



Prediction can be safely used as a proxy for explanation in causally consistent Bayesian generalized linear models

Maximilian Scholz & Paul-Christian Bürkner

To cite this article: Maximilian Scholz & Paul-Christian Bürkner (21 Jan 2025): Prediction can be safely used as a proxy for explanation in causally consistent Bayesian generalized linear models, Journal of Statistical Computation and Simulation, DOI: [10.1080/00949655.2024.2449534](https://doi.org/10.1080/00949655.2024.2449534)

To link to this article: <https://doi.org/10.1080/00949655.2024.2449534>

 View supplementary material 

 Published online: 21 Jan 2025.

 Submit your article to this journal 

 View related articles 

 View Crossmark data 



Prediction can be safely used as a proxy for explanation in causally consistent Bayesian generalized linear models

Maximilian Scholz  and Paul-Christian Bürkner 

Cluster of Excellence SimTech, University of Stuttgart, Stuttgart, Germany

ABSTRACT

Bayesian modeling provides a principled approach to quantifying uncertainty and has seen a surge of applications in recent years. Within the context of a Bayesian workflow, we are concerned with model selection for the purpose of finding models that best explain the data or underlying data generating process. Since insight into the true process is rare, what remains is incomplete causal knowledge and model predictions of the data. This leads to the important question of when the use of prediction as a proxy for explanation for the purpose of model selection is valid. We approach this question by means of large-scale simulations of Bayesian generalized linear models where we investigate various causal and statistical misspecifications. Our results indicate that the use of prediction as proxy for explanation is valid and safe if the models under consideration are sufficiently consistent with the underlying causal structure of the true data generating process.

ARTICLE HISTORY

Received 19 July 2023

Accepted 17 December 2024


KEYWORDS


Bayesian workflow; causal inference; explanation; prediction; generalized linear models; simulation study

1. Introduction

Probabilistic modeling provides a principled approach to quantifying uncertainty in model parameters and model structure. In recent years, it has seen a surge of applications in almost all quantitative sciences and industrial areas [1–3]. To support the principled application of Bayesian methods, [3] proposed an overarching workflow to conduct Bayesian data analysis. In short, the workflow asks users to pick an initial model and iteratively refine it, performing various checks on the way to ensure that probabilistic assumptions are sensible, computations are valid, and model results are trustworthy for the intended purposes. This basic model building loop is repeated until either the user's requirements are satisfied or no satisfactory model can be found with the available resources. Within this overarching workflow, there are still many unknowns when it comes to making optimal decisions for each iterative step or sub-workflow.

The current work is concerned with the decision-making process during model building iterations, namely model selection for the purpose of finding models that best *explain*

CONTACT Maximilian Scholz  research.scholz@mailbox.org  University of Stuttgart – Cluster of Excellence SimTech, Universitätsstraße 32, 70569 Stuttgart, Deutschland

 Supplemental data for this article can be accessed online at <http://dx.doi.org/10.1080/00949655.2024.2449534>.

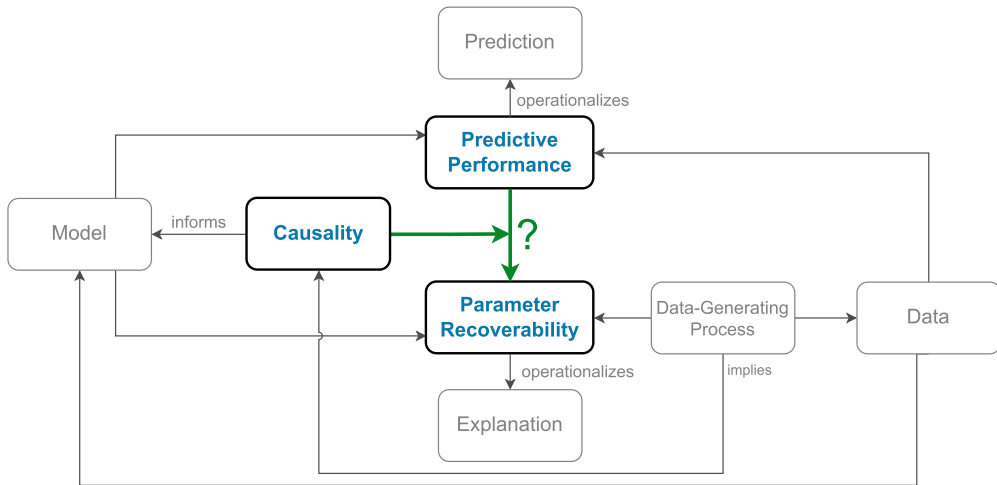


Figure 1. High-level overview of our conceptual framework that connects explanation, prediction and causality in statistical modeling. Parameter recovery compares model parameters and the data-generating process to operationalize explanation. Predictive performance compares model predictions and observed data to operationalize prediction. Causality is implied by the data-generating process and can inform model building. The three-way relationship between predictive performance, parameter recovery, and causality is unclear and subject of this paper. See Section 1.1 for details.

the data, that is, help us understand the underlying data generating process. Unfortunately, explanation is an illusive concept in practice as, for real data, we have no direct access to the underlying data generating process. Thus, all we are left with during real-world analyses is incomplete *causal* knowledge from sources outside of the current data and model *predictions* of said data. The distinction between explanation (of the underlying process) and prediction (of observable data) has a long history in many fields of sciences engaged in model-based reasoning and inference (see Section 1.1). However, the statistical relationship of explanation and prediction is still not sufficiently understood, which is surprising given the importance of these concepts to statistical theory and practice.

To this end, we study the validity of using prediction as a proxy for explanation in Bayesian statistical models (see Figure 1 for a conceptual overview to be detailed in Section 1.1). Our main conclusion can be summarized as follows: *Using prediction as a proxy for explanation is valid and safe only when the considered models are sufficiently consistent with the underlying causal structure of the true data generating process.* Specifically, our contributions are:

- (i) a conceptual introduction and overview of the relationship of explanation and prediction as well as their connection to causality;
- (ii) large-scale simulations of Bayesian generalized-linear models to study said relationship under various causal and statistical misspecifications;
- (iii) initial evidence that causality is indeed the missing link that connects prediction and explanation when comparing statistical models; and
- (iv) a set of R packages [4–6] that facilitate simulation studies of Bayesian models.

1.1. Explanation vs. prediction

Douglas [7] provides a historical overview of the relationship of prediction and explanation from the perspective of philosophy of science. She concludes that ‘explanations are the means that help us think our way through to the next testable prediction’ [7, p. 462, Section 6]. In other words, and more statistically speaking, explanations help us understand the inner workings of a process by means of a model, while prediction concerns the comparison of the model’s output to real-world observables. Explanation and prediction are not statistical concepts themselves, though. Instead, we have to operationalize them in terms of mathematical quantities computable from a statistical model. Following [8], we call the statistical operationalization of explanation *parameter recoverability* and the operationalization of prediction *predictive performance*:

Parameter Recoverability (PR) refers to a model’s ability to recover the latent (i.e. not directly observable) parameters of an assumed true data-generating process (DGP), for example, the true effect size of a treatment for a medical condition [8]. This implies that assessing PR requires concrete knowledge of the true DGP including its true parameter values. Accordingly, it can only be studied directly in artificial settings where the true DGP is known, with the hope that the real data of interest fulfills the assumptions of these artificial settings sufficiently well.

Predictive Performance (PP) describes the ability of a model to accurately predict new observations from existing observations, for example, predicting a diagnosis based on symptoms and medical history [9,10]. It is one of the most prominent utilities for gauging model performance in data analysis and a central tool for model comparisons [1,11,12]. In most cases, one is interested in out-of-sample PP, as predicting unseen data is almost always the primary goal of modeling [11]. PP is conceptually similar to PR in that both target the accurate estimation of model-implied quantities [8], with the main difference that the former targets quantities that are observable at real data inference time (i.e. outcome variables), which allows to derive estimates that are agnostic with respect to the true DGP [11].

In the following, we present some of the common perspectives on prediction and explanation as well as their connections with causality. When talking about these concepts from statistical perspectives, we use them and their operationalizations interchangeably.

1.1.1. Explanation only

The statistical field that is probably most interested in explanation in the form of unbiased estimators is that of causality [13]. The *do*-calculus [14] and equivalently the potential outcomes framework [15,16] offer sound theories to identify models that are able to provide an unbiased estimate of a parameter of choice. We call a model ‘causally consistent’ if it is consistent with a theoretically justified causal graph of its contributing variables [8]. This consistency is necessary for models to provide at least asymptotically unbiased estimates for latent parameters.

While the ability to find an unbiased estimator is highly valuable, it is not the only metric we need to pay attention to. As argued for example by Shmueli [17], optimizing the overall bias-variance trade-off is a sensible alternative goal, which also requires to minimize the variance of a parameter estimator in addition to minimizing its bias. Examples in the causal literature that attempt to additionally reduce variance are the works of [18,19],

which use graphical criteria to find lower variance estimators [see also 20, for a more general overview]. However, we found that the reduction of variance is only considered after unbiasedness has been established. Thus, no actual bias-variance trade-off is allowed to take place from this perspective. Additionally, predictive performance considerations are missing altogether from the general discussion in the causal literature. We go into more details about how causal consistency serves as the foundation of the presented study in Section 2.1.

