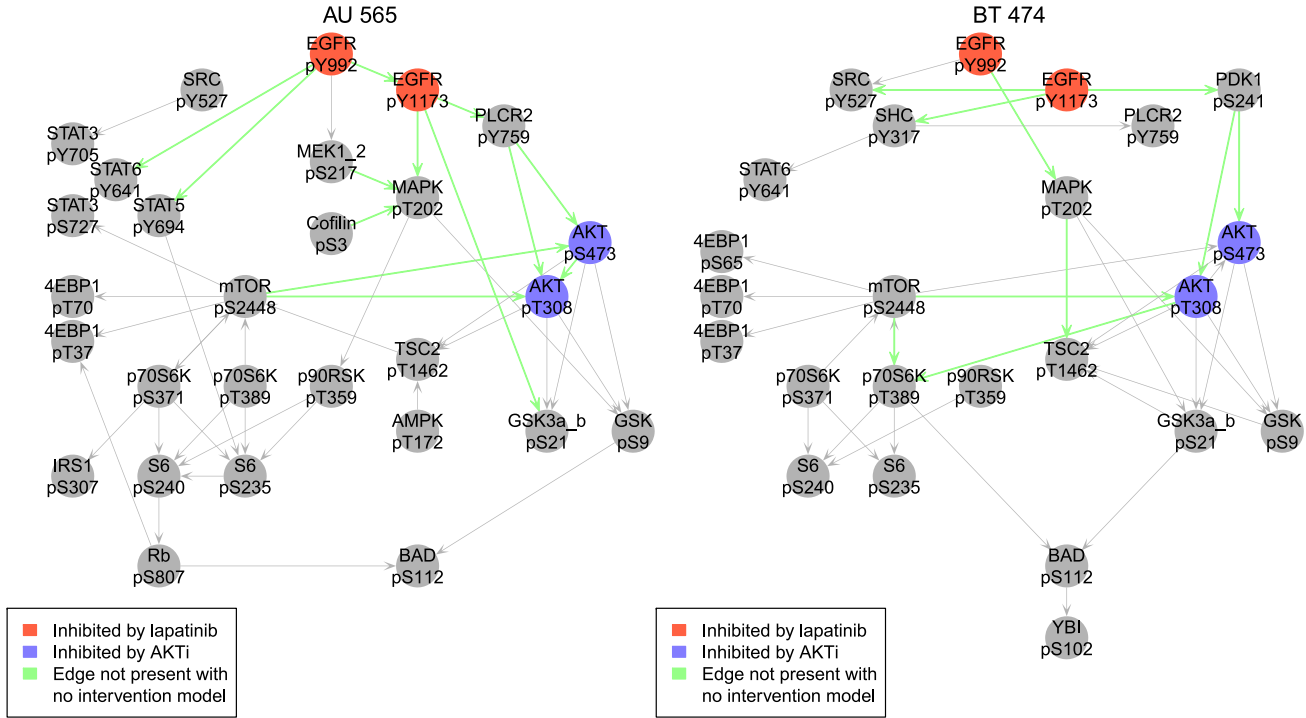


FIG. 4. Estimated networks for cell lines AU565 and BT474. Data were analyzed using a CDBN with informative prior and perfect fixed effect out interventions; all edges with posterior probability greater than 0.5 are shown. Highlighted edges are not present when no intervention model is used. Proteins with no edges are not shown.

say that  $k$  is a descendant of AKT in the underlying causal graph. Testing each node for change under AKTi (this was done using a paired  $t$ -test at the 5% level) gave a set  $D_{\text{AKT}}$  of nodes downstream of AKT that could be compared against the corresponding set of descendants from the inferred networks. This was done in an ROC-sense in the following way for each cell line. First, we thresholded posterior edge probabilities at  $\tau$  to obtain a network  $\hat{G}_\tau$ , which then gave an estimated set of descendants of AKT,  $\hat{D}_{\text{AKT}}(\tau)$ . The number of true and false positives at threshold  $\tau$  are then  $|\hat{D}_{\text{AKT}}(\tau) \cap D_{\text{AKT}}|$  and  $|\hat{D}_{\text{AKT}}(\tau) \setminus D_{\text{AKT}}|$ , respectively. Varying threshold  $\tau$  then gave an ROC curve assessing ability to recover causal descendance across the full range of thresholds. To ensure that our inference approach could uncover direct associations and that our conclusions did not depend solely on the form of this analysis, we also compared the set of significant downstream nodes  $D_{\text{AKT}}$  with the inferred set of direct children of AKT for each posterior edge probability threshold. These variants of the first two ROC plots are shown in Figure 5(c) and (d).

