Assessing the statistical significance of income inequality differences

Stephen P. Jenkins
(LSE, EUI, and IZA)
Email: s.jenkins@lse.ac.uk

Joint work with Nicolas Hérault (Université de Bordeaux)

Preliminary draft paper available on request

LSE

Background and motivations

- You might think that it is straightforward to assess whether income inequality trends over time are statistically significant …
    - Simply derive estimates of indices and their standard errors for each year using standard methods and then do t-tests of inequality differences for pairs of years
    - Just as you might do for assessing the statistical significance of trends in other socioeconomic indicators
- However, this strategy is flawed …

LSE

Background and motivations

- Conventional approaches to assessing the statistical significance of inequality differences have long been criticized as having poor statistical performance even in large samples

- See, e.g., Cowell and Flachaire (2007, 2015), Davidson (2012), Davidson and Flachaire (2007), Schluter (2012), and Schluter and van Garderen (2009)

Background and motivations

- Tests based on standard asymptotic (linearization) methods and standard bootstrap methods are not reliable
  - ‘*t*-statistics’ do not tend towards *N*(0,1) but towards skewed distributions
  - Examples from Schluter (2012) for GE(0.05), GE(1.05), GE(2) indices →
- The source of the problem is that income distributions are typically *heavy-tailed*
  - Skewed long right-hand tail with Pareto-like shape at the top, and also …
  - Right-hand tail in survey data is sparse and may contain influential outlier observations

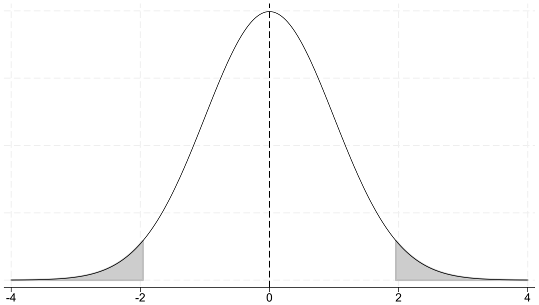

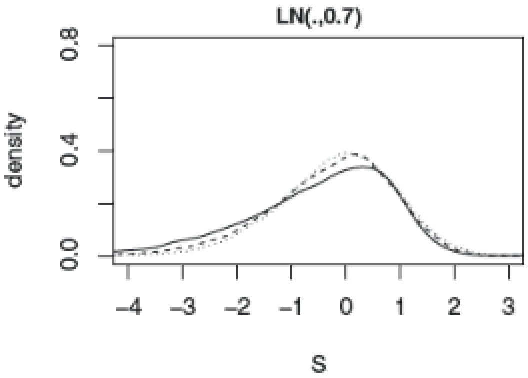

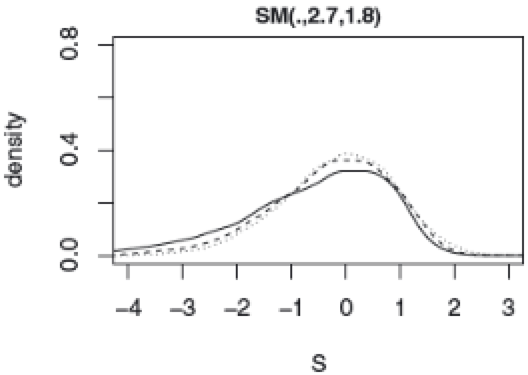

LSE

Proposals for better inference

[All use simulation analysis to verify properties. More details later]

Davidson & Flachaire (*J E'metrics*, 2007), Cowell & Flachaire (*J E'metrics*, 2007)

- Semiparametric asymptotic (= first stage of …)
- Semiparametric percentile-*t* bootstrap
- M out of N ('Moon') bootstrap (less favoured by them)

Ibragimov et al. (2010, 2021), also evaluated and applied by Midões & de Crombrugghe (*JEI*, 2023)

- Student-*t* tests based on sample splitting (problematic?)

Dufour, Flachaire & Khalaf (*JBES*, 2019)

- Permutation tests (but don't work with weighted data)

Schluter & van Garderen (*J E'metrics*, 2009)

- Normalizing transformations (only application we've seen)

Our main contribution

- Comparison of the improved inferential methods in a real-world setting, using yearly UK household survey data for 1977–2018 (42 years), not SimData world
  - Has inequality changed in a statistically significant sense?
  - Do the methods provide similar or different conclusions?
  - Recommendations about which method to use because some are more complex to implement than others
- Important because income inequality statistics are key social indicators
  - Inequality levels and trends are the subject of public attention
  - Relatively long time series are available, as well as the unit record survey data used to derive them

LSE

UK inequality changes – are they significant?

- Unit record data from UK ONS (equivalized household income among individuals): more later

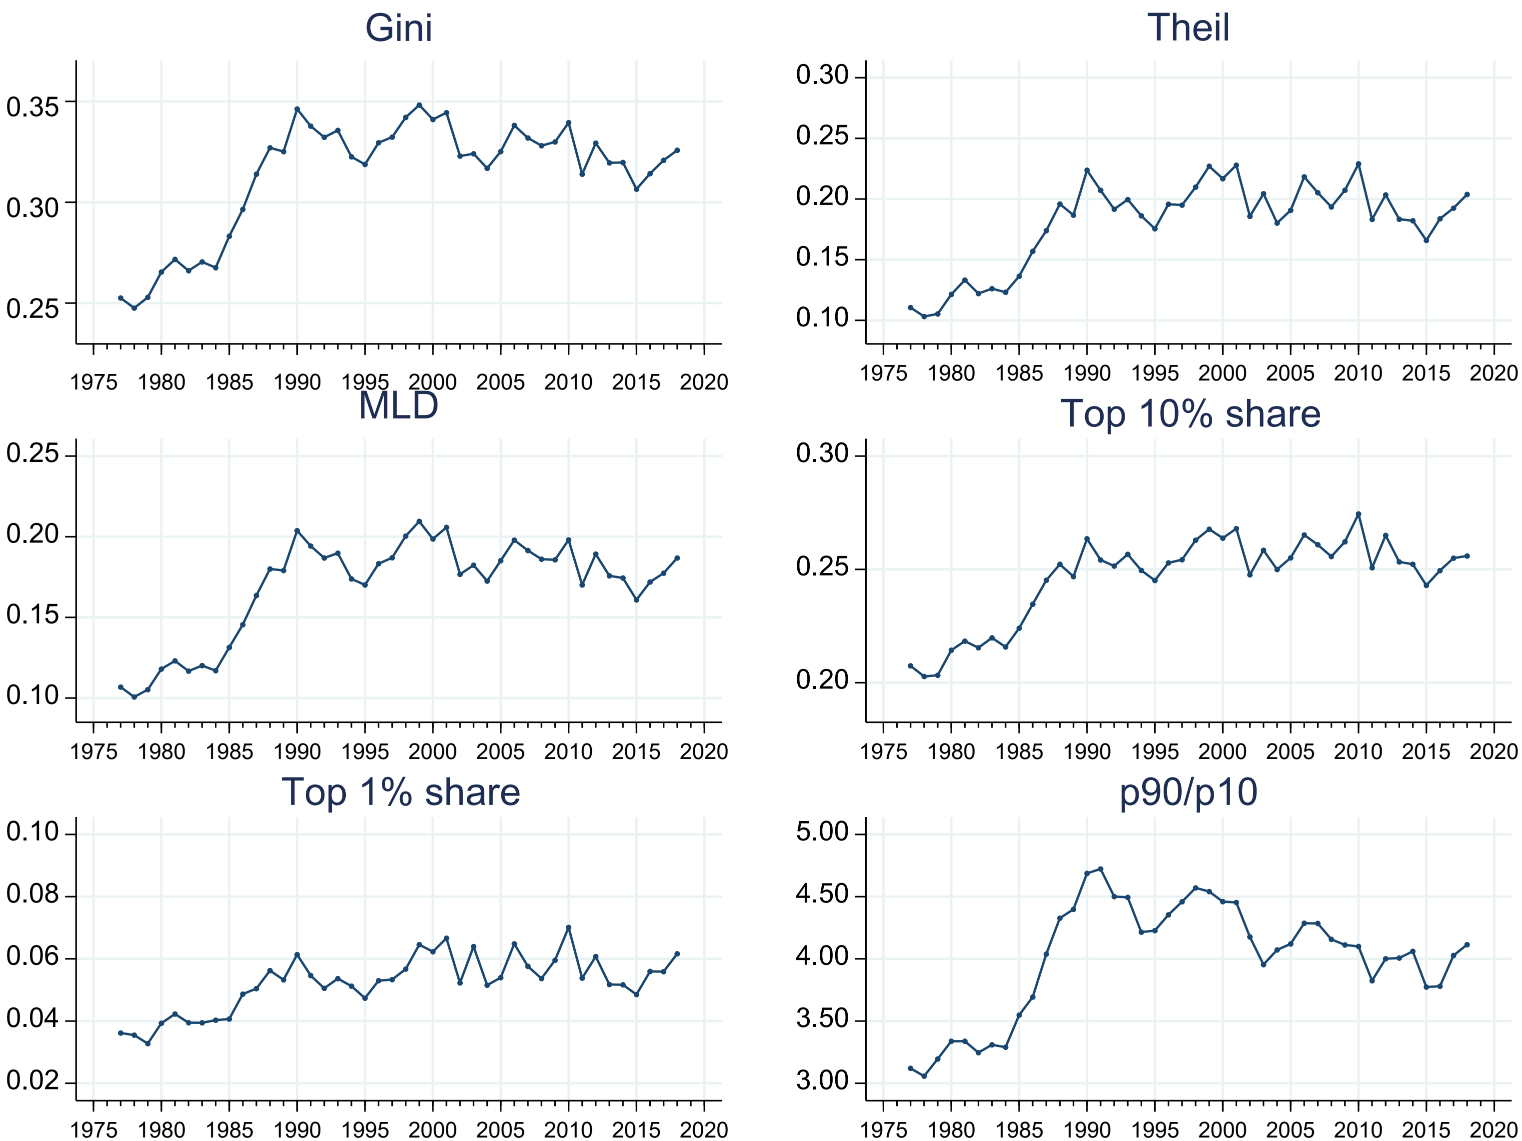

Go to
GE(−1)
GE(2)

Additional contributions

- We extend the proposed inferential methods to increase their useability
- Unit record data approach to estimation rather than employing method of moments estimators as Davidson and Flachaire (2007) and Cowell and Flachaire (2007), which means …

1. You can incorporate survey weights
   - Weights are ignored by all the improved inference papers even though virtually all household surveys contain weighting variables
   - Not using the weights when deriving inequality indices (or other descriptive statistics) leads to biased estimates
   - Could also incorporate survey clustering and stratification – (relevant for SEs) but information not available in our datasets

LSE

Additional contributions

- We extend the proposed inferential methods to increase their useability
- Unit record data approach to estimation rather than employing method of moments estimators as Davidson and Flachaire (2007) and Cowell and Flachaire (2007), which means …

2. You can easily use a wide(r) range of inequality indices
    - Almost all research proposing revised inference methods has focused on the Theil index or a few other Generalized Entropy indices
    - Our approach: almost any index you like, including the Gini, *p*90/*p*10, top income shares (as in earlier slide)
        - Useful to be able to consider the different judgements built into different indices (do inequality trend assessments depend on the index chosen?)

LSE

Our paper versus others …

- Closest in spirit to Alfons, Templ, and Filzmoser (*JRSSC*, 2013)
  - *Like them:* unit record data approach and incorporation of survey weights, plus use of robust estimators of Pareto distribution shape parameters
  - *But* Alfons et al. only consider semiparametric asymptotic estimators and not also semiparametric percentile-t estimators, they measure inequality using the Gini coefficient alone, their substantive application examines only two countries (Belgium and Austria) for two years (2005, 2006), and they do not formally test for inequality differences
  - *We* compare a larger portfolio of estimators and inequality indices and undertake formal tests of inequality differences on a relatively large scale (for 42 years of UK survey data)
- Ibragimov et al. (arXiv, 2021) and Midões and Crombrugghe (*JEI*, 2023) also use some real-world data …
  - They assess inequality differences (for one year) across Russian regions, and between two Russian surveys, and compare findings of several inferential methods albeit with a Student-t focus
  - *But* only for the Gini coefficient and Theil index respectively; and they use unweighted survey data
  - *We* make more extensive comparisons, use wider range of indices, weights

Headline findings (1)

1. All 3 of the semiparametric methods we consider yield similar conclusions about the statistical significance of inequality changes in the UK
   - All 3 differ from the conventional asymptotic approach which provides less precise estimators
   - Evidence suggests practitioners can use the semiparametric asymptotic approaches for inference rather than the more complex semiparametric bootstrap variants

LSE

Headline findings (2)

2. The large inequality rise over the 1980s was statistically significant, as expected, but …
3. Application of semiparametric approaches to inference substantially increases the number of statistically significant inequality differences for pairs of years in the 30-year period following the late-1980s
   - Comparing pairs of Gini coefficients rounded to two decimal places (as in UK DWP *HBAI* reports) is a relatively conservative approach to assessing the statistical significance of inequality changes

LSE

Outline

1. Five approaches to estimation and inference
2. Data
3. Results
4. Conclusions

plus

5. Additional slides

Approaches to estimation and inference

1. Conventional asymptotic
2. Semiparametric asymptotic
3. Semiparametric bootstrap percentile-t: DF
4. Semiparametric bootstrap percentile-t: CF
5. Student-t

Conventional asymptotic method

Given unit record data from a household survey:

Estimation of inequality index (point estimate)

1. Direct calculation from unit record data (most packages), or
2. Method of moments for distribution F
   - E.g., Theil index $T(F) = (v_F/\mu_F) - \log(\mu_F)$,

     where $v_F = E_F[Y\log(Y)]$ and $\mu_F = E_F[Y]$

     Replace moments v_F, μ_F by their sample counterparts (from data)

Variance of estimate

1. Via linearization/influence function method applied to unit record data
2. Method of moments: expression via delta method; variance estimate is function of covariance matrix for moments

- Both consistent, asymptotically normal, given standard assumptions (top-income problems aside)

Conventional asymptotic method: testing

- To test the hypothesis of equality for pair of years *A* and *B*, the Studentized (t-type) test statistic W^* is given by:

$$W^* = \frac{\hat{I}_B - \hat{I}_A}{[\hat{V}(\hat{I}_A) + \hat{V}(\hat{I}_B)]^{0.5}}$$

where the $\hat{I}$ and $\hat{V}$ are the inequality index and variance estimates (Davidson and Flachaire, 2007, eqn. 16)

- The *p*-value for the null hypothesis of no difference in inequality according to index *I* is:

$$P^* = 2N(-|W^*|)$$

where *N*(.) is standard normal CDF

LSE

Semiparametric asymptotic method

1. Fit Pareto distribution to the top *ptail* fraction of the distribution, i.e., fraction of obs with $R = 1$, where R means 'Rich'
   - Addresses heavy-tail and outlier/sparsity issues by 'filling in' at the top
2. Income distribution analyzed is a combination of observed incomes for $R = 0$ cases and Pareto for $R = 1$ cases
3. Estimation of index and its variance can be done in 2 ways:

   a) Method of moments (D&F/C&F): Derive moments for full distribution by combining empirical moments for $R = 0$ cases and, for $R = 1$ cases, the moments implied by the fitted Pareto distribution
      - C&F have index-specific formulae for GE family of indices but only for the case of unweighted data
      - If have survey weights, formulae exist in principle but very complicated (not considered by D&F/C&F)

   b) Unit record data method (what we use) Observed data for $R = 0$ but, for $R = 1$, take random draws from the fitted Pareto distribution, then allocate to the r^{th} richest imputed value the weight of the r^{th} richest observed value
      - This is essentially the Alfons et al. (*JRSSA*, 2013) 'Replacement of non-representative outliers' method'
      - But we also improve coverage of top by drawing $M = 10$ for each R=1 obs and modify the survey weights accordingly (cf. Blanchet, Flores, Morgan, *JEI*, 2023)
      - Because we have a unit record distribution of incomes (and weights), we can apply the conventional asymptotic methods for unit record data (previous slide)
      - Easy to calculate any index and its SE, i.e., not restricted to only e.g. GE family (D&F/C& focus) or Gini (Alfons et al.'s focus)

LSE

Semiparametric asymptotic method: testing

- Test statistic W^* and the p-value for the hypothesis of no inequality difference are calculated as for conventional asymptotic approach (above) …
- … Except that the inequality and variance estimates now refer to estimates derived from semiparametric distributions

Semiparametric percentile-t bootstrap method: D&F → DF (our variant)

