Protest from Above: online appendix*

Jonathan Schoots† Johan Fourie‡

Contents

S. Displays cited from the main manuscript

A. Non-parametric tests of wealth and household differences

B. Pre-1779 wealth trend test

C. Network analysis: logit specifications and permutation tests

D. Impact of signatories outside the common wealth sample

E. Heemraad reconstruction and linkage under alternative specifications

*This is the supplementary material for Schoots, Jonathan, and Johan Fourie. 2026. "Protest from Above: Elite Decoupling and Public Petitioning in the VOC Cape Colony." Working Paper, Department of Economics, Stellenbosch University. This paper was created with the help of Anthropic's Claude Code and OpenAI's Codex. Cite the paper, not this supplement.

†Department of Sociology, University of the Witwatersrand. Email: jonathan.schoots@wits.ac.za.

‡Department of Economics, Stellenbosch University. Email: johanf@sun.ac.za.

Displays cited from the main manuscript

The main text cites the following displays, which are reported here to stay within the journal's page limit. Numbering follows the citations in the manuscript.

Online Appendix Table S1: Match rate by decile of signature position

Decile (1 = first)	Signatories	Matched	Match rate
D1	41	11	26.8%
D2	41	4	9.8%
D3	41	3	7.3%
D4	41	1	2.4%
D5	40	12	30.0%
D6	40	26	65.0%
D7	40	34	85.0%
D8	40	33	82.5%
D9	40	30	75.0%
D10	40	11	27.5%

Note: Matches of petition signatories to the Stellenbosch–Drakenstein tax censuses, by decile of position on the signature list. Matches in the first decile include several high officeholders at the front of the petition. The low rates in deciles two to four are consistent with recruitment beginning in the Cape District, which is not covered by our linked census data; the decline in the final decile may reflect Swellendam signers or Cape signers added at the end.

Online Appendix Table S2: Officeholding among petition signatories by institutional domain

Office	English title	Domain	Current	Past	Total
Civic administration					
Burgerraad	Burgher councillor	Civic	3	0	3
Heemraad	District official	Civic	4	8	12
Commissaris	Commissioner	Civic	1	1	2
Brandmeester	Fire warden	Civic	1	0	1
Civic subtotal			9	9	18 (38%)
Military					
Capitein	Captain	Military	2	2	4
Luijtenant	Lieutenant	Military	5	1	6
Ritmeester	Cavalry captain	Military	1	0	1
Cornet	Cornet	Military	2	1	3
Vaandrig	Ensign	Military	3	0	3
Military subtotal			13	4	17 (36%)
Church					
Ouderling	Elder	Church	2	2	4
Diacon	Deacon	Church	1	7	8
Church subtotal			3	9	12 (26%)
Total			25	22	47

Note: The table records the office title transcribed for each officeholding signatory. Many officeholders may have held additional offices over the life course. The 47 officeholders represent 11.6% of the 404 petition signatories. The source roster in Baartman (2011) identifies Godlieb Rudolph Opperman as *Oud Cornet*; the supplied spreadsheet had misplaced *Oud* in his name. This table restores that past-office classification.

Online Appendix Table S3: Logistic regression of petition signing on network centrality

	(1) Degree only	(2) Betweenness only	(3) Both
Log degree	4.97***		0.62
	(0.705)		(0.088)
Log betweenness		0.79***	0.71*
		(0.112)	(0.101)

