

How Districts Can Align AI with Students' Civil Rights: Findings from the Urban AI Unlocked Project

Patricia Burch, Junmin Byon, Alvin Makori, Shelby Smith, Eric Larsen, Dedrick
McCord

Rossier School of Education & USC EdPolicy Hub

June 2026

Notes: This project was supported by the Google Foundation. All opinions expressed in this report represent those of the authors and not our funder nor the institution with which we are affiliated. All errors in this report are solely the responsibility of the authors.

The authors would like to thank Dr. Jerry Alamanderez (former Superintendent) Santa Ana Unified School District, California; Dr. Maurice R Draggan, Chief Information Officer, Palm County Public Schools, Florida and Dr. Tarek Israwi, Assistant Director of Technology & Information, Downey Unified School District, CA for expert feedback on survey design and report recommendations.

Executive Summary

Urban AI Unlocked is a national study of how large urban districts are buying and using artificial intelligence (AI tools), framed explicitly as a civil-rights issue rather than simply an adoption and implementation story. Here, *civil rights refer to students' concrete entitlements to access, nondiscrimination, privacy, fair decision-making, language access, disability-related protections, and meaningful voice in decisions that affect them*. The report asks whether AI helps districts meet those obligations or shifts new risks onto the communities that schools have historically marginalized.

Using a civil-rights-first framework, we examine where AI is being deployed, how decisions are made, who profits, and what safeguards are in place, and then translate those findings into six policy recommendations for big-city systems.

The analysis combines transaction-level purchasing data from thousands of districts, a survey of leaders in 27 Council of the Great City Schools districts, National Center for Education Statistics (NCES) demographic data, and market data on major K–12 AI vendors and their investors. AI is already becoming part of how large districts make decisions about teaching, support, placement, and access, especially in systems serving students with the greatest historical exposure to educational inequity. The central challenge is no longer whether districts will encounter AI, but whether they will govern it in ways that protect students' civil rights in practice rather than only in principle.

The report offers six civil-rights-grounded strategies: treat high-risk AI as an exception requiring explicit justification; make equity an enforceable condition of adoption; make bias and performance checks minimum compliance duties; use contracts to prevent AI lock-in and protect student rights; treat who gets AI as an equity decision; and build governance structures that shift power toward students, families, and educators in high-impact schools.

Executive Summary	2
Introduction	4
Our Aims	4
Why a Civil-Rights Frame for AI in Districts	5
AI Use Where Civil Rights Concerns Are Especially Sharp	5
A Civil-Rights-First Lens on AI Procurement	6
Report Findings	6
What the Data Show About Adoption and Spending.....	6
Policy Implications and Recommendations	8
What Concrete Steps Can a District Take?.....	8
1. Treat high-risk AI as a civil-rights exception, not the default.....	9
2. Make equity an enforceable condition for AI, not just a stated value	10
3. Make bias and performance checks a minimum compliance duty	11
4. Use contracts to prevent AI lock-in and protect student rights.....	12
5. Make “who gets AI” an equity decision, not a convenience decision	13
6. Build AI governance that shifts power, not just structure	14
How to Read These Recommendations in Practice.....	15
Appendices.....	16
Research Design	16

Introduction

This report shares early findings from Urban AI Unlocked, a national study of how large urban school districts are buying and using AI tools, and what that means for students' civil rights. It is written for district leaders, educators, families, and advocates who want AI to support, not undermine, more just public education.

The core question is simple: not “Is AI innovative?” but “Does AI help fulfill students' rights, or does it shift new risks onto communities that schools have historically marginalized?” To answer that, the project looks closely at where AI is being used, how decisions are being made, who owns and profits from the tools, and what safeguards districts have in place. The report closes with six practical policy recommendations that districts can act on now.

In this report, “high-stakes AI” does not mean that all AI is high stakes. It refers to AI uses in areas such as risk scoring, placement, discipline, language access, and access to tutoring or interventions, where the tool can shape students' educational opportunities, services, or treatment. Likewise, “civil rights” is used here in the education-specific sense developed throughout this report: students' concrete entitlements to access, nondiscrimination, privacy, fair decision-making, language access, disability-related protections, and meaningful voice in decisions that affect them.

Our Aims

The report is aligned with national concerns about bias, privacy, and evidence, but it is distinct in treating AI procurement itself as a civil-rights process, not a technology or innovation issue with equity layered on later. This report does the following to evaluate AI procurement as a civil-rights process:

- Documents *who is exposed to AI tools and where the stakes are highest*, then builds recommendations specifically for those systems rather than for an abstract average district.
- *Quantifies a gap between what districts have on paper and what they do in practice* for equity, then targets that gap with specific strategies.
- 1. Ties civil rights risks not just district policy, but also *financial models and consolidation on the supply side*.
- *Translates findings into six plain-language strategies*: guardrails on high-risk uses, equity impact requirements, bias checks, contract protections, targeting investments to high-need students, and standing AI governance teams.
- Shows how these *strategies map onto actions already being taken* and that fit the realities of big-city systems.