Steps:

1. Derive an income distribution using the semiparametric asymptotic approach (as described earlier) and compute each index and its variance
2. Construct a bootstrap sample, b, and calculate bootstrap sample estimates:
   - D&F and DF: Each bootstrap sample combines a standard bootstrap sample of $R = 0$ obs plus, for $R = 1$ obs, random draws from the Pareto distribution fitted at step #1
   - Calculate indices and variance estimates from bootstrap sample b
     - D&F use the method of moments
     - DF: we use the unit record data (allowing for weights as above)
3. Repeat Step 2 B times and calculate bootstrap test statistics
4. Repeat Steps 1–3 for each year's data, and
5. Undertake tests for pairs of years (see below)

Semiparametric percentile-t bootstrap method C&F → CF (our variant)

- C&F/CF #1: derive semiparametric estimate and variance (as for D&F), including fitting of a Pareto distribution to $R = 1$ cases
  - But the next steps (C&F and CF) also allow for the uncertainty introduced by having to estimate the Pareto shape parameter (unlike D&F/DF)
  - Each bootstrap sample uses a different Pareto distribution to model the data for the $R = 1$ cases (rather than the same one)
- **C&F #2**: Take a bootstrap sample b of *all* cases
  - (i) Fit a Pareto distribution to the $R = 1$ obs in bootstrap sample b; construct a semiparametric income distribution combining observed sample data for $R = 0$ obs and, for sample $R = 1$ obs, draws from Pareto distribution with shape parameter $\widehat{\theta}_b$;
  - (ii) calculate index and variance estimates using method of moments
- **CF #2: As C&F #2, except …**
  - (i) $\widehat{\theta}_b$ is a random draw from the Pareto shape parameter fitted at Step #1, with distribution assumed ~ $N(\widehat{\theta}, \mathrm{SE}(\widehat{\theta}))$
  - (ii) calculate index and variance estimates using unit record data (as per DF)
  - Substep (i) allows for uncertainty (as C&F) but is faster and more robust, we found, as not refitting at each replication and fewer convergence issues
- C&F/CF #3, #4, #5 : as for D&F/DF

Semiparametric percentile-t bootstrap methods: testing

- The Studentized test statistic for each pair of bootstrap samples, *b*, is:

$$W_b^* = \frac{\left[\left(\hat{I}_{Bb} - \tilde{I}_B\right) - \left(\hat{I}_{Ab} - \tilde{I}_A\right)\right]}{\left[\hat{V}\left(\hat{I}_{Ab}\right) + \hat{V}\left(\hat{I}_{Bb}\right)\right]^{0.5}}$$

  - The "~" estimates are the ones from Step #1
  - "[T]he numerator is recentred so that the statistic tests a hypothesis that is true for the bootstrap samples" (Davidson & Flachaire, 2007, p. 146)
- The bootstrap *p*-value for the test of no difference in equality is the proportion of bootstrap samples for which the bootstrap statistic is more extreme than the one calculated from the original (semiparametric) data, W^*:

$$P^* = \left(\frac{1}{B}\right)\sum_{b=1}^{B} \iota\left(\left|W_b^*\right| > \left|W^*\right|\right)$$

where $\iota(.)$ is the indicator function (Cowell & Flachaire, 2007, eqn. 18)

Researcher choices (i)

Pareto: definition of the $R = 1$ group to which model fitted?

- Trading off bias and variance: literature on methods for choosing k richest obs for fitting (and hence lower threshold defining Pareto)
- Instead of choosing k for each year separately, we fix ptail because we have 42 years of data and survey weights
  - Consistent, economical rule for large number of comparisons
  - Atkinson & Jenkins (2022), Jenkins (2017), and C&F/D&F also use rules of thumb, e.g. D&F: $k = \text{sqrt}(N)$
- We focus on case of ptail = 5%, in which case min(k) = 239 and max(k) = 339 households
  - We also have results for ptail = 1% (k much smaller), with more larger swings in CIs from one year to next, and smaller p-values; conclusions similar anyway

Researcher choices (ii)

Pareto: which fitting method?

- D&F/C&F and others use the Maximum Likelihood (Hill) estimator
- We use a robust (Optimal B-Robust Estimator, OBRE) method (i.e., modified ML) though, in practice, estimates differ little from ML
  - Robustness refers to how influenced by ‘contamination’ (high leverage outliers at the top)
  - Brzezinski (2016) Monte-Carlo study: OBRE performs well compared to 4 competitors including that used by Alfons et al. (2013)

LSE

Researcher choices (iii)

Which inequality indices?

- We want a wide range to allow a range of value judgements as well as those in common use (e.g., Gini, $p90/p10$)
  - *D&F*: focus entirely on Theil = GE(1)
  - *C&F*: focus on GE(α) indices, bottom- to top-sensitive: GE(–1), GE(0), GE(0.5), GE(1), GE(2)
  - *Alfons et al*. and *Ibragimov et al.:* Gini
  - *MdeC*: Theil
- Our derivations use: Gini, GE(0) = MLD, GE(1) = Theil, income share of top 10%, income share of top 1%, plus $p90/p10$, GE(–1), GE(2) = half-CV^2
  - But not all are reported in the main text because …

LSE

Researcher choices (iv)

Heavy-tailed distributions and their moments

- Existence of Pareto distribution moments depends on $\hat{\theta}$
  - The more tail heavy the distribution, the more moments that don't exist, e.g. if $\hat{\theta} < 2$, infinite variance; if $\hat{\theta} < 1$, infinite mean
  - For index GE(α), require $\hat{\theta} \geq \max(2, 2\alpha)$ to calculate SEs via method of moments
  - Hence top-sensitive indices like GE(2) problematic – require $\hat{\theta} \geq 4$ – and this rarely the case
  - And these top-sensitive indices very affected by high-income outliers anyway (C&F)
- Problems with a unit record data approach for related reasons: some moments based on fourth powers
  - We had computational problems/non-robustness implementing the C&F method, motivating some of our modifications of it described earlier
  - So, we do not report GE(2) estimates in the main text, nor GE(–1) estimates (they're reported in the appendices at the end)

LSE

Researcher choices (v)

Number of bootstrap repetitions, B

- D&F/C&F/MdeC: $B = 199$
- We use $B = 999$
  - Trade-off between reducing bootstrap uncertainty and increasing computational burden

LSE

Student-*t* approach (Ibragimov et al., MdeC)

For each year's data

1. Split sample randomly into q groups, with $q \geq 2$
2. Calculate each inequality index I separately for each group
   - Can use any inequality index (but they focus on Gini, Theil)
3. Derive overall estimate of I as simple group-average of the $\hat{I}$, and its variance as sample variance of $\hat{I}$
4. Use these components as inputs to usual pairwise t-tests

- MdeC: (a) "simple, intuitive, and computationally cheap", and also claim (b) has good properties (from simulation evidence about test size and power)
- Researcher choice of q: MdeC highlight the case of $q = 8$
- We've applied the approach but do not show full results
  - We have questions about the approach's robustness to data sort order, and the consistency of estimates in practice (see below)

LSE

Data

UK, 1977–2018 (42 years)

Data: UK, 1977–2018

- Same data as used in ONS official statistics: 'Effects of taxes and benefits on household income' (ETB) statistics and report (annual)
  - Our dataset is the version ONS used prior to introduction of top-income adjustments; long consistent series up to onset of Covid-19
- Data from annual Living Costs and Food Survey (LCFS) and its predecessors
  - Nationally representative survey of UK private household population; independent samples per year
  - Survey weights adjust for household size, and non-response and are calibrated to population totals
    - Non-response/grossing-up aspect of weights introduced in 1996 dataset
  - No PSU or strata information available in the public-use file
  - $N \approx$ 6,500 households per year on average; N ranges from c. 4,900 to 7,500 (more on numbers in appendix table at end)

LSE

Data: definitions, selections

- Income: disposable/net income = gross income (from market, transfers) – direct taxes (income taxes, employee National Insurance Contributions, local taxes), equivalised using modified-OECD scale
  - Standard 'Canberra Group' definition, also used by Eurostat/SILC
- Unit of analysis = individual (each person attributed with equivalised income of his/her household)
  - Survey weights include household size
- Trimmed distributions – we drop a small number of incomes: (a) zero, negative, positive but less than one, and (b) highest income(s) if more than twice as large as the next highest income (only 9 households over 42 years)
  - Fraction of households dropped averages 0.23% yearly (never more than 0.48%): for more details, see appendix table at end
  - Analysis repeated using untrimmed data show effects of trimming most noticeable for GE(−1) and GE(2), as expected

Empirical results

Empirical results: running order

1. How (point) estimates differ by method
2. How precision of estimates differs by method
   - With some evidence to show why we don't proceed further with the Student-*t* approach
3. Tests of inequality differences between year-pairs (*A*, *B*)
   - Focus on test results for Gini and Theil for brevity
     - Results for other indices available in the additional slides at end
   - 42 years ⇒ 861 possible pairwise tests per index! So, we show:

   a) Count of number of years with *p*-value < 5% for test of no inequality difference between year *A* and every other year (*B*), by method and index (maximum count = 41)

   b) *p*-values for tests of no inequality difference between year *A* and every other year (*B*), for each $A \in \{1977, 1990, 2006, 2018\}$, by method and index

- NB in what follows, 'DF', 'CF' refer to our variants of D&F's and C&F's approaches

Inequality estimates, by index

- Semiparametric asymptotic series very similar to the conventional asymptotic estimates (red versus blue)
- Student-t series out of line with the asymptotic ones (green versus blue/red)

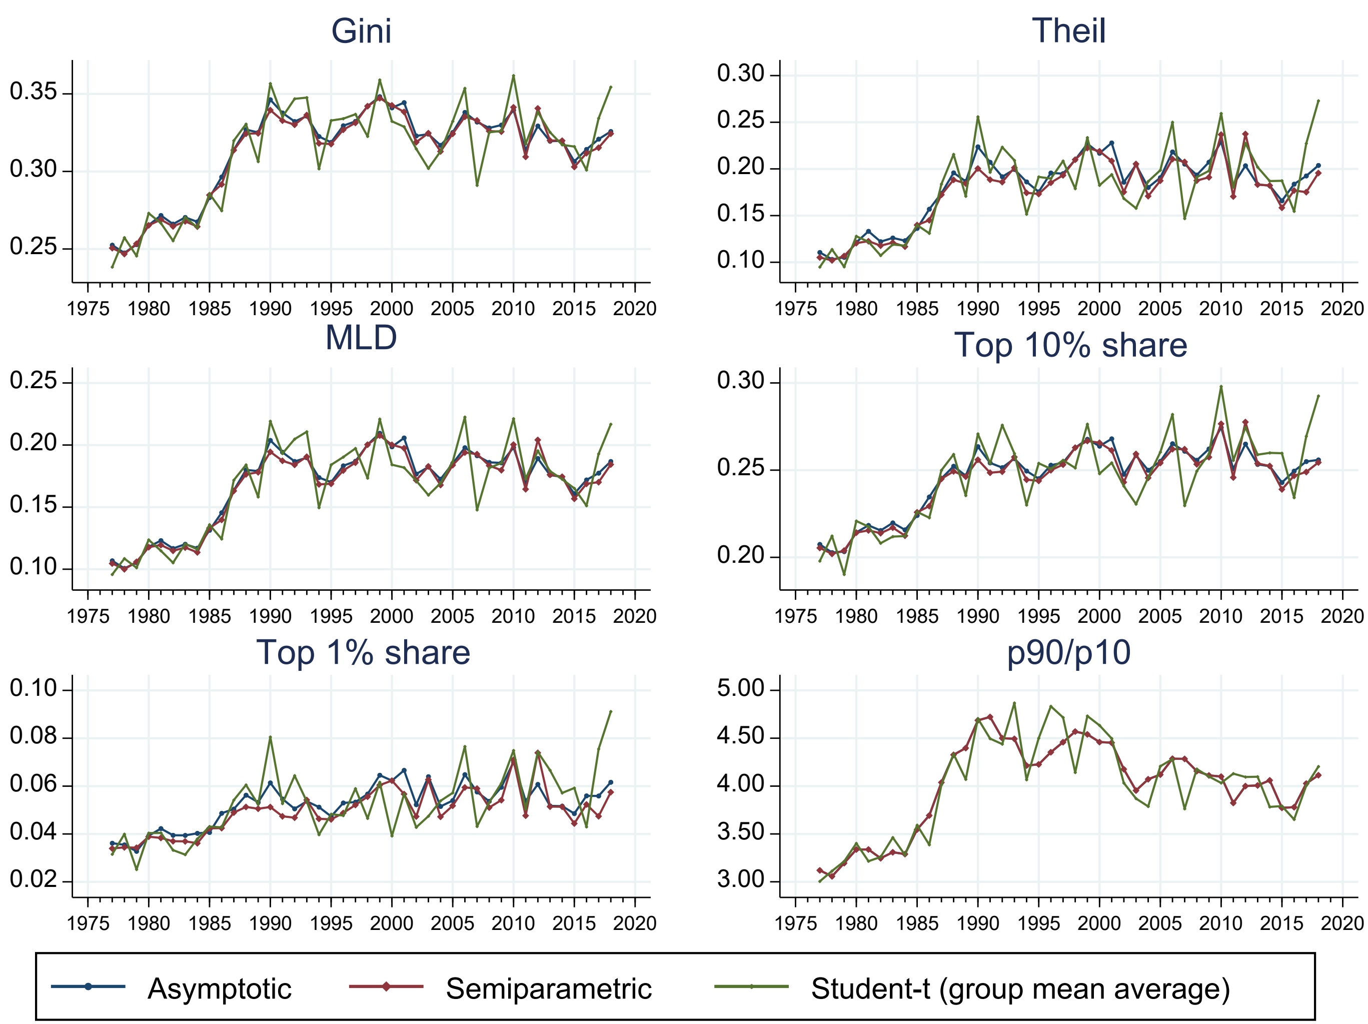

Go to
GE(–1)
GE(2)

Indicative precision: 95% CI 'spans' (CI as % estimate), 42-year average

- Larger values mean less precise
- Greater precision from semiparametric approaches relative to asymptotic
- Much lower precision for Student-t approach
- Much lower precision for GE(-1), GE(2) regardless of approach
- *p*90/*p*10: precision ~ same for all (except Student-t approach)

Index	Asymptotic (1)	Semiparametric asymptotic (2)	Semiparametric bootstrap (DF) (3)	Semiparametric bootstrap (CF) (4)	Student-t (8 groups) (5)
Gini	6.07	3.18	2.86	4.12	10.11
MLD = GE(0)	13.14	7.34	6.94	9.53	21.70
Theil = GE(1)	20.39	7.99	7.22	13.55	33.35
Top 10% share	7.58	3.10	2.60	4.98	12.50
Top 1% share	29.97	8.93	8.11	19.60	48.18
GE(−1)	28.88	27.31	49.68	49.21	50.98
GE(2)	46.63	19.49	27.03	44.15	80.56
*p*90/*p*10	6.88	6.18	6.16	6.23	10.72

Indicative precision: 95% CI 'spans': year by year

- Ordering is consistently as per Table
  - except for Theil index in 2010
- Semiparametric asymptotic and DF are the most precise

