

A causal viewpoint on prediction model performance under changes in case-mix: discrimination and calibration respond differently for prognosis and diagnosis predictions

Wouter van Amsterdam

September 4, 2024

Abstract

Prediction models inform important clinical decisions, aiding in diagnosis, prognosis, and treatment planning. The predictive performance of these models is typically assessed through discrimination and calibration. However, changes in the distribution of the data impact model performance. In health-care, a typical change is a shift in case-mix: for example, for cardiovascular risk management, a general practitioner sees a different mix of patients than a specialist in a tertiary hospital.

This work introduces a novel framework that differentiates the effects of case-mix shifts on discrimination and calibration based on the causal direction of the prediction task. When prediction is in the causal direction (often the case for prognosis predictions), calibration remains stable under case-mix shifts, while discrimination does not. Conversely, when predicting in the anti-causal direction (often with diagnosis predictions), discrimination remains stable, but calibration does not.

A simulation study and empirical validation using cardiovascular disease prediction models demonstrate the implications of this framework. This framework provides critical insights for evaluating and deploying prediction models across different clinical settings, emphasizing the importance of understanding the causal structure of the prediction task.

1 Introduction

Clinicians use prediction models for medical decisions, for example when making a diagnosis, estimating a patient’s prognosis, or when making triaging or treatment decisions. When basing a medical decision on a prediction model it is important to know how reliable the model’s predictions are, i.e. what is the model’s *predictive performance*, typically measured with *discrimination* and *calibration* in the case of binary outcomes. Discrimination measures how well a prediction model separates *positive* cases from *negative* cases, whereas *calibration* measures how well predicted probabilities align with observed event rates.


An issue with predictive performance is that there may be important changes between when a model’s predictive performance was last evaluated, and when and where it is used, meaning that the underlying *data distribution* may have changed. No model can have good predictive performance under all arbitrary changes in the data distribution, but we may consider one important class of changes in distribution described with the term ‘case-mix’. For instance, when comparing the general practitioner (GP) setting with a tertiary hospital setting, the frequency of certain comorbidities and risk factors will be different across settings. Typically the tertiary center will encounter more high risk patients, so their ‘case-mix’ is different than in the GP setting. Another change in case-mix is the frequency of myocardial infarction in patients presenting with chest pain at either the GP or in those referred to acute cardiac care centers.

Changes in the data distribution necessarily lead to changes in predictive performance. However, we introduce a new framework that shows that when the prediction task is in the *causal* direction, which is often the case in prognosis predictions, versus in the *anti-causal* direction, often the case when predicting a diagnosis, a change in case-mix has a different interpretation. Importantly, depending on the prediction direction, either calibration *or* discrimination is preserved under shifts in case-mix, but never both. We achieve this by looking at diagnosis and prognosis through a causal lens. Inferring a diagnosis is typically prediction in the *anti-causal* direction, meaning predicting


the cause (=underlying diagnosis) based on its effects (=symptoms), and here a change in case-mix is a change in distribution of the prediction target (the diagnosis). In contrast, prognosis is typically prediction in the *causal* direction, meaning a future outcome predicted from current patient characteristics, and here a shift in case-mix is a change in the distribution of patient characteristics. The crucial insight of our framework is that a prediction model's *discrimination* depends on the distribution of the features given the outcome (X given Y) and is thereby invariant to changes in the marginal distribution of the outcome. Conversely, *calibration* depends on the distribution of the outcome given the features (Y given X) and is thereby invariant to changes in the marginal distribution of the features. See Figure 1 for a schematic overview.

Our new perspective sheds light on changes in predictive performance under shifts in case-mix. Our result shows that the causal direction of the prediction has important implications for the evaluation and deployment of prediction models. For example, when evaluating a model used for prognosis across different settings, changes in discrimination are expected under shifts in case-mix, but changes in calibration are not, and vice-versa for diagnostic models. When re-evaluating a prognostic model in a different setting, a change in discrimination is expected and thus no cause for concern. However, a marked change in calibration may warrant further research. This perspective helps developers and guideline makers judge where and when a prediction model has dependable predictive performance. Also, depending on the task and whether discrimination or calibration is more important, prediction model developers may improve the robustness of their model to changes in case-mix by only including variables in the prediction model that are either all causal or all anti-causal but not mixed when possible.


