

Model Uncertainty and Correctability for Directed Graphical Models*

Panagiota Birmpa[†], Jinchao Feng[‡], Markos A. Katsoulakis[†], and Luc Rey-Bellet[†]

Abstract. Probabilistic graphical models are a fundamental tool in probabilistic modeling, machine learning, and artificial intelligence. They allow us to integrate in a natural way expert knowledge, physical modeling, heterogeneous and correlated data, and quantities of interest. For exactly this reason, multiple sources of model uncertainty are inherent within the modular structure of the graphical model. In this paper we develop information-theoretic, robust uncertainty quantification methods and nonparametric stress tests for directed graphical models to assess the effect and the propagation through the graph of multisourced model uncertainties to quantities of interest. These methods allow us to rank the different sources of uncertainty and correct the graphical model by targeting its most impactful components with respect to the quantities of interest. Thus, from a machine learning perspective, we provide a mathematically rigorous approach to correctability that guarantees a systematic selection for improvement of components of a graphical model while controlling potential new errors created in the process in other parts of the model. We demonstrate our methods in two physicochemical examples, namely, quantum scale-informed chemical kinetics and materials screening to improve the efficiency of fuel cells.

Key words. Bayesian networks, uncertainty quantification, sensitivity analysis, stress tests, information inequalities, correctability

MSC codes. 62H22, 62P30, 68T37, 80A30, 93B35, 94A17

DOI. 10.1137/21M1434453

1. Introduction. Data-informed, structured probability models are typically constructed by combining expert-based mathematical models with available data, the latter often being heterogeneous, i.e., from multiple sources and scales, and possibly sparse or imperfect. Typically such structured models are formulated as probabilistic graphical models (PGMs), which in turn are generally classified into Markov random fields (MRFs) over undirected graphs and Bayesian networks over a directed acyclic graph (DAG) [58, 48], as well as mixtures of those

* Received by the editors July 19, 2021; accepted for publication (in revised form) May 16, 2022; published electronically October 31, 2022. The U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. Copyright is owned by SIAM to the extent not limited by these rights.

<https://doi.org/10.1137/21M1434453>

Funding: The research of the first author was supported by the Air Force Office of Scientific Research (AFOSR) under grants FA-9550-18-1-0214 and FA9550-21-1-0354. The research of the second author was partially supported by the Defense Advanced Research Projects Agency (DARPA) EQUIPS program under grant W911NF1520122. The research of the third and fourth authors was partially supported by the Air Force Office of Scientific Research (AFOSR) under grants FA-9550-18-1-0214 and FA9550-21-1-0354 and by the National Science Foundation (NSF) under grants NSF TRIPODS CISE-1934846 and DMS-2008970.

[†] Department of Mathematics and Statistics, University of Massachusetts, Amherst, MA 01003 USA (birmpa@math.umass.edu, markos@math.umass.edu, luc@math.umass.edu).

[‡] Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD USA (jfeng34@jhu.edu).

two classes [31]. In this paper we focus on Bayesian networks. DAGs are graphs with directed edges and without cycles, where individual vertices correspond to different model components or data inputs, while the directed edges encode conditional dependencies between vertices. Formally, a Bayesian network over a DAG is defined as a pair $\{G, P\}$ consisting of the following ingredients: $G = \{V, E\}$ is a DAG with n vertices denoted by $V = \{1, \dots, n\}$, $n \in \mathbb{N}$, along with directed, connecting edges $E \in V \times V$. In addition, we define a set of random variables $X_V = \{X_1, \dots, X_n\}$ over V with probability distribution P with density

$$(1.1) \quad p(x) = \prod_{i=1}^n p(x_i | x_{\pi_i}),$$

where $x_{\pi_i} = \{x_{i_1}, \dots, x_{i_m}\} \subset \{x_1, \dots, x_n\}$ denotes the values of parents π_i of each vertex i (see Figure 1), and $p(x_i | x_{\pi_i})$ is the conditional probability density (CPD) for the conditional distribution $P_{i|\pi_i}$ with parents π_i . In such models we are typically interested in quantities of interest (QoIs) $f(X_A)$ that involve one or more vertices $A \subset V$ and the corresponding random variables $X_A \subset X_V$.

The general mathematical formulation of PGMs was developed in foundational works in [58, 59] and are widely used in many real-world applications of artificial intelligence, like medical diagnostics, natural language processing, computer vision, robotics, computational biology, and cognitive science, to name a few; see, e.g., [26, 42, 3, 51, 50, 27]. Recently PGMs were built as computationally tractable surrogates for multiscale/multiphysics models (e.g., from quantum to molecular to engineering scales), such as in porous media and energy storage [69, 40]. Such models often have hidden correlations in data used in their construction [68] or include physical constraints in parameters [69], necessitating conditional relations between model components and thus giving rise to CPDs such as the ones in (1.1). Finally, PGMs can be used as the mathematical foundation for building digital twins used for control and optimization tasks of real systems [61]. Some examples include Bayesian networks for fuel cells [25] and hidden Markov models (a time-dependent special case of Bayesian networks) for unmanned aerial vehicles [46]. The structured probabilistic nature of such models allows them to be continuously improved, e.g., based on available real-time data [46] or through targeted data acquisition [25].

A. Model uncertainty in Bayesian networks. Bayesian networks will typically have multiple sources of uncertainty due to modeling choices or learning from imperfect data in the process of building the graph G and each one of the CPDs in (1.1). These uncertainties will propagate (and occasionally not propagate—see section 7) through the directed graph structure to the targeted QoIs. Uncertainties in probabilistic models are broadly classified in two categories: *aleatoric*, due to the inherent stochasticity of probabilistic models such as (1.1) and *model uncertainties* (also known as “epistemic”) [21, 65]. In this paper we primarily focus on model uncertainties which stem from the inability to accurately model one or more of the components of a Bayesian network $\{G, P\}$: omitting or simplifying model components as is often the case in multiscale systems, learning from sparse data, or using approximate inference methods to build CPDs in (1.1). Next, we discuss more concretely these challenges in the context of two physicochemical examples that we analyze further using model uncertainty methods developed here.

First, we consider a Langmuir bimolecular adsorption model (see section 6) that describes the chemical kinetics with competitive dissociative adsorption of hydrogen and oxygen on a catalyst surface [62]. It is a multiscale system of random differential equations with correlated dependencies in their parameters (kinetic coefficients), arising from quantum-scale computational data calculated using density functional theory (DFT) (i.e., quantum computations) for actual metals. The combination of chemical kinetics with parameter dependencies, correlations, and DFT data gives rise naturally to a Bayesian network. The QoIs are the equilibrium hydrogen and oxygen coverages computed as the stationary solutions of a system of mean-field differential equations. Here the Bayesian network allows us to incorporate data and correlations from a different scale to the parameters of an established chemical kinetics (differential equations) model. However, the limited availability of the quantum-scale data creates significant model uncertainties in the distributions of kinetic coefficients (see, for example, Figure 5(a)) and the need to be accounted towards obtaining reliable predictions for the QoI.

In a second example analyzed in section 7 we build suitable Bayesian networks for trustworthy screening of materials to increase the efficiency of chemical reactions. Here we consider the oxygen reduction reaction (ORR), which is a known performance bottleneck in fuel cells [63]. This electrochemistry mechanism involves two reactions which are typically slow. Thus, we seek new materials that speed up these two slowest reactions. For this reason here we focus only on the thermodynamics of these reactions described by the volcano curve of Sabatier's principle [62]. Based on Sabatier's principle, the optimal oxygen binding energy is the natural descriptor for discovering new materials, and hence it has to be our QoI. Starting from this QoI we build a Bayesian network that includes expert knowledge (volcano curves), as well as various available experimental and computational data and their correlations or conditional independence. Model uncertainties enter in the construction of the Bayesian network due to lack of complete knowledge of physics and sparse, expensive, multisourced experimental and/or computational data; see, for example, Figure 7(c)–(g).

Both these examples illustrate how PGMs (here Bayesian networks) allow us to (a) organize in a natural way expert knowledge, modeling, heterogeneous and correlated data, and QoIs; and (b) study the propagation of all related model uncertainties to the QoI through the graph. Practically these PGMs are built around the QoI so that it contains all available sources of information that may influence QoI predictions.

B. Mathematical results. In this article, we focus on quantifying the impact of model uncertainties on QoIs in Bayesian networks. Our analysis builds on earlier uncertainty quantification (UQ) information inequalities for QoIs of high-dimensional probabilistic models and stochastic processes [15, 21, 39, 9, 8] (see also Appendix A). However here we demonstrate that, due to the directed nature of the graph structure of Bayesian networks as well as the factorization into CPDs of the joint density given by (1.1), we can obtain a highly detailed and computationally feasible understanding of model uncertainty. First, by defining structured localized ambiguity sets and by using the chain rule of the Kullback–Leibler (KL) divergence, uncertainties can be localized on the graph, while we can quantify their propagation across the graph to the QoIs by (1.1) developing corresponding model uncertainty and model sensitivity indices. Also given a baseline Bayesian network, our uncertainty bounds are proven to be tight within this family of graphical models in the sense that Q^\pm given by (2.6) and (3.9),

respectively, are Bayesian networks as well. The latter finding shows that these worst-case scenarios are “realistic” in the sense that they can materialize by a feasible alternative model that is still a Bayesian network.

First, we refer to an already constructed Bayesian network $\{G, P\}$ as the *baseline model*. We will describe mathematically the model uncertainty of the baseline by considering alternative models Q in a suitably defined neighborhood of P referred to as the *ambiguity set*,

$$(1.2) \quad \mathcal{D}^\eta := \{\text{all Bayesian networks } Q : d(Q, P) \leq \eta\}.$$

The two primary ingredients for constructing ambiguity sets are the choice of a divergence or probabilistic metric d between the baseline Bayesian network P and an alternative model Q and its size η called *model misspecification*, which essentially describes the level of uncertainty in the model. Next, given an ambiguity set \mathcal{D}^η , we define the *model uncertainty indices* for our QoI f as

$$(1.3) \quad I^\pm(f, P; \mathcal{D}^\eta) := \sup/\inf_{Q \in \mathcal{D}^\eta} \{\mathbb{E}_Q[f] - \mathbb{E}_P[f]\}.$$

We view these indices as a *nonparametric stress test* on the baseline P for the QoI f within the ambiguity set \mathcal{D}^η , since they provide the corresponding worst-case scenarios. Furthermore, the ambiguity set is nonparametric, allowing us to test the *robustness* of the baseline against a broader set of scenarios than just a fixed parametric family.

Here we will define ambiguity sets using the KL divergence as it allows us to obtain easily computable and scalable model uncertainty indices $I^\pm(f, P; \mathcal{D}^\eta)$. Indeed, the KL chain rule allows us to break down the calculation of any KL distance between different Bayesian network components, i.e., in terms of conditional KL divergences between distinct CPDs, as well as to isolate the uncertainty impact on QoIs from multiple Bayesian network components and data sources. A discussion on other natural choices of divergences and metrics can be found in section 8. On the other hand, the model misspecification η can be selected in two ways. First, by the user adjusting the stress test on the QoI, for example, when available data are too sparse or absent. Otherwise, η can be estimated as the KL divergence between the model and the available data. Thus, we can consider user-determined or data-informed stress tests, respectively.

Next, we design different stress tests by adjusting the ambiguity set (1.2) to account for global or local perturbations/uncertainties of the baseline model (1.1).

Model uncertainty indices (perturbing the entire model). Let $f(X_A)$ be a QoI defined on any set of random variables $X_A \subset X_V$. The ambiguity set (1.2) in this case contains all the possible alternative Bayesian networks Q η -close to the baseline Bayesian network P in the KL divergence for some model misspecification η . In [Theorem 2.1](#), we demonstrate that the model uncertainty indices for $f(X_A)$ (1.3) can be rewritten only in terms of the baseline P through the one-dimensional optimization

$$(1.4) \quad I^\pm(f(X_A), P; \mathcal{D}^\eta) = \pm \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_A} \left[e^{\pm c f(X_A)} \right] + \frac{\eta}{c} \right] = \mathbb{E}_{Q^\pm} [f] - \mathbb{E}_P [f],$$

where P_A is the marginal distribution of X_A and $\bar{f}(X_A)$ is the centered QoI with respect to P . Furthermore, there exist optimizers Q^\pm (last equality in (1.4)) that are Bayesian networks that can be computed explicitly. We note that although the optimization in (1.3) is infinite-dimensional and thus essentially computationally intractable, formula (1.4) gives rise to a computable one-dimensional optimization involving only the baseline Bayesian network P . This significant advantage will be exploited throughout the paper to provide practical quantification of model uncertainty and model sensitivity for PGMs.

Next we quantify the robustness of the baseline against perturbations of individual components of (1.1). We intend to use these methods to explore the relative sensitivity of the baseline on different Bayesian network components; hence we will refer to the corresponding indices as model sensitivity indices.

Model sensitivity indices (perturbing a model component). Let $f(X_k)$ be a QoI with $k \in V$. We examine two ambiguity sets depending on the manner in which individual model components are perturbed. The first ambiguity set consists of all Bayesian networks (1.1) with the same CPDs except for the CPD at a specific vertex $l \in V$; the structure/parents of the component l can be different; however, the alternative CPDs are η_l -“close” to P at the l th component in KL divergence for some model misspecification η_l ; see Figure 2. The second ambiguity set consists of all Bayesian networks with the parents of the vertex l being fixed and only the CPD of l varying. Even if the latter set is a subset of the first ambiguity set, such graph-based constraints allow us to focus on uncertainties arising from a given CPD of the network. For these ambiguity sets, we derive explicit formulas for the corresponding sensitivity indices that are tight and practically computable similarly to (1.4); see Theorems 3.2 and 3.3.

C. Model sensitivity for ranking and correctability. Model sensitivity indices are used here to rank the impact of different sources of uncertainty, from least to most influential, in the prediction of QoIs for Bayesian networks. From a machine learning perspective, such rankings are a systematic form of *interpretability*, i.e., the ability to identify cause and effect in a model, [19, 52, 13], and *explainability*, i.e., the ability to explain model outputs through the modeling and data choices made in the construction of the baseline predictive model; see [1] and references therein. In the ORR model discussed earlier, we compute model misspecification parameters η_l from data, we implement the ranking procedure for each graph component of the ORR Bayesian network (i.e., solvation, DFT, experiment and parameter correlation), and we reveal the least and the most influential parts of the Bayesian network in the prediction of the optimal oxygen binding energy QoI; see Figure 9 and section 7.

Lastly we leverage model uncertainty and model sensitivity indices to improve the baseline with either targeted data acquisition or improved modeling of CPDs and graph G in (1.1). We target for correction underperforming components of the baseline, i.e., those inducing the most uncertainty on the QoI in the ranking above. Again from a machine learning perspective such a strategy is a step towards the *correctability* of PGMs, namely, the ability to correct targeted components of a (baseline) model without creating new errors in other parts of the model in the process [1]. Indeed, in the ORR model, we correct the baseline Bayesian network in two distinct ways: by including targeted high quality data and by increasing the model complexity, e.g., considering richer CPD classes or more complex PGMs, as discussed in section 7. This is an example of closing the “data-model-predictions loop” by iteratively

improving the model while taking into account trade-offs between model complexity, data, and predictive guarantees on QoIs.

D. Related work. The robust perspective in (1.3) for general probabilistic models is known in the operations research literature as *distributionally robust optimization* (DRO) and includes different choices for divergences or metrics in (1.2); see, for example, [17, 33, 74, 44, 28, 49, 53, 75, 12]. Related work is also encountered in macroeconomics; we refer the reader to the book by Hansen and Sargent [41]. Stress testing via a DRO perspective was developed in the context of insurance risk analysis in [11]. Finally, [57] and [37] developed robust uncertainty quantification methods using different combinations of concentration inequalities and/or information divergences. Regarding sensitivity analysis, we note that existing methods, e.g., gradient and ANOVA-based methods [65], are suitable for parametric uncertainties, and thus cannot handle model uncertainty. Furthermore, it is not immediately obvious how they can take advantage of the graphical structure in Bayesian networks such as conditional independence. Here, our mathematical methods broadly rely on UQ information inequalities for QoIs of high-dimensional probabilistic models and stochastic processes [15, 21, 39, 9, 8] (see also Appendix A). The mathematical novelty of our results lies in extending these earlier works on directed graphs by developing the proper model uncertainty and model sensitivity framework for general Bayesian networks and studying the propagation of multiple uncertainties through the network to the QoIs. Earlier work on building a predictive chemistry-based PGM with quantified model uncertainty for the resulting Gaussian Bayesian network was carried out in [25] and demonstrated in materials design for speeding up the oxygen reduction reaction in fuel cells. Model uncertainty for PGMs over undirected graphs, also known as Markov random fields (MRFs), was recently studied in [5]. An MRF is a unifying model for statistical mechanics (Gibbs measures) and machine learning (Boltzmann machines), while the special case of Gibbs measures was studied earlier in [47]. We note that for MRFs the robust perspective is less flexible as we cannot fully take advantage of the KL chain rule due to the undirected structure of the graphs.

E. Organization. The main mathematical results are presented in section 2 (model uncertainty) and section 3 (model sensitivity). In sections 4 and section 5, we discuss ranking and correctability for Bayesian networks. In section 6 we discuss a DFT-informed Langmuir model, while in section 7 we analyze the ORR model arising in fuel cells. In section 8 we discuss some outstanding issues and directions. Supporting material is included in the appendices.

2. Model uncertainty indices for Bayesian networks. In this section, we develop model uncertainty methods and associated indices for Bayesian networks. We start with the key ingredients needed to state and prove the main result (Theorem 2.1), namely, the ambiguity set, QoIs, and the definition of the model uncertainty indices. First we define the ambiguity set with model misspecification η by

$$(2.1) \quad \mathcal{D}^\eta := \{\text{all PGMs } Q : R(Q\|P) \leq \eta\},$$

where $R(Q\|P) = \mathbb{E}_Q \left[\log \frac{dQ}{dP} \right]$ denotes the KL divergence, i.e., we perturb the baseline model P to any alternative model $Q \in \mathcal{D}^\eta$, altering both the structure of the graph and the CPDs. Examples of models Q included in \mathcal{D}^η can be Bayesian networks defined on a smaller number of vertices than P , or same number of vertices with some of them having extra parents, or




Figure 1. (a) Example of the graph structure of a baseline Bayesian network P and the corresponding random variables $X = \{X_1, \dots, X_8\}$. (b) Example of the graph structure of a Bayesian network $Q \in \mathcal{D}^n$ defined on a set with one vertex less than the baseline Bayesian network P . (c) An example of an alternative Bayesian network $Q \in \mathcal{D}^n$ with the same number of vertices while X_2 and X_8 have extra new parents (in yellow). (d) An example of a PGM $Q \in \mathcal{D}^n$ with a new undirected edge (in blue).

same number of vertices and parents but different CPDs; see Figure 1(b)–(c). Furthermore, \mathcal{D}^n can include PGMs which are not necessarily Bayesian networks, for example, when some of the edges between vertices are not directed [31]; see Figure 1(d). In particular, in Appendix C, we construct an example of a *mixed* graphical model, i.e., with both directed and undirected edges; cf. Figure 1(d), which belongs to ambiguity set (2.1) when P is a Gaussian Bayesian network (see the definition in section 2.1). For a baseline Bayesian network P we define the model uncertainty indices as

$$(2.2) \quad I^\pm(f(X_A), P; \mathcal{D}^n) = \sup_{Q \in \mathcal{D}^n} / \inf_{Q \in \mathcal{D}^n} \mathbb{E}_Q[f(X_A)] - \mathbb{E}_P[f(X_A)]$$

for a QoI f , which is a function of some subset of A vertices in the graph, i.e.,

$$(2.3) \quad \text{for } f(X_A) = f(X_{i_1}, \dots, X_{i_m}), \text{ with } A = \{i_1, \dots, i_m\} \subseteq V.$$

In the next theorem, we characterize the optimizers Q^\pm in (2.2), which turn out to be Bayesian networks of the form (1.1), and we provide their CPDs explicitly.

Notation. Before we state our results let us fix some notation. For a Bayesian network, $\{G, P\}\pi_i^P$ denotes the set of indices of all the parents of vertex i and ρ_i^P denotes the set of indices of all the ancestors for i (we may omit the superscript “ P ” if only one Bayesian network is involved). Without loss of generality we assume that all Bayesian networks are topologically ordered as we can always relabel the DAG so that $j < i$ for all $j \in \pi_i$ by topological sorting [48].

The random vector $X = (X_1, \dots, X_n)$, indexed by the vertices $V = \{1, \dots, n\}$, takes values $X = x = (x_1, \dots, x_n) \in \mathcal{X}$. The joint probability distribution of X is denoted by P with density $p(x)$; the results are presented when the joint probability distribution P is continuous, but all results hold when p is a discrete distribution as well. For any subset $A = \{i_1, \dots, i_m\} \subset V$ we denote $X_A = (X_{i_1}, \dots, X_{i_m})$, which takes values $X_A = x_A = (x_{i_1}, \dots, x_{i_m}) \in \mathcal{X}_A$ and we denote by P_A its marginal distribution.

We denote by $P_{i|\pi_i^P}$ the conditional distribution of X_i given parent values $X_{\pi_i^P} = x_{\pi_i}$, i.e., $P_{i|\pi_i^P}(dx_i) = P(dx_i|x_{\pi_i})$ with corresponding CPD $p(x_i|x_{\pi_i})$. In the Bayesian network literature it is always implicitly assumed that all joints and marginals exist. Thus, we consider the next marginals. The marginal distribution P_A of X_A has the form $P_A(dx_A) = \int_{\mathcal{X}_{\rho_A}} \prod_{i \in A} P_{i|\pi_i^P}(dx_i) \prod_{i \in \rho_A^P} P_{i|\pi_i^P}(dx_i)$, where ρ_A^P is the set of indices of all the ancestors of A ,

i.e., $\rho_A^P = \cup_{i \in A} \rho_i^P$. Furthermore the density p_A of P_A is $p_A(x_A) = \int_{\mathcal{X}_{\rho_A^P}} \prod_{i \in A} p(x_i | x_{\pi_i}) \prod_{i \in \rho_A^P} p(x_i | x_{\pi_i}) dx_{\rho_A^P}$. Two special cases are the marginals of X_k , $P_{\{k\}}(dx_k) = \int_{x_{\rho_k}} \prod_{i \in \{k \cup \rho_k\}} P(dx_i | x_{\pi_i})$ and the marginal of X_{ρ_A} , $P_{\rho_A}(dx_{\rho_A}) = \prod_{i \in \rho_A} P(dx_i | x_{\pi_i})$.

Finally for $l_1 < \dots < l_k$ and any QoI f and for $j \in \{1, \dots, k\}$ we define the notation

$$(2.4) \quad \mathbb{E}_{P_{l_j | \pi_{l_j}^P, \dots, P_{l_k} | \pi_{l_k}^P}} [f] := \mathbb{E}_{P_{l_j | \pi_{l_j}^P}} \left[\mathbb{E}_{P_{l_{j+1} | \pi_{l_{j+1}}^P}} \left[\dots \mathbb{E}_{P_{l_k | \pi_{l_k}^P}} [f] \right] \right].$$

Theorem 2.1. *Let $\{G, P\}$ be a Bayesian network with density defined as in (1.1), and let $f(X_A)$ be a QoI given in (2.3), $f(X_A) = f(X_{i_1}, \dots, X_{i_m})$. Let also $\bar{f}(X_A) := f(X_A) - \mathbb{E}_P[f(X_A)]$ be the centered QoI with finite moment generating function (MGF), $\mathbb{E}_P[e^{c\bar{f}(X_A)}]$, in a neighborhood of the origin.*

(a) **Tightness.** *For the model uncertainty indices defined in (2.2), there exist $0 < \eta_{\pm} \leq \infty$, such that for any $\eta \leq \eta_{\pm}$,*

$$(2.5) \quad \begin{aligned} I^{\pm}(f(X_A), P; \mathcal{D}^{\eta}) &= \pm \inf_{c > 0} \left[\frac{1}{c} \log \mathbb{E}_{P_A} \left[\pm e^{c\bar{f}(X_A)} \right] + \frac{\eta}{c} \right] \\ &= \mathbb{E}_{Q^{\pm}} [f(X_A)] - \mathbb{E}_P [f(X_A)], \end{aligned}$$

where P_A is the marginal distribution of X_A with respect to P and is defined in the notation above, and $Q^{\pm}(\cdot) \equiv Q^{\pm}(\cdot; \pm c_{\pm}) \in \mathcal{D}^{\eta}$ are Bayesian networks (1.1) that depend on η and are given by

$$(2.6) \quad \frac{dQ^{\pm}}{dP} = \frac{e^{\pm c_{\pm} f(x_A)}}{\mathbb{E}_P [e^{\pm c_{\pm} f(X_A)}]},$$

where $c_{\pm} \equiv c_{\pm}(\eta)$ are the unique solutions of the equation

$$(2.7) \quad R(Q^{\pm} \| P) = \eta.$$

(b) **Graph Structure of Q^{\pm} .** *Let L be all vertices that include A and all its ancestors, i.e., $L = \cup_{j \in A} \rho_j^P \cup A = \{l_1, \dots, l_{k+1}\}$, where $l_1 < \dots < l_{k+1}$ and $l_{k+1} = i_m$. Then the CPDs of Q^{\pm} are given by*

$$(2.8) \quad q^{\pm}(x_i | x_{\pi_i^{Q^{\pm}}}) = \begin{cases} p(x_i | x_{\pi_i^P}), & i \notin L, \\ \frac{\mathbb{E}_{P_{l_{j+1} | \pi_{l_{j+1}}^P, \dots, P_{l_k} | \pi_{l_k}^P}} [e^{\pm c_{\pm} f(X_A)}]}{\mathbb{E}_{P_{l_j | \pi_{l_j}^P, \dots, P_{l_k} | \pi_{l_k}^P}} [e^{\pm c_{\pm} f(X_A)}]} p(x_{l_j} | x_{\pi_{l_j}^P}), & i = l_j, j \in A, \end{cases}$$

where $l_j \in L$ and $\mathbb{E}_{P_{l_j | \pi_{l_j}^P, \dots, P_{l_k} | \pi_{l_k}^P}}$ is given by (2.4). The parents/structure is given by $\pi_i^{Q^{\pm}} \equiv \pi_i^P, i \notin L$, and $\pi_{l_j}^{Q^{\pm}} \subset \pi_{l_j}^P \cup (\pi_{l_{j+1}}^{Q^{\pm}} \setminus l_j)$.

