

# 000 BEYOND OVERCONFIDENCE: RETHINKING CALIBRATION 001 IN LARGE-SCALE VISION MODELS

002  
003  
004  
005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT

009  
010  
011 Reliable uncertainty calibration is crucial for the safe deployment of deep neu-  
012 ral networks in high-stakes settings. While these networks are known to exhibit  
013 systematic overconfidence, especially under distribution shifts, the calibration of  
014 large-scale vision models, such as ConvNeXt, EVA, and BEiT, has remained un-  
015 derexplored. We comprehensively examine their calibration behavior, uncovering  
016 evidence that challenges well-established assumptions. We find that these models  
017 are underconfident on in-distribution data, which results in increased calibration  
018 error, yet exhibit improved calibration under distribution shifts. This phenomenon  
019 is primarily driven by modern training techniques, including massive pretrain-  
020 ing and sophisticated regularization and augmentation methods, rather than ar-  
021 chitectural innovations alone. We also demonstrate that these large-scale models  
022 are highly responsive to post-hoc calibration techniques in the in-distribution set-  
023 ting, enabling practitioners to mitigate underconfidence bias effectively. However,  
024 these methods become progressively less reliable under severe distribution shifts  
025 and can occasionally produce counterproductive effects. Our findings highlight  
026 the complex, non-monotonic effects of architectural and training innovations on  
027 calibration, challenging established narratives of continuous improvement.

## 028 1 INTRODUCTION

029  
030 Deep neural networks deployed in high-stakes applications require not only high predictive accu-  
031 racy but also reliable uncertainty estimates. In safety-critical domains, such as medical diagnosis,  
032 autonomous driving, and financial decision-making, the consequences of incorrect predictions ac-  
033 companied by misleadingly high confidence scores can be severe. Model calibration – aligning  
034 predicted confidence with empirical accuracy – provides a formal framework for assessing the reli-  
035 ability of these uncertainty estimates (Guo et al., 2017). In a well-calibrated model, predictions made  
036 with 80% confidence should be correct approximately 80% of the time.

037  
038 A fundamental challenge in model calibration is that deep neural networks are typically poorly  
039 calibrated. In particular, they tend to exhibit systematic overconfidence, assigning probabilities to  
040 predictions that exceed their actual accuracy (see, e.g., Guo et al. (2017); Hendrycks et al. (2021);  
041 Lakshminarayanan et al. (2017); Rahaman & Thiery (2021); Cheng & Vasconcelos (2024); Wang  
042 et al. (2021)). This calibration error becomes even more pronounced when models encounter dis-  
043 tribution shifts (see, e.g., (Ovadia et al., 2019; Hendrycks & Dietterich, 2019)). To address the  
044 challenge of miscalibration, post-hoc calibration methods (see, e.g., Guo et al. (2017); Zhang et al.  
045 (2020b); Gupta et al. (2021); Tomani et al. (2022)) are promising since they can be applied directly  
046 to trained models and can therefore be used as a lightweight post-processing step to recalibrate  
047 the model’s outputs. Additionally, Minderer et al. (2021) demonstrated that architectural innova-  
048 tions available at the time (such as Vision Transformers) have inherently well-calibrated outputs  
049 and improved robustness to distribution shifts, suggesting that miscalibration is more pronounced in  
traditional models than in then-current state-of-the-art models.

050  
051 Most recent advances in deep learning have catalyzed the emergence of models featuring large-  
052 scale training regimes, characterized by massive-scale pre-training using novel training techniques  
053 (e.g., masked image modeling (He et al., 2022; Bao et al., 2021)) and sophisticated regularization  
and augmentation techniques (e.g., CutMix (Yun et al., 2019), MixUp (Zhang et al., 2017), label  
smoothing (Zhang et al., 2020a; Lukasik et al., 2020), and RandAugment (Cubuk et al., 2020)).

054 While models trained with these approaches – such as ConvNeXt, EVA, and BEiT – achieve state-  
 055 of-the-art accuracy and have led to a broad adoption by practitioners, the implications for model  
 056 calibration properties remain insufficiently explored. Specifically, it is unclear whether exposure  
 057 to diverse, web-scale training data improves calibration by providing broader coverage of the input  
 058 distribution, or if it introduces new calibration issues due to inherent dataset biases and complex  
 059 regularization schemes. Furthermore, despite significant advances in post-hoc calibration techniques  
 060 for traditional neural architectures, the efficacy of these techniques when applied to these large-scale  
 061 models has not been adequately investigated.

062 In this paper, we systematically benchmark the quality of predictive uncertainty of large-scale vision  
 063 models and make the following key contributions:

- 065 Through a systematic benchmark, we demonstrate that large-scale models (ConvNeXt,  
 066 EVA, and BEiT) exhibit significant in-distribution calibration errors, characterized by sys-  
 067 tematic **underconfidence** in predictive probabilities. This finding contrasts with the well-  
 068 documented overconfidence bias observed in traditional deep neural networks.
- 069 Our analysis reveals that this systematic underconfidence in large-scale models results from  
 070 the combination of pretraining on extensive datasets and advanced regularization strategies,  
 071 rather than architectural design choices.
- 072 We further show that large-scale models maintain calibration quality under both synthetic  
 073 and real-world distribution shifts. This finding contrasts with traditional neural archi-  
 074 tectures, which exhibit a monotonic increase in calibration error as the magnitude of the dis-  
 075 tribution shift increases.
- 076 We demonstrate that post-hoc calibration methods can significantly improve the calibration  
 077 of large-scale models for in-distribution predictions. However, their benefits diminish under  
 078 distribution shift.

## 080 2 RELATED WORK

082 **Empirical Studies of Model Calibration** Over the past decade, research into neural network cal-  
 083ibration has established a strong empirical foundation. The seminal work by Guo et al. (2017) first  
 084 documented that neural architectures used at the time, such as ResNets and DenseNets, typically  
 085 produce overconfident predictions. Several subsequent studies have corroborated this finding (Thu-  
 086 lasidasan et al., 2019; Hendrycks et al., 2021; Lakshminarayanan et al., 2017; Rahaman & Thiery,  
 087 2021).

088 Distribution shift conditions exacerbate these calibration issues: Ovadia et al. (2019) demonstrated  
 089 through a comprehensive evaluation that “along with accuracy, the quality of uncertainty consis-  
 090 tently degrades with increasing dataset shift.” Hendrycks & Dietterich (2019) further validated  
 091 this phenomenon, whose ImageNet-C benchmark revealed a direct correlation between corruption  
 092 severity and increasing calibration error. Similarly, Recht et al. (2019) demonstrated that temporal  
 093 distribution drift in ImageNet-V2 negatively impacts both predictive performance and calibration  
 094 metrics.

095 Recent architectural advancements have challenged these established patterns. Minderer et al.  
 096 (2021) documented improved calibration in Vision Transformers and MLP-Mixers compared to  
 097 previous generations of models. They noted that these models were “well calibrated compared  
 098 to past models and their performance is more robust to distribution shift.” They also emphasized  
 099 the importance of model architecture in determining calibration quality, suggesting that the most  
 100 recent architectural innovations may improve calibration quality inherently. Tao et al. (2024) further  
 101 substantiated this architectural dependency. Their large-scale calibration benchmark, which used  
 102 NAS-searched architectures, demonstrated a strong correlation between the design choices of neu-  
 103 ral networks and their calibration properties. However, their investigation was limited to models  
 104 with conventional training regimes and did not examine large-scale models pre-trained on massive  
 105 datasets.

106 Recent work has emphasized decomposing predictive uncertainty to understand miscalibration  
 107 sources beyond aggregate metrics. Perez-Lebel et al. Perez-Lebel et al. (2023) show that even  
 perfectly calibrated classifiers can exhibit grouping loss—samples with identical confidence but

108 different true probabilities. This builds on proper scoring rule theory Murphy (1973a); Bröcker  
 109 (2009); Kull & Flach (2015), which decomposes predictive errors into calibration, refinement, and  
 110 irreducible components.  
 111

112 **Post-hoc Calibration Techniques** Post-hoc calibration methods represent a computationally ef-  
 113 ficient framework for enhancing the reliability of neural network confidence estimates without re-  
 114 quiring architectural modifications or extensive retraining procedures. These approaches operate  
 115 by learning mapping functions that transform a model’s raw outputs into recalibrated probability  
 116 distributions, thereby optimizing the correspondence between predictive confidence and empirical  
 117 accuracy. The re-calibration process typically leverages a held-out validation set to estimate the  
 118 parameters of these transformations while maintaining the model’s discriminative capabilities.  
 119

120 The literature has explored various approaches to post-hoc calibration, each of which is character-  
 121 ized by a distinct set of trade-offs between functional expressivity, parameter efficiency, and gener-  
 122 alization properties Guo et al. (2017); Zhang et al. (2020b); Gupta et al. (2021); Kull et al. (2019).  
 123 While post-hoc calibration methods achieve strong performance on in-distribution data, their reli-  
 124 ability degrades substantially under distribution shift. Recent advances address this limitation through  
 125 density-aware approaches that ensure calibration not only globally but also within local regions of  
 126 the input space Xiong et al. (2023); Tomani et al. (2023).  
 127

### 3 PROBLEM DEFINITION AND NOTATION

128 In this paper, we systematically benchmark the quality of predictive uncertainty of multi-class neural  
 129 network models across different architecture and training paradigms. A neural network parameter-  
 130 izes a prediction function  $f$  that maps input  $\mathbf{x} \in \mathbb{R}^D$  to a probability vector  $\mathbf{p} \in [0, 1]^C$  over  $C$   
 131 classes. These predictions reside in the  $(C - 1)$ -dimensional probability simplex:  $\Delta = \{\mathbf{p} \in [0, 1]^C \mid \sum_{c=1}^C p_c = 1\}$ , where  $p_c$  denotes the  $c$ -th component of the probability vector  $\mathbf{p}$ .  
 132

133 A model  $f$  is perfectly calibrated (Bröcker, 2009) if and only if:

$$\forall \mathbf{p} \in \Delta : \mathbb{P}(y = c \mid f(\mathbf{x}) = \mathbf{p}) = p_c. \quad (1)$$

134 Throughout this paper, we focus on the weaker notion of *top-label calibration* (Guo et al., 2017),  
 135 requiring that predictions made with maximum confidence  $p^* = \max f(\mathbf{x})$  are correct with proba-  
 136 bility  $p^*$ :  
 137

$$\forall p^* \in [0, 1] : \mathbb{P}(y \in \arg \max f(\mathbf{x}) \mid \max f(\mathbf{x}) = p^*) = p^*. \quad (2)$$