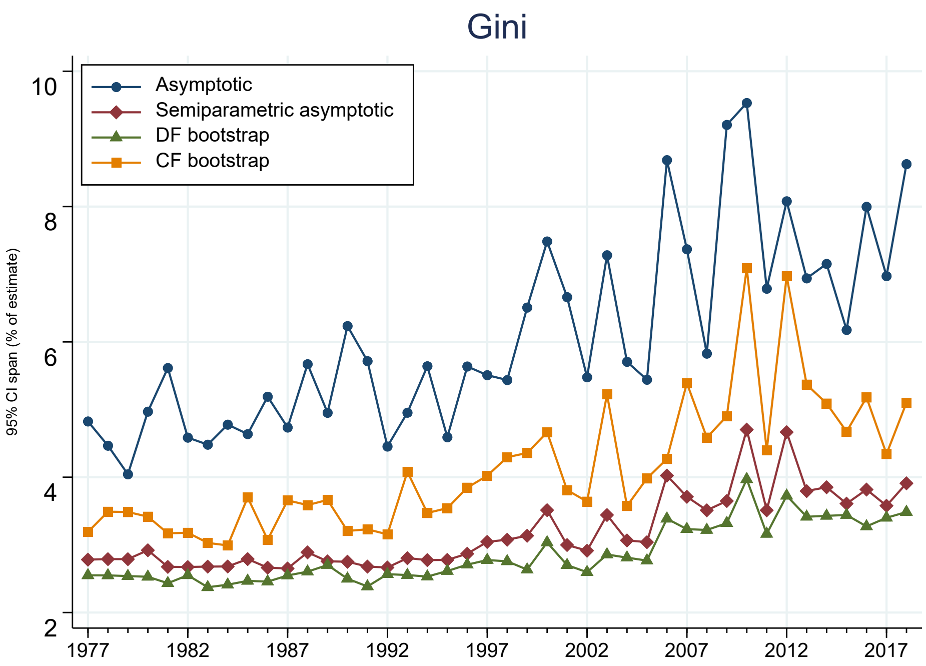

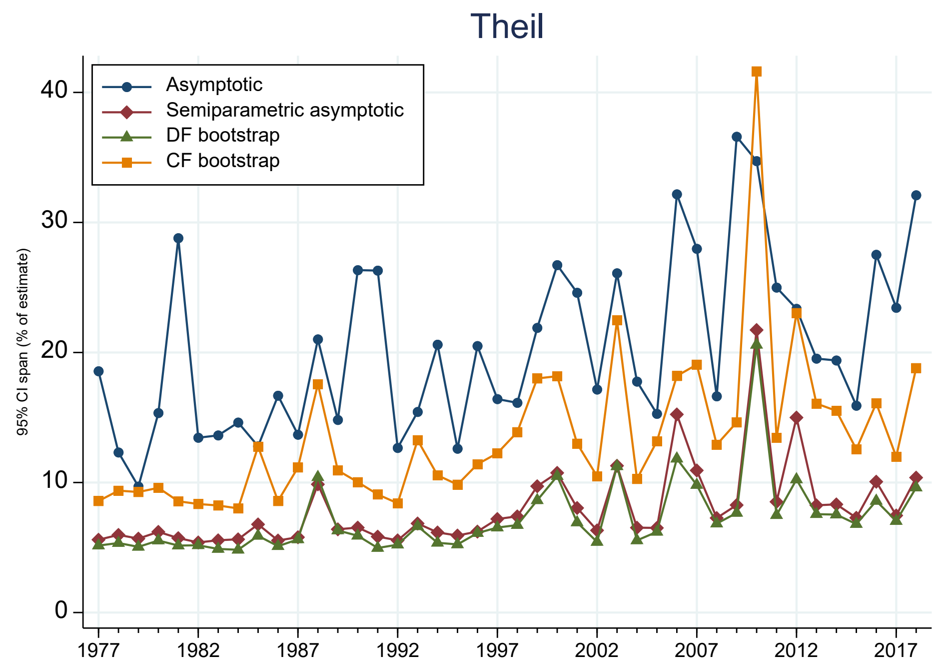

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: Gini

- DF/CF labels refer to our modifications to D&F/C&F

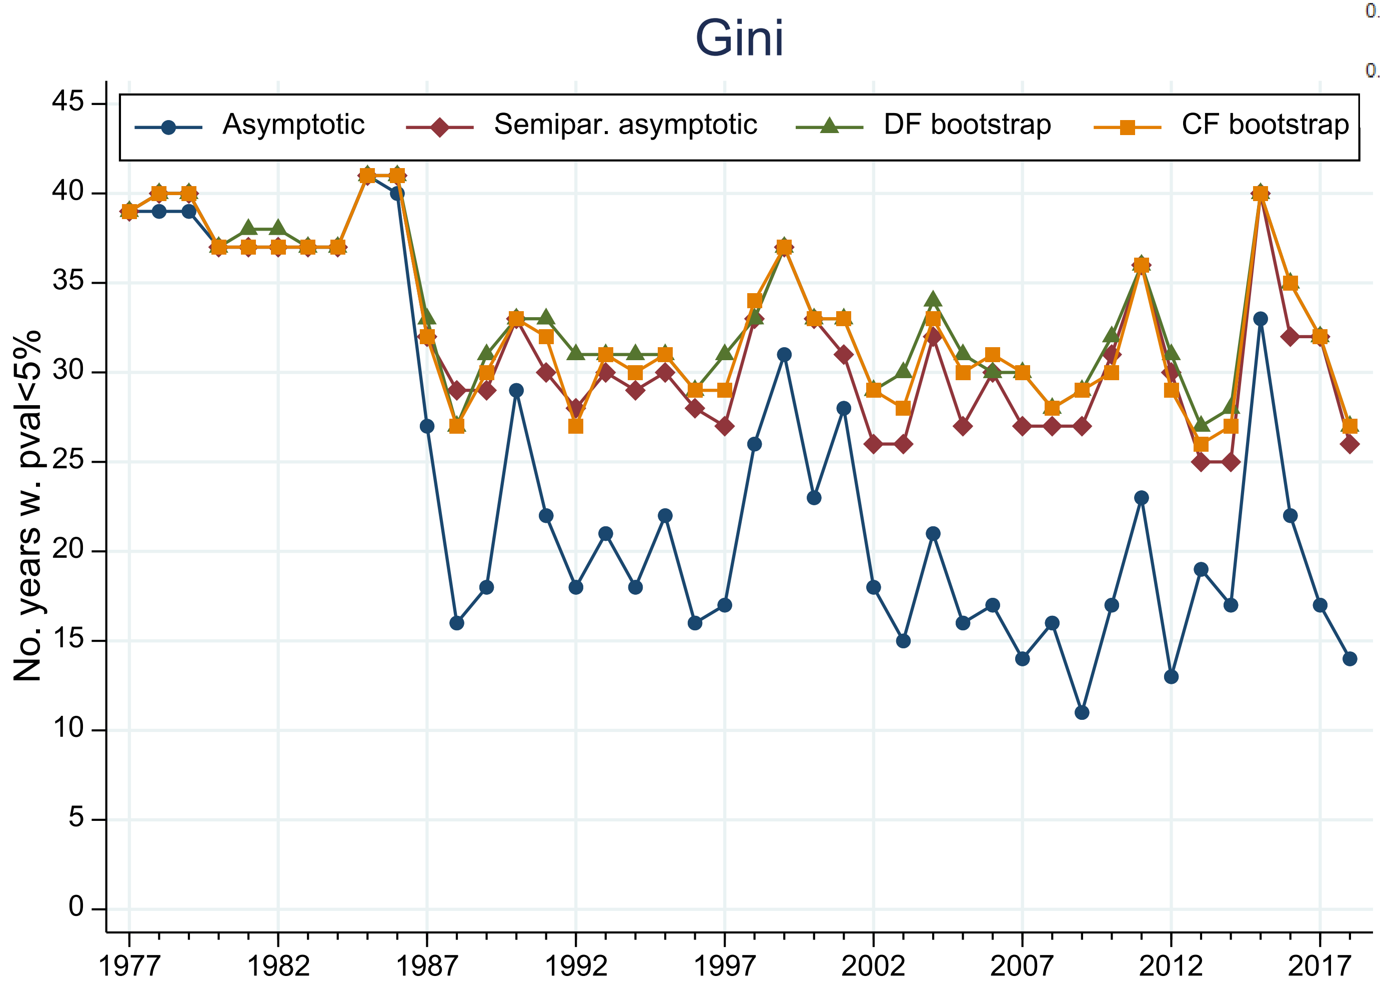

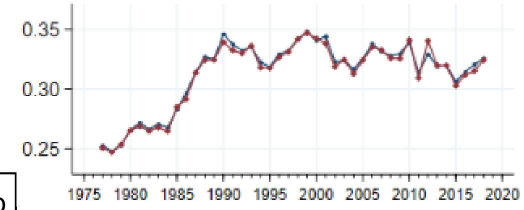

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: Theil

- DF/CF labels refer to our modifications to D&F/C&F
- Go to counts for GE(0), top 10% share, top 1% share, *p*90/*p*10, GE(−1), GE(2)