Remark 2.2. **Theorem 2.1** readily implies that we can severely restrict the ambiguity set (2.1) to a subclass of Bayesian networks, yielding the exact same index,

$$(2.9) \quad I^{\pm}(f(X_A), P; \mathcal{D}^{\eta}) = I^{\pm}(f(X_A), P; \mathcal{D}_{\rho_A}^{\eta}),$$

where a set of Bayesian networks $\mathcal{D}_{\rho_A}^\eta$ is defined as

$$(2.10) \quad \mathcal{D}_{\rho_A}^\eta := \left\{ \begin{array}{l} \text{all Bayesian networks } Q : R(Q_{A \cup \rho_A} \| P_{A \cup \rho_A}) \leq \eta \\ \text{and } q(x_i | x_{\pi_i^Q}) \equiv p(x_i | x_{\pi_i^P}) \text{ with } \pi_i^Q \equiv \pi_i^P \text{ for all } i \notin A \cup \rho_A \end{array} \right\}.$$

This follows from the formula

$$\mathbb{E}_P [f(X_A)] = \int_{\mathcal{X}} f(x_A) \prod_{i=1}^n p(x_i | x_{\pi_i}) dx = \int_{\mathcal{X}_{A \cup \rho_A}} f(x_A) \prod_{x_i \in A \cup \rho_A} p(x_i | x_{\pi_i}) dx_A dx_{\rho_A},$$

which implies that only the perturbation of $P_{A \cup \rho_A}$ affects the prediction of the QoI. A similar calculation for the MGF of \bar{f} implies that the optimal Q^\pm has the same CPDs as P for all $X_i, i \notin \{A\} \cup \rho_A$, as shown in [Theorem 2.1](#).

Remark 2.3. We illustrate [Theorem 2.1](#) in the special case $A = \{k\}$, i.e., QoIs defined on one vertex through the example in [Figure 2](#); see also [Appendix D](#) for more details.

Remark 2.4. We note that given a QoI f , there is a natural parametric family within the nonparametric ambiguity set (2.1), given by $dP^c \propto e^{cf} dP$, that contains the baseline and the extreme probability distributions Q^\pm ; see the blue curve in [Figure 12](#) (left).

Proof of Theorem 2.1.

- (a.) The existence of Q^\pm and (2.6) is a direct consequence of (A.3) with $f(X) = f(X_A)$. For $p(x) = \prod_{i=1}^n p(x_i | x_{\pi_i})$, we further compute

$$(2.11) \quad \begin{aligned} \sup/\inf_{Q \in \mathcal{D}^\eta} \mathbb{E}_Q [f(X_A)] - \mathbb{E}_P [f(X_A)] &= \pm \inf_{c > 0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{\pm c \bar{f}(X_A)} \right] + \frac{\eta}{c} \right] \\ &= \pm \inf_{c > 0} \left[\frac{1}{c} \log \int_{\mathcal{X}} e^{\pm c \bar{f}(X_A)} \prod_{i=1}^n P(dx_i | x_{\pi_i}^P) + \frac{\eta}{c} \right] \\ &= \pm \inf_{c > 0} \left[\frac{1}{c} \log \mathbb{E}_{P_A} \left[e^{\pm c \bar{f}(X_A)} \right] + \frac{\eta}{c} \right], \end{aligned}$$

where p_A is given in the notation before the theorem.




Figure 2. (a) Example of graph structure of a baseline Bayesian network P . (b) The structure of the optimizers Q^\pm in [Theorem 2.1](#) (b) with QoI $f(X_6)$ is highlighted (in green). In contrast to the CPDs of the vertices involved in the QoI and their ancestors, the CPD of any other vertex does not change. The new parents of X_4, X_2 are connected (in yellow), i.e., X_3 is a new parent for X_4 and X_1 is a new parent for X_2 . (c) Structure of the optimizers Q^\pm in [Theorem 2.1](#) (b) for a QoI of the type $f(X_3, X_6, X_7)$; see (B.6)–(B.13).

(b.) We use (2.6) and factorize q^\pm as follows:

$$\begin{aligned}
 q^\pm(x) &= \frac{e^{\pm c_\pm f(x_A)}}{\mathbb{E}_P [e^{\pm c_\pm f(X_A)}]} \prod_{i=1}^n p(x_i | x_{\pi_i^P}) \\
 &= \frac{1}{\mathbb{E}_{P_A} [e^{\pm c_\pm f(X_A)}]} \prod_{i \notin \{l_1, \dots, l_{k+1}\}} p(x_i | x_{\pi_i^P}) \cdot e^{\pm c_\pm f(x_A)} p(x_{i_m} | x_{\pi_{i_m}^P}) \\
 (2.12) \quad &\times \prod_{i \in \{l_1, \dots, l_k\}} p(x_i | x_{\pi_i^P}),
 \end{aligned}$$

where $\pm c_\pm$ are the unique solutions of $R(P^{\pm c_\pm} \| P) = \eta$. Formula (2.12) is not factorized yet into CPDs as in (1.1) due to the normalization factor at the denominator. The following analysis provides the steps for expressing (2.12) in a product of certain CPDs: Assuming that $i_1 < \dots < i_m$, we start with the CPD of X_{i_m} as its index is the largest among the elements of A . Based on (2.12),

$$(2.13) \quad q^\pm(x_{i_m} | x_{\pi_{i_m}^{Q^\pm}}) \propto e^{\pm c_\pm f(x_A)} p(x_{i_m} | x_{\pi_{i_m}^P}).$$

We normalize the left-hand side (LHS) of (2.12) by dividing by

$$(2.14) \quad \mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}]$$

and by conditioning to $x_{\pi_{i_m}^P}$ and $x_{A \setminus i_m}$. Therefore, the CPD of X_{i_m} and its parents $\pi_{i_m}^{Q^\pm}$ are given by

$$q^\pm(x_{i_m} | x_{\pi_{i_m}^{Q^\pm}}) = \frac{e^{\pm c_\pm f(x_A)}}{\mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}]} \cdot p(x_{i_m} | x_{\pi_{i_m}^P}) \text{ and } \pi_{i_m}^{Q^\pm} \subset \pi_{i_m}^{Q^\pm} = \pi_{i_m}^P \cup (A \setminus i_m).$$

Such a consideration provides the new edges in the graph of Q^\pm . In particular, X_{i_m} has the same parents as in the P model and possibly new parents specified by $x_{A \setminus i_m}$, e.g., if $A \setminus i_m \neq \pi_{i_m}^P$. Next, we compute the CPD of X_{l_k} since $l_k < l_{k+1} = i_m$: As we divided by (2.14) to normalize the LHS of (2.12), we keep (2.12) the same if we also multiple $q^\pm(x)$ by (2.14). Hence, we pair (2.12) and $p(x_{l_k} | x_{\pi_{l_k}^P})$ so that

$$(2.15) \quad q^\pm(x_{l_k} | x_{\pi_{l_k}^{Q^\pm}}) \propto \mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}] p(x_{l_k} | x_{\pi_{l_k}^P}).$$

As before, we normalize the LHS of (2.15) by dividing by

$$(2.16) \quad \mathbb{E}_{P_{l_k | \pi_{l_k}^P}} \left[\mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}] \right],$$

and by conditioning to $x_{\pi_{l_k}^P}$ and $x_{\pi_{i_m}^{Q^\pm} \setminus l_k}$, we obtain

$$q^\pm(x_{l_k} | x_{\pi_{l_k}^{Q^\pm}}) = \frac{\mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{l_k | \pi_{l_k}^P}} \left[\mathbb{E}_{P_{i_m | \pi_{i_m}^P}} [e^{\pm c_\pm f(X_A)}] \right]} \cdot p(x_{i_m} | x_{\pi_{i_m}^P})$$

and $\pi_{l_k}^{Q^\pm} \subset \pi_{l_k}^{Q^\pm} = \pi_{l_k}^P \cup (\pi_{i_m}^{Q^\pm} \setminus l_k)$. The latter shows the new edges that the associated graph to Q^\pm may have. In this way, we obtain the remaining CPDs given by the second part of (2.8). It is straightforward that the random variables indexed differently than $\{l_1, \dots, l_{k+1}\}$ inherit the corresponding CPDs of P , and thus (2.8) is obtained. ■

2.1. Gaussian Bayesian networks. In this subsection, we focus on Gaussian Bayesian networks, which is a special class of Bayesian networks commonly used in natural and social sciences with the CPDs as in (1.1) being linear and Gaussian [48, 64, 36, 34, 35]. More specifically, for a Gaussian Bayesian network consisting of variables X , each vertex X_i is a linear Gaussian of its parents, i.e.,

$$(2.17) \quad \begin{aligned} p(x_i|x_{\pi_i}) &= \mathcal{N}(\beta_{i0} + \beta_i^T x_{\pi_i}, \sigma_i^2), \quad \text{equivalently} \\ X_i &= \beta_{i0} + \beta_i^T X_{\pi_i} + \epsilon_i, \quad \text{with } \epsilon_i \sim \mathcal{N}(0, \sigma_i^2) \end{aligned}$$

for some β_0 , σ_i , and $\beta_i = [\beta_{ii_1}, \dots, \beta_{ii_{|\pi_i|}}]$. By the conjugacy properties of Gaussians, the joint distribution P becomes $p(x) = \mathcal{N}(\mu, \mathcal{C})$, i.e., it is also a Gaussian with parameters μ , \mathcal{C} , which can be calculated from β_{i0} , β_i , and σ_i [10].

Theorem 2.5. *Let P be a Gaussian Bayesian network that satisfies (2.17), and let $f(X_k) = aX_k + b$ be a QoI that depends only on X_k linearly.*

(a) *Then for the model uncertainty indices defined in (2.2), we have*

$$(2.18) \quad I^\pm(f(X_k), P; \mathcal{D}^\eta) = \pm \sqrt{2a^2 \mathcal{C}_{kk} \eta},$$

where \mathcal{C}_{kk} is the variance for the marginal distribution of X_k .

(b) *Furthermore, the optimizers $Q^\pm = Q^\pm(\eta) \in \mathcal{D}^\eta$ are given by (2.8) in Theorem 2.1 and are also Gaussian Bayesian networks with the same graph structure as P .*

Proof.

(a.) The distribution of X_k denoted by $P_{\{k\}}$ is Gaussian with variance

$$\mathcal{C}_{kk} = \sigma_k^2 + \beta_k^T \mathcal{C}_{\rho_k} \beta_k,$$

where \mathcal{C}_{ρ_k} is the variance of the joint distribution of the random variables $\{X_i : i \in \rho_k\}$ [49, Theorem 7.3]. By a straightforward computation, the MGF of $\bar{f}(X_k)$ is given by

$$(2.19) \quad \begin{aligned} \mathbb{E}_{P_{\{k\}}} \left[e^{\pm c \bar{f}(X_k)} \right] &= \exp(a^2 c^2 \beta_k^T \mathcal{C}_{\rho_k} \beta_k), \\ I^\pm(f(X_k), P; \mathcal{D}^\eta) &= \pm \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{\{k\}}} \left[e^{\pm c \bar{f}(X_k)} \right] + \frac{\eta}{c} \right] = \pm \inf_{c>0} \left[a^2 c \mathcal{C}_{kk} + \frac{\eta}{c} \right]. \end{aligned}$$

Then the optimal c is given by $c = \sqrt{\frac{\eta}{a^2 \mathcal{C}_{kk}}}$, which in turn proves (2.18).

(b.) Next, we show that the graph structure of Q^\pm is the same as that of P . For any $j > k$, by Theorem 2.1, $q(x_j|x_{\pi_j^{Q^\pm}}) = p(x_j|x_{\pi_j^P})$. For $j = k$, we compute

$$(2.20) \quad \begin{aligned} q^\pm(x_k|x_{\pi_k^{Q^\pm}}) &= \frac{e^{\pm c_\pm f(x_k)}}{\mathbb{E}_{P_{k|\pi_k^P}} \left[e^{\pm c_\pm f(X_k)} \right]} \cdot p(x_k|x_{\pi_k^P}) \\ &= \frac{\exp \left\{ -\frac{(x_k - \beta_{k0} - \beta_k^T x_{\pi_k^P} \mp c_\pm a \sigma_k^2)^2}{2\sigma_k^2} \pm c_\pm a (\beta_k^T x_{\pi_k^P}) \right\}}{\int_{\mathcal{X}_k} \exp \left\{ -\frac{(x_k - \beta_{k0} - \beta_k^T x_{\pi_k^P} \mp c_\pm a \sigma_k^2)^2}{2\sigma_k^2} \pm c_\pm a (\beta_k^T x_{\pi_k^P}) \right\} dx_k} \\ &= \mathcal{N}(\beta_{k0} + \beta_k^T x_{\pi_k^P} \pm c_\pm a \sigma_k^2, \sigma_k^2). \end{aligned}$$

Thus $\pi_k^{Q^\pm} = \pi_k^P$ since $c_\pm a(\beta_k^T x_{\pi_k})$ of the numerator and denominator are canceled out. Let k_m be the maximum element of $\pi_k^P = \{k_1, \dots, k_m : k_1 < \dots < k_{m-1} < k_m\}$ and $\beta_k^{m-1} := [\beta_{kk_1}, \dots, \beta_{kk_{m-1}}]$. Then

(2.21)

$$\begin{aligned}
 q^\pm(x_{k_m} | x_{\pi_{k_m}^{Q^\pm}}) &= \frac{\mathbb{E}_{P_{k|\pi_k^P}} [e^{\pm c_\pm f(X_k)}]}{\mathbb{E}_{P_{k_m|\pi_{k_m}^P}} [\mathbb{E}_{P_{k|\pi_k^P}} [e^{\pm c_\pm f(X_k)}]]} \cdot p(x_{k_m} | x_{\pi_{k_m}^P}) \\
 &= \frac{\exp \left\{ \pm c_\pm a \beta_k^T x_{\pi_k^P} - \frac{(x_{k_m} - \beta_{k_m 0} - \beta_{k_m}^T x_{\pi_{k_m}^P})^2}{2\sigma_m^2} \right\}}{\int_{\mathcal{X}_{k_m}} \exp \left\{ \pm c_\pm a \beta_k^T x_{\pi_k^P} - \frac{(x_{k_m} - \beta_{k_m 0} - \beta_{k_m}^T x_{\pi_{k_m}^P})^2}{2\sigma_{k_m}^2} \right\} dx_{k_m}} \\
 &= \frac{\exp \left\{ \pm c_\pm a (\beta_k^{m-1})^T x_{\pi_{k_m}^P \setminus k_m} - \frac{(x_{k_m} - \beta_{k_m 0} - \beta_{k_m}^T x_{\pi_{k_m}^P} \mp c_\pm a \beta_{kk_m} \sigma_{k_m}^2)^2}{2\sigma_{k_m}^2} \right\}}{\int_{\mathcal{X}_{k_m}} \exp \left\{ \pm c_\pm a (\beta_k^{m-1})^T x_{\pi_{k_m}^P \setminus k_m} - \frac{(x_{k_m} - \beta_{k_m 0} - \beta_{k_m}^T x_{\pi_{k_m}^P} \mp c_\pm a \beta_{kk_m} \sigma_{k_m}^2)^2}{2\sigma_{k_m}^2} \right\} dx_{k_m}} \\
 &= \mathcal{N} \left(\beta_{k_m 0} + \beta_{k_m}^T x_{\pi_{k_m}^P} \pm c_\pm a \beta_{kk_m} \sigma_{k_m}^2, \sigma_{k_m}^2 \right).
 \end{aligned}$$

Again, $\pi_{k_m}^{Q^\pm} = \pi_{k_m}^P$ as the factor $\exp \{ \pm c_\pm a (\beta_k^{m-1})^T x_{\pi_{k_m}^P \setminus k_m} \}$ in the numerator and denominator are canceled out. The CPD of the remaining vertices in π_k^P are computed in the same way, which further implies that their parents do not change. Therefore, the factors in CPDs of Q^\pm that could create new directed edges appear in both numerator and denominator and are finally canceled out. We demonstrate (2.20) and (2.21) as it applies in Example D.1 in Appendix D. ■

3. Model sensitivity indices for Bayesian networks. In this section, we develop a non-parametric sensitivity analysis for Bayesian networks by refining the concepts of model uncertainty indices introduced in section 2. This is accomplished through designing localized ambiguity sets suitable for model uncertainty/perturbations in specific components of the graphical model such as a single CPD.

Notation. For the notation of this section, we refer the reader to section 2. Moreover, we denote $\bar{\rho}_k^P := \rho_k^P \cup \{k\}$.

Let $f = f(X_k)$ be a QoI depending only on vertex $k \in V$, and let $l \in V$ be another vertex. The first ambiguity set $\mathcal{D}_l^{\eta_l}$ consists of all Bayesian networks Q that differ from the baseline P only in the CPD at the vertex l while also allowing for the parents π_l^P at l to change. Namely,

$$(3.1) \quad \mathcal{D}_l^{\eta_l} = \left\{ \begin{array}{l} \text{all Bayesian networks } Q : R(Q_{l|\pi_l^Q} \| P_{l|\pi_l^P}) \leq \eta_l \text{ for all } x_{\pi_l^P} \cup x_{\pi_l^Q}, \\ Q_{j|\pi_j} \equiv P_{j|\pi_j} \text{ for all } j \neq l \end{array} \right\},$$

where the parents π_l^Q in model Q may differ from the parents π_l^P in model P .

The second ambiguity set $\mathcal{D}_{l,P}^{\eta_l}$ consists of all Bayesian networks Q that differ from the baseline P only in the CPD at the vertex l ; however, here we require that $\pi_l^Q = \pi_l^P = \pi_l$, i.e., parents are not allowed to change:

$$(3.2) \quad \mathcal{D}_{l,P}^{\eta_l} = \{ \text{all Bayesian networks } Q : R(Q_{l|\pi_l} \| P_{l|\pi_l}) \leq \eta_l \text{ for all } x_{\pi_l}, Q_{j|\pi_j} \equiv P_{j|\pi_j} \text{ for all } j \neq l \}.$$

Note that

$$(3.3) \quad \mathcal{D}_{l,P}^{\eta_l} \subset \mathcal{D}_l^{\eta_l}.$$

We accordingly define the model sensitivity indices of the QoI $f(X_k)$ as

$$(3.4) \quad I^{\pm}(f(X_k), P; \mathcal{Q}_{\eta_l}) = \sup_{Q \in \mathcal{Q}_{\eta_l}} / \inf_{Q \in \mathcal{Q}_{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)],$$

where $\mathcal{Q}_{\eta_l} = \mathcal{D}_l^{\eta_l}$ or $\mathcal{D}_{l,P}^{\eta_l}$ given by (3.1) and (3.2), respectively.

The evaluation of these model sensitivity indices will necessarily depend on the relative graph position of vertices $k, l \in V$ and in particular if l is an ancestor of k . In particular we have the following:

Lemma 3.1. *Let $Q \in \mathcal{Q}_{\eta_l}$, where $\mathcal{Q}_{\eta_l} = \mathcal{D}_l^{\eta_l}$ or $\mathcal{D}_{l,P}^{\eta_l}$. Then*

$$(3.5) \quad \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] = \begin{cases} \mathbb{E}_{P_{\rho_l^P}} \left[\mathbb{E}_{Q_{l|\pi_l^Q}} [F] - \mathbb{E}_{P_{l|\pi_l^P}} [F] \right], & l \in \bar{\rho}_k^P, \\ 0, & l \notin \bar{\rho}_k^P, \end{cases}$$

where

$$(3.6) \quad \begin{aligned} F &:= F(x_l, x_{\rho_l^P}) = \int_{\mathcal{X}_{\bar{\rho}_k^P \setminus \rho_l^P \cup \{l\}}} f(x_k) \prod_{i \in \bar{\rho}_k^P \setminus \rho_l^P \cup \{l\}} P(dx_i | x_{\pi_i^P}) \\ &= \mathbb{E}_{P_{\{k\}|\bar{\rho}_l^P}} [f(X_k)], \end{aligned}$$

and the last expectation is with respect to the conditional distribution of X_k given by $X_{\bar{\rho}_l^P} = x_{\bar{\rho}_l^P}$.

The proof of Lemma 3.1 is a direct calculation of the difference between the expectations of $f(X_k)$ and is based on a rearrangement between the CPDs of $X_{\rho_k^P \cup \{k\}}$, $X_{\rho_l^P}$, and X_l with respect to P and Q (see Appendix F), while a concrete computation of F is given in Appendix D.2 for the Bayesian network of Example D.1.

Next, following the structure of Theorem 2.1 and using Lemma 3.1, we present our results on tightness and optimal distributions over $\mathcal{D}_l^{\eta_l}$ and $\mathcal{D}_{l,P}^{\eta_l}$ as stated in Theorems 3.2 and 3.3, respectively. Theorem 3.3 could be thought of as a subcase of Theorem 3.2 due to (3.3); however, tightness on $\mathcal{D}_{l,P}^{\eta_l}$ cannot be accomplished unless the additional condition (3.19) is assumed. All these results are summarized in a schematic in Figure 15.

Theorem 3.2 (model sensitivity indices-vary graph structure and CPD). *Let P be a Bayesian network with density defined as in (1.1), and let $f(X_k)$ be a QoI that only depends on X_k . Let also $\bar{f}(X_k)$ be the centered QoI with finite moment generating function (MGF), $\mathbb{E}_P \left[e^{c\bar{f}(X_k)} \right]$, in a neighborhood of the origin.*

(a) **Tightness.** For the model sensitivity indices defined in (3.4), there exist $0 < \eta_{\pm} \leq \infty$, such that for any $\eta \leq \eta_{\pm}$,

$$\begin{aligned}
 (3.7) \quad I^{\pm}(f(X_k), P; \mathcal{D}_l^{\eta}) &= \sup_{Q \in \mathcal{D}_l^{\eta}} / \inf \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] \\
 &= \begin{cases} \pm \mathbb{E}_{P_{\rho_l^P}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|\pi_l^P}} \left[e^{\pm c \bar{F}} \right] + \frac{\eta}{c} \right] \right], & l \in \bar{\rho}_k^P, \\ 0, & l \notin \bar{\rho}_k^P, \end{cases} \\
 &= \mathbb{E}_{Q^{\pm}} [f(X_k)] - \mathbb{E}_P [f(X_k)],
 \end{aligned}$$

where \bar{F} is the centered function of F defined in (3.6), $\eta_l \equiv \eta$, and $Q^{\pm}(\cdot) \equiv Q^{\pm}(\cdot; \pm c_{\pm}) \in \mathcal{D}_l^{\eta}$ are Bayesian networks of the form (1.1) that depend on η_l with $f(X_k)$ and $c_{\pm} \equiv c_{\pm}(x_{\rho_l^P}; \eta)$ being functions of $x_{\rho_l^P}$, depending on η_l , and determined by the equations

$$(3.8) \quad R(Q_{l|\pi_l^{Q^{\pm}}} \| P_{l|\pi_l^P}) = \eta_l.$$

(b) **Graph Structure of Q^{\pm} .** The optimal distributions Q^{\pm} are the probability measures with densities given by

$$(3.9) \quad q^{\pm}(x_i | x_{\pi_i^{Q^{\pm}}}) = \begin{cases} p(x_i | x_{\pi_i^P}), & i \neq l, \\ \frac{e^{\pm c_{\pm} F(x_i, x_{\rho_l^P})}}{\mathbb{E}_{P_{l|\pi_l^P}} \left[e^{\pm c_{\pm} F(X_l, x_{\rho_l^P})} \right]} p(x_l | x_{\pi_l^P}), & i = l. \end{cases}$$

The structure of the first and second parts of (3.9) satisfies $\pi_i^{Q^{\pm}} \equiv \pi_i^P$ and $\pi_l^P \subset \pi_l^{Q^{\pm}} \subset \rho_l^P = \rho_l^{Q^{\pm}}$, respectively.

Proof. The proofs of (a) and (b) are worked together and split into two main steps.