138 To quantify top-label calibration error, we compute the Expected Calibration Error (ECE), which  
 139 measures the expected discrepancy between the two sides of Eq. 2 and is defined as  
 140

$$\mathbb{E}[|p^* - \mathbb{P}(y \in \arg \max f(\mathbf{x}) \mid \max f(\mathbf{x}) = p^*)|]. \quad (3)$$

141 Due to the continuous-valued probability space, direct estimation of Eq. 3 is intractable. Therefore,  
 142 a binning approach is typically employed by partitioning the prediction space into  $m$  equally spaced  
 143 bins  $B_1, \dots, B_m$ . Given  $n$  i.i.d. samples  $(\mathbf{x}_i, y_i)_{i=1}^n$  drawn from the joint distribution  $\mathbb{P}(\mathbf{x}, y)$ , we  
 144 assign each  $i \in \{1, \dots, n\}$  to a bin  $B_j$  based on  $\max f(\mathbf{x}_i)$ .  
 145

146 Then, we compute for each bin  $B_j$  the mean top-level confidence  $\text{conf}(B_j) = \frac{1}{|B_j|} \sum_{i \in B_j} \max f(\mathbf{x}_i)$  and the mean accuracy  $\text{acc}(B_j) = \frac{1}{|B_j|} \sum_{i \in B_j} \mathbf{1}(\arg \max f(\mathbf{x}_i) = y_i)$  and  
 147 finally compute the Expected Calibration Error according to  
 148

$$\text{ECE} = \sum_{j=1}^m \frac{|B_j|}{n} |\text{acc}(B_j) - \text{conf}(B_j)|. \quad (4)$$

149 In addition to ECE, we quantify Brier score and the negative log likelihood as proper scoring rules,  
 150 capturing both model calibration and model sharpness Murphy (1973b); Popordanoska et al. (2024).  
 151 Formal definitions of these metrics are provided in the supplementary material.  
 152

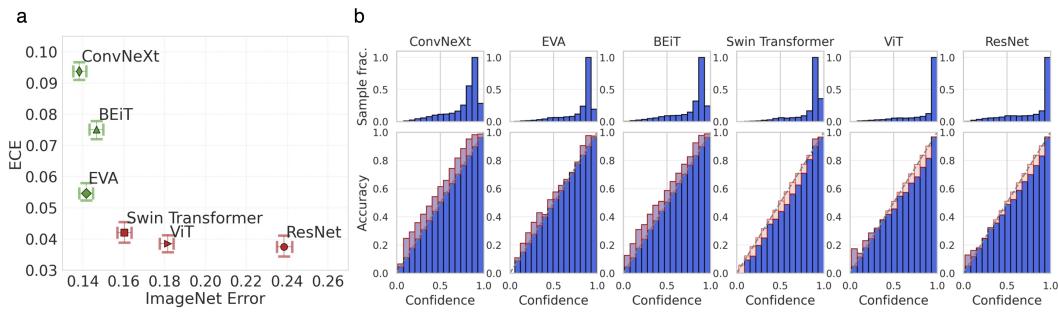


Figure 1: (a) Inverse relationship between ImageNet classification error and Expected Calibration Error (ECE). Green markers represent large-scale models (from 2022), while red markers represent traditional models (until 2021). Error bars show 95% bootstrap confidence intervals ( $n=100$ ). Despite their superior classification performance, large-scale models consistently exhibit significantly higher calibration errors. (b) Reliability diagrams showing the systematic underestimation of predictive confidence in large-scale models (ConvNeXt, EVA, and BEiT), contrasting with the overconfidence observed in traditional models (Swin Transformer, ViT, and ResNet).

## 4 EMPIRICAL EVALUATION

### 4.1 EXPERIMENTAL SETUP

**Models Under Evaluation.** Throughout the paper we consistently evaluate six neural networks. We distinguish between traditional training paradigms (ResNet-50, ViT-B/16, Swin-S3-B) and contemporary large-scale training regimes characterized by massive pretraining combined with sophisticated regularization and augmentation techniques (BEiT-B/16, EVA-S/14, ConvNeXt-B). Note that this distinction is based on training methodology rather than model architecture or pretraining dataset scale alone. To disentangle the effects of architecture from training methodology, we additionally evaluate ViT and ResNet variants trained with large-scale techniques while maintaining traditional architectures. Detailed model specifications are provided in Appendix B.

**Datasets.** We evaluate accuracy and calibration error on the **ImageNet-1k** dataset Deng et al. (2009) and the following distributed-shifted benchmarks:

1. **ImageNet-C** (Hendrycks & Dietterich (2019)), which augments the standard ImageNet-1k dataset by introducing 19 distinct types of synthetic corruptions, each applied at 5 severity levels.
2. **ImageNet-V2** (Recht et al. (2019)), comprising 10,000 temporally shifted real-word samples collected using the original ImageNet-1k sampling protocol.
3. **ImageNet-A** (Hendrycks et al. (2021)), containing 7,500 natural adversarial examples specifically selected for their ability to induce misclassification in standard ResNet-50 models.

To optimize post-hoc calibration methods, 10% of the ImageNet-1k validation set is randomly selected for parameter tuning. All metrics are reported on the remaining 90% of the validation set to ensure methodological consistency. For ImageNet-C, we ensure that the images used for tuning the post-hoc calibration parameters are excluded from their corresponding corrupted versions, to prevent data leakage.

**Post-hoc Calibration Techniques.** To systematically evaluate the effectiveness of post-hoc calibration techniques across various models, we evaluate for post-hoc calibration methods: Temperature Scaling (TS), Ensemble Temperature Scaling (ETS), Isotonic Regression (IRM), and Spline Calibration (SPL). Details are provided in Appendix C.

**Calibration Metrics.** Throughout our analysis, we follow Minderer et al. (2021) and estimate the Expected Calibration Error (ECE) using 15 equal-mass bins as our primary calibration metric. To

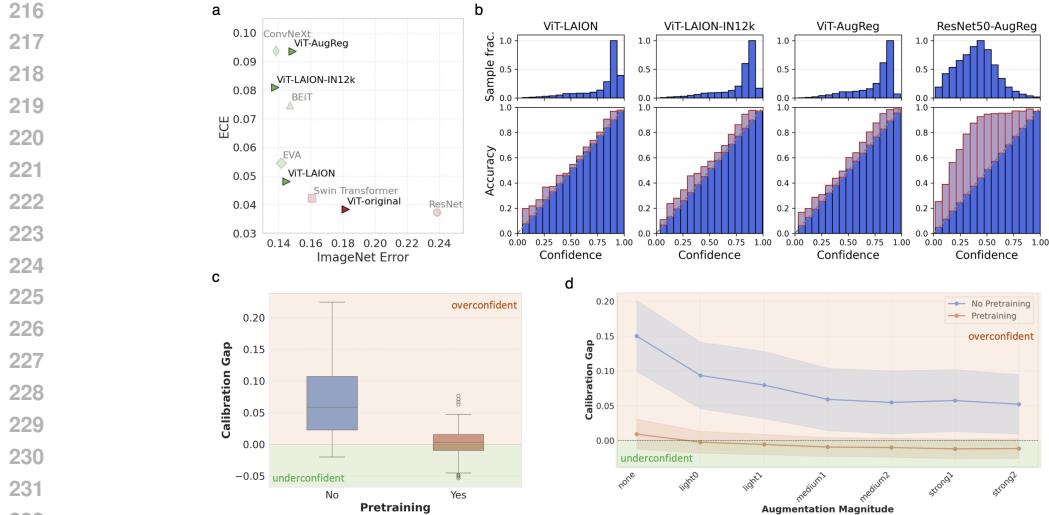


Figure 2: (a) ECE vs. classification error shows that the same ViT architecture trained with modern training techniques (green markers) have higher calibration errors than the original ViT model (red marker). Grey markers represent models from previous experiments that are included for comparison. (b) Reliability diagrams reveal systematic underconfidence across models pretrained on large datasets (ViT-LAION and ViT-LAION-IN12k) or aggressive regularization and augmentation techniques (ViT-AugReg and ResNet50-AugReg). (c) Box plot of calibration gap (difference between mean predicted confidence and accuracy) across 1008 ViT models confirms that pretraining systematically shifts models toward underconfidence. (d) Calibration gap as a function of augmentation magnitude demonstrates that increasing augmentation magnitude results in more underconfident models.

provide a more comprehensive analysis, we present additional results in the supplementary material, using different bin sizes and alternative metrics (e.g. Brier score and negative log-likelihood as proper scoring rules).

#### 4.2 LARGE-SCALE MODELS EXHIBIT SYSTEMATIC IN-DISTRIBUTION UNDERCONFIDENCE

First, we investigate the inherent in-distribution calibration properties of neural networks, before applying any post-hoc calibration techniques. Unlike the findings of Minderer et al. (2021), who reported concurrent improvements in accuracy and calibration for then-current models, our investigation reveals a significant divergence in this relationship for contemporary large-scale models (Figure 1a). While recent model innovations have substantially improved classification performance, they have also demonstrated an increasing calibration error, showing an emerging trade-off between these performance aspects.

Closer examination of the reliability diagrams (Figure 1b) reveals that the increased ECE of the large-scale models is due to a systematic underconfidence in in-distribution predictions – a notable departure from the widespread overconfidence documented in previous calibration literature.

Although this underconfidence increases the overall calibration error, it indicates a different calibration regime that could be advantageous for deployment in high-stakes domains.

#### 4.3 EXPLORING FACTORS INFLUENCING CALIBRATION BEHAVIOR

While our previous experiments reveal systematic underconfidence in large-scale models, the underlying mechanisms driving these phenomena remain unclear. To gain insights into these mechanisms, we conduct a controlled experiment that isolated the influence of training methodology from architectural design.

In this experiment, we utilize the same Vision Transformer architecture as in the previous experiments, while systematically varying the training settings. Specifically, we first explore different

270 pre-training pathways by pre-training the Vision Transformer on the large-scale LAION dataset  
 271 (Cherti et al., 2023). Then the model is further trained either by (1) direct fine-tuning on ImageNet-1k  
 272 (ViT-LAION) or (2) sequential fine-tuning on ImageNet-12k and then fine-tuned on ImageNet-1k  
 273 (ViT-LAION-IN12k). Second, we investigate the impact of aggressive augmentation and regularization  
 274 techniques for the ViT (ViT-AugReg) and the ResNet50 architecture (ResNet50-AugReg) while  
 275 maintaining the same dataset for training as the original models (Steiner et al., 2022).