1.1.2. *Explanation first*

Arguably the most prominent perspective on the relationship of explanation and prediction is coming from the philosophy of science [as exemplarily presented by 17,21–25] in cognitive-, social-, and life-sciences. The main discussion point is that, even if explanation is the main research goal, the value of good prediction shouldn't simply be ignored. Yarkoni and Westfall [24] and Breiman [25] are examples for the call to put more focus on predictions, even if explanation is the main goal. The underlying argument is that a model with bad predictive performance has a higher chance of misrepresenting or missing relevant features of the true DGP. One should thus be careful when relying on the explanations of a model with weak predictive performance. Similarly, [17] makes the argument that predictive power cannot be inferred from explanatory power, that is, good explanatory models do not automatically lead to good predictions. A common proposal is to use a predictive model (i.e. a model that maximizes predictive performance) as a benchmark [17,26] to aim for with an explanatory model (i.e. a model that maximizes parameter recoverability). Douglas [7] gives an illustrative example for this approach with the discovery of Pluto. The paths of Uranus and Neptune did not follow their predicted paths which implied a problem with the explanatory model. The solution was the discovery of Pluto, which explained the observed deviation and lead to an updated model. Finally, while there are sound high level arguments for why models that provide better explanations should generally provide better predictions as well, we are not aware of other works that investigate this relationship further besides common textbook examples of improved predictions despite bad parameter recovery in the face of causal misspecification [2].

1.1.3. *Prediction first*

Machine-learning models are sometimes also called black box models due to the fact that they solely focus on providing the best possible predictions while their inner workings are hard or impossible to understand and do not directly map to any assumed DGP. The lack of interpretability, and thus potential to provide explanations, is a common critique of machine learning techniques, more recently more work has been done to improve this aspect [see for example 27,28]. While predictive performance remains the main goal of these models, requirements regarding responsible reporting and accountability as well as improved performance for small sample sizes led to the incorporation of causal assumptions, which are a requirement for reliable explanations [29–31]. And while not necessarily using the same causal vocabulary, the idea of physics-informed neural networks [32,33] is closely related, as the partial or ordinary differential equations describing physical processes can be understood as causal constraints that have to be met.

However, with the increasing requirements for explainable AI [34] in fields like medicine, aspects such as feature importance [35,36] become more relevant to understand

how machine learning models make decisions. Such aspects can be seen as extending the concept of explanation in some sense, although this has not been explored systematically to our knowledge and, at least in the stricter sense of recovering (latent) parameters of an assumed DGP, explanation is not really a concern in machine learning in general.

1.1.4. Combining explanation, prediction, and causality

Based on the presented perspectives, there is an apparent gap in the joint consideration of explanation, prediction, and causality. We illustrate this gap in Figure 1, where we provide an overview of these concepts' relationships in a statistical modeling workflow. In this paper, we approach their joint consideration specifically from the perspective of Bayesian model selection.

When explanation is the goal, prediction reduces to a conveniently available supporting utility that *ideally* helps to select better explaining models at real data inference time [8]. In practice at least, and despite the theoretical arguments for caution, this assumption is very commonly (and often implicitly) made whenever explanatory model choices are based on out-of-sample posterior predictive metrics or their approximations, such as AIC [e.g., 2,37], DIC [38], WAIC [12,39] or ELPD-LOO [11,12]. As we know from counterexamples [40,41], this assumption cannot hold in general, but it remains unclear under which conditions it is actually justified. The goal of this paper is to investigate under which circumstances this proxy use might be valid.

1.2. Generalized linear models

We focus our studies on the class of Bayesian generalized linear models (GLMs). This choice is motivated by the fact that they represent the minimal general class of models that enable us to investigate the utility of using prediction as a proxy for explanation. GLMs allow us to represent causal DAGs while adding the ability to change model aspects that are not DAG-dependent (i.e. the likelihood family and link function). Despite (or perhaps because of) their simplicity, GLMs make up a big part of all statistical data analyses. Their success can be explained by several factors, involving the ease of interpretation of their additive structure, rich mathematical theory, and (relatively) straightforward estimation [42–45].

More specifically, we consider GLMs of the form $y \sim \text{likelihood}(\text{link}^{-1}(\mu), \phi)$ with a univariate response variable y that is assumed to follow a parametric likelihood distribution, often called a likelihood family [46,47], one main location parameter $\mu = \alpha + X\beta$ that is predicted, as well as zero or more auxiliary distributional parameters ϕ that are assumed to be constant over observations [44,45].

When setting up GLMs, the four important choices the analyst has to make are (i) the likelihood family, (ii) the link function, (iii) the linear predictor term, and (iv) whether and how to regularize, that is, for Bayesian GLMs, what prior distributions to use. All of these choices are mutually related [48], but specifically (i) and (ii) are closely intertwined, as the choice of link function depends on the support of μ and thus on the chosen likelihood. We discuss our specific choices of likelihoods and link functions in Section 2.2, the specific models we generate data from in Section 2.3 and the models we fit on the generated data in Section 2.4.

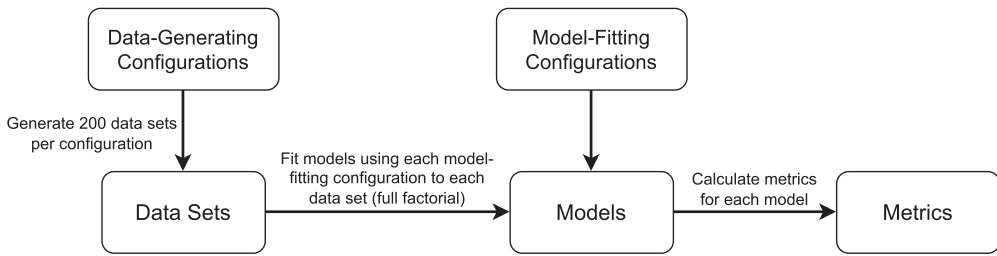


Figure 2. Conceptual simulation architecture. We generated 200 data sets for each entry from a list of data-generating configurations (see Table 1) for a total of 28,800 datasets. We then fit a model for each entry from a list of model-fitting configurations (see Table 2) to each data set for a total of 1,728,000 models. Finally, we calculated metrics for all fitted models.

1.3. Aims and scope

As laid out in Section 1.1, there is a gap in the literature regarding the relationship of (the statistical operationalizations of) prediction and explanation under varying causal assumptions. The aim of this paper is not to fill said gap entirely but rather to be a starting point for further research. As explained above, direct measures of explanation are rarely accessible during real-world modeling and prediction provides a convenient (i.e. easily available) alternative for model selection. Even though the validity of the practice has so far remained unclear, prediction is commonly used as a proxy for explanation during practical model selection in the context of Bayesian or other statistical workflows. It is the main aim of this paper to empirically investigate the use of prediction as a proxy for explanation under several causal and statistical misspecification mechanisms in Bayesian GLMs. Further, the methods introduced in this paper could serve as a basis for further work, some of which we propose in Section 4.

While this paper focuses on Bayesian models, we assume that the presented results also hold for their frequentist counterparts, as the mechanisms of causality are agnostic to the specific method of estimation (under reasonable equivalence of estimators; see also Section 2.4).

2. Methods

As explained in Section 1.1, studying parameter recoverability (PR; the statistical operationalization of explanation) requires knowledge of the true data-generating process (DGP). Combined with the fact that we aim to study Bayesian GLMs, this renders an analytical approach infeasible and we thus resort to extensive simulations as the method of choice.

In this section, we start with an explanation of the causal foundation of the simulations and how we used it throughout the process in Section 2.1. We also provide an overview of the considered likelihood families and link functions in Section 2.2. Next, we walk through the study's design (see Figure 2 for a conceptual overview), with data generation described in Section 2.3, model fitting in Section 2.4, and metric calculation in Section 2.5. Finally, we discuss the statistical analysis of the results of the simulations in Section 2.6.

Table 1. Overview of data generation configurations.

Factor	Levels
Double-bounded likelihoods	beta, Kumaraswamy, simplex, transformed-normal
Double-bounded links	logit, cauchit, cloglog
Double-bounded shapes	symmetric, asymmetric, bathtub
Lower-bounded likelihoods	gamma, Weibull, transformed-normal, Fréchet, beta-prime, Gompertz
Lower-bounded links	log, softplus
Lower-bounded shapes	ramp, heavy tail, thin tail
True β_{xy}	zero, positive

Table 2. Overview of model fit configurations.

Factor	Levels
Double-bounded Likelihoods	beta, Kumaraswamy, simplex, transformed-normal
Double-bounded Links	logit, cauchit, cloglog
Single-bounded Likelihoods	gamma, Weibull, transformed-normal, Fréchet, beta-prime, Gompertz
Single-bounded Links	log, softplus
Formulas (right-hand side)	$x + z_1 + z_2$, $x + z_2$, $x + z_1$, $x + z_1 + z_2 + z_3$, $x + z_1 + z_2 + z_4$

Many aspects of this study are not only dependent on the methods and algorithms used, but also on their software implementations. The simulation was implemented in R [49] using Stan [50,51] and brms [47]. Our software packages bayesim [4], bayesfam [5], and bayeshear [6] are available online. A more thorough discussion of the included likelihoods and links as well as additional results are available in the supplemental material. Said discussion as well as the code for the simulation configurations, the simulation data and corresponding analyses are available in our online Appendix [52].

2.1. Causal foundation

In practice, we usually cannot know if a model actually represents the true DGP, not only in terms of the distributional assumptions of likelihood and link, but also in terms of the structural and causal assumptions of the included (causal) covariates and their (non-)linear combination on the latent space. This uncertainty about covariates and their causal effects is a focus point for the causal literature. As explained in Section 1.1.1, causal consistency is a necessity for good parameter recoverability and thus a fundamental aspect of this study’s methods.

Here, we use a single, prototypical causal directed acyclic graph [DAG; 13] to guide both data generation and causal model assumptions (see Figure 3). The full DAG displayed on the left-hand side of Figure 3 consists of an outcome y , a treatment x , and four additional variables z_1, z_2, z_3, z_4 that correspond to four common, qualitatively different types of controls [20]. With respect to the effect of x on y , z_1 is a fork, z_2 is an ancestor of y , z_3 is an ancestor of x , and z_4 is a collider. These archetypes represent the majority of controls occurring in reality [13,20]. The only common control missing from this list is the pipe ($x \rightarrow z \rightarrow y$), which was not included here due to its similarity with the fork.