Why a Civil-Rights Frame for AI in Districts

This report uses the language of civil rights on purpose, and more narrowly than the way equity is often used in ed-tech conversations. Here, civil rights do not mean a general commitment to fairness. It refers to students' specific entitlements under education, nondiscrimination, language-access, disability, and privacy laws — for example, the right to non-discriminatory access to rigorous coursework, the right to appropriate services and accommodations, the right to learn in and through their home language, and the right to have their data and dignity protected.

A civil-rights frame shifts the starting question from “How can school systems use AI responsibly?” to “Can school systems still meet their existing obligations to students if they use AI in this way?” Instead of beginning with what AI tools can do and then asking whether they seem fair, the report begins with four basic obligations districts already recognize: providing access and nondiscrimination, protecting privacy and data governance, ensuring fair and explainable decision-making, and creating meaningful voice and agency for students, families, and educators. The core test is whether AI makes it easier or harder to uphold those obligations in practice.

This framing is especially important for large urban districts because per-pupil spending on AI is highest in high-enrollment school systems serving many racially marginalized students. Some of the fastest-growing uses of AI such as predictive risk scores, placement tools, tutoring, and intervention platforms sit directly on top of rights-sensitive decisions about who is identified as at risk, who is placed in which program or track, and who gets extra support. When AI enters those decisions, it does not replace the district's civil-rights duties; it changes how those duties should be carried out and monitored.

While equity is the broader goal of closing opportunity gaps and repairing historical harms, a civil-rights frame focuses on the minimum, non-negotiable protections that must be in place as districts adopt AI, especially in high-stakes domains. The recommendations that follow are designed as conditions under which AI adoption can be squared with those obligations in big-city school systems, not as an endorsement of AI use in every context.

AI Use Where Civil Rights Concerns Are Especially Sharp

Below are three domains where AI applications raise particularly sharp civil-rights risks.

Predictive risk systems and lowered expectations. Commercial early warning and risk models often use race, discipline history, and other sensitive data to predict who is “at risk,” which can reinforce bias under the label of data-driven insight. Students tagged as high risk may quietly be steered away from rigorous courses, with little chance to challenge the label.

Emergent bilingual students and language rights. AI translation and placement tools are often trained on English-dominant data, which can misread bilingual students' abilities and push them toward lower tracks. Protections for language rights rarely appear as enforceable clauses in vendor contracts.

AI tutoring in under-resourced schools. Wealthier districts are more likely to get higher-quality AI support, while high-need schools often lack priority access and implementation support. Without explicit equity commitments in procurement, personalized tutoring can become another layer of unequal opportunity.

Taken together, these examples show that AI is being normalized in areas that shape access to courses, services, and fair treatment in school, not just back-end operations.

A Civil-Rights-First Lens on AI Procurement

Instead of starting from what AI can do, the project starts from what students are legally and morally entitled to. The team uses a four-part civil-rights-first framework to evaluate AI procurement and use.

- **Access and nondiscrimination** ask whether AI tools are helping to open doors to courses and services or quietly closing them for some groups.
- **Privacy and data governance** asks whether AI systems treat students' data and identities with at least the same care districts already expect under FERPA and other privacy laws.
- **Fairness in decision-making** asks whether AI-supported decisions are open to scrutiny and free from systematic bias, especially when they affect students' lives.
- **Voice and agency** ask whether students, families, and educators have real power to shape, question, and contest AI use.

To apply this framework, the study combines transaction-level purchasing data for thousands of districts, a survey of leaders in 27 large urban systems, demographic data from NCES, and information about the major K–12 AI vendors and their investors. This makes it possible to see both where AI money is going and how prepared districts are to govern it. See the appendix for details on research design.

Report Findings

What the Data Show About Adoption and Spending

2. AI is spreading quickly in K–12, and especially in large, high-need systems.

Since 2018, thousands of districts have made their first AI-related purchases, and many have gone on to buy AI tools in multiple years, making AI a recurring line item rather than a one-off experiment. Even in states where only a minority of districts have AI contracts, those districts tend to be large, so a much larger share of students is in AI-using systems.

Districts serving mostly students of color spend more per pupil on AI than majority white districts, meaning AI tools are thickest on the ground where civil-rights protections are most crucial.

Within the CGCS sample, governance infrastructure tends to be somewhat stronger in medium and large districts than in the smallest ones, but equity-oriented practice does not consistently line up with higher-need contexts.

3. Governance capacity exists, but equity practice lags.

Governance capacity exists, but equity practice is less consistently institutionalized: governance scores cluster in the middle-to-high range, while equity-oriented practice varies more widely across districts. (See Appendix Table 4) About 70 percent of surveyed districts say they conduct formal needs assessments before AI-related purchases, and over 80 percent report including four or more privacy and data-governance protections in their AI contracts. Roughly three-quarters explicitly mention equity or fairness in their evaluation criteria, yet fewer than one in five districts regularly prioritize AI investments for high-need students in a systematic way. In most surveyed districts (65%), teachers participate in AI-related procurement at minimal or no levels, and fewer than four in ten districts have formal mechanisms for teacher, student, or family input on AI equity impacts. Taken together, these patterns suggest that while many districts have policies and infrastructure on paper, far fewer have built-in routines that make equity a binding constraint on day-to-day AI decisions.