276 As illustrated in Fig. 2a, both methodological variations substantially improve  
 277 the classification accuracy for the ViT models while concurrently increasing ECE.  
 278 This accuracy-calibration trade-off is even  
 279 more pronounced for ResNet50-AugReg,  
 280 where modern training techniques boost ac-  
 281 curacy from 76.2% to 80.4% but cause ECE  
 282 to increase from 0.037 to 0.408, a tenfold  
 283 increase in miscalibration that exemplifies how  
 284 contemporary training practices can severely  
 285 compromise model calibration. The reliability  
 286 diagrams in Figure 2b reveal that these modern  
 287 training methods induce systematic undercon-  
 288 fidence in both architectures, transforming even  
 289 traditional models like ResNet50 from their  
 290 typical overconfident behavior to the undercon-  
 291 fident regime characteristic of contemporary  
 292 large-scale models.

292 To further strengthen our analysis on the causes  
 293 underlying the observed underconfidence, we  
 294 conducted a large-scale ablation study with 1,008 ViT models of different sizes as well as differ-  
 295 ent pretraining, regularization, augmentation and fine-tuning strategies. Detailed parameter config-  
 296 urations and the complete ablation design are provided in Appendix B. This ablation study further  
 297 supported our conclusions and revealed a consistent mechanism behind underconfidence: stronger  
 298 augmentations induce underconfidence (see Fig. 2d), and pretraining consistently amplifies this  
 299 effect across all ViT variants (see Fig. 2c).

300 Beyond the empirical results, Vicinal Risk Minimization (VRM) provides a theoretical framework  
 301 to explain the underconfidence caused by stronger augmentations and pretraining. VRM minimizes  
 302 the expected loss not only on the empirical samples  $x$  but also on a neighborhood  $v(x)$  induced by  
 303 data augmentation:

$$\mathcal{L}_{\text{VRM}} = \mathbb{E}_{(x,y) \sim \mathcal{D}} \mathbb{E}_{x' \sim v(x)} \ell(f(x'), y) \quad (5)$$

306 Therefore, the model also assigns probability mass to neighborhoods around each training sample.  
 307 When augmentations become stronger or more diverse (e.g., RandAugment, Mixup), the vicinal  
 308 distribution becomes increasingly smoother and more dispersed, effectively enlarging the support  
 309 region around each labeled example. This has two key implications relevant to the observed under-  
 310 confidence: First, a wider vicinal distribution forces the classifier to assign similar probabilities to  
 311 a broader set of augmented variants, which encourages smoother logits. This naturally pushes pre-  
 312 dictions toward underconfidence rather than overconfidence. Intuitively, the model learns to hedge  
 313 its predictions across the enlarged neighborhood  $v(x)$ , resulting in lower confidence values even  
 314 for correctly classified samples. Second, pretraining amplifies VRM’s smoothing effect. Pretrained  
 315 models already encode broader invariances and feature smoothness learned across large datasets.  
 316 When combined with strong augmentations during fine-tuning, the model effectively samples from  
 317 an even more dispersed vicinal distribution. This compounds the smoothing effect on logits, leading  
 318 to stronger underconfidence - exactly what we observe empirically in our ablation study. Models  
 319 pretrained on LAION-400M and subsequently fine-tuned with strong augmentation exhibit the most  
 320 pronounced underconfidence, consistent with this compounding effect.

321 This theoretical viewpoint helps explain the universal trends across all 1,008 ViT models we eval-  
 322 uated and supports the conclusion that underconfidence emerges primarily from properties of the  
 323 training strategy rather than from architectural biases. The systematic nature of our findings - span-  
 ning different model scales, augmentation strategies, and pretraining regimes - demonstrates that the

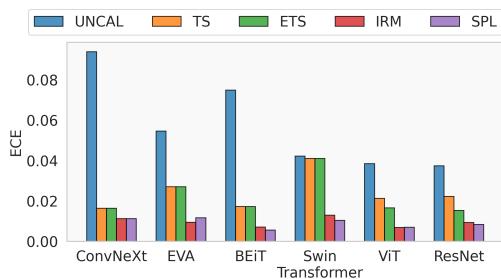


Figure 3: Comparative evaluation of post-hoc calibration methods shows that simple temperature scaling (TS) is sufficient to align the calibration performance of large-scale models with that of traditional models.

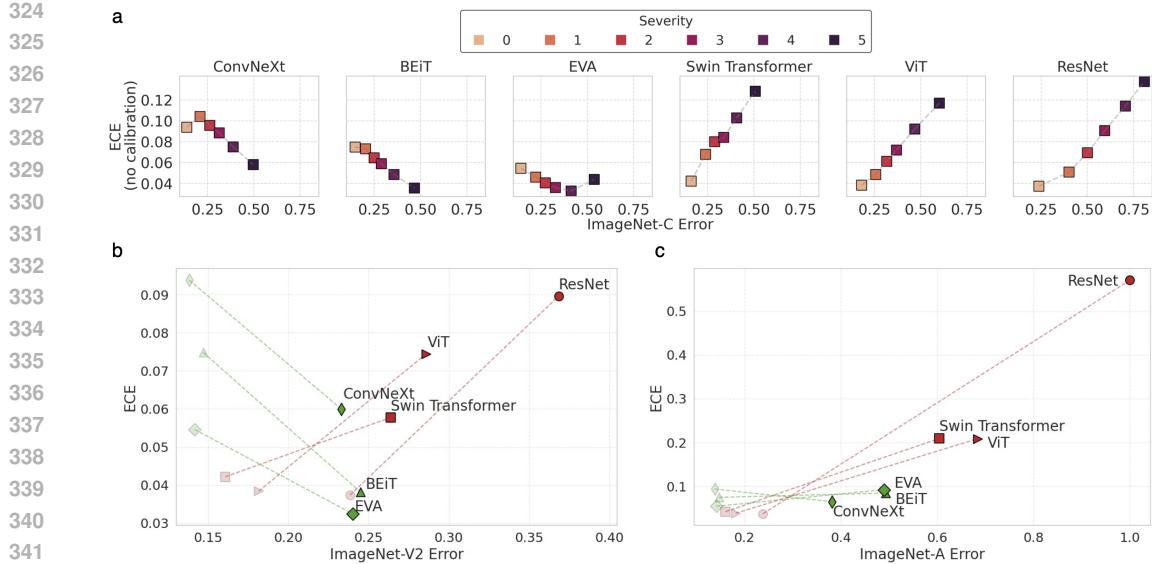


Figure 4: (a) Classification error and Expected Calibration Error (ECE) on ImageNet-C for the uncalibrated models. Severity 0 refers to the clean ImageNet test set. The calibration of large-scale models is more robust to distribution shift than past models. (b) Classification error and ECE on ImageNet-V2 and (c) on ImageNet-A, comparing large-scale models (green markers) with traditional models (red markers). For comparison, the performance of the same models on the clean (in-distribution) ImageNet test set is shown by grey markers of the same type, allowing the distribution shift effect to be visualised. Dotted lines connect each model’s in-distribution result with its corresponding out-of-distribution result.

VRM perspective provides a unifying framework for understanding calibration behavior in modern vision models.

#### 4.4 POST-HOC CALIBRATION FOR IN-DISTRIBUTION PREDICTIONS

Next, we evaluate the effectiveness of post-hoc calibration techniques in addressing the in-distribution miscalibration (Fig. 3).

Remarkably, the simple TS approach is sufficient to align the calibration performance of large-scale models with that of traditional models. However, while ETS theoretically offers greater flexibility by incorporating ensemble-based transformation, it provides no measurable benefits for large-scale models compared to simple TS, despite its higher expressive power.

Among the evaluated methods, isotonic regression (IRM) and spline calibration (SPL) achieve the best calibration quality, outperforming temperature-based approaches consistently across all architectures. This is likely due to their ability to learn a more flexible, nonlinear transformation of the confidence scores.

#### 4.5 CALIBRATION UNDER DISTRIBUTION SHIFT

Building on our observation that large-scale models exhibit systematic underconfidence on in-distribution data, we now investigate how their calibration properties change when they are faced with synthetic and real-world distribution shifts.

**Synthetic Distribution Shifts** We first analyze model calibration under controlled synthetic corruptions using ImageNet-C (Figure 4a). While traditional models follow the well-documented pattern of decreasing accuracy and increasing calibration error as corruption severity increases, large-scale models exhibit a fundamentally different behavior. As the severity of corruption increases, these models show the expected drop in classification accuracy and a decrease in ECE values.

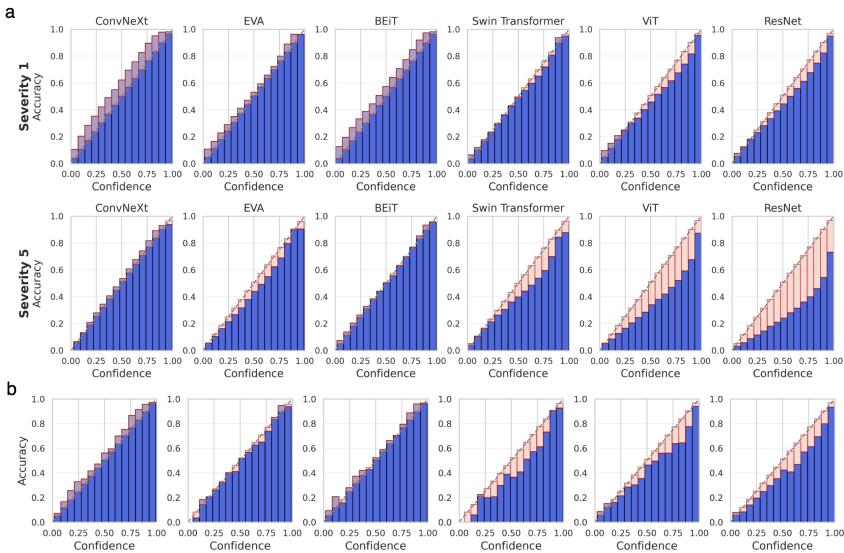


Figure 5: Reliability diagrams comparing predicted confidence with empirical accuracy for (a) artificial distribution shifts of ImageNet-C at severity levels 1 and 5, averaged across all 19 corruption types, and (b) real-world distribution shift using ImageNet-V2. Reliability diagrams for intermediate severity levels and ImageNet-A are provided in the supplementary material.

This counterintuitive improvement in calibration metrics can be explained by examining the underlying dynamics of confidence. Previous studies have shown that distribution shifts typically lead to an increase in model confidence compared to accuracy. For large-scale models that start with underconfidence, this shift-induced increase in confidence acts as a corrective mechanism, bringing predictions closer to actual accuracy levels.

