Extracting Value from Coworkers:

Strategic Collaboration and Skill Formation

John Fallon*

February 5, 2026

Abstract

Workers learn from their peers, but which collaborations build lasting skills and which do not? I study co-teaching, where a general education teacher and a special education teacher share responsibility for the same classroom, using administrative data on all teachers in Indiana from 2012-2019. Comparing teachers before and after co-teaching, I find that pairing with highly experienced partners (16+ years) raises student achievement by 0.10σ after teachers return to solo instruction, while pairing with inexperienced partners produces no gains. Longer collaborations provide no additional benefit, and teachers with 2-3 years of prior experience gain more than novices. Strategic pairing of junior teachers with experienced mentors could accelerate skill development.

I thank Andrew Bacher-Hicks, Kevin Lang, Josh Goodman, Marcus Winters, James Feigenbaum, Olivia Chi, Seth Gershenson, Roddy Theobald, Meagan Comb, Sidrah Baloch, Lindsay Kaler, Josefina Senese, Tashnuva Shaheen, Michael Briskin, Danielle Williamson, and Jesse Meredith for helpful comments and suggestions. I also thank participants in the PREREQ seminar. Any remaining errors are my own.

*I declare that I have no relevant or material financial interests that relate to the research described in this paper. IRB approval was not obtained for this study. This project does not involve any interaction or intervention with subjects. All data sources provided by the Indiana Department of Education do not include any access to identifiable personal information.

1 Introduction

Peer effects in the workplace are well documented, but how should organizations leverage them to maximize workforce quality? I study this question using co-teaching partnerships, where a general education teacher and a special education teacher share instructional responsibility for a single classroom. This setting allows me to observe the same workers' productivity before, during, and after collaboration. Co-teaching assignment depends on special education enrollment and scheduling constraints rather than teacher quality, providing variation in who collaborates with whom. Teachers paired with highly experienced partners (16+ years) show 0.10σ higher student achievement after returning to solo instruction, while those paired with inexperienced colleagues show no improvement. Longer collaborations do not amplify these gains and may even reduce them. Organizations can substantially accelerate workforce development through strategic pairing, but current practice leaves most of these gains unrealized.

The title phrase "Extracting Value from Coworkers" is intentionally provocative, reframing workplace collaboration in starkly transactional terms. Yet this framing captures a fundamental question for organizational design: how and when should firms strategically leverage peer learning opportunities? My findings show that organizations leave substantial productivity gains on the table. Current practice allows institutional constraints to drive partnership formation, frequently pairing inexperienced workers together and maintaining partnerships longer than necessary. The answers to "how" and "when" matter enormously. Pairing with experienced colleagues generates lasting benefits while pairing novices together does not. Brief collaborations suffice while extended ones may crowd out valuable independent experience. Some baseline competence is necessary before workers can absorb what experienced partners offer. Strategic pairing represents smart human capital investment, but only when the right workers pair at the right time for the right duration.

Existing research establishes that coworkers affect each other's productivity but provides little guidance on how to maximize these effects. Some studies document skill transfer:

teachers improve when surrounded by better colleagues (Jackson and Bruegmann, 2009; Nix, 2020; Gershenson et al., 2023). Others identify social pressure or monitoring as the mechanism: supermarket cashiers speed up when they can observe fast peers (Mas and Moretti, 2009), and workers reduce absenteeism when their colleagues have low absence rates (Ichino and Maggi, 2000). Productivity spillovers appear across diverse settings through various channels (De Grip and Sauermann, 2012; Battu et al., 2003; Azoulay et al., 2010; Guryan et al., 2009). These studies document that peer effects exist but do not address how long collaborations should last or when in a worker's career collaboration is most valuable. I focus on these design questions because they determine whether organizations can translate peer effects into systematic workforce improvement.

Co-teaching offers an ideal setting to study optimal collaboration design. Under co-teaching arrangements, a general education teacher provides content expertise while a special education teacher supports a classroom containing both students with disabilities and those without.[1] This structure enables schools to satisfy dual federal mandates. The Individuals with Disabilities Education Act requires instruction in the least restrictive environment, while the No Child Left Behind Act mandates highly qualified content instruction. Neither teacher needs dual certification under this arrangement. The setting forces two teachers into close daily proximity, creating stronger peer influence opportunities than typical same-school interactions. Three features make this setting valuable for identification. I observe productivity before, during, and after collaboration. Assignment depends on institutional factors largely independent of teacher quality. High partnership dissolution rates (over 60% separate after one year) generate substantial variation in collaboration experiences, and I present evidence that dissolution is driven by institutional factors rather than match quality. A large literature examines student outcomes in co-taught classrooms (Austin, 2001; Jones and Winters, 2024; Jones et al., 2025; Ballis and Heath, 2021; Bessette, 2008; Cook et al., 2017; De Backer et al., 2023; Andersen et al., 2020; Conderman and Hedin, 2014), but I focus instead on how co-teaching affects teachers themselves.