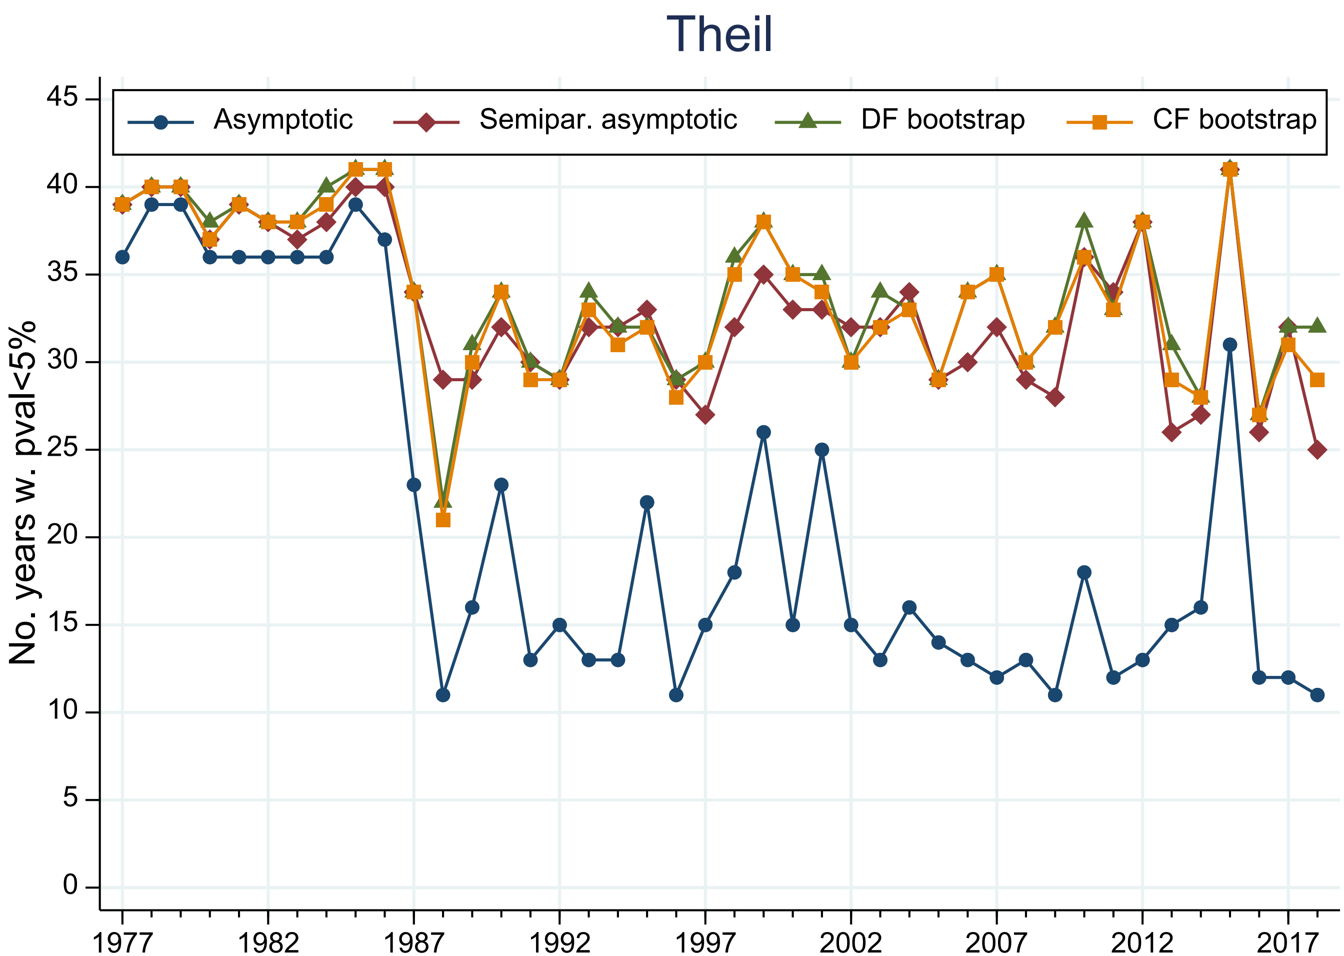

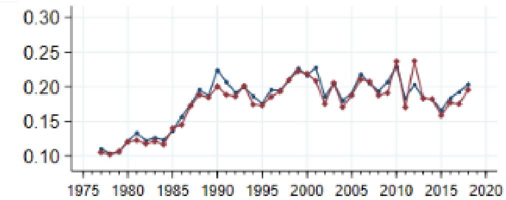

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **Gini**

- DF/CF labels refer to our modifications to D&F/C&F

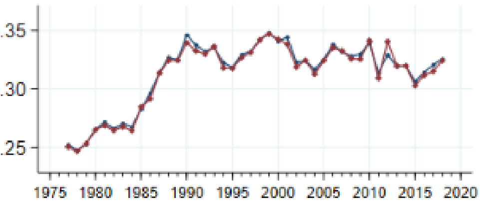

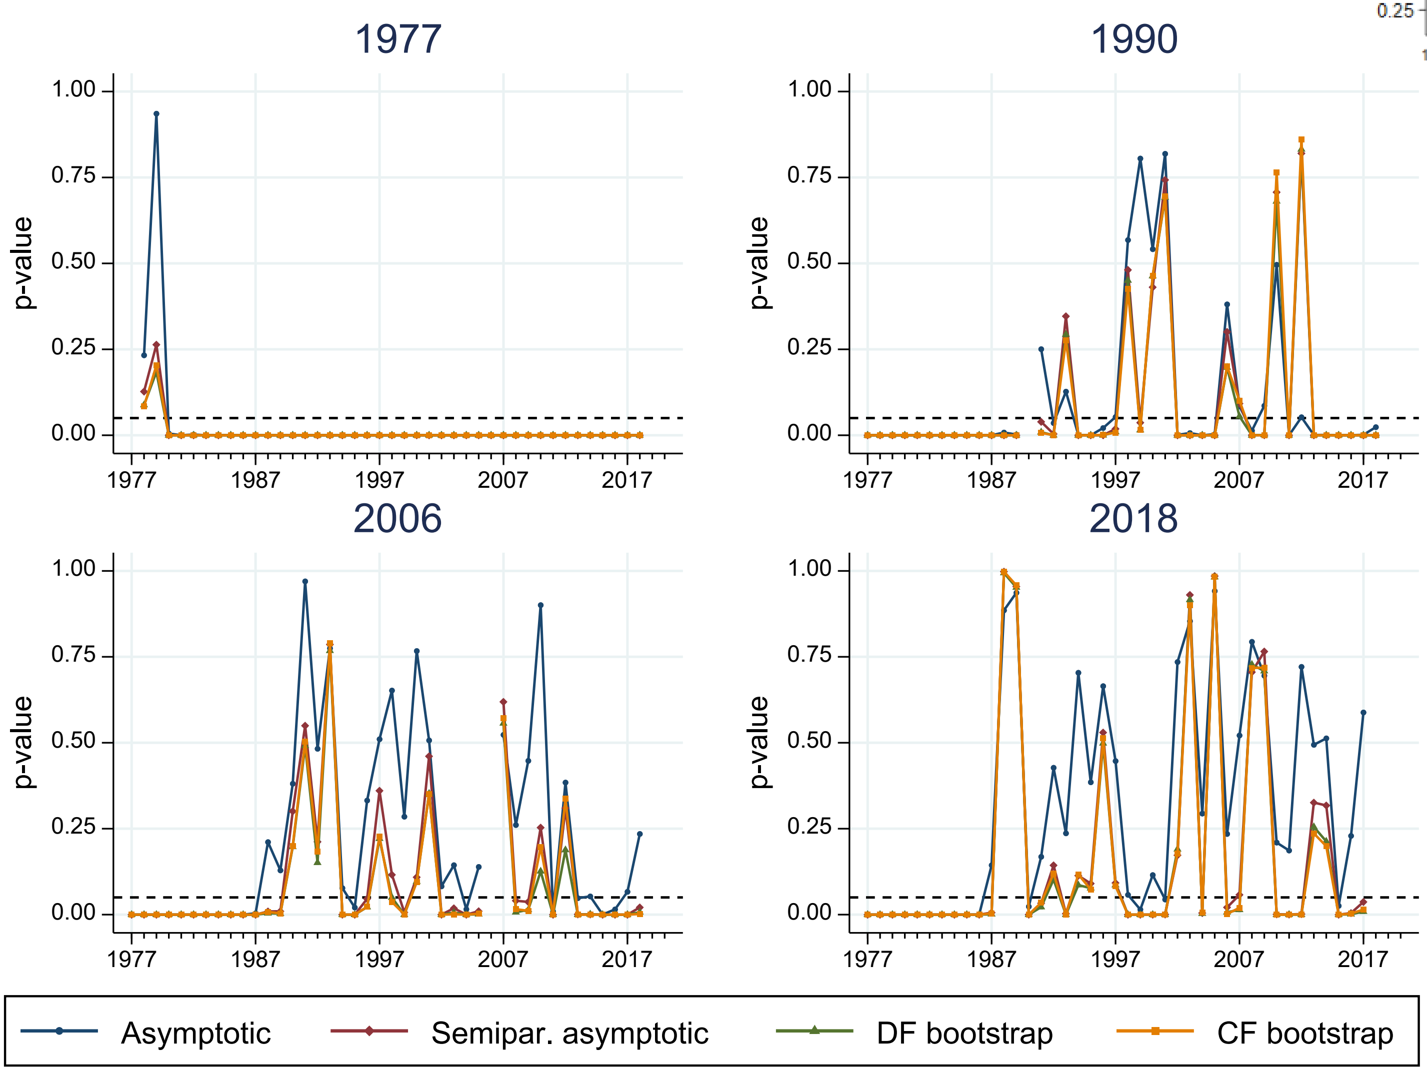

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **Theil**

- DF/CF labels refer to our modifications to D&F/C&F

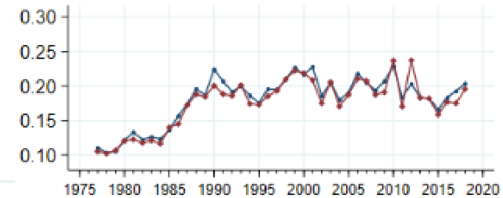

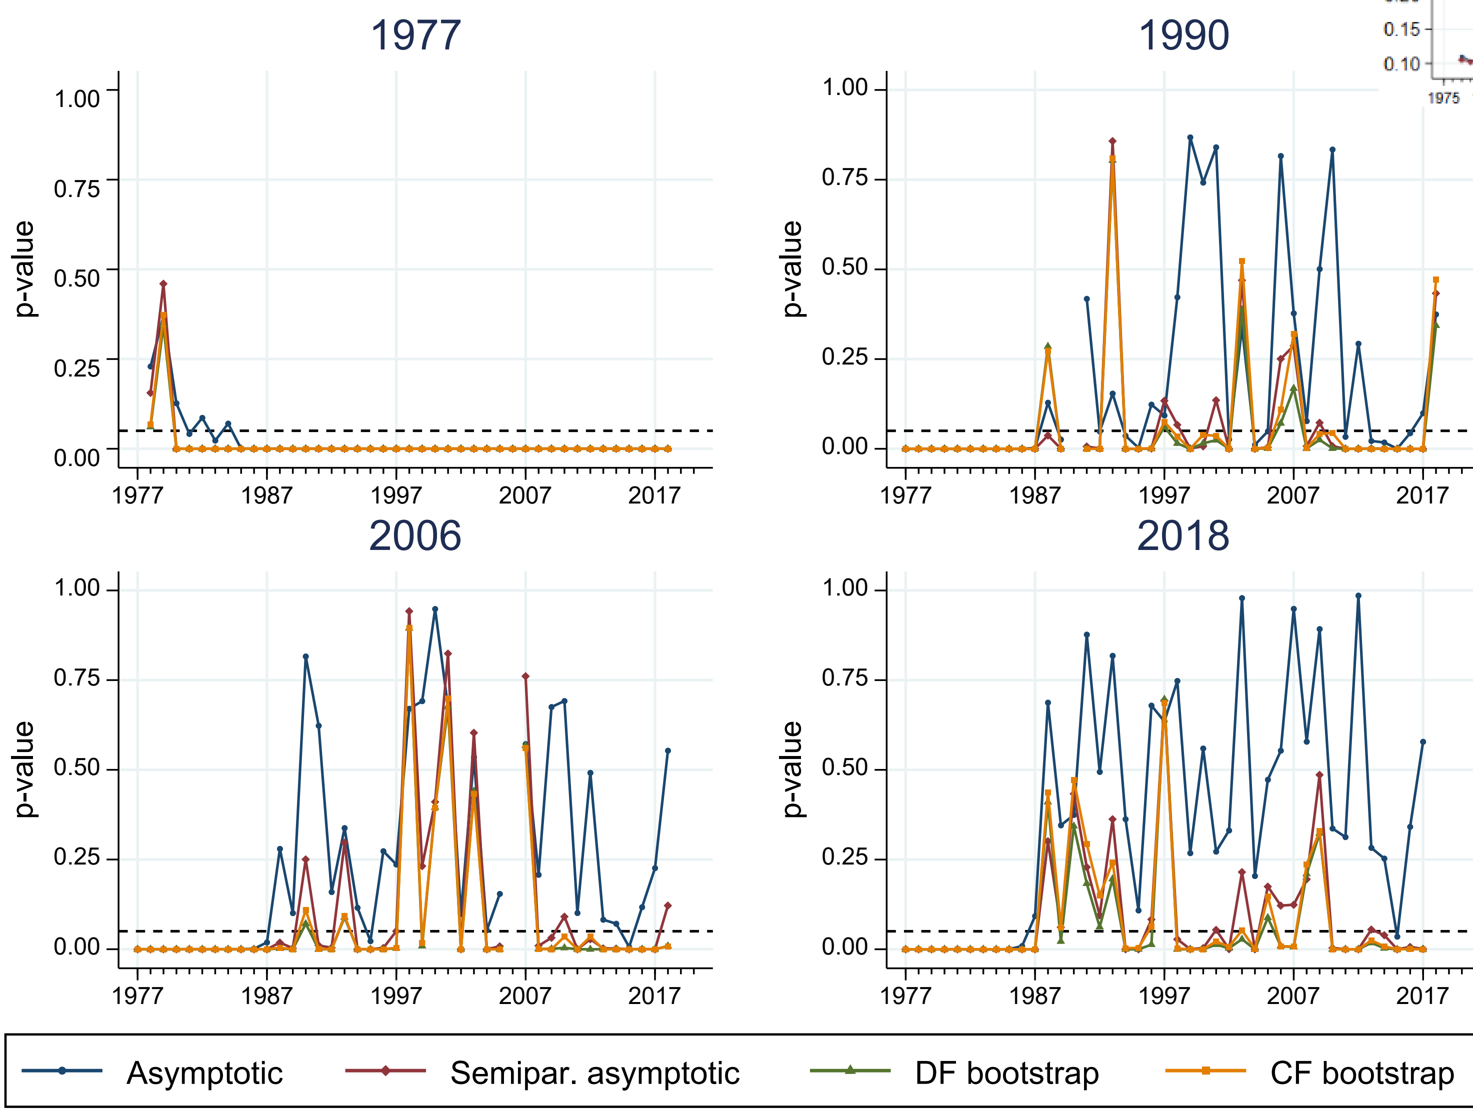

Go to p-values for GE(0), top 10% share, top 1% share $p90/p10$, GE(−1), GE(2)

Lessons for statistical agencies and practitioners? (1)

1. Statistical agencies and other researchers typically use the conventional asymptotic approach to calculate SEs … but none account for heavy tail issues

- Our analysis shows that semiparametric asymptotic could be used: performs well and easy to implement
  - And can be easily extended to account for clustering and stratification

Lessons for statistical agencies and practitioners? (2)

2. UK DWP's HBAI reports only publish Gini estimates rounded to 2 d.p. – does this provide a cheap and transparent way to guard against unwarranted claims of statistically significant inequality differences?

- In the UK context, changes in the Gini (and Theil) index rounded to the nearest percent may not be statistically significant, but there are multiple instances where they are
  - E.g., consider differences in Gini relative to the 1990 value (0.3396, which rounds to 34%): there are statistically significant differences between the 1990 Gini and the Ginis for 1996, 1997, 2008 and 2009 (rounded Gini = 33%) according to the semiparametric approaches
  - Similarly, for the Theil index (not reported by the ONS ETB report)
- So, from a sampling variability perspective, the presentational choice in the DWP's HBAI reports is conservative

LSE

Conclusions

Conclusions (1)

- All 3 semiparametric methods provide similar results, and results that differ from the conventional asymptotic method (which yields larger SEs; fewer significant differences)
  - Improvements relative to conventional asymptotic method particularly marked from late 1980s onwards
    - E.g., using semiparametric methods, for each year from 1987, around 30 of 41 pairwise comparisons with each other year is statistically significant for Gini and Theil (cf. ~ 15–20 for conventional asymptotic)
  - Similarity between semiparametric asymptotic and semiparametric percentile-*t* results – noted by Cowell & Flachaire (2007) from simulation analysis, but we've shown similarity occurs in the real world too! Similarly, …
  - Inference about GE(2) differences not improved using the semiparametric methods – problematic for all methods – and, to a lesser extent GE(–1): influential outliers at top (and bottom) … even after light trimming of data
  - *p*90/*p*10 inference/test outcomes similar for all methods including conventional asymptotic – unsurprising given nature of index?

LSE

Conclusions (2)

- Switching from methods of moments approaches to unit record data approaches has benefits
  - Can consider almost any inequality index you like
  - Can incorporate sample design features: weighting (also clustering, stratification)
- Need more simulation evidence about the performance of the different methods for two-sample tests
  - Davidson and Flachaire (2007) and Cowell and Flachaire (2007), and other researchers proposing improved inference, focus on one-sample tests when assessing performance of methods, and yet two-sample tests the most relevant for applied researchers
- Ongoing issues relating to whether/how income data should be pre-screened to remove egregious outliers
  - Are outliers random 'dirt' or genuine (cf. Cowell/Flachaire) – influential in either case
  - We've used a conservative trimming approach (more conservative than many)
  - NB using data in which we do not trim outliers (according to our rule) doesn't change results much – except for very top- or bottom sensitive inequality indices
- And … finally …

Conclusions (3)

- New inference methods are also required to address new developments
- The income inequality data published by the UK ONS now include a top-income adjustment (backdated to 2001), and the Department for Work and Pensions (Family Resources Survey), has included a top-income ('SPI') adjustment since 1992
- These top-income adjustments use *external information about top-incomes taken from income tax administrative data*
  - Our current paper and the 'improved inference' literature uses no external information (Pareto distributions for top incomes estimated from the survey data to hand)
- How to assess statistical significance of differences in inequality calculated from survey data that have been top-income adjusted using external data is an open issue: we're beginning to work on this using the revised ONS data series
- But NB suitable administrative record data on incomes not available for most countries in the world; they must continue to rely on the household survey data that are available, and the findings of this paper remain relevant for this common situation

LSE

Additional slides

1. Sample sizes (and numbers trimmed)
2. Test *p*-values (1977, 1990, 2006, 2018 versus every other year) for additional indices: MLD, top 10% share, top 1% share, GE(–1), GE(2)
3. Test *p*-values (Gini, Theil) for ptail = 1%

Sample sizes

and numbers of trimmed observations:

Number of outliers at bottom (income < 1) and

Number of outliers at top (top income > 2× next highest income)

Return to Data section

Year	Number of households in LCFS	Negative or zero income	Income between 0 and 1	Top-income outlier	(2) + (3) + (4) as % of (1)	Number of households in analysis sample	Weighted number of individuals
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1977	7,193	11	3	0	0.19	7,179	19,797
1978	6,996	7	2	0	0.13	6,987	18,939
1979	6,768	8	0	0	0.12	6,760	18,202
1980	6,942	9	3	0	0.17	6,930	18,799
1981	7,520	8	1	0	0.12	7,511	20,462
1982	7,420	7	2	1	0.13	7,410	19,900
1983	6,964	9	0	0	0.13	6,955	18,412
1984	7,077	9	0	0	0.13	7,068	18,498
1985	7,007	6	0	0	0.09	7,001	18,140
1986	7,175	7	0	1	0.11	7,167	18,285
1987	7,395	1	0	0	0.01	7,394	18,723
1988	7,264	7	1	1	0.12	7,255	18,244
1989	7,406	6	1	0	0.09	7,399	18,526
1990	7,038	12	0	0	0.17	7,026	17,333
1991	7,054	6	0	0	0.09	7,048	17,063
1992	7,417	10	2	0	0.16	7,405	18,144
1993	6,975	7	2	0	0.13	6,966	17,236
1994	6,849	16	0	0	0.23	6,833	16,550
1995	6,794	11	0	0	0.16	6,783	16,532
1996	6,413	9	0	0	0.14	6,404	57,712
1997	6,409	6	0	0	0.09	6,403	58,089
1998	6,629	14	1	0	0.23	6,614	58,215
1999	7,096	26	2	0	0.39	7,068	58,449
2000	6,634	20	0	0	0.30	6,614	58,691
2001	7,466	31	1	0	0.43	7,434	58,790
2002	6,926	31	1	1	0.48	6,893	57,700
2003	7,047	22	2	1	0.35	7,022	57,945
2004	6,794	15	1	0	0.24	6,778	58,059
2005	6,778	11	1	0	0.18	6,766	58,014
2006	6,387	18	2	1	0.33	6,366	58,457
2007	6,108	15	3	1	0.31	6,089	59,300
2008	5,764	14	5	1	0.35	5,744	60,307
2009	5,575	18	0	0	0.32	5,557	60,550
2010	5,253	10	1	0	0.21	5,242	61,408
2011	5,672	13	0	0	0.23	5,659	61,335
2012	5,456	12	1	1	0.26	5,442	62,805
2013	5,089	22	0	0	0.43	5,067	63,155
2014	5,095	19	0	0	0.37	5,076	63,514
2015	4,912	20	0	0	0.41	4,892	63,704
2016	5,041	21	0	0	0.42	5,020	64,243
2017	5,407	23	0	0	0.43	5,384	64,350
2018	5,473	24	2	0	0.48	5,447	64,801
Average	6,540	13.60	0.95	0.21	0.23	6,525	41,366

Results for additional indices

MLD

Share of top 10%, share of top 1%

$p90/p10$

GE(–1), GE(2)

Return to main

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: **MLD**

Back to main indices

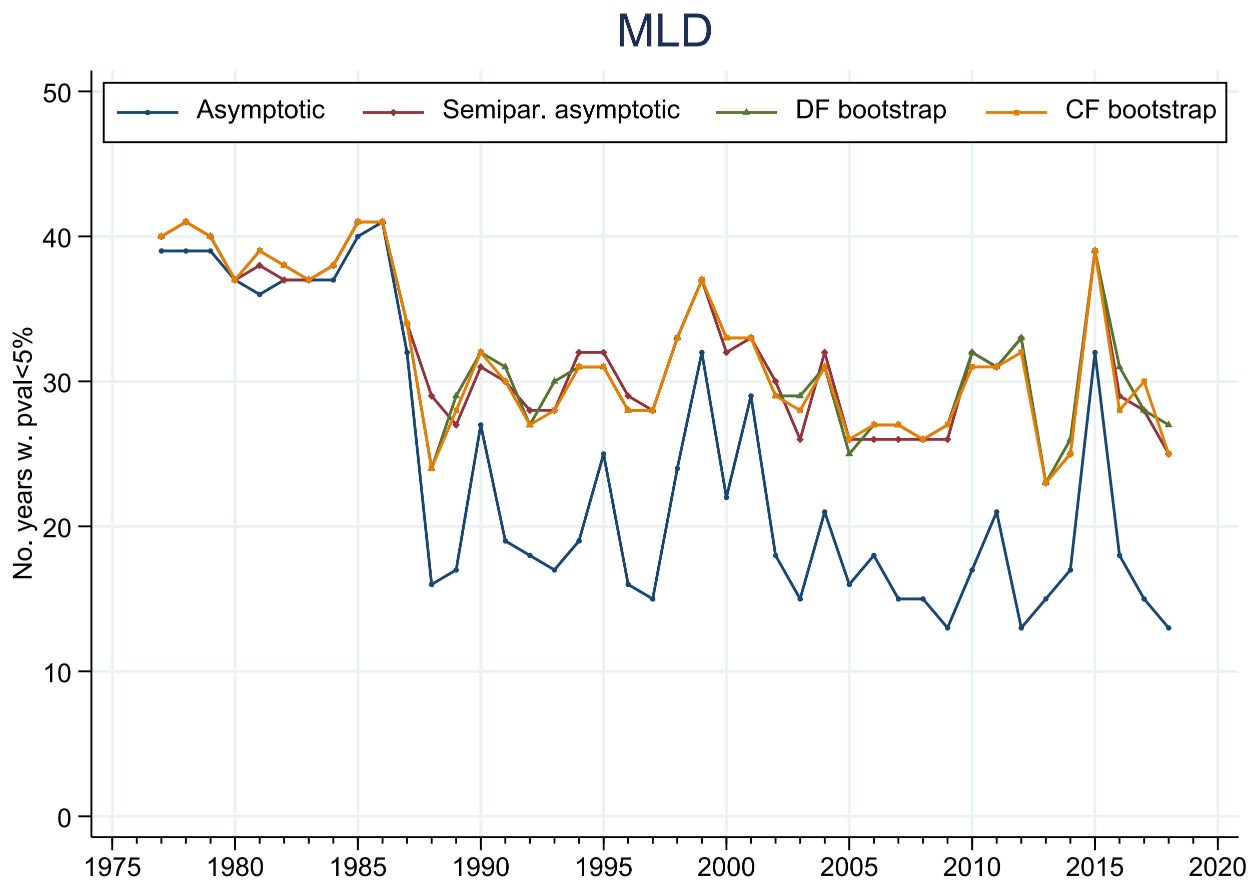

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **MLD**

Back to main indices

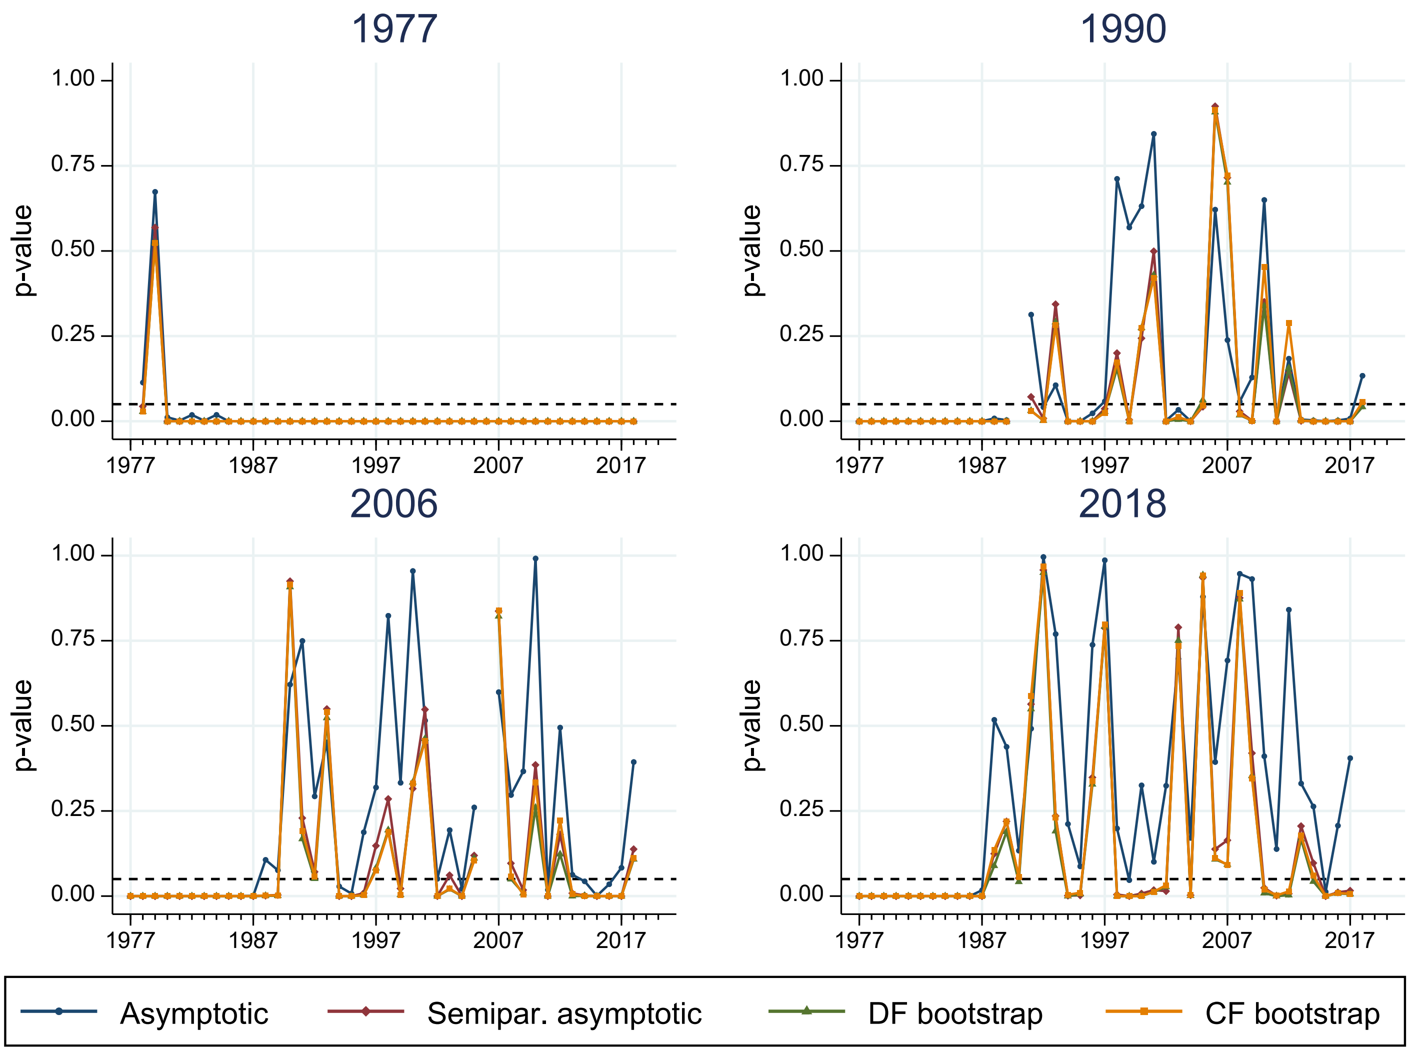

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: **Share top 10%**

Back to main indices

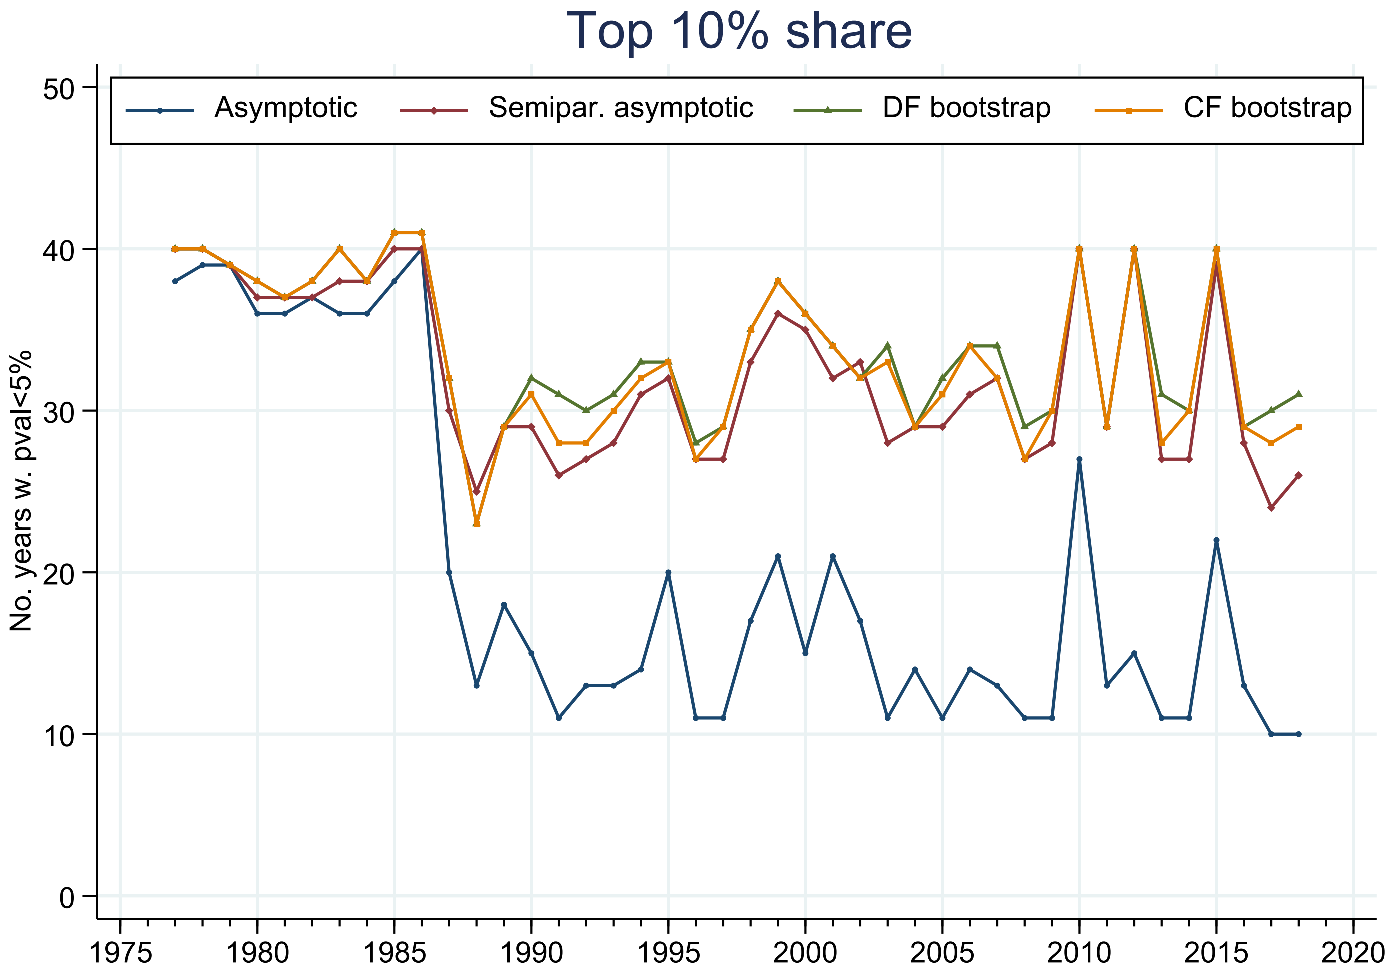

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **share of top 10%**

Back to main indices

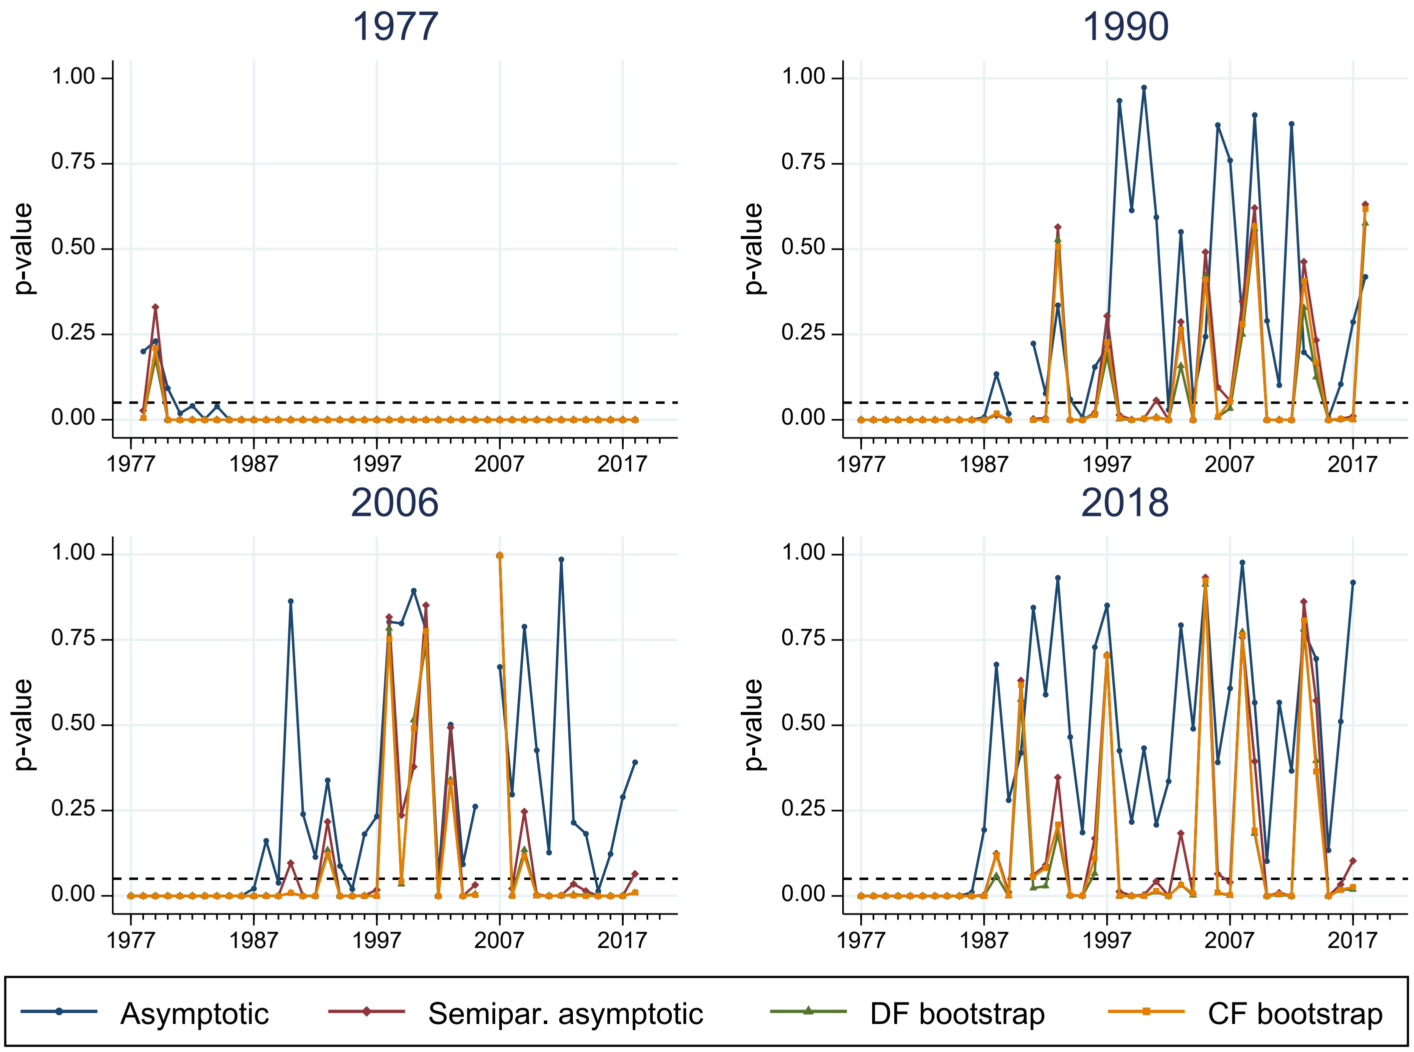

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: **Share top 1%**

Back to main indices