The goal of our statistical models will be to estimate the causal effect β_{xy} of x on y as accurately as possible. For this purpose, z_1 constitutes a ‘good control’ in the sense that controlling for it decreases bias in the estimation, while z_4 is a ‘bad control’, as controlling for it increases bias. In contrast, z_2 and z_3 are ‘neutral controls’: controlling for them

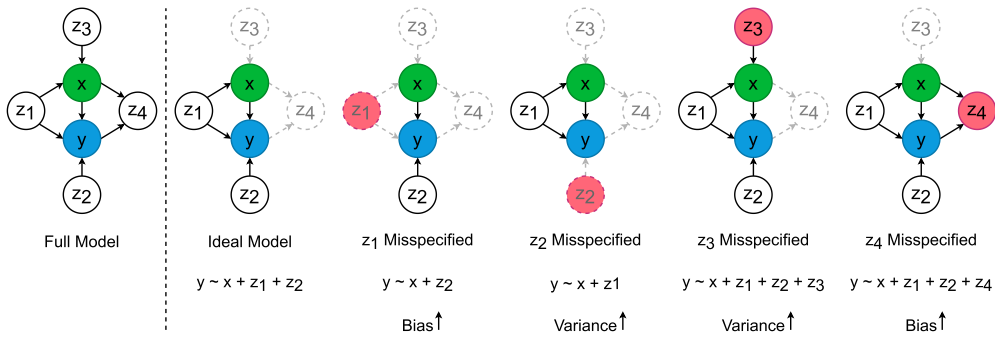


Figure 3. Full data-generating DAG and the resulting models. The ideal model is the subset of the full model that optimally estimates β_{xy} . Misspecifying with respect to each of the z_n variables leads to an additional model, where excluding z_1 or including z_4 increases bias, while excluding z_2 or including z_3 increases variance of the estimation.

does not influence bias. However, controlling for z_2 increases precision (decreases variance) while controlling for z_3 decreases precision of the estimation of β_{xy} . As a result, to estimate β_{xy} , the *ideal (causal) model* only controls for z_1 and z_2 (see the second graph in Figure 3). From there, we can obtain *misspecified (causal) models* by including or excluding controls. Based on four controls, $2^4 - 1 = 15$ misspecified models are conceivable, but for simplicity, we focus only on a subset of four of them, each of which deviates from the ideal model in exactly one control (see the third to sixth graph in Figure 3). Two of these misspecified models (i.e. when incorrectly excluding z_2 or incorrectly including z_3) still yield (asymptotically) unbiased estimates of β_{xy} . We will call them and the ideal causal model ‘causally unbiased’, while we call (asymptotically) biased models ‘causally biased’ [8]. Using this terminology, we can formulate our primary research questions more precisely as:

- (i) Within a set of statistical models that all share the same underlying causal model, can prediction be reliably used as a proxy for explanation?
- (ii) Does the answer to (i) depend on whether or not the causal model is biased?

Focusing on these questions implies that we do not aim to compare models with differing linear predictor terms (i.e. causal assumptions as per Figure 3). For one, there already are well known examples of cross-formula comparisons where the alignment of prediction and explanation doesn’t hold [e.g., adding a collider improves prediction but worsens explanation compared to an ideal unbiased model; 2]. Still, cross-formula comparison includes a lot of uncharted territory, which we think is worth studying but out of scope of the present paper. Of course, as a result of this choice, we need other aspects to vary among the compared statistical models. In this study, these aspects will be the likelihood and link functions of Bayesian GLMs as detailed next.

2.2. Likelihoods and link functions

The range of practically relevant likelihood classes is extensive and encompasses, among others, likelihoods for unbounded, lower-bounded, and double-bounded continuous data,

as well as binary, categorical, ordinal, count, and proportional (sum-to-one) data [53–57]. Studying all of them at once would be too large of a scope for a single paper. Here, we focus our efforts on GLMs with lower-bounded or double-bounded continuous likelihoods. Within these classes, we not only have several qualitatively different (non-nested) likelihood options, but can also study both main classes of non-identity link functions.

Below, we give a short overview of the likelihood and link functions included in the simulations. For consistency, we use mean parameterizations of likelihoods throughout or, if those are unavailable, median parameterizations. A more detailed review of the considered options, our inclusion criteria, and the used parameterizations are available in the supplemental material and online appendix [52].

2.2.1. Likelihoods and links for double-bounded responses

Without loss of generality, any double-bounded response can be linearly transformed to the unit interval. Accordingly, it is sufficient to focus on likelihoods for unit interval data. We included the beta [58], Kumaraswamy [59], simplex [60], and transformed-normal [61,62] likelihoods. The transformed-normal likelihoods arise from applying the link functions to the response variable y (e.g. resulting in the logit-normal likelihood), instead of to the location parameter μ as would be usual in standard GLMs. All of these likelihoods have two distributional parameters, one location (mean or median) and one scale/shape parameter, only the former being predicted as per the causal graph structures. Figure 4 shows some example densities for each likelihood, illustrating qualitatively different shapes they can accommodate. The three distinct shapes are unimodal symmetric and asymmetric shapes as well as a bimodal bathtub shape. For the remainder of the paper, we will refer to these shapes as symmetric, asymmetric, and bathtub, respectively. As link functions, we included the logit, cloglog, and cauchit links, each of them having qualitatively different properties [45,63–70]. The logit link is based on the symmetric, light-tailed logistic distribution, the cloglog link is based on the asymmetric Gumbel distribution, while the cauchit link is based on the symmetric, heavy tailed Cauchy distribution. Since logit and probit yield almost indistinguishable results due to the similar shapes of the logistic and normal distributions [45,67,68], we decided against including the probit link despite its prominence.

2.2.2. Likelihoods and links for lower-bounded responses.

Without loss of generality, any continuous lower-bounded response can be linearly transformed to have a lower bound of zero. Accordingly, it is sufficient to focus on likelihoods for strictly positive data. We included the gamma, Weibull, Fréchet, inverse Gaussian, beta prime, Gompertz, and transformed-normal likelihoods¹ all of whom have two distributional parameters, namely location (mean or median) and scale or shape). Figure 5 shows example densities for each likelihood, illustrating qualitatively different kinds of shapes they can accommodate. The three distinct shapes are unimodal thin tail and heavy tail shapes as well as a ramp shape. For the remainder of the paper we will refer to these shapes as thin tail, heavy tail, and ramp, respectively. As link functions, we included the log and the softplus link. In contrast to the multiplicative log link, softplus approaches the identity for larger values, thus approximating additive behavior of regression terms while enforcing positive predictions at the same time [71,72].

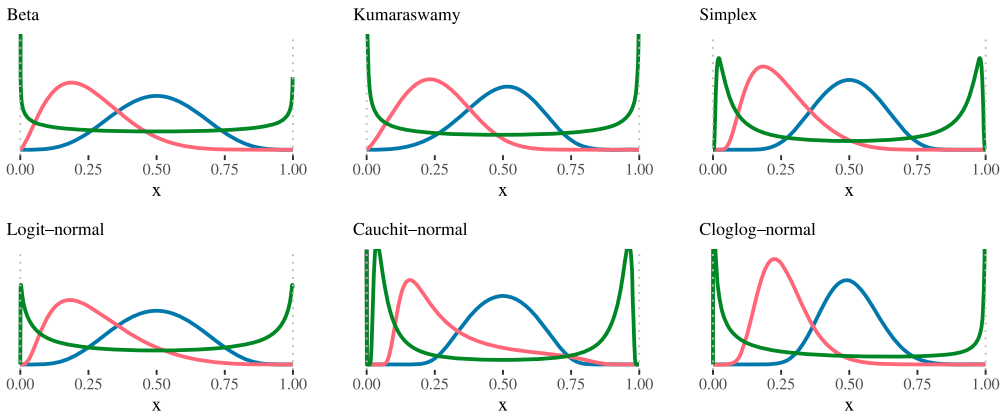


Figure 4. Example illustrations of all included double-bounded densities each with three different shapes. The shapes result from different distributional parameters (detailed in the online Appendix [52]). The parameters were specifically chosen to produce the three qualitatively different shapes accommodated by the different likelihoods, namely a symmetric, an asymmetric and a bimodal bathtub shape. The y-axis is truncated at 5 from above for better visibility of the different shapes.