To introduce the framework, we first review the concepts of discrimination and calibration and then we define changes in case-mix from a causal viewpoint. Next we put the two pieces together in a new framework and answer: when to expect what changes in predictive performance? We illustrate the result with a simulation study and test the framework empirically in a systematic review of 1382 prediction models, where we find that prognostic models indeed have more variance in discrimination when tested in external validation studies. Finally we discuss how this theory can be applied in practice.




(a) DAG for prediction models predicting in the *causal* direction, as in many *prognosis* settings (e.g. predict future heart attacks based on current age and cholesterol levels)



(b) DAG for prediction models predicting in the *anti-causal* direction as in many *diagnosis* settings (e.g. predict presence of a current heart attack based on the presence of chest pain and electrocardiography abnormalities).



(c) Between the training data and testing data, the *calibration* remains the same (upper facet), but the *discrimination* changes (lower facet)



(d) Between the training data and testing data, the *calibration* changes (upper facet), but the *discrimination* remains the same (lower facet)

Figure 1: Overview of main results. Depending on the causal direction of the prediction, a shift in ‘case-mix’ may be defined as either a shift in the marginal distribution of the *features* X for *causal* prediction (1a) or a shift in the marginal distribution of the *outcome* Y for *anti-causal* prediction (1b). With these definitions, for models predicting in the *causal* direction, the *calibration* will remain constant under case-mix shifts between the training data and the testing data but not the *discrimination* (1c). For models predicting in the *anti-causal* direction the reverse is true (1d). The calibration facets are calibration curves with on the horizontal axis the predicted probability and on the vertical axis the actual probability. The discrimination facets are receiver-operating-curves with on the horizontal axis 1 minus specificity and on the vertical axis sensitivity. DAG: directed acyclic graph

		outcome (Y)	
		1	0
prediction ($f(X) > \tau$)	1	true positive	false positive
	0	false negative	true negative
		sensitivity: $P(f(X) > \tau Y = 1)$	specificity: $P(f(X) \leq \tau Y = 0)$

Table 1: Confusion table. By specifying a threshold $0 < \tau < 1$ for a prediction model $f : \mathcal{X} \rightarrow [0, 1]$ and tabulating the results against the ground truth outcome $Y \in \{0, 1\}$, one gets the confusion table and can calculate metrics of discrimination such a sensitivity and specificity.

2 Notation and Recap of predictive performance: discrimination and calibration

We consider prediction models of a binary *outcome* Y using *features* X with a prediction model $f : \mathcal{X} \rightarrow [0, 1]$. The features can come from an arbitrary (multi-)dimensional distribution but we assume the model’s predictions $f(X)$ can reasonably be considered as continuous. We will denote environments with an environment variable E where for example $E = 0$ may be a GP setting and $E = 1$ a community hospital, $E = 2$ a university medical center and so on[1]. With $P(\cdot)$ we will denote (conditional) distributions or densities over random variables, for example $P(Y|X)$ denotes the distribution of outcome Y given features X .

2.1 Discrimination: sensitivity, specificity and AUC

Typical metrics of discrimination are sensitivity (sometimes called recall), specificity and AUC. Sensitivity is the ratio of true positives over the total number of positive cases. Specificity is the ratio of true negatives over the total number of negative cases. To calculate sensitivity and specificity, we need to choose a threshold $0 < \tau < 1$ for the output of $f(X)$ and label all $f(X) > \tau$ as *positive* predicted cases and $f(X) \leq \tau$ as *negative* predicted cases. This results in a 2 by 2 table with predicted versus actual outcomes (sometimes called the ‘confusion table’), see Table 1. By varying τ between 0 and 1 we get a range of values for specificity and sensitivity. Plotting these in the receiver-operating-curve and calculating the area under this curve we get the popular AUC metric or c-statistic. Note that for calculating sensitivity we only need the *positive* cases ($Y = 1$), and for specificity we only need the *negative* cases ($Y = 0$). *Measures of discrimination depend on the distribution of the prediction (and thus the features) given the outcome.* This immediately implies that if we were to only change the ratio of positive and negative cases through some hypothetical intervention, the sensitivity and specificity will remain unchanged, and thus the resulting AUC. Therefore it is sometimes said that sensitivity and specificity are prevalence independent.

