Addendum to "Screens and AI in K-5 Classrooms: Research Brief for ARPS Parents"

October 6, 2026. New material since Draft 2, for integration into the working document. Each item names the brief section it belongs in. Boston section corrected against DESE data.

A1. PISA 2025: the first large-scale data tying student AI use to lower scores (for Section 6b)

The OECD's PISA 2025 results, released September 8, 2026, cover about 760,000 fifteen-year-olds in 91 countries. US reading scores fell 14 points between 2022 and 2025, to 490, the lowest on record; across the OECD, reading fell 28 points and math 22 points over the decade from 2015. Forty-six percent of OECD students use AI chatbots at least weekly. Students who reported using AI for specific schoolwork tasks scored lower in science than students who never used it: daily use for summarizing was associated with roughly 27 to 30 points lower, drafting text about 24 lower, and preliminary research about 12 lower, while use "to help me learn" was associated with about 4 points higher. The OECD equates 20 points to roughly a year of schooling and describes the AI findings as associations, not causes, noting that social media and screen time are among the confounders (PISA 2025 Results, Volume I; PDF; summary of the AI-use breakdown). This is 15-year-olds, not K-5, and correlational; its value for the brief is that it is the largest dataset on the question and that the pattern (offloading tasks hurts, using AI to learn does not) matches the Bastani RCT already in Section 6b.

A2. The WSJ investigation and podcast (for Sections 6a and 6b)

"Google Pushed Gemini AI Into Schools, and Even Students Say It's Gone Too Far," Wall Street Journal, September 30, 2026, by Shalini Ramachandran, Georgia Wells, and Molly Reinmann, based on interviews with more than 60 teachers, school officials, parents, children, and clinicians, current and former Google employees, and more than a dozen Google research papers (WSJ, paywalled; summaries at GIGAZINE, Tech Brew, Business Model Analyst). Companion podcast: The Journal, "Google Put AI in Schools. Test Scores Plummeted.", October 5, 2026, host Ryan Knutson with Ramachandran (transcript; audio). The full claims-and-sources breakdown is in the companion doc "WSJ Google AI in Schools, Research Claims and Sources." The items below are the ones that belong in the brief.

Google's own researchers warned first. In April 2024 more than 50 Google DeepMind researchers and ethicists published "The Ethics of Advanced AI Assistants," warning that chatty, anthropomorphized AI poses particular risks to children, who struggle to distinguish AI from human interaction, and could cause "diminished intellectual engagement," "reduced sense of personal competence," and "social degradation" (Gabriel et al., 2024). The same year Google's LearnLM team, 67 authors, wrote that general-purpose helpfulness "may often be at odds with

pedagogy and learning," that models give away answers quickly and promote over-reliance, that sycophancy makes them poor at flagging learner mistakes, and that easy expert answers create "an illusion of mastery before it has been achieved" (Jurenka et al., 2024). Two years later Google turned Gemini on by default for every grade.

Google's counter-evidence is on trial. Google's spokesperson told the WSJ that its studies show AI paired with teacher-led instruction improves learning, citing a 2025 Sierra Leone trial in which some students gained more than a year of math progress in eight weeks using Gemini's tutoring mode. The trial was a pre-registered RCT of 1,763 junior secondary students in 12 schools, Google-funded with Google-affiliated evaluation, with an effect of +0.258 SD and a wide confidence interval (0.027 to 0.488); Google's own write-up says the strongest students benefited most and that "new tools frequently widen achievement gaps rather than close them" (Google DeepMind, June 2026; technical report). Ramachandran's summary on the podcast: "there isn't any evidence, is the truth ... all of these tools have been put in front of kids without any really evidence backing that it helps improve learning."

Boston, checked against DESE. The WSJ reported that Boston Public Schools promoted Gemini for teenagers starting in 2025 and that 2026 scores fell. DESE's district profiles confirm the declines (grade 10 ELA 40% to 33% meeting or exceeding expectations, math 38% to 33%, science 33% to 27%) but show the whole state fell almost identically the same year (grade 10 ELA 51% to 45%, science 46% to 41%, math 45% to 44%), and Boston's grades 3-8 held flat while the state slipped. 2026 was the first year the grade 10 MCAS carried no graduation stakes after voters removed the requirement in November 2024, and state officials and researchers attributed the drop to pandemic recovery and reduced testing stakes, with no mention of AI. The Boston figures therefore do not support a Gemini effect and should not be cited as one; they are a useful example of why single-district score correlations need a statewide comparison before they go in front of a committee (DESE 2026 release; Boston profile 2026 and 2025; WBUR, Sept 22, 2026). The same caution applies to the WSJ's Denver and New York figures, which were not checked here. Miami-Dade, the largest US district AI deployment (Gemini to 105,000+ high schoolers in 2025), supplied the article's teacher accounts of A and B students failing exams after a year on AI; those are anecdotes, not data.

Expert quotes usable in the brief. Wharton's Hamsa Bastani (author of the PNAS RCT): students face "this constant choice between doing the easy thing and doing the hard thing," and the problem is worse for younger children. Former Google research scientist Mikel Rodriguez: "We're doing this experiment on an entire generation on cognitive development, with no control group." MIT Media Lab's Nataliya Kosmyna (author of "Your Brain on ChatGPT," a visiting researcher at Google): "We learned nothing from social media."