4. Vendor and Investor Incentives May Conflict with Civil Rights

Major vendors such as GoGuardian, Securly, Frontline, HMH, and PowerSchool are owned or backed by large private-equity firms that seek rapid earnings growth, often through consolidation and pricing changes. This can encourage recurring, bundled contracts that are hard for districts to unwind and may not reward investments in strong equity safeguards. At the same time, none of the highest-revenue companies in the sample holds the strongest level of research backing for its flagship products, which suggests that market dominance reflects procurement relationships and infrastructure lock-in more than clear proof that the tools reliably improve learning.

Policy Implications and Recommendations

What Concrete Steps Can a District Take?

Our recommendations are organized around two central principles. *First, they build on what urban districts generally already do: needs assessments, board review, privacy clauses, pilot processes, annual reporting, and cross-functional coordination. Second, they focus on a small number of high-stakes tools, where risk and impact are greatest, rather than every AI feature or tool.* The point is not to build a new department or a new compliance layer. It is to aim familiar district processes more precisely at access, nondiscrimination, privacy, fair decision-making, and meaningful voice when AI enters rights-sensitive decisions.

1. Treat high-risk AI as a civil-rights exception, not the default

Why This Is Needed

AI is spreading fastest in rights-relevant domains like risk scoring, placement, and tutoring, especially in large, high-need districts, while governance protections lag. That combination raises the risk that AI will quietly lower expectations or restrict access to rigorous coursework for historically marginalized students.

New York City already [says AI cannot be used for high-stakes decisions](#) like grading or discipline. The recommendation takes that idea a step further: whenever a district wants to use AI in those kinds of high-stakes areas, it should be required to spell out, in writing, how that use is consistent with its civil-rights obligations. In practical terms, the district should be able to say not only why the tool is justified, but what improvement it expects to see, for whom, and on what timeline before renewal or expansion is considered.

Recommendations

- Classify AI used for risk scoring, placement (including for emergent bilingual students), discipline, and access to tutoring or interventions as rights-relevant systems requiring heightened review under the district's existing civil rights obligations.
- Require a civil-rights justification memo before adoption or expansion: why AI is needed instead of non-AI alternatives, what specific student entitlement or district problem it is meant to address, how success will be measured, how the district will verify that it does not worsen disparities, and who is the named accountable decision-maker for that use case.
- For high-stakes AI procurements, districts should consider a limited outcome-based contracting approach that ties continuation, renewal, or scale-up to clearly defined educational goals and documented performance against them. This does not replace civil-rights review; it reinforces it by requiring districts and vendors to specify in advance what necessity and success mean in practice.
- Require that every high-stakes AI use have one accountable district official attached to it, on the record, before the tool is activated. The accountable official may consult a committee, but the decision may not rest with a committee alone.
- Prohibit vendor-defined default uses in rights-relevant domains; the district must explicitly authorize each high-stakes use case through its own process.
- High-risk AI should be treated as an exception to be justified, and every such use should have a named human being standing behind it.

2. Make equity an enforceable condition for AI, not just a stated value

Why This Is Needed

Most surveyed districts explicitly name equity or fairness in their evaluation criteria, but far fewer prioritize AI investments for high-need students or use formal bias protocols. This is the governance-equity gap: strong language, thinner practice. Chicago Public Schools' AI Principles and Vision explicitly commit to equitable access to AI opportunities across students, educators, and families. CPS has already set out what equity with AI should look like; our proposal is that those principles become gatekeeping criteria for approving and renewing high-stakes AI tools, not just values statements that sit alongside the decision.

Recommendations

- Require an equity impact statement as part of the decision record for every high-stakes AI procurement and renewal, specifying which student groups are expected to benefit first, which groups carry the clearest risks, and how access to rigorous coursework and key services will be protected or expanded.
- Tie board approval to a small set of measurable equity commitments, for example, subgroup outcome measures, protections for advanced coursework access, or service access indicators that are reviewed at defined intervals, such as every six months for early implementation and again at renewal.
- If equity cannot be described concretely and monitored, the AI is not civil-rights ready. Put differently, districts should define the set points in advance: what will be measured, when it will be reviewed, and what result would trigger continuation, revision, or non-renewal.

3. Make bias and performance checks a minimum compliance duty

Why This Is Needed

About 70–80 percent of districts report strong privacy and data-governance clauses, but only about half have any protocol for assessing algorithmic bias, even though tools are used in high-stakes decisions. Capacity is part of the equity problem: the districts most exposed to these tools are often the least staffed to evaluate them independently, which increases the risk that vendor self-audits will substitute for public accountability. In some districts, this work will also require building a baseline that does not yet exist, because comparable information on placement, intervention, discipline, or service access may not be tracked consistently even without AI. Orange County Public Schools in Florida is working toward an [AI policy](#) that, among other guardrails, prohibits entering student data into AI programs and emphasizes the use of “closed” AI environments to better protect student information.