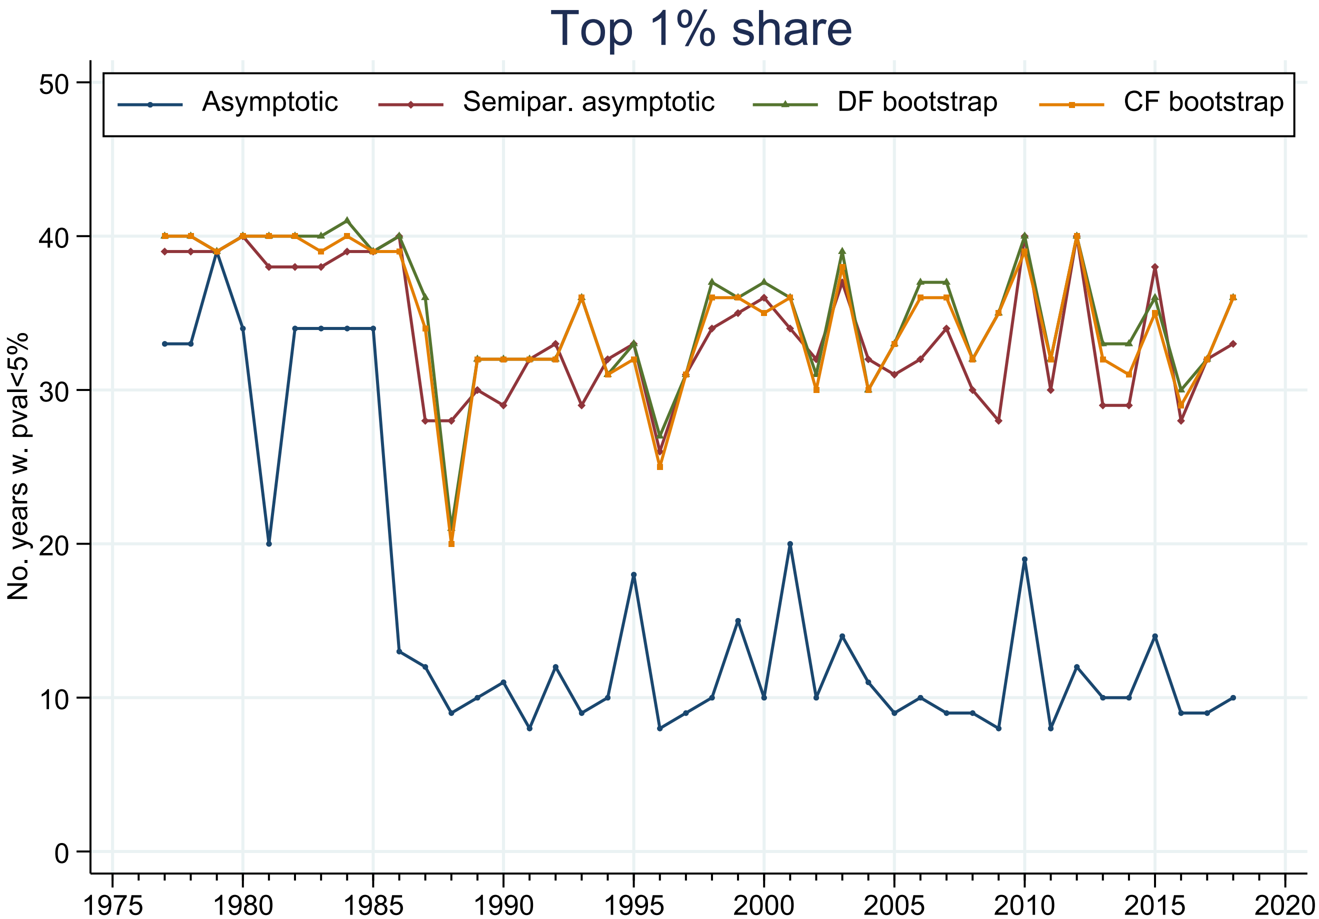

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **share of top 1%**

Back to main indices

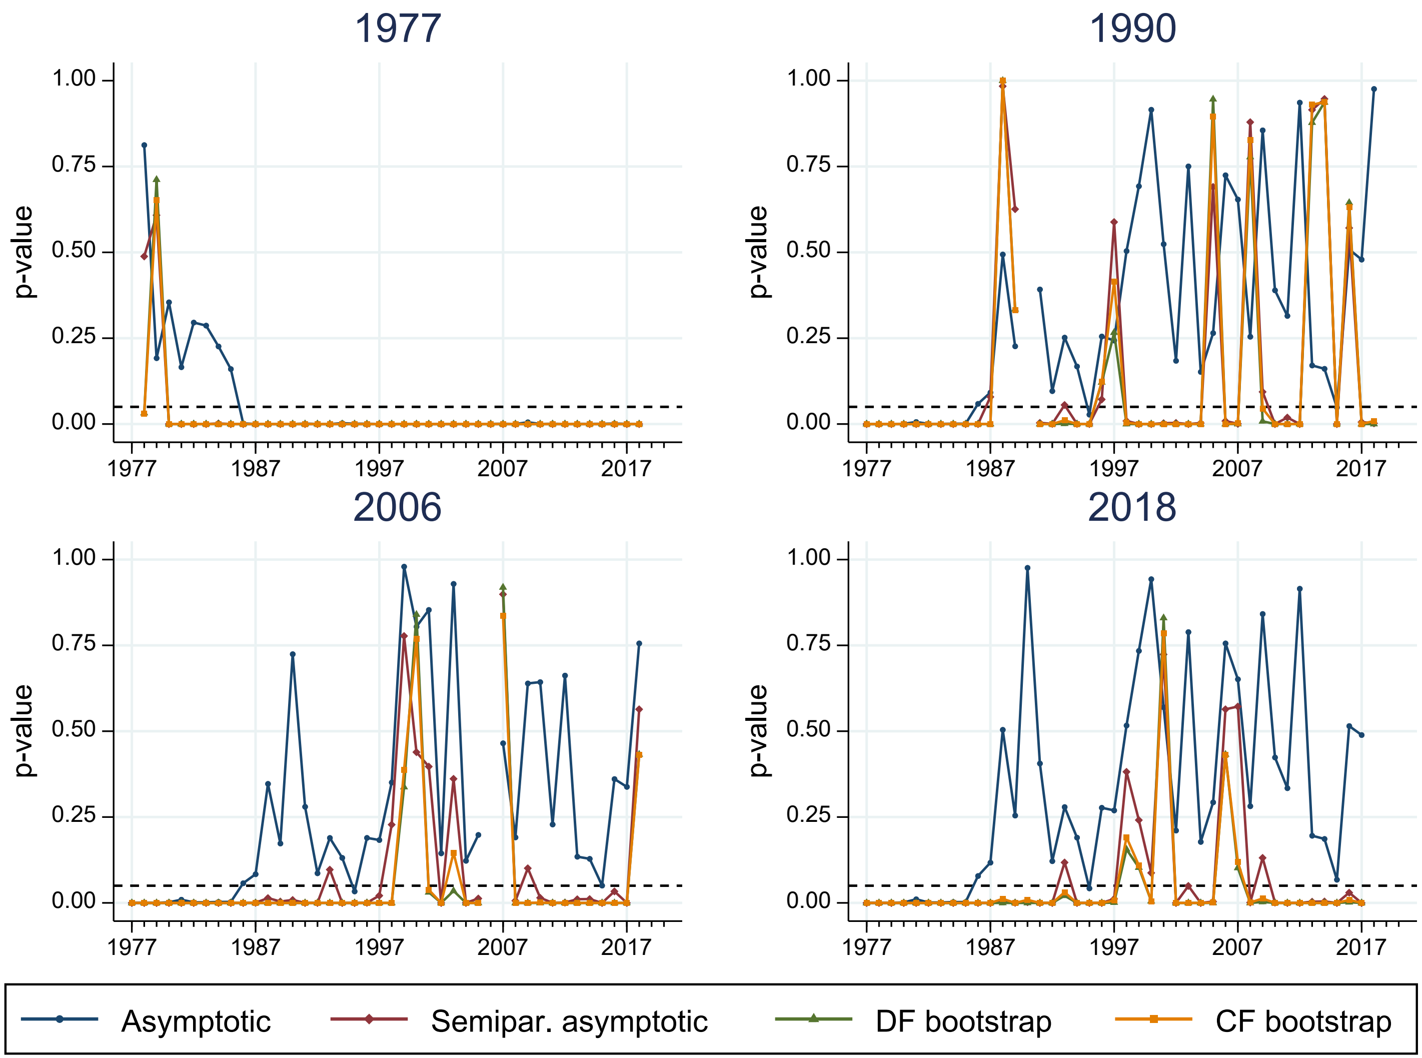

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: ***p*90/*p*10**

Back to main indices

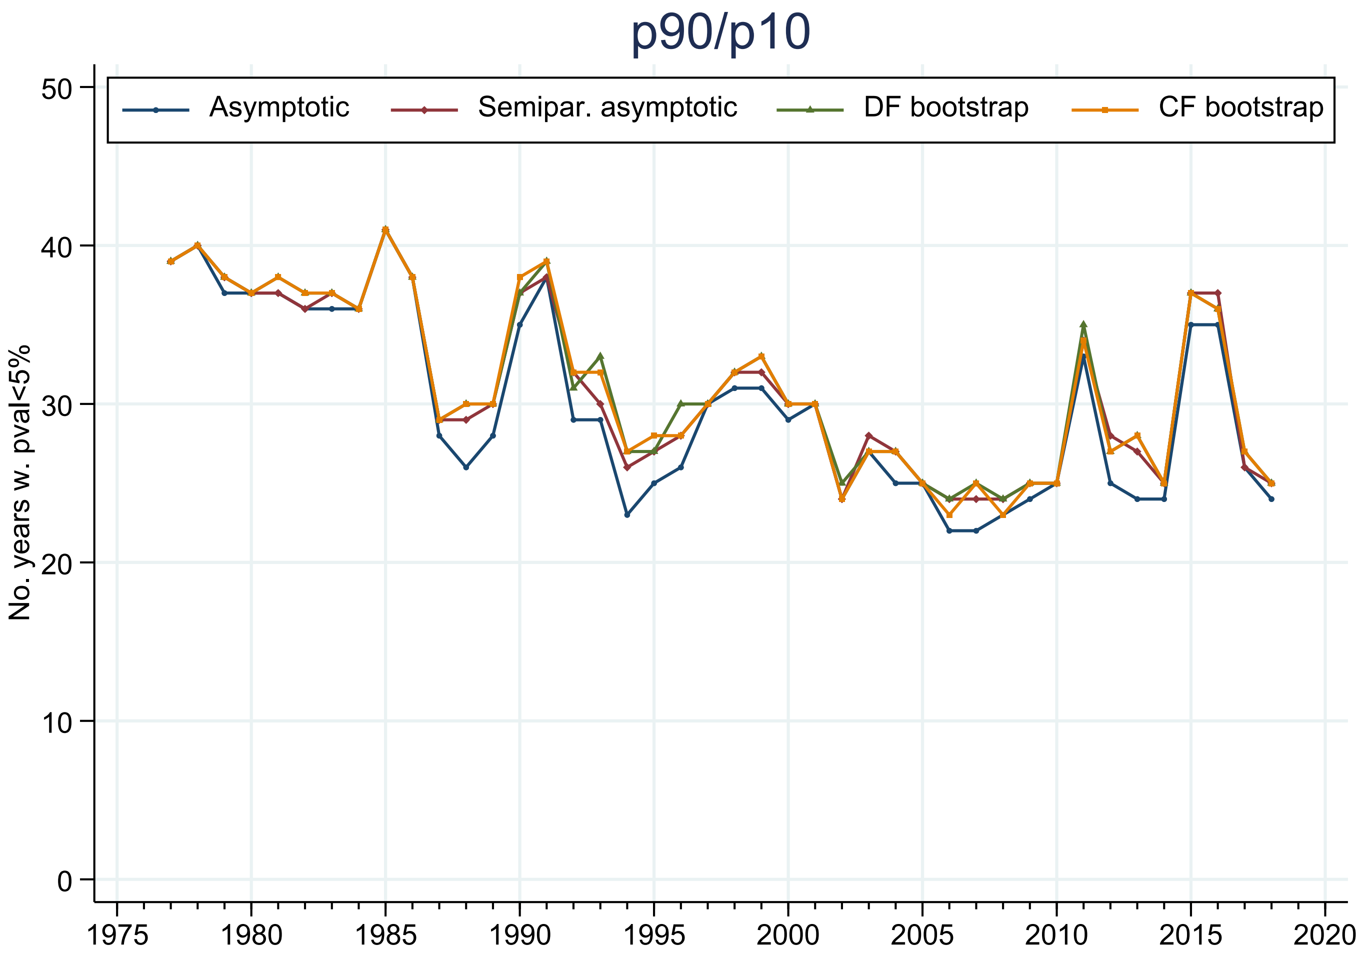

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: ***p*90/*p*10**

Back to main indices

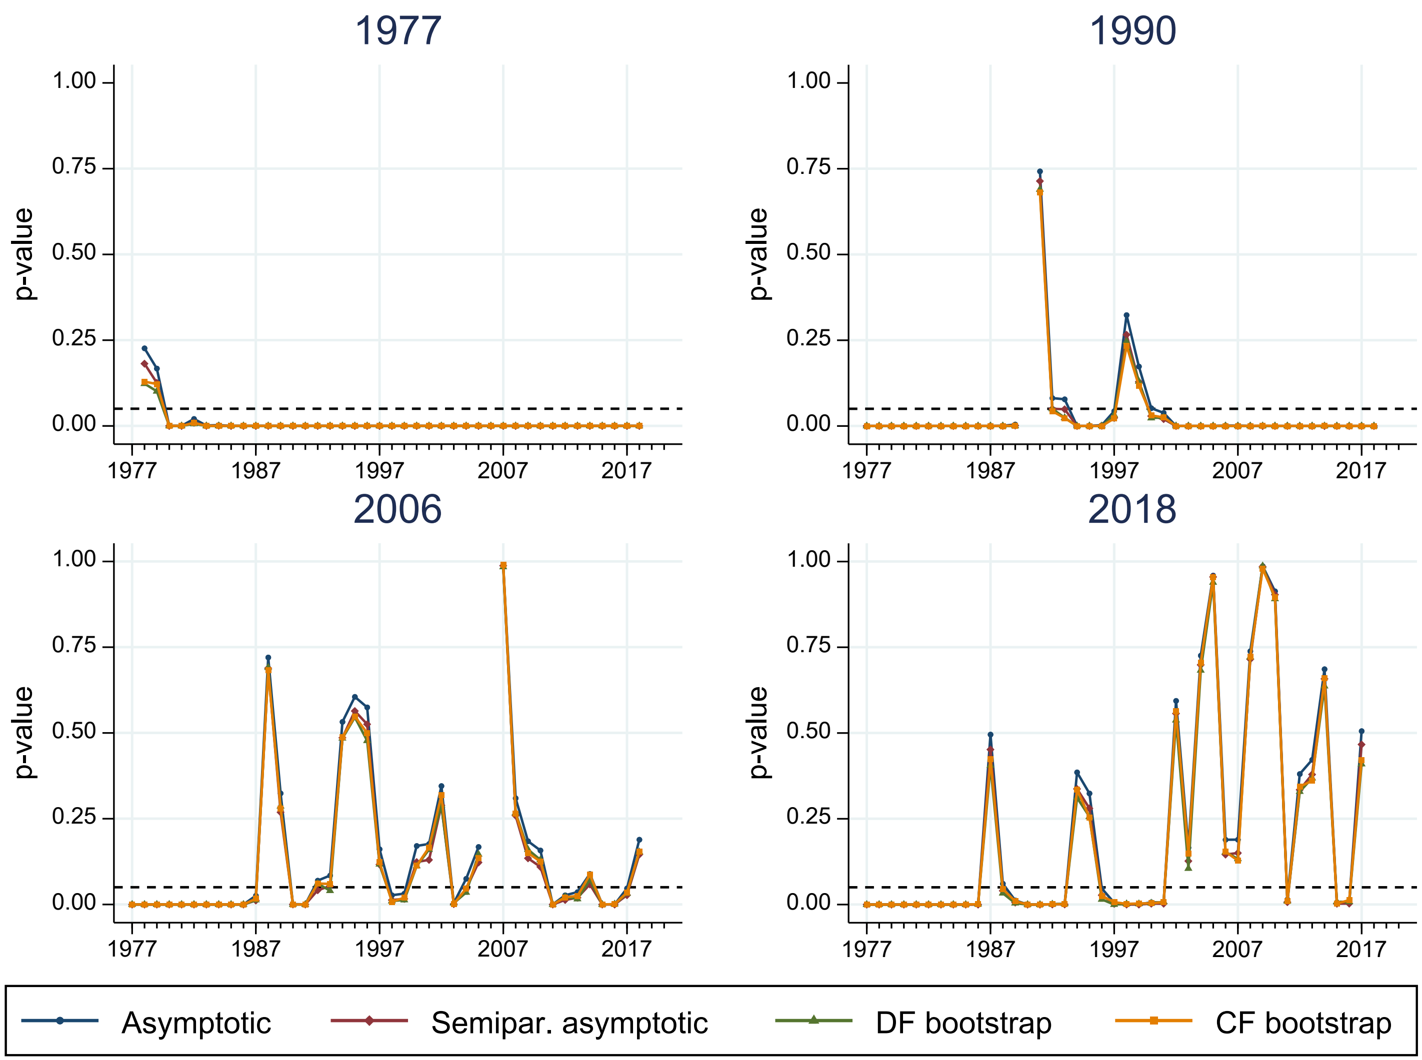

GE(–1), by year and method

Back to estimates #1, #2

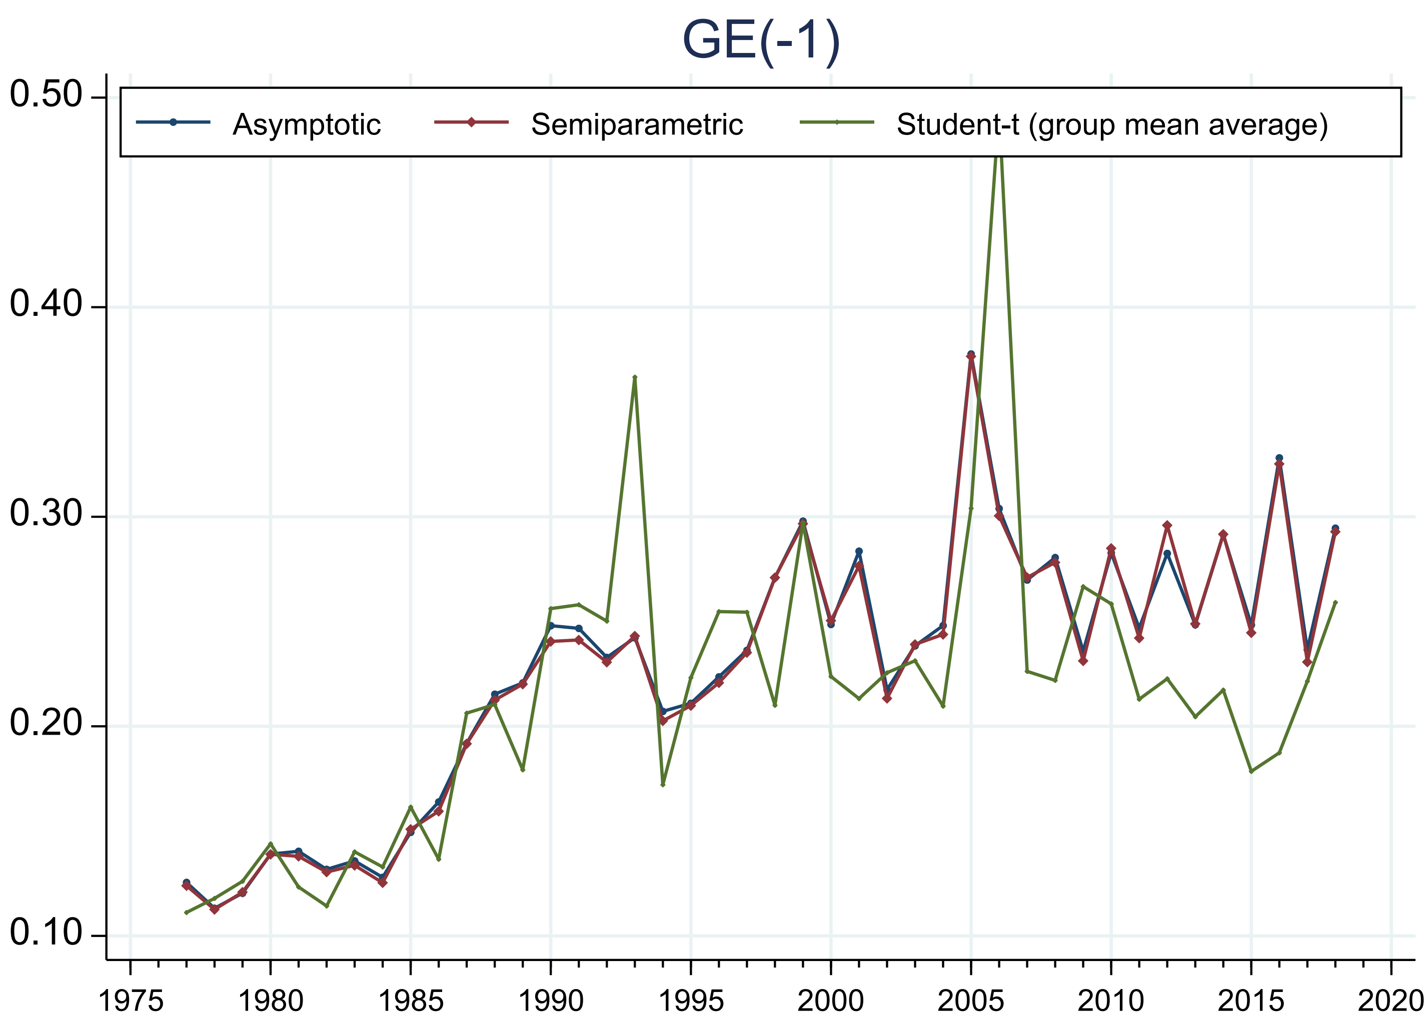

Count of *p*-values < 5% for test of no inequality difference between year *A* and every other year (*B*), by method: **GE(–1)**

