HD: Consistent Targets, Uniform-Compatible Quotas, and Theorem 2 Rounding

Computational replication for Section 6.4

10 October 2026

1 Results at a glance

The supplied HD data contain 239 respondents, 33 traits, and seven dimensions. Using the group size of 30 specified in the manuscript, we constructed both requested consistent instances and drew 10,000 groups per case from the sampling procedure underlying Theorem 2. Every saved group has 30 distinct members.

Case 1 keeps the original quotas and calibrates profile-dependent probabilities by least squares relative to the uniform vector. **Case 2** uses probability 30/239 for every respondent and replaces quotas by the floor and ceiling of uniform expected representation. The calibration objective, rounding rule, and outcome definitions match the CCA experiment.

Statistic	Case 1	Case 2
Draws	10,000	10,000
Mean quota violation (33 traits)	0.269539	0.215327
Approx. 95% Monte Carlo interval	[0.267892, 0.271187]	[0.213666, 0.216989]
Mean of group maximum miss	1.6556	1.3907
Largest observed quota miss	4	4
Theoretical bound for every group	6	6
Mean number of violated traits	7.8613	6.6170
Mean total shortfall	4.4075	3.2287
Mean total excess	4.4873	3.8771

Quota misses are measured in seats. All 33 traits are nonempty, so including or excluding empty traits gives the same average. The intervals describe Monte Carlo uncertainty in estimation of the mean.

Group maximum miss	Case 1		Case 2	
	Groups	Share	Groups	Share
0	3	0.03%	20	0.20%
1	3,746	37.46%	6,286	62.86%
2	5,946	59.46%	3,471	34.71%
3	302	3.02%	213	2.13%
4	3	0.03%	10	0.10%
5	0	0.00%	0	0.00%
6	0	0.00%	0	0.00%

The observed maxima concern these samples. Theorem 2 guarantees a quota miss of at most six in every outcome of its exact-arithmetic construction. The saved draws are Monte Carlo samples, not exhaustive lottery supports.

2 Data and consistent input construction

2.1 Source and validation

This experiment follows Section 6.4 and Theorem 2 of the supplied manuscript, *Meritocratic Lotteries under Intersectional Diversity Quotas*, by Yuval Heller, Ron Peretz, and Amnon Schreiber. The new attachments are `respondents(1).csv` and `categories(1).csv`; byte-for-byte copies are saved under the standard names `input/respondents.csv` and `input/categories.csv`. Their SHA-256 hashes appear in `run_manifest.json`.

All 239 identifiers are unique, every trait field is filled, and every response matches a category row. Each respondent has one trait in each of a, b, c, d, e, f, g. These dimensions contain 9, 2, 7, 3, 6, 2, and 4 traits, respectively, totaling 33. All traits are populated. There are 202 distinct complete profiles.

The selected-group size is $k = 30$, as stated for HD in Section 6.4. Unlike the CCA input, the HD interval quotas alone do not determine a unique group size. We therefore record the manuscript as the source of this parameter. The maximum degree is $D = 7$ because every respondent affects seven trait quotas; the applicable Theorem 2 bound is $D - 1 = 6$.

2.2 Case 1: keep quotas and calibrate probabilities

Section 6.4 does not select a unique consistent target vector. We use the same explicit choice as for CCA. Let T be the set of complete profiles, with multiplicity m_t for profile t. Choose its common selection probability p_t by solving the strictly convex program

$$\begin{aligned} \min_{(p_t)_{t\in T}} \quad & \frac{1}{2}\sum_{t\in T} m_t \left(p_t - \frac{30}{239}\right)^2, \\ \text{subject to} \quad & 0 \le p_t \le 1, \qquad \sum_{t\in T} m_t p_t = 30, \\ & l_e \le \sum_{t:\, t \text{ contains } e} m_t p_t \le u_e \quad \text{for every trait } e. \end{aligned}$$

Positive multiplicities imply uniqueness of the solution. Weighting by m_t minimizes distance over respondents, not over equally weighted profiles. All original lower and upper quotas are preserved. Interval quotas remain interval constraints; the optimization does not replace them by their midpoints.

