Educational losses and damages attributable to anthropogenic climate change: End-to-end attribution Evidence from South Africa

EGU26 Presentation

Misheck Tussle Mundowa, Jacobus Cilliers, Christopher Trisos, Nicholas Simpson, Harald Winkler, Stephen Taylor

Presentation Supplementary Material

S1. Data

The primary education dataset comprises approximately ~8.4 million learner-level records from the National Senior Certificate (NSC) examinations across ~9000 high schools, covering the period 2008 to 2024, obtained from the South African Department of Basic Education. All records are retained at the individual learner level. Learners are matched to schools using the Education Management Information System (EMIS) unique school identifier. The outcome variable is attainment of a bachelor pass — the threshold qualification required for university enrolment in South Africa.

Temperature data are drawn from ERA5-Land (0.1° resolution) which serves as the primary dataset for the main independent variable. We controlled for precipitation, and used the CHIRPS dataset. Our preferred counterfactual dataset is from the pattern scaling approach from the ERA5-land which we are currently processing at the time of the EGU26 GA. In the interim, the W5E5 dataset (GSWP3-W5E5 obsclim, ISIMIP3a; 0.5° resolution) is used for the impact attribution analysis to ensure dataset consistency with the ATTRICI v1.1 counterfactual climate product. Each school is linked to the nearest ERA5-Land or W5E5 grid cell centroid. Cumulative pre-exam rainfall is sourced from CHIRPS v2.0. We have geolocations for each school allowing us to merge the learner and climate data at school level.

The pre-exam exposure window is defined as all official school days from 1 January to 15 October of the examination year, constructed from the official South African school calendar published annually by the Department of Basic Education. All temperature variables are computed over this window only, capturing the thermal conditions to which learners are exposed during the school year leading to their final examinations. The Tx5 variable is defined as the mean temperature of the five hottest consecutive pre-exam school days within this window, capturing the extreme heat intensity.

S2. Econometric Model and Identification Strategy

Our primary specification follows the Callahan and Mankin (2022) framework adapted for education outcomes and is estimated as a linear probability model:

$$Y_{ist} = \alpha_1 T_{st} + \alpha_2 T_{st}^2 + \beta_1 Tx5_{st} + b_2 T_{st}.Tx5_{st} + \boldsymbol{\gamma}'\boldsymbol{X_{st}} + \mu_s + \lambda_t + \varepsilon_{ist}$$

where Y_{ist} is a binary indicator equal to 1 if learner *i* at school *s* in year t attained a bachelor pass; T_{st} is the mean daily maximum temperature over pre-exam school days; $Tx5_{st}$ is the mean temperature of the five hottest consecutive pre-exam school days; $\boldsymbol{X_{st}}$ is a vector of controls including cumulative pre-exam rainfall (divided by 100); individual-level controls (gender, age); and μ_s and λ_t are school and year fixed effects respectively. Standard errors are clustered at the school level.

The interaction term $T_{st}\#Tx5_{st}$ is central to the specification. It captures the key insight that the damage from an extreme heat event (measured by Tx5) depends on the ambient heat to which learners are adapted: the marginal effect of Tx5 is ${\partial Y}/{\partial Tx5} = \beta_1 + \beta_2 T$, which is positive at low mean temperatures and turns harmful above a zero-crossing temperature determined by the ratio $-\beta_1/\beta_2$.

Our identification relies on within-school, across-year variation in temperature after controlling for school and year fixed effects. The identifying assumption is that year-to-year temperature variations within a school are plausibly random and exogenous to other determinants of exam performance (Garg et al., 2020; Park et al., 2021; Zivin et al., 2020). The potential threats to the this identification is that learners may choose schools based on temperature-sensitive characteristics and school-level adaptation responses to hot years. School fixed effects absorb all time-invariant confounders, including chronic local climate, infrastructure quality, and socioeconomic context; year fixed effects remove national-level shocks common to all schools. Given the granularity of the fixed effects and the inherently unpredictable nature of within-school temperature variations, systematic bias from these threats is unlikely.

S3. Impact Model Results: ERA5-Land (2008–2024)

Table S1 presents results from the ERA5-Land specification across five model specifications, progressively adding controls. Column (3) is the preferred model, combining the full T + T^2 + Tx5 + T×Tx5 interaction structure with demographic controls and rainfall. The Tx5 coefficient is positive, reflecting a positive effect of heat extreme at low baseline temperatures (0°C); the negative interaction term implies this positive impact diminishes and reverses at higher mean annual temperatures. The zero-crossing, above which an additional degree of Tx5 is harmful, is approximately 25.7°C (Figure S1, Panel a).

Table S1. Effect of Temperature on Pr(Bachelor Pass): ERA5-Land, 1-Year Window (2008–2024)

	(1)	(2)	(3)*	(4)	(5)
Mean pre-exam temperature, T (°C)	-0.0278***	-0.0250***	-0.0280***	-0.0277***	-0.0072
	(0.0084)	(0.0082)	(0.0085)	(0.0085)	(0.0084)
T^2	0.0014***	0.0013***	0.0013***	0.0013***	0.0001
	(0.0003)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Hottest 5-day mean, Tx5 (°C)	0.0385***	0.0342***	0.0330***	0.0325***	
	(0.0049)	(0.0047)	(0.0047)	(0.0047)	
T × Tx5	-0.0015***	-0.0013***	-0.0013***	-0.0013***	
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	
Female		-0.0380***	-0.0380***	-0.0380***	-0.0380***
		(0.0014)	(0.0014)	(0.0014)	(0.0014)
Age (years)		-0.1182***	-0.1182***	-0.1182***	-0.1182***
		(0.0006)	(0.0006)	(0.0006)	(0.0006)
Cumulative rainfall (÷100, mm)			-0.0012**	-0.0013**	-0.0018***
			(0.0006)	(0.0006)	(0.0006)
SD pre-exam temperature (°C)				-0.0017	
				(0.0017)	
Observations	8,420,833	8,069,247	8,069,247	8,069,247	8,069,247
R-squared	0.167	0.211	0.211	0.211	0.211
Outcome mean	0.360	0.360	0.360	0.360	0.360
Fixed effects and controls					
School FE	X	X	X	X	X
Year FE	X	X	X	X	X
Clustered SE (school)	X	X	X	X	X
Interaction (T×Tx5)	X	X	X	X	
Demographics		X	X	X	X
Rainfall			X	X	X
Temperature SD				X	

Notes: Linear probability model. Temperature from ERA5-Land (0.1° nearest grid cell). 1-year exposure window.
** Col (3) = preferred interaction model. Cols (1)(2)(4)(5) = sensitivity specifications.*
SE clustered at school level. Tx5 = mean of 5 hottest consecutive pre-exam school days.
** p<0.10 ** p<0.05 *** p<0.01*

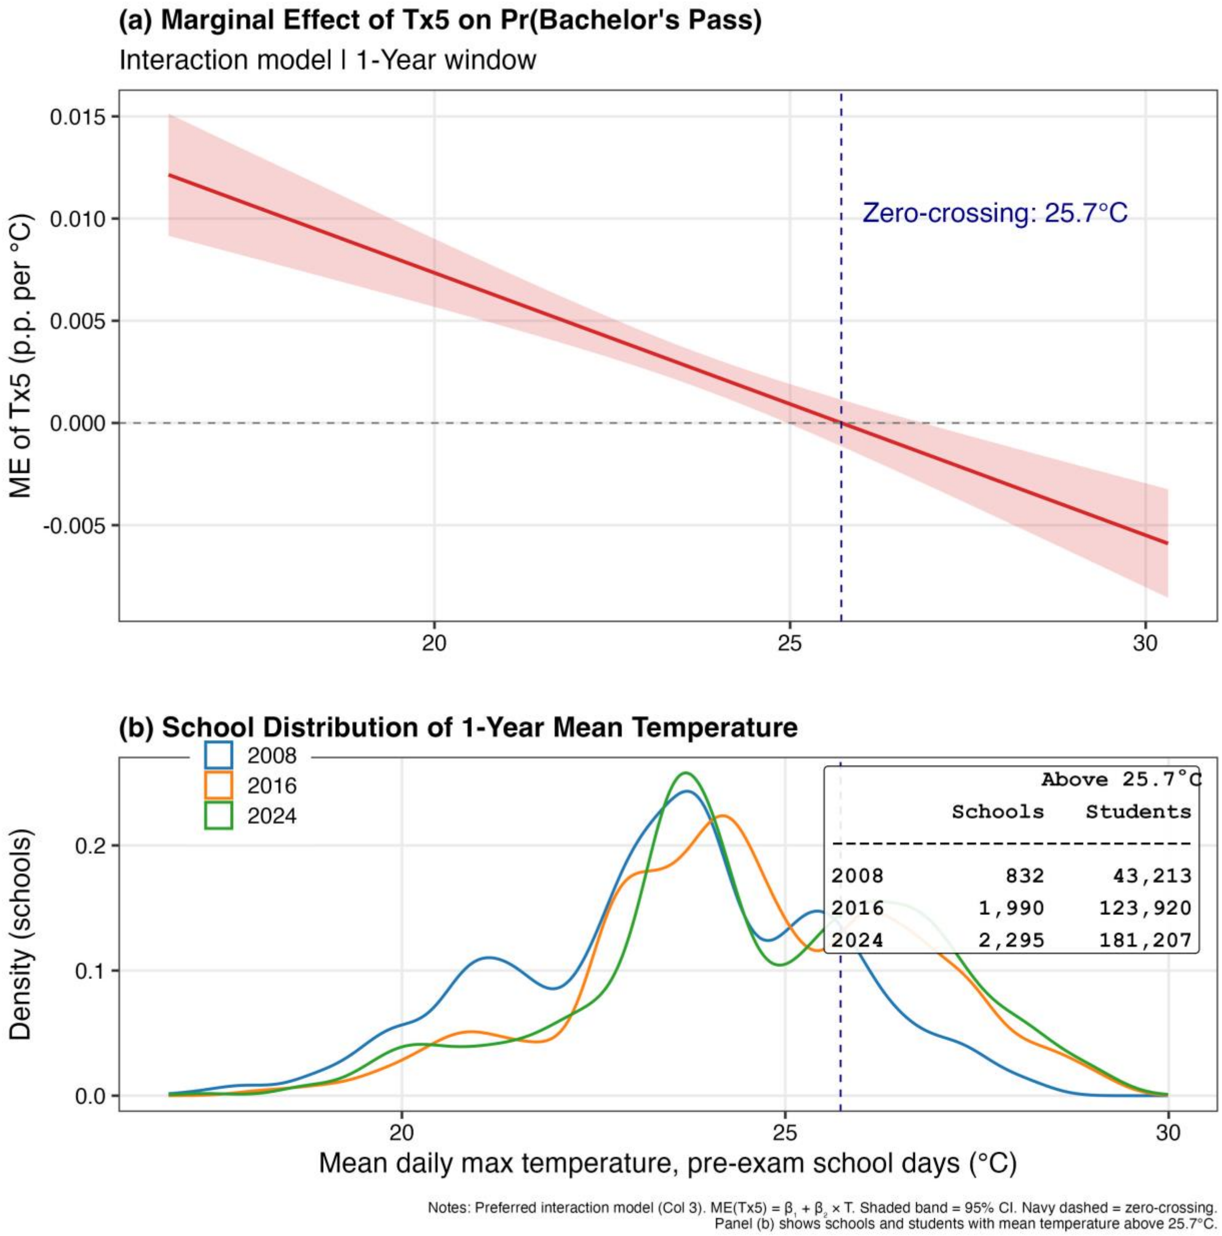

Figure S1. Marginal effect of Tx5 on Pr(Bachelor Pass) by mean pre-exam temperature (ERA5-Land), and distribution of school-level mean temperatures across selected years. Zero-crossing at 25.7°C.

S4. Preliminary Attribution Analysis: Results (W5E5/ATTRICI, 2008–2019)

To attribute examination outcomes (bachelor passes) to anthropogenic climate change, the impact model is re-estimated using W5E5 temperature data (sample: 2008–2019) to ensure dataset consistency with the ATTRICI v1.1 counterfactual, which simulates temperatures under a world without anthropogenic greenhouse gas forcing. The re-estimated W5E5 damage function (Table S2, Figure S2) is qualitatively similar to the ERA5-Land results, with a zero-crossing at 22.2°C. The lower threshold reflects the warmer temperature distribution implied by W5E5 relative to ERA5-Land at South African school locations (compare Panel (b) of Figures S1 and S2). These results should be interpreted as preliminary: the planned final analysis will derive the counterfactual directly from ERA5-Land using the Callahan and Mankin (2025) pattern-scaling approach, maintaining full consistency with the primary impact model.

Table S2. Effect of Temperature on Pr(Bachelor Pass): W5E5 Factual Climate, 1-Year Window (2008–2019)

	(1)	(2)	(3)*	(4)	(5)
Mean pre-exam temperature, T (°C)	-0.0986***	-0.0846***	-0.0828***	-0.0829***	-0.0835***
	(0.0141)	(0.0137)	(0.0137)	(0.0137)	(0.0139)
T^2	0.0027***	0.0023***	0.0022***	0.0022***	0.0018***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Hottest 5-day mean, Tx5 (°C)	0.0226***	0.0161***	0.0154***	0.0153**	
	(0.0060)	(0.0059)	(0.0060)	(0.0060)	
T × Tx5	-0.0010***	-0.0007***	-0.0007***	-0.0007***	
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	
Female		-0.0400***	-0.0400***	-0.0400***	-0.0400***
		(0.0015)	(0.0015)	(0.0015)	(0.0015)
Age (years)		-0.1059***	-0.1059***	-0.1059***	-0.1059***
		(0.0006)	(0.0006)	(0.0006)	(0.0006)
Cumulative rainfall (÷100, mm)			0.0017***	0.0016***	0.0018***
			(0.0006)	(0.0006)	(0.0006)
SD pre-exam temperature (°C)				-0.0010	
				(0.0018)	
Observations	5,117,076	5,117,076	5,117,076	5,117,076	5,117,076
R-squared	0.189	0.231	0.231	0.231	0.231
Outcome mean	0.317	0.317	0.317	0.317	0.317
Fixed effects and controls					
School FE	X	X	X	X	X
Year FE	X	X	X	X	X
Clustered SE (school)	X	X	X	X	X
Interaction (T×Tx5)	X	X	X	X	
Demographics		X	X	X	X
Rainfall			X	X	X
Temperature SD				X	

*Notes: Linear probability model. Temperature from W5E5 (GSWP3-W5E5 obsclim, ISIMIP3a; 0.5° nearest grid cell). Sample: 2008–2019 (constrained by W5E5/ATTRICI availability). * Col (3) = preferred model for attribution. SE clustered at school level. Tx5 = mean of 5 hottest consecutive pre-exam school days.*
** p<0.10 ** p<0.05 *** p<0.01*

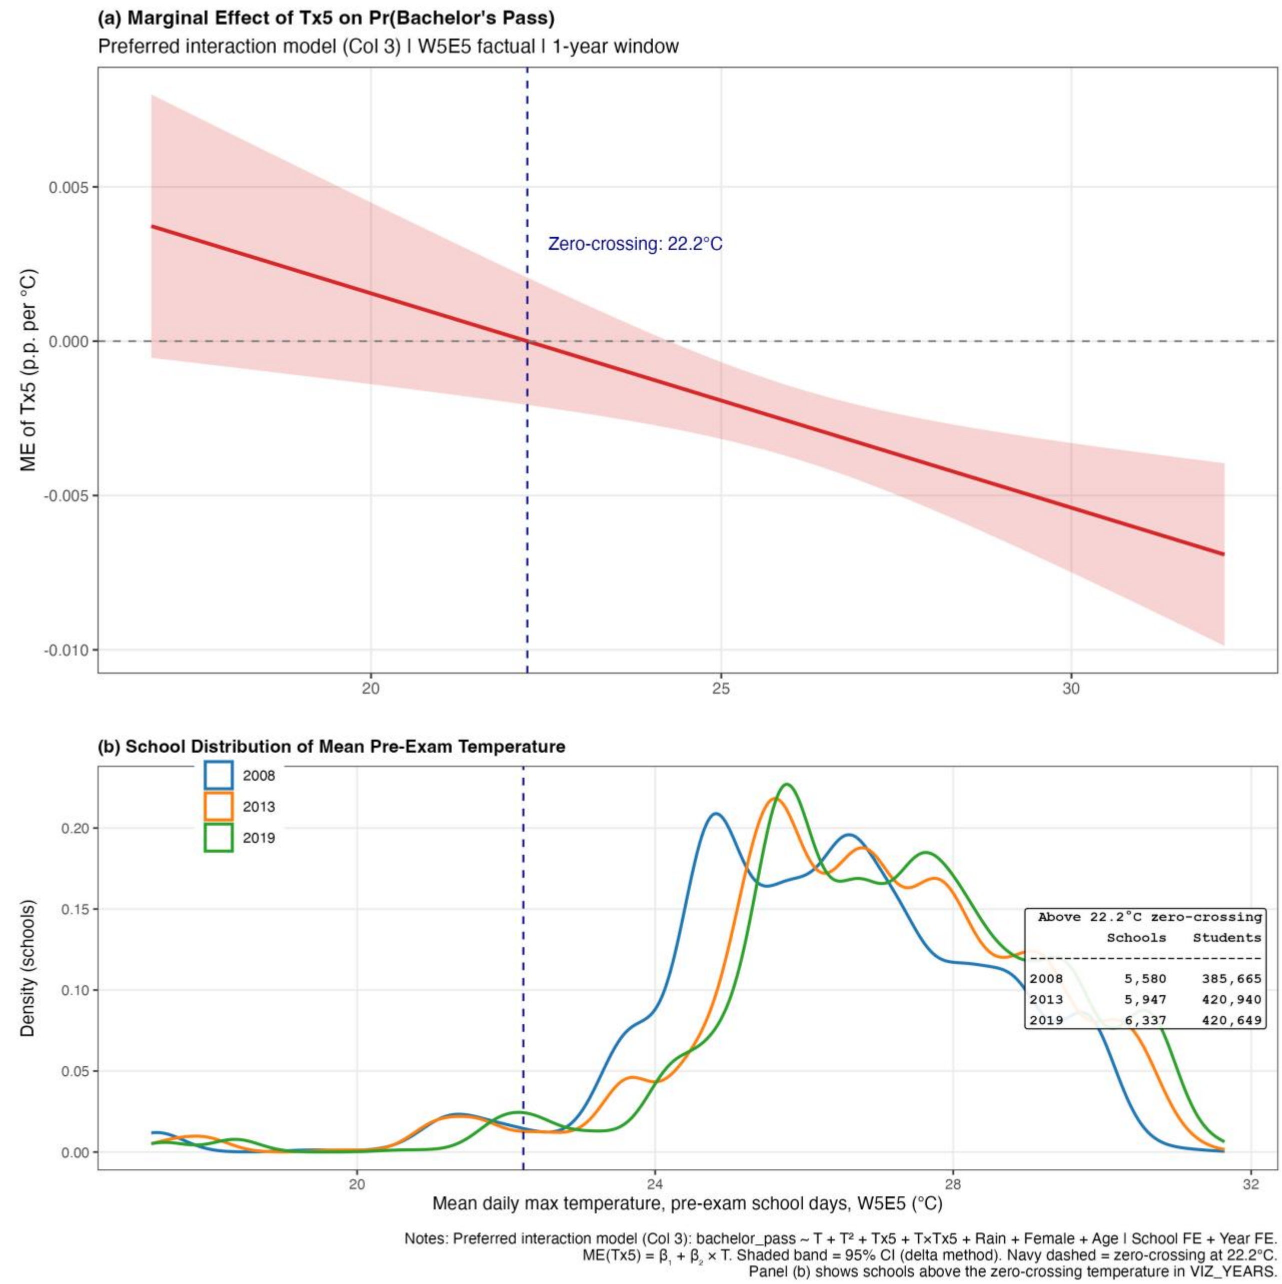

Figure S2. Marginal effect of Tx5 on Pr(Bachelor Pass) by mean pre-exam temperature (W5E5), and distribution of school-level mean temperatures across selected years. Zero-crossing at 22.2°C.

The attribution pipeline applies the estimated W5E5 damage function coefficients to school-level factual (W5E5) and counterfactual (ATTRICI) temperature time series. For each school-year observation, the predicted pass probability is computed under factual and counterfactual temperatures while holding all other covariates at their sample means across the period. The per-student attribution signal is $\Delta\hat{y}_{st} = \hat{y}(T^f, Tx5^f) - \hat{y}(T^{cf}, Tx5^{cf})$, and the national count of additional bachelor passes is $\Delta N_t = \Sigma_s n_{st} \times \Delta\hat{y}_{st}$. Uncertainty intervals are obtained by school-clustered bootstrap (50 replications – but to increase to 1000 when using our preferred counterfactual data) applied to the damage function coefficients.

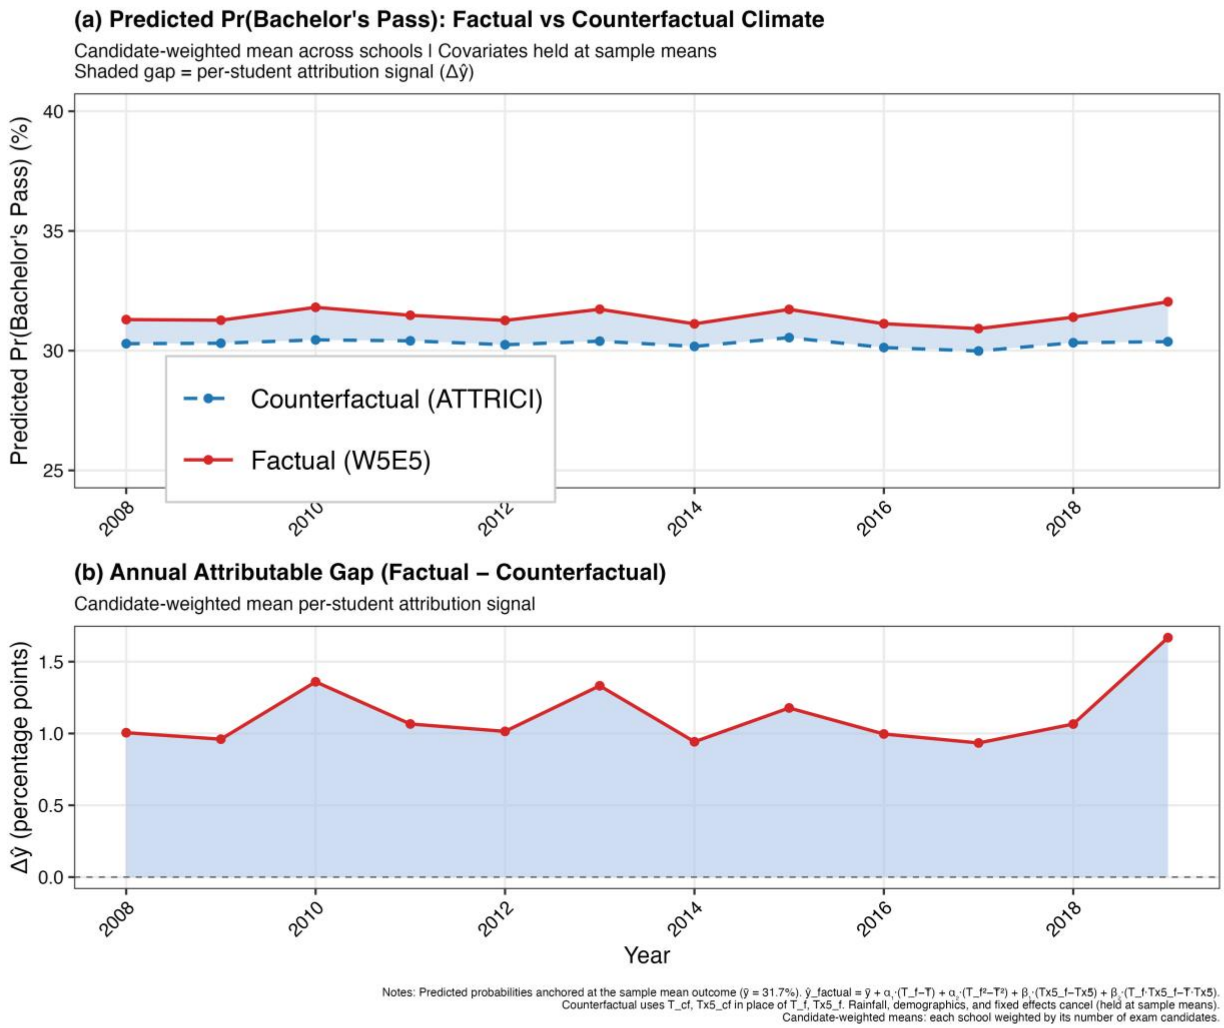

Figure S3. National-level predicted Pr(Bachelor Pass) under factual (W5E5) and counterfactual (ATTRICI) climates (Panel a), and the annual per-student attributable gap Δŷ (Panel b). Candidate-weighted means; covariates held at sample means.

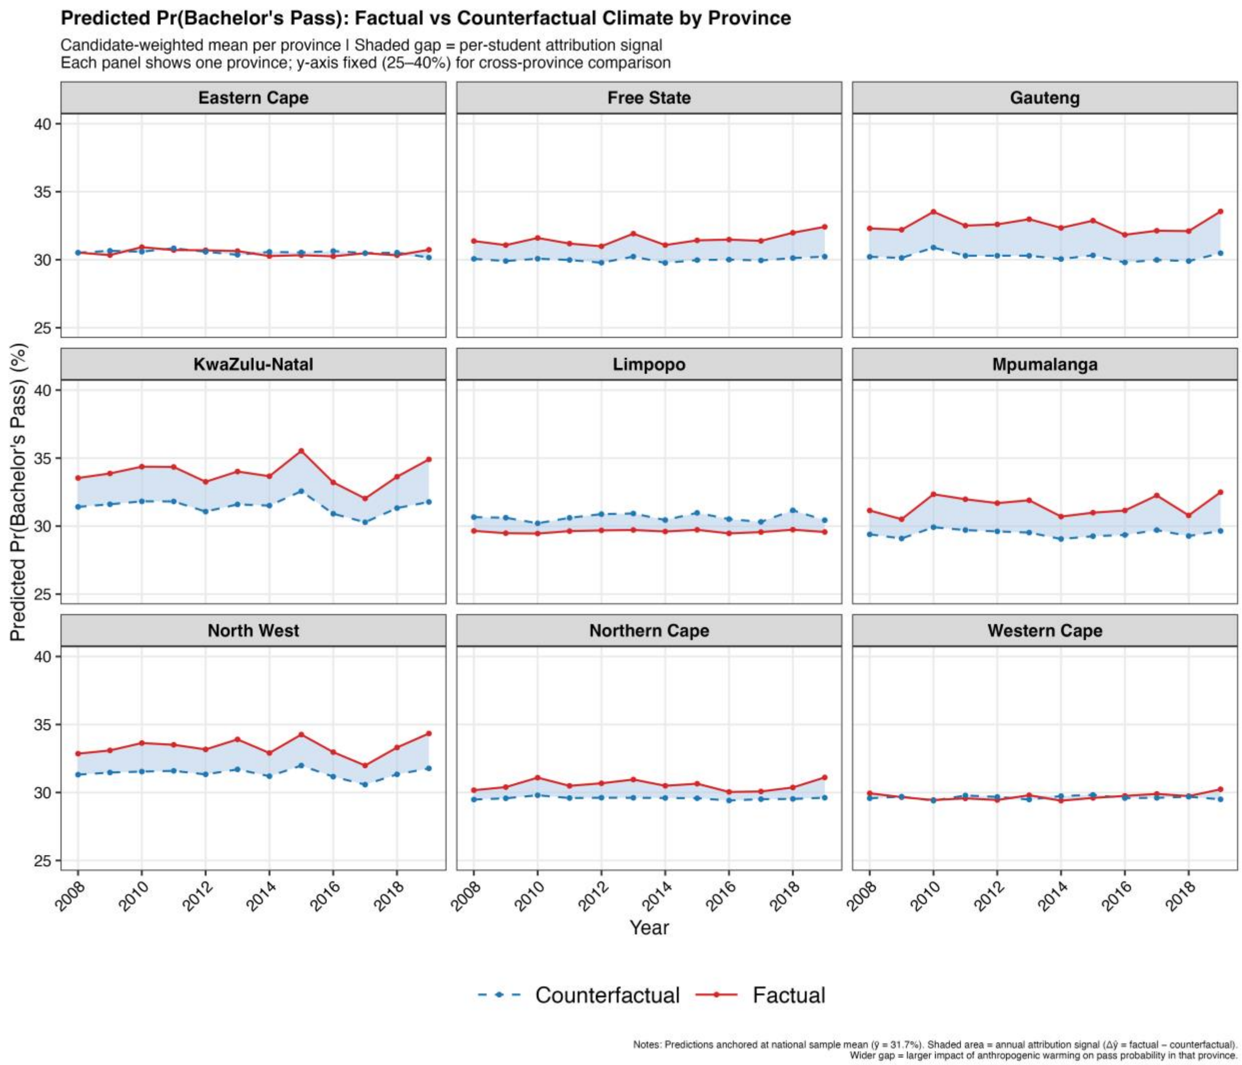

Figure S4. Provincial predicted Pr(Bachelor Pass) under factual and counterfactual climates (2008–2019). The shaded area in each panel represents the province-level annual attribution signal. Y-axis is fixed at 25–40% across provinces to facilitate comparison.

On average, South Africa is benefiting from anthropogenic climate change as the factual bachelor pass rate trend in Figure S3 is consistently above the counterfactual. We disaggregate these attributable impacts to per province level with results shown in Figure S4. One of the poorest province (Limpopo) which already has the least bachelor pass rates is consistently harmed by anthropogenic climate change during the 2008-2019 period. This reflects inequalities in the climate change impacts and energy insulation poverty in South Africa. Eastern Cape and Western Cape have mixed anthropogenic climate impact signals and the rest of the provinces are gaining from anthropogenic climate change. Our interpretation is that the Limpopo province is already showing the future of all provinces if the world continues to warm. We have shown that more and more schools and students are entering the harmful zone every year.

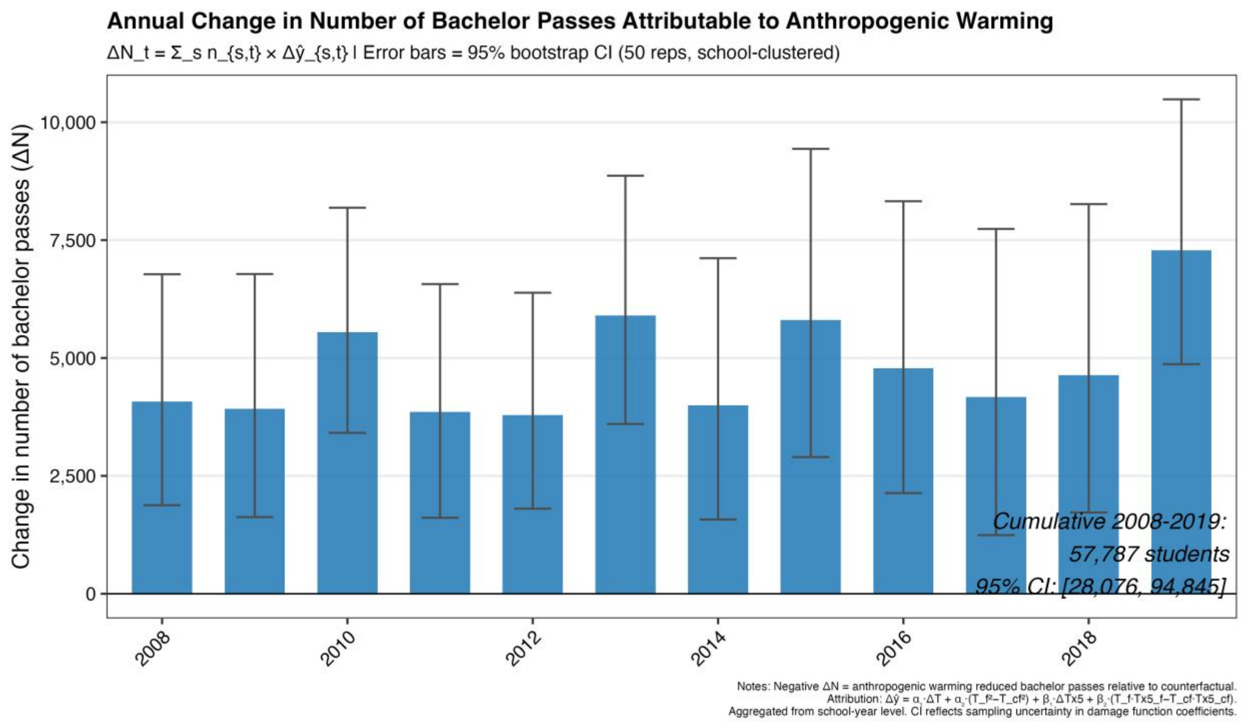

Figure S5. Annual change in the number of bachelor passes attributable to anthropogenic warming (ΔN_t), with 95% bootstrap confidence intervals. Cumulative estimate: 57,787 additional passes over 2008–2019 (95% CI: [28,076; 94,845]).

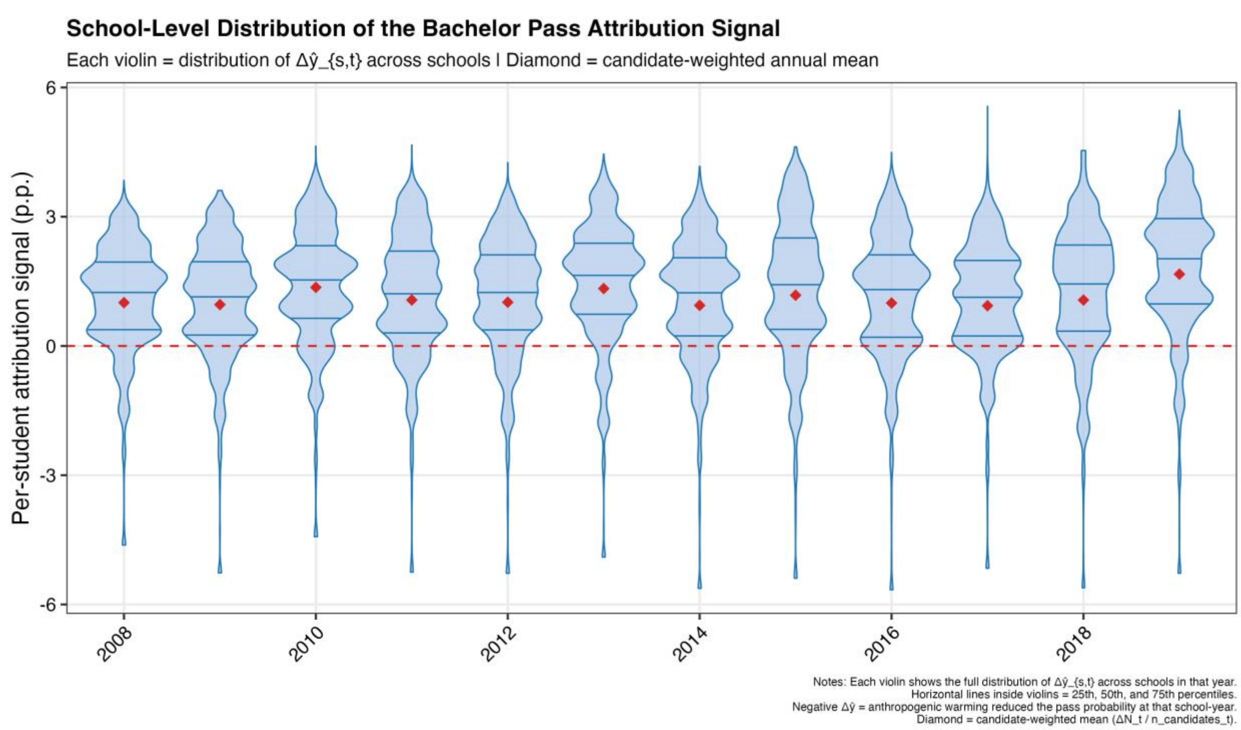

Figure S6. School-level distribution of the per-student attribution signal ($\Delta\hat{y}_{st}$) for each year (2008–2019). Each violin shows the full cross-school distribution; the diamond denotes the

candidate-weighted annual mean. Negative values indicate schools where anthropogenic warming reduced pass probability.

S5. Monetary Valuation (Preliminary)

We convert the attribution results for the Limpopo province into an economic value (Net present Value). The formular for the conversion is given by. We take into account the probability of enrolling into a university degree and completing it. We made a number of assumptions and we will update the valuation methods.