The reliability diagrams in Figure 5a confirm these patterns: traditional models transition to extreme overconfidence, while large-scale models improve calibration while maintaining mild underconfidence.

**Real-World Distribution Shift** To validate whether these findings generalize beyond synthetic perturbations, we examine calibration behavior under real-world distribution shifts using ImageNet-V2 and ImageNet-A (Figure 4b and 4c).

On ImageNet-V2, large-scale models demonstrate consistent improvements in calibration, with ECE decreasing between 36% and 49%. In contrast, traditional models suffer from an increase in ECE ranging from 40% to 140%, reflecting their strong overconfidence under distribution shift (Figure 5b).

The severe distribution shift in ImageNet-A further amplifies these differences. While traditional models experience substantial calibration degradation, large-scale models maintain relatively stable calibration. Notably, ConvNeXt achieves a slight ECE improvement despite the highly challenging nature of the shift.

#### 4.6 POST-HOC CALIBRATION UNDER DISTRIBUTION SHIFT

Finally, we evaluate the performance of post-hoc calibration methods under varying levels of distribution shift. As shown in Figure 6, temperature scaling consistently reduces the ECE for ViT and ResNet50 compared to the uncalibrated baseline as expected. However, the recalibration behavior for large-scale models is different, with its effectiveness dependent on the severity of the distribution shift. While these methods can significantly improve calibration under in-distribution conditions and mild corruptions (severity levels 1-2), their effectiveness decreases as the severity of the distribution shift increases. We observe that the performance of post-hoc calibration methods can degrade under severe distribution shifts to levels worse than those of uncalibrated models. However, overall the

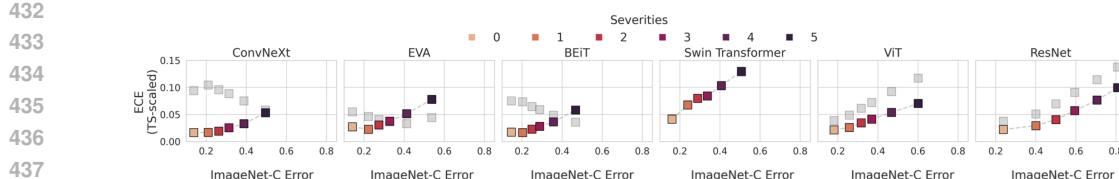


Figure 6: Analysis of classification error and Expected Calibration Error (ECE) across five severity levels of synthetic distribution shifts in ImageNet-C. While temperature scaling perform well under in-distribution conditions, its effectiveness declines with increasing shift magnitude, becoming even counterproductive. The gray markers represent uncalibrated results for comparison.

calibration error of large-scale models under distribution shift remains comparable or smaller than the CE of traditional models.

One underlying cause of this phenomenon is that large-scale models are calibrated on underconfident predictions from the in-distribution validation set, resulting in an increased global confidence. The subsequent application of TS to partially overconfident out-of-distribution samples then further exacerbates this overconfidence. However, observed behaviour of re-calibrated large scale models under distribution shift cannot be attributed solely to underconfidence. Our results reveal model-dependent responses to post-hoc calibration that follow patterns more complex than simple underconfidence correction would predict. EVA, for instance, exhibits distinct temperature scaling behavior at low shift severities that cannot be explained solely by its underconfidence pattern, demonstrating that post-hoc calibration effectiveness depends on model-specific confidence mechanisms beyond global underconfidence.

These findings are consistently observed across additional calibration metrics (ETS, IR, and SPL) as well as on real-world distribution shift datasets including ImageNet-V2 and ImageNet-A (see Appendix).

## 5 LIMITATIONS AND TAKEAWAYS

While our analysis provides robust evidence of systematic underconfidence in large-scale models, our focus in this paper is primarily on characterizing these phenomena rather than fully investigating their underlying causes. Although we conducted experiments in Section 4.3 to begin exploring these factors, these initial investigations could be extended in the future. Furthermore, our study focuses on vision-only classification models. Extending this analysis to vision-language models (e.g., LLaVA, Qwen-VL) represents an important direction for future work, though such models require different evaluation protocols due to their generative nature.

Based on our findings, we offer the following key insights for researchers and practitioners:

- **Underconfidence:** Large-scale models exhibit systematic underconfidence, which provides a practical advantage in safety-critical applications, as conservative uncertainty estimates reduce the risk of errors arising from overconfidence.
- **Distribution Shift Robustness:** Large-scale models demonstrate robust calibration under distribution shifts, ensuring that practitioners can deploy them in dynamic environments where data distributions evolve over time.
- **Limitations of Recalibration Techniques:** Post-hoc calibration methods can lead to counterproductive results for severe distribution shifts in large-scale models, so practitioners should exercise caution when applying these techniques in dynamic environments.
- **Determinants of Calibration Properties:** Calibration properties are predominantly determined by the specifics of the training procedures, rather than by architectural design choices.
- **Best Practices for Model Selection:** Among the evaluated models, ConvNeXt shows the most favorable trade-off between accuracy and calibration, providing comparatively reliable uncertainty estimates after recalibration, even under distribution shifts.

486 REFERENCES  
487

488 Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers.  
489 *arXiv preprint arXiv:2106.08254*, 2021.

490 Jochen Bröcker. Reliability, sufficiency, and the decomposition of proper scores. *Quarterly Journal*  
491 *of the Royal Meteorological Society*, 135(643):1512–1519, July 2009. ISSN 1477-870X.

492

493 Jiacheng Cheng and Nuno Vasconcelos. Towards calibrated multi-label deep neural networks. *2024*  
494 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

495 Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gor-  
496 don, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for con-  
497 trastive language-image learning. In *2023 IEEE/CVF Conference on Computer Vision and Pattern*  
498 *Recognition (CVPR)*, pp. 2818–2829. IEEE, June 2023. doi: 10.1109/cvpr52729.2023.00276.

499

500 Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. RandAugment: Practical Automated  
501 Data Augmentation with a Reduced Search Space. *Advances in Neural Information Processing*  
502 *Systems*, 33:18613–18624, 2020.

503 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale  
504 hierarchical image database. *Conference on Computer Vision and Pattern Recognition*, pp. 248–  
505 255, 2009.

506 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
507 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszko-  
508 reit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at  
509 scale. *International Conference on Learning Representations*, 2021.

510

511 Yuxin Fang, Wen Wang, Binhui Xie, Quan-Sen Sun, Ledell Yu Wu, Xinggang Wang, Tiejun Huang,  
512 Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning  
513 at scale. *Conference on Computer Vision and Pattern Recognition*, pp. 19358–19369, 2022.

514 Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural  
515 networks. *International conference on machine learning*, pp. 1321–1330, 2017.

516

517 Kartik Gupta, Amir Rahimi, Thalaiyasingam Ajanthan, Thomas Mensink, Cristian Sminchisescu,  
518 and Richard Hartley. Calibration of neural networks using splines. *International Conference on*  
519 *Learning Representations*, 2021.

520 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
521 nition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
522 770–778, 2016.

523 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked au-  
524 toencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*  
525 *vision and pattern recognition*, pp. 16000–16009, 2022.

526

527 Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common cor-  
528 ruptions and perturbations. *International Conference on Learning Representations*, 2019.

529

530 Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial  
531 examples. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,  
532 pp. 15262–15271, 2021.

533 Meelis Kull and Peter Flach. Novel decompositions of proper scoring rules for classification:  
534 Score adjustment as precursor to calibration. In *Machine Learning and Knowledge Discovery*  
535 *in Databases (ECML PKDD)*, volume 9284 of *Lecture Notes in Computer Science*, pp. 68–85.  
536 Springer, 2015. doi: 10.1007/978-3-319-23528-8\_5.

537 Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter  
538 Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with  
539 dirichlet calibration. In *Advances in Neural Information Processing Systems*, pp. 12295–12305,  
2019.

540 Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predic-  
 541 tive uncertainty estimation using deep ensembles. *Advances in Neural Information Processing*  
 542 *Systems*, 30, 2017.

543

544 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.  
 545 Swin transformer: Hierarchical vision transformer using shifted windows. *Proceedings of the*  
 546 *IEEE/CVF International Conference on Computer Vision*, pp. 10012–10022, October 2021.

547

548 Zhuang Liu, Hanzi Mao, Chao Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie.  
 549 A convnet for the 2020s. *Conference on Computer Vision and Pattern Recognition*, pp. 11966–  
 550 11976, 2022.

551

552 Michal Lukasik, Srinadh Bhojanapalli, A. Menon, and Surinder Kumar. Does label smoothing  
 553 mitigate label noise? *ArXiv*, March 2020. URL <https://www.semanticscholar.org/paper/82c77a88969ac0e3a4e55c9a7dc5ced4afee0225>.

554

555 Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby,  
 556 Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. *Advances*  
 557 *in Neural Information Processing Systems*, 34:15682–15694, 2021.

558

559 Allan H. Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*,  
 560 12(4):595–600, 1973a. doi: 10.1175/1520-0450(1973)012<0595:ANVPOT>2.0.CO;2.

561

562 Allan H Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*  
 563 and *Climatology*, 12(4):595–600, 1973b.

564

565 Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua  
 566 Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty?  
 567 Evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing*  
 568 *Systems*, pp. 13991–14002, 2019.

569

570 Alexandre Perez-Lebel, Marine Le Morvan, and Gaël Varoquaux. Beyond calibration: estimat-  
 571 ing the grouping loss of modern neural networks. In *The Eleventh International Conference on*  
 572 *Learning Representations (ICLR)*, Kigali, Rwanda, 2023. URL <https://arxiv.org/abs/2210.16315>.

573

574 Teodora Popordanoska, Sebastian Gregor Gruber, Aleksei Tiulpin, Florian Buettner, and Matthew B  
 575 Blaschko. Consistent and asymptotically unbiased estimation of proper calibration errors. *Inter-  
 576 national Conference on Artificial Intelligence and Statistics*, pp. 3466–3474, 2024.

577

578 Rahul Rahaman and Alexandre Thiery. Uncertainty Quantification and Deep Ensembles. *Advances*  
 579 *in Neural Information Processing Systems*, 34:20063–20075, 2021.

580

581 Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers  
 582 generalize to imagenet? In *International conference on machine learning*, pp. 5389–5400. PMLR,  
 583 2019.

584

585 Andreas Peter Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and  
 586 Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers.  
 587 *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.

588

589 Linwei Tao, Younan Zhu, Haolan Guo, Minjing Dong, and Chang Xu. A benchmark study on  
 590 calibration. In *The Twelfth International Conference on Learning Representations*, 2024.