[1] Models of how teachers split responsibility vary. I am agnostic about what model the teachers follow.

The main results identify which collaborations build workforce quality and which do not. Teachers paired with highly experienced colleagues (16+ years) show 0.10σ higher student achievement after returning to solo instruction. This effect falls within the range of intensive coaching interventions (Kraft et al., 2018) but may be more scalable. Teachers paired with moderately experienced colleagues (8-15 years) show smaller gains. Teachers paired with inexperienced colleagues (0-7 years) show no benefit. This gradient implies that experienced workers possess tacit knowledge that can transfer through close collaboration. Such knowledge includes classroom management techniques, pedagogical strategies, and approaches to differentiation. Inexperienced workers have little to offer their partners beyond mutual support.

Longer partnerships do not produce larger gains. Comparing teachers who co-taught for one, two, or three years, I find no evidence that additional collaboration time amplifies skill transfer. If anything, extended partnerships may reduce development by substituting for solo teaching experience that builds independent capability.

Career timing also matters. Teachers with 2-3 years of prior independent experience benefit most from collaboration, while first-year teachers gain less. Some baseline competence may be necessary to absorb what experienced partners can offer.

Research on learning-by-doing suggests that experience improves productivity, but the mechanisms remain unclear (Arrow, 1962; Chen, 2021; Haggag et al., 2017; Ost, 2014; Shaw and Lazear, 2008; Thompson, 2010, 2012). My results show that collaboration can accelerate this process when partners are chosen strategically. Experimental evidence from Papay et al. (2020) demonstrates that strategically pairing teachers improves outcomes and examines how partner skill dimensions affect gains. I complement this work by identifying when in a teacher's career collaboration is most valuable and how long partnerships should last.

I conduct a policy simulation to quantify the implications of these findings. The simulation compares strategic pairing to current practice and tracks teachers over their careers. It accounts for the opportunity cost when experienced mentors leave their solo classrooms to co-teach. Under baseline assumptions, strategic pairing generates the

equivalent of $\$16,400$ in additional lifetime earnings for every student in a school. Positive benefits emerge within six years and prove robust across conservative scenarios. Even in the most pessimistic case, cumulative effects turn positive within eight years. The break-even opportunity cost is twice as large as empirical estimates. Strategic pairing represents a high-return investment in teacher quality across a wide range of assumptions.

These findings speak directly to personnel policy. The 0.10σ effect of pairing with highly experienced teachers represents a meaningful and persistent productivity gain. Scaled across a teaching workforce, strategic pairing could substantially improve student achievement at minimal cost. Yet current practice rarely exploits this opportunity. Current assignment practices appear driven by scheduling convenience rather than developmental potential, and novice-novice pairings remain common despite producing no lasting gains. Organizations should treat collaboration as a deliberate human capital investment, assigning experienced workers to mentor junior colleagues through brief, strategically-timed partnerships.

2 Data

I use longitudinal administrative data from the Indiana Department of Education covering all students and teachers from 2012 to 2019. A unique classroom identifier links students to teachers. Student data include annual demographics, supplemental service classifications (special education, English language learners), and standardized test scores in math and ELA, which I standardize by grade and year.

I define co-taught classrooms as classes with exactly two teachers: one flagged as "Team-Teacher" or "Co-teacher" (the special education teacher) and one general education teacher. For my analysis of co-teaching spillovers, I restrict the sample to teachers identified with a "First Year" flag to ensure accurate experience measures, as the administrative "Years of Experience" field ignores career breaks and would misclassify returning teachers.