SLSQP first identifies which probability bounds and trait inequalities are active. A rational linear solve reconstructs the resulting candidate optimum exactly. Because active trait rows can be linearly dependent, a separate dual feasibility solve finds multipliers with the correct inequality signs. Those multipliers are also reconstructed rationally. We then verify, in exact arithmetic, all probability bounds, all quota inequalities, the total of 30, stationarity, dual signs, and complementarity. These conditions certify the unique global optimum. The archive supplies the fractions and multipliers for independent checking.

This calibration gives **51 respondents probability zero** (37 profiles) and **no respondent probability one**. The smallest positive probability is 0.0053948240, and the largest is 0.5357170111. The sum of squared deviations from 30/239 is 5.0294131945. The zero probabilities are a feature of this calibration objective; we do not assert that each zero is required by consistency alone. The constructed targets are not estimates of observed merit, prior priorities, or participation response rates.

2.3 Case 2: uniform probabilities and adjusted quotas

Give every respondent probability $p_v = 30/239 \simeq 0.125523$. For a trait containing N_e respondents, set

$$l_e^{\text{unif}} = \left\lfloor \frac{30N_e}{239} \right\rfloor, \qquad u_e^{\text{unif}} = \left\lceil \frac{30N_e}{239} \right\rceil.$$

Floor and ceiling are evaluated using exact rational arithmetic, ensuring consistency. This changes all 33 quota rows. For example, $c1$ has 20 respondents, so its original interval $[8, 10]$ becomes $[2, 3]$. Trait $e6$ has 200 respondents, so its interval changes from $[16, 18]$ to $[25, 26]$. The full comparison, including population counts and exact expected counts, is in `quota_audit.csv`.

3 Rounding implementation and numerical checks

For each draw, start again at $x = p$. Freeze every integral coordinate. Let F denote the fractional coordinates. Preserve the total-sum equation together with each trait-sum equation containing more than $D = 7$ members of F. Find a nonzero null direction h supported on F. Let $\alpha, \beta > 0$ be the maximal feasible step lengths in its two directions. Move to

$$x + \alpha h \quad \text{with probability } \frac{\beta}{\alpha + \beta}, \qquad x - \beta h \quad \text{with probability } \frac{\alpha}{\alpha + \beta}.$$

The step preserves the current vector in expectation and makes at least one additional coordinate integral. The total-sum equation preserves group size. When a trait equation is released, at most seven fractional coordinates remain in that trait, giving a strict count deviation below seven from its original consistent expectation. Integer quotas therefore yield quota miss at most six. The exact-arithmetic induced lottery implements the prescribed marginals.

We use the same C++ sampling code as for CCA. Null directions are selected by incremental column elimination in original respondent order. The total row comes first, followed by active trait rows in their original file order. Elimination stops at the first column dependency. New pivots use the largest-magnitude residual entry, with the first row breaking ties. The direction is normalized to maximum absolute entry one. Neither directions nor seeds were chosen to improve observed quota statistics.

Each case uses 10,000 fresh paths. Seeds are 20261012 and 20261013. The generator is standard C++ `mt19937_64`; its top 53 output bits are multiplied by 2^{-53} to obtain a number in $[0, 1)$. **No support compression is performed.** Thus the averages are estimated over the sampling law. Equal weights on the saved draws do not generally give exact marginal implementation. The additional explicit-distribution construction of Theorem 2, which could reduce support to at most 239 groups while preserving marginals, is not what is being empirically enumerated here.