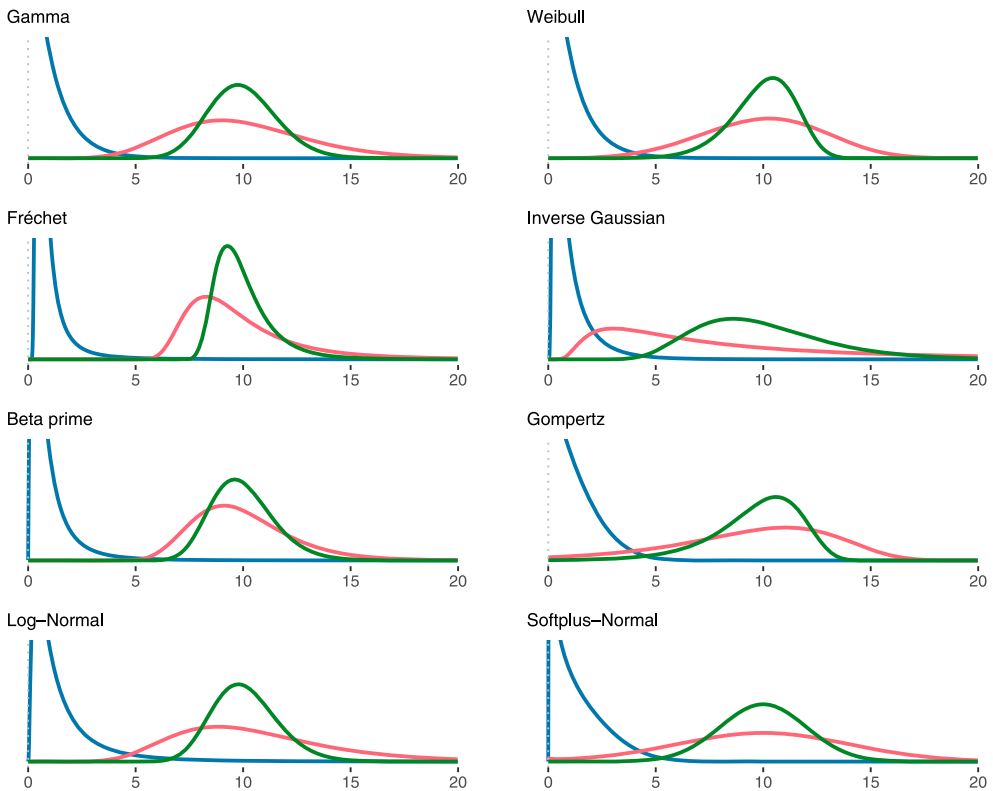


Figure 5. Example illustrations of all included lower-bounded densities, each with three different shapes. The shapes result from different distributional parameters (detailed in the online Appendix [52]). The parameters were specifically chosen to produce the three qualitatively different shapes accommodated by the different likelihoods, namely a more symmetric, a more asymmetric and a ramp shape. The y-axis is truncated at 0.4 from above for better visibility of the different shapes.

2.3. Data generation

For each of the combinations of likelihood family, likelihood shape, and link presented in Section 2.2, we simulated data sets based on the full DAG (Figure 3 left-hand side) as follows:

$$\begin{aligned} z_1 &\sim \text{normal}(0, \sigma_{z_1}), & z_2 &\sim \text{normal}(0, \sigma_{z_2}), & z_3 &\sim \text{normal}(0, \sigma_{z_3}) \\ x &\sim \text{normal}(\beta_{z_1x}z_1 + \beta_{z_3x}z_3, \sigma_x) \\ y &\sim \text{likelihood}(\text{link}^{-1}(\alpha_y + \beta_{xy}x + \beta_{z_1y}z_1 + \beta_{z_2y}z_2), \phi) \\ z_4 &\sim \text{normal}(\beta_{xz_4}x + \beta_{yz_4}y, \sigma_{z_4}). \end{aligned}$$

Here, ϕ denotes the second distributional (scale or shape) parameter of the specified likelihood family and link^{-1} denotes the inverse link (or response-) function. We chose normal distributions to generate all data variables besides y to limit the scope of our simulations. Each simulated data set contained 100 observations. Since the fitted models only have between 4 and 6 parameters, they are simple enough to be well identified on the basis of 100 observations alone. The parameters of the true DGP were set to fixed values, rather than being drawn from a prior distribution [73], to avoid blurring the effect of the induced causal misspecifications and different likelihood shapes. This decision was made due to the extreme data sets that prior distributions over parameters could create, which is a general challenge for Bayesian simulations [74,75].

Each likelihood's second distributional parameter ϕ as well as the intercept α_y were chosen to produce the shapes presented in Section 2.2. The individual coefficients for x and the z_i were calibrated so that the parameter recovery was imperfect for the ideal model while also preventing the causally misspecified models from consistently failing (see also Section 2.4). The true causal effect β_{xy} of x on y was either fixed to zero or set to a non-zero value that was calibrated together with all other coefficients. To prevent response values from under- or overflowing to the lower- or upper boundaries numerically, we truncated them near the boundaries around the 6th digit.

For each data generation configuration implied by fully crossing the design factors (see Table 1), we generated 200 data sets, which resulted in 14,400 data sets each for the double- and single-bounded scenarios.

2.4. Model fitting

On each generated data set, we fitted all models resulting from the fully crossed combination of likelihoods and links (see Section 2.2) as well as the five different linear predictor terms implied by the ideal and misspecified causal models (see Section 2.1). In reference to R formula syntax, we will also refer to the different linear predictor terms as *formulas* in the following. Accordingly, given a data set generated from a double-bounded likelihood, we fitted models on that data set using all combinations of double-bounded likelihoods, links, and DAG-based formulas (see Figure 3). The same approach was followed for the lower-bounded data. An overview of the model fit configurations is given in Table 2.

The fully-crossed design results in 60 fit configurations for the double-bounded and lower-bounded models. Multiplied with 14,400 data sets each, this leads to a total of 864,000 double-bounded and single-bounded models each fitted in our simulations.

Contrary to what we would recommend in practical applications of Bayesian models, we used flat priors for all model parameters, as it is not clear to us how one would specify equivalent priors for the different auxiliary parameters ϕ across all likelihood. Additionally, different links imply different latent scales, which render the regression coefficients' scales incomparable across (assumed) links and thus further complicates equivalent prior specification. In a real-world analysis, we would prefer to use at least weakly-informative priors [1,2,50]. Accordingly, the current choice is to be understood only in the context of the present simulation study aiming to ensure comparability across likelihoods and links. In pilot experiments (not shown here), we have confirmed that the differences in posteriors as well as the implied prediction metrics between models with flat vs. weakly informative priors are minimal for the models under investigation. We argue that such minimal differences do not justify extensive evaluation of different prior choices, since this is not the focus of the present paper. Additionally, the use of flat priors results in model estimations very similar to maximum likelihood estimation, which is another reason why we expect the results of this study to generalize to frequentist models as well.

All models were fitted using Stan [50,51] via brms [47] with two chains, 500 warmup- and 2000 post-warmup samples, which resulted in 4000 total post-warmup posterior samples per model. We used an initialization range of 0.1 around the origin on the unconstrained parameter space to avoid occasional initialization failures. For all other MCMC hyperparameters, we applied the brms defaults [47].

2.5. Model-based metrics

To measure parameter recovery and predictive performance of each fitted model, we used multiple metrics as detailed below. Implementations of these metrics are provided in the R packages loo [76], posterior [77], bayesim [4], and bayeshear [6].

2.5.1. Parameter recoverability

To date, the primary metric for causal parameter recoverability (operationalizing explanation) remains the estimation bias. However, as argued in Section 1.1, it may not be the only sensible metric. Instead, in order to allow for a bias-variance trade-off in the causal estimate's evaluation, the RMSE may be a worthy alternative metric in non-asymptotic regimes.

More precisely, given a true parameter value θ and a set of corresponding posterior samples $\{\theta^{(s)}\}$, we compute the sampling-based posterior bias and RMSE as

$$\text{bias}(\theta^{(s)}) := \frac{1}{S} \sum_{s=1}^S (\theta^{(s)}) - \theta, \quad (1)$$

$$\text{RMSE}(\theta^{(s)}) := \sqrt{\frac{1}{S} \sum_{s=1}^S ((\theta^{(s)} - \theta)^2)} = \sqrt{\text{bias}(\theta^{(s)})^2 + \text{Var}(\theta^{(s)})}, \quad (2)$$

where $\text{Var}(\theta^{(s)})$ denotes the variance over the posterior samples. Our analysis focuses on the true causal effect β_{xy} and hence all metrics were computed for $\theta = \beta_{xy}$. Furthermore, as we are interested in the size of the bias but not in its direction, we will present the absolute bias in our results.

The above are reasonable measures for comparing models only if the assumed link coincides with the true link of the DGP. This is because the link determines the scale of the linear predictor and thus the comparability of the posterior samples $\{\beta_{xy}^{(s)}\}$ with the true parameter value β_{xy} . To enable a comparison of models using different links, we also calculated the false positive rate (FPR; i.e. Type I-error rate) and the true positive rate (TPR; i.e. statistical power; inverse of the Type II-error rate) implied by the central 95% credible interval of $\{\beta_{xy}^{(s)}\}$. These metrics can be inferred from our simulations, because the true causal effect β_{xy} was set to zero in some conditions (to study FPR) and to non-zero values in others (to study TPR).

For a deeper dive into posterior accuracy and calibration under model misspecification see [78].

2.5.2. Predictive performance

In the absence of any case-specific arguments for a particular predictive metric, log-probability scores are recommended as a general-purpose choice [12]. For this reason, we computed the expected log pointwise predictive density [ELPD; 11,12] as our main metric for predictive performance (operationalizing prediction), which is defined as

$$\text{ELPD}(y^*) := \sum_{i=1}^{N^*} \log p(y_i^* | \tilde{y}) = \sum_{i=1}^{N^*} \log \left(\frac{1}{S} \sum_{s=1}^S p(y_i^* | \theta^{(s)}) \right), \quad (3)$$

where \tilde{y} denotes the training data and y^* denotes $N^* = 100$ independent test data points (previously unseen by the model) simulated from the same DGP configuration as the training data. We refer to the above metric as $\text{ELPD}_{\text{test}}$. In addition, we also calculated ELPD via leave-one-out cross-validation (LOO-CV) as approximated via Pareto-smoothed importance sampling [PSIS; 12,76]. We refer to this metric as ELPD_{loo} . Approximate LOO-CV metrics have the advantage that they are readily and efficiently available also when analysing real data, at the expense of being only an approximation that might fail to estimate out-of-sample predictive performance accurately if there are influential observations [12]. As we have the true data-generating process available during simulations, we used $\text{ELPD}_{\text{test}}$ to validate the results of ELPD_{loo} , but used the latter for our analyses as ELPD_{loo} is available also in real-world scenarios.