2.2 Calibration

Calibration measures how well predicted probabilities align with actual event rates. In words, assume we take a particular value for the predicted probability of the outcome, say $\alpha = 10\%$. Then if we gather all cases for which $f(X) = \alpha$, then the model is calibrated for that value of α when the fraction of positive outcomes in this subset is exactly α . A prediction model is perfectly calibrated when this holds for all unique values that $f(X)$ attains. For a formal definition, see Definition 1 in the Appendix 5. Unfortunately, measuring discrimination with a single metric is much harder than measuring discrimination for practical [2, 3] and theoretical reasons [4], a problem we will ignore. However, fundamentally, calibration measures the alignment between $f(X)$ and the probability of the outcome given X . *Measures of calibration are thus measures of the distribution of the outcome given the features (Y given X).*

3 A causal framework for predictive performance under changes in case-mix

Since discrimination depends on the distribution of the features given the outcome (X given Y) but calibration on the distribution of the outcome given the features (Y given X), we may expect metrics of discrimination and calibration to respond differently when changes occur in the marginal distribution of X or Y . In this section we first formalize the notion of a shift in *case-mix* and how this depends on whether a prediction is in the *causal* direction (future outcome given features) or the *anti-causal* direction (disease given symptoms). Then we will draw the connection between the two insights leading to our main result.

3.1 A shift in case-mix is a change in the marginal distribution of the cause variable

We define a shift in *case-mix* between different environments (e.g. GP versus hospital setting) as a change in the marginal distribution of the *cause* variable. When the prediction is in the *causal* direction, a shift in case-mix is a change in the marginal distribution of the features X , whereas when the prediction is in the *anti-causal* direction it is a shift in the marginal distribution of the outcome Y . Finally, the prediction problem could be neither causal or anti-causal, but *confounded* by another variable Z , in that case the shift is in the distribution of the confounder Z . See Figure 2 for directed acyclic graphs (DAGs) depicting these situations and Table 2 for an overview with examples. We give a formal definition in the Appendix 2.

Table 2: different prediction settings

	anti-causal	causal	confounded
shifted distribution	Y	X	Z
typical setting	diagnosis	prognosis	prognosis
example outcome	pneumonia	survival	lung cancer diagnosis
example features	temperature	age	yellow fingers
Figure	2a	2b	2c




Figure 2: directed acyclic graphs for 2-variable prediction problems with a shift in case-mix, meaning the environment variable only affects the marginal distribution of one variable but not a conditional distribution. The prediction is always made from X to Y . X denotes an arbitrary set of features with arbitrary feature types.

Each of the DAGs in Figure 2 encodes different conditional independencies. Specifically the DAG in the causal direction (Figure 2a) implies that Y is independent of E given X . This entails that the distribution $P(Y|X)$ is *transportable* across environments, so for different environments $E = 0, 1, \dots$, $P(Y|X, E = 0) = P(Y|X, E = 1) = P(Y|X)$, whereas in the anti-causal direction (Figure 2b) the distribution $P(Y|X)$ is transportable, meaning $P(X|Y, E) = P(X|Y)$ [1]. In the confounded DAG (Figure 2c) neither $P(Y|X)$ or $P(X|Y)$ is transportable.

In the DAGs in Figure 2 the environment variable influences the cause variable (X, Y or Z) but not the effect variable (Y or X). Why exclude arrows from the environment to the effect variable in the definition of a shift in case-mix? If there is an arrow from environment to the effect variable, neither $P(Y|X)$ or $P(X|Y)$ are transportable across environments so nothing can be said regarding the calibration and discrimination of a prediction model on an unseen environment based on data from the observed environments only. Also, based on concrete clinical settings it may be reasonable based on patient selection mechanisms to assume that *at least* the distribution of the cause variable

differs between environments, but maybe not the effect given the cause. See the Appendix 5.1 for several concrete medical examples where these assumptions may hold and where they are violated.

3.2 Main result: discrimination and calibration respond differently to changes in case-mix depending on the causal direction of the prediction

With the DAGs describing the different possible shifts in case-mix under consideration and the definitions of discrimination and calibration we can now state our main result, of which the formal versions are presented in the Appendix 5.