Step 1: Model sensitivity indices. For $l \in \bar{\rho}_k^P$, we denote $\pi_l := \pi_l^Q \cup \pi_l^P$, and $\rho_i := \rho_i^Q \cup \rho_i^P$ for all i . We define

$$(3.10) \quad Q(dx_l | x_{\pi_l}) := Q(dx_l | x_{\pi_l^Q}) \text{ for all } x_{\pi_l}, \quad P(dx_l | x_{\pi_l}) := P(dx_l | x_{\pi_l^P}) \text{ for all } x_{\pi_l}.$$

We now use Lemma 3.1 and further bound the right-hand side of the first part of (3.7) as follows:

$$\begin{aligned}
 (3.11) \quad \sup_{Q \in \mathcal{D}_l^{\eta}} \mathbb{E}_{P_{\rho_l}} \left[\mathbb{E}_{Q_{l|\pi_l}} [F] - \mathbb{E}_{P_{l|\pi_l}} [F] \right] &\leq \mathbb{E}_{P_{\rho_l}} \left[\sup_{Q \in \mathcal{D}_l^{\eta}} \mathbb{E}_{Q_{l|\pi_l}} [F] - \mathbb{E}_{P_{l|\pi_l}} [F] \right] \\
 &= \mathbb{E}_{P_{\rho_l}} \left[\sup_{Q_l \in \mathcal{E}_l^{\eta}} \mathbb{E}_{Q_{l|\pi_l}} [F] - \mathbb{E}_{P_{l|\pi_l}} [F] \right],
 \end{aligned}$$

where \mathcal{E}_l^{η} is the ambiguity set for CPDs at l defined as

$$(3.12) \quad \mathcal{E}_l^{\eta} := \{ \text{all CPD } Q_{l|\pi_l} : R(Q_{l|\pi_l} \| P_{l|\pi_l}) \leq \eta_l \text{ for all } x_{\pi_l} = x_{\pi_l^P} \cup x_{\pi_l^Q} \}.$$

By using Lemma A.1, for any given $X_{\rho_l} = x_{\rho_l}$, we have

$$(3.13) \quad \sup_{Q_l \in \mathcal{E}_l^{\eta_l}} \mathbb{E}_{Q_l|\pi_l} [F] - \mathbb{E}_{P_l|\pi_l} [F] \leq \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_l|\pi_l} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right].$$

Hence (3.11) implies that

$$(3.14) \quad \sup_{Q \in \mathcal{D}_l^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] \leq \mathbb{E}_{P_{\rho_l}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_l|\pi_l} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right] \right].$$

Step 2: Tightness of the bounds. As in Theorem 2.1, for any given $x_{\rho_l^P}$, we can consider the conditional measure $P_{l|\rho_l^P}^{c+}$ defined by

$$(3.15) \quad \frac{dP_{l|\rho_l^P}^{c+}}{dP_{l|\pi_l^P}} = \frac{e^{c_+(x_{\rho_l^P})F(x_l, x_{\rho_l^P})}}{\mathbb{E}_{P_{l|\pi_l^P}} \left[e^{c_+(x_{\rho_l^P})F(X_l, x_{\rho_l^P})} \right]},$$

where $c_+(x_{\rho_l^P})$ is a function of $x_{\rho_l^P}$ determined by $R(P_{l|\rho_l^P}^{c+} \| P_{l|\pi_l^P}) = \eta_l$. By using Lemma A.2, we define

$$(3.16) \quad q_l^+(x_l | x_{\pi_l^Q}) := P_{l|\rho_l^P}^{c+} \propto e^{c_+(x_{\rho_l^P})F(x_l, x_{\rho_l^P})} p(x_l | x_{\pi_l^P}) \quad \text{for all } x_{\pi_l^Q}.$$

Note that π_l^{Q+} depends on π_l^P and $F(x_l, x_{\rho_l^P})$, hence $\pi_l^P \subset \pi_l^{Q+} \subset \rho_l^P$ and $\rho_l^{Q+} = \rho_l^P$. Therefore, using the same notation as in Step 1, for $\pi_l = \pi_l^{Q+}$, $\rho_l = \rho_l^{Q+}$, we have

$$(3.17) \quad \mathbb{E}_{Q_l^+|\pi_l} [F] - \mathbb{E}_{P_l|\pi_l} [F] = \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_l|\pi_l} \left[e^{c\bar{F}} \right] + \frac{\eta_l}{c} \right].$$

Furthermore, $R(Q_l^+ \| P_l|\pi_l) \leq \eta_l$ for all x_{π_l} and hence $Q_l^+ \in \mathcal{E}_l^{\eta_l}$. Let $q^+(x) = q_l^+(x_l | x_{\pi_l}) \prod_{i \neq l} p(x_i | x_{\pi_i})$. Then $Q^+ \in \mathcal{D}_l^{\eta_l}$, and

$$\mathbb{E}_{Q^+} [f(X_k)] - \mathbb{E}_P [f(X_k)] = \mathbb{E}_{P_{\rho_l}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_l|\pi_l} \left[e^{c\bar{F}} \right] + \frac{\eta_l}{c} \right] \right],$$

and thus (3.7) is proved. The calculations for $\inf_{Q \in \mathcal{D}_l^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)]$ are similar.

We turn next to the ambiguity set $\mathcal{D}_{l,P}^{\eta_l}$ defined as in (3.2) and its corresponding index. Due to Theorem 3.2 and (3.3), the following uncertainty bound holds for $\mathcal{D}_{l,P}^{\eta_l}$:

$$(3.18) \quad \begin{aligned} I^+(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l}) &= \sup_{Q \in \mathcal{D}_{l,P}^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] \\ &\leq \mathbb{E}_{P_{\rho_l}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_l|\pi_l} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right] \right] \end{aligned}$$

for any $l \in \bar{\rho}_k^P$; see also Figure 15. A similar bound holds for $I^-(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l})$. However, the next theorem provides a condition on the Bayesian network P that implies equality in (3.18); see (3.19) and Figure 4.

Theorem 3.3 (model sensitivity indices-only vary CPD). *Let P be a Bayesian network with density defined as in (1.1), and let $f(X_k)$ be a QoI that only depends on X_k with its centered QoI $\bar{f}(X_k)$ having finite MGF, $\mathbb{E}_P \left[e^{c\bar{f}(X_k)} \right]$, in a neighborhood of the origin.*

- (a) For $l \notin \bar{\rho}_k^P$, $\mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] = 0$ for any $Q \in \mathcal{D}_{l,P}^{\eta_l}$.
- (b) For $l \in \bar{\rho}_k^P$ satisfying the condition

$$(3.19) \quad X_k \perp X_{\rho_l \setminus \pi_l} | X_{\pi_l},$$

i.e., X_k is independent of all the ancestors of X_l given the parents of X_l , there exist probability measures $Q^\pm = Q^\pm(\eta) \in \mathcal{D}_{l,P}^{\eta_l}$ given by (3.8)–(3.9) such that

$$(3.20) \quad \mathbb{E}_{Q^\pm} [f(X_k)] - \mathbb{E}_P [f(X_k)] = \sup/\inf_{Q \in \mathcal{D}_{l,P}^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)].$$

- (c) For $l \in \bar{\rho}_k^P$ such that (3.19) is not satisfied, (3.18) holds.

Proof. Parts (a) and (c) are straightforward consequences of Lemma 3.1 and (3.18), respectively. The proof of part (b) is as follows: For $l \in \bar{\rho}_k^P$ with $X_k \perp X_{\rho_l \setminus \pi_l} | X_{\pi_l}$, we have $F(x_l, x_{\rho_l^P}) = F(x_l, x_{\pi_l})$. Then the proof is the same as the proof of Theorem 3.2. Indeed, let

$$(3.21) \quad q_l^+(x_l | x_{\pi_l^{Q^+}}) = \frac{e^{\pm c_+ F(x_l, x_{\rho_l^P})}}{\mathbb{E}_{P_{l|\pi_l^P}} \left[e^{c_\pm F(X_l, x_{\pi_l^P})} \right]} p(x_l | x_{\pi_l^P}) \quad \text{for all } x_{\pi_l^{Q^+}} \text{ and } \pi_l^{Q^+} = \pi_l^P,$$

where $c_+ \equiv c_+(x_{\pi_l^P}; \eta_l)$ are functions of x_{π_l} (since F only depends on x_l and x_{π_l}), depend on η_l , and are determined by the equations $R(Q_{l|\pi_l^{Q^+}}^+ \| P_{l|\pi_l^P}) = \eta_l$. Therefore, the density of Q^+ is given by $q^+(x) = q_l^+(x_l | x_{\pi_l^{Q^+}}) \prod_{i \neq l} p(x_i | x_{\pi_i})$. Thus, $Q_l^+ \in \mathcal{D}_{l,P}^{\eta_l}$ makes (3.21) an equality. Therefore we can conclude that

$$(3.22) \quad \sup_{Q \in \mathcal{D}_{l,P}^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] = \mathbb{E}_{P_{\rho_l}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|\pi_l}} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right] \right].$$

The case of $\inf_{Q \in \mathcal{D}_{l,P}^{\eta_l}} \mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)]$ is treated similarly. By Lemma 3.1, for $l \notin \bar{\rho}_k^P$ and $Q \in \mathcal{D}_{l,P}^{\eta_l}$, $\mathbb{E}_Q [f(X_k)] - \mathbb{E}_P [f(X_k)] = 0$. ■

Remark 3.4. The condition $X_k \perp X_{\rho_l \setminus \pi_l} | X_{\pi_l}$ can be satisfied when $\rho_l \cap \rho_i \subset \pi_l$ for all $i \in \bar{\rho}_k \setminus \bar{\rho}_l$, i.e., any path from $X_{\rho_l \setminus \pi_l}$ to X_k must go through X_{π_l} , for instance, all Markov chains, tree/polytree structure model, etc. Two simple examples where the assumption is satisfied or violated are shown in Figure 4. This condition is also satisfied by the baseline Bayesian network discussed in section 7.

Remark 3.5. Note that for the model sensitivity indices shown in (3.7) in Theorem 3.2 or the uncertainty bounds shown in (3.18) in Theorem 3.3, sometimes it might be practically difficult to find the infimum for every conditioning ρ_l . However, we can use an alternative looser bound by Jensen’s inequality, i.e.,

$$(3.23) \quad \begin{aligned} I^+(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l}) &\leq \mathbb{E}_{P_{\rho_l}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|\pi_l}} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right] \right] \\ &\leq \inf_{c>0} \left[\mathbb{E}_{P_{\rho_l}} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|\pi_l}} \left[e^{c\bar{F}(X_l, X_{\rho_l})} \right] \right] + \frac{\eta_l}{c} \right], \end{aligned}$$




Figure 3. Example of the structure of a Bayesian network baseline model. The QoI is given by $f(X_7) = X_7$ (blue). We fix X_7 and perturb one vertex at a time, e.g., X_3 (left) and X_6 (right) in green. The vertices involved in the graph can be classified into $l \in \bar{\rho}_7^P = \rho_7^P \cup \{7\} = \{1, 2, 3, 4, 5, 6, 7\}$ (vertices in the dashed area) and $\{8, 9\}$ which are not in $\bar{\rho}_7^P$ (vertices outside of the dashed area); see left and right figures. Based on these figures and Lemma 3.1, the model sensitivity indices (3.4) over $\mathcal{D}_l^{\eta_l}$ and $\mathcal{D}_{l,P}^{\eta_l}$ is 0 for $l = 8, 9$, meaning that perturbations on vertices which are not ancestors of 7 do not affect the QoI, while perturbations on those vertices in $\bar{\rho}_7^P$ affect the QoI.




Figure 4. (Left) Two examples of the structure of a baseline Bayesian network. The QoI is $f(X_7)$ in yellow (thus $k = 7$) and $l = 6$ in purple. (Right) Schematic of relationships between the ambiguity sets $\mathcal{D}_6^{\eta_6}, \mathcal{D}_{6,P}^{\eta_6}$. They share the same boundary, and thus we represent $\mathcal{D}_6^{\eta_6}$ as a sphere in blue, while $\mathcal{D}_{6,P}^{\eta_6}$ is shown as an embedded disc in brown. The yellow curve in both figures demonstrates the parametric family of Bayesian networks P^c with $dP_{l|\pi_l}^c = dP_{l|\pi_l}$ for $l \neq 6$ and $dP_{6|\pi_6}^c \propto \exp\{cF(x_6, x_{\rho_6^P})\}dP_{6|\pi_6}$. The top graph does not satisfy condition (3.19) since X_1 is not conditionally independent of X_7 given X_{π_6} . This is illustrated through the path $X_1 \rightarrow X_5 \rightarrow X_7$ in black. The function F given by (3.6) depends on x_1 and x_6 , which makes the parents of X_6 in the optimizers Q^\pm different from its parents in P and thus $Q^\pm \notin \mathcal{D}_{6,P}^{\eta_6}$ (in general), as illustrated in the top left picture. The bottom graph could achieve the equality in (3.7) since it satisfies condition (3.19) ($X_{\rho_6 \setminus \pi_6} = \{X_1, X_2\}$ are connected with X_7 only through $X_3 \in X_{\pi_6}$). The function F depends on x_3 and x_6 and $\pi_6^Q = \{3, 4\} = \pi_6$, which makes $Q^\pm \in \mathcal{D}_{6,P}^{\eta_6}$; see bottom right picture.

and the model sensitivity index $I^-(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l})$ can be treated analogously. Moreover, the corresponding bounds for $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$ are similar. In addition, if $\rho_l^P = \emptyset$, then expectation $\mathbb{E}_{P, \rho_l^P}[\cdot]$ does not enter into the overall calculations, and hence

$$I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) = \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|\pi_l^P}} \left[e^{\pm c \bar{F}} \right] + \frac{\eta_l}{c} \right], \quad l \in \bar{\rho}_k^P,$$

e.g., for $l \in \{1, 2, 4\}$ and $k = 7$ as illustrated in Figure 3. This is a special case, however; it is used in the computation of the model sensitivity indices for the materials design problem in section 7.

3.1. Gaussian Bayesian networks. Next, we develop model sensitivity indices $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$, when P is a Gaussian Bayesian network, and $f(X_k)$ depends on X_k linearly. We first use Theorem 3.6, along with the fact that each model component is a linear Gaussian of its parents, and compute F and \bar{F} explicitly. We show that \bar{F} depends only on the l th component and its parents π_l^P . Then, to implement Theorem 3.2, we calculate the MGF of \bar{F} with respect to $P_{l|\pi_l^P}$. We prove that it no longer depends on π_l^P , due to cancellations between the terms involving π_l^P . Thus, the expectation $\mathbb{E}_{P_{\rho_l^P}}$ does not enter the overall computation of (3.7). Finally, we prove that $Q^\pm \in \mathcal{D}_{l,P}^{\eta_l}$, i.e., Q^\pm are Gaussian Bayesian networks with the same structure as P , without requiring condition (3.19) be satisfied, as explained in the proof of the theorem.

Theorem 3.6 (model sensitivity indices for Gaussian Bayesian networks). *Let P be a Gaussian Bayesian network satisfying (2.17), and let $f(X_k) = aX_k + b$ be a QoI that depends only on X_k linearly. Then the following hold:*

(a.) For the model sensitivity indices defined in (3.4), we have that

$$(3.24) \quad I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) \equiv I^\pm(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l})$$

and the optimizer $Q^\pm = Q^\pm(\eta) \in \mathcal{D}_{l,P}^{\eta_l} \subset \mathcal{D}_l^{\eta_l}$ given by (3.9)–(3.8) are also Gaussian Bayesian networks with the same graph structure as P . Furthermore, for $l \in \pi_k^P$ and $l \notin \rho_{\pi_j}^P$ for all $j \in \pi_k$, $j \neq l$, we have

$$(3.25) \quad I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) = \pm |\beta_{kl}| \sqrt{2a^2 \sigma_l^2 \eta_l}.$$

(b.) Moreover, for any $l \in \rho_k^P$, we also have

$$(3.26) \quad I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) = \pm |\tilde{\beta}_{kl}| \sqrt{2a^2 \sigma_l^2 \eta_l}$$

for a computable constant $\tilde{\beta}_{kl}$.

Proof. Let $f(X_k) = aX_k + b$ and $l \in \bar{\rho}_k$. Then by a straightforward calculation, F given by (3.6) can be expressed as

$$(3.27) \quad F(X_l, X_{\rho_l}) = a\tilde{\beta}_{k_0} + a \sum_{j \in \rho_l} \tilde{\beta}_{kj} X_j + a\tilde{\beta}_{kl} X_l + b$$

for some computable $\tilde{\beta}_{k_0}, \tilde{\beta}_{kj}$ with $j \in \rho_l$ (see Example D.3, where we compute β_{kl} and $\tilde{\beta}_{kl}$). Furthermore, by using (2.17), we obtain the centered F denoted by \bar{F} ,

$$(3.28) \quad \bar{F}(X_l, X_{\pi_l}) = \tilde{\beta}_{kl} a(X_l - \beta_{l0} - \beta_l^T X_{\pi_l}),$$

and thus the MGF of \bar{F} with respect to $P_{l|\pi_l}$ in the second equality of (3.7) is

$$(3.29) \quad \mathbb{E}_{P_{l|\pi_l}} \left[e^{\pm c \bar{F}(X_l, X_{\pi_l})} \right] = \int_{\mathcal{X}_l} e^{\pm c \pm \tilde{\beta}_{kl} a x_l} e^{\mp c \pm \tilde{\beta}_{kl} a (\beta_{l0} + \beta_l^T x_{\pi_l})} dx_l$$

$$(3.30) \quad \begin{aligned} &= e^{\pm c \tilde{\beta}_{kl} a (\beta_{l0} + \beta_l^T x_{\pi_l}) + c^2 \tilde{\beta}_{kl}^2 a^2 \frac{\sigma_l^2}{2}} e^{\mp c \tilde{\beta}_{kl} a (\beta_{l0} + \beta_l^T x_{\pi_l})} \\ &= e^{c^2 \tilde{\beta}_{kl}^2 a^2 \frac{\sigma_l^2}{2}}. \end{aligned}$$

We compute the minimization problem of (3.7) by following the steps given in the proof of Theorem 2.5.

Regarding the structure of Q^\pm , $Q^\pm \in \mathcal{D}_{l,P}^{\eta_l}$, i.e., the graph of Q^\pm is the same as P , as proved in Theorem 2.5 (see also Example D.1), where we showed that due to cancellations that may occur in the derivation of CPDs q^\pm the graph remains the same. ■

4. Stress tests, ranking, and correctability. Based on the model sensitivity indices discussed in section 3 we build an iterative approach that ranks the Bayesian network components of the baseline P according to their model sensitivity indices and subsequently improve its predictive ability for specific QoIs. The model misspecification η_l of the ambiguity sets can be either set up by the user, e.g., when the data for component l are very sparse or absent, or estimated from data, building a data-informed ambiguity set. Once η_l 's are specified, we rank the sensitivity indices $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$ for all vertices l based on their relative size. Here the largest indices correspond to the most “sensitive” CPDs in the sense that they have the largest effect on the uncertainty of the QoI. From a machine learning perspective, such a ranking procedure is a form of *interpretability*, i.e., the ability to identify the cause and effect in a model [19, 52, 13], and *explainability*, i.e., the ability to explain model outputs based on modeling and data choices made during the learning of the baseline [1].

Once the ranking is completed, we turn to correcting the most influential components of a baseline Bayesian network, a task also referred to as *correctability* in machine learning, namely, the ability to correct predictive errors without introducing or (tightly) controlling any newly created errors (see Theorem 5.1 for Gaussian Bayesian networks) [1, 38, 13]. To this end we need to assess the impact of limited data, seek additional data targeting specific model components, or update some of the CPDs or the graph of the baseline Bayesian $\{G, P\}$. All these elements can be organized in a 4-step strategy discussed next, while they are implemented in an example in materials design for fuel cells in section 7.

Notation. We recall that $P_{l|\pi_l}$ is the conditional distribution of X_l with the given parent values $X_{\pi_l} = x_{\pi_l}$. However, we write $P_{l|X_{\pi_l}^P}$ when X_{π_l} is still a random variable and $P_{l|X_{\pi_l} = x_{\pi_l}}$ when we simply emphasize the dependence on given parents; see Step 1 below and the KL chain rule in Appendix G. Finally, for each vertex $l \in V$ we use the notation $\pi_l := \pi_l^Q \cup \pi_l^P$ when we consider simultaneously the parents for both models.

Step 1: Stress tests and model sensitivity. In this step, we determine the level of model misspecification η_l for each component $l \in V$ of the baseline using data-informed or user-determined stress tests.

A. *Data-informed stress tests.* For Bayesian networks (or parts thereof) for which there is a reasonable amount of data, here we construct data-informed ambiguity sets (2.1), (3.1), and (3.2), respectively. The corresponding levels of model misspecification η, η_l are computed as distances between the baseline P and the data distribution Q ; the latter can be selected as a histogram or a Kernel density estimation (KDE). In that sense, we provide surrogate values for the model misspecifications η or η_l , taking into account the “real” model which is accessible only through the available data. In these calculations we are taking full advantage of the graph structure of the models. First, we discuss the model uncertainty ambiguity set \mathcal{D}^η in (2.1). Using the chain rule of KL divergence for Bayesian networks (Appendix G) we define a data-informed misspecification η as

$$(4.1) \quad \eta := R(Q||P) = \sum_{l=1}^n \mathbb{E}_Q [\eta_l^{\pi_l}], \quad Q \in \mathcal{D}^\eta,$$

where $\eta_l^{\pi_l}$ is a function of X_{π_l} given by

$$(4.2) \quad \eta_l^{\pi_l} = \mathbb{E}_Q \left[R(Q_{l|X_{\pi_l}^Q} || P_{l|X_{\pi_l}^P}) \right] = \int_{\mathcal{X}_l} \log \frac{Q_{l|X_{\pi_l}^Q}}{P_{l|X_{\pi_l}^P}} Q_{l|X_{\pi_l}^Q} dx_l.$$

Second, for the case of model sensitivity, definition (4.1) reduces to

$$(4.3) \quad \eta_l = R(Q||P) = \mathbb{E}_Q [\eta_l^{\pi_l}], \quad Q \in \mathcal{Q}_{\eta_l},$$

where \mathcal{Q}_{η_l} is given by (3.1) or (3.2); to obtain this simplification of (4.1) we used the structure of the ambiguity sets \mathcal{Q}_{η_l} where all CPDs are identical except for the one on the l th vertex.

We now turn to the estimation of (4.1) and (4.3). We note that due to the graphical structure of Bayesian networks their estimation reduces to focusing on individual model components. Related recent ideas using subadditivity for divergences or probability metrics of PGMs, instead of a full chain rule, were explored for statistical learning in [18]; such an approach could be also used here in an uncertainty quantification context. Lastly, we can simplify the estimation of (4.1) or (4.3) by using an upper bound, $\eta_l \leq \sup_{x_{\pi_l}} R(Q_{l|X_{\pi_l}=x_{\pi_l}} || P_{l|X_{\pi_l}=x_{\pi_l}})$. Under certain conditions we can also show that using KDE gives rise to consistent statistical estimator; see (I.1)–(I.4) for a Gaussian Bayesian network baseline. Finally, we note that significant literature on statistical estimators for divergences includes nonparametric estimators [54], statistical estimators based on variational representations of divergences [56, 4], density-estimator-based methods for estimating divergences in low dimensions [45], estimators of divergence based on nearest-neighbor distances [71, 70, 60], and statistical estimators for Rényi divergences [7].

B. *User-determined stress tests.* Here we use $\eta_l \geq 0$ as a parameter to be tuned by hand to explore how different levels of uncertainty will affect the QoI; for instance, when we have very sparse or missing data and η_l 's are set by a user. This is a form of nonparametric sensitivity analysis and in spirit is reminiscent of the stress tests used in finance and actuarial science (see, e.g., [11]) to protect against sudden changes and extreme uncertainty under various scenarios. In our Bayesian network context, individual model misspecification $\eta_l, l \in V$,

for the model sensitivity indices $I^\pm(f(X_k), P; \mathcal{Q}_{\eta_l})$ can take arbitrary fixed values that correspond to model perturbations associated with local sensitivity analysis (small η_l) or global sensitivity analysis (larger η_l). Both local and global sensitivity analyses are conducted in the same mathematical framework, and therefore we have the flexibility to explore combinations of small/large model perturbations at different vertices of the Bayesian network. From a practical point of view, these sensitivity computations can be done using only one fixed constructed Bayesian network (the baseline), yielding guarantees for entire neighborhoods of models.

Step 2: Ranking of model sensitivities. Once η_l 's are specified in Step 1 for each vertex l , we calculate the model sensitivity indices $I^\pm(f(X_k), P; \mathcal{Q}_{\eta_l})$ using Theorems 3.2 and 3.3, where $\mathcal{Q}_{\eta_l} = \mathcal{D}_l^{\eta_l}$ or $\mathcal{D}_{l,P}^{\eta_l}$ are defined in (3.1) and (3.2). Subsequently we rank them according to their relative contributions

$$(4.4) \quad \frac{I^+(f(X_k), P; \mathcal{Q}_{\eta_l})}{\sum_j I^+(f(X_k), P; \mathcal{Q}_{\eta_j})}$$

See also the example in Figure 9.

