

---

# Calibration tests in multi-class classification: A unifying framework

---

**David Widmann**

Department of Information Technology  
Uppsala University, Sweden  
david.widmann@it.uu.se

**Fredrik Lindsten**

Division of Statistics and Machine Learning  
Linköping University, Sweden  
fredrik.lindsten@liu.se

**Dave Zachariah**

Department of Information Technology  
Uppsala University, Sweden  
dave.zachariah@it.uu.se

## Abstract

In safety-critical applications a probabilistic model is usually required to be calibrated, i.e., to capture the uncertainty of its predictions accurately. In multi-class classification, calibration of the most confident predictions only is often not sufficient. We propose and study calibration measures for multi-class classification that generalize existing measures such as the expected calibration error, the maximum calibration error, and the maximum mean calibration error. We propose and evaluate empirically different consistent and unbiased estimators for a specific class of measures based on matrix-valued kernels. Importantly, these estimators can be interpreted as test statistics associated with well-defined bounds and approximations of the p-value under the null hypothesis that the model is calibrated, significantly improving the interpretability of calibration measures, which otherwise lack any meaningful unit or scale.

## 1 Introduction

Consider the problem of analyzing microscopic images of tissue samples and reporting a tumour grade, i.e., a score that indicates whether cancer cells are well-differentiated or not, affecting both prognosis and treatment of patients. Since for some pathological images not even experienced pathologists might all agree on one classification, this task contains an inherent component of uncertainty. This type of uncertainty that can not be removed by increasing the size of the training data set is typically called aleatoric uncertainty (Kiureghian and Ditlevsen, 2009). Unfortunately, even if the ideal model is among the class of models we consider, with a finite training data set we will never obtain the ideal model but we can only hope to learn a model that is, in some sense, close to it. Worse still, our model might not even be close to the ideal model if the model class is too restrictive or the number of training data is small—which is not unlikely given the fact that annotating pathological images is expensive. Thus ideally our model should be able to express not only aleatoric uncertainty but also the uncertainty about the model itself. In contrast to aleatoric uncertainty this so-called epistemic uncertainty can be reduced by additional training data.

Dealing with these different types of uncertainty is one of the major problems in machine learning. The application of our model in clinical practice demands “meaningful” uncertainties to avoid doing harm to patients. Being too certain about high tumour grades might cause harm due to unneeded aggressive therapies and overly pessimistic prognoses, whereas being too certain about low tumour grades might result in insufficient therapies. “Proper” uncertainty estimates are also crucial if the

model is supervised by a pathologist that takes over if the uncertainty reported by the model is too high. False but highly certain gradings might incorrectly keep the pathologist out of the loop, and on the other hand too uncertain gradings might demand unneeded and costly human intervention.

Probability theory provides a solid framework for dealing with uncertainties. Instead of assigning exactly one grade to each pathological image, so-called probabilistic models report subjective probabilities, sometimes also called confidence scores, of the tumour grades for each image. The model can be evaluated by comparing these subjective probabilities to the ground truth.

One desired property of such a probabilistic model is sharpness (or high accuracy), i.e., if possible, the model should assign the highest probability to the true tumour grade (which maybe can not be inferred from the image at hand but only by other means such as an additional immunohistochemical staining). However, to be able to trust the predictions the probabilities should be calibrated (or reliable) as well (DeGroot and Fienberg, 1983; Murphy and Winkler, 1977). This property requires the subjective probabilities to match the relative empirical frequencies: intuitively, if we could observe a long run of predictions  $(0.5, 0.1, 0.1, 0.3)$  for tumour grades 1, 2, 3, and 4, the empirical frequencies of the true tumour grades should be  $(0.5, 0.1, 0.1, 0.3)$ . Note that accuracy and calibration are two complementary properties: a model with over-confident predictions can be highly accurate but miscalibrated, whereas a model that always reports the overall proportion of patients of each tumour grade in the considered population is calibrated but highly inaccurate.

Research of calibration in statistics and machine learning literature has been focused mainly on binary classification problems or the most confident predictions: common calibration measures such as the expected calibration error (ECE) (Naeini et al., 2015), the maximum calibration error (MCE) (Naeini et al., 2015), and the kernel-based maximum mean calibration error (MMCE) (Kumar et al., 2018), and reliability diagrams (Murphy and Winkler, 1977) have been developed for binary classification. This is insufficient since many recent applications of machine learning involve multiple classes. Furthermore, the crucial finding of Guo et al. (2017) that many modern deep neural networks are miscalibrated is also based only on the most confident prediction.

Recently Vaicenavicius et al. (2019) suggested that this analysis might be too reduced for many realistic scenarios. In our example, a prediction of  $(0.5, 0.3, 0.1, 0.1)$  is fundamentally different from a prediction of  $(0.5, 0.1, 0.1, 0.3)$ , since according to the model in the first case it is only half as likely that a tumour is of grade 3 or 4, and hence the subjective probability of missing out on a more aggressive therapy is smaller. However, commonly in the study of calibration all predictions with a highest reported confidence score of 0.5 are grouped together and a calibrated model has only to be correct about the most confident tumour grade in 50% of the cases, regardless of the other predictions. Although the ECE can be generalized to multi-class classification, its applicability seems to be limited since its histogram-regression based estimator requires partitioning of the potentially high-dimensional probability simplex and is asymptotically inconsistent in many cases (Vaicenavicius et al., 2019).

## 2 Our contribution

In this work, we propose and study a general framework of calibration measures for multi-class classification. We show that this framework encompasses common calibration measures for binary classification such as the expected calibration error (ECE), the maximum calibration error (MCE), and the maximum mean calibration error (MMCE) by Kumar et al. (2018). In more detail we study a class of measures based on vector-valued reproducing kernel Hilbert spaces, for which we derive consistent and unbiased estimators. The statistical properties of the proposed estimators are not only theoretically appealing, but also of high practical value, since they allow us to address two main problems in calibration evaluation.

As discussed by Vaicenavicius et al. (2019), all calibration error estimates are inherently random, and comparing competing models based on these estimates without taking the randomness into account can be very misleading, in particular when the estimators are biased (which, for instance, is the case for the commonly used histogram-regression based estimator of the ECE). Even more fundamentally, all commonly used calibration measures lack a meaningful unit or scale and are therefore not interpretable as such (regardless of any finite sample issues).

The consistency and unbiasedness of the proposed estimators facilitate comparisons between competing models, and allow us to derive multiple statistical tests for calibration that exploit these properties. Moreover, by viewing the estimators as *calibration test statistics*, with well-defined bounds and approximations of the corresponding p-value, we give them an interpretable meaning.

We evaluate the proposed estimators and statistical tests empirically and compare them with existing methods. To facilitate multi-class calibration evaluation we provide the Julia packages `ConsistencyResampling.jl` (Widmann, 2019c), `CalibrationErrors.jl` (Widmann, 2019a), and `CalibrationTests.jl` (Widmann, 2019b) for consistency resampling, calibration error estimation, and calibration tests, respectively.

### 3 Background

We start by shortly summarizing the most relevant definitions and concepts. Due to space constraints and to improve the readability of our paper, we do not provide any proofs in the main text but only refer to the results in the supplementary material, which is intended as a reference for mathematically precise statements and proofs.

#### 3.1 Probabilistic setting

Let  $(X, Y)$  be a pair of random variables with  $X$  and  $Y$  representing inputs (features) and outputs, respectively. We focus on classification problems and hence without loss of generality we may assume that the outputs consist of the  $m$  classes  $1, \dots, m$ .

Let  $\Delta^m$  denote the  $(m - 1)$ -dimensional probability simplex  $\Delta^m := \{z \in \mathbb{R}_{\geq 0}^m : \|z\|_1 = 1\}$ . Then a *probabilistic model*  $g$  is a function that for every input  $x$  outputs a prediction  $g(x) \in \Delta^m$  that models the distribution

$$(\mathbb{P}[Y = 1 | X = x], \dots, \mathbb{P}[Y = m | X = x]) \in \Delta^m$$

of class  $Y$  given input  $X = x$ .

#### 3.2 Calibration

##### 3.2.1 Common notion

The common notion of calibration, as, e.g., used by Guo et al. (2017), considers only the most confident predictions  $\max_y g_y(x)$  of a model  $g$ . According to this definition, a model is calibrated if

$$\mathbb{P}[Y = \arg \max_y g_y(X) | \max_y g_y(X)] = \max_y g_y(X) \quad (1)$$

holds almost always. Thus a model that is calibrated according to Eq. (1) ensures that we can *partly trust* the uncertainties reported by its predictions. As an example, for a prediction of  $(0.4, 0.3, 0.3)$  the model would only guarantee that in the long run inputs that yield a most confident prediction of 40% are in the corresponding class 40% of the time.<sup>1</sup>

##### 3.2.2 Strong notion

According to the more general calibration definition of Bröcker (2009); Vaicenavicius et al. (2019), a probabilistic model  $g$  is calibrated if for almost all inputs  $x$  the prediction  $g(x)$  is equal to the distribution of class  $Y$  given prediction  $g(X) = g(x)$ . More formally, a calibrated model satisfies

$$\mathbb{P}[Y = y | g(X)] = g_y(X) \quad (2)$$

almost always for all classes  $y \in \{1, \dots, m\}$ . As Vaicenavicius et al. (2019) showed, for multi-class classification this formulation is stronger than the definition of Zadrozny and Elkan (2002) that only demands calibrated marginal probabilities. Thus we can *fully trust* the uncertainties reported by the predictions of a model that is calibrated according to Eq. (2). The prediction  $(0.4, 0.3, 0.3)$  would actually imply that the class distribution of the inputs that yield this prediction is  $(0.4, 0.3, 0.3)$ . To emphasize the difference to the definition in Eq. (1), we call calibration according to Eq. (2) *calibration in the strong sense* or *strong calibration*.

<sup>1</sup>This notion of calibration does not consider for which class the most confident prediction was obtained.

To simplify our notation, we rewrite Eq. (2) in vectorized form. Equivalent to the definition above, a model  $g$  is calibrated in the strong sense if

$$r(g(X)) - g(X) = 0 \quad (3)$$

holds almost always, where

$$r(\xi) := (\mathbb{P}[Y = 1 | g(X) = \xi], \dots, \mathbb{P}[Y = m | g(X) = \xi])$$

is the distribution of class  $Y$  given prediction  $g(X) = \xi$ .

The calibration of certain aspects of a model, such as the five largest predictions or groups of classes, can be investigated by studying the strong calibration of models induced by so-called calibration lenses. For more details about evaluation and visualization of strong calibration we refer to Vaicenavicius et al. (2019).

### 3.3 Matrix-valued kernels

The miscalibration measure that we propose in this work is based on matrix-valued kernels  $k: \Delta^m \times \Delta^m \rightarrow \mathbb{R}^{m \times m}$ . Matrix-valued kernels can be defined in a similar way as the more common real-valued kernels, which can be characterized as symmetric positive definite functions (Berlinet and Thomas-Agnan, 2004, Lemma 4).

**Definition 1 (Micchelli and Pontil (2005, Definition 2); Caponnetto et al. (2008, Definition 1)).** We call a function  $k: \Delta^m \times \Delta^m \rightarrow \mathbb{R}^{m \times m}$  a *matrix-valued kernel* if for all  $s, t \in \Delta^m$   $k(s, t) = k(t, s)^\top$  and it is positive semi-definite, i.e., if for all  $n \in \mathbb{N}$ ,  $t_1, \dots, t_n \in \Delta^m$ , and  $u_1, \dots, u_n \in \mathbb{R}^m$

$$\sum_{i,j=1}^n u_i^\top k(t_i, t_j) u_j \geq 0.$$

There exists a one-to-one mapping of such kernels and reproducing kernel Hilbert spaces (RKHSs) of vector-valued functions  $f: \Delta^m \rightarrow \mathbb{R}^m$ . We provide a short summary of RKHSs of vector-valued functions on the probability simplex in Appendix D. A more detailed general introduction to RKHSs of vector-valued functions can be found in the publications by Caponnetto et al. (2008); Carmeli et al. (2010); Micchelli and Pontil (2005).

Similar to the scalar case, matrix-valued kernels can be constructed from other matrix-valued kernels and even from real-valued kernels. Very simple matrix-valued kernels are kernels of the form  $k(s, t) = \tilde{k}(s, t) \mathbf{I}_m$ , where  $\tilde{k}$  is a scalar-valued kernel, such as the Gaussian or Laplacian kernel, and  $\mathbf{I}_m$  is the identity matrix. As Example D.1 shows, this construction can be generalized by, e.g., replacing the identity matrix with an arbitrary positive semi-definite matrix.

An important class of kernels are so-called universal kernels. Loosely speaking, a kernel is called universal if its RKHS is a dense subset of the space of continuous functions, i.e., if in the neighbourhood of every continuous function we can find a function in the RKHS. Prominent real-valued kernels on the probability simplex such as the Gaussian and the Laplacian kernel are universal, and can be used to construct universal matrix-valued kernels of the form in Example D.1, as Lemma D.3 shows.

## 4 Unification of calibration measures

In this section we introduce a general measure of strong calibration and show its relation to other existing measures.

### 4.1 Calibration error

In the analysis of strong calibration the discrepancy in the left-hand side of Eq. (3) lends itself naturally to the following calibration measure.

**Definition 2.** Let  $\mathcal{F}$  be a non-empty space of functions  $f: \Delta^m \rightarrow \mathbb{R}^m$ . Then the calibration error (CE) of model  $g$  with respect to class  $\mathcal{F}$  is

$$\text{CE}[\mathcal{F}, g] := \sup_{f \in \mathcal{F}} \mathbb{E} [(r(g(X)) - g(X))^\top f(g(X))].$$

A trivial consequence of the design of the CE is that the measure is zero for every model that is calibrated in the strong sense, regardless of the choice of  $\mathcal{F}$ . The converse statement is not true in general. As we show in Theorem C.2, the class of continuous functions is a space for which the CE is zero if and only if model  $g$  is strongly calibrated, and hence allows to characterize calibrated models. However, since this space is extremely large, for every model the CE is either 0 or  $\infty$ .<sup>2</sup> Thus a measure based on this space does not allow us to compare miscalibrated models and hence is rather impractical.

## 4.2 Kernel calibration error

Due to the correspondence between kernels and RKHSs we can define the following kernel measure.

**Definition 3.** Let  $k$  be a matrix-valued kernel as in Definition 1. Then we define the kernel calibration error (KCE) with respect to kernel  $k$  as  $\text{KCE}[k, g] := \text{CE}[\mathcal{F}, g]$ , where  $\mathcal{F}$  is the unit ball in the RKHS corresponding to kernel  $k$ .

As mentioned above, a RKHS with a universal kernel is a dense subset of the space of continuous functions. Hence these kernels yield a function space that is still large enough for identifying strongly calibrated models.

**Theorem 1 (cf. Theorem C.1).** Let  $k$  be a matrix-valued kernel as in Definition 1, and assume that  $k$  is universal. Then  $\text{KCE}[k, g] = 0$  if and only if model  $g$  is calibrated in the strong sense.

From the supremum-based Definition 2 it might not be obvious how the KCE can be evaluated. Fortunately, there exists an equivalent kernel-based formulation.

**Lemma 1 (cf. Lemma E.2).** Let  $k$  be a matrix-valued kernel as in Definition 1. If  $\mathbb{E}[||k(g(X), g(X))||] < \infty$ , then

$$\text{KCE}[k, g] = \left( \mathbb{E} [(e_Y - g(X))^\top k(g(X), g(X')) (e_{Y'} - g(X'))] \right)^{1/2}, \quad (4)$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$  and  $e_i$  denotes the  $i$ th unit vector.

## 4.3 Expected calibration error

The most common measure of calibration error is the expected calibration error (ECE). Typically it is used for quantifying calibration in a binary classification setting but it generalizes to strong calibration in a straightforward way. Let  $d: \Delta^m \times \Delta^m \rightarrow \mathbb{R}_{\geq 0}$  be a distance measure on the probability simplex. Then the expected calibration error of a model  $g$  with respect to  $d$  is defined as

$$\text{ECE}[d, g] = \mathbb{E}[d(r(g(X)), g(X))]. \quad (5)$$

If  $d(p, q) = 0 \Leftrightarrow p = q$ , as it is the case for standard choices of  $d$  such as the total variation distance or the (squared) Euclidean distance, then  $\text{ECE}[d, g]$  is zero if and only if  $g$  is strongly calibrated as per Eq. (3).

The ECE with respect to the cityblock distance, the total variation distance, or the squared Euclidean distance, are special cases of the calibration error CE, as we show in Lemma I.1.

## 4.4 Maximum mean calibration error

Kumar et al. (2018) proposed a kernel-based calibration measure, the so-called maximum mean calibration error (MMCE), for training (better) calibrated neural networks. In contrast to their work, in our publication we do not discuss how to obtain calibrated models but focus on the evaluation of calibration and on calibration tests. Moreover, the MMCE applies only to a binary classification setting whereas our measure quantifies strong calibration and hence is more generally applicable. In fact, as we show in Example I.1, the MMCE is a special case of the KCE.

<sup>2</sup>Assume  $\text{CE}[\mathcal{F}, g] < \infty$  and let  $f_1, f_2, \dots$  be a sequence of continuous functions with  $\text{CE}[\mathcal{F}, g] = \lim_{n \rightarrow \infty} \mathbb{E} [(r(g(X)) - g(X))^\top f_n(g(X))]$ . From Remark C.2 we know that  $\text{CE}[\mathcal{F}, g] \geq 0$ . Moreover,  $\tilde{f}_n := 2f_n$  are continuous functions with  $2\text{CE}[\mathcal{F}, g] = \lim_{n \rightarrow \infty} \mathbb{E} [(r(g(X)) - g(X))^\top \tilde{f}_n(g(X))] \leq \sup_{f \in \mathcal{F}} \mathbb{E} [(r(g(X)) - g(X))^\top f(g(X))] = \text{CE}[\mathcal{F}, g]$ . Thus  $\text{CE}[\mathcal{F}, g] = 0$ .

## 5 Calibration error estimators

Consider the task of estimating the calibration error of model  $g$  using a validation set  $\mathcal{D} = \{(X_i, Y_i)\}_{i=1}^n$  of  $n$  i.i.d. random pairs of inputs and labels that are distributed according to  $(X, Y)$ .

From the expression for the ECE in Eq. (5), the natural (and, indeed, standard) approach for estimating the ECE is as the sample average of the distance  $d$  between the predictions  $g(X)$  and the calibration function  $r(g(X))$ . However, this is problematic since the calibration function is not readily available and needs to be estimated as well. Typically, this is addressed using histogram-regression, see, e.g., Guo et al. (2017); Naeini et al. (2015); Vaicenavicius et al. (2019), which unfortunately leads to inconsistent and biased estimators in many cases (Vaicenavicius et al., 2019) and can scale poorly to large  $m$ . In contrast, for the KCE in Eq. (4) there is no explicit dependence on  $r$ , which enables us to derive multiple consistent and also unbiased estimators.

Let  $k$  be a matrix-valued kernel as in Definition 1 with  $\mathbb{E}[\|k(g(X), g(X))\|] < \infty$ , and define for  $1 \leq i, j \leq n$

$$h_{i,j} := (e_{Y_i} - g(X_i))^\top k(g(X_i), g(X_j)) (e_{Y_j} - g(X_j)).$$

Then the estimators listed in Table 1 are consistent estimators of the squared kernel calibration error  $\text{SKCE}[k, g] := \text{KCE}^2[k, g]$  (see Lemmas F.1 to F.3). The subscript letters “q” and “l” refer to the quadratic and linear computational complexity of the unbiased estimators, respectively.

Table 1: Three consistent estimators of the SKCE.

Notation	Definition	Properties	Complexity
$\widehat{\text{SKCE}}_{\text{b}}$	$n^{-2} \sum_{i,j=1}^n h_{i,j}$	biased	$O(n^2)$
$\widehat{\text{SKCE}}_{\text{uq}}$	$\binom{n}{2}^{-1} \sum_{1 \leq i < j \leq n} h_{i,j}$	unbiased	$O(n^2)$
$\widehat{\text{SKCE}}_{\text{ul}}$	$\lfloor n/2 \rfloor^{-1} \sum_{i=1}^{\lfloor n/2 \rfloor} h_{2i-1, 2i}$	unbiased	$O(n)$

## 6 Calibration tests

In general, calibration errors do not have a meaningful unit or scale. This renders it difficult to interpret an estimated non-zero error and to compare different models. However, by viewing the estimates as test statistics with respect to Eq. (3), they obtain an interpretable probabilistic meaning.

Similar to the derivation of the two-sample tests by Gretton et al. (2012), we can use the consistency and unbiasedness of the estimators of the SKCE presented above to obtain bounds and approximations of the p-value for the null hypothesis  $H_0$  that the model is calibrated, i.e., for the probability that the estimator is greater than or equal to the observed calibration error estimate, if the model is calibrated. These bounds and approximations do not only allow us to perform hypothesis testing of the null hypothesis  $H_0$ , but they also enable us to transfer unintuitive calibration error estimates to an intuitive and interpretable probabilistic setting.

As we show in Theorems H.2 to H.4, we can obtain so-called distribution-free bounds without making any assumptions about the distribution of  $(X, Y)$  or the model  $g$ . A downside of these uniform bounds is that usually they provide only a loose bound of the p-value.

**Lemma 2 (Distribution-free bounds (see Theorems H.2 to H.4)).** *Let  $k$  be a matrix-valued kernel as in Definition 1, and assume that  $K_{p,q} := \sup_{s,t \in \Delta^m} \|k(s, t)\|_{p,q} < \infty$  for some  $1 \leq p, q \leq \infty$ .<sup>3</sup> Let  $t > 0$  and  $B_{p,q} := 2^{1+1/p-1/q} K_{p,q}$ , then for the biased estimator we can bound*

$$\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{b}} \geq t \mid H_0 \right] \leq \exp \left( - \frac{1}{2} \left( \max \left\{ 0, \sqrt{nt/B_{p,q}} - 1 \right\} \right)^2 \right),$$

and for either of the unbiased estimators  $T \in \{\widehat{\text{SKCE}}_{\text{uq}}, \widehat{\text{SKCE}}_{\text{ul}}\}$ , we can bound

$$\mathbb{P} \left[ T \geq t \mid H_0 \right] \leq \exp \left( - \frac{\lfloor n/2 \rfloor t^2}{2B_{p,q}^2} \right).$$

<sup>3</sup>For a matrix  $A$  we denote by  $\|A\|_{p,q}$  the induced matrix norm  $\sup_{x \neq 0} \|Ax\|_q / \|x\|_p$ .

Asymptotic bounds exploit the asymptotic distribution of the test statistics under the null hypothesis, as the number of validation data points goes to infinity. The central limit theorem implies that the linear estimator is asymptotically normally distributed.

**Lemma 3 (Asymptotic distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  (see Corollary G.1)).** *Let  $k$  be a matrix-valued kernel as in Definition 1, and assume that  $\mathbb{E}[\|k(g(X), g(X))\|] < \infty$ . If  $\text{Var}[h_{i,j}] < \infty$ , then*

$$\mathbb{P}\left[\sqrt{\lfloor n/2 \rfloor} \widehat{\text{SKCE}}_{\text{ul}} \geq t \mid H_0\right] \rightarrow 1 - \Phi\left(\frac{t}{\hat{\sigma}}\right) \quad \text{as } n \rightarrow \infty,$$

where  $\hat{\sigma}$  is the sample standard deviation of  $h_{2i-1,2i}$  ( $i = 1, \dots, \lfloor n/2 \rfloor$ ) and  $\Phi$  is the cumulative distribution function of the standard normal distribution.

In Theorem G.2 we derive a theoretical expression of the asymptotic distribution of  $n \widehat{\text{SKCE}}_{\text{uq}}$ , under the assumption of strong calibration. This limit distribution can be approximated, e.g., by bootstrapping (Arcones and Giné, 1992) or Pearson curve fitting, as discussed by Gretton et al. (2012).

## 7 Experiments

We conduct experiments to confirm the derived theoretical properties of the proposed calibration error estimators empirically and to compare them with a standard histogram-regression based estimator of the ECE, denoted by  $\widehat{\text{ECE}}$ .<sup>4</sup>

We construct synthetic data sets  $\{(g(X_i), Y_i)\}_{i=1}^{250}$  of 250 labeled predictions for  $m = 10$  classes from three generative models. For each model we first sample predictions  $g(X_i) \sim \text{Dir}(0.1, \dots, 0.1)$ , and then simulate corresponding labels  $Y_i$  conditionally on  $g(X_i)$  from

$$\mathbf{M1}: \text{Cat}(g(X_i)), \quad \mathbf{M2}: 0.5 \text{Cat}(g(X_i)) + 0.5 \text{Cat}(1, 0, \dots, 0), \quad \mathbf{M3}: \text{Cat}(0.1, \dots, 0.1),$$

where **M1** gives a calibrated model, and **M2** and **M3** are uncalibrated. In Appendix J.2 we investigate the theoretical properties of these models in more detail.