When predicting in the anti-causal direction (often with diagnosis predictions), a shift in case-mix across environments means a shift in the marginal distribution of the outcome, and discrimination remains stable but not calibration. Conversely, when predicting in the causal direction (often with prognosis predictions), a shift in case-mix across environments means a shift in marginal distribution of the features, and calibration remains stable, but not discrimination.

When f is perfectly calibrated on an environment, it will remain perfectly calibrated under shifts of the marginal distribution of the features (see Theorem 1 in the Appendix). Note that when f is not perfectly calibrated and this mis-calibration depends on X , in general the average calibration will also change when predicting in the causal direction.

An important implication of this result is that *discrimination or calibration may be preserved under changes in case-mix, but never both*

As a remark, we note that perfectly calibrated models obviously cannot be better calibrated in other environments, so any change in calibration necessarily implies a worsening of calibration. For discrimination, this is not automatically the case. In fact, models can show better discrimination in other environments when the distribution of outcome probabilities is less concentrated around 50%.

4 Simulation and empirical evaluation

4.1 Illustrative simulation

Our main result has important implications when interpreting changes in predictive performance across environments. To illustrate our result we now present a simulation study. Consider two prediction models, one is a prognostic model predicting in the causal direction, the other a diagnostic model predicting in the anti-causal direction. Denoting $\sigma^{-1}(p) = \log \frac{p}{1-p}$ as the logit function and \mathcal{N} the Gaussian distribution, the data-generating mechanisms are:

prognosis: $P_y \sim \text{Beta}(\alpha_e, \beta_e)$ $x = \sigma^{-1}(P_y)$ $y \sim \text{Bernoulli}(P_y)$	diagnosis: $y \sim \text{Bernoulli}(P_e)$ $x \sim \mathcal{N}(y, 1)$
---	--

We evaluate both models in three hypothetical environments: a screening environment (with low outcome prevalence), a GP setting (with intermediate prevalence) and a hospital setting (high prevalence). For the prognosis model, the marginal distribution of X depends on the environment through α_e, β_e , but not the distribution of Y given X . For the diagnosis model, the marginal distribution of Y depends on the environment through P_e , but not the distribution of X given Y . The different values for these parameters are given in Table 3

In Figure 3 we show the results of training a prediction model in the screening environment and evaluating it either in the same environment (=internal validation) or in a different environment (=external validation). For the prognostic model the calibration remains the same across environments (though some values of $P(Y = 1|X)$ become very rare because of the shift in distribution of X). The discrimination changes across environments. For the diagnostic model, the reverse is true.

By repeating this process for each of the three environments, each time training on one environment and evaluating on all environments for both the causal prediction model and the anti-causal model, we get in total six models, each evaluated three times. We measure discrimination with AUC for discrimination, and calibration error as the absolute difference between the predicted outcome




Figure 3: Overview figure of illustrative simulation experiment of a model trained on data from a screening environment, and evaluated on either the screening environment (=internal validation) or the GP environment or the hospital environment (=external validation), with increasing outcome probabilities. For models predicting in the *anti-causal* direction (e.g. diagnostic models), a shift in case-mix entails an intervention on the distribution of the predicted outcome, so discrimination remains the same but calibration changes. For models predicting in the *causal* direction (e.g. prognosis models), a shift in case-mix entails an intervention on the distribution of the features, so calibration remains the same but the discrimination changes. The discrimination facets are receiver-operating-curves with on the horizontal axis 1 minus specificity and on the vertical axis sensitivity. The calibration facets are calibration-curves with on the horizontal axis the predicted probability and on the vertical axis the actual probability.

task	parameter	screening	GP	hospital
prognosis	α	2	5	10
	β	20	10	20
diagnosis	p	0.2	1/3	0.5

Table 3: Values for different simulation parameters in three hypothetical environments. GP = general practitioner

probability and the actual outcome probability for each observation: $\frac{1}{N} \sum_i^N |P(Y = 1|X = x_i) - f(x_i)|$ (analogous to the Integrated Calibration Error defined in [5]). Plotting these 18 points on 6 lines in 2 dimensions leads to an interesting pattern, where the models predicting in the causal direction are easily discernible from those predicting in the anti-causal direction (Figure 4). In the Appendix 5.2 we provide visualizations of $P(Y|X)$ and $P(X|Y)$ for the different environments and tasks.