Step 3: Assessing the baseline. After we have ranked the model sensitivities in Step 2, we focus on the most impactful model components and assess their impact on the QoI $\mathbb{E}_P[f(X_k)]$. Specifically, if the relative model uncertainty is less than an application-dependent tolerance TOL ,

$$(4.5) \quad I^+(f(X_k), P; \mathcal{Q}_{\eta_l}) \leq TOL,$$

then we decide to “trust” the model component l . If there are model components that do not satisfy (4.5), we proceed to the next step in order to correct the baseline model P . This is a form of interpretability, since we can systematically identify underperforming parts of the model. A related quantity that can also be used in (4.5) is the relative model sensitivity

$$(4.6) \quad \frac{I^+(f(X_k), P; \mathcal{Q}_{\eta_l})}{\mathbb{E}_P[f(X_k)]};$$

see, for example, Figure 10.

Step 4: Model correctability. Once Steps 2 and 3 are completed, we turn to correcting the most influential components of the baseline Bayesian network P , a task also referred to as correctability in machine learning. We formulate mathematically this procedure in section 5; however, practically we aim at reducing the index $I^+(f(X_k), P; \mathcal{D}_l^{\eta_l})$ for each vertex $l \in V$ that violates (4.5). This can be accomplished, for instance, by either acquiring additional data or updating the CPD of these specific vertices. However, as we correct these targeted model components of the baseline, we also need to guarantee that we do not introduce new, bigger errors in the remaining components of the Bayesian network that would violate (4.5). Section 5 provides both theory and related practical implementation strategies to this end.

5. Mathematical analysis of correctability in Bayesian networks. In this section, we focus on the mathematical formulation of correctability in Bayesian networks outlined in Step 4 of section 4. Our methods are motivated by “correcting” a baseline model by either acquiring

targeted high-quality data, or updating the CPDs of the most underperforming components (see Step 3 of section 4), or correcting the graph G itself. We demonstrate these scenarios, their combinations, and our mathematical methods on a materials screening problem for fuel cells in section 7.

The intuition behind our correctability analysis lies in the model sensitivity results for the Gaussian case. By [Theorem 3.6](#), the model sensitivity indices of a baseline P for a targeted l^* th CPD component are given by

$$(5.1) \quad I^\pm(f(X_k), P; \mathcal{D}_{l^*, P}^{\eta_{l^*}}) = \pm |\tilde{\beta}_{kl^*}| \sqrt{2\sigma_{l^*}^2 \eta_{l^*}}.$$

Therefore, additional/better data or an improved CPD for the l^* th vertex could allow l^* th CPD with mean zero with a corresponding new Gaussian Bayesian model \tilde{P} that is otherwise identical to P . Indeed, if we could guarantee a combination of

$$\tilde{\sigma}_{l^*}^2 < \sigma_{l^*}^2 \quad \text{and/or} \quad \tilde{\eta}_{l^*} < \eta_{l^*}$$

for the new model \tilde{P} , then we can quantify the improvement of the baseline P using (5.1) and show that the indices of \tilde{P} at l^* would decrease.

In general, we seek to correct the targeted l^* th vertex of the baseline P to obtain a new Bayesian network \tilde{P} such that

$$(5.2) \quad I^\pm(f(X_k), \tilde{P}; \mathcal{D}_l^{\eta_l}) \leq I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) \quad \text{for all } l \neq l^*$$

and

$$(5.3) \quad I^\pm(f(X_k), \tilde{P}; \mathcal{D}_{l^*, P}^{\eta_{l^*}}) < I^\pm(f(X_k), P; \mathcal{D}_{l^*, P}^{\eta_{l^*}}).$$

In particular, (5.2) and (5.3) would imply that we can improve the CPD of the l^* vertex, and at the same time we do not decrease the performance of the rest of the Bayesian network. The next theorem demonstrates that we can achieve (5.2) when P is a Gaussian Bayesian network. Moreover, when P is a general Bayesian network, we prove that new errors that may violate (5.2) can only be created in the descendant components of l^* ; see also [Remark 5.2](#).

Theorem 5.1.

(a) **(Gaussian Bayesian network)** Consider $f(X_k) = aX_k + b$ to be a QoI that only depends on X_k linearly. Let also P be a Gaussian Bayesian network satisfying (2.17). Suppose now that we construct a new Bayesian network \tilde{P} by only updating the CPD $p(x_{l^*} | x_{\pi_{l^*}})$ for some $l^* \in \rho_k$ as follows: we change the distribution of ϵ_{l^*} in (2.17) from Gaussian to another mean zero distribution denoted by $\tilde{p}(x_{l^*} | x_{\pi_{l^*}})$. Note that the graph structure of \tilde{P} is the same as P . Then

$$(5.4) \quad I^\pm(f(X_k), \tilde{P}; \mathcal{D}_l^{\eta_l}) = I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) \quad \text{for all } l \neq l^*,$$

where $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$ is given by (3.25)–(3.26). Moreover, for the relative model sensitivity (4.6) the following holds:

$$(5.5) \quad \frac{I^\pm(f(X_k), \tilde{P}; \mathcal{D}_l^{\eta_l})}{\mathbb{E}_{\tilde{P}}[f]} = \frac{I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})}{\mathbb{E}_P[f]} \quad \text{for all } l \neq l^*.$$

(b) **(Non-Gaussian Bayesian network)** Let $f(X_k)$ be a QoI that only depends on X_k . Let also P be a non-Gaussian Bayesian network. Let us suppose that we construct a new Bayesian network \tilde{P} with the same structure as P by only updating the CPD $p(x_{l^*}|x_{\pi_{l^*}})$ for some $l^* \in \rho_k^P$. Then

$$(5.6) \quad I^\pm(f(X_k), \tilde{P}; \mathcal{D}_{l^*,P}^n) = I^\pm(f(X_k), P; \mathcal{D}_{l^*,P}^n) \quad \text{for all } l \in \rho_k^P \text{ with } l < l^*,$$

while the model sensitivity indices for any $l \in \rho_k^P$ with $l \geq l^*$ (descendant components) change and are given by Theorem 3.3.

Proof.

- (a.) First, updating $p(x_{l^*}|x_{\pi_{l^*}})$ with $l^* \in \rho_k$ does not affect the computation of F defined in (3.6). This is straightforward by (3.6). In the case of a Gaussian Bayesian network, F is given by (3.27). Second, if $l \neq l^*$, by (3.27), the MGF of \bar{F} with respect to $\tilde{P}_{l|\pi_l}$ is always the same with the MGF computed with respect to $P_{l|\pi_l}$ (since $\tilde{P}_{l|\pi_l} = P_{l|\pi_l}$); see (3.29). However, it only changes when $l = l^*$. Moreover, the relative model sensitivity with respect to model \tilde{P} satisfies (5.5), since $\tilde{P}_{l^*|\pi_{l^*}} : X_{l^*} = \beta_{l^*0} + \beta_{l^*}^T X_{\pi_{l^*}} + \tilde{\epsilon}_{l^*}$, with $\tilde{\epsilon}_{l^*}$ another mean zero distribution. Thus the expected values of $f(X_k)$ with respect to P and \tilde{P} are equal.
- (b.) It is enough to observe that for any $l \in \rho_k^P$ with $l < l^*$, the MGF in Theorem 3.2 computed with the respect to $\tilde{P}_{l|\pi_l}^P = P_{l|\pi_l}^P$ and $P_{l|\pi_l}^P$ are equal and both depend on the ancestors of ρ_l^P , where $l^* \notin \rho_l^P$. Hence, (3.7) for both models is the same. Similarly, we prove the case $l \in \rho_k^P$ with $l \geq l^*$ (descendant components). Note that this time $l^* \in \rho_l^P$, and thus (3.7) is different for the two Bayesian networks. ■

Both developed approaches are implemented in section 7.3. For example, Theorem 5.1 (a) is applied when we update a CPD of the baseline Gaussian Bayesian network by using a kernel-based (KDE) method; see Figure 10. We refer the reader to section 7.3 for full details.

Remark 5.2. Even if the conditions of Theorem 5.1 are not applicable, the ranking procedure of Steps 2 and 3 in section 4 can always identify the best candidates among the components of the graphical model for improvement relative to a QoI. Once we correct the component l^* selected through ranking we need to recompute the relative model uncertainties in (4.5) for all vertices $l \in V$ and then determine the suitability of the corrected model. In fact, due to Theorem 5.1 (b) we only need to compute (4.5) for just the vertices l in the descendants of l^* since all the remaining ones are not affected by the model correction.

6. DFT-informed Langmuir model. In this section, we consider the Langmuir bimolecular adsorption model that describes the chemical kinetics with competitive dissociative adsorption of hydrogen and oxygen on a catalyst surface [62]. It is a multiscale system of random differential equations with correlated dependencies in their parameters (kinetic coefficients), arising from quantum-scale computational data calculated using density functional theory (DFT) (i.e., quantum computations) for actual metals. The combination of chemical kinetics with parameter dependencies, correlations, and DFT data gives rise naturally to a Bayesian network. However, the limited availability of the quantum-scale data creates significant model uncertainties in both the distributions of kinetic coefficients and their correlations; see, for example, Figure 5(a). Thus, we will quantify the ensuing model uncertainties by implementing

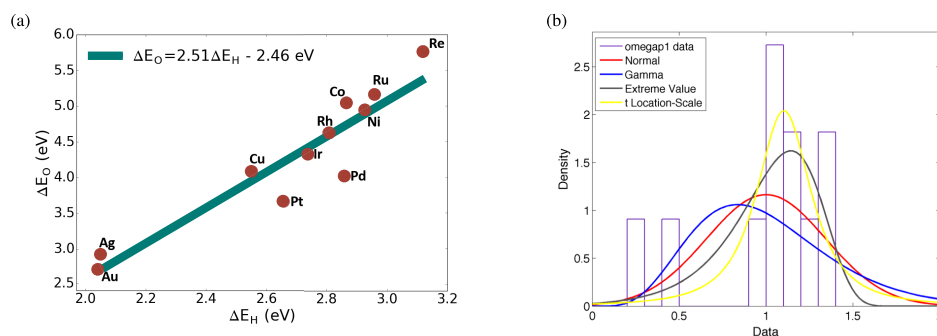


Figure 5. (a) Correlation between oxygen and hydrogen adsorption energies on metal surfaces as defined in (6.6), (b) Fit of ω in (6.6) with various parametric distributions.

our analysis in sections 2 and 3. Here, the equilibrium hydrogen and oxygen coverages are our QoIs and can be calculated by the dynamics of the chemical reaction network described by the following system of random ODEs with random (correlated) coefficients:

$$(6.1) \quad \frac{dC_{H^*}}{dt} = k_{H_2}^{ads} P_{H_2} (1 - C_{H^*} - C_{O^*})^2 - k_{H_2}^{des} C_{H^*}^2, \quad C_{H^*}^0 = C_{H^*}(0),$$

$$(6.2) \quad \frac{dC_{O^*}}{dt} = k_{O_2}^{ads} P_{O_2} (1 - C_{H^*} - C_{O^*})^2 - k_{O_2}^{des} C_{O^*}^2, \quad C_{O^*}^0 = C_{O^*}(0),$$

where C_{H^*} and C_{O^*} represent the hydrogen and oxygen coverages. P_{H_2} and P_{O_2} are the partial pressures of the gas phase species and are fixed.

Then the steady state solution of (6.1)–(6.2), which constitute our QoIs, is given by

$$(6.3) \quad \hat{C}_{H^*} = \frac{(K_{H_2} P_{H_2})^{\frac{1}{2}}}{1 + (K_{H_2} P_{H_2})^{\frac{1}{2}} + (K_{O_2} P_{O_2})^{\frac{1}{2}}}, \quad \hat{C}_{O^*} = \frac{(K_{O_2} P_{O_2})^{\frac{1}{2}}}{1 + (K_{H_2} P_{H_2})^{\frac{1}{2}} + (K_{O_2} P_{O_2})^{\frac{1}{2}}}.$$

Here $K_i = \frac{k_i^{ads}}{k_i^{des}}$ for $i = H_2, O_2$, and for each species they are related to electronic structure (DFT) calculations through an Arrhenius law [24]:

$$(6.4) \quad K_{H_2} = e^{-\frac{G_{H_2}}{k_B T}} (P_{H_2} + P_{O_2})^{-1}, \quad G_{H_2} \propto -2\Delta E_H,$$

$$(6.5) \quad K_{O_2} = e^{-\frac{G_{O_2}}{k_B T}} (P_{H_2} + P_{O_2})^{-1}, \quad G_{O_2} \propto -2\Delta E_O.$$

The constants k_B and T are the Boltzmann constant and the temperature, respectively. In the above formulas, G_{H_2} and G_{O_2} are the hydrogen and oxygen Gibbs free energies of adsorption. Therefore, the coverages \hat{C}_{H^*} and \hat{C}_{O^*} are nonlinear functions of ΔE_H and ΔE_O . We refer the reader to for the chemistry background and analysis of the model. In, the authors have estimated the two binding energies for various metal catalyst surfaces via DFT calculations as illustrated in Figure 5(a). Furthermore, correlations between ΔE_O and ΔE_H are captured by a statistical linear model

$$(6.6) \quad \Delta E_O = a\Delta E_H + b + \omega,$$

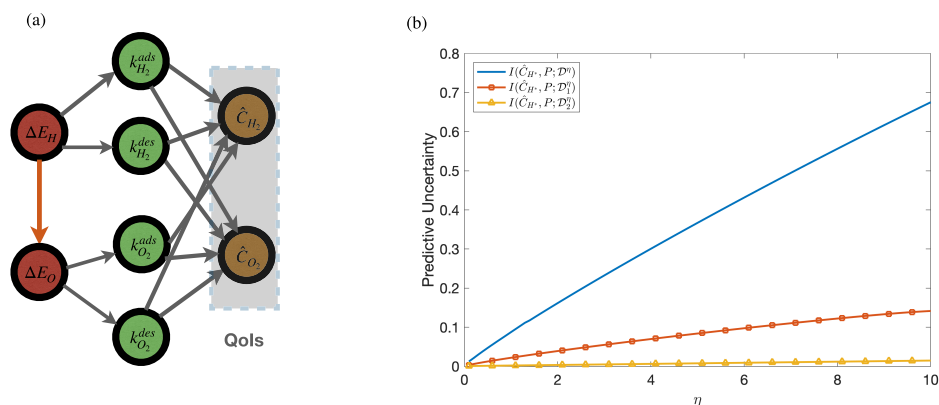


Figure 6. (a) Graph structure of the baseline Bayesian network P in (6.9) built by data (6.7) and (6.8), physics knowledge (6.4), (6.5), (6.8), and the steady state of the ODEs given by (6.3). (b) We consider the QoIs (6.3) of (6.9); the blue line represents the model uncertainty index $I^\pm(f, P; \mathcal{D}^\eta)$ as a function of η (Theorem 2.1); the red and yellow lines are, respectively, the model sensitivity indices $I^\pm(f, P; \mathcal{D}_1^\eta)$ and $I^\pm(f, P; \mathcal{D}_2^\eta)$ (Theorem 3.2) for $p(\Delta E_H)$ and $p(\Delta E_O|\Delta E_H)$, where \mathcal{D}_1^η indicates the perturbation on $p(\Delta E_H)$ and \mathcal{D}_2^η for $p(\Delta E_O|\Delta E_H)$.

where ω is a random variable. The distribution of ω can be determined by fitting the residual data from linear regression using maximum likelihood estimation (MLE); see Figure 5(b). In (6.6) we select a Gaussian distribution for ω (red line in Figure 5(b)) as the baseline CPD for the correlation in Figure 5:

$$(6.7) \quad p(\Delta E_O|\Delta E_H) = \mathcal{N}(a \Delta E_H + b, \sigma_\omega^2).$$

Next we model the distribution of the prior $p(\Delta E_H)$. Based on physical constraints (e.g., positivity of the random variable without physical upper bound), in the distribution of ΔE_H was selected to be a gamma distribution with mean x_H with standard deviation given by the difference between experiment and DFT, $(x_H - y_H)$,

$$(6.8) \quad p(\Delta E_H) = \frac{1}{b_H^{a_H} \Gamma(a_H)} \Delta E_H^{a_H-1} \exp\left(-\frac{\Delta E_H}{b_H}\right) \quad \text{for } \Delta E_H > 0,$$

where $a_H = x_H^2 / (x_H - y_H)^2$ and $b_H = (x_H - y_H)^2 / x_H$. This is a case with very little data (x_H, y_H) and only some reasonable physical constraints without any further knowledge on the model, and therefore model uncertainty in (6.8) is evident.

We now build the baseline Bayesian network P by combining the following ingredients: data through (6.7) and (6.8), physics and expert knowledge in (6.4), (6.5), (6.8), and the steady state of the ODEs (QoI) given by (6.3); see also Figure 6. We obtain the following Bayesian network and the corresponding CPDs:

$$(6.9) \quad p(x) = \underbrace{p(\hat{C}_{H^*}, \hat{C}_{O^*} | K_{H_2}, K_{O_2})}_{(6.3)} \prod_{i=H_2, O_2} \underbrace{p(K_i | \Delta E_i)}_{(6.4), (6.5)} \underbrace{p(\Delta E_O | \Delta E_H)}_{(6.7)} \underbrace{p(\Delta E_H)}_{(6.8)}.$$

In the above formula, $p(\hat{C}_{H^*}, \hat{C}_{O^*} | K_{H_2}, K_{O_2})$ and $p(K_i | \Delta E_i)$ are deterministic, while the only random parts in P are $p(\Delta E_O | \Delta E_H)$ and $p(\Delta E_H)$.

In the process of building the baseline model P above, the sparse data in Figure 5 for (6.7) and the lack of both knowledge and (almost any) data in (6.8) create model uncertainties for the prediction of the QoIs in (6.3). We quantify these uncertainties by implementing the model uncertainty index of Theorem 2.1 and the model sensitivity indices of Theorem 3.2; see Figure 6(b), where we readily see how the indices change for different values η ; the implementation of the indices was carried out through Monte Carlo simulation of the moment generating functions. Moreover, we observe that for the QoIs (6.3) the impact of uncertainties in the prior $p(\Delta E_H)$ are significantly higher than in the correlation $p(\Delta E_O|\Delta E_H)$ when we perturb with same model misspecification η . Finally, we note that, due to the lack of data in (6.8), we elected to perform the user-determined stress tests of Step 1.B of section 4 where the user selects various levels of model misspecification η .

7. Model uncertainty for Sabatier's principle. We study Bayesian networks built for trustworthy prediction of materials screening to increase the efficiency of chemical reactions in catalysis. Our starting point is Sabatier's principle, which describes the efficiency of a catalyst [62] through the so-called "volcano curve," e.g., the black curve in Figure 7(c). The volcano curve suggests that high catalytic activity is exhibited when the binding interaction between reactants and catalysts is neither too strong nor too weak, i.e., the peak of the volcano marked by a star in Figure 7(c). For this reason Sabatier's principle is widely viewed as an important criterion for screening materials for increased efficiency in catalysis. Our ultimate goal here is to understand how various uncertainties can affect the shape and position of the volcano curve and its peak.

Here we consider the oxygen reduction reaction (ORR), which is a known performance bottleneck in fuel cells [63]. The ORR depends on the formation of surface hydroperoxyl (OOH^*) from molecular oxygen (O_2), and water (H_2O) from surface hydroxide (OH^*) [67]. The complete mechanism [14, 2, 43] involves four electron exchange steps with reactions (R1) and (R4) being slow; see Figure 7(a). Therefore, the discovery of new materials will have to rely on speeding up the two slowest reactions in order to accelerate the entire ORR mechanism. Furthermore, such a physicochemical system has hidden correlations between variables which have emerged after statistical analysis of data [25]. In particular, the corresponding Gibbs energies of reactions (1) and (4) denoted by $-\Delta G_4 \equiv y_1$ and $-\Delta G_1 \equiv y_2$ are computed as linear combinations of free energies of species and are regressed versus the oxygen binding energy $\Delta G_O \equiv x$ calculated by DFT calculations. The oxygen binding energy x is chosen as a descriptor in [25] since it is the natural coordinate arising from Sabatier's principle. The principle is graphically represented by the volcano curve, i.e., the solid black lines in Figure 7(c), which is a function of the descriptor. Therefore, the QoI considered here is the optimal oxygen binding energy ΔG_O denoted by $x_{O^*}^P$ and identified as the maximum of the volcano curve:

$$(7.1) \quad x_{O^*}^P := \operatorname{argmax}_{x_0} [\min\{\mathbb{E}_P[y_1|x_0], \mathbb{E}_P[y_2|x_0]\}] .$$

Starting from this QoI we build a Bayesian network in Figure 7(b) that includes expert knowledge (volcano curves), as well as various available experimental and computational data and their correlations or conditional independence.

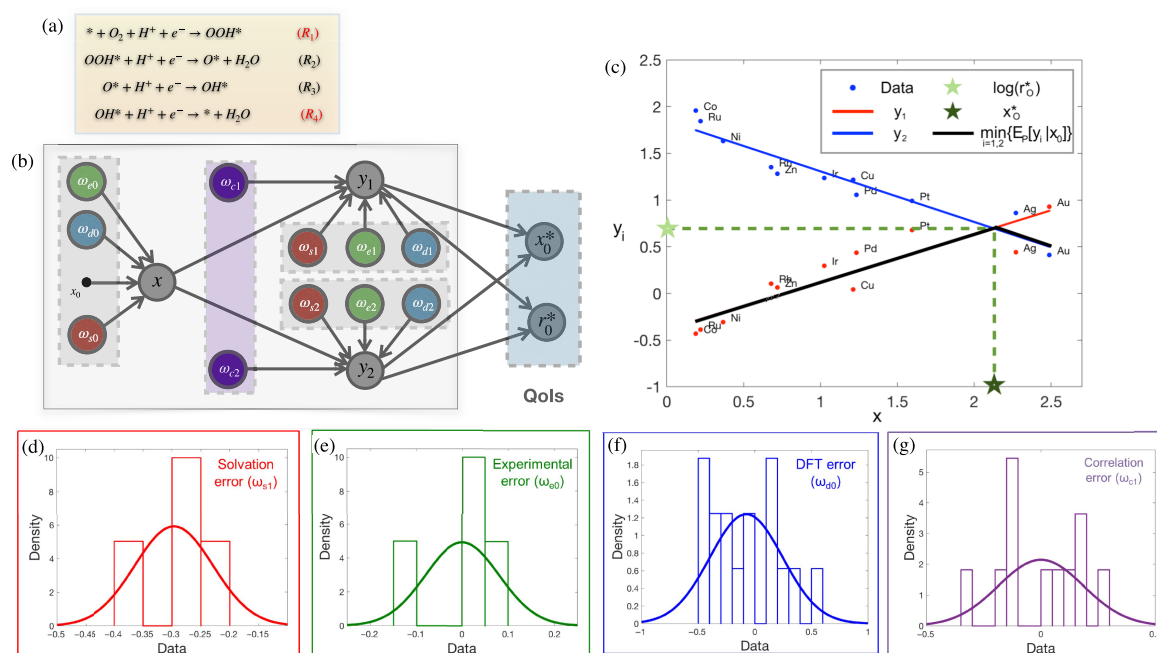


Figure 7. (a) ORR steps (R_1 to R_4) in hydrogen fuel cells, and (b) Bayesian network for ORR. The construction of the Bayesian network (section 7.1) is based on expert knowledge, physicochemical modeling, and statistical analysis of data. We include these random variables into the Bayesian network and build the directional relationships (connection/arrows) between corresponding random variable x or y_i . We build a Gaussian Bayesian network, i.e., all CPDs are Gaussians, which are fitted to available data using MLE (see histogram approximations in (d)–(g)). Note the conditional independence between the y -variables, assumed based on expert knowledge. (c) The QoI of the ORR model is the optimal oxygen binding energy $x_{O^*}^P$ and is identified when the two reaction energies are equal by physical modeling (marked with a star). (d)–(f) Here we model different kinds of errors in x and y_i , given expert knowledge.

7.1. Construction of the ORR Bayesian network for the QoI (7.1). First, we relate the QoI with the y_i 's and then we include errors from different sources in x and y_i 's.

- [Graph] We first build the directed graph for the Bayesian network. The first selected vertices in the graph are the QoIs $x_{O^*}^P$, $r_{O^*}^P$, as well as y_i 's and x ; see gray vertices in Figure 7(b). Subsequently, we have the following:

(1a) Through the statistical independence test [73], we learn that y_1 and y_2 depend on x and are conditionally independent given x as illustrated in Figure 7(b).

(1b) The construction of x comes from the DFT data (using quantum calculations) for the oxygen binding energy given the real unknown value x_0 . As mentioned in the beginning of the section, x is also selected to be the descriptor by expert knowledge (see also the supplementary material of [25]) and justifies the conditional relationships between x and y_i 's.

(1c) The evaluations of the QoIs depend on the values of y_i 's for each x_0 due to the volcano curve of Sabatier's principle.

Overall, in (1) we built part of the network structure for x , y_1 , y_2 , and the QoI using a constraint-based method [66], which selects a desired structure based on constraints of dependency among variables.

2. [CPD] Next, we build the individual CPDs on the graph constructed above.

(2a) We include statistical correlations between DFT (quantum calculation) data for x and y_i ; see data in [25]. We model the residual using a linear model with a random correlation error denoted by ω_{ci} ; see (7.3).

(2b) We model as random variables and incorporate in the Bayesian network different kinds of errors in x and y_i 's from the following sources: ω_{ei} is the error in experimental data, ω_{di} is the error between quantum and experimental values, and ω_{si} is the error due to solvation effects; all are calculated by DFT; see the corresponding data in [25]. See (7.3).

