HESHBON

Center on Data and Jewish Life
heshbon.org

חשבון
ḥeshbon · ‘reckoning’

METHODOLOGY AND DATA COMPANION

OPEN ACCESS

The Build-a-School Model

Methodology and data companion to the working paper

Gavriel Brown, PhD

A full white paper, with the complete introduction, related literature, data and design, results, discussion, limitations, conclusion, and references, is available on request.

At a glance	
Research question	What does a Jewish day school actually cost, decision by decision, and can tuition cover it?
Data	Merged NAIS, PRIZMAH, and SAIS day-school benchmarking exports; an IRS Form 990 head-of-school compensation panel hand-corrected to one record per person; and a hierarchical per-pupil analysis panel.
Unit of analysis	A hypothetical school assembled by the user, priced per pupil and compared against real school distributions.
Statistical method	A bottom-up, decision-driven cost-simulation model anchored to benchmarking percentiles, not a hypothesis test.
Headline finding	At typical settings the model produces about $21,000 per pupil, and at fully resourced settings about $27,400, within several hundred dollars of the real median of $27,434.
What it shows	A transparent, auditable translation of benchmarking data into per-pupil cost, with each figure disclosed as benchmarked or estimated.

The question behind the paper

Tuition at a Jewish day school is a published number; the cost it is meant to cover is not. A family, a board member, or a prospective funder has little way to see what a tuition dollar actually buys, whether teachers are paid at the regional floor or well above it, whether there is a dedicated development professional, or how a fundraising gap traces back to specific staffing and program choices.

The benchmarking data to answer this largely exists. What has been missing is a way to turn a percentile table into an answer to the question a board actually asks, which is what a given set of staffing and program choices costs per child, and whether tuition can cover it. This paper documents a model built to do exactly that.

Where the data come from

Benchmarking exports. Merged NAIS, PRIZMAH, and SAIS day-school benchmarking data supply full-time teacher base salaries by percentile, administrator salaries, operating expenses, the facilities, technology, and general administrative lines, enrollment, and gross and net tuition revenue by peer group. These anchor the teacher pay tiers, the fixed-overhead base, and the financial-aid discount tiers.

Head-of-school compensation panel. An IRS Form 990 panel of head-of-school compensation, hand-corrected to one record per canonical individual, is cross-tabulated by a revenue-based school-size tier and an experience tier. It anchors the head-of-school pay ladder, and by extension the relational principal and assistant-principal chain, which the source data do not report directly.

Per-pupil analysis panel. A hierarchical per-pupil panel, covering non-Orthodox Jewish, Orthodox, and general independent schools, supplies the real per-pupil cost distribution the model compares against, along with the ratio of net tuition to operating expense used to check

the revenue side. Two corrections preceded any use of these sources: the benchmarking net operating surplus field defaults to zero for non-reporting schools, so every ratio here uses total operating income and expense directly, and the compensation fields are per-person snapshots treated as cross-sectional evidence about pay levels rather than a longitudinal series.

The statistical method, named and explained

This is a methods paper rather than an empirical study, and it reports no hypothesis test. Its object is a bottom-up, decision-driven cost-simulation model: a user assembles a hypothetical PreK4 school through twenty-nine discrete decisions across seven domains, and each decision resolves into one of thirteen structural cost categories that price out per child. Every category contributes zero dollars until its governing decision is made, so the running total reflects the user's own choices rather than a designer's default. Wherever a figure could be tied to a reported benchmark it was, and wherever it could not, as with principal pay, program costs, and the economies-of-scale exponents, it is disclosed as an informed estimate. Fixed overhead is scaled for school size using a square-root rule and administrative staffing with a shallower exponent, both stated as reasonable functional forms rather than fitted curves.

Validation is a convergence check rather than a claim of external validity. Run at typical settings the model produces about $21,000 per pupil, just below the real 25th percentile; pushed to its most resourced tier it reaches about $27,400, within several hundred dollars of the real median of $27,434; and at its leanest it falls to about $17,400. None of the three figures was targeted in advance.

Key results from the paper

Table 1. The three underlying datasets.

Dataset	Role in the model
NAIS, PRIZMAH, and SAIS benchmarking exports (merged peer-group panel)	Anchors teacher pay tiers, the fixed-overhead base (facilities, technology, general administration), and the financial-aid discount tiers.
IRS Form 990 head-of-school compensation panel	Anchors the head-of-school pay ladder, and the principal and assistant-principal chain defined as a fixed ratio of it.
Hierarchical per-pupil analysis panel	Supplies the real per-pupil comparison distribution and the tuition-to-cost ratio used to validate the revenue side.

Table 2. Selected cost anchors drawn from the benchmarking data.

Anchor	Value
Teacher pay: regional minimum / competitive / leading	$48,000 / $66,000 / $82,000
Head-of-school pay: small first-time to large veteran	$175,000 to $485,000
Facilities median (259-student school)	$730,907
Technology median	$121,419
General administrative median	$418,350

Note. The overhead medians are scaled by school size; the teacher and head-of-school figures anchor the pay ladders.

Figure 1. The convergence check: modeled cost per pupil at three settings.

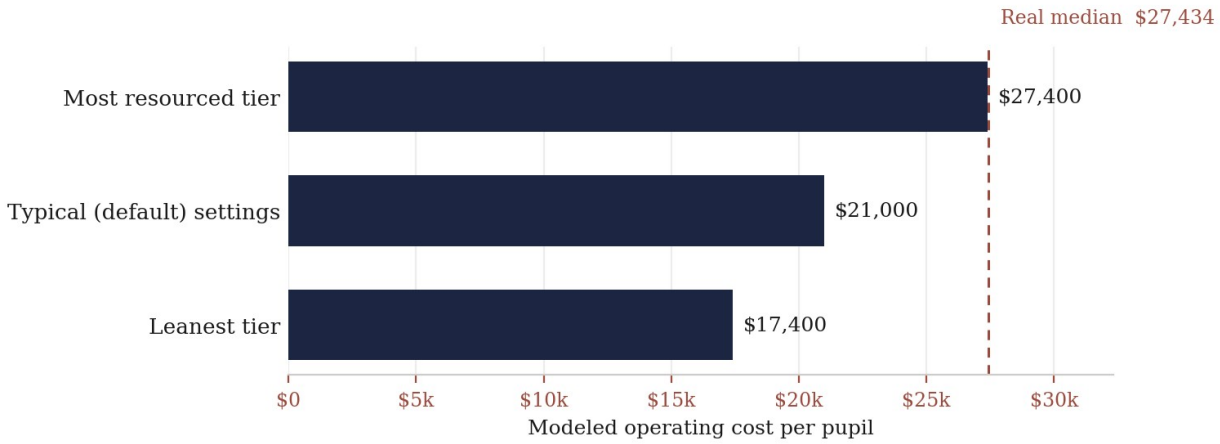

Note. The real per-pupil distribution runs from $19,523 (25th percentile) to $34,335 (75th), with a median of $27,434. None of the three modeled figures was targeted in advance.

How to read the results

Table 1 sets out what each dataset contributes, and Table 2 lists the concrete anchors the model reads from them, from a three-tier teacher pay scale to a head-of-school ladder running from $175,000 to $485,000. Figure 1 shows the convergence check: the modeled cost per pupil at leanest, typical, and most resourced settings, against the real per-pupil distribution whose median is $27,434. The proximity of the independently built figures to that distribution is offered as evidence that the model's cost categories are, in aggregate, of the right order of magnitude and not merely internally self-consistent.

What the numbers do not claim

The model shows what a school decides to cost, not what any real school will cost, and it makes no causal claim. Program-level costs such as special education, security, and Israel programs are planning estimates built from typical staffing ratios rather than reported benchmarks, and the economies-of-scale exponents and the high-school enrollment share are assumptions rather than fitted values. Some subgroups are small, with the Orthodox peer group contributing on the order of six to ten schools to certain ratios, and the model is a single-year snapshot that has not yet been checked against real school budgets.

Reproducing the analysis

The model's grounding is public or reconstructable. The benchmarking distributions come from NAIS, PRIZMAH, and SAIS day-school data; the leadership ladder rests on head-of-school compensation compiled from IRS Form 990 filings and hand-corrected to one record per person; and the per-pupil comparison band comes from the hierarchical per-pupil panel. Every cost category discloses whether it is benchmarked or estimated, so the build can be audited line by line and re-run as better data becomes available.