In order to compare multiple models with respect to their ELPD performance, we calculated the ELPD difference to the best performing model from a set $M := \{m_1, \dots, m_K\}$ of K models all fit to the same training data:

$$\Delta \text{ELPD}(y^* | m_i) := \text{ELPD}(y^* | m_i) - \max(\text{ELPD}(y^* | M)). \quad (4)$$

As explained in Section 2.1, we always restricted comparisons (i.e. M) to models that used the same linear predictor. Additionally, for investigating posterior bias and RMSE performance, only models assuming the correct link function were compared.

2.6. Statistical analysis

To analyse the relationship between predictive performance, as measured by $\Delta \text{ELPD}_{\text{loo}}$ and $\Delta \text{ELPD}_{\text{test}}$, and parameter recoverability, as measured by absolute bias, RMSE, FPR, and TPR, we modeled the simulation results via Bayesian multilevel-models (BMMs).

Simulation results of models that did not converge were excluded from the analysis. Specifically, we treated models as converged if the posterior samples $\{\beta_{xy}^{(s)}\}$ yielded $\widehat{R} < 1.01$ and $\text{ESS} > 400$ [for details on these thresholds, see 79]. We further required models to have less than 10 divergent transitions out of a total of 4000 post-warmup iterations. Ideally, we would like all models to converge with no divergent transitions, but this would have required extensive manual intervention to resolve all individual sampling problems, which is practically infeasible in our large simulation setup. Within the set of convergent models, PSIS diagnostics indicated good ELPD_{100} approximations such that the latter can be treated as a trustworthy estimate of out-of-sample predictive performance [12]. Finally, we set a lower threshold $\Delta\text{ELPD}_{100} > -N (= -100)$ to exclude models that made exceptionally bad predictions compared to the best performing model on a given data set. This resulted in approximately 7% of models being filtered from the results. Most of the filtered models (i.e. of the the approximately 7%) from the double-bounded results were fit on data generated from a bathtub link (73%) and filtered due to missing the ΔELPD_{100} threshold (94%). For the lower-bounded models, most filtered models failed to meet our convergence criteria (74%) and were fit using a softplus link (73%) and Fréchet likelihood.

In the linear predictor term of the BMMs, we included overall (‘fixed’) effects of ΔELPD in interaction with formula (i.e. the linear predictor term), (true) data generating link, and data generating shape, which we refer to as the *global slopes* in the results in Section 3. Additionally, to account for the dependency of results obtained from the same simulated data set, we included varying (‘random’) effects across data sets of ΔELPD in interaction with formula which we refer to as the *varying slopes* in the results in Section 3. For the continuous positive metrics of parameter recovery, that is, absolute bias and RMSE, we assumed a log-normal likelihood, which enabled the use of highly optimized (log-)linear regression functions in Stan [50]. For the binary metrics, that is, FPR and TPR, we assumed a canonical Bernoulli likelihood with logit link. The BMMs were estimated using a single MCMC chain with 500 warmup and 1000 post-warmup samples to keep the estimation time manageable (between one and three days wall-clock time per BMM). All convergence metrics indicated sufficient convergence. As we were mainly interested in the qualitative patterns, rather than high resolution numerical results, we considered 1000 post-warmup samples as sufficient, leading to effective sample sizes of a few hundred.

3. Results

Recall, that this paper consists of two levels of Bayesian models. The first level are the models from the simulation study, that were fitted on simulated data as described in Section 2.4. The second level are meta-models that we fit on the metrics we extracted from the first level as described in Section 2.6. All results presented in this section are based on the meta-models to learn about trends across the entire simulation study. Due to the vast number of simulation conditions, the main text illustrates selected results which are representative of the key patterns in the full simulation. The full results are available in the supplementary material and in the online Appendix [52]. In preliminary comparisons of ELPD_{100} and $\text{ELPD}_{\text{test}}$ results, we found both to be highly similar. Additionally, Pareto- \hat{k} values of the ELPD_{100} approximation were generally low indicating good approximation accuracy [12]. For these reasons, we only present results of ELPD_{100} below due to its availability during

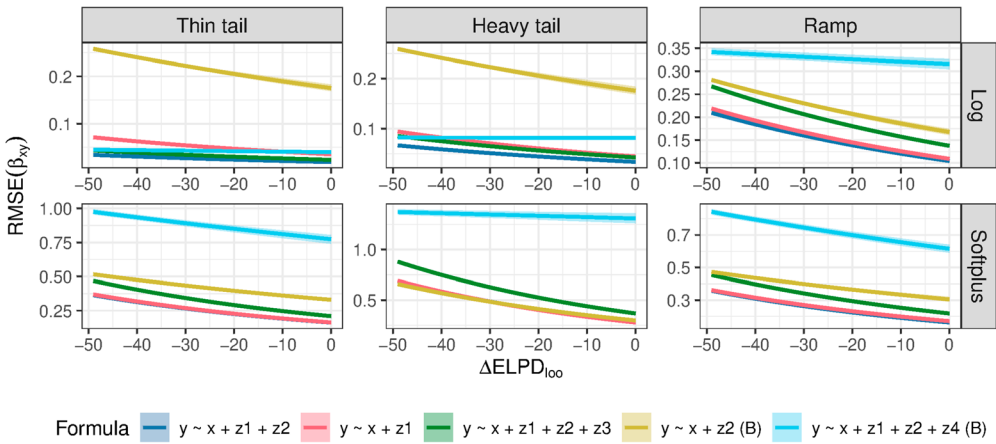


Figure 6. Global slopes of ΔELPD_{100} on $\text{RMSE}(\beta_{xy}^{(s)})$ for lower-bounded data and models. The shaded area represents the 95% credible intervals of the posterior. The results are split by data generating link and shape as well as colored by formula, where (B) indicates the causally biased formulas. The x-axis is truncated at -50 to prevent unreasonable extrapolation to sparse data spaces. The negative slopes for causally unbiased formulas have fully negative 95% credible intervals and indicate improved RMSE for better predicting models.

real-world inference. Further, we do not present results for $|\text{bias}(\beta_{xy}^{(s)})|$ here, as they show the same qualitative trends as $\text{RMSE}(\beta_{xy}^{(s)})$. Similarly, we will generally only show examples for either the double- or the lower-bounded data when the conclusions drawn from them match. Some of the presented results show floor or ceiling effects, which point to a suboptimal hyperparameter calibration, discussed in more detail in Section 4.

3.1. RMSE

Figure 6 displays the global slopes of ΔELPD_{100} on $\text{RMSE}(\beta_{xy}^{(s)})$ for the lower-bounded data, grouped by formula (i.e. linear predictor), data generating link, and shape as explained in Section 2.6. The most important observation is that all slopes of unbiased formulas were negative, often very strongly so. This means that, within a set of models fitted to the same data set and assuming the same causally unbiased formula, models that have higher (better) ΔELPD_{100} can be expected to have substantially lower (better) $\text{RMSE}(\beta_{xy}^{(s)})$. While the exact size of the improvement is dependent on the simulation’s hyperparameters, $\text{RMSE}(\beta_{xy}^{(s)})$ reduces by around 50% over the displayed span of ΔELPD_{100} , an improvement of the same order of magnitude as the true effect β_{xy} . Results of the double-bounded models showed the same behavior. For both data types, some of the unbiased formulas have rather flat slopes. This results from a floor effect in the RMSE, as slopes for formulas that generally have a small error have little room to decrease further. The causally biased formulas are less consistent, with some scenarios even showing positive slopes as exemplarily shown in Figure 7 for double-bounded models fit on logit data. Additionally, even though the generally higher $\text{RMSE}(\beta_{xy}^{(s)})$ of biased formulas would allow for steeper negative slopes, the biased slopes are often similar or even less negative than the unbiased ones.




Figure 7. Global slopes of ΔELPD_{100} on $\text{RMSE}(\beta_{xy}^{(s)})$ for double-bounded models and data generated with a logit link. The shaded area represents the 95% credible intervals of the posterior. The results are split by data generating shape and colored by formula, where (B) indicates the causally biased formulas. The x-axis is truncated at -50 to prevent unreasonable extrapolation to sparse data spaces. The negative slopes for causally unbiased formulas have fully negative 95% credible intervals and indicate improved RMSE for better predicting models.

Table 3. Proportion of negative varying slopes of ΔELPD_{100} on $\text{RMSE}(\beta_{xy}^{(s)})$ for individual data sets (higher is better).

	Log			Softplus		
	<i>Thin Tail</i>	<i>Heavy Tail</i>	<i>Ramp</i>	<i>Thin Tail</i>	<i>Heavy Tail</i>	<i>Ramp</i>
Unbiased	0.95	0.96	0.96	0.97	0.99	0.96
Biased	0.69	0.59	0.67	0.76	0.71	0.79

	Logit			Cauchit			Cloglog		
	<i>Sym</i>	<i>Asym</i>	<i>BT</i>	<i>Sym</i>	<i>Asym</i>	<i>BT</i>	<i>Sym</i>	<i>Asym</i>	<i>BT</i>
Unbiased	0.63*	0.69	0.76	0.9	0.88	0.91	0.83	0.78	0.94
Biased	0.51	0.66	0.64	0.9	0.89	0.92	0.73	0.73	0.93

* Floor effects in the global slopes were present for at least one of the respective formulas.

In Table 3, we show the proportion of negative varying slopes across individual data sets. Similarly to the global slopes, the varying slopes of the unbiased formulas were largely negative (between 95% and 99% of the varying slopes were negative for the lower-bounded data and between 63% and 94% were negative for the double-bounded data). The exact proportion of negative slopes differed somewhat between ground-truth scenarios, as a result of specific hyperparameters choices in the calibration as well as random noise induced by the simulations. However, the generally high proportions of negative varying slopes for the unbiased formulas indicates that the trend of lower (better) $\text{RMSE}(\beta_{xy}^{(s)})$ for higher (better) ΔELPD_{100} is consistent across data sets. For the causally biased formulas, the proportion of negative varying slopes were generally lower than for the unbiased formulas for each ground truth. In addition to those overall trends, the proportion of negative slopes for the double-bounded models showed more variability between ground truths than the lower-bounded data. Especially the logit link had a smaller proportion than the other scenarios as a result of floor effects.