Figure 4: Combined results of the simulation experiment. Each model is connected by a line.

4.2 Empirical Study

As an empirical validation we re-used data from a published systematic review on prediction models in cardiovascular disease which included 2030 external validations of 1382 prediction models [6] and whose data is publicly available at <https://www.pacecpmregistry.org>. The review investigated changes in model performance when comparing the original publication with later external validation studies. The authors classified the prediction models as either ‘diagnostic’ or ‘prognostic’ (indicated by a follow-up time of less than 3 months, 3-6 months or more than 6 months). Selecting only prediction models with one or more validations and information on AUC in both the original study and validation study, and with information on model type (diagnostic versus prognostic), 1170 validation studies remained of 342 prediction models, 16 of which were validation studies of

11 diagnostic models. Comparing the AUC in the original study (AUC_0) with external validation studies (AUC_1), we calculated the relative difference in AUC as suggested by the authors:

$$\delta := \frac{(AUC_1 - 0.5) - (AUC_0 - 0.5)}{AUC_0 - 0.5}.$$

Our framework predicts that for diagnostic models that predict in the anti-causal direction, the AUC remains the same so $AUC_0 = AUC_1$, thus $\text{VAR}(\delta) = 0$, but not for prognosis models that predict in the causal direction. The studies in this systematic review are likely not perfectly *causal* or *anti-causal*, and because of varying finite sample sizes, variation in AUC will occur. Still we expect the variance of δ to be higher for prognosis models than for diagnosis models. In these data, using an F-test to compare the to variances this was indeed the case, with $\text{VAR}(\delta_{\text{diagnostic}}) = 0.019 \approx 0.122 * \text{VAR}(\delta_{\text{prognostic}})$, p-value < 0.001. Unfortunately the review provided no quantitative measures of calibration so a similar comparison of the variance of changes in calibration could not be made.

5 Discussion

We present a novel causal framework for understanding changes in prediction model performance under shifts in case-mix, by defining a shift in case-mix as a change in the marginal distribution of the cause variable. This leads to a new understanding of why in certain situations the discrimination of a model may be relatively stable when evaluated in a different setting, but not the calibration, and vice-versa.

Prior work noted that prediction models that are calibrated in multiple environments are provably free from anti-causal predictors [7]. Our focus is in the reverse direction: when to expect stable calibration across environments. Bareinboim and Pearl describe the data fusion problem and discuss when a causal effect may be transportable across settings and with what data [1]. Our work describes when certain specific functionals from the distributions are stable across settings, tailored to typical needs in the (medical) prediction model setting.

Our framework also provides a new perspective on the results of the study by Fehr et al [8]. They experimented with prediction models that contained either causal factors of the outcome (related to our *prognosis* models), anti-causal factors (related to our *diagnosis* models), or a combination of both. The performance of different prediction models was evaluated under different shifts in variables that were at the same time a direct cause of the outcome and a cause of other variables. Fehr et al find that for models predicting only with cause variables, the calibration is stable under interventions on only cause variables, as directly explained by our framework. When predicting with anti-causal factors, they observed that under interventions on the cause variables, the calibration degrades for models that are well calibrated on the training data. This setting is the closest to our *diagnostic* setting, though technically it is a mix of the anti-causal DAG 2b and confounding 2c.

Limitations are that the definition of a shift in case-mix is an abstraction and pure interventions on only either the features or the outcome may be unrealistic in practice. Many diagnostic prediction models may contain features that have a causal path to the diagnosis (e.g. age), or ‘risk factors’ for the disease that are not caused by the presence or absence of the disease. Systematic reviews of diagnostic models indeed show variation in sensitivity and specificity with variation in disease prevalence, a phenomenon also referred to as the *spectrum-effect* [9]. Still, when compared with prognostic models, diagnostic models had lower variability in discrimination in our empirical study. The current empirical validation was limited, and classifying diagnostic models as anti-causal and prognostic models as causal may be too crude. Also, no quantitative data on calibration were available to test whether calibration was more stable for prognostic models. Future empirical studies of externally evaluated prediction models will shed more light on how this theory pans out in practice. Mis-calibration may occur when variables not included in the model are also shifted between environments.