Figure 5(a) shows the ROC curves for each of the intervention approaches considered, combined across both inhibited proteins (EGFR and AKT) and cell lines (AU565 and BT474). The fixed-effect approach has the highest ROC curve area (0.830), marginally ahead of perfect fixed effect (0.823), which has the posterior median network with the highest true positive rate. The perfect fixed-effect model is compared with other methods in the literature in Figure 5(b), including DDEPN [Bender et al. (2010)] and a Gaussian process-based method due to Äijö and Lähdesmäki (2009), which does not explicitly model interventions. The influence of the prior network is also explored.

Due to the limited size of the data set, it is not feasible to leave out the inhibitor data used to produce the ROC curves and still carry out network inference. The results shown should therefore be regarded as an assessment of “causal fit” rather than a validation of causal links. It is noteworthy that the existing approaches, including classical DBNs, are not able to correctly estimate the causal dependencies even for the interventions and protein levels that are present in the training data itself.

**4. Discussion.** Recently, there has been interesting work on explicitly causal methods for networks, including linear [Maathuis, Kalisch and Bühlmann (2009)] and nonlinear [Oates and Mukherjee (2012)] models. Interventional data are important for elucidation of causal links. However, standard DBNs are not appropriate for modeling interventional data. By modeling interventions explicitly we were able to extend DBNs in a causal direction. We discussed and illustrated the issues of confounding that can arise in network inference. As we showed using real data, nodes not linked in terms of regulation can nonetheless exhibit statistical association and thereby easily lead network inference astray. We showed how such confounding can present a concern even with only three nodes, but the issue becomes rapidly more severe in higher dimensions.



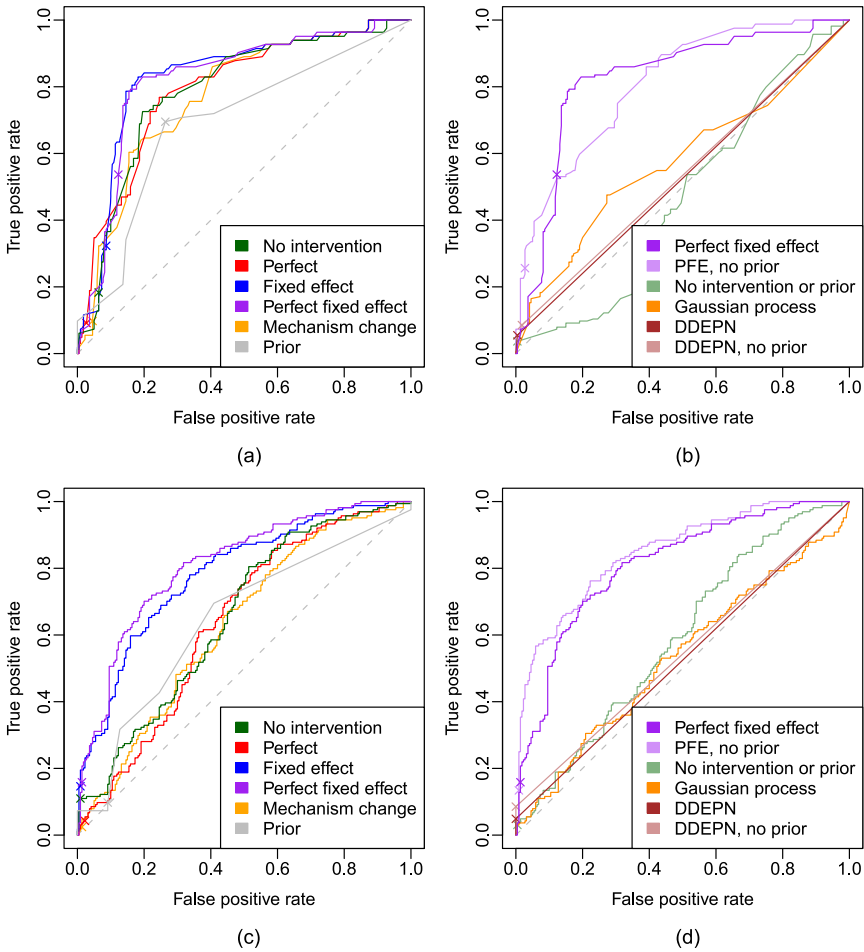


FIG. 5. *Real data results. Receiver Operating Characteristic (ROC) curves showing agreement of estimated networks with changes observed under experimental intervention; the crosses correspond to the point estimate of the network obtained by thresholding marginal posterior edge probabilities at 1/2. Upper panels show results based on analysis of descendency in estimated networks; lower panels show corresponding results using direct children in estimated networks (see text for details). Left panels compare CDBNs using various interventional models against each other; right panels compare selected CDBNs with several existing approaches (as detailed in text).*