3.2. FPR and TPR

Figure 8 shows the global slopes of ΔELPD_{100} on FPR and TPR for the lower-bounded data grouped by formula, data generating link, and shape as explained in Section 2.6. Starting

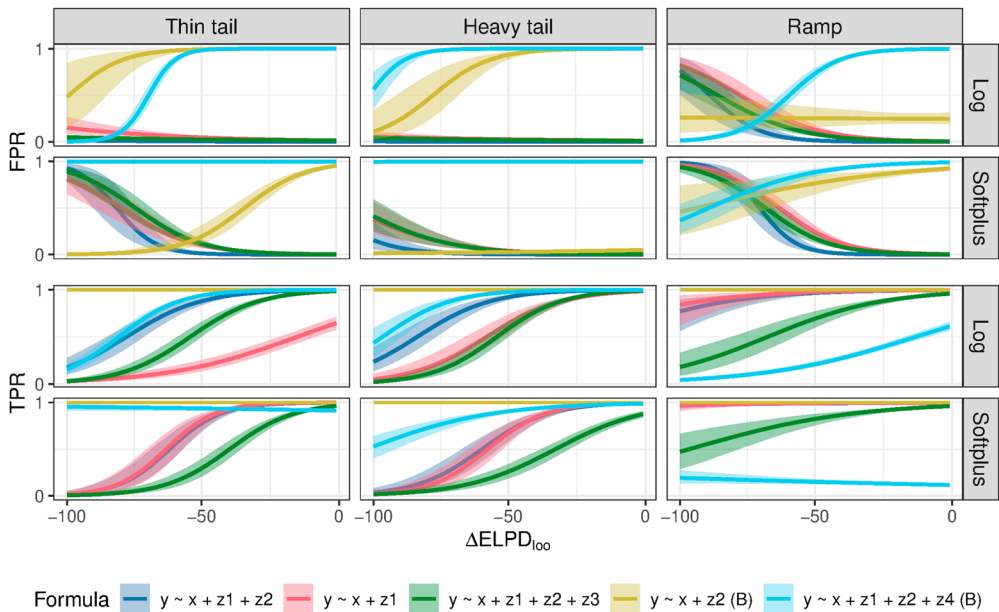


Figure 8. Global slopes of ΔELPD_{100} on the false positive and true positive rates (FPR and TPR) for lower-bounded data. The shaded area represents the 95% credible intervals of the posterior. The results are split by data generating link and shape as well as colored by formula, where (B) indicates the causally biased formulas. The low FPR and positive TPR slopes of the unbiased formulas indicate better Calibration for better predicting models.

with the FPR slopes, there is a clear distinction between the biased and unbiased formulas. Where the unbiased formulas have negative slopes or even stay constant near 0.05 (nominal Type-I error rate), the biased formulas have positive slopes or stay constant near 1 (i.e. 100% Type-I error rate) in all cases. The results for the double-bounded scenarios show the same pattern. Acceptable FPR (i.e. values close to or less than the desired Type-I error rate) is a requirement for using TPR as a measure of parameter recovery, because it is trivial to achieve high TPR if FPR is allowed to be high at the same time. Accordingly, while the low FPR for the unbiased formulas implies the usefulness of TPR in these cases, the high FPR for the biased formulas makes their TPR results much less meaningful.

In the TPR results, we found a clear pattern of positive global slopes for the unbiased formulas whenever there are no ceiling effects (i.e. $\text{TPR} = 1$ regardless of ΔELPD_{100}). This indicates, that within models fitted to the same data set and using the same causally unbiased formula, models with higher (better) ΔELPD_{100} can be expected to have substantially higher (better) TPR. The biased formulas again are less consistent in their slopes with cases of negative global slopes or shallower slopes even for generally low TPR, such that ceiling effects are not an issue. Results of the double-bounded models showed the same behavior with some additional cases of floor effects.

In Table 4, we show the proportion of positive varying slopes across individual data sets for the TPR. For the causally unbiased formulas, the majority of the varying slopes are strongly positive, such that within a set of models fitted to the same data set and assuming the same causally unbiased formula, models that ranked higher in ΔELPD_{100} can be

Table 4. Proportion of positive varying slopes of ΔELPD_{100} on TPR (higher is better).

	Log			Softplus		
	<i>Thin Tail</i>	<i>Heavy Tail</i>	<i>Ramp</i>	<i>Thin Tail</i>	<i>Heavy Tail</i>	<i>Ramp</i>
Unbiased	0.78	0.86	0.71	0.92	0.87	0.61*
Biased	0.7*	0.59*	0.51*	0.54*	0.65*	0.25*

	Logit			Cauchit			Cloglog		
	<i>Sym</i>	<i>Asym</i>	<i>BT</i>	<i>Sym</i>	<i>Asym</i>	<i>BT</i>	<i>Sym</i>	<i>Asym</i>	<i>BT</i>
Unbiased	0.94	0.93	0.24*	0.95	0.93	0.78*	0.86*	0.94	0.74*
Biased	0.94*	0.94*	0.81	0.96*	0.88*	0.87*	0.48*	0.93*	0.55

*Floor or ceiling effects in the global slopes were present for at least one of the respective formulas.

expected to have a higher TPR. This effect is rather consistent with positive proportions between 86% and 95% for different ground truths, besides the ramp and bathtub shape. The lower proportion of positive slopes for the ramp and bathtub shapes coincide with the floor and ceiling effects visible in Figure 8.

The biased formulas showed less consistency of the varying slope signs in the lower-bounded scenarios and similar consistency in the double-bounded scenarios.

4. Discussion

In this paper, we studied the use of prediction as a proxy for explanation under several causal and statistical misspecification mechanisms in Bayesian GLMs. We split our research question into two parts:

- (i) Within a set of statistical models that all share the same underlying causal model, can prediction be reliably used as a proxy for explanation?
- (ii) Does the answer to (i) depend on whether or not the causal model is biased?

Within our results, we observed one consistent trend over almost all investigated scenarios: When comparing statistical models sharing the same underlying *unbiased* causal model, all considered measures of explanation (the causal parameter's absolute bias, RMSE, true positive rate and false positive rate) improved with improving out-of-sample prediction (measured by the expected log predictive density). Further, in almost all scenarios, the trends were also highly consistent across individual data sets, which indicates that the proxy can be reliably applied in practice, where inference usually concerns a single data set only. In the few causally unbiased cases without a clear trend, the relationship was effectively zero due to floor and ceiling effects (further discussed below). While better predicting models did not necessarily provide better explanation in those cases, prediction was (on average across datasets) not worsening explanation either. For statistical models sharing the same underlying *biased* causal model, the trends were inconsistent with examples for both improving and degrading explanation with improving prediction. We conclude that, given a set of GLMs that all share the same unbiased causal model, prediction can be safely used as a proxy for explanation.

Our findings also offer empirical support for the more conceptual arguments of [24,25] that called for more attention to better predicting models, even if explanation is the

main goal. In that way, causality seems to be the missing link that connects the statistical relationship of explanation and prediction.

As is typical of simulation studies, their generalizability beyond the studied scenarios should be treated with caution. While an analytic investigation supplementing the simulations would have been desirable, the lack of available closed-form posteriors outside of a few conjugate cases prevented such an investigation (as does the iterative nature of maximum likelihood in the models' frequentist counterparts). We understand this paper as a promising starting point for further work on the relationship between explanation and prediction under different causal and statistical assumptions. Below, we offer our perspective on subsequent research questions.

As discussed in Section 1.2, we limited the scope of this paper to models with simple linear predictor terms. We would expect to observe the same general trends for more complicated additive predictor terms where individual components may be non-linear, as the DAG-based misspecifications are agnostic to the details of the statistical model structure [80]. However, investigating the implications of (mis-)specifying individual terms of the linear predictor (e.g. modeling an effect as linear while it is truly non-linear in some way) might be interesting. This would more strongly play into the question of overfitting in sparse scenarios with (relatively) small data yet complicated non-linear relationships to be approximated. Similar questions would arise when adding multilevel structure especially in sparse data scenarios where the (true) relevance of multilevel terms need to be balanced against their weak identification implied by the given data [8,81].

With regard to causal assumptions, we chose our data-generating processes (DGPs) and the four misspecified linear predictors in an attempt to span the most important and common classes of controls [20]. Still, there remain many other possible DGPs and misspecifications, for example, unobserved variables or variables with measurement error [82,83]. An extension of this work to different kinds of DGPs and misspecifications would ultimately add to the generalizability of results. That said, we currently do not see how an unbiased causal model would have to look like in order to render the relationship of prediction and explanation negative (i.e. better prediction implying worse explanation), as long as all compared models all shared the same unbiased causal model.

In contrast, when comparing models with *different* underlying causal models, the relationship between prediction and explanation is still not fully understood. Even though it is known, for example, that including colliders improves prediction but worsens explanation [2], the situation for a set of different but all unbiased causal models has not yet been systematically explored to our knowledge. That is, when all compared statistical models have underlying *unbiased* causal models, it is unclear if improved prediction implies improved explanation if the causal models are allowed to vary across the statistical models. We believe that this research direction would be one of the most important extensions of the present work.