For simplicity, we use the matrix-valued kernel  $k(x, y) = \exp(-\|x - y\|/\nu) \mathbf{I}_{10}$ , where the kernel bandwidth  $\nu > 0$  is chosen by the median heuristic. The total variation distance is a common distance measure of probability distributions and the standard distance measure for the ECE (Bröcker and Smith, 2007; Guo et al., 2017; Vaicenavicius et al., 2019), and hence it is chosen as the distance measure for all studied calibration errors.

### 7.1 Calibration error estimates

In Fig. 1 we show the distribution of  $\widehat{\text{ECE}}$  and of the three proposed estimators of the SKCE, evaluated on  $10^4$  randomly sampled data sets from each of the three models. The true calibration error of these models, indicated by a dashed line, is calculated theoretically for the ECE (see Appendix J.2.1) and empirically for the SKCE using the sample mean of all unbiased estimates of  $\widehat{\text{SKCE}}_{\text{uq}}$ .

We see that the standard estimator of the ECE exhibits both negative and positive bias, whereas  $\widehat{\text{SKCE}}_{\text{b}}$  is theoretically guaranteed to be biased upwards. The results also confirm the unbiasedness of  $\widehat{\text{SKCE}}_{\text{ul}}$ .

### 7.2 Calibration tests

We repeatedly compute the bounds and approximations of the p-value for the calibration error estimators that were derived in Section 6 on  $10^4$  randomly sampled data sets from each of the three models. More concretely, we evaluate the distribution-free bounds for  $\widehat{\text{SKCE}}_{\text{b}}(\mathbf{D}_{\text{b}})$ ,  $\widehat{\text{SKCE}}_{\text{uq}}(\mathbf{D}_{\text{uq}})$ , and  $\widehat{\text{SKCE}}_{\text{ul}}(\mathbf{D}_{\text{ul}})$  and the asymptotic approximations for  $\widehat{\text{SKCE}}_{\text{uq}}(\mathbf{A}_{\text{uq}})$  and  $\widehat{\text{SKCE}}_{\text{ul}}(\mathbf{A}_{\text{ul}})$ , where the former is approximated by bootstrapping. We compare them with a previously proposed hypothesis test for the standard ECE estimator based on consistency resampling (C), in which

<sup>4</sup>The implementation of the experiments is available online at <https://github.com/devmotion/CalibrationPaper>.

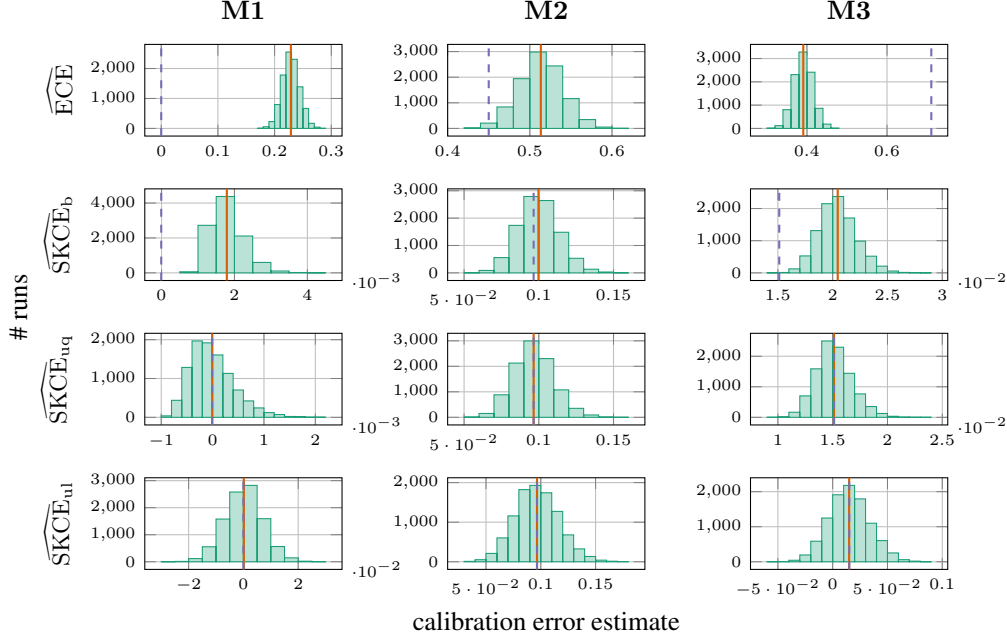


Figure 1: Distribution of calibration error estimates of  $10^4$  data sets that are randomly sampled from the generative models **M1**, **M2**, and **M3**. The solid line indicates the mean of the calibration error estimates, and the dashed line displays the true calibration error.

data sets are resampled under the assumption that the model is calibrated by sampling labels from resampled predictions (Bröcker and Smith, 2007; Vaicenavicius et al., 2019).

For a chosen significance level  $\alpha$  we compute from the p-value approximations  $p_1, \dots, p_{10^4}$  the empirical test error

$$\frac{1}{10^4} \sum_{i=1}^{10^4} \mathbb{1}_{[0, \alpha]}(p_i) \quad (\text{for M1}) \quad \text{and} \quad \frac{1}{10^4} \sum_{i=1}^{10^4} \mathbb{1}_{(\alpha, 1]}(p_i) \quad (\text{for M2 and M3}).$$

In Fig. 2 we plot these empirical test errors versus the significance level.

As expected, the distribution-free bounds seem to be very loose upper bounds of the p-value, resulting in statistical tests without much power. The asymptotic approximations, however, seem to estimate the p-value quite well on average, as can be seen from the overlap with the diagonal in the results for the calibrated model **M1** (the empirical test error matches the chosen significance level). Additionally, calibration tests based on asymptotic distribution of these statistics, and in particular of  $\widehat{\text{SKCE}}_{\text{uq}}$ , are quite powerful in our experiments, as the results for the uncalibrated models **M2** and **M3** show. For the calibrated model, consistency resampling leads to an empirical test error that is not upper bounded by the significance level, i.e., the null hypothesis of the model being calibrated is rejected too often. This behaviour is caused by an underestimation of the p-value on average, which unfortunately makes the calibration test based on consistency resampling for the standard ECE estimator unreliable.

### 7.3 Additional experiments

In Appendix J.2.3 we provide additional results for varying number of classes and a non-standard ECE estimator with data-dependent bins. We observe that the bias of  $\widehat{\text{ECE}}$  becomes more prominent with increasing number of classes, showing high calibration error estimates even for calibrated models. The estimators of the SKCE are not affected by the number of classes in the same way. In some experiments with 100 and 1000 classes, however, the distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  shows multi-modality.

The considered calibration measures depend only on the predictions and the true labels, not on how these predictions are computed. Hence directly specifying the predictions allows a clean numerical



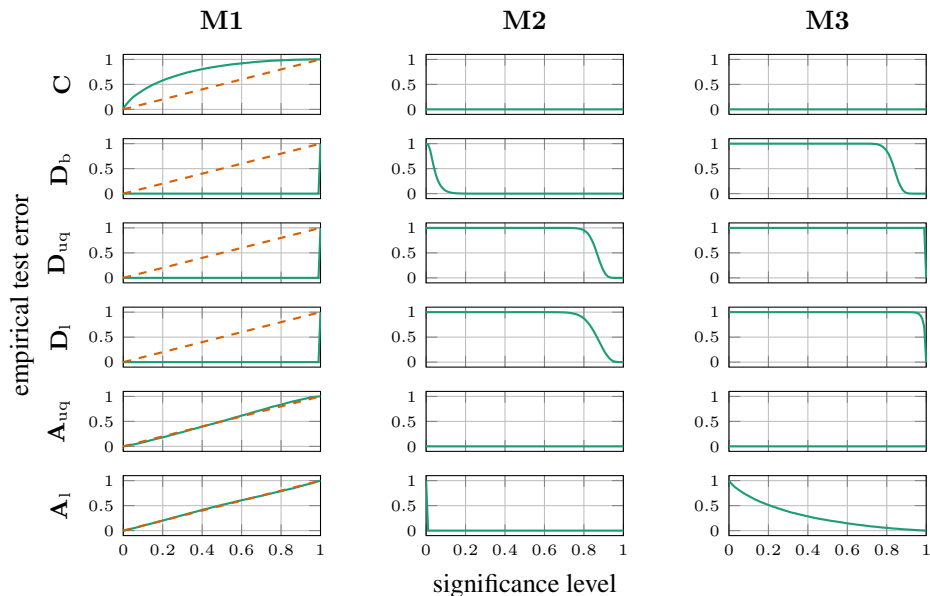


Figure 2: Empirical test error versus significance level for different bounds and approximations of the p-value, evaluated on  $10^4$  data sets that are randomly sampled from the generative models **M1**, **M2**, and **M3**. The dashed line highlights the diagonal of the unit square.

evaluation and enables comparisons of the estimates with the true calibration error. Nevertheless, we provide a more practical evaluation of calibration for several modern neural networks in Appendix J.3.

## 8 Conclusion

We have presented a unified framework for quantifying the calibration error of probabilistic classifiers. The framework encompasses several existing error measures and enables the formulation of a new kernel-based measure. We have derived unbiased and consistent estimators of the kernel-based error measures, which are properties not readily enjoyed by the more common and less tractable ECE. The impact of the kernel and its hyperparameters on the estimators is an important question for future research. We have refrained from investigating it in this paper, since it deserves a more exhaustive study than, what we felt, would have been possible in this work.

The calibration error estimators can be viewed as test statistics. This confers probabilistic interpretability to the error measures. Specifically, we can compute well-founded bounds and approximations of the p-value for the observed error estimates under the null hypothesis that the model is calibrated. We have derived distribution-free bounds and asymptotic approximations for the estimators of the proposed kernel-based error measure, that allow reliable calibration tests in contrast to previously proposed tests based on consistency resampling with the standard estimator of the ECE.

## Acknowledgements

We thank the reviewers for all the constructive feedback on our paper. This research is financially supported by the Swedish Research Council via the projects *Learning of Large-Scale Probabilistic Dynamical Models* (contract number: 2016-04278) and *Counterfactual Prediction Methods for Heterogeneous Populations* (contract number: 2018-05040), by the Swedish Foundation for Strategic Research via the project *Probabilistic Modeling and Inference for Machine Learning* (contract number: ICA16-0015), and by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

## References

- M. A. Arcones and E. Giné. On the bootstrap of  $U$  and  $V$  statistics. *The Annals of Statistics*, 20(2): 655–674, 1992.
- N. Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–337, 3 1950.
- P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer US, 2004.
- J. Bröcker. Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal of the Royal Meteorological Society*, 135(643):1512–1519, 7 2009.
- J. Bröcker and L. A. Smith. Increasing the reliability of reliability diagrams. *Weather and Forecasting*, 22(3):651–661, 6 2007.
- A. Caponnetto, C. A. Micchelli, M. Pontil, and Y. Ying. Universal multi-task kernels. *Journal of Machine Learning Research*, 9:1615–1646, 6 2008.
- C. Carmeli, E. De Vito, A. Toigo, and V. Umanità. Vector valued reproducing kernel Hilbert spaces and universality. *Analysis and Applications*, 08(01):19–61, 01 2010.
- A. Christmann and I. Steinwart. *Support Vector Machines*. Information Science and Statistics. Springer New York, 2008.
- M. H. DeGroot and S. E. Fienberg. The comparison and evaluation of forecasters. *The Statistician*, 32(1/2):12, 3 1983.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 3 2012.
- C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330, 2017.
- W. Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 9 1948.
- W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- A. D. Kiureghian and O. Ditlevsen. Aleatory or epistemic? Does it matter? *Structural Safety*, 31(2): 105–112, 3 2009.
- A. Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- A. Kumar, S. Sarawagi, and U. Jain. Trainable calibration measures for neural networks from kernel mean embeddings. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2805–2814, 2018.
- C. A. Micchelli and M. Pontil. On learning vector-valued functions. *Neural Computation*, 17(1): 177–204, 2005.
- A. H. Murphy and R. L. Winkler. Reliability of subjective probability forecasts of precipitation and temperature. *Applied Statistics*, 26(1):41, 1977.
- M. P. Naeni, G. Cooper, and M. Hauskrecht. Obtaining well calibrated probabilities using Bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- M. Pastell. Weave.jl: Scientific reports using Julia. *The Journal of Open Source Software*, 2(11):204, Mar. 2017.

- H. Phan. PyTorch-CIFAR10, 2019. URL <https://github.com/huyvnphan/PyTorch-CIFAR10/>.
- W. Rudin. *Real and complex analysis*. McGraw-Hill, New York, 3 edition, 1986.
- R. J. Serfling, editor. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, Inc., 11 1980.
- J. Vaicenavicius, D. Widmann, C. Andersson, F. Lindsten, J. Roll, and T. B. Schön. Evaluating model calibration in classification. In *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 3459–3467, 2019.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. Springer New York, 1996.
- D. Widmann. devmotion/CalibrationErrors.jl: v0.1.0, Sept. 2019a. URL <https://doi.org/10.5281/zenodo.3457945>.
- D. Widmann. devmotion/CalibrationTests.jl: v0.1.0, Oct. 2019b. URL <https://doi.org/10.5281/zenodo.3514933>.
- D. Widmann. devmotion/ConsistencyResampling.jl: v0.2.0, May 2019c. URL <https://doi.org/10.5281/zenodo.3232854>.
- B. Zadrozny and C. Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD*. ACM Press, 2002.

## A Notation

We introduce some additional notation to keep the following discussion concise.

Let  $U$  be a compact metric space and  $V$  be a Hilbert space. In this paper we only consider  $U = [0, 1]$  and  $U = \Delta^m$ . By  $\mathcal{C}(U, V)$  we denote the space of continuous functions  $f: U \rightarrow V$ .

For  $1 \leq p < \infty$ ,  $L^p(U, \mu; V)$  is the space of (equivalence classes of) measurable functions  $f: U \rightarrow V$  such that  $\|f\|^p$  is  $\mu$ -integrable, equipped with norm  $\|f\|_{\mu,p} = (\int_U \|f(x)\|^p \mu(dx))^{1/p}$ ; for  $p = \infty$ ,  $L^\infty(U, \mu; V)$  is the space of  $\mu$ -essentially bounded measurable functions  $f: U \rightarrow V$  with norm  $\|f\|_{\mu,\infty} = \mu\text{-ess sup}_{x \in \mathcal{X}} \|f(x)\|$ . We denote the closed unit ball of the space  $L^p(U, \mu; V)$  by  $K^p(U, \mu; V) := \{f \in L^p(U, \mu; V): \|f\|_{\mu,p} \leq 1\}$ .

If the norm  $\|\cdot\|_V$  on  $V$  is not clear from the context we indicate it by writing  $L^p(U, \mu; V, \|\cdot\|_V)$ ,  $K^p(U, \mu; V, \|\cdot\|_V)$ , and  $\|\cdot\|_{\mu,p;\|\cdot\|_V}$ . If  $V \subset \mathbb{R}^d$ , all possible norms  $\|\cdot\|_V$  are equivalent and hence we choose  $\|\cdot\|_p$  on  $V$  to simplify our calculations, if not stated otherwise. Moreover, if  $V \subset \mathbb{R}^d$  and  $\|\cdot\|_V = \|\cdot\|_q$  for some  $1 \leq q \leq \infty$ , for convenience we write  $\|\cdot\|_{\mu,p;q} = \|\cdot\|_{\mu,p;\|\cdot\|_V}$ .

Let  $W$  be another Hilbert space. Then  $\mathcal{L}(V, W)$  denotes the space of bounded linear operators  $T: V \rightarrow W$ ; if  $W = V$ , we write  $\mathcal{L}(V) := \mathcal{L}(V, V)$ . The induced operator norm on  $\mathcal{L}(V, W)$  is defined by

$$\begin{aligned} \|T\|_{\|\cdot\|_V; \|\cdot\|_W} &= \inf\{c \geq 0: \|Tv\|_W \leq c\|v\|_V \text{ for all } v \in V\} \\ &= \sup_{v \in V: \|v\|_V \leq 1} \|Tv\|_W. \end{aligned}$$

If  $W = V$ , we write  $\|\cdot\|_{\|\cdot\|_V} = \|\cdot\|_{\|\cdot\|_V; \|\cdot\|_V}$ . Moreover, for convenience if  $V \subset \mathbb{R}^d$  and  $\|\cdot\|_V = \|\cdot\|_p$  for some  $1 \leq p \leq \infty$ , we use the notation  $\|\cdot\|_{p; \|\cdot\|_W} = \|\cdot\|_{\|\cdot\|_V; \|\cdot\|_W}$ . In the same way, if  $W \subset \mathbb{R}^d$  and  $\|\cdot\|_W = \|\cdot\|_q$  for some  $1 \leq q \leq \infty$ , we write  $\|\cdot\|_{\|\cdot\|_V; q} = \|\cdot\|_{\|\cdot\|_V; \|\cdot\|_W}$ .

By  $\mathcal{B}(T)$  we denote the Borel  $\sigma$ -algebra of a topological space  $T$ .

We write  $\mu_A$  for the distribution of a random variable  $A$ , i.e., the pushforward measure it induces, if  $A$  is defined on a probability space with probability measure  $\mu$ .

## B Probabilistic setting

Let  $(\Omega, \mathcal{A}, \mu)$  be a probability space. Let  $m \in \mathbb{N}$  and define the random variables  $X: (\Omega, \mathcal{A}) \rightarrow (\mathcal{X}, \Sigma_X)$  and  $Y: (\Omega, \mathcal{A}) \rightarrow (\{1, \dots, m\}, 2^{\{1, \dots, m\}})$ , such that  $\Sigma_X$  contains all singletons. We denote a version of the regular conditional distribution of  $Y$  given  $X$  by  $\mu_{Y|X}(\cdot|x)$  for all  $x \in \mathcal{X}$ .

We consider the problem of learning a measurable function  $g: (\mathcal{X}, \Sigma_X) \rightarrow (\Delta^m, \mathcal{B}(\Delta^m))$  that returns the prediction  $g_y(x)$  of  $\mu_{Y|X}(\{y\}|x)$  for all  $x \in \mathcal{X}$  and  $y \in \{1, \dots, m\}$ . We define the random variable  $G: (\Omega, \mathcal{A}) \rightarrow (\Delta^m, \mathcal{B}(\Delta^m))$  as  $G := g(X)$ .

In the same way as above, we denote a version of the regular conditional distribution of  $Y$  given  $G$  by  $\mu_{Y|G}(\cdot|t)$  for all  $t \in \Delta^m$ . The function  $\delta: (\Delta^m, \mathcal{B}(\Delta^m)) \rightarrow (\mathbb{R}^m, \mathcal{B}(\mathbb{R}^m))$ , given by

$$\delta(t) := \begin{pmatrix} \mu_{Y|G}(1|t) - t \\ \vdots \\ \mu_{Y|G}(m|t) - t \end{pmatrix}$$

for all  $t \in \Delta^m$ , gives rise to another random variable  $\Delta: (\Omega, \mathcal{A}) \rightarrow (\mathbb{R}^m, \mathcal{B}(\mathbb{R}^m))$  that is defined by  $\Delta := \delta(G)$ .

Using the newly introduced mathematical notation, we can reformulate strong calibration in a compact way. A model  $g$  is calibrated in the strong sense if  $\mu_{Y|G}(\cdot|G) = G$  almost surely, or equivalently if  $\Delta = 0$  almost surely.

## C Calibration error

**Definition C.1 (Calibration error).** Let  $\mathcal{F} \subset L^1(\Delta^m, \mu_G; \mathbb{R}^m)$  be non-empty. Then the calibration error CE of model  $g$  with respect to class  $\mathcal{F}$  is

$$\text{CE}[\mathcal{F}, g] := \sup_{f \in \mathcal{F}} \mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}].$$

*Remark C.1.* Note that  $\|\Delta\|_\infty \leq 1$  almost surely, and thus  $\|\delta\|_{\mu_G, \infty} \leq 1$ . Hence by Hölder's inequality for all  $f \in \mathcal{F}$  we have

$$|\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]| \leq \mathbb{E}[|\langle \Delta, f(G) \rangle_{\mathbb{R}^m}|] \leq \|\delta\|_{\mu_G, \infty} \|f\|_{\mu_G, 1} \leq \|f\|_{\mu_G, 1} < \infty.$$

However, it is still possible that  $\text{CE}[\mathcal{F}, g] = \infty$ .

*Remark C.2.* If  $\mathcal{F}$  is symmetric in the sense that  $f \in \mathcal{F}$  implies  $-f \in \mathcal{F}$ , then  $\text{CE}[\mathcal{F}, g] \geq 0$ .

The measure highly depends on the choice of  $\mathcal{F}$  but strong calibration always implies a calibration error of zero.

**Theorem C.1 (Strong calibration implies zero error).** Let  $\mathcal{F} \subset L^1(\Delta^m, \mu_G; \mathbb{R})$ . If model  $g$  is calibrated in the strong sense, then  $\text{CE}[\mathcal{F}, g] = 0$ .

*Proof.* If model  $g$  is calibrated in the strong sense, then  $\Delta = 0$  almost surely. Hence for all  $f \in \mathcal{F}$  we have  $\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}}] = 0$ , which implies  $\text{CE}[\mathcal{F}, g] = 0$ .  $\square$

Of course, the converse statement is not true in general. A similar result as above shows that the class of continuous functions, albeit too large and impractical, allows to identify calibrated models.

**Theorem C.2.** Let  $\mathcal{F} = \mathcal{C}(\Delta^m, \mathbb{R}^m)$ . Then  $\text{CE}[\mathcal{F}, g] = 0$  if and only if model  $g$  is calibrated in the strong sense.

*Proof.* Note that  $\mathcal{F}$  is well defined since  $\mathcal{F} \subset L^1(\Delta^m, \mu_G; \mathbb{R}^m)$ .

If model  $g$  is calibrated in the strong sense, then  $\text{CE}[\mathcal{F}, g] = 0$  by Theorem C.1.

If model  $g$  is not calibrated in the strong sense, then  $\Delta = 0$  does not hold almost surely. In particular, there exists  $s \in \{-1, 1\}^m$  such that  $\langle \Delta, s \rangle_{\mathbb{R}^m} \leq 0$  does not hold almost surely. Define the function  $f_s: \Delta^m \rightarrow \mathbb{R}^m$  by  $f_s := \langle \delta(\cdot), s \rangle_{\mathbb{R}^m}$  and let  $A_s := f_s^{-1}((0, \infty))$ . Then  $A_s \in \mathcal{B}(\Delta^m)$  since  $f_s$  is Borel measurable, and  $\mu_G(A_s) > 0$ . Hence we know that

$$\alpha_s := \mathbb{E}[\langle \Delta, s \mathbb{1}_{A_s}(G) \rangle_{\mathbb{R}^m}] > 0.$$

Since  $\mu_G$  is a Borel probability measure on a compact metric space, it is regular and hence there exist a compact set  $K$  and an open set  $U$  such that  $K \subset A_s \subset U$  and  $\mu_G(U \setminus K) < \alpha_s/4$  (Rudin, 1986, Theorem 2.17). Thus by Urysohn's lemma applied to the closed sets  $K$  and  $U^c$ , there exists a continuous function  $h \in \mathcal{C}(\Delta^m, \mathbb{R})$  such that  $\mathbb{1}_K \leq h \leq \mathbb{1}_U$ . By defining  $f = sh \in \mathcal{C}(\Delta^m, \mathbb{R}^m)$  and applying Hölder's inequality we obtain

$$\begin{aligned} \mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] &= \mathbb{E}[\langle \Delta, s \mathbb{1}_{A_s}(G) \rangle_{\mathbb{R}^m}] \\ &\quad + \mathbb{E}[\langle \Delta, f(G) - s \mathbb{1}_{A_s}(G) \rangle_{\mathbb{R}^m}] \\ &\geq \alpha_s - \mathbb{E}[|(h(G) - \mathbb{1}_{A_s}(G)) \langle \Delta, s \rangle_{\mathbb{R}^m}|] \\ &\geq \alpha_s - \mathbb{E}[|h(G) - \mathbb{1}_{A_s}(G)| |\langle \Delta, s \rangle_{\mathbb{R}^m}|] \\ &\geq \alpha_s - \mathbb{E}[\mathbb{1}_{U \setminus K}(G) \|\Delta\|_1 \|s\|_\infty] \geq \alpha_s - 2\mu_G(U \setminus K) \\ &> \alpha_s - \alpha_s/2 = \alpha_s/2 > 0. \end{aligned}$$

This implies  $\text{CE}[\mathcal{F}, g] > 0$ .  $\square$

## D Reproducing kernel Hilbert spaces of vector-valued functions on the probability simplex

**Definition D.1 (Micchelli and Pontil (2005, Definition 1)).** Let  $\mathcal{H}$  be a Hilbert space of vector-valued functions  $f: \Delta^m \rightarrow \mathbb{R}^m$  with inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ . We call  $\mathcal{H}$  a reproducing kernel Hilbert space (RKHS), if for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$  the functional  $E_{t,u}: \mathcal{H} \rightarrow \mathbb{R}$ ,  $E_{t,u}f := \langle u, f(t) \rangle_{\mathbb{R}^m}$ , is a bounded (or equivalently continuous) linear operator.

Riesz's representation theorem ensures that there exists a unique function  $k: \Delta^m \times \Delta^m \rightarrow \mathbb{R}^{m \times m}$  such that for all  $t \in \Delta^m$  the function  $k(\cdot, t)$  is a linear map from  $\Delta^m$  to  $\mathcal{H}$  and for all  $u \in \mathbb{R}^m$  it satisfies the so-called reproducing property

$$\langle u, f(t) \rangle_{\mathbb{R}^m} = E_{t,u} f = \langle k(\cdot, t)u, f \rangle_{\mathcal{H}}. \quad (\text{D.1})$$

It can be shown that function  $k$  is self-adjoint<sup>5</sup> and positive semi-definite, and hence a kernel according to Definition 1 (Micchelli and Pontil, 2005, Proposition 1). Similar to the scalar-valued case, conversely by the Moore-Aronszajn theorem (Aronszajn, 1950) to every kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  there exists a unique RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  with  $k$  as reproducing kernel (Micchelli and Pontil, 2005, Theorem 1).

Other useful properties are summarized in Lemma D.1. Micchelli and Pontil (2005) considered only the Euclidean norm on  $\mathbb{R}^m$ , corresponding to  $p = q = 2$  in our statement. For convenience we use the notation  $\|\cdot\|_{p;\mathcal{H}} = \|\cdot\|_{\cdot} \|\cdot\|_{p;\mathcal{H}}$ .

**Lemma D.1 (Micchelli and Pontil (2005, Proposition 1)).** *Let  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  be a RKHS with kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$ . Let  $1 \leq p, q \leq \infty$  with Hölder conjugates  $p'$  and  $q'$ , respectively.*

1. For all  $t \in \Delta^m$

$$\|k(\cdot, t)\|_{p;\mathcal{H}} = \|k(t, t)\|_{p;p'}^{1/2}. \quad (\text{D.2})$$

2. For all  $s, t \in \Delta^m$

$$\|k(s, t)\|_{p;q} \leq \|k(s, s)\|_{q';q}^{1/2} \|k(t, t)\|_{p;p'}^{1/2}. \quad (\text{D.3})$$

3. For all  $f \in \mathcal{H}$  and  $t \in \Delta^m$

$$\|f(t)\|_p \leq \|f\|_{\mathcal{H}} \|k(\cdot, t)\|_{p';\mathcal{H}} = \|f\|_{\mathcal{H}} \|k(t, t)\|_{p;p'}^{1/2}. \quad (\text{D.4})$$

*Proof.* Let  $t \in \Delta^m$ . From the reproducing property, Hölder's inequality, and the definition of the operator norm, we obtain for all  $u \in \mathbb{R}^m$

$$\|k(\cdot, t)u\|_{\mathcal{H}}^2 = \langle k(\cdot, t)u, k(\cdot, t)u \rangle_{\mathcal{H}} = \langle u, k(t, t)u \rangle_{\mathbb{R}^m} \leq \|u\|_p \|k(t, t)u\|_{p'} \leq \|u\|_p^2 \|k(t, t)\|_{p;p'},$$

which implies that

$$\|k(\cdot, t)\|_{p;\mathcal{H}} = \sup_{u \in \mathbb{R}^m \setminus \{0\}} \frac{\|k(\cdot, t)u\|_{\mathcal{H}}}{\|u\|_p} \leq \|k(t, t)\|_{p;p'}^{1/2}. \quad (\text{D.5})$$

On the other hand, for all  $u, v \in \mathbb{R}^m$  it follows from the reproducing property, the Cauchy-Schwarz inequality, and the definition of the operator norm that

$$\langle u, k(t, t)v \rangle_{\mathbb{R}^m} = \langle k(\cdot, t)u, k(\cdot, t)v \rangle_{\mathcal{H}} \leq \|k(\cdot, t)u\|_{\mathcal{H}} \|k(\cdot, t)v\|_{\mathcal{H}} \leq \|u\|_p \|v\|_p \|k(\cdot, t)\|_{p;\mathcal{H}}^2.$$

Since the  $\ell_p$ -norm is the dual norm of the  $\ell_{p'}$ -norm, it follows that

$$\|k(t, t)v\|_{p'} = \sup_{u \in \mathbb{R}^m: \|u\|_p \leq 1} \langle u, k(t, t)v \rangle_{\mathbb{R}^m} \leq \|v\|_p \|k(\cdot, t)\|_{p;\mathcal{H}}^2,$$

which implies that

$$\|k(t, t)\|_{p;p'} = \sup_{v \in \mathbb{R}^m \setminus \{0\}} \frac{\|k(t, t)v\|_{p'}}{\|v\|_p} \leq \|k(\cdot, t)\|_{p;\mathcal{H}}^2. \quad (\text{D.6})$$

Equation (D.2) follows from Eqs. (D.5) and (D.6).

Let  $s, t \in \Delta^m$ . From the reproducing property, the Cauchy-Schwarz inequality, and the definition of the operator norm, we get for all  $u, v \in \mathbb{R}^m$

$$\begin{aligned} \langle u, k(s, t)v \rangle_{\mathbb{R}^m} &\leq \langle k(\cdot, s)u, k(\cdot, t)v \rangle_{\mathcal{H}} \leq \|k(\cdot, s)u\|_{\mathcal{H}} \|k(\cdot, t)v\|_{\mathcal{H}} \\ &\leq \|u\|_{q'} \|v\|_p \|k(\cdot, s)\|_{q';\mathcal{H}} \|k(\cdot, t)\|_{p;\mathcal{H}}. \end{aligned}$$

<sup>5</sup>Let  $U, V$  be two Hilbert spaces. Then the adjoint of a linear operator  $T \in \mathcal{L}(U, V)$  is the linear operator  $T^* \in \mathcal{L}(V, U)$  such that for all  $u \in U, v \in V$   $\langle Tu, v \rangle_V = \langle u, T^*v \rangle_U$ .

Thus we obtain

$$\|k(s, t)v\|_q = \sup_{u \in \mathbb{R}^m: \|u\|_{q'} \leq 1} \langle u, k(s, t)v \rangle_{\mathbb{R}^m} \leq \|v\|_p \|k(\cdot, s)\|_{q'; \mathcal{H}} \|k(\cdot, t)\|_{p; \mathcal{H}},$$

which implies

$$\|k(s, t)\|_{p; q} = \sup_{v \in \mathbb{R}^m \setminus \{0\}} \frac{\|k(s, t)v\|_q}{\|v\|_p} \leq \|k(\cdot, s)\|_{q'; \mathcal{H}} \|k(\cdot, t)\|_{p; \mathcal{H}}.$$

Hence from Eq. (D.2) we obtain Eq. (D.3).

For the third statement, let  $f \in \mathcal{H}$  and  $t \in \Delta^m$ . From the reproducing property, the Cauchy-Schwarz inequality, and the definition of the operator norm, we obtain for all  $u \in \mathbb{R}^m$

$$\langle u, f(t) \rangle_{\mathbb{R}^m} = \langle k(\cdot, t)u, f \rangle_{\mathcal{H}} \leq \|k(\cdot, t)u\|_{\mathcal{H}} \|f\|_{\mathcal{H}} \leq \|u\|_{p'} \|k(\cdot, t)\|_{p'; \mathcal{H}} \|f\|_{\mathcal{H}}.$$

Thus the duality of the  $\ell_p$ - and the  $\ell_{p'}$ -norm implies

$$\|f(t)\|_p = \sup_{u \in \mathbb{R}^m: \|u\|_{p'} \leq 1} \langle u, f(t) \rangle_{\mathbb{R}^m} \leq \|k(\cdot, t)\|_{p'; \mathcal{H}} \|f\|_{\mathcal{H}},$$

which together with Eq. (D.2) yields Eq. (D.4).  $\square$

If  $\mu$  is a measure on  $\Delta^m$ , we define for  $1 \leq p, q \leq \infty$   $\|k\|_{\mu, p; q} := \|\tilde{k}_q\|_{\mu, p}$  where  $\tilde{k}_q: \Delta^m \rightarrow \mathbb{R}$  is given by  $\tilde{k}_q(t) := \|k(\cdot, t)\|_{q; \mathcal{H}} = \|k(t, t)\|_{q; q}^{1/2}$ . We omit the value of  $q$  if it is clear from the context or does not matter, since all norms on  $\mathbb{R}^m$  are equivalent.

It is possible to construct certain classes of matrix-valued kernels from scalar-valued kernels, as the following example shows.

**Example D.1 (Micchelli and Pontil (2005); Caponnetto et al. (2008, Example 1))**

For all  $i \in \{1, \dots, n\}$ , let  $k_i: \Delta^m \times \Delta^m \rightarrow \mathbb{R}$  be a scalar-valued kernel and  $A_i \in \mathbb{R}^{m \times m}$  be a positive semi-definite matrix. Then the function

$$k: \Delta^m \times \Delta^m \rightarrow \mathbb{R}^{m \times m}, \quad k(s, t) := \sum_{i=1}^n k_i(s, t) A_i, \quad (\text{D.7})$$

is a matrix-valued kernel.

We state a simple result about measurability of functions in the considered RKHSs. The result can be formulated in a much more general fashion and is similar to a result by Christmann and Steinwart (2008, Lemma 4.24).

**Lemma D.2 (Measurable RKHS).** *Let  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  be a RKHS with kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$ . Then all  $f \in \mathcal{H}$  are measurable if and only if  $k(\cdot, t)u \in (\mathbb{R}^m)^{\Delta^m}$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ .*

*Proof.* If all  $f \in \mathcal{H}$  are measurable, then  $k(\cdot, t)u \in \mathcal{H}$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ .

If  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , then all functions in  $\mathcal{H}_0 := \text{span} \{k(\cdot, t)u: t \in \Delta^m, u \in \mathbb{R}^m\} \subset \mathcal{H}$  are measurable.

Let  $f \in \mathcal{H}$ . Since  $\mathcal{H} = \overline{\mathcal{H}_0}$  (see, e.g., Carmeli et al., 2010), there exists a sequence  $(f_n)_{n \in \mathbb{N}} \subset \mathcal{H}_0$  such that  $\lim_{n \rightarrow \infty} \|f - f_n\|_{\mathcal{H}} = 0$ . For all  $t \in \Delta^m$ , since the operator  $k^*(\cdot, t)$  is continuous, by the reproducing property we obtain  $\lim_{n \rightarrow \infty} f_n(t) = f(t)$ . Thus  $f$  is measurable.  $\square$

By definition (see, e.g., Carmeli et al., 2010, Definition 1), a RKHS with a continuous kernel is a subspace of the space of continuous functions. The following equivalent formulation is an immediate consequence of the result by Carmeli et al. (2010).

**Corollary D.1 (Carmeli et al. (2010, Proposition 1)).** *A kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathbb{R}^m$  is continuous if for all  $t \in \Delta^m$   $t \mapsto \|k(t, t)\|$  is bounded and for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$   $k(\cdot, t)u$  is a continuous function from  $\Delta^m$  to  $\mathbb{R}^m$ .*

An important class of continuous kernels are so-called universal kernels, for which the corresponding RKHS is a dense subset of the space of continuous functions with respect to the uniform norm. A result by Caponnetto et al. (2008) shows under what assumptions matrix-valued kernels of the form in Example D.1 are universal.

**Lemma D.3 (Caponnetto et al. (2008, Theorem 14)).** *For all  $i \in \{1, \dots, n\}$ , let  $k_i: \Delta^m \times \Delta^m \rightarrow \mathbb{R}$  be a universal scalar-valued kernel and  $A_i \in \mathbb{R}^{m \times m}$  be a positive semi-definite matrix. Then the matrix-valued kernel defined in Eq. (D.7) is universal if and only if  $\sum_{i=1}^n A_i$  is positive definite.*

## E Kernel calibration error

The one-to-one correspondence between matrix-valued kernels and RKHSs of vector-valued functions motivates the introduction of the kernel calibration error (KCE) in Definition 3. For certain kernels we are able to identify strongly calibrated models.

**Theorem E.1.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a universal continuous kernel. Then  $\text{KCE}[k, g] = 0$  if and only if model  $g$  is calibrated in the strong sense.*

*Proof.* Let  $\mathcal{F}$  be the unit ball in the RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  corresponding to kernel  $k$ . Since kernel  $k$  is continuous, by definition  $\mathcal{H} \subset \mathcal{C}(\Delta^m, \mathbb{R}^m)$  (Carmeli et al., 2010, Definition 1). Thus  $\mathcal{F}$  is well defined since  $\mathcal{F} \subset \mathcal{C}(\Delta^m, \mathbb{R}^m) \subset L^1(\Delta^m, \mu_G; \mathbb{R}^m)$ .

If  $g$  is calibrated in the strong sense, it follows from Theorem C.1 that  $\text{KCE}[k, g] = \text{CE}[\mathcal{F}, g] = 0$ .

Assume that  $\text{KCE}[k, g] = \text{CE}[\mathcal{F}, g] = 0$ . This implies  $\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] = 0$  for all  $f \in \mathcal{F}$ . Let  $f \in \mathcal{C}(\Delta^m, \mathbb{R}^m)$ . Since  $\mathcal{H}$  is dense in  $\mathcal{C}(\Delta^m, \mathbb{R}^m)$  (Carmeli et al., 2010, Theorem 1), for all  $\epsilon > 0$  there exists a function  $h \in \mathcal{H}$  with  $\|f - h\|_\infty < \epsilon/2$ . Define  $\tilde{h} \in \mathcal{F}$  by  $\tilde{h} := h/\|h\|_{\mathcal{H}}$  if  $\|h\|_{\mathcal{H}} \neq 0$  and  $\tilde{h} := h$  otherwise. Since

$$\mathbb{E}[\langle \Delta, h(G) \rangle_{\mathbb{R}^m}] = \|h\|_{\mathcal{H}} \mathbb{E}[\langle \Delta, \tilde{h}(G) \rangle_{\mathbb{R}^m}] = 0,$$

by Hölder's inequality we obtain

$$\begin{aligned} |\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]| &= |\mathbb{E}[\langle \Delta, f(G) - h(G) \rangle_{\mathbb{R}^m}]| \\ &\leq \mathbb{E}[|\langle \Delta, f(G) - h(G) \rangle_{\mathbb{R}^m}|] \\ &\leq \|\delta\|_{\mu_G, 1} \|f - h\|_{\mu_G, \infty} \\ &\leq 2\|f - h\|_\infty < \epsilon. \end{aligned}$$

Thus  $\text{CE}[\mathcal{C}(\Delta^m, \mathbb{R}^m), g] = 0$ , and hence  $g$  is calibrated in the strong sense by Theorem C.2.  $\square$

Similar to the maximum mean discrepancy (Gretton et al., 2012), if we consider functions in a RKHS, we can rewrite the expectation  $\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]$  as an inner product in the Hilbert space.

**Lemma E.1 (Existence and uniqueness of embedding).** *Let  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  be a RKHS with kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$ , and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_G, 1} < \infty$ .*

*Then there exists a unique embedding  $\mu_g \in \mathcal{H}$  such that for all  $f \in \mathcal{H}$*

$$\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] = \langle f, \mu_g \rangle_{\mathcal{H}}.$$

*The embedding  $\mu_g$  satisfies for all  $t \in \Delta^m$  and  $y \in \mathbb{R}^m$*

$$\langle y, \mu_g(t) \rangle_{\mathbb{R}^m} = \mathbb{E}[\langle \Delta, k(G, t)y \rangle_{\mathbb{R}^m}].$$

*Proof.* By Lemma D.2 all  $f \in \mathcal{H}$  are measurable. Moreover, by Eq. (D.4) for all  $f \in \mathcal{H}$  we have

$$\begin{aligned} \int_{\Delta^m} \|f(t)\|_1 \mu_G(dt) &\leq \|f\|_{\mathcal{H}} \int_{\Delta^m} \|k(t, t)\|_{\infty; 1}^{1/2} \mu_G(dt) \\ &= \|f\|_{\mathcal{H}} \|k\|_{\mu_G, 1; \infty} < \infty, \end{aligned}$$

and thus  $\mathcal{H} \subset L^1(\Delta^m, \mu_G; \mathbb{R}^m)$ . Hence from Remark C.1 (with  $\mathcal{F} = \mathcal{H}$ ) we know that for all  $f \in \mathcal{H}$  the expectation  $\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]$  exists and is finite.



Define the linear operator  $T_g: \mathcal{H} \rightarrow \mathbb{R}$  by  $T_g f := \mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]$  for all  $f \in \mathcal{H}$ . In the same way as above, for all functions  $f \in \mathcal{H}$  Hölder's inequality and Eq. (D.4) imply

$$\begin{aligned} |T_g f| &= |\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]| \leq \mathbb{E}[|\langle \Delta, f(G) \rangle_{\mathbb{R}^m}|] \\ &\leq \|\delta\|_{\mu_G, \infty} \|f\|_{\mu_G, 1} \\ &\leq \|f\|_{\mu_G, 1} \leq \|f\|_{\mathcal{H}} \|k\|_{\mu_G, 1; \infty} < \infty. \end{aligned}$$

Thus  $T_g$  is a continuous linear operator, and therefore it follows from Riesz's representation theorem that there exists a unique function  $\mu_g \in \mathcal{H}$  such that

$$\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] = T_g f = \langle f, \mu_g \rangle_{\mathcal{H}}$$

for all  $f \in \mathcal{H}$ . This implies that for all  $t \in \Delta^m$  and  $y \in \mathbb{R}^m$

$$\langle y, \mu_g(t) \rangle_{\mathbb{R}^m} = \langle k(\cdot, t)y, \mu_g \rangle_{\mathcal{H}} = \mathbb{E}[\langle \Delta, k(G, t)y \rangle_{\mathbb{R}^m}]. \quad \square$$

Lemma E.1 allows us to rewrite  $\text{KCE}[k, g]$  in a more explicit way.

**Lemma E.2 (Explicit formulation).** *Let  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  be a RKHS with kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$ , and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$  and  $\|k\|_{\mu_G, 1} < \infty$ . Then*

$$\text{KCE}[k, g] = \|\mu_g\|_{\mathcal{H}},$$

where  $\mu_g$  is the embedding defined in Lemma E.1. Moreover,

$$\text{SKCE}[k, g] := \text{KCE}^2[k, g] = \mathbb{E}[\langle e_Y - g(X), k(g(X), g(X'))(e_{Y'} - g(X')) \rangle_{\mathbb{R}^m}],$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$  and  $e_i$  denotes the  $i$ th unit vector.

*Proof.* Let  $\mathcal{F}$  be the unit ball in the RKHS  $\mathcal{H}$ . From Lemma E.1 we know that for all  $f \in \mathcal{F}$  the expectation  $\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}]$  exists and is given by

$$\mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] = \langle f, \mu_g \rangle_{\mathcal{H}},$$

where  $\mu_g$  is the embedding defined in Lemma E.1. Thus the definition of the dual norm yields

$$\text{KCE}[k, g] = \text{CE}[\mathcal{F}, g] = \sup_{f \in \mathcal{F}} \mathbb{E}[\langle \Delta, f(G) \rangle_{\mathbb{R}^m}] = \sup_{f \in \mathcal{F}} \langle f, \mu_g \rangle_{\mathcal{H}} = \|\mu_g\|_{\mathcal{H}}.$$

Thus from the reproducing property and Lemma E.1 we obtain

$$\begin{aligned} \text{SKCE}[k, g] &= \text{KCE}^2[k, g] = \langle \mu_g, \mu_g \rangle_{\mathcal{H}} = \mathbb{E}[\langle \Delta, \mu_g(G) \rangle_{\mathbb{R}^m}] \\ &= \mathbb{E}[\mathbb{E}[\langle \Delta', k(G', G)\Delta \rangle_{\mathbb{R}^m} | G]] \\ &= \mathbb{E}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m}], \end{aligned}$$

where  $G'$  is an independent copy of  $G$  and  $\Delta' := \delta(G')$ .

By rewriting

$$\Delta = \mathbb{E}[e_Y | G] - G = \mathbb{E}[e_Y - G | G]$$

and  $\Delta'$  in the same way, we get

$$\text{SKCE}[k, g] = \mathbb{E}[\langle e_Y - G, k(G, G')(e_{Y'} - G') \rangle_{\mathbb{R}^m}].$$

Plugging in the definitions of  $G$  and  $G'$  yields

$$\text{SKCE}[k, g] = \mathbb{E}[\langle e_Y - g(X), k(g(X), g(X'))(e_{Y'} - g(X')) \rangle_{\mathbb{R}^m}]. \quad \square$$

## F Estimators

Let  $\mathcal{D} = \{(X_i, Y_i)\}_{i=1}^n$  be a set of random variables that are i.i.d. as  $(X, Y)$ . Regardless of the space  $\mathcal{F}$  the plug-in estimator of  $\text{CE}[\mathcal{F}, g]$  is

$$\widehat{\text{CE}}[\mathcal{F}, g, \mathcal{D}] := \sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \langle \delta(X_i, Y_i), f(g(X_i)) \rangle_{\mathbb{R}^m}.$$

If  $\mathcal{F}$  is the unit ball in a RKHS, i.e., for the kernel calibration error, we can calculate this estimator explicitly.

**Lemma F.1 (Biased estimator).** Let  $\mathcal{F}$  be the unit ball in a RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  with kernel  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$ . Then

$$\widehat{\text{CE}}[\mathcal{F}, g, \mathcal{D}] = \frac{1}{n} \left[ \sum_{i,j=1}^n \langle \delta(X_i, Y_i), k(g(X_i), g(X_j)) \delta(X_j, Y_j) \rangle_{\mathbb{R}^m} \right]^{1/2}.$$

*Proof.* From the reproducing property and the definition of the dual norm it follows that

$$\widehat{\text{CE}}[\mathcal{F}, g, \mathcal{D}] = \sup_{f \in \mathcal{F}} \left\langle \frac{1}{n} \sum_{i=1}^n k(\cdot, g(X_i)) \delta(X_i, Y_i), f \right\rangle_{\mathcal{H}} = \frac{1}{n} \left\| \sum_{i=1}^n k(\cdot, g(X_i)) \delta(X_i, Y_i) \right\|_{\mathcal{H}}.$$

Applying the reproducing property yields the result.  $\square$

Since we can uniquely identify the unit ball  $\mathcal{F}$  with the matrix-valued kernel  $k$  and the plug-in estimator in Lemma F.1 does not depend on  $\mathcal{F}$  explicitly, we introduce the notation

$$\widehat{\text{KCE}}[k, g, \mathcal{D}] := \widehat{\text{CE}}[\mathcal{F}, g, \mathcal{D}] \quad \text{and} \quad \widehat{\text{SKCE}}_{\text{b}}[k, g, \mathcal{D}] := \widehat{\text{KCE}}^2[k, g, \mathcal{D}],$$

where  $\mathcal{F}$  is the unit ball in the RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  corresponding to kernel  $k$ . By removing the terms involving the same random variables we obtain an unbiased estimator.

**Lemma F.2 (Unbiased estimator).** Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_{G,1}} < \infty$ . Then

$$\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] := \frac{1}{n(n-1)} \sum_{\substack{i,j=1, \\ i \neq j}}^n \langle \delta(X_i, Y_i), k(g(X_i), g(X_j)) \delta(X_j, Y_j) \rangle_{\mathbb{R}^m}$$

is an unbiased estimator of  $\text{SKCE}[k, g]$ .

*Proof.* The assumptions of Lemma E.2 are satisfied, and hence we know that

$$\text{SKCE}[k, g] = \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X')) \delta(X', Y') \rangle_{\mathbb{R}^m}],$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$ . Since  $(X, Y)$ ,  $(X', Y')$ , and  $(X_i, Y_i)$  are i.i.d., we have

$$\begin{aligned} \mathbb{E}[\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}]] &= \frac{1}{n(n-1)} \sum_{\substack{i=1, \\ i \neq j}}^n \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X')) \delta(X', Y') \rangle_{\mathbb{R}^m}] \\ &= \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X')) \delta(X', Y') \rangle_{\mathbb{R}^m}] \\ &= \text{SKCE}[k, g], \end{aligned}$$

which shows that  $\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}]$  is an unbiased estimator of  $\text{SKCE}[k, g]$ .  $\square$

There exists an unbiased estimator with higher variance that scales not quadratically but only linearly with the number of samples.

**Lemma F.3 (Linear estimator).** Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_{G,1}} < \infty$ . Then

$$\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] := \frac{1}{\lfloor n/2 \rfloor} \sum_{i=1}^{\lfloor n/2 \rfloor} \langle \delta(X_{2i-1}, Y_{2i-1}), k(g(X_{2i-1}), g(X_{2i})) \delta(X_{2i}, Y_{2i}) \rangle_{\mathbb{R}^m}$$

is an unbiased estimator of  $\text{SKCE}[k, g]$ .

*Proof.* The assumptions of Lemma E.2 are satisfied, and hence we know that

$$\text{SKCE}[k, g] = \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}],$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$ . Since  $(X, Y)$ ,  $(X', Y')$ , and  $(X_i, Y_i)$  are i.i.d., we have

$$\begin{aligned} \mathbb{E}[\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}]] &= \frac{1}{\lfloor n/2 \rfloor} \sum_{i=1}^{\lfloor n/2 \rfloor} \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}] \\ &= \mathbb{E}[\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}] \\ &= \text{SKCE}[k, g], \end{aligned}$$

which shows that  $\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}]$  is an unbiased estimator of  $\text{SKCE}[k, g]$ .  $\square$

## G Asymptotic distributions

In this section we investigate the asymptotic behaviour of the proposed estimators. We start with a simple but very useful statement.

**Lemma G.1.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_G, 2} < \infty$ .*

*Then  $\text{Var}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m}] < \infty$ , where  $G'$  is an independent copy of  $G$  and  $\Delta' := \delta(G')$ .*

*Proof.* From the Cauchy-Schwarz inequality and the definition of the operator norm we obtain

$$\mathbb{E}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m}^2] \leq \mathbb{E}[\|\Delta\|_2^2 \|k(G, G')\|_{2;2}^2 \|\Delta\|_2^2] \leq 4 \mathbb{E}[\|k(G, G')\|_{2;2}^2].$$

Hence by Eq. (D.3)

$$\begin{aligned} \mathbb{E}[\langle \Delta', k(G, G')\Delta' \rangle_{\mathbb{R}^m}^2] &\leq 4 \mathbb{E}[\|k(G, G)\|_{2;2} \|k(G', G')\|_{2;2}] \\ &= 4 \mathbb{E}[\|k(G, G)\|_{2;2}] \mathbb{E}[\|k(G', G')\|_{2;2}] = 4(\mathbb{E}[\|k(G, G)\|_{2;2}])^2, \end{aligned}$$

which implies

$$\mathbb{E}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m}^2] < \infty$$

since by assumption  $\|k\|_{\mu_G, 2;2} < \infty$ .  $\square$

Since the unbiased estimator  $\widehat{\text{SKCE}}_{\text{uq}}$  is a U-statistic, we know that the random variable  $\sqrt{n}(\widehat{\text{SKCE}}_{\text{uq}} - \text{SKCE})$  is asymptotically normally distributed under certain conditions.

**Theorem G.1.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_G, 1} < \infty$ .*

*If  $\text{Var}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m}] < \infty$ , then*

$$\sqrt{n} \left( \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \right) \xrightarrow{d} \mathcal{N}(0, 4\zeta_1),$$

where

$$\zeta_1 := \mathbb{E}[\langle \Delta, k(G, G')\Delta' \rangle_{\mathbb{R}^m} \langle \Delta, k(G, G'')\Delta'' \rangle_{\mathbb{R}^m}] - \text{SKCE}^2[k, g],$$

where  $G'$  and  $G''$  are independent copies of  $G$  and  $\Delta' := \delta(G')$  and  $\Delta'' := \delta(G'')$ .

*Proof.* The statement follows immediately from van der Vaart (Theorem 12.3 1998).  $\square$

If model  $g$  is strongly calibrated, then  $\zeta_1 = 0$ , and hence  $\widehat{\text{SKCE}}_{\text{uq}}$  is a so-called degenerate U-statistic (see, e.g., Section 12.3 van der Vaart, 1998).

**Theorem G.2.** Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_{G,2}} < \infty$ .

If  $g$  is strongly calibrated, then

$$n \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] \xrightarrow{d} \sum_{i=1}^{\infty} \lambda_i (Z_i^2 - 1),$$

where  $Z_i$  are independent standard normally distributed random variables and  $\lambda_i$  with  $\sum_{i=1}^{\infty} \lambda_i^2 < \infty$  are eigenvalues of the integral operator

$$Kf(\xi, y) := \int \langle e_y - \xi, k(\xi, \xi')(e_{y'} - \xi') \rangle_{\mathbb{R}^m} f(\xi', y') \mu_{G \times Y}(\mathrm{d}(\xi', y'))$$

on the space  $L^2(\Delta^m \times \{1, \dots, m\}, \mu_{G \times Y})$ .

*Proof.* From Lemma G.1 we know that

$$\text{Var}[\langle \Delta, k(G, G') \Delta' \rangle_{\mathbb{R}^m}] < \infty.$$

Moreover, since  $g$  is strongly calibrated,  $\Delta = 0$  almost surely and by Theorem C.1  $\text{KCE}[k, g] = 0$ . Thus we obtain

$$\mathbb{E}[\langle \Delta, k(G, G') \Delta' \rangle_{\mathbb{R}^m} \langle \Delta, k(G, G'') \Delta'' \rangle_{\mathbb{R}^m}] - \text{SKCE}^2[k, g] = 0.$$

The statement follows from Serfling (Theorem in Section 5.5.2 1980).  $\square$

As discussed by Gretton et al. (2012) in the case of two-sample tests, a natural idea is to find a threshold  $c$  such that  $\mathbb{P}[n \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] > c \mid H_0] \leq \alpha$ , where  $H_0$  is the null hypothesis that the model is strongly calibrated. The desired quantile can be estimated by fitting Pearson curves to the empirical distribution by moment matching (Gretton et al., 2012), or alternatively by bootstrapping (Arcones and Giné, 1992), both computed and performed under the assumption that the model is strongly calibrated.

If model  $g$  is strongly calibrated we know  $\mathbb{E}[\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}]] = 0$ . Moreover, it follows from Hoeffding (p. 299 1948) that

$$\mathbb{E} \left[ \widehat{\text{SKCE}}_{\text{uq}}^2[k, g, \mathcal{D}] \right] = \frac{2}{n(n-1)} \mathbb{E} \left[ (\langle e_Y - g(X), k(g(X), g(X'))(e_{Y'} - g(X')) \rangle_{\mathbb{R}^m})^2 \right],$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$ . By some tedious calculations we can retrieve higher-order moments as well. If model  $g$  is strongly calibrated, we know from Serfling (1980, Lemma B, Section 5.2.2) that for  $r \geq 2$

$$\mathbb{E} \left[ \widehat{\text{SKCE}}_{\text{uq}}^r[k, g, \mathcal{D}] \right] = O(n^{-r})$$

as the number of samples  $n$  goes to infinity, provided that

$$\mathbb{E} \left[ |\langle e_Y - g(X), k(g(X), g(X'))(e_{Y'} - g(X')) \rangle_{\mathbb{R}^m}|^r \right] < \infty.$$

Alternatively, as discussed by Arcones and Giné (Section 5 1992), we can estimate  $c$  by using quantiles of the bootstrap statistic

$$T = 2n^{-1} \sum_{1 \leq i < j \leq n} \left[ h((X_{*,i}, Y_{*,i}), (X_{*,j}, Y_{*,j})) - n^{-1} \sum_{k=1}^n h((X_{*,i}, Y_{*,i}), (X_k, Y_k)) \right. \\ \left. - n^{-1} \sum_{k=1}^n h((X_k, Y_k), (X_{*,j}, Y_{*,j})) + n^{-2} \sum_{k,l=1}^n h((X_k, Y_k), (X_l, Y_l)) \right],$$

where

$$h((x, y), (x', y')) := \langle \delta(x, y), k(g(x), g(x')) \delta(x', y') \rangle_{\mathbb{R}^m}$$

and  $(X_{*,1}, Y_{*,1}), \dots, (X_{*,n}, Y_{*,n})$  are sampled with replacement from the data set  $\mathcal{D}$ . Then asymptotically

$$\mathbb{P} \left[ n \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] > c \mid H_0 \right] \approx \mathbb{P}[T > c \mid \mathcal{D}].$$

For the linear estimator, the asymptotical behaviour follows from the central limit theorem (e.g., Theorem A in Section 1.9 Serfling, 1980).

**Corollary G.1.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_G, 1} < \infty$ .*

*If  $\sigma^2 := \text{Var}[\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}] < \infty$ , where  $(X', Y')$  is an independent copy of  $(X, Y)$ , then*

$$\sqrt{[n/2]} \left( \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \right) \xrightarrow{d} \mathcal{N}(0, \sigma^2).$$

As noted in the following statement, the variance  $\sigma^2$  is finite if  $t \mapsto \|k(t, t)\|$  is  $L^2$ -integrable with respect to measure  $\mu_G$ .

**Corollary G.2.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $\|k\|_{\mu_G, 2} < \infty$ .*

*Then  $\sigma^2 := \text{Var}[\langle \delta(X, Y), k(G, G')\delta(X', Y') \rangle_{\mathbb{R}^m}] < \infty$ , where  $(X', Y')$  is an independent copy of  $(X, Y)$  with  $G' = g(X')$ , and*

$$\sqrt{[n/2]} \left( \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \right) \xrightarrow{d} \mathcal{N}(0, \sigma^2).$$

*Proof.* The statement follows from Corollary G.1. □

The weak convergence of  $\widehat{\text{SKCE}}_{\text{ul}}$  yields the following asymptotic test.

**Corollary G.3.** *Let the assumptions of Corollary G.1 be satisfied.*

*A one-sided statistical test with test statistic  $\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}]$  and asymptotic significance level  $\alpha$  has the acceptance region*

$$\sqrt{[n/2]} \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] < z_{1-\alpha} \hat{\sigma},$$

*where  $z_{1-\alpha}$  is the  $(1-\alpha)$ -quantile of the standard normal distribution and  $\hat{\sigma}$  is a consistent estimator of the standard deviation of  $\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}$ .*

## H Distribution-free bounds

First we prove a helpful bound.

**Lemma H.1.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $K_{p;q} := \sup_{s,t \in \Delta^m} \|k(s, t)\|_{p;q} < \infty$  for some  $1 \leq p, q \leq \infty$ . Then*

$$\sup_{x, x' \in \mathcal{X}, y, y' \in \{1, \dots, m\}} |\langle \delta(x, y), k(g(x), g(x'))\delta(x', y') \rangle_{\mathbb{R}^m}| \leq 2^{1+1/p-1/q} K_{p;q} =: B_{p;q}.$$

*Proof.* By Hölder's inequality and the definition of the operator norm for all  $s, t \in \Delta^m$  and  $u, v \in \mathbb{R}^m$

$$|\langle u, k(s, t)v \rangle_{\mathbb{R}^m}| \leq \|u\|_{q'} \|k(s, t)v\|_q \leq \|u\|_{q'} \|v\|_p \|k(s, t)\|_{p;q} \leq K_{p;q} \|u\|_{q'} \|v\|_p.$$

The result follows from the fact that  $\max_{s,t \in \Delta^m} \|s - t\|_p = 2^{1/p}$  and  $\max_{s,t \in \Delta^m} \|s - t\|_{q'} = 2^{1/q'} = 2^{1-1/q}$ . □

Unfortunately, the tightness of the bound in Lemma H.1 depends on the choice of  $p$  and  $q$ , as the following example shows.

**Example H.1**

Let  $k = \phi \mathbf{I}_m$ , where  $\phi: \Delta^m \times \Delta^m \rightarrow \mathbb{R}$  is a scalar-valued kernel and  $\mathbf{I}_m \in \mathbb{R}^{m \times m}$  is the identity matrix. Assume that  $\Phi := \sup_{s,t \in \Delta^m} |\phi(s,t)| < \infty$ . One can show that for all  $s, t \in \Delta^m$

$$\|k(s,t)\|_{p;q} = \begin{cases} \phi(s,t) & \text{if } p \leq q, \\ m^{1/q-1/p} \phi(s,t) & \text{if } p > q, \end{cases}$$

which implies that

$$K_{p;q} = \begin{cases} \Phi & \text{if } p \leq q, \\ m^{1/q-1/p} \Phi & \text{if } p > q. \end{cases}$$

Thus the bound  $B_{p;q}$  in Lemma H.1 is

$$B_{p;q} = \begin{cases} 2^{1+1/p-1/q} \Phi & \text{if } p \leq q, \\ 2^{1+1/p-1/q} m^{1/q-1/p} \Phi & \text{if } p > q, \end{cases}$$

which attains its smallest value  $\min_{1 \leq p, q \leq \infty} B_{p;q} = 2\Phi$  if and only if  $p = q$  or  $m = 2$  and  $p > q$ . Thus for any other choice of  $p$  and  $q$  Lemma H.1 provides a non-optimal bound.

**Theorem H.1.** Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $K_{p;q} := \sup_{s,t \in \Delta^m} \|k(s,t)\|_{p;q} < \infty$  for some  $1 \leq p, q \leq \infty$ . Then for all  $\epsilon > 0$

$$\mathbb{P} \left[ \left| \widehat{\text{KCE}}[k, g, \mathcal{D}] - \text{KCE}[k, g] \right| \geq 2(B_{p;q}/n)^{1/2} + \epsilon \right] \leq \exp \left( -\frac{\epsilon^2 n}{2B_{p;q}} \right).$$

*Proof.* Let  $\mathcal{F}$  be the unit ball in the RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  corresponding to kernel  $k$ . We consider the random variable

$$F := \left| \widehat{\text{KCE}}[k, g, \mathcal{D}] - \text{KCE}[k, g] \right|.$$

The randomness of  $F$  is due to the randomness of the data points  $(X_i, Y_i)$ , and by Lemmas E.2 and F.1 we can rewrite  $F$  as

$$F = n^{-1} \left\| \left\| \sum_{i=1}^n k(\cdot, g(X_i)) \delta(X_i, Y_i) \right\|_{\mathcal{H}} - n \|\mu_g\|_{\mathcal{H}} \right\| =: f((X_1, Y_1), \dots, (X_n, Y_n)),$$

where  $\mu_g$  is the embedding defined in Lemma E.1. The triangle inequality implies that for all  $z_i = (x_i, y_i) \in \mathcal{X} \times \{1, \dots, m\}$

$$\begin{aligned} f(z_1, \dots, z_n) &= n^{-1} \left\| \left\| \sum_{i=1}^n k(\cdot, g(x_i)) \delta(x_i, y_i) \right\|_{\mathcal{H}} - n \|\mu_g\|_{\mathcal{H}} \right\| \\ &\leq n^{-1} \left\| \sum_{i=1}^n (k(\cdot, g(x_i)) \delta(x_i, y_i) - \mu_g) \right\|_{\mathcal{H}} =: h(z_1, \dots, z_n), \end{aligned} \tag{H.1}$$

where  $h: (\mathcal{X} \times \{1, \dots, m\})^n \rightarrow \mathbb{R}$  is measurable and hence induces a random variable  $H := h((X_1, Y_1), \dots, (X_n, Y_n))$ .

By the reproducing property and Lemma H.1, for all  $x, x' \in \mathcal{X}$  and  $y, y' \in \{1, \dots, m\}$  we have

$$\begin{aligned} \|k(\cdot, g(x)) \delta(x, y) - k(\cdot, g(x')) \delta(x', y')\|_{\mathcal{H}}^2 &= \langle \delta(x, y), k(g(x), g(x)) \delta(x, y) \rangle_{\mathbb{R}^m} \\ &\quad - \langle \delta(x, y), k(g(x), g(x')) \delta(x', y') \rangle_{\mathbb{R}^m} \\ &\quad - \langle \delta(x', y'), k(g(x'), g(x)) \delta(x, y) \rangle_{\mathbb{R}^m} \\ &\quad + \langle \delta(x', y'), k(g(x'), g(x')) \delta(x', y') \rangle_{\mathbb{R}^m} \\ &\leq 4B_{p;q}. \end{aligned}$$

Thus for all  $i \in \{1, \dots, m\}$  the triangle inequality implies

$$\begin{aligned} &\sup_{z, z', z_j (j \neq i)} |h(z_1, \dots, z_{i-1}, z, z_{i+1}, \dots, z_n) - h(z_1, \dots, z_{i-1}, z', z_{i+1}, \dots, z_n)| \\ &\leq \sup_{x, x', y, y'} n^{-1} \|k(\cdot, g(x)) \delta(x, y) - k(\cdot, g(x')) \delta(x', y')\|_{\mathcal{H}} \leq \frac{2B_{p;q}^{1/2}}{n}. \end{aligned}$$

Hence we can apply McDiarmid's inequality to the random variable  $H$ , which yields for all  $\epsilon > 0$

$$\mathbb{P}[H \geq \mathbb{E}[H] + \epsilon] \leq \exp\left(-\frac{\epsilon^2 n}{2B_{p;q}}\right). \quad (\text{H.2})$$

In the final parts of the proof we bound the expectation  $\mathbb{E}[H]$ . By Lemmas E.1 and F.1, we know that

$$\begin{aligned} H &= h((X_1, Y_1), \dots, (X_n, Y_n)) \\ &= \sup_{f \in \mathcal{F}} n^{-1} \left| \sum_{i=1}^n \left( \langle \delta(X_i, Y_i), f(g(X_i)) \rangle_{\mathbb{R}^m} - \mathbb{E}[\langle \delta(X, Y), f(g(X)) \rangle_{\mathbb{R}^m}] \right) \right| \\ &= \sup_{f \in \mathcal{F}_0} n^{-1} \left| \sum_{i=1}^n f(X_i, Y_i) - \mathbb{E}[f(X, Y)] \right|, \end{aligned}$$

where  $\mathcal{F}_0 := \{f: \mathcal{X} \times \{1, \dots, m\} \rightarrow \mathbb{R}, (x, y) \mapsto \langle \delta(x, y), \tilde{f}(g(x)) \rangle_{\mathbb{R}^m} : \tilde{f} \in \mathcal{F}\}$  is a class of measurable functions. As Gretton et al. (2012), we make use of symmetrization ideas (van der Vaart and Wellner, 1996, p. 108). From van der Vaart and Wellner (1996, Lemma 2.3.1) it follows that

$$\mathbb{E}[H] = \mathbb{E} \left[ \sup_{f \in \mathcal{F}_0} n^{-1} \left| \sum_{i=1}^n f(X_i, Y_i) - \mathbb{E}[f(X, Y)] \right| \right] \leq 2 \mathbb{E} \left[ \sup_{f \in \mathcal{F}_0} \left| n^{-1} \sum_{i=1}^n \epsilon_i f(X_i, Y_i) \right| \right],$$

where  $\epsilon_1, \dots, \epsilon_n$  are independent Rademacher random variables. Similar to Bartlett and Mendelson (2002, Lemma 22), we obtain

$$\begin{aligned} \mathbb{E}[H] &\leq 2n^{-1} \mathbb{E} \left[ \sup_{f \in \mathcal{F}} \left| \sum_{i=1}^n \epsilon_i \langle \delta(X_i, Y_i), f(g(X_i)) \rangle_{\mathbb{R}^m} \right| \right] \\ &= 2n^{-1} \mathbb{E} \left[ \sup_{f \in \mathcal{F}} \left| \left\langle \sum_{i=1}^n \epsilon_i k(\cdot, g(X_i)) \delta(X_i, Y_i), f \right\rangle_{\mathcal{H}} \right| \right] \\ &= 2n^{-1} \mathbb{E} \left[ \left\| \sum_{i=1}^n \epsilon_i k(\cdot, g(X_i)) \delta(X_i, Y_i) \right\|_{\mathcal{H}} \right] \\ &= 2n^{-1} \mathbb{E} \left[ \left( \sum_{i,j=1}^n \epsilon_i \epsilon_j \langle k(\cdot, g(X_i)) \delta(X_i, Y_i), k(\cdot, g(X_j)) \delta(X_j, Y_j) \rangle_{\mathcal{H}} \right)^{1/2} \right]. \end{aligned}$$

By Jensen's inequality we get

$$\begin{aligned} \mathbb{E}[H] &\leq 2n^{-1} \left( \sum_{i,j=1}^n \mathbb{E} [\epsilon_i \epsilon_j \langle k(\cdot, g(X_i)) \delta(X_i, Y_i), k(\cdot, g(X_j)) \delta(X_j, Y_j) \rangle_{\mathcal{H}}] \right)^{1/2} \\ &= 2n^{-1/2} \left( \mathbb{E} [\langle k(\cdot, g(X)) \delta(X, Y), k(\cdot, g(X)) \delta(X, Y) \rangle_{\mathcal{H}}] \right)^{1/2} \\ &\leq 2(B_{p;q}/n)^{1/2}. \end{aligned} \quad (\text{H.3})$$

All in all, from Eqs. (H.1) to (H.3) we obtain for all  $\epsilon > 0$

$$\begin{aligned} \mathbb{P} \left[ \left| \widehat{\text{KCE}}[k, g, \mathcal{D}] - \text{KCE}[k, g] \right| \geq 2(B_{p;q}/n)^{1/2} + \epsilon \right] &= \mathbb{P}[F \geq 2(B_{p;q}/n)^{1/2} + \epsilon] \\ &\leq \mathbb{P}[H \geq 2(B_{p;q}/n)^{1/2} + \epsilon] \\ &\leq \mathbb{P}[H \geq \mathbb{E}[H] + \epsilon] \\ &\leq \exp\left(-\frac{\epsilon^2 n}{2B_{p;q}}\right), \end{aligned}$$

which concludes our proof.  $\square$

If model  $g$  is calibrated in the strong sense, we can improve the bound.

**Theorem H.2.** Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $K_{p;q} := \sup_{s,t \in \Delta^m} \|k(s,t)\|_{p;q} < \infty$  for some  $1 \leq p, q \leq \infty$ . Define

$$B_1 := n^{-1/2} [\mathbb{E} [\langle \delta(X, Y), k(g(X), g(X))\delta(X, Y) \rangle_{\mathbb{R}^m}]]^{1/2}, \quad \text{and}$$

$$B_2 := (B_{p;q}/n)^{1/2}.$$

Then  $B_1 \leq B_2$ , and for all  $\epsilon > 0$  and  $i \in \{1, 2\}$

$$\mathbb{P} \left[ \widehat{\text{KCE}}[k, g, \mathcal{D}] \geq B_i + \epsilon \right] \leq \exp \left( -\frac{\epsilon^2 n}{2B_{p;q}} \right),$$

if  $g$  is calibrated in the strong sense.

*Proof.* Let  $\mathcal{F}$  be the unit ball in the RKHS  $\mathcal{H} \subset (\mathbb{R}^m)^{\Delta^m}$  corresponding to kernel  $k$ . Lemma H.1 implies

$$\begin{aligned} B_1 &= n^{-1/2} [\mathbb{E} [\langle \delta(X, Y), k(g(X), g(X))\delta(X, Y) \rangle_{\mathbb{R}^m}]]^{1/2} \\ &\leq n^{-1/2} [\mathbb{E}[B_{p;q}]]^{1/2} = (B_{p;q}/n)^{1/2} = B_2. \end{aligned} \quad (\text{H.4})$$

Let  $H$  be defined as in the proof of Theorem H.1. Since  $g$  is strongly calibrated, it follows from Theorem C.1 and Lemma E.2 that  $\mu_g = 0$ , and thus by Lemma F.1

$$H = n^{-1} \left\| \sum_{i=1}^n k(\cdot, g(x_i))\delta(x_i, y_i) \right\|_{\mathcal{H}} = \widehat{\text{KCE}}[k, g, \mathcal{D}].$$

Thus Eq. (H.2) implies

$$\mathbb{P} \left[ \widehat{\text{KCE}}[k, g, \mathcal{D}] \geq \mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] + \epsilon \right] \leq \exp \left( -\frac{\epsilon^2 n}{2B_{p;q}} \right). \quad (\text{H.5})$$

Next we bound  $\mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]]$ . From Lemma F.1 we get

$$\mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] = \frac{1}{n} \mathbb{E} \left[ \left( \sum_{i,j=1}^n \langle \delta(X_i, Y_i), k(g(X_i), g(X_j))\delta(X_j, Y_j) \rangle_{\mathbb{R}^m} \right)^{1/2} \right],$$

and hence by Jensen's inequality we obtain

$$\begin{aligned} \mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] &\leq \frac{1}{n} \left( \mathbb{E} \left[ \sum_{i,j=1}^n \langle \delta(X_i, Y_i), k(g(X_i), g(X_j))\delta(X_j, Y_j) \rangle_{\mathbb{R}^m} \right] \right)^{1/2} \\ &= \frac{1}{n} \left( n \mathbb{E} [\langle \delta(X, Y), k(g(X), g(X))\delta(X, Y) \rangle_{\mathbb{R}^m}] \right. \\ &\quad \left. + n(n-1) \mathbb{E} [\langle \delta(X, Y), k(g(X), g(X'))\delta(X', Y') \rangle_{\mathbb{R}^m}] \right)^{1/2}, \end{aligned}$$

where  $(X', Y')$  denotes an independent copy of  $(X, Y)$ . From Lemma E.2 it follows that

$$\mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] \leq \left( \frac{1}{n} \mathbb{E} [\langle \delta(X, Y), k(g(X), g(X))\delta(X, Y) \rangle_{\mathbb{R}^m}] + \left( 1 - \frac{1}{n} \right) \text{SKCE}[k, g] \right)^{1/2}.$$

If model  $g$  is calibrated in the strong sense, we know from Theorem C.1 that  $\text{SKCE}[k, g] = 0$ . Thus we obtain

$$\mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] \leq B_1. \quad (\text{H.6})$$



All in all, from Eqs. (H.4) to (H.6) it follows that for all  $\epsilon > 0$  and  $i \in \{1, 2\}$

$$\begin{aligned} \mathbb{P} \left[ \widehat{\text{KCE}}[k, g, \mathcal{D}] \geq B_i + \epsilon \right] &\leq \mathbb{P} \left[ \widehat{\text{KCE}}[k, g, \mathcal{D}] \geq B_1 + \epsilon \right] \\ &\leq \mathbb{P} \left[ \widehat{\text{KCE}}[k, g, \mathcal{D}] \geq \mathbb{E}[\widehat{\text{KCE}}[k, g, \mathcal{D}]] + \epsilon \right] \\ &\leq \exp \left( -\frac{\epsilon^2 n}{2B_{p;q}} \right), \end{aligned}$$

if  $g$  is calibrated in the strong sense.  $\square$

Thus we obtain the following distribution-free hypothesis test.