Districts like Palm Beach County Schools have already set [clear ‘no-go’ zones for AI](#), such as standardized assessments; our proposal is to pair those bright lines with a simple annual disparity check wherever AI is allowed to influence high stakes decisions, so districts can see if any student group is being flagged, placed, or served differently in ways that undermine their rights.

Recommendations

- Treat basic disparity and performance checks as a minimum compliance duty wherever AI affects placement, discipline, identification for services, or access to interventions.
- At a minimum, once a year, districts should document how the tool is used, report usage and key outcomes by race/ethnicity, income, disability, and language status, compare those patterns against an appropriate baseline where one exists, and flag large differences in who is labeled, flagged, placed, or served.
- Where no reliable baseline exists, the district should establish one as part of implementation by documenting pre-AI process patterns or the earliest period of comparable use, so later disparity and performance reviews have a meaningful point of comparison.
- Where districts lack internal capacity to conduct these checks, county offices of education, regional service agencies, or comparable intermediaries should be named as entities that can perform, pool, or support these reviews across districts.
- Embed this duty in vendor contracts: vendors must provide the data access, documentation, and technical cooperation needed for district or regional reviews. A

vendor that will not supply the information needed to run the check should be treated as failing a basic governance condition.

- Specify in policy that if disparity thresholds are exceeded, the district will suspend, narrow, or pause the AI use until the issue is addressed.
- To keep the work manageable, districts can focus initially on a small number of high-stakes AI tools with the greatest impact on student opportunities, rather than trying to monitor every AI-enabled feature at once.
- Bias checks are not a nice-to-have; they are the basic way a district demonstrates it is meeting nondiscrimination duties when AI is involved.

4. Use contracts to prevent AI lock-in and protect student rights

Why This Is Needed

A small set of vendors, many backed by private equity, have consolidated control over core platforms and often sell bundled, recurring contracts. The evidence base for their flagship products is mixed, and contract structures can make it hard for districts to unwind harmful tools. Districts often learn this too late: every protection they want has to be negotiated before the contract is signed, and after signature the district generally lives with the terms it accepted. This risk may be especially acute when long-standing district platforms begin layering in AI features, because institutional familiarity with the base product can make new rights-relevant capabilities seem like routine updates rather than material changes requiring review.

Recommendation

- Add a standard AI civil-rights addendum to contracts covering high-stakes uses, specifying access and nondiscrimination commitments, data minimization, strict limits on reuse, sharing, and model training, and transparency and appeal rights for students and families in AI-assisted decisions.
- State plainly that these protections must be negotiated before contract execution, not after implementation problems emerge.
- Prohibit automatic feature rollups in which new AI capabilities are activated inside existing platforms without explicit district review and board acknowledgment of rights implications.
- Districts should treat AI additions inside established platforms as a fresh governance event, not as a minor extension of an already approved tool.
- Include exit and suspension clauses that allow the district to suspend AI components if disparity checks, incident reviews, or community feedback raise serious concerns, without having to tear out the entire system.
- Treat exit rights as non-negotiable for high-stakes AI tools, on par with data privacy, because the ability to walk away is often the district's final line of protection when a tool harms a student group.

The Los Angeles County Office of Education is already guiding districts to vet vendors and monitor AI risks. Its' [AI Education Guidelines](#) that ask districts to consider vendor selection, risk assessment, and ongoing monitoring when they implement AI tools, with specific attention to data privacy and alignment with district goals. That is a county-level example of treating vendor choice and oversight as a governance issue, not just a purchasing detail. Our recommendation is to make this protection concrete in contracts by adding explicit civil-rights clauses and clear exit rights for high-stakes AI tools.

5. Make “who gets AI” an equity decision, not a convenience decision

Why This Is Needed

The data show AI spending is thickest in districts serving historically marginalized students, but equity-oriented implementation does not consistently track with higher need. There is a risk that AI follows where capacity is highest, not where stakes are highest. This is also where the report's vendor consolidation and financialization findings matter at the district level. In practice, some vendors repackage capabilities that district staff may already be able to access through lower-cost general-purpose AI tools, then sell them back to districts at enterprise prices by bundling them as turnkey solutions. Districts with experienced, curious staff are more likely to challenge those claims; districts without that internal capacity are more exposed to overpaying for convenience, especially in systems serving high-need students. That is not just a market inefficiency; it is a governance-equity gap with a budget consequence.

Recommendations

- Treat allocation of AI supports as an explicit equity decision: before purchasing, require a simple plan that answers “who first, and why,” with priority given to high-poverty schools and students systematically underserved by existing interventions.
- Require internal coverage, benefit, and alternatives reviews, not just uptake reviews: once a year, show which schools and student groups are using each AI support, what outcome trends look like across groups, and whether comparable lower-cost alternatives could meet the same need.
- Make expansion contingent on evidence that early deployments are closing, not widening, gaps in access to rigorous courses, credit accumulation, language access, or needed services.
- Encourage districts to build baseline staff's familiarity with widely available AI tools before procurement decisions are made, so someone in the room can test vendor's claims against existing consumer capability.

In this sense, where AI goes first and whether the district needed this product at all are both equity questions, not just purchasing questions.