In addition to the influence of the true and assumed causal models, we observed a noticeable effect of the true likelihood shapes assumed in our simulations. The mechanisms by which different shapes imply different downstream results, including the extend of the relationship between prediction and explanation, remain somewhat unclear to us. While the investigation of additional DGPs would generally offer potential for improving the generalizability of our findings, an extension of this work that would especially focus on different data-generating shapes and their properties more systematically could improve

said understanding. This directly ties into the problem of choosing reasonable ground-truth model parameters for the simulations. Fully Bayesian simulations tend to struggle when drawing parameters from prior distributions as this can easily lead to extreme and unrealistic data sets [74,75]. Our solution of carefully crafting scenarios that would represent specific, distinct, but still realistic, ground truths comes with the challenge of choosing an overwhelming amount of parameters. This did not fully succeed in all cases, as shown by the floor and ceiling effects that were present in a few scenarios, specifically for the true and false positive rates where causal relationships were sometimes too easy or too hard to detect. While results influenced by those effects still support our overall conclusions, they reduce what we can learn from the affected scenarios.

Finally, in this study, we followed a Bayesian perspective on model building, despite using flat priors. The latter choice was primarily made to avoid some models having an ‘unfair’ advantage due to incidentally more suitable prior choices (see Section 2.4). As a side effect of this choice, we think that our main results are likely to hold as well in case of frequentist (maximum likelihood) estimation of all models, and corresponding model-based metrics. We see it as unlikely that our conclusions would have changed in the light of applying (weakly-)informative priors [1,2,50]. However, there may be one subtle place where priors (or regularization more generally) may influence the relationship between prediction and explanation: Cross-validation measures of out-of-sample prediction produce correlated folds due to overlap in training data and have an intricate relationship with different true quantities they can be seen to approximate [84]. These properties can be altered via regularization, especially in sparse regimes where the number of model parameters is high relative to the number of observations in the data [84]. As such, the relationship between prediction and explanation may also vary in such scenarios, which would be interesting to study in the future.

Notes

1. In preliminary analysis, we observed consistently bad posterior sampling behavior of inverse-Gaussian models, with very slow sampling overall and more than half of the models failing to converge. For these reasons we decided to exclude the inverse-Gaussian likelihood from the final analysis.

Acknowledgments

We also want to thank Marvin Schmitt and Stefan Radev for their feedback on earlier versions of this manuscript as well as Yannick Dzubba for his help during development of the software used for our simulation study.

Data availability statement

The data that support the findings of this study are openly available in OSF at <http://doi.org/10.17605/OSF.IO/XGKZV>.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2075 - 390740016 (the Stuttgart Cluster of Excellence SimTech). This work was performed on the computational resource bwUniCluster funded by the Ministry of Science, Research and the Arts Baden-Württemberg and the Universities of the State of Baden-Württemberg, Germany, within the framework program bwHPC. The authors gratefully acknowledge the support and funding.

ORCID

Maximilian Scholz  <http://orcid.org/0000-0002-4486-1393>

Paul-Christian Bürkner  <http://orcid.org/0000-0001-5765-8995>

References

- [1] Gelman A, Carlin JB, Stern HS, et al. Bayesian data analysis. 3rd ed. New York: Chapman and Hall/CRC; 2013.
- [2] McElreath R. Statistical rethinking: a bayesian course with examples in R and stan. 2nd ed. New York: Chapman and Hall/CRC; 2020.
- [3] Gelman A, Vehtari A, Simpson D, et al. Bayesian Workflow. 2020 Nov. arXiv:201101808 [stat].
- [4] Scholz M, Bürkner PC. Bayesim: framework for bayesian simulation studies. 2022. R package version 0.29.5.9000; Available from: <https://github.com/sims1253/bayesim>.
- [5] Scholz M, Dzubba Y, Bürkner PC. Bayesfam: custom families for brms. 2022. R package version 0.2.4; Available from: <https://github.com/sims1253/bayesfam>.
- [6] Scholz M, Bürkner PC. Bayeshear: toolbox of bayesian model metrics. 2022. R package version 0.2.0.9000; Available from: <https://github.com/sims1253/bayeshear>.
- [7] Douglas HE. Reintroducing prediction to explanation. *Philos Sci.* 2009;76(4):444–463. doi: [10.1086/648111](https://doi.org/10.1086/648111)
- [8] Bürkner PC, Scholz M, Radev ST. What makes a good bayesian model? Towards a unified model taxonomy. preprint, 2022. arXiv.
- [9] Hastie T, Tibshirani R, Friedman JH, et al. The elements of statistical learning: data mining, inference, and prediction. Vol. 2. Heidelberg: Springer; 2009.
- [10] Gabaix X, Laibson D. The seven properties of good models. In: The foundations of positive and normative economics: A handbook. Cambridge: Oxford University Press; 2008.
- [11] Vehtari A, Ojanen J. A survey of Bayesian predictive methods for model assessment, selection and comparison. *Stat Surv.* 2012 Jan;6:142–228. doi: [10.1214/12-SS102](https://doi.org/10.1214/12-SS102)
- [12] Vehtari A, Gelman A, Gabry J. Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Stat Comput.* 2017;27(5):1413–1432. doi: [10.1007/s11222-016-9696-4](https://doi.org/10.1007/s11222-016-9696-4)
- [13] Pearl J. Causality. Cambridge: Cambridge University Press; 2009.
- [14] Pearl J. The do-calculus revisited. preprint, 2012. arXiv:12104852.
- [15] Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol.* 1974;66(5):688–701. doi: [10.1037/h0037350](https://doi.org/10.1037/h0037350)
- [16] Rubin DB. Bayesian inference for causal effects: the role of randomization. *Ann Stat.* 1978;6(1):34–58. doi: [10.1214/aos/1176344064](https://doi.org/10.1214/aos/1176344064)
- [17] Shmueli G. To explain or to predict? *Stat Sci.* 2010 Aug;25(3):289–310. doi: [10.1214/10-STS330](https://doi.org/10.1214/10-STS330)
- [18] Henckel L, Perković E, Maathuis MH. Graphical criteria for efficient total effect estimation via adjustment in causal linear models. *J R Stat Soc Ser B.* 2022;84(2):579–599. doi: [10.1111/rssb.12451](https://doi.org/10.1111/rssb.12451)
- [19] Rotnitzky A, Snucler E. Efficient adjustment sets for population average causal treatment effect estimation in graphical models. *J Mach Learn.* 2020;21(1):7642–7727.
- [20] Cinelli C, Forney A, Pearl J. A crash course in good and bad controls. *Sociol Methods Res.* 2020;53(3):1071–1104.

- [21] Ramspek CL, Steyerberg EW, Riley RD, et al. Prediction or causality? a scoping review of their conflation within current observational research. *Eur J Epidemiol.* 2021;36:889–898. doi: [10.1007/s10654-021-00794-w](https://doi.org/10.1007/s10654-021-00794-w)
- [22] Navarro DJ. Between the devil and the deep blue sea: tensions between scientific judgement and statistical model selection. *Comput Brain Behav.* 2019 Mar;2(1):28–34. doi: [10.1007/s42113-018-0019-z](https://doi.org/10.1007/s42113-018-0019-z)
- [23] Hernán MA, Hsu J, Healy B. A second chance to get causal inference right: a classification of data science tasks. *Chance.* 2019;32(1):42–49. doi: [10.1080/09332480.2019.1579578](https://doi.org/10.1080/09332480.2019.1579578)
- [24] Yarkoni T, Westfall J. Choosing prediction over explanation in psychology: lessons from machine learning. *Perspect Psychol Sci.* 2017 Nov;12(6):1100–1122. doi: [10.1177/1745691617693393](https://doi.org/10.1177/1745691617693393)
- [25] Breiman L. Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Stat Sci.* 2001 Aug;16(3):199–231. doi: [10.1214/ss/1009213726](https://doi.org/10.1214/ss/1009213726)
- [26] Cranmer SJ, Desmarais BA. What can we learn from predictive modeling? *Polit Anal.* 2017;25(2):145–166. doi: [10.1017/pan.2017.3](https://doi.org/10.1017/pan.2017.3)
- [27] Gilpin LH, Bau D, Yuan BZ, et al. Explaining explanations: an overview of interpretability of machine learning. In: 2018 IEEE 5th International Conference on data science and advanced analytics (DSAA); IEEE; 2018. p. 80–89.
- [28] Carvalho DV, Pereira EM, Cardoso JS. Machine learning interpretability: A survey on methods and metrics. *Electronics.* 2019;8(8):832. doi: [10.3390/electronics8080832](https://doi.org/10.3390/electronics8080832)
- [29] Li J, Jia X, Yang H, et al. Incorporating causal effects into deep learning predictions on ehr data. preprint, 2022. arXiv:201105466.
- [30] Richens JG, Lee CM, Johri S. Improving the accuracy of medical diagnosis with causal machine learning. *Nat Commun.* 2020;11(1):3923. doi: [10.1038/s41467-020-17419-7](https://doi.org/10.1038/s41467-020-17419-7)
- [31] Schölkopf B. Causality for machine learning. 1st ed. New York (NY): Association for Computing Machinery; 2022. p. 765–804. doi: [10.1145/3501714.3501755](https://doi.org/10.1145/3501714.3501755).
- [32] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys.* 2019;378:686–707. doi: [10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045)
- [33] Karniadakis GE, Kevrekidis IG, Lu L, et al. Physics-informed machine learning. *Nat Rev Phys.* 2021;3(6):422–440. doi: [10.1038/s42254-021-00314-5](https://doi.org/10.1038/s42254-021-00314-5)
- [34] Arrieta AB, Díaz-Rodríguez N, Del Ser J, et al. Explainable artificial intelligence (xai): concepts, taxonomies, opportunities and challenges toward responsible ai. *Inf Fusion.* 2020;58:82–115. doi: [10.1016/j.inffus.2019.12.012](https://doi.org/10.1016/j.inffus.2019.12.012)
- [35] Scott M, Su-In Lee. A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst.* 2017;30:4765–4774.
- [36] Ribeiro MT, Singh S, Guestrin C. ‘why should I trust you?’: Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining; 2016 Aug 13–17; San Francisco (CA); 2016. p. 1135–1144.
- [37] Watanabe S. Algebraic geometry and statistical learning theory. Vol. 25. Cambridge: Cambridge University Press; 2009.
- [38] Spiegelhalter DJ, Best NG, Carlin BP. Bayesian deviance, the effective number of parameters, and the comparison of arbitrarily complex models. Citeseer; 1998.
- [39] Watanabe S, Opper M. Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *J Mach Learn Res.* 2010;11(12).
- [40] Betancourt M. Towards a principled bayesian workflow. 2020. Available from: https://betanalpha.github.io/assets/case_studies/principled_bayesian_workflow.html.
- [41] Hagerty MR, Srinivasan V. Comparing the predictive powers of alternative multiple regression models. *Psychometrika.* 1991;56(1):77–85. doi: [10.1007/BF02294587](https://doi.org/10.1007/BF02294587)
- [42] Gelman A, Hill J, Vehtari A. Regression and other stories. Cambridge: Cambridge University Press; 2020.
- [43] Harrell FE. Regression modeling strategies: with applications to linear models, logistic and ordinal regression, and survival analysis. Heidelberg: Springer; 2015.