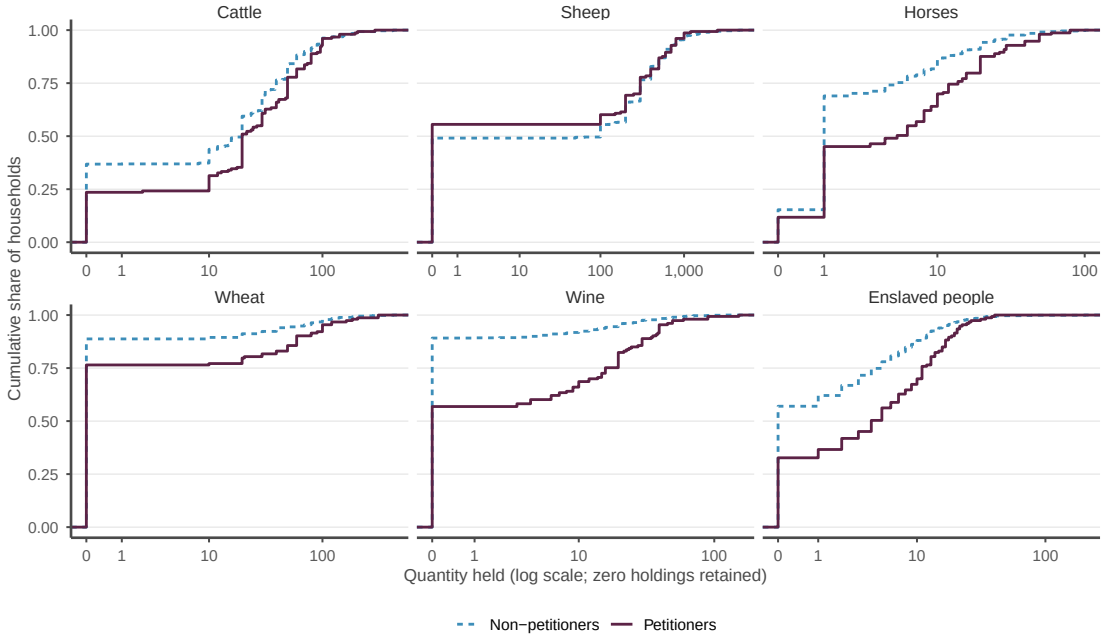
Note: Auction network, largest connected component. Average marginal effects from logit, reported as percentage points per one-log-unit change. Underlying logit coefficients in parentheses. N = 2,225 (petitioners = 172, non-petitioners = 2,053). Variance inflation factors in col. 3 are 3.6 for both regressors. Firth-penalised and residualised specifications are in Section C1 below. * $p < 0.05$, *** $p < 0.001$.

Online Appendix Table S4: Network centrality — full test statistics

	Mean non-pet.	Mean pet.	t	p		% diff.
Degree	2.03	2.74	−3.25	.001	**	34.9
Closeness	0.234	0.244	−6.37	<.001	***	4.3
Betweenness	1,327	2,754	−2.07	.04	*	107.6
Non-parametric tests (Mann–Whitney U)						
	Med. non-pet.	Med. pet.	U	p		Cliff's Δ
Degree	1	2	140,400	<.001	***	0.205
Closeness	0.238	0.244	133,934	<.001	***	0.241
Betweenness	0	171	138,625	<.001	***	0.215

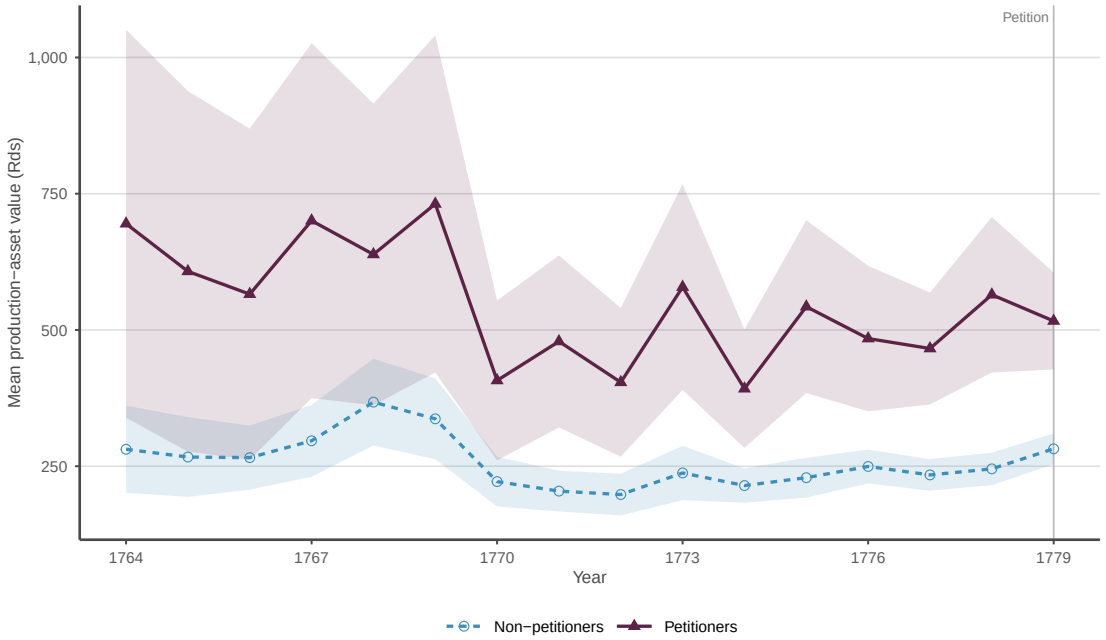
Note: Full statistics underlying Table 5 of the main text. Auction purchase network, 1775–1779, largest connected component. Non-petitioners n = 2,053; petitioners n = 172. U is reported for the non-petitioner group; Cliff's Δ is oriented so that positive values indicate petitioners ranking higher. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Online Appendix Figure S1: Asset distributions