Still, the framework is useful as it explains prior observations and provides guidance to evaluators of prediction models on what to expect and how to explain certain changes in prediction model performance.

Acknowledgments

The author kindly acknowledges Anne de Hond, Oisín Ryan and Valentijn de Jong for fruitful discussions on this framework.

References

- [1] Elias Bareinboim and Judea Pearl. “Causal Inference and the Data-Fusion Problem”. In: *Proceedings of the National Academy of Sciences* 113.27 (July 5, 2016), pp. 7345–7352. DOI: 10.1073/pnas.1510507113. URL: <https://www.pnas.org/doi/abs/10.1073/pnas.1510507113> (visited on 11/03/2023).
- [2] Ben Van Calster et al. “Calibration: The Achilles Heel of Predictive Analytics”. In: *BMC Medicine* 17.1 (Dec. 16, 2019), p. 230. ISSN: 1741-7015. DOI: 10.1186/s12916-019-1466-7. URL: <https://doi.org/10.1186/s12916-019-1466-7> (visited on 02/05/2024).
- [3] Ben Van Calster et al. “A Calibration Hierarchy for Risk Models Was Defined: From Utopia to Empirical Data”. In: *Journal of Clinical Epidemiology* 74 (June 1, 2016), pp. 167–176. ISSN: 0895-4356. DOI: 10.1016/j.jclinepi.2015.12.005. URL: <https://www.sciencedirect.com/science/article/pii/S0895435615005818> (visited on 11/15/2023).
- [4] Jarosław Błasiok et al. *A Unifying Theory of Distance from Calibration*. Mar. 31, 2023. DOI: 10.48550/arXiv.2211.16886. arXiv: 2211.16886 [cs]. URL: <http://arxiv.org/abs/2211.16886> (visited on 02/05/2024). Pre-published.
- [5] Peter C. Austin and Ewout W. Steyerberg. “The Integrated Calibration Index (ICI) and Related Metrics for Quantifying the Calibration of Logistic Regression Models”. In: *Statistics in Medicine* 38.21 (Sept. 20, 2019), pp. 4051–4065. ISSN: 0277-6715. DOI: 10.1002/sim.8281. pmid: 31270850. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6771733/> (visited on 06/11/2024).
- [6] Benjamin S. Wessler et al. “External Validations of Cardiovascular Clinical Prediction Models: A Large-Scale Review of the Literature”. In: *Circulation: Cardiovascular Quality and Outcomes* 14.8 (Aug. 2021), e007858. DOI: 10.1161/CIRCOUTCOMES.121.007858. URL: <https://www.ahajournals.org/doi/full/10.1161/CIRCOUTCOMES.121.007858> (visited on 08/23/2024).
- [7] Yoav Wald et al. “On Calibration and Out-of-Domain Generalization”. In: *Advances in Neural Information Processing Systems*. Vol. 34. Curran Associates, Inc., 2021, pp. 2215–2227. URL: https://papers.nips.cc/paper_files/paper/2021/hash/118bd558033a1016fcc82560c65cca5f-Abstract.html (visited on 08/28/2023).
- [8] Jana Fehr et al. “Assessing the Transportability of Clinical Prediction Models for Cognitive Impairment Using Causal Models”. In: *BMC Medical Research Methodology* 23.1 (Aug. 19, 2023), p. 187. ISSN: 1471-2288. DOI: 10.1186/s12874-023-02003-6. URL: <https://bmcmmedresmethodol.biomedcentral.com/articles/10.1186/s12874-023-02003-6> (visited on 10/03/2023).
- [9] Mariska M.G. Leeflang et al. “Variation of a Test’s Sensitivity and Specificity with Disease Prevalence”. In: *Canadian Medical Association Journal* 185.11 (Aug. 6, 2013), E537–E544. ISSN: 0820-3946, 1488-2329. DOI: 10.1503/cmaj.121286. URL: <http://www.cmaj.ca/lookup/doi/10.1503/cmaj.121286> (visited on 08/30/2024).