The posterior edge probabilities that we report are not truly causal quantities. In principle, it could be possible to instead consider causal coefficients calculated via the do-calculus [as in Maathuis, Kalisch and Bühlmann (2009)]. However, since interventions in time-course experiments are in fact compound do operations (applying to multiple time points and therefore multiple nodes in the unrolled DAG), calculation of causal coefficients is more complicated than in the static DAG case, and we are not aware of a simple way to proceed in this setting, even for the lin-

ear models considered here. On the other hand, in contrast to static DAGs as in Maathuis, Kalisch and Bühlmann (2009), for feed-forward DBNs of the type considered here, the underlying DAG is identifiable (i.e., the equivalence class always contains exactly one graph). We showed empirically that the posterior edge probabilities we reported provide useful information on causal edges, but we do not currently fully understand the relationship between such measures [as used here and in most mainstream Bayesian approaches for biological network inference, including, among others, Husmeier (2003), Hill et al. (2012)] and the corresponding causal coefficients, and further work in this area would be valuable. We reiterate that causal interpretation of CDBNs requires additional assumptions that go beyond those needed to justify conditional independence statements. However, as noted by Dawid (2007) in the context of static DAGs, such assumptions are generally difficult, if not impossible, to check. Therefore, empirical validation of causal inference remains a crucial direction for future work.

Results on simulated data suggested that the “perfect-fixed-effect-out” intervention scheme we proposed represents a good default choice for kinase inhibition experiments. We conjecture that “perfect-fixed-effect-in” interventions may represent a good default approach for analysis of gene expression time-course data obtained under knockouts and RNAi knockdowns, but we did not explore such data here. We recommend that an intervention model should be chosen in line with the mechanism of the intervention under consideration. In situations where biochemical knowledge is insufficient, it may be possible to treat the choice of interventional regime as a model selection problem, but we did not explore this possibility here.

The networks shown in Figure 4 reflect several features that are typical of protein signaling, including a cascade-type structure originating from the receptor EGFR. The edges highlighted in green show the changes in the network that are induced by modeling the interventions, and the improvement in the ROC curve in Figure 5(a) suggests that using the perfect fixed-effect intervention model has produced a more accurate network, particularly around the inhibited proteins. Since these edges (in green) are inferred only when the inhibition is taken into account, they may be more likely to reflect causal information. There are more differences between the two cell lines than might have been expected. These differences may be real or may be due to some of the many limitations inherent to biological network analyses, including experimental caveats and limitations of the inference approach and causal models. However, experimental validation of the networks is beyond the scope of this paper.

The ROC curves in Figure 5 show the perfect fixed-effect model performs better than several other approaches. However, the poor performance of the no intervention DBN model (which is identical except for the modeling of interventions) demonstrates that this success is not based on the network inference scheme or the prior, but on the appropriate handling of interventions. Surprisingly, the Gaussian process method performs better overall than DDEPN, even though the latter models interventions. This may be due to the fact that the inference is conducted over

a relatively large network (48 proteins) and DDEPN suggests a very small set of potential edges.

The self-edge (the edge that connects a protein to itself) has two roles in the model. First, it can represent statistical autocorrelation in the protein time course. Second, it can represent a (negative or positive) feedback loop, possibly via some additional unmeasured variables. Since we integrate out the regression coefficient to obtain the marginal likelihood, the posterior signaling network does not give any indication of the sign of any feedback, nor the role of the self-edge. In future work we hope to differentiate between inhibition and activation effects in the signaling network, helping to clarify the role of the self-edges.

**Acknowledgments.** We would like to thank the Editor and anonymous referees for constructive input that we think improved the paper. We would like to thank Jim Korkola and Joe Gray (OHSU Knight Cancer Institute, Portland, OR) for the proteomic data set used in this paper and for a productive, ongoing collaboration.

#### SUPPLEMENTARY MATERIAL

**Supplement to “Inferring network structure from interventional time-course experiments”** (DOI: [10.1214/15-AOAS806SUPP](https://doi.org/10.1214/15-AOAS806SUPP); .pdf). Additional technical information about orthogonalization, the experimental procedure and the intervention models, including a toy example. Supplementary figures showing the prior network, the “true” network used for simulations and the posterior signaling networks without interventions.

#### REFERENCES

- ÄIJÖ, T. and LÄHDESMÄKI, H. (2009). Learning gene regulatory networks from gene expression measurements using non-parametric molecular kinetics. *Bioinformatics* **25** 2937–2944.
- AKBANI, R. et al. (2014). A pan-cancer proteomic perspective on the Cancer Genome Atlas. *Nature Communications* **5** 3887.
- BANSAL, M., GATTA, G. D. and DI BERNARDO, D. (2006). Inference of gene regulatory networks and compound mode of action from time course gene expression profiles. *Bioinformatics* **22** 815–822.
- BENDER, C., HENJES, F., FRÖHLICH, H., WIEMANN, S., KORF, U. and BEISSBARTH, T. (2010). Dynamic deterministic effects propagation networks: Learning signalling pathways from longitudinal protein array data. *Bioinformatics* **26** i596–i602.
- DAWID, A. P. (2007). Fundamentals of statistical causality. Research Report No. 279, Dept. Statistical Science, Univ. College, London.
- DENISON, D. G. T. et al. (2002). *Bayesian Methods for Non-Linear Classification and Regression*. Wiley, Chichester, UK.
- EATON, D. and MURPHY, K. (2007). Exact Bayesian structure learning from uncertain interventions. *Journal of Machine Learning Research: Workshop and Conference Proceedings* **2** 107–114.
- FRIEDMAN, N., MURPHY, K. and RUSSELL, S. (1998). Learning the structure of dynamic probabilistic networks. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence* 139–147. Morgan Kaufmann, San Francisco, CA.