More specifically, after conducting independence tests on the corresponding data, and also based on expert knowledge or intuition [25], we assume that the random variables ω are independent. Based on the graph construction above we obtain the Bayesian network

$$(7.2) \quad p(\mathbf{x}|x_0) = \prod_{i=1,2} p(y_i|x, \omega_{ei}, \omega_{di}, \omega_{si}, \omega_{ci}) \cdot p(x|\omega_{e0}, \omega_{d0}, \omega_{s0}, x_0) \cdot \prod_{\substack{j=e_k, d_k, s_k, c_1, c_2 \\ k=0,1,2}} p(\omega_j),$$

where $\mathbf{x} = (x, y_1, y_2, \omega_{e0}, \omega_{d0}, \omega_{s0}, \omega_{e1}, \omega_{d1}, \omega_{s1}, \omega_{c1}, \omega_{e2}, \omega_{d2}, \omega_{s2}, \omega_{c2})$. The baseline CPDs in (7.2) are constructed as linear Gaussian models, namely, for $i = 1, 2$,

$$(7.3) \quad y_i = \beta_{y_i,0} + \beta_{y_i,x}x + \omega_{ei} + \omega_{di} + \omega_{si} + \omega_{ci} \quad \text{and} \quad x = x_0 + \omega_{e0} + \omega_{d0} + \omega_{s0}.$$

The CPDs for each vertex are selected as

$$(7.4) \quad p(y_i|x, \omega_{ei}, \omega_{di}, \omega_{si}, \omega_{ci}) = \mathcal{N}(\beta_{y_i,0} + \beta_{y_i,x}x + \omega_{ei} + \omega_{di} + \omega_{si} + \omega_{ci}, 0),$$

$$(7.5) \quad p(x|\omega_{e0}, \omega_{d0}, \omega_{s0}, x_0) = \mathcal{N}(x_0 + \omega_{e0} + \omega_{d0} + \omega_{s0}, 0),$$

$$(7.6) \quad p(\omega_j) = \mathcal{N}(\beta_{j,0}, \sigma_j^2),$$

where $i = 1, 2$, and $j = e0, d0, s0, e1, d1, s1, c1, e2, d2, s2, c2$. Then the resulting baseline model (7.2) is a Gaussian Bayesian network. Subsequently we use the global likelihood decomposition method [48] to learn the parameters $\beta_{y_i,0}$, $\beta_{y_i,x}$, and σ_j . The outcomes are given in Table 1. This approach is essentially a maximum likelihood estimation (MLE) on PGMs (see [48], Chapter 17.2), which exploits a fundamental scalability property that allows us to “divide and conquer” the parameter inference problem on the graph. We can also employ a Bayesian approach instead of MLE; see, for instance, [48] for the case of PGMs.

7.2. Model sensitivity, stress tests and ranking. Here, we implement the four-step strategy of section 4 to the ORR model by using data-informed stress tests or user-determined stress tests (Steps 1.A and 1.B of section 4). The primary goal is to quantify and rank the impact of model uncertainties from each component of the Bayesian network through the model sensitivity indices in section 3. Next, we compute these model sensitivity indices for the QoI $x_{O^*}^P$ in (7.1), namely,

$$(7.7) \quad \sup/\inf_{Q \in \mathcal{D}_{i,P}^n} \left\{ x_{O^*}^Q - x_{O^*}^P \right\}$$




Figure 8. Typical model uncertainty bounds $I^\pm(y_i, P; \mathcal{D}_{l,P}^{\eta_l})$, $i = 1, 2$, computed by (7.8). The model uncertainty for the QoI $x_{O^*}^P$ (see Figure 7(c)) is computed by (7.9)–(7.10) and demonstrated in yellow for model misspecification η_l in $P(\omega_l)$: (a) for $l = e1, d1, s1, c1$, $I^\pm(y_1, P; \mathcal{D}_{l,P}^{\eta_l}) = \pm\sqrt{2\sigma_l^2\eta_l}$; (b) for $l = e2, d2, s2, c2$, $I^\pm(y_2, P; \mathcal{D}_{l,P}^{\eta_l}) = \pm\sqrt{2\sigma_l^2\eta_l}$; (c) for $l = e0, d0, s0$, $I^\pm(y_i, P; \mathcal{D}_{l,P}^{\eta_l}) = \pm|\beta_{y_i,x}|\sqrt{2\sigma_l^2\eta_l}$, $i = 1, 2$.

for $l \in \{ei, di, si, ci, e0, d0, s0\}$ with $i = 1, 2$. To this end, we first use Theorem 3.6 for $i = 1, 2$ to obtain

$$(7.8) \quad I^\pm(y_i, P; \mathcal{D}_{l,P}^{\eta_l}) = \pm|\tilde{\beta}_{y_i,\omega_l}|\sqrt{2\sigma_l^2\eta_l}$$

with σ_l and $\tilde{\beta}_{y_i,\omega_l}$ given in (7.4) and Table 2, respectively. Subsequently we solve the optimization problem for $x_O = x_{O^*}^P$ and obtain the bounds for $x_{O^*}^Q - x_{O^*}^P$ as shown in Figure 8 and given by

$$(7.9) \quad \frac{-\sqrt{2\sigma_l^2\eta_l}}{\beta_{y_1,x} - \beta_{y_2,x}} \leq x_{O^*}^Q - x_{O^*}^P \leq \frac{\sqrt{2\sigma_l^2\eta_l}}{\beta_{y_1,x} - \beta_{y_2,x}}$$

for $l = ei, di, si, ci$ and $i = 1, 2$; note that the model uncertainty of ω_l only affects y_i according to the ORR Bayesian network. Furthermore,

$$(7.10) \quad \frac{-\left(|\beta_{y_1,x}| + |\beta_{y_2,x}|\right)\sqrt{2\sigma_l^2\eta_l}}{\beta_{y_1,x} - \beta_{y_2,x}} \leq x_{O^*}^Q - x_{O^*}^P \leq \frac{\left(|\beta_{y_1,x}| + |\beta_{y_2,x}|\right)\sqrt{2\sigma_l^2\eta_l}}{\beta_{y_1,x} - \beta_{y_2,x}}$$

for $l = e0, d0, s0$ as the model uncertainty of ω_l affects both y_1 and y_2 . Here $\beta_{y_i,x}$ are the coefficients given by the first CPD in (7.4). The complete algebraic calculation of (7.9) and (7.10) is given in Appendix J.1. Then by implementing Step 2 of section 4, we rank the model components as demonstrated in Figure 9. There we plot (4.4) as a pie chart, where the most impactful components are depicted.

Remark 7.1 (propagation/non-propagation of uncertainties to the QoIs). The discrepancies in the propagation of model misspecification to the QoI between different Bayesian network components is depicted in Figure 9. In particular, in Figure 9 (left) the same user-selected model misspecification η_l is applied on all ORR Bayesian network vertices. However, not all propagate and affect the same QoI. See also the example in Figure 16.

Remark 7.2. The construction of the ORR Bayesian network and its model uncertainty was carried out in [25] for the optimal oxygen binding energy defined differently than (7.1), that

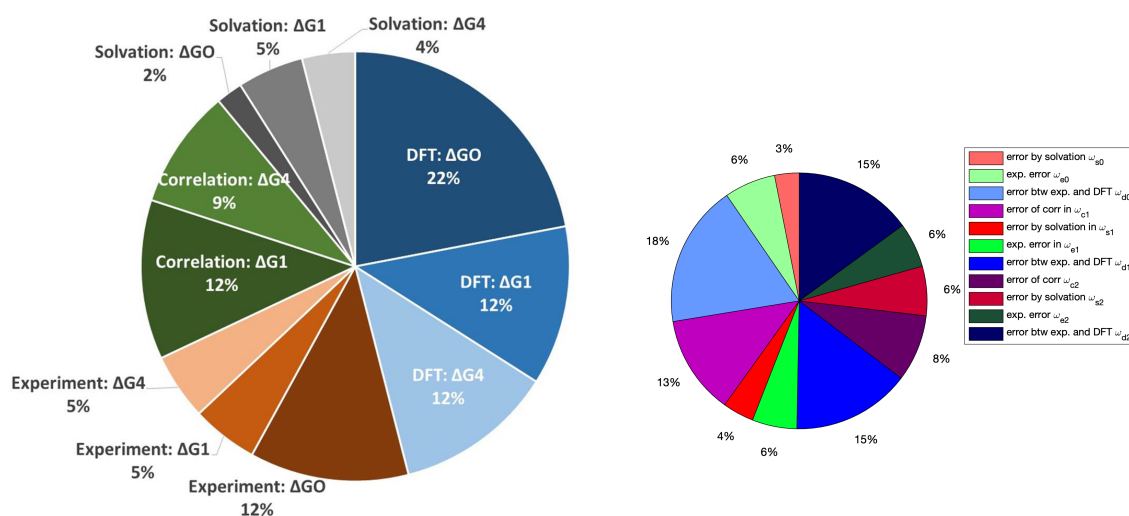


Figure 9. Relative model sensitivities (4.4) for the QoI $x_{O^*}^P$ in each ORR Bayesian network mechanism in Figure 7(b). (Left) User-determined stress test (Step 1.B in section 4); η_l has a fixed value for all l ; the particular value does not matter since it is canceled out by the ratio in (4.4). (Right) Data-informed stress test (Step 1.A in section 4); $\eta_l = R(\text{data}||P_l)$ selected as a distance of each CPD from the available data.

is, as $\text{argmax}_{x_0} \mathbb{E}_P [y|x_0]$ with $y|x_0 = \min\{y_1|x_0, y_2|x_0\}$. This is an alternative mathematical description of the same concept; however, (7.1) allows us to explicitly calculate the model sensitivity indices given by (7.8)–(7.10) and provide clear insights into what model elements and uncertainties affect them the most. On the other hand, in [25] the model sensitivity indices provided by Theorem 3.2 can only be calculated computationally.

7.3. Correctability of the ORR Bayesian network. Here we use the earlier model uncertainty/sensitivity analysis to first identify and then correct the most impactful components in several ways as discussed in Step 4 of section 4 and in the theoretical results on correctability in section 5.

1. *Including targeted high-quality data.* We seek data that lead to the reduction of the variance $\sigma_{l^*}^2$ for some $l^* \in L$ (see Step 3 of section 4), while the model misspecification η_{l^*} does not increase or the increment is much smaller than the reduction of $\sigma_{l^*}^2$. Notice that in this case the model remains a Gaussian Bayesian network. For the ORR Bayesian network, it turns out that we can add more data using DFT calculations for bimetallics to reduce the relative error for the correlation errors ω_{ci} , σ_{ci}^2 ; see the bimetallics data set in [25]. Then the model sensitivity indices of y_i on ω_{ci} , $I^\pm(y_i, P; \mathcal{D}_{l,P}^{\eta_l})$, $l = \omega_{ci}$, given by (7.8) and the model misspecification $\eta_{\omega_{ci}}$ are reduced. Consequently, the model sensitivity indices of $x_{O^*}^P$ do so as well; see (7.9). The relative predictive uncertainty (4.6) of such an updated model is demonstrated in Figure 10 (center), updated model 2.
2. *Increasing the complexity of CPDs.* We reduce the model misspecification η_{l^*} by picking a better model \tilde{P}_{l^*} than the baseline model P_{l^*} for the l^* component. The new model should represent the (fixed) available data more accurately by using a




Figure 10. (Left) DFT-computed data for reaction energies with respect to different metals/oxygen binding energies. Here bimetallics data are also included in addition to the single metals in Figure 7(c). (Center) Different relative model sensitivities (4.6) when we only perturb the model of ω_{c1} by $\eta_{c1} = R(\text{data}||P_{c1})$ when P_{c1} is Gaussian with the original single-metal data; or using a KDE given by (1.3) with the original data (updated model 1); or using a Gaussian with the additional bimetallics data (updated model 2); or using both KDE and bimetallics data (updated model 3). (Right) Baseline model (Gaussian) of ω_{c1} (red curve) and the updated model (normal-kernel density estimation, blue curve) and additional bimetallics data in this figure (Left).

kernel-based method. In this case the new model is a mixture of Gaussian and kernel-based networks [48]. For example, we replace the linear, Gaussian model for ω_{c1} demonstrated in Figure 7(g) with a linear, kernel-based model as shown in Figure 10 (right). Then we can reduce the model sensitivity indices by decreasing the model misspecification η_* without introducing new errors into the remaining components of the Bayesian network as proved in Theorem 5.1 (a). Moreover, we can combine the approaches above to reduce the model sensitivity indices. For example, after adding more bimetallics data, we first reduce the model sensitivity indices for the correlation errors ω_{ci} . Then we further reduce the indices of ω_{c1} by replacing the corresponding component of the baseline model for ω_{c1} (Gaussian model) by normal kernel density estimator without increasing the indices of the remaining nodes (see Theorem 5.1 (a)). The new model is the updated model 3 in Figure 10 (center). We can compute the model sensitivity indices for the updated mixed model, where P_i could be KDE or another distribution, using Theorem 3.3 and in particular (3.18).

3. *Increasing the complexity of the graph.* Here, we discuss how model sensitivity indices can investigate the change in graph structure. The available data for solvation energies in Figure 11 (left) indicate that there might be a linear dependence between ω_{s1} and ω_{s2} . We represent such a connection as a directed edge $\omega_{s1} \rightarrow \omega_{s2}$, and thus the new graph has an extra edge, illustrated in orange in Figure 11 (right). The CPDs of the new Bayesian Network Q are given by

$$(7.11) \quad q(\omega_{s1}) := p(\omega_{s1}) = \mathcal{N}(\beta_{s1,0}, \sigma_{s1}^2),$$

$$(7.12) \quad q(\omega_{s2}|\omega_{s1}) := \mathcal{N}(\omega_{s1} + \beta_{s2,0}, \sigma_{s2}^2),$$

and all the remaining ones (i.e., x, y_1, y_2 and all ω_j with $j \neq s_2$) are the same and given by (7.4)–(7.6). The correlation parameters $\beta_{s1,0}, \beta_{s2,0}$ as well as $\sigma_{s1}^2, \sigma_{s2}^2$ are




Figure 11. (left) DFT data for solvation energies ω_{s_0} and ω_{s_i} of x and y_i , respectively, with different water layers. (right) Based on the left figure, a potential correlation between ω_{s_i} is found. We incorporate such a correlation into the graph by adding a new edge between ω_{s_i} (orange edge) into the existing graph in Figure 7 (b). The two energies y_1 and y_2 are now not conditionally independent given x . However, by using (4.6), the model sensitivity indices $I^\pm(x_{O^*}^*, P; \mathcal{D}_{s_2}^{\eta_{s_2}})$ are very small compared to the QoI $\mathbb{E}_P[x_{O^*}^*]$ (here the index in (4.5) is normalized by the QoI), implying that we can ignore the proposed graph connection.

learned by using the global likelihood decomposition method mentioned earlier. The KL divergence between the Gaussian Bayesian networks P and Q of Figures 7(b) and Figure 11, respectively, is given by

$$(7.13) \quad R(Q\|P) = \int \log \frac{q(\omega_{s_2}|\omega_{s_1})}{p(\omega_{s_2})} q(\omega_{s_2}|\omega_{s_1}) q(\omega_{s_1}) ds_2 ds_1$$

and serves as a surrogate for the model misspecification η_{s_2} . Using Gaussianity $\eta_{s_2} = 0.9173$, and by Theorem 3.2, $I^\pm(x_{O^*}^*, P; \mathcal{D}_{s_2}^{\eta_{s_2}}) = \pm 0.0928$. The latter value is very small compared to the QoI $\mathbb{E}_P[x_{O^*}^*] = 2.0434$. Thus, we may safely ignore the correlation between ω_{s_1} and ω_{s_2} . Therefore, no further model improvement is necessary, and we can retain the (simpler) baseline Bayesian network of Figure 7(b).

8. Other divergences and probabilistic metrics. A question that arises naturally is the selection of the probabilistic metric or divergence in the formulation of robust uncertainty quantification, e.g., in the definition of model uncertainty indices (1.3). In this paper we selected the KL divergence to define the ambiguity sets (1.2) as well as structured localized ambiguity sets given by (3.1) and (3.2) since they allowed us to obtain easily computable and scalable model uncertainty indices. In fact, we developed UQ methods that build on UQ information inequalities for QoIs of high-dimensional probabilistic models and stochastic processes [15, 21, 39, 9, 8], taking advantage of the directed nature of the graph structure of Bayesian networks as well as the factorization into CPDs of the joint density given by (1.1). Furthermore, the chain rule of KL divergence allowed us to break down the calculation of KL divergence in terms of conditional KL divergences and thus isolate the uncertainty impact on QoIs from multiple model components as well as to quantify their propagation across the graph to the QoIs.

However, for Bayesian networks with vastly different graphical structures, e.g., an alternative model with more vertices than the baseline, the choice of KL is not suitable due to the lack of absolute continuity between the baseline and the alternative model. In such cases, new divergences could be considered, e.g., Wasserstein metrics already studied in the DRO literature [53, 12] or their integral probability metrics (IPM) generalization [55]; alternatively we can consider various interpolations of divergences and IPMs studied recently in the machine

learning literature such as [29, 23, 6, 32, 30] and references therein. For instance, the recently introduced (f, Γ) -divergences [6] are interpolations of f -divergences and IPMs that combine advantageous features of both, such as the capability to handle heavy-tailed data (property inherited from f -divergences) and to compare nonabsolutely continuous distributions (inherited from IPMs). Also, in [32] they propose an interpolation between KL divergence and maximum mean discrepancy (MMD) termed KALE which inherits the sensitivity to mismatch in the support of the distributions from KL and is well-defined for mutually singular distributions, while in [30] an interpolation between p -Wasserstein and MMD has been proposed. Even though these new divergences seem promising for the reasons we mentioned above, the corresponding UQ methods have not been developed yet, with the exception of the Wasserstein DRO methods [53, 12].

An additional issue that we touched upon here when we discussed model sensitivity indices is the need for divergences to be able to isolate sources of uncertainty on localized parts of the graphical model in the spirit of “divide and conquer.” In that respect concepts of subadditivity of divergences, e.g., subadditivity of f -divergences, total variation, Jensen–Shannon divergence, and Wasserstein, for PGMs [18, 23] could also prove useful.

Appendix A. Background on model uncertainty.

A.1. Mathematical formulation of model uncertainty. We can formulate mathematically model uncertainty by constructing (nonparametric) families \mathcal{Q} of *alternative models* Q to compare to a *baseline model* P which is computationally tractable and inferred from data, and believed to be a good approximation for the physical model of X , while the “true,” intractable, partially unknown model Q^* should belong to \mathcal{Q} ; for this reason we refer to \mathcal{Q} as the *ambiguity set*, typically defined as a neighborhood of models around the baseline P :

$$(A.1) \quad \mathcal{Q} = \mathcal{D}^\eta = \{Q : d(Q, P) \leq \eta\},$$

where $\eta > 0$ corresponds to the size of the ambiguity set and $d = d(Q, P)$ denotes a probability metric or divergence (see Figure 12 (left) for the schematic depiction where d is the KL divergence (a.k.a. relative entropy) $R(Q\|P)$ [16]). The next natural mathematical goal is to assess the baseline model and understand the resulting biases for QoIs f when we use P for predictions instead of the true model $Q^* \in \mathcal{Q}$. As we see later, the free energies $f = -\Delta G_i$ are considered as QoIs for the ORR PGM (see section 7).

We define the *predictive uncertainty* (or bias) for the QoI f when using the baseline model P instead of any alternative model $Q \in \mathcal{Q}$ as the two worst-case scenarios:

$$(A.2) \quad I^\pm(f, P; \mathcal{Q}) := \sup/\inf_{Q \in \mathcal{Q}} \{\mathbb{E}_Q[f] - \mathbb{E}_P[f]\},$$

where $\mathbb{E}_Q[f]$ denotes the expected value of the QoI f . Therefore, (A.2) provides a *robust performance guarantee* for the predictions of the baseline model P for the QoI f within the ambiguity set \mathcal{Q} . This robust perspective for general probabilistic models P is also known in operations research as *distributionally robust optimization* (DRO); see, e.g., [17, 33, 74, 44, 28, 49, 53, 75, 12], where optimal-transport (Wasserstein) metrics were recently proposed for (A.1). Note that the predictive uncertainty represents the robustness of the model P with respect to \mathcal{Q} , i.e., all the biases between the predictions of f with $Q \in \mathcal{Q}$ and P are bounded by the predictive uncertainty.

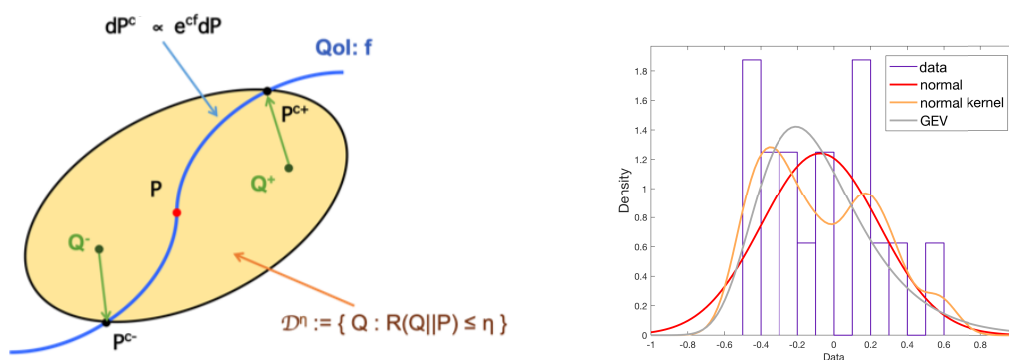


Figure 12. (Left) The schematic illustration of the ambiguity set (nonparametric family of models) given by (A.1) with d being the KL Divergence $R(Q||P)$; the blue line represents a parametric family; Q^\pm are the probability measures that the UQ indices/bounds I^\pm with respect to QoI f are attained and are provided by (A.4), i.e., tightness of the bounds. (Right) Three probabilistic models with different CPDs for sparse data of an ORR PGM vertex ω_{a0} : the red curve is used to build a baseline Gaussian model denoted by P , the gray curve is another parametric model (generalized extreme value (GEV) distribution) which fits the data better, and the yellow curve is a nonparametric model (kernel density estimation (KDE) with normal kernel).

A.2. Existing results on model uncertainty. While the definition (A.2) is rather natural and intuitive, at least based on the model uncertainty challenge depicted in Figure 12 (left), it is not obvious that it is computable in practice. However, it becomes tractable if we use for metric d in (A.1) the KL divergence $R(Q||P)$. Accordingly, η is a measure of the confidence we put in the baseline model P measured using KL divergence. In recent work [15, 21, 37, 47], it has been shown that $I^\pm(f, P; \mathcal{D}^\eta)$ (an infinite-dimensional optimization problem) can be directly computable by a one-dimensional optimization problem:

$$(A.3) \quad I^\pm(f, P; \mathcal{D}^\eta) = \pm \inf_{c>0} \left[\frac{1}{c} \log \int e^{\pm c(f - \mathbb{E}_P[f])} P(dx) + \frac{\eta}{c} \right] = \mathbb{E}_{Q^\pm} [f] - \mathbb{E}_P [f],$$

which is derived by using the Gibbs variational principle [21] for KL divergence. In the first equality of this formula we recognize two ingredients: η is model uncertainty from (A.1), while the moment generating function (MGF) $\int e^{\pm cf} P(dx)$ encodes the QoI f at the baseline model P . In [21, 37, 47] techniques were developed to compute (exactly or approximately via asymptotics [21]) as well as provide explicitly upper and lower bounds on $I^\pm(f, P; \mathcal{D}^\eta)$ in terms of concentration inequalities [37]. A key point in (A.3) is that the parameter η is not necessarily small, allowing *global and nonparametric* sensitivity analysis.

Moreover, in [37] the authors have proven that the second equality of (A.3) holds. In fact, this shows that $I^\pm(f, P; \mathcal{D}^\eta)$ is also tight, i.e., when the sup and inf in (A.2) are attained by appropriate measures Q^\pm . Formally, the authors have shown that there exist $0 < \eta_\pm \leq \infty$, such that for any $\eta \leq \eta_\pm$, $Q^\pm(\cdot) = Q^\pm(\cdot; \pm c_\pm)$ depend on η and are given by

$$(A.4) \quad dQ^\pm = \frac{e^{\pm c_\pm f}}{\mathbb{E}_P [e^{\pm c_\pm f}]} dP,$$

where $c_\pm \equiv c_\pm(\eta)$ are the unique solutions of

$$(A.5) \quad R(Q^\pm || P) = \eta.$$

A.3. Some fundamental lemmas. In this subsection, we include Lemmas A.1 and A.2 for completeness of the background presentation. These results were proved in [21, 22, 37], and we present them here for the convenience of the reader.