Null directions use double precision; states, sums, and endpoint distances use C++ `longdouble`. Coordinates within 10^{-15} of zero or one are snapped to that endpoint. One remaining fractional coordinate is impossible in exact arithmetic under the integer total equation; a numerical residue is corrected only if within 10^{-9} of the unique required integer. The observed maxima were:

Numerical diagnostic	Case 1	Case 2
Null-direction equality residual	7.72e-15	7.11e-15
Group-size equality drift	2.90e-15	1.18e-14
Active-trait equality drift	2.63e-15	9.95e-15
Final numerical correction	2.80e-15	1.16e-14
Groups requiring final correction	1,160	9,798

The implementation checks the total and active trait equations at every step. A separate Python program independently checks the final groups and all integer trait counts, exact target feasibility, the

calibration certificate, and the reported summary statistics. These checks passed for all 20,000 groups. The ideal mathematical guarantee and measured floating-point residuals are reported separately; numerical errors are not claimed to be identically zero.

4 Outcome definitions and interpretation

For group S and trait e, the seat count and quota miss are

$$c_e(S) = |S \cap e|, \qquad \delta_e(S) = \max\{l_e - c_e(S),\ c_e(S) - u_e,\ 0\}.$$

A count inside its quota interval has zero miss. Shortfalls and excesses are not canceled. For every realized group we record

$$A(S) = \frac{1}{33}\sum_e \delta_e(S), \qquad M(S) = \max_e \delta_e(S).$$

The first-page average is $R^{-1}\sum_{r=1}^{R} A(S_r)$, with $R = 10{,}000$. The histogram separately records $M(S_r)$ for every draw. Approximate 95% Monte Carlo intervals are the average plus or minus $1.96 s_A/\sqrt{R}$. The standard deviation s_A is calculated across groups; the 33 traits within a group are not treated as independent observations.

The group maximum miss was one or two seats in 96.92% and 97.57% of the groups, respectively. A four-seat maximum occurred in 3 groups in Case 1 and 10 in Case 2. The first-page maximum and histogram give the values for these runs. The original-quota run has maximum 4; the earlier maximum of three in the manuscript should therefore be updated when describing this calibration and run.

The largest absolute empirical marginal errors were 0.012700 and 0.012223 in the respective cases. These are maxima across 239 respondents, not individual confidence limits. Each respondent's target, number of selections, empirical frequency, frequency error, and binomial standard error $\sqrt{p_v(1-p_v)/R}$ are archived in `marginal_diagnostics.csv`. There were 10,000 and 10,000 distinct sampled groups; these numbers do not enumerate or estimate the full support.

The two cases have different quota systems and target vectors. Their mean misses describe those different instances under the same rounding rule. Other consistent targets or null-direction choices can produce different averages. This run provides reproducible HD numbers for the specified choices; it does not assume the numerical statements previously appearing in Section 6.4 as inputs. No LEXIMIN comparison is computed in this experiment.

5 Replication files and commands

The package includes the original CSVs, both modified instances, all selected groups, every group's requested metrics and all 33 counts, exact probabilities, the optimality certificate, and independent verification code.

- The Case 1 files are in `case1_original_quotas/`. The respondent file adds a column named `selection_probability`. The quota file is unchanged. Exact fractions appear in `profile_targets.csv`; the decimal probability column uses a 17-digit numerical representation.
- `case2_uniform/categories.csv` contains the adjusted quotas. Its companion `respondents.csv` gives the common probability 30/239.
- Both folders list each group's members in `panel_members.csv` and its requested metrics and all trait counts in `panel_metrics.csv`. Aggregates appear in `trait_results.csv` and `summary.json`; individual frequency diagnostics appear in `marginal_diagnostics.csv`.

- `run_experiment.py` prepares the inputs and calls `theorem2_sampler.cpp`. The standard-library script `verify_results.py` checks the saved results independently. `run_manifest.json` records versions, seeds, and hashes; `verification.json` records the verification outcome.