6. Build AI governance that shifts power, not just structure

Why This Is Needed

Decision-making is centralized in many districts; roughly two thirds of districts report not involving teachers in AI procurement, with even thinner representation from students and families. At the same time, policy shifts are more often driven by new products than by mandates, suggesting vendor dynamics are shaping the agenda. Governance must also sit close enough to the real work of implementation to catch domain-specific nuance in high-stakes settings.

Recommendations

- Establish a standing AI governance council by expanding the remit and membership of an existing cross-functional committee where possible, rather than building a separate body from scratch.
- For especially sensitive domains, districts may create targeted review groups or subcommittees connected to the standing council—for example, for placement, multilingual learner services, special education, discipline, or tutoring access—so that governance remains close to the practice setting while still reporting through the district's formal accountability structure.
- Anchor the council in board-adopted policy that defines its charge, its scope, and the categories of AI proposals, incidents, and renewals that must move to the superintendent and board.
- Name the superintendent as the accountability point between the council's work and the board's oversight, so the council is situated inside, not besides, the district's formal governance structure.
- Define what reaches the board, and when, including thresholds for notification outside the regular reporting cycle when a serious AI incident or civil-rights concern emerges.
- Require representation from procurement, IT, curriculum, special education/ELL, legal or civil-rights staff, school leaders, and designated teacher, student, and family members from high-poverty and high-impact schools.
- Give the council formal authority to review high-stakes AI proposals and renewals, request disparity data and contract changes, recommend community engagement, and recommend suspension or narrowing of use when rights risks emerge. This allows districts to avoid a false choice between a single centralized body that is too distant from implementation and a set of isolated committees that lack formal authority.
- Governance should not just coordinate across offices; it should place AI review inside the district's actual chain of democratic accountability while staying close enough to the work to catch consequential nuance.

How to Read These Recommendations in Practice

These recommendations are meant to sharpen and amplify work that districts are already doing, not to build a new layer of administration from scratch. In most cases, the shift is in *how existing board memos, RFPs, contracts, committees, and reporting routines are used when AI is involved*, rather than in creating entirely new structures.

For civil-rights memos and equity impact statements, districts already prepare board memos, RFPs, and decision notes when they adopt new tools. The recommendation is to tighten what goes into those existing documents for AI-related purchases, for example by adding a short section that explains how a proposed tool aligns with access, nondiscrimination, language rights, and due process rather than inventing a new document type.

For bias and disparity checks, the bar is intentionally modest: a once-per-year check focused on a small number of high-stakes tools, not an audit of every AI-enabled feature in the district. A single-page summary showing who is flagged, placed, or served by these tools, broken out by major student groups, is enough to start acting on patterns.

For contracts, most districts already use standard data privacy and security language in their technology agreements. The recommendation is to add one short AI civil-rights addendum that specifies access, nondiscrimination, data-use limits, and exit rights for high-stakes AI, rather than rewriting contracts from the ground up.

For governance councils, many districts already have cross-functional committees for instructional technology, data governance, or equity. The recommendation is to make AI a standing agenda item with clear authority over high-stakes AI decisions and to ensure representation from teachers, students, and families most affected rather than creating an entirely separate governance body.

AI is already becoming part of how large districts make decisions about teaching, support, placement, and access, especially in systems serving students with the greatest historical exposure to educational inequity. The central challenge is no longer whether districts will encounter AI, but whether they will govern it in ways that protect students' civil rights in practice rather than only in principle. The findings in this report suggest that many districts have the beginnings of that infrastructure, but far fewer have built the routines, accountability mechanisms, and power-sharing structures needed to make equity a binding condition of everyday AI use. The recommendations are intended to help districts close that gap by translating existing commitments around access, nondiscrimination, privacy, and voice into concrete governance actions that can be used now.

Appendices

Research Design

Data Sources

This paper draws preliminary findings from Urban AI Unlocked: Transforming District Systems for Equitable AI Integration, a national project led by USC Rossier's EdPolicy Hub. The study combines (1) district-level purchasing data on AI-enabled instructional technologies, (2) a survey of administrators in large urban districts, (3) administrative data on district composition, (4) financial and market data on suppliers, and (5) case studies involving interviews and focus groups with district administrators, teachers, and families (the case studies are in progress and are not reported here).

First, we analyzed AI-related purchasing using transaction-level data from GovSpend, a national database of government procurement. We compiled purchase order records from U.S. educational agencies between the 2018–19 and 2024–25 school years and identified transactions for AI-enabled instructional technologies, doing so using a vendor classification system and a set of inclusion and exclusion search terms. We developed the system and search terms through a systematic review of product descriptions, company websites, and industry reports. We focused on local education agencies designated by NCES as regular school districts, as well as aggregate spending to the district level. We then matched GovSpend records to the NCES Common Core of Data (CCD) using NCES district IDs and manual name matching, which allowed us to merge information on enrollment size, grade span, student demographics, and locale.

Second, we fielded a survey of procurement officials and technology leaders in Council of the Great City Schools (CGCS) member districts. The survey was organized around six constructs aligned to our civil-rights-first frame: governance infrastructure; equity-oriented practices; access and nondiscrimination; data and governance; fairness in decision-making; and student, teacher, and school leader voice and agency. Respondents received information about the study purpose, anonymity protections, voluntary participation, and contact information for the principal investigator and university human subjects office. Data were stored on password-protected systems accessible only to the research team. In some districts, we received multiple responses, which we reconciled by using non-missing responses where available, taking the union of options for “select all that apply” items, and, when single-response items conflicted, adopting the most generous interpretation of district practice. The full paper provides a detailed rubric mapping the items to the constructs.

Third, we used administrative data from the NCES Elementary/Secondary Information System on race/ethnicity and National School Lunch Program eligibility for students enrolled in CGCS districts in 2023–24. We supplemented this administrative data with disability and English learner data for 2022–23. Excluding Puerto Rico from these analyses, we grouped districts into quantiles based on enrollment and key demographic indicators to support descriptive comparisons.

Fourth, the market share and financial analyses drew on a commercial proprietary database that synthesizes data from multiple underlying sources, including SEC filings, earnings call transcripts, press releases, credit and equity analyst reports, expert interviews, and industry news for major K–12 vendors selling AI products and services to school districts. This database aggregates quantitative indicators such as revenue estimates, margins, debt levels, alongside qualitative material on strategy, product mix, and risk, allowing triangulation of vendor market position and business models. The reports produced through this system also integrate expert calls to inform understanding of competitive dynamics, pricing power, and risk, among other datapoints. We treated figures and claims as synthesized secondary data and cross-checked them against the cited primary sources that they reference (for example, 10-Ks, press releases, and S&P reports).


Measures and Analytic Approach

To characterize district AI spending, we construct two per-pupil measures using aggregated contract amounts and 2023–24 enrollment—Total Expenditure and Population-Adjusted—to capture cumulative contracted investment in AI-enabled instructional tools over the study period relative to district size. Average Annual Per-Pupil Expenditure divides this total by seven to approximate average yearly investment intensity. These measures should be interpreted as contracted spending rather than realized annual disbursements and assume that district enrollment is relatively stable over the period.

We describe patterns of adoption and intensity among districts that have procured AI-enabled tools, including the timing of initial adoption, the number of years with AI purchases, and variation across states and district characteristics. To examine associations between district composition and AI-related spending, we estimate linear regression models with each expenditure measure as the outcome and predictors for student race/ethnicity, grade span, locale, and, in some specifications, the percentage of students eligible for free or reduced-price lunch. Secondary models add state fixed effects and cluster standard errors at the state level to account for unobserved, time-invariant differences across states.

For the survey, we used the civil-rights-first framework to construct indices for each of the six constructs. Items were scored 0 or 1 based on a rubric specifying what counts as a rights-advancing practice (for example, whether procurement processes require equity statements tied to specific student groups or whether contracts include explicit limits on data use and retention). Scores were summed within constructs to generate district-level indices. We reported distributions of these scores, along with selected item-level patterns, to illustrate how governance infrastructure and equity-oriented practices vary among large urban districts (See Appendix Tables 4–6).

Map 1. Percentage of School Districts in Each State That Have Purchased AI Tools



Map 2. Percentage of Students in Each State Enrolled in Districts That Have Purchased AI Tools





Figure 1. Variation in Contracted Expenditure by School District Composition

**Table 1. Key Findings**

Topic	What the report finds	Why it matters
Where AI is spreading	AI adoption is growing quickly in K-12 and is especially concentrated in large, high-need districts serving many students of color.	The districts with the greatest civil-rights obligations are also among those facing the most AI exposure.
Governance vs. practice	Many districts report governance structures such as needs assessments, privacy clauses, and equity language, but much weaker day-to-day equity practice.	Formal rules alone do not guarantee that AI decisions will protect access, fairness, or student rights in practice.
Vendor power	A small set of vendors, many backed by private equity, hold strong positions in district AI markets through bundled and recurring contracts.	Districts may face lock-in, weaker bargaining power, and fewer opportunities to exit tools that create harm.
High-stakes uses	AI is entering rights-sensitive areas such as risk	These are not low-stakes experiments; they can shape

	scoring, placement, language access, and tutoring.	students' opportunities, services, and treatment.
Equity risk	Equity is often named in policy language but less often enforced through routines such as bias checks, targeted investment, or community decision-making power.	This creates a governance-equity gap between what districts say and what they consistently do.

Table 2. What Districts are Doing Now

Current district practice	What the data suggest	Key implication
Formal needs assessments	About 70 percent of surveyed districts say they conduct formal needs assessments before AI-related purchases.	Districts already have a planning structure that can be strengthened rather than replaced.
Contract protections	More than 80 percent report including four or more privacy and data-governance protections in AI contracts.	Privacy clauses are becoming common, but they are not enough on their own.
Equity language	Roughly three-quarters say they explicitly mention equity or fairness in evaluation criteria.	The challenge is turning equity language into enforceable decision rules.
Targeting high-need students	Fewer than one in five districts systematically prioritize AI investments for high-need students.	Districts often lack a routine for making AI allocation an equity decision.