Appendix

Definitions

Definition 1 (calibration). Let $P(X, Y)$ be a joint distribution over feature X and binary outcome Y , and $f : X \rightarrow [0, 1]$ a prediction model. f is perfectly calibrated with respect to $P(X, Y)$ if, for all $\alpha \in [0, 1]$ in the range of f , $\mathbb{E}_{X, Y \sim P(X, Y)}[Y | f(X) = \alpha] = \alpha$.

Definition 2 (case mix). Let Z, X, Y be random variables and E an environment variable. Assume one of the three following causal directed acyclic graphs labeled causal, anti-causal and fork (shown also in Figure 2):

1. causal: $E \rightarrow X \rightarrow Y$
2. anti-causal: $E \rightarrow Y \rightarrow X$
3. fork: $E \rightarrow Z \rightarrow X; Z \rightarrow Y$

Let $P_E(\cdot)$ denote the distribution of variable \cdot in environment E . A shift in case-mix across environments $e, e' \in \mathcal{E}$ is a shift in the distribution of the direct child of E in the DAG, meaning a shift in:

1. $P_E(X)$ when DAG = causal
2. $P_E(Y)$ when DAG = anti-causal
3. $P_E(Z)$ when DAG = fork

Remark 1 (conditional independencies). The causal DAGs enumerated in Definition 2 imply the following conditional independencies regarding random variables X, Y :

	causal	anti-causal	fork
$P_E(X)$ vs $P(X)$	=	≠	≠
$P_E(Y)$ vs $P(Y)$	≠	=	≠
$P_E(Y X)$ vs $P(Y X)$	=	≠	≠
$P_E(X Y)$ vs $P(X Y)$	≠	=	≠

Theorems

We now describe our main result.

Theorem 1 (perfectly calibrated models remain perfectly calibrated under marginal shifts in X). Given prediction model $f : \mathcal{X} \rightarrow [0, 1]$ and environment $E \in \{\text{train}, \text{test}\}$. Assume $P_E(Y, X) = P_E(X)P(Y|X)$ (shift in $P(X)$ but not $P(Y|X)$ between environments). Furthermore assume $\text{supp}(P_{\text{test}}(X)) \subseteq \text{supp}(P_{\text{train}}(X))$ (marginal distribution of X in training data is non-zero wherever the test distribution is non-zero). Define the miscalibration of f under P_E for value $X = x$ as:

$$\xi(x) := |f(x) - P(Y|X = x)| \quad (1)$$

Then the integrated calibration index (ICI) on distribution P_E is [5]:

$$ICI_E = \mathbb{E}_{X \sim P_E(X)} \xi(x) \quad (2)$$

Theorem statement: a perfectly calibrated model (i.e. $\xi(x) = 0, \forall x \in \text{supp}(P_{\text{train}}(X))$) remains perfectly calibrated under marginal shifts in X , i.e.

$$ICI_{\text{train}} = ICI_{\text{test}} = 0 \quad (3)$$

proof of theorem 1. By assumption we have $ICI_{\text{train}} = 0$ By definition of ICI, denoting $\phi_{\text{test}}(x)$ the density of X in distribution P_{test} .

$$ICI_{\text{test}} = \mathbb{E}_{X \sim P_{\text{test}}(X)} \xi(x) \quad (4)$$

$$= \int_{\text{supp}P_{\text{test}}(X)} \phi_{\text{test}}(x)\xi(x) \quad (5)$$

$$= \int_{\text{supp}P_{\text{test}}(X)} \phi_{\text{test}}(x) * 0 \quad (6)$$

$$= 0 \int_{\text{supp}P_{\text{test}}(X)} \phi_{\text{test}}(x) \quad (7)$$

$$= 0 * 1 \quad (8)$$

$$= 0 \quad (9)$$

□

Theorem 2 (discrimination is constant under marginal shifts in Y). *Given prediction model $f : \mathcal{X} \rightarrow [0, 1]$ and environment $E \in \{\text{train}, \text{test}\}$. Assume $P_E(Y, X) = P_E(Y)P(X|Y)$ (marginal shift of Y but not $X|Y$). Furthermore assume $0 < P_E(Y = 1) < 1$ (marginal distribution of Y in both distributions is non-deterministic). Then for all thresholds τ :*

$$\text{sens}_{\text{test}}(\tau) = \text{sens}_{\text{train}}(\tau) \quad (10)$$

$$\text{spec}_{\text{test}}(\tau) = \text{spec}_{\text{train}}(\tau) \quad (11)$$