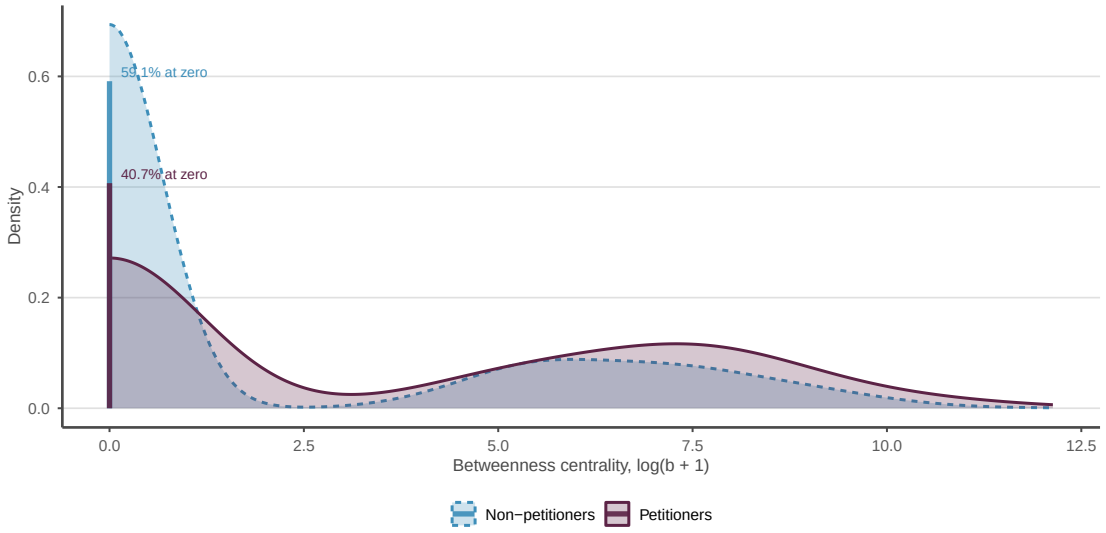
Note: Cumulative distributions of asset quantities for petitioners and non-petitioners. Each curve plots the share of households holding at or below the quantity on the horizontal axis, so a curve farther to the right indicates systematically larger holdings. Households holding none of an asset are retained and appear as the height of the curve at zero. Petitioners tend to rank higher on all series except sheep, but the curves cross; these comparisons do not establish first-order stochastic dominance.

Online Appendix Figure S2: Average household wealth, 1764–1779



Note: Average household production-asset wealth for petitioners and non-petitioners, 1764–1779. Means are calculated over available observations for households in the common 1779 sample; sample composition varies by year. The plot does not exclude economically meaningful differences in pre-1779 trends. Section B reports the uncertainty.

Online Appendix Figure S3: Betweenness centrality by petition status



Note: Density distributions of betweenness centrality by petition status. Petitioners are shifted rightward and are less clustered at zero.