Common Sense Media, two new items. The 2026 census of AI use by tweens and teens (1,200+ children ages 9 to 17, fieldwork March 2026) finds 86% have used generative AI and

75% have used AI search summaries, the most widely experienced form of AI for kids (report). A July 2026 risk assessment of Google Search's AI Overviews and AI Mode, using accounts configured for ages 11 and 15 across 2,600+ searches, found the AI completed all 180 math and 180 humanities homework assignments it was handed and missed 29% of explicit suicide references; rated "unacceptable risk." Google called the queries "narrow" and "contrived" (assessment; PBS NewsHour).

A3. What a district can and cannot switch off (for Section 6a and the Information Systems questions)

Google Admin can disable the Gemini app, Gemini Notebook, the Gemini tab in Classroom, and Gemini in Docs, Slides, and Gmail, per organizational unit (Google admin help). AI Overviews and AI Mode in Google Search cannot be disabled by an administrator; the only Google-provided option is blocking Google Search entirely. Photomath (Google-owned, step-by-step answers from a screenshot) and Google Lens are also built into or one click from a Chromebook. Common Sense reports that several districts learned Gemini had been enabled for all ages only after the fact (Common Sense). Per the WSJ, Los Angeles and Fairfax County built browser plug-ins to strip AI results from Search, Madison, Wisconsin blocked Gemini, and Chicago dropped plans for high school Gemini access.

Public tools. Neither LA nor Fairfax has published its plug-in. A free, open equivalent exists: the technology director of Niles Community Schools, Michigan, published a Chrome extension that removes AI Overviews and AI Mode from Google search results, written for K-12, with instructions for force-installing it to managed Chromebooks through the Google Admin console so students cannot remove it (GitHub). Similar extensions are on the Chrome Web Store (one, another), and school IT staff compare approaches in a Jamf community thread. Caveat: these work by hiding page elements, so a Google layout change can break them and a district has to maintain the extension.

Revised questions for Information Systems. The first question in Section 6 and in the memo should become two. Is Gemini (the app, the Classroom tab, and Gemini in Docs, Slides, and Gmail) disabled for the elementary organizational units, and since when? And is anything in place to remove AI Overviews and AI Mode from Google Search on elementary Chromebooks, given that Google provides no administrative switch for those; if not, would the district force-install a filtering extension such as the Niles tool to the elementary OU?

A4. Fairfax County, Virginia (for Section 5c)

On July 23, 2026, the Fairfax County School Board, the tenth-largest district in the country, prohibited generative AI for all elementary students; banned district-issued tablets and laptops in pre-K and kindergarten, with exceptions for IEP and 504 accommodations, translation, and multilingual support; required devices to be stored when not in use, with the superintendent to develop a storage plan; ordered YouTube blocked at the network level except for

teacher-embedded content; approved a whitelist-style firewall pilot at six schools, which a board member said required no new contracts because "our existing platform has this capability built in, and it's a simple system configuration"; set guardrails rather than a ban for secondary students; and required the superintendent to deliver a comprehensive AI inventory by January 14, 2027 (FFXnow; district AI page FCPS; board materials January 2026 Academic Matters). Of all the 2026 policies in Section 5, Fairfax's is the closest to a complete template for an Amherst School Committee action under goal 5: it combines an elementary AI prohibition, a K-2 device rule with carve-outs, a network-level YouTube block, a pilot using existing tools, and an inventory deadline.

A5. Seattle (already added to Section 5c on Sept 23; recorded here for completeness)

Superintendent Ben Shuldiner announced September 21 that Seattle Public Schools will end one-to-one iPads and cap classroom screen time at 30 minutes a day in K-2 from January 4, 2027, with shared carts, by administrative decision; the parent group Seattle Families for Intentional Technology, a Schools Beyond Screens chapter, partnered on it (KUOW).

A6. Correction to Section 5c, New York City

The NYC entry should distinguish the two provisions. The AI piece is a one-year moratorium on student-facing generative AI from pre-K through 8th grade, with exceptions for assistive technology, multilingual learners, and career-readiness programs, plus a ban on companion chatbots at every grade. The screen-time piece is separate guidance: school-issued devices restricted in pre-K through 2 (the city's release says "restricted 1:1 screen time"; Education Week reports a ban), and "recommended caps," not mandates, of 30 minutes a day in grades 3-5 and 45 in grades 6-8 (NYC; EdWeek). Los Angeles's September 3 action was a K-12 moratorium on student generative AI on district devices; its minute caps come from the June 23 screen policy, which also removed district devices from preschool through grade 1.

A7. Tech leaders on their own children (for Section 7 or a sidebar)

A separate doc in this folder, "Tech Leaders on Their Own Kids and Screens, Verified Quotes," holds about 30 sourced quotes, twelve poster picks, and a list of commonly fabricated ones. The most useful for the brief are from 2026: Sam Altman on when his son may use AI, "I'd rather be on the late end of what's reasonable there, not the early end" (Mostly Human, April 2026), and Anthropic co-founder Jack Clark, "I have the classic Californian technology executive view of not having that much technology around for children" (Ezra Klein Show, Feb 2026). The NYT ran a survey of the pattern on August 18, 2026, "Silicon Valley Executives Are Tech Fans. Just Not for Their Own Kids." (NYT).

Flags

PISA causation. Any use of the 20-point figure should keep the OECD's "association" wording

and the "to help me learn" exception. WSJ article text. wsj.com was not reachable; article content comes from the podcast transcript and six independent summaries, checked against primary sources. A widely shared Substack summary adds four studies the WSJ did not cite (Stanford sycophancy, Microsoft/CMU, Harvard physics, a China study); cite those on their own if at all. Denver and New York figures. Reported by the WSJ; not checked against state data here, and the Boston check suggests they need the same statewide comparison. LA and Fairfax plug-ins. Reported by the WSJ, not published by the districts.