Bias review	Only about half have a formal or developing protocol for assessing algorithmic bias.	Many districts still do not treat bias checks as a basic compliance function
Shared governance	Teachers, students, and families are rarely given structured roles or formal power in procurement and implementation.	Those most affected by AI decisions often have the least influence over them.

Table 3. Six District Actions

Current district practice	What the data suggest	Key implication
Formal needs assessments	About 70 percent of surveyed districts say they conduct formal needs assessments before AI-related purchases.	Districts already have a planning structure that can be strengthened rather than replaced.
Contract protections	More than 80 percent report including four or more privacy and data-governance protections in AI contracts.	Privacy clauses are becoming common, but they are not enough on their own.
Equity language	Roughly three-quarters say they explicitly mention equity or fairness in evaluation criteria.	The challenge is turning equity language into enforceable decision rules.
Targeting high-need students	Fewer than one in five districts systematically prioritize AI investments for high-need students.	Districts often lack a routine for making AI allocation an equity decision.
Bias review	Only about half have a formal or developing protocol for assessing algorithmic bias.	Many districts still do not treat bias checks as a basic compliance function.

Shared governance	Teachers, students, and families are rarely given structured roles or formal power in procurement and implementation.	Those most affected by AI decisions often have the least influence over them.
-------------------	---	---

Governance Infrastructure and Equity-Oriented Practices Tables

Table 4. Score Distribution: Governance vs. Equity-Oriented Practices

Band (% of scale max)	GI raw scores	EOP raw scores	GI count	GI share	EOP count	EOP share
0 (0%)	0	0	0	0.0%	0	0.0%
>0%–25% of scale	1–2 (9%–18% of max)	1 (14% of max)	4	14.8%	7	25.9%
>25%–50% of scale	3–5 (27%–45% of max)	2–3 (29%–43% of max)	11	40.7%	9	33.3%
>50%–75% of scale	6–8 (55%–73% of max)	4–5 (57%–71% of max)	11	40.7%	7	25.9%
>75% of scale (top quarter)	9–11 (82%–100% of max)	6–7 (86%–100% of max)	1	3.7%	4	14.8%

Table 5. Upper-Half Share by District Characteristics

Upper-half cut points: GI = scores 6-11 on a 0-11 scale; EOP = scores 4-7 on a 0-7 scale.

District characteristic	Subgroup	GI upper-half share	EOP upper-half share
Total PK-12 Enrollment	Small (< 38,178)	25.0%	25.0%
Total PK-12 Enrollment	Medium (31,878-68,568)	66.7%	55.6%
Total PK-12 Enrollment	Large (> 68,568)	50.0%	50.0%
% Eligible for NSLP	Q1 (< 64%)	66.7%	55.6%
% Eligible for NSLP	Q2 (64%-76%)	20.0%	20.0%
% Eligible for NSLP	Q3 (76%-83%)	50.0%	66.7%
% Eligible for NSLP	Q4 (> 83%)	28.6%	14.3%
% Hispanic or Latinx	Q1 (< 17%)	50.0%	16.7%
% Hispanic or Latinx	Q2 (17%-32%)	57.1%	42.9%
% Hispanic or Latinx	Q3 (32%-49%)	33.3%	55.6%
% Hispanic or Latinx	Q4 (> 49%)	40.0%	40.0%
% Black or African-American	Q1 (< 16%)	60.0%	40.0%
% Black or African-American	Q2 (16% – 27%)	30.0%	40.0%
% Black or African-American	Q3 (27% – 49%)	66.7%	83.3%
% Black or African-American	Q4 (> 49%)	33.3%	0.0%
% White	Q1 (< 10%)	33.3%	33.3%
% White	Q2 (10% – 17%)	14.3%	42.9%
% White	Q3 (17% – 25%)	62.5%	37.5%
% White	Q4 (> 25%)	55.6%	44.4%

Table 3=6. GI and EOP Share by District Characteristics

District characteristic	Subgroup	GI upper-half share	EOP upper-half share
Total PK-12 Enrollment	Small (< 38,178)	36.4%	39.3%
Total PK-12 Enrollment	Medium (31,878-68,568)	50.5%	49.1%
Total PK-12 Enrollment	Large (> 68,568)	45.5%	52.4%
% Eligible for NSLP	Q1 (< 64%)	52.5%	57.1%
% Eligible for NSLP	Q2 (64%-76%)	32.7%	31.4%
% Eligible for NSLP	Q3 (76%-83%)	43.9%	54.7%
% Eligible for NSLP	Q4 (> 83%)	37.6%	32.7%
% Hispanic or Latinx	Q1 (< 17%)	43.9%	33.3%
% Hispanic or Latinx	Q2 (17%-32%)	44.2%	44.9%
% Hispanic or Latinx	Q3 (32%-49%)	40.4%	47.6%
% Hispanic or Latinx	Q4 (> 49%)	45.5%	57.1%
% Black or African-American	Q1 (< 16%)	54.5%	45.7%
% Black or African-American	Q2 (16% – 27%)	34.5%	45.7%
% Black or African-American	Q3 (27% – 49%)	51.5%	61.9%
% Black or African-American	Q4 (> 49%)	39.4%	28.6%
% White	Q1 (< 10%)	33.4%	42.9%
% White	Q2 (10% – 17%)	35.1%	44.9%
% White	Q3 (17% – 25%)	52.3%	48.3%
% White	Q4 (> 25%)	44.5%	44.4%