Back to main indices

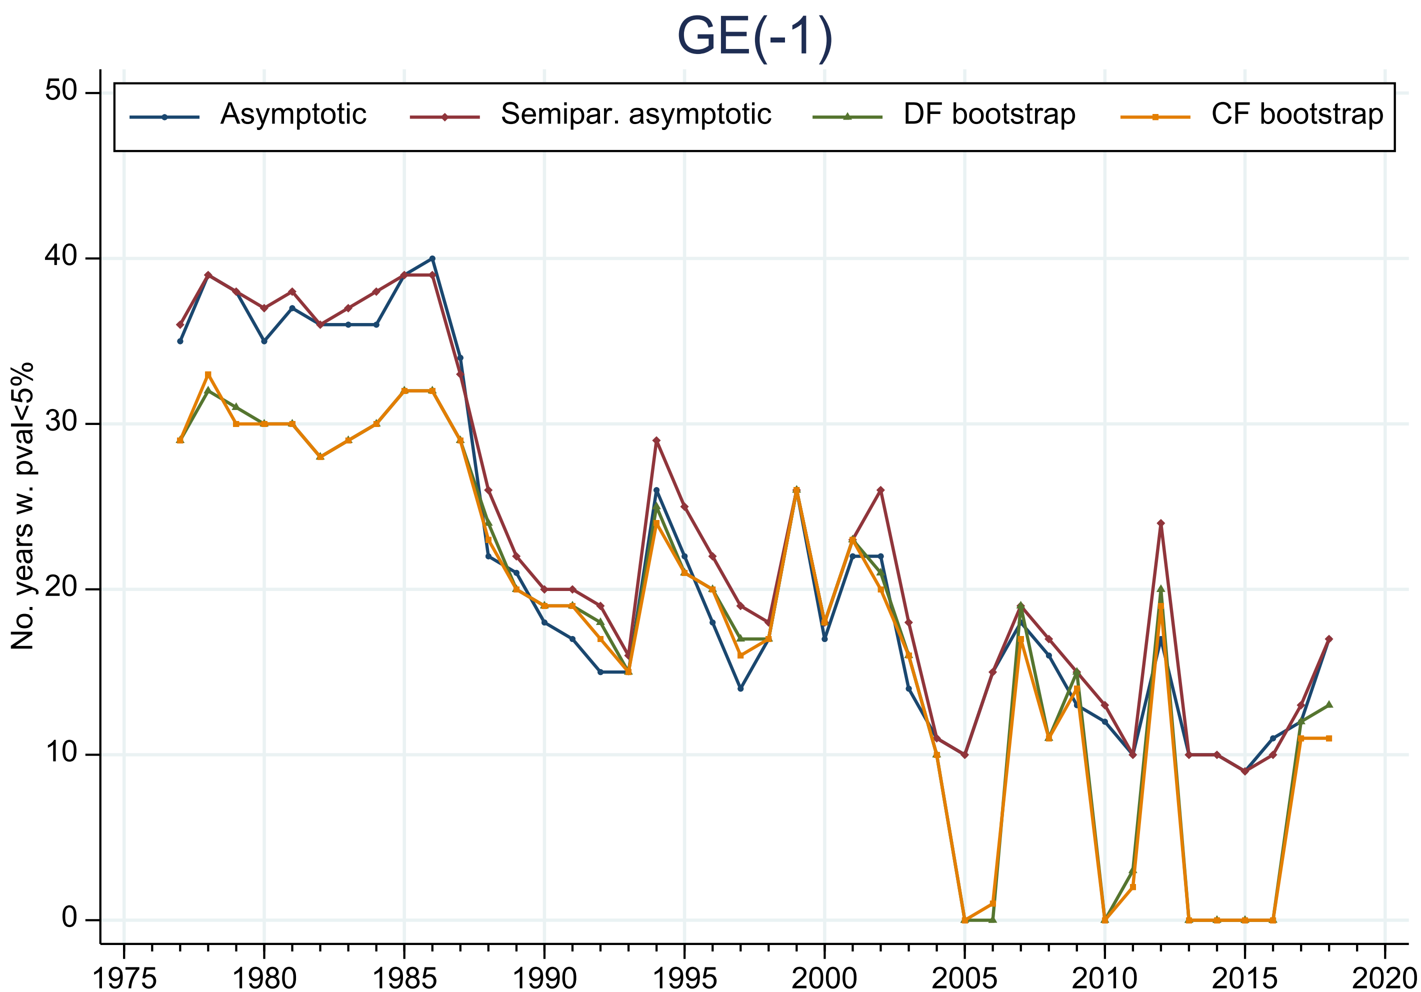

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **GE(–1)**

Back to main indices

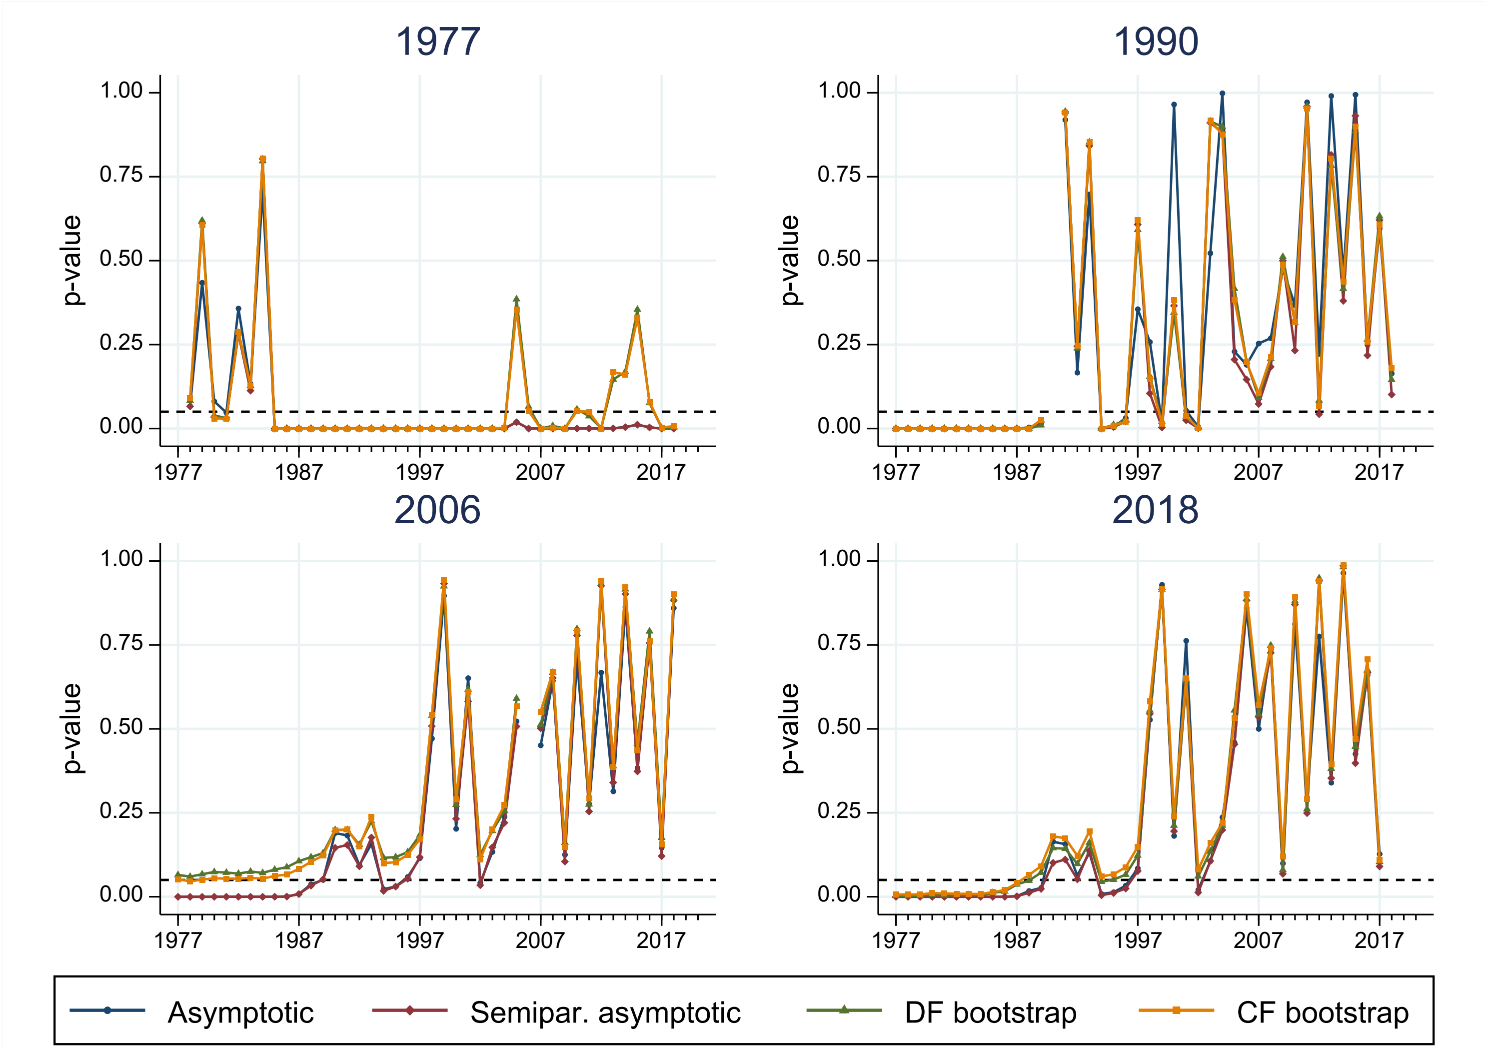

LSE

GE(2), by year and method

Back to estimates #1, #2

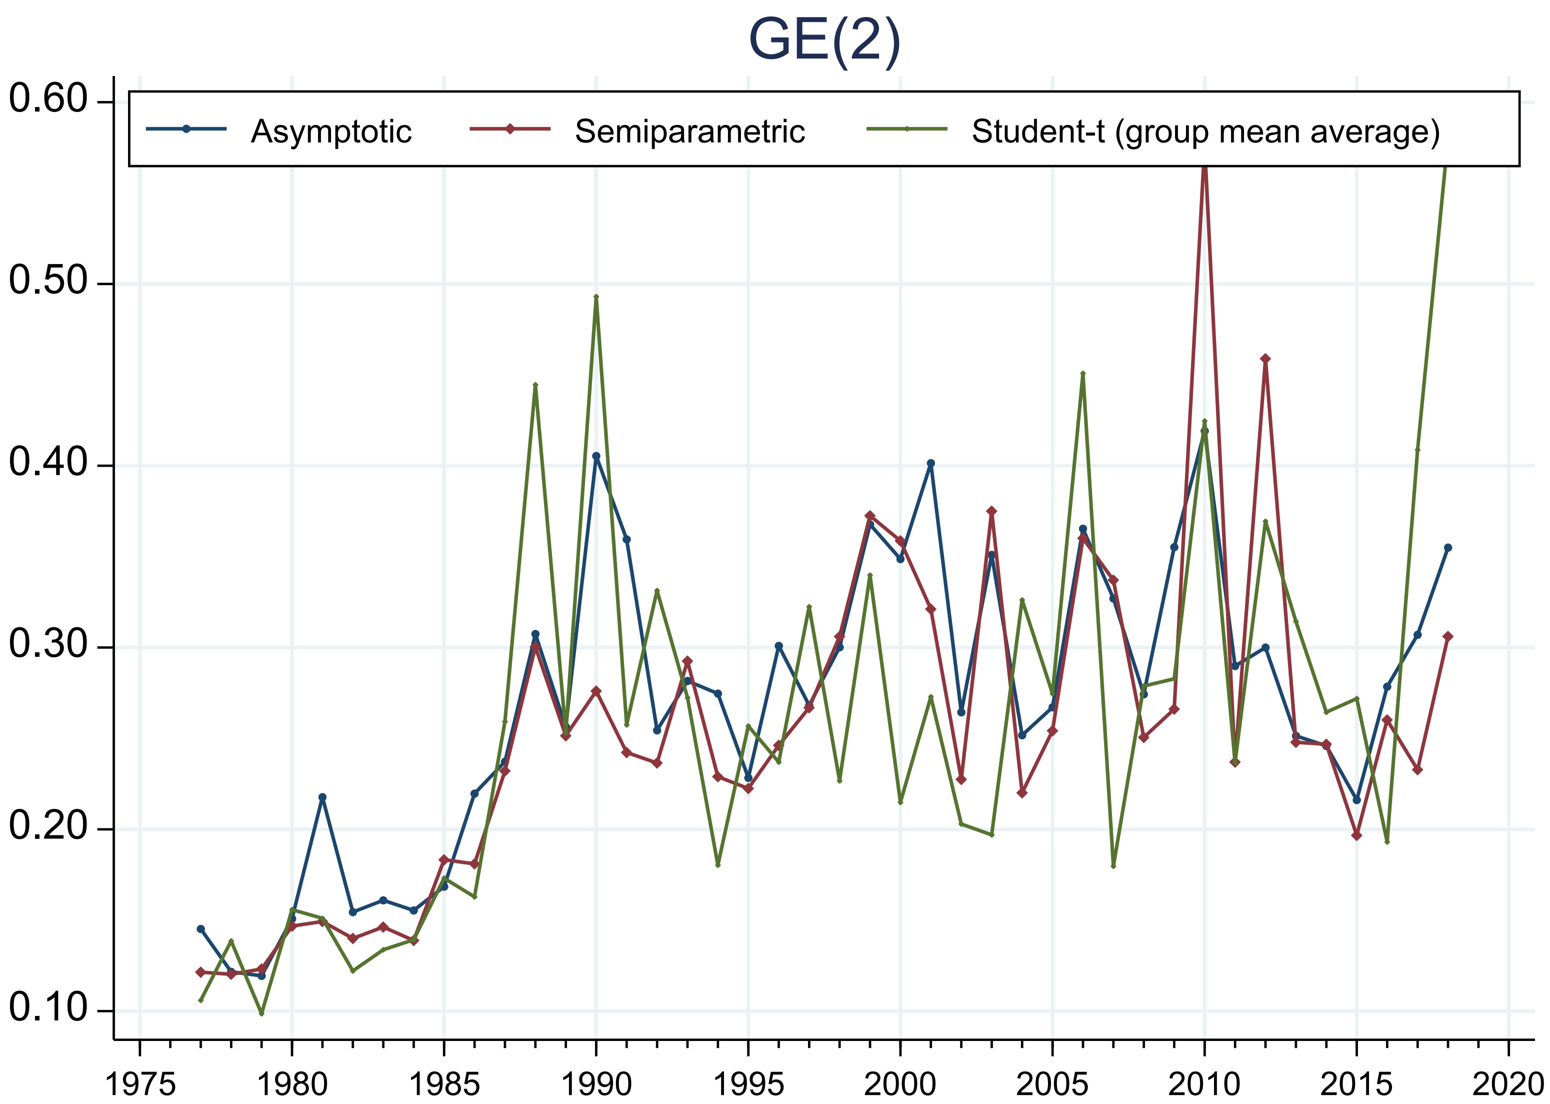

LSE

Count of p-values < 5% for test of no inequality difference between year A and every other year (B), by method: **GE(2)**

Back to main indices

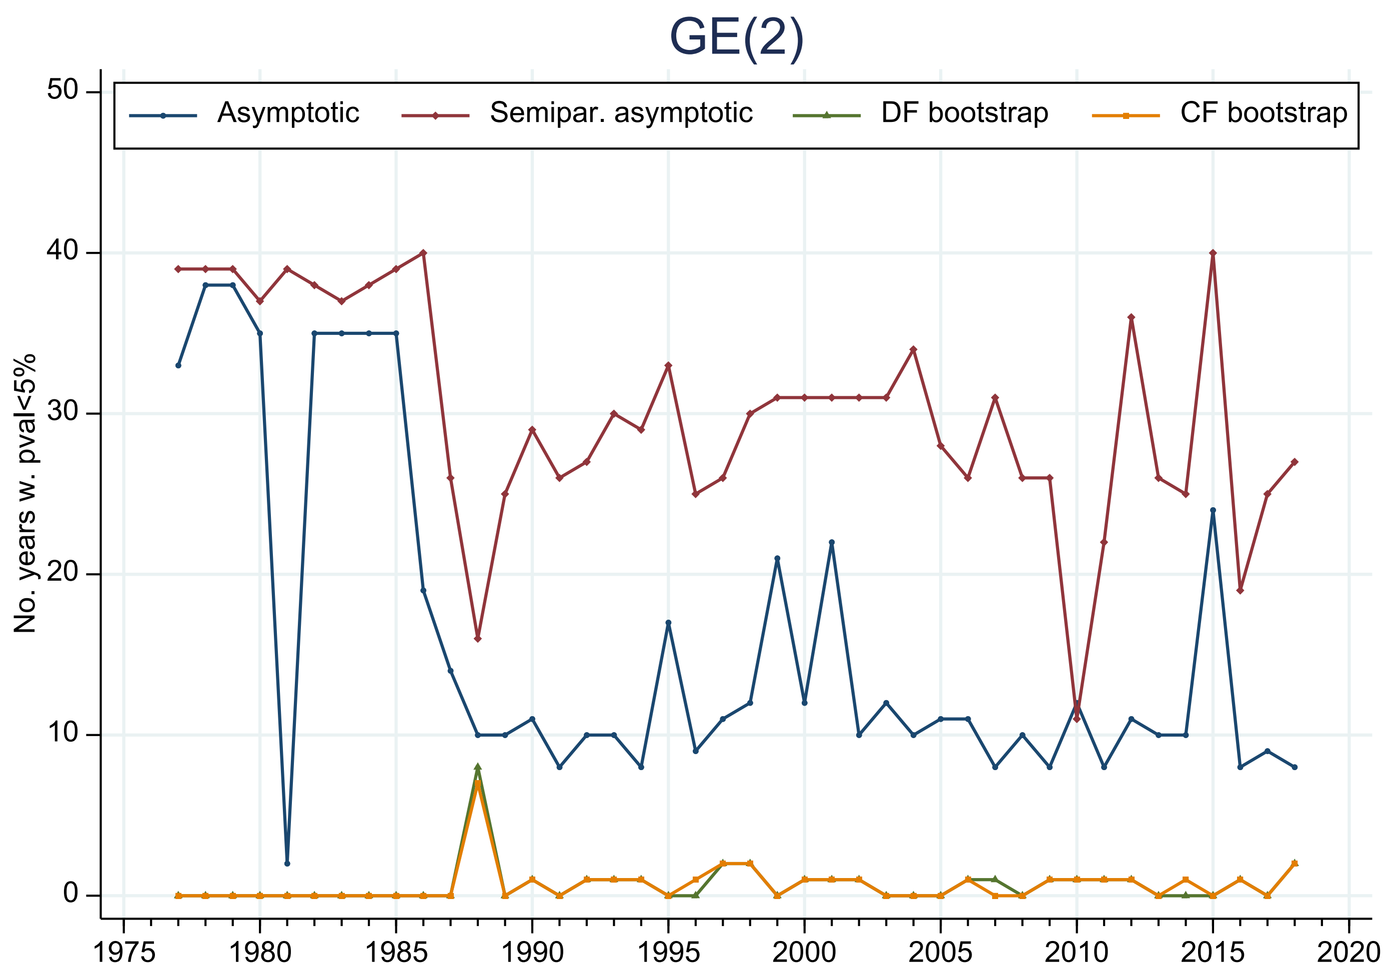

p-values for test of inequality difference between year $A \in \{1977, 1990, 2006, 2018\}$ and every other year (B), by method: **GE(2)**

