



ROCKWOOL Foundation Berlin

Institute for the Economy and the Future of Work (RFBerlin)

DISCUSSION PAPER SERIES

195/26

AI Diffusion Gaps: Unequal Integration of AI Across K-12 Schools

Christopher Campos, John Singleton

AI Diffusion Gaps: Unequal Integration of AI Across K-12 Schools

Authors

Christopher Campos, John Singleton

Reference

JEL Codes: O33, I21, I24, D23

Keywords: AI, education, diffusion, technology

Recommended Citation: Christopher Campos, John Singleton (2026): AI Diffusion Gaps: Unequal Integration of AI Across K-12 Schools. RFBerlin Discussion Paper No. 195/26

Access

Papers can be downloaded free of charge from the RFBerlin website: <https://www.rfberlin.com/discussion-papers>

Discussion Papers of RFBerlin are indexed on RePEc: <https://ideas.repec.org/s/crm/wpaper.html>

Disclaimer

Opinions and views expressed in this paper are those of the author(s) and not those of RFBerlin. Research disseminated in this discussion paper series may include views on policy, but RFBerlin takes no institutional policy positions. RFBerlin is an independent research institute.

RFBerlin Discussion Papers often represent preliminary or incomplete work and have not been peer-reviewed. Citation and use of research disseminated in this series should take into account the provisional nature of the work. Discussion papers are shared to encourage feedback and foster academic discussion.

All materials were provided by the authors, who are responsible for proper attribution and rights clearance. While every effort has been made to ensure proper attribution and accuracy, should any issues arise regarding authorship, citation, or rights, please contact RFBerlin to request a correction.

These materials may not be used for the development or training of artificial intelligence systems.

Imprint

RFBerlin
ROCKWOOL Foundation Berlin –
Institute for the Economy
and the Future of Work

Gormannstrasse 22, 10119 Berlin
Tel: +49 (0) 151 143 444 67
E-mail: info@rfberlin.com
Web: www.rfberlin.com



AI Diffusion Gaps: Unequal Integration of AI Across K-12 Schools*

Christopher Campos

University of Chicago & NBER

John Singleton

University of Rochester & NBER

June 2026

Abstract

Although use of generative AI tools has quickly become widespread in education settings, emerging evidence suggests that effects on learning will depend on how that use is supported and guided. This paper reports findings from an original national survey of K–12 school principals designed to measure institutional integration of AI in schools through policies, teacher training, guidance for student use, leadership engagement, and the availability of AI-enabled tools. We find that AI use has spread rapidly across schools, largely as a productivity aid. Students mainly use AI for homework help and writing, while educators primarily use it for lesson planning and administrative tasks. The development of teacher training, guidance, and school policies has lagged adoption. We next document two *diffusion gaps* across schools: First, lower AI integration is associated with a higher share of disadvantaged students (a one standard deviation increase in disadvantage is associated with a $0.07\text{--}0.11\sigma$ lower score on an index of AI integration); Second, private and charter schools score $0.23\text{--}0.44\sigma$ lower on the AI integration index than traditional public schools. Although several surveyed school-level factors strongly predict AI integration, they do little to explain these gaps. Differences in district size account for roughly one-third of the disadvantage gap between public schools. These findings suggest that the factors associated with greater AI integration differ from those needed to narrow disparities in how schools support and guide AI use.

*Campos: Booth School of Business, University of Chicago; Christopher.Campos@chicagobooth.edu; Singleton: Department of Economics, University of Rochester; John.Singleton@rochester.edu. Connor Fogal provided outstanding research assistance. Campos acknowledges support from the John E. Jeuck Faculty Fellowship.

1 Introduction

While generative AI tools have spread rapidly through K–12 schools, their educational consequences remain uncertain and contested. Like other general-purpose technologies, AI’s value and impacts are likely to depend on complementary investments in skills, routines, governance, and organizational practices (Bresnahan and Trajtenberg, 1995, Bresnahan et al., 2002). In schools, such complements include guidance for student use, teacher training, written policies, AI-enabled tutoring and instructional tools, and changes to how instruction is organized and delivered. This distinction between use and integration matters because emerging evidence suggests that unguided AI use can undermine learning (Lehmann et al. 2024, Bastani et al. 2025, Lira et al. 2025). As a result, AI’s effects on teaching and learning—whether beneficial or harmful—will depend on whether and how schools integrate the technology and how widely integration diffuses across schools.

This paper studies AI integration in K-12 schools using original survey data from school principals nationwide. Principals provide a unique perspective on how AI is used across the full K-12 landscape, including traditional public, charter, and private schools. To measure integration, our survey captures whether and how AI is used and supported, including written policies, attitudes toward student use, AI-enabled tools, professional development, and principals’ beliefs about effects on learning and achievement gaps. Guided by a simple framework in which AI integration reflects complementary investments, we organize the empirical analysis around four themes: the extent to which integration lags adoption; how students and educators use AI; how integration varies across schools; and which factors predict integration levels versus explain integration gaps.

We first document that AI adoption in K-12 schools has risen rapidly to near-universal levels. 90% of principals report that teachers in their school were using AI as of June 2025, compared to about 20% just two years earlier, consistent with prior surveys documenting adoption in the K-12 setting (Diliberti et al., 2024, Kaufman et al., 2025). We further show that this diffusion has occurred in a decentralized, bottom-up manner, rather than as the result of coordinated policy initiatives. Schools have largely set their own approaches, overwhelmingly allowing AI use with guidelines rather than banning it. Last, we find that use by teachers has outpaced professional development and training on AI tools, an initial piece of evidence demonstrating that institutional integration is not equivalent to passive adoption.

We next examine how AI is used by students and educators. Students primarily use AI for homework help and writing, with little evidence that it substitutes for classroom instruction. Although concerns about inappropriate use are real, their incidence is limited; only half of principals report one or more instances of student misuse. Views on effectiveness are mixed: 30% of principals report that AI has not improved student learning, while 25% believe it already has. Teachers and principals use AI primarily as a productivity tool—such as for lesson planning, instructional support, and administrative tasks—consistent with evidence that AI can raise productivity in writing-intensive and knowledge-work tasks (Dell’Acqua et al., 2023, Noy and Zhang, 2023). Overall, AI use is widespread but remains largely supplemental to, rather than embedded in, instructional practices.

We then compare AI integration across schools. To do so, we construct a composite index,

motivated by the idea of complementary investments, that incorporates multiple dimensions of integration, including policies, educator use, encouragement of student use, future investment plans, professional development, and AI-enabled tutoring. Using this index, we document two robust *diffusion gaps* in integration across schools: first, schools serving more economically disadvantaged students have systematically lower levels of AI integration. A one standard deviation increase in disadvantage is associated with a $0.07\text{--}0.11\sigma$ reduction in integration. This pattern holds across alternative measures of disadvantage, including test scores, and is robust to a host of controls. Second, we find that charter and private schools lag traditional public schools by substantial margins, with adjusted gaps of roughly 0.33σ for charter schools and 0.39σ for private schools. We leverage randomized monetary incentives to test for bias from selection into responding and find that these gaps are stable across incentive groups, suggesting that our results generalize to the broader population of schools.

We then distinguish predictors of integration levels from mediators of integration gaps. Principals’ positive beliefs about AI, teachers who champion the technology, budget availability, and targeted Elementary and Secondary School Emergency Relief (ESSER) use strongly predict integration levels, but these factors do not account for the socioeconomic or sectoral gaps.¹ This distinction is central to the framework. A factor can raise integration within schools without explaining why disadvantaged, charter, or private schools lag if that factor is not distributed differently across groups in the relevant direction. Instead, we find that district scale—rather than per-pupil or administrative spending—explains roughly 36% of the disadvantage gap, pointing to system-level implementation capacity as an important factor.

Our findings have three implications. First, although AI use has taken off rapidly in schools, passive adoption has advanced faster than the investment needed for the technology to be effectively supported. School adoption is widespread for supplemental uses, such as homework help and administrative tasks, but policies, training, guidance, and AI-conscious instructional routines remain less developed. Second, the relevant inequality is not access to AI tools, but schools’ capacity to structure, supervise, and govern their use. Schools serving more disadvantaged students, as well as charter and private schools, have lower levels of integration. Third, because predictors of integration levels do not mediate integration gaps, policies that raise AI integration may have limited impact on reducing across school disparities. While evidence on AI’s educational effects continues to emerge, differences in integration will contribute to heterogeneity in those effects, making them an important consideration for research and policy moving forward.

Our paper contributes to the literature measuring generative AI diffusion. We apply the logic of general-purpose technologies, whose impacts depend on downstream complementary investment (Bresnahan and Trajtenberg, 1995, Bresnahan et al., 2002, Brynjolfsson et al., 2021), to study AI in K–12 schools. This framework draws a distinction between adoption—which prior work has shown has grown quickly among educators and students (e.g. Diliberti et al. 2024, Kaufman et al.

¹We find some suggestive evidence that the presence of strict cell-phone restrictions is predictive of lower teacher-based integration, but not too predictive of other components of AI integration. As other factors, it is unable to explain the set of diffusion gaps we document.

2025, Gallup and Walton Family Foundation 2025)—and AI integration, which our original survey measures through questions about policies, training, guidance, tools, and routines. The persistent diffusion gaps we document connect with other work showing unequal adoption of AI across workers (Humlum and Vestergaard, 2025, Bick et al., 2026, Bloom and Makridis, 2026).

Our paper also connects to recent work on how generative AI affects student learning. Existing evidence shows mixed effects of AI-assisted learning, with impacts that are sensitive to how the technology is supported and incorporated into instruction (Lehmann et al., 2024, Bastani et al., 2025, Contractor and Reyes, 2025, Lira et al., 2025). Our framework implies that such heterogeneity is to be expected. This perspective motivates measuring AI integration across schools because widespread but uneven and weakly supported AI use may widen educational disparities. Our findings suggest that narrowing those gaps will likely require policies that expand system-level implementation support.

2 Conceptual Framework

In this section, we present a simple framework that motivates our survey and helps organize the empirical analysis. The framework interprets generative AI as a general purpose technology whose educational consequences depend on complementary investment by schools and districts (Bresnahan and Trajtenberg, 1995). Schools are downstream units through which AI diffuses and the key complementary investment in the framework is integration. AI integration stands in for policies, teacher training, guidance for student use, leadership engagement, AI-enabled tools, and changes to instructional practice.

2.1 Adoption versus Integration

Let z denote the exogenous quality of AI. School j in district or region g first decides whether to adopt AI, $A_j \in \{0, 1\}$, and, conditional on adoption, how intensively to integrate it into school practice, $I_j = I(x_j; S_g)$, which in turn depends on school inputs, $x_j = (x_{1j}, \dots, x_{Kj})$, that include but are not limited to teacher training, school policy, guidance for student use, and district scale, S_g . This distinction matters because schools may obtain standalone value from low-friction AI use, denoted $\Lambda_j(z)$, such as lesson planning or administrative tasks.

The benefits of deeper integration are $B_j(I_j(x_j); z)$, where x_j captures school choices and inputs that we can measure in a survey. Integration is costly, $C_j(x_j; w_j, q_j, S_{g(j)})$, where w_j is the shadow cost of teacher time, q_j can represent school managerial capacity among other factors that affect relative costs of integration, and district scale can effectively lower school-level integration costs through scale externalities S_g if, for example, policy is more streamlined.

The school’s payoff depends on whether it adopts AI and, conditional on adoption, whether it invests in the complementary inputs that generate integration,

$$V(A_j, x_j) = A_j \left[\Lambda_j(z) + B_j(I_j(x_j); z) - C_j(x_j; w_j, q_j, S_{g(j)}) \right],$$

where $\Lambda_j(z)$ captures the standalone value of AI use absent school-level integration. Let x_j^0 denote the no-integration input bundle, normalized so that $I_j(x_j^0) = 0$, $C_j(x_j^0; w_j, q_j, S_{g(j)}) = 0$, and $B_j(I(x_j); z) = 0$.

This formulation separates adoption from integration. A school may adopt AI without meaningfully integrating it if the standalone value of adoption is positive, $\Lambda_j(z) > 0$, but the marginal return to investing in each integration input is below its marginal cost at the no-integration bundle:

$$\frac{\partial B_j}{\partial I_j} \frac{\partial I_j}{\partial x_{kj}} < \frac{\partial C_j}{\partial x_{kj}} \quad \text{at } x_j = x_j^0 \quad \text{for all } k.$$

In this case, AI use diffuses into the school, but the organizational complements needed to turn use into sustained practice do not. Otherwise, the school chooses integration level $I_j^* = I(x_j^*)$ such that

$$\frac{\partial B_j}{\partial I_j} \frac{\partial I_j}{\partial x_{kj}} = \frac{\partial C_j}{\partial x_{kj}} \quad \text{for all } x_{kj}, \tag{1}$$

or in other words, where the marginal integration benefit of each input is equal to the marginal cost. Thus, it is clear that differences in marginal benefits and marginal costs will drive heterogeneity in integration among schools that adopt.

The central complementarity in the framework is between the quality of the AI and school-level integration. Following Bresnahan and Trajtenberg (1995), the value of improvements in the general-purpose technology is larger when downstream units have made complementary investments. In our setting, this means

$$\frac{\partial^2 B_j(I_j; z)}{\partial I_j \partial z} > 0. \tag{2}$$

This suggests that the marginal value of AI quality is increasing in integration, or equivalently, the return to integration is higher when the underlying AI technology is more capable. Thus, AI's consequences depend not only on access, but on complementary school practices.²

The framework also clarifies why adoption and integration are distinct margins. Adoption can become widespread when AI provides standalone value through low-friction uses that require little organizational change. Integration, in contrast, requires complementary investments in schools. Because the costs and returns to these investments vary across schools and system support, S_g also varies, widespread adoption need not imply uniform integration. The complementarity between AI quality and integration then explains why this distinction matters: improvements in the technology are more valuable in schools that have made the complementary investments needed to use AI.

2.2 Diffusion Gaps

The framework also clarifies how to interpret integration differences. Let D_j indicate membership in a group of interest, such as high-disadvantage, charter, or private schools. The integration gap

²A longer-run implication is that rising adoption and integration may change input demand within schools, including demand for teachers and administrators with skills that are complementary to AI-enabled instruction and management (Bresnahan et al., 2002).

is

$$\Delta_D I^* = E[I_j^* \mid D_j = 1] - E[I_j^* \mid D_j = 0].$$

Because I_j^* is determined by Equation 1, differences in I_j^* can arise from differences in either marginal benefits or costs. A first order approximation gives

$$\Delta_D I^* \approx \kappa [\Delta_D MB - \Delta_D MC],$$

where $\kappa > 0$ depends on the curvature of net marginal benefits, $\Delta_D MB$ denotes group differences in perceived marginal benefits, and $\Delta_D MC$ denotes group differences in marginal implementation costs. Thus, a factor may predict integration by either shifting marginal benefits or costs, but it need not explain integration gaps if it does not vary between groups significantly.

2.3 Connections to Policy Evaluation

The complementary nature of integration is a theoretically important driver of policy treatment effect heterogeneity on student outcomes, potentially contributing to mixed findings of AI policies in education. Let P denote an AI-related policy or intervention. School-level outcomes depend on integration as input choices contribute to the education production function (Todd and Wolpin, 2003), so then policy effects can be written:

$$\tau_j(P, I_j) \equiv Y_j(P, I_j) - Y_j(0, I_j),$$

where $Y_j(P, I_j)$ denotes the outcome produced when policy P is implemented in a school with integration level I_j . If $\partial \tau_j(P, I_j) / \partial I_j \neq 0$, then the same AI policy can represent different effective treatments across schools. Therefore, differences in estimated policy effects across studies may reflect differences in integration rather than differences in the effectiveness of the policies.

Bastani et al. (2025) illustrates this point. Students were randomized to different forms of access to generative AI-assisted learning, with one arm providing access to a GPT-based tutor, while another added guardrails designed to discourage substitution away from learning. In our framework, this is a contrast between access with $I_j = 0$ and access bundled with an instructional complement, $I_j > 0$. The guardrails change the treatment effect by shaping how students use the technology.

This distinction matters because AI interventions are often treated as portable objects, even though their effects may depend on local complements such as guidance, training, governance, and instructional routines. Evidence from high-integration settings may not translate to low-integration settings, and effects from simple access may not predict effects under more structured use.

2.4 Empirical Roadmap

Our framework organizes the empirical analysis into four exercises. First, we document the diffusion process, whether AI use has become widespread and whether institutional complements such as

professional development have kept pace. Second, we characterize how students and educators use AI, distinguishing low-friction productivity and supplemental uses from deeper instructional integration. Third, we construct an AI integration index as an empirical proxy for I_j^* and assess factors that predict integration and document disadvantage and sector gaps. Last, we assess if school-level factors (x_j) or system-level factors ($S_{g(j)}$) are more important determinants of unequal integration.

3 Survey Design and Data

We survey school principals to obtain a comprehensive picture of generative AI integration in K-12 schools in the US. Survey modules aim to measure both elements that contribute to integration I_j^* —principal use, policy and governance, professional development, AI-enabled tutoring—and other factors in x_j —student, teacher, perceived drivers and barriers of AI use, and principal beliefs about AI’s effects. Table A.1 summarizes the survey’s full contents. The survey required around 15 minutes to complete and was administered to school principals between December 2025 and February 2026.

Respondents were drawn from a sample of over 70,000 K-12 schools—public and private—constructed from the NCES Common Core of Data (CCD) and Private School Survey (PSS) universes. Principals were contacted by email and offered compensation for their time through a combination of guaranteed payments and a raffle for several \$600 gift cards. Our final sample contains 1,254 principals representing 1,085 public schools (including 87 charter schools), 143 private schools, and 26 schools of other or unclassifiable type.

The survey incorporates two design features. First, invitations were uniquely linked to principals, allowing us to merge responses with school-level demographic and neighborhood data and to construct nationally representative estimates. Second, we randomized guaranteed monetary incentives for survey completion. This randomization generates exogenous variation in responding, which we use to assess the extent to which our estimates may be affected by non-response bias.

3.1 Sample Construction and Representativeness

We merge the survey with data from the National Longitudinal School Database (NLSD) (Carroll et al., 2023), test score information from the Stanford Education Data Archive (SEDA), and Census block group characteristics.³ Because information about subsidized lunch eligibility and performance are not available for private schools, we rely on block group household income to capture disadvantage. Importantly, linked district size and fiscal data proxy for system-level support $S_{g(j)}$, including scale and spending capacity.

Appendix Table A.2 summarizes the characteristics of schools served by respondents and compares them with those of non-respondents. Respondents and non-respondents differ modestly on

³For schools not linked to NLSD, including many charter and private schools, we filled missing covariates using CCD enrollment and lunch program files (2023–24), PSS enrollment data (2021–22), or, where necessary, by imputing means among observed schools in the same LEA, same county, or full invitation frame.

observables. To produce nationally representative estimates, we use weights throughout our analyses that adjust for these differences. Appendix Table A.2 confirms that the weighted respondent means closely approximate the full sample means across all variables.⁴

The randomized response incentives will provide evidence that non-response is unlikely to bias these estimates. As Appendix Table A.3 reports, the incentives increase response rates by 35–47% relative to the no-incentive control group. The key idea is that the incentives draw in *marginal* respondents—principals who would not otherwise have participated—allowing for testing how our key conclusions change as the respondent pool broadens to include more reluctant participants. We return to this idea in Section 6.3 when assessing robustness of findings.


4 Rapid and Decentralized Diffusion

We begin by examining survey responses to characterize the speed and nature of AI adoption in schools. We asked principals whether teachers were using generative AI tools as of June in each of the past four years (2022–2025). Figure 1 plots the share reporting teacher use over time: just 6% report use in June 2022, rising to 90%, indicative of virtually universal adoption, by June 2025. This is consistent evidence that AI diffusion has been rapid by historical standards (Kalyani et al., 2025). We also ask principals whether staff had received any professional development training on generative AI at the same time points. Figure 1 likewise shows a dramatic increase in such training (to about 64% of schools, as of June 2025). Figure 1 therefore documents adoption outpacing one key institutional complement. We return to broader integration in Section 6.

That use by teachers leads formal training suggests a decentralized, bottom-up process of diffusion. This is consistent with other survey findings: First, AI adoption does not appear to be driven by institutional mechanisms: most principals (58%) report that a written policy governing generative AI is not currently in place (Appendix Figure B.4); most schools (76%) do not have a license or contract for AI-powered tutoring (Appendix Figure B.7); and most schools (67%) did not use pandemic-era ESSER funds for AI purchases (Appendix Figure B.3). Second, principals’ reported stance toward student AI use is largely permissive: nearly half report actively encouraging use within rules (48%), while others discourage but stop short of banning (46%) (Appendix Figure B.8).

⁴These weights are formed from a logit model, estimated on the full sample with percent economically disadvantaged, block group poverty rate, median household income, block group shares White, Black, and Asian, and a private school indicator. The weights are computed as the reciprocal of the in-sample school’s predicted propensity, trimmed at the 1st and 99th percentiles and normalized to sum to the number of respondents.

Figure 1: AI Adoption and Professional Development over Time



Notes: This figure reports teacher AI adoption and teacher AI professional development trends for 2022–2025. Panel (a) reports the share of principals who replied “yes” to whether “Teachers at my school were using Generative AI tools for teaching-related tasks on or before” June 2022, June 2023, June 2024, and June 2025. Panel (b) reports the share responding “yes” to whether “Staff at my school had received any professional development training on generative AI” by those same June dates. Both panels stratify the series by economic disadvantage, which is based on the free/reduced-price lunch share of students split at the median. All means are IPW-weighted adoption shares by year. The stars reported above each pair of points correspond to the significance level of a difference in means test between high- and low-disadvantage school rates in a given year: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Figure 1 plots the time series separately for schools above and below the median in economic disadvantage, showing that teacher adoption and professional development are consistently higher in low disadvantage schools.⁵ This suggests a persistent diffusion gap across schools, a question we investigate in Section 6 by constructing an index of school-level AI integration.

5 Student and Educator Use

We next examine how students and educators are using AI. This margin matters because emerging evidence suggests that generative AI can support learning when it provides structured assistance, but can undermine learning when students rely on it without guardrails (e.g. Bastani et al. 2025). The key descriptive question is therefore whether AI use in schools looks like guided instructional integration, informal supplemental use, or replacement of traditional instruction. We find that student use is concentrated in homework and writing, while educators use AI primarily as a productivity aid. Overall, the patterns point to broad but shallow integration that is only weakly embedded in instructional practices.

⁵Appendix Figure B.2 instead stratifies the series by type of school, documenting gaps across sectors.

Figure 2: Student and Educator Use of AI



Notes: This figure reports principal responses about generative AI use by students and educators. Panel (a) reports principal responses to the question: “When AI tools are used in class, how frequently do they replace direct whole-class instruction?” Panel (b) reports principal responses to the question: “What percent of students at your school do you think use Generative AI for the following tasks?” Tasks include homework help and worked examples; essay drafting and writing feedback; brainstorming and project ideation; personalized tutoring; and study guides, quizzes, and flashcards. Panel (c) reports principal responses to the question: “What percent of teachers at your school do you think use Generative AI for the following tasks?” Tasks include any use, lesson planning, grading, detecting academic-integrity violations, assessment and feedback, and administrative tasks. Panel (d) reports the share of principals answering Yes to the question: “Do you use Generative AI for any of the following tasks?” Tasks include offloading administrative busywork; streamlining data analysis and reporting; designing professional development; and prepping meetings and presentations. Estimates are computed among principals with non-missing responses to the relevant item and are IPW-weighted to be nationally representative.

5.1 Student Use

The vast majority (88%) of principals report that AI “never” or “very rarely” replaces traditional classroom instruction (Figure 2, Panel (a)) and describe student use of AI tools as primarily supporting out-of-class work. Figure 2, Panel (b), shows the five most frequently cited student tasks are homework help (37%), essay drafting (35%), brainstorming (31%), making study guides or

practice quizzes (30%), and personalized tutoring or explanation (25%).

Principals’ perceptions of AI’s effects on student learning are mixed. While most principals agreed that AI has no effect on within-school achievement gaps (55%), there was disagreement about its effects on widening of achievement gaps (Appendix Figure B.9). Another 25% of principals believe that generative AI tools have already produced improvements in student learning at their school. But 31% believe it probably or definitely has not (Appendix Figure B.10).

Last, we asked principals about student misuse. Although AI policies, where they exist, nearly universally cover issues of academic honesty and plagiarism (see Appendix Figure B.5), disciplinary actions appear limited: nearly half of all principals report that no disciplinary actions have taken place (Appendix Figure B.6).

5.2 Educator Use

For teachers, the most common tasks reported by principals are administrative tasks (40%), lesson planning (37%), and grading and assessment (28%), with student integrity monitoring and grading rounding out the list (Figure 2, Panel (c)). Some use cases that have received substantial public attention—automated grading and AI-driven differentiation—are not the most common.

Principals report using AI tools in ways that enhance their productivity. 84% of respondents report using AI for at least one of: offloading administrative busywork, streamlining data analysis and reporting, designing professional development, and prepping meetings and presentations. The most common use tasks are administrative busywork (72%), meeting preparation (71%), with data analysis and designing professional development closely following. Appendix Figure E.4 shows that male principals were more likely to be early adopters of AI tools, but that this gap has since closed. In contrast, there was initially no racial gap, while current reported use is significantly higher among white principals.

Among AI tools and platforms, ChatGPT is by far the dominant, used by about 90% of teachers and principals. Around half of each group also uses Google Gemini. There is a long tail of other single-digit alternatives that includes Microsoft Copilot, Claude, and education-specific tools such as MagicSchool and SchoolAI (Appendix Figures B.11 and B.12).

Taken together, educator use resembles the knowledge-work tasks where generative AI has been shown to reduce time costs and improve output quality, including writing, planning, analysis, and communication (Noy and Zhang, 2023, Dell’Acqua et al., 2023). Our survey does not directly estimate productivity effects, but the task mix suggests that AI is currently being used far more to augment educator workflow than to automate instruction or deliver personalized learning. In terms of the framework, these patterns are consistent with schools realizing the standalone value $\Lambda_j(z)$ from low-friction AI use, while deeper integration I_j^* remains limited.

6 AI Diffusion Gaps

Although AI use is widespread, we next examine differences in AI integration: the policies, guidance, educator routines, training, and tools that translate AI use into sustained practice. We document AI diffusion gaps across K-12 schools along two distinct dimensions, economic disadvantage and school type. The economic disadvantage gap reveals that schools with higher shares of economically disadvantaged students are integrating AI at lower levels. The sectoral gaps reveal that both charter and private schools are integrating AI at lower rates than traditional public schools. We show that these gaps are robust to accounting for a variety of factors and that the headline results are not biased by non-response.

6.1 The AI Integration Index

We construct an AI integration index as our empirical proxy for I_j^* . The index combines seven standardized pillars: written AI policy and policy depth, encouragement of student use, teacher adoption, principal-use breadth, expected future integration, AI-tutoring availability, and AI professional development. Following Kling et al. (2007), we standardize each pillar using inverse-probability weights, average the standardized pillars, and re-standardize the aggregate to mean zero and standard deviation one. Appendix A.1 provides the full coding rules.

An important question is whether this index captures conditions that principals perceive as making AI more educationally valuable. To assess perceived complementarities (Equation 2), we examine how principals’ beliefs about AI’s impact on student learning vary with the index. Appendix Figure A.2 shows that principals in more integrated schools are substantially more likely to report positive expected effects of AI on learning: a one standard deviation increase in the AI integration index is associated with an 11.7 percentage point increase in the probability of reporting positive impacts. This relationship is not driven by any single component of the index; positive beliefs are especially strongly associated with student use guidance, principal use breadth, expected future integration, and availability of AI-enabled tutoring. These patterns are consistent with principals perceiving AI as more effective in schools where complementary policies, practices, and tools are more developed.

6.2 Findings

We begin by relating the AI integration index to the free and reduced-price lunch (FRPL) share across public (and charter) schools. Panel (a) of Figure 3 overlays a regression fit over a binned scatterplot of the underlying data, showing a clear socioeconomic gradient: public and charter schools serving more disadvantaged students exhibit lower levels of AI integration. Panels (b) and (c) investigate the sensitivity of this relationship to how disadvantage is measured: panel (b) replaces FRPL share with average test scores while panel (c) instead uses log median household income (which is also available for private schools). Greater disadvantage continues to be associated with lower AI integration. Panel (d) reports slope coefficients after standardizing; across all measures, a

one standard deviation increase in economic disadvantage is associated with 0.07-0.11 σ lower index value.⁶ In Appendix Table C.1, we examine how these results change when controls are introduced. Notably, the disadvantage diffusion gap is not attenuated—and is actually stronger—after accounting for school racial demographics and state effects.


We next turn to examining how AI integration varies across sectors—public, charter, and private. Figure 4 reports public-charter and public-private differences in the AI integration. We find that integration in both charter schools and private schools lags integration in public schools. Charter schools score 0.23–0.33 σ lower than traditional public schools on the AI integration index; private schools score 0.39–0.44 σ lower than (non-charter) public schools on the AI integration index. As in the case of the disadvantage gap, these sector gaps remain after accounting for state effects, neighborhood income, and racial demographics. Appendix Table C.2 reports the corresponding regression output. These sector gaps are notable given that charter and private schools’ autonomy could in principle accelerate integration; their lower integration suggests autonomy may not substitute for district infrastructure and vendor reach, though it may also reflect different institutional priorities. The public-private gap also echoes earlier evidence that private schools lagged in the diffusion of older school technologies such as advanced telecommunications (Parsad et al., 2001).

6.3 Robustness

Although the statistics we report are weighted for representation, there may be imbalance between sample respondents and non-respondents on unobservable attributes and these differences may affect our conclusions about the diffusion gaps across K–12 schools. We use the randomized response incentives to assess whether the headline gaps are driven by selection into the survey. Appendix Table A.3 shows that the incentives substantially increased response rates, allowing us to compare estimates across increasingly marginal respondents. Appendix Figure C.2 reports the disadvantage, charter, and private-school gaps by incentive group; the estimates are stable, and we cannot reject equality of the coefficients across groups. Appendix Figure C.3 shows the same pattern using alternative disadvantage measures. Appendix Figure E.1 reports qualitatively similar selection-corrected results. Appendix Figures C.4–C.6 show that the gaps are robust to alternative integration measures, including a PCA-weighted index, leave-one-component-out indices, and current-use and institutional-support subindices.


⁶In Appendix Figure C.1, we examine which index components especially drive the gradient.

Figure 3: AI Diffusion Gap by Disadvantage



Notes: This figure reports evidence of the AI diffusion gap using alternative measures of school disadvantage. The outcome is a standardized AI integration index that aggregates seven dimensions of school-level AI integration: written AI policy; stance toward student AI use; teacher AI adoption; principal AI use; expected future integration; AI tutoring availability; and AI professional development. Panels (a)-(c) plot binned scatterplots relating the index to free/reduced-price lunch (FRPL) shares, SEDA achievement, and log neighborhood median income, respectively. Panels (a) and (b) are restricted to public schools because FRPL and SEDA variables are unavailable for private schools; Panel (c) includes both public and private schools. Panel (d) reports standardized regression slope coefficients and 95 percent confidence intervals, where each disadvantage measure is normalized to mean zero and standard deviation one. All regressions are IPW-weighted and use heteroskedasticity-robust standard errors.

Figure 4: AI Diffusion Gap by Sector



Notes: This figure reports charter and private school gaps in the standardized AI integration index relative to traditional public schools. The AI integration index is the seven-dimension measure. Each bar reports the coefficient on a charter school or private school indicator from a separate regression of the standardized AI integration index on the sector indicator. Unadjusted bars (maroon) report unconditional differences. Adjusted bars (grey) additionally control for log neighborhood median income, neighborhood Black, Hispanic, and Asian shares, missingness indicators for each control, and state fixed effects. Charter school regressions are estimated on the public school sample (traditional public versus charter schools), while private school regressions are estimated on the non-charter sample (traditional public versus private schools). All regressions are IPW-weighted. Whiskers denote 95 percent confidence intervals based on heteroskedasticity-robust standard errors.

7 Predictors of AI Integration and Mediators of Diffusion Gaps

We now use our survey and auxiliary data to distinguish between factors that predict integration levels and factors associated with diffusion gaps. In Section 2, observed school-level factors x_j may raise integration by increasing the marginal return to complementary investment or lowering its cost, but these factors need not mediate any differences in integration across groups. Therefore, Panel (a) of Figure 5 asks which observed factors predict the integration index, while Panel (b) asks whether conditioning on those factors attenuates the disadvantage and sector gaps.

We organize the observed factors into four commonly emphasized channels: beliefs, leadership, budgets, and barriers. Beliefs capture principals' perceived value of AI and concerns about equity effects. Leadership captures principal discretion and teacher champions. Budgets capture available

and flexible resources, including targeted ESSER use. Barriers capture reported frictions such as cost, teacher skill gaps, student misuse, infrastructure limits, parent concerns, and policy restrictions. We also examine district scale and fiscal capacity as proxies for system-level support $S_{g(j)}$ and assess the role of cell-phone restriction policies.


7.1 What Factors Predict AI Integration?

We begin the mechanisms analysis by examining associations between AI integration and surveyed factors. For each of ten surveyed factors, spread across the four channels discussed above, we regress the AI integration index on the predictors while controlling for disadvantage, demographic missingness, and state fixed effects. Because the district/system capacity measures are available only for traditional public schools, Panel (a) of Figure 5 focuses on the common traditional-public-school sample. Across all factors, principals’ positive perceptions about AI’s impact is by far the strongest predictor of AI integration. Principals that agree that AI is likely to have a positive impact score 0.5σ higher on the AI integration index. Schools with teacher champions, a factor related to leadership, also score roughly 0.2σ higher on the index. Principal discretion over AI purchasing and adoption and budget availability (determined as having at least \$5,000 available for AI purchases) also strongly predict AI integration. District scale is positively associated with AI integration, though imprecisely estimated, while log current spending per pupil and log administrative spending per pupil do not strongly predict integration.

We also find that targeted funding appears more informative than general spending. Roughly 13% of schools report using ESSER funds for AI-related purchases and these schools score about 0.25σ higher on the integration index. This pattern is consistent with targeted resources helping schools convert general capacity into AI-specific complements, even though overall district spending and administrative spending are not strongly predictive. Last, we find that a standardized cell-phone restriction index, included as a proxy for schools’ broader technology environment, is not strongly associated with the full integration index, though its negative association is concentrated in the teacher-adoption pillar (see Appendix Figure D.2 and Appendix A.2 for additional details).

Perhaps the most striking feature of Panel (a) of Figure 5 is the asymmetry between facilitators and perceived barriers. Although principals cite several kinds of barriers to integrating, including costs, student misuse, teacher skill gaps, and policy, none appears to predict AI integration very well. Overall, principals can readily name the obstacles they encounter, but those self-reported obstacles do not appear to translate into observable differences in integration. This suggests that moving integration may require more than removing reported obstacles or increasing general spending.

Figure 5: Factors that Predict AI Integration and Explain Diffusion Gaps



Notes: This figure reports regression estimates assessing predictors of AI integration and the extent to which those predictors explain the diffusion gaps. The outcome is the standardized AI integration index. Panel (a) reports coefficients on school-level predictors from a single regression estimated on the traditional public school sample that controls for standardized FRPL share, demographic missingness indicators, and state fixed effects. The predictors are grouped into six categories: principal beliefs, leadership and discretion, budget flexibility and ESSER use, perceived barriers to adoption, cell-phone restrictions, and district/system capacity. The displayed coefficients in Panel (a) include whether the principal believes AI has improved student learning (Beliefs); whether the principal credits teacher champions for enabling AI adoption and has authority to purchase instructional software (Leadership); and whether the school has sufficient budget and flexibility to direct funds toward AI purchases (Budget). Panel (b) reports coefficients on each diffusion gap across sequential specifications that cumulatively add the predictor groups included in Panel (a). Disadvantage coefficients are estimated on the traditional public school sample; charter school coefficients are estimated on the public school sample (traditional public versus charter schools); private school coefficients are estimated on the non-charter sample (traditional public versus private schools). District/system capacity specifications are estimated only for the FRPL-gap analysis because district scale and fiscal capacity measures are not defined comparably for charter and private schools. All regressions are IPW-weighted. Whiskers denote 95 percent confidence intervals based on heteroskedasticity-robust standard errors.

7.2 What Explains Diffusion Gaps?

What can we learn about the mechanisms driving diffusion gaps? Appendix Figure D.1 shows why the observed school-level channels are unlikely to mediate the gaps. Beliefs differ more by sector than by disadvantage; leadership and budget flexibility often run against the observed gaps; and commonly reported barriers, including teacher skill gaps and student misuse, are relatively flat across disadvantage groups.

Panel (b) of Figure 5 confirms this implication. Conditioning on beliefs, leadership, budget, ESSER use, and barriers leaves the disadvantage and sector gaps essentially unchanged. For example, the coefficient on log neighborhood income moves from -0.13σ in the baseline specification to -0.14σ in the fully saturated model. Likewise, the charter gap only moves from -0.28σ to -0.37σ and the private school gap moves from -0.40σ to -0.49σ when accounting for the surveyed factors. We also add a the standardized cell-phone restriction index; the lower-income and charter

gaps are effectively unchanged, and the private-school gap becomes slightly more negative, moving from -0.49σ after the surveyed mechanisms to -0.51σ after adding the cell-phone index. Thus, although stronger cell-phone restrictions predict slower teacher adoption, they do not explain the private-school gap or the other diffusion gaps. None of these movements are statistically meaningful. Appendix Figure E.3 reports selection corrected estimates and are qualitatively similar.

The one margin that does attenuate the disadvantage gap is system scale, the closest observed proxy for $S_{g(j)}$ in the framework. Adding district scale reduces the gap from about -0.11σ to about -0.07σ , roughly 36% of the gap. Adding the fiscal measures on top of district scale produces little additional attenuation, suggesting that the relevant capacity margin is organizational scale rather than per-pupil spending.

Taken together, Figure 5 shows that the factors associated with higher integration levels are not necessarily the factors that explain the gaps. Our results suggest that unequal diffusion arises from differences in the conditions and choices that enable integration. This suggests that policies that raise AI integration levels—e.g., by targeting principals’ beliefs or budgets—may not reduce disparities across schools. Instead, the evidence points to the relative importance of factors that shape how access to AI tools and implementation support are distributed across schools: the relative importance of $S_{g(j)}$ over x_j . These include differences in district capacity, vendor reach, and other system-level features.

8 Conclusion

Generative AI has entered K-12 schools with remarkable speed. In less than three years, use within schools rose from near zero to near universal, and schools have largely chosen to manage student use rather than prohibit it. But adoption is not integration. Student use remains concentrated in homework and writing, educators use AI mainly for planning and administrative support, and institutional complements such as professional development, written guidance, and AI-enabled instructional infrastructure remain less developed.

If AI is a general-purpose technology whose value depends on complementary investments, then adoption alone is an incomplete measure of diffusion. The relevant question is not only whether schools use AI, but whether they have built the policies, routines, training, tools, and guidance to shape its use. Measuring those complements reveals a new form of digital inequality in K-12 schools: not unequal access to AI tools, but unequal development of the policies, routines, training, and infrastructure that govern school-level use. Schools serving more disadvantaged students, as well as charter and private schools, have lower levels of AI integration across policy, training, use, tutoring infrastructure, and future plans.

Our evidence on factors facilitating integration and mediating gaps reinforce our interpretation. Principal beliefs, teacher champions, budget availability, and targeted ESSER use predict which schools are more integrated, but they do not explain the socioeconomic or sectoral gaps. The factors that raise integration levels are therefore not necessarily the factors that close integration

gaps. District scale, by contrast, explains part of the disadvantage gap, suggesting that system-level supports may matter more than school-level willingness or general spending alone.

These findings reveal a narrower and harder policy problem than simply encouraging AI adoption. If AI's effects depend on local complements, then policies that increase use may leave disparities intact. Closing AI integration gaps may therefore require strengthening the school and district capacity that converts widespread access into directed practice, while recognizing that differences in integration may also reflect schools' underlying models and priorities rather than purely constraints. Whether these diffusion gaps translate into achievement gaps remains an important and open question, but the evidence here shows that the conditions for unequal effects are already present.

References

- Abdulkadiroğlu, Atila, Joshua D Angrist, Susan M Dynarski, Thomas J Kane, and Parag A Pathak, “Accountability and flexibility in public schools: Evidence from Boston’s charters and pilots,” *The Quarterly Journal of Economics*, 2011, 126 (2), 699–748.
- Allcott, Hunt, E. Jason Baron, Thomas Dee, Angela L Duckworth, Matthew Gentzkow, and Brian Jacob, “The Effects of School Phone Bans: National Evidence from Lockable Pouches,” Working Paper 35132, National Bureau of Economic Research April 2026.
- Arkansas Department of Education, “Arkansas Awarded \$25 Million in Federal Grants to Study Literacy and Tutoring Programs,” 2026.
- Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman, “Generative AI without guardrails can harm learning: Evidence from high school mathematics,” *Proceedings of the National Academy of Sciences*, 2025, 122 (26), e2422633122.
- Bick, Alexander, Adam Blandin, and David J Deming, “The rapid adoption of generative AI,” *Management Science*, 2026.
- Bloom, Nicholas and Christos Makridis, “The Politics of AI,” Technical Report, National Bureau of Economic Research 2026.
- Bresnahan, Timothy F and Manuel Trajtenberg, “General purpose technologies ‘Engines of growth’?,” *Journal of econometrics*, 1995, 65 (1), 83–108.
- , Erik Brynjolfsson, and Lorin M Hitt, “Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence,” *The quarterly journal of economics*, 2002, 117 (1), 339–376.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, “The productivity J-curve: How intangibles complement general purpose technologies,” *American Economic Journal: Macroeconomics*, 2021, 13 (1), 333–372.
- Bulman, George and Robert W. Fairlie, “Technology and education: Computers, software, and the internet,” *Handbook of the Economics of Education*, 2016, 5, 239–280.
- Carroll, Jamie M., Douglas N. Harris, Anjana Nair, and Emilia Nordgren, “National Longitudinal School Database: Data Description,” Technical Report, Tulane University, National Center for Research on Education Access and Choice, New Orleans, LA 2023.
- Contractor, Zara and Germán Reyes, “Generative AI in higher education: Evidence from an elite college,” *Higher Education: Evidence from an Elite College. IZA Discussion Paper*, 2025, (18055).
- Dell’Acqua, Fabrizio, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraye, François Cadelon, and Karim R Lakhani, “Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality,” *Harvard business school technology & operations mgt. Unit working paper*, 2023, (24-013).
- Diliberti, Melissa Kay, Heather L Schwartz, Sy Doan, Anna Shapiro, Lydia R Rainey, and Robin J Lake, “Using Artificial Intelligence Tools in K-12 Classrooms. American School District Panel. Research Report. RR-A956-21.,” *RAND Corporation*, 2024.
- EAB, “Voice of the Superintendent: Executive Brief,” 2024.
- Ertmer, Peggy A., “Addressing first- and second-order barriers to change: Strategies for technology integration,” *Educational Technology Research and Development*, 1999, 47 (4), 47–61.

- , **Anne T. Ottenbreit-Leftwich, Olgun Sadik, Emine Sendurur, and Polat Sendurur**, “Teacher beliefs and technology integration practices: A critical relationship,” *Computers & Education*, 2012, *59* (2), 423–435.
- Escueta, Maya, Vincent Quan, Andre Joshua Nickow, and Philip Oreopoulos**, “Education technology: An evidence-based review,” *Journal of Economic Literature*, 2020, *58* (4), 897–996.
- Figlio, David N and Umut Özek**, “The Impact of Cellphone Bans in Schools on Student Outcomes: Evidence from Florida,” Technical Report, National Bureau of Economic Research 2025.
- FutureEd**, “A Primer on Sustaining High-Impact Tutoring After ESSER,” 2024.
- , “Research Notes: Two Emerging Strategies for Using AI in Tutoring,” 2024.
- Gallup and Walton Family Foundation**, “Teaching for Tomorrow: Unlocking Six Weeks a Year With AI,” Report, Gallup 2025.
- Heckman, James J**, “Sample selection bias as a specification error,” *Econometrica: Journal of the econometric society*, 1979, pp. 153–161.
- Hew, Khe Foon and Thomas Brush**, “Integrating technology into K-12 teaching and learning: Current knowledge gaps and recommendations for future research,” *Educational technology research and development*, 2007, *55* (3), 223–252.
- Humlum, Anders and Emilie Vestergaard**, “The unequal adoption of ChatGPT exacerbates existing inequalities among workers,” *Proceedings of the National Academy of Sciences*, 2025, *122* (1), e2414972121.
- Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun**, “The diffusion of new technologies,” *The Quarterly Journal of Economics*, 2025, *140* (2), 1299–1365.
- Kaufman, Julia H, Ashley Woo, Joshua Eagan, Sabrina Lee, and Emma B Kassan**, *Uneven adoption of artificial intelligence tools among US teachers and principals in the 2023-2024 school year*, RAND, 2025.
- Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz**, “Experimental analysis of neighborhood effects,” *Econometrica*, 2007, *75* (1), 83–119.
- Lehmann, Matthias, Philipp B Cornelius, and Fabian J Sting**, “AI meets the classroom: When do large language models harm learning?,” *arXiv preprint arXiv:2409.09047*, 2024.
- Lira, Benjamin, Todd Rogers, Daniel G Goldstein, Lyle Ungar, and Angela L Duckworth**, “Coach not crutch: AI assistance can enhance rather than hinder skill development,” *arXiv preprint arXiv:2502.02880*, 2025.
- Mirror Indy**, “Indiana Schools Embrace AI but Seek to “Keep Humans in the Loop”,” 2025.
- Noy, Shakked and Whitney Zhang**, “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 2023, *381* (6654), 187–192.
- Parsad, Basmat, Rebecca R Skinner, Elizabeth Farris, and Shelley Burns**, *Advanced Telecommunications in US Private Schools, 1998-99*, US Department of Education, 2001.
- Renta-Davids, Ana-Inés, Marta Camarero-Figuerola, and Mar Camacho**, “Navigating the challenges and opportunities of artificial intelligence in educational leadership: A scoping review,” *Review of Education*, 2025, *13* (2), e70101.
- Shi, Ying and Francisco Villarroel**, “The Consequences of Cellphone Restrictions in Classrooms,” Technical Report, IZA Network@ LISER 2026.

- Todd, Petra E and Kenneth I Wolpin**, “On the specification and estimation of the production function for cognitive achievement,” *The Economic Journal*, 2003, *113* (485), F3–F33.
- U.S. Department of Education**, “U.S. Department of Education Issues Guidance on Artificial Intelligence Use in Schools,” 2024.
- Vigdor, Jacob L, Helen F Ladd, and Erika Martinez**, “Scaling the digital divide: Home computer technology and student achievement,” *Economic inquiry*, 2014, *52* (3), 1103–1119.

Online Appendix

AI Diffusion Gaps: Unequal Integration of AI Across K-12 Schools

Christopher Campos John Singleton

Table of Contents

A	Data Appendix	1
A.1	AI Integration Index	4
A.2	Cell-Phone Policy Data	6
B	Rapid and Decentralized Diffusion — Supporting Evidence	8
C	AI Diffusion Gaps — Supporting Evidence	17
D	What Mechanisms Drive the Diffusion Gaps? — Supporting Evidence	25
D.1	Cell-Phone Restrictions and AI Integration	26
E	Additional Supporting Evidence	27
E.1	Selection	27
E.2	Principal Characteristics	30

A Data Appendix

Table A.1: Survey Modules

Module	Title	Content
0	School & Principal Profile	School type (public, charter, private), grades served, years as principal
1	Decision Authority & Budget	Purchasing authority for instructional software, discretionary technology budget, use of federal pandemic-relief (ESSER) funds for AI, ease of redirecting funds toward AI
2	Policy, Governance & Ethics	Existence and scope of written generative AI policy (district- vs. school-level), topics covered (privacy, academic honesty, accessibility), disciplinary actions for AI misuse
3	Teacher Professional Development	Timeline of teacher AI adoption (2022–2025), estimated share of teachers using AI for specific tasks (lesson planning, grading, feedback), AI professional development delivery and coverage
4	Tutoring & AI-Assisted Teaching	AI-tutoring platform licensing, student usage rates, frequency of AI replacing direct instruction, principal beliefs about measurable learning improvements from AI
5	Student Adoption & Guidance	Estimated student AI usage by task (homework, essay drafting, study tools), school stance toward student AI use (encouraged, discouraged, prohibited), beliefs about AI and achievement gaps
6	Barriers, Enablers & Outlook	Top two barriers to AI adoption (cost, policy, teacher skill gaps, infrastructure, parent concerns, student misuse), most important enabling factor (teacher champions, vendor support, district mandate), expected AI integration three years ahead, principal’s own AI use
7	Demographics	Race/ethnicity, gender, education level, years as educator

Notes: This table summarizes the eight modules of the survey. Each row lists the module number, title, and a brief description of the topics covered.

Table A.2: Comparison of Respondents to Non-Respondents

	Respondents			Non-respondents		Difference	
	Mean	SD	IPW mean	Mean	SD	Raw	IPW
FRPL Share	0.59	0.27	0.54	0.55	0.26	0.04***	-0.01
SEDA Achievement	-0.09	0.42	-0.03	-0.03	0.40	-0.06***	-0.00
School Enrollment	590.80	416.76	593.40	593.38	447.69	-2.57	0.02
Neighborhood Median Income	73.18	36.42	75.26	74.75	34.81	-1.57	0.51
Neighborhood White Share	64.47	26.70	69.05	68.60	22.85	-4.13**	0.46
Neighborhood Black Share	15.82	23.18	15.12	15.12	19.51	0.70	0.00
Neighborhood Hispanic Share	22.78	25.64	19.79	20.15	22.67	2.62**	-0.36
Neighborhood Asian Share	6.23	11.22	4.93	5.06	8.30	1.17*	-0.12
Charter School	0.08	0.27	0.07	0.09	0.28	-0.01	-0.02
Private School	0.12	0.33	0.15	0.15	0.35	-0.02	0.00
N principals		1,254		71,985		73,239	


Notes: This table compares observable school and neighborhood characteristics for survey respondents with non-respondents in the sample. The sample contains 73,239 K–12 school contacts assembled from the NCES Common Core of Data and Private School Survey. Columns report unweighted respondent means and standard deviations, IPW-weighted respondent means, and unweighted non-respondent means and standard deviations. The difference columns report coefficients from regressions of each characteristic on a respondent indicator, estimated without and with IPW weights, respectively, using state-clustered standard errors. IPW weights are inverse response propensities, trimmed at the first and ninety-ninth percentiles and normalized to sum to the respondent count. The logit propensity model includes seven predictors: FRPL share, neighborhood poverty rate, log neighborhood median income, neighborhood White share, neighborhood Black share, neighborhood Asian share, and private school indicator. Private schools are excluded from the FRPL share and SEDA achievement rows because those measures are unavailable for private schools. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Randomization Balance and Incentive Effects

	Control mean	\$5	\$10+
Responded (%)	1.31	0.46*** (0.12)	0.61*** (0.11)
FRPL Share	55.14	-0.08 (0.28)	-0.05 (0.26)
SEDA Achievement	-0.03	0.01 (0.00)	0.00 (0.00)
Log School Enrollment	6.14	-0.01 (0.01)	0.00 (0.01)
Log Neighborhood Median Income	4.22	0.00 (0.00)	-0.00 (0.00)
Neighborhood White Share	68.60	0.44* (0.23)	0.40* (0.21)
Neighborhood Black Share	15.19	-0.39** (0.20)	-0.21 (0.18)
Neighborhood Hispanic Share	19.83	0.13 (0.23)	0.17 (0.21)
Neighborhood Asian Share	5.11	0.03 (0.08)	-0.11 (0.08)
Joint test p -value		0.164	0.366
Joint Wald test (both arms) p -value		0.239	

Notes: The table uses the invitation data of 73,239 schools. Each row reports coefficients from a regression of the listed outcome on indicators for assignment to the \$5 guaranteed-incentive arm and to either the \$10 or \$20 guaranteed-incentive arm, with the no-incentive arm omitted and randomization-block fixed effects included. The response-rate and FRPL-share outcomes are reported in percentage points; all other outcomes are pre-survey school or neighborhood characteristics from the cleaned analysis sample. The FRPL share and SEDA achievement rows are restricted to public schools because those measures are unavailable for private schools. We report p -values of tests that jointly test balance across all covariates for each arm separately and a version that combines both arms. Heteroskedasticity-robust standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A.1: Respondents by State



Notes: This figure plots a choropleth map of the geographic distribution of retained survey respondents across the contiguous United States. States are shaded by the number of retained survey respondents in each state, binned into quintiles using the empirical distribution of respondent counts across all states with at least one response. State assignments are taken from the school’s NCES state FIPS code where available and otherwise from the principal’s email domain or manual school identification. Color shading transitions from sand (lowest quintile) to maroon (highest quintile); states with no respondents are shown in light grey. The map describes raw respondent counts and is not weighted by IPW.

A.1 AI Integration Index


We use the survey responses to build an AI integration index that aggregates information on formal policies, encouragement of student use, leadership engagement, teacher adoption and breadth of classroom use, expected future integration, and the availability of AI-based instructional tools and professional development. This index is our empirical proxy for I_j^* outlined in the conceptual framework. The integration index is constructed as the average of seven standardized pillars:

1. **AI policy:** School has a written generative AI policy at the district or school level and depth of policy. Coded as zero if there is no policy, one if there is a policy, and two if there is a policy and at least three of the five listed items are covered by the policy.
2. **Student use encouraged** (binary): School actively encourages student AI use with clear rules, rather than discouraging or prohibiting it.
3. **Teacher adoption:** Mean of four binary indicators for whether teachers had adopted AI tools by June 2022, 2023, 2024, and 2025. This component ranges from 0 to 1.
4. **Principal-use breadth:** Share of four administrative tasks—offloading administrative busy-work, streamlining data analysis and reporting, designing professional development, and prepping meetings and presentations—for which the principal personally uses generative AI. This component ranges from 0 to 1.

5. **Future integration** (binary): Principal expects extensive or embedded AI integration within three years.
6. **AI tutoring license** (binary): School currently has a license or contract for an AI-powered tutoring platform.
7. **AI professional development**: Staff received generative AI PD by June 2025 and perceptions about PD effectiveness. Coded as zero if no PD, one if received by June 2025, and two if received PD and perceived as effective.

We follow Kling et al. (2007) and aggregate the pillars by first standardizing each one (to the inverse-probability-weighted survey distribution), then averaging the standardized values across pillars, and finally re-standardizing the resulting aggregate to mean zero and standard deviation one.

Figure A.2: AI Integration and Beliefs about AI’s Impact on Student Learning



Notes: This figure reports coefficients from linear probability models regressing principals’ positive student-learning-impact belief on the standardized AI integration index and each standardized index pillar. The outcome is an indicator equal to one when the principal answered “probably yes” or “definitely yes” to “Do you believe Generative AI tools have produced measurable improvements in student learning at your school?” The full index is the seven-pillar AI integration index defined in Appendix A.1; component rows use the corresponding standardized pillars. Unadjusted estimates include no controls or fixed effects. Adjusted estimates control for standardized FRPL disadvantage, school racial composition, missingness indicators, and state fixed effects. All regressions are estimated on the common traditional-public-school sample, use IPW weights, and report 95 percent confidence intervals based on heteroskedasticity-robust standard errors.

A.2 Cell-Phone Policy Data

We supplement the principal survey with school-level measures of student cell-phone restrictions. The policy data were collected from publicly available sources, including student handbooks, codes of conduct, board policy manuals, district or school web pages, and related parent-student policy documents. We searched for current policies governing student possession, storage, and use of cell phones or personal electronic devices during the school day. After the initial collection, we conducted an audit and recoding process in which coders reviewed the original source information and, when necessary, conducted an independent search to verify whether the policy applied to the relevant school and school year.

We code each school's policy along a restrictiveness scale from 0 to 5. A value of 0 indicates that no cell-phone restriction was found. A value of 1 indicates a weak or limited restriction, such as rules that silence phones or restrict disruptive use but do not clearly prohibit use during class. A value of 2 indicates a classroom-only restriction, where phones are prohibited during instructional time but may be allowed during passing periods, lunch, or other non-instructional parts of the day. A value of 3 indicates a full-day use restriction, where students may possess phones but may not use them during the school day. A value of 4 indicates a full-day storage restriction, where phones must be kept away, such as in backpacks, lockers, or designated storage. A value of 5 indicates the strongest restrictions, such as policies that prohibit possession during the school day or require phones to be locked in pouches, lockers, or another controlled storage system. Appendix Table A.4 reports summary statistics of the various measures we collect separately for each sector.


Table A.4: Cell-Phone Policy Measures by Sector

	Public	Charter	Private
Any policy found (%)	95.9	93.8	57.0
No policy found (%)	3.9	6.1	42.2
Weak/silenced-only policy (%)	1.0	0.0	0.6
Classroom-only restriction (%)	20.3	14.2	5.5
School-day restriction or stronger (%)	74.5	79.6	50.7
Full-day use only (%)	22.8	24.7	8.7
Storage restriction or stronger (%)	51.6	54.9	41.9
Full-day storage only (%)	46.4	34.3	20.9
Possession/locked-storage restriction (%)	5.3	20.6	21.0
0–5 restriction index (mean)	3.22	3.43	2.26
0–5 restriction index (SD)	1.12	1.30	2.12
Standardized restriction index (mean)	0.09	0.24	-0.60

Notes: This table reports student cell-phone and personal-device policy measures by school sector, coded from publicly available non-survey sources such as student handbooks, codes of conduct, board policy manuals, district or school web pages, and parent-student policy documents. Public denotes traditional public schools. Traditional public schools are assigned the district-level policy; charter and private schools are assigned school-, network-, or governing-organization-level policies when available. Percentages, means, and standard deviations are IPW-weighted. Any policy found indicates that we identified an applicable student cell-phone or personal-device policy; no policy found indicates that no applicable policy was identified. Weak/silenced-only policies require phones to be silenced or prohibit disruptive use but do not clearly restrict classroom or school-day use. Classroom-only restrictions prohibit use during instructional time but allow or do not clearly prohibit use during non-instructional time. School-day restriction or stronger is a cumulative indicator for policies that prohibit use throughout the school day, require storage throughout the day, ban possession, or require locked storage. Full-day use only is the mutually exclusive category for policies that prohibit use throughout the school day while allowing possession. Storage restriction or stronger is a cumulative indicator for policies requiring phones to be stored away, banning possession, or requiring locked storage. Full-day storage only is the mutually exclusive category for policies requiring storage throughout the day without banning possession or requiring locked storage. Possession/locked-storage restrictions ban possession during the school day or require phones to be locked in pouches, lockers, or another controlled storage system. The 0–5 restriction index ranges from 0 for no restriction found to 5 for possession bans or locked-storage requirements, with higher values indicating more restrictive student cell-phone policies. The standardized restriction index standardizes this 0–5 index using IPW weights.


B Rapid and Decentralized Diffusion — Supporting Evidence

Figure B.1: Teacher AI Adoption and Professional Development Timelines




Notes: This figure documents the rapid diffusion of generative AI in U.S. K–12 schools between 2022 and 2025. Panel (a) plots the IPW-weighted share of schools for which principals reported that teachers were using Generative AI tools for teaching-related tasks on or before June of each year, 2022–2025. Panel (b) plots the IPW-weighted share of schools for which principals reported that staff had received any professional development training on generative AI by the same June dates. Both panels show annual point estimates connected by lines on the full sample of 1,254 responding principals; no additional controls or fixed effects are included. The outcomes are binary indicators for any reported teacher AI adoption and any reported AI professional development, respectively. All estimates are IPW-weighted to be nationally representative of U.S. K–12 schools.

