DISTRICT FORESIGHT

# Forecast Methodology & Measured Performance

Enrollment projection for public school districts · Technical report · July 6, 2026

## Executive summary

District Foresight projects K–12 enrollment for public school districts five to ten years ahead. This report documents the model and, more importantly, measures it: we projected every Pennsylvania school district from every base year 1991–2023, at horizons of one to ten years, using only information available at each projection date, and scored 141,134 projections against realized enrollment.

| 10.3% | –25.6% | –23.5% | 61.6% |
|---|---|---|---|
| mean 10-year forecast error (MAPE), all districts and base years | error vs. 5-year trend extrapolation (13.8%) at 10 years | error vs. constant-enrollment baseline (13.5%) at 10 years | of 10-year head-to-heads won against the constant baseline |

The model beats both baselines at every horizon, with the advantage widening as the horizon grows. Half of all ten-year projections land within 7.9% of the eventual figure — a measured, defensible planning range. The same experiment quantifies how new housing converts to students (a rising multi-year yield curve, single-family above multi-family) and yields an honest decomposition of where the model's edge comes from.

## 1. The experiment

Accuracy claims in this report come from a **walk-forward backtest**, the same standard used to evaluate trading strategies: simulate the forecast as it would actually have been made, then score it against what happened.

1. **Panel.** All 499 Pennsylvania public school districts (charter and technical schools excluded), fall 1986 through fall 2024, enrollment by district and grade.
2. **Protocol.** For each district and each base year B from 1991 to 2023: fit every model using only data through B; project K–12 totals for B+1 … B+10; score each horizon h against realized enrollment at

B+h.

3. **No lookahead.** Model components that are estimated across districts (the housing regression, recalibration) are re-fit at each base year using only outcomes already observed by that year. A projection made "from 2005" uses nothing that happened after 2005.
4. **Eligibility.** A district enters a window when it has five consecutive years of complete grade-level data through B and observed enrollment at B+h; all models are scored on identical district-year pairs.
5. **Metrics.** Absolute percentage error (APE) per projection; we report the mean (MAPE), median, and interquartile range by horizon, plus head-to-head win rates on identical pairs.

## 2. Data

| SERIES | SOURCE | COVERAGE |
|---|---|---|
| District enrollment by grade | NCES Common Core of Data (Urban Institute harmonized files) | 1986–2024, national |
| Residential building permits by municipality (1-unit / 2–4 / 5+ structures) | U.S. Census Bureau, Building Permits Survey | 1980–2025, national |
| District ↔ municipality crosswalk | NCES EDGE geographic relationship files | current vintage |
| October-1 enrollment; official district projections | Pennsylvania Department of Education | 2004–2025; projections to 2034 |

Permits are joined to districts through each district's constituent municipalities (county subdivisions); 83% of Pennsylvania's permitted units map cleanly to a district. Pre-2007 permit records lack FIPS codes and are joined by county and municipality name.

## 3. The models

### Baselines (the models a district gets for free)

**Constant enrollment** projects the base-year K–12 total unchanged. **Trend extrapolation** fits an ordinary least-squares line through the last five annual totals and extends it. Both are scored on exactly the same district-years as the model.

### Cohort survival (grade progression)

The core engine. For each grade g, a progression ratio $r_g$ is the average, over the last four observed transitions, of enrollment in grade g divided by enrollment in grade g−1 the year before — capturing retention, in/out-migration, and public/private shifts as the district actually experiences them. Ratios are

capped at 2.0 to guard tiny-feeder artifacts; genuinely small ratios are preserved, because a cohort that largely departs — an exit grade in a K–8 district, or students leaving to a regional high school — is structure, not noise. **Entry grades** — kindergarten always, plus any grade whose feeder grade is structurally empty (grade 7 in a 7–12 regional district, for example) — are projected from the mean of their own last three cohorts. This one rule makes the same engine correct across district structures nationally. Cohorts are rolled forward grade by grade; the projected total is the sum over K–12. This is the same model family the Pennsylvania Department of Education uses for its official projections.

### Housing signal

Building permits enter as *deviations from the district's own norm*: units permitted in the three years ending at the base year, minus three years at the district's trailing five-year rate (measured over the preceding window), per enrolled student, floored at zero, separately for single-family and multi-family structures. The floor reflects a measured asymmetry: construction booms add students, but a construction *pause* does not subtract them — existing homes keep their children. A pooled regression maps these deviations to enrollment above or below the cohort-survival path at each horizon, re-fit walk-forward with only observed outcomes.

# 4. Results

**Mean forecast error by horizon**

MAPE across 499 districts, base years 1991–2023, walk-forward.

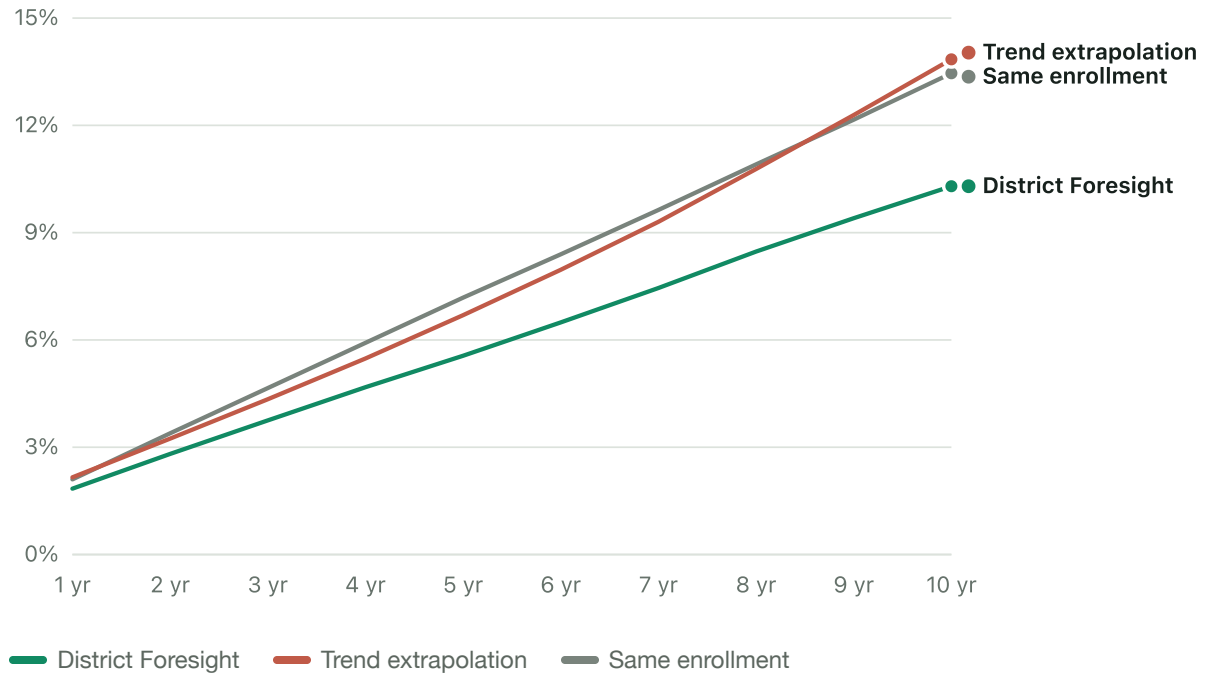

| HORIZON | DF MODEL MAPE | MEDIAN | IQR | COHORT-ONLY MAPE | TREND MAPE | CONSTANT MAPE | WIN VS CONSTANT | WIN VS TREND | N |
|---|---|---|---|---|---|---|---|---|---|
| 1 yr | 1.84% | 1.29% | 0.6%–2.4% | 1.85% | 2.16% | 2.10% | 57.1% | 59.0% | 16,353 |
| 2 yr | 2.81% | 2.08% | 1.0%–3.7% | 2.83% | 3.24% | 3.39% | 58.3% | 58.5% | 15,855 |
| 3 yr | 3.75% | 2.84% | 1.3%–5.0% | 3.80% | 4.34% | 4.65% | 59.7% | 59.2% | 15,358 |
| 4 yr | 4.68% | 3.56% | 1.7%–6.2% | 4.80% | 5.48% | 5.92% | 60.0% | 59.4% | 14,860 |
| 5 yr | 5.56% | 4.22% | 2.0%–7.4% | 5.79% | 6.69% | 7.18% | 61.1% | 61.0% | 14,362 |
| 6 yr | 6.49% | 4.98% | 2.3%–8.6% | 6.81% | 7.96% | 8.39% | 60.9% | 61.2% | 13,864 |
| 7 yr | 7.45% | 5.76% | 2.7%–9.9% | 7.89% | 9.30% | 9.63% | 60.8% | 61.2% | 13,366 |
| 8 yr | 8.47% | 6.52% | 3.0%–11.3% | 9.08% | 10.78% | 10.91% | 60.4% | 61.7% | 12,869 |
| 9 yr | 9.40% | 7.15% | 3.4%–12.6% | 10.22% | 12.29% | 12.16% | 60.9% | 62.7% | 12,372 |
| 10 yr | 10.30% | 7.89% | 3.7%–13.8% | 11.40% | 13.84% | 13.46% | 61.6% | 63.8% | 11,875 |

Three observations. First, the model dominates both baselines at every horizon, and the margin grows with distance: at ten years the error reduction is 25.6% against the trend line and 23.5% against constant enrollment. Second, trend extrapolation is barely better than doing nothing at short horizons and *worse* than doing nothing at ten years (13.84% vs 13.46%) — recent direction is not a forecast. Third, the win rates show the improvement is broad, not an artifact of a few districts.

**Error distribution by horizon (District Foresight model)**

Median absolute error with the interquartile band — the measured planning range.

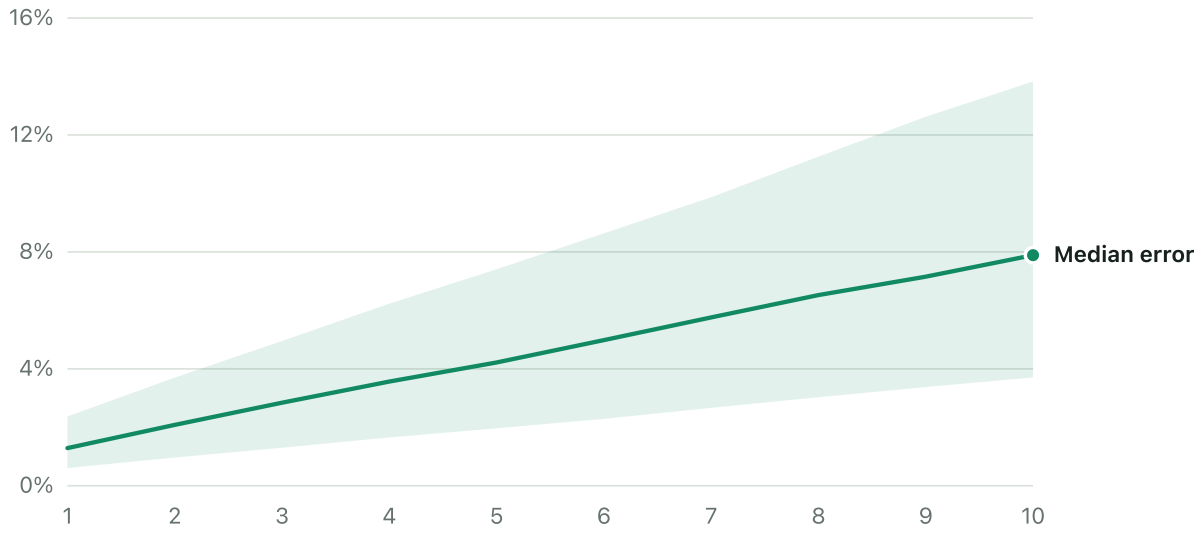


## 5. Robustness across eras

A model that only worked in one macro regime would be a liability. Grouping the backtest by the decade of the projection date:

| HORIZON | BASE YEARS | DF MODEL MAPE | TREND MAPE | CONSTANT MAPE | N |
|---|---|---|---|---|---|
| 5 yr | 1990s | 5.46% | 6.94% | 6.67% | 4,453 |
| 5 yr | 2000s | 5.52% | 6.72% | 7.60% | 4,957 |
| 5 yr | 2010s | 5.67% | 6.44% | 7.23% | 4,952 |
| 10 yr | 1990s | 11.43% | 14.82% | 12.50% | 4,453 |
| 10 yr | 2000s | 9.45% | 13.65% | 14.20% | 4,950 |
| 10 yr | 2010s | 9.94% | 12.48% | 13.69% | 2,472 |

The advantage holds through the 1990s expansion, the mid-2000s housing boom, the 2008 crash, and the post-2010 statewide enrollment decline.

## 6. The housing signal, measured

**Students per excess housing unit permitted, by years elapsed**

Pooled regression coefficients on cohort-survival residuals, full sample.

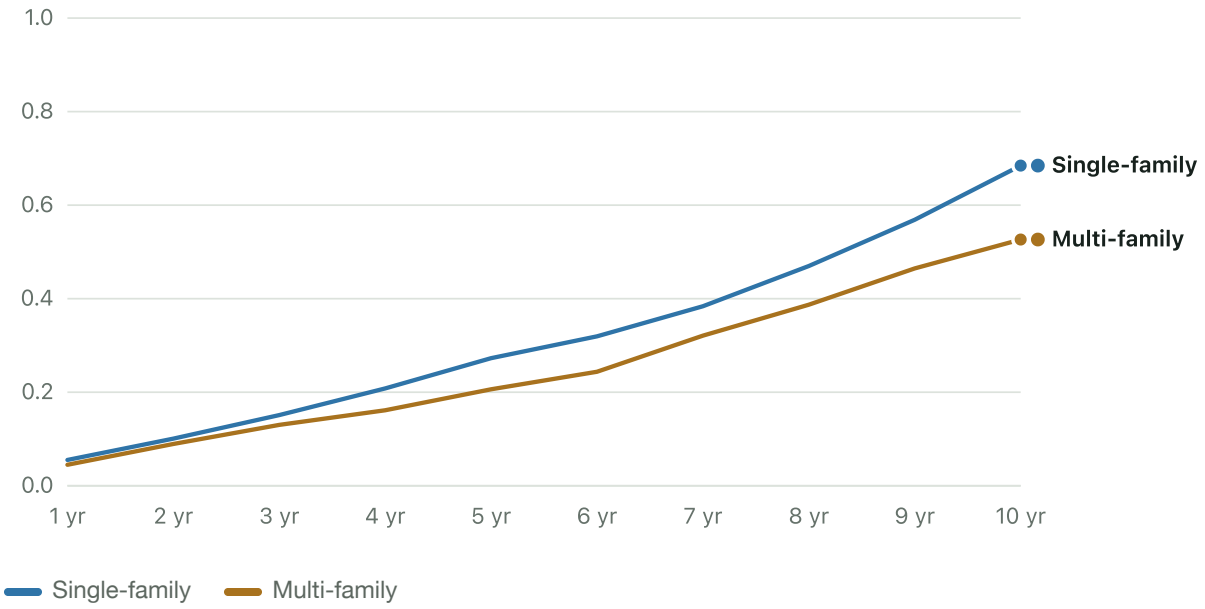


| YEARS AFTER PERMITS | STUDENTS / EXCESS SF UNIT | STUDENTS / EXCESS MF UNIT | RECALIBRATION INTERCEPT | N |
|---|---|---|---|---|
| 1 | 0.055 | 0.045 | -0.003 | 16,353 |
| 2 | 0.101 | 0.089 | -0.007 | 15,855 |
| 3 | 0.151 | 0.130 | -0.012 | 15,358 |
| 4 | 0.208 | 0.161 | -0.018 | 14,860 |
| 5 | 0.273 | 0.206 | -0.025 | 14,362 |
| 6 | 0.320 | 0.244 | -0.031 | 13,864 |
| 7 | 0.384 | 0.321 | -0.038 | 13,366 |
| 8 | 0.470 | 0.387 | -0.046 | 12,869 |
| 9 | 0.569 | 0.465 | -0.055 | 12,372 |
| 10 | 0.684 | 0.527 | -0.064 | 11,875 |

The yield curve rises monotonically with elapsed time — new homes fill with school-age children over years, not months — and single-family units out-yield multi-family units throughout, consistent with published residential demographic multipliers (e.g., the Rutgers Center for Urban Policy Research series). The negative intercepts reflect the statewide enrollment decline over the sample: cohort-survival ratios estimated from recent history lag a persistent downtrend, and the walk-forward recalibration corrects for it.

> **An honest decomposition.** Two findings temper the housing story. First, a district's own grade-progression ratios absorb most of a housing boom within a few years of its start — the ratios are measured from the very years in which new units were filling with students — so the explicit permits term adds the most information *early* in a boom, before absorption. Second, in aggregate, most of the model's improvement over plain cohort survival comes from the walk-forward recalibration rather than the permits features: an intercept-only variant scores within a tenth of a point of the full model on statewide MAPE. The permits term earns its place in the minority of district-years with genuinely unusual construction — which are exactly the situations in which districts seek outside help — and in scenario analysis, where a proposed development can be folded directly into the projection through the measured yield curve.

## 7. Limitations

- **District totals, not schools.** This backtest scores district-wide K–12 totals. School-level and grade-band projections inherit additional boundary and program noise; engagement work models them explicitly.
- **Pennsylvania panel.** Results are measured on all 499 PA districts. The method is national — every input series exists for every state — but published numbers here are PA-specific.
- **Charter interaction.** Charter enrollment shifts are captured implicitly through progression ratios, not modeled structurally.
- **Permits ≠ completions.** Some permitted units are never built; the fitted yield curve absorbs average realization rates but not project-level outcomes. Engagement work replaces the average pipeline with the district's actual development schedule.
- **Structural breaks.** No statistical model anticipates a district consolidation, a major program change, or a pandemic. The measured error bands include such events; point forecasts do not predict them.

## 8. Reproducibility

Every input is public: NCES Common Core of Data enrollment (via the Urban Institute's harmonized extracts), the Census Bureau Building Permits Survey, NCES EDGE geographic relationship files, and

Pennsylvania Department of Education enrollment files and official projections. The full pipeline — download, database construction, models, and this backtest — is a single auditable codebase, and any figure in this report can be traced to its source records.

## Appendix A. Definitions

- **APE** — absolute percentage error: |projected − actual| / actual.
- **MAPE** — mean APE over a set of projections.
- **IQR** — interquartile range: the 25th–75th percentile span of APE.
- **Win rate** — share of identical district/base-year/horizon pairs where a model's APE is strictly lower than the comparison baseline's.
- **Horizon** — years between the projection date (base year) and the projected fall.
- **Base year** — the fall whose data is the last a projection may use.