

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DOES WEAK-TO-STRONG GENERALIZATION HAPPEN UNDER SPURIOUS CORRELATIONS?

Anonymous authors

Paper under double-blind review

ABSTRACT

We initiate a unified theoretical and algorithmic study of a key problem in weak-to-strong (W2S) generalization: when fine-tuning a strong pre-trained student with pseudolabels from a weaker teacher on a downstream task with spurious correlations, does W2S happen, and how to improve it upon failures? We consider two sources of spurious correlations caused by group imbalance: (i) a weak teacher fine-tuned on group-imbalanced labeled data with a minority group of fraction η_ℓ , and (ii) a group-imbalanced unlabeled set pseudolabeled by the teacher with a minority group of fraction η_u . Theoretically, a precise characterization of W2S gain at the proportional asymptotic limit shows that W2S always happens with sufficient pseudolabels when $\eta_u = \eta_\ell$ but may fail when $\eta_u \neq \eta_\ell$, where W2S gain diminishes as $(\eta_u - \eta_\ell)^2$ increases. Our theory is corroborated by extensive experiments on various spurious correlation benchmarks and teacher-student pairs. To boost W2S performance upon failures, we further propose a simple, effective algorithmic remedy that retrains the strong student on its high-confidence data subset after W2S fine-tuning. Our algorithm is group-label-free and achieves consistent, substantial improvements over vanilla W2S fine-tuning.

1 INTRODUCTION

Traditional learning paradigms like supervised learning and knowledge distillation (Hinton et al., 2015) learn from training data generated by strong teachers, *e.g.*, human experts. In contrast, contemporary foundation models encode encyclopedic knowledge through astronomical-scale pre-training, thereby achieving comparable or even superior performance to human experts in various domains via light post-training adaptation like fine-tuning (Brown et al., 2020; Achiam et al., 2023). This motivates the question on *superalignment* (Leike & Sutskever, 2023): can models with superhuman intelligence learn from weaker human supervision? *Weak-to-strong (W2S) generalization* (Burns et al., 2024) provides an encouraging answer for this question: a strong pre-trained student fine-tuned with pseudolabels generated by a weaker teacher can often outperform its teacher.

Since the first introduction of W2S by Burns et al. (2024), its mechanism has been extensively studied empirically (Guo et al., 2024; Liu & Alahi, 2024; Guo & Yang, 2024; Yang et al., 2024; 2025; Goel et al., 2025), and theoretically (Lang et al., 2024; Charikar et al., 2024; Wu & Sahai, 2025; Ildiz et al., 2025; Mulgund & Pabbaraju, 2025; Dong et al., 2025; Medvedev et al., 2025). While existing works on W2S generally assume access to clean downstream data, in practice, both the weak teacher and the unlabeled data for weak supervision often carry systematic biases, such as spurious correlations tied to demographic or acquisition factors (Arjovsky et al., 2019; Sagawa et al., 2020).

This challenge is especially relevant in the very settings that motivate W2S: a student broadly pre-trained on general data is fine-tuned for a specialized task where labeled samples are scarce and imperfect. In medicine, labels may be biased toward certain patient groups (Gupta et al., 2016) or imaging devices (Zech et al., 2018); in law, datasets may overrepresent particular jurisdictions or case types (Chalkidis et al., 2022); in autonomous driving, sensor data may be skewed toward specific weather or traffic conditions (Liu et al., 2024). For these specialized downstream tasks, one usually cannot interfere with the data acquisition process, nor obtain additional balanced data. It is therefore crucial to understand *whether W2S can remain effective under spurious correlations* caused by group imbalance—*when it succeeds, when it fails, and how its procedure can be improved*.

Our contributions. We initiate a systematic study of W2S under spurious correlations, providing (i) a theoretical analysis that answers the “when” question comprehensively by characterizing the

054 impact of spurious correlations on W2S precisely in the proportional asymptotic limit, as well as (ii)
 055 a simple, effective remedy for the failure of W2S under spurious correlations inspired by our theory,
 056 toward answering the “how” question. Concretely, our contributions are as follows.
 057

- 058 • **A theory of W2S under spurious correlations.** In Section 2, we conduct a systematic analysis in
 059 the ridgeless regression setting with zero approximation error, where W2S happens due to different
 060 estimation errors (*i.e.*, efficiency in utilizing data). At the proportional asymptotic limit, we provide
 061 *precise characterizations for the generalization errors of both teacher and student*. Consider
 062 using a weak teacher fine-tuned on labeled samples with a minority fraction η_e to pseudolabel N
 063 unlabeled samples with a minority fraction η_u for W2S fine-tuning. We show that (i) *W2S always*
 064 *happens with sufficiently large N when $\eta_e = \eta_u$ and improves when the teacher and student have*
 065 *distinct representations*; whereas (ii) *when $\eta_e \neq \eta_u$, W2S can fail even with $N \rightarrow \infty$, and W2S*
 066 *gain tends to diminish as $(\eta_u - \eta_e)^2$ increases*. Our theory is validated with extensive experiments
 067 on synthetic regression problems and real classification tasks (Section 3).
- 068 • **An algorithmic enhancement for W2S when $\eta_e \neq \eta_u$.** In Section 4, we propose a simple, effective
 069 algorithm that retrains the strong student on its high-confidence data subset after W2S fine-tuning
 070 via the generalized cross-entropy loss (Zhang & Sabuncu, 2018). Our method requires no group
 071 annotations and improves W2S when the gap between η_u and η_e is large. We conduct extensive
 072 experiments on assorted spurious correlation benchmarks (e.g., Waterbirds (Sagawa et al., 2020),
 073 BFFHQ (Lee et al., 2021), and ImageNet-9 (Xiao et al., 2020)), across 10 different teacher–student
 074 model pairs. The results show that our algorithm achieves consistent and substantial gains over
 075 vanilla W2S.

076 1.1 RELATED WORKS

077 **W2S generalization.** Empirically, many methods have been developed to validate/enhance W2S
 078 across various vision and natural language modeling tasks, including adjustable loss functions (Guo
 079 et al., 2024), multi-teacher algorithms (Liu & Alahi, 2024), data refinement strategies (Guo & Yang,
 080 2024; Yang et al., 2024), and the use of weak models for data filtering (Li et al., 2024). Theoretical
 081 work on W2S is also rapidly expanding, offering various mechanistic explanations from first principles,
 082 including the perspectives of neighborhood expansion (Lang et al., 2024), data overlap density (Shin
 083 et al., 2025), transfer learning (Somerstep et al., 2024), teacher-student disagreement (Charikar et al.,
 084 2024; Mulgund & Pabbaraju, 2025; Yao et al., 2025; Xu et al., 2025), benign overfitting (Wu &
 085 Sahai, 2025; Xue et al., 2025), knowledge distillation (Ildiz et al., 2025), low intrinsic dimension
 086 of fine-tuning (Aghajanyan et al., 2021; Dong et al., 2025), regularization (Medvedev et al., 2025;
 087 Moniri & Hassani, 2025), and feature learning with different inductive biases (Oh et al., 2025).

088 **Group robustness in knowledge distillation.** When transferring knowledge from a strong teacher
 089 to a weaker student, knowledge distillation (Hinton et al., 2015) has been shown to harm the minority
 090 group performance (Lukasik et al., 2021; Vilouras et al., 2023; Wang et al., 2023; Lee & Lee, 2023;
 091 Kenfack et al., 2024). To address this issue, different methods have been proposed, including adaptive
 092 mixing weights and per-class margins (Lukasik et al., 2021), distributionally robust optimization
 093 (Wang et al., 2023; Vilouras et al., 2023), last-layer transplantation (Lee & Lee, 2023), and gradient-
 094 based reweighting (Kenfack et al., 2024)¹. Our work differs from these approaches in three key
 095 aspects: (a) W2S generalization, where a weak teacher supervises a stronger student, is fundamentally
 096 distinct from classical knowledge distillation, (b) we explicitly consider the impact of mismatched
 097 minority group proportions between teacher and student, and (c) our method for improving W2S
 098 performance does not require any auxiliary information such as group annotations.

099 2 A THEORY OF W2S UNDER SPURIOUS CORRELATION

100 **Notations.** For any $p, q \in \mathbb{N}$, $p \geq q$, let $\text{Stiefel}(p, q) = \{\mathbf{Q} \in \mathbb{R}^{p \times q} \mid \mathbf{Q}^\top \mathbf{Q} = \mathbf{I}_q\}$ be the Stiefel
 101 manifold. $\mathbf{A} \otimes \mathbf{B} \in \mathbb{R}^{mp \times nq}$ denotes the Kronecker product of $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times q}$; when
 102 $n = q$, let $[\mathbf{A}; \mathbf{B}] \in \mathbb{R}^{(m+p) \times n}$ be the vertical stack; when $m = p$, let $[\mathbf{A}, \mathbf{B}] \in \mathbb{R}^{m \times (n+q)}$ be the
 103 horizontal stack. For any $\mathbf{w} \in \mathbb{R}^d$ and $i \in [d]$ or $\mathcal{I} \subseteq [d]$, let w_i and $[\mathbf{w}]_{\mathcal{I}}$ denote the i -th entry and
 104 the subvector of \mathbf{w} indexed by \mathcal{I} . For any $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $i \in [m]$, $j \in [n]$, let $A_{i,j}$ denote the
 105 (i, j) -th entry; $[\mathbf{A}]_{i,:} \in \mathbb{R}^n$ denotes the i -th row; $[\mathbf{A}]_{:,j} \in \mathbb{R}^m$ denotes the j -th column; and index
 106 subsets $\mathcal{I} \subseteq [m]$, $\mathcal{J} \subseteq [n]$ pick the corresponding submatrices.

107 ¹Group robustness under spurious correlations in supervised learning has been extensively studied and is out
 108 of scope of this work. We defer more discussions to Appendix B.

108
109

2.1 PROBLEM SETUP: REGRESSION UNDER SPURIOUS CORRELATION

110

Downstream task. Consider a downstream regression task characterized by a distribution $\mathcal{D}(\eta) : \mathcal{X} \times \mathcal{Y} \times \mathcal{G} \rightarrow [0, 1]$ where \mathcal{X} is the input space, $\mathcal{Y} = \mathbb{R}$ is the label space, and $\mathcal{G} = \{0, 1\}$ contains group labels (i.e., 1 for minority and 0 for majority). The fraction of the minority group in the population is controlled by $\eta \in [0, \frac{1}{2}]$ such that $\Pr[g = 1] = 1 - \Pr[g = 0] = \eta$.

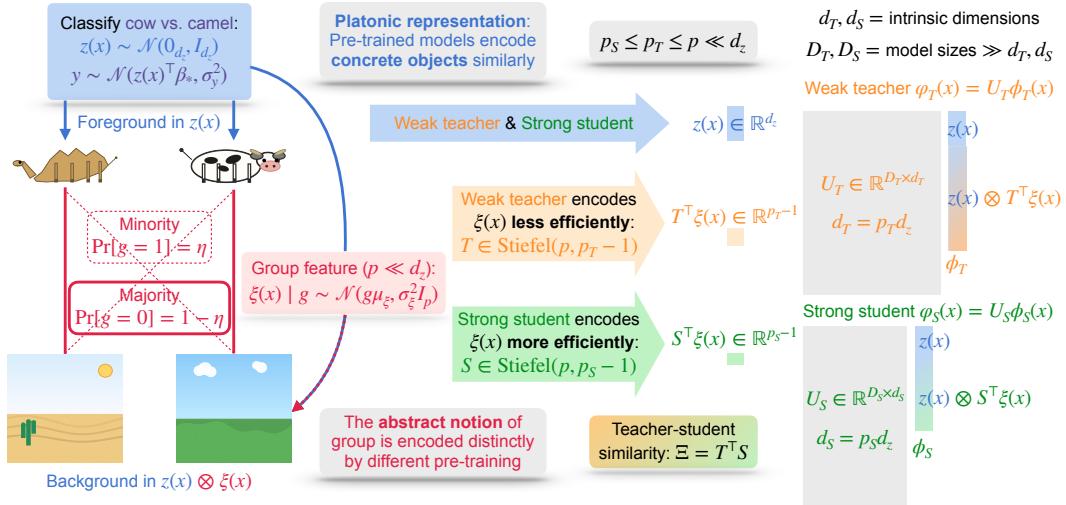
111

Definition 1 (Regression under spurious correlations). Let \mathcal{D}_x be the marginal distribution of $\mathbf{x} \in \mathcal{X}$; $\mathcal{D}_{\mathbf{x}|g}$ be the conditional distribution of \mathbf{x} given g ; and $\mathcal{D}_{y|\mathbf{x}}$ be the conditional distribution of y given \mathbf{x} satisfying $y = f_*(\mathbf{x}) + \epsilon$ for unknown $f_* : \mathcal{X} \rightarrow \mathbb{R}$ and i.i.d. label noise $\epsilon \sim \mathcal{N}(0, \sigma_y^2)$ independent of \mathbf{x} . Consider two feature maps: (i) the core feature $\mathbf{z} : \mathcal{X} \rightarrow \mathbb{R}^{d_z}$ determines the label y through $\mathbf{z}(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_{d_z}, \mathbf{I}_{d_z})$ and $f_*(\mathbf{x}) = \mathbf{z}(\mathbf{x})^\top \beta_*$ for fixed $\beta_* \in \mathbb{R}^{d_z}$; while (ii) the group feature $\xi : \mathcal{X} \rightarrow \mathbb{R}^p$ ($2 < p < \infty$) determines the group label g through $\xi(\mathbf{x}) \sim \mathcal{N}(g\mu_\xi, \sigma_\xi^2 \mathbf{I}_p)$ for fixed $\mu_\xi \in \mathbb{R}^p$ with dimension-independent $\|\mu_\xi\|_2, \sigma_\xi^2 \asymp 1$.

112

122 Here, $\mathbf{z}(\mathbf{x})$ encodes the core information for predicting y that is invariant across groups, typically rich
123 in semantics and therefore hard to learn (high-dimensional); while $\xi(\mathbf{x})$ is a latent feature controlling
124 which group \mathbf{x} belongs to, typically simpler to represent and therefore low-dimensional.

125



126

127 Figure 1: Visualization of the theoretical setup in Definitions 1 and 2 through Example 1.
128

129

Weak vs. strong models. We consider two pre-trained models that provide reasonably high-quality features for the downstream task: a weak teacher model $f_T : \mathcal{X} \rightarrow \mathbb{R}$ and a strong student model $f_S : \mathcal{X} \rightarrow \mathbb{R}$. Adapting the setting in Dong et al. (2025), we model fine-tuning in the kernel regime (Jacot et al., 2018; Malladi et al., 2023) with low intrinsic dimensions (Aghajanyan et al., 2021). In particular, we consider learning overparametrized linear layers $\theta_T \in \mathbb{R}^{D_T}$ and $\theta_S \in \mathbb{R}^{D_S}$ over high-dimensional pre-trained representations $\varphi_T : \mathcal{X} \rightarrow \mathbb{R}^{D_T}$ and $\varphi_S : \mathcal{X} \rightarrow \mathbb{R}^{D_S}$, respectively. When fine-tuning lies in the kernel regime, φ_T, φ_S correspond to the gradients of the tunable parameters in f_T, f_S at the pre-trained initialization, respectively, where D_T, D_S stand for the large tunable parameter counts. The difference between φ_T, φ_S that separates the weak and strong models on the downstream task with spurious correlations is pivotal in this setting:

130

Definition 2 (Weak vs. strong models). (i) *The weak teacher representation φ_T heavily entangles the core and group features: there exists $\mathbf{U}_T \in \text{Stiefel}(D_T, d_T)$ ($d_T \ll D_T$) such that $\varphi_T(\mathbf{x}) = \mathbf{U}_T \phi_T(\mathbf{x})$ and $\phi_T(\mathbf{x}) = \mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x})$, where $\mathbf{w}(\mathbf{x}) = [1; \mathbf{T}^\top \xi(\mathbf{x})] \in \mathbb{R}^{p_T}$ ($2 \leq p_T \leq p$) for a fixed $\mathbf{T} \in \text{Stiefel}(p, p_T-1)$ that projects $\xi(\mathbf{x})$ to a lower dimension (i.e., $\phi_T(\mathbf{x}) = [\mathbf{z}(\mathbf{x}); \mathbf{z}(\mathbf{x}) \otimes (\mathbf{T}^\top \xi(\mathbf{x}))] \in \mathbb{R}^{d_T}$). We note that $d_T = p_T d_z$. Let $\mu_T = \mathbf{T}^\top \mu_\xi \in \mathbb{R}^{p_T-1}$.*

131

(ii) *A strong student representation φ_S partially disentangles the core and group features: there exists $\mathbf{U}_S \in \text{Stiefel}(D_S, d_S)$ ($d_S \ll D_S$) such that $\varphi_S(\mathbf{x}) = \mathbf{U}_S \phi_S(\mathbf{x})$ and $\phi_S(\mathbf{x}) = \mathbf{z}(\mathbf{x}) \otimes \psi(\mathbf{x})$, where $\psi(\mathbf{x}) = [1; \mathbf{S}^\top \xi(\mathbf{x})] \in \mathbb{R}^{p_S}$ ($2 \leq p_S \leq p_T$) for a fixed $\mathbf{S} \in \text{Stiefel}(p, p_S-1)$ that projects*

162 $\xi(\mathbf{x})$ to a much lower dimension, $p_S \ll p$ (i.e., $\phi_S(\mathbf{x}) = [\mathbf{z}(\mathbf{x}); \mathbf{z}(\mathbf{x}) \otimes (\mathbf{S}^\top \xi(\mathbf{x}))] \in \mathbb{R}^{d_S}$). We
 163 note that $d_S = p_S d_z$. Let $\mu_S = \mathbf{S}^\top \mu_\xi \in \mathbb{R}^{p_S-1}$.²

164 Definition 2 formalizes the intuitions that compared to φ_T , the stronger φ_S (i) represents the information required for the downstream task more efficiently ($d_S \leq d_T$) and (ii) partially disentangles the core and group features, bringing robustness to spurious correlations. Notice that with $\mathbf{z}(\mathbf{x})$ prepending in both $\varphi_T(\mathbf{x})$ and $\varphi_S(\mathbf{x})$, the teacher and student both have zero approximation error (i.e., both pre-trained models are expressive enough for the downstream task), and W2S happens due to different estimation errors (i.e., the student is more sample efficient than its teacher).

170 **Example 1.** We take the well-known analogy of classifying cows (often in pastures) vs. camels (often
 171 in deserts) (Arjovsky et al., 2019) as an example (see Figure 1). With $\mathbf{z}(\mathbf{x})$ encoding the foreground
 172 of cows/camels, $\xi(\mathbf{x})$ represents whether the background is typical or not, while $\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x})$ and
 173 $\mathbf{z}(\mathbf{x}) \otimes \psi(\mathbf{x})$ correspond to the representations of background from the weak and strong models.

174 While the Platonic representation hypothesis (Huh et al., 2024) suggests that different pre-trained
 175 models tend to represent similar concrete objects similarly (with the same $\mathbf{z}(\mathbf{x})$), different model
 176 capacities can lead to distinct representations of a “typical” group in $\xi(\mathbf{x})$. For instance, a strong
 177 model that has learned the natural habitat of cows/camels during pre-training can encode typical
 178 samples as those with their respective backgrounds, leading to a simple, low-dimensional $\psi(\mathbf{x})$;
 179 whereas a weaker model without such knowledge have to rely on more complicated mechanisms to
 180 represent typical samples (e.g., counting), resulting in a more complex, higher-dimensional $\mathbf{w}(\mathbf{x})$.

181 Analogous to Dong et al. (2025), a key quantity that controls W2S gain is the similarity between the
 182 weak teacher and strong student representations, φ_T and φ_S , as formalized in Definition 3.

183 **Definition 3** (Teacher-student similarity). Under Definition 2, we define a similarity matrix $\Xi =$
 184 $\mathbf{T}^\top \mathbf{S} \in \mathbb{R}^{(p_T-1) \times (p_S-1)}$. Notice that $\|\Xi\|_F^2 \leq p_S - 1$ and $\|\Xi\|_2 \leq 1$.

185 Ξ measures the similarity of group features extracted by φ_T, φ_S , e.g., $\|\Xi\|_F^2 \rightarrow 0$ means $\mathbf{w}(\mathbf{x})$ and
 186 $\psi(\mathbf{x})$ are orthogonal, while $\|\Xi\|_F^2 \rightarrow p_S - 1$ means $\mathbf{w}(\mathbf{x})$ and $\psi(\mathbf{x})$ are highly aligned.