And also

$$\text{AUC}_{\text{train}} = \text{AUC}_{\text{test}} \quad (12)$$

proof of theorem 2. Theorem 2 follows directly from the fact that sensitivity: $P(f(X) > \tau | Y = 1)$, specificity: $P(f(X) \leq \tau | Y = 0)$ and $P_{\text{train}}(f(X) | Y = y) = P_{\text{test}}(f(X) | Y = y)$. \square

5.1 Examples

Clinical examples of when the definition of a shift in case-mix as in Definition 2 may apply across different environments.

5.1.1 Examples in the causal direction

Prediction of the occurrence of a cardiovascular event in the coming 10 years based on age and the presence of diabetes at baseline.

1. train environment: general practitioner
2. test environment: a diabetes out-patient clinic

5.1.2 Examples in the anti-direction

Example 1: prediction of the presence of a stroke based on computed tomography imaging of the brain:

1. train environment: secondary care hospital
2. test environment: stroke center where patients are referred when they have stroke symptoms

Example 2: Diagnosing sexually transmittable disease (see Figure 5).




Figure 5: Example setting of diagnosing a sexually transmittable disease (STD, = Y) with a blood test (= X) in either general public setting (5a) or in a HIV-positive clinic (5b). Patients with previous STDs such as HIV (Y_0) have a higher risk of future STDs, summarized with the arrow from Y_0 to Y . $Y_0 = 1$ is a selection criterion for the HIV-clinic, meaning that only patients with a prior STD get seen at the HIV-clinic. Treating Y_0 as not observed (thus marginalizing it out) results in the DAG in 5b

5.2 Additional figures of simulation study




Figure 6: Marginal distribution $P(Y = 1)$ in different environments, given by beta-distributions with parameters listed in Table 3.

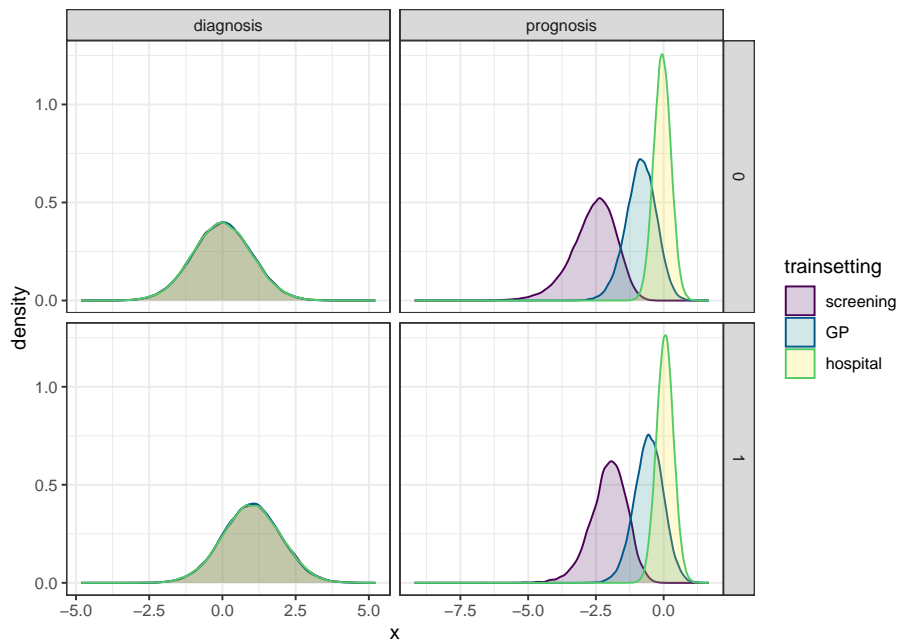


Figure 7: Conditional density of $P(X|Y = y)$ in the diagnosis or the prognosis simulation setting




Figure 8: Conditional distribution of $P(Y = 1|X = x)$ in the diagnosis or the prognosis simulation setting.