References

- Balfanz, R., Herzog, L., & Mac Iver, D. J. (2007). Preventing student disengagement and keeping students on the graduation path in urban middle-grades schools: Early identification and effective interventions. *Educational Psychologist*, 42(4), 223–235.
- Barocas, S., & Selbst, A. D. (2016). Big data’s disparate impact. *California Law Review*, 104(3), 671–732.
- Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning*. fairmlbook.org.
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the New Jim Code*. Polity.
- Bird, K. A., Castleman, B. L., & Song, Y. (2024). Are algorithms biased in education? Exploring racial bias in predicting community college student success. *Journal of Policy Analysis and Management*, 43(1), 220–253.
<https://doi.org/10.1002/pam.22569>
- Burch, P. (2009). *Hidden markets: The new education privatization*. Routledge.
- Callahan, R. (2005). Tracking and high school English learners: Limiting opportunity to learn. *American Educational Research Journal*, 42(2), 305–328.
- Center for American Progress. (2024). *Scaling up high-dosage tutoring is crucial to students’ academic success*. <https://www.americanprogress.org/article/scaling-up-high-dosage-tutoring-is-crucial-to-students-academic-success/>
- Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.
- García, O., & Kleifgen, J. A. (2018). *Educating emergent bilinguals: Policies, programs, and*

- practices for English learners* (2nd ed.). Teachers College Press.
- Kieffer, M. J., Lesaux, N. K., Rivera, M., & Francis, D. J. (2009). Accommodations for English language learners taking large-scale assessments: A meta-analysis on effectiveness and validity. *Review of Educational Research*, 79(3), 1168–1201.
- Kraft, M. A., & Falken, G. (2021). A review of the evidence on high-dosage Tutoring (EdWorkingPaper No. 21-384). Annenberg Institute at Brown University.
- Lau v. Nichols, 414 U.S. 563 (1974).
- Menken, K. (2008). *English learners left behind: Standardized testing as language policy*. Multilingual Matters.
- Morris, E. W., & Perry, B. L. (2017). Girls behaving badly? Race, gender, and subjective evaluation in the discipline of African American girls. *Sociology of Education*, 90(2), 127–148.
- Oakes, J. (2005). *Keeping track: How schools structure inequality* (2nd ed.). Yale University Press.
- Poell, T., Nieborg, D. B., & van Dijck, J. (2019). Platformisation of society and the transformation of public values. *Internet Policy Review*, 8(2), 1–16.
- Scott, J., & Jabbar, H. (2014). The hub and the spokes: Foundations, intermediary organizations, incentivist reforms, and the politics of research evidence. *Educational Policy*, 28(2), 233–257.
- <https://doi.org/10.1177/0895904813515327>

- Skiba, R. J., Michael, R. S., Nardo, A., & Peterson, R. (2002). The color of discipline: Sources of racial and gender disproportionality in school punishment. *The Urban Review*, 34(4), 317–342. <https://doi.org/10.1023/A:1021320817372>
- U.S. Department of Education, Office for Civil Rights. (2014). *Civil Rights Data Collection: Data snapshot—School discipline*.
- U.S. Department of Education, Office for Civil Rights. (2014). *Civil Rights Data Collection: Data snapshot—College and career readiness*.
- U.S. Department of Education, Office for Civil Rights. (2025). *Equal education opportunities for English learners*.
- U.S. Department of Education & U.S. Department of Justice. (2015). *Dear colleague letter: English learner students and limited English proficient parents*.
- Vandell, D. L., Reisner, E. R., & Pierce, K. M. (2007). *Outcomes linked to high-quality afterschool programs: Longitudinal findings from the study of promising afterschool programs*. Policy Studies Associates.
- Veale, M., & Edwards, L. (2018). Clarity, surprises, and further questions in the Article 29 Working Party draft guidance on automated decision-making and profiling. *Computer Law & Security Review*, 34(2), 398–404. <https://doi.org/10.1016/j.clsr.2017.12.002>
- Verger, A., Lubienski, C., & Steiner-Khamisi, G. (Eds.). (2016). *World yearbook of education 2016: The global education industry*. Routledge.
- Williamson, B. (2017). *Big data in education: The digital future of learning, policy and practice*. Sage.

Als for Education. (2026). *State AI guidance for education*.

<https://www.aiforeducation.io/ai-resources/state-ai-guidance>

SchoolDay. (2025, December 9). *Data governance in K–12: Building trust through*

transparency. <https://www.schoolday.com/data-governance-in-k-12-building-trust-through-transparency/>

U.S. Department of Education, Office for Civil Rights. (2024). *Fact sheet: Protecting the rights of English learner students with disabilities in public elementary and secondary schools*.