591

592 Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Micha-  
 593 lak. On Mixup Training: Improved Calibration and Predictive Uncertainty for Deep Neural Net-  
 594 works. *Advances in Neural Information Processing Systems*, 32, 2019.

595

596 Christian Tomani, Daniel Cremers, and Florian Buettner. Parameterized temperature scaling for  
 597 boosting the expressive power in post-hoc uncertainty calibration. *European Conference on Com-  
 598 puter Vision*, pp. 555–569, 2022.

594 Christian Tomani, Futa Waseda, Yuesong Shen, and Daniel Cremers. Beyond in-domain scenarios:  
 595 Robust density-aware calibration. In *Proceedings of the 40th International Conference on*  
 596 *Machine Learning*, volume 202, pp. 34344–34368. PMLR, 2023.

597

598 Deng-Bao Wang, Lei Feng, and Min-Ling Zhang. Rethinking calibration of deep neu-  
 599 ral networks: Do not be afraid of overconfidence. *Neural Information Process-  
 600 ing Systems*, January 2021. URL <https://www.semanticscholar.org/paper/2c6ee2692351ae09ad9e933b0ba1eee0fd04c4cc>.

601

602 Miao Xiong, Ailin Deng, Pang Wei Koh, Jiaying Wu, Shen Li, Jianqing Xu, and Bryan Hooi.  
 603 Proximity-informed calibration for deep neural networks. In *Advances in Neural Information*  
 604 *Processing Systems*, volume 36, pp. 1–24, 2023.

605

606 Sangdoo Yun, Dongyoon Han, Sanghyuk Chun, Seong Joon Oh, Youngjoon Yoo, and Junsuk Choe.  
 607 CutMix: Regularization Strategy to Train Strong Classifiers With Localizable Features. *Inter-  
 608 national Conference on Computer Vision*, pp. 6022–6031, 2019.

609

610 Chang-Bin Zhang, Peng-Tao Jiang, Qibin Hou, Yunchao Wei, Qi Han, Zhen Li, and Ming-Ming  
 611 Cheng. Delving deep into label smoothing. *IEEE Transactions on Image Processing*, November  
 2020a. URL <https://ieeexplore.ieee.org/document/9464693/>.

612

613 Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical  
 614 risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.

615

616 Jize Zhang, Bhavya Kailkhura, and T Yong-Jin Han. Mix-n-match: Ensemble and compositional  
 617 methods for uncertainty calibration in deep learning. *International conference on machine learn-  
 618 ing*, pp. 11117–11128, 2020b.

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 APPENDIX  
649650 A SUMMARY  
651652 To support and expand upon our core findings, we provide additional metrics, experimental results,  
653 and technical details.  
654655 Section B provides comprehensive technical details of all models evaluated in this study, including  
656 the six primary benchmark models and the 1,008 ViT models trained for our ablation study. This  
657 section includes timm specifiers for reproducibility and model complexity metrics.  
658659 Section C describes the post-hoc calibration methods investigated in our study: Temperature Scaling  
660 (TS), Ensemble Temperature Scaling (ETS), Accuracy-Preserving Isotonic Regression (IRM), and  
661 Spline Calibration (SPL).  
662663 Section D introduces calibration metrics that complement the Expected Calibration Error (ECE)  
664 used in the main paper. These include Maximum Calibration Error (MCE), Root Mean Square  
665 Calibration Error (RMSCE), Root Brier Score (RBS), and Negative Log-Likelihood (NLL). These  
666 metrics capture different aspects of calibration quality and help validate the robustness of our find-  
667 ings.  
668669 Section E presents further in-distribution calibration results, demonstrating that the inverse rela-  
670 tionship between classification and calibration errors holds across different ECE configura-  
671 tions, bin resolutions, and alternative calibration metrics. This section also includes additional results for post-  
672 hoc calibration techniques under distribution shift.  
673674 Section F presents reliability diagrams illustrating model calibration under varying degrees of dis-  
675 tribution shift, including ImageNet-C at intermediate severity levels (2, 3, and 4), ImageNet-A, and  
676 individual reliability diagrams for all 19 synthetic corruptions at severity levels 3 and 5.  
677678 Section G evaluates the effectiveness of post-hoc calibration methods under real-world distribu-  
679 tion shifts using the ImageNet-V2 and ImageNet-A datasets, confirming the patterns observed on  
680 ImageNet-C.  
681682 Section H provides a statement on LLM usage during manuscript preparation.  
683684 Sections I and J cover the availability of the code and the publicly available datasets that can be used  
685 to reproduce the experimental results and provide opportunities for further research.  
686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **B MODEL SPECIFICATIONS**  
703704 This section provides technical details for all models evaluated in this study: our primary benchmark  
705 models and the 1,008 ViT models trained for our ablation study.  
706707 **B.1 BENCHMARK MODELS**  
708709 Throughout the paper, we consistently analyze the following six models:  
710

- 711 • **ResNet-50** (He et al. (2016)): A widely used baseline convolutional architecture, trained  
712 on ImageNet-1k.
- 713 • **ViT-B/16** (Dosovitskiy et al. (2021)): A pure transformer architecture pretrained on  
714 ImageNet-21k with supervised learning and fine-tuned on ImageNet-1k. Previous studies  
715 have demonstrated that this architecture has strong calibration properties.
- 716 • **Swin-S3-B** (Liu et al. (2021)): A hierarchical transformer model with shifted window par-  
717 titioning, trained on ImageNet-1k.
- 718 • **BEiT-B/16** (Bao et al. (2021)): A transformer leveraging self-supervised masked image  
719 modeling, pretrained on ImageNet-22k and fine-tuned on ImageNet-1k.
- 720 • **EVA-S/14** (Fang et al. (2022)): A scaled transformer model pretrained on ImageNet-22k  
721 with self-supervised masked image modeling and subsequently fine-tuned on ImageNet-1k.
- 722 • **ConvNeXt-B** (Liu et al. (2022)): A convolutional architecture that integrates transformer-  
723 inspired design principles, pretrained on ImageNet-22k and fine-tuned on ImageNet-1k.  
724

725 To investigate the factors contributing to underconfidence, we additionally evaluate several variants  
726 of ViT and ResNet models that employ traditional architectures but are trained with contemporary  
727 methodologies:  
728

- 729 • **ViT-B/16-LAION**: A Vision Transformer pretrained on the LAION-2B dataset (2 bil-  
730 lion image-text pairs) using contrastive learning objectives, subsequently fine-tuned on  
731 ImageNet-1k for classification (Cherti et al. (2023)).
- 732 • **ViT-B/16-LAION-IN12k**: A Vision Transformer pretrained on LAION-2B through con-  
733 trastive learning, followed by sequential fine-tuning first on ImageNet-12k and then on  
734 ImageNet-1k (Cherti et al. (2023)).
- 735 • **ViT-B/16-AugReg**: A Vision Transformer initially pretrained on ImageNet-21k, then fine-  
736 tuned on ImageNet-1k using extensive data augmentation strategies and advanced regular-  
737 ization techniques including Mixup, Cutmix, and dropout scheduling (Steiner et al. (2022)).
- 738 • **ResNet50-AugReg**: The classical ResNet-50 architecture trained on ImageNet-1k, incor-  
739 porating advanced regularization and augmentation techniques such as CutMix, MixUp,  
740 Label Smoothing, and Random Erasing.  
741

742 This section provides technical details about the neural network models used in our calibration study  
743 (Table 1 and Table 2). All models can be accessed through the PyTorch Image Models (timm) library  
744 (see Table 1 and <https://github.com/huggingface/pytorch-image-models>),  
745 with corresponding checkpoints available on the Hugging Face model hub  
746 (<https://huggingface.co/>). This allows for direct reproducibility of our results. In-  
747 put preprocessing followed the standard procedures specified in each model’s documentation,  
748 including normalization with ImageNet statistics and appropriate resizing. For additional technical  
749 specifications beyond what is provided here, we refer to the model cards available on Hugging Face  
750 under the corresponding model identifiers and the documentation of the timm library.  
751752 **B.2 ABLATION STUDY MODELS**  
753754 Our comprehensive ablation study comprises 1,008 independently trained Vision Transformer mod-  
755 els based on the training protocol from Steiner et al. (2022). We evaluate six architectural variants  
(ViT-Ti/16, ViT-S/16, ViT-S/32, ViT-B/16, ViT-B/32, ViT-L/16) and systematically vary training  
parameters to isolate factors influencing calibration behavior (Table 3).  
756

756

757

Table 1: Timm specifiers of the models used in the study.

758

759

Model	timm Specifier
ResNet-50	resnet50.tv_in1k
ViT-B/16	vit_base_patch16_224.orig_in21k_ft_in1k
Swin-S3-B	swin_s3_base_224.ms_in1k
BEiT-B/16	beit_base_patch16_224.in22k_ft_in22k_in1k
EVA-S/14	eva02_small_patch14_336.mim_in22k_ft_in1k
ConvNeXt-B	convnext_base.fb_in22k_ft_in1k
ViT-LAION	vit_base_patch16_clip_224.laion2b_ft_in1k
ViT-LAION-IN12k	vit_base_patch16_clip_224.laion2b_ft_in12k_in1k
ViT-AugReg	vit_base_patch16_224.augreg2_in21k_ft_in1k

760

761

762

763

764

765

766

767

768

769

770

771

Table 2: Complexities of the models used in the study.

Model	Params (M)	GMACs
ResNet-50	25.6	4.1
ViT-B/16	86.6	16.9
Swin-S3-B	71.1	13.7
BEiT-B/16	86.5	17.6
EVA-S/14	22.1	15.5
ConvNeXt-B	88.6	15.4
ViT-LAION	86.6	16.9
ViT-LAION-IN12k	86.6	16.9
ViT-AugReg	86.6	16.9

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810

811

Table 3: Search space of training parameters for the ablation study.

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

## C POST-HOC CALIBRATION METHODS

832

We investigate the following commonly used post-hoc calibration techniques:

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

Table 4: Augmentation configurations.