- HILL, S. M. et al. (2012). Bayesian inference of signaling network topology in a cancer cell line. *Bioinformatics* **28** 2804–2810.
- HUSMEIER, D. (2003). Sensitivity and specificity of inferring genetic regulatory interactions from microarray experiments with dynamic Bayesian networks. *Bioinformatics* **19** 2271–2282.
- HYTTINEN, A., EBERHARDT, F. and HOYER, P. O. (2013). Experiment selection for causal discovery. *J. Mach. Learn. Res.* **14** 3041–3071. [MR3138909](#)
- IDEKER, T. and KROGAN, N. J. (2012). Differential network biology. *Mol. Syst. Biol.* **8** 565.
- KOHN, R., SMITH, M. and CHAN, D. (2001). Nonparametric regression using linear combinations of basis functions. *Stat. Comput.* **11** 313–322. [MR1863502](#)
- MAATHUIS, M. H., KALISCH, M. and BÜHLMANN, P. (2009). Estimating high-dimensional intervention effects from observational data. *Ann. Statist.* **37** 3133–3164. [MR2549555](#)
- MAHER, B. (2012). ENCODE: The human encyclopaedia. *Nature* **489** 46–48.
- MUKHERJEE, S. and SPEED, T. P. (2008). Network inference using informative priors. *Proc. Natl. Acad. Sci. USA* **105** 14313–14318.
- MURPHY, K. P. (2002). Dynamic Bayesian networks: Representation, inference and learning. Ph.D. thesis, Univ. California, Berkeley. [MR2704368](#)
- NEVE, R. M. et al. (2006). A collection of breast cancer cell lines for the study of functionally distinct cancer subtypes. *Cancer Cell* **10** 515–527.
- OATES, C. J. and MUKHERJEE, S. (2012). Network inference and biological dynamics. *Ann. Appl. Stat.* **6** 1209–1235. [MR3012527](#)
- PEARL, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge Univ. Press, Cambridge. [MR1744773](#)
- PEARL, J. (2009). Causal inference in statistics: An overview. *Stat. Surv.* **3** 96–146. [MR2545291](#)
- PEARL, J. and BAREINBOIM, E. (2014). External validity: From do-calculus to transportability across populations. *Statist. Sci.* **29** 579–595. [MR3300360](#)
- SCOTT, J. G. and BERGER, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *Ann. Statist.* **38** 2587–2619. [MR2722450](#)
- SMITH, M. and KOHN, R. (1996). Nonparametric regression using Bayesian variable selection. *J. Econometrics* **75** 317–344.
- SPENCER, S. E. F., HILL, S. M. and MUKHERJEE, S. (2015). Supplement to “Inferring network structure from interventional time-course experiments.” DOI:10.1214/15-AOAS806SUPP.
- WERHLI, A. V. and HUSMEIER, D. (2007). Reconstructing gene regulatory networks with Bayesian networks by combining expression data with multiple sources of prior knowledge. *Stat. Appl. Genet. Mol. Biol.* **6** 47 pp. (electronic). [MR2349908](#)

S. E. F. SPENCER  
 DEPARTMENT OF STATISTICS  
 UNIVERSITY OF WARWICK  
 COVENTRY  
 CV4 7AL  
 UNITED KINGDOM  
 E-MAIL: [s.e.f.spencer@warwick.ac.uk](mailto:s.e.f.spencer@warwick.ac.uk)

S. M. HILL  
 MRC BIostatISTICS UNIT  
 CAMBRIDGE  
 CB2 0SR  
 UNITED KINGDOM  
 E-MAIL: [steven.hill@mrc-bsu.cam.ac.uk](mailto:steven.hill@mrc-bsu.cam.ac.uk)

S. MUKHERJEE  
 MRC BIostatISTICS UNIT  
 AND UNIVERSITY OF CAMBRIDGE  
 CAMBRIDGE  
 CB2 0SR  
 UNITED KINGDOM  
 E-MAIL: [sach@mrc-bsu.cam.ac.uk](mailto:sach@mrc-bsu.cam.ac.uk)