**Corollary H.1.** *Let the assumptions of Theorem H.2 be satisfied.*

A statistical test with test statistic  $\widehat{\text{KCE}}[k, g, \mathcal{D}]$  and significance level  $\alpha$  for the null hypothesis of model  $g$  being calibrated in the strong sense has the acceptance region

$$\widehat{\text{KCE}}[k, g, \mathcal{D}] < (B_{p;q}/n)^{1/2} (1 + \sqrt{-2 \log \alpha}).$$

A distribution-free bound for the deviation of the unbiased estimator can be obtained from the theory of U-statistics.

**Theorem H.3.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $K_{p;q} := \sup_{s,t \in \Delta^m} \|k(s, t)\|_{p;q}$  for some  $1 \leq p, q \leq \infty$ . Then for all  $t > 0$*

$$\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \geq t \right] \leq \exp \left( -\frac{[n/2]t^2}{2B_{p;q}^2} \right).$$

The same bound holds for  $\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \leq -t \right]$ .

*Proof.* By Lemma F.2,  $\mathbb{E}[\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}]] = \text{SKCE}[k, g]$ . Moreover, by Lemma H.1 we know that

$$\sup_{x, x' \in \mathcal{X}, y, y' \in \{1, \dots, m\}} |\langle \delta(x, y), k(g(x), g(x'))\delta(x', y') \rangle_{\mathbb{R}^m}| \leq B_{p;q}.$$

Thus the result follows from the bound on U-statistics by Hoeffding (1963, p. 25).  $\square$

We can derive a hypothesis test using the unbiased estimator.

**Corollary H.2.** *Let the assumptions of Theorem H.3 be satisfied.*

A one-sided statistical test with test statistic  $\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}]$  and significance level  $\alpha$  for the null hypothesis of model  $g$  being calibrated in the strong sense has the acceptance region

$$\widehat{\text{SKCE}}_{\text{uq}}[k, g, \mathcal{D}] < \frac{B_{p;q}}{\sqrt{[n/2]}} \sqrt{-2 \log \alpha}.$$

Analogously we can obtain a bound for the linear estimator.

**Theorem H.4.** *Let  $k: \Delta^m \times \Delta^m \rightarrow \mathcal{L}(\mathbb{R}^m)$  be a kernel, and assume that  $k(\cdot, t)u$  is measurable for all  $t \in \Delta^m$  and  $u \in \mathbb{R}^m$ , and  $K_{p;q} := \sup_{s,t \in \Delta^m} \|k(s, t)\|_{p;q}$  for some  $1 \leq p, q \leq \infty$ . Then for all  $t > 0$*

$$\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \geq t \right] \leq \exp \left( -\frac{[n/2]t^2}{2B_{p;q}^2} \right).$$

The same bound holds for  $\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \leq -t \right]$ .

*Proof.* By Lemma F.3,  $\mathbb{E}[\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}]] = \text{SKCE}[k, g]$ . Moreover, by Lemma H.1 we know that

$$\sup_{x, x' \in \mathcal{X}, y, y' \in \{1, \dots, m\}} |\langle \delta(x, y), k(g(x), g(x'))\delta(x', y') \rangle_{\mathbb{R}^m}| \leq B_{p;q}.$$

Thus by Hoeffding's inequality (Hoeffding, 1963, Theorem 2) for all  $t > 0$

$$\mathbb{P} \left[ \widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] - \text{SKCE}[k, g] \geq t \right] \leq \exp \left( -\frac{[n/2]t^2}{2B_{p;q}^2} \right). \quad \square$$

Obviously this results yields another distribution-free hypothesis test.

**Corollary H.3.** *Let the assumptions of Theorem H.4 be satisfied.*

*A one-sided statistical test with test statistic  $\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}]$  and significance level  $\alpha$  for the null hypothesis of model  $g$  being calibrated in the strong sense has the acceptance region*

$$\widehat{\text{SKCE}}_{\text{ul}}[k, g, \mathcal{D}] < \frac{B_{p;q}}{\sqrt{\lfloor n/2 \rfloor}} \sqrt{-2 \log \alpha}.$$

## I Comparisons

### I.1 Expected calibration error and maximum calibration error

For certain spaces of bounded functions the calibration error CE turns out to be a form of the ECE. In particular, the ECE with respect to the cityblock distance, the total variation distance, and the squared Euclidean distance are special cases of CE. Choosing  $p = 1$  in the following statement corresponds to the special case of the MCE.

**Lemma I.1 (ECE and MCE as special cases).** *Let  $1 \leq p \leq \infty$  with Hölder conjugate  $p'$ . If  $\mathcal{F} = K^p(\Delta^m, \mu_G; \mathbb{R}^m)$ , then  $\text{CE}[\mathcal{F}, g] = \|\delta\|_{\mu_G, p'}$ .*

*Proof.* Note that  $\mathcal{F}$  is well defined since  $\mathcal{F} \subset L^1(\Delta^m, \mu_G; \mathbb{R}^m)$ .

The statement follows from the extremal case of Hölder’s inequality. More explicitly, let  $\nu$  denote the counting measure on  $\{1, \dots, m\}$ . Since both  $\mu_G$  and  $\nu$  are  $\sigma$ -finite measures, the product measure  $\mu_G \otimes \nu$  on the product space  $B := \Delta^m \times \{1, \dots, m\}$  is uniquely defined and  $\sigma$ -finite. Define  $\tilde{\delta}(t, k) := \delta_k(t)$  for all  $(t, k) \in B$ . Then we can rewrite

$$\begin{aligned} \text{CE}[\mathcal{F}, g] &= \sup_{f \in K^p(\Delta^m, \mu_G; \mathbb{R}^m)} \int_{\Delta^m} \langle \delta(x), f(x) \rangle_{\mathbb{R}^m} \mu_G(\text{d}x) \\ &= \sup_{f \in K^p(B, \mu_G \otimes \nu; \mathbb{R}^m)} \int_B |\tilde{\delta}(x, k) f(x, k)| (\mu_G \times \nu)(\text{d}(x, k)) \\ &= \|\tilde{\delta}\|_{\mu_G \otimes \nu, p'} = \|\delta\|_{\mu_G, p'}, \end{aligned}$$

to make the reasoning more apparent. Since  $\mu_G \otimes \nu$  is  $\sigma$ -finite the statement holds even for  $p = 1$ .  $\square$

### I.2 Maximum mean calibration error

The so-called “correctness score” (Kumar et al. (2018))  $c(x, y)$  of an input  $x$  and a class  $y$  is defined as  $c(x, y) = \mathbb{1}_{\{\arg \max_{y'} g_{y'}(x)\}}(y)$ . It is 1 if class  $y$  is equal to the class that is most likely for input  $x$  according to model  $g$ , and 0 otherwise. Let  $k: [0, 1] \times [0, 1] \rightarrow \mathbb{R}$  be a scalar-valued kernel. Then the maximum mean calibration error  $\text{MMCE}[k, g]$  of a model  $g$  with respect to kernel  $k$  is defined<sup>6</sup> as

$$\begin{aligned} \text{MMCE}[k, g] &= \left( \mathbb{E}[(c(X, Y) - g_{\max}(X))(c(X', Y') - g_{\max}(X'))k(g_{\max}(X), g_{\max}(X'))] \right)^{1/2}, \end{aligned}$$

where  $(X', Y')$  is an independent copy of  $(X, Y)$ .

Example I.1 shows that the KCE allows exactly the same analysis of the common notion of calibration as the MMCE proposed by Kumar et al. (2018) by applying it to a model that is reduced to the most confident predictions.

#### Example I.1 (MMCE as special case)

Reduce model  $g$  to its most confident predictions by defining a new model  $\tilde{g}$  with  $\tilde{g}(x) := (g_{\max}(x), 1 - g_{\max}(x))$ . The predictions  $\tilde{g}(x)$  of this new model can be viewed as a model of

<sup>6</sup>For illustrative purposes we present a variation of the original definition of the MMCE by Kumar et al. (2018).

the conditional probabilities ( $\mathbb{P}[\tilde{Y} = 1 \mid X = x], \mathbb{P}[\tilde{Y} = 2 \mid X = x]$ ) in a classification problem with inputs  $X$  and classes  $\tilde{Y} := 2 - c(X, Y)$ .<sup>7</sup>

Let  $k: [0, 1] \times [0, 1] \rightarrow \mathbb{R}$  be a scalar-valued kernel. Define a matrix-valued function  $\tilde{k}: \Delta^2 \times \Delta^2 \rightarrow \mathbb{R}^{2 \times 2}$  by

$$\tilde{k}((p_1, p_2), (q_1, q_2)) = \frac{k(p_1, q_1)}{2} \mathbf{I}_2.$$

Then by Caponnetto et al. (2008, Example 1 and Theorem 14)  $\tilde{k}$  is a matrix-valued kernel and, if it is continuous, it is universal if and only if  $k$  is universal. By construction  $e_{\tilde{Y}} - \tilde{g}(X) = (c(X, Y) - g_{\max}(X))(1, -1)$ , and hence

$$\begin{aligned} \text{SKCE}[\tilde{k}, \tilde{g}] &= \mathbb{E}[(e_{\tilde{Y}} - \tilde{g}(X))^{\top} \tilde{k}(\tilde{g}(X), \tilde{g}(X'))(e_{\tilde{Y}} - \tilde{g}(X'))] \\ &= \mathbb{E}[(c(X, Y) - g_{\max}(X))(c(X', Y') - g_{\max}(X'))k(g_{\max}(X), g_{\max}(X'))] \\ &= \text{MMCE}^2[k, g], \end{aligned}$$

where  $(X', \tilde{Y}')$  and  $(X', Y')$  are independent copies of  $(X, \tilde{Y})$  and  $(X, Y)$ , respectively.

---

<sup>7</sup>In the words of Vaicenavicius et al. (2019),  $\tilde{g}$  is induced by the maximum calibration lens.

## J Experiments

The Julia implementation for all experiments is available online at <https://github.com/devmotion/CalibrationPaper>. The code is written and documented with the literate programming tool Weave.jl Pastell (2017) and exported to HTML files that include results and figures.

### J.1 Calibration errors

In our experiments we evaluate the proposed estimators of the SKCE and compare them with two estimators of the ECE.

#### J.1.1 Expected calibration error

As commonly done (Bröcker and Smith, 2007; Guo et al., 2017; Vaicenavicius et al., 2019), we study the ECE with respect to the total variation distance.

The standard histogram-regression estimator of the ECE is based on a partitioning of the probability simplex (Guo et al., 2017; Vaicenavicius et al., 2019). In our experiments we use two different partitioning schemes. The first scheme is the commonly used partitioning into bins of uniform size, based on splitting the predictions of each class into 10 bins. The other partitioning is data-dependent: the data set is split iteratively along the median of the class predictions with the highest variance as long as the number of samples in a bin is at least 10.

#### J.1.2 Kernel calibration error

We consider the matrix-valued kernel  $k(x, y) = \exp(-\|x - y\|/\nu)\mathbf{I}_m$  with kernel bandwidth  $\nu > 0$ . Analogously to the ECE, we take the total variation distance as distance measure. Moreover, we choose the bandwidth adaptively with the so-called median heuristic. The median heuristic is a common heuristic that proposes to set the bandwidth to the median of the pairwise distances of samples in a, not necessarily separate, validation data set (see, e.g., Gretton et al., 2012).

### J.2 Generative models

Since the considered calibration errors depend only on the predictions and labels, we specify generative models of labeled predictions  $(g(X), Y)$  without considering  $X$ . Instead we only specify the distribution of the predictions  $g(X)$  and the conditional distribution of  $Y$  given  $g(X) = g(x)$ . This setup allows us to design calibrated and uncalibrated models in a straightforward way, which enables clean numerical evaluations with known calibration errors.

We study the generative model

$$\begin{aligned}g(X) &\sim \text{Dir}(\alpha), \\Z &\sim \text{Ber}(\pi), \\Y | Z = 1, g(X) = \gamma &\sim \text{Cat}(\beta), \\Y | Z = 0, g(X) = \gamma &\sim \text{Cat}(\gamma),\end{aligned}$$

with parameters  $\alpha \in \mathbb{R}_{>0}^m$ ,  $\beta \in \Delta^m$ , and  $\pi \in [0, 1]$ . The model is calibrated if and only if  $\pi = 0$ , since for all labels  $y \in \{1, \dots, m\}$  we obtain

$$\mathbb{P}[Y = y | g(X)] = \pi\beta_y + (1 - \pi)g_y(X),$$

and hence  $\Delta = \pi(\beta - g(X)) = 0$  almost surely if and only if  $\pi = 0$ .

By setting  $\alpha = (1, \dots, 1)$  we can model uniformly distributed predictions, and by decreasing the magnitude of  $\alpha$  we can push the predictions towards the edges of the probability simplex, mimicking the predictions of a trained model (cf., e.g., Vaicenavicius et al., 2019).

### J.2.1 Theoretical expected calibration error

For the considered model, the ECE with respect to the total variation distance is

$$\begin{aligned} \text{ECE}[\|\cdot\|_{\text{TV}}, g] &= \mathbb{E}[\|\Delta\|_{\text{TV}}] = \pi \mathbb{E}[\|\beta - g(X)\|_{\text{TV}}] = \pi/2 \sum_{i=1}^m \mathbb{E}[|\beta_i - g_i(X)|] \\ &= \frac{\pi}{2} \sum_{i=1}^m \left( \left( \frac{\alpha_i}{\alpha_0} - \beta_i \right) \left( 1 - \frac{2B(\beta_i; \alpha_i, \alpha_0 - \alpha_i)}{B(\alpha_i, \alpha_0 - \alpha_i)} \right) + \frac{2\beta_i^{\alpha_i}(1 - \beta_i)^{\alpha_0 - \alpha_i}}{\alpha_0 B(\alpha_i, \alpha_0 - \alpha_i)} \right), \end{aligned}$$

where  $\alpha_0 := \sum_{i=1}^m \alpha_i$  and  $B(x; a, b)$  denotes the incomplete Beta function  $\int_0^x t^{a-1}(1-t)^{b-1} dt$ . By exploiting  $\sum_{i=1}^m \beta_i = 1$ , we get

$$\text{ECE}[\|\cdot\|_{\text{TV}}, g] = \frac{\pi}{\alpha_0} \sum_{i=1}^m \frac{(\alpha_0 \beta_i - \alpha_i) B(\beta_i; \alpha_i, \alpha_0 - \alpha_i) + \beta_i^{\alpha_i} (1 - \beta_i)^{\alpha_0 - \alpha_i}}{B(\alpha_i, \alpha_0 - \alpha_i)}.$$

Let  $I(x; a, b) := B(x; a, b)/B(a, b)$  denote the regularized incomplete Beta function. Due to the identity  $x^a(1-x)^b/B(a, b) = a(I(x; a, b) - I(x; a+1, b))$ , we obtain

$$\text{ECE}[\|\cdot\|_{\text{TV}}, g] = \pi \sum_{i=1}^m \left( \beta_i I(\beta_i; \alpha_i, \alpha_0 - \alpha_i) - \frac{\alpha_i}{\alpha_0} I(\beta_i; \alpha_i + 1, \alpha_0 - \alpha_i) \right).$$

If  $\alpha = (a, \dots, a)$  for some  $a > 0$ , then

$$\text{ECE}[\|\cdot\|_{\text{TV}}, g] = \pi \sum_{i=1}^m (\beta_i I(\beta_i; a, (m-1)a) - m^{-1} I(\beta_i; a+1, (m-1)a)).$$

If  $\beta = e_j$  for some  $j \in \{1, \dots, m\}$  we get

$$\begin{aligned} \text{ECE}[\|\cdot\|_{\text{TV}}, g] &= \pi (I(1; a, (m-1)a) - m^{-1} I(1; a+1, (m-1)a)) \\ &= \pi(1 - m^{-1}) = \frac{\pi(m-1)}{m}, \end{aligned}$$

whereas if  $\beta = (1/m, \dots, 1/m)$  we obtain

$$\begin{aligned} \text{ECE}[\|\cdot\|_{\text{TV}}, g] &= \pi (I(m^{-1}; a, (m-1)a) - I(m^{-1}; a+1, (m-1)a)) \\ &= \pi \frac{m^{-a}(1-m^{-1})^{(m-1)a}}{aB(a, (m-1)a)} = \frac{\pi}{aB(a, (m-1)a)} \left( \frac{(m-1)^{m-1}}{m^m} \right)^a. \end{aligned}$$

We see that, as the number of classes goes to infinity, the ECE with respect to the total variation distance tends to  $\pi$  and  $\pi \exp(-a)a^{a-1}/\Gamma(a)$ , respectively.

### J.2.2 Mean total variation distance

For the considered generative models, we can compute the mean total variation distance  $\mathbb{E}[\|X - X'\|_{\text{TV}}]$ , which does not depend on the number of available samples (but, of course, is usually not available). If  $X$  and  $X'$  are i.i.d. according to  $\text{Dir}(\alpha)$  with parameter  $\alpha \in \mathbb{R}_{>0}^m$ , then their mean total variation distance is

$$\begin{aligned} \mathbb{E}[\|X - X'\|_{\text{TV}}] &= 1/2 \sum_{i=1}^m \mathbb{E}[|X_i - X'_i|] \\ &= \sum_{i=1}^m \mathbb{E}[X_i - X'_i | X_i > X'_i] \\ &= \frac{2B(\alpha_0, \alpha_0)}{\alpha_0} \sum_{i=1}^m [B(\alpha_i, \alpha_i) B(\alpha_0 - \alpha_i, \alpha_0 - \alpha_i)]^{-1}, \end{aligned}$$

where  $\alpha_0 := \sum_{i=1}^m \alpha_i$ . We conduct additional experiments in which we set the kernel bandwidth to the mean total variation distance.

### J.2.3 Distribution of estimates: Additional figures

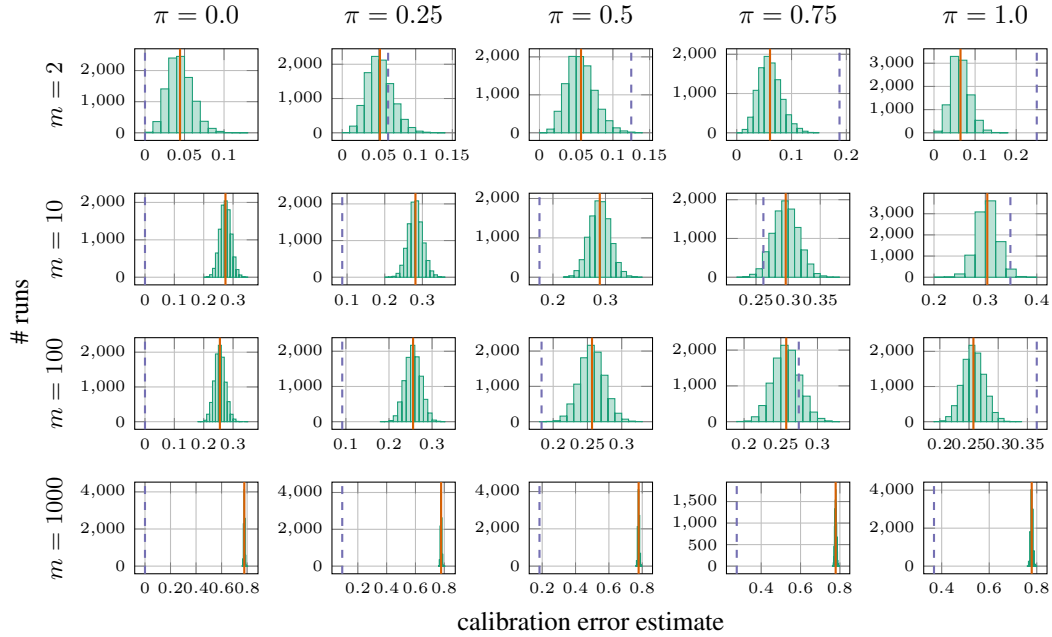


Figure 3: Distribution of  $\widehat{\text{ECE}}$  with bins of uniform size, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

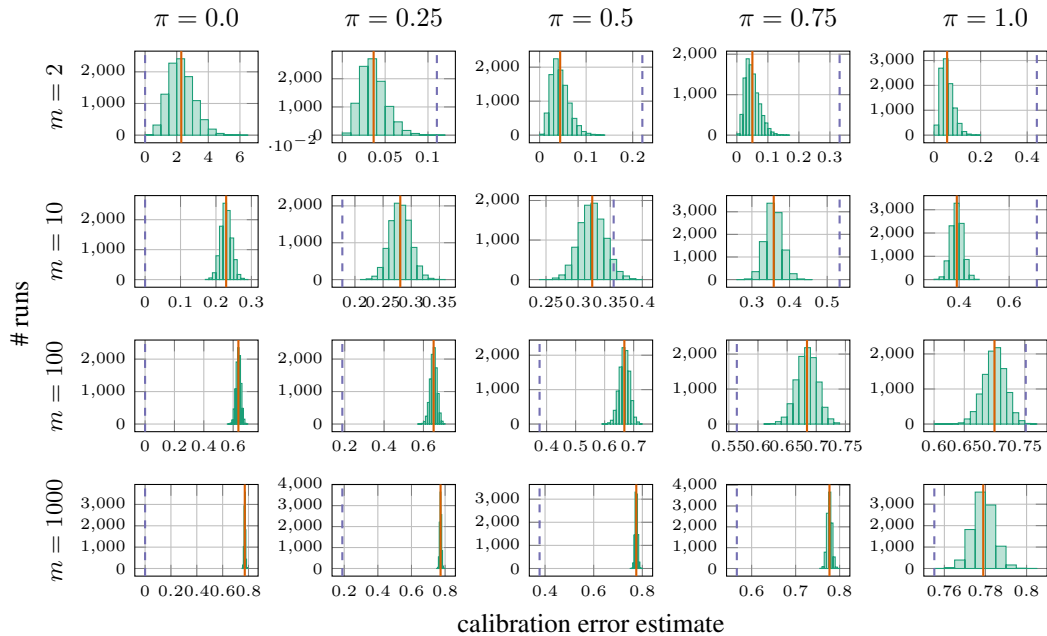


Figure 4: Distribution of  $\widehat{\text{ECE}}$  with bins of uniform size, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