Setup	$l$	$m$	$\alpha$
none	0	0	0
light0	2	0	0
light1	2	10	0.2
medium0	2	15	0.2
medium1	2	15	0.5
strong0	2	20	0.5
strong1	2	20	0.8

---

## 864 D ADDITIONAL CALIBRATION METRICS AND THEIR DEFINITIONS 865

866 To validate the robustness of our findings presented in the main paper, we extend our analysis of  
867 model calibration using different bin configurations and complementary calibration metrics. While  
868 our primary investigation focused on Expected Calibration Error (ECE) with 15 bins, we demon-  
869 strate here that our conclusions hold consistently across the following calibration metrics:  
870

871 1. **Maximum Calibration Error (MCE)** quantifies the worst-case miscalibration scenario  
872 by measuring the maximum discrepancy between confidence and accuracy across all bins:  
873

$$874 \quad MCE = \max_j (acc(B_j) - conf(B_j)).$$

875 2. **Root Mean Square Error (RMSCE)** penalizes larger calibration errors more heavily than  
876 ECE by using squared differences:  
877

$$878 \quad RMSCE = \sqrt{\sum_j^m \frac{|B_j|}{n} (acc(B_j) - conf(B_j))^2}.$$

881 3. **Root Brier Score (RBS)** measures the accuracy of probabilistic predictions:  
882

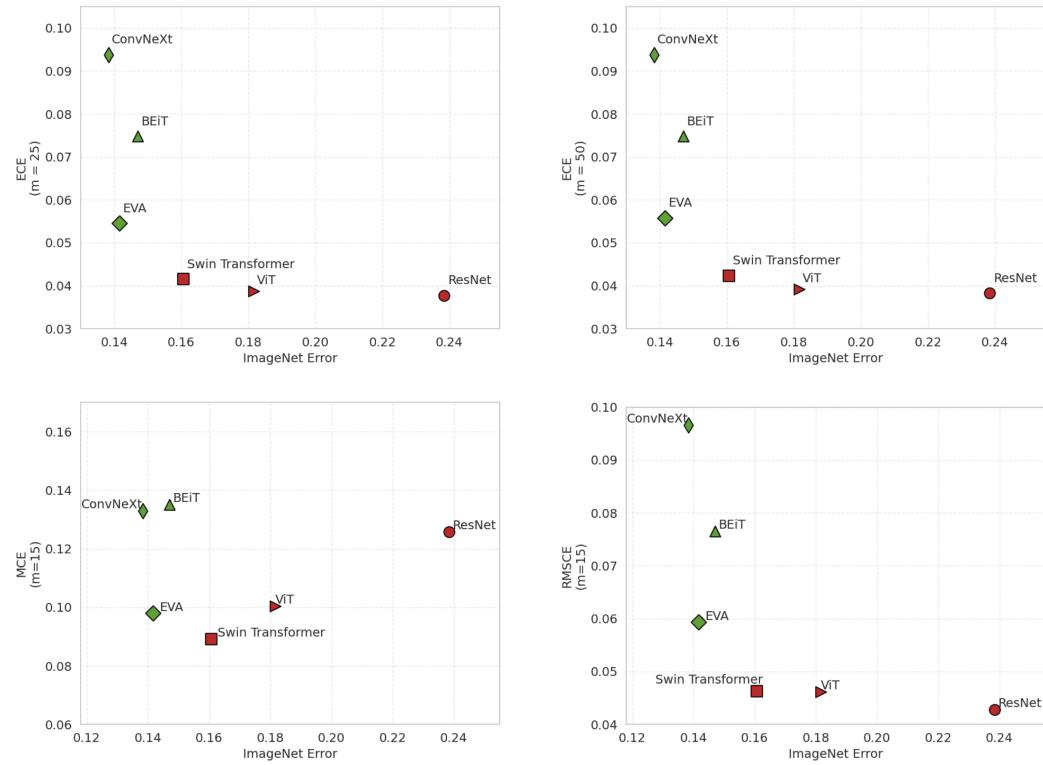
$$883 \quad BS = \sqrt{\frac{1}{n} \sum_{i=1}^n \sum_{c=1}^C (p_{i,c} - y_{i,c})^2},$$

887 where  $p_{i,c}$  represents the predicted probability for class  $c$  of sample  $i$ , and  $y_{i,c}$  is the corre-  
888 sponding one-hot encoded ground truth label.  
889

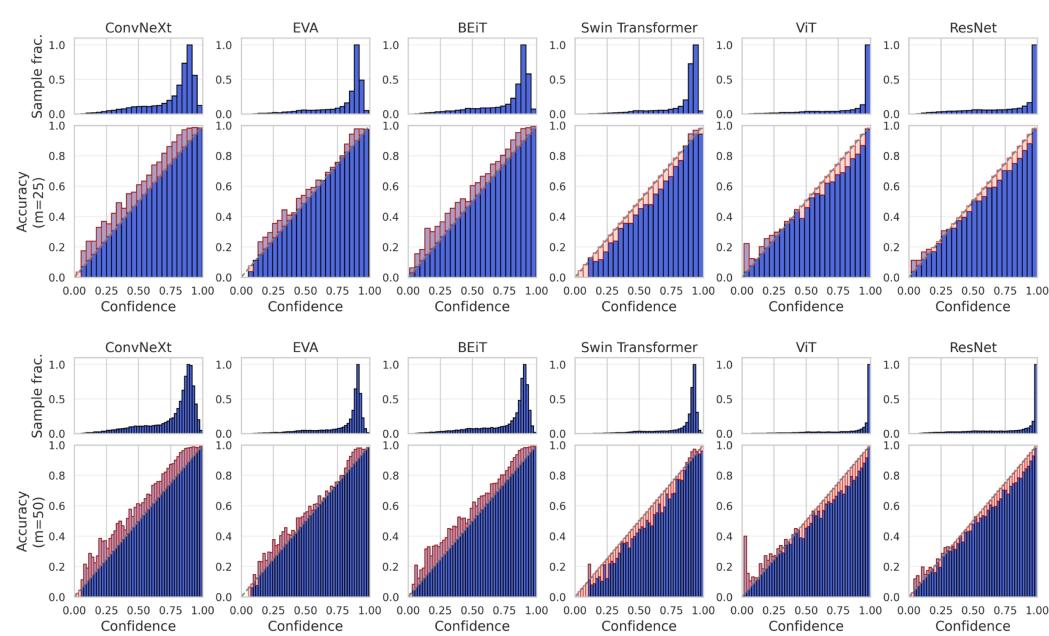
890 4. **Negative Log-Likelihood (NLL)** evaluates the quality of probabilistic predictions by mea-  
891 suring the likelihood of the true labels under the model’s predicted distributions:  
892

$$893 \quad NLL = -\frac{1}{n} \sum_{i=1}^n \log(p_{i,y_i}),$$

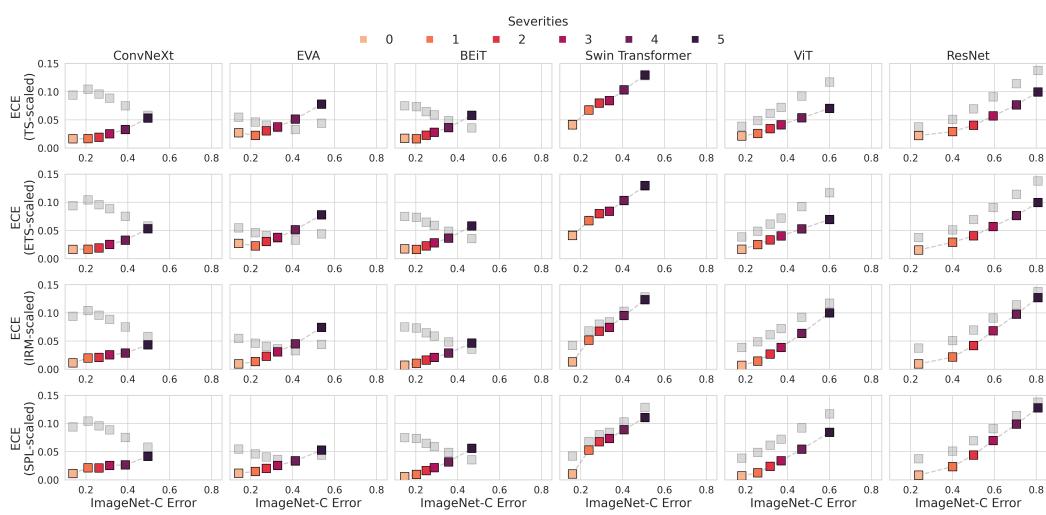
894 where  $p_{i,y_i}$  is the predicted probability for the true class  $y_i$  of sample  $i$ .  
895

918 E FURTHER RESULTS ON IN-DISTRIBUTION CALIBRATION  
919920 E.1 RESULTS FOR IN-DISTRIBUTION CALIBRATION FOR DIFFERENT ECE  
921 CONFIGURATIONS AND TYPES OF CALIBRATION ERROR  
922

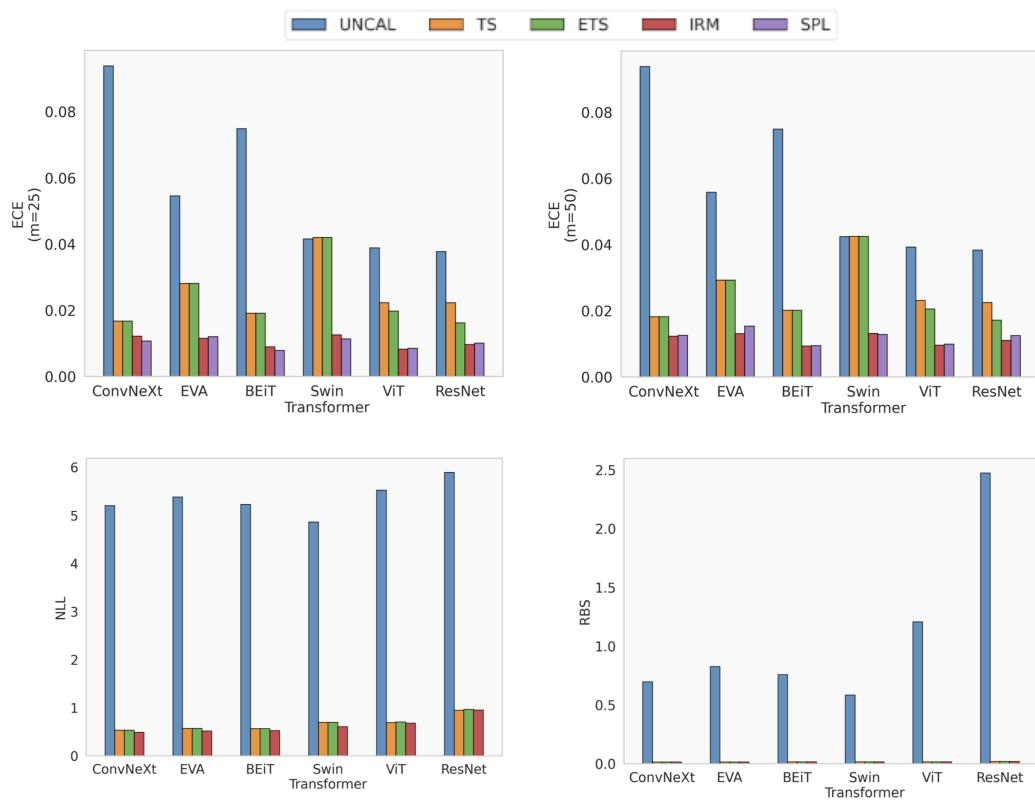
948  
949 Figure 7: Scatter plot illustrating the inverse relationship between ImageNet classification error and  
950 calibration error. The results presented in the main part of the paper hold true for ECE across varying  
951 number of bins ( $m=25$  and  $m=50$ ) and for different types of calibration error, such as Maximum  
952 Calibration Error (MCE) and Root Mean Square Calibration Error (RMSCE).