Back to main indices

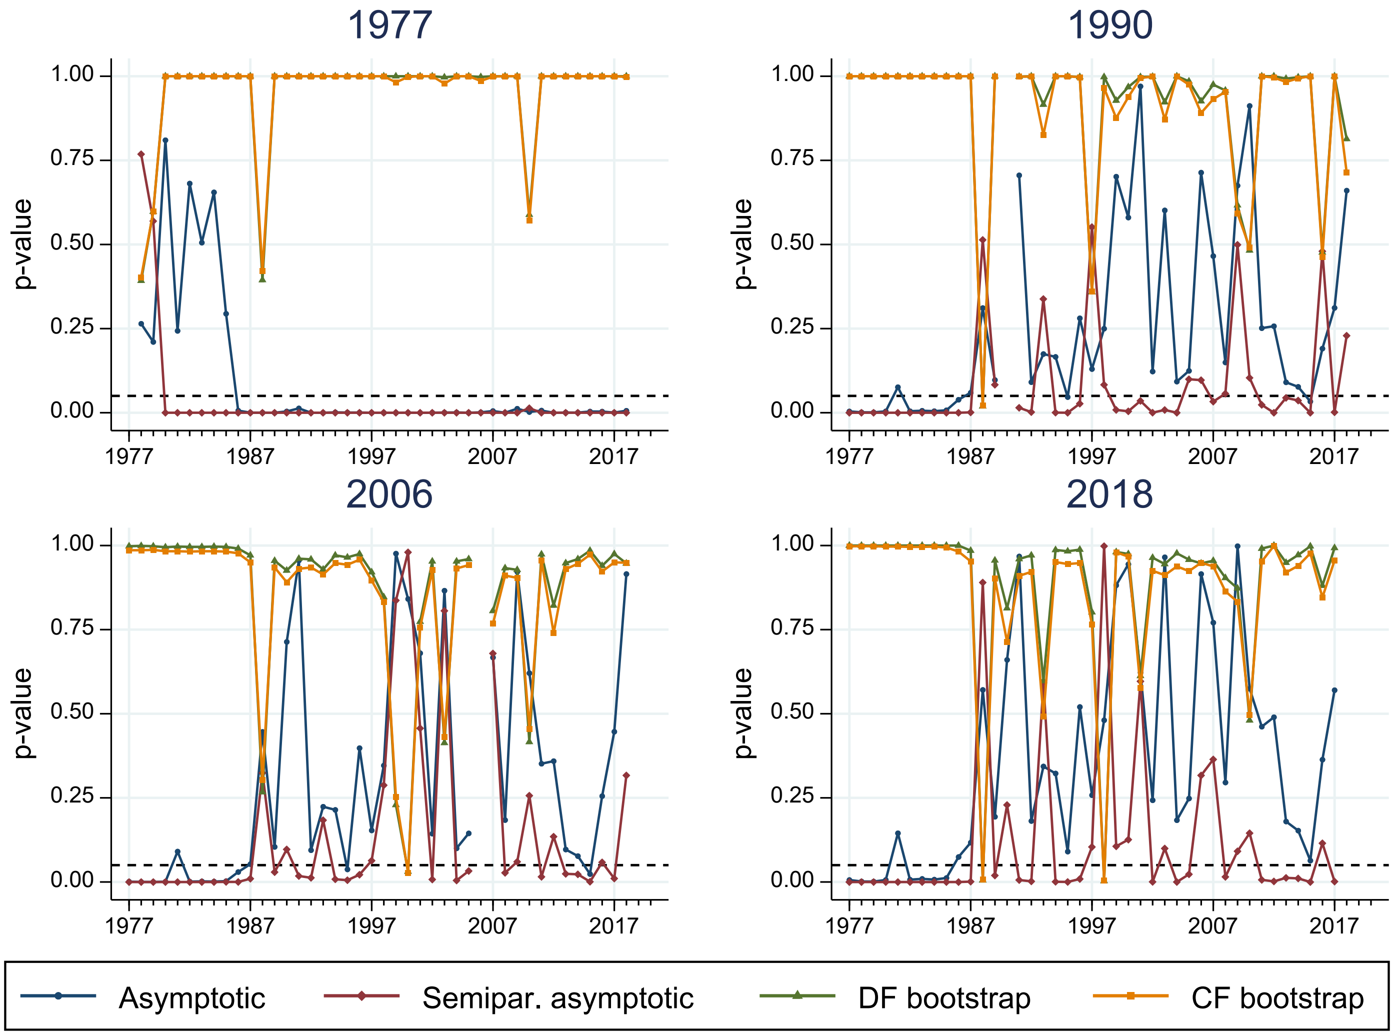

Selected results based on fitting Pareto distributions with ptail = 1%

[Cf. main text: ptail = 5%]

p-values for test of inequality difference between year *A* (1977, 1990, 2006, 2018) and every other year (*B*), by method: **Gini ptail = 1%**

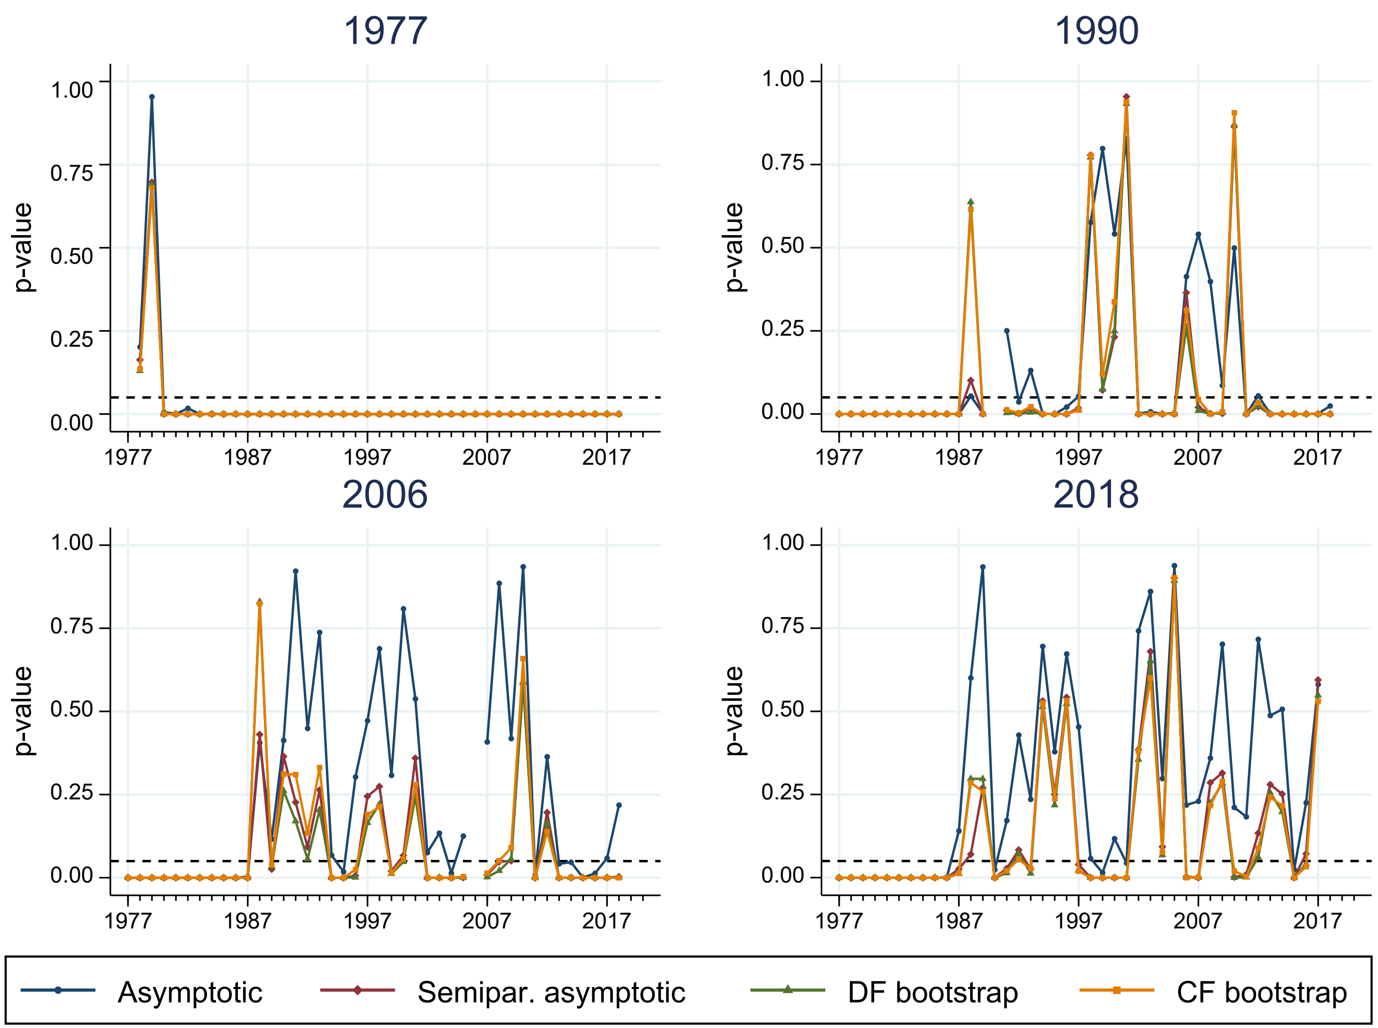

LSE

p-values for test of inequality difference between year *A* (1977, 1990, 2006, 2018) and every other year (*B*), by method: **Theil ptail = 1%**