Table 1: Summary Statistics

	Co-Taught Sample			Event Study Sample		
Panel A: Sample Size						
Number of Teachers	10,100			804		
Number of Students	152,847			61,652		
Number of Classes	28,588			7,361		
Panel B: Student Characteristics						
	Obs	Mean	SD	Obs	Mean	SD
Test Z-Score	455,759	−0.387	0.965	171,843	−0.141	0.986
Special Education	630,310	0.329		171,843	0.227	
English Language Learner	630,310	0.097		171,843	0.089	
Free/Reduced-Price Lunch	630,310	0.606		171,843	0.506	

Notes: This table presents descriptive statistics for two samples: the full co-taught sample (all teachers and students in co-taught classrooms from 2012-2019) and the restricted event study sample (teachers with specific experience profiles used in the main analysis).

3 Co-Teaching Assignment Patterns

Co-teaching assignment follows predictable institutional patterns. Special education enrollment, scheduling constraints, and certification requirements shape which teachers co-teach and with whom. I present three pieces of evidence characterizing these assignment patterns.

3.1 Entry and Exit Patterns Reflect Institutional Constraints

Teachers enter co-teaching early in their careers when scheduling flexibility is highest. Both special education and general education teachers are most likely to begin co-teaching in their first few years (Figure 1). This pattern reflects the fact that junior teachers have less seniority-based control over their assignments.

Figure 1: Hazard Rate for Entering Co-Teaching

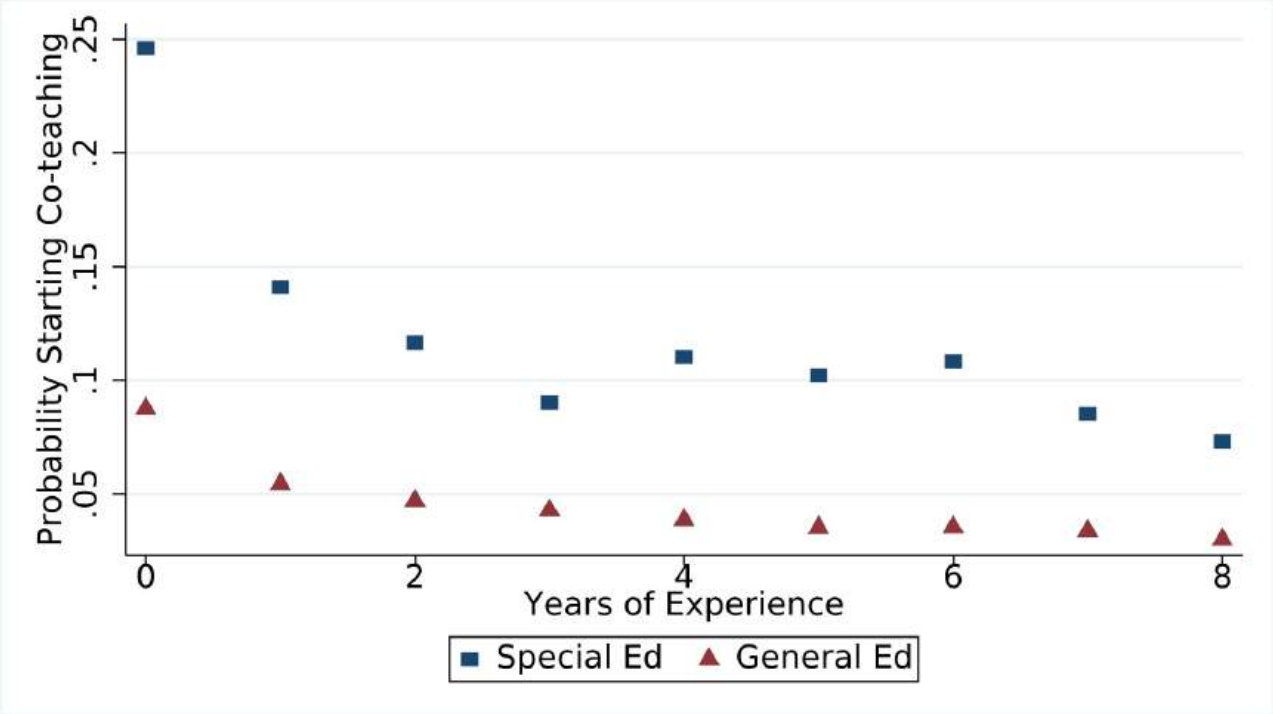

This figure shows the probability that a teacher begins co-teaching for the first time, conditional on their years of experience and not having co-taught previously, separately for special education and general education teachers.

Partnership dissolution rates are high. Over 60% of pairs separate after just one year together, and separation probability remains above 40% even after five years (Figure 2; see also Figures A.1 and A.2). These rates are consistent with enrollment fluctuations and scheduling changes shaping partnership formation and dissolution.