188 **W2S fine-tuning pipeline.** We consider two training sets with *i.i.d.* samples: (i) a small labeled
 189 training set $\tilde{\mathcal{S}} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i) \mid i \in [n]\} \sim \mathcal{D}(\eta_\ell)^n$ that is privately available only to the weak teacher, φ_T ,
 190 and (ii) a large unlabeled training set $\mathcal{S}_x = \{\mathbf{x}_i \mid i \in [N]\}$ from $\mathcal{S} = \{(\mathbf{x}_i, y_i) \mid i \in [N]\} \sim \mathcal{D}(\eta_u)^N$
 191 with hidden labels that is privately available only to the strong student, φ_S , where $\eta_\ell, \eta_u \in [0, \frac{1}{2}]$.
 192 The W2S fine-tuning pipeline consists of two stages: (i) Supervised fine-tuning (SFT) of $f_T(\cdot) =$
 193 $\varphi_T(\cdot)^\top \theta_T$ on $\tilde{\mathcal{S}}$ via ridgeless regression: assuming $n > d_T$,

$$\theta_T = \underset{\theta \in \mathbb{R}^{D_T}}{\operatorname{argmin}} \|\theta\|_2^2 \quad \text{s.t.} \quad \theta \in \underset{\theta' \in \mathbb{R}^{D_T}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n (\varphi_T(\tilde{\mathbf{x}}_i)^\top \theta' - \tilde{y}_i)^2, \quad (1)$$

197 (ii) W2S fine-tuning of $f_S(\cdot) = \varphi_S(\cdot)^\top \theta_S$ on \mathcal{S}_x labeled by f_T via ridgeless regression:

$$\theta_S = \underset{\theta \in \mathbb{R}^{D_S}}{\operatorname{argmin}} \|\theta\|_2^2 \quad \text{s.t.} \quad \theta \in \underset{\theta' \in \mathbb{R}^{D_S}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N (\varphi_S(\mathbf{x}_i)^\top \theta' - f_T(\mathbf{x}_i))^2, \quad (2)$$

202 Following Burns et al. (2024), in this W2S fine-tuning pipeline, we assume the weak teacher after SFT
 203 is fixed and not trainable, accessible in the W2S fine-tuning stage only through inference. Moreover,
 204 the labeled training set, $\tilde{\mathcal{S}}$, is only accessible in the first, SFT stage to the weak teacher, whereas the
 205 unlabeled set \mathcal{S}_x is only accessible in the second, W2S fine-tuning stage to the strong student.

206 **Remark 1** (Why ridgeless regression provides sufficient regularization?). We note that under Definition 2 where both $\varphi_T(\mathbf{x})$ and $\varphi_S(\mathbf{x})$ are constrained in low-dimensional subspaces, $\text{Range}(\mathbf{U}_T)$ and
 207 $\text{Range}(\mathbf{U}_S)$, ridgeless regression provides nearly optimal regularization to avoid overfitting (Wu &
 208 Xu, 2020; Hastie et al., 2022), which is essential for W2S generalization (Burns et al., 2024). When
 209 $\varphi_T(\mathbf{x})$ and $\varphi_S(\mathbf{x})$ are concentrated (in contrast to constrained) in low-dimensional subspaces with
 210 tails evenly distributed in the orthogonal complement, explicit regularization (Moniri & Hassani,
 211 2025; Dong et al., 2025) or early stopping (Burns et al., 2024; Medvedev et al., 2025) becomes
 212 necessary to prevent the student from overfitting to noisy teacher labels. Nevertheless, analogous to

213
 214 ²For both $\mathbf{w}(\mathbf{x})$ and $\psi(\mathbf{x})$, the first entry 1 effectively prepends the core feature $\mathbf{z}(\mathbf{x})$ in $\varphi_T(\mathbf{x})$ and $\varphi_S(\mathbf{x})$,
 215 which is essential to ensure that both teacher and student have negligible approximation error. Intuitively,
 pre-trained models have sufficient expressivity to learn the downstream task over population.

216 Dong et al. (2025), extending our ridgeless analysis to ridge regression does not alter our key insights
 217 on spurious correlations. Therefore, we focus on the ridgeless setting for clarity of exposition.
 218

219 The generalization performance is evaluated over a test distribution $\mathcal{D}(\eta_t)$ for some $\eta_t \in [0, 1]$: with
 220 the test risk $\mathcal{R}_{\eta_t}(f) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_{\mathbf{x}, y}(\eta_t)}[(f(\mathbf{x}) - y)^2]$, we consider the excess risk

$$221 \quad \mathbf{ER}_{\eta_t}(f) := \mathcal{R}_{\eta_t}(f) - \mathcal{R}_{\eta_t}(f_*) = \mathcal{R}_{\eta_t}(f) - \sigma_y^2. \quad (3)$$

223 In particular, $\eta_t = \frac{1}{2}$ corresponds to the average test risk; $\eta_t = 0$ corresponds to the majority test risk;
 224 and $\eta_t = 1$ corresponds to the minority test risk.

225 2.2 W2S GENERALIZATION UNDER SPURIOUS CORRELATION

227 With the problem setup, we are ready to present our main theorems regarding the effect of spurious
 228 correlations on W2S generalization. First, to characterize the excess risks of f_T and f_S (and thereby
 229 the W2S generalization gain) precisely, we push the problem to the proportional asymptotic limit:

230 **Assumption 1** (Proportional asymptotic limit). *We consider $d_z, n, N \rightarrow \infty$ with $d_z/n \rightarrow \gamma_z \in$
 231 $(0, p_T^{-1})$ (i.e., $n > d_T$), $d_z/N \rightarrow \nu_z \in (0, p_S^{-1})$ (i.e., $N > d_S$), whereas $2 \leq p_S \leq p_T \leq p$ are fixed.*

233 We highlight that in practice, the unlabeled samples are typically much more affordable than the
 234 labeled ones, leading to $\nu_z \ll \gamma_z$. Now, we characterize the excess risks of the weak teacher after
 235 SFT and the strong student after W2S fine-tuning, respectively, in Theorems 1 and 2.

236 **Theorem 1** (SFT of weak teacher (Appendix D.1)). *Under Assumption 1, (1) satisfies*

$$237 \quad \mathbb{E}_{\mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_T)] \xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \left(\underbrace{\frac{1}{p_T}}_{\mathcal{V}_T^{(0)} \text{ from label noise}} + \underbrace{\frac{\|(\eta_t - \eta_\ell) \boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2}}_{\mathcal{V}_T^{(1)} \text{ from spurious correlations}} \right).$$

238 **Theorem 2** (W2S, formally in Theorem 3). *Under Assumption 1, (2) satisfies*

$$239 \quad \mathbb{E}_{\mathcal{D}(\eta_u)^N, \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_S)] \xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \left(\underbrace{\frac{1}{p_{T \wedge S}}}_{\mathcal{V}_S^{(0)} \leq \mathcal{V}_T^{(0)}} + \underbrace{\frac{\|(\eta_u - \eta_\ell) \boldsymbol{\mu}_T + (\eta_t - \eta_u) \boldsymbol{\Xi} \boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2}}_{\mathcal{V}_S^{(1)} \leq \mathcal{V}_T^{(1)} \text{ when } \eta_u = \eta_\ell} + \underbrace{\Theta(\nu_z)}_{\mathcal{E}_S \ll 1} \right),$$

240 where $p_{T \wedge S} = 1 + \|\boldsymbol{\Xi}\|_F^2 \in [1, p_S]$ is the effective group feature dimension learned by the strong
 241 student from the weak teacher controlled by the similarity between φ_T and φ_S in encoding group
 242 features (see Definition 3)—less similar teacher-student pairs enjoy lower $p_{T \wedge S}$; $\mathcal{V}_S^{(0)}$ and $\mathcal{V}_T^{(0)}$ are
 243 generalization errors of f_S and f_T from noisy labels; $\mathcal{V}_S^{(1)}$ and $\mathcal{V}_T^{(1)}$ are generalization errors of f_S
 244 and f_T induced by spurious correlations, $\eta_u, \eta_\ell \neq \eta_t$; and the higher-order term \mathcal{E}_S , formalized in
 245 Theorem 3, becomes negligible when $\nu_z \ll 1$.

246 It is worth noting that the proportional asymptotic limit (Assumption 1) assumed in Theorems 1
 247 and 2 can be relaxed to incorporate finite-sample cases via standard edge fluctuation analysis (see
 248 e.g., Hastie et al. (2022); Cheng & Montanari (2024)). We omit such extensions here since they do
 249 not bring additional insights to Theorems 1 and 2.

250 As a special case, without spurious correlations ($\eta_\ell = \eta_u = \eta_t$ or $\boldsymbol{\mu}_\xi = \mathbf{0}_p$), Theorems 1 and 2 exactly
 251 recover the results in Dong et al. (2025) at the proportional asymptotic limit: $\mathbb{E}[\mathbf{ER}_{\eta_t}(f_T)] \rightarrow \sigma_y^2 \gamma_z p_T$
 252 and $\mathbb{E}[\mathbf{ER}_{\eta_t}(f_S)] \rightarrow \sigma_y^2 \gamma_z (p_{T \wedge S} + \Theta(\nu_z))$, where with a small $\nu_z \ll 1$, the W2S gain is larger when
 253 the teacher and student representations are less aligned (i.e., lower $p_{T \wedge S}$). Meanwhile, Theorems 1
 254 and 2 together reveal insights regarding the effect of spurious correlations on the W2S gain,

$$255 \quad \Delta \mathcal{R}_{\eta_t} := \mathbb{E}_{\mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_T)] - \mathbb{E}_{\mathcal{D}(\eta_u)^N, \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_S)], \quad (4)$$

256 as discussed in Remark 2, where W2S generalization happens whenever $\Delta \mathcal{R}_{\eta_t} > 0$.

257 **Remark 2** (Does W2S happen under spurious correlations?). *Theorems 1 and 2 provide a mixed
 258 answer to this question conditioned on various factors, including the teacher-student similarity, the
 259 separation between groups, and the choice of η_u for given η_ℓ^3 , as summarized below:*

260 ³In practice, η_ℓ is typically fixed and known (e.g., given a weak teacher fine-tuned on the Waterbirds training
 261 set), while η_u can be controlled by the practitioner when collecting unlabeled data for W2S fine-tuning.

(a) **W2S happens whenever $\eta_u = \eta_\ell$ and ν_z is small**, e.g., when $\|\Xi\|_F^2 \approx 0$ and $\nu_z \ll 1$, $\Delta\mathcal{R}_{\eta_t} > 0$ is optimized at $\eta_u \approx \eta_\ell$ (Figure 2). We highlight that when $\eta_u = \eta_\ell$, in addition to the W2S gain from variance reduction $\mathcal{V}_T^{(0)} - \mathcal{V}_S^{(0)} = p_T - p_{T \wedge S} \geq 0$, **the strong student improves upon its teacher in handling spurious correlations**: $\mathcal{V}_T^{(1)} - \mathcal{V}_S^{(1)} = ((\eta_t - \eta_\ell)^2 / \sigma_\xi^2) (\|\mu_T\|_2^2 - \|\Xi\mu_S\|_2^2) \geq 0$, where the gain increases as the teacher-student similarity decreases.

(b) For fixed Ξ , assume $\mu_T \neq \Xi\mu_S$, when $\nu_z \ll 1$, the optimal η_u that maximizes W2S gain is $\eta_u^* = \frac{\eta_\ell \|\mu_T\|_2^2 - (\eta_t + \eta_\ell) \mu_T^\top \Xi \mu_S + \eta_t \|\Xi \mu_S\|_2^2}{\|\mu_T - \Xi \mu_S\|_2^2}$, e.g., when $\eta_\ell = \frac{1}{2}$, $\eta_u^* = \frac{1}{2}$; when $\|\Xi \mu_S\|_2 \ll \|\mu_T\|_2$, $\eta_u^* \approx \eta_\ell$; with $\|\Xi \mu_S\|_2 \neq 0$, η_u^* tends to increase with $\|\mu_S\|_2^2$ and deviate from η_ℓ when $\|\mu_S\|_2^2 \approx \|\mu_T\|_2^2$ (Figure 3 left).

(c) W2S gain increases as the teacher-student similarity $\|\Xi\|_F^2$ decreases (Figure 3 right).

(d) **W2S may not happen if $\eta_u \neq \eta_\ell$, even when $\nu_z \ll 1$ and $\|\Xi\|_F^2 = 0$** , e.g., when $\eta_\ell = 0.4$ but $\eta_u = 0.1$, with $\|\Xi\|_F^2 = 0$, W2S does not happen if the majority and minority groups are well separated: $\Delta\mathcal{R}_{1/2} < 0$ for any ν_z if $\|\mu_T\|_2^2 / \sigma_\xi^2 > 12.5(p_T - 1)$. More generally, for $\|\Xi\|_F^2 = 0$, $\mathcal{V}_S^{(1)}$ increases proportionally to $(\eta_u - \eta_\ell)^2$, and thus $\Delta\mathcal{R}_{\eta_t}$ diminishes as the gap increases.

In Appendix F, we further discuss implications of Theorems 1 and 2 on the fairness of W2S.

2.3 SYNTHETIC EXPERIMENTS

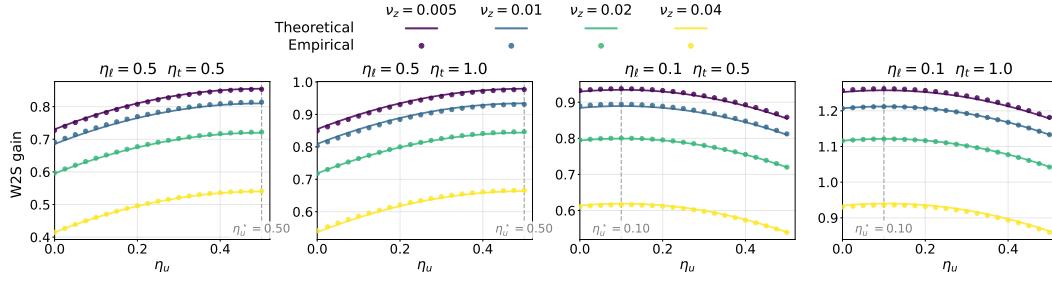


Figure 2: W2S gains across different combinations of η_ℓ and η_t . Each panel shows theoretical (solid lines) and empirical (circles) results for W2S gain as a function of η_u , across different ν_z values. Here we fix μ_T , μ_S , Ξ , and d_z with $\|\mu_T\|_2^2 = 10.0$, $\|\mu_S\|_2^2 = 0.1$, $\|\Xi\|_F^2 = 0.1p_S$. Vertical dashed lines indicate the theoretical optimal η_u^* values that maximize W2S gain.

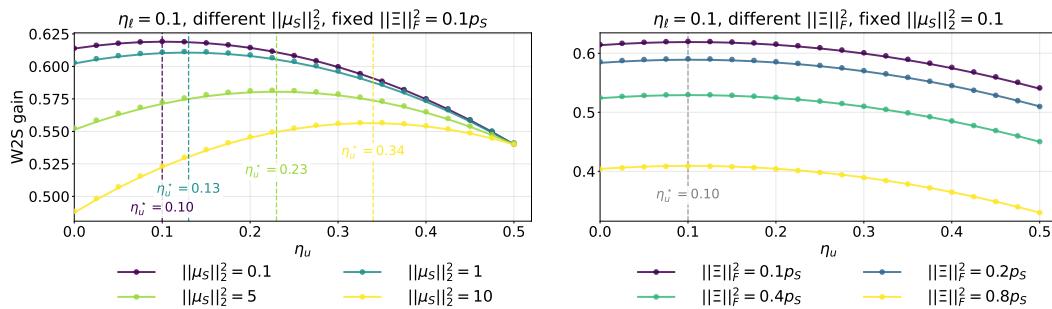


Figure 3: Impact of μ_S and Ξ on W2S gain. Both panels show theoretical (solid lines) and empirical (circles) results for W2S gain as a function of η_u . Fixed parameters: $\eta_\ell = 0.1$, $\eta_t = 0.5$, $\nu_z = 0.04$, $\|\mu_T\|_2^2 = 10.0$. Left: varying $\|\mu_S\|_2^2$ with fixed $\|\Xi\|_F^2 = 0.1p_S$. Right: varying $\|\Xi\|_F^2$ with fixed $\|\mu_S\|_2^2 = 0.1$. Dashed lines indicate the theoretical optimal η_u^* values that maximize W2S gain.

Figures 2 and 3 validate the theory in Section 2.2 through synthetic Gaussian experiments, with fixed $d_z = 2048$ in all experiments. We begin by examining how varying η_u affects W2S gains under different values of η_ℓ . As shown in Figure 2, when $\|\Xi\|_F^2$ is small (a distinct teacher-student pair), W2S gains are maximized at $\eta_u \approx \eta_\ell$ for both balanced ($\eta_\ell = 0.5$) and highly spurious ($\eta_\ell = 0.1$) unlabeled data. This holds for both the average test risk and the minority test risk, consistent with Remark 2(a). Moreover, the magnitude of the W2S gain decreases as ν_z increases, reflecting the

324 role of \mathcal{E}_S in Theorem 2. Figure 3 left shows that as $\|\mu_S\|_2^2$ increases so that $\|\Xi\mu_S\|_2^2$ becomes non-
 325 negligible compared to $\|\mu_T\|_2^2$, the optimal value η_u^* gradually shifts away from η_ℓ . This indicates
 326 that in some special cases η_u^* may not lie near η_ℓ , consistent with Remark 2(b). Figure 3 right
 327 illustrates that the W2S gain decreases as the teacher-student similarity $\|\Xi\|_F^2$ increases, consistent
 328 with Remark 2(c).

330 3 REAL-WORLD EVALUATION

331 Now we extend our theoretical understanding of W2S under spurious correlation to real-world tasks.
 332 We first leverage the theoretical framework to interpret our findings on how spurious correlations
 333 affect W2S performance across real-world benchmarks.

335 3.1 MODEL AND DATASET SETUP

336 **Pre-trained models.** Our weak teachers and strong students are drawn from a diverse set of
 337 pre-trained vision backbones that differ in architecture and training paradigm. Specifically, we
 338 consider ResNet-18 (ResNet18) (He et al., 2016), CLIP ViT-B/32 (Clipb32) (Radford et al., 2021),
 339 ConvNeXt-L (ConvNeXt) (Liu et al., 2022), DINOv2 ViT-L/14 (DINOv2) (Oquab et al., 2023), and
 340 MAE ViT-B/16 (MAE) (He et al., 2022). For each experiment on a given dataset, we include all
 341 teacher–student pairs whose relative strength (measured by accuracy) remains stable when we vary
 342 parameters including η_ℓ , η_u , N , or n . We freeze all backbone parameters, view the pre-trained feature
 343 for the teacher and the student as φ_T and φ_S , and only finetune the classification head.

344 **Datasets.** From both theoretical and practical perspectives, effective W2S requires that the pre-
 345 trained weak teacher and strong student have learned feature representations that are useful for the
 346 downstream task. Therefore, we evaluate W2S performance on widely used spurious correlation
 347 benchmarks that are relatively close to the pre-training distribution. These include Waterbirds (Sagawa
 348 et al., 2020), BFFHQ (Lee et al., 2021), and ImageNet-9 (Xiao et al., 2020), which contain spurious
 349 correlations between background and bird labels, age and gender labels, and background and object
 350 labels, respectively. We further provide a self-generated dataset, BG-COCO, by creating spurious
 351 correlations between cats/dogs from COCO (Lin et al., 2014) and indoor/outdoor scenes from
 352 Places (Zhou et al., 2017). We note that in all four datasets, the spurious correlation arises from highly
 353 imbalanced group proportions between the majority and minority groups. We denote the minority
 354 group proportion in the original training set of each dataset as η_o , which equals 0.05, 0.005, 0, and
 355 0.05 for Waterbirds, BFFHQ, ImageNet-9, and BG-COCO, respectively. Detailed configurations of
 356 the dataset splits are provided in Appendix E.1.

357 3.2 INTERPRETING W2S UNDER SPURIOUS CORRELATIONS

358 We investigate how the proportion of the minority group in the unlabeled data affects W2S per-
 359 formance when the labeled data are either group-balanced or group-imbalanced. Specifically, we
 360 fix $\eta_\ell = 0.5$ and $\eta_\ell = \eta_o$, respectively, and vary η_u while recording the change in W2S gain.⁴
 361 Figure 4 presents the average W2S gain across all teacher–student pairs on all four datasets. More
 362 comprehensive results are provided in Appendix E.2.

363 Our results show that, for both average accuracy ($\eta_t = 0.5$) and worst group accuracy ($\eta_t = 1$),
 364 increasing the minority proportion in the unlabeled data improves W2S performance when the weak
 365 teacher is free of spurious correlation ($\eta_\ell = 0.5$). Moreover, when the weak teacher itself encodes
 366 spurious correlation ($\eta_\ell = \eta_o$), the W2S gain is consistently positive across all four datasets at
 367 $\eta_u = \eta_o$, but surprisingly decreases for more balanced data as η_u increases from η_o to 0.5. Overall,
 368 W2S gain is negatively affected as the gap between η_u and η_ℓ increases. These observations echo
 369 our theory and synthetic experiments (see Theorems 1 and 2, Remark 2, and Figure 2), showing
 370 that our theoretical findings on regression extends naturally to broader, real-world classification
 371 problems.