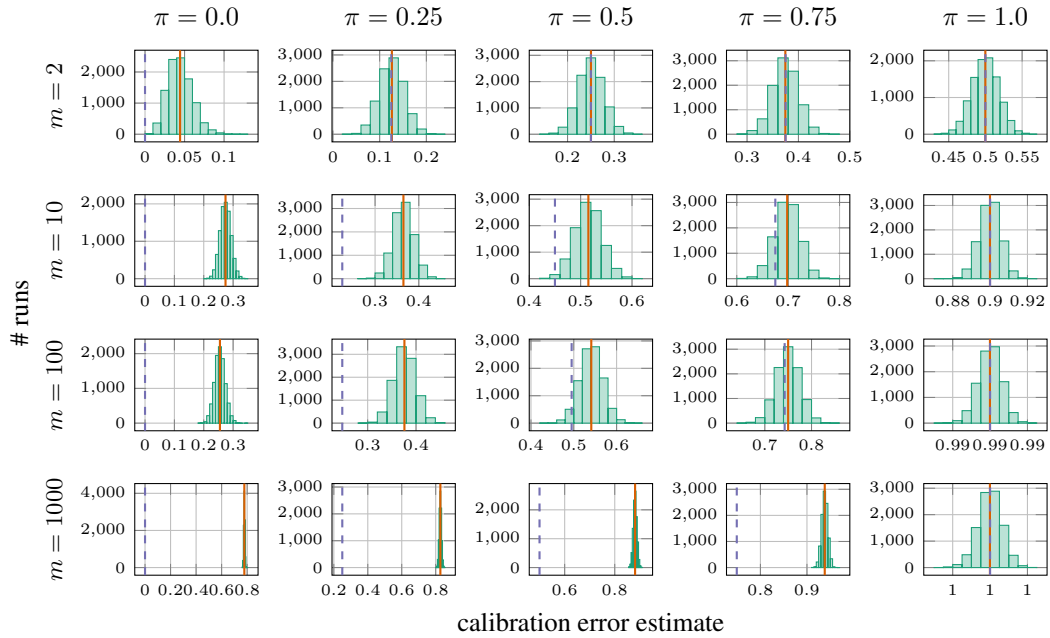


Figure 5: Distribution of  $\widehat{\text{ECE}}$  with bins of uniform size, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

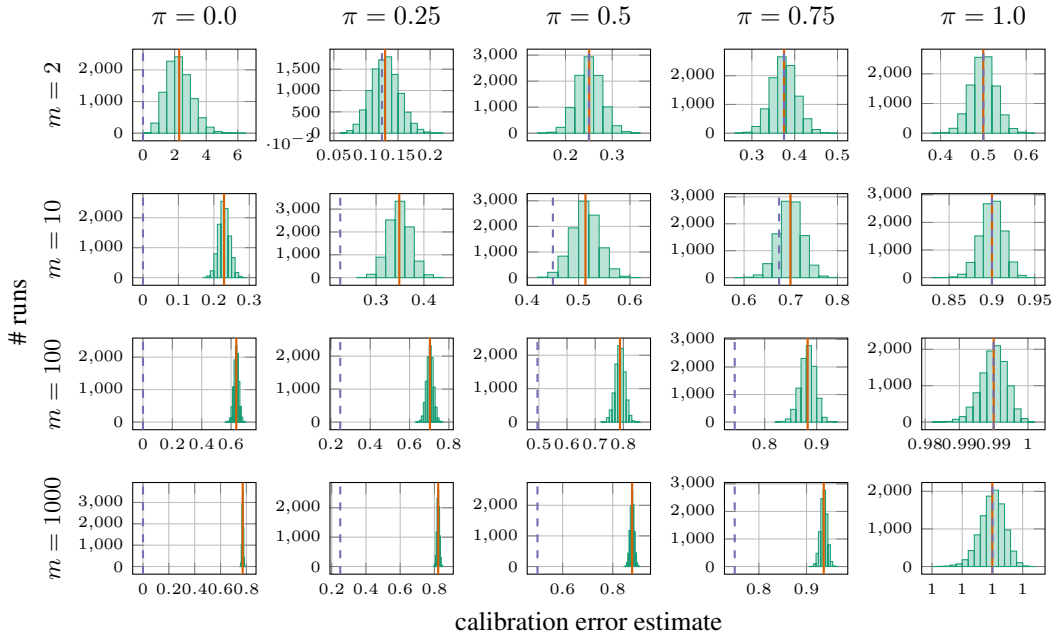


Figure 6: Distribution of  $\widehat{\text{ECE}}$  with bins of uniform size, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

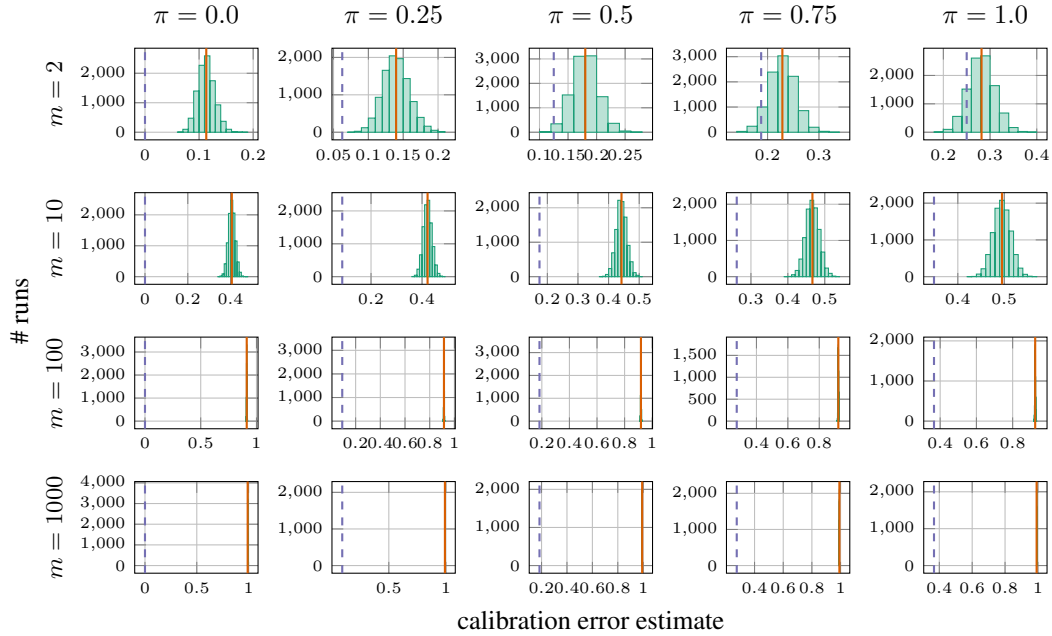


Figure 7: Distribution of  $\widehat{\text{ECE}}$  with data-dependent bins, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

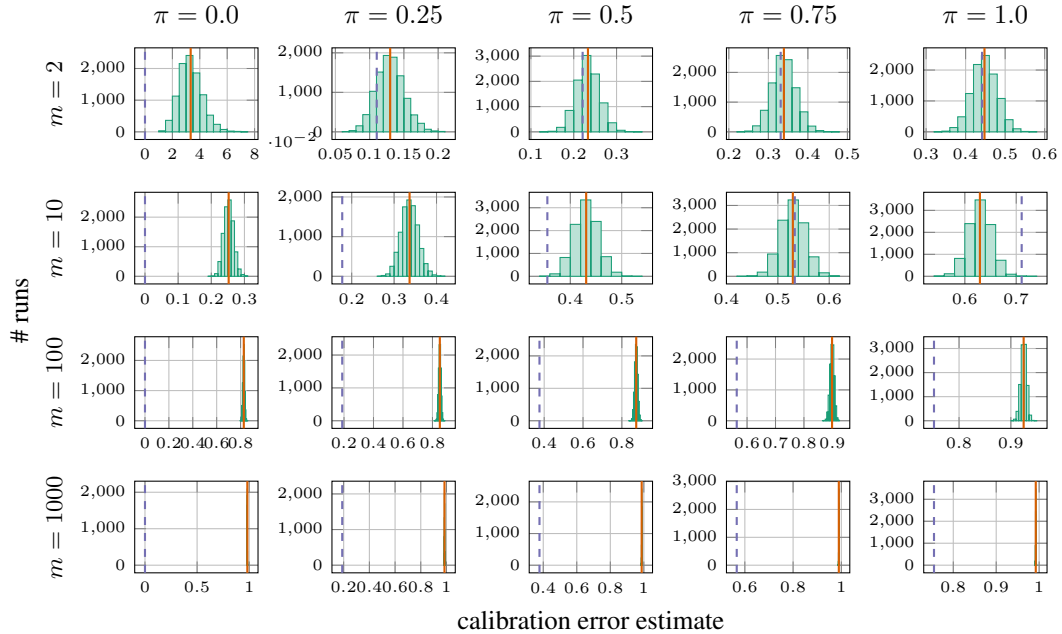


Figure 8: Distribution of  $\widehat{\text{ECE}}$  with data-dependent bins, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .



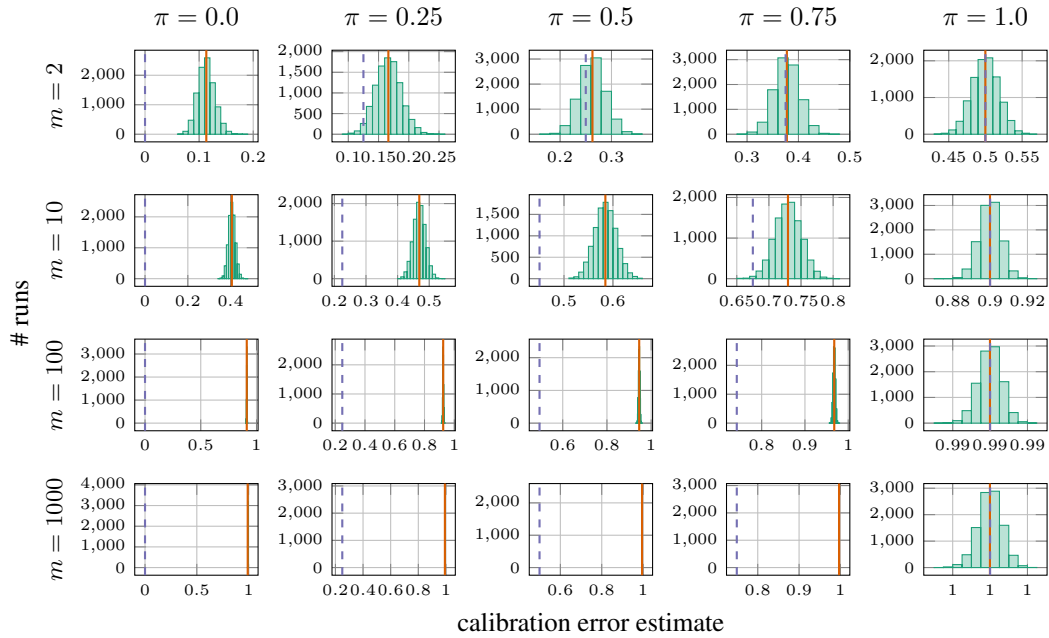


Figure 9: Distribution of  $\widehat{\text{ECE}}$  with data-dependent bins, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

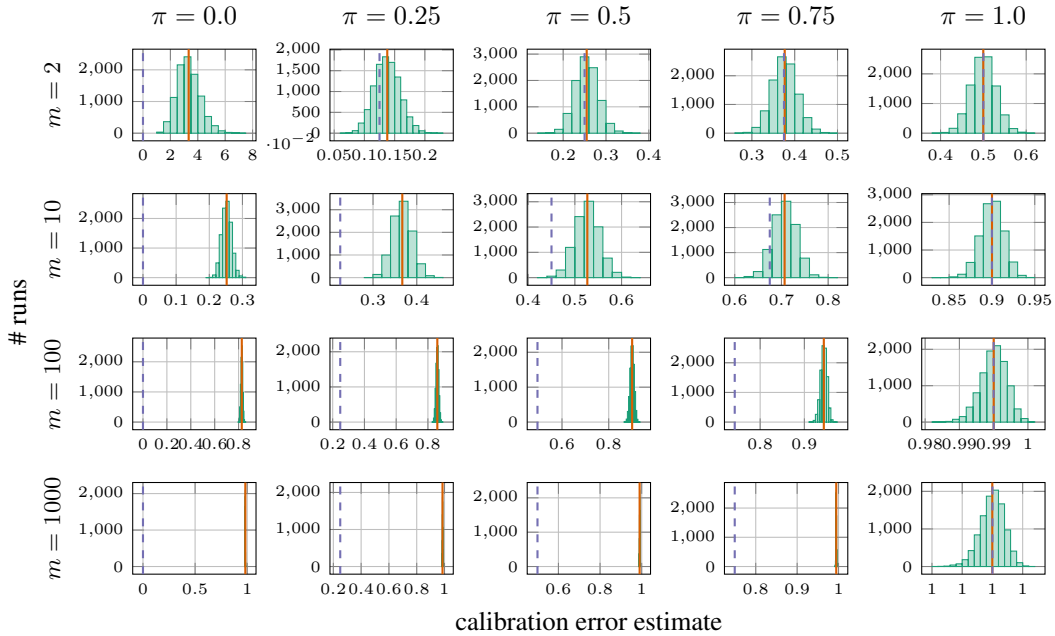


Figure 10: Distribution of  $\widehat{\text{ECE}}$  with data-dependent bins, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

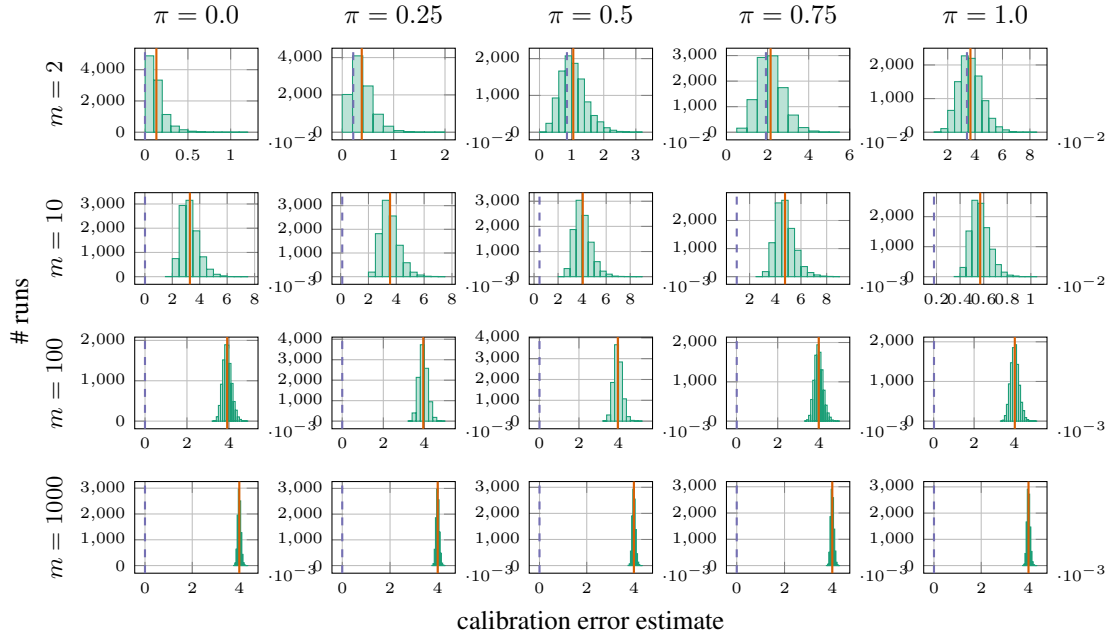


Figure 11: Distribution of  $\widehat{\text{SKCE}}_b$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

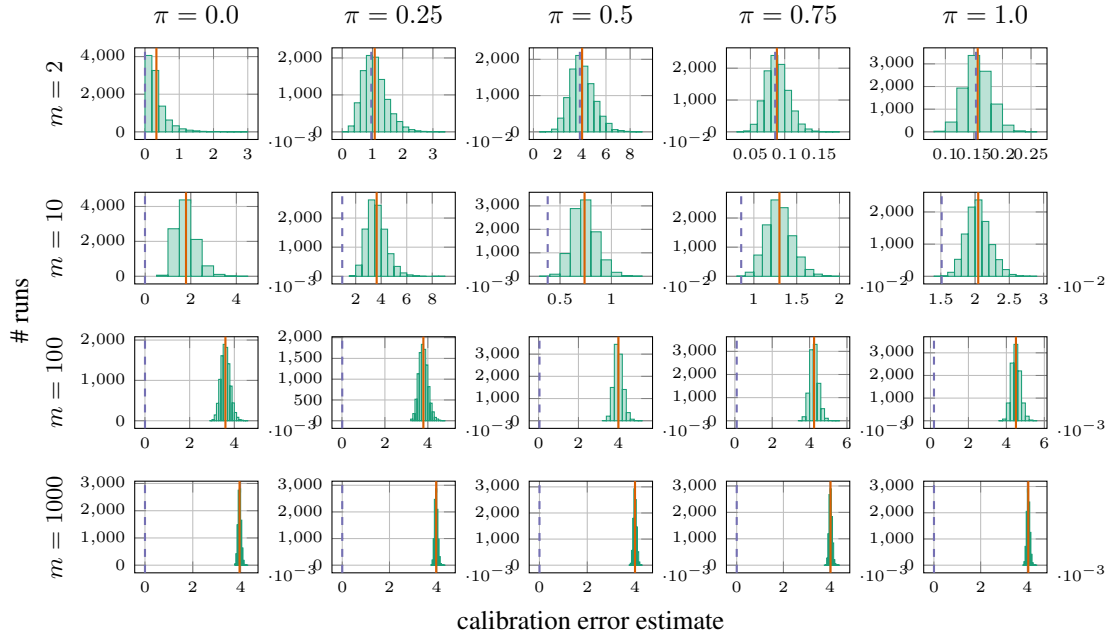


Figure 12: Distribution of  $\widehat{\text{SKCE}}_b$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

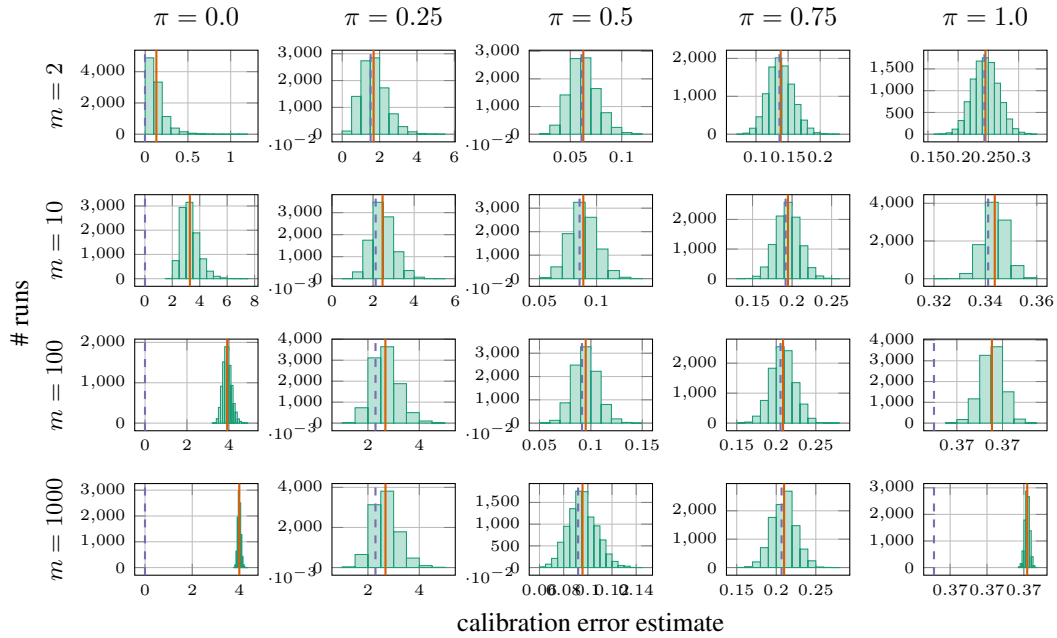


Figure 13: Distribution of  $\widehat{\text{SKCE}}_b$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

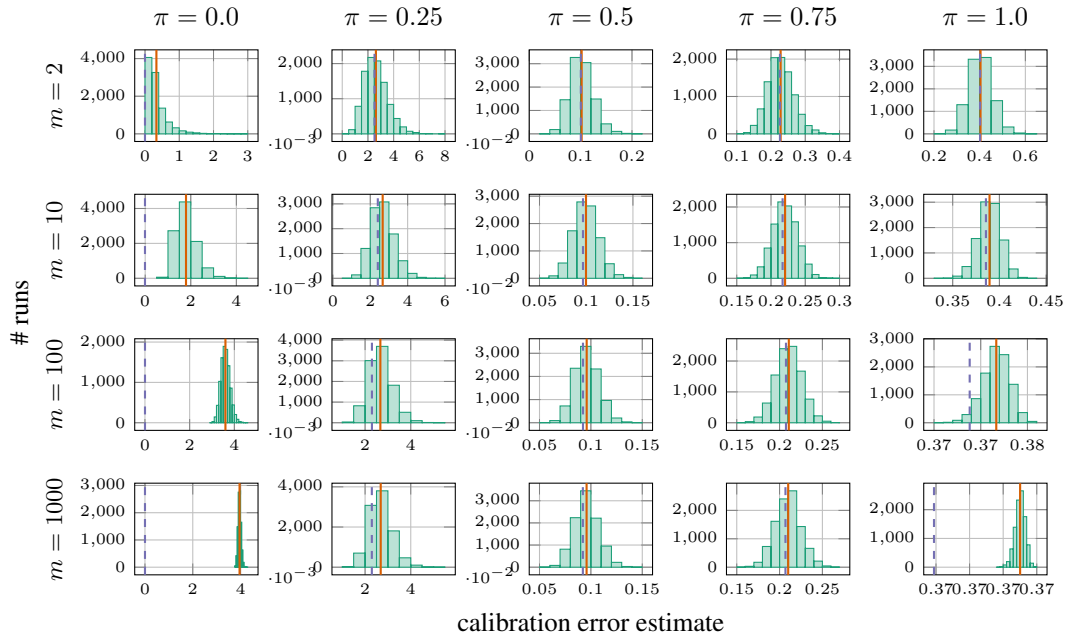


Figure 14: Distribution of  $\widehat{\text{SKCE}}_b$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

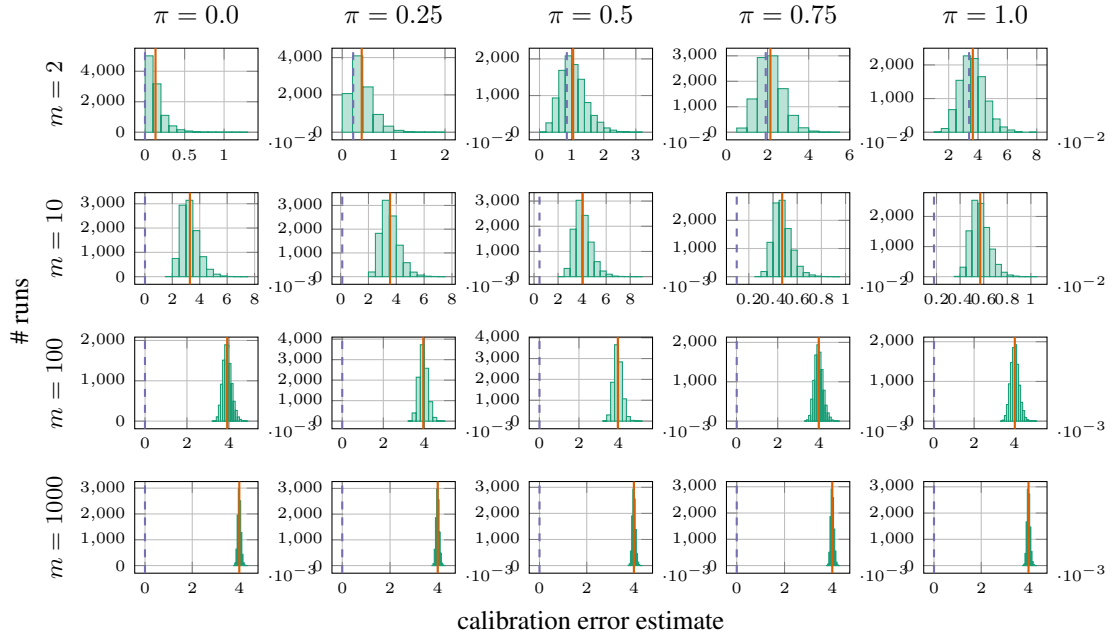


Figure 15: Distribution of  $\widehat{\text{SKCE}}_b$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

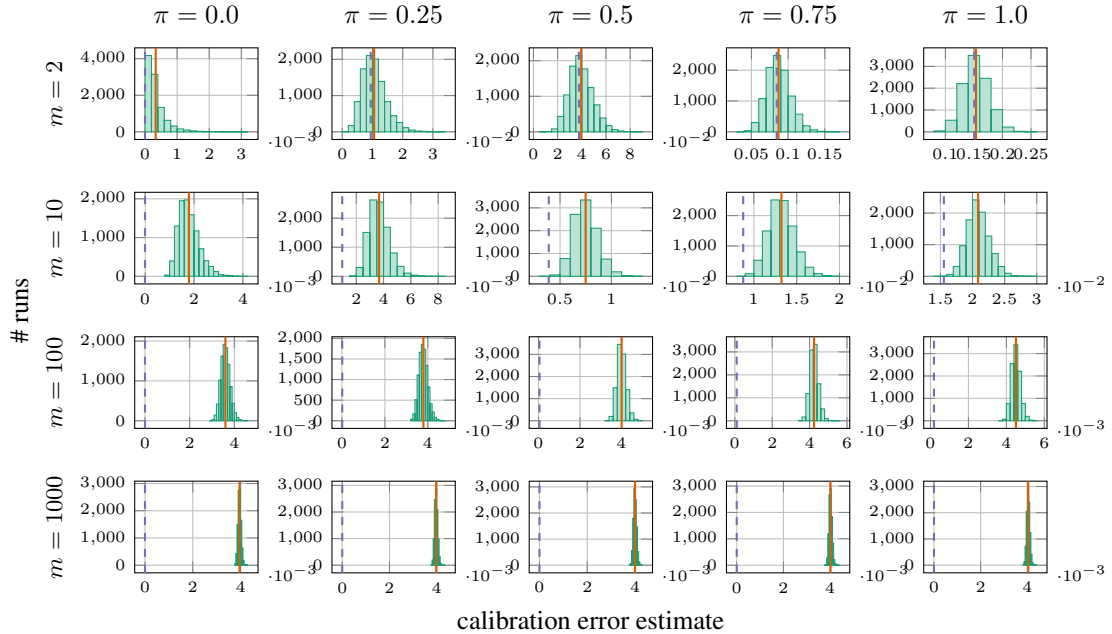


Figure 16: Distribution of  $\widehat{\text{SKCE}}_b$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

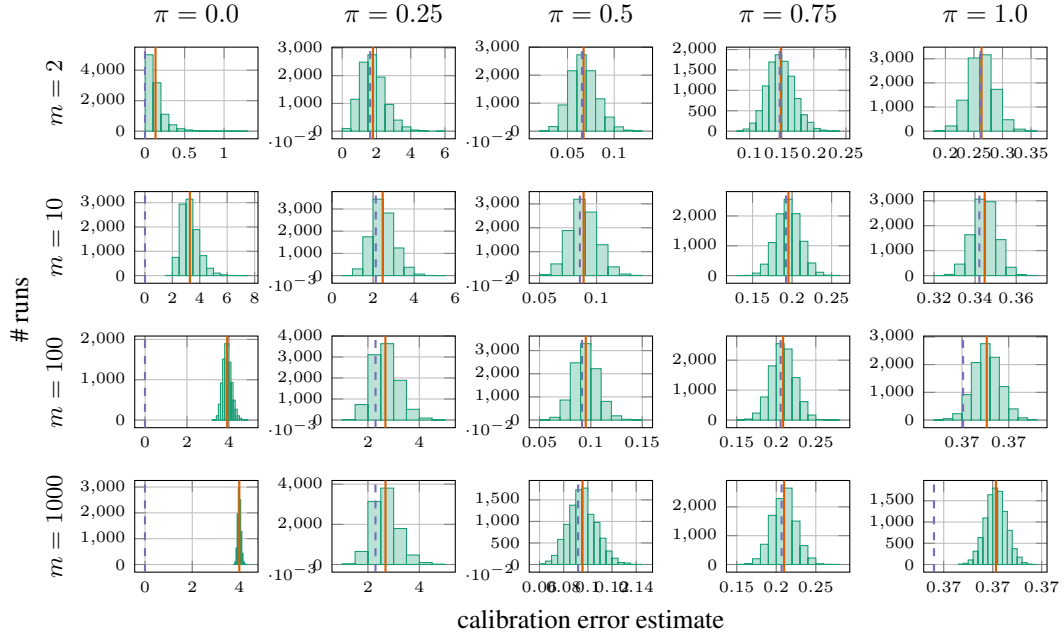


Figure 17: Distribution of  $\widehat{\text{SKCE}}_b$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

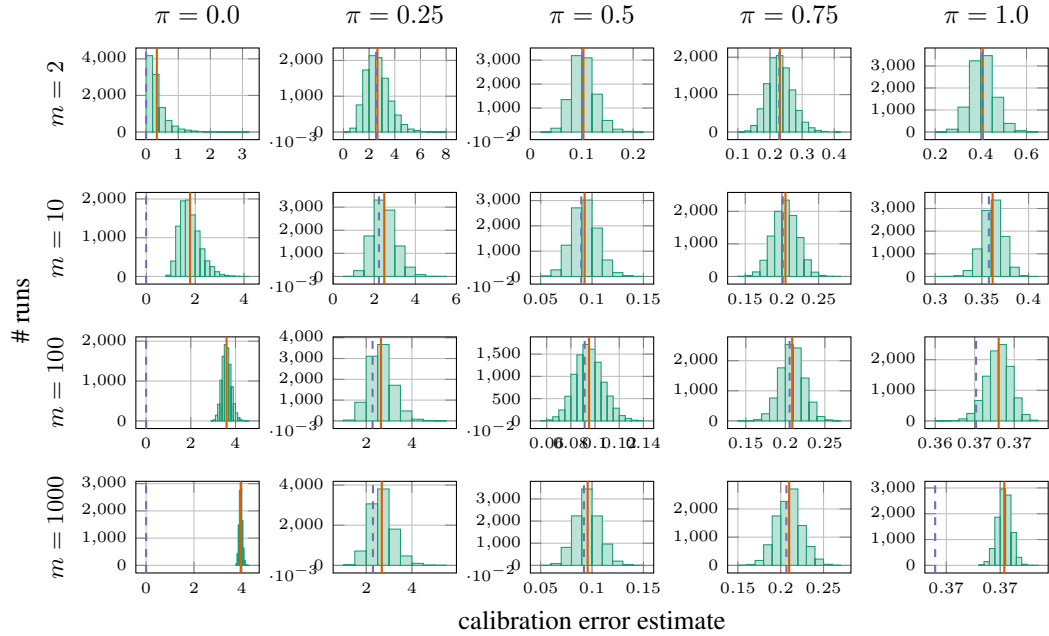


Figure 18: Distribution of  $\widehat{\text{SKCE}}_b$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

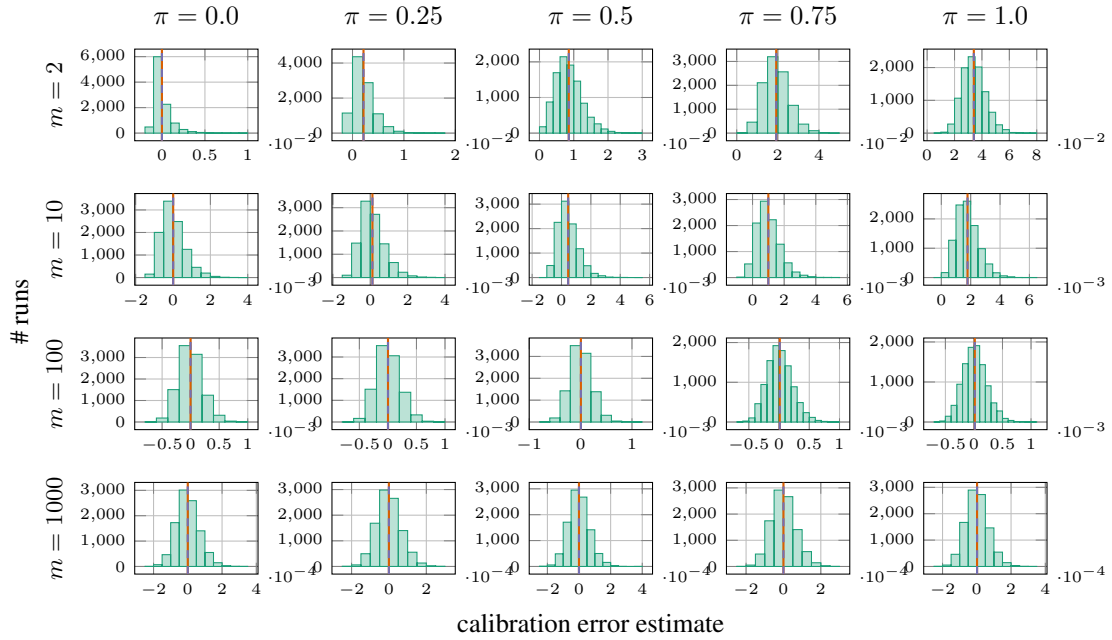


Figure 19: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

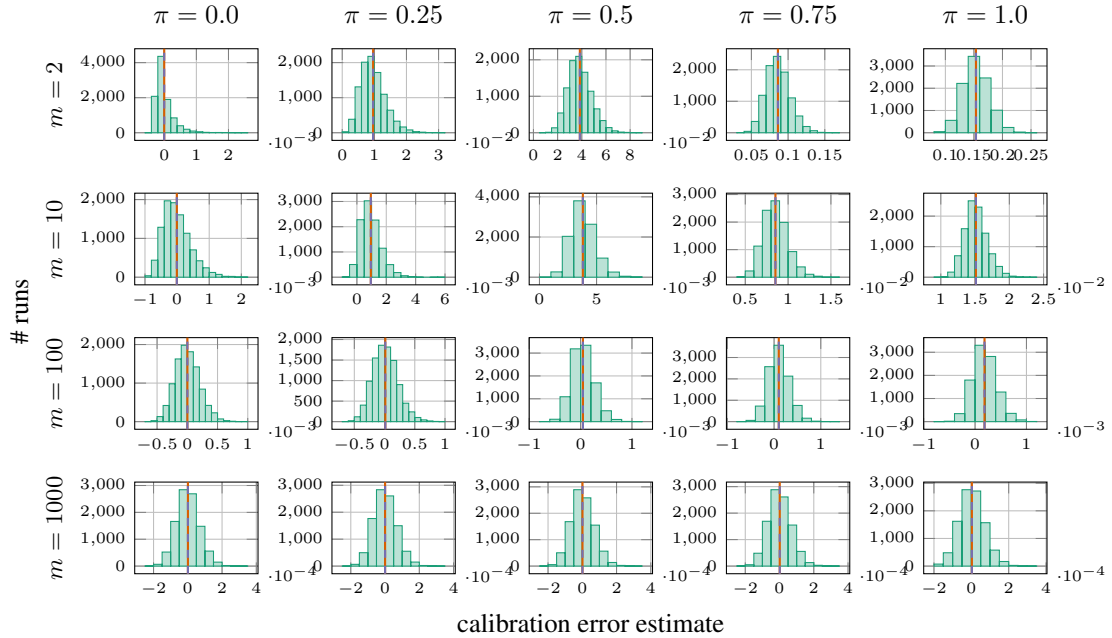


Figure 20: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

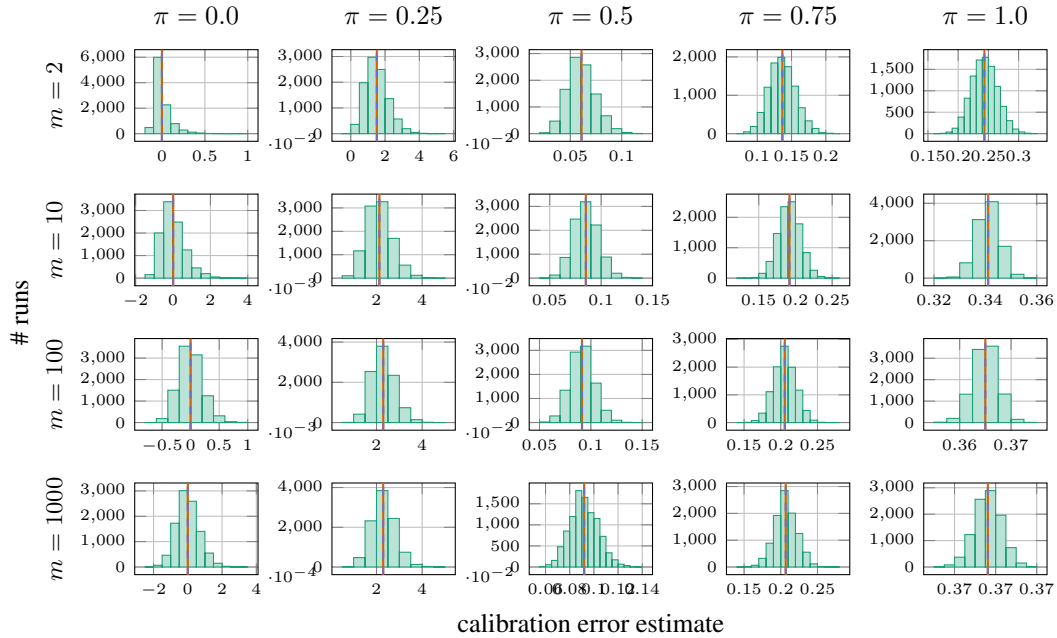


Figure 21: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

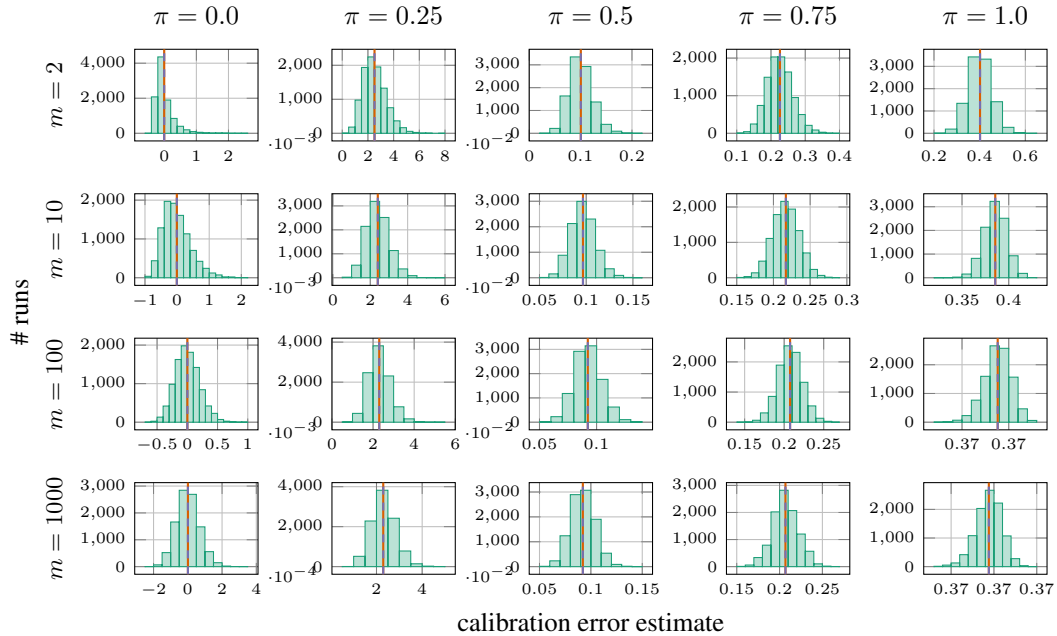


Figure 22: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

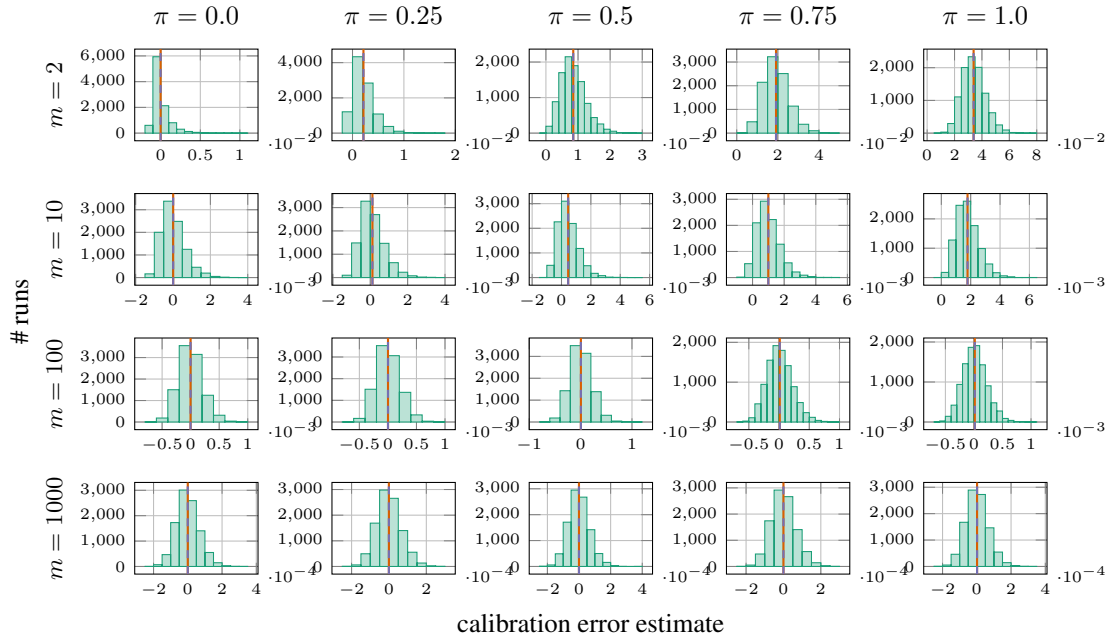


Figure 23: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

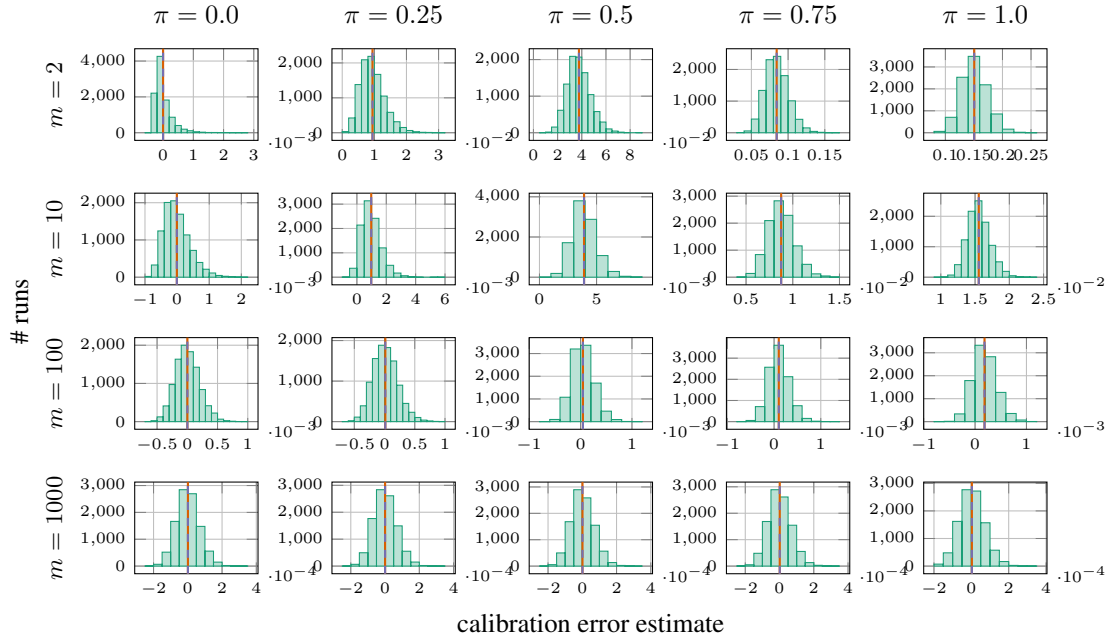


Figure 24: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .



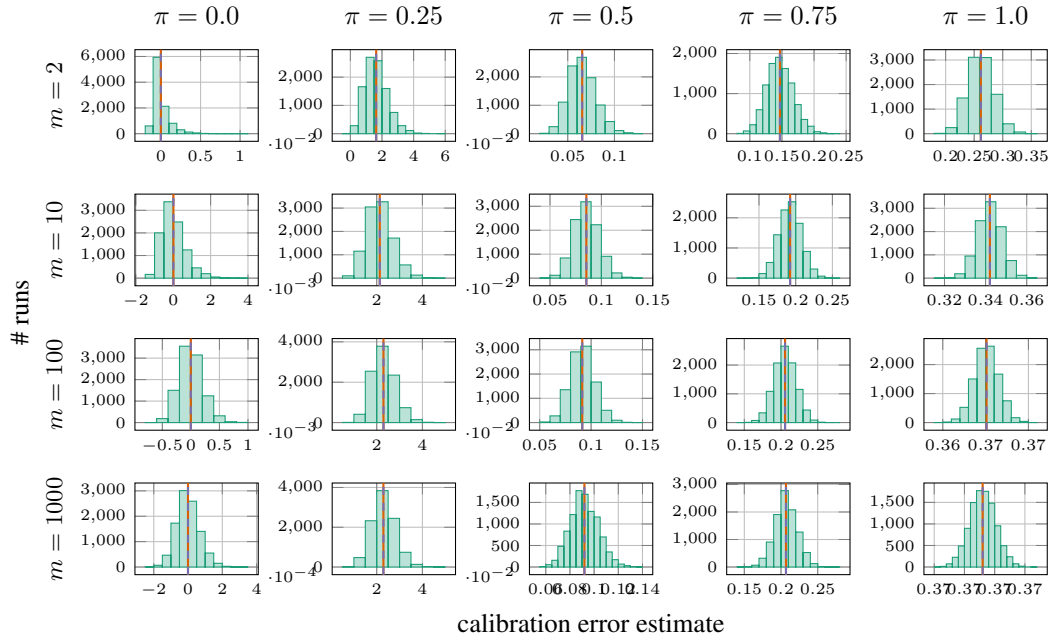


Figure 25: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

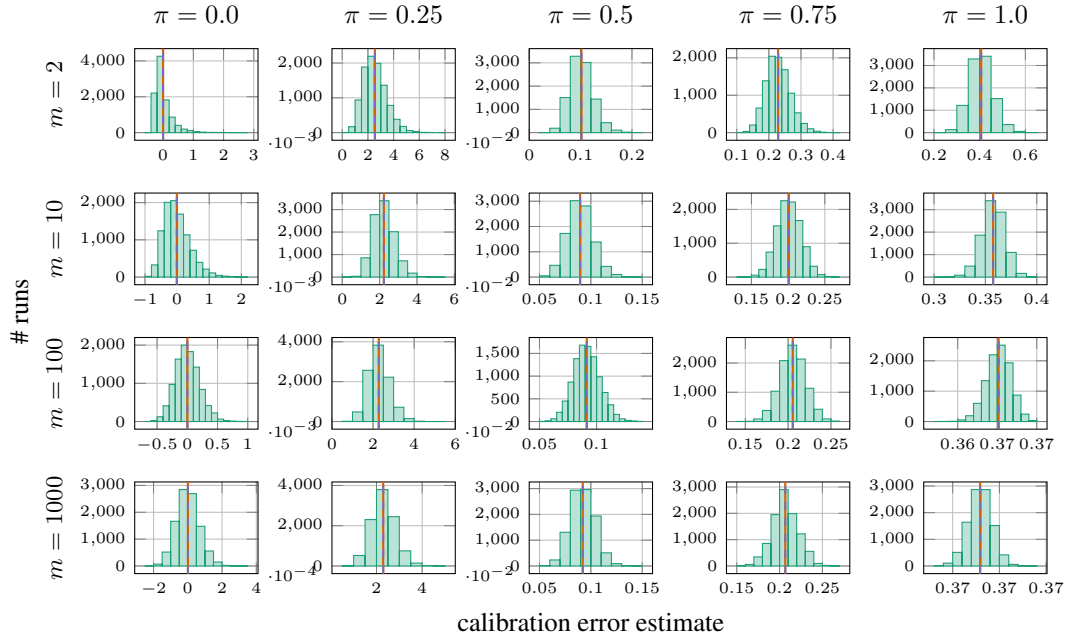


Figure 26: Distribution of  $\widehat{\text{SKCE}}_{\text{uq}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

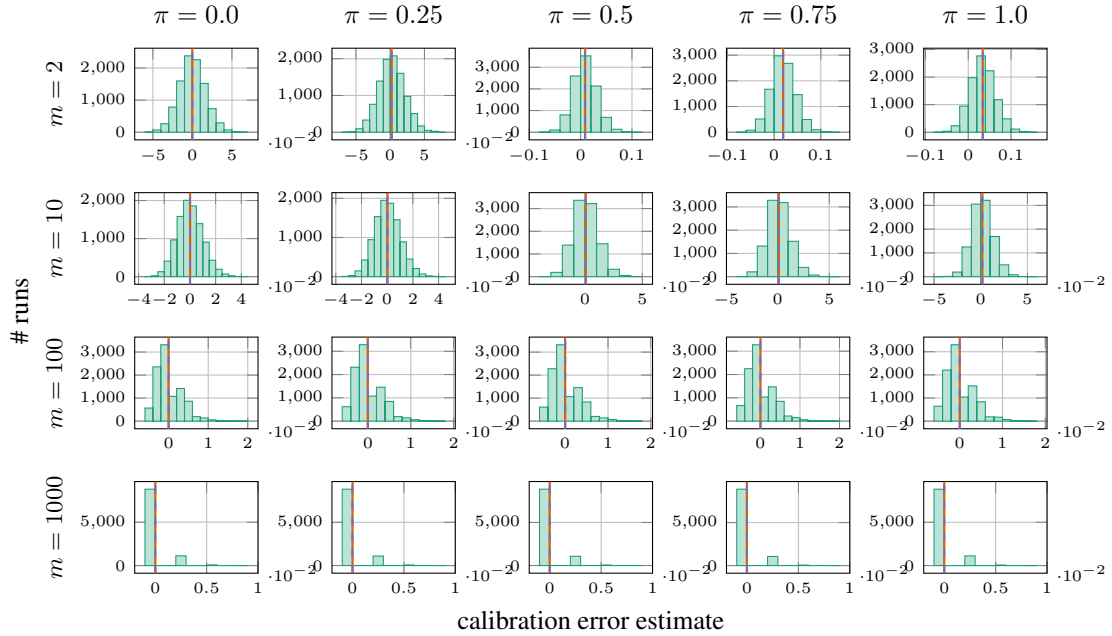


Figure 27: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

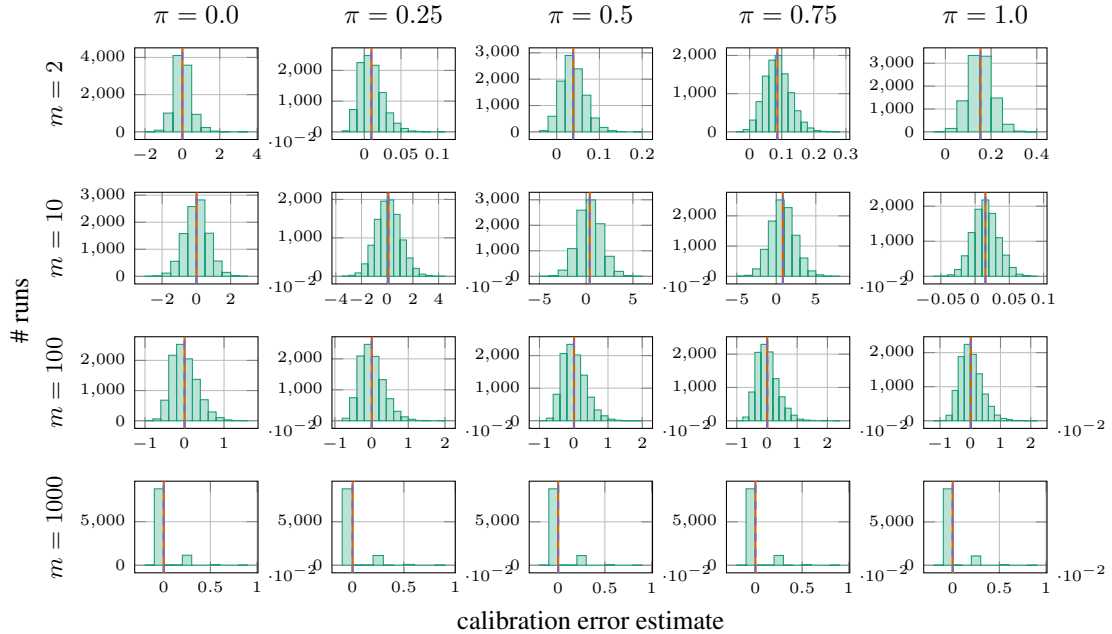


Figure 28: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

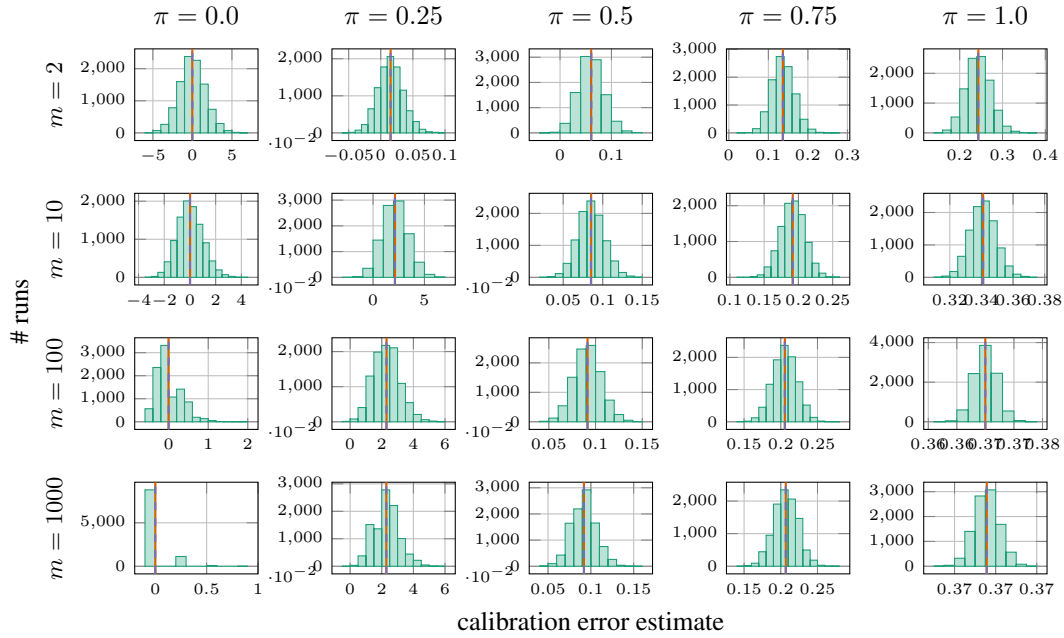


Figure 29: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

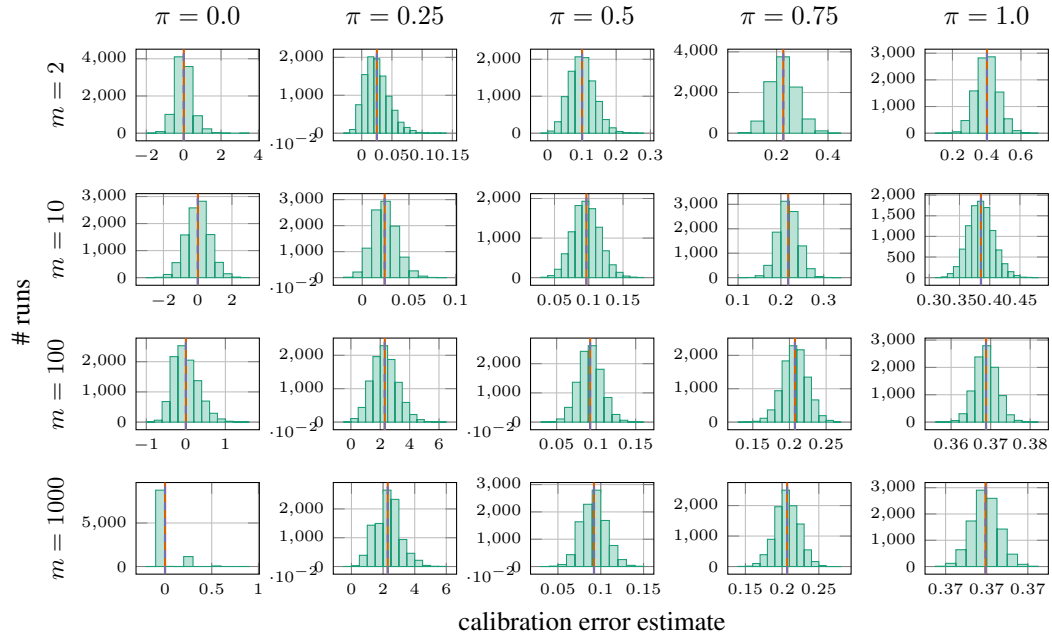


Figure 30: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the median heuristic, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

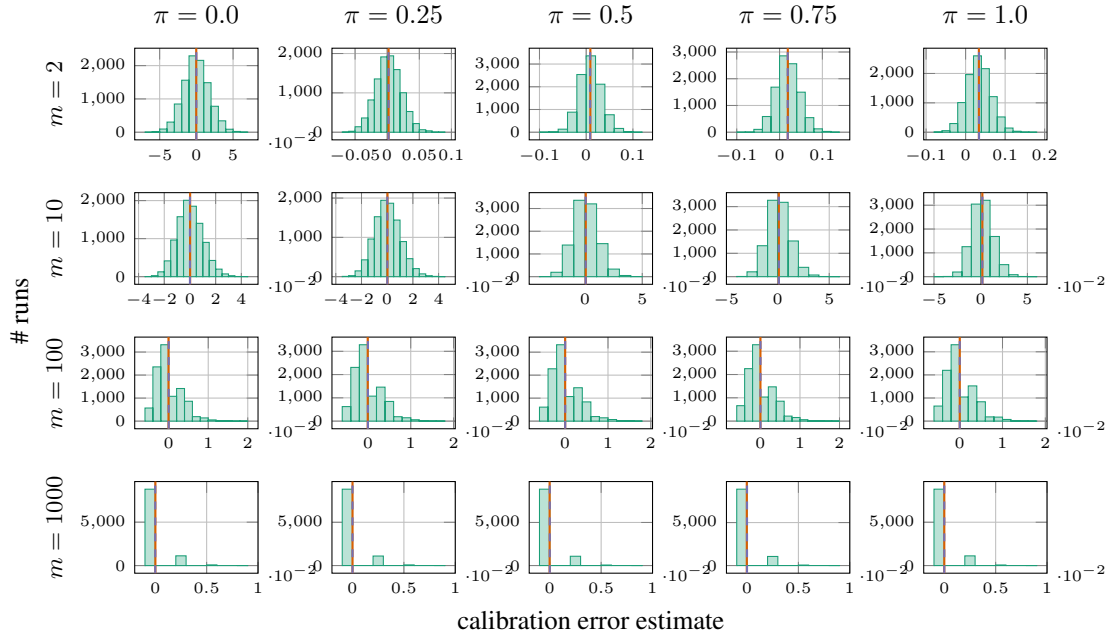


Figure 31: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1/m, \dots, 1/m)$ .

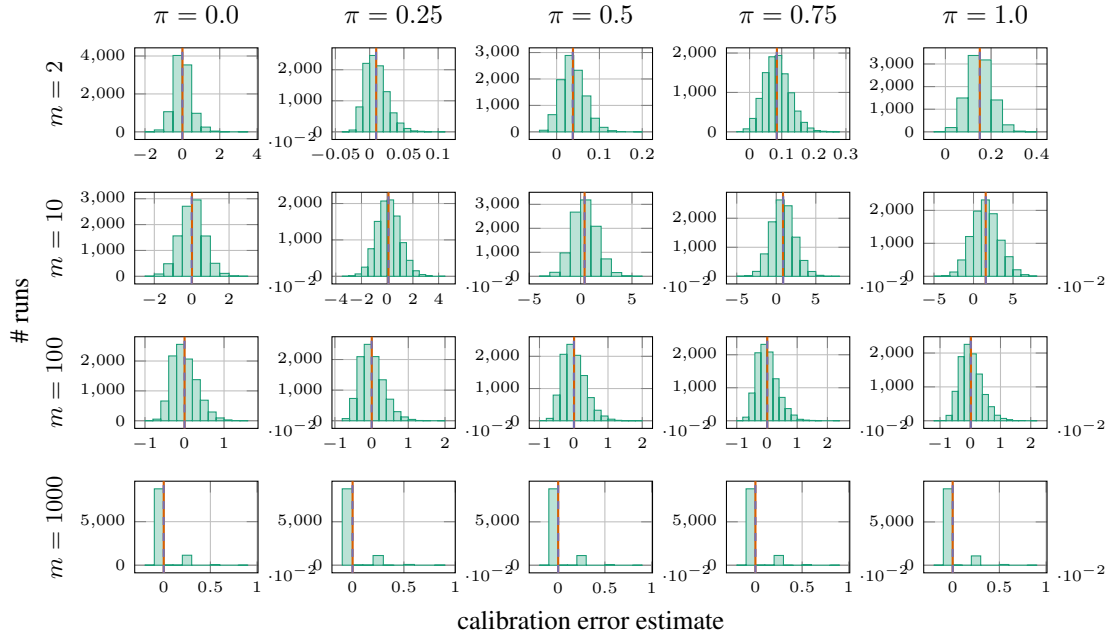


Figure 32: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1/m, \dots, 1/m)$ .

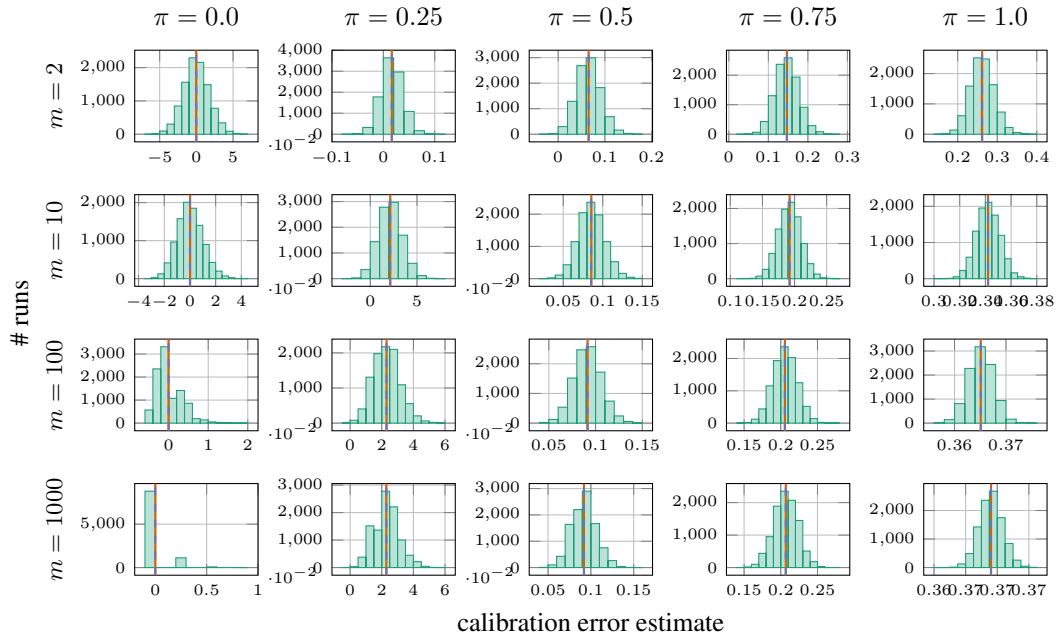


Figure 33: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (1, \dots, 1)$  and  $\beta = (1, 0, \dots, 0)$ .

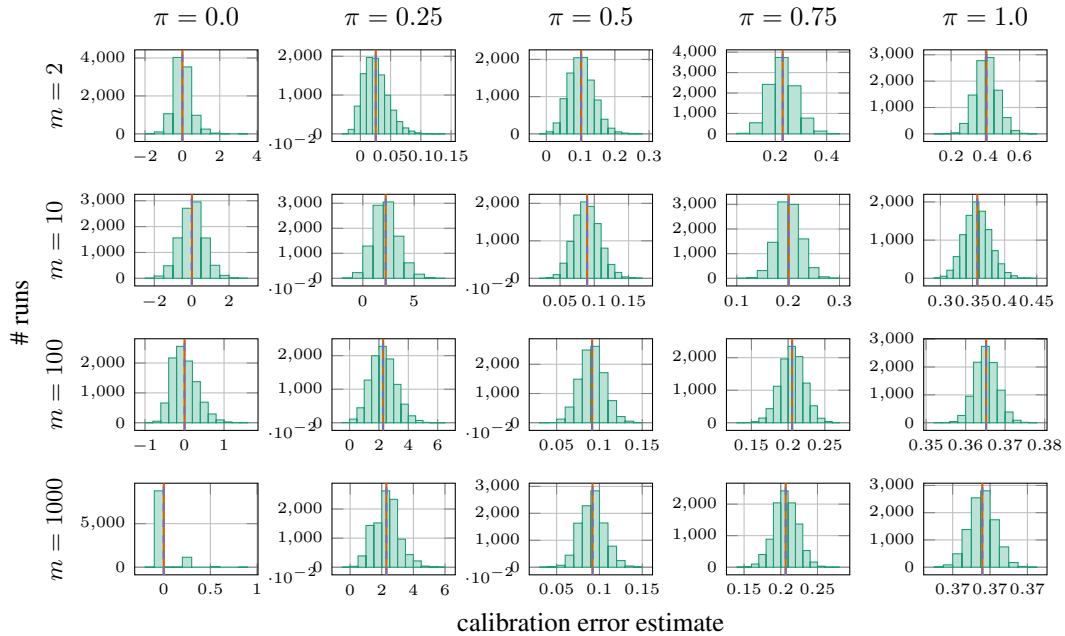


Figure 34: Distribution of  $\widehat{\text{SKCE}}_{\text{ul}}$  with the mean total variation distance, evaluated on  $10^4$  data sets of 250 labeled predictions that are randomly sampled from generative models with  $\alpha = (0.1, \dots, 0.1)$  and  $\beta = (1, 0, \dots, 0)$ .

### J.3 Modern neural networks

In the main experiments of our paper discussed in Appendix J.2 we focus on an experimental confirmation of the derived theoretical properties of the kernel-based estimators and their comparison with the commonly used ECE. In contrast to Guo et al. (2017), neither the study of the calibration of different neural network architectures nor the re-calibration of uncalibrated models are the main goal of our paper. The calibration measures that we consider only depend on the predictions and the true labels, not on how these predictions are computed. We therefore believe that directly specifying the predictions in a “controlled way” results in a cleaner and more informative numerical evaluation.

That being said, we recognize that this approach might result in an unnecessary disconnect between the results of the paper and a practical use case. We therefore conduct additional evaluations with different modern neural networks as well. We consider pretrained ResNet, DenseNet, VGGNet, GoogLeNet, MobileNet, and Inception neural networks Phan (2019) for the classification of the CIFAR-10 image data set Krizhevsky (2009). The CIFAR-10 data set is a labeled data set of  $32 \times 32$  colour images and consists of 50000 training and 10000 test images in 10 classes. The calibration of the neural network models is estimated from their predictions on the CIFAR-10 test data set. We use the same calibration error estimators and p-value approximations as for the generative models above; however, the minimum number of samples per bin in the data-dependent binning scheme of the ECE estimator is increased to 100 to account for the increased number of data samples.




Figure 35: Calibration error estimates of modern neural networks for classification of the CIFAR-10 image data set.

The computed calibration error estimates are shown in Fig. 35. As we argue in our paper, the raw calibration estimates are not interpretable and can be misleading. The results in Fig. 35 endorse this opinion. The estimators rank the models in different order (also the two estimators of the ECE), and it is completely unclear if the observed calibration error estimates (in the order of  $10^{-2}$  and  $10^{-4}$ !) actually indicate that the neural network models are not calibrated.

Hence to obtain an interpretable measure, we consider different bounds and approximations of the p-value for the calibration error estimators, assuming the models are calibrated. More concretely, we estimate the p-value by consistency resampling of the standard ( $C_{\text{uniform}}$ ) and the data-dependent




Figure 36: Bounds and approximations of the p-value of modern neural networks for classification of the CIFAR-10 image data set for different calibration error estimators, assuming the models are calibrated.

( $C_{\text{data-dependent}}$ ) estimator of the ECE, evaluate the distribution-free bounds of the p-value for the estimators  $\widehat{SKCE}_b$  ( $D_b$ ),  $\widehat{SKCE}_{uq}$  ( $D_{uq}$ ), and  $\widehat{SKCE}_{ul}$  ( $D_l$ ) of the SKCE, and approximate the p-value using the asymptotic distribution of the estimators  $\widehat{SKCE}_{uq}$  ( $A_{uq}$ ) and  $\widehat{SKCE}_{ul}$  ( $A_l$ ). The results are shown in Fig. 36.

The approximations obtained by consistency resampling are almost always zero, apart from GoogLeNet for the ECE estimator with partitions of uniform size. However, since our controlled experiments with the generative models showed that consistency resampling might underestimate the p-value of calibrated models on average, these approximations could be misleading. On the contrary, the bounds and approximations of the p-value for the estimators of the SKCE are theoretically well-founded. In our experiments with the generative models, the asymptotic distribution of the estimator  $\widehat{SKCE}_{uq}$  seemed to allow to approximate the p-value quite accurately on average and yielded very powerful tests. For all studied neural network models these p-value approximations are zero, and hence for all models we would always reject the null hypothesis of calibration. The p-value approximations based on the asymptotic distribution of the estimator  $\widehat{SKCE}_{ul}$  vary between around 0.18 for the ResNet18 and 0.91 for the GoogLeNet model. The higher p-value approximations correspond to the increased empirical test errors with the uncalibrated generative models compared to the tests based on the asymptotic distribution of the estimator  $\widehat{SKCE}_{uq}$ . Most distribution-free bounds of the p-value are between 0.99 and 1, indicating again that these bounds are quite loose.

All in all, the evaluations of the modern neural networks seem to match the theoretical expectations and are consistent with the results we obtained in the experiments with the generative models. Moreover, the p-value approximations of zero are consistent with Guo et al. (2017)’s finding that modern neural networks are often not calibrated.

#### J.4 Computational time

The computational time, although dependent on our Julia implementation and the hardware used, might provide some insights to the interested reader in addition to the algorithmic complexity.

However, in our opinion, a fair comparison of the proposed calibration error estimators should take into account the error of the calibration error estimation, similar to work precision diagrams for numerical differential equation solvers.

A simple comparison of the computational time for the calibration error estimators used in the experiments with the generative models in Appendix J.2 on our computer (3.6 GHz) shows the expected scaling of the computational time with increasing number of samples. As Fig. 37 shows, even for 1000 samples and 1000 classes the estimators  $\widehat{\text{SKCE}}_b$  and  $\widehat{\text{SKCE}}_{uq}$  with the median heuristic can be evaluated in around 0.1 seconds. Moreover, the simple benchmark indicates that the evaluation of the estimator of the ECE with data-dependent bins is much slower than of the one with partitions of uniform size.

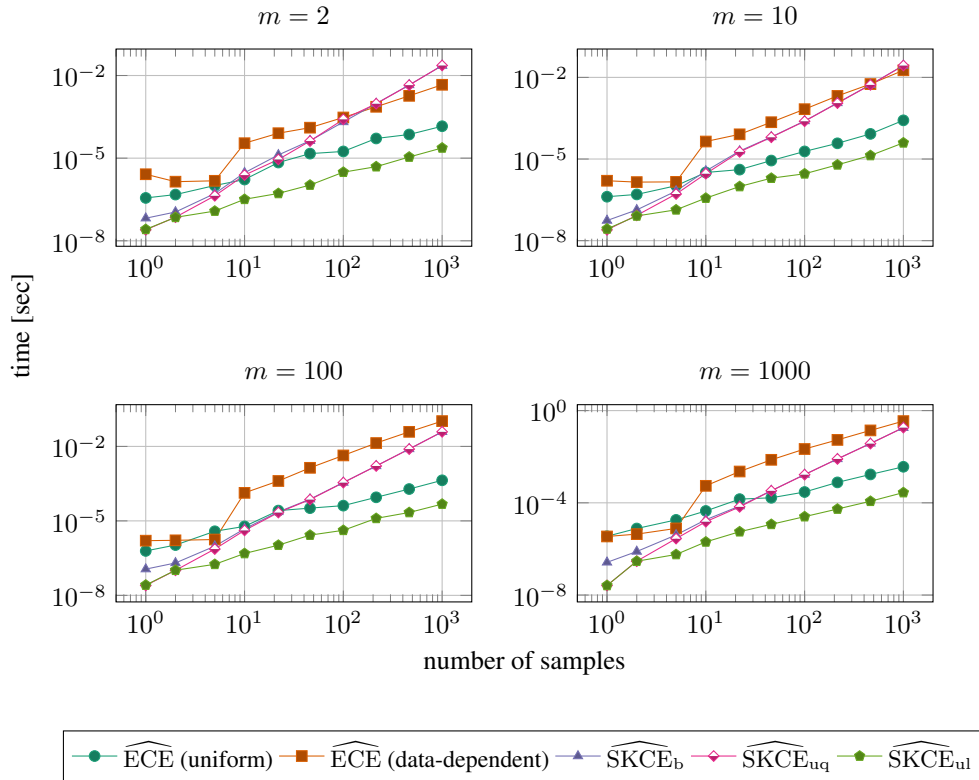


Figure 37: Computational time for the evaluation of calibration error estimators on data sets with different number of classes versus number of data samples.