972 E.2 EFFECT OF BIN RESOLUTION ON RELIABILITY DIAGRAM  
973995 Figure 8: Reliability diagrams calculated with different bin resolutions (25 and 50 bins). The diagrams demonstrate that the observed calibration patterns remain consistent across different bin  
996 counts, supporting the robustness of our findings.  
997

998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

1026  
1027 E.3 RESULTS FOR POST-HOC CALIBRATION TECHNIQUES UNDER DISTRIBUTION SHIFT  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
10431044  
1045 Figure 9: Additional results for the performance of ETS, IR and SPL under distribution shift.  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

1080  
 1081 **E.4 RESULTS FOR POST-HOC CALIBRATION TECHNIQUES FOR DIFFERENT CALIBRATION**  
 1082 **METRICS**



1099 **Figure 10: Comparison of post-hoc calibration effectiveness across multiple evaluation metrics:**  
 1100 **ECE with varying bin resolutions (25, and 50 bins), Root Brier Score (RBS), and Negative Log-**  
 1101 **Likelihood (NLL).**

1134 F FURTHER RESULTS ON RELIABILITY DIAGRAMS UNDER DISTRIBUTION  
1135 SHIFT  
1136

1137 To provide a comprehensive view of model calibration behavior under varying distribution shifts,  
1138 we additionally present the reliability diagrams for ImageNet-C at intermediate severity levels (2, 3,  
1139 and 4) and for ImageNet-A (Figure 11)

1140 These visualizations reveal the progressive changes in calibration behavior as distribution shift sever-  
1141 ity increases, demonstrating how large-scale models maintain their calibration advantage over tra-  
1142 ditional architectures. For traditional architectures (ResNet-50, ViT, and Swin), we observe a con-  
1143 sistent pattern of increasing overconfidence as severity levels progress from 2 to 4. In contrast,  
1144 foundation models (ConvNeXt, EVA, and BEiT) demonstrate remarkable robustness across these  
1145 intermediate severity levels. Their initial underconfidence on in-distribution data gradually dimin-  
1146 ishes as severity increases.

1147 The reliability diagrams for ImageNet-A complement our analysis of ImageNet-V2 presented in the  
1148 main text and provide insights into calibration behavior under particularly challenging conditions.  
1149 On ImageNet-A, traditional architectures exhibit extreme overconfidence across all confidence bins,  
1150 with dramatic gaps between predicted probabilities and actual accuracy rates. Foundation mod-  
1151 els demonstrate significantly better calibrated predictions on ImageNet-A. In particular, ConvNeXt  
1152 maintains relatively well-calibrated predictions across most confidence bins. These reliability di-  
1153 agrams further substantiate our findings that foundation models fundamentally alter the traditional  
1154 calibration paradigm, maintaining better alignment between confidence and accuracy under chal-  
1155 lenging distribution shifts compared to traditional architectures.

1156 Figures 12 and 13 provide a detailed view of calibration behavior across all 19 individual corruption  
1157 types from ImageNet-C. Figure 12 presents reliability diagrams for each corruption type at severity  
1158 level 3, while Figure 13 shows the corresponding diagrams at severity level 5.

1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187

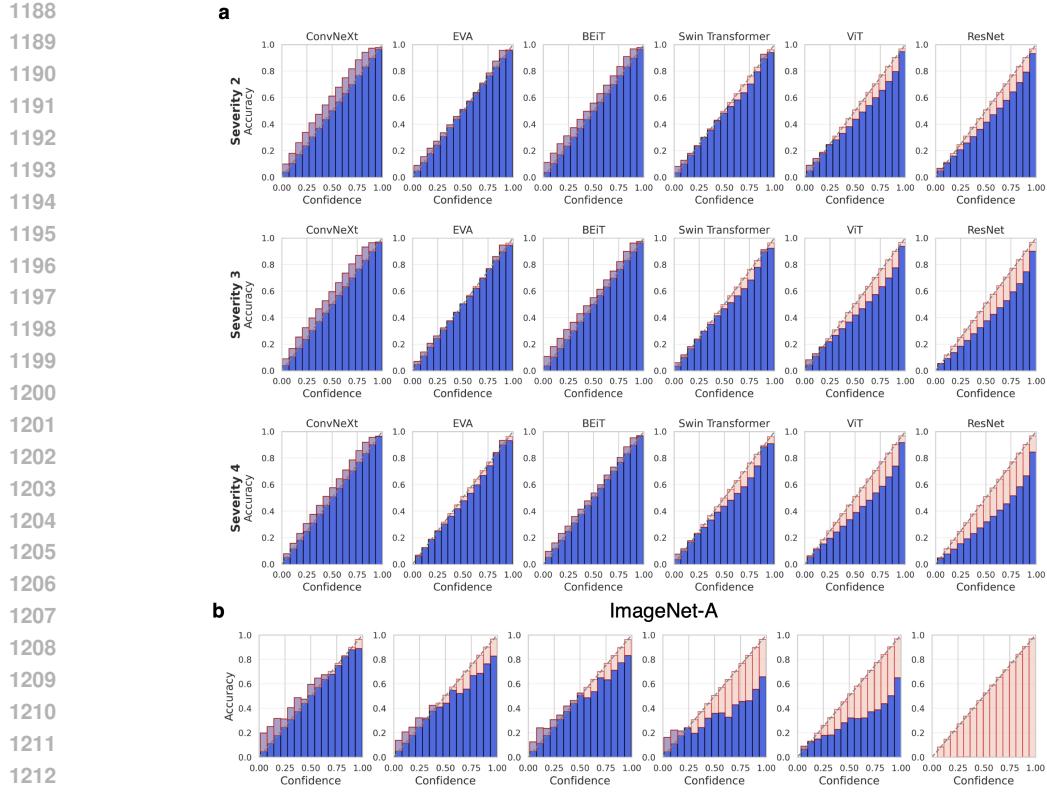


Figure 11: Reliability diagrams ( $m=15$  bins) illustrating model calibration under (a) synthetic distribution shift induced by ImageNet-C corruptions at different severity levels (averaged over all corruptions), and (b) real-world distribution shift as represented by ImageNet-A.

## G EVALUATING POST-HOC CALIBRATION PERFORMANCE UNDER REAL-WORLD DISTRIBUTION SHIFT

To validate our findings beyond synthetic corruptions, we extend our analysis to real-world distribution shifts using the ImageNet-V2 and ImageNet-A datasets (Figure 14). These benchmarks provide complementary perspectives on model robustness: ImageNet-V2 represents a moderate temporal distribution shift, while ImageNet-A introduces severe natural adversarial examples.

The ConvNeXt model demonstrates strong recalibration performance on ImageNet-V2, with post-hoc methods achieving significant ECE reductions. However, as the distribution shift becomes more severe on ImageNet-A, this effectiveness diminishes substantially, with recalibration methods yielding ECE values comparable to or exceeding those of the uncalibrated baseline. This pattern mirrors our observations with synthetic corruptions, where recalibration performance degraded with increasing severity.

EVA exhibits even more pronounced calibration challenges. Even under moderate shifts (ImageNet-V2), post-hoc calibration methods not only fail to improve calibration but actively increase ECE compared to the uncalibrated model. This aligns with trends observed under synthetic corruptions, where EVA’s recalibration performance began deteriorating at lower severity levels than other foundation models. On ImageNet-A, all methods produce substantially higher ECE values relative to the uncalibrated models.

In contrast, traditional architectures demonstrate more consistent responses to calibration techniques. ResNet-50 and ViT show calibration improvements across both benchmarks, though the magnitude of improvement is notably higher on ImageNet-V2 than ImageNet-A. This reflects the increasing challenge of calibration under severe distribution shifts. Nevertheless, the absolute ECE values remain lower for foundation models compared to traditional architectures, even under se-

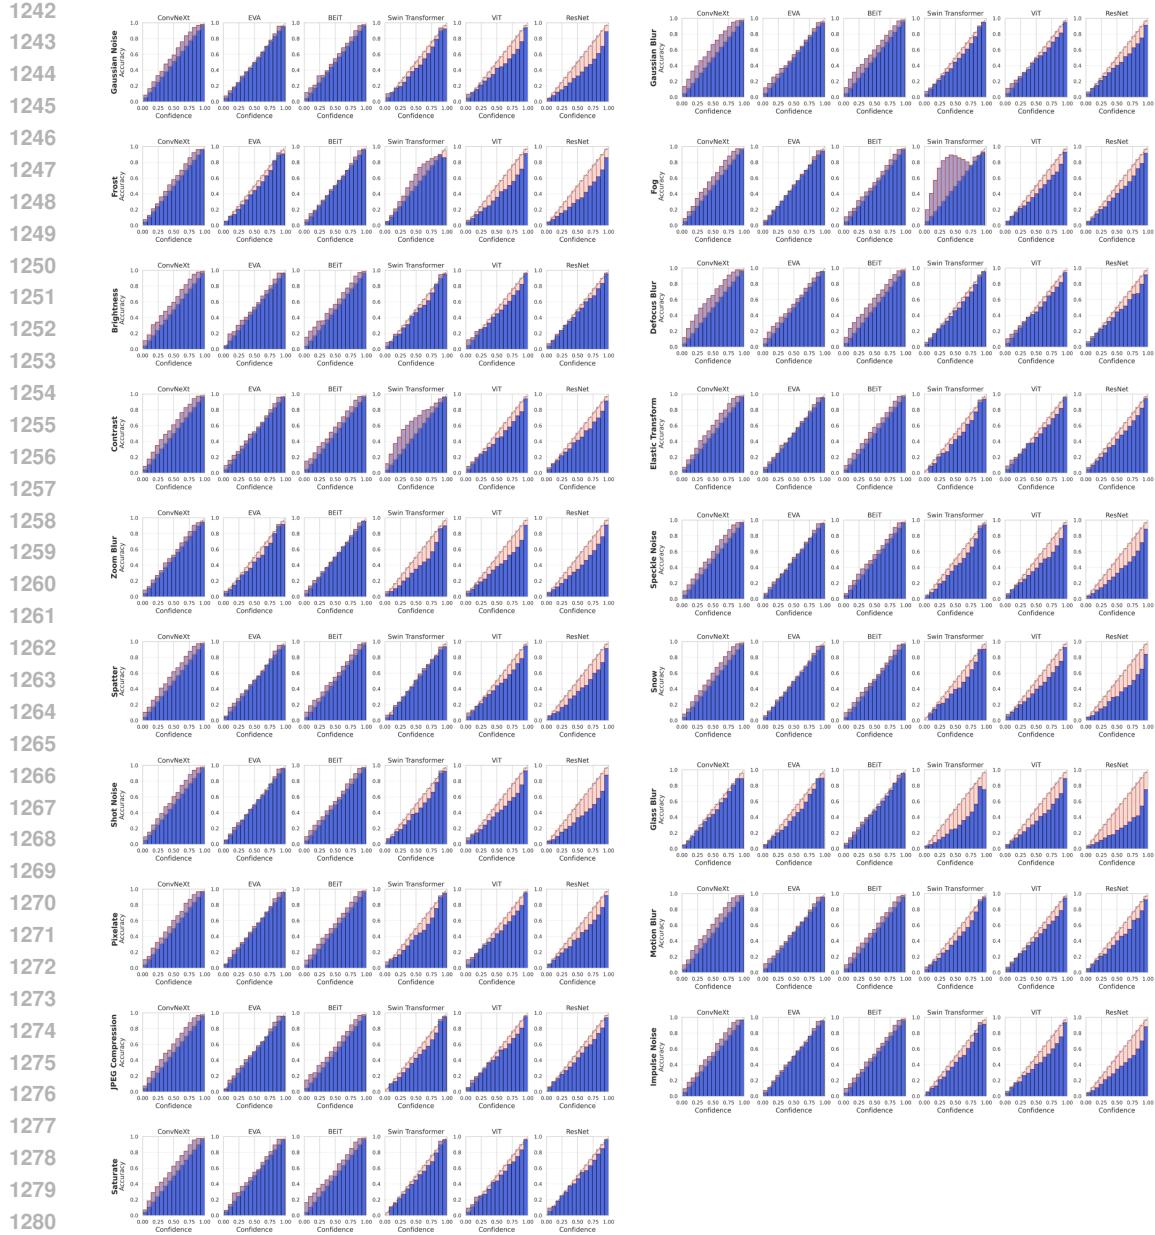
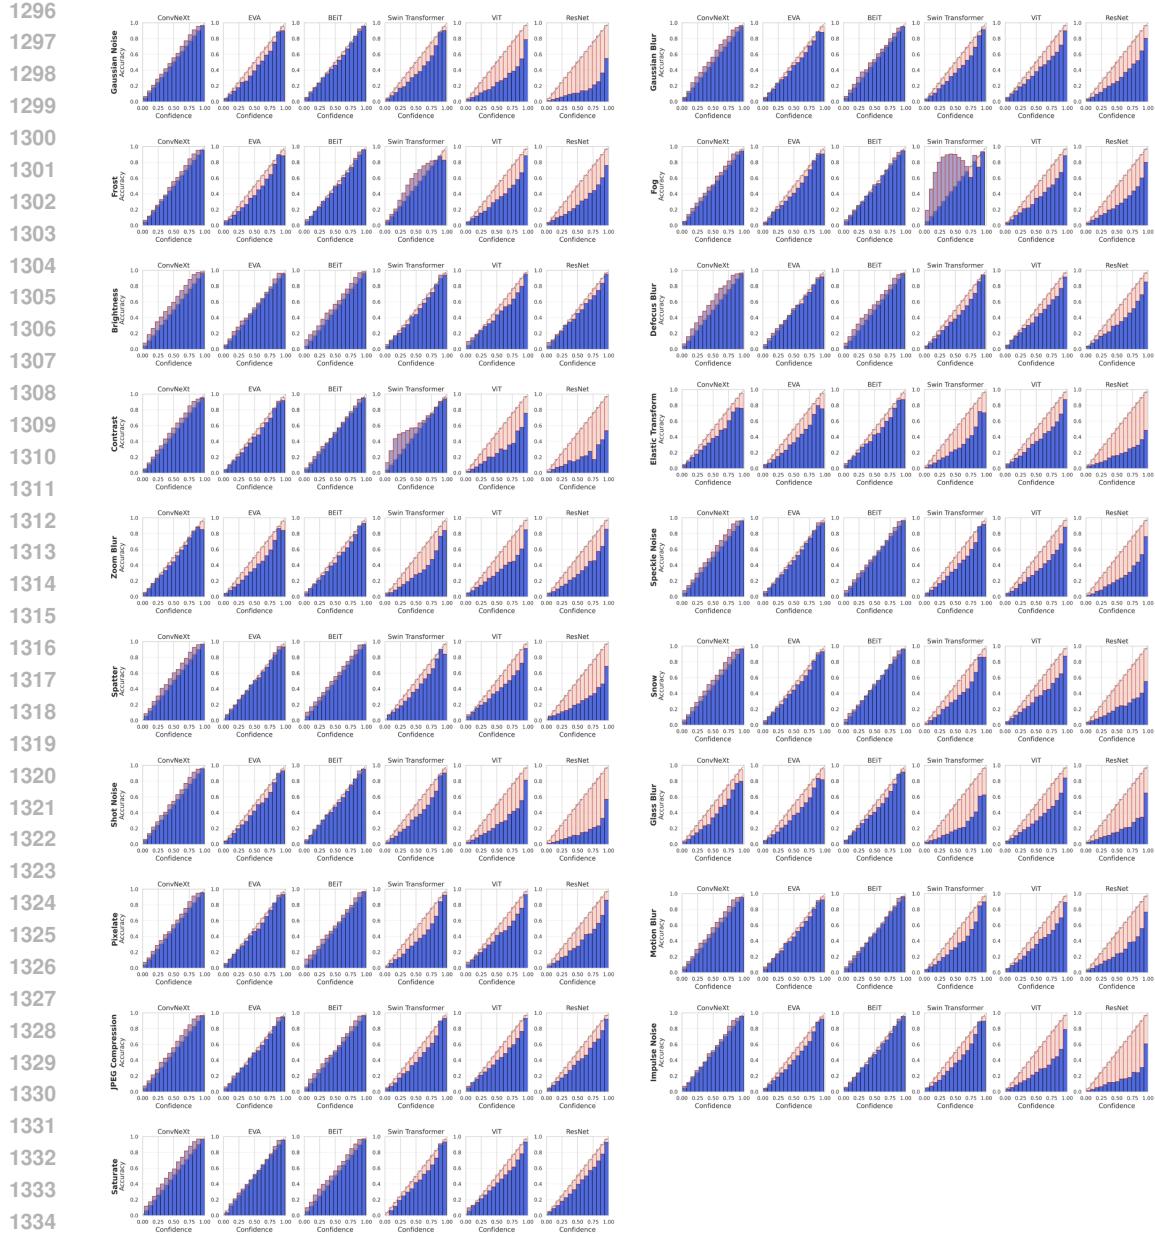


Figure 12: Individual reliability diagrams for all 19 synthetic corruptions of ImageNet-C for severity 3.

vere shifts. Interestingly, the Swin Transformer also exhibits negligible responsiveness to post-hoc calibration across real-world distribution shifts, reaffirming the pattern observed in in-distribution scenarios. This consistent behavior suggests architectural characteristics that fundamentally limit the effectiveness of post-hoc calibration techniques.



1336 Figure 13: Individual reliability diagrams for all 19 synthetic corruptions of ImageNet-C for severity  
1337 5.

## H LLM USAGE STATEMENT

1345 Large Language Models (LLMs) were used exclusively for linguistic refinement and proofreading  
1346 of this manuscript. Specifically, we employed LLMs to improve grammar, sentence structure, and  
1347 overall readability of the text. No LLMs were used for research design, hypothesis generation,  
1348 data analysis, interpretation of results, or the development of core ideas presented in this work. All  
1349 scientific contributions, experimental designs, theoretical insights, and conclusions are entirely the  
product of the authors' original research.

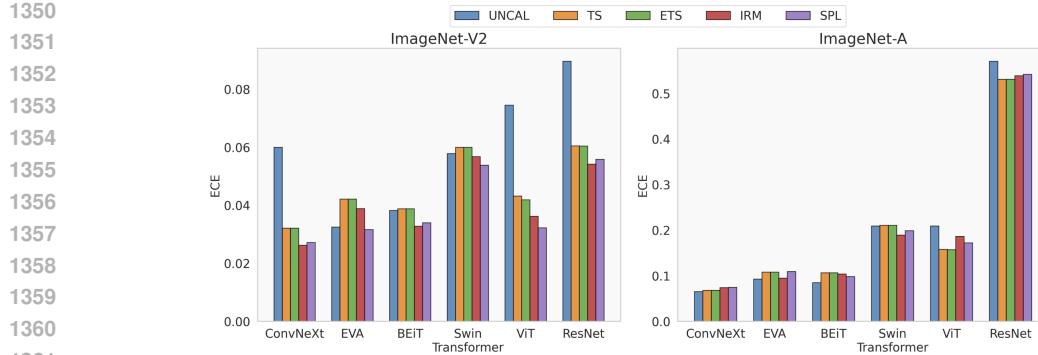


Figure 14: Analysis of post-hoc calibration methods under real-world distribution shift.

## I CODE AVAILABILITY

We implement and analyze the post-hoc calibration methods introduced in the previous section within a newly developed Python package, called ModelTransformer. This package provides a unified framework inspired by the design principles of scikit-learn. The package offers consistent interfaces for fitting and transforming data, enabling parameter estimation on validation datasets and subsequent application to test sets. The complete implementation is available at <https://github.com/XXX/XXX>.

All the code used to generate the analysis and figures in this paper is publicly available at <https://github.com/XXX/XXX>. This repository contains the code that enables the complete reproduction of our experimental results and graphical representations.

## J DATA AVAILABILITY

The complete set of raw and recalibrated model outputs used in this paper is publicly available at <https://doi.org/XX.XXXX/zenodo.XXXXXXXX>.

This extensive collection of datasets enables the full reproduction of our calibration analysis, as well as providing opportunities for researchers to conduct further investigations beyond the scope of this work.