372 4 ENHANCED-W2S METHOD

373 Inspired by the theory and observations in Sections 2 and 3, we introduce a simple retraining
 374 method based on student confidence and generalized cross-entropy to strengthen W2S under spurious

375 ⁴For classification tasks, W2S gain refers to the improvement in test accuracy achieved by the strong student
 376 after W2S fine-tuning over its weak teacher.

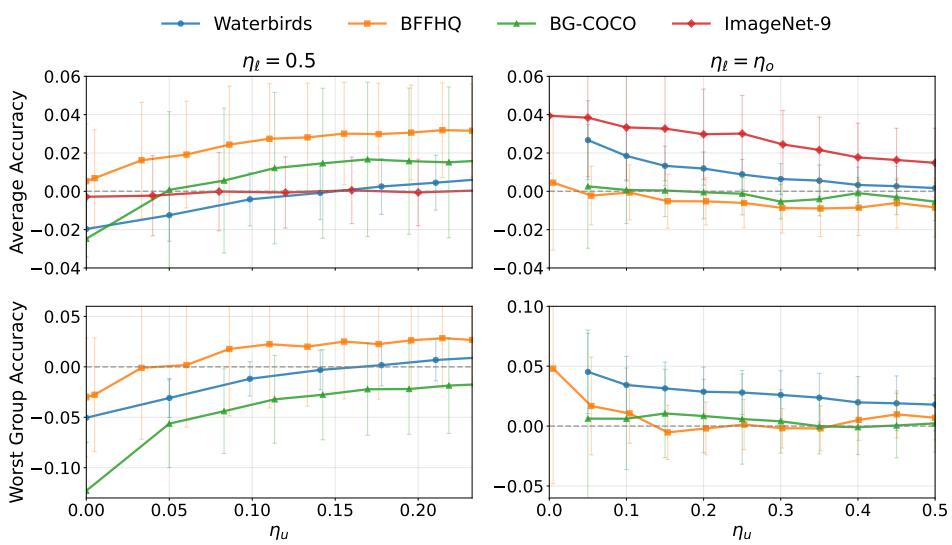


Figure 4: Average W2S gain across all teacher-student pairs as a function of η_u on all four datasets. Top row: average accuracy; bottom row: worst group accuracy. Left column fixes $\eta_\ell = 0.5$; right column fixes $\eta_\ell = \eta_o$. For $\eta_\ell = 0.5$, curves are plotted over a shared η_u interval aligned across datasets (bounded by minority group sample availability) to enable direct comparability. For $\eta_\ell = \eta_o$, each dataset is plotted from its own η_o (0.05, 0.005, 0.05, and 0 for Waterbirds, BFFHQ, BG-COCO, and ImageNet-9, respectively) up to 0.5. ImageNet-9 does not have a clearly defined worst group and is therefore omitted from the bottom panels.

correlations. We show that this approach remarkably outperforms vanilla W2S across multiple datasets and large pre-trained backbones, without requiring any group annotations.

Method. Both our theoretical results and empirical findings indicate that W2S gain is noticeably reduced when there is a large discrepancy between the minority proportions of the unlabeled data and the labeled data. Therefore, we propose a simple method that requires no group label annotations and is capable of improving W2S gain in two particularly important settings: one where the labeled data are heavily affected by spurious correlation while the unlabeled data are free of it ($\eta_\ell = \eta_o, \eta_u = 0.5$), and the other where the unlabeled data suffer from spurious correlation while the labeled data are balanced ($\eta_\ell = 0.5, \eta_u = \eta_o$).

Formally, let the unlabeled data be $\hat{\mathcal{S}} = \{(\mathbf{x}_i, \hat{y}_i) \mid i \in [N]\}$, where \hat{y}_i is the pseudolabel given by the weak teacher on $\mathbf{x}_i \in \mathcal{S}_x$. Our method enhances W2S gain by retraining the strong student after W2S fine-tuning, based on two components: (i) selecting a fraction $p \in (0, 1]$ of $\hat{\mathcal{S}}$ consisting of the samples with the highest student confidence (equivalently, the lowest entropy), and (ii) applying the generalized cross-entropy (GCE) loss (Zhang & Sabuncu, 2018) with parameter $q \in (0, 1]$ to each $(\mathbf{x}_i, \hat{y}_i)$ in this subset:

$$\mathcal{L}_{\text{GCE}}(\mathbf{x}_i, \hat{y}_i; q) = \frac{1 - \mathbf{p}_{\hat{y}_i}(\mathbf{x}_i)^q}{q},$$

where $\mathbf{p}_{\hat{y}_i}(\mathbf{x}_i)$ is the softmax probability that the student assigns to the pseudolabel \hat{y}_i for \mathbf{x}_i .

In both settings ($\eta_\ell = \eta_o, \eta_u = 0.5$ and $\eta_\ell = 0.5, \eta_u = \eta_o$), selecting a high-confidence subset of the student’s predictions filters for samples where all relevant features are clearly expressed and effectively used during prediction, thus preventing the strong student from over-relying on any single (potentially spurious) feature. Moreover, unlike the CE loss which imposes overly strong penalties on high-confidence but incorrect pseudolabels, applying the GCE loss to the selected subset mitigates the negative impact of pseudolabel noise from the weak teacher.⁵ More importantly, for the case $\eta_\ell = \eta_o, \eta_u = 0.5$, confidence-based selection further provides a crucial benefit. As shown in

⁵In our Enhanced-W2S method, the role of GCE loss is analogous to its original use in Zhang & Sabuncu (2018) for handling noisy labels. Different from the setting studied in Nam et al. (2020), where GCE loss on ground-truth labeled datasets with spurious correlations was observed to amplify bias, in our method GCE loss

432
 433
 434
 435
 436
 437
 438
 439
 440
 441
 442
 443
 444
 445
 446
 447
 448
 449
 450
 451
 452
 453
 454
 455
 456
 457
 458
 459
 460
 461
 462
 463
 464
 465
 466
 467
 468
 469
 470
 471
 472
 473
 474
 475
 476
 477
 478
 479
 480
 481
 482
 483
 484
 485
 Appendix E.3, the high-confidence subset tends to filter out a larger fraction of minority samples to effectively reduce the new η_u during retraining. This observation aligns with our theoretical prediction that when $\eta_\ell = \eta_o$, decreasing η_u from 0.5 leads to improved W2S gain.

Dataset	η_ℓ	η_u	Teacher–Student pair									
			DINOv2 ConvNeXt	DINOv2 Clipb32	DINOv2 ResNet18	DINOv2 MAE	ConvNeXt Clipb32	ConvNeXt ResNet18	ConvNeXt MAE	Clipb32 ResNet18	Clipb32 MAE	ResNet18 MAE
Waterbirds	0.5	η_o	6.60	11.29	7.34	16.68	3.79	2.05	6.28	—	2.07	0.77
	η_o	0.5	7.19	13.86	11.73	11.62	2.85	2.02	4.33	—	1.32	14.54
BFFHQ	0.5	η_o	6.85	2.75	8.42	4.93	4.05	—	—	6.54	5.12	—
	η_o	0.5	3.92	8.53	2.02	4.56	2.09	—	—	2.06	-1.37	—
BG-COCO	0.5	η_o	5.38	13.40	12.88	24.01	9.82	6.49	15.25	3.39	12.43	2.05
	η_o	0.5	10.21	16.99	12.25	-3.52	3.41	1.21	-3.07	3.48	0.31	3.70
ImageNet-9	0.5	η_o	—	6.03	7.45	24.11	4.74	5.30	18.49	4.22	21.73	17.98
	η_o	0.5	—	8.21	11.28	22.00	3.77	1.81	10.50	4.51	23.24	15.76

Table 1: Relative improvement of Enhanced-W2S over vanilla W2S (%), measured by average accuracy) across all datasets and teacher–student pairs. Each entry reports the mean improvement over all N, n combinations. For each model pair in the table header, the assignment of weak teacher and strong student depends on the dataset. We report for each dataset only those model pairs whose relative strength relationship remains consistent across different (η_ℓ, η_u) settings within that dataset.

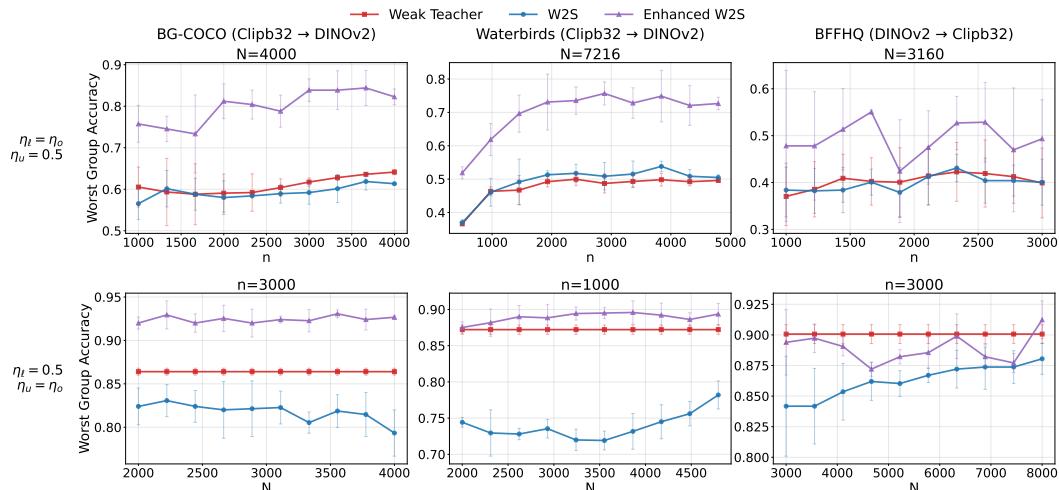


Figure 5: Comparison of Enhanced-W2S and original W2S for the (Clipb32, DINOv2) pair on BG-COCO, Waterbirds, and BFFHQ. Top row: worst group accuracy with $\eta_\ell = \eta_o$, $\eta_u = 0.5$ (fixed N , varying n). Bottom row: worst group accuracy with $\eta_\ell = 0.5$, $\eta_u = \eta_o$ (fixed n , varying N).

Main results. We evaluate our Enhanced-W2S method across all four datasets. Table 1 reports the relative improvement of Enhanced-W2S over vanilla W2S for each teacher–student pair. Figure 5 further visualizes the performance of Enhanced-W2S versus vanilla W2S for a representative model pair. Overall, for both average accuracy and worst group accuracy, Enhanced-W2S achieves consistent and substantial improvements over vanilla W2S under both (η_ℓ, η_u) settings. Specifically, Table 1 shows that 67 out of 70 model pairs exhibit a positive gain (measured by average accuracy), with the mean relative improvements across all pairs reaching 7.02% (Waterbirds), 4.32% (BFFHQ), 7.50% (BG-COCO), and 11.73% (ImageNet-9). In addition, the relative improvement of Enhanced-W2S in terms of worst group accuracy across all pairs is 21.14% (Waterbirds), 3.81% (BFFHQ), and 7.73% (BG-COCO). Further details are provided in Appendix E.3.

is applied to a pseudolabeled dataset restricted to a high-confidence subset, and thus serves a fundamentally different role.

486 487 488 489	Comparison	Metric ($\times 10^{-2}$)	Waterbirds		BFFHQ		BG-COCO		ImageNet-9	
			(η_ℓ, η_u)		(η_ℓ, η_u)		(η_ℓ, η_u)		(η_ℓ, η_u)	
			$(0.5, \eta_o)$	$(\eta_o, 0.5)$						
490	Enhanced W2S	Average	3.27	1.00	3.51	0.46	3.41	0.52	1.00	0.51
491	– Enhanced W2S ($q \rightarrow 0$)	Worst	5.15	2.39	4.34	1.90	3.79	1.46	—	—
492	Enhanced W2S	Average	2.79	4.84	2.77	2.69	4.97	3.37	4.51	5.08
493	– Enhanced W2S ($p = 1$)	Worst	3.57	8.58	2.75	3.73	6.26	5.44	—	—

Table 2: Ablation across four datasets: improvements of Enhanced-W2S over two baselines, using only the CE loss (i.e., the $q \rightarrow 0$ limit of GCE) and using all unlabeled data ($p = 1$), in terms of either average accuracy or worst group accuracy. For each dataset, improvements are computed as the mean across all model pairs whose relative strength relationship remains consistent under different (η_ℓ, η_u) settings. ImageNet-9 has no well-defined worst group, so those entries are omitted.

Ablation study. We conduct controlled ablations to examine the contribution of the two key components of our Enhanced-W2S methods, namely the use of the GCE loss and confidence-based selection. Specifically, we conduct two separate ablations: (i) replacing the GCE loss with the standard CE loss, and (ii) replacing confidence-based selection with using the full unlabeled dataset. We then retrain the model under each variant and compare the results with the original Enhanced-W2S method. Table 2 shows that under both (η_ℓ, η_u) settings, the GCE loss and confidence-based selection each play a positive role in improving W2S gain. When $\eta_\ell = \eta_o, \eta_u = 0.5$, the impact of using CE loss is consistently smaller than that of using the full unlabeled dataset; whereas under $\eta_\ell = 0.5, \eta_u = \eta_o$, the effects of the two ablations are more comparable. This suggests that filtering out minority group samples via high-confidence selection plays a more critical role in the former setting, while in the latter setting the contributions of GCE and high-confidence selection are comparable.

5 CONCLUSIONS, DISCUSSIONS, AND FUTURE DIRECTIONS

In this work, we start with a theoretical framework that models W2S generalization under spurious correlations induced by group imbalance. Within this framework, we precisely characterize how different factors, such as the proportions of minority groups in labeled and unlabeled data and the teacher-student similarity, affect W2S, which is validated through extensive synthetic experiments and on diverse real-world tasks. Inspired by our analysis, we proposed Enhanced-W2S, a confidence-based retraining algorithm that does not require any group labels and substantially improves W2S gains when the labeled or unlabeled data are highly group-imbalanced. The effectiveness of this approach is demonstrated across assorted real-world datasets.

It is important to emphasize that spurious correlations in W2S constitute a critical issue that deserves closer attention. Beyond standard benchmarks, such correlations can pose substantial risks: in socially sensitive domains they may reflect demographic biases, and in safety-critical applications they can degrade reliability under rare but high-stakes conditions. While our experiments focus on public computer vision benchmarks, the mechanisms we analyze are broadly relevant. Our algorithm provides the first attempt to improve W2S in this setting, and we hope this work will inspire further efforts toward more reliable and efficient W2S methods.

Meanwhile, from a more technical perspective, another exciting future direction is to investigate W2S generation (with or without spurious correlation) beyond the kernel regime by taking the training dynamics of the teacher and student models, conditioned on their pre-trained initializations, into consideration.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Armen Aghajanyan, Sonal Gupta, and Luke Zettlemoyer. Intrinsic dimensionality explains the effectiveness of language model fine-tuning. In *Proceedings of the 59th Annual Meeting of the*

540 *Association for Computational Linguistics and the 11th International Joint Conference on Natural*
 541 *Language Processing (Volume 1: Long Papers)*, pp. 7319–7328, 2021.

542

543 Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization.
 544 *arXiv preprint arXiv:1907.02893*, 2019.

545

546 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,
 547 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are
 548 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

549

550 Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner,
 551 Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, et al. Weak-to-strong generalization:
 552 eliciting strong capabilities with weak supervision. In *Proceedings of the 41st International*
 553 *Conference on Machine Learning*, pp. 4971–5012, 2024.

554

555 Ilias Chalkidis, Tommaso Pasini, Sheng Zhang, Letizia Tomada, Sebastian Felix Schwemer, and
 556 Anders Søgaard. Fairlex: A multilingual benchmark for evaluating fairness in legal text processing.
 557 *arXiv preprint arXiv:2203.07228*, 2022.

558

559 Moses Charikar, Chirag Pabbaraju, and Kirankumar Shiragur. Quantifying the gain in weak-to-strong
 560 generalization. *Advances in neural information processing systems*, 37:126474–126499, 2024.

561

562 Chen Cheng and Andrea Montanari. Dimension free ridge regression. *The Annals of Statistics*, 52(6):
 563 2879–2912, 2024.

564

565 Yihe Deng, Yu Yang, Baharan Mirzasoleiman, and Quanquan Gu. Robust learning with progressive
 566 data expansion against spurious correlation. *Advances in neural information processing systems*,
 567 36:1390–1402, 2023.

568

569 Yijun Dong, Kevin Miller, Qi Lei, and Rachel Ward. Cluster-aware semi-supervised learning:
 570 relational knowledge distillation provably learns clustering. *Advances in Neural Information*
 571 *Processing Systems*, 36, 2024.

572

573 Yijun Dong, Yicheng Li, Yunai Li, Jason D Lee, and Qi Lei. Discrepancies are virtue: Weak-to-strong
 574 generalization through lens of intrinsic dimension. In *Forty-second International Conference on*
 575 *Machine Learning*. PMLR, 2025.

576

577 Shashwat Goel, Joschka Strüber, Ilze Amanda Auzina, Karuna K Chandra, Ponnurangam Kumaraguru,
 578 Douwe Kiela, Ameya Prabhu, Matthias Bethge, and Jonas Geiping. Great models think alike and
 579 this undermines ai oversight. In *Forty-second International Conference on Machine Learning*,
 580 2025.

581

582 Jianyuan Guo, Hanting Chen, Chengcheng Wang, Kai Han, Chang Xu, and Yunhe Wang. Vision
 583 superalignment: Weak-to-strong generalization for vision foundation models. *arXiv preprint*
 584 *arXiv:2402.03749*, 2024.

585

586 Yue Guo and Yi Yang. Improving weak-to-strong generalization with reliability-aware alignment.
 587 *arXiv preprint arXiv:2406.19032*, 2024.

588

589 Alpana K Gupta, Mausumi Bharadwaj, and Ravi Mehrotra. Skin cancer concerns in people of color:
 590 risk factors and prevention. *Asian Pacific journal of cancer prevention: APJCP*, 17(12):5257,
 591 2016.

592

593 Yujin Han and Difan Zou. Improving group robustness on spurious correlation requires preciser
 594 group inference. *arXiv preprint arXiv:2404.13815*, 2024.

595

596 Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J Tibshirani. Surprises in high-
 597 dimensional ridgeless least squares interpolation. *Annals of statistics*, 50(2):949, 2022.

598

599 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
 600 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
 601 pp. 770–778, 2016.

594 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked
 595 autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer*
 596 *vision and pattern recognition*, pp. 16000–16009, 2022.

597

598 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv*
 599 *preprint arXiv:1503.02531*, 2015.

600 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation
 601 hypothesis. *arXiv preprint arXiv:2405.07987*, 2024.

602

603 Muhammed Emrullah Ildiz, Halil Alperen Gozeten, Ege Onur Taga, Marco Mondelli, and Samet
 604 Oymak. High-dimensional analysis of knowledge distillation: Weak-to-strong generalization and
 605 scaling laws. In *ICLR*, 2025.

606

607 Pavel Izmailov, Polina Kirichenko, Nate Gruver, and Andrew G Wilson. On feature learning in
 608 the presence of spurious correlations. *Advances in Neural Information Processing Systems*, 35:
 609 38516–38532, 2022.

610

611 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and
 612 generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

613

614 Patrik Kenfack, Ulrich Aïvodji, and Samira Ebrahimi Kahou. Adaptive group robust ensemble
 615 knowledge distillation. *arXiv preprint arXiv:2411.14984*, 2024.

616

617 David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai
 618 Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation
 619 (rex). In *International conference on machine learning*, pp. 5815–5826. PMLR, 2021.

620

621 Tyler LaBonte, Vidya Muthukumar, and Abhishek Kumar. Towards last-layer retraining for group
 622 robustness with fewer annotations. *Advances in Neural Information Processing Systems*, 36, 2024.

623

624 Hunter Lang, David Sontag, and Aravindan Vijayaraghavan. Theoretical analysis of weak-to-strong
 625 generalization. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*,
 626 2024.

627

628 Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection.
 629 *Annals of statistics*, pp. 1302–1338, 2000.

630

631 Jiwoon Lee and Jaeho Lee. Debiased distillation by transplanting the last layer. *arXiv preprint*
 632 *arXiv:2302.11187*, 2023.

633

634 Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased
 635 representation via disentangled feature augmentation. *Advances in Neural Information Processing*
 636 *Systems*, 34:25123–25133, 2021.

637

638 Jan Leike and Ilya Sutskever. Introducing superalignment, July 2023. Accessed: 2025-09-24.

639

640 Ming Li, Yong Zhang, Shuai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi
 641 Zhou. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. In *Proceedings*
 642 *of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*
 643 *Papers)*, pp. 14255–14273, 2024.

644

645 Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr
 646 Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European*
 647 *conference on computer vision*, pp. 740–755. Springer, 2014.

648

649 Chenruo Liu, Hongjun Liu, Zeyu Lai, Yiqiu Shen, Chen Zhao, and Qi Lei. Superclass-guided repre-
 650 sentation disentanglement for spurious correlation mitigation. *arXiv preprint arXiv:2508.08570*,
 651 2025a.

652

653 Chenruo Liu, Kenan Tang, Yao Qin, and Qi Lei. Bridging distribution shift and ai safety: Conceptual
 654 and methodological synergies. *arXiv preprint arXiv:2505.22829*, 2025b.

648 Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa,
 649 Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training
 650 group information. In *International Conference on Machine Learning*, pp. 6781–6792. PMLR,
 651 2021.

652 Lydia T Liu, Max Simchowitz, and Moritz Hardt. The implicit fairness criterion of unconstrained
 653 learning. In *International Conference on Machine Learning*, pp. 4051–4060. PMLR, 2019.

654 Mingyu Liu, Ekim Yurtsever, Jonathan Fossaert, Xingcheng Zhou, Walter Zimmer, Yuning Cui,
 655 Bare Luka Zagar, and Alois C Knoll. A survey on autonomous driving datasets: Statistics,
 656 annotation quality, and a future outlook. *IEEE Transactions on Intelligent Vehicles*, 2024.

657 Yuejiang Liu and Alexandre Alahi. Co-supervised learning: Improving weak-to-strong generalization
 658 with hierarchical mixture of experts. *arXiv preprint arXiv:2402.15505*, 2024.

659 Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie.
 660 A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and*
 661 *pattern recognition*, pp. 11976–11986, 2022.

662 Michal Lukasik, Srinadh Bhojanapalli, Aditya Krishna Menon, and Sanjiv Kumar. Teacher’s pet:
 663 understanding and mitigating biases in distillation. *arXiv preprint arXiv:2106.10494*, 2021.

664 Fan Ma, Deyu Meng, Xuanyi Dong, and Yi Yang. Self-paced multi-view co-training. *Journal of*
 665 *Machine Learning Research*, 21(57):1–38, 2020.

666 Sadhika Malladi, Alexander Wettig, Dingli Yu, Danqi Chen, and Sanjeev Arora. A kernel-based
 667 view of language model fine-tuning. In *International Conference on Machine Learning*, pp.
 668 23610–23641. PMLR, 2023.

669 Marko Medvedev, Kaifeng Lyu, Dingli Yu, Sanjeev Arora, Zhiyuan Li, and Nathan Srebro. Weak-to-
 670 strong generalization even in random feature networks, provably. In *Forty-second International*
 671 *Conference on Machine Learning*. PMLR, 2025.

672 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
 673 on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6):1–35, 2021.

674 Behrad Moniri and Hamed Hassani. On the mechanisms of weak-to-strong generalization: A
 675 theoretical perspective. *arXiv preprint arXiv:2505.18346*, 2025.

676 Abhijeet Mulgund and Chirag Pabbaraju. Relating misfit to gain in weak-to-strong generalization
 677 beyond the squared loss. In *Forty-second International Conference on Machine Learning*, 2025.

678 Vaishnavh Nagarajan, Aditya K Menon, Srinadh Bhojanapalli, Hossein Mobahi, and Sanjiv Ku-
 679 mar. On student-teacher deviations in distillation: does it pay to disobey? *Advances in Neural*
 680 *Information Processing Systems*, 36:5961–6000, 2023.

681 Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure:
 682 De-biasing classifier from biased classifier. *Advances in Neural Information Processing Systems*,
 683 33:20673–20684, 2020.

684 Junsoo Oh, Jerry Song, and Chulhee Yun. From linear to nonlinear: Provable weak-to-strong
 685 generalization through feature learning. In *High-dimensional Learning Dynamics 2025*, 2025.

686 Utkarsh Ojha, Yuheng Li, Anirudh Sundara Rajan, Yingyu Liang, and Yong Jae Lee. What knowledge
 687 gets distilled in knowledge distillation? *Advances in Neural Information Processing Systems*, 36:
 688 11037–11048, 2023.

689 Luca Oneto and Silvia Chiappa. Fairness in machine learning. In *Recent trends in learning from*
 690 *data: Tutorials from the inns big data and deep learning conference (innsbddl2019)*, pp. 155–196.
 691 Springer, 2020.

692 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov,
 693 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning
 694 robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.

702 Hoang Phan, Andrew Gordon Wilson, and Qi Lei. Controllable prompt tuning for balancing group
 703 distributional robustness. *arXiv preprint arXiv:2403.02695*, 2024.

704

705 Mary Phuong and Christoph Lampert. Towards understanding knowledge distillation. In *International
 706 conference on machine learning*, pp. 5142–5151. PMLR, 2019.

707

708 Aahlad Puli, Lily H Zhang, Eric K Oermann, and Rajesh Ranganath. Out-of-distribution generaliza-
 709 tion in the presence of nuisance-induced spurious correlations. *arXiv preprint arXiv:2107.00520*,
 710 2021.

711

712 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 713 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 714 models from natural language supervision. In *International conference on machine learning*, pp.
 8748–8763. PmLR, 2021.

715

716 Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust
 717 neural networks. In *International Conference on Learning Representations*, 2020.

718

719 Jessica Schrouff, Alexis Bellot, Amal Rannen-Triki, Alan Malek, Isabela Albuquerque, Arthur
 720 Gretton, Alexander D’Amour, and Silvia Chiappa. Mind the graph when balancing data for
 721 fairness or robustness. *Advances in Neural Information Processing Systems*, 37:29913–29947,
 722 2024.

723

724 Changho Shin, John Cooper, and Frederic Sala. Weak-to-strong generalization through the data-
 centric lens. In *The Thirteenth International Conference on Learning Representations*, 2025.

725

726 Seamus Somerstep, Felipe Maia Polo, Moulinath Banerjee, Yaakov Ritov, Mikhail Yurochkin, and
 727 Yuekai Sun. A statistical framework for weak-to-strong generalization. In *ICML 2024 Next
 Generation of AI Safety Workshop*, 2024.

728

729 Samuel Stanton, Pavel Izmailov, Polina Kirichenko, Alexander A Alemi, and Andrew G Wilson.
 730 Does knowledge distillation really work? *Advances in Neural Information Processing Systems*, 34:
 6906–6919, 2021.

731

732 Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational
 733 mathematics*, 12(4):389–434, 2012.

734

735 Victor Veitch, Alexander D’Amour, Steve Yadlowsky, and Jacob Eisenstein. Counterfactual invariance
 736 to spurious correlations: why and how to pass stress tests. In *Proceedings of the 35th International
 737 Conference on Neural Information Processing Systems*, pp. 16196–16208, 2021.

738

739 Konstantinos Vilouras, Xiao Liu, Pedro Sanchez, Alison Q O’Neil, and Sotirios A Tsaftaris. Group
 740 distributionally robust knowledge distillation. In *International Workshop on Machine Learning in
 Medical Imaging*, pp. 234–242. Springer, 2023.

741

742 Serena Wang, Harikrishna Narasimhan, Yichen Zhou, Sara Hooker, Michal Lukasik, and Aditya Kr-
 743 ishna Menon. Robust distillation for worst-class performance: on the interplay between teacher
 744 and student objectives. In *Uncertainty in Artificial Intelligence*, pp. 2237–2247. PMLR, 2023.

745

746 Tao Wen, Zihan Wang, Quan Zhang, and Qi Lei. Elastic representation: Mitigating spurious
 747 correlations for group robustness. *arXiv preprint arXiv:2502.09850*, 2025.

748

749 Robert Williamson and Aditya Menon. Fairness risk measures. In *International conference on
 750 machine learning*, pp. 6786–6797. PMLR, 2019.

751

752 David Xing Wu and Anant Sahai. Provable weak-to-strong generalization via benign overfitting. In
 753 *The Thirteenth International Conference on Learning Representations*, 2025.

754

755 Denny Wu and Ji Xu. On the optimal weighted ℓ_2 regularization in overparameterized linear
 756 regression. *Advances in Neural Information Processing Systems*, 33:10112–10123, 2020.

Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image
 757 backgrounds in object recognition. *ArXiv preprint arXiv:2006.09994*, 2020.

756 Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student
 757 improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision*
 758 and pattern recognition, pp. 10687–10698, 2020.

759

760 Gengze Xu, Wei Yao, Ziqiao Wang, and Yong Liu. On the emergence of weak-to-strong generalization:
 761 A bias-variance perspective. *arXiv preprint arXiv:2505.24313*, 2025.

762 Yihao Xue, Jiping Li, and Baharan Mirzasoleiman. Representations shape weak-to-strong generaliza-
 763 tion: Theoretical insights and empirical predictions. In *Forty-second International Conference on*
 764 *Machine Learning*, 2025.

765

766 Wenkai Yang, Shiqi Shen, Guangyao Shen, Wei Yao, Yong Liu, Gong Zhi, Yankai Lin, and Ji-
 767 Rong Wen. Super (ficial)-alignment: Strong models may deceive weak models in weak-to-strong
 768 generalization. In *The Thirteenth International Conference on Learning Representations*, 2025.

769

770 Yuqing Yang, Yan Ma, and Pengfei Liu. Weak-to-strong reasoning. In *Findings of the Association*
 771 *for Computational Linguistics: EMNLP 2024*, pp. 8350–8367, 2024.

772

773 Wei Yao, Wenkai Yang, Ziqiao Wang, Yankai Lin, and Yong Liu. Understanding the capabilities
 774 and limitations of weak-to-strong generalization. In *Scaling Self-Improving Foundation Models*
 775 *without Human Supervision*, 2025.

776

777 Sriram Yenamandra, Pratik Ramesh, Viraj Prabhu, and Judy Hoffman. Facts: First amplify correla-
 778 tions and then slice to discover bias. In *Proceedings of the IEEE/CVF International Conference on*
 779 *Computer Vision*, pp. 4794–4804, 2023.

780

781 Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao, and Chao Zhang. Fine-tuning pre-
 782 trained language model with weak supervision: A contrastive-regularized self-training approach.
 783 In *Proceedings of the 2021 Conference of the North American Chapter of the Association for*
 784 *Computational Linguistics: Human Language Technologies*, pp. 1063–1077, 2021.

785

786 John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl
 787 Oermann. Variable generalization performance of a deep learning model to detect pneumonia in
 788 chest radiographs: a cross-sectional study. *PLoS medicine*, 15(11):e1002683, 2018.

789

790 Jingzhao Zhang, Aditya Menon, Andreas Veit, Srinadh Bhojanapalli, Sanjiv Kumar, and Suvrit Sra.
 791 Coping with label shift via distributionally robust optimisation. *arXiv preprint arXiv:2010.12230*,
 792 2020.

793

794 Michael Zhang, Nimit S Sohoni, Hongyang R Zhang, Chelsea Finn, and Christopher Ré. Correct-n-
 795 contrast: A contrastive approach for improving robustness to spurious correlations. *arXiv preprint*
 796 *arXiv:2203.01517*, 2022.

797

798 Min-Ling Zhang and Zhi-Hua Zhou. Cotrade: Confident co-training with data editing. *IEEE*
 799 *Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 41(6):1612–1626, 2011.

800

801 Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks
 802 with noisy labels. *Advances in neural information processing systems*, 31, 2018.

803

804 Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10
 805 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine*
 806 *Intelligence*, 2017.

807

808

809

810 Appendices

811	A Usage of Large Language Models	16
812		
813		
814	B Additional Related Works	16
815		
816	C Additional Notations	17
817		
818	D Proofs for Section 2.2	17
819		
820	D.1 SFT of Weak Teacher	17
821	D.2 W2S Fine-tuning of Strong Student	21
822		
823	E Additional Experimental Details	26
824		
825	E.1 Dataset Statistics	26
826	E.2 Results for Interpreting W2S under Spurious Correlations	27
827	E.3 Results for Enhanced W2S	28
828		
829	F Group Fairness in W2S Generalization	30
830		
831		
832		
833		
834		
835	A USAGE OF LARGE LANGUAGE MODELS	
836	Large language models were used in a limited manner to (i) search for related literature, (ii) check	
837	grammar/phrasing, and (iii) make stylistic adjustments.	
838		
839	B ADDITIONAL RELATED WORKS	
840		
841	Knowledge distillation. Knowledge distillation (KD) (Hinton et al., 2015) is closely related to W2S	
842	but with the roles reversed: KD transfers knowledge from a larger teacher model to a smaller student	
843	model. A series of works has analyzed when and why a distilled student generalizes (Phuong &	
844	Lampert, 2019; Stanton et al., 2021; Ojha et al., 2023; Nagarajan et al., 2023; Dong et al., 2024; Ildiz	
845	et al., 2025). Analytically, W2S departs from traditional KD because "weak" vs. "strong" is defined	
846	relative to pretraining, so W2S is naturally studied as fine-tuning on pseudolabeled data.	
847	Group robustness to spurious correlation. Extensive efforts have been devoted to mitigating	
848	spurious correlation for robust and safe generalization to unseen test domains (Arjovsky et al., 2019;	
849	Sagawa et al., 2020; Krueger et al., 2021; Deng et al., 2023; Phan et al., 2024; Wen et al., 2025;	
850	Liu et al., 2025b). Among these studies, a subset of work specifically targets spurious correlation	
851	arising from group imbalance. When group labels are available, canonical approaches include	
852	reweighting minority groups (Sagawa et al., 2020), downsampling majority groups (Deng et al.,	
853	2023), distributionally robust optimization (Sagawa et al., 2020; Zhang et al., 2020), and progressive	
854	data expansion (Deng et al., 2023). Since obtaining group annotations in training data can be	
855	costly or even infeasible, alternative strategies aim to identify biased samples without explicit group	
856	supervision (Nam et al., 2020; Liu et al., 2021; Zhang et al., 2022; Yenamandra et al., 2023; Han &	
857	Zou, 2024), or leverage auxiliary signals such as knowledge of spurious attributes (Puli et al., 2021),	
858	class annotations (LaBonte et al., 2024), and superclass-level information (Liu et al., 2025a).	
859	Multi-round retraining and confidence-based selection. Multi-round retraining and confidence-	
860	based data selection are widely adopted ideas that have been leveraged, independently, to improve	
861	model performance for W2S generalization (Burns et al., 2024; Liu & Alahi, 2024), mitigate spurious	
862	correlation (Liu et al., 2021; Nam et al., 2020), and in the broader literature on self-training (Xie et al.,	
863	2020; Yu et al., 2021) and co-training (Zhang & Zhou, 2011; Ma et al., 2020). Our theory provides a	
864	principled motivation to combine these practical techniques, bringing an effective algorithmic remedy	
865	for the failures of W2S under spurious correlation.	

864 **C ADDITIONAL NOTATIONS**

865
 866 For any $n \in \mathbb{N}$, let $[n] = \{1, 2, \dots, n\}$. \mathbf{e}_i is the i -th canonical basis of a conformable dimension.
 867 We adapt the standard big-O notations for functions of multiple variables: for two functions $f, g : \mathbb{N}^k \rightarrow \mathbb{R}_{\geq 0}$, $f(\mathbf{n}) = O(g(\mathbf{n}))$ means that there exists a constant $C > 0$ such that $f(\mathbf{n}) \leq Cg(\mathbf{n})$ for
 868 all $\mathbf{n} \in \mathbb{N}^k$; $f(\mathbf{n}) = \Omega(g(\mathbf{n}))$ means that there exists a constant $c > 0$ such that $f(\mathbf{n}) \geq cg(\mathbf{n})$ for
 869 all $\mathbf{n} \in \mathbb{N}^k$; $f(\mathbf{n}) = \Theta(g(\mathbf{n}))$ means that $f(\mathbf{n}) = O(g(\mathbf{n}))$ and $f(\mathbf{n}) = \Omega(g(\mathbf{n}))$; $f(\mathbf{n}) = o(g(\mathbf{n}))$
 870 means that $\liminf_{n_i \rightarrow \infty} f(\mathbf{n})/g(\mathbf{n}) = 0$. For a scalar quantity $f(n) \geq 0$ depending on $n \in \mathbb{N}$,
 871 $f(n) = O_{\mathbb{P}}(g(n))$ means that for any $\delta \in (0, 1)$, there exists a constant $C(\delta) > 0$, independent of n ,
 872 and a natural number $N(\delta) \in \mathbb{N}$ such that $\Pr[f(n) \leq C(\delta)g(n)] \geq 1 - \delta$ for all $n \geq N(\delta)$.
 873

874 **D PROOFS FOR SECTION 2.2**

875 We first observe that when $n > d_T$, (1) can be equivalently written as $\boldsymbol{\theta}_T = \mathbf{U}_T \boldsymbol{\beta}_T$ where

876

$$\boldsymbol{\beta}_T = \underset{\boldsymbol{\beta} \in \mathbb{R}^{d_T}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n (\phi_T(\tilde{\mathbf{x}}_i)^\top \boldsymbol{\beta} - \tilde{y}_i)^2. \quad (5)$$

877 Analogously, (2) can be equivalently written as $\boldsymbol{\theta}_S = \mathbf{U}_S \boldsymbol{\beta}_S$ where

878

$$\boldsymbol{\beta}_S = \underset{\boldsymbol{\beta} \in \mathbb{R}^{d_S}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N (\phi_S(\mathbf{x}_i)^\top \boldsymbol{\beta} - f_T(\mathbf{x}_i))^2. \quad (6)$$

880 **D.1 SFT OF WEAK TEACHER**

881 We start by considering the population-optimal linear predictor over the weak teacher representation,
 882 $\phi_T(\cdot)$: as $n \rightarrow \infty$, (5) converges to

883

$$\boldsymbol{\beta}_T^\infty = \underset{\boldsymbol{\beta} \in \mathbb{R}^{d_T}}{\operatorname{argmin}} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_{\mathbf{x}, y}(\eta_\ell)} [(\phi_T(\mathbf{x})^\top \boldsymbol{\beta} - y)^2]. \quad (7)$$

884 **Lemma 1** (Population SFT of weak teacher). *When supervisedly fine-tuned over the population, the*
 885 *weak teacher from (7) satisfies $f_T^\infty(\mathbf{x}) = \phi_T(\mathbf{x})^\top \boldsymbol{\beta}_T^\infty = \mathbf{z}(\mathbf{x})^\top \boldsymbol{\beta}_* = f_*$.*

886 *Proof of Lemma 1.* Notice that (7) admits a closed-form solution

887

$$\boldsymbol{\beta}_T^\infty = \boldsymbol{\Sigma}_{\phi_T, \eta_\ell}^{-1} \boldsymbol{\mu}_{\phi_T, \eta_\ell}, \quad \boldsymbol{\Sigma}_{\phi_T, \eta_\ell} = \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\phi_T(\mathbf{x}) \phi_T(\mathbf{x})^\top], \quad \boldsymbol{\mu}_{\phi_T, \eta_\ell} = \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\phi_T(\mathbf{x}) y].$$

888 Since $\mathbf{z}(\mathbf{x})$ and $\mathbf{w}(\mathbf{x})$ are independent, we have

889

$$\begin{aligned} \boldsymbol{\Sigma}_{\phi_T, \eta_\ell} &= \mathbb{E}_{\mathcal{D}(\eta_\ell)} [(\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x})) (\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x}))^\top] \\ &= \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\mathbf{z}(\mathbf{x}) \mathbf{z}(\mathbf{x})^\top] \otimes \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\mathbf{w}(\mathbf{x}) \mathbf{w}(\mathbf{x})^\top] \\ &= \mathbf{I}_{d_z} \otimes \mathbf{C}_T(\eta_\ell), \end{aligned} \quad (8)$$

890 where

891

$$\mathbf{C}_T(\eta_\ell) = \begin{bmatrix} 1 & \eta_\ell \boldsymbol{\mu}_T^\top \\ \eta_\ell \boldsymbol{\mu}_T & \sigma_\xi^2 \mathbf{I}_{p_T-1} + \eta_\ell^2 \boldsymbol{\mu}_T \boldsymbol{\mu}_T^\top \end{bmatrix} = \operatorname{diag} (0, \sigma_\xi^2 \mathbf{I}_{p_T-1}) + \begin{bmatrix} 1 \\ \eta_\ell \boldsymbol{\mu}_T \end{bmatrix} [1 \quad \eta_\ell \boldsymbol{\mu}_T], \quad (9)$$

892 whose inverse can be computed via block matrix inversion as

893

$$\mathbf{C}_T(\eta_\ell)^{-1} = \begin{bmatrix} 1 + \sigma_\xi^{-2} \|\eta_\ell \boldsymbol{\mu}_T\|_2^2 & -\sigma_\xi^{-2} \eta_\ell \boldsymbol{\mu}_T^\top \\ -\sigma_\xi^{-2} \eta_\ell \boldsymbol{\mu}_T & \sigma_\xi^{-2} \mathbf{I}_{p_T-1} \end{bmatrix}. \quad (10)$$

894 Meanwhile, by the independence of $\mathbf{z}(\mathbf{x})$ and $\mathbf{w}(\mathbf{x})$, we have

895

$$\begin{aligned} \boldsymbol{\mu}_{\phi_T, \eta_\ell} &= \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x}) (\mathbf{z}(\mathbf{x})^\top \boldsymbol{\beta}_* + \epsilon)] \\ &= (\mathbb{E}_{\mathcal{D}(\eta_\ell)} [\mathbf{z}(\mathbf{x}) \mathbf{z}(\mathbf{x})^\top] \boldsymbol{\beta}_*) \otimes \mathbb{E}_{\mathcal{D}(\eta_\ell)} [\mathbf{w}(\mathbf{x})] \\ &= \boldsymbol{\beta}_* \otimes \begin{bmatrix} 1 \\ \eta_\ell \boldsymbol{\mu}_T \end{bmatrix}. \end{aligned}$$

918 Therefore, the population-optimal linear predictor over ϕ_T is given by
919

$$\begin{aligned} 920 \beta_T^\infty &= \Sigma_{\phi_T, \eta_\ell}^{-1} \mu_{\phi_T, \eta_\ell} = (\mathbf{I}_{d_z} \otimes \mathbf{C}_T(\eta_\ell))^{-1} \left(\beta_* \otimes \begin{bmatrix} 1 \\ \eta_\ell \mu_T \end{bmatrix} \right) = \beta_* \otimes \left(\mathbf{C}_T(\eta_\ell)^{-1} \begin{bmatrix} 1 \\ \eta_\ell \mu_T \end{bmatrix} \right) \\ 921 &= \beta_* \otimes \left(\begin{bmatrix} 1 + \sigma_\xi^{-2} \|\eta_\ell \mu_T\|_2^2 & -\sigma_\xi^{-2} \eta_\ell \mu_T^\top \\ -\sigma_\xi^{-2} \eta_\ell \mu_T & \sigma_\xi^{-2} \mathbf{I}_{p_T-1} \end{bmatrix} \begin{bmatrix} 1 \\ \eta_\ell \mu_T \end{bmatrix} \right) = \beta_* \otimes \mathbf{e}_1, \\ 922 \\ 923 \\ 924 \end{aligned}$$

925 where $\mathbf{e}_1 \in \mathbb{R}^{p_T}$ is the first canonical basis. \square
926

927 While the $f_T^\infty = f_*$ achieves the optimal population risk $\mathcal{R}(f_T^\infty) = \mathcal{R}(f_*) = \sigma_y^2$, the inefficient repre-
928 sentation ($d_T = d_z p_T \gg d_z$) and the entangled features of ϕ_T make the finite-sample generalization
929 challenging, especially under spurious correlations, as we will show next.
930

Theorem 1 (SFT of weak teacher (Appendix D.1)). *Under Assumption 1, (1) satisfies*

$$\begin{aligned} 931 \mathbb{E}_{\mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_T)] &\xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \left(\underbrace{\mathbf{P}_T}_{\mathcal{V}_T^{(0)} \text{ from label noise}} + \underbrace{\frac{\|(\eta_t - \eta_\ell) \mu_T\|_2^2}{\sigma_\xi^2}}_{\mathcal{V}_T^{(1)} \text{ from spurious correlations}} \right). \\ 932 \\ 933 \\ 934 \end{aligned}$$

935 *Proof of Theorem 1.* For a small labeled set $\tilde{\mathcal{S}} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i) \mid i \in [n]\} \sim \mathcal{D}(\eta_\ell)^n$, the SFT in (1)
936 admits a closed-form solution
937

$$\beta_T = (\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \tilde{\Phi}_T^\top \tilde{\mathbf{y}}, \quad (11)$$

938 where $\tilde{\Phi}_T = [\phi_T(\tilde{\mathbf{x}}_1), \dots, \phi_T(\tilde{\mathbf{x}}_n)]^\top \in \mathbb{R}^{n \times d_T}$ and $\tilde{\mathbf{y}} = [\tilde{y}_1, \dots, \tilde{y}_n]^\top \in \mathbb{R}^n$. Since Lemma 1
939 shows that the population-optimal linear predictor over ϕ_T is $f_T^\infty(\mathbf{x}) = \phi_T(\mathbf{x})^\top \beta_T^\infty = f_*(\mathbf{x})$, we
940 have $\tilde{\mathbf{y}} = \tilde{\Phi}_T \beta_T^\infty + \tilde{\epsilon}$ where $\tilde{\epsilon} \sim \mathcal{N}(\mathbf{0}_n, \sigma_y^2 \mathbf{I}_n)$. Therefore, we observe that
941

$$\beta_T - \beta_T^\infty = (\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \tilde{\Phi}_T^\top \tilde{\epsilon}.$$

942 Since the excess risk over the test distribution $\mathcal{D}(\eta_t)$ is given by
943

$$\begin{aligned} 944 \mathbf{ER}_{\eta_t}(f_T) &= \mathbb{E}_{\mathcal{D}(\eta_t)}[(f_T(\mathbf{x}) - f_*(\mathbf{x}))^2] = \mathbb{E}_{\mathcal{D}(\eta_t)}[(\phi_T(\mathbf{x})^\top (\beta_T - \beta_T^\infty))^2] \\ 945 &= \|\beta_T - \beta_T^\infty\|_{\Sigma_{\phi_T, \eta_t}}^2, \end{aligned}$$

946 where $\Sigma_{\phi_T, \eta_t} = \mathbb{E}_{\mathcal{D}(\eta_t)}[\phi_T(\mathbf{x}) \phi_T(\mathbf{x})^\top]$, let $\tilde{\Sigma}_n = \frac{1}{n} \tilde{\Phi}_T^\top \tilde{\Phi}_T \in \mathbb{R}^{d_T \times d_T}$ be the sample covariance
947 matrix of $\phi_T(\tilde{\mathbf{x}})$ over $\tilde{\mathbf{x}} \sim \mathcal{D}_x(\eta_\ell)$, we have
948

$$\begin{aligned} 949 \mathbb{E}_{\tilde{\mathcal{S}} \sim \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_T)] &= \text{tr} \left(\mathbb{E}_{\tilde{\mathcal{S}} \sim \mathcal{D}(\eta_\ell)^n} \left[\Sigma_{\phi_T, \eta_t} (\beta_T - \beta_T^\infty) (\beta_T - \beta_T^\infty)^\top \right] \right) \\ 950 &= \sigma_y^2 \text{tr} \left(\Sigma_{\phi_T, \eta_t} \mathbb{E}_{\tilde{\mathcal{S}} \sim \mathcal{D}(\eta_\ell)^n} \left[(\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \right] \right) \\ 951 &= \frac{\sigma_y^2}{n} \text{tr} \left(\Sigma_{\phi_T, \eta_t} \tilde{\Sigma}_n^{-1} \right). \end{aligned} \quad (12)$$

952 Recall $\phi_T(\mathbf{x}) = \mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x})$ and notice that for any $\eta \in [0, 1]$,
953

$$\mathbb{E}_{\mathcal{D}(\eta)}[\phi_T(\mathbf{x})] = \mathbb{E}_{\mathcal{D}(\eta)}[\mathbf{z}(\mathbf{x})] \otimes \mathbb{E}_{\mathcal{D}(\eta)}[\mathbf{w}(\mathbf{x})] = \mathbf{0}_{d_T},$$

954 while the derivation of (8) suggests that for any $\eta \in [0, 1]$,
955

$$\Sigma_{\phi_T, \eta} = \mathbb{E}_{\mathcal{D}(\eta)}[\phi_T(\mathbf{x}) \phi_T(\mathbf{x})^\top] = \mathbf{I}_{d_z} \otimes \mathbf{C}_T(\eta). \quad (13)$$

956 However, we notice that $\phi_T(\mathbf{x})$ is not multivariate Gaussian due to the non-Gaussianity of products
957 of independent Gaussian variables and the dependence of entries in $\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x})$. Therefore, $\tilde{\Sigma}_n$
958 cannot be directly computed using inverse Wishart. Instead, we leverage the concentration of $\tilde{\Sigma}_n$ in
959 the proportional asymptotic limit (Assumption 1). In particular, Lemma 2 and (13) implies that as
960 $d_z, n \rightarrow \infty$ with $d_z/n \rightarrow \gamma_z \in (0, p_T^{-1})$,

$$\frac{\sigma_y^2}{n} \text{tr} \left(\Sigma_{\phi_T, \eta_t} \tilde{\Sigma}_n^{-1} \right) \xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \text{tr} \left(\mathbf{C}_T(\eta_t) \mathbf{C}_T(\eta_t)^{-1} \right). \quad (14)$$

972 Leveraging the derivation of (9) and (10), we observe that
 973

$$\begin{aligned} \text{tr}(\mathbf{C}_T(\eta_t)\mathbf{C}_T(\eta_\ell)^{-1}) &= p_T + \sigma_\xi^{-2}\eta_\ell^2\|\boldsymbol{\mu}_T\|_2^2 - 2\sigma_\xi^{-2}\eta_t\eta_\ell\|\boldsymbol{\mu}_T\|_2^2 + \sigma_\xi^{-2}\eta_t^2\|\boldsymbol{\mu}_T\|_2^2 \\ &= p_T + \sigma_\xi^{-2}(\eta_t - \eta_\ell)^2\|\boldsymbol{\mu}_T\|_2^2 = p_T + (\eta_t - \eta_\ell)^2\frac{\|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2}, \end{aligned} \quad (15)$$

978 and therefore, (14) becomes
 979

$$\begin{aligned} \frac{\sigma_y^2}{n} \text{tr}(\boldsymbol{\Sigma}_{\phi_T, \eta_t} \tilde{\boldsymbol{\Sigma}}_n^{-1}) &\xrightarrow{\mathbb{P}} \sigma_y^2 \frac{d_z}{n} \left(p_T + (\eta_t - \eta_\ell)^2 \frac{\|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2} \right) \\ &= \sigma_y^2 \gamma_z \left(p_T + (\eta_t - \eta_\ell)^2 \frac{\|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2} \right). \end{aligned}$$

985 Plugging the above into (12) completes the proof. \square
 986

987 **Lemma 2.** For fixed $p_T \geq 2$, let $\mathbf{C} \in \mathbb{R}^{p_T \times p_T}$ be any fixed symmetric matrix with $\|\mathbf{C}\|_2 < \infty$.
 988 Recall $\tilde{\boldsymbol{\Sigma}}_n = \frac{1}{n} \sum_{i=1}^n \phi_T(\tilde{\mathbf{x}}_i) \phi_T(\tilde{\mathbf{x}}_i)^\top$ where $\tilde{\mathbf{x}}_i \sim \mathcal{D}_{\mathbf{x}}(\eta_\ell)$ i.i.d. for all $i \in [n]$. As $d_z, n \rightarrow \infty$ with
 989 $d_z/n \rightarrow \gamma_z \in (0, p_T^{-1})$,

$$\frac{1}{n} \text{tr}((\mathbf{I}_{d_z} \otimes \mathbf{C}) \tilde{\boldsymbol{\Sigma}}_n^{-1}) \xrightarrow{\mathbb{P}} \gamma_z \text{tr}(\mathbf{C} \mathbf{C}_T(\eta_\ell)^{-1}).$$

990 **Proof of Lemma 2.** We observe that $\phi_T(\mathbf{x})\phi_T(\mathbf{x})^\top = (\mathbf{z}(\mathbf{x})\mathbf{z}(\mathbf{x})^\top) \otimes (\mathbf{w}(\mathbf{x})\mathbf{w}(\mathbf{x})^\top)$, and the sample
 991 covariance matrix $\tilde{\boldsymbol{\Sigma}}_n$ can be partitioned into $d_z \times d_z$ blocks of size $p_T \times p_T$:
 992

$$\tilde{\boldsymbol{\Sigma}}_n = \left[\tilde{\boldsymbol{\Sigma}}_n^{(k,l)} \right]_{k,l \in [d_z]} \quad \text{where} \quad \tilde{\boldsymbol{\Sigma}}_n^{(k,l)} = \frac{1}{n} \sum_{i=1}^n z_k(\tilde{\mathbf{x}}_i) z_l(\tilde{\mathbf{x}}_i) \cdot \mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top \in \mathbb{R}^{p_T \times p_T},$$

993 where for any $\mathbf{x} \in \mathcal{X}$, $z_k(\mathbf{x})$ is the k -th entry of $\mathbf{z}(\mathbf{x})$. Notice that $\mathbb{E}[\tilde{\boldsymbol{\Sigma}}_n^{(k,l)}] = \delta_{k,l} \mathbf{C}_T(\eta_\ell)$ where
 994 $\delta_{k,l}$ is the Kronecker delta since $\mathbf{z}(\mathbf{x})$ and $\mathbf{w}(\mathbf{x})$ are independent, and $\mathbb{E}[z_k(\mathbf{x}) z_l(\mathbf{x})] = \delta_{k,l}$ given
 995 $\mathbf{z}(\mathbf{x}) \sim \mathcal{N}(\mathbf{0}_{d_z}, \mathbf{I}_{d_z})$.

1002 **Off-diagonal blocks are negligible.** For any $k, l \in [d_z]$ with $k \neq l$, we define $p_T \times p_T$ self-adjoint
 1003 matrices,
 1004

$$\mathbf{Y}_i^{(k,l)} := \frac{1}{n} z_k(\tilde{\mathbf{x}}_i) z_l(\tilde{\mathbf{x}}_i) (\mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top - \mathbf{C}_T(\eta_\ell)) \quad \text{and} \quad \mathbf{R}_i^{(k,l)} := \frac{1}{n} z_k(\tilde{\mathbf{x}}_i) z_l(\tilde{\mathbf{x}}_i) \mathbf{C}_T(\eta_\ell),$$

1007 where recall from the derivation of (8) that $\mathbf{C}_T(\eta_\ell) = \mathbb{E}_{\mathcal{D}(\eta_\ell)}[\mathbf{w}(\mathbf{x}) \mathbf{w}(\mathbf{x})^\top]$. Since $\mathbb{E}[z_k(\mathbf{x}) z_l(\mathbf{x})] =$
 1008 $\delta_{k,l}$, we have $\mathbb{E}[\mathbf{Y}_i^{(k,l)}] = \mathbb{E}[\mathbf{R}_i^{(k,l)}] = \mathbf{0}_{p_T \times p_T}$ for $k \neq l$, and
 1009

$$\tilde{\boldsymbol{\Sigma}}_n^{(k,l)} = \sum_{i=1}^n \mathbf{Y}_i^{(k,l)} + \mathbf{R}_i^{(k,l)}, \quad \mathbb{E}[\tilde{\boldsymbol{\Sigma}}_n^{(k,l)}] = \mathbf{0}_{p_T \times p_T}.$$

1012 Let $L_n := 4\sqrt{\log(n)}$ and consider the event
 1013

$$E_n := \left\{ \max_{i \in [n]} \|\mathbf{z}(\tilde{\mathbf{x}}_i)\|_\infty^2 \leq L_n^2 \right\} \wedge \left\{ \max_{i \in [n]} \|\mathbf{w}(\tilde{\mathbf{x}}_i)\|_2^2 \leq L_n^2 \right\}.$$

1017 First, for $\mathbf{z}(\tilde{\mathbf{x}}_i) \sim \mathcal{N}(\mathbf{0}_{d_z}, \mathbf{I}_{d_z})$, the union bound and Gaussian tail bound imply that
 1018

$$\begin{aligned} \Pr \left[\max_{i \in [n]} \|\mathbf{z}(\tilde{\mathbf{x}}_i)\|_\infty > \frac{L_n}{\sqrt{2}} \right] &= \Pr \left[\max_{i \in [n], k \in [d_z]} |z_k(\tilde{\mathbf{x}}_i)| > \frac{L_n}{\sqrt{2}} \right] \\ &\leq 2nd_z \exp \left(-\frac{L_n^2}{4} \right) = \frac{2d_z}{n} \cdot n^{-2} = o(n^{-1}). \end{aligned}$$

1023 Meanwhile, we observe that for $\mathbf{g}_i \sim \mathcal{N}(\mathbf{0}_{p_T-1}, \mathbf{I}_{p_T-1})$,
 1024

$$\|\mathbf{w}(\tilde{\mathbf{x}}_i)\|_2^2 = 1 + \left\| [\mathbf{w}(\tilde{\mathbf{x}}_i)]_{2:p_T} \right\|_2^2 \leq 2\eta_\ell^2 \|\boldsymbol{\mu}_T\|_2^2 + 2\sigma_\xi^2 \|\mathbf{g}_i\|_2^2,$$

1026 while the Laurent-Massart χ^2 tail bound (Laurent & Massart, 2000) implies that
 1027

$$1028 \Pr \left[\|\mathbf{g}_i\|_2^2 > p_T - 1 + 2\sqrt{(p_T - 1)t} + 2t \right] \leq e^{-t}, \quad \forall t > 0.$$

1029 Then, for fixed and finite p_T and $\|\mathbf{\mu}_T\|_2$, with a sufficiently large n , there exists a constant $a_n > 1/8$
 1030 such that applying the union bound with $t = a_n L_n^2$ yields that
 1031

$$1032 \Pr \left[\max_{i \in [n]} \|\mathbf{w}(\tilde{\mathbf{x}}_i)\|_2^2 > L_n^2 \right] \leq n \exp(-a_n L_n^2) = o(n^{-1}).$$

1033 Applying the union bound again, we get $\Pr[E_n^c] = o(n^{-1})$ for sufficiently large n . Meanwhile,
 1034 conditioned on E_n , we have for any $i \in [n]$,
 1035

$$1036 \left\| \mathbf{Y}_i^{(k,l)} \right\|_2 \leq \frac{1}{n} \|\mathbf{z}(\tilde{\mathbf{x}}_i)\|_\infty^2 \|\mathbf{w}(\tilde{\mathbf{x}}_i)\|_2^2 \leq \frac{L_n^4}{n}, \quad \left\| \mathbf{R}_i^{(k,l)} \right\|_2 \leq \frac{1}{n} \|\mathbf{z}(\tilde{\mathbf{x}}_i)\|_\infty^2 \|\mathbf{C}_T(\eta_\ell)\|_2 \lesssim \frac{L_n^2}{n},$$

1037 which implies that
 1038

$$1039 \left\| \sum_{i=1}^n \mathbb{E} \left[(\mathbf{Y}_i^{(k,l)})^2 \mid E_n \right] \right\|_2 \leq \sum_{i=1}^n \mathbb{E} \left[\left\| \mathbf{Y}_i^{(k,l)} \right\|_2^2 \mid E_n \right] \leq n \cdot \frac{L_n^8}{n^2} = \frac{L_n^8}{n},$$

$$1040 \left\| \sum_{i=1}^n \mathbb{E} \left[(\mathbf{R}_i^{(k,l)})^2 \mid E_n \right] \right\|_2 \leq \sum_{i=1}^n \mathbb{E} \left[\left\| \mathbf{R}_i^{(k,l)} \right\|_2^2 \mid E_n \right] \lesssim n \cdot \frac{L_n^4}{n^2} = \frac{L_n^4}{n}.$$

1041 Then, applying the matrix Bernstein inequality (Tropp, 2012, Theorem 1.6) to $\sum_{i=1}^n \mathbf{Y}_i^{(k,l)}$ and
 1042 $\sum_{i=1}^n \mathbf{R}_i^{(k,l)}$ and a union bound over all $k, l \in [d_z]$ off-diagonal blocks with $k \neq l$ yields that
 1043

$$1044 \Pr \left[\max_{k \neq l} \left\| \tilde{\Sigma}_n^{(k,l)} \right\|_2 \gtrsim \sqrt{\frac{\log(n)}{n}} \right] \leq \Pr [E_n^c] + \Pr \left[\max_{k \neq l} \left\| \tilde{\Sigma}_n^{(k,l)} \right\|_2 \gtrsim \sqrt{\frac{\log(n)}{n}} \mid E_n \right] \\ 1045 \leq o(n^{-1}) + 2p_T n^{-\Omega\left(\frac{n^2}{\log^4(n)}\right)} \cdot d_z^2 = o(1),$$

1046 and therefore, the off-diagonal blocks are negligible:
 1047

$$1048 \max_{k \neq l} \left\| \tilde{\Sigma}_n^{(k,l)} \right\|_2 = O_{\mathbb{P}} \left(\sqrt{\frac{\log(n)}{n}} \right).$$

1049 **Diagonal blocks are concentrated.** Consider the k -th diagonal block $\tilde{\Sigma}_n^{(k,k)}$ for any fixed $k \in [d_z]$:
 1050

$$1051 \tilde{\Sigma}_n^{(k,k)} = \frac{1}{n} \sum_{i=1}^n z_k(\tilde{\mathbf{x}}_i)^2 \cdot \mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top \\ 1052 = \frac{1}{n} \sum_{i=1}^n \mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top + \frac{1}{n} \sum_{i=1}^n (z_k(\tilde{\mathbf{x}}_i)^2 - 1) \cdot \mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top,$$

1053 where we denote $\hat{\mathbf{C}}_{T,n} = \frac{1}{n} \sum_{i=1}^n \mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top$. Let
 1054

$$1055 s_k := \frac{1}{n} \sum_{i=1}^n z_k(\tilde{\mathbf{x}}_i)^2. \quad (16)$$

1056 Then,
 1057

$$1058 \tilde{\Sigma}_n^{(k,k)} - s_k \mathbf{C}_T(\eta_\ell) = (\hat{\mathbf{C}}_{T,n} - \mathbf{C}_T(\eta_\ell)) + \frac{1}{n} \sum_{i=1}^n (z_k(\tilde{\mathbf{x}}_i)^2 - 1) \cdot (\mathbf{w}(\tilde{\mathbf{x}}_i) \mathbf{w}(\tilde{\mathbf{x}}_i)^\top - \mathbf{C}_T(\eta_\ell)),$$

1059 where both terms are sums of independent random matrices with zero mean. Leveraging the same
 1060 argument as for the off-diagonal blocks using the same event E_n , the matrix Bernstein inequality
 1061 (Tropp, 2012, Theorem 1.6), and a union bound over all $k \in [d_z]$, we have for sufficiently large
 1062 n ,

$$1063 \max_{k \in [d_z]} \left\| \tilde{\Sigma}_n^{(k,k)} - s_k \mathbf{C}_T(\eta_\ell) \right\|_2 = O_{\mathbb{P}} \left(\sqrt{\frac{\log(n)}{n}} \right). \quad (17)$$

1080 Also, the χ^2 concentration (Laurent & Massart, 2000) implies that for any fixed $\epsilon \in (0, 1/2)$, as
 1081 $d_z, n \rightarrow \infty$ with $d_z/n \rightarrow \gamma_z \in (0, p_T^{-1})$,
 1082

$$\begin{aligned} \Pr \left[\min_{k \in [d_z]} s_k < 1 - \epsilon \right] &\leq d_z \Pr [s_k < 1 - \epsilon] \leq d_z \exp(-\Theta(n\epsilon^2)) = o(1), \\ \Pr \left[\max_{k \in [d_z]} s_k > 1 + \epsilon \right] &\leq d_z \Pr [s_k > 1 + \epsilon] \leq d_z \exp(-\Theta(n\epsilon^2)) = o(1), \end{aligned} \quad (18)$$

1088 so that all s_k 's are close to 1 with high probability.

1091 **Concentration of $\tilde{\Sigma}_n^{-1}$.** Let $\mathbf{D}_n = \text{diag}(s_1 \mathbf{C}_T(\eta_\ell), \dots, s_{d_z} \mathbf{C}_T(\eta_\ell))$ be the block-diagonal matrix
 1092 with k -th diagonal block $s_k \mathbf{C}_T(\eta_\ell)$ for all $k \in [d_z]$; and $\mathbf{E}_n = \tilde{\Sigma}_n - \mathbf{D}_n$ be the fluctuations around
 1093 \mathbf{D}_n . Since $\mathbf{C}_T(\eta_\ell)$ is positive definite, (18) implies that $\|\mathbf{D}_n^{-1}\|_2 < \infty$ with high probability for
 1094 sufficiently large n . Then, the resolvent identity implies that
 1095

$$\tilde{\Sigma}_n^{-1} = (\mathbf{D}_n + \mathbf{E}_n)^{-1} = \mathbf{D}_n^{-1} - \mathbf{D}_n^{-1} \mathbf{E}_n (\mathbf{D}_n + \mathbf{E}_n)^{-1}. \quad (19)$$

1096 In particular, the block matrix inversion formula implies that for any $k \in [d_z]$, the k -th diagonal block
 1097 of $\tilde{\Sigma}_n^{-1}$, denoted as $(\tilde{\Sigma}_n^{-1})^{(k,k)} \in \mathbb{R}^{p_T \times p_T}$ is concentrated around the k -th diagonal block of \mathbf{D}_n^{-1} ,
 1098 $(s_k \mathbf{C}_T(\eta_\ell))^{-1}$:

$$\left\| (\tilde{\Sigma}_n^{-1})^{(k,k)} - (s_k \mathbf{C}_T(\eta_\ell))^{-1} \right\|_2 \lesssim \|\mathbf{E}_n\|_2 = O_{\mathbb{P}} \left(\sqrt{\frac{\log(n)}{n}} \right), \quad (20)$$

1102 **Concentration of the trace.** Finally, notice that the trace of interest, $\frac{1}{n} \text{tr} \left((\mathbf{I}_{d_z} \otimes \mathbf{C}) \tilde{\Sigma}_n^{-1} \right)$, de-
 1103 pends only on the diagonal blocks of $\tilde{\Sigma}_n^{-1}$. Then, (20) implies that
 1104

$$\begin{aligned} \frac{1}{n} \text{tr} \left((\mathbf{I}_{d_z} \otimes \mathbf{C}) \tilde{\Sigma}_n^{-1} \right) &= \frac{1}{n} \sum_{k=1}^{d_z} \text{tr} \left(\mathbf{C} (\tilde{\Sigma}_n^{-1})^{(k,k)} \right) \\ &= \left(\frac{1}{n} \sum_{k=1}^{d_z} \frac{1}{s_k} \right) \text{tr} (\mathbf{C} \mathbf{C}_T(\eta_\ell)^{-1}) + O_{\mathbb{P}} \left(\sqrt{\frac{\log(n)}{n}} \right) \\ &= \left(\frac{1}{d_z} \sum_{k=1}^{d_z} \frac{1}{s_k} \right) \cdot \frac{d_z}{n} \text{tr} (\mathbf{C} \mathbf{C}_T(\eta_\ell)^{-1}) + o_{\mathbb{P}}(1). \end{aligned}$$

1120 Since $\{ns_k\}_{k=1}^{d_z}$ are independent and χ_n^2 distributed, for any fixed $n > 2$, $\mathbb{E}[s_k^{-1}] = \frac{n}{n-2}$. Then, the
 1121 weak law of large numbers implies that as $d_z, n \rightarrow \infty$,
 1122

$$\frac{1}{d_z} \sum_{k=1}^{d_z} \frac{1}{s_k} \xrightarrow{\mathbb{P}} \frac{n}{n-2} \xrightarrow{n \rightarrow \infty} 1.$$

1123 Putting everything together with $d_z/n \rightarrow \gamma_z \in (0, p_T^{-1})$ completes the proof. \square
 1124

D.2 W2S FINE-TUNING OF STRONG STUDENT

1131 **Theorem 3** (W2S fine-tuning of strong student (formal restatement of Theorem 2)). *Under Assumption 1, as $d_z, n, N \rightarrow \infty$ with $d_z/n \rightarrow \gamma_z \in (0, p_T^{-1})$ and $d_z/N \rightarrow \nu_z \in (0, p_S^{-1})$,*

1134 $f_S(\mathbf{x}) = \varphi_S(\mathbf{x})^\top \boldsymbol{\theta}_S = \phi_S(\mathbf{x})^\top \boldsymbol{\beta}_S$ from (2) satisfies
 1135

1136 $\mathbb{E}_{\mathcal{S}_x \sim \mathcal{D}(\eta_u)^N, \tilde{\mathcal{S}} \sim \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_{\eta_t}(f_S)] \xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \left(\frac{(p_{T \wedge S})}{\nu_z} + \right. +$
 1137 $\left. \mathcal{V}_S^{(0)} \leq \mathcal{V}_T^{(0)} \right)$
 1138

1139 $\frac{\|(\eta_u - \eta_\ell)\boldsymbol{\mu}_T + (\eta_t - \eta_u)\boldsymbol{\Xi}\boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2} +$
 1140

1141 $\mathcal{V}_S^{(1)} \leq \mathcal{V}_T^{(1)} \text{ when } \eta_u = \eta_\ell$
 1142

1143

1144 $\nu_z(p_T - p_{T \wedge S}) \left(p_S + (\eta_t - \eta_u)^2 \frac{\|\boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2} \right) \right),$
 1145

1146

1147 $\mathcal{E}_S = \Theta(\nu_z) \ll 1 \text{ negligible when } \nu_z \ll 1$
 1148

1149 where $p_{T \wedge S} = 1 + \|\boldsymbol{\Xi}\|_F^2 \in [1, p_S]$ quantifies the effective dimension of group features learned by
 1150 the strong student from the weak teacher.

1151 Notice that in Theorem 3, $\mathcal{V}_S^{(0)} + \mathcal{V}_S^{(1)}$ is dominant and will be small if \mathbf{T}, \mathbf{S} are nearly orthogonal
 1152 (i.e., $\|\boldsymbol{\Xi}\|_2 \approx 0$) and $\eta_u \approx \eta_t$; whereas \mathcal{E}_S tends to be much smaller than $\mathcal{V}_S^{(0)} + \mathcal{V}_S^{(1)}$, especially
 1153 since unlabeled data is usually abundant compared to labeled data (i.e., $\nu_z \ll \gamma_z$).
 1154

1155

1156 *Proof of Theorems 2 and 3.* We first introduce some helpful notions for the proof. Recall $\phi_S(\mathbf{x}) =$
 1157 $\mathbf{z}(\mathbf{x}) \otimes \boldsymbol{\psi}(\mathbf{x})$. Let $\boldsymbol{\Sigma}_{\phi_S, \eta} = \mathbb{E}_{\mathcal{D}(\eta)}[\phi_S(\mathbf{x})\phi_S(\mathbf{x})^\top]$ for any $\eta \in [0, 1]$ and observe that
 1158

1159 $\boldsymbol{\Sigma}_{\phi_S, \eta} = \mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta), \quad \mathbf{C}_S(\eta) = \mathbb{E}_{\mathcal{D}(\eta)}[\boldsymbol{\psi}(\mathbf{x})\boldsymbol{\psi}(\mathbf{x})^\top] = \begin{bmatrix} 1 & \eta\boldsymbol{\mu}_S^\top \\ \eta\boldsymbol{\mu}_S & \sigma_\xi^2 \mathbf{I}_{p_S-1} + \eta^2 \boldsymbol{\mu}_S \boldsymbol{\mu}_S^\top \end{bmatrix}. \quad (21)$
 1160

1161 The block matrix inversion formula implies that
 1162

1163 $\mathbf{C}_S(\eta)^{-1} = \begin{bmatrix} 1 + \sigma_\xi^{-2}\eta^2\|\boldsymbol{\mu}_S\|_2^2 & -\sigma_\xi^{-2}\eta\boldsymbol{\mu}_S^\top \\ -\sigma_\xi^{-2}\eta\boldsymbol{\mu}_S & \sigma_\xi^{-2}\mathbf{I}_{p_S-1} \end{bmatrix}. \quad (22)$
 1164

1165 Meanwhile, the cross covariance of the student-teacher representations under $\mathcal{D}(\eta)$ is given by
 1166

1167 $\boldsymbol{\Sigma}_{\phi_S, \phi_T, \eta} = \mathbb{E}_{\mathcal{D}(\eta)}[\phi_S(\mathbf{x})\phi_T(\mathbf{x})^\top] = \mathbb{E}_{\mathcal{D}(\eta)}[(\mathbf{z}(\mathbf{x}) \otimes \boldsymbol{\psi}(\mathbf{x}))(\mathbf{z}(\mathbf{x}) \otimes \mathbf{w}(\mathbf{x}))^\top]$
 1168
 1169 $= \mathbb{E}_{\mathcal{D}(\eta)}[\mathbf{z}(\mathbf{x})\mathbf{z}(\mathbf{x})^\top] \otimes \mathbb{E}_{\mathcal{D}(\eta)}[\boldsymbol{\psi}(\mathbf{x})\mathbf{w}(\mathbf{x})^\top] = \mathbf{I}_{d_z} \otimes \mathbf{A}(\eta),$
 1170

1171 where

1172 $\mathbf{A}(\eta) = \mathbb{E}_{\mathcal{D}(\eta)}[\boldsymbol{\psi}(\mathbf{x})\mathbf{w}(\mathbf{x})^\top] = \mathbb{E}_{\mathcal{D}(\eta)} \left[\begin{bmatrix} 1 \\ \mathbf{S}^\top \boldsymbol{\xi}(\mathbf{x}) \end{bmatrix} [1 \quad \boldsymbol{\xi}(\mathbf{x})^\top \mathbf{T}] \right] \quad (23)$
 1173

1174
 1175 $= \begin{bmatrix} 1 & \eta\boldsymbol{\mu}_T^\top \\ \eta\boldsymbol{\mu}_S & \sigma_\xi^2 \mathbf{S}^\top \mathbf{T} + \eta^2 \boldsymbol{\mu}_S \boldsymbol{\mu}_T^\top \end{bmatrix} \in \mathbb{R}^{p_S \times p_T}. \quad (24)$
 1176

1177 **Close-form solution and population-optimal predictor of W2S fine-tuning.** Given the equivalence
 1178 between (2) and (6), we consider the latter throughout the proof. Adapting the notion from the
 1179 proof of Theorem 1, given the labeled set $\tilde{\mathcal{S}} = \{(\tilde{\mathbf{x}}_i, \tilde{y}_i) \mid i \in [n]\} \sim \mathcal{D}(\eta_\ell)^n$ and the unlabeled set
 1180 $\mathcal{S} = \{(\mathbf{x}_i, y_i) \mid i \in [N]\} \sim \mathcal{D}_\mathbf{x}(\eta_u)^N$ with unknown y_i 's, we denote
 1181

1182 $\tilde{\boldsymbol{\Phi}}_T = [\phi_T(\tilde{\mathbf{x}}_1), \dots, \phi_T(\tilde{\mathbf{x}}_n)]^\top \in \mathbb{R}^{n \times d_T}, \quad \tilde{\mathbf{y}} = [\tilde{y}_1, \dots, \tilde{y}_n]^\top \in \mathbb{R}^n,$
 1183
 1184 $\boldsymbol{\Phi}_S = [\phi_S(\mathbf{x}_1), \dots, \phi_S(\mathbf{x}_N)]^\top \in \mathbb{R}^{N \times d_S}, \quad \boldsymbol{\Phi}_T = [\phi_T(\mathbf{x}_1), \dots, \phi_T(\mathbf{x}_N)]^\top \in \mathbb{R}^{N \times d_T}.$
 1185

1186 Then, since $n > d_T$ and $N > d_S$ by Assumption 1, (6) admits a unique closed-form solution
 1187

1188 $\boldsymbol{\beta}_S = (\boldsymbol{\Phi}_S^\top \boldsymbol{\Phi}_S)^{-1} \boldsymbol{\Phi}_S^\top \boldsymbol{\Phi}_T \boldsymbol{\beta}_T \quad \text{where} \quad \boldsymbol{\beta}_T = (\tilde{\boldsymbol{\Phi}}_T^\top \tilde{\boldsymbol{\Phi}}_T)^{-1} \tilde{\boldsymbol{\Phi}}_T^\top \tilde{\mathbf{y}}$

from (11). Recall from Lemma 1 that the population-optimal linear predictor over ϕ_T in Lemma 1 is $f_T^\infty(\mathbf{x}) = \phi_T(\mathbf{x})^\top \beta_T^\infty = f_*(\mathbf{x})$ with $\beta_T^\infty = \beta_* \otimes \mathbf{e}_1$. Conditioned on $f_T^\infty(\mathbf{x})$, the population-optimal linear predictor over ϕ_S is given by

$$\begin{aligned}\beta_S^\infty &= \mathbb{E}_{\mathcal{D}(\eta_u)} [\phi_S(\mathbf{x}) \phi_S(\mathbf{x})^\top]^{-1} \mathbb{E}_{\mathcal{D}(\eta_u)} [\phi_S(\mathbf{x}) \phi_T(\mathbf{x})^\top] \beta_T^\infty \\ &= (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_u))^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{A}(\eta_u)) \beta_T^\infty \\ &= (\mathbf{I}_{d_z} \otimes (\mathbf{C}_S(\eta_u)^{-1} \mathbf{A}(\eta_u))) (\beta_* \otimes \mathbf{e}_1) \\ &= \beta_* \otimes (\mathbf{C}_S(\eta_u)^{-1} \mathbf{A}(\eta_u) \mathbf{e}_1) = \beta_* \otimes \mathbf{e}_1,\end{aligned}$$

which implies that $f_S^\infty(\mathbf{x}) = \phi_S(\mathbf{x})^\top \beta_S^\infty = f_*(\mathbf{x})$, *i.e.*, a strong student W2S fine-tuned with pseudolabels from the Bayes-optimal weak teacher over the population is also Bayes-optimal. Therefore, the student estimator in (6) differs from β_S^∞ by

$$\begin{aligned}\beta_S - \beta_S^\infty &= (\Phi_S^\top \Phi_S)^{-1} \Phi_S^\top \Phi_T (\beta_T - \beta_T^\infty) \\ &= (\Phi_S^\top \Phi_S)^{-1} \Phi_S^\top \Phi_T (\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \tilde{\Phi}_T^\top \tilde{\epsilon},\end{aligned}$$

and the estimation error of W2S fine-tuning is given by

$$\mathbf{ER}_{\eta_t}(f_S) = \mathbb{E}_{\mathcal{D}(\eta_t)} [(f_S(\mathbf{x}) - f_*(\mathbf{x}))^2] = \mathbb{E}_{\mathcal{D}(\eta_t)} [(\phi_S(\mathbf{x})^\top (\beta_S - \beta_S^\infty))^2] = \|\beta_S - \beta_S^\infty\|_{\Sigma_{\phi_S, \eta_t}}^2.$$

Then, conditioned on $\tilde{\Phi}_T$ and Φ_S, Φ_T , the excess risk can be expressed as

$$\begin{aligned}\mathbb{E}_{\tilde{\epsilon}} [\mathbf{ER}_{\eta_t}(f_S) | \tilde{\Phi}_T, \Phi_S, \Phi_T] &= \mathbb{E}_{\tilde{\epsilon}} \left[\left\| (\Phi_S^\top \Phi_S)^{-1} \Phi_S^\top \Phi_T (\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \tilde{\Phi}_T^\top \tilde{\epsilon} \right\|_{\Sigma_{\phi_S, \eta_t}}^2 | \tilde{\Phi}_T, \Phi_S, \Phi_T \right] \\ &= \sigma_y^2 \operatorname{tr} \left(\Sigma_{\phi_S, \eta_t} (\Phi_S^\top \Phi_S)^{-1} \Phi_S^\top \Phi_T (\tilde{\Phi}_T^\top \tilde{\Phi}_T)^{-1} \tilde{\Phi}_T^\top \Phi_S (\Phi_S^\top \Phi_S)^{-1} \right).\end{aligned}\tag{25}$$

Concentration of sample covariance matrices. Define the sample (cross) covariance matrices

$$\hat{\Sigma}_{S,N} = \frac{1}{N} \Phi_S^\top \Phi_S, \quad \hat{\Sigma}_{S,T,N} = \frac{1}{N} \Phi_S^\top \Phi_T, \quad \tilde{\Sigma}_{T,n} = \frac{1}{n} \tilde{\Phi}_T^\top \tilde{\Phi}_T.$$

Then, taking the expectation of (25) over $\tilde{\mathcal{S}}$ and \mathcal{S}_x yields

$$\mathbb{E}_{\tilde{\mathcal{S}}, \mathcal{S}_x} [\mathbf{ER}_{\eta_t}(f_S)] = \frac{\sigma_y^2}{n} \operatorname{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\hat{\Sigma}_{S,T,N}^\top \hat{\Sigma}_{S,N}^{-1} \Sigma_{\phi_S, \eta_t} \hat{\Sigma}_{S,N}^{-1} \hat{\Sigma}_{S,T,N} \right] \mathbb{E}_{\tilde{\mathcal{S}}} [\tilde{\Sigma}_{T,n}^{-1}] \right).\tag{26}$$

At the proportional asymptotic limit, Lemma 3 below shows that

$$\begin{aligned}\frac{1}{n} \operatorname{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\hat{\Sigma}_{S,T,N}^\top \hat{\Sigma}_{S,N}^{-1} \Sigma_{\phi_S, \eta_t} \hat{\Sigma}_{S,N}^{-1} \hat{\Sigma}_{S,T,N} \right] \mathbb{E}_{\tilde{\mathcal{S}}} [\tilde{\Sigma}_{T,n}^{-1}] \right) \\ \xrightarrow{\mathbb{P}} \gamma_z \operatorname{tr} (\mathbf{C}_{T,S}(\eta_t, \eta_u) \mathbf{C}_T(\eta_\ell)^{-1}) + \gamma_z \nu_z (p_T - p_{T \wedge S}) \operatorname{tr} (\mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1}),\end{aligned}$$

Leveraging (21), (22), and (23), we have

$$\begin{aligned}\operatorname{tr} (\mathbf{C}_{T,S}(\eta_t, \eta_u) \mathbf{C}_T(\eta_\ell)^{-1}) &= \operatorname{tr} (\mathbf{A}(\eta_u)^\top \mathbf{C}_S(\eta_u)^{-1} \mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1} \mathbf{A}(\eta_u) \mathbf{C}_T(\eta_\ell)^{-1}) \\ &= 1 + \|\boldsymbol{\Xi}\|_F^2 + \frac{\|(\eta_u - \eta_\ell) \boldsymbol{\mu}_T + (\eta_t - \eta_u) \boldsymbol{\Xi} \boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2} \\ &= p_{T \wedge S} + \frac{\|(\eta_u - \eta_\ell) \boldsymbol{\mu}_T + (\eta_t - \eta_u) \boldsymbol{\Xi} \boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2},\end{aligned}$$

while an analogous derivation as in (15) implies that

$$\operatorname{tr} (\mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1}) = p_S + (\eta_t - \eta_u)^2 \frac{\|\boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2}.$$

1242 Overall, plugging everything back to (26) yields
1243

$$\begin{aligned} \mathbb{E}_{\tilde{\mathcal{S}}, \mathcal{S}_x} [\mathbf{ER}_{\eta_t}(f_S)] &\xrightarrow{\mathbb{P}} \sigma_y^2 \gamma_z \left(p_{T \wedge S} + \frac{\|(\eta_u - \eta_t) \boldsymbol{\mu}_T + (\eta_t - \eta_u) \boldsymbol{\Xi} \boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2} \right) \\ &\quad + \sigma_y^2 \gamma_z \nu_z (p_T - p_{T \wedge S}) \left(p_S + (\eta_t - \eta_u)^2 \frac{\|\boldsymbol{\mu}_S\|_2^2}{\sigma_\xi^2} \right), \end{aligned}$$

□

1250 **Lemma 3.** *In the proof of Theorem 2, at the proportional asymptotic limit,*
1251

$$\begin{aligned} \frac{1}{n} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\widehat{\boldsymbol{\Sigma}}_{S,T,N}^\top \widehat{\boldsymbol{\Sigma}}_{S,N}^{-1} \boldsymbol{\Sigma}_{\phi_S, \eta_t} \widehat{\boldsymbol{\Sigma}}_{S,N}^{-1} \widehat{\boldsymbol{\Sigma}}_{S,T,N} \right] \mathbb{E}_{\tilde{\mathcal{S}}} \left[\widetilde{\boldsymbol{\Sigma}}_{T,n}^{-1} \right] \right) \\ \xrightarrow{\mathbb{P}} \gamma_z \text{tr} \left(\mathbf{C}_{T,S}(\eta_t, \eta_u) \mathbf{C}_T(\eta_t)^{-1} \right) + \gamma_z \nu_z (p_T - p_{T \wedge S}) \text{tr} \left(\mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1} \right), \end{aligned}$$

1255 where $\mathbf{C}_{T,S}(\eta_t, \eta_u) \in \mathbb{R}^{p_T \times p_T}$ is defined as

$$\mathbf{C}_{T,S}(\eta_t, \eta_u) = \mathbf{A}(\eta_u)^\top \mathbf{C}_S(\eta_u)^{-1} \mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1} \mathbf{A}(\eta_u),$$

1258 *Proof of Lemma 3.* The proof mostly follows the same argument as in Lemma 2, with the key
1259 difference being a careful treatment of the off-diagonal blocks in the sample (cross) covariance
1260 matrices $\widehat{\boldsymbol{\Sigma}}_{S,N}$ and $\widehat{\boldsymbol{\Sigma}}_{S,T,N}$, which are still small but with an additional non-negligible higher-order
1261 moment in the proportional asymptotic limit.

1262 Following the proof of Lemma 2, “Separation of block diagonals and off-diagonals”, we first partition
1263 $\widehat{\boldsymbol{\Sigma}}_{S,N}$ and $\widehat{\boldsymbol{\Sigma}}_{S,T,N}$ into $d_z \times d_z$ blocks:
1264

$$\widehat{\boldsymbol{\Sigma}}_{S,N} = \left[\widehat{\boldsymbol{\Sigma}}_{S,N}^{(k,l)} \right]_{k,l=1}^{d_z}, \quad \widehat{\boldsymbol{\Sigma}}_{S,T,N} = \left[\widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(k,l)} \right]_{k,l=1}^{d_z},$$

1267 where $\widehat{\boldsymbol{\Sigma}}_{S,N}^{(k,l)} \in \mathbb{R}^{p_S \times p_S}$ and $\widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(k,l)} \in \mathbb{R}^{p_S \times p_T}$ are given by

$$\widehat{\boldsymbol{\Sigma}}_{S,N}^{(k,l)} = \frac{1}{N} \sum_{i=1}^N z_k(\mathbf{x}_i) z_l(\mathbf{x}_i) \cdot \boldsymbol{\psi}(\mathbf{x}_i) \boldsymbol{\psi}(\mathbf{x}_i)^\top,$$

$$\widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(k,l)} = \frac{1}{N} \sum_{i=1}^N z_k(\mathbf{x}_i) z_l(\mathbf{x}_i) \cdot \boldsymbol{\psi}(\mathbf{x}_i) \mathbf{w}(\mathbf{x}_i)^\top.$$

1274 We denote

$$s_{kl} = \frac{1}{N} \sum_{i=1}^N z_k(\mathbf{x}_i) z_l(\mathbf{x}_i) \quad \text{for all } k, l \in [d_z],$$

1278 and observe that for $k \neq l$, $\mathbb{E}[s_{kl}] = 0$, and for $k = l$, $\mathbb{E}[s_{kk}] = 1$. We therefore observe and denote
1279 that

$$\mathbf{D}_S := \mathbb{E}_{\mathcal{D}(\eta_u)} \left[\widehat{\boldsymbol{\Sigma}}_{S,N} \right] = \mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_u), \quad \mathbf{D}_{S,T} := \mathbb{E}_{\mathcal{D}(\eta_u)} \left[\widehat{\boldsymbol{\Sigma}}_{S,T,N} \right] = \mathbf{I}_{d_z} \otimes \mathbf{A}(\eta_u). \quad (27)$$

1280 We further define the reminder fluctuation matrices around \mathbf{D}_S and $\mathbf{D}_{S,T}$:

$$\mathbf{E}_S = \widehat{\boldsymbol{\Sigma}}_{S,N} - \mathbf{D}_S, \quad \mathbf{E}_{S,T} = \widehat{\boldsymbol{\Sigma}}_{S,T,N} - \mathbf{D}_{S,T}, \quad (28)$$

1284 where

$$\begin{aligned} \mathbf{E}_S = \left[\mathbf{E}_S^{(k,l)} \right]_{k,l=1}^{d_z} &= \begin{bmatrix} \widehat{\boldsymbol{\Sigma}}_{S,N}^{(1,1)} - \mathbf{C}_S(\eta_u) & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(1,2)} & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(1,d_z)} \\ \widehat{\boldsymbol{\Sigma}}_{S,N}^{(2,1)} & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(2,2)} - \mathbf{C}_S(\eta_u) & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(2,d_z)} \\ \vdots & \vdots & \ddots & \vdots \\ \widehat{\boldsymbol{\Sigma}}_{S,N}^{(d_z,1)} & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(d_z,2)} & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,N}^{(d_z,d_z)} - \mathbf{C}_S(\eta_u) \end{bmatrix}, \\ \mathbf{E}_{S,T} = \left[\mathbf{E}_{S,T}^{(k,l)} \right]_{k,l=1}^{d_z} &= \begin{bmatrix} \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(1,1)} - \mathbf{A}(\eta_u) & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(1,2)} & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(1,d_z)} \\ \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(2,1)} & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(2,2)} - \mathbf{A}(\eta_u) & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(2,d_z)} \\ \vdots & \vdots & \ddots & \vdots \\ \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(d_z,1)} & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(d_z,2)} & \cdots & \widehat{\boldsymbol{\Sigma}}_{S,T,N}^{(d_z,d_z)} - \mathbf{A}(\eta_u) \end{bmatrix}. \end{aligned} \quad (29)$$

Again, following the proof of Lemma 2, the resolvent identity implies that

$$\begin{aligned} \widehat{\Sigma}_{S,N}^{-1} &= (\mathbf{D}_S + \mathbf{E}_S)^{-1} = \mathbf{D}_S^{-1} - \mathbf{D}_S^{-1} \mathbf{E}_S (\mathbf{D}_S + \mathbf{E}_S)^{-1} \\ &= \mathbf{D}_S^{-1} - \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} + \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} \mathbf{E}_S (\mathbf{D}_S + \mathbf{E}_S)^{-1}. \end{aligned} \quad (30)$$

Then, since $\mathbb{E}[\mathbf{E}_S] = \mathbf{0}_{d_S \times d_S}$ and $\mathbb{E}[\mathbf{E}_{S,T}] = \mathbf{0}_{d_S \times d_T}$, we have

$$\begin{aligned} &\mathbb{E}_{\mathcal{S}_x} \left[\widehat{\Sigma}_{S,T,N}^\top \widehat{\Sigma}_{S,N}^{-1} \Sigma_{\phi_S, \eta_t} \widehat{\Sigma}_{S,N}^{-1} \widehat{\Sigma}_{S,T,N} \right] \\ &= \mathbb{E} [(\mathbf{D}_{S,T} + \mathbf{E}_{S,T})^\top (\mathbf{D}_S + \mathbf{E}_S)^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) (\mathbf{D}_S + \mathbf{E}_S)^{-1} (\mathbf{D}_{S,T} + \mathbf{E}_{S,T})] \\ &= \mathbf{D}_{S,T}^\top \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{D}_{S,T} + \mathbf{R}_N \\ &\quad + \mathbb{E} [\mathbf{E}_{S,T}^\top \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{E}_{S,T}] \quad (=: \mathbf{R}_{S,T}) \\ &\quad + \mathbb{E} [\mathbf{D}_{S,T}^\top \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} \mathbf{D}_{S,T}] \quad (=: \mathbf{R}_{S,S}) \\ &\quad - \mathbb{E} [\mathbf{D}_{S,T}^\top \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{E}_{S,T}] \quad (=: \mathbf{R}_{S,S,T}) \\ &\quad - \mathbb{E} [\mathbf{E}_{S,T}^\top \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{E}_S \mathbf{D}_S^{-1} \mathbf{D}_{S,T}], \quad (=: \mathbf{R}_{S,T,S}) \end{aligned} \quad (31)$$

where $\|\mathbf{R}_N\|_2 = o_{\mathbb{P}}(1)$ for sufficiently large N ; \mathbf{E}_S and $\mathbf{E}_{S,T}$ are averages over N i.i.d. random matrices with $d_z \times d_z$ independent blocks. Therefore, when taking expectation for the second moments of \mathbf{E}_S and $\mathbf{E}_{S,T}$, the off-diagonal blocks in $\mathbf{R}_{S,T}, \mathbf{R}_{S,S}, \mathbf{R}_{S,S,T}, \mathbf{R}_{S,T,S} \in \mathbb{R}^{d_T \times d_T}$ vanish due to independence, and only the diagonal blocks remain, which are i.i.d. across $k \in [d_z]$. Notice that (27) implies that

$$\mathbf{D}_{S,T}^\top \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{D}_{S,T} = \mathbf{I}_{d_z} \otimes \mathbf{C}_{T,S}(\eta_t, \eta_u).$$

Also, we recall that the fourth moment of any Gaussian random vector $\mathbf{g} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ satisfies for any fixed matrix $\mathbf{M} \in \mathbb{R}^{d \times d}$,

$$\mathbb{E} [(\mathbf{g} \mathbf{g}^\top)^2] = (d+2) \mathbf{I}_d, \quad \mathbb{E} [(\mathbf{g} \mathbf{g}^\top) \mathbf{M} (\mathbf{g} \mathbf{g}^\top)] = \text{tr}(\mathbf{M}) \mathbf{I}_d + \mathbf{M} + \mathbf{M}^\top. \quad (32)$$

Define a function $g : \mathbb{R}^{d_T \times d_T} \rightarrow \mathbb{R}$ as

$$g(\mathbf{A}) = \frac{1}{n} \text{tr} \left(\mathbf{A} \mathbb{E}_{\tilde{\mathcal{S}}} \left[\tilde{\Sigma}_{T,n}^{-1} \right] \right).$$

Then, we have

$$\begin{aligned} &\frac{1}{n} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\widehat{\Sigma}_{S,T,N}^\top \widehat{\Sigma}_{S,N}^{-1} \Sigma_{\phi_S, \eta_t} \widehat{\Sigma}_{S,N}^{-1} \widehat{\Sigma}_{S,T,N} \right] \mathbb{E}_{\tilde{\mathcal{S}}} \left[\tilde{\Sigma}_{T,n}^{-1} \right] \right) \\ &= g(\mathbf{I}_{d_z} \otimes \mathbf{C}_{T,S}(\eta_t, \eta_u)) + g(\mathbf{R}_{S,T}) + g(\mathbf{R}_{S,S}) - g(\mathbf{R}_{S,S,T}) - g(\mathbf{R}_{S,T,S}) + o_{\mathbb{P}}(1), \end{aligned}$$

where given the $\mathbf{I}_{d_z} \otimes \mathbf{C}_{T,S}(\eta_t, \eta_u)$ structure, Lemma 2 then implies that under the proportional asymptotic limit,

$$g(\mathbf{I}_{d_z} \otimes \mathbf{C}_{T,S}(\eta_t, \eta_u)) \xrightarrow{\mathbb{P}} \gamma_z \text{tr} (\mathbf{C}_{T,S}(\eta_t, \eta_u) \mathbf{C}_T(\eta_\ell)^{-1}).$$

Let $\mathbf{M} := \mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1} \in \mathbb{R}^{p_S \times p_S}$ and $\mathbf{M}' := [\mathbf{M}]_{2:p_S, 2:p_S}$. Recall from (31) that $\mathbf{R}_{S,T} = \mathbb{E} [\mathbf{E}_{S,T}^\top \mathbf{D}_S^{-1} (\mathbf{I}_{d_z} \otimes \mathbf{C}_S(\eta_t)) \mathbf{D}_S^{-1} \mathbf{E}_{S,T}]$, (32) and (15), along with the proof of Lemma 2 imply that

$$g(\mathbf{R}_{S,T}) \xrightarrow{\mathbb{P}} \gamma_z \nu_z \left(p_T \text{tr}(\mathbf{M}) + \frac{2}{\sigma_\xi^2} \text{tr}(\mathbf{M}' \boldsymbol{\Xi}^\top \boldsymbol{\Xi}) + C_S \frac{(\eta_u - \eta_\ell)^2 \|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2} \right),$$

for some constant $C_S > 0$ independent of d_z, N . Analogously, we have

$$g(\mathbf{R}_{S,S}) \xrightarrow{\mathbb{P}} \gamma_z \nu_z \left(p_{T \wedge S} \text{tr}(\mathbf{M}) + \frac{2}{\sigma_\xi^2} \text{tr}(\mathbf{M}' \boldsymbol{\Xi}^\top \boldsymbol{\Xi}) + C_S \frac{(\eta_u - \eta_\ell)^2 \|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2} \right),$$

$$g((\mathbf{R}_{S,S,T})) \xrightarrow{\mathbb{P}} \gamma_z \nu_z \left(p_{T \wedge S} \text{tr}(\mathbf{M}) + \frac{2}{\sigma_\xi^2} \text{tr}(\mathbf{M}' \boldsymbol{\Xi}^\top \boldsymbol{\Xi}) + C_S \frac{(\eta_u - \eta_\ell)^2 \|\boldsymbol{\mu}_T\|_2^2}{\sigma_\xi^2} \right)$$

$$g((\mathbf{R}_{S,T,S})) = g((\mathbf{R}_{S,S,T})).$$

1350 Overall, at the proportional asymptotic limit,

$$\begin{aligned}
 & \frac{1}{n} \text{tr} \left(\mathbb{E}_{\mathcal{S}_x} \left[\widehat{\Sigma}_{S,T,N}^\top \widehat{\Sigma}_{S,N}^{-1} \boldsymbol{\Sigma}_{\phi_S, \eta_t} \widehat{\Sigma}_{S,N}^{-1} \widehat{\Sigma}_{S,T,N} \right] \mathbb{E}_{\widetilde{\mathcal{S}}} \left[\widetilde{\Sigma}_{T,n}^{-1} \right] \right) \\
 & \xrightarrow{\mathbb{P}} \gamma_z \text{tr} \left(\mathbf{C}_{T,S}(\eta_t, \eta_u) \mathbf{C}_T(\eta_t)^{-1} \right) + \gamma_z \nu_z (p_T - p_{T \wedge S}) \text{tr} \left(\mathbf{C}_S(\eta_t) \mathbf{C}_S(\eta_u)^{-1} \right).
 \end{aligned}$$

□

E ADDITIONAL EXPERIMENTAL DETAILS

E.1 DATASET STATISTICS

In this work, we construct three distinct splits for each of the four datasets, Waterbirds (Sagawa et al., 2020), BFFHQ (Lee et al., 2021), ImageNet-9 (Xiao et al., 2020), and BG-COCO. Specifically, each dataset is partitioned into a group-imbalanced training set \mathcal{D}_1 , a group-balanced training set \mathcal{D}_2 , and a group-balanced test set \mathcal{D}_3 . The minority group proportion in \mathcal{D}_1 , \mathcal{D}_2 , and \mathcal{D}_3 is η_o , 0.5, and 0.5, respectively. Across different real-world experiments in the paper, we vary the group proportions (η_e and η_u) as well as the sample sizes (n and N). We first summarize the dataset statistics for each benchmark, and then describe how \mathcal{D}_1 to \mathcal{D}_3 are utilized in different experimental setups.

Waterbirds statistics. The Waterbirds (Sagawa et al., 2020) dataset is designed to capture spurious correlations between natural backgrounds and bird labels, with $\eta_o = 0.05$. Table 3 reports the detailed group distributions across \mathcal{D}_1 to \mathcal{D}_3 . Following Sagawa et al. (2020), we supplement additional samples for the minority groups (waterbird, land) and (landbird, water) in the same manner as the original dataset, due to the limited size of the raw data.

Split	(waterbird, water)	(waterbird, land)	(landbird, water)	(landbird, land)	Total
\mathcal{D}_1	1,057	56	184	3,498	4,795
\mathcal{D}_2	1,804	1,804	1,804	1,804	7,216
\mathcal{D}_3	451	451	451	451	1,804

Table 3: Dataset statistics for Waterbirds. Each column corresponds to a group, and the last column gives the total sample count.

BFFHQ statistics. The BFFHQ (Lee et al., 2021) dataset is designed to capture spurious correlations between age and gender labels, with $\eta_o = 0.005$. Table 4 reports the detailed group distributions across \mathcal{D}_1 to \mathcal{D}_3 . Due to the limited size of the minority groups in the raw data, our splits are constructed from de-duplicated samples across multiple BFFHQ subsets.

Split	(young, female)	(young, male)	(old, female)	(old, male)	Total
\mathcal{D}_1	9,552	48	48	9,552	19,200
\mathcal{D}_2	790	790	790	790	3,160
\mathcal{D}_3	198	198	198	198	792

Table 4: Dataset statistics for BFFHQ. Each column corresponds to a group, and the last column gives the total sample count.

ImageNet-9 statistics. The ImageNet-9 (Xiao et al., 2020) dataset is designed to capture spurious correlations between object and background labels. Different from Waterbirds and BFFHQ, ImageNet-9 is a 9-class classification task over categories dog, bird, wheeled vehicle, reptile, carnivore, insect, musical instrument, primate, and fish. The original dataset provides two variants, mixed-same and mixed-rand. In the mixed-same version, each image background is replaced with a background from an image of the same class, thus preserving spurious correlations; in the mixed-rand version, the background is randomized and contains no information about the true label. These two variants correspond to minority group proportions of 0 and 0.5, respectively. Table 5 reports the dataset statistics across \mathcal{D}_1 to \mathcal{D}_3 . Based on this table, we set $\eta_o = 0$. Note that ImageNet-9 does not have a well-defined group structure under either the mixed-same or mixed-rand settings. Therefore, we do not report worst-group accuracy for this dataset.

	Split	mixed-same	mixed-rand	Total	Per-class
1404	\mathcal{D}_1	4,050	0	4,050	450
1405	\mathcal{D}_2	0	3,240	3,240	360
1406	\mathcal{D}_3	0	810	810	90
1407					
1408					

Table 5: Dataset statistics for ImageNet-9. Within each split, the nine classes have identical counts.

BG-COCO statistics. The BG-COCO dataset is a self-generated benchmark designed to capture spurious correlations between cats/dogs from COCO (Lin et al., 2014) and indoor/outdoor scenes from Places (Zhou et al., 2017). Specifically, we define the indoor/outdoor split as living room, dining room (indoor) and park (outdoor). By construction, cats are aligned with indoor scenes and dogs with outdoor scenes. Table 6 reports the detailed group distributions across \mathcal{D}_1 to \mathcal{D}_3 . Based on this table, we set $\eta_o = 0.05$.

	Split	(cat, indoor)	(cat, outdoor)	(dog, indoor)	(dog, outdoor)	Total
1419	\mathcal{D}_1	1,900	100	100	1,900	4,000
1420	\mathcal{D}_2	1,000	1,000	1,000	1,000	4,000
1421	\mathcal{D}_3	250	250	250	250	1,000
1422						
1423						

Table 6: Dataset statistics for BG-COCO. Each column corresponds to a group, and the last column gives the total sample count.

Across all four datasets, we construct training and evaluation splits as follows. When either η_ℓ or η_u is fixed η_o , samples are drawn from \mathcal{D}_1 with the desired size N (for unlabeled data) or n (for labeled data). When either η_ℓ or η_u is fixed to 0.5, balanced samples are instead drawn from \mathcal{D}_2 . Several experiments involve fixing η_ℓ while varying η_u . In this setting, if necessary, we keep the labeled data unchanged and supplement the unlabeled data with additional samples independently drawn from \mathcal{D}_1 or \mathcal{D}_2 , while ensuring that the total unlabeled sample size N remains constant across different η_u . The balanced dataset \mathcal{D}_3 is reserved for testing, and when required, we further split 20% of \mathcal{D}_3 as a separate validation set.

E.2 RESULTS FOR INTERPRETING W2S UNDER SPURIOUS CORRELATIONS

Dataset	# of model pairs with increased W2S gain	
	Average accuracy	Worst group accuracy
Waterbirds	10/10	10/10
BFFHQ	10/10	9/10
BG-COCO	9/10	10/10
ImageNet-9	7/10	—

Table 7: Proportion of teacher-student pairs that exhibit an increase in W2S gain as η_u increases from 0 to the maximum feasible value of η_u (Waterbirds: 0.5, BFFHQ: 0.23, BG-COCO: 0.5, ImageNet-9: 0.4) when $\eta_\ell = 0.5$, summarized across all datasets. ImageNet-9 has no well-defined worst group, so only average accuracy is reported.

In Section 3.2, we primarily presented how the average W2S gain across all teacher–student pairs varies with increasing η_u on each dataset. Here, we further provide results for individual model pairs. Specifically, Figure 6 compares the difference in W2S gain between the group-balanced ($\eta_\ell = 0.5$) and group-imbalanced ($\eta_\ell = \eta_o$) settings on selected datasets. Table 7 summarizes, for $\eta_\ell = 0.5$, the proportion of model pairs that exhibit an increase in W2S gain as η_u increases from 0 across all datasets. Table 8 summarizes, for $\eta_\ell = \eta_o$, the proportion of model pairs that exhibit a decrease in W2S gain as η_u increases from η_o across all datasets. These results further validate our theoretical analysis in Section 2.2, which predicts that in most cases the larger the gap between η_u and η_ℓ , the smaller the resulting W2S gain.

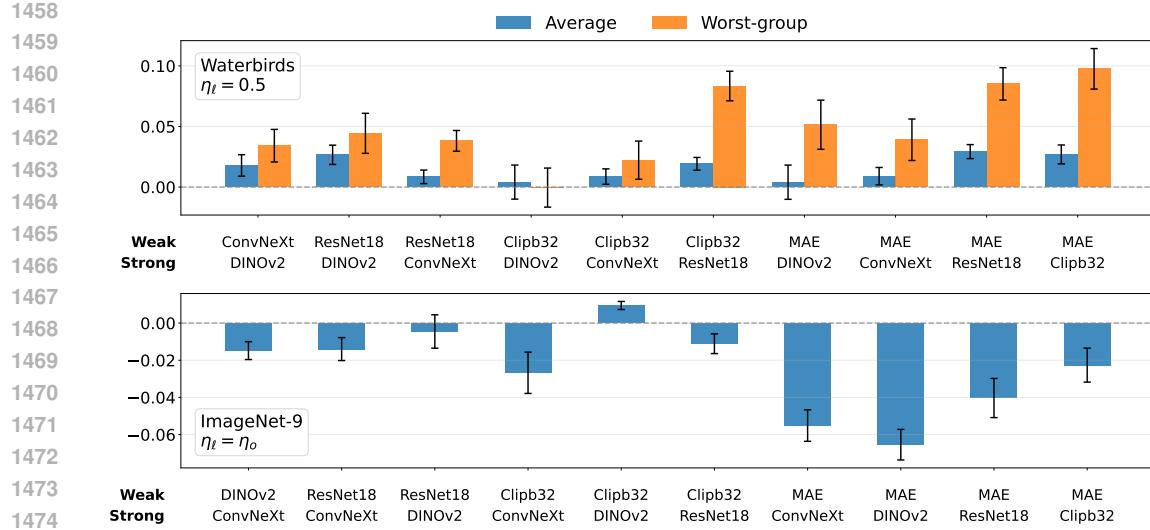


Figure 6: Top: On the Waterbirds dataset, the change in W2S gain (value at $\eta_u = 0.5$ minus value at $\eta_u = 0$) across all teacher-student pairs with fixed $\eta_\ell = 0.5$. Bottom: On the ImageNet-9 dataset, the change in W2S gain (value at $\eta_u = 0.5$ minus value at $\eta_u = \eta_o$) across all teacher-student pairs with fixed $\eta_\ell = \eta_o$. ImageNet-9 does not have a clearly defined worst group and is therefore omitted from the bottom panel.

Dataset	# of model pairs with decreased W2S gain	
	Average accuracy	Worst group accuracy
Waterbirds	7/10	8/10
BFFHQ	8/10	7/10
BG-COCO	8/10	7/10
ImageNet-9	9/10	—

Table 8: Proportion of teacher-student pairs that exhibit a decrease in W2S gain as η_u increases from η_o to 0.5 when $\eta_\ell = \eta_o$, summarized across all datasets. ImageNet-9 has no well-defined worst group, so only average accuracy is reported.

E.3 RESULTS FOR ENHANCED W2S

Model training. Enhanced-W2S improves upon vanilla W2S by retraining the strong student after the initial W2S fine-tuning. First, we select a fraction $p \in (0, 1]$ of $\hat{\mathcal{S}}$ consisting of those samples for which the student exhibits the lowest prediction entropy. Second, we apply the GCE loss $\mathcal{L}_{\text{GCE}}(\mathbf{x}_i, \hat{y}_i; q)$ with parameter $q \in (0, 1]$ to each selected sample $(\mathbf{x}_i, \hat{y}_i)$. We tune the hyperparameters by grid search over $p \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ and $q \in \{0, 0.2, 0.7\}$, where $q = 0$ corresponds to the CE loss (i.e., the $q \rightarrow 0$ limit of GCE). To avoid a trivial overlap with the vanilla W2S baseline, $(p, q) = (1, 0)$ is excluded from the Enhanced-W2S search space. In the case of $(\eta_\ell, \eta_u) = (\eta_o, 0.5)$, we further restrict the subset ratio to $p \in \{0.2, 0.4, 0.6\}$ to emphasize the role of high-confidence subsets in filtering for the majority group. Each run of Enhanced-W2S is repeated with multiple random seeds, and the reported results are obtained by averaging across seeds.

Role of confidence-based selection. When $(\eta_\ell, \eta_u) = (\eta_o, 0.5)$, Figure 7 shows that samples with high student confidence (i.e., low predictive entropy) after W2S fine-tuning are almost exclusively drawn from the majority group, and furthermore are nearly always assigned the correct pseudolabels by both the weak teacher and the strong student. At the same time, Theorem 2 predicts that reducing η_u from 0.5 directly increases the W2S gain. These two observations together suggest that confidence-based selection provides significant benefits for improving W2S performance in the setting $(\eta_\ell, \eta_u) = (\eta_o, 0.5)$.

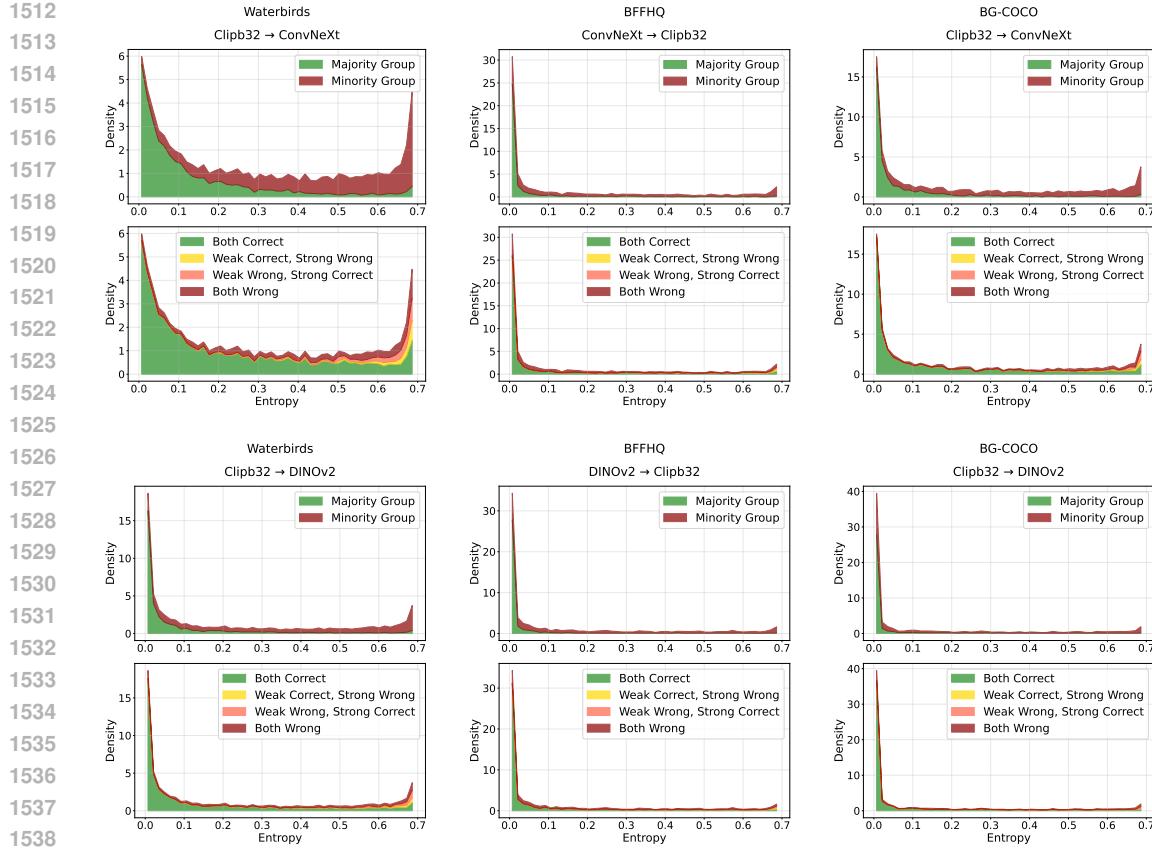


Figure 7: Student confidence on unlabeled data as stacked density plots of predictive entropy ($\eta_\ell = \eta_o, \eta_u = 0.5$). Each panel shows the student’s predictive entropy (after W2S fine-tuning), visualized as two stacked density plots: (top) split by group (majority vs. minority) and (bottom) split by prediction correctness of the weak teacher and the strong student. Columns correspond to datasets (Waterbirds, BFFHQ, BG-COCO). Rows correspond to model pairs: (ConvNeXt, Clipb32) and (Clipb32, DINOv2).

Dataset	η_ℓ	η_u	Mean relative improvement (%)	
			Average Accuracy	Worst Group Accuracy
Waterbirds	0.5	η_o	6.32	10.12
	η_o	0.5	7.72	32.15
BFFHQ	0.5	η_o	5.52	4.06
	η_o	0.5	3.12	3.57
BG-COCO	0.5	η_o	10.51	11.76
	η_o	0.5	4.50	3.71
ImageNet-9	0.5	η_o	12.23	—
	η_o	0.5	11.23	—

Table 9: Mean relative improvement (%) of Enhanced-W2S over vanilla W2S, averaged across selected teacher–student pairs, for both average accuracy and worst group accuracy. For each dataset, we select all model pairs whose relative strength relationship remains consistent across different (η_ℓ, η_u) settings.

Mean relative gains. Table 9 summarizes the mean relative improvement of Enhanced-W2S over vanilla W2S, averaged across all teacher–student pairs. Consistent with the main text, our method achieves clear gains under both average accuracy and worst group accuracy. On the Waterbirds dataset,

we further compare the performance of Enhanced-W2S with the auxiliary confidence loss proposed in (Burns et al., 2024), which was also designed to improve the generalization ability of W2S. Specifically, we perform a grid search over the auxiliary confidence loss weight $\alpha \in \{0.2, 0.4, 0.6, 0.8\}$, and Table 10 reports, for each (η_ℓ, η_u) configuration, the mean relative improvement of Enhanced-W2S minus the mean relative improvement obtained with the auxiliary confidence loss. Our method yields larger gains in all cases, confirming that it is motivated by our theoretical analysis (see Section 4) and is specifically tailored to address W2S under spurious correlation.

η_ℓ	η_u	Difference in mean relative improvement (%)	
		Average Accuracy	Worst Group Accuracy
0.5	η_o	5.22	3.88
η_o	0.5	5.90	3.41

Table 10: Difference in mean relative improvement (%) Waterbirds, computed as Enhanced-W2S minus the auxiliary confidence loss baseline, averaged across selected teacher–student pairs for each (η_ℓ, η_u) configuration.

F GROUP FAIRNESS IN W2S GENERALIZATION

Ensuring that algorithmic decisions do not exhibit systematic bias against certain attributes (e.g., race, gender, age) has long been a central objective in fair machine learning (Liu et al., 2019; Oneto & Chiappa, 2020; Mehrabi et al., 2021). At the same time, when the data contains spurious correlations caused by group imbalance, unfairness across different groups is likely to arise, as the model tends to rely on spurious features when making predictions (Izmailov et al., 2022). This lack of group fairness is particularly concerning when groups are defined by sensitive attributes. Therefore, a line of work on mitigating spurious correlation has explicitly targeted group robustness, and the evaluation metrics adopted in this literature (e.g., worst group accuracy) can be interpreted as a measure of fairness. In parallel, several works have more directly studied the relationship between spurious correlations and formal notions of group fairness (Veitch et al., 2021; Schrouff et al., 2024).

In this section, we extend our analysis of W2S under spurious correlation to incorporate the notion of group fairness. Under W2S, the minority group proportions in both the labeled dataset (η_ℓ) and the unlabeled dataset (η_u) jointly influence the extent to which the strong student preserves group-level parity after the W2S process.

Definition 4 (Group risk disparity). *Under Definitions 1 and 2, we define the group risk disparity of the strong student after W2S fine-tuning as*

$$\Delta_{\text{grp}}(f_S) := \left| \mathbb{E}_{\mathcal{D}(\eta_u)^N, \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_0(f_S)] - \mathbb{E}_{\mathcal{D}(\eta_u)^N, \mathcal{D}(\eta_\ell)^n} [\mathbf{ER}_1(f_S)] \right|, \quad (33)$$

where $\mathbf{ER}_0(f_S)$ and $\mathbf{ER}_1(f_S)$ denote the excess risks of the student on the majority ($\eta_t = 0$) and minority ($\eta_t = 1$) groups, respectively.

In Definition 4, we quantify the group fairness through the absolute difference between the student’s excess risk on the majority group and the minority group. It is important to note that our definition of group risk disparity is directly aligned with the notion of perfect fairness (also referred to as risk parity) in the group-fairness literature (Williamson & Menon, 2019; Liu et al., 2025b). In particular, the condition $\Delta_{\text{grp}}(f_S) = 0$ is equivalent to achieving perfect fairness (risk parity).

Corollary 1 (Group risk disparity of W2S). *Under Definitions 1 and 2 and assumption 1, the group risk disparity of the strong student after W2S fine-tuning satisfies*

$$\Delta_{\text{grp}}(f_S) \xrightarrow{\mathbb{P}} \frac{\sigma_y^2 \gamma_z}{\sigma_\xi^2} \left| 2(\eta_\ell - \eta_u) \boldsymbol{\mu}_T^\top \boldsymbol{\Xi} \boldsymbol{\mu}_S - (1 - 2\eta_u) \left(\|\boldsymbol{\Xi} \boldsymbol{\mu}_S\|_2^2 + \nu_z (p_T - p_{T \wedge S}) \|\boldsymbol{\mu}_S\|_2^2 \right) \right|$$

Corollary 1 follows directly from the precise asymptotic characterization of the strong student excess risk in Theorem 3, providing a precise quantification of the group risk disparity in the proportional asymptotic limit.

We outline several key insights from Corollary 1 below:

- (a) **Low teacher-student similarity ($p_{T \wedge S} = 1$) brings robustness of group fairness $\Delta_{\text{grp}}(f_S)$ to teacher bias $\eta_\ell < 0.5$, where W2S is fair if the unlabeled training set is balanced $\eta_u = 0.5$.**

1620 Notably, when $p_{T \wedge S} = 1$ (i.e., $\|\Xi\|_F^2 = 0$), $\Delta_{\text{grp}}(f_S)$ becomes independent of η_ℓ and is only
 1621 affected by η_u through the $(1 - 2\eta_u)$ factor. When $\eta_u = 0.5$ further, we have $1 - 2\eta_u = 0$, and
 1622 therefore $\Delta_{\text{grp}}(f_S) \xrightarrow{\mathbb{P}} 0$, i.e., the strong student from W2S fine-tuning is fair, even though a
 1623 biased weak teacher fine-tuned with $\eta_\ell < \eta_u$ can still hurt the generalization of the strong student.
 1624

1625 (b) **Low teacher-student similarity ($p_{T \wedge S} = 1$) induces group fairness of W2S,** $\Delta_{\text{grp}}(f_S) \rightarrow 0$,
 1626 as $\nu_z \rightarrow 0$. While $p_{T \wedge S} = 1$ (i.e., $\|\Xi\|_F^2 = 0$) alone does not guarantee fairness, it provides

$$1627 \Delta_{\text{grp}} \xrightarrow{\mathbb{P}} (1 - 2\eta_u)\gamma_z\nu_z \frac{\sigma_y^2}{\sigma_\xi^2} (p_T - p_{T \wedge S}) \|\boldsymbol{\mu}_S\|_2^2 \asymp (1 - 2\eta_u)\gamma_z\nu_z,$$

1630 where we have $\Delta_{\text{grp}}(f_S) \xrightarrow{\mathbb{P}} 0$ as $d_z, n, N \rightarrow \infty$ if $\nu_z \rightarrow 0$, i.e., when N is large enough
 1631 compared to d_z , low teacher-student similarity induces fairness of W2S.
 1632

1633 (c) **For high teacher-student similarity ($p_{T \wedge S} \rightarrow p_S$), group fairness of the student** $\Delta_{\text{grp}}(f_S)$ **is**
 1634 **influenced by the fairness of the teacher** η_ℓ , **with the dependence determined by** $\boldsymbol{\mu}_T^\top \Xi \boldsymbol{\mu}_S$. In
 1635 particular, when $p_{T \wedge S} \rightarrow p_S$ so that $|\boldsymbol{\mu}_T^\top \Xi \boldsymbol{\mu}_S|$ is non-negligible, W2S is fair (i.e., $\Delta_{\text{grp}}(f_S) \xrightarrow{\mathbb{P}} 0$)
 1636 when

$$1637 \eta_\ell^{\text{fair}} = \eta_u + (1 - 2\eta_u) \frac{\|\Xi\boldsymbol{\mu}_S\|_2^2 + \nu_z(p_T - p_{T \wedge S}) \|\boldsymbol{\mu}_S\|_2^2}{\boldsymbol{\mu}_T^\top \Xi \boldsymbol{\mu}_S},$$

1640 assuming it falls in the range $\eta_\ell^{\text{fair}} \in [0, 0.5]$, while the group fairness gets worse (i.e., $\Delta_{\text{grp}}(f_S)$
 1641 increases) as η_ℓ deviates from η_ℓ^{fair} . Notably,

- 1642 (1) if $\eta_\ell^{\text{fair}} < 0$, the group fairness gets worse as η_ℓ increases, best when $\eta_\ell = 0$; while
- 1643 (2) if $\eta_\ell^{\text{fair}} > 0.5$ the group fairness gets worse as η_ℓ decreases, best when $\eta_\ell = 0.5$.

1644
 1645
 1646
 1647
 1648
 1649
 1650
 1651
 1652
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660
 1661
 1662
 1663
 1664
 1665
 1666
 1667
 1668
 1669
 1670
 1671
 1672
 1673