Figure 2: Hazard Rate for Co-teaching Pair Separation

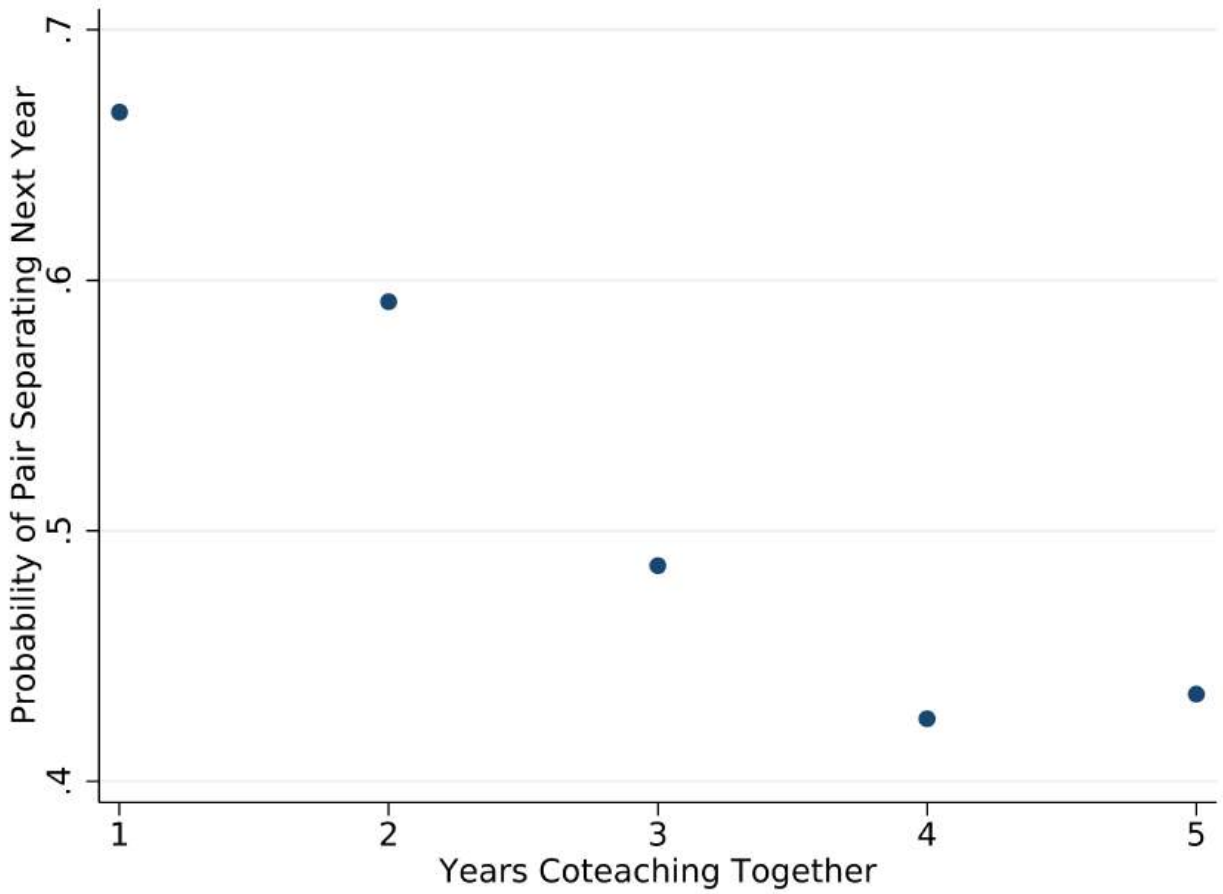

This figure shows the probability that co-teaching pairs separate in the following year, conditional on the number of years they have worked together as a pair.

General education and special education teachers follow different paths after partnerships end (Figures A.3 and A.4). General education teachers typically return to solo instruction (76%), while special education teachers cycle to new co-teaching partners (47%).

3.2 Student Assignment Changes Predictably

Co-teaching teachers receive different students than solo teachers. I test this using difference-in-differences comparing student assignment before and after co-teaching begins (Table 2).

General education teachers who co-teach receive 2 percentage points more students with

disabilities and 2.4 percentage points more students eligible for free or reduced-price lunch after beginning co-teaching. These changes are mechanical: co-taught classrooms exist specifically to serve students with disabilities, and disability status correlates with socioeconomic disadvