

# Estimating turning points and testing for climate adaptation \*

Tom Bearpark

University of Exeter

Jordan Rosenthal-Kay

Federal Reserve Bank of San Francisco

February 25, 2026

## Abstract

Estimating the turning points of nonlinear response functions, such as the peak of the temperature–GDP relationship or the minimum-mortality temperature, is central to environmental economics. For example, turning point estimates are used to parameterize damage functions, and to test for adaptation by comparing estimated optima to local climate. We show that the standard turning point estimator suffers from the same pathologies as weak-instrument IV: it is a ratio estimator with a fat-tailed sampling distribution and severe median bias. When curvature is weak, the sampling distribution of this estimator is approximately Cauchy, centered at the regressor mean, generating spurious correlations between estimated optima and local climate that mimic adaptation. We demonstrate the empirical relevance of our analysis in temperature–GDP and temperature–mortality applications. Conventional confidence intervals substantially understate uncertainty and cross-sectional patterns previously interpreted as adaptation arise as statistical artifacts. We propose two solutions: (i) robust sign-restricted confidence sets for turning points, and (ii) modeling curvature heterogeneity via dimensionality reduction and hierarchical pooling across units to stabilize estimates and reduce median bias.

---

\*Authors' emails: [T.Bearpark@exeter.ac.uk](mailto:T.Bearpark@exeter.ac.uk), [Jordan.Rosenthal-Kay@sf.frb.org](mailto:Jordan.Rosenthal-Kay@sf.frb.org). The views in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. We thank Gregory Casey, Antonio Gasparrini, Manuel Linsenmeier, Eric Qian, and Guglielmo Zappalà for helpful comments. We thank Gracie Ellsworth for excellent research assistance. All errors are our own.

# 1 Introduction

An influential empirical literature estimates dose–response relationships between environmental variables, such as temperature, and socioeconomic outcomes, such as GDP growth and mortality (see Hsiang (2025) and Lemoine et al. (2025) for recent reviews). As both low and high extremes of many environmental variables are likely to be harmful, researchers typically estimate models that allow for curvature, using, for example, quadratic, spline, or binned regression models to approximate the underlying dose-response functions (Hogan and Schlenker, 2024). In these models, the turning point is a central estimand. Estimated optima are used in integrated assessment models to parameterize damage functions (Kotlikoff et al., 2021; Krusell and Smith, 2022; Cruz and Rossi-Hansberg, 2024; Bourany, 2025). They also underpin empirical claims about the extent of climate adaptation (Linsenmeier, 2024). For example, in environmental epidemiology, researchers estimate location-specific minimum mortality temperatures (MMTs) and compare them to local climate normals. Proximity between the two is interpreted as evidence that populations are well adapted to their thermal environment (Gasparri et al., 2015; Yin et al., 2019; Tobías et al., 2021; Carleton et al., 2022).

In this paper, we show that inference on turning points is statistically fragile for two reasons. First, when curvature is weakly identified, the estimate of the optimum constructed from OLS regression coefficients has a *fat-tailed sampling distribution*, analogous to an IV estimator under weak instruments, so conventional delta-method confidence intervals can severely understate uncertainty. Moreover, this estimator is *median biased* toward the sample mean of the regressor, which can generate spurious cross-sectional patterns that resemble those that would be found in a world with strong adaptation. As a result, existing studies may both overstate the precision of estimated optima and mistake statistical artifacts for evidence of adaptation.

To illustrate the basic problem, consider outcomes  $y_{it}$  in location  $i$  and period  $t$  generated by  $y_{it} = f_i(T_{it}) + u_{it}$ , with  $\mathbb{E}[u_{it} | T_{it}] = 0$ , where  $T_{it}$  denotes realized temperature and  $f_i(\cdot)$  is twice continuously differentiable and strictly concave. A turning point  $\tau_{0i}$  exists, defined by  $f'_i(\tau_{0i}) = 0$  when  $f''_i(\tau_{0i}) \neq 0$ . Empirically, researchers estimate  $f_i$  and recover an implied optimum  $\hat{\tau}_i$  from the first-order condition  $\hat{f}'_i(\hat{\tau}_i) = 0$ . A local expansion around  $\tau_{0i}$  yields  $\hat{\tau}_i - \tau_{0i} \approx -\frac{\hat{f}'_i(\tau_{0i})}{\hat{f}''_i(\tau_{0i})}$ , where  $\hat{f}'_i(\tau_{0i}) \rightarrow_p f'_i(\tau_{0i})$ . Thus, the sampling behavior of  $\hat{\tau}_i$  depends critically on curvature. When  $f''_i(\tau_{0i})$  is weakly identified,  $\hat{\tau}_i$  is effectively a ratio estimator and inherits a fat-tailed sampling distribution, analogous to instrumental variables estimators under a weak first stage. Two implications follow.

**Implication 1: Conventional inference understates uncertainty under weak curvature.** The turning point is a ratio estimand. Standard delta-method approximations rely on the denominator (the curvature term) being estimated precisely. When the curvature term is imprecisely estimated,  $\hat{\tau}$  inherits the same pathologies as ratio estimators in the weak-instrument literature: the sampling distribution is heavy-tailed, confidence intervals based on local normal approximations undercover, and the appropriate uncertainty can be unbounded.<sup>1</sup> Under the null that curvature is zero, the sampling distribution of  $\hat{\tau}$  is well approximated by a Cauchy centered at the mean of the regressor. When an interior optimum  $\tau_0$  does exist but curvature is weak, the distribution of  $\hat{\tau}$  remains non-Gaussian and is well-approximated by a mixture of (i) this Cauchy component and (ii) a regular component that is approximately normal and centered at  $\tau_0$ . The mixture weights are governed by the strength of curvature identification, conveniently summarized by the  $t$ -statistic on the curvature term.

**Implication 2: Median bias and spurious adaptation patterns.** Following the climate impacts literature, we say location  $i$  is *fully adapted* to its climate if marginal damages vanish at the prevailing climate, so that  $f'_i(\mu_i) = 0$ . When curvature is weakly identified, turning-point estimators exhibit a systematic bias toward the sample mean of the regressor. In this regime, cross-sectional patterns that resemble climate adaptation can arise mechanically, even when the true underlying response function is identical across locations. Figure 1 illustrates the point. Panel (a) plots country-specific estimates of the optimal temperature,  $\hat{\tau}_i$ , against each country's long-run mean temperature,  $\mu_i$ , using canonical panel GDP data and a quadratic response function (Burke et al., 2015, see details of the estimation in Appendix A). The figure reproduces a striking empirical regularity: hotter countries appear to have higher estimated optima, and the relationship lies close to the 45-degree line that would obtain under full adaptation. However, many estimated turning points reflect minima, and as Panel (b) shows, the same pattern arises in a placebo exercise in which we simulate from a data generating process in which there is no impact of temperature on GDP. These issues are not unique to a quadratic specification: Appendix Figure A1 repeats this exercise using higher-order polynomials, splines, and binned estimators of the response function, finding the same result. The implication is that correlations between estimated optima and climate are not, by themselves, informative about adaptation.

These implications are amplified when curvature is allowed to be heterogeneous, since the effective sample size for each estimated optimum is reduced. Researchers face a tradeoff in specifying the degree of heterogeneity that they allow for. Pooling across locations (i.e. imposing  $\tau_{0i} = \tau_0$  for all  $i$ ) yields a more precisely estimated common optimum, since curvature is identified from the full

---

<sup>1</sup>The analogy is to IV under a weak first stage, where the Wald estimator is a ratio of approximately normal statistics and conventional asymptotics fail.

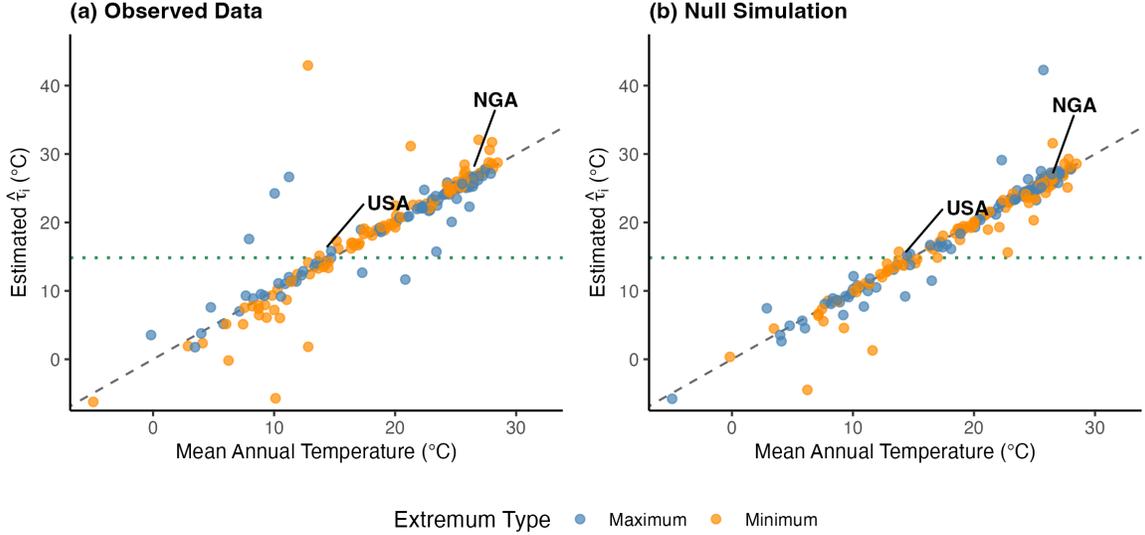


Figure 1: **Spurious adaptation from median bias in turning-point estimates.**

Panel (a) plots country-specific estimates of the optimal temperature  $\hat{\tau}_i$  against each country's long-run mean temperature  $\mu_i$  using panel GDP data. Details of the estimation procedure used here are in Appendix A.1. Panel (b) repeats the same exercise in which temperature has no effect on GDP, while matching the empirical temperature process and residual variation. Points are colored by whether the turning point is a maximum or a minimum. Details of the simulation procedure used here are in Appendix A.1. The dashed line is the 45-degree line ( $\hat{\tau}_i = \mu_i$ ), corresponding to full adaptation, and the dotted horizontal line is the pooled optimum estimate.

sample, but rules out heterogeneity and precludes the possibility that all units are locally adapted to their climate. Estimating  $\tau_{0i}$  separately for each unit is more general: it nests adaptation ( $\tau_{0i} = \mu_i$ ), a common optimum ( $\tau_{0i} = \tau_0$ ), and intermediate structures such as  $\tau_{0i} = g(X_i)$  for exogenous covariate or grouping variable  $X_i$ . Thus, accounting for heterogeneity allows the data to discriminate between different models of adaptation, but this generality requires that curvature be well-identified in each unit individually.

## Contributions

This paper makes three contributions.

**(1) Turning-point statistics under weak curvature.** We characterize the finite-sample behavior of turning-point (argmax/argmin) estimators when curvature is weakly identified. In a simple parametric setting, the estimator is a ratio statistic with nonregular sampling behavior: it is well-approximated by a mixture of (i) a Cauchy component centered at the regressor mean and (ii) an

approximately Gaussian component centered at the true optimum, with mixture weights governed by identification strength. We complement the analytic results with simulations showing that the same qualitative pathologies—fat tails, potential unbounded uncertainty, and pronounced median bias toward the regressor mean—arise in semi- and nonparametric implementations commonly used in applied work. While related ratio-estimator issues have been noted in specific applications (e.g., environmental Kuznets curve peaks and elasticity calculations, Hinkley, 1969; Plassmann and Khanna, 2007; Hirschberg et al., 2008; Lind and Mehlum, 2010; Bernard et al., 2015), our contribution is to provide a unified characterization that (i) isolates the *median-bias channel* as a source of spurious cross-sectional ‘adaptation’ patterns and (ii) connects turning-point estimation directly to weak-identification logic for ratio estimators (Staiger and Stock, 1997).

Moreover, our results have value for inference on turning points in other contexts, for example, in education economics (estimating optimal class size, Kedagni et al., 2021), or urban economics (estimating optimal city size, Au and Henderson, 2006), macroeconomics (estimating the turning point in debt-to-GDP/growth relationship, Pattillo, Ricci, et al., 2011; Checherita-Westphal and Rother, 2012; Eberhardt and Presbitero, 2015) and bear upon key findings in industrial organization (Aghion et al., 2005, find that the inverted-U between innovation and competition peaks at the median level of competition).<sup>2</sup>

**(2) Practical solutions for inference and estimation.** Building on our diagnosis, we propose two complementary remedies. For inference, we construct identification-robust confidence sets for turning points by inverting tests of the first-order condition in the spirit of Anderson and Rubin (1949), delivering correct coverage regardless of curvature strength.<sup>3</sup> We also develop an economically-motivated likelihood ratio test that assumes the sign of curvature is known. Our tests are therefore closely related to the literature on weak IV with a known first-stage sign (Andrews and Armstrong, 2017; Angrist and Kolesár, 2024), but are, to our knowledge, a novel formulation.

For estimation under heterogeneity, we propose methods that explicitly manage the heterogeneity–identification tradeoff: (i) projection approaches that model optima as functions of observables (e.g.,  $\tau_{0i} = g(X_i)$ ), and (ii) a hierarchical Bayesian estimator that shrinks location-specific turning points toward a pooled mean with shrinkage intensity disciplined by local identification strength. Intuitively, heterogeneity is preserved where the data strongly identify curvature, and suppressed where local optima are weakly identified; this yields a disciplined compromise between fully pooled

---

<sup>2</sup>The location of the competition-innovation peak, and how it changes with respect to the technological gap between competitors can be used as a test of Aghion et al. (2005)’s theory; see Hashmi (2013). We return to this example in Appendix C.2.

<sup>3</sup>This parallels identification-robust inference for weak-IV settings, where Wald-type approximations can severely undercover (Staiger and Stock, 1997).

and fully heterogeneous specifications. Conceptually, this approach complements recent Bayesian work on heterogeneous weather effects (Goh et al., 2025).

**(3) Climate applications and the econometrics of ‘adaptation’ tests.** We apply the framework to canonical temperature–GDP (Burke et al., 2015) and temperature–mortality (Carleton et al., 2022) relationships. Empirically, robust confidence sets for optimal temperatures are often substantially wider than conventional delta-method intervals when curvature is weak, and shrinkage-based turning-point estimates can materially alter projected damages relative to standard pooled estimators. More broadly, we provide statistical foundations for using estimated turning points to test for adaptation. This speaks to emerging ‘optimum–climate’ adaptation diagnostics, like estimating country-specific peaks in spline models (Linsenmeier, 2024), and heterogeneous minimum mortality temperatures common in environmental epidemiology (Gasparrini et al., 2015; Yin et al., 2019; Tobías et al., 2021). It is also related to common two-step exercises that regress estimated response-function features on long-run climate to make inferences about adaptation (Heutel et al., 2021; Auffhammer, 2022).

**Outline.** The paper proceeds as follows. Section 2 characterizes the sampling distribution of turning-point estimators under weak curvature in a simple parametric setting, and provides simulation evidence that the key pathologies, fat tails and median bias, extend to nonparametric estimators and more realistic data generating processes. We propose robust confidence sets that deliver correct coverage regardless of identification strength, and compare them to standard delta-method intervals in temperature–GDP and temperature–mortality applications. Section 3 addresses the heterogeneity–identification tradeoff, developing projection and hierarchical Bayesian approaches that pool information across units while allowing for heterogeneous optima where the data support them. Section 4 concludes.

## 2 Inference on turning points

In this section, we characterise the econometrics of turning-point estimators. Section 2.1 derives the sampling distribution in a simple parametric setting, showing that it is governed by a curvature-strength parameter  $\theta$  and ranges from approximately Cauchy (weak curvature) to approximately normal (strong curvature). Section 2.2 provides simulation evidence that these pathologies extend to heteroskedastic settings and to nonparametric estimators. Section 2.3 proposes identification-robust confidence sets that deliver correct coverage regardless of  $\theta$ , and illustrates the practical difference from delta-method intervals in two canonical applications.

## 2.1 The sampling distribution of turning point estimators

We begin with the simplest parametric environment in which the distribution of the turning point estimator can be characterized analytically: i.i.d. Gaussian  $x_i$  and a quadratic regression with homoskedastic errors, where  $\hat{\tau} = -\hat{\beta}_1/(2\hat{\beta}_2)$ .<sup>4</sup>

**Proposition 2.1** (Sampling distribution of turning points in quadratic models). *Let  $\{(x_i, u_i)\}_{i=1}^n$  be i.i.d. with*

$$x_i \sim \mathcal{N}(\mu, \sigma_x^2), \quad u_i \sim \mathcal{N}(0, \sigma_u^2), \quad x_i \perp u_i,$$

*and suppose  $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + u_i$  with  $\beta_2 \neq 0$ . Define  $\hat{\tau} = -\hat{\beta}_1/(2\hat{\beta}_2)$  and  $\tau_0 = -\beta_1/(2\beta_2)$ .*

1. **Sampling distribution (Hinkley)** For sample size  $n$ :

$$\hat{\tau} \sim \mu - \frac{\sigma_x}{\sqrt{2}} \cdot \frac{Z_1 + \sqrt{2} \left( \frac{\mu - \tau_0}{\sigma_x} \right) \theta}{Z_2 + \theta}, \quad Z_1, Z_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1),$$

where

$$\theta = 2\beta_2 \frac{\sigma_x}{\sqrt{2}} \sqrt{n} \frac{\sigma_x}{\sigma_u}.$$

2. **Behavior indexed by  $|\theta|$**

- (a) Well-identified ( $|\theta|$  large):  $\hat{\tau}$  is normal, centered at  $\tau_0$ . Delta-method inference is valid.
- (b) Poorly-identified ( $|\theta|$  small): Heavy-tailed distribution. When both  $|\theta|$  and  $\frac{\mu - \tau_0}{\sigma_x}$  are small,  $\hat{\tau}$  approaches Cauchy( $\mu, \sigma_x/\sqrt{2}$ ), symmetric about  $\mu$  (not  $\tau_0$ ). Delta-method inference fails, and the estimator is median biased towards  $\mu$ .

3. **Observable diagnostic** For large  $n$ , the  $t$ -statistic on  $\beta_2$ ,  $t_{\beta_2} \sim \mathcal{N}(\theta, 1)$ , so  $t_{\beta_2}$  estimates  $\theta$ . (Under normal errors,  $t_{\beta_2}$  is exactly noncentral- $t$  with  $n - 3$  d.f.; the normal approximation holds for large  $n$ .)

*Proof:* See Appendix D.1.

Proposition 2.1 shows that the sampling behavior of the turning-point estimator is governed by a parameter,  $\theta$ , which measures the strength of the curvature signal in the data. The distribution of  $\hat{\tau}$  is the Fieller–Hinkley distribution for the ratio of two normals, which can be viewed as ranging continuously from approximately Cauchy when  $|\theta|$  is small to approximately normal when  $|\theta|$  is

<sup>4</sup>Note that the presence of fixed effects or additional covariates does not affect the arguments here. We can think of  $x_i, x_i^2$ , and  $y_i$  as being the residual variation after partialling out such controls.

large.<sup>5</sup> Appendix D.3 fully derives this distribution in closed form for all values of  $|\theta|$ . In Appendix D.2, we provide a complementary asymptotic formulation using local-to-zero sequences for  $\beta_2$ , which delivers the same Hinkley distribution as an exact limit under weaker assumptions (finite fourth moments rather than normality of errors).

The appearance of  $\mu$  in the weak-curvature limit may seem surprising: why should the mean of  $x$  matter for a ratio of regression coefficients? Intuitively, with little information about curvature, many quadratic fits achieve nearly identical mean-squared error. OLS resolves this near-indeterminacy by placing the extremum where the data are densest, which is the center of the  $x$  distribution. The distribution also depends on  $(\tau_0 - \mu)/\sigma_x$ , which governs skewness and the speed at which the normal approximation becomes adequate.

Figure 2 illustrates both features. When curvature is weak (left column), the Hinkley distribution is skewed and heavy-tailed, with the direction and magnitude of the skewness governed by  $(\tau_0 - \mu)/\sigma_x$ . When the peak is far from climate (bottom left), the distribution is bimodal and the delta-method normal is visibly displaced from the bulk of the density. When curvature is strong (right column), the normal approximation improves, but convergence is slow when  $|\tau_0 - \mu|/\sigma_x$  is large: the turning point lies where data are sparse, generating residual skewness even at  $\theta = 3$  (bottom right). The implication for adaptation inference is nuanced. On the one hand, the normal approximation works best when the peak is near climate and curvature is strong — precisely the setting consistent with well-identified adaptation. On the other hand, when curvature is weak, the estimator is most severely biased toward  $\mu$  when the peak is genuinely near climate, making it impossible to distinguish true adaptation from statistical artifact without accounting for the non-normality of the sampling distribution.

In Appendix Figure A2, we display median bias and confidence interval width as a function of  $\theta$  for the values of  $(\tau_0 - \mu)/\sigma_x$  used in Figure 2. The Hinkley median is a smooth, sigmoid-like function in  $|\theta|$ , and the Hinkley distribution is median-biased towards  $\mu$  until  $\theta$  exceeds 2. The delta method undercovers for  $\theta > 0.5$  and has incorrect coverage when  $\theta$  is small. When  $\theta$  is small, the Hinkley distribution is bimodal, so the high-density region is not convex. The confidence interval

---

<sup>5</sup>Hinkley (1969) derived the distribution for the ratio of two normals on which our proposition is based. Prior to Hinkley, Geary (1930) found a transformation of the ratio of two correlated normals such that the resulting random variable itself was normal, and along the way derives that the ratio of two mean-zero normals is Cauchy distributed, though writes that “this particular case has probably no practical value.” Marsaglia (1965) derived the distribution of two normals with unit variance and nonzero means, but incorrectly characterized the ratio of two arbitrary normals. The 1969 paper by Hinkley corrects Marsaglia and notes that his work draws on Fieller (1932), who bounded the error of a normal approximation for the ratio of two normals. Concurrently, Press (1969) derived the density for the ratio of two arbitrary  $t$ -distributed random variables. The distribution derived in Hinkley (1969) for the ratio of two arbitrary normals,  $f(x)$ , is heavy tailed in the sense that as  $x \rightarrow \infty$ ,  $f(x) \approx cx^{-2}$ , which implies that it has no finite moments.

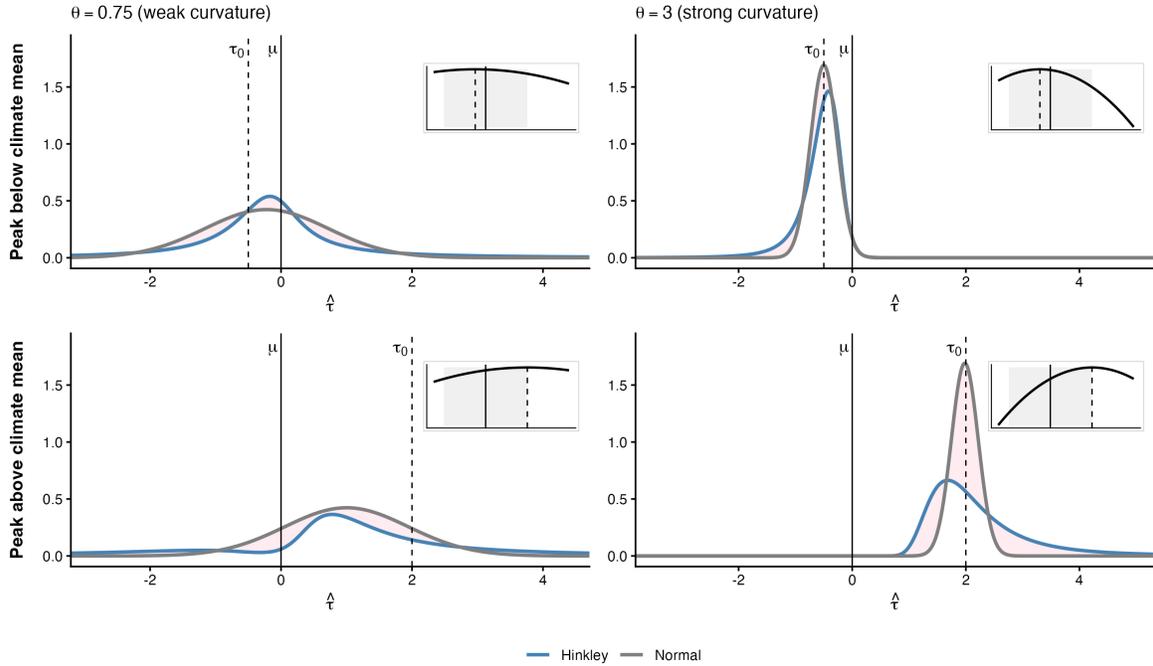


Figure 2: **Sampling distribution of the turning-point estimator.**

Each panel plots the exact Hinkley distribution (blue) against a normal distribution (grey) centered at the Hinkley median with delta-method variance; any discrepancy reflects the failure of the normal shape approximation. Pink shading highlights where the two distributions diverge. Columns vary curvature strength:  $\theta = 0.75$  (left, weak) and  $\theta = 3$  (right, strong). Rows vary the distance between the true optimum and local climate:  $\tau_0 = \mu - \sigma_x/2$  (top, peak near climate) and  $\tau_0 = \mu + 2\sigma_x$  (bottom, peak far from climate). Vertical lines mark  $\mu$  (solid) and  $\tau_0$  (dashed). Insets show the underlying response function; the shaded region indicates  $\mu \pm 2\sigma_x$ : the approximate support of the data. Parameters:  $\mu = 0$ ,  $\sigma_x = 1$ .

width shown in the figure therefore does not correspond to a valid confidence set: a single interval necessarily excludes some regions of the support. For this reason bootstrapped confidence intervals, like those reported in Newell et al. (2021), will also provide incorrect coverage when curvature is weak.

The extent to which the delta method fails, and the role  $(\tau_0 - \mu)/\sigma_x$  plays, can be further understood through power calculations. In the left panel of Appendix Figure A3, we plot power curves for the test  $H_0 : \tau_0 = \mu$ , as a function of  $(\tau_0 - \mu)/\sigma_x$ . The delta method performs poorly: when  $\tau_0$  is close to  $\mu$ , even when  $|\theta|$  is large, the test is undersized, and rarely rejects the null. As  $(\tau_0 - \mu)/\sigma_x$  becomes large, the test becomes oversized, and rejects the null far too frequently, even when  $|\theta|$  is large.

## 2.2 Robustness and extensions

Proposition 2.1 relies on i.i.d. Gaussian regressors, homoskedastic errors, and a correctly specified quadratic model. These restrictions yield closed-form expressions, but are implausible in real world applications. In this subsection, we provide simulation evidence that the core pathologies (median bias toward  $\mu$ , fat tails, and the failure of normal approximations under weak curvature) survive when these assumptions are relaxed. We then show that they extend to nonparametric estimators, where finite-sample performance is even worse.

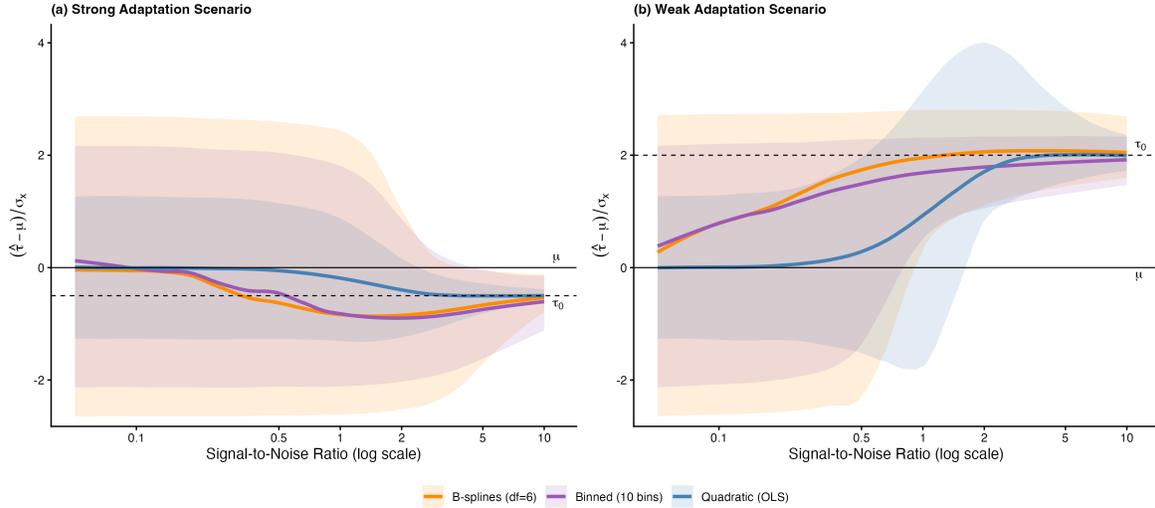
### Relaxing the DGP assumptions

We simulate  $\hat{\tau}$  estimates from the DGP of Proposition 2.1, and then sequentially relax two assumptions: first, we introduce heteroskedasticity; second, we replace the normal distribution for  $x_i$  with a skewed distribution, holding the mean and variance fixed at  $\mu$  and  $\sigma_x^2$ . Appendix Figure A5 reports the results. Introducing heteroskedasticity (middle panel) worsens median bias and generates bimodality and regions of no support when  $\tau_0$  differs from  $\mu$ . Even when the curvature signal is strong ( $|t_{\beta_2}| \approx 4$ ), the sampling distribution remains skewed, with the direction depending on the sign of  $\tau_0 - \mu$ . Introducing skewness in  $x_i$  (bottom panel) has little additional effect: the performance of  $\hat{\tau}$  is essentially unchanged. Overall, the qualitative patterns from Proposition 2.1 are robust to these departures from the baseline DGP.

### Extension to nonparametric estimators

The intuition behind Proposition 2.1 extends beyond the quadratic model. Consider a general estimator  $\hat{\tau} = \operatorname{argmax}_{x \in \Omega} \hat{f}(x)$ , where  $\hat{f}$  minimizes a least-squares criterion over some function class  $\mathcal{F}$ . Provided  $f$  is twice differentiable, a Taylor expansion around  $\tau_0$  gives  $\hat{\tau} - \tau_0 \approx -\frac{\hat{f}'(\tau_0)}{\hat{f}''(\tau_0)}$ , so  $\hat{\tau}$  is again a ratio estimator, and weak curvature ( $\hat{f}''(\tau_0) \approx 0$ ) generates fat tails even when  $\hat{f}'(\tau_0) \rightarrow 0$ . For parametric estimators with well-behaved  $\mathcal{F}$ , the argmax continuous mapping theorem delivers  $n^{-1/2}$  consistency. However, for nonparametric estimators, where  $\hat{f} \rightarrow f$  in a function-space norm, the argmax can converge at rates as slow as  $n^{-1/3}$  (J. Kim and Pollard, 1990), suggesting substantially worse finite-sample performance.

To examine this directly, we compare three commonly used estimators under a quadratic DGP: OLS with a quadratic specification, a cubic B-spline estimator (6 basis functions), and a binned estimator (10 bins). For each, we simulate the distribution of  $\hat{\tau}$  across a grid of signal-to-noise



**Figure 3: Neither B-splines nor binning fix weak identification.**

Each panel plots the median estimate of  $\hat{\tau}$  (lines) and 68% Monte Carlo confidence bands (shaded areas) across signal-to-noise ratios, defined by  $\sigma_u/\sigma_x$  (x-axis), for three estimators: quadratic OLS, cubic B-spline (6 d.f.), and binned (10 bins). Left:  $\tau_0 = \mu - \sigma_x/2$  (peak near climate). Right:  $\tau_0 = \mu + 2\sigma_x$  (peak far from climate). Horizontal dashed lines mark  $\tau_0$  and  $\mu$ . As the signal weakens (moving right), all three estimators converge toward  $\mu$ , confirming that median bias is not an artifact of the quadratic specification. Nonparametric estimators exhibit higher variance but no less bias.  $n = 200$ , 15,000 replications,  $\sigma_x = 1$ .

ratios, holding  $\sigma_x$  fixed and varying  $\sigma_u$ . Figure 3 reports the results for two scenarios: a ‘strong adaptation’ case where  $\tau_0$  is near  $\mu$  (left panel,  $\tau_0 = \mu - \sigma_x/2$ ) and a ‘weak adaptation’ case where  $\tau_0$  is far from  $\mu$  (right panel,  $\tau_0 = \mu + 2\sigma_x$ ). All three estimators converge to  $\mu$  as the signal-to-noise ratio deteriorates, confirming that the bias toward  $\mu$  is not an artifact of the quadratic specification. The quadratic estimator degrades smoothly, transitioning from  $\tau_0$  to  $\mu$  as curvature weakens. The B-spline and binned estimators are more erratic: when  $\tau_0$  is far from  $\mu$ , their finite support acts as a form of implicit regularization, bounding  $\hat{\tau}$  within the range of the data. But when  $\tau_0$  is near  $\mu$  and within the data support, the nonparametric estimators exhibit substantially higher variance than OLS. In neither case do flexible methods resolve the weak-identification problem. They merely trade the smooth parametric bias for noisier, less predictable behavior.

## Misspecified DGP

Appendix Figure A6 repeats the exercise under misspecified DGPs — an asymmetric quadratic and an asymmetric Gaussian peak — where the true response function is not quadratic. The semiparametric estimators exhibit less bias but substantially higher variance than the quadratic estimator when the model is misspecified, and all estimators continue to show bias toward  $\mu$  under weak

curvature, consistent with the patterns in Proposition 2.1.

### 2.3 Robust confidence sets for turning points

The results in Section 2.1 imply that delta-method confidence intervals for  $\tau_0$  can severely under-cover when curvature is weak. We now develop identification-robust alternatives. We first construct Anderson-Rubin inspired confidence sets by inverting a test of the linear restriction  $\beta_1 + 2\beta_2\tau_0 = 0$ , which provides correct coverage regardless of  $\theta$ , extending the weak-IV toolkit of Staiger and Stock (1997) to our setting. We then develop a sign-restricted likelihood-ratio (SLR) test that imposes the economically motivated restriction that  $\tau_0$  is a maximum (or minimum), i.e. we treat the sign of  $\beta_2$  as known. Appendix C extends both constructions to a sign-restricted projection method (SAR) that is more conservative, but has an intuitive geometric interpretation.

#### Anderson-Rubin inspired confidence sets (AR)

The key observation is that, for any hypothesized value  $\tau$ , the first-order condition  $f'(\tau) = 0$  is equivalent to the linear restriction  $H_0(\tau) : \beta_1 + 2\beta_2\tau = 0$ . Testing this restriction requires only a standard test on a linear combination of coefficients, which is valid regardless of whether curvature is well-identified. Let  $\hat{\Sigma}$  denote a consistent estimator of the variance-covariance matrix of  $(\hat{\beta}_1, \hat{\beta}_2)'$ . An identification-robust confidence set for the turning point is then obtained by collecting all values of  $\tau$  that are not rejected:

$$\text{CS}_\alpha^{\text{AR}} = \left\{ \tau : \frac{(\hat{\beta}_1 + 2\hat{\beta}_2\tau)^2}{r(\tau)' \hat{\Sigma} r(\tau)} \leq \chi_{1,1-\alpha}^2 \right\},$$

where  $r(\tau) = (1, 2\tau)$  and  $\chi_{1,1-\alpha}^2$  is the  $(1 - \alpha)$ -quantile of a chi-squared distribution with one degree of freedom. This construction inverts the test of a linear hypothesis for each candidate  $\tau_0$ , in the spirit of Anderson and Rubin (1949).

Under the conditions of Proposition 2.1, the set  $\text{CS}_\alpha^{\text{AR}}$  defined above has exact coverage  $1 - \alpha$  for  $\tau_0$ , regardless of the value of  $\theta$ . We confirm this using power calculations in the right panel of Appendix Figure A3, which shows that the Anderson-Rubin test is correctly sized regardless of  $\theta$  and  $(\tau_0 - \mu)/\sigma_x$ . When  $|\theta|$  is large,  $\text{CS}_\alpha$  reduces approximately to the standard delta-method interval. When  $|\theta|$  is small, the set can be a strict subset of the real line (wide but bounded), the entire real line (reflecting the impossibility of locating  $\tau_0$ ), or the complement of a bounded

interval (an ‘exclusive’ set indicating where  $\tau_0$  is unlikely to be).<sup>6</sup> Note, this test can be constructed for as long as we have a consistent variance-covariance matrix, which could be cluster- and/or heteroskedasticity-robust.

### Sign-restricted likelihood-ratio (SLR) confidence sets

Anderson-Rubin confidence sets provide correct coverage for a turning point, but do not impose whether  $\tau_0$  reflects a maximum or minimum of the response function. When curvature is weak, this can be problematic: the confidence set may include values of  $\tau$  corresponding to  $\beta_2 > 0$ , which are economically irrelevant if the researcher knows the response is concave. Imposing  $\beta_2 < 0$  should exclude these regions and sharpen inference. We now develop a likelihood-ratio test that imposes this restriction under both the null and alternative.<sup>7</sup> The SLR test asks whether the best-fitting concave quadratic passing through  $\tau$  fits significantly worse than the best-fitting concave quadratic overall.

In particular, define the restricted parameter space,  $\Omega = \{\beta : \beta_2 \leq 0\}$  and define the parameter space under the null  $\Omega_0(\tau) = \{\beta : \beta_1 + 2\tau\beta_2 = 0, \beta_2 \leq 0\}$ . The likelihood ratio test statistic for a test of  $\tau = \tau_0$  under  $\beta_2 \leq 0$  is,

$$\omega(\tau) = \min_{\beta \in \Omega_0(\tau)} (\hat{\beta} - \beta)' \hat{\Sigma}^{-1} (\hat{\beta} - \beta) - \min_{\beta \in \Omega} (\hat{\beta} - \beta)' \hat{\Sigma}^{-1} (\hat{\beta} - \beta)$$

and so,

$$CS_\alpha^{SLR} = \{\tau : \omega(\tau) \leq \chi_{1,1-\alpha}^2\}.$$

When  $\hat{\beta}_2 < 0$ , the inequality constraints do not bind and  $\omega(\tau)$  reduces to the AR statistic, which is  $\chi_1^2$  under the null. When  $\hat{\beta}_2 \geq 0$ , the unrestricted minimum in  $\Omega$  is attained at the boundary  $\beta_2 = 0$ , and  $\omega(\tau)$  is bounded above by the AR statistic, so the  $\chi_1^2$  critical value remains conservative. Formally, the null distribution of  $\omega(\tau)$  is  $\bar{\chi}_{0:1}^2$ , a mixture of  $\chi_0^2$  and  $\chi_1^2$  (Self and Liang, 1987), so the  $\chi_1^2$  critical

---

<sup>6</sup>The confidence set  $CS_\alpha$  is closely related to the classical Fieller interval for the ratio of two normal means (Fieller, 1954). Fieller’s interval for  $\beta_1/(2\beta_2)$  inverts the same quadratic inequality that defines  $CS_\alpha$ , yielding an equivalent set when the joint distribution of  $(\hat{\beta}_1, \hat{\beta}_2)$  is normal. The connection to Anderson–Rubin underscores that the key requirement is linearity of the null in the original parameters, not normality of  $\hat{\tau}$  itself. The three geometric cases (i.e. the bounded interval, the entire real line, or complementary set) correspond to the taxonomy of Fieller intervals, governed by whether  $g = (c_\alpha/\hat{\theta})^2$  is less than, equal to, or greater than one. When  $g > 1$ , the curvature estimate is statistically indistinguishable from zero at the  $\alpha$  level, and the confidence set is unbounded: no finite turning point can be ruled out.

<sup>7</sup>In Appendix C.1 we develop an alternative procedure which has an intuitive geometric interpretation: the ‘sign restricted Anderson Rubin test’ (SAR). However, the likelihood ratio based interval described in the main text has superior power.

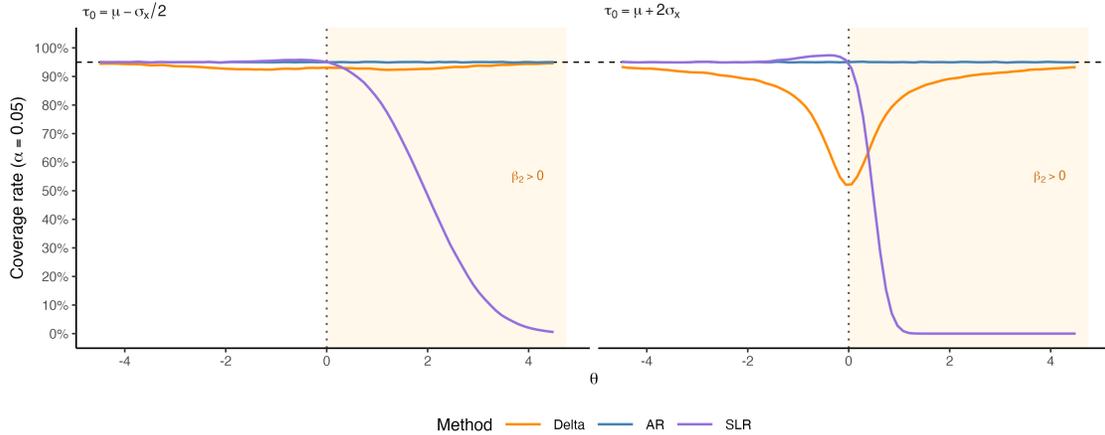


Figure 4: **Coverage of 95% confidence sets for the turning point.**

Each panel plots empirical coverage of nominal 95% confidence sets for the turning point  $\tau_0$  as a function of curvature strength  $\theta$ , under the maintained restriction  $\beta_2 < 0$ . The horizontal dashed line marks the nominal 95% level; departures below indicate undercoverage. The shaded region ( $\theta > 0$ ) corresponds to  $\beta_2 > 0$ , where the sign restriction is violated and sign-restricted tests should reject. Left:  $\tau_0 = \mu - \sigma_x/2$  (peak near climate). Right:  $\tau_0 = \mu + 2\sigma_x$  (peak far from climate). The delta method (orange) undercovers when  $\theta$  is small and the peak is far from climate. The Anderson–Rubin set (AR, blue) achieves correct coverage throughout. The sign-restricted likelihood ratio set (SLR, lilac) matches AR when  $\beta_2 < 0$  and rejects when the sign restriction is violated. Based on 20,000 Monte Carlo draws from the DGP of Proposition 2.1.

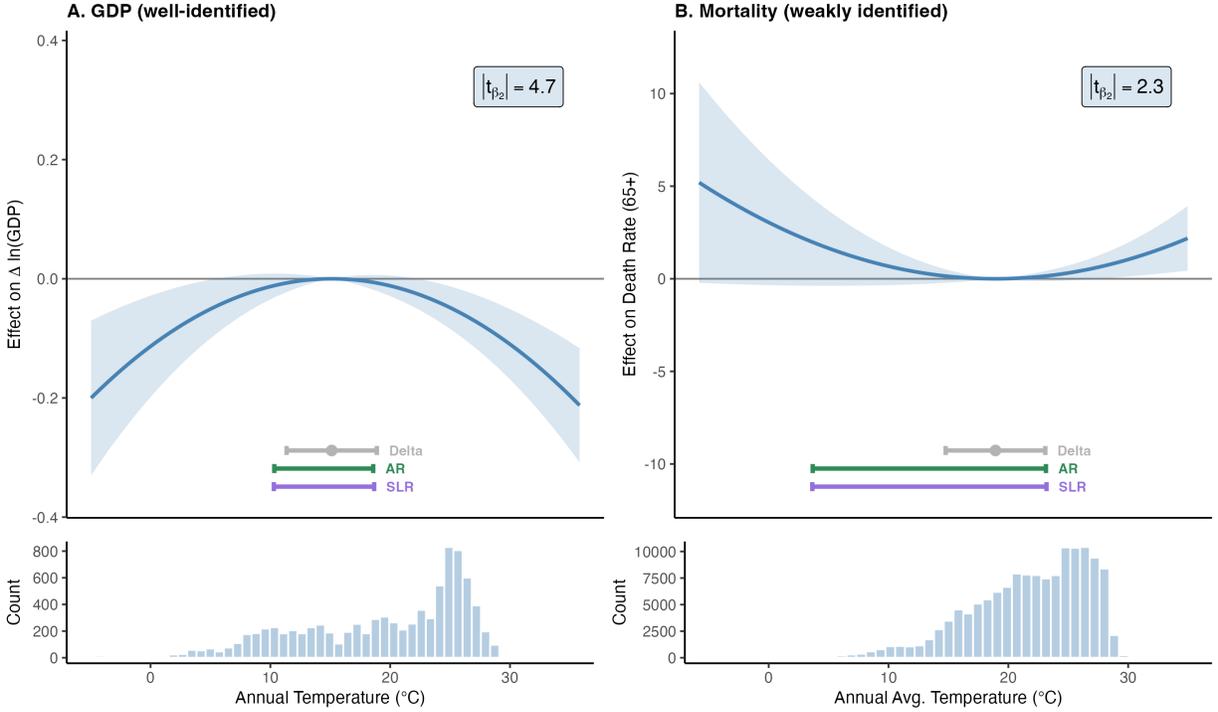
value provides valid (weakly conservative) coverage.

### Comparison of confidence sets

Figure 4 plots empirical coverage of nominal 95% confidence sets as a function of curvature strength  $\theta$ , under the maintained restriction  $\beta_2 < 0$ . The delta method undercovers when curvature is weak, and severely so when the peak is far from climate (right panel). The AR set achieves correct coverage throughout, regardless of  $\theta$ . The SLR set matches AR when the sign restriction is correctly imposed ( $\theta < 0$ ), with slight conservatism near  $\theta = 0$ , and coverage goes to zero as  $\theta$  grows positive — correctly rejecting turning points that are not maxima.

### Applications

Figure 5 illustrates the practical implications in two canonical climate applications. Burke et al. (2015) estimate a temperature–GDP relationship with strong curvature ( $|t_{\hat{\beta}_2}| > 4$ ); Carleton et al. (2022) estimate a temperature–mortality relationship with weaker curvature ( $|t_{\hat{\beta}_2}| \approx 2.3$ ). When



**Figure 5: Delta-method versus weak identification-robust inference on turning points.**

*Top row.* Estimated response functions for GDP growth (left, Burke et al., 2015) and 65+ age-group mortality rates (right, Carleton et al., 2022), with pointwise 95% confidence bands. Horizontal bars report three 95% confidence sets for the turning point  $\tau_0$ : the delta-method interval (grey); the Anderson–Rubin interval (AR, green), which inverts a  $\chi^2_1$  test of  $H_0: \tau_0 = \tau$ ; and the sign-restricted likelihood-ratio interval (SLR, purple), which additionally imposes the maintained assumption that the response is concave (GDP) or convex (mortality), recovering a  $\chi^2_1$  critical value. *Bottom row.* Histograms of temperature data included in the regressions. Further details on data and estimation in Appendix A.

curvature is strong, the AR and SLR intervals are only modestly wider than the delta-method interval, and all three largely agree. When curvature is weaker, the gap is stark: the AR and SLR intervals are approximately 2.3 times wider than the delta-method interval, spanning much of the temperature support.

Thus, in this example, we cannot reject the hypothesis of full adaptation for most locations in the estimating sample. For the GDP application, where curvature is well-identified ( $|t_{\beta_2}| > 4.7$ ), the delta method and AR confidence intervals are similar in width, and 8 countries fall in the gap between them. For mortality, where curvature is weakly identified ( $|t_{\beta_2}| = 2.3$ ) 31 of the 40 countries in the sample, including most of Northern and Central Europe have average temperatures that fall outside the delta method confidence interval but inside the AR interval. That is, the delta method falsely rejects proximity to the estimated optimum for the majority of the sample.

Appendix Figure A7 visualizes how the confidence intervals are constructed, and additionally presents the confidence interval derived from the SAR procedure described in Appendix C.1. For each hypothesized  $\tau_0$ , we plot the marginal effect  $\hat{\beta}_1 + 2\hat{\beta}_2\tau_0$  and its confidence band; the AR set collects all  $\tau_0$  for which this band includes zero. When curvature is weak, the marginal effect is nearly flat, so zero cannot be excluded over a wide range and the delta-method interval substantially understates uncertainty about the location of the optimum.

Appendix C.2 shows that the same patterns arise outside climate economics: using data from Aghion et al. (2005), we examine the innovation-maximising level of competition across UK manufacturing industries. Industry-level turning points cluster near mean competitiveness, consistent with median bias, and the AR confidence set is approximately 10% wider than the delta-method interval. This example uses a Poisson regression, rather than the OLS estimator discussed in the main text.

Identification-robust confidence sets tell the researcher which values of  $\tau_0$  are consistent with the data, but they do not resolve the median-bias problem identified in Section 2.1: the point estimate  $\hat{\tau}$  remains pulled toward  $\mu$  when curvature is weak. A natural response is to test adaptation directly but such tests have poor power under weak curvature, since detecting departures from adaptation requires precisely estimating  $\beta_2$ .<sup>8</sup> When  $\theta$  is small, the data cannot distinguish true adaptation from the absence of any curvature. This limitation motivates the estimation approaches developed in Section 3, which address median bias directly by pooling information across units.

### 3 Reducing median bias

The AR confidence sets developed in Section 2.3 deliver valid inference regardless of curvature strength, but they do not resolve the point-estimation problem:  $\hat{\tau}$  remains median-biased toward  $\mu$  when  $\theta$  is small, and this bias cannot be corrected analytically because the Hinkley distribution has no finite moments. The natural alternative is to pool information across units, trading some heterogeneity for improved precision. In this section, we consider two approaches: modeling  $\tau_i$  as a function of covariates (Section 3.1), which reduces dimensionality but inherits the same weak-curvature pathology; and Bayesian shrinkage (Section 3.2), which regularizes noisy estimates toward a prior. Pooling helps when it is guided by a sensible prior or a genuinely predictive covariate, but data-driven approaches struggle because the Hinkley likelihoods are too diffuse to discipline an

---

<sup>8</sup>E.g., testing  $H_0 : \tau_0 = \mu$ , which reduces to the linear restriction  $\beta_1 + 2\beta_2\mu = 0$ . Both AR and SLR confidence sets are likely to cover  $\mu$  when curvature is weakly identified.

empirical prior.

### 3.1 Modelling $\tau_i$ as a function of covariates

One way to reduce the dimensionality of  $\tau_i$ , and potentially improve its precision, is by modeling it as a function of covariates, rather than allowing it to be fully heterogeneous across units. In this section, we consider linear models of the form,  $\tau_i = \bar{\tau} + w_i' \eta$  where  $w_i$  is a standardized (mean zero, unit variance) covariate. For example,  $w_i$  could capture the level of economic development: e.g., income per capita, as in Carleton et al. (2022), or indicators for development status (Dell et al., 2012; Burke et al., 2015), or could be used to separate weather effects on agricultural and non-agricultural output, as in Conte et al. (2021). This motivates regressions of the form,

$$y_i = \beta_0 + (\beta_1 + \eta w_i)x_i + \beta_2 x_i^2 + u_i. \quad (1)$$

Equation (1) inherits the same weak curvature pitfalls as estimating  $\tau_i$  with unit-specific heterogeneity: in particular, under the null of no effect,  $\tau(w_i)$  estimates will be centered near  $\mathbb{E}[x | w]$ . Corollary 3.1 explores this property in the same restricted parametric setup as Proposition 2.1.

**Corollary 3.1.** *Let,*

$$\begin{pmatrix} x_i \\ w_i \end{pmatrix} \sim \mathcal{N} \left( \begin{pmatrix} \mu \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_x^2 & \rho\sigma_x \\ \rho\sigma_x & 1 \end{pmatrix} \right), \quad i = 1, \dots, n.$$

*Consider a regression of some random variable  $y_i$ ,  $y_i = \beta_0 + (\beta_1 + \eta_1 w_i)x_i + \beta_2 x_i^2 + u_i$  and suppose  $\mathbb{E}[u | x, w] = 0$  and  $Var(u | x, w) = \sigma_u^2$ . The statistic  $\tau(w) = -\frac{1}{2} \frac{\hat{\beta}_1 + \hat{\eta} w}{\hat{\beta}_2}$  has the following property for OLS estimators of  $\hat{\beta}, \hat{\eta}$ : When  $\beta_1 = \beta_2 = \eta_1 = 0$ ,*

$$\tau(w) \sim \text{Cauchy}(\check{\mu}(w), \check{\sigma}(w))$$

where,

$$\begin{aligned} \check{\mu}(w) &= \lambda \frac{\mathbb{E}[x | w]}{1 + \rho^2} + (1 - \lambda) \mu, \\ \check{\sigma}(w) &= \frac{\sigma_x}{\sqrt{2}} \times \frac{1}{1 + \rho^2} \times \sqrt{1 + \rho^2(1 - 2\lambda)} \times \sqrt{1 + \rho^2 + \lambda \left( w - \rho \frac{\mu}{\sigma_x} \right)^2} \end{aligned}$$

and  $\lambda = \frac{1 + \rho^2}{(\mu/\sigma_x)^2(1 - \rho^2) + (1 + \rho^2)}$ .

*Proof.* See Appendix D.4. □

The main idea of Corollary 3.1 is that modeling the peak as a linear function of a covariate  $w_i$  can improve precision on  $\beta_2$ , but under weak curvature, the sampling distribution inherits the same unattractive features of the distribution outlined in Proposition 2.1: it is fat-tailed and median biased. However, it is median biased towards a convex combination of a  $\rho$ -adjusted  $\mathbb{E}[x | w]$  and  $\mu$ , with weights dependent on the correlation  $\rho$ .

Often, researchers include interactions between  $w_i$  and  $x_i^2$  in the regression as well, so implicitly the peak is,  $\tau(w) = -\frac{1}{2} \frac{\beta_1 + \eta_1 w}{\beta_2 + \eta_2 w}$  (where  $\eta_2$  is the coefficient on  $x_i^2 \cdot w_i$ ). Likewise, this is Cauchy under the null, but its location parameter is not an easily interpretable function of  $\mathbb{E}[x | w]$ . Appendix Figure A9 plots the Cauchy median  $\check{\mu}(w)$  for models in which both the peak  $\tau_i$  and the slope  $\gamma_i$  are modeled as linear functions of  $w_i$ .<sup>9</sup> The plot shows that when including the  $w_i x_i^2$  interaction,  $\check{\mu}(w)$  behaves like  $\mathbb{E}[x | w]$  for over most of the support of  $w_i$ , but outlier  $w$ , the bias tends towards  $\mu$ . When there is no relationship between  $w_i$  and  $x_i$ , i.e.,  $\rho = 0$ ,  $\check{\mu}(w) = \mu$ , regardless of the model.

We implement this approach in which  $w_i$  represents the standardized mean log GDP per capita at the country level (for temperature-GDP) and at the admin 1 level (for mortality). Appendix Table A1 reports regression results in which  $w_i$  is interacted with  $x_i$ , its square, or  $w_i$  is replaced with a dummy equal to one if  $w_i$  is in the bottom tercile of the estimation sample. Appendix Figure A10 visualizes the implied peaks from these regressions, plotting  $\Delta\tau(w) = \hat{\tau}(w) - \mathbb{E}[x | w]$ , with 95% AR confidence bands.

In the temperature-GDP relationship, we find that estimated  $\hat{\tau}(w)$  look no different from the pooled estimator, and Anderson-Rubin confidence bands are wide:  $\hat{\tau}(w)$  is only statistically distinguishable from  $\mathbb{E}[x | w]$  for very poor countries. Including the interaction  $x_i^2 w_i$  significantly widens the confidence sets for  $\hat{\tau}(w)$ . However, for mortality, including  $w_i$  introduces meaningful  $\tau_i$  heterogeneity that is statistically distinguishable from  $\mathbb{E}[x | w]$  (approximately the null) and the pooled estimator for poorer countries. In both datasets, we find that including an indicator for ‘poor’ ( $w_i$  in the bottom tercile) performs as well as including a linear interaction in  $w_i$ . The delta method undercovers or provides incorrect coverage even for well-identified  $\hat{\tau}$  using the interaction with the ‘poor’ dummy, failing to capture the skewness of the Anderson-Rubin confidence sets.

## 3.2 Bayesian approaches

A complementary approach is to treat  $\tau_i$  as drawn from a distribution and shrink unit-specific estimates toward a common prior. In this subsection, we consider two implementations of this ap-

---

<sup>9</sup>Modeling the slope as a linear function means including  $w_i x_i^2$  as a regressor.

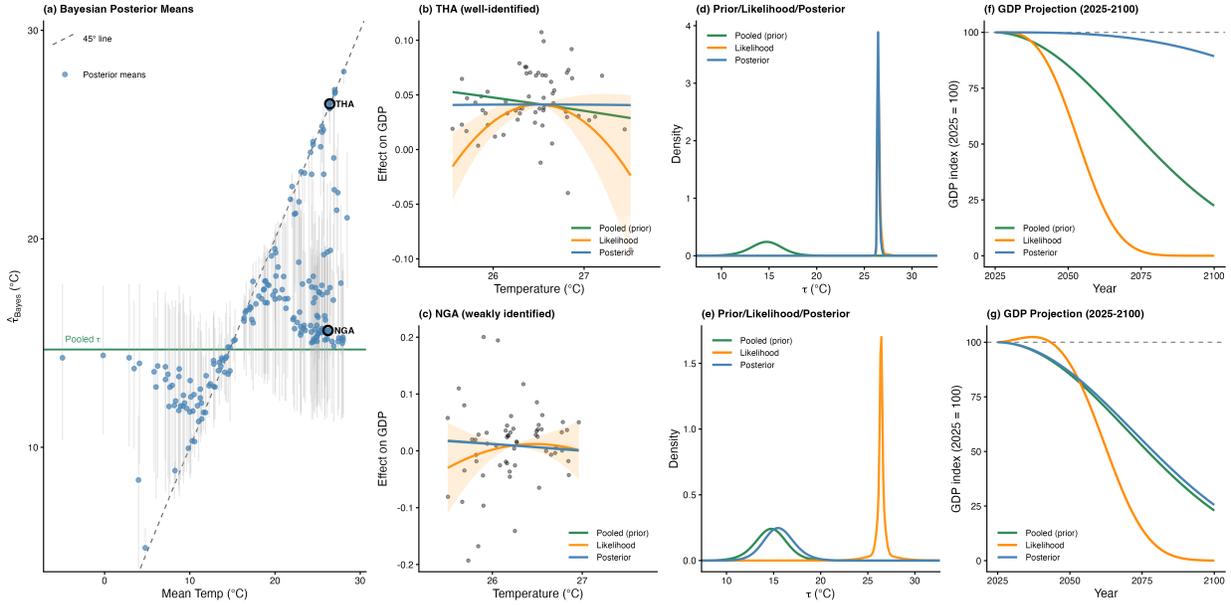
proach. We focus on a Bayesian approach in which we shrink towards the pooled  $\tau$  baseline and a modern (NPML) empirical Bayes approach in which we non-parametrically estimate the prior.

### Fixed prior: the pooled sampling distribution

We begin with a simple approach in which we treat the pooled (across all units) estimate of  $\tau$  as a prior. We estimate the pooled model with country and year fixed effects, draw 50,000 parametric bootstrap samples of  $(\hat{\beta}_1, \hat{\beta}_2)$  from the asymptotic joint normal distribution, compute  $\hat{\tau} = -\frac{1}{2}\hat{\beta}_1/\hat{\beta}_2$  for each draw, and smooth the resulting sampling distribution using a Gaussian mixture. This smoothed density serves as a common prior for each unit’s  $\tau_i$ . We use this bootstrapped prior, as opposed to the analytic Hinkley, as the data do not match the assumptions of Proposition 2.1 exactly: the errors are not homoskedastic (we use heteroskedasticity-consistent standard errors) and time-series temperature data is non-normal. Regardless, the prior resembles a near-normal Hinkley (here,  $|\theta| \approx 4.5$ ). For each country, we similarly bootstrap from the country-level joint coefficient distribution to approximate the Hinkley-like likelihood, smooth it with a Gaussian mixture, and obtain the posterior by multiplying prior and likelihood on a fine grid of candidate peaks, normalizing numerically. We report posterior means and 95% credible intervals.<sup>10</sup> We additionally shrink  $\hat{\gamma}_i$  estimates towards the pooled estimate  $\gamma$ . As  $\gamma_i$  is asymptotically normal, we use normal-normal shrinkage.

Figure 6 reports the results for the temperature–GDP relationship. Panel (a) plots posterior means against each country’s mean temperature. The estimator shrinks weakly identified countries—especially those with extreme temperatures—toward the pooled estimate, while well-identified countries retain peaks closer to their OLS estimates. The 45-degree line (the approximate null under weak curvature, where  $\hat{\tau} \approx \bar{x}_i$ ) and the pooled  $\tau$  line provide visual benchmarks: posterior means cluster nearer the pooled line than the 45-degree line, indicating effective shrinkage away from the spurious correlation. Appendix Figure A11 visualizes the shrinkage country by country: for each country ordered by mean temperature, it plots the OLS and posterior mean  $\hat{\tau}_i$ , connected by segments indicating the direction and magnitude of the shift toward the pooled estimate. Panels (b)–(c) illustrate the mechanism for two example countries. For Thailand (well-identified,  $|t_{\beta_2}| > 2$ ), the country-level likelihood is concentrated and dominates the prior, so the posterior closely tracks the likelihood (panel d). For Nigeria (weakly identified,  $|t_{\beta_2}| < 1.5$ ), the likelihood is diffuse and the

<sup>10</sup>We estimate 7-component mixtures and fit the empirical density well. However, the ability to take means here is not an artifact of Gaussian smoothing. The regularization comes from the Bayesian updating itself: while Hinkley distributions have no finite moments, the product of two Hinkleys has tails that behave like  $x^{-4}$ , which allows for the existence of a first (and second) moment.



**Figure 6: Bayesian shrinkage toward the pooled estimate in the temperature–GDP relationship.**

*Panel (a):* posterior mean  $\hat{\tau}_{\text{Bayes}}$  against country mean temperature, with 95% posterior credible intervals (grey bars). The dashed line is the 45-degree line ( $\hat{\tau} = \bar{x}$ , the approximate null under weak curvature); the solid green line is the pooled  $\hat{\tau}$ . Posterior means cluster near the pooled line rather than the 45-degree line, indicating effective shrinkage away from the spurious correlation. Circled points are the example countries in panels (b)–(g). *Panels (b)–(c):* scatter of GDP growth against temperature with pooled (green) and country-specific (orange) quadratic fits for a well-identified country (Thailand,  $|t_{\beta_2}| > 2$ ) and a weakly identified country (Nigeria,  $|t_{\beta_2}| < 1.5$ ). *Panels (d)–(e):* prior (pooled sampling distribution, green), likelihood (country-specific Hinkley likelihood, orange), and posterior (blue) densities for the same two countries. For Thailand, the likelihood dominates; for Nigeria, the prior dominates. *Panels (f)–(g):* projected GDP index (2025 = 100) under uniform  $+3^\circ\text{C}$  warming from 2025–2100, comparing posterior (blue) and pooled (green) damage estimates.

prior dominates, pulling the posterior toward the pooled estimate (panel e). Panels (f)–(g) translate these estimates into GDP projections under a uniform  $+3^\circ\text{C}$  warming scenario from 2025–2100: for well-identified countries, the posterior and pooled projections diverge meaningfully, reflecting heterogeneity in peaks; for weakly identified countries, the posterior projection collapses toward the pooled estimate. GDP projections based solely on the country-specific OLS estimates imply implausibly large damages that are considerably attenuated through Bayesian shrinkage.

We repeat this exercise for the temperature–mortality relationship, estimating admin-1-level heterogeneity in the minimum-mortality temperature for the 65+ age group. Appendix Figure A12 reports the results. The same pattern emerges: well-identified admin-1 units retain posterior means close to their OLS estimates, while weakly identified units are pulled toward the pooled minimum-mortality temperature, attenuating the spurious correlation between  $\hat{\tau}_i$  and local climate normals.

## Data-driven prior: nonparametric empirical Bayes

The fixed-prior approach imposes a specific view about the distribution of  $\tau_i$  across units. An alternative is to let the data determine the prior using empirical Bayes methods.<sup>11</sup> Standard normal-likelihood, normal-prior models are inappropriate here because the sampling distribution of  $\hat{\tau}_i$  is heavily skewed. Instead, we use nonparametric maximum likelihood estimation (NPMLE) of the prior  $G$ , modeling it as a discrete mixture over a fine mesh of candidate peaks:  $G(\tau) = \sum_{j=1}^J \phi_j \delta(\tau - \tau_j)$ , where the mixture weights  $\phi_j \geq 0$  are recovered by solving a convex program. Appendix E details this approach.

The estimated NPMLE prior is essentially uninformative, placing nearly uniform mass across the observed temperature support, for both GDP and mortality. Because each unit’s Hinkley likelihood is so diffuse under weak curvature, the data provide little information about the cross-sectional distribution of  $\tau_i$ . The NPMLE cannot distinguish genuine heterogeneity in peaks from the noise inherent in weakly identified ratio estimators. Consequently, the posterior means under the NPMLE prior exhibit minimal shrinkage, and the median bias toward  $\mu_i$  persists.

The contrast between the two Bayesian approaches underscores the identification–estimation trade-off at the heart of this paper. When the researcher is willing to impose structure, such as a fixed prior that pools toward a common peak, shrinkage is effective and produces economically meaningful estimates of partial adaptation. When the prior is left to the data, the weak-curvature likelihoods are too diffuse to learn from, and the bias remains. This parallels the finding in Section 3.1: covariates help when they genuinely predict  $\tau_i$ , but of course cannot overcome the fundamental precision problem created by weak curvature.

## 4 Conclusion

This paper has shown that turning-point estimators in quadratic regressions inherit the pathologies of ratio statistics when curvature is weakly identified: fat-tailed sampling distributions, median bias toward the regressor mean, and conventional confidence intervals that dramatically understate uncertainty. The strength of these pathologies is governed by a single parameter,  $\theta$ , which is consistently estimated by  $t_{\beta_2}$ , the  $t$ -statistic on the quadratic term. These findings are not artifacts of

---

<sup>11</sup>There are many interpretations of these estimators: e.g., to estimate a hierarchical model, to solve for decision rules that minimize compound risk, or to mimic an oracle. See Efron (2019) and Koenker and Gu (2024) or Walters (2024) for exposition.

parametric assumptions: they persist under heteroskedasticity, non-normal regressors, and semi-parametric estimation, for both quadratic and other DGPs.

Our results carry direct implications for applied work. At a minimum, researchers should report  $t_{\beta_2}$  alongside any turning-point estimate as a diagnostic for identification strength, and should use identification-robust confidence sets (which we construct by inverting tests of the first-order condition) rather than delta-method or bootstrapped confidence intervals. When heterogeneous optima are of interest, shrinkage toward a pooled estimate can reduce median bias, but only when the researcher is willing to impose structure on the prior; purely data-driven approaches cannot overcome the fundamental precision problem created by weak curvature. These recommendations apply not only in climate economics but wherever turning-point estimation arises: optimal class size, optimal city size, the debt-growth relationship, or the competition-innovation relationship of Aghion et al. (2005) (see discussion in Appendix C.2).

In the temperature–GDP application, the pooled relationship is well identified ( $|t_{\beta_2}| \approx 4.7$ ), and the pooled optimal temperature is relatively precisely estimated. But once curvature is allowed to vary across countries or subgroups, most units have weakly identified curvature, and the striking cross-sectional pattern in which hotter countries appear to have higher optimal temperatures (Figure 1) is statistically indistinguishable from a placebo exercise in which temperature has no effect on GDP. Interacting the peak with GDP per capita does not resolve this: estimated  $\hat{\tau}(w)$  cannot be distinguished from the pooled estimate for most of the income distribution. Bayesian shrinkage toward the pooled peak is more successful, producing estimates consistent with partial adaptation, but the data lack the power to determine the degree of heterogeneity without imposing prior structure. The temperature–mortality application tells a broadly similar story, with one important additional finding: even the pooled minimum-mortality temperature is only moderately identified ( $|t_{\beta_2}| \approx 2.3$ ), so the Anderson–Rubin confidence set for the pooled optimum spans much of the temperature support. This is a cautionary finding for the large environmental epidemiology literature that estimates location-specific minimum-mortality temperatures and interprets their proximity to local climate as evidence of adaptation.

Several directions for future research remain open. The Bayesian shrinkage estimator we outlined in Section 3.2 uses a simple fixed prior derived from the pooled model; richer hierarchical specifications may improve performance, and covariates and Bayesian approaches may be combined. Our empirical applications to temperature–GDP and temperature–mortality are intended to illustrate the statistical issues rather than to deliver definitive damage estimates. More broadly, developing formal decision rules for when heterogeneous turning-point estimation is feasible, analogous to pre-test or first-stage  $F$ -statistic thresholds in the weak-instruments literature, would provide useful guidance

for applied researchers. In the case of quadratic regression, we have shown that this object is the  $t$ -statistic on the quadratic term. It would be valuable to define the analogous object for semi- and non-parametric estimators. Such tools would be valuable not only in climate economics but in any empirical setting where turning points are central estimands, since the weak-curvature pathologies we document apply whenever the curvature term is imprecisely estimated. Finally, our results suggest that the climate impacts literature may benefit from greater formalism about what constitutes evidence for climate adaptation. Throughout this paper, our definition of adaptation more closely resembles acclimatisation (the optimal temperature is the average) than, e.g., resilience (reduced sensitivity to variance). Formalizing the hypotheses being tested, the estimands of interest, and the identification requirements for distinguishing adaptation from statistical artifacts would complement the substantive insights of this literature and strengthen the empirical foundations for climate impacts research.

## References

- Aghion, Philippe et al. (2005). “Competition and innovation: An inverted-U relationship”. In: *The Quarterly Journal of Economics* 120.2, pp. 701–728.
- Anderson, T. W. and Herman Rubin (1949). “Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations”. In: *The Annals of Mathematical Statistics* 20.1, pp. 46–63. DOI: [10.1214/aoms/1177730090](https://doi.org/10.1214/aoms/1177730090).
- Andrews, Isaiah and Timothy B Armstrong (2017). “Unbiased instrumental variables estimation under known first-stage sign”. In: *Quantitative Economics* 8.2, pp. 479–503.
- Angrist, Joshua and Michal Kolesár (2024). “One instrument to rule them all: The bias and coverage of just-ID IV”. In: *Journal of Econometrics* 240.2, p. 105398.
- Au, Chun-Chung and J Vernon Henderson (2006). “Are Chinese cities too small?” In: *The Review of Economic Studies* 73.3, pp. 549–576.
- Auffhammer, Maximilian (2022). “Climate Adaptive Response Estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption”. In: *Journal of Environmental Economics and Management* 114, p. 102669.
- Bernard, Jean-Thomas et al. (2015). “Environmental Kuznets curve: Tipping points, uncertainty and weak identification”. In: *Environmental and Resource Economics* 60.2, pp. 285–315.
- Bourany, Thomas (2025). “The optimal design of climate agreements”. Working paper.
- Burke, Marshall, Solomon Hsiang, and Edward Miguel (2015). “Global non-linear effect of temperature on economic production”. In: *Nature* 527.7577, pp. 235–239.
- Carleton, Tamma et al. (2022). “Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits”. In: *The Quarterly Journal of Economics* 137.4, pp. 2037–2105.
- Checherita-Westphal, Cristina and Philipp Rother (2012). “The impact of high government debt on economic growth and its channels: An empirical investigation for the euro area”. In: *European Economic Review* 56.7, pp. 1392–1405.
- Conte, Bruno et al. (2021). “Local sectoral specialization in a warming world”. In: *Journal of Economic Geography* 21.4, pp. 493–530.
- Cruz, José-Luis and Esteban Rossi-Hansberg (2024). “The economic geography of global warming”. In: *Review of Economic Studies* 91.2, pp. 899–939.
- Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2012). “Temperature shocks and economic growth: Evidence from the last half century”. In: *American Economic Journal: Macroeconomics* 4.3, pp. 66–95.

- Dempster, Arthur P, Nan M Laird, and Donald B Rubin (1977). “Maximum likelihood from incomplete data via the EM algorithm”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 39.1, pp. 1–22.
- Eberhardt, Markus and Andrea F Presbitero (2015). “Public debt and growth: Heterogeneity and non-linearity”. In: *Journal of International Economics* 97.1, pp. 45–58.
- Efron, Bradley (2019). “Bayes, oracle Bayes and empirical Bayes”. In: *Statistical Science* 34.2, pp. 177–201.
- Fieller, Edgar C (1932). “The distribution of the index in a normal bivariate population”. In: *Biometrika* 24.3/4, pp. 428–440.
- (1954). “Some problems in interval estimation”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 16.2, pp. 175–185.
- Gasparrini, Antonio et al. (2015). “Mortality risk attributable to high and low ambient temperature: a multicountry observational study”. In: *The Lancet* 386.9991, pp. 369–375.
- Geary, Robert Charles (1930). “The frequency distribution of the quotient of two normal variates”. In: *Journal of the Royal Statistical Society* 93.3, pp. 442–446.
- Goh, Gyuhyeong et al. (2025). “Bayesian space–time varying coefficient modeling for climate econometrics: A spatial–temporal Gaussian process approach”. In: *Journal of Econometrics*, p. 106075.
- Hashmi, Aamir Rafique (2013). “Competition and innovation: The inverted-U relationship revisited”. In: *Review of Economics and Statistics* 95.5, pp. 1653–1668.
- Heutel, Garth, Nolan H Miller, and David Molitor (2021). “Adaptation and the mortality effects of temperature across US climate regions”. In: *Review of Economics and Statistics* 103.4, pp. 740–753.
- Hinkley, David V (1969). “On the ratio of two correlated normal random variables”. In: *Biometrika* 56.3, pp. 635–639.
- Hirschberg, JG, JN Lye, and DJ Slottje (2008). “Inferential methods for elasticity estimates”. In: *Journal of Econometrics* 147.2, pp. 299–315.
- Hogan, Dylan and Wolfram Schlenker (2024). “Chapter 2 - Empirical approaches to climate change impact quantification”. In: ed. by Lint Barrage and Solomon Hsiang. Vol. 1. *Handbook of the Economics of Climate Change*. North-Holland, pp. 53–111. DOI: [10.1016/bs.hesec.2024.10.006](https://doi.org/10.1016/bs.hesec.2024.10.006).
- Hsiang, Solomon (2025). *The Global Economic Impact of Climate Change: An Empirical Perspective*. Working Paper 34357. National Bureau of Economic Research. DOI: [10.3386/w34357](https://doi.org/10.3386/w34357).
- Kedagni, Desire et al. (2021). “Does class size matter? How, and at what cost?” In: *European Economic Review* 133, p. 103664.

- Kim, Jeankyung and David Pollard (1990). “Cube root asymptotics”. In: *The Annals of Statistics*, pp. 191–219.
- Kim, Youngseok et al. (2020). “A fast algorithm for maximum likelihood estimation of mixture proportions using sequential quadratic programming”. In: *Journal of Computational and Graphical Statistics* 29.2, pp. 261–273.
- Koenker, Roger and Jiaying Gu (2024). “Empirical Bayes for the reluctant frequentist”. In: *arXiv preprint arXiv:2404.03422*.
- Koenker, Roger and Ivan Mizera (2014). “Convex optimization, shape constraints, compound decisions, and empirical Bayes rules”. In: *Journal of the American Statistical Association* 109.506, pp. 674–685.
- Kotlikoff, Laurence et al. (2021). “Making Carbon Taxation a Generational Win Win”. In: *International Economic Review* 62.1, pp. 3–46. doi: [10.1111/iere.12483](https://doi.org/10.1111/iere.12483).
- Krusell, Per and Anthony A. Smith (2022). *Climate change around the world*. Working Paper.
- Lemoine, Derek, Catherine Hausman, and Jeffrey G Shrader (2025). *A Guide to Climate Damages*. Working Paper 34348. National Bureau of Economic Research. doi: [10.3386/w34348](https://doi.org/10.3386/w34348).
- Lind, Jo Thori and Halvor Mehlum (2010). “With or without U? The appropriate test for a U-shaped relationship”. In: *Oxford bulletin of economics and statistics* 72.1, pp. 109–118.
- Linsenmeier, Manuel (2024). “Global variation in the preferred temperature for recreational outdoor activity”. In: *Journal of Environmental Economics and Management* 127, p. 103032.
- Marsaglia, George (1965). “Ratios of normal variables and ratios of sums of uniform variables”. In: *Journal of the American Statistical Association* 60.309, pp. 193–204.
- Newell, Richard G, Brian C Prest, and Steven E Sexton (2021). “The GDP-temperature relationship: implications for climate change damages”. In: *Journal of Environmental Economics and Management* 108, p. 102445.
- Pattillo, Catherine, Luca Antonio Ricci, et al. (2011). “External debt and growth”. In: *Review of Economics and Institutions* 2.3, p. 30.
- Plassmann, Florenz and Neha Khanna (2007). “Assessing the precision of turning point estimates in polynomial regression functions”. In: *Econometric Reviews* 26.5, pp. 503–528.
- Press, S James (1969). “The t-ratio distribution”. In: *Journal of the American Statistical Association* 64.325, pp. 242–252.
- Self, Steven G and Kung-Yee Liang (1987). “Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions”. In: *Journal of the American Statistical Association* 82.398, pp. 605–610.
- Silva, JMC Santos and Silvana Tenreiro (2006). “The log of gravity”. In: *The Review of Economics and Statistics* 88.4, pp. 641–658.

- Staiger, Douglas and James H. Stock (1997). “Instrumental Variables Regression with Weak Instruments”. In: *Econometrica* 65.3, pp. 557–586. doi: [10.2307/2171753](https://doi.org/10.2307/2171753).
- Tobías, Aurelio et al. (2021). “Geographical variations of the minimum mortality temperature at a global scale: a multicountry study”. In: *Environmental Epidemiology* 5.5, e169.
- Walters, Christopher (2024). “Empirical Bayes methods in labor economics”. In: *Handbook of Labor Economics*. Vol. 5. Elsevier, pp. 183–260.
- Yin, Qian et al. (2019). “Mapping the increased minimum mortality temperatures in the context of global climate change”. In: *Nature Communications* 10.1, p. 4640.

# Online Appendix

<b>A Data appendix</b>	<b>28</b>
A.1 GDP . . . . .	28
A.2 Mortality . . . . .	29
<b>B Additional figures and tables</b>	<b>31</b>
B.1 Tables . . . . .	43
<b>C Robust confidence sets: extensions</b>	<b>44</b>
C.1 Sign-restricted Anderson-Rubin (SAR) confidence sets . . . . .	44
C.2 Extension to Poisson estimators . . . . .	46
<b>D Proofs and derivations</b>	<b>49</b>
D.1 Proof of Proposition 2.1 . . . . .	49
D.2 Asymptotic distribution under local-to-zero curvature . . . . .	51
D.3 Derivation of the analytic Hinkley distribution . . . . .	53
D.4 Proof of Corollary 3.1 . . . . .	61
<b>E Empirical Bayes implementation details</b>	<b>65</b>

# A Data appendix

## A.1 GDP

### A.1.1 Data

We use an updated version of the panel dataset from Burke et al. (2015), which covers 166 countries observed annually from 1960 to 2019. The outcome variable is annual GDP growth,  $\Delta \ln y_{it}$ , computed as the first difference of log real GDP. Temperature is measured as the population-weighted annual average in degrees Celsius, constructed from gridded climate reanalysis data (ERA5). The panel is unbalanced; after dropping observations with missing values for  $\Delta \ln y_{it}$ , temperature, and temperature squared, the estimation sample contains approximately 8,800 country-year observations.

### A.1.2 Estimation

We estimate a quadratic temperature–growth relationship in first differences:

$$\Delta \ln y_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \alpha_i + \lambda_t + \gamma_i t + u_{it}, \quad (\text{A1})$$

where  $\Delta T_{it}$  and  $\Delta T_{it}^2$  denote first differences of annual average temperature and its square,  $\alpha_i$  are country fixed effects,  $\lambda_t$  are year fixed effects, and  $\gamma_i t$  are country-specific linear time trends. The implied turning point is  $\hat{\tau} = -\hat{\beta}_1 / (2\hat{\beta}_2)$ .

Our specification differs from Burke et al. (2015) in two respects. First, we do not control for precipitation here (previous studies have found negligible impacts of doing so). Second, we include country-specific linear rather than quadratic time trends, to reduce the risk of overfitting low-frequency variation in the outcome.

The pooled model (Figures 1 and 5) estimates common coefficients  $(\beta_1, \beta_2)$  across all countries. The heterogeneous model (Figure 1) allows fully country-specific  $(\beta_{1i}, \beta_{2i})$  by interacting the temperature regressors with country indicators.

### A.1.3 Null simulation

The placebo exercise in Figure 1 and Appendix Figure A1 generates data under the null hypothesis that temperature has no effect on GDP growth, while matching the empirical panel structure, residual variation, and temperature process. We proceed in three steps.

*Step 1: Calibrate the DGP.* We regress  $\Delta \ln y_{it}$  on country fixed effects  $\alpha_i$ , year fixed effects  $\lambda_t$ , and country-specific linear time trends  $\gamma_i t$ , omitting all temperature terms. From this regression we extract  $\hat{\alpha}_i$ ,  $\hat{\lambda}_t$ , and  $\hat{\gamma}_i$ , as well as the country-specific residual standard deviation  $\hat{\sigma}_i = \text{sd}(\hat{u}_{it})$ .

*Step 2: Simulate under the null.* For each country-year observation in the estimation sample, we draw

$$\Delta \ln y_{it}^{\text{sim}} = \hat{\alpha}_i + \hat{\lambda}_t + \hat{\gamma}_i t + u_{it}^{\text{sim}}, \quad u_{it}^{\text{sim}} \sim \mathcal{N}(0, \hat{\sigma}_i^2).$$

The simulated outcome preserves cross-country differences in growth levels, global time trends, and country-specific trends, but contains no temperature effect by construction. Temperature variables ( $\Delta T_{it}$ ,  $\Delta T_{it}^2$ ) are held at their empirical values.

*Step 3: Estimate.* We fit the country-level model of equation (A1) to the simulated data and compute  $\hat{\tau}_i^{\text{null}} = -\hat{\beta}_{1i}/(2\hat{\beta}_{2i})$  for each country. The resulting correlation between  $\hat{\tau}_i^{\text{null}}$  and mean temperature is nearly as large as in the real data, demonstrating that the adaptation-like pattern arises mechanically from weak identification.

## A.2 Mortality

### A.2.1 Data

We use the global mortality panel of Carleton et al. (2022), which links subnational death rates to temperature and precipitation across multiple countries. The unit of observation is an admin-2 region (e.g., county or district) observed annually. Temperature and precipitation are measured using the Global Meteorological Forcing Dataset (GMFD) and enter the regression as population-weighted annual averages. Death rates are winsorized at the 99th percentile within each country–age–group cell to limit the influence of outliers. We use the publicly available sample, which does not include all countries that were used in the original article.

We restrict the sample to the 65+ age group (agegroup = 3) since this is the sample for which impacts are most acute. Following the original paper, we restrict to years through 2010. The esti-

mation is population-weighted, with weights normalized within each year.

### A.2.2 Estimation

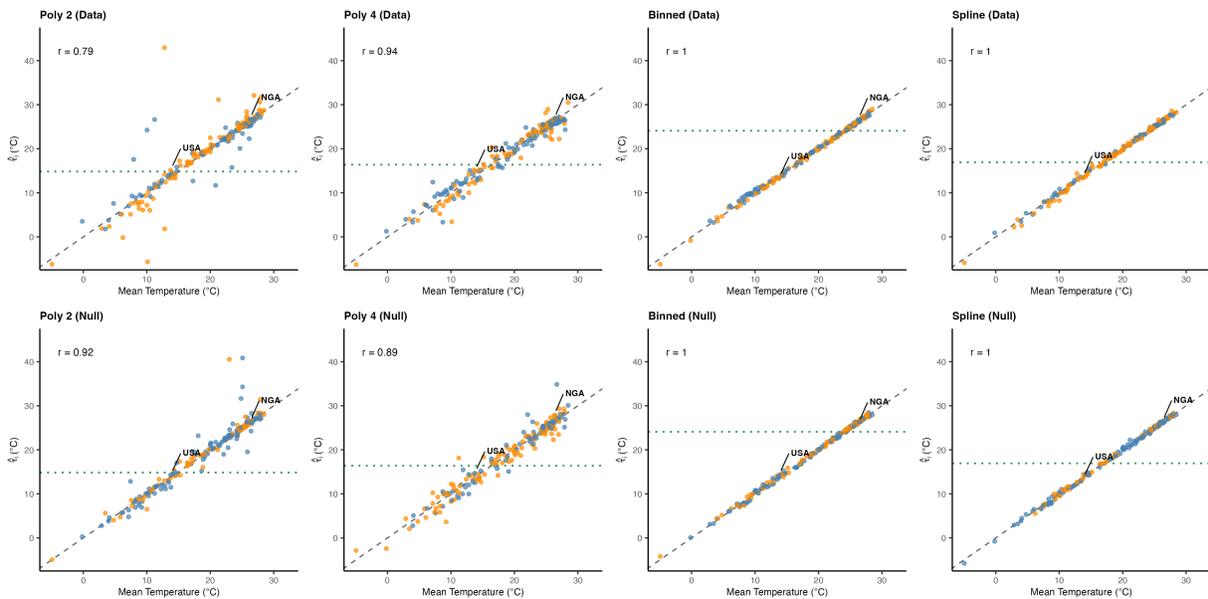
We estimate a quadratic temperature–mortality relationship:

$$\text{deathrate}_{it} = \beta_1 T_{it} + \beta_2 T_{it}^2 + \mathbf{X}'_{it} \boldsymbol{\delta} + \alpha_i + \lambda_{ct} + u_{it}, \quad (\text{A2})$$

where  $T_{it}$  is annual average temperature,  $\mathbf{X}_{it}$  includes country-specific precipitation controls (linear and quadratic),  $\alpha_i$  are admin-2  $\times$  age-group fixed effects, and  $\lambda_{ct}$  are country  $\times$  year  $\times$  age-group fixed effects. The latter absorb all country-year-age-specific shocks, so identification comes from within-country, cross-region temperature variation in a given year.

The pooled model (Figure 5) estimates common  $(\beta_1, \beta_2)$  across all regions, with population weights. The admin-1-level heterogeneous model (Appendix Figure A12) allows  $(\beta_{1j}, \beta_{2j})$  to vary across admin-1 units.

## B Additional figures and tables



**Figure A1: Spurious adaptation persists across functional-form specifications.**

Each column estimates country-specific turning points  $\hat{\tau}_i$  using a different specification of the temperature–GDP response function: quadratic (Poly 2), quartic polynomial (Poly 4), 8-bin indicator model (Binned), and natural spline with 2 interior knots (Spline). The top row uses actual GDP panel data; the bottom row uses a placebo simulation in which temperature has no effect on GDP. Points are colored by whether the turning point is a maximum (blue) or minimum (orange). The dashed line is the 45-degree line ( $\hat{\tau}_i = \mu_i$ ); the dotted horizontal line is the pooled optimum. Correlation coefficients ( $r$ ) between  $\hat{\tau}_i$  and mean temperature are reported in each panel. The spurious correlation between estimated optima and local climate persists regardless of specification, confirming that the pattern is driven by weak identification rather than functional-form choice.

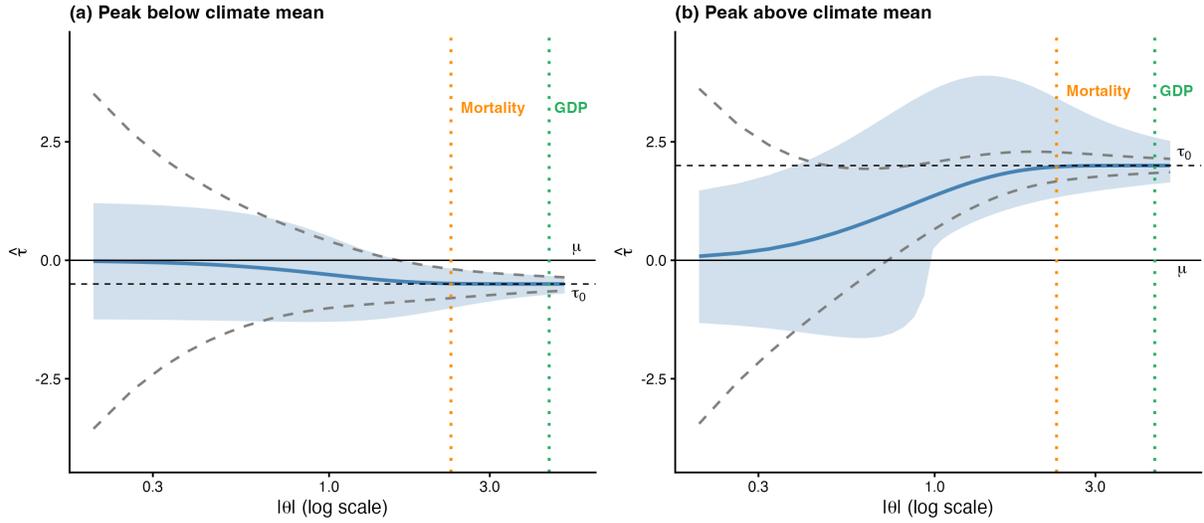


Figure A2: **Median bias and confidence-interval width as a function of curvature strength.** Each panel plots the Hinkley median (solid blue line) and 68% quantile interval (blue band, 16th–84th percentiles) of the turning-point estimator  $\hat{\tau}$  as a function of  $|\theta|$  (log scale). Grey dashed lines show the corresponding 68% delta-method confidence interval, centered at the Hinkley median. Panel (a):  $\tau_0 = \mu - \sigma_x/2$  (peak below climate mean). Panel (b):  $\tau_0 = \mu + 2\sigma_x$  (peak above climate mean). Horizontal reference lines mark  $\tau_0$  (dashed) and  $\mu$  (solid). Vertical dotted lines mark the empirical  $|\theta|$  for the GDP (green) and mortality (orange) applications. The Hinkley median transitions smoothly from  $\mu$  (under weak curvature) to  $\tau_0$  (under strong curvature), and the delta-method interval diverges from the Hinkley interval when  $|\theta|$  is small. Parameters:  $\mu = 0, \sigma_x = 1$ .

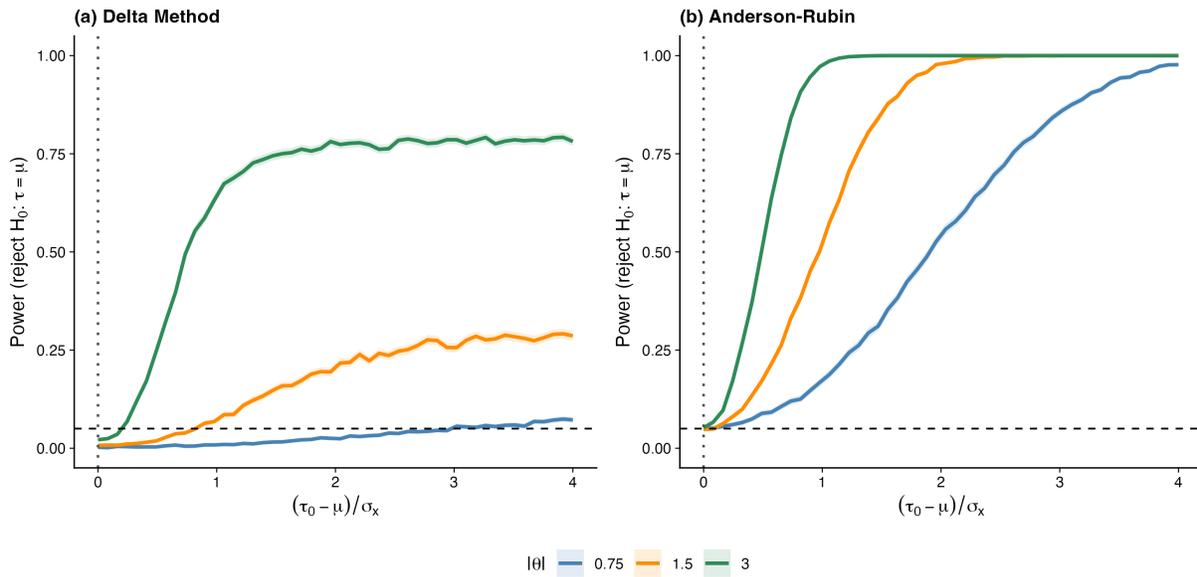
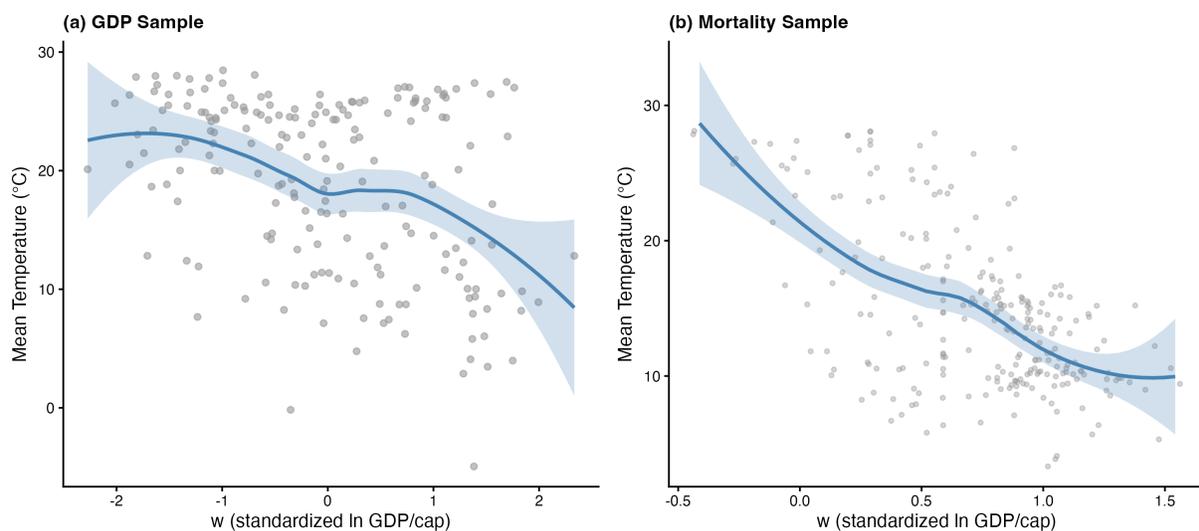


Figure A3: **Power of delta-method and Anderson–Rubin tests for  $H_0: \tau = \mu$ .**

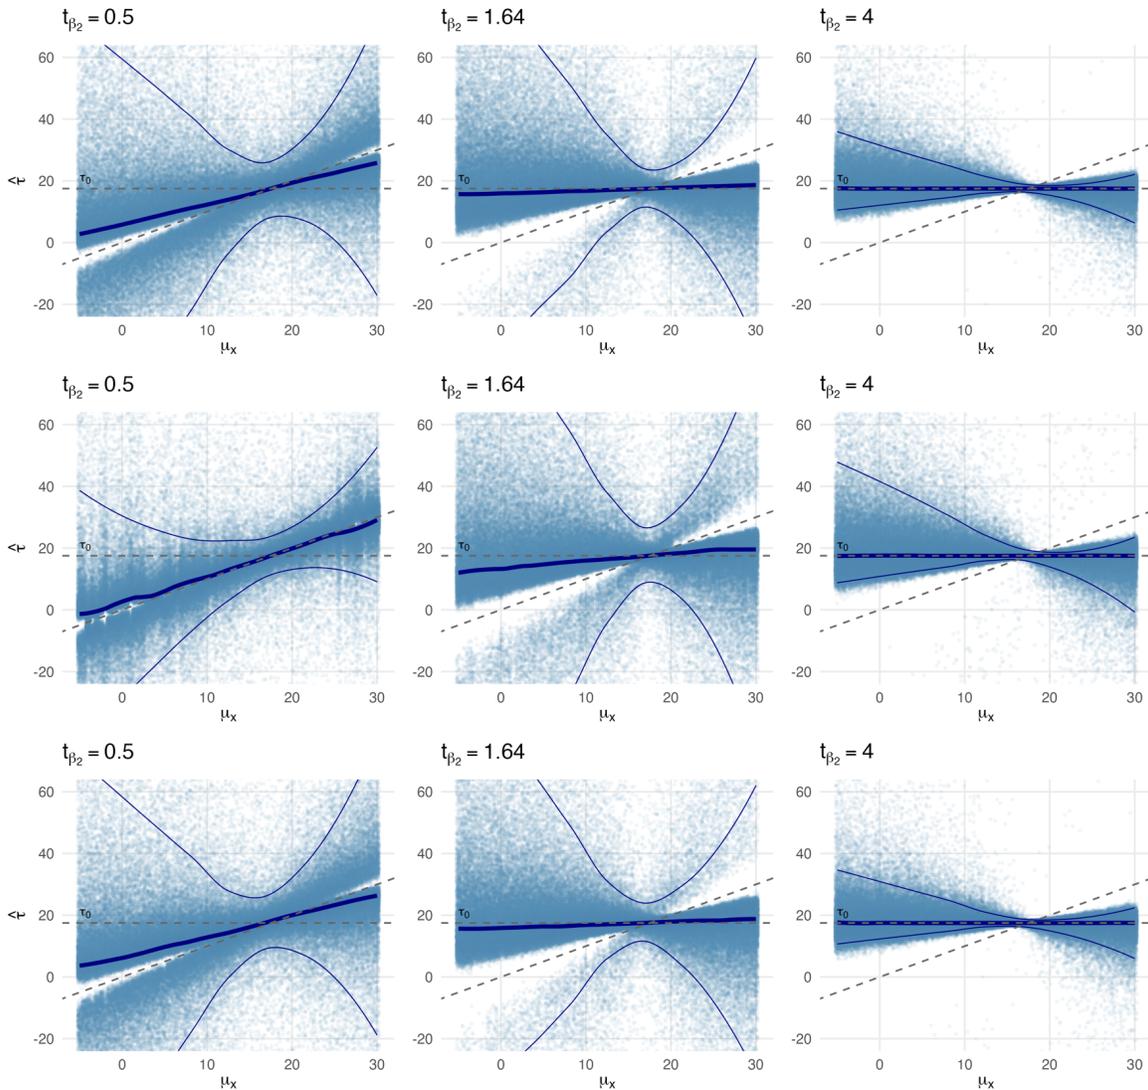
Each panel plots rejection rates (power) as a function of the true  $(\tau_0 - \mu) / \sigma_x$  for three values of curvature strength:  $|\theta| = 0.75$  (blue, weak),  $|\theta| = 1.5$  (orange, borderline), and  $|\theta| = 3$  (green, strong). Panel (a): delta-method test. Panel (b): Anderson–Rubin (AR) test. Shaded bands show 95% simulation confidence intervals around the estimated power curves. The dashed horizontal line marks the nominal 5% significance level. At  $(\tau_0 - \mu) / \sigma_x = 0$  (the null), the AR test is correctly sized across all  $\theta$  values, while the delta method is severely undersized. Under the alternative, the AR test has substantially higher power, particularly for weak and borderline curvature.  $n = 100$ ; 5,000 simulation replications.



**Figure A4: Average temperature versus economic development in both datasets.**

Panel (a): temperature–GDP sample. Each point is a country; the horizontal axis is standardized (mean zero, unit variance) average log GDP per capita. Panel (b): temperature–mortality sample. Each point is an admin-1 unit. Blue line: LOESS smoother with 95% pointwise confidence band for  $\mathbb{E}[x \mid w]$ . Richer units (higher  $w$ ) tend to have lower mean temperatures in both samples, establishing that the covariate  $w$  is informative about  $x$ —the key requirement for the projection approach in Section 3.1.

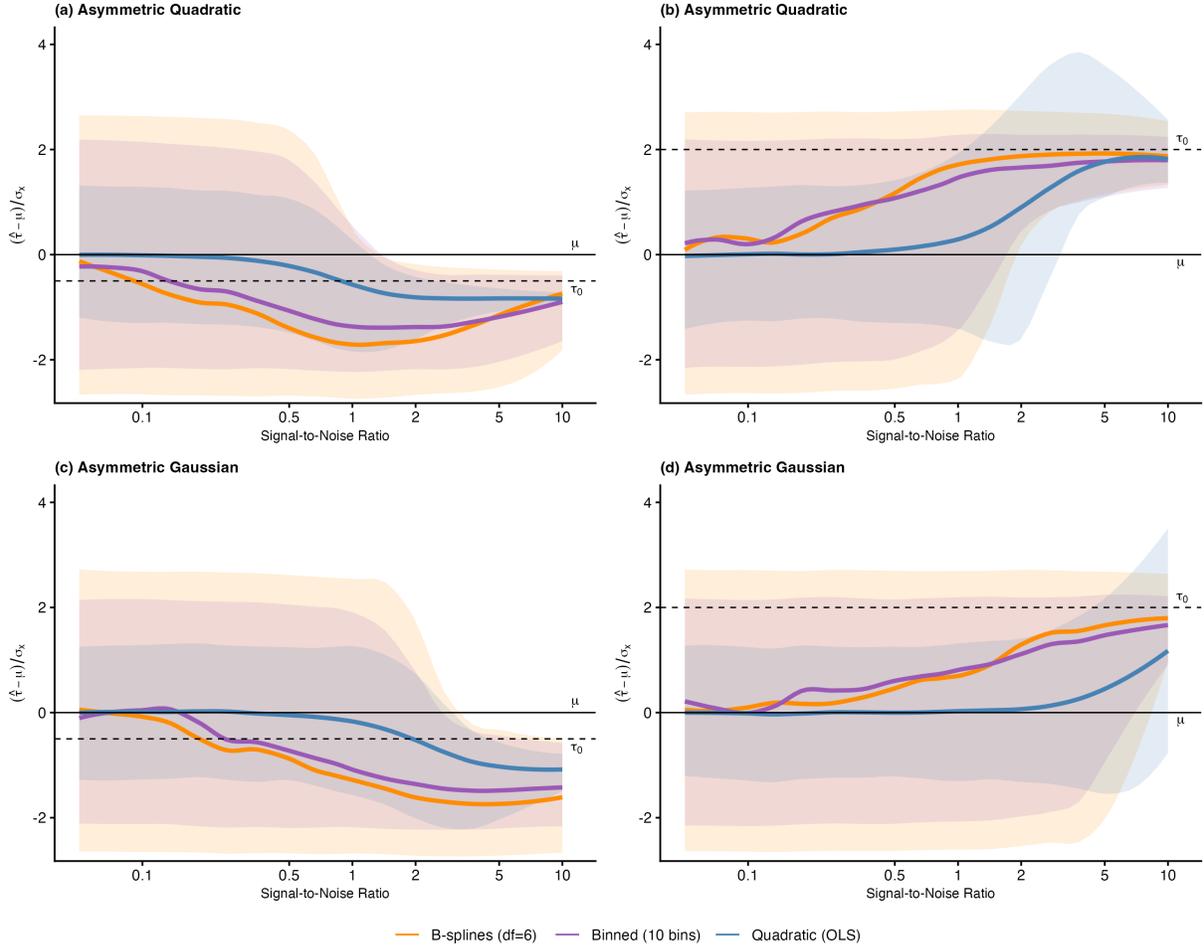
Sampling distribution of  $\hat{\tau}$  vs regressor mean  $\mu_x$



Notes: Each point is a simulation draw ( $n = 100$ ). Solid line: median. Dashed lines: 5th/95th percentiles. Horizontal dashed: true  $\tau_0$ . Diagonal dashed: 45-degree line.

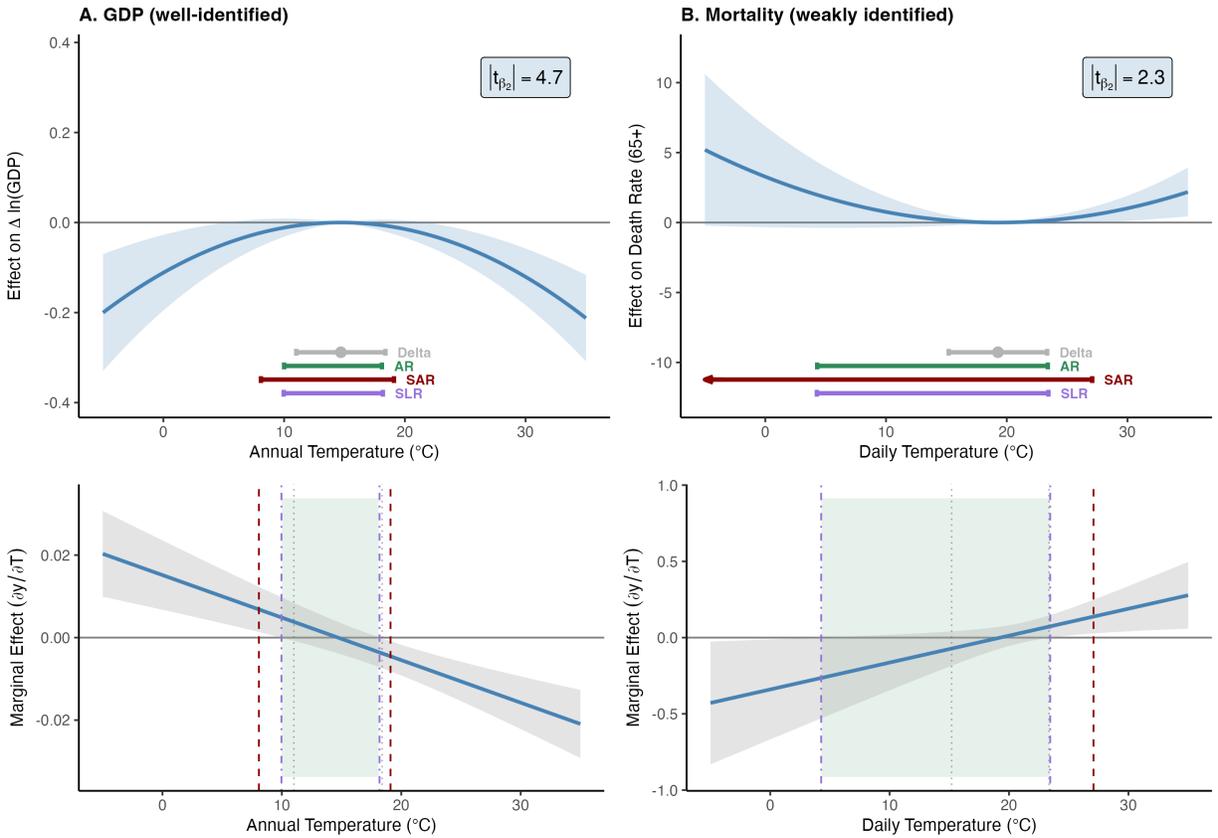
**Figure A5: Sampling distribution of  $\hat{\tau}$  under alternative data-generating processes.**

Each panel shows bootstrap samples of  $\hat{\tau}$  against the regressor mean  $\mu_x$  (transparent points), the bootstrap median (dark blue line), and 95% confidence band (thin blue lines). Columns vary the target curvature strength  $|t_{\beta_2}|$ ; rows vary the DGP. Top row: baseline DGP from Proposition 2.1 (Gaussian  $x$ , homoskedastic errors). Middle row: heteroskedastic errors with  $\text{Var}(u_i | x_i) = \sigma_u^2(1 + ((x - \mu)/\sigma_x)^2)$ . Bottom row:  $x_i$  drawn from a skewed normal distribution, holding the mean and variance fixed at  $\mu$  and  $\sigma_x^2$ . The dashed diagonal line is the 45-degree line ( $\hat{\tau} = \mu_x$ ); the horizontal dashed line marks the true  $\tau_0$ . The qualitative pathologies of Proposition 2.1—median bias toward  $\mu$ , fat tails, and bimodality under weak curvature—are robust to heteroskedasticity and non-normal regressors.  $n = 100$ ;  $\sigma_x = 2$ .



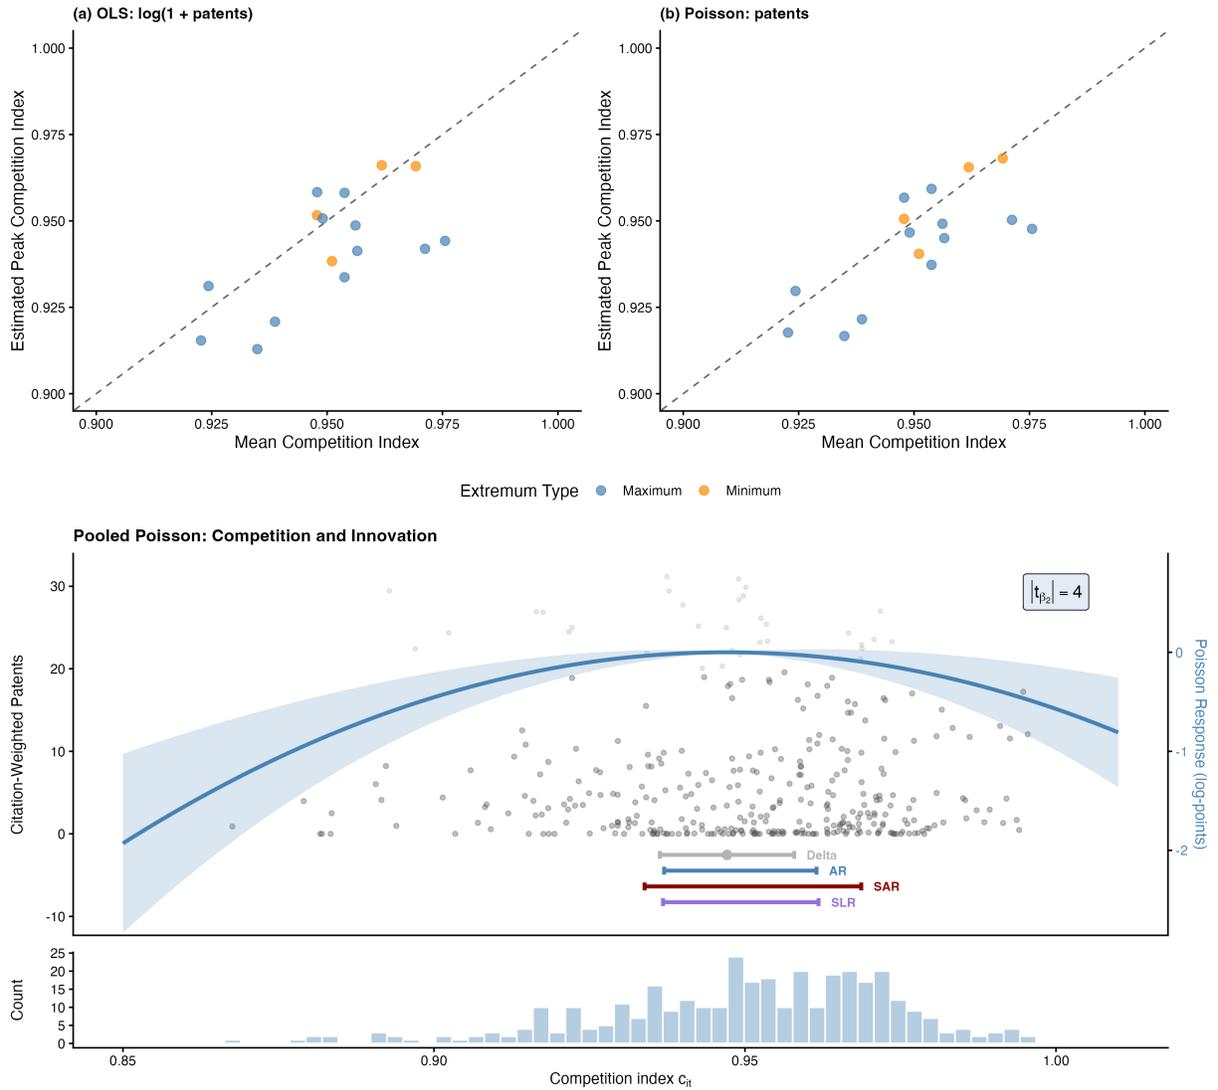
**Figure A6: Weak identification persists under misspecified DGPs.**

Comparison of quadratic (blue), cubic B-spline with 6 d.f (orange), and binned (10 bins, purple) estimators for  $\hat{\tau}$  under two non-quadratic response functions. Top row: asymmetric piecewise-quadratic (gentle cold side, steep hot side). Bottom row: asymmetric Gaussian (non-polynomial exponential peak with different spreads on each side). Left column:  $\tau_0 = \mu - \sigma_x/2$  (peak near climate). Right column:  $\tau_0 = \mu + 2\sigma_x$  (peak far from climate). Lines show the median of  $(\hat{\tau} - \mu)/\sigma_x$ ; shaded bands show 68% Monte Carlo confidence intervals. All three estimators converge toward  $\mu$  as the signal-to-noise ratio deteriorates, confirming that the weak-identification problem is fundamental rather than an artifact of functional-form assumptions. The semi-parametric estimators exhibit less median bias but substantially higher variance than OLS when the quadratic model is misspecified.  $n = 200$ ; 2,000 replications;  $\sigma_x \approx 2$ .

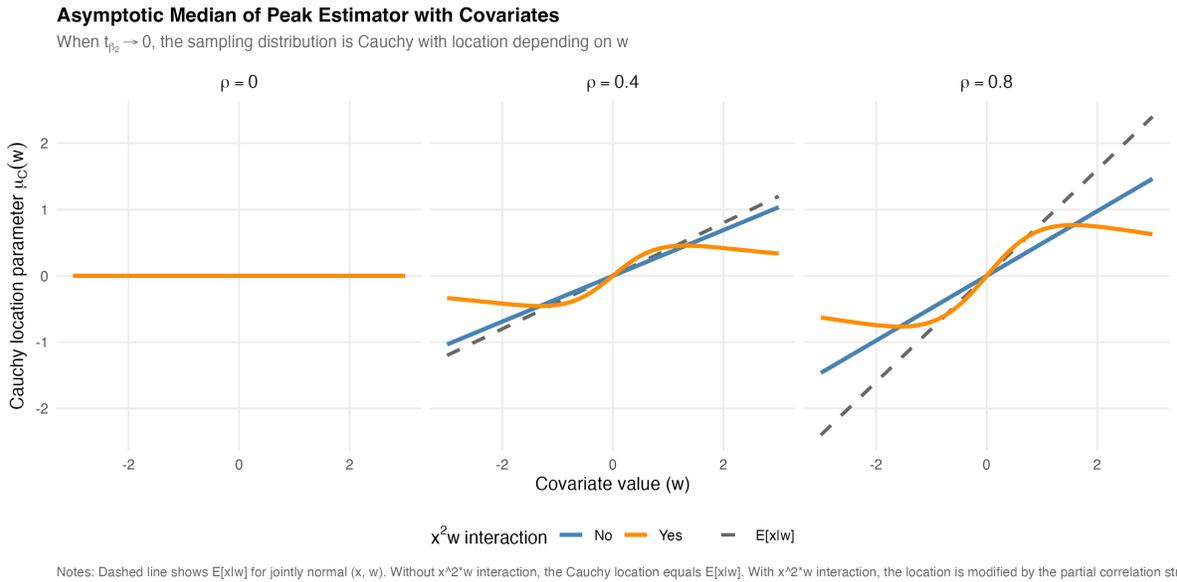


**Figure A7: Delta-method versus weak identification-robust inference on turning points.**

*Top row.* Estimated response functions for GDP growth (left, Burke et al., 2015) and 65 age-group mortality rates (right, Carleton et al., 2022), with pointwise 95% confidence bands. Horizontal bars report three 95% confidence sets for the turning point  $\tau_0$ : the delta-method interval (grey); the Anderson–Rubin interval (AR, green), which inverts a  $\chi_1^2$  test of  $H_0: \tau_0 = \tau$ ; and the sign-restricted likelihood-ratio interval (SLR, purple), which additionally imposes the maintained assumption that the response is concave (GDP) or convex (mortality), recovering a  $\chi_1^2$  critical value. *Bottom row.* AR inversion: for each hypothesized  $\tau_0$ , the marginal effect  $\hat{\beta}_1 + 2\hat{\beta}_2\tau_0$  and its delta-method confidence band (grey ribbon) are plotted; the AR confidence set (green shading) collects all  $\tau_0$  for which this band includes zero. Vertical lines mark the delta-method point estimate (grey dotted), the AR bounds (green shaded area), SAR bounds (maroon) and the SLR bounds (purple dot-dash). Further details on data and estimation in Appendix A.



**Figure A8: Competition and innovation: industry-level peaks and pooled response function.** Data from Aghion et al. (2005) (UK manufacturing industries, 1973-1994). The outcome is citation-weighted patents; the regressor of interest is the Lerner competition index  $c_{it}$  for industry  $i$  in year  $t$ . Panels (a) and (b) plot industry-level estimates of the peak competition index  $\hat{c}_i^*$  against mean industry  $i$  competition  $\bar{c}_i$ , estimated by OLS on  $\log(1 + \text{patents}_{it})$  and by Poisson pseudo-maximum likelihood, respectively. Blue: industries where  $\hat{\beta}_{2i} < 0$  (maximum); orange:  $\hat{\beta}_{2i} > 0$  (minimum). Dashed line: 45-degree benchmark. Panel (c) shows the pooled Poisson response function with industry and year fixed effects and HC1 standard errors. Grey points are the raw data (darker points lie between the 10th and 90th percentiles of the patent distribution, replicating the data in Figure 1 of Aghion et al., 2005). The blue curve and shaded region show the estimated response function and 95% confidence band. Horizontal bars at the bottom display the delta-method and Anderson–Rubin 95% confidence intervals for the peak  $c^*$ ; the AR interval is  $1.1\times$  wider.



**Figure A9: Cauchy median of  $\hat{\tau}(w)$  under covariate-augmented models.**

Each panel plots the Cauchy location parameter  $\check{\mu}(w)$ —the asymptotic median of  $\hat{\tau}(w)$  under the null of no temperature effect—for models in which both the peak  $\tau_i$  and curvature  $\gamma_i$  are linear functions of  $w_i$ . Panels vary the correlation  $\rho$  between  $x_i$  and  $w_i$ . Blue: model without  $x_i^2w_i$  interaction; orange: model including the  $x_i^2w_i$  interaction. Grey dashed line:  $E[x | w]$  for jointly normal  $(x, w)$ . When  $\rho = 0$ ,  $\check{\mu}(w) = \mu$  regardless of the model. When  $\rho > 0$ , the median bias tracks  $E[x | w]$  over most of the support, especially when the nonlinear interaction is included. Both  $x_i$  and  $w_i$  are standardized to zero mean and unit variance.

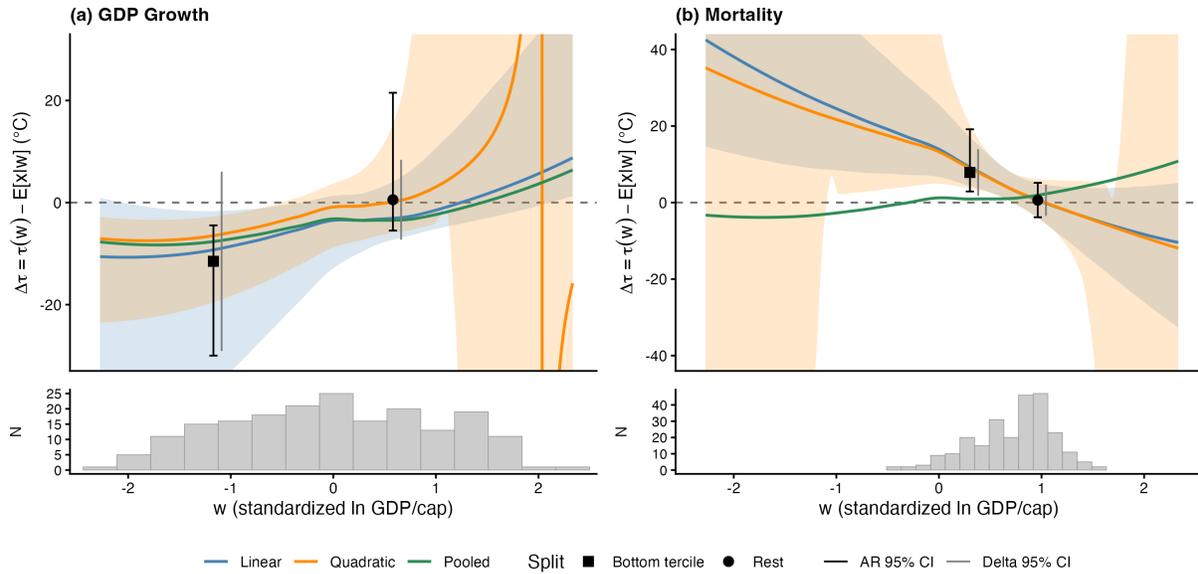
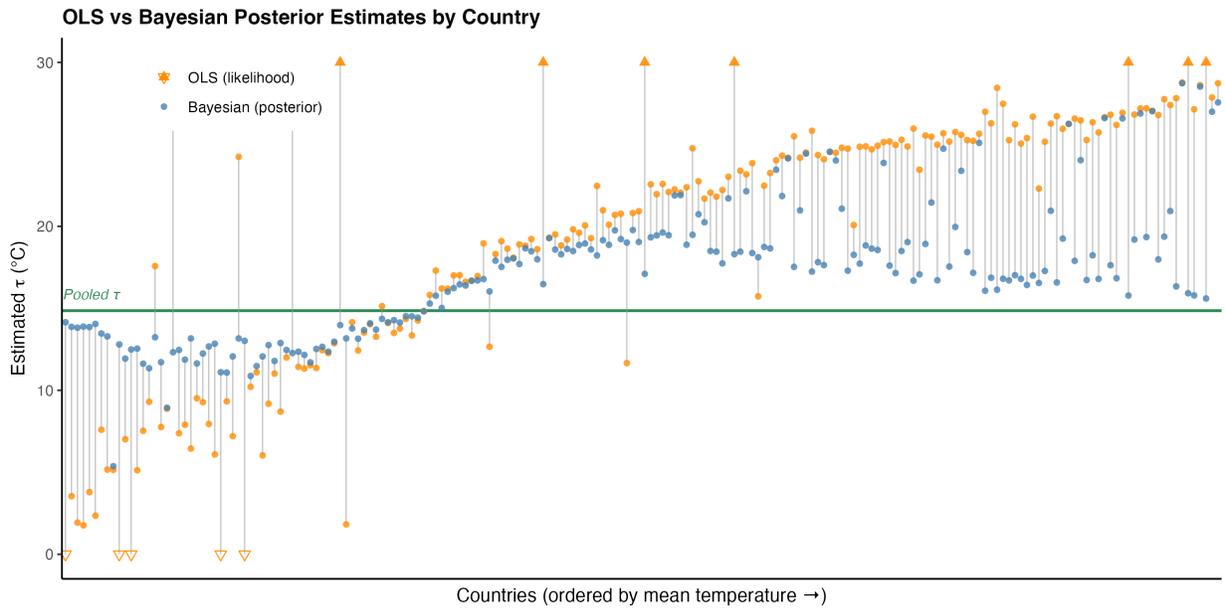


Figure A10: **Covariate-adjusted turning points and Anderson–Rubin confidence bands.**

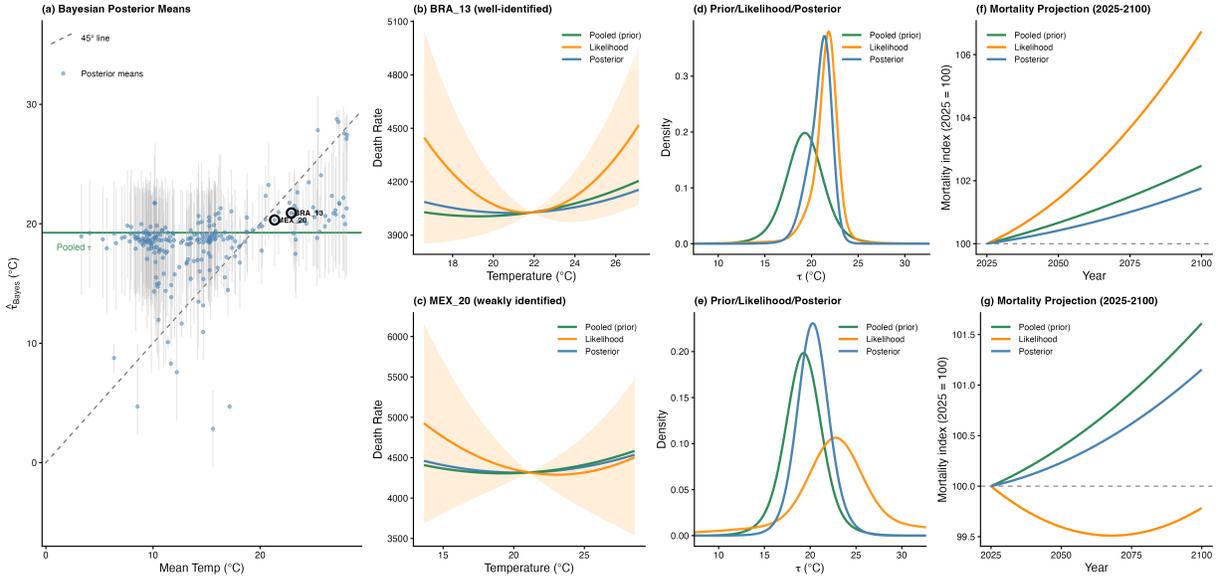
Each panel plots  $\hat{\tau}(w) - \mathbb{E}[x | w]$  as a function of standardized mean log GDP per capita ( $w$ ), so that zero corresponds to the approximate null under weak curvature. Panel (a): temperature–GDP data. Panel (b): temperature–mortality data. Lines show implied peaks from three specifications: pooled (no interaction, green), linear interaction  $w_i \times x_i$  only (blue), and full interaction including  $w_i \times x_i^2$  (orange). Shaded bands are 95% Anderson–Rubin confidence sets. Points and error bars show  $\hat{\tau}$  from a model replacing  $w_i$  with an indicator for the bottom tercile. Histograms show the estimation-sample distribution of  $w$ . For GDP, covariate-adjusted peaks are indistinguishable from the pooled estimate over most of the income distribution; for mortality, poorer regions show statistically significant departures. Regression estimates in Appendix Table A1;  $\mathbb{E}[x | w]$  visualized in Appendix Figure A4.



**Figure A11: Country-level shrinkage from OLS to Bayesian posterior in the temperature–GDP sample.**

OLS and Bayesian posterior estimates of  $\hat{\tau}_i$  for each country, ordered by mean temperature (coldest at left). Orange points are OLS estimates ( $-\frac{1}{2}\hat{\beta}_{1i}/\hat{\beta}_{2i}$ ); blue points are posterior means under the pooled prior described in Section 3.2. Grey segments connect each country’s OLS and Bayesian estimates, visualizing the direction and magnitude of shrinkage. The solid green line marks the pooled  $\hat{\tau}$ . Triangles indicate estimates clipped to the displayed  $(0, 30)^\circ\text{C}$  range. The posterior pulls noisy OLS estimates—particularly those far from the pooled value—toward the common prior, while countries with precisely estimated curvature retain OLS-like peaks.

**Mortality: Bayesian Shrinkage**



**Figure A12: Bayesian shrinkage toward the pooled estimate in the temperature–mortality relationship.**

Mortality (65+ age group), estimated at the admin-1 level. Panel (a): posterior mean  $\hat{\tau}_{\text{Bayes}}$  against admin-1 mean temperature, with 95% posterior credible intervals (grey bars). The dashed line is the 45-degree line ( $\hat{\tau} = \bar{x}$ ); the solid green line is the pooled  $\hat{\tau}$ . Circled points are the example admin-1 units shown in panels (b)–(g). Panels (b)–(c): scatter of death rates against temperature with pooled (green) and admin-1-specific (orange) quadratic fits and 95% confidence bands for a well-identified and a weakly identified unit. Panels (d)–(e): prior (pooled sampling distribution, green), likelihood (admin-1-specific Hinkley likelihood, orange), and posterior (blue) densities for the same two units. Panels (f)–(g): projected mortality index (2025 = 100) under uniform  $+3^\circ\text{C}$  warming from 2025–2100, comparing posterior (blue) and pooled (green) damage estimates.

## B.1 Tables

	GDP Growth				Mortality			
	No- $w$	Linear	Full	Tercile	No- $w$	Linear	Full	Tercile
<i>(A) Coefficient Estimates</i>								
$T$	0.0091 (0.0023)	0.0083 (0.0037)	0.0123 (0.0049)	0.0082 (0.0024)	-0.3401 (0.0855)	-0.6158 (0.1144)	-0.6291 (0.1368)	-0.3497 (0.0881)
$T \times w$	—	0.0007 (0.0015)	-0.0049 (0.0033)		—	0.2803 (0.0927)	0.2958 (0.1459)	
$T \times \mathbf{1}_{\text{poor}}$	—	—	—	-0.0005 (0.0085)	—	—	—	0.0088 (0.0491)
$T^2$	-0.0003 (0.0001)	-0.0003 (0.0001)	-0.0004 (0.0001)	-0.0002 (0.0001)	0.0088 (0.0020)	0.0096 (0.0021)	0.0100 (0.0034)	0.0097 (0.0022)
$T^2 \times w$	—	—	0.0002 (0.0001)		—	—	-0.0005 (0.0042)	
$T^2 \times \mathbf{1}_{\text{poor}}$	—	—	—	-0.0001 (0.0002)	—	—	—	-0.0032 (0.0016)
<i>(B) Fixed Effects</i>								
Unit FE	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Unit $\times$ Year Trend	✓	✓	✓	✓				
Unit $\times$ Precip.					✓	✓	✓	✓
Unit $\times$ Year $\times$ Age					✓	✓	✓	✓
Observations	8,805	8,786	8,786	8,786	137,078	137,078	137,078	137,078
Adj. $R^2$	0.138	0.138	0.139	0.138	0.764	0.764	0.764	0.764
<i>(C) Peak Temperature (at <math>w = 0</math> / base group for tercile)</i>								
$\hat{\tau}$	14.8	14.6	17.2	17.8	19.3	32.0	31.3	18.0
Delta CI (95%)	[10.9, 18.8]	[9.6, 19.5]	[12.5, 21.9]	[10.0, 25.6]	[15.3, 23.2]	[23.5, 40.4]	[19.0, 43.6]	[13.9, 22.1]
AR CI (95%)	[10.7, 19.4]	[5.1, 19.4]	[9.2, 21.8]	[11.7, 38.7]	[14.9, 23.6]	[24.7, 43.7]	[22.9, 64.0]	[13.5, 22.5]
$ t_{\beta_2} $	4.49	2.83	3.04	2.70	4.38	4.67	2.92	4.46

**Table A1: Quadratic temperature regressions with covariate interactions.**

Panel (A) reports coefficient estimates from quadratic temperature regressions for GDP growth (columns 1–4) and mortality (columns 5–8).  $T$  is temperature in degrees Celsius and  $w$  is standardized (mean zero, unit variance) average log GDP per capita in constant 2015 dollars from the World Development Indicators. “No- $w$ ”: pooled baseline with no covariate interaction. “Linear”:  $w$  interacted with  $T$  only. “Full”:  $w$  interacted with both  $T$  and  $T^2$ . “Tercile”:  $w$  replaced with an indicator  $\mathbf{1}_{\text{poor}}$  for being in the bottom tercile of the estimation sample. Panel (B) lists fixed effects: the GDP specification includes unit, year, and unit-specific linear time trends; the mortality specification includes unit, year, unit-specific precipitation controls, and unit  $\times$  year  $\times$  age-group effects. Panel (C) reports the implied peak temperature  $\hat{\tau}$  evaluated at  $w = 0$  (or at the base group for the tercile model), with both conventional delta-method and identification-robust Anderson–Rubin 95% confidence intervals, and the absolute  $t$ -statistic on the curvature term  $|t_{\beta_2}|$ . Standard errors in parentheses, clustered at the unit level.

## C Robust confidence sets: extensions

### C.1 Sign-restricted Anderson-Rubin (SAR) confidence sets

Suppose  $\tau_0$  reflects a maximum, as in the temperature-GDP relationship (it is straightforward to extend all our theory to the case in which  $\tau_0$  is a minimum). Anderson-Rubin confidence sets provide correct coverage for a turning point  $\tau_0$ , but do not impose the nature of the turning point. This can be problematic, as when curvature is weak, Anderson-Rubin confidence sets can be disjoint, covering economically irrelevant cases. Intuitively, when curvature is weak, imposing  $\beta_2 < 0$  should rule out some parts of the Anderson-Rubin confidence set, but at the cost of expanding the width of the remaining interval.

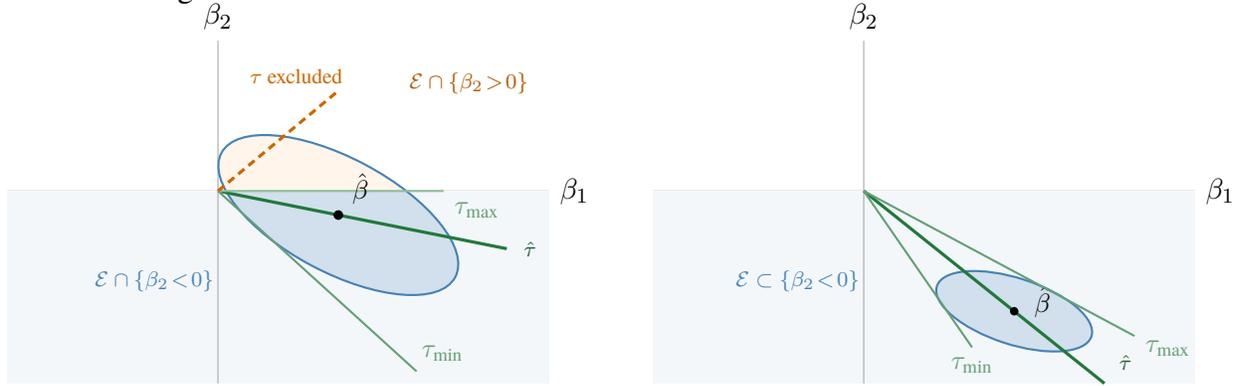


Figure A13: Geometric intuition for the sign-restricted Anderson-Rubin (SAR) confidence set. Green rays from the origin trace values of  $\tau = -\beta_1/(2\beta_2)$  through the joint confidence ellipse  $\mathcal{E}_2$ ; the darkest ray marks  $\hat{\tau}$ , while  $\tau_{\min}$  and  $\tau_{\max}$  bound the set. Orange dashed rays correspond to  $\tau$  values excluded by the sign restriction  $\beta_2 < 0$ . *Left:*  $\hat{\beta}$  lies near  $\beta_2 = 0$ , so the ellipse straddles the axis and the confidence set is wide. *Right:*  $\hat{\beta}$  is well below  $\beta_2 = 0$ , so the ellipse lies entirely in  $\{\beta_2 < 0\}$  and the confidence set is compact.

To capture this intuition, we work in the space of  $(\beta_1, \beta_2)$ . Letting,  $\hat{\beta} \sim \mathcal{N}(\beta, \Sigma)$ , we can consider the Wald ellipse for confidence level  $\alpha$ ,

$$\mathcal{E} = \left\{ \beta : (\hat{\beta} - \beta)' \Sigma^{-1} (\hat{\beta} - \beta) \leq \chi_{2,1-\alpha}^2 \right\}.$$

where  $\chi_{2,1-\alpha}^2$  is the  $\alpha$ -critical value of a chi-squared distribution with two degrees of freedom. We use two degrees of freedom, as here the confidence set reflects a joint hypothesis on  $\hat{\tau}$  and  $\beta_2$ .<sup>12</sup> Any turning point  $\tau_0 = -\beta_1/(2\beta_2)$  can be represented in this space as a ray through the origin parameterized by  $s$ ,  $\mathcal{R}(\tau_0) = (-2\tau_0 s, s)$ . For  $\beta_2 < 0$ , this is the ray segment  $s < 0$ . A candidate  $\tau$

<sup>12</sup>Note that at the boundary ( $\beta_2 = 0$ ), the chi-bar distribution provides correct critical values (Self and Liang, 1987).

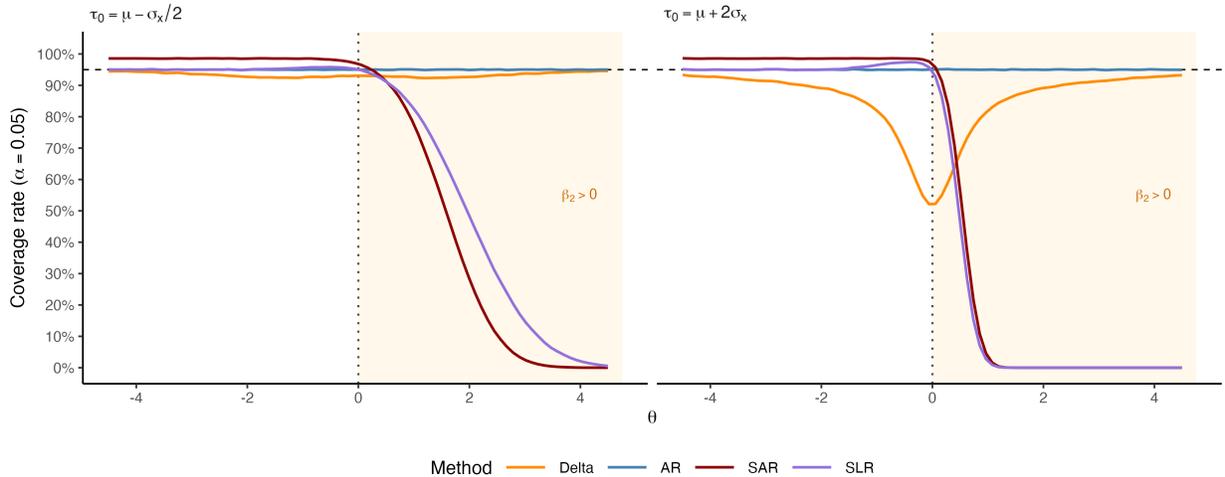
is included in the confidence set if  $R(\tau)$  intersects  $\mathcal{E}$ , so

$$\text{CS}_\alpha^{\text{SAR}} = \{\tau : \mathcal{E} \cap \mathcal{R}(\tau) \neq \emptyset\}.$$

Figure A13 provides geometric intuition for this set. The figure on the left represents the construction of an SAR confidence set when  $\beta_2$  is weakly identified and  $\mathcal{E}$  contains values  $\beta_2 > 0$ . Green rays  $\mathcal{R}(\tau)$  are pictured coming out of the origin, intersecting the ‘Wald’ ellipse  $\mathcal{E}$ . The dark orange dashed line represents a  $\tau$  value that would be included in a traditional Anderson-Rubin confidence set, but is excluded from our SAR set as it does not correspond to a maximum. In the right panel,  $\beta_2$  is strongly identified, and the confidence set is not affected by the sign restriction.

### Comparison of confidence sets

Figure A14 extends the coverage comparison in Figure 4 to include the sign-restricted Anderson–Rubin (SAR) confidence set. The SAR set is conservative relative to SLR, overcovering when the sign restriction  $\beta_2 < 0$  is correctly imposed, but shares the desirable property that coverage goes to zero when the restriction is violated.



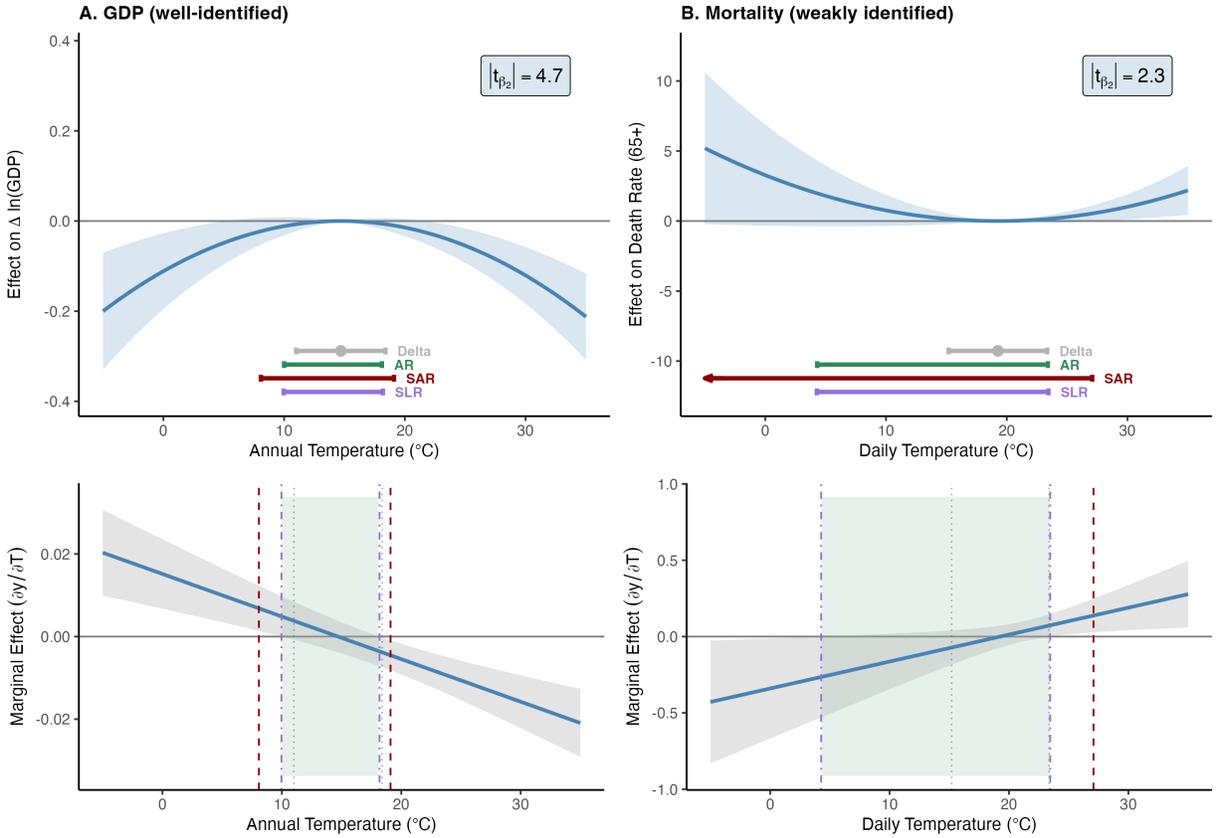
**Figure A14: Coverage of sign-restricted confidence sets, including the projection method.**

Each panel plots empirical coverage of nominal 95% confidence sets for the turning point  $\tau_0$  as a function of curvature strength  $\theta$ , under the maintained restriction  $\beta_2 < 0$ . The horizontal dashed line marks the nominal 95% level; departures below indicate undercoverage. The shaded region ( $\theta > 0$ ) corresponds to  $\beta_2 > 0$ , where the sign restriction is violated. Left:  $\tau_0 = \mu - \sigma_x/2$  (peak near climate). Right:  $\tau_0 = \mu + 2\sigma_x$  (peak far from climate). The delta method (orange) under-covers when  $\theta$  is small and the peak is far from climate. The Anderson–Rubin set (AR, blue) achieves correct coverage throughout. The sign-restricted projection set (SAR, maroon) is conservative, over-covering when  $\theta < 0$ . The sign-restricted likelihood ratio set (SLR, lilac) matches AR when  $\beta_2 < 0$  and is less conservative than SAR. Coverage for both sign-restricted tests goes to zero as  $\theta$  grows positive, consistent with correctly rejecting turning points that are not maxima. Based on 20,000 Monte Carlo draws from the DGP of Proposition 2.1.

Figure A15 reproduces the applications in Figure 5, adding the sign-restricted Anderson–Rubin (SAR) confidence set. The SAR set is wider than the SLR set in both applications, reflecting the conservatism of the projection approach. When curvature is strongly identified (GDP, left), the difference is modest. When curvature is weaker (mortality, right), the SAR set is substantially wider than both the AR and SLR sets.

## C.2 Extension to Poisson estimators

In many contexts, the left-hand side variable is given in logs. If the conditional mean is misspecified; e.g., exogeneity between  $x_i$  and  $u_i$  is given by  $\mathbb{E}[\exp(u_i) | x_i] = 1$ , rather than  $\mathbb{E}[u_i | x_i] = 0$ , then taking logs and running OLS can generate biased estimates for  $\beta_1$  and  $\beta_2$  and Poisson regression may be more appropriate (Silva and Tenreyro, 2006). We examine whether this bias can affect estimates of the innovation-maximising level of competition across U.K. manufacturing industries using data from Aghion et al. (2005). Appendix Figure A8 estimates industry-level turning points



**Figure A15: Delta-method versus identification-robust inference on turning points, including the sign-restricted projection method.**

*Top row.* Estimated response functions for GDP growth (left, Burke et al., 2015) and mortality rates (right, Carleton et al., 2022), with pointwise 95% confidence bands. Horizontal bars report four 95% confidence sets for the turning point  $\tau_0$ : the delta-method interval (grey); the Anderson–Rubin interval (AR, green), which inverts a  $\chi^2_1$  test of  $H_0: \tau_0 = \tau$ ; the sign-restricted AR interval (SAR, dark red), which additionally imposes the maintained assumption that the response is concave (GDP) or convex (mortality) using a  $\chi^2_2$  critical value; and the sign-restricted likelihood-ratio interval (SLR, purple), which imposes the same restriction but recovers a  $\chi^2_1$  critical value. The SAR set is wider than SLR, reflecting the conservatism of the projection approach. *Bottom row.* AR inversion: for each hypothesized  $\tau_0$ , the marginal effect  $\hat{\beta}_1 + 2\hat{\beta}_2\tau_0$  and its delta-method confidence band (grey ribbon) are plotted; the AR confidence set (green shading) collects all  $\tau_0$  for which this band includes zero. Vertical lines mark the delta-method point estimate (grey dotted), the AR bounds (green shaded area), the SAR bounds (dark red dashed), and the SLR bounds (purple dot-dash). Further details on data and estimation in Appendix A.

for the effect of a competition index on citation weighted patents, in both a log-plus-one OLS model and a Poisson model. Estimated turning points are near average levels of industry competitiveness for both models. A naive reading of the data suggests that level of competition is at its innovation-maximising level across manufacturing industries in the U.K.. The pooled peak is well-identified,  $|\hat{t}_{\beta_2}| \approx 4$ , and its Anderson-Rubin confidence set is about 10% wider than that implied by the

delta-method.

## D Proofs and derivations

In this Appendix, we begin by proving Proposition 2.1 (D.1). We then provide an alternative statement of the Proposition by considering sequences of  $\beta_{2,n} = h_2/\sqrt{n}$ ,  $h_2 \neq 0$ , and  $\tau_0$  held fixed, and relax assumptions on the normality of the regression errors, and prove it using the same approach (Appendix D.2). Appendix D.3 laboriously derives the Hinkley distribution that characterises the sampling distribution of  $\hat{\tau}$  outside of the limiting cases. We rely on the closed form to generate figures like Figure 2, or Appendix Figure A2. In Appendix D.4, we prove Corollary 3.1.

### D.1 Proof of Proposition 2.1

To prove the proposition, we begin by ‘recentering’ the regressors, writing the regression in terms of,

$$y_i = \alpha_0 + \alpha_1[x_i - \mu] + \alpha_2[(x_i - \mu)^2 - \sigma_x^2] + u_i$$

Now,  $\alpha_1 = \beta_1 + 2\mu\beta_2$  and  $\alpha_2 = \beta_2$ .

The ‘recentered’ regressors are orthogonal, as  $\mathbb{E}[(x_i - \mu)((x_i - \mu)^2 - \sigma_x^2)] = 0$  by Isserelis’s theorem: odd moments of centered normal random variables are zero. We can re-parameterize the turning point now as,

$$\hat{\tau} = \mu - \frac{\hat{\alpha}_1}{2\hat{\alpha}_2}.$$

Constructing the design matrix for this regression where,  $X = [1, x_i - \mu, (x_i - \mu)^2 - \sigma_x^2]$ , we have that,

$$\mathbb{E}[XX'] = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \sigma_x^2 & 0 \\ 0 & 0 & 2\sigma_x^4 \end{pmatrix}$$

The asymptotic variance of the coefficients is given by  $\sigma_u^2 \mathbb{E}[XX']^{-1}$ , where the inverse design matrix is simple as it is diagonal. Then, as  $\beta_1 = -2\beta_2\tau_0$ , we have,

$$\hat{\alpha}_1 = -2\beta_2\tau_0 + 2\mu\beta_2 + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sigma_x} Z_1, \quad \hat{\alpha}_2 = \beta_2 + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sqrt{2}\sigma_x^2} Z_2, \quad Z_1, Z_2 \sim^{\text{iid}} \mathcal{N}(0, 1).$$

Simplifying,

$$\frac{\hat{\alpha}_1}{\hat{\alpha}_2} = \frac{2\beta_2(\mu - \tau_0) + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sigma_x} Z_1}{\beta_2 + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sqrt{2}\sigma_x^2} Z_2}$$

Multiplying the numerator and denominator by  $\sqrt{2n}\sigma_x^2/\sigma_u$ ,

$$\begin{aligned}\frac{1}{2} \frac{\hat{\alpha}_1}{\hat{\alpha}_2} &= \frac{1}{2} \frac{2\beta_2\sqrt{2n}\frac{\sigma_x^2}{\sigma_u}(\mu - \tau_0) + \sqrt{2}\sigma_x Z_1}{\beta_2\sqrt{2n}\frac{\sigma_x^2}{\sigma_u} + Z_2} \\ &= \frac{\sigma_x\sqrt{2}}{2} \frac{2\beta_2\sqrt{2n}\frac{\sigma_x^2}{\sigma_u}\left(\frac{\mu-\tau_0}{\sqrt{2}\sigma_x}\right) + Z_1}{\beta_2\sqrt{2n}\frac{\sigma_x^2}{\sigma_u} + Z_2} \\ &= \frac{\sigma_x}{\sqrt{2}} \frac{2\beta_2\sqrt{n}\frac{\sigma_x^2}{\sigma_u}\left(\frac{\mu-\tau_0}{\sigma_x}\right) + Z_1}{\beta_2\sqrt{2n}\frac{\sigma_x^2}{\sigma_u} + Z_2}\end{aligned}$$

Defining  $\theta = 2\beta_2\frac{\sigma_x}{\sqrt{2}}\sqrt{n}\frac{\sigma_x}{\sigma_u}$ , we have,

$$\hat{\tau} = \mu - \frac{1}{2} \frac{\hat{\alpha}_1}{\hat{\alpha}_2} = \mu - \frac{\sigma_x}{\sqrt{2}} \frac{Z_1 + \sqrt{2} \left(\frac{\mu-\tau_0}{\sigma_x}\right) \theta}{Z_2 + \theta}, \quad Z_1, Z_2 \sim^{\text{iid}} \mathcal{N}(0, 1).$$

Under the null,  $\theta = 0$ , so,

$$\hat{\tau}|_{\theta=0} \sim \mu - \frac{\sigma_x}{\sqrt{2}} \frac{Z_1}{Z_2}, \quad Z_1, Z_2 \sim^{\text{iid}} \mathcal{N}(0, 1).$$

i.e.,

$$\hat{\tau} \sim \text{Cauchy}\left(\mu, \frac{\sigma_x}{\sqrt{2}}\right).$$

Rewriting  $\hat{\tau}$ , we can see that as  $\sqrt{n} \rightarrow \infty, \theta \rightarrow \infty$ ,

$$\begin{aligned}\lim_{\sqrt{n} \rightarrow \infty} \hat{\tau} &= \lim_{\theta \rightarrow \infty} \mu - \frac{\sigma_x}{\sqrt{2}} \frac{Z_1/\theta + \sqrt{2} \left(\frac{\mu-\tau_0}{\sigma_x}\right)}{Z_2/\theta + 1} \\ &= \mu - \frac{\sigma_x}{\sqrt{2}} \sqrt{2} \left(\frac{\mu - \tau_0}{\sigma_x}\right) \\ &= \tau_0\end{aligned}$$

So  $\hat{\tau}$  is root- $n$  consistent for  $\tau_0$ . Consistency allows us to take distributional limits: provided  $\beta_2 \neq 0$  and  $\tau_0$  exists, the statistic is continuous in the parameters in a neighborhood of  $\tau_0$  and the continuous mapping theorem applies. To get the asymptotic variance (i.e., delta method variance), we note,

$$\sqrt{n} \begin{pmatrix} \hat{\alpha}_1 - \alpha_1 \\ \hat{\alpha}_2 - \alpha_2 \end{pmatrix} \rightarrow^d \mathcal{N}\left(0, \begin{pmatrix} \sigma_u^2 & 0 \\ 0 & \frac{1}{2}\sigma_x^{-2} \end{pmatrix}\right).$$

Let  $g(\alpha_1, \alpha_2) = \mu - \frac{\alpha_1}{2\alpha_2}$ . Its gradient evaluated at the true  $(\alpha_1, \alpha_2)$  is

$$\nabla g = \begin{pmatrix} -\frac{1}{2\beta_2} \\ \frac{\mu - \tau_0}{\beta_2} \end{pmatrix}.$$

So the delta method variance is,

$$\begin{aligned} \begin{pmatrix} -\frac{1}{2\beta_2} \\ \frac{\mu - \tau_0}{\beta_2} \end{pmatrix}' \frac{\sigma_u^2}{\sigma_x^2} \begin{pmatrix} 1 & 0 \\ 0 & \frac{1}{2}\sigma_x^{-2} \end{pmatrix} \begin{pmatrix} -\frac{1}{2\beta_2} \\ \frac{\mu - \tau_0}{\beta_2} \end{pmatrix} &= \frac{\sigma_u^2}{\beta_2^2} \left( \frac{1}{4\sigma_x^2} + \frac{(\mu - \tau_0)^2}{2\sigma_x^4} \right) \\ &= \frac{1}{4} \frac{\sigma_u^2}{\beta_2^2 \sigma_x^2} \left( 1 + \left( \sqrt{2} \frac{\mu - \tau_0}{\sigma_x} \right)^2 \right) \end{aligned}$$

So,

$$\sqrt{n}(\hat{\tau} - \tau_0) \rightarrow^d \mathcal{N}\left(0, \frac{1}{4} \frac{\sigma_u^2}{\beta_2^2 \sigma_x^2} \left( 1 + \left( \sqrt{2} \frac{\mu - \tau_0}{\sigma_x} \right)^2 \right)\right)$$

## D.2 Asymptotic distribution under local-to-zero curvature

Proposition 2.1 in the main text derives the Hinkley distribution under fixed parameters and normal errors. We now provide a complementary asymptotic formulation using local-to-zero sequences for  $\beta_2$ , which delivers the same distribution as an *exact* limiting result under weaker assumptions on the error distribution (finite fourth moments rather than normality).

**Proposition D.1** (Asymptotic distribution under local-to-zero curvature). *Let  $\{(x_i, u_i)\}_{i=1}^n$  be i.i.d. with*

$$x_i \sim \mathcal{N}(\mu, \sigma_x^2), \quad \mathbb{E}[u_i | x_i] = 0, \quad \mathbb{E}[u_i^2 | x_i] = \sigma_u^2, \quad \mathbb{E}[u_i^4] < \infty,$$

*and suppose  $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + u_i$ . Define  $\hat{\tau} = -\hat{\beta}_1 / (2\hat{\beta}_2)$  and  $\tau_0 = -\beta_1 / (2\beta_2)$  for  $\beta_2 \neq 0$ .*

1. **Local-to-zero asymptotics (Hinkley distribution).** *Under  $\beta_{2,n} = h_2 / \sqrt{n}$  for fixed  $h_2 \neq 0$ , with  $\tau_0$  held fixed (so  $\beta_{1,n} = -2\tau_0 \beta_{2,n}$ ):*

$$\hat{\tau} \Rightarrow \mu - \frac{\sigma_x}{\sqrt{2}} \cdot \frac{Z_1 + \left( \sqrt{2} \left( \frac{\mu - \tau_0}{\sigma_x} \right) \theta \right)}{Z_2 + \theta}, \quad Z_1, Z_2 \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1),$$

where

$$\theta = h_2 \sigma_x^2 \sqrt{2} / \sigma_u$$

indexes curvature signal-to-noise.

2. **Limiting behavior indexed by  $|\theta|$ .**

(a) Large  $|\theta|$ : The Hinkley limit concentrates at  $\tau_0$  with variance  $\approx \sigma_x^2/(2\theta^2)$ ; approximately normal.

(b) Small  $|\theta|$  ( $(\mu - \tau_0)/\sigma_x$  bounded): Approaches  $\text{Cauchy}(\mu, \sigma_x/\sqrt{2})$ , symmetric about  $\mu$  (not  $\tau_0$ ). Delta-method inference fails, and the estimator is median biased.

Remark (Strong identification). If  $\beta_2 \neq 0$  is fixed (not local-to-zero), then  $\sqrt{n}(\hat{\tau} - \tau_0) \Rightarrow \mathcal{N}(0, V_\tau)$ : the standard delta-method result, valid when  $|\theta| \rightarrow \infty$  with  $n$ .

3. **Observable diagnostic.**  $t_{\beta_2} \Rightarrow Z_2 + \theta \sim \mathcal{N}(\theta, 1)$ , so  $|t_{\beta_2}|$  estimates  $|\theta|$ .

*Proof.* The reparameterization used in the proof for Proposition 2.1 holds, though we no longer have exactly normal coefficients,

$$\sqrt{n} \begin{pmatrix} \hat{\alpha}_1 - \alpha_1 \\ \hat{\alpha}_2 - \alpha_2 \end{pmatrix} \rightarrow^d \mathcal{N} \left( 0, \sigma_u^2 \begin{pmatrix} \sigma_x^2 & 0 \\ 0 & 2\sigma_x^4 \end{pmatrix} \right)$$

and,

$$\begin{aligned} \hat{\alpha}_1 &= \alpha_1 + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sigma_x} Z_1 + o_p(n^{-1/2}) \\ \hat{\alpha}_2 &= \alpha_2 + \frac{1}{\sqrt{n}} \frac{\sigma_u}{\sqrt{2}\sigma_x^2} Z_2 + o_p(n^{-1/2}) \end{aligned}$$

with  $Z_1, Z_2 \sim^{\text{iid}} \mathcal{N}(0, 1)$ .

Now, with  $\beta_{2,n} = h_2/\sqrt{n}$  and  $\tau_0$  fixed, the local coefficients are  $\alpha_{1,n} = 2(\mu - \tau_0)h_2/\sqrt{n} \equiv h_1/\sqrt{n}$ , and  $\alpha_{2,n} = h_2/\sqrt{n}$ . Therefore,

$$\sqrt{n}\hat{\alpha}_1 \rightarrow h_1 + \frac{\sigma_u}{\sigma_x} Z_1, \quad \sqrt{n}\hat{\alpha}_2 \rightarrow h_2 + \frac{\sigma_u}{\sqrt{2}\sigma_x^2} Z_2.$$

We now define  $\theta = h_2\sigma_x^2\sqrt{2}/\sigma_u$ . Then the numerator signal-to-noise ratio is  $h_1\sigma_x/\sigma_u = \sqrt{2}(\mu - \tau_0)/\sigma_x \cdot \theta$  (using  $h_1 = 2(\mu - \tau_0)h_2$ ). So now we have,

$$\frac{\hat{a}_1}{\hat{a}_2} = \frac{\sqrt{n}\hat{\alpha}_1}{\sqrt{n}\hat{\alpha}_2} \rightarrow \sqrt{2}\sigma_x \frac{Z_1 + \sqrt{2} \left( \frac{\mu - \tau_0}{\sigma_x} \right) \theta}{Z_2 + \theta},$$

as in Proposition 2.1. Then, as we still have,  $\hat{\tau} = \mu - \frac{1}{2} \frac{\hat{\alpha}_1}{\hat{\alpha}_2}$ , plugging in we obtain,

$$\hat{\tau} \rightarrow \mu - \frac{\sigma_x}{\sqrt{2}} \cdot \frac{Z_1 + \sqrt{2} \left( \frac{\mu - \tau_0}{\sigma_x} \right) \theta}{Z_2 + \theta}.$$

This formulation is identical to that in Proposition 2.1, so we have the same limits: when  $|\theta| \rightarrow \infty$ :  $Z_2 + \theta \approx \theta$ , giving  $\hat{\tau} \approx \tau_0 - (\sigma_x/\sqrt{2}\theta)Z_1 \sim \mathcal{N}(\tau_0, \sigma_x^2/(2\theta^2))$ . Taking  $\theta \rightarrow 0$  (with  $(\mu - \tau_0)/\sigma_x$  bounded),  $\hat{\tau} \Rightarrow \mu - (\sigma_x/\sqrt{2})(Z_1/Z_2) \sim \text{Cauchy}(\mu, \sigma_x/\sqrt{2})$ , using the fact that the ratio of two independent standard normals is Cauchy(0, 1).

The  $t$ -statistic remains a useful diagnostic  $t_{\beta_2} = \hat{\beta}_2/\widehat{\text{se}}(\hat{\beta}_2) = (\sqrt{2n}\sigma_x^2/\sigma_u)\hat{\beta}_2 + o_p(1)$ , where the  $o_p(1)$  term reflects that  $\widehat{\text{se}} \xrightarrow{p} \text{se}$ . Under  $\beta_{2,n} = h_2/\sqrt{n}$ :  $t_{\beta_2} \Rightarrow Z_2 + \theta \sim \mathcal{N}(\theta, 1)$ .  $\square$

### D.3 Derivation of the analytic Hinkley distribution

To fully analytically characterise the Hinkley distribution, we work with the standard regressors. Since  $x_i \sim \mathcal{N}(\mu, \sigma_x^2)$ , when  $X = [1, x_i, x_i^2]$ , the design matrix

$$\mathbb{E}[X X'] = \begin{pmatrix} 1 & \mu & \mu^2 + \sigma_x^2 \\ \mu & \mu^2 + \sigma_x^2 & \mu^3 + 3\mu\sigma_x^2 \\ \mu^2 + \sigma_x^2 & \mu^3 + 3\mu\sigma_x^2 & \mu^4 + 6\mu^2\sigma_x^2 + 3\sigma_x^4 \end{pmatrix}.$$

As the errors are homoscedastic,

$$V(\beta) = \sigma_u^2 \mathbb{E}[X' X]^{-1} = \frac{\sigma_u^2}{n} \frac{1}{2\sigma_x^4} \begin{pmatrix} \mu^4 + 3\sigma_x^4 & -2\mu^3 & \mu^2 - \sigma_x^2 \\ -2\mu^3 & 2(2\mu^2 + \sigma_x^2) & -2\mu \\ \mu^2 - \sigma_x^2 & -2\mu & 1 \end{pmatrix}.$$

First, we verify Proposition 2.1: We know the ratio of two centered normals is Cauchy. What remains is to derive the parameters of the sampling distribution. The Cauchy distribution has location parameter,  $-\frac{1}{2}\rho\beta\sigma_{\beta_1}/\sigma_{\beta_2}$ ,

$$\rho = \frac{[V_\beta]_{23}}{\sqrt{[V_\beta]_{22}} \cdot \sqrt{[V_\beta]_{33}}}, \quad \frac{\sigma_{\beta_1}}{\sigma_{\beta_2}} = \frac{\sqrt{[V_\beta]_{22}}}{\sqrt{[V_\beta]_{33}}}$$

so, the location is,

$$-\frac{1}{2} \frac{[V_\beta]_{23}}{[V_\beta]_{33}} = -\frac{1-2\mu}{2 \cdot 1} = \mu$$

and for the scale  $\frac{1}{2}\sigma_{\beta_1}/\sigma_{\beta_2}\sqrt{1-\rho^2}$ ,

$$\begin{aligned} \sqrt{1-\rho^2} &= \sqrt{1 - \frac{([V_\beta]_{23})^2}{[V_\beta]_{22}[V_\beta]_{33}}} \\ &= \sqrt{1 - \frac{4\mu^2}{4\mu^2 + 2\sigma_x^2}} \\ &= \sqrt{\frac{2\sigma_x^2}{4\mu^2 + 2\sigma_x^2}} \end{aligned}$$

So the scale is,

$$\frac{1}{2} \sqrt{4\mu^2 + 2\sigma_x^2} \sqrt{\frac{2\sigma_x^2}{4\mu^2 + 2\sigma_x^2}} = \frac{\sqrt{2}}{2} \sigma_x$$

As  $\sqrt{2}/2 = 1/\sqrt{2}$

$$-\frac{1}{2} \frac{\hat{\beta}_1}{\hat{\beta}_2} \sim \text{Cauchy} \left( \mu, \frac{\sigma_x}{\sqrt{2}} \right)$$

Now, when the coefficients are not mean zero, the distribution is neither Cauchy nor Gaussian. As made explicit in Hinkley (1969), the density for the ratio of two normals  $r = \beta_1/\beta_2$  is,

$$\begin{aligned} f(r) &= \frac{b(r)d(r)}{\sqrt{2\pi}\sigma_{\beta_1}\sigma_{\beta_2}a(r)^3} \left[ 2\Phi \left( \frac{b(r)}{\sqrt{1-\rho^2}a(r)} \right) - 1 \right] \\ &+ \frac{\sqrt{1-\rho^2}}{\pi\sigma_{\beta_1}\sigma_{\beta_2}a(r)^2} \exp \left( -\frac{c}{2(1-\rho^2)} \right) \end{aligned}$$

where,

$$\begin{aligned} a(r) &= \left( \frac{r^2}{\sigma_{\beta_1}^2} - \frac{2\rho r}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{1}{\sigma_{\beta_2}^2} \right)^{\frac{1}{2}} \\ b(r) &= \frac{\beta_1 r}{\sigma_{\beta_1}^2} - \frac{\rho(\beta_1 + \beta_2 r)}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{\beta_2}{\sigma_{\beta_2}^2} \\ c &= \frac{\beta_1^2}{\sigma_{\beta_1}^2} - \frac{2\rho\beta_1\beta_2}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{\beta_2^2}{\sigma_{\beta_2}^2} \\ d(r) &= \exp \left( \frac{b(r)^2 - ca(r)^2}{2(1-\rho^2)a(r)^2} \right) \end{aligned}$$

In our case, we are interested in the density for  $\tau = -\frac{1}{2}r$ . The change of variables rule says,  $f_{\hat{\tau}}(x) = f_r(g^{-1}(x)) \cdot \left| \frac{d}{dx}g^{-1}(x) \right|$  where here  $g(x) = -x/2$  so,  $f_{\tau}(x) = f(-2r)2$ .

We begin by noting some useful formulas:

$$\begin{aligned}\sigma_{\beta_1}\sigma_{\beta_2} &= \frac{n\sigma_u^2}{2\sigma_x^4}\sqrt{2(2\mu^2 + \sigma_x^2)} \\ \frac{\rho}{\sigma_{\beta_1}\sigma_{\beta_2}} &= \frac{-2\mu}{2\mu^2 + \sigma_x^2} \frac{\sigma_x^4}{n\sigma_u^2} \\ 1 - \rho^2 &= \frac{\sigma_x^2}{2\mu^2 + \sigma_x^2}.\end{aligned}$$

**Deriving  $a(x)^2$**  Using our change of variables, and relabelling  $n\sigma_u^2$  to just  $\sigma_u^2$ ,

$$\begin{aligned}a(x)^2 &= \frac{4x^2}{\sigma_{\beta_1}^2} + \frac{4\rho x}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{1}{\sigma_{\beta_2}^2} \\ &= \frac{4x^2}{\sigma_{\beta_1}^2} + 4x \left( \frac{-2\mu\sigma_x^2}{2\mu^2 + \sigma_x^2} \frac{\sigma_x^2}{\sigma_u^2} \right) + \frac{1}{\sigma_{\beta_2}^2} \\ &= \frac{4x^2}{\sigma_{\beta_1}^2} + 4x \left( \frac{-2\mu\sigma_x^2}{2\mu^2 + \sigma_x^2} \frac{\sigma_x^2}{\sigma_u^2} \right) + 2\sigma_x^2 \frac{\sigma_x^2}{\sigma_u^2} \\ &= \frac{4x^2}{2(2\mu^2 + \sigma_x^2)} \frac{2\sigma_x^4}{\sigma_u^2} + 4x \left( \frac{-2\mu\sigma_x^2}{2\mu^2 + \sigma_x^2} \frac{\sigma_x^2}{\sigma_u^2} \right) + 2\sigma_x^2 \frac{\sigma_x^2}{\sigma_u^2} \\ &= \frac{2x^2}{(2\mu^2 + \sigma_x^2)} \frac{2\sigma_x^4}{\sigma_u^2} + 4x \left( \frac{-2\mu\sigma_x^2}{2\mu^2 + \sigma_x^2} \frac{\sigma_x^2}{\sigma_u^2} \right) + 2\sigma_x^2 \frac{\sigma_x^2}{\sigma_u^2} \\ &= \frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (4x^2 - 8\mu x + 2(2\mu^2 + \sigma_x^2)) \\ &= \frac{2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2x^2 - 4\mu x + 2\mu^2 + \sigma_x^2)\end{aligned}$$

We now complete the square: noting that  $2x^2 - 4\mu x + 2\mu^2 = 2(x - \mu)^2$ , so that, undoing the relabel,

$$a(x)^2 = \frac{2\sigma_x^4}{n\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(x - \mu)^2 + \sigma_x^2)$$

**Deriving  $b(x)$**  We now derive  $b(x)$ . Here we use  $\beta_1 = -2\tau\beta_2$ , and again ignore  $n$ ,

$$\begin{aligned}
b(x) &= -2\frac{\beta_1}{\sigma_{\beta_1}^2}x - \frac{\rho(\beta_1 - 2\beta_2x)}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{\beta_2}{\sigma_{\beta_2}^2} \\
&= -2\frac{-2\tau\beta_2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}x - (-2\tau\beta_2 - 2\beta_2x)\frac{-2\mu\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} + \frac{\beta_2}{\sigma_{\beta_2}^2} \\
&= \frac{4\tau\beta_2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}x - (\tau + x)\beta_2\frac{4\mu\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} + \frac{\beta_2}{\sigma_{\beta_2}^2} \\
&= \frac{4\tau\beta_2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}x - (\tau + x)\beta_2\frac{4\mu\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} + 2\beta_2\frac{\sigma_x^4}{\sigma_u^2} \\
&= \beta_2\frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}(4\tau x - (\tau + x)4\mu + 2(2\mu^2 + \sigma_x^2)) \\
&= 2\beta_2\frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}(2\tau x - (\tau + x)2\mu + 2\mu^2 + \sigma_x^2)
\end{aligned}$$

Noting,  $2\tau x - (\tau + x)2\mu + 2\mu^2 = 2(\tau - \mu)(x - \mu)$ , we have,

$$b(x) = 2\beta_2\frac{\sigma_x^4}{n\sigma_u^2(2\mu^2 + \sigma_x^2)}(2(\tau - \mu)(x - \mu) + \sigma_x^2)$$

**Deriving  $c$**  To obtain the constant  $c$  in the formula, ignore  $n$  again and,

$$\begin{aligned}
c &= \frac{\beta_1^2}{\sigma_{\beta_1}^2} - \frac{2\rho\beta_1\beta_2}{\sigma_{\beta_1}\sigma_{\beta_2}} + \frac{\beta_2^2}{\sigma_{\beta_2}^2} \\
&= \frac{4\tau^2\beta_2^2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} - 4\tau\beta_2^2\frac{-2\mu\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} + \frac{\beta_2^2}{\sigma_{\beta_2}^2} \\
&= \frac{4\tau^2\beta_2^2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} - 8\tau\beta_2^2\frac{\mu\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} + \beta_2^2\frac{\sigma_x^4}{\sigma_u^2} \\
&= 2\beta_2^2\frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}(2\tau^2 - 4\tau\mu + 2\mu^2 + \sigma_x^2)
\end{aligned}$$

As  $2\tau^2 - 4\tau\mu + 2\mu^2 = 2(\tau - \mu)^2$ , we have,

$$c = 2\beta_2^2\frac{\sigma_x^4}{n\sigma_u^2(2\mu^2 + \sigma_x^2)}(2(\tau - \mu)^2 + \sigma_x^2)$$

**Deriving**  $d(x)$  So, first define,

$$K = \frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)}$$

and again subsuming  $n$  in  $\sigma_u^2$  for the algebra,

$$\begin{aligned} d(x) &= \exp\left(\frac{b(x)^2 - ca(x)^2}{2(1 - \rho^2)a(x)^2}\right) \\ &= \exp\left(\frac{[2\beta_2 K (2(\tau - \mu)(x - \mu) + \sigma_x^2)]^2 - [2\beta_2^2 K (2(\tau - \mu)^2 + \sigma_x^2)] [2K (2(x - \mu)^2 + \sigma_x^2)]}{2(1 - \rho^2) [2K (2(x - \mu)^2 + \sigma_x^2)]}\right) \\ &= \exp\left(2\beta_2^2 K \frac{[(2(\tau - \mu)(x - \mu) + \sigma_x^2)]^2 - [(2(\tau - \mu)^2 + \sigma_x^2)] [(2(x - \mu)^2 + \sigma_x^2)]}{2(1 - \rho^2) [2(2(x - \mu)^2 + \sigma_x^2)]}\right) \end{aligned}$$

The numerator simplifies,

$$[(2(\tau - \mu)(x - \mu) + \sigma_x^2)]^2 - [(2(\tau - \mu)^2 + \sigma_x^2)] [(2(x - \mu)^2 + \sigma_x^2)] = -2\sigma_x^2(\tau - x)^2$$

And note that  $1 - \rho^2 = \frac{\sigma_x^2}{2\mu^2 + \sigma_x^2} = K\sigma_u^2/\sigma_x^2$  so we have,

$$d(x) = \exp\left(\beta_2^2 K \frac{-2\sigma_x^2(\tau - x)^2}{K\sigma_u^2/\sigma_x^2 [2K (2(x - \mu)^2 + \sigma_x^2)]}\right)$$

which gives, undoing the  $\sigma_u^2$  relabel,

$$d(x) = \exp\left(-2\beta_2^2 \frac{\sigma_x^4}{n\sigma_u^2} \frac{(\tau - x)^2}{[(2(x - \mu)^2 + \sigma_x^2)]}\right)$$

Taking stock, we have,

$$\begin{aligned} a(x)^2 &= \frac{2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(x - \mu)^2 + \sigma_x^2) \\ b(x) &= 2\beta_2 \frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)(x - \mu) + \sigma_x^2) \\ c &= 2\beta_2^2 \frac{\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)^2 + \sigma_x^2) \\ d(x) &= \exp\left(-2\beta_2^2 \frac{\sigma_x^4}{\sigma_u^2} \frac{(\tau - x)^2}{[(2(x - \mu)^2 + \sigma_x^2)]}\right) \end{aligned}$$

Now, noting that  $\theta^2 = 2\beta_2^2 \frac{n\sigma_x^4}{\sigma_u^2}$ , we can write these as,

$$\begin{aligned} a(x)^2 &= \frac{2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(x - \mu)^2 + \sigma_x^2) \\ b(x) &= \theta^2 \frac{1}{\beta_2(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)(x - \mu) + \sigma_x^2) \\ c &= \theta^2 \frac{1}{(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)^2 + \sigma_x^2) \\ d(x) &= \exp\left(-\theta^2 \frac{(\tau - x)^2}{[(2(x - \mu)^2 + \sigma_x^2)]}\right) \end{aligned}$$

Returning to the Hinkley distribution, we can simplify the  $\rho$  terms, and the  $\sigma$  terms to obtain,

$$\begin{aligned} f_\tau(x) &= \frac{2b(x)d(x)}{\sqrt{2\pi} \left(n \frac{\sigma_u^2}{2\sigma_x^4} \sqrt{2(2\mu^2 + \sigma_x^2)}\right) a(x)^3} \left[2\Phi\left(\frac{b(x)}{\sqrt{1 - \rho^2}a(x)}\right) - 1\right] \\ &+ 2 \frac{\sqrt{\frac{\sigma_x^2}{2\mu^2 + \sigma_x^2}}}{\pi \left(n \frac{\sigma_u^2}{2\sigma_x^4} \sqrt{2(2\mu^2 + \sigma_x^2)}\right) a(x)^2} \exp\left(-\frac{c}{2} \left(1 + 2\left(\frac{\mu}{\sigma_x}\right)^2\right)\right) \end{aligned}$$

Cleaning up,

$$\begin{aligned} f_\tau(x) &= \underbrace{\frac{2b(x)d(x)}{\sqrt{2\pi} \left(\frac{\sigma_u^2}{2\sigma_x^4} \sqrt{2(2\mu^2 + \sigma_x^2)}\right) a(x)^3} \left[2\Phi\left(\frac{b(x)}{\sqrt{1 - \rho^2}a(x)}\right) - 1\right]}_{\text{term 1}} \\ &+ \underbrace{\frac{\sigma_x/\sqrt{2}}{\pi \left(\frac{\sigma_u^2}{2\sigma_x^4} (2\mu^2 + \sigma_x^2)\right) a(x)^2} \exp\left(-\frac{c}{2} \left(1 + 2\left(\frac{\mu}{\sigma_x}\right)^2\right)\right)}_{\text{term 2}} \end{aligned}$$

Plugging in for  $c$  and  $a(x)^2$  in the second term,

$$\text{term 2} = 2 \frac{\sigma_x/\sqrt{2}}{\pi(2(x - \mu)^2 + \sigma_x^2)} \exp\left(-\frac{1}{2}\theta^2 \frac{1}{(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)^2 + \sigma_x^2) \left(1 + 2\left(\frac{\mu}{\sigma_x}\right)^2\right)\right)$$

Cleaning up the second term,

$$\text{term 2} = \frac{\sigma_x/\sqrt{2}}{\pi((x - \mu)^2 + \sigma_x^2/2)} \exp\left(-\theta^2 \frac{1}{(\mu^2 + \sigma_x^2/2)} (2(\tau - \mu)^2 + \sigma_x^2) \left(1 + 2\left(\frac{\mu}{\sigma_x}\right)^2\right)\right)$$

Simplifying further,

$$\begin{aligned} \text{term 2} &= \frac{\sigma_x/\sqrt{2}}{\pi((x-\mu)^2 + \sigma_x^2/2)} \exp\left(-\theta^2 \frac{2}{2\mu^2 + \sigma_x^2} (2(\tau - \mu)^2 + \sigma_x^2) \left(\frac{\sigma_x^2 + 2\mu^2}{\sigma_x^2}\right)\right) \\ &= \frac{\sigma_x/\sqrt{2}}{\pi((x-\mu)^2 + \sigma_x^2/2)} \exp\left(-\frac{\theta^2}{2} \left(1 + \left(\frac{2(\tau - \mu)}{\sigma_x}\right)^2\right)\right) \end{aligned}$$

Define,

$$\kappa^2 = \theta^2 \left(1 + \left(\sqrt{2} \frac{\tau - \mu}{\sigma_x}\right)^2\right)$$

So,

$$\text{term 2} = \frac{\sigma_x/\sqrt{2}}{\pi((x-\mu)^2 + \sigma_x^2/2)} \exp(-\kappa^2/2)$$

Now reexamine term 1,

$$\text{term 1} = \frac{2b(x)}{\sqrt{2\pi} \left(\frac{\sigma_u^2}{2\sigma_x^4} \sqrt{2(2\mu^2 + \sigma_x^2)}\right) a(x)^3} \left[2\Phi\left(\frac{b(x)}{\sqrt{1 - \rho^2} a(x)}\right) - 1\right] d(x)$$

Consider  $b(x)/a(x)^2$ . In our context, the sign of  $b(x)$  is undone by the odd function  $2\Phi() - 1$  so we can put absolute values over any signed quantities,

$$\begin{aligned} b(x)/a(x)^2 &= \frac{\theta^2 \frac{1}{|\beta_2|(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)(x - \mu) + \sigma_x^2)}{\frac{2\sigma_x^4}{\sigma_u^2(2\mu^2 + \sigma_x^2)} (2(x - \mu)^2 + \sigma_x^2)} \\ &= \frac{\theta^2 \frac{1}{(2\mu^2 + \sigma_x^2)} (2(\tau - \mu)(x - \mu) + \sigma_x^2)}{\sqrt{2} |\beta_2| \frac{\sigma_x^2}{\sigma_u} \frac{\sqrt{2}\sigma_x^2}{\sigma_u(2\mu^2 + \sigma_x^2)} (2(x - \mu)^2 + \sigma_x^2)} \\ &= \frac{\theta^2 (2(\tau - \mu)(x - \mu) + \sigma_x^2)}{|\theta| \frac{\sqrt{2}\sigma_x^2}{\sigma_u} (2(x - \mu)^2 + \sigma_x^2)} \\ &= |\theta| \frac{\sigma_u}{\sqrt{2}\sigma_x^2} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(x - \mu)^2 + \sigma_x^2)} \end{aligned}$$

Now note that,

$$a(x) = \frac{\sqrt{2}\sigma_x^2}{\sigma_u \sqrt{(2\mu^2 + \sigma_x^2)}} \sqrt{(2(x - \mu)^2 + \sigma_x^2)}$$

So,

$$\begin{aligned}
b(x)/a(x)^3 &= |\theta| \frac{\sigma_u}{\sqrt{2}\sigma_x^2} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(x - \mu)^2 + \sigma_x^2)} \frac{1}{\frac{\sqrt{2}\sigma_x^2}{\sigma_u\sqrt{(2\mu^2 + \sigma_x^2)}} \sqrt{(2(x - \mu)^2 + \sigma_x^2)}} \\
&= |\theta| \frac{1}{2} \frac{\sigma_u^2}{\sigma_x^4} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(x - \mu)^2 + \sigma_x^2)} \frac{\sqrt{(2\mu^2 + \sigma_x^2)}}{\sqrt{(2(x - \mu)^2 + \sigma_x^2)}}
\end{aligned}$$

So then,

$$\begin{aligned}
&\frac{2b(x)}{\sqrt{2\pi} \left( \frac{\sigma_u^2}{2\sigma_x^4} \sqrt{2(2\mu^2 + \sigma_x^2)} \right) a(x)^3} \\
&= \frac{2\sigma_x^4}{\sqrt{\pi}\sigma_u^2\sqrt{2\mu^2 + \sigma_x^2}} \frac{b(x)}{a(x)^3} \\
&= \frac{2\sigma_x^4}{\sqrt{\pi}\sigma_u^2\sqrt{2\mu^2 + \sigma_x^2}} |\theta| \frac{1}{2} \frac{\sigma_u^2}{\sigma_x^4} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(x - \mu)^2 + \sigma_x^2)} \frac{\sqrt{(2\mu^2 + \sigma_x^2)}}{\sqrt{(2(x - \mu)^2 + \sigma_x^2)}} \\
&= \frac{1}{\sqrt{\pi}} |\theta| \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(x - \mu)^2 + \sigma_x^2)^{\frac{3}{2}}}
\end{aligned}$$

Or,

$$= \frac{1}{\sqrt{\pi}} \frac{\sigma_x}{\kappa} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{(2(\tau - \mu)^2 + \sigma_x^2)^{\frac{1}{2}} (2(x - \mu)^2 + \sigma_x^2)^{\frac{3}{2}}}$$

Now  $d(x)$  can be written as,

$$\begin{aligned}
d(x) &= \exp\left(-\theta^2 \frac{(\tau - x)^2}{2(x - \mu)^2 + \sigma_x^2}\right) \\
&= \exp\left(-\theta^2 \left(1 + 2 \left(\frac{\tau - \mu}{\sigma_x}\right)^2\right) \frac{(\tau - x)^2}{2(x - \mu)^2 + \sigma_x^2}\right) \\
&= \exp\left(-\frac{1}{2} \frac{\frac{1}{2\kappa^2} (\tau - x)^2}{\left(1 + 2 \left(\frac{\tau - \mu}{\sigma_x}\right)^2\right) (2(x - \mu)^2 + \sigma_x^2)}\right) \\
&= \exp\left(-\frac{1}{2} \frac{(\tau - x)^2}{\frac{\sigma_x^2}{2\kappa^2} \left(1 + \left(\sqrt{2} \frac{\tau - \mu}{\sigma_x}\right)^2\right)^2}\right) \\
&= \exp\left(-\frac{1}{2} \left(\frac{(\tau - x)}{\frac{\sigma_x}{\sqrt{2}|\theta|} \sqrt{1 + \left(\sqrt{2} \frac{\tau - \mu}{\sigma_x}\right)^2}}\right)^2\right)
\end{aligned}$$

So the entire term 1 becomes,

$$\text{term 1} = \frac{1}{\sqrt{\pi}} \frac{(2(\tau - \mu)(x - \mu) + \sigma_x^2)}{\frac{\sigma_x}{\kappa} (2(\tau - \mu)^2 + \sigma_x^2)^{\frac{1}{2}} (2(x - \mu)^2 + \sigma_x^2)^{\frac{3}{2}}} \exp\left(-\frac{1}{2} \left(\frac{(\tau - x)}{\frac{\sigma_x}{\sqrt{2}|\theta|} \sqrt{1 + \left(\sqrt{2} \frac{\tau - \mu}{\sigma_x}\right)^2}}\right)^2\right)$$

where I have omitted the  $2\Phi(\cdot) - 1$  term. This peaks at  $\tau$  and is approximately normal with,

$$\kappa \rightarrow \infty \implies \hat{\tau} \sim \mathcal{N}\left(\tau, \frac{\sigma_x^2}{2} \theta^{-2} \left(1 + \sqrt{2} \frac{\tau - \mu}{\sigma_x}\right)\right)$$

## D.4 Proof of Corollary 3.1

We begin with recentering the regressors so the regression is as follows,

$$y_i = a_0 + \underbrace{(\beta_1 + 2\mu\beta_2)}_{\equiv \alpha_1} [x_i - \mu] + \beta_2 [(x - \mu)^2 - \sigma_x^2] + \eta_1 x_i w_i + u_i$$

and now the first two regressors are orthogonal. Then, for  $X_i = (1, x_i - \mu, (x - \mu)^2 - \sigma_x^2, x_i w_i)$ ,

$$\mathbb{E}[X_i X_i'] = \begin{pmatrix} 1 & 0 & 0 & \rho\sigma_x \\ 0 & \sigma_x^2 & 0 & \mu\rho\sigma_x \\ 0 & 0 & 2\sigma_x^4 & 2\rho\sigma_x^3 \\ \rho\sigma_x & \mu\rho\sigma_x & 2\rho\sigma_x^3 & \mu^2 + \sigma_x^2(1 + 2\rho^2) \end{pmatrix}$$

which relies on Isserlis's theorem for the higher moments. Consider the off-diagonal entries. First,

$$\begin{aligned} \text{Cov}(x - \mu, xw) &= \mathbb{E}[(x - \mu)(x - \mu + \mu)w] \\ &= \mathbb{E}[(x - \mu)^2 w] + \mu\mathbb{E}[(x - \mu)w] \\ &= \mu\rho\sigma_x \end{aligned}$$

where  $\mathbb{E}[(x - \mu)^2 w] = 0$  by Isserlis's theorem. For,

$$\begin{aligned} \text{Cov}(x - \mu)^2 - \sigma_x^2, xw) &= \mathbb{E}[((x - \mu)^2 - \sigma_x^2)(x - \mu + \mu)w] \\ &= \mathbb{E}[(x - \mu)^2 - \sigma_x^2)(x - \mu)w] + \mu\mathbb{E}[(x - \mu)^2 - \sigma_x^2)w] \\ &= \mathbb{E}[(x - \mu)^3 w] - \sigma_x^2\mathbb{E}[(x - \mu)w] \\ &= \mathbb{E}[(x - \mu)^3 w] - \sigma_x^3\rho \\ &= 2\rho\sigma_x^3 \end{aligned}$$

as  $\mathbb{E}[(x - \mu)^2 - \sigma_x^2)w] = 0$  by Isserlis's theorem; likewise, for the last line, we can compute all the pairings for even moments to conclude  $\mathbb{E}[(x - \mu)^3 w] = 3\rho\sigma_x^3$ . Finally,

$$\begin{aligned} \mathbb{E}[x_i^2 w_i^2] &= \mathbb{E}[(x_i - \mu + \mu)^2 w_i^2] \\ &= \mathbb{E}[(x - \mu)^2 w^2 + 2\mu(x - \mu)w^2 + \mu^2 w^2] \\ &= \mathbb{E}[(x - \mu)^2 w^2] + \mu^2 \end{aligned}$$

as  $2\mu\mathbb{E}[(x - \mu)w^2] = 0$  again by Isserlis's. Isserlis also gives,

$$\begin{aligned} \mathbb{E}[(x - \mu)^2 w^2] &= \mathbb{E}[(x - \mu)^2]\mathbb{E}[w^2] + 2\mathbb{E}[(x - \mu)w]^2 \\ &= \sigma_x^2 + 2(\rho\sigma_x)^2 \\ &= \sigma_x^2(1 + 2\rho^2) \end{aligned}$$

This completes our characterization of the design matrix. Using a symbolic mathematical toolkit (sympy, in Python), we invert the design matrix and derive the variance-covariance matrix for the OLS coefficients,

$$\sigma_u^2 \mathbb{E}[X_i X_i']^{-1} = \begin{pmatrix} \frac{\sigma_u^2(\mu^2 \rho^2 - \mu^2 - \sigma_x^2)}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & -\frac{\mu \rho^2 \sigma_u^2}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & -\frac{\rho^2 \sigma_u^2}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & \frac{\rho \sigma_u^2 \sigma_x}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} \\ -\frac{\mu \rho^2 \sigma_u^2}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & \frac{\sigma_u^2(-\mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)}{\sigma_x^2(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & -\frac{\mu \rho^2 \sigma_u^2}{\sigma_x^2(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & \frac{\mu \rho \sigma_u^2}{\sigma_x(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} \\ -\frac{\rho^2 \sigma_u^2}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & -\frac{\mu \rho^2 \sigma_u^2}{\sigma_x^2(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & \frac{\sigma_u^2(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)}{2\sigma_x^4(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & \frac{\rho \sigma_u^2}{\sigma_x(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} \\ \frac{\rho \sigma_u^2 \sigma_x}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} & \frac{\mu \rho \sigma_u^2}{\sigma_x(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & -\frac{\rho \sigma_u^2}{\sigma_x(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} & -\frac{\sigma_u^2}{\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2} \end{pmatrix}$$

Now,

$$\hat{\tau}(w) = -\frac{1}{2} \frac{\hat{\alpha}_1 - 2\mu\hat{\beta}_2 + \hat{\eta}_1 w}{\hat{\beta}_2}$$

To construct the null, define matrix,

$$V = \begin{pmatrix} 0 & 1 & -2\mu & w \\ 0 & 0 & -2 & 0 \end{pmatrix}$$

Under the null, these are distributed,

$$\begin{pmatrix} \hat{\alpha}_1 - 2\mu\hat{\beta}_2 + \hat{\eta}_1 w \\ -2\hat{\beta}_2 \end{pmatrix} \sim \mathcal{N}(0, \Sigma), \quad \Sigma = V \sigma_u^2 \mathbb{E}[X_i X_i']^{-1} V'$$

where,

$$\begin{aligned} \Sigma_{11} &= \frac{\sigma_u^2(2\mu^4 \rho^2 - 2\mu^4 + 2\mu^2 \rho^2 \sigma_x^2 - 3\mu^2 \sigma_x^2 - 2\mu \rho \sigma_x^3 w + \rho^2 \sigma_x^4 - \sigma_x^4 w^2 - \sigma_x^4)}{\sigma_x^4(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} \\ \Sigma_{12} &= \frac{2\sigma_u^2(\mu^3 \rho^2 - \mu^3 - \mu \sigma_x^2 - \rho \sigma_x^3 w)}{\sigma_x^4(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} \\ \Sigma_{22} &= \frac{2\sigma_u^2(\mu^2 \rho^2 - \mu^2 - \rho^2 \sigma_x^2 - \sigma_x^2)}{\sigma_x^4(\mu^2 \rho^2 - \mu^2 + \rho^2 \sigma_x^2 - \sigma_x^2)} \end{aligned}$$

Using these, and recalling,  $\mathbb{E}[x | w] = \mu + \rho\sigma_x w$  we have,

$$\begin{aligned}
\check{\mu}(w) &= \frac{\Sigma_{12}}{\Sigma_{22}} \\
&= \frac{\mu^3\rho^2 - \mu^3 - \mu\sigma_x^2 - \rho\sigma_x^3 w}{\mu^2\rho^2 - \mu^2 - \rho^2\sigma_x^2 - \sigma_x^2} \\
&= \frac{-\mu^3\rho^2 + \mu^3 + \mu\sigma_x^2 + \rho\sigma_x^3 w}{-\mu^2\rho^2 + \mu^2 + \rho^2\sigma_x^2 + \sigma_x^2} \\
&= \frac{\mu(\mu^2(1 - \rho^2)) + \mu\sigma_x^2 + \sigma_x^2(\mathbb{E}[x | w] - \mu)}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)} \\
&= \frac{(\mu^2(1 - \rho^2))}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)}\mu + \frac{\sigma_x^2\mathbb{E}[x | w]}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)}
\end{aligned}$$

Defining,

$$\lambda = \frac{\sigma_x^2(1 + \rho^2)}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)}$$

Note that  $\lambda$  can be written as,

$$\lambda = \frac{(1 + \rho^2)}{(\mu/\sigma_x)^2(1 - \rho^2) + (1 + \rho^2)}$$

which is clearly governed by how strong the  $w$  signal is ( $\rho^2$ ) relative to the strength of the baseline mean after conditioning on  $w$ .

Now, we have,

$$\check{\mu}(w) = \lambda \frac{\mathbb{E}[x | w]}{1 + \rho^2} + (1 - \lambda)\mu$$

For the scale parameter, we must solve for,

$$\check{\sigma}(w) = \frac{\Sigma_{11}}{\Sigma_{22}} \sqrt{1 - \frac{(\Sigma_{12})^2}{\Sigma_{11}\Sigma_{22}}}$$

We derive this with sympy,

$$\check{\sigma}(w)^2 = \frac{\sigma_x^2}{2} \cdot \frac{\mathbb{E}[x^2](1 - \rho^2)}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)} \cdot \left(1 + \frac{(\sigma_x w - \mu\rho)^2}{\mu^2(1 - \rho^2) + \sigma_x^2(1 + \rho^2)}\right)$$

We can write this in terms of  $\lambda$ ,

$$\check{\sigma}(w) = \frac{\sigma_x}{\sqrt{2}} \cdot \frac{1}{1 + \rho^2} \cdot \sqrt{1 + \rho^2(1 - 2\lambda)} \cdot \sqrt{1 + \rho^2 + \lambda \left(w - \rho \frac{\mu}{\sigma_x}\right)^2}$$

## E Empirical Bayes implementation details

This appendix details the empirical Bayes procedure we implement in Section 3.2.

In contrast to the Bayesian estimates of the preceding section which simply relied on a fixed prior, *empirical* Bayes methods use the data to estimate the prior, which we refer to as  $G()$ . There are many interpretations of these estimators: e.g., to estimate a hierarchical model, to solve for decision rules that minimize their compound risk, to mimic an ‘oracle.’ See, e.g., Efron (2019) and Koenker and Gu (2024) or Walters (2024) for exposition.

In our setting, standard Empirical Bayes estimators, like a normal-likelihood, normal-prior model are inappropriate because the sampling distribution of the estimators  $\hat{\tau}_i$  is considerably skew, and we have no reason to believe the empirical prior on  $\tau$  ought to be normal. Instead, we proceed by following state-of-the-art methods in nonparametric maximum likelihood estimation of the prior  $G()$ . We model  $G()$  as a discrete mixture over a fine mesh of potential peaks,  $\tau_j$

$$G(\tau) = \sum_{j=1}^J \phi_j \delta(\tau - \tau_j)$$

where  $\delta()$  is the Dirac delta function. This is sometimes referred to as a ‘spiky’ prior.

We recover mixture weights  $\phi_j > 0$  by solving,

$$\min_{\lambda_j} - \sum_i \log \left( \sum_j \lambda_j L_{ij} \right) \quad \text{subject to} \quad \lambda_j > 0, \sum_j \lambda_j = 1$$

The idea is that is substantially easier to fit high-dimensional mixtures on a discretized grid using tools from convex optimization (Koenker and Mizera, 2014) by recovering mixture weights  $\phi_j$  than it is to fit low-dimensional mixture models by, e.g., the expectation-maximisation (EM) algorithm (Dempster et al., 1977). To solve this optimization problem, we rely on the toolkit from Youngseok Kim et al. (2020). In practice, we use a  $J = 200$  point grid for  $\tau$  on  $(-2, 42)$  in degrees C.

The elements of  $L_{ij}$  are the likelihoods  $p_i(\hat{\tau}_i - \tau_j)$ , where we fit  $p_i(\hat{\tau}_i - \tau)$  by fitting a mixture of 5 gaussian distributions,  $\hat{p}_i()$ , by EM, onto the residuals  $R_i$  of bootstrap draws  $b$  from the likelihood for  $\hat{\tau}_i$ ;  $R_i = \tau_i^b - \hat{\tau}_i$ .

Appendix Figure A16 shows the estimated prior for the temperature-GDP data against the bootstrap distribution of  $\hat{\tau}_i$  (left panel). The right panel shows the posterior means of  $\tau_i$ . The sampling

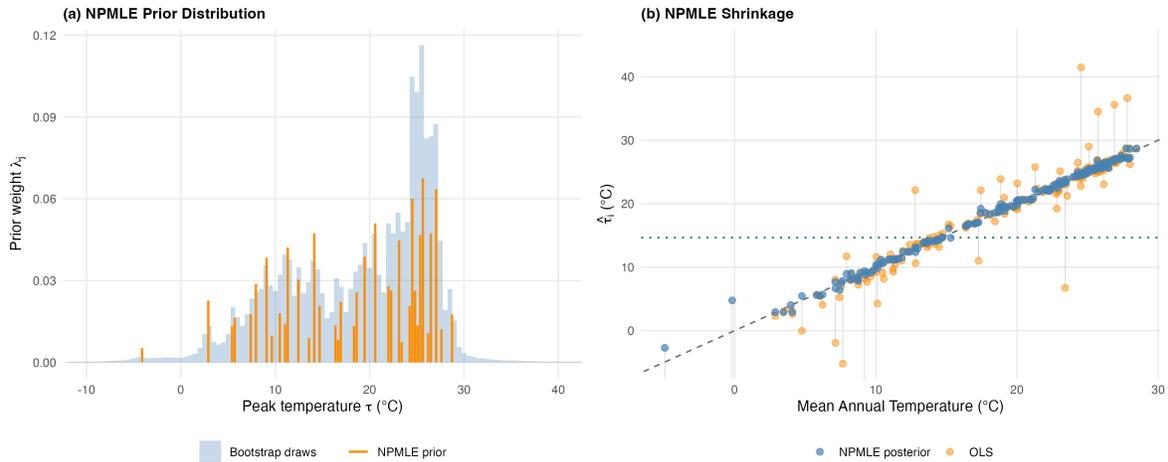


Figure A16: **Estimation of the NPMLE  $\tau_i$  prior and resulting posterior means for GDP**

Panel (a) shows the parametric bootstrap distribution of  $\tau_i$  draws across all  $(\hat{\beta}_{1i}, \hat{\beta}_{2i})$  in blue, and the estimated prior  $G(\cdot)$ . The right panel shows posterior means (blue) and OLS point estimates (orange). Little shrinkage occurs.

distributions for each  $\hat{\tau}_i$  are far too diffuse to discipline a prior. The resulting spiky prior estimate is diffuse across the distribution of  $\hat{\tau}_i$  and mirrors its density. Subsequently, posterior means essentially align with the medians of the likelihood, simply shrinking OLS outliers towards the 45 degree line.

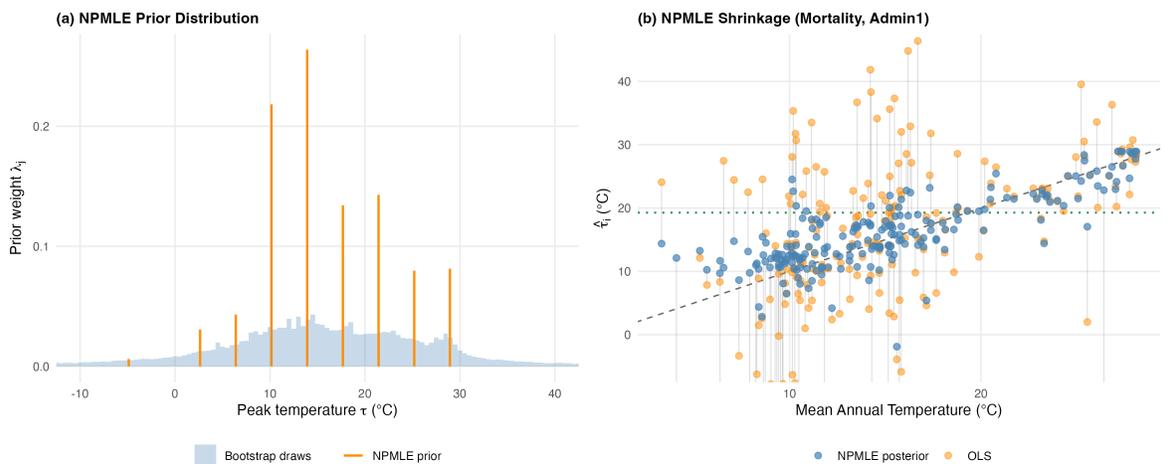


Figure A17: **Estimation of the NPMLE  $\tau_i$  prior and resulting posterior means for mortality**

Panel (a) shows the parametric bootstrap distribution of  $\tau_i$  draws across all  $(\hat{\beta}_{1i}, \hat{\beta}_{2i})$  in blue, and the estimated prior  $G(\cdot)$ . The right panel shows posterior means (blue) and OLS point estimates (orange). Little shrinkage occurs.

Appendix Figure A17 shows the same plot for mortality. The estimated prior performs equally as poorly as it did for the GDP data.