The archived environment used Python 3.12.14, NumPy 2.3.5, SciPy 1.17.0, and GCC 13.3.0 on x86_64 Linux. From the extracted package directory:

```
python3 -m pip install -r requirements.txt
OPENBLAS_NUM_THREADS=1 OMP_NUM_THREADS=1 \
  python3 run_experiment.py --draws 10000 --seed 20261012
python3 verify_results.py
python3 build_report.py
pdflatex -interaction=nonstopmode -halt-on-error HD-results-report.tex
pdflatex -interaction=nonstopmode -halt-on-error HD-results-report.tex
```

Add `--recalibrate` to reconstruct the probability vector and certificate from the raw inputs, rather than loading the archived rational targets. A C++17 compiler is needed for sampling and LaTeX only for rebuilding the PDF. The verifier requires only Python's standard library. Different compiler or architecture floating-point behavior can change branching decisions; the saved panel files are the exact empirical record underlying this report.

A All-trait results

"Pop." is the trait population. Each quota is lower–upper. "Target 1" is the Case 1 expected count (rounded to four decimals here; exact in the audit). The Case 2 expected count is always $30N_e/239$. "Mean" is the mean quota miss over sampled groups and "Max" its largest observed value for that trait.

		Case 1: original quotas				Case 2: uniform		
Trait	Pop.	Quota	Target 1	Mean	Max	Quota	Mean	Max
a1	44	2–4	4.0000	0.2821	3	5–6	0.2708	3
a2	33	2–4	4.0000	0.3158	3	4–5	0.3512	3
a3	20	2–4	3.3629	0.1446	3	2–3	0.2557	3
a4	22	2–4	4.0000	0.3504	4	2–3	0.2924	4
a5	36	2–4	4.0000	0.3089	3	4–5	0.3032	4
a6	15	2–4	2.0119	0.3118	2	1–2	0.2803	3
a7	33	2–4	4.0000	0.2271	3	4–5	0.3280	3
a8	16	2–4	2.0000	0.3461	2	2–3	0.3646	3
a9	20	2–4	2.6252	0.1456	2	2–3	0.2754	3
b1	112	13–15	15.0000	0.2645	2	14–15	0.2291	2
b2	127	15–17	15.0000	0.2645	2	15–16	0.2291	2
c1	20	8–10	8.0000	0.3824	3	2–3	0.2342	3
c2	30	4–6	4.0000	0.3503	3	3–4	0.2886	4
c3	39	4–6	4.8339	0.0529	2	4–5	0.3126	4
c4	39	2–4	4.0000	0.3036	3	4–5	0.2881	4
c5	45	2–3	3.0000	0.3161	3	5–6	0.1979	2
c6	40	3–5	3.9305	0.0439	2	5–6	0.3743	3
c7	26	2–4	2.2356	0.1310	1	3–4	0.2767	2
d1	133	13–15	14.3959	0.0600	2	16–17	0.1615	2
d2	100	13–15	13.6041	0.0504	1	12–13	0.0968	1
d3	6	2–2	2.0000	0.1353	1	0–1	0.0078	1
e1	7	2–2	2.0000	0.7807	4	0–1	0.2097	3
e2	4	2–2	2.0000	0.7164	2	0–1	0.0729	2
e3	11	4–6	4.0000	0.4131	3	1–2	0.2390	3
e4	13	2–4	2.0000	0.2945	2	1–2	0.1885	3
e5	4	2–2	2.0000	0.8892	2	0–1	0.0797	2
e6	200	16–18	18.0000	0.2444	2	25–26	0.1562	2
f1	35	4–6	5.6500	0.0004	1	4–5	0.0008	1
f2	204	24–26	24.3500	0.0004	1	25–26	0.0008	1
g1	5	2–3	2.0000	0.1703	2	0–1	0.0807	2
g2	19	4–6	4.0000	0.2697	2	2–3	0.2388	3
g3	73	11–13	12.0000	0.0552	2	9–10	0.2563	3
g4	142	10–12	12.0000	0.2732	2	17–18	0.1641	2