- [44] Nelder JA, Wedderburn RW. Generalized linear models. *J R Stat Soc Ser A (General)*. 1972;135(3):370–384. doi: [10.2307/2344614](https://doi.org/10.2307/2344614)
- [45] Gill J, Gill JM, Torres M, et al. Generalized linear models: a unified approach. Vol. 134. California: Sage Publications; 2001.
- [46] Bates D, Mächler M, Bolker B, et al. Fitting linear mixed-effects models using `\pkg{lme4}`. *J Stat Softw*. 2015;67(1):1–48. doi: [10.18637/jss.v067.i01](https://doi.org/10.18637/jss.v067.i01)
- [47] Bürkner PC. brms: an r package for bayesian multilevel models using stan. *J Stat Softw*. 2017;80:1–28.
- [48] Gelman A, Simpson D, Betancourt M. The prior can often only be understood in the context of the likelihood. *Entropy*. 2017;19(10):555–567. doi: [10.3390/e19100555](https://doi.org/10.3390/e19100555)
- [49] R Core Team. R: a language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing; 2018. Available from: <https://www.R-project.org/>.
- [50] Stan Development Team. Stan modeling language users guide and reference manual, 2.30.0. 2022. Available from: <https://mc-stan.org>.
- [51] Carpenter B, Gelman A, Hoffman MD, et al. Stan: A probabilistic programming language. *J Stat Softw*. 2017;76(1). doi: [10.18637/jss.v076.i01](https://doi.org/10.18637/jss.v076.i01)
- [52] Scholz M, Bürkner PC. Prediction can be safely used as a proxy for explanation in causally consistent bayesian generalized linear models. 2022. Available from: <https://osf.io/xgkzv/>.
- [53] Johnson NL, Kotz S, Balakrishnan N. Continuous univariate distributions. Vol. 2. New Jersey: John Wiley & Sons; 1995.
- [54] Johnson NL, Kotz S, Kemp AW. Univariate discrete distributions. New Jersey: John Wiley & Sons; 2005.
- [55] Stasinopoulos DM, Rigby RA. Generalized additive models for location scale and shape (GAMLSS) in R. *J Stat Softw*. 2007;23(7):1–46. doi: [10.18637/jss.v023.i07](https://doi.org/10.18637/jss.v023.i07)
- [56] Yee TW. The VGAM package for categorical data analysis. *J Stat Softw*. 2010;32(10):1–34. doi: [10.18637/jss.v032.i10](https://doi.org/10.18637/jss.v032.i10)
- [57] Bürkner PC. Bayesian item response modelling in R with brms and stan. *J Stat Softw*. 2021;100(5):1–54.
- [58] Espinheira PL, Ferrari SL, Cribari-Neto F. On beta regression residuals. *J Appl Stat*. 2008;35(4):407–419. doi: [10.1080/02664760701834931](https://doi.org/10.1080/02664760701834931)
- [59] Kumaraswamy P. A generalized probability density function for double-bounded random processes. *J Hydrol (Amst)*. 1980 Mar;46(1):79–88. doi: [10.1016/0022-1694\(80\)90036-0](https://doi.org/10.1016/0022-1694(80)90036-0)
- [60] Barndorff-Nielsen OE, Jørgensen B. Some parametric models on the simplex. *J Multivar Anal*. 1991 Oct;39(1):106–116. doi: [10.1016/0047-259X\(91\)90008-P](https://doi.org/10.1016/0047-259X(91)90008-P)
- [61] Atchison J, Shen S. Logistic-normal distributions: Some properties and uses. *Biometrika*. 1980 Jan;67(2):261–272. doi: [10.1093/biomet/67.2.261](https://doi.org/10.1093/biomet/67.2.261)
- [62] Kim S, Heath E, Heilbrun L. Sample size determination for logistic regression on a logit-normal distribution. *Stat Methods Med Res*. 2017 Jun;26(3):1237–1247. doi: [10.1177/0962280215572407](https://doi.org/10.1177/0962280215572407)
- [63] Yin S, Dey DK, Valdez EA, et al. Skewed link regression models for imbalanced binary response with applications to life insurance. 2020 Jul. arXiv:200715172 [stat].
- [64] Jiang X, Dey DK, Prunier R, et al. A new class of flexible link functions with application to species co-occurrence in cape floristic region. *Ann Appl Stat*. 2013;7(4):2180–2204. doi: [10.1214/13-AOAS663](https://doi.org/10.1214/13-AOAS663)
- [65] Lemonte AJ, Bazán JL. New links for binary regression: an application to coca cultivation in Peru. *Test*. 2018 Sep;27(3):597–617. doi: [10.1007/s11749-017-0563-1](https://doi.org/10.1007/s11749-017-0563-1)
- [66] Damisa S, Musa F, Sani S. On the comparison of some link functions in binary response analysis. *J Biomed Stat Inform*. 2017;2(4):145–149.
- [67] Fahrmeir L, Tutz G, Hennevogl W, et al. Multivariate statistical modelling based on generalized linear models. Vol. 425. Heidelberg: Springer; 1994.
- [68] Powers D, Xie Y. Statistical methods for categorical data analysis. Leeds: Emerald Group Publishing; 2008.
- [69] Morgan BJ, Smith D. A note on Wadley’s problem with overdispersion. *J R Stat Soc Ser C (Appl Stat)*. 1992;41(2):349–354.

- [70] Koenker R, Yoon J. Parametric links for binary choice models: A Fisherian–Bayesian colloquy. *J Econom.* 2009 Oct;152(2):120–130. doi: [10.1016/j.jeconom.2009.01.009](https://doi.org/10.1016/j.jeconom.2009.01.009)
- [71] Zheng H, Yang Z, Liu W, et al. Improving deep neural networks using softplus units. In: 2015 International Joint Conference on Neural Networks (IJCNN); IEEE; 2015. p. 1–4.
- [72] Dugas C, Bengio Y, Bélisle F, et al. Incorporating second-order functional knowledge for better option pricing. *Adv Neural Inf Process Syst.* 2000;13
- [73] Talts S, Betancourt M, Simpson D, et al. Validating bayesian inference algorithms with simulation-based calibration. preprint, 2018. arXiv:180406788.
- [74] Gabry J, Simpson D, Vehtari A, et al. Visualization in Bayesian workflow. *J R Stat Soc Ser A (Stat Soc).* 2019 Feb;182(2):389–402. doi: [10.1111/rssa.12378](https://doi.org/10.1111/rssa.12378)
- [75] Mikkola P, Martin OA, Chandramouli S, et al. Prior knowledge elicitation: the past, present, and future. preprint, 2021. arXiv:211201380.
- [76] Vehtari A, Gabry J, Magnusson M, et al. loo: Efficient leave-one-out cross-validation and waic for bayesian models. 2022. R package version 2.5.0; Available from: <https://mc-stan.org/loo/>.
- [77] Bürkner PC, Gabry J, Kay M, et al. Posterior: tools for working with posterior distributions. 2022. R package version 1.3.0; Available from: <https://mc-stan.org/posterior/>.
- [78] Scholz M, Bürkner PC. Posterior accuracy and calibration under misspecification in bayesian generalized linear models. preprint, 2023. arXiv:231109081.
- [79] Vehtari A, Gelman A, Simpson D, et al. Rank-normalization, folding, and localization: an improved \hat{R} for assessing convergence of MCMC. *Bayesian Anal.* 2021 Jun;16(2):667–718. doi: [10.1214/20-BA1221](https://doi.org/10.1214/20-BA1221).
- [80] Pearl J. Causal inference in statistics: an overview. *Stat Surv.* 2009 Jan;3:96–146. doi: [10.1214/09-SS057](https://doi.org/10.1214/09-SS057)
- [81] Gelman A, Hill J. Data analysis using regression and multilevel/hierarchical models. Cambridge: Cambridge University Press; 2006.
- [82] Kuroki M, Pearl J. Measurement bias and effect restoration in causal inference. *Biometrika.* 2014;101(2):423–437. doi: [10.1093/biomet/ast066](https://doi.org/10.1093/biomet/ast066)
- [83] Tchetgen EJT, Ying A, Cui Y, et al. An introduction to proximal causal learning. preprint, 2020. arXiv:200910982.
- [84] Bates S, Hastie T, Tibshirani R. Cross-validation: what does it estimate and how well does it do it? preprint, 2021. arXiv:210400673.