Figure B.2: Adoption and Professional Development by Sector




Notes: This figure reports teacher AI adoption and AI professional development trends by school sector for 2022–2025. Panel (a) reports the IPW-weighted share of schools for which principals reported that teachers were using Generative AI tools for teaching-related tasks on or before June of each year, separately for traditional public, charter, and private schools. Panel (b) reports the analogous share of schools for which principals reported that staff had received any professional development training on generative AI by the same June dates. The stars reported above charter and private points correspond to the significance level of unadjusted differences relative to traditional public schools in a given year: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All estimates are IPW-weighted.

Figure B.3: ESSER Funds and AI Adoption




Notes: Panel (a) reports the IPW-weighted share of schools whose principals reported that their school had used federal pandemic-relief funds (e.g., ESSER) for AI-related purchases, separately for traditional public, charter, and private schools. Principals could answer yes, no, or not sure; the plotted indicator codes yes as reported ESSER use and no or not sure as not reported ESSER use. Panel (b) plots the IPW-weighted share of schools for which principals reported that teachers were using Generative AI tools for teaching-related tasks on or before June of each year, 2022–2025, separately by this reported ESSER-use indicator. The vertical axes show percentages.

Figure B.4: Generative AI Policy Status




Notes: This figure reports the IPW-weighted distribution of schools' formal generative AI policy status as of June 2025. Principals were asked whether their school has a written policy specifically addressing generative AI use. The bars show the share of schools in each of four mutually exclusive response categories: a written policy in place at the district level, a written policy in place at the school level, a policy currently in draft, or no written policy. The figure includes all 1,254 responding principals and uses IPW weights to be nationally representative.

Figure B.5: Topics Covered by AI Policy




Notes: This figure reports the IPW-weighted share of schools reporting that a district-level generative AI policy, school-level generative AI policy, or draft generative AI policy explicitly covers each topic area. Principals who reported a district-level policy, school-level policy, or draft in progress were asked, “Which topics are explicitly covered in that policy? (Select all.)” Response options are “student data privacy,” “academic honesty/plagiarism,” “special-education accommodations,” “accessibility,” “content filtering/safety,” and “other.” The sample is restricted to principals with a non-missing response to this policy-topics item. All estimates are IPW-weighted to be nationally representative of schools answering the item.

Figure B.6: Disciplinary Actions for AI Misuse




Notes: This figure reports the IPW-weighted distribution of the number of disciplinary actions for AI misuse in the most recent academic year, separately for students, teachers, and other staff, among schools reporting a district-level generative AI policy, school-level generative AI policy, or draft in progress. Principals who reported a district-level policy, school-level policy, or draft in progress were asked, “Based on existing policy at your school, what number of disciplinary actions for AI misuse occurred at your school in the most recent academic year for the following parties:” students, teachers, and other staff. Response bins are 0, 1–2, 3–5, and 6+. The sample is restricted to principals with a non-missing response to the relevant group item. All estimates are IPW-weighted to be nationally representative of schools answering the item.

Figure B.7: AI-Powered Tutoring License or Subscription




Notes: This figure displays whether schools hold an active license or contract for an AI-powered tutoring platform. Principals were asked whether their school currently has a license or contract for any AI-powered tutoring platform. Response options are yes and no. All estimates are IPW-weighted to be nationally representative.

Figure B.8: Stance Toward Student Use of Generative AI Tools




Notes: This figure reports schools' official stance toward student use of generative AI tools. Principals were asked the current stance communicated to students regarding Generative AI use. The bars show the IPW-weighted share of schools in each of three response categories: actively encouraged for learning, including studying and homework assistance, with clear rules; discouraged but not banned; or prohibited. The figure includes all 1,254 responding principals and uses IPW weights to be nationally representative.

Figure B.9: Generative AI Widening Achievement Gaps




Notes: This figure reports the IPW-weighted distribution of principal agreement with the statement, “Generative AI is widening achievement gaps among students in my school.” Responses are on a five-point agreement scale: “strongly disagree,” “somewhat disagree,” “neither agree nor disagree,” “somewhat agree,” and “strongly agree.” The figure includes principals with non-missing responses to this item. All estimates are IPW-weighted to be nationally representative.

Figure B.10: Generative AI Produced Improvements in Learning




Notes: This figure reports the IPW-weighted distribution of principal responses to whether Generative AI tools have produced measurable improvements in student learning at their school. Principals were asked, “Do you believe Generative AI tools have produced measurable improvements in student learning at your school?” Response options are “definitely not,” “probably not,” “might or might not,” “probably yes,” and “definitely yes.” The figure includes principals with non-missing responses to this item. All estimates are IPW-weighted to be nationally representative.

Figure B.11: Generative AI Services Used by Principals



Notes: This figure reports the IPW-weighted share of principals frequently using each generative AI service. Principals were asked, “Which of the following services do you frequently use? Check all that apply,” with response options for OpenAI–ChatGPT, Google–Gemini, Anthropic–Claude, X–Grok, DeepSeek–R1, and Other. Free-text responses originally reported as Other were reclassified into named service categories where identifiable; residual unclassified responses are reported in the Other bar. The sample is restricted to principals who selected at least one listed or reclassified service. All estimates are IPW-weighted to be nationally representative.

Figure B.12: Generative AI Services Used by Teachers



Notes: This figure reports the IPW-weighted share of schools where principals report teachers use each generative AI service. Principals were asked, “Which of the following services do you know teachers at your school use? Check all that apply,” with response options for OpenAI–ChatGPT, Google–Gemini, Anthropic–Claude, X–Grok, DeepSeek–R1, and Other. Free-text responses originally reported as Other were reclassified into named service categories where identifiable; residual unclassified responses are reported in the Other bar. The sample is restricted to schools where principals selected at least one listed or reclassified teacher service. All estimates are IPW-weighted to be nationally representative.


C AI Diffusion Gaps — Supporting Evidence

Table C.1: AI Diffusion Gap by Disadvantage

	(1)	(2)	(3)	(4)	(5)	(6)
	FRPL		SEDA		Income	
	Baseline	Controls	Baseline	Controls	Baseline	Controls
FRPL Share	-0.359*** (0.110)	-0.368** (0.156)				
SEDA Achievement			0.269*** (0.072)	0.302*** (0.108)		
Log Neighborhood Median Income					0.170*** (0.054)	0.126* (0.067)
Neighborhood Black Share		0.146 (0.192)		0.222 (0.207)		0.036 (0.153)
Neighborhood Hispanic Share		0.085 (0.190)		0.108 (0.185)		-0.006 (0.150)
Neighborhood Asian Share		-0.192 (0.300)		-0.282 (0.306)		0.143 (0.293)
Observations	932	932	932	932	1228	1228
State FE	No	Yes	No	Yes	No	Yes
Sample	Public + charter schools				All schools	

Notes: This table reports the association between AI integration and economic disadvantage using three alternative disadvantage measures. The dependent variable is the AI integration index; Appendix Section A.1 describes its construction. Columns are grouped by the focal disadvantage measure. The FRPL Share and SEDA Achievement columns use the public school sample, including charter schools, because these measures are unavailable for private schools. The Income columns use Log Neighborhood Median Income and include all sectors. Baseline specifications include the focal disadvantage measure and its missingness indicator. Controls specifications additionally include neighborhood Black, Hispanic, and Asian shares, missingness indicators for those controls, and state fixed effects. All regressions use IPW weights and report heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.1: Component-Level FRPL Gaps




Notes: Each point is the coefficient on FRPL share from a separate regression of the indicated AI index component on FRPL, with state fixed effects, IPW weights, and missingness controls. Components are: Policy, written district- or school-level generative AI policy/topic depth, coded 0 for no written policy or draft only, 1 for a written policy, and 2 for a written policy covering at least three of five listed topics (student data privacy, academic honesty/plagiarism, special-education accommodations, accessibility, and content filtering/safety); Student policy, whether the current stance communicated to students is “Actively encouraged for learning, including studying and homework assistance, with clear rules”; Teacher adopt, the mean of four indicators for teacher use of generative AI tools for teaching-related tasks on or before June 2022, June 2023, June 2024, and June 2025; Principal use, the mean of four indicators for principal use to offload administrative busywork, streamline data analysis/reporting, design professional development, and prepare meetings/presentations; Future plans, expected “Extensive—most instruction” or “Embedded—routine in all classes” integration three years ahead; Tutoring, a current license or contract for any AI-powered tutoring platform; and PD, staff generative AI professional development by June 2025, with an additional point when at least 25% of instructional staff completed at least two hours of such professional development. Error bars are 95% CIs.

Table C.2: AI Diffusion Gap by Sector

	(1)	(2)	(3)	(4)
	Charter		Private	
	Baseline	Controls	Baseline	Controls
Charter School	-0.236** (0.112)	-0.327*** (0.124)		
Private School			-0.438*** (0.086)	-0.392*** (0.100)
Log Neighborhood Median Income		0.072 (0.072)		0.100 (0.070)
Neighborhood Black Share		-0.087 (0.169)		-0.095 (0.164)
Neighborhood Hispanic Share		-0.213 (0.170)		-0.123 (0.160)
Neighborhood Asian Share		-0.212 (0.284)		-0.080 (0.307)
Observations	1085	1085	1129	1129
State FE	No	Yes	No	Yes
Sample	Public + charter		Public + private	


Notes: This table reports the association between AI integration and charter school and private school status. The dependent variable is the AI integration index; Appendix Section A.1 describes its construction. The charter columns use the traditional public and charter school sample; the private columns use the traditional public and private school sample. Controls specifications include neighborhood Black, Hispanic, and Asian shares, missingness indicators for those controls, and state fixed effects. All regressions use IPW weights and report heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.2: Robustness of Diffusion Gaps to Selection into Survey




Notes: This figure reports diffusion gaps estimated on the full sample and by randomized response-incentive tier (no incentive, \$5, and \$10/\$20). The outcome is the standardized AI integration index, which aggregates seven standardized survey pillars: written generative AI policy and topics explicitly covered; student use actively encouraged for learning, including studying and homework assistance, with clear rules; teacher use of generative AI tools for teaching-related tasks by June 2022, 2023, 2024, and 2025; principal use for offloading administrative busywork, streamlining data analysis/reporting, designing professional development, and prepping meetings/presentations; expected extensive or embedded integration three years ahead; a license or contract for any AI-powered tutoring platform; and staff generative AI professional development by June 2025 plus at least 25% of instructional staff completing at least two hours of such professional development. Each reported estimate corresponds to the coefficient on economic disadvantage/FRPL share (Panel a), charter indicator (Panel b), or private indicator (Panel c), estimated in a regression of the standardized AI index on the focal measure, state indicators, racial demographics, missing indicators for schools with missing demographics, and log neighborhood income for the sector panels. All regressions are IPW-weighted, and heteroskedasticity-robust standard errors are used to construct the reported 95 percent confidence intervals. The reported p -value tests the joint null that the focal coefficient is equal across incentive tiers.