Lemma A.1. *Let P be a probability measure, and let $f(X)$ be such that its MGF is finite in a neighborhood of the origin. Then for any Q with $R(Q||P) < \infty$, we have*

$$(A.6) \quad -\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{-c\bar{f}(X)} \right] + \frac{\eta}{c} \right] \leq \mathbb{E}_Q [f(X)] - \mathbb{E}_P [f(X)] \leq \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{c\bar{f}(X)} \right] + \frac{\eta}{c} \right].$$

Proof of Lemma A.1. For any general QoI $f(X)$ which has finite MGF, $\mathbb{E}_P \left[e^{\pm c\bar{f}(X)} \right] := \mathbb{E}_P \left[e^{c(f(X) - \mathbb{E}_P[f(X)])} \right]$, in a neighborhood of the origin, there is a known fact in statistics and large deviation theory [20, 21] that

$$(A.7) \quad \log \mathbb{E}_P \left[e^{f(X)} \right] = \sup_{Q \ll P} \{ \mathbb{E}_Q [f(X)] - R(Q||P) \} .$$

Changing $f(X)$ to $c\bar{f}(X) = c(f(X) - \mathbb{E}_P[f(X)])$, we get

$$(A.8) \quad \mathbb{E}_P \left[e^{\pm c\bar{f}(X)} \right] = \sup_{Q \ll P} \{ c(\mathbb{E}_Q [f(X)] - \mathbb{E}_P [f(X)]) - R(Q||P) \},$$

which gives us the following upper and lower bounds with $c > 0$:

$$(A.9) \quad -\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{-c\bar{f}(X)} \right] + \frac{\eta}{c} \right] \leq \mathbb{E}_Q [f(X)] - \mathbb{E}_P [f(X)] \leq \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{c\bar{f}(X)} \right] + \frac{\eta}{c} \right],$$

where $\eta = R(Q||P)$. ■

Lemma A.2. *Suppose (d_-, d_+) is the largest open set such that the cumulant generating function $\Lambda(c) = \log \mathbb{E}_P \left[e^{c\bar{f}(X)} \right] < \infty$ for all $c \in (d_-, d_+)$.*

1. For any $\eta \geq 0$ the optimization problems

$$\inf_{c>0} \frac{\Lambda(\pm c) + \eta}{c}$$

have unique minimizers $c_{\pm} \in [0, \pm d_{\pm}]$. Let η_{\pm} be defined by

$$\eta_{\pm} = \lim_{c \nearrow \pm d_{\pm}} \pm c\Lambda'(\pm c) - \Lambda(\pm c).$$

Then the minimizers $c_{\pm} = c_{\pm}(\eta)$ are finite for $\eta < \eta_{\pm}$ and $c_{\pm}(\eta) = \pm d_{\pm}$ if $\eta \geq \eta_{\pm}$.

2. If $c_{\pm}(\eta) < \pm d_{\pm}$, then

$$(A.10) \quad \frac{\Lambda(\pm c_{\pm}) + \eta}{c_{\pm}} = \inf_{c>0} \frac{\Lambda(\pm c) + \eta}{c} = \pm \Lambda'(\pm c_{\pm}) = \pm \left(\mathbb{E}_{P_{\pm c_{\pm}}} [f] - \mathbb{E}_P [f] \right),$$

where $c_{\pm}(\eta)$ is strictly increasing in η and is determined by the equation

$$(A.11) \quad R(P_{\pm c_{\pm}} || P) = \eta.$$

3. η_{\pm} is finite in two distinct cases.

(a) If $\pm d_{\pm} < \infty$ (in which case g must be unbounded above/below), η_{\pm} is finite if $\lim_{c \rightarrow \pm d_{\pm}} \Lambda(\pm c) := \Lambda(d_{\pm}) < \infty$ and $\lim_{c \rightarrow \pm d_{\pm}} \pm \Lambda'(\pm c) := \pm \Lambda'(d_{\pm}) < \infty$, and for $\eta \geq \eta_{\pm}$ we have

$$(A.12) \quad \inf_{c>0} \frac{\Lambda(\pm c) + \eta}{c} = \frac{\Lambda(d_{\pm}) + \eta}{\pm d_{\pm}} = \pm \left(\mathbb{E}_{P_{d_{\pm}}} [f] - \mathbb{E}_P [f] \right) + \frac{\eta - \eta_{\pm}}{\pm d_{\pm}}.$$

(b) If $\pm d_{\pm} = \infty$ and η_{\pm} is finite, then f is P -a.s. bounded above/below and for $\eta \geq \eta_{\pm}$ we have

$$(A.13) \quad \inf_{c>0} \frac{\Lambda(\pm c) + \eta}{c} = \text{ess sup}_{x \in \mathcal{X}} \{ \pm (f(x) - \mathbb{E}_P [f(X)]) \}.$$

Proof of the Lemma A.2. For notational ease, in the proof, let us set $\Lambda(c) = \log \mathbb{E}_P [e^{c\tilde{f}(X)}]$ so that the UQ indices are

$$I^{\pm}(f(X), P; \mathcal{D}^{\eta}) = \inf_{c>0} \left\{ \frac{\Lambda(\pm c) + \eta}{c} \right\}.$$

Note that $\Lambda(c)$ is a convex function which we assume to be finite on an interval (d_-, d_+) with $d_- < 0 < d_+$. On that interval $\Lambda(c)$ is infinitely differentiable and strictly convex. Since we centered the QoI we have $\Lambda(0) = \Lambda'(0) = 0$ and $\Lambda''(0) = \text{Var}_P(f)$.

First note that it is enough to prove the result for $\Lambda(c)$ since the result for $\Lambda(-c)$ is obtained by replacing f by $-f$. We also use the notation $\tilde{f}_+ = \text{ess sup} \{ f(x) - \mathbb{E}_P [f(X)] \}$.

We first claim that automatically

$$\Lambda(d_+) = \lim_{c \nearrow d_+} \Lambda(c),$$

where $\Lambda(d_+)$ may be infinite. By monotone convergence

$$\mathbb{E}_P [1_{\{\tilde{f} \geq 0\}} e^{c\tilde{f}}] \nearrow \mathbb{E}_P [1_{\{\tilde{f} \geq 0\}} e^{d_+\tilde{f}}]$$

as $c \nearrow d_+$. By dominated convergence

$$\mathbb{E}_P [1_{\{\tilde{f} < 0\}} e^{c\tilde{f}}] \searrow \mathbb{E}_P [1_{\{\tilde{f} < 0\}} e^{d_+\tilde{f}}]$$

as $c \nearrow d_+$, and the claim follows. A very similar argument shows that $\Lambda'(c)$ also has a limit as $c \nearrow d_+$.

Let

$$(A.14) \quad B(c; \eta) = \frac{\Lambda(c) + \eta}{c}.$$

We divide this into cases.

1. $\tilde{f}_+ < \infty$. In this case $\Lambda'(c) \nearrow \tilde{f}_+ < \infty$ as $c \rightarrow \infty$ and $\Lambda'(0) < \tilde{f}_+$. If $\eta = 0$, then the infimum is $\Lambda'(0)$ and attained at $c_+ = 0$ since $\Lambda(c)/c$ is an increasing function. If $\eta > 0$, then

$$B'(c; \eta) = \frac{c\Lambda'(c) - \Lambda(c) - \eta}{c^2}$$

for $c \geq 0$. The function $c\Lambda'(c) - \Lambda(c)$ strictly increases from 0 at $c = 0$ to some limit $\eta_+ > 0$ at $c = \infty$, and the minimizer is at the unique finite root of $c\Lambda'(c) - \Lambda(c) = \eta$ for $\eta < \eta_+$ and $c_+ = \infty$ for $\eta \geq \eta_+$.

2. $\tilde{f}_+ = \infty$. In this case there are two subcases.
 - (a) $d_+ = \infty$. In this case since $\tilde{f}_+ = \infty$ we have $\Lambda'(c) \nearrow \infty$ as $c \rightarrow \infty$ and $c\Lambda'(c) - \Lambda(c) \rightarrow \infty$ as $c \rightarrow \infty$. Since $0\Lambda'(0) - \Lambda(0) = 0$, in all cases of $\eta \geq 0$ there is a unique root to $c\Lambda'(c) - \Lambda(c) = \eta$ and hence a unique minimizer.
 - (b) $d_+ < \infty$. We know that $\Lambda'(c)$ converges as $c \nearrow d_+$ to a well-defined left-hand limit which we call $\Lambda'(d_+)$ (note that this value could be ∞). Thus we have that $c\Lambda'(c) - \Lambda(c)$ ranges from 0 at $c = 0$ to $\eta_+ = d_+\Lambda'(d_+) - \Lambda(d_+)$. For $\eta \in [0, \eta_+)$ there is a unique minimizer in $[0, d_+)$. For $\eta \geq \eta_+$ the unique minimizer is at $c_+ = d_+$.

To conclude the proof we note that if $c_+ < d_+$, then an easy computation shows that

$$c_+\Lambda'(c_+) - \Lambda(c_+) = R(P_{c_+} || P) = \eta,$$

and thus

$$B(c_+, \eta) = \Lambda'(c_+) = \mathbb{E}_{P_{c_+}}[f] - \mathbb{E}_P[f(X)],$$

which proves (A.10) and (A.5). Finally if $d_+ = \infty$ and f is P -a.s. bounded above, then the infimum is equal to $\lim_{c \rightarrow \infty} \frac{\Lambda(c)}{c}$ and this establishes (A.13). If $d_+ < \infty$ and $\eta_+ < \infty$, then the bound takes the form (A.12). ■

Appendix B. A simple example for Bayesian networks.

Example B.1. In this example, we focus on the construction of the graph structure and CPDs of the optimal distributions provided by Theorem 2.1 (b) following the strategy of its proof. Note that in the next subsection by assuming that each X_i is linear Gaussian of its parents, we also compute the model uncertainty indices given by (2.5) in Theorem 2.1 (a). Let us consider a Bayesian network as shown in Figure 2(a), with density given by

$$(B.1) \quad p(x) = p(x_1)p(x_2)p(x_3|x_2, x_1)p(x_4)p(x_5|x_3)p(x_6|x_4, x_3)p(x_7|x_6, x_5)p(x_8|x_6).$$

For a QoI $f(X_6)$, the optimizers in Theorem 2.1 (b) are obtained when the CPDs of $X_5, X_7,$ and X_8 are the same with the corresponding CPDs of P as these vertices are not ancestors of X_6 while

$$(B.2) \quad q^\pm(x_6|x_{\pi_6^\pm}) = \frac{e^{\pm c_\pm f(x_6)}}{\mathbb{E}_{P_{\delta\{4,3\}}}[e^{\pm c_\pm f(X_6)}]} \cdot p(x_6|x_4, x_3),$$

where $\pi_6^{Q^\pm} \equiv \pi_6^P = \{4, 3\}$, then for $i \in \rho_6 = \{1, 2, 3, 4\}$

$$(B.3) \quad q^\pm(x_4|x_{\pi_4^{Q^\pm}}) = \frac{\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm f(X_6)}]}{\mathbb{E}_{P_4} [\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm f(X_6)}]]} p(x_4)$$

since both normalization factors on the numerator and denominator depend on $X_{\pi_6} = \{X_4, X_3\}$. So in general we have $\pi_4^{Q^\pm} = \pi_4^P \cup \{3\} = \{3\}$, i.e., there is a new connection $X_3 \rightarrow X_4$ in Q^\pm , and

$$(B.4) \quad q^\pm(x_3|x_{\pi_3^{Q^\pm}}) = \frac{\mathbb{E}_{P_4} [\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm f(X_6)}]]}{\mathbb{E}_{P_{3|\{2,1\}}} [\mathbb{E}_{P_4} [\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm f(X_6)}]]]} p(x_3|x_2, x_1),$$

where $\pi_3^{Q^\pm} \equiv \pi_3^P = \{2, 1\}$ since the normalization factors do not contain other variables. We can similarly do the same for X_2 and X_1 to get the entire structure of Q^\pm which has another new connection $X_1 \rightarrow X_2$, and the results are shown in Figure 2(b).

For $A = \{3, 6, 7\}$, we consider a QoI $f(X_A) = f(X_3, X_6, X_7)$, and by Theorem 2.1 (a) the following holds:

$$(B.5) \quad \begin{aligned} I^\pm(f(X_3, X_6, X_7), P; \mathcal{D}^\eta) &= \pm \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_A} [e^{\pm c f(X_A)}] + \frac{\eta}{c} \right] \\ &= \mathbb{E}_{Q^\pm} [f(X_A)] - \mathbb{E}_P [f(X_A)], \end{aligned}$$

where Q^\pm are the optimizers with CPDs given by (B.6)–(B.13) and

$$\mathbb{E}_{P_A} [e^{\pm c f(X_A)}] = \int e^{\pm c f(x_3, x_6, x_7)} \prod_{i=1}^7 p(x_i|x_{\pi_i}) dx_i.$$

We recall (2.8) of Theorem 2.1 (b), and we obtain the CPDs of Q^\pm and the new parents of each vertex as follows:

$$(B.6) \quad q^\pm(x_8|x_{\pi_8^{Q^\pm}}) \equiv p(x_8|x_{\pi_8^{Q^\pm}}) \equiv p(x_8|x_6),$$

$$(B.7) \quad q^\pm(x_7|x_{\pi_7^{Q^\pm}}) = \frac{e^{\pm c_\pm f(x_7, x_6, x_3)}}{\mathbb{E}_{P_{7|\{6,5\}}} [e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_7|x_6, x_5)$$

with $\pi_7^P \subset \pi_7^{Q^\pm} = \pi_7^P \cup \{3\} = \{6, 5, 3\}$. Let $\{l_1, \dots, l_6\} \equiv \rho_3^P \cup \rho_6^P \cup \rho_7^P \cup \{3, 6\} = \{1, 2, 3, 4, 5, 6\}$. We start with X_6 as it is indexed by the $\max\{l_j : j \in 1, \dots, 6\}$

$$(B.8) \quad q^\pm(x_6|x_{\pi_6^{Q^\pm}}) = \frac{\mathbb{E}_{P_{7|\pi_7}} [e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{6|\pi_6}, P_{7|\pi_7}} [e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_6|x_4, x_3)$$

with $\pi_6^P \subset \pi_6^{Q^\pm} \subset \pi_6^P \cup \{5\} = \{3, 4, 5\}$. Similarly, the CPDs of X_1, \dots, X_5 are given by

$$(B.9) \quad q^\pm(x_5|x_{\pi_5^{Q^\pm}}) = \frac{\mathbb{E}_{P_{6|\pi_6}, P_{7|\pi_7}}[e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{5|\pi_5}, P_{6|\pi_6}, P_{7|\pi_7}}[e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_5|x_3),$$

$$(B.10) \quad q^\pm(x_4|x_{\pi_4^{Q^\pm}}) = \frac{\mathbb{E}_{P_{5|\pi_5}, P_{6|\pi_6}, P_{7|\pi_7}}[e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{4|\pi_4}, P_{5|\pi_5}, P_{6|\pi_6}, P_{7|\pi_7}}[e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_4),$$

$$(B.11) \quad q^\pm(x_3|x_{\pi_3^{Q^\pm}}) = \frac{\mathbb{E}_{P_{4|\pi_4}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{3|\pi_3}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_3|x_2, x_1),$$

$$(B.12) \quad q^\pm(x_2|x_{\pi_2^{Q^\pm}}) = \frac{\mathbb{E}_{P_{3|\pi_3}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{2|\pi_2}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_2|x_1),$$

$$(B.13) \quad q^\pm(x_1|x_{\pi_1^{Q^\pm}}) = \frac{\mathbb{E}_{P_{2|\pi_2}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_A)}]}{\mathbb{E}_{P_{1|\pi_1}, \dots, P_{7|\pi_7}}[e^{\pm c_\pm f(X_7, X_6, X_3)}]} \cdot p(x_1),$$

where the expectations involved in the above formulas are given by (2.4). The corresponding structures are

$$(B.14) \quad \pi_5^P \subset \pi_5^{Q^\pm} \subset \pi_5^P \cup \{4\} = \{3, 4\},$$

$$(B.15) \quad \pi_4^P \subset \pi_4^{Q^\pm} \subset \pi_4^P \cup \{3\} = \{3\},$$

$$(B.16) \quad \pi_3^P \subset \pi_3^{Q^\pm} \subset \pi_3^P = \{1, 2\},$$

$$(B.17) \quad \pi_2^P \subset \pi_2^{Q^\pm} \subset \pi_2^P \cup \{1\} = \{1\},$$

$$(B.18) \quad \pi_1^{Q^\pm} = \pi_1^P = \emptyset.$$

As a result, the structure of the associated graph to Q^\pm may change and, in particular, the vertices—with potentially extra parents—are X_2, X_4, X_5, X_6 , and X_7 as illustrated in Figure 2(c).

Appendix C. A simple example on nondirected graphical model which belongs to ambiguity set (2.1). Let us assume that P is a Gaussian Bayesian network defined on $X = \{X_1, \dots, X_8\}$ with graph structure given in Figure 13(a) and CPDs denoted by $p(x_i|x_{\pi_i^P})$ defined in section 2.1. An alternative mixed graphical model Q with structure as shown in Figure 13(b) with joint density given by

$$(C.1) \quad q(x) = q(x_1, x_2) \prod_{i=3}^8 p(x_i|x_{\pi_i})$$

is considered here. In (C.1), we keep the CPDs $p(x_i|x_{\pi_i})$ for $i = 3, \dots, 8$ the same as the baseline Gaussian Bayesian network, while $p(x_1|x_{\pi_1})p(x_2|x_{\pi_2})$ is replaced by a bivariate Gaussian $q(x_1, x_2)$, with $\rho \neq 0$ being the correlation between X_1 and X_2 , and corresponds to this new nondirected part of the graph, i.e.,

$$(C.2) \quad q(x_1, x_2) = \frac{1}{Z(\tilde{\sigma}_1, \tilde{\sigma}_2)} e^{-\frac{1}{2(1-\rho^2)} \left[\left(\frac{x_1 - \tilde{\mu}_1}{\tilde{\sigma}_1} \right)^2 + \left(\frac{x_2 - \tilde{\mu}_2}{\tilde{\sigma}_2} \right)^2 - 2\rho \frac{x_1 - \tilde{\mu}_1}{\tilde{\sigma}_1} \frac{x_2 - \tilde{\mu}_2}{\tilde{\sigma}_2} \right]},$$

where $(\tilde{\mu}_i)_{i=1,2}$ is the mean vector and $\tilde{\sigma}_i > 0$ for $i = 1, 2$. Also $Z(\tilde{\sigma}_1, \tilde{\sigma}_2) = 2\pi\tilde{\sigma}_1\tilde{\sigma}_2\sqrt{1-\rho^2}$.

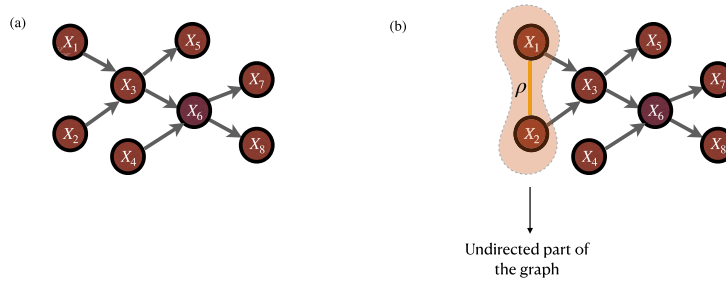


Figure 13. (a) Example of the graph structure of a baseline Gaussian Bayesian network P and the corresponding random variables $X = \{X_1, \dots, X_8\}$. The CPDs $p(x_i|x_{\pi_i^P})$ are defined in section 2.1. (b) An example of a PGM $Q \in \mathcal{D}^n$ with a new undirected edge is highlighted in orange. The joint probability density between X_1 and X_2 , $q(x_1, x_2)$, is a bivariate Gaussian with correlation $\rho \neq 0$. Note that the correlation in the corresponding part of $p(x)$ of the baseline model is 0.

We now compute the likelihood ratio between the models, i.e.,

$$\frac{dQ}{dP} = \frac{q(x_1, x_2)}{p(x_1|x_{\pi_1})p(x_2|x_{\pi_2})},$$

and thus the KL divergence between the models is given by

$$\begin{aligned} R(Q\|P) &= E_Q \left[\log \left(\frac{q(x_1, x_2)}{p(x_1|x_{\pi_1})p(x_2|x_{\pi_2})} \right) \right] \\ &= \frac{1}{2} \left(\frac{\tilde{\sigma}_1}{\sigma_1} \right)^2 + \frac{1}{2} \left(\frac{\tilde{\sigma}_2}{\sigma_2} \right)^2 + \frac{1}{2} \left(\frac{\tilde{\mu}_1 - \beta_{10}}{\sigma_1} \right)^2 + \frac{1}{2} \left(\frac{\tilde{\mu}_2 - \beta_{20}}{\sigma_2} \right)^2 - 1 \\ (C.3) \quad & - \frac{1}{2} \log \left(\frac{\tilde{\sigma}_1}{\sigma_1} \right) - \frac{1}{2} \log \left(\frac{\tilde{\sigma}_2}{\sigma_2} \right) - \frac{1}{2} \log(1 - \rho^2). \end{aligned}$$

Given a small $\eta > 0$, we choose $\tilde{\sigma}_i = \sigma_i$, $i = 1, 2$, and thus (C.3) takes the form

$$(C.4) \quad R(Q\|P) = \frac{1}{2} \left(\frac{\tilde{\mu}_1 - \beta_{10}}{\sigma_1} \right)^2 + \frac{1}{2} \left(\frac{\tilde{\mu}_2 - \beta_{20}}{\sigma_2} \right)^2 - \frac{1}{2} \log(1 - \rho^2).$$

Additionally, an appropriately small perturbation around β_{i0} resulting in a small $\tilde{\mu}_i - \beta_{i0}$ can lead to $R(Q\|P) \leq \eta$ under a suitable choice of σ_i, ρ , $i = 1, 2$. Therefore, we constructed an example where for small η 's an ambiguity set can include a model with nondirected structure, i.e., with both directed and undirected edges.

Appendix D. A simple example for Gaussian Bayesian networks.

Example D.1 (continuation of Example B.1). We assume that CPDs of Example B.1 with graph structure as in Figure 2(a) are given by

$$\begin{aligned} p(x_8|x_6) &= \mathcal{N}(\beta_{80} + \beta_{86}x_6, \sigma_8^2), & p(x_4) &= \mathcal{N}(\beta_{40}, \sigma_4^2), \\ p(x_7|x_6, x_5) &= \mathcal{N}(\beta_{70} + \beta_{76}x_6 + \beta_{75}x_5, \sigma_7^2), & p(x_3|x_2, x_1) &= \mathcal{N}(\beta_{30} + \beta_{32}x_2 + \beta_{31}x_1, \sigma_3^2), \\ p(x_6|x_4, x_3) &= \mathcal{N}(\beta_{60} + \beta_{64}x_4 + \beta_{63}x_3, \sigma_6^2), & p(x_2) &= \mathcal{N}(\beta_{20}, \sigma_2^2), \\ p(x_5|x_3) &= \mathcal{N}(\beta_{50} + \beta_{53}x_3, \sigma_5^2), & p(x_1) &= \mathcal{N}(\beta_{10}, \sigma_1^2). \end{aligned}$$

Then $p(x) = \mathcal{N}(\mu, \mathcal{C})$ [49, Theorem 7.3]. As before, the QoI depends on X_6 , and for simplicity we consider $f(X_6) = X_6$. For $c > 0$, we compute the MGF of f with respect to P :

$$\mathbb{E}_P \left[e^{\pm c f(X_6)} \right] = \exp \left\{ \frac{c^2}{2} (\sigma_6^2 + \beta_{64}^2 \sigma_4^2 + \beta_{63}^2 \sigma_3^2 + \beta_{63}^2 \beta_{32}^2 \sigma_2^2 + \beta_{63}^2 \beta_{31}^2 \sigma_1^2) \right\} \equiv e^{\frac{c^2}{2} \mathcal{C}_{66}},$$

where $\mathcal{C}_{66} = \sigma_6^2 + \beta_{64}^2 \sigma_4^2 + \beta_{63}^2 \sigma_3^2 + \beta_{63}^2 \beta_{32}^2 \sigma_2^2 + \beta_{63}^2 \beta_{31}^2 \sigma_1^2$ since $\mathbb{E}_P [X_6] = \beta_{60} + \beta_{64} \beta_{40} + \beta_{63} \beta_{30} + \beta_{63} \beta_{32} \beta_{20} + \beta_{63} \beta_{31} \beta_{10}$. We minimize (2.5) in Theorem 2.1 with respect to c :

$$I^\pm(f(X_6), P; \mathcal{D}^\eta) = \pm \inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_P \left[e^{\pm c f(X_6)} \right] + \frac{\eta}{c} \right] = \pm \inf_{c>0} \left[c \mathcal{C}_{66} + \frac{\eta}{c} \right],$$

which in turn gives us the optimizer $c = \sqrt{\frac{\eta}{\mathcal{C}_{66}}}$, and thus

$$(D.1) \quad I^\pm(f(X_6), P; \mathcal{D}^\eta) = \pm \sqrt{2 \mathcal{C}_{66} \eta} = \pm \sqrt{2(\sigma_6^2 + \beta_{64}^2 \sigma_4^2 + \beta_{63}^2 \sigma_3^2 + \beta_{63}^2 \beta_{32}^2 \sigma_2^2 + \beta_{63}^2 \beta_{31}^2 \sigma_1^2) \eta}.$$

By (2.8), the optimizers in Theorem 2.1 are obtained when

$$(D.2) \quad q^\pm(x_i | x_{\pi_i^{Q^\pm}}) = p(x_i | x_{\pi_i}) = \mathcal{N}(\beta_{i0} + \beta_i^T x_{\pi_i}, \sigma_i^2)$$

for $i = 5, 7, 8$, since they are not ancestors of X_6 , and by recalling (B.2)–(B.4) we further compute the CPDs of the remaining vertices as follows: Since $f(X_6) = X_6$ is linear and all random variables are depend linearly on their parents, we appropriately pair the factor $e^{\pm c_\pm x_6}$ with the exponential of the Gaussian CPD $p(x_6 | x_{\pi_6^P})$ and get a new quadratic term in the exponential as well as a term which linearly depends on the parents of X_6 . The latter term is canceled out with the corresponding one in the normalizing factor $\mathbb{E}_{P_{6|\pi_6^P}} [e^{\pm c_\pm X_6}]$ as the parents of X_6 are given. Precisely,

$$\begin{aligned} q^\pm(x_6 | x_{\pi_6^{Q^\pm}}) &= \frac{e^{\pm c_\pm x_6}}{\mathbb{E}_{P_{6|\pi_6^P}} [e^{\pm c_\pm X_6}]} \cdot p(x_6 | x_{\pi_6^P}) \\ &= \frac{\exp \left\{ -\frac{(x_6 - \beta_{60} - \beta_{64} x_4 - \beta_{63} x_2 \mp c_\pm \sigma_6^2)^2}{2\sigma_6^2} \pm c_\pm (\beta_{64} x_4 + \beta_{63} x_3) \right\}}{\int_{\mathcal{X}_6} \exp \left\{ -\frac{(x_6 - \beta_{60} - \beta_{64} x_4 - \beta_{63} x_2 \mp c_\pm \sigma_6^2)^2}{2\sigma_6^2} \pm c_\pm (\beta_{64} x_4 + \beta_{63} x_3) \right\} dx_6}. \end{aligned}$$

Thus,

$$(D.3) \quad q^\pm(x_6 | x_{\pi_6^{Q^\pm}}) = \mathcal{N}(\beta_{60} + \beta_{64} x_4 + \beta_{63} x_3 \pm c_\pm \sigma_6^2, \sigma_6^2), \quad \pi_6^{Q^\pm} \equiv \pi_6^P = \{4, 3\}.$$

Similarly, for $4 \in \rho_6 = \{4, 3, 2, 1\}$

$$q^\pm(x_4 | x_{\pi_4^{Q^\pm}}) = \frac{\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm X_6}]}{\mathbb{E}_{P_4} [\mathbb{E}_{P_{6|\{4,3\}}} [e^{\pm c_\pm X_6}]]} p(x_4) = \frac{e^{-\frac{(x_4 - \beta_{40} \mp c_\pm \beta_{43} \sigma_4^2)^2}{2\sigma_4^2}} e^{\pm c_\pm (\beta_{63} x_3)}}{\int_{\mathcal{X}_4} e^{-\frac{(x_4 - \beta_{40} \mp c_\pm \beta_{43} \sigma_4^2)^2}{2\sigma_4^2}} dx_4 e^{\pm c_\pm (\beta_{63} x_3)}}.$$

Using same the same argument as before, we get

$$(D.4) \quad q^\pm(x_4|x_{\pi_4^{Q^\pm}}) = \mathcal{N}(\beta_{40} \pm c_\pm \sigma_4^2, \sigma_4^2), \quad \pi_4^{Q^\pm} \equiv \pi_4^P = \emptyset.$$

Furthermore,

$$(D.5) \quad q^\pm(x_3|x_{\pi_3^{Q^\pm}}) = \mathcal{N}(\beta_{30} + \beta_{32}x_2 + \beta_{31}x_1 \pm c_\pm \sigma_3^2, \sigma_3^2), \quad \pi_3^{Q^\pm} \equiv \pi_3^P = \{1, 2\},$$

$$(D.6) \quad q^\pm(x_2|x_{\pi_2^{Q^\pm}}) = \mathcal{N}(\beta_{20} \pm c_\pm \sigma_2^2, \sigma_2^2), \quad \pi_2^{Q^\pm} \equiv \pi_2^P = \emptyset,$$

$$(D.7) \quad q^\pm(x_1|x_{\pi_1^{Q^\pm}}) = \mathcal{N}(\beta_{10} \pm c_\pm \sigma_1^2, \sigma_1^2), \quad \pi_1^{Q^\pm} \equiv \pi_1^P = \emptyset.$$

By using the equation $\pm c_\pm \mathbb{E}_{Q^\pm}[X_6] - \log \mathbb{E}_P[e^{\pm c_\pm X_6}] = \eta$, the parameters c_\pm are given by

$$(D.8) \quad c_\pm = \pm \sqrt{\frac{2\eta}{C_{66}}} = \pm \sqrt{\frac{2\eta}{\sigma_6^2 + \beta_{64}^2 \sigma_4^2 + \beta_{63}^2 \sigma_3^2 + \beta_{63}^2 \beta_{32}^2 \sigma_2^2 + \beta_{63}^2 \beta_{31}^2 \sigma_1^2}}.$$

Example D.2 (computation of F for Example D.1). For Example D.1, we compute $F(x_3, \rho_3)$ with $f(X_6) = X_6$ and $l = 3$ (and thus $\rho_6 = \{1, 2, 3, 4\}$ and $\rho_3 = \{1, 2\}$) as

$$(D.9) \quad \begin{aligned} F(x_3, x_{\rho_3}^P) &\equiv F(x_3, x_2, x_1) = \int_{\mathcal{X}_{\{4,6\}}} x_6 p(x_6|x_4, x_3) p(x_4) dx_6 dx_4 \\ &= \beta_{60} + \beta_{64} \beta_{40} + \beta_{63} x_3 = F(x_3). \end{aligned}$$

Example D.3 (computation of β_{kl} and $\tilde{\beta}_{kl}$ for Example D.1). Let us now revisit Example D.1 and compute β_{kl} and $\tilde{\beta}_{kl}$ of Theorem 3.6 when $l \in \pi_k^P$, e.g., $l = 3$, and when $l \in \rho_k \setminus \pi_k$, e.g., $l = 2$, respectively. In the first case, $P_{3|\pi_3}$ is perturbed under the constraint $R(Q_{3|\pi_3}^Q \| P_{3|\pi_3}) \leq \eta_3$ or $R(Q_{3|\pi_3} \| P_{3|\pi_3}) \leq \eta_3$, i.e., consider $Q \in \mathcal{D}_3^{\eta_3}$ or $\mathcal{D}_{3,P}^{\eta_3}$ and $f(X_6) = X_6$. $F(x_3, x_{\rho_3})$ is given by (D.9) and by Theorems 3.2 and 3.3 and (3.7), we can conclude that

$$(D.10) \quad I^\pm(f(X_6), P; \mathcal{D}_3^{\eta_3}) = I^\pm(f(X_6), P; \mathcal{D}_{3,P}^{\eta_3}) = \pm |\beta_{63}| \sqrt{2\sigma_3^2 \eta_3}.$$

In the second case, $P_{2|\pi_2}$ is perturbed under the constraint $R(Q_{2|\pi_2}^Q \| P_{2|\pi_2}) \leq \eta_2$ or $R(Q_{2|\pi_2} \| P_{3|\pi_2}) \leq \eta_2$. We compute $F(x_2, x_{\rho_2}) = \beta_{60} + \beta_{64} \beta_{40} + \beta_{63} \beta_{30} + \beta_{63} \beta_{32} x_2 + \beta_{63} \beta_{31} \beta_{10} = F(x_2)$ and $\tilde{\beta}_{62} = \beta_{63} \beta_{32}$.

Appendix E. Model uncertainty for inhomogeneous Markov chains. We consider the Markov chain models shown in Figure 14 and the QoI $f(X_k)$. Then we only perturb $P_{l|l-1}$ with $l \leq k$, under the constraint $R(Q_{l|\pi_l^Q} \| P_{l|l-1}) \leq \eta_l$. The function $F(x_l, x_{\rho_l^P})$ defined in (3.6) depends only on x_l , and by Theorem 3.2 we have

$$(E.1) \quad I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) = \pm \mathbb{E}_{P_{l|l-1}} \left[\inf_{c>0} \left[\frac{1}{c} \log \mathbb{E}_{P_{l|l-1}} \left[e^{\pm c \bar{F}(X_l, X_{\rho_l})} \right] + \frac{\eta_l}{c} \right] \right].$$



Figure 14. An inhomogeneous Markov chain consists of $X = \{X_1, X_2, \dots, X_n\}$ with $p(x) = \prod_{i=1}^n p(x_i|x_{i-1})$.

Since $F(x_l, x_{\rho_l^P}) = F(x_l)$ the condition on **Theorem 3.3** is satisfied, and therefore we have $I^\pm(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l}) = I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$. To obtain the optimizers in both **Theorem 3.2** and **Theorem 3.3**, we use (3.8)–(3.9) and thus

$$(E.2) \quad q^\pm(x_i|x_{i-1}) \equiv p(x_i|x_{i-1}) \quad \text{for all } i \neq l$$

and

$$(E.3) \quad q^\pm(x_l|x_{l-1}) = \frac{e^{\pm c_\pm(x_{l-1})F(x_l)}}{\mathbb{E}_P [e^{\pm c_\pm(x_{l-1})F(X_l)}|x_{l-1}]} p(x_l|x_{l-1}),$$

where $c_\pm(x_{l-1})$ are the unique solutions of $R(P_{l|l-1}^{\pm c_\pm} \| P_{l|l-1}) = \eta_l$ for all x_{l-1} . Moreover, by perturbing $P_{l|l-1}$, $l > k$, with the constraint

$$R(Q_{l|\pi_l^Q} \| P_{l|l-1}) \leq \eta_l \text{ or } R(Q_{l|l-1} \| P_{l|l-1}) \leq \eta_l,$$

and by **Theorems 3.2** and **3.3**, we have $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l}) = I^\pm(f(X_k), P; \mathcal{D}_{l,P}^{\eta_l}) = 0$. Note that when the ambiguity set is given by (2.1), it includes also Q 's that are *non-Markovian*. However, the optimizers are inhomogeneous Markov chains and are provided by **Theorem 2.1**.

Appendix F. Proof of Lemma 3.1. Since for any $Q \in \mathcal{D}_l^{\eta_l}$, we have $\pi_j^Q \equiv \pi_j^P = \pi_j$ and $Q_{j|\pi_j} \equiv P_{j|\pi_j}$ for all $j \neq l$, therefore, we can rewrite the bias $\mathbb{E}_Q[f(X_k)] - \mathbb{E}_P[f(X_k)]$ as

$$(F.1) \quad \begin{aligned} &= \int_{\mathcal{X}} f(x_k) \prod_{i=1}^n Q(dx_i|x_{\pi_i^Q}) - \int_{\mathcal{X}} f(x_k) \prod_{i=1}^n P(dx_i|x_{\pi_i^P}) \\ &= \int_{\mathcal{X}_k} \int_{\mathcal{X}_{\rho_k^Q}} f(x_k) \prod_{i \in \rho_k^Q \cup \{k\}} Q(dx_i|x_{\pi_i^Q}) - \int_{\mathcal{X}_k} \int_{\mathcal{X}_{\rho_k^P}} f(x_k) \prod_{i \in \rho_k^P \cup \{k\}} P(dx_i|x_{\pi_i^P}) \\ &= \mathbb{E}_{Q_{\{k\}}} [f(X_k)] - \mathbb{E}_{P_{\{k\}}} [f(X_k)]. \end{aligned}$$

If $l \notin \rho_k^P$, we have $\pi_i^Q \equiv \pi_i^P =: \pi_i$ and $Q(dx_i|x_{\pi_i}) \equiv P(dx_i|x_{\pi_i})$ for all $i \in \bar{\rho}_k$; therefore $Q_{\{k\}} \equiv P_{\{k\}}$, and thus $\mathbb{E}_Q[f(X_k)] - \mathbb{E}_P[f(X_k)] = 0$. Based on this calculation for $Q \in \mathcal{D}_l^{\eta_l}$, we stress that our indices capture the graph structure correctly, e.g., perturbations on disconnected

vertices do not affect the QoI $f = f(X_k)$. Since $Q(dx_j|x_{\pi_j}) \equiv P(dx_j|x_{\pi_j})$ for all $j \neq l$, (F.1) equals

$$\begin{aligned}
(F.2) \quad \mathbb{E}_Q[f(X_k)] - \mathbb{E}_P[f(X_k)] &= \mathbb{E}_{Q_{\{k\}}}[f(X_k)] - \mathbb{E}_{P_{\{k\}}}[f(X_k)] \\
&= \int_{\mathcal{X}_{\rho_k^Q \cup \{k\}}} f(x_k) \prod_{i \in \rho_k^Q \cup \{k\} \setminus \rho_l^Q \cup \{l\}} Q(dx_i|x_{\pi_i^Q}) \cdot Q(dx_l|x_{\pi_l^Q}) \cdot \prod_{i \in \rho_l^Q} Q(dx_i|x_{\pi_i^Q}) \\
&\quad - \int_{\mathcal{X}_{\rho_k^P \cup \{k\}}} f(x_k) \prod_{i \in \rho_k^P \setminus \rho_l^P \cup \{l\}} P(dx_i|x_{\pi_i^P}) \cdot P(dx_l|x_{\pi_l^P}) \cdot \prod_{i \in \rho_l^P} P(dx_i|x_{\pi_i^P}) \\
&= \int_{\mathcal{X}_{\rho_k^P \cup \{k\}}} f(x_k) \prod_{i \in \bar{\rho}_k^P \setminus \rho_l^P \cup \{l\}} P(dx_i|x_{\pi_i^P}) \cdot Q(dx_l|x_{\pi_l^Q}) \cdot \prod_{i \in \rho_l^P} P(dx_i|x_{\pi_i^P}) \\
&\quad - \int_{\mathcal{X}_{\rho_k^P \cup \{k\}}} f(x_k) \prod_{i \in \bar{\rho}_k^P \setminus \rho_l^P \cup \{l\}} P(dx_i|x_{\pi_i^P}) \cdot P(dx_l|x_{\pi_l^P}) \cdot \prod_{i \in \rho_l^P} P(dx_i|x_{\pi_i^P}) \\
&= \int_{\mathcal{X}_{\rho_l^P}} \left[\int_{\mathcal{X}_l} F(x_l, x_{\rho_l^P}) Q(dx_l|x_{\pi_l^Q}) - \int_{\mathcal{X}_l} F(x_l, x_{\rho_l^P}) P(dx_l|x_{\pi_l^P}) \right] \prod_{i \in \rho_l^P} P(dx_i|x_{\pi_i^P}) \\
&= \mathbb{E}_{P_{\rho_l^P}} \left[\mathbb{E}_{Q_{l|\pi_l^Q}} [F(X_l, X_{\rho_l^P})] - \mathbb{E}_{P_{l|\pi_l^P}} [F(X_l, X_{\rho_l^P})] \right].
\end{aligned}$$

Appendix G. KL divergence chain rule for Bayesian networks. In this subsection, we discuss the KL chain rule [16] in the context of Bayesian networks as it paves the way for considering suitable ambiguity sets (different from (2.1)) and applying model sensitivity analysis to each component on a baseline Bayesian network. We recall that $P_{i|\pi_i^P}$ is the conditional distribution of X_i with given parents $X_{\pi_i^P} = x_{\pi_i}$, i.e., $P_{i|\pi_i^P}(dx_i) = P(dx_i|x_{\pi_i})$, and for purposes of clarity and to stress the given values, we write $P_{i|X_{\pi_i^P}=x_{\pi_i}}(dx_i)$ instead.

Definition G.1. Let P and Q be two PGMs with densities p and q , respectively, defined as in (1.1). For each $i \in \{1, \dots, n\}$, we define the conditional KL divergence between $Q_{i|X_{\pi_i^Q}}$ and $P_{i|X_{\pi_i^P}}$ with given $X_{\pi_i^Q} = x_{\pi_i^Q}$ and $X_{\pi_i^P} = x_{\pi_i^P}$ as

$$(G.1) \quad R(Q_{i|X_{\pi_i^Q}=x_{\pi_i^Q}} \| P_{i|X_{\pi_i^P}=x_{\pi_i^P}}) = \int_{\mathcal{X}_i} \log \frac{Q(dx_i|x_{\pi_i^Q})}{P(dx_i|x_{\pi_i^P})} Q(dx_i|x_{\pi_i^Q}).$$

Lemma G.2 (chain rule of relative entropy for PGMs). For any two PGMs P and Q with densities $p(x) = \prod_{i=1}^n p(x_i|x_{\pi_i^P})$ and $q(x) = \prod_{i=1}^n q(x_i|x_{\pi_i^Q})$, the KL divergence can be expressed as

$$(G.2) \quad R(Q \| P) = \sum_{i=1}^n \mathbb{E}_{Q_{\pi_i^Q \cup \pi_i^P}} \left[R(Q_{i|X_{\pi_i^Q}} \| P_{i|X_{\pi_i^P}}) \right],$$

where $R(Q_{i|X_{\pi_i^Q}} \| P_{i|X_{\pi_i^P}})$ is the conditional KL divergence given in Definition G.1 and $\mathbb{E}_{Q_{\pi_i^Q \cup \pi_i^P}}$ is the expectation with respect to Q_A defined in section 2 with $A = \pi_i^Q \cup \pi_i^P$.

Proof.

$$\begin{aligned}
 \text{(G.3)} \quad R(Q\|P) &= \int_{\mathcal{X}} \sum_{i=1}^n \log \frac{Q(dx_i|x_{\pi_i^Q})}{P(dx_i|x_{\pi_i^P})} \prod_{j=1}^n Q(dx_j|x_{\pi_j^Q}) \\
 &= \sum_{i=1}^n \int_{\mathcal{X}_{\rho_i^P \cup \rho_i^Q}} \int_{\mathcal{X}_i} \log \frac{Q(dx_i|x_{\pi_i^Q})}{P(dx_i|x_{\pi_i^P})} Q(dx_i|x_{\pi_i^Q}) \prod_{j \in \{\rho_i^Q \cup \rho_i^P\}} Q(dx_j|x_{\pi_j^Q}) \\
 &= \sum_{i=1}^n \mathbb{E}_{Q_{\pi_i^Q \cup \pi_i^P}} \left[R(Q_i|X_{\pi_i^Q} \| P_i|X_{\pi_i^P}) \right]. \quad \blacksquare
 \end{aligned}$$

Appendix H. Schematic for model sensitivity indices. The schematic in Figure 15 refers to the main theorems of section 3.

Appendix I. Data-informed stress tests for Gaussian Bayesian networks. In this section, we explain in detail data-informed stress test analysis when the baseline model P is a Gaussian Bayesian network. Let P be a Gaussian Bayesian network with conditional probability densities $p(x_i|x_{\pi_i})$ satisfying $p(x_i|x_{\pi_i}) = \mathcal{N}(\beta_{i0} + \beta_i^T x_{\pi_i}, \sigma_i^2)$ for some β_{i0} , β_i , and σ_i^2 , i.e., $P_i|_{\pi_i}$ is the conditional distribution of $X_i = \beta_{i0} + \beta_i^T X_{\pi_i} + \epsilon_i$. The random variable ϵ_i has density $p_{\epsilon_i}(x) = \mathcal{N}(0, \sigma_i^2)$ and comes from fitting data with maximum likelihood estimation. Let us consider alternative models to P such as

$$\text{(I.1)} \quad Q_i|_{\pi_i} : \quad X_i = \beta_{i0} + \beta_i^T X_{\pi_i} + \tilde{\epsilon}_i,$$

where $\tilde{\epsilon}_i$ follows another approximate distribution of the data with density $q_{\tilde{\epsilon}_i}(x)$. For instance, we can consider $Q_{\tilde{\epsilon}_i}$ with density $q_{\tilde{\epsilon}_i}$ as the histogram, that is,

$$\text{(I.2)} \quad q_{\tilde{\epsilon}_i}^{hist}(x) = \sum_{k=1}^m \frac{\nu_k}{nh} I(x \in B_k),$$

where B_1, \dots, B_m are the histogram bins, h is the bin width, n is the number of observations, and ν_k is the number of observations in B_k . Alternatively, we can consider the model $Q_{\tilde{\epsilon}_i}$ given by a KDE viewed here as a high-resolution but smooth approximation of the histogram, namely,

$$\text{(I.3)} \quad q_{\tilde{\epsilon}_i}^{KDE}(x) = \sum_{k=1}^n \frac{1}{nh} K\left(\frac{x - x_i}{h}\right),$$

where $K(\cdot)$ is the normal kernel smoothing function with bin width h , and (x_1, \dots, x_n) are the samples of ϵ_i . Other KDE kernels can be considered here (see [72]) or any other probabilistic representations of the data in the histogram. Therefore, for given x_{π_i} , we have

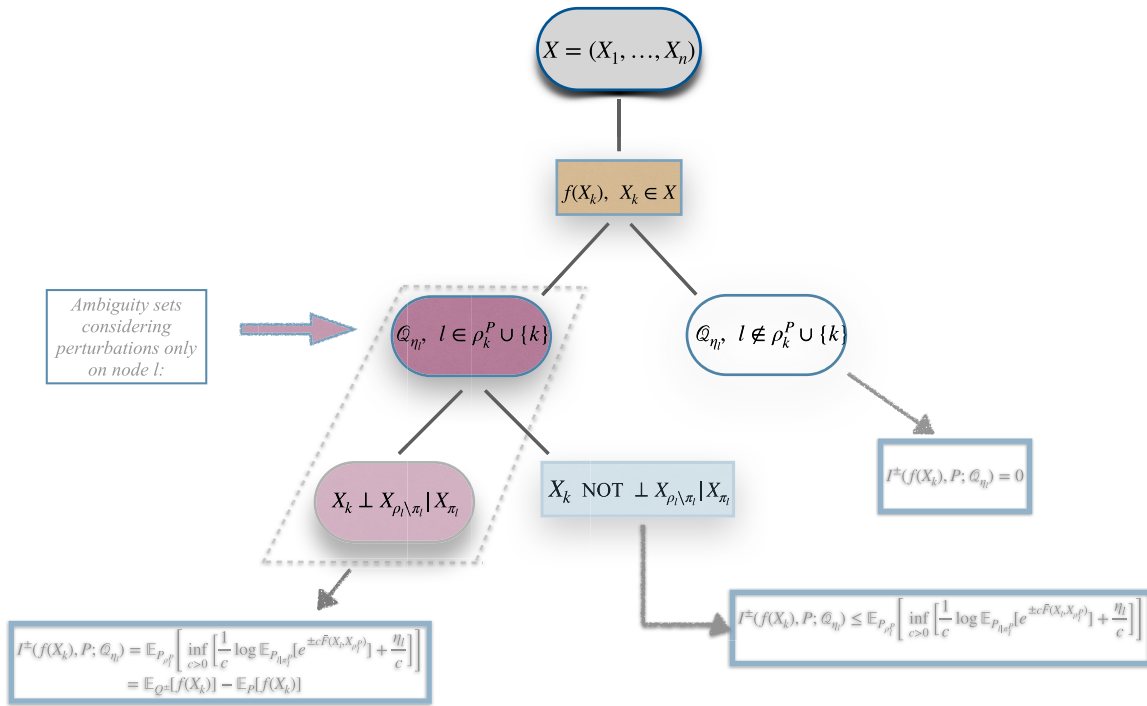


Figure 15. A schematic representation of how the set of vertices of a graph can be decomposed according to the relative position of vertex k that corresponds to the QoI $f(X_k)$, and a vertex l such that perturbations of $P_{l|\pi_l}$ are considered. By Lemma 3.1 and Theorems 3.2 and 3.3, the predictive uncertainty given by (3.4) with l being in different parts of the decomposition varies: The set of vertices is first split as $X = (\bar{\rho}_k^P) \cup (\bar{\rho}_k^P)^c$. The predictive uncertainty $I^\pm(f(X_k), P; \mathcal{D}_l^{\eta_l})$ over models with only perturbed $P_{l|\pi_l}$ when $l \in \bar{\rho}_k^P$ is given by (3.7) and is tight in $\mathcal{D}_l^{\eta_l}$, while perturbation on $P_{l|\pi_l}$ when $l \notin \bar{\rho}_k^P$ does not affect the QoI and thus $I^\pm(f(X_k), P; \mathcal{Q}_{\eta_l}) = 0$. We then decompose the set of vertices $\bar{\rho}_k^P$ into $\{l : X_k \perp X_{\rho_l \setminus \pi_l} | X_{\pi_l}\}$ and its complement. The predictive uncertainty over $\mathcal{D}_{l,P}^{\eta_l}$ with l in the former set is the same as the one for $l \in \bar{\rho}_k^P$, with the difference that it is tight on $\mathcal{D}_{l,P}^{\eta_l}$, contrary to the one in the latter set where the bound is not attained and is also not tight.

$$\begin{aligned}
 \eta_i^{\pi_i} &= \int \log \frac{q(x_i | x_{\pi_i})}{p(x_i | x_{\pi_i})} q(x_i | x_{\pi_i}) dx_i \\
 &= \int \log \frac{q(x_i - \beta_{i0} - \beta_i^T x_{\pi_i} | x_{\pi_i})}{p(x_i - \beta_{i0} - \beta_i^T x_{\pi_i} | x_{\pi_i})} q(x_i - \beta_{i0} - \beta_i^T x_{\pi_i} | x_{\pi_i}) dx_i \\
 \text{(I.4)} \quad &= \int \log \frac{q_{\bar{\epsilon}_i}(x)}{p_{\epsilon_i}(x)} q_{\bar{\epsilon}_i}(x) dx.
 \end{aligned}$$

Based on the above computation, $\eta_i^{\pi_i}$ is independent of π_i and hence $\eta_i^{\pi_i} \equiv \eta_i$.

Appendix J. Model sensitivity indices for the ORR Bayesian network.

J.1. Calculation of model sensitivity indices. We recall that the optimal oxygen binding energy is defined as

$$x_{O^*}^P = \operatorname{argmax}_{x_0} [\min\{\mathbb{E}_P [y_1 | x_0], \mathbb{E}_P [y_2 | x_0]\}].$$

We compute $\mathbb{E}_P [y_i|x_0]$ for $i = 1, 2$ by using (7.3) and (7.6) as follows:

$$\begin{aligned} \mathbb{E}_P [y_i|x_0] &= \beta_{y_i,0} + \beta_{y_i,x}(x_0 + \beta_{s0,0} + \beta_{e0,0} + \beta_{d0,0}) \\ &\quad + \beta_{ci,0} + \beta_{si,0} + \beta_{ei,0} + \beta_{di,0}. \end{aligned}$$

Then

$$(J.1) \quad x_{O^*}^P = \frac{\beta_{y_2,0} + \bar{\beta}_2 - \beta_{y_1,0} - \bar{\beta}_1}{\beta_{y_1,x} - \beta_{y_2,x}},$$

where

$$(J.2) \quad \bar{\beta}_i = \beta_{y_i,x}(\beta_{s0,0} + \beta_{e0,0} + \beta_{d0,0}) + \beta_{ci,0} + \beta_{si,0} + \beta_{ei,0} + \beta_{di,0}.$$

It is a straightforward calculation that for $i = 1, 2$ and $l = e0, d0, s0, e1, d1, s1, c1, e2, d2, s2,$ and $c2,$

$$\beta_{y_i,0} + \bar{\beta}_i = \mathbb{E}_{P_{\omega_l}} [F_{l,i}] \quad \text{for any } i \text{ and } l,$$

where $F_{l,i} = \mathbb{E}_{P_{y_i|\omega_l}} [y_i|x_0]$ and $p(y_i|\omega_l, x_0) = \mathcal{N}(\tilde{\beta}_{y_i,0} + \tilde{\beta}_{y_i,\omega_l}\omega_l, \tilde{\sigma}_{y_i}^2)$. Hence (J.1) equals

$$x_{O^*}^P = \frac{\mathbb{E}_{P_{\omega_l}} [F_{l,2}] - \mathbb{E}_{P_{\omega_l}} [F_{l,1}]}{\beta_{y_1,x} - \beta_{y_2,x}}.$$

Note that since we compute the model sensitivity indices over $\mathcal{D}_{l,P}^n$ for any $l \in \{e0, d0, s0, e1, d1, s1, c1, e2, d2, s2\}$, the alternative Bayesian networks Q have the same structure as that given by (7.3), and the same CPDs as the Bayesian network P except from the CPD of ω_l . Let us denote its conditional distribution by Q_{ω_l} (since $\rho_{\omega_l} = \pi_{\omega_l} = \emptyset$) and its CPD by q_{ω_l} . Then

$$(J.3) \quad x_{O^*}^Q - x_{O^*}^P = \frac{\mathbb{E}_{Q_{\omega_l}} [F_{l,2}] - \mathbb{E}_{P_{\omega_l}} [F_{l,2}] - (\mathbb{E}_{Q_{\omega_l}} [F_{l,2}] - \mathbb{E}_{P_{\omega_l}} [F_{l,1}])}{\beta_{y_1,x} - \beta_{y_2,x}},$$

which further gives us

$$\frac{I^\mp(y_2, P; \mathcal{D}_{l,P}^n) - I^\pm(y_1, P; \mathcal{D}_{l,P}^n)}{\beta_{y_1,x} - \beta_{y_2,x}} \leq x_{O^*}^Q - x_{O^*}^P \leq \frac{I^\pm(y_2, P; \mathcal{D}_{l,P}^n) - I^\mp(y_1, P; \mathcal{D}_{l,P}^n)}{\beta_{y_1,x} - \beta_{y_2,x}}.$$

In the above inequality, by combining Theorem 3.6 and Table 2 we get (7.9) and (7.10).

Table 1

Outcomes of MLE for the parameters involved in (7.4)–(7.6) for the ORR Bayesian network in section 7.

$\beta_{y_1,0} = 0.0595$	$\beta_{e0,0}, \beta_{ei,0} = 0$
$\sigma_{e0}^2 = 0.0329$	$\sigma_{ei}^2 = 0.0065$
$\beta_{y_2,0} = 1.8231$	$\beta_{d0,0} = -0.0754$
$\beta_{di,0} = -0.0222$	$\sigma_{di}^2 = 0.0354$
$\beta_{y_1,x_0} = 0.5111$	$\sigma_{d0}^2 = 0.1032$
$\beta_{s1,0} = -0.2967$	$\sigma_{s1}^2 = 0.0046$
$\beta_{y_2,x_0} = -0.5564$	$\beta_{s0,0} = 0.0067$
$\beta_{s2,0} = -0.1209$	$\sigma_{s2}^2 = 0.0054$
$\beta_{ci,0} = 0$	$\sigma_{s0}^2 = 0.0010$
$\sigma_{c1}^2 = 0.0347$	$\sigma_{c2}^2 = 0.0204$

Table 2

The values of $\tilde{\beta}_{y_i, \omega_l}$ involved in $p(y_i | \omega_l, x_0) = \mathcal{N}(\tilde{\beta}_{y_i, 0} + \tilde{\beta}_{y_i, \omega_l} \omega_l, \tilde{\sigma}_{y_i}^2)$. They are used to evaluate the model sensitivity indices $I^\pm(y_i, P; \mathcal{D}_{i,P}^\eta)$, $i = 1, 2$, provided by Theorem 3.6.

	$\omega_l = \omega_{e0}, \omega_{d0}, \omega_{s0}$	$\omega_l = \omega_{e1}, \omega_{d1}, \omega_{s1}, \omega_{c1}$	$\omega_l = \omega_{e2}, \omega_{d2}, \omega_{s2}, \omega_{c2}$
$f = y_1$	$\tilde{\beta}_{y_1, \omega_l} = \beta_{y_1, x}$	$\tilde{\beta}_{y_1, \omega_l} = 1$	$\tilde{\beta}_{y_1, \omega_l} = 0$
$f = y_2$	$\tilde{\beta}_{y_2, \omega_l} = \beta_{y_2, x}$	$\tilde{\beta}_{y_2, \omega_l} = 0$	$\tilde{\beta}_{y_2, \omega_l} = 1$




Figure 16. Propagation vs. nonpropagation of model misspecification of the Bayesian network vertices ω_{d0} and ω_{e1} , respectively, to the predictions of the QoI x_{O*}^P ; model misspecification is set to $\eta = 1$ for both Bayesian network vertices. First, note that $I^+(y_2, P; \mathcal{D}_{\omega_{e1}}^\eta) = 0$, i.e., the model misspecification of ω_{e1} only affects the prediction of y_1 , but not y_2 (see Figure 8); therefore the uncertainty of ω_{e1} only propagates to x_{O*}^P through y_1 , while $I^+(y_1, P; \mathcal{D}_{\omega_{e1}}^\eta)$ is small since ω_{e1} has a lower variance which is associated with more informative available data. Thus, it results in a small corresponding uncertainty in x_{O*}^P . Meanwhile, the uncertainty of ω_{d0} propagates to x_{O*}^P through both y_1 and y_2 , (i.e., the model misspecification of ω_{d0} affects the predictions of both y_1 and y_2), and $I^+(y_i, P; \mathcal{D}_{\omega_{d0}}^\eta)$ is larger since ω_{d0} has a higher variance (due to insufficient informative data available). Therefore we have a larger corresponding uncertainty in x_{O*}^P predictions, as shown in the figure.

J.2. Propagation of model uncertainties to the QoIs. We note the discrepancies in the propagation of model misspecification to the QoI between different Bayesian network components, as demonstrated in Figure 9. In particular, in Figure 9 (left) the same uncertainty (described by model misspecification η_l) is applied on all ORR Bayesian network vertices. However, not all propagate and affect the same QoI; see Figure 16 for examples of propagation (22%) and nonpropagation (5% and 0%) of model misspecification to the QoI.

REFERENCES

- [1] A. ADADI AND M. BERRADA, *Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)*, IEEE Access, 6 (2018), pp. 52138–52160, <https://doi.org/10.1109/ACCESS.2018.2870052>.
- [2] O. ANTOINE, Y. BULTELE, AND R. DURAND, *Oxygen reduction reaction kinetics and mechanism on platinum nanoparticles inside Nafion*, J. Electroanal. Chem., 499 (2001), pp. 85–94.
- [3] D. ARONSKY AND P. J. HAUG, *Diagnosing community-acquired pneumonia with a Bayesian network*, in Proceedings of the AMIA Symposium, American Medical Informatics Association, 1998, pp. 632–636.
- [4] M. I. BELGHAZI, A. BARATIN, S. RAJESHWAR, S. OZAIR, Y. BENGIO, A. COURVILLE, AND D. HJELM, *Mutual information neural estimation*, in Proceedings of the 35th International Conference on Machine Learning, PMLR, 2018, pp. 531–540, <http://proceedings.mlr.press/v80/belghazi18a.html>.
- [5] P. BIRMPA AND M. A. KATSOULAKIS, *Uncertainty quantification for Markov random fields*, SIAM/ASA J. Uncertain. Quantif., 9 (2021), pp. 1457–1498, <https://doi.org/10.1137/20M1374614>.
- [6] J. BIRRELL, P. DUPUIS, M. A. KATSOULAKIS, Y. PANTAZIS, AND L. REY-BELLET, *(f, Γ) -Divergences: Interpolating between f -divergences and integral probability metrics*, J. Mach. Learn. Res., 23 (2022), pp. 1–70.
- [7] J. BIRRELL, P. DUPUIS, M. A. KATSOULAKIS, L. REY-BELLET, AND J. WANG, *Variational Representations and Neural Network Estimation of Rényi Divergences*, preprint, <https://arxiv.org/abs/2007.03814>, 2021.
- [8] J. BIRRELL, M. KATSOULAKIS, AND L. REY-BELLET, *Robustness of Dynamical Quantities of Interest via Goal-Oriented Information Theory*, preprints, <https://arxiv.org/abs/1906.09282>, 2019.
- [9] J. BIRRELL AND L. REY-BELLET, *Uncertainty quantification for Markov processes via variational principles and functional inequalities*, SIAM/ASA J. Uncertain. Quantif., 8 (2020), pp. 539–572, <https://doi.org/10.1137/19M1237429>.
- [10] C. M. BISHOP, *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag, New York, 2006.
- [11] J. BLANCHET, H. LAM, Q. TANG, AND Z. YUAN, *Robust actuarial risk analysis*, North Amer. Actuar. J., 23 (2019), pp. 33–63, <https://doi.org/10.1080/10920277.2018.1504686>.
- [12] J. BLANCHET AND K. MURTHY, *Quantifying distributional model risk via optimal transport*, Math. Oper. Res., 44 (2019), pp. 565–600, <https://doi.org/10.1287/moor.2018.0936>.
- [13] C. BURGESS, T. HART, Z. YANG, S. CUCERZAN, R. W. WHITE, A. PASTUSIAK, AND J. LEWIS, *A Base Camp for Scaling AI*, preprint, <https://arxiv.org/abs/1612.07896>, 2016.
- [14] F. CALLE-VALLEJO, J. TYMOCZKO, V. COLIC, Q. H. VU, M. D. POHL, K. MORGENSTERN, D. LOFFREDA, P. SAUTET, W. SCHUHMAN, AND A. S. BANDARENKA, *Finding optimal surface sites on heterogeneous catalysts by counting nearest neighbors*, Science, 350 (2015), pp. 185–189.
- [15] K. CHOWDHARY AND P. DUPUIS, *Distinguishing and integrating aleatoric and epistemic variation in uncertainty quantification*, ESAIM Math. Model. Numer. Anal., 47 (2013), pp. 635–662.
- [16] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, John Wiley & Sons, 2012.
- [17] E. DELAGE AND Y. YE, *Distributionally robust optimization under moment uncertainty with application to data-driven problems*, Oper. Res., 58 (2010), pp. 595–612, <https://doi.org/10.1287/opre.1090.0741>.
- [18] M. DING, C. DASKALAKIS, AND S. FEIZI, *GANs with conditional independence graphs: On subadditivity of probability divergences*, in Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR, 2021.
- [19] F. DOSHI-VELEZ AND B. KIM, *Towards a Rigorous Science of Interpretable Machine Learning*, preprint, <https://arxiv.org/abs/1702.08608>, 2017.
- [20] P. DUPUIS AND R. S. ELLIS, *A Weak Convergence Approach to the Theory of Large Deviations*, John Wiley & Sons, 2011.
- [21] P. DUPUIS, M. A. KATSOULAKIS, Y. PANTAZIS, AND P. PLECHÁČ, *Path-space information bounds for uncertainty quantification and sensitivity analysis of stochastic dynamics*, SIAM/ASA J. Uncertain. Quantif., 4 (2016), pp. 80–111, <https://doi.org/10.1137/15M1025645>.
- [22] P. DUPUIS, M. A. KATSOULAKIS, Y. PANTAZIS, AND L. REY-BELLET, *Sensitivity Analysis for Rare Events Based on Rényi Divergence*, preprint, <https://arxiv.org/abs/1805.06917>, 2018.
- [23] P. DUPUIS AND Y. MAO, *Formulation and Properties of a Divergence Used to Compare Probability Measures without Absolute Continuity*, preprint, <https://arxiv.org/abs/1911.07422>, 2019.

- [24] J. FENG, J. LANSFORD, A. MIRONENKO, D. B. POURKARGAR, D. G. VLACHOS, AND M. A. KATSOULAKIS, *Non-parametric correlative uncertainty quantification and sensitivity analysis: Application to a Langmuir bimolecular adsorption model*, AIP Adv., 8 (2018), 035021.
- [25] J. FENG, J. L. LANSFORD, M. A. KATSOULAKIS, AND D. G. VLACHOS, *Explainable and trustworthy artificial intelligence for correctable modeling in chemical sciences*, Sci. Adv., 6 (2020), <https://doi.org/10.1126/sciadv.abc3204>.
- [26] N. FRIEDMAN, M. LINIAL, I. NACHMAN, AND D. PE'ER, *Using Bayesian networks to analyze expression data*, J. Comput. Biol., 7 (2000), pp. 601–620.
- [27] R. FUNG AND B. DEL FAVERO, *Applying Bayesian networks to information retrieval*, Commun. ACM, 38 (1995), pp. 42–ff.
- [28] R. GAO AND A. J. KLEYWEGT, *Distributionally Robust Stochastic Optimization with Wasserstein Distance*, preprint, <https://arxiv.org/abs/1604.02199>, 2016
- [29] A. GENEVAY, G. PEYRÉ, AND M. CUTURI, *GAN and VAE from an Optimal Transport Point of View*, preprint, <https://arxiv.org/abs/1706.01807>, 2017
- [30] A. GENEVAY, G. PEYRÉ, AND M. CUTURI. *Learning generative models with Sinkhorn divergences*, in International Conference on Artificial Intelligence and Statistics, PMLR, 2018, pp. 1608–1617.
- [31] Z. GHAHRAMANI, *Probabilistic machine learning and artificial intelligence*, Nature, 521 (2015), pp. 452–459, <https://doi.org/10.1038/nature14541>.
- [32] P. GLASER, M. ARBEL, AND A. GRETTON, *KALE Flow: A Relaxed KL Gradient Flow for Probabilities with Disjoint Support*, preprint, <https://arxiv.org/abs/2106.08929>, 2019
- [33] J. GOH AND M. SIM, *Distributionally robust optimization and its tractable approximations*, Oper. Res., 58 (2010), pp. 902–917, <https://doi.org/10.1287/opre.1090.0795>.
- [34] M. Á. GÓMEZ-VILLEGAS, P. MAIN, H. NAVARRO, AND R. SUSI, *Evaluating the difference between graph structures in Gaussian Bayesian networks*, Expert Syst. Appl., 38 (2011), pp. 12409–12414.
- [35] M. A. GÓMEZ-VILLEGAS, P. MAÍN, AND R. SUSI, *Sensitivity analysis in Gaussian Bayesian networks using a divergence measure*, Commun. Stat. Theory Methods, 36 (2007), pp. 523–539.
- [36] M. A. GÓMEZ-VILLEGAS, P. MAIN, AND R. SUSI, *The effect of block parameter perturbations in Gaussian Bayesian networks: Sensitivity and robustness*, Inform. Sci., 222 (2013), pp. 439–458.
- [37] K. GOURGOULIAS, M. KATSOULAKIS, L. REY-BELLET, AND J. WANG, *How biased is your model? Concentration inequalities, information and model bias*, IEEE Trans. Inform. Theory, 66 (2020), pp. 3079–3097.
- [38] D. GUNNING AND D. AHA, *Darpa's explainable artificial intelligence (XAI) program*, AI Mag., 40 (2019), pp. 44–58, <https://doi.org/10.1609/aimag.v40i2.2850>.
- [39] E. J. HALL AND M. A. KATSOULAKIS, *Robust information divergences for model-form uncertainty arising from sparse data in random PDE*, SIAM/ASA J. Uncertain. Quantif., 6 (2018), pp. 1364–1394, <https://doi.org/10.1137/17M1143344>.
- [40] E. J. HALL, S. TAVERNIERS, M. A. KATSOULAKIS, AND D. M. TARTAKOVSKY, *Ginns: Graph-informed neural networks for multiscale physics*, J. Comput. Phys., 433 (2021), 110192, <https://doi.org/10.1016/j.jcp.2021.110192>.
- [41] L. P. HANSEN AND T. J. SARGENT, *Robustness*, Princeton University Press, 2008, <http://www.jstor.org/stable/j.ctt1dr35gx>.
- [42] D. HECKERMAN, E. HORVITZ, AND B. N. NATHWANI, *Toward normative expert systems, Part I, Methods Inf. Med.*, 31 (2016), pp. 90–105.
- [43] A. HOLEWINSKI AND S. LINIC, *Elementary mechanisms in electrocatalysis: Revisiting the ORR Tafel slope*, J. Electrochem. Soc., 159 (2012), H864.
- [44] R. JIANG AND Y. GUAN, *Data-driven chance constrained stochastic program*, Math. Program., 158 (2016), pp. 291–327, <https://doi.org/10.1007/s10107-015-0929-7>.
- [45] K. KANDASAMY, A. KRISHNAMURTHY, B. PÓCZOS, L. WASSERMAN, AND J. M. ROBINS, *Nonparametric von Mises estimators for entropies, divergences and mutual informations*, in Proceedings of the 28th International Conference on Neural Information Processing Systems, MIT Press, Cambridge, MA, 2015, pp. 397–405.
- [46] M. G. KAPTEYN, J. V. R. PRETORIUS, AND K. E. WILLCOX, *A Probabilistic Graphical Model Foundation for Enabling Predictive Digital Twins at Scale*, preprint, <https://arxiv.org/abs/2012.05841>, 2020.

- [47] M. A. KATSOUKAKIS, L. REY-BELLET, AND J. WANG, *Scalable information inequalities for uncertainty quantification*, J. Comput. Phys., 336 (2017), pp. 513–545.
- [48] D. KOLLER AND N. FRIEDMAN, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, 2009.
- [49] H. LAM, *Recovering Best Statistical Guarantees via the Empirical Divergence-Based Distributionally Robust Optimization*, preprint, <https://arxiv.org/abs/1605.09349>, 2016.
- [50] E. LAZKANO, B. SIERRA, A. ASTIGARRAGA, AND J. M. MARTINEZ-OTZETA, *On the use of Bayesian networks to develop behaviours for mobile robots*, Rob. Auton. Syst., 55 (2007), pp. 253–265.
- [51] T. S. LEVITT, J. M. AGOSTA, AND T. O. BINFORD, *Model-based influence diagrams for machine vision*, Mach. Intell. Pattern Recogn., 10 (1990), pp. 371–388, Elsevier.
- [52] T. MILLER, *Explanation in artificial intelligence: Insights from the social sciences*, Artif. Intell., 267 (2019), pp. 1–38, <https://doi.org/10.1016/j.artint.2018.07.007>.
- [53] P. MOHAJERIN ESFAHANI AND D. KUHN, *Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations*, Math. Program., 171 (2018), pp. 115–166, <https://doi.org/10.1007/s10107-017-1172-1>.
- [54] K. R. MOON, K. SRICHARAN, AND A. O. HERO, *Ensemble estimation of mutual information*, in IEEE International Symposium on Information Theory (ISIT), 2017, pp. 3030–3034, <https://doi.org/10.1109/ISIT.2017.8007086>.
- [55] A. MÜLLER, *Integral probability metrics and their generating classes of functions*, Adv. Appl. Probab., 29 (1997), pp. 429–443, <http://www.jstor.org/stable/1428011>.
- [56] X. NGUYEN, M. J. WAINWRIGHT, AND M. I. JORDAN, *Estimating divergence functionals and the likelihood ratio by convex risk minimization*, IEEE Trans. Inform. Theory, 56 (2010), pp. 5847–5861.
- [57] H. OWHADI, C. SCOVEL, T. J. SULLIVAN, M. MCKERNS, AND M. ORTIZ, *Optimal uncertainty quantification*, SIAM Rev., 55 (2013), pp. 271–345, <https://doi.org/10.1137/10080782X>.
- [58] J. PEARL, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Reasoning*, Morgan Kaufmann, Los Altos, CA, 1988.
- [59] J. PEARL, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Elsevier, 2014.
- [60] F. PÉREZ-CRUZ, *Estimation of information theoretic measures for continuous random variables*, in Advances in Neural Information Processing Systems, D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, eds., Curran Associates, 2009, <https://proceedings.neurips.cc/paper/2008/file/ccb0989662211f61edae2e26d58ea92f-Paper.pdf>.
- [61] A. RASHEED, O. SAN, AND T. KVAMSDAL, *Digital twin: Values, challenges and enablers from a modeling perspective*, IEEE Access, 8 (2020), pp. 21980–22012, <https://doi.org/10.1109/ACCESS.2020.2970143>.
- [62] G. ROTHENBERG, *Catalysis*, John Wiley & Sons, 2008, <https://doi.org/10.1002/9783527621866.ch2>.
- [63] B. P. SETZLER, Z. ZHUANG, J. A. WITTKOPF, AND Y. YAN, *Activity targets for nanostructured platinum-group-metal-free catalysts in hydroxide exchange membrane fuel cells*, Nat. Nanotechnol., 11 (2016), pp. 1020–1025.
- [64] R. D. SHACHTER AND C. R. KENLEY, *Gaussian influence diagrams*, Manag. Sci., 35 (1989), pp. 527–550, <http://www.jstor.org/stable/2632102>.
- [65] R. C. SMITH, *Uncertainty Quantification: Theory, Implementation, and Applications*, SIAM, Philadelphia, 2013.
- [66] P. SPIRITES, C. N. GLYMOUR, R. SCHEINES, D. HECKERMAN, C. MEEK, G. COOPER, AND T. RICHARDSON, *Causation, Prediction, and Search*, MIT Press, 2000.
- [67] N.-T. SUEN, S.-F. HUNG, Q. QUAN, N. ZHANG, Y.-J. XU, AND H. M. CHEN, *Electrocatalysis for the oxygen evolution reaction: Recent development and future perspectives*, Chem. Soc. Rev., 46 (2017), pp. 337–365.
- [68] J. E. SUTTON, W. GUO, M. A. KATSOUKAKIS, AND D. G. VLACHOS, *Effects of correlated parameters and uncertainty in electronic-structure-based chemical kinetic modelling*, Nat. Chem., 8 (2016), pp. 331–337, <https://doi.org/10.1038/NCHEM.2454>.
- [69] K. UM, E. J. HALL, M. A. KATSOUKAKIS, AND D. M. TARTAKOVSKY, *Causality and bayesian network pdes for multiscale representations of porous media*, J. Comput. Phys., 394 (2019), pp. 658–678, <https://doi.org/10.1016/j.jcp.2019.06.007>.

- [70] Q. WANG, S. R. KULKARNI, AND S. VERDU, *A nearest-neighbor approach to estimating divergence between continuous random vectors*, in 2006 IEEE International Symposium on Information Theory, 2006, pp. 242–246, <https://doi.org/10.1109/ISIT.2006.261842>.
- [71] Q. WANG, S. R. KULKARNI, AND S. VERDU, *Divergence estimation for multidimensional densities via k -nearest-neighbor distances*, IEEE Trans. Inform. Theory, 55 (2009), pp. 2392–2405, <https://doi.org/10.1109/TIT.2009.2016060>.
- [72] L. WASSERMAN, *All of Nonparametric Statistics*, Springer, New York, 2006.
- [73] L. WASSERMAN, *All of Statistics: A Concise Course in Statistical Inference*, Springer Science & Business Media, 2013.
- [74] W. WIESEMANN, D. KUHN, AND M. SIM, *Distributionally robust convex optimization*, Oper. Res., 62 (2014), pp. 1358–1376, <https://doi.org/10.1287/opre.2014.1314>.
- [75] W. XIE AND S. AHMED, *On deterministic reformulations of distributionally robust joint chance constrained optimization problems*, SIAM J. Optim., 28 (2018), pp. 1151–1182, <https://doi.org/10.1137/16M1094725>.