Appendix A: Non-parametric tests of wealth and household differences

Table A1: Kolmogorov-Smirnov tests with Benjamini–Hochberg correction

Asset	D	Raw p	BH-adjusted p
Wine	0.325	1.4×10^{-12}	1.2×10^{-11}
Total index	0.287	$< 10^{-9}$	$< 10^{-8}$
Slaves	0.265	$< 10^{-7}$	$< 10^{-6}$
Horses	0.251	$< 10^{-6}$	$< 10^{-5}$
Settlers	0.174	$< 10^{-3}$	$< 10^{-2}$
Cattle	0.145	.007	.010
Wheat	0.124	.033	.038
Sheep	0.065	.637	.637

Note: R's default `ks.test()` was used; the function emits the warning "p-value will be approximate in the presence of ties" on all of these discrete series. Wheat moves from raw-marginal (0.033) to slightly-more-marginal (0.038) under BH. Only sheep remains non-significant.

Table A2: Mann-Whitney U tests with Benjamini–Hochberg correction

Asset	Pet. median	Non-pet. median	U	BH-adjusted p	r_b
Wine	0	0	99,971	$< .001$	0.33
Total wealth	394	147	99,249	$< .001$	0.32
Slaves	4	0	98,423	$< .001$	0.30
Horses	5	1	94,450	$< .001$	0.25
Settlers	4	2	90,351	$< .001$	0.20
Cattle	20	20	87,218	.002	0.16
Wheat	0	0	84,896	$< .001$	0.13
Sheep	0	100	71,460	.262	-0.05

Note: Mann–Whitney U tests compare petitioners and non-petitioners. Reported p-values are Benjamini–Hochberg-adjusted. r_b is the rank-biserial correlation; positive values indicate that petitioners tend to rank above non-petitioners, while negative values indicate the reverse. All variables except sheep are significant.

Appendix B: Pre-1779 wealth trend test

Table 1 reports the within-household pre-trend model discussed in §6.3 of the main text. The sample is restricted to pre-petition observations (1764–1778) for households in the common 1779 wealth sample. We demean both wealth and year within household, so the interaction term captures whether future petitioners were on a different pre-1779 trajectory from other observed households. Standard errors are clustered at the household level.

Table 1: Within-household pre-1779 wealth trend test.

Term	Estimate	SE	p
Year (demeaned within household)	5.18	1.46	$<.001$***
Future petitioner × year	−1.59	7.27	0.827

Note: Outcome is total production-asset value in rixdollars. Household fixed effects are implemented through within transformation. The positive year term reflects overall growth in observed asset value in the pre-petition panel; the interaction is imprecisely estimated. Its approximate normal 95% confidence interval is $[-15.83, 12.65]$ Rds/year, allowing relative decline as well as improvement. Two supplementary growth comparisons also fail to reject equality. Defining growth as the change in a household's production-asset value between its first and last observed pre-1779 year, a t-test on mean growth yields $p = .43$ and a Mann–Whitney test on the same quantity yields $p = .52$. A normal-approximation calculation at 80 per cent power and 5 per cent two-sided size gives a minimum detectable slope of $(1.96 + 0.842) \times 7.27 = 20.4$ Rds/year. Thus non-rejection is not evidence that economically meaningful trends are absent. The year-by-year counterpart is plotted in Figure 1 and tabulated in Table 2.

Table 2 and Figure 1 report the year-specific counterpart: household wealth is regressed on future-petitioner status interacted with year indicators, with household and year fixed effects and 1778 as the reference year. The CR2 standard errors from `estimatr::lm_robust` differ from the `sandwich::vcovCL` estimator used for the linear trend, although both cluster by household. Five of the fourteen interactions reach individual significance at 5 per cent. Their negative signs indicate a smaller petitioner wealth premium in those years than in 1778, not year-on-year declines or gains. None survives a Benjamini–Hochberg correction across the fourteen contrasts at 5 per cent. The fluctuating estimates do not establish a sustained differential trajectory; the linear trend estimate likewise remains inconclusive.

Table 2: Event-study specification of the pre-1779 panel.

Future petitioner × year	Estimate	SE	95% CI	p
1764	34.8	127.7	$(-223.6, 293.1)$	.787
1765	−35.8	112.1	$(-261.6, 189.9)$	.751
1766	−29.6	118.1	$(-265.9, 206.6)$	.803
1767	7.9	95.4	$(-183.6, 199.4)$	.934
1768	−97.4	77.2	$(-251.7, 56.8)$	.212
1769	28.1	81.5	$(-134.6, 190.9)$	.731
1770	−173.6	81.4	$(-336.1, -11.2)$	.037
1771	−133.3	59.8	$(-252.2, -14.4)$	.028
1772	−210.6	72.8	$(-355.2, -66.0)$	.005
1773	−45.4	74.7	$(-193.8, 103.0)$	.545
1774	−163.7	63.4	$(-289.4, -37.9)$	.011
1775	−31.1	52.3	$(-134.8, 72.6)$	.553
1776	−78.4	45.1	$(-167.7, 10.8)$	.084
1777	−90.8	45.3	$(-180.5, -1.0)$	.047

Note: 5,558 household-year observations on 786 households (106 future petitioners, 680 non-petitioners), 1764–1778. Household fixed effects implemented through within transformation and year fixed effects included via dummies; SEs cluster-robust (CR2) at household level (`estimatr::lm_robust`).

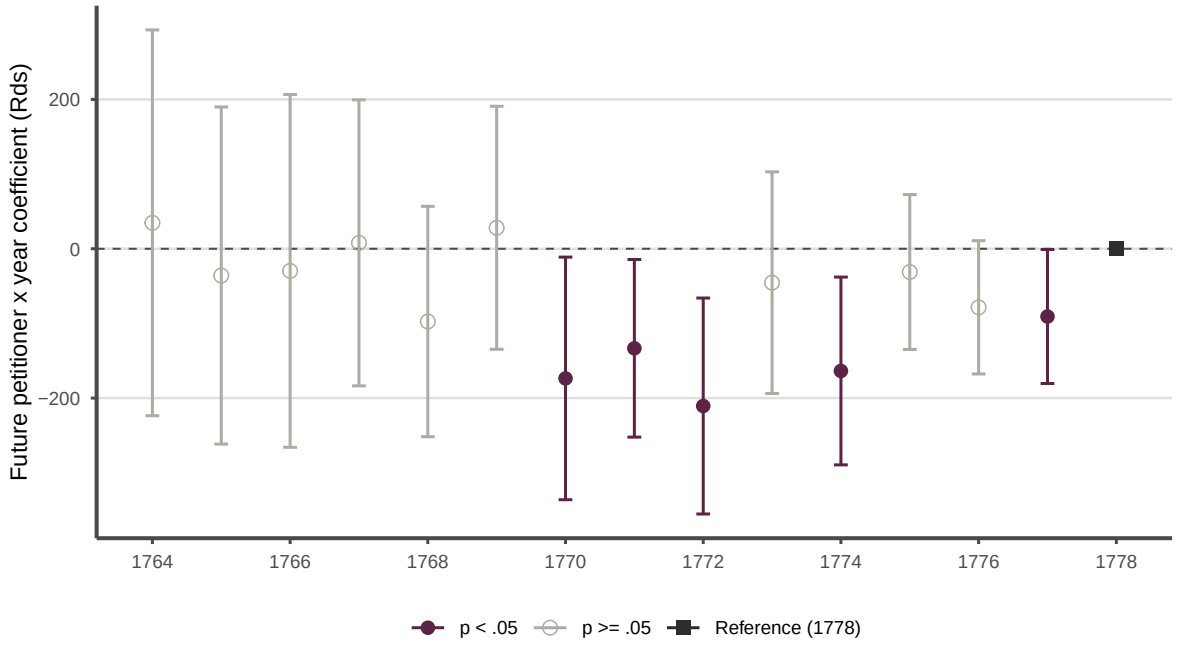


Figure 1: Event-study coefficients from the within-household pre-1779 wealth regression, 1764–1778. Points are $\hat{\beta}_t$ on Future petitioner $\times$ Year t, with 1778 as the reference year and household fixed effects absorbed through within transformation. Whiskers are 95% confidence intervals based on CR2 household-clustered standard errors. Plum points mark years at which the coefficient is individually significant at $p < .05$.

Appendix C: Network analysis

C1: Network logit

Table 3: Network logit — MLE, Firth penalisation, and residualised betweenness

	MLE				*Firth penalisation*	
	Coef.	SE	p	AME (pp)	AME (pp)	VIF
Log betweenness	0.101	0.041	.014	0.71	0.72	3.60
Log degree	0.088	0.299	.768	0.62	0.68	3.60
Residualised specification						
Betweenness $\perp$ degree (residual)	0.101	0.041	.014	0.71	—	—
Log degree (in residualised spec.)	0.790	0.161	$< .001$	5.56	—	—

Note: Sample: 172 signatories of 2,225 nodes (7.7%). Firth-penalised likelihood (Firth, 1993) yields AMEs within 0.07 pp of the MLE. The residualised specification regresses log betweenness on log degree by OLS and enters the residual alongside log degree in a logit. This is an invertible reparameterisation of the same joint model: fitted probabilities and likelihood are unchanged. Its degree coefficient combines the direct degree term with the component of betweenness predictable from degree. Firth penalisation addresses likelihood bias, not dependence among connected observations. The reported conventional standard errors do not model network dependence.

C2: Network permutation tests

Table 4 reports two randomisation comparisons. Both compute betweenness on the binary person projection: two people are connected if they purchased at the same auction. This differs from the bipartite centrality used in the descriptive table and logistic regressions. Panel A reallocates 172 petition labels uniformly across 2,225 people, keeping the network fixed (B = 2,000). Panel B fixes labels and preserves every person's and auction's degree through double-edge swaps on the simple bipartite graph ($B = 500$). Each draw starts from the observed graph and makes five times the number of edges in attempted swaps, rejecting self-pairs and proposals that duplicate an edge. This is an exact-degree edge-swap procedure, not the Stochastic Degree Sequence Model, which is a different null specification. Finite swapping does not by itself establish independent uniform draws from all admissible graphs.

The observed mean log-betweenness gap is 1.443. Uniform label reallocations yield a mean gap close to zero. Under the fixed-degree rewiring procedure, the mean is 1.370: about 95 per cent of the observed gap is reproduced by this benchmark. The remaining 0.073 lies in its upper tail ($p = .006$). The corresponding observed top-quintile share gap is 0.155, against a rewired mean of 0.123 ($p = .004$). These comparisons indicate a smaller residual brokerage advantage conditional on the chosen degree sequence and sampler. The ratio of null mean to observed gap is a descriptive comparison, not the share of an effect causally explained by degree. Monte Carlo resolution is approximately 0.002 for Panel B, and the results remain conditional on adequate exploration of the

admissible networks.

A diagnostic using 25 draws at each of five, twenty and fifty times the number of edges in attempted swaps yields mean log-betweenness gaps of 1.382, 1.367 and 1.368, respectively. Proposal acceptance is approximately 85 per cent at each length. The longer runs give similar means to the 500-draw benchmark, but this small diagnostic neither establishes uniform sampling nor precisely checks its tail probabilities.

Table 4: Permutation tests on the realised auction network.

Metric	Observed gap	Perm. mean	Perm. SD	p
Panel A. Uniform-label permutation ($B = 2{,}000$)				
Mean log betweenness gap	1.443	0.002	0.283	≤ 0.0005***
Top-broker share gap	0.155	−0.0003	0.032	≤ 0.0005***
Panel B. Fixed-degree edge-swap comparison ($B = 500$)				
Mean log betweenness gap	1.443	1.370	0.032	0.006**
Top-broker share gap	0.155	0.123	0.012	0.004**

Note: "Observed gap" is the difference in the statistic between petitioners and non-petitioners on the realised network, computed by a single procedure shared across both panels (betweenness from the bipartite one-mode projection, log-transformed as $\log(b+1)$; top-broker share equal to membership in the top quintile). Panel A permutes petition labels uniformly at random while holding the edge set fixed; Panel B holds labels fixed and swaps edges while preserving the bipartite degree sequence. Each reported p equals $(1+\#\{|T_b| \geq |T_{obs}|\})/(B+1)$. For Panel B the positive null gaps make this an upper-tail comparison, not a two-sided test centred on the nonzero null mean. With only 500 draws, tail probabilities have appreciable Monte Carlo uncertainty.

Appendix D: Signatories outside the common wealth sample

Of 404 signatories, 153 enter the positive-asset wealth sample, 12 match households with recorded zero values on the five priced series, and 239 are unmatched. We retain the twelve recorded zeros and vary the assets of the 239 unmatched signers. These are imputation scenarios, not formal bounds. They compare a hypothetical mean for all petitioners with the observed mean among positive-asset rural non-petitioners; they do not supply the missing urban comparison group.

Table D1: Imputation scenarios for 239 unmatched signatories

Unmatched signers' assumed assets	Pet. mean	Rural non-pet. mean	Gap
Same as observed petitioners	501	282	+219
Same as non-petitioners	362	282	+80
All hold nothing	196	282	−86

Note: Values in Rds. Petitioner mean is $(153 \times 516.452 + 12 \times 0 + 239 \times \mu)/404$, where μ is the assumed mean for unmatched signers. The comparison mean is 281.940 Rds for the 986 positive-asset rural non-petitioners in main Tables 1 and 2. No inferential test accompanies the imputed means.

Table D2: Break-even grid for the assumed Cape District component

Cape share	Assumed Cape n	Break-even mean (Rds)	% of rural non-pet. mean
0.1	23.9	−1078	−382%
0.2	47.8	−398	−141%
0.3	71.7	−171	−61%
0.4	95.6	−58	−21%
0.5	119.5	10	4%
0.6	143.4	55	20%
0.7	167.3	88	31%
0.8	191.2	112	40%
0.9	215.1	131	46%
1.0	239.0	146	52%

Note: The Cape share applies only to the 239 unmatched signers. The remaining unmatched signers are assigned the positive-asset rural non-petitioner mean; twelve matched zeros remain zero. The break-even Cape mean equates the combined mean for 404 petitioners with that rural benchmark. Negative values cannot be attained with nonnegative assets. Fractional counts describe a scenario, not identified residence.

If all 239 unmatched signers formed the Cape component, their break-even mean would be approximately 146 Rds. When half are assigned to it, the break-even mean is approximately 10 Rds. The signatory sequence does not identify residence well enough to assign a probability to either scenario. Moreover, little recorded agricultural wealth need not imply low urban wealth or status.

A separate check addresses the positive-asset filter without imputing any unmatched signer's assets. Including all 1,375 observed 1779 rural households gives mean production-asset values of 479 Rds for 165 petitioners and 230 Rds for 1,210 non-petitioners. The difference is 249 Rds (Welch 95% interval $[161, 337]$). Thus the rural mean difference survives the inclusion of observed zero-asset households, while the wealth composition of all 404 signers remains unidentified.

Appendix E: Heemraad reconstruction and linkage under alternative specifications

The two heemraad tables in the main text use different designs. Table 3 starts from officeholders and asks what share signed. Comparison with the district-labelled source workbook identifies the first four entries in each annual cohort as Stellenbosch–Drakenstein appointments and the final two as Swellendam appointments. Table 3 retains the first four. Late-year appointments begin service the following January and last two years, so the 1777–1778 appointment cohorts give the eight men serving in 1779. The inherited hand crosswalk consolidates spelling variants while preserving distinguishing suffixes. It links 4 of 8 current officeholders, 7 of 17 men whose terms began during 1775–1779, and 11 of 39 whose terms began during 1763–1779 to titled petitioners. The longer rows exclude men whose terms began earlier but overlapped the first year. A proposed match between the twelfth titled signer, François du Toit de Oude, and a 1722 appointee remains biographically unverified and is not used in these counts. The longer source roll omits appointment years 1746–1749, limiting the prior-service indicator; this gap precedes all three cohorts in main Table 3. The replication files record 68 appointment slots and the identity crosswalk; reproducing these joins does not independently verify every historical identity.

Table 4 instead starts from matched 1779 households and asks whether petitioners were more likely than non-petitioners to have prior heemraad service. Its indicator is necessarily asymmetric. For petitioners we use the Rank field of the 1779 signatory roster, which records each signatory's civic office at the time of signing; this is office attribution transcribed in the supplied roster. For non-petitioners no such field exists, so service must be inferred by matching normalised names against the heemraden roll. Table 5 reports the differential under the specifications that this asymmetry makes available.

Table 5: Prior heemraad service: petitioner–non-petitioner differential under alternative linkage rules.

Specification	N	Pet. %	Non-pet. %	Diff. (pp)	p
Roster title, continuity match (Table 4)	1,135	7.2	2.4	+4.7	.004
Continuity match on both sides	1,135	6.5	2.4	+4.1	.017
Random draw within collision groups	1,139	5.3	2.8	+2.4	—
Name match, collisions ignored	1,139	14.4	5.2	+9.2	<.001
Collision cases dropped	1,081	0.8	1.5	−0.7	1.00
One-office-per-name stress test					
every collision resolved against	1,139	2.6	3.2	−0.6	1.00
every collision resolved in favour	1,139	8.5	2.3	+6.2	<.001

Note: Common 1779 wealth sample. Rows differ only in how a normalised heemraad name matching several 1779 household heads is resolved. "Continuity" assigns the office to the candidate whose panel tenure best fits the recorded service years; "collisions ignored" flags every candidate, which double-counts populous lineages such as de Villiers and du Toit; "collisions dropped" removes them from both groups. Fisher exact two-sided p; the random-draw row reports means over Monte Carlo replications and has no single p (its difference column is the mean of per-replication differences, which need not equal the difference of the two rounded shares). The final two rows are a stress test, not bounds on the reported estimate: they assume exactly one real officeholder per normalised name, an assumption the roster itself contradicts – two distinct Jan de Villiers signatories are recorded as heemraad and oud heemraad – and they discard the petitioner-side roster titles that the reported specification uses. A lineage-level reading yields +22.6pp but changes the unit of analysis from household to lineage and is not comparable with these rows.

Three points follow. First, the specification reported in the main text is not the most favourable available: resolving every ambiguous case in favour gives +6.2pp, and the naive match that ignores namesake collisions gives +9.2pp. Second, the differential is not robust to every rule. Dropping the ambiguous cases, or resolving them all against, produces small negative differences indistinguishable from zero. Both discard or reassign much of the title-identified signal: the conservative rule removes 58 households and, between exclusion and reclassification, retains only one of the eleven petitioners the roster identifies as heemraden. They are not thereby illegitimate, and we report them for that reason.

Third, the continuity rule is weaker than a summary description suggests. Of the twenty-one collision groups, eight are settled by an actual overlap between the recorded service years and a candidate's panel tenure. Eleven have no overlap at all and fall back on assigning the office to the longest-tenured candidate, a heuristic whose error direction we cannot sign. Two are unresolvable and are dropped. We prefer this rule because it uses information the alternatives throw away, but the officeholding differential rests on a linkage judgement in a way the wealth and network results do not.

The petitioner side allows one direct check. Against the twelve signatories the roster identifies as current or former heemraden, the simple name match flags twenty-one, of which seven are correct:

precision 0.33, recall 0.58. It is tempting to read the over-inclusion as conservative, on the grounds that false positives inflate the non-petitioner rate and so shrink the gap. We do not make that claim. The check validates the simple match rather than the continuity-corrected one used in Table 4; recall of 0.58 implies substantial false negatives pushing in the opposite direction; and the two Jan de Villiers show that a single normalised name can conceal two genuine officeholders, which the continuity rule will undercount. The net direction of the error is not identified.

References

Baartman, Teunis. 2011. "Fighting for the Spoils: Cape Burgerschap and Faction Disputes in Cape Town in the 1770s." PhD diss. University of Cape Town.

Firth, David. 1993. "Bias Reduction of Maximum Likelihood Estimates." *Biometrika*, 80(1): 27–38.