Figure C.3: Robustness to Selection into Survey Responses - Alternative Measures of Disadvantage




Notes: This figure replicates Figure C.2 using two alternative measures of school disadvantage. The outcome is the same standardized AI integration index, aggregating written generative AI policy and topics explicitly covered; student use actively encouraged for learning, including studying and homework assistance, with clear rules; teacher use of generative AI tools for teaching-related tasks by June 2022, 2023, 2024, and 2025; principal use for offloading administrative busywork, streamlining data analysis/reporting, designing professional development, and prepping meetings/presentations; expected extensive or embedded integration three years ahead; a license or contract for any AI-powered tutoring platform; and staff generative AI professional development by June 2025 plus at least 25% of instructional staff completing at least two hours of such professional development. Panel (a) reports the coefficient on standardized SEDA average test scores (1 SD) across the full sample and randomized response-incentive tiers, estimated on the public-school sample with controls for racial demographics, missingness indicators, and state fixed effects. Panel (b) reports the coefficient on log median neighborhood household income across the same displayed rows, estimated on the full sample of public and private schools with the same controls. Each estimate is IPW-weighted with heteroskedasticity-robust standard errors used to construct the reported 95 percent confidence intervals. The reported p -value tests the joint null that the coefficient is equal across incentive tiers; failure to reject this null suggests that the diffusion gaps are robust to selection into survey response generated by the randomized incentive arms.

Figure C.4: Diffusion Gaps Using a PCA-Based Index




Notes: The PCA index is the first principal component of the seven standardized baseline components, oriented to be positively correlated with the equal-weight baseline index. The components are written generative AI policy/topic depth; student use actively encouraged for learning with clear rules; teacher use of generative AI tools for teaching-related tasks by June 2022–2025; principal use for administrative busywork, data analysis/reporting, professional-development design, and meetings/presentations; expected extensive or embedded integration three years ahead; an AI-powered tutoring platform license or contract; and staff generative AI professional development by June 2025 plus at least 25% of instructional staff completing at least two hours of such professional development. PC1 explains 32% of component variance and correlates 1.00 with the baseline index. Specifications, samples, and inference match Figure C.5.

Figure C.5: Robustness of Diffusion Gaps to Leave-One-Component-Out Indexes



Notes: Each row corresponds to a re-estimation of the three headline diffusion gaps using an alternative AI integration index that drops one of the seven baseline components: written generative AI policy/topic depth; student use actively encouraged for learning with clear rules; teacher use of generative AI tools for teaching-related tasks by June 2022–2025; principal use for administrative busywork, data analysis/reporting, professional-development design, and meetings/presentations; expected extensive or embedded integration three years ahead; an AI-powered tutoring platform license or contract; and staff generative AI professional development by June 2025 plus at least 25% of instructional staff completing at least two hours of such professional development. The disadvantage gap is the coefficient on standardized FRPL share (1 SD) on the common public-school sample with missingness controls; charter and private gaps are estimated using the adjusted specifications from Figure 4 (state fixed effects, log neighborhood income, racial demographics, and missingness controls). All regressions are IPW-weighted with heteroskedasticity-robust 95% confidence intervals.


Figure C.6: Diffusion Gaps for Current-Use vs. Institutional-Support Sub-Indexes



Notes: The current-use index combines the student-encouragement, teacher-adoption, and principal-use components: student use actively encouraged for learning, including studying and homework assistance, with clear rules; teacher use of generative AI tools for teaching-related tasks by June 2022–2025; and principal use for offloading administrative busywork, streamlining data analysis/reporting, designing professional development, and prepping meetings/presentations. The institutional-support index combines the written-policy, AI-tutoring-license, and professional-development components: a written policy specifically addressing generative AI use plus topics explicitly covered in that policy; a license or contract for any AI-powered tutoring platform; and staff generative AI professional development by June 2025 plus at least 25% of instructional staff completing at least two hours of such professional development. The expected-future-integration component is included only in the baseline all-seven index. Each index is standardized using IPW weights. Specifications, samples, and inference match Figure C.5.

D What Mechanisms Drive the Diffusion Gaps? — Supporting Evidence


Figure D.1: Factors that May Explain the Diffusion Gap



Notes: This figure reports seven principal-reported mechanism variables by school disadvantage and sector. The left column splits public schools at the median FRPL share, with above-median schools classified as high disadvantage; the right column compares traditional public, charter, and private schools. Variables cover belief, leadership, budget, and professional-development channels. Belief measures indicate whether principals report measurable student-learning improvements from generative AI and whether they agree that generative AI is widening achievement gaps. Leadership measures indicate principal authority to purchase instructional software and teacher champions as the top enabler of current AI efforts. Budget measures indicate a discretionary technology budget of at least \$5,000 and whether redirecting funds toward AI is easy or very easy. The professional-development measure indicates agreement that AI-related PD improved lesson-planning efficiency, conditional on the school having received AI PD. All shares are IPW-weighted.

D.1 Cell-Phone Restrictions and AI Integration

Figure D.2: Cell-Phone Restrictions and Teacher AI Adoption




Notes: This figure reports coefficients from separate regressions of the standardized teacher-adoption pillar on each recoded cell-phone restriction measure. Cell-phone measures are coded from public district- or school-level policy sources, not the principal survey; the 0–5 index codes 0 for no restriction found, 1 for weak or silenced-only policies, 2 for classroom-only restrictions, 3 for full-day use restrictions, 4 for full-day storage restrictions, and 5 for possession bans or locked-storage requirements. The teacher-adoption pillar is the mean of four yes/no principal reports of whether teachers at their school were using Generative AI tools for teaching-related tasks on or before June 2022, June 2023, June 2024, and June 2025, standardized using IPW weights. The 0–5 restriction index is standardized using IPW weights, so its coefficient reports the association with a one standard deviation increase in policy restrictiveness. The remaining measures are binary indicators for school-day restrictions, storage-or-stronger restrictions, and possession or locked-storage restrictions. All regressions control for school sector, log neighborhood median household income, school racial composition, school level, log enrollment, missingness indicators, and state fixed effects. Regressions are IPW-weighted and use heteroskedasticity-robust standard errors to construct the reported 95 percent confidence intervals.

E Additional Supporting Evidence


E.1 Selection

Figure E.1: Headline Diffusion Gaps with Heckman Selection Correction




Notes: This figure reports the three headline gap estimates with and without the Heckman selection correction. The FRPL disadvantage row is the coefficient on standardized FRPL share in the common public-school sample, controlling for FRPL missingness. The charter row is the coefficient on a charter-school indicator among public schools, and the private row is the coefficient on a private-school indicator among non-charter schools; both sector specifications control for log neighborhood income, income missingness, racial-composition controls and missingness indicators, and state fixed effects. The first-stage probit is fit on the sample used in Tables A.2 and A.3, regressing a respondent indicator on cleaned sample covariates: FRPL, neighborhood income, neighborhood White, Black, and Asian shares, a private-school indicator, and randomized incentive indicators for \$5, \$10, and \$20 offers, with no incentive as the omitted group. The inverse Mills ratio is computed for each respondent and added as an additional control in the Heckman-corrected specifications. All regressions are IPW-weighted with heteroskedasticity-robust 95% confidence intervals.

Figure E.2: Mechanism Coefficients with Heckman Selection Correction



Notes: This figure reports mechanism coefficients with and without the Heckman selection correction. The outcome is the standardized AI integration index. Surveyed mechanism rows include positive student-learning-impact beliefs; teacher champions as the factor that most enabled current AI efforts; principal authority to purchase new instructional software; discretionary technology budget of at least \$5,000; redirecting previously allocated funds toward AI initiatives and AI-related tools as easy or very easy; ESSER use for AI-related purchases; and cost, student misuse, teacher skill gap, and policy restriction barriers selected from the up-to-two-barriers item. Baseline specifications control for log neighborhood income, income missingness, and state fixed effects. Heckman-corrected specifications additionally include the inverse Mills ratio from the first-stage probit described in Figure E.1. All regressions are IPW-weighted with heteroskedasticity-robust 95% confidence intervals.


Figure E.3: Gap Stability Across Mechanism Specifications — Heckman Correction



Notes: This figure reports gap stability across sequential mechanism specifications, comparing baseline and Heckman-corrected coefficients. The disadvantage estimate is the coefficient on standardized FRPL share among public schools; the charter estimate is the coefficient on a charter school indicator among public schools; and the private estimate is the coefficient on a private school indicator among non-charter schools. Sequential specifications cumulatively add positive-impact beliefs, leadership controls, budget controls, ESSER use, and barrier controls. The disadvantage specifications control for FRPL missingness; the sector specifications control for log neighborhood income and income missingness. Heckman-corrected specifications additionally include the inverse Mills ratio from the first-stage probit described in Figure E.1. All regressions are IPW-weighted with state fixed effects and heteroskedasticity-robust 95% confidence intervals.

E.2 Principal Characteristics

Figure E.4: Adoption by Principal Characteristics



Notes: This figure reports IPW-weighted shares of schools where principals answered yes that teachers at their school were using Generative AI tools for teaching-related tasks on or before June of each year, by principal gender, race/ethnicity, and years served as principal at this school. Panel (a) uses the gender item; Panel (b) uses the race/ethnicity item and reports White, Black, and Hispanic or Latino groups; Panel (c) groups reported years served as principal at this school into 1–3, 4–7, and 8+ years. Stars denote significant pairwise differences relative to the reference group (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). Reference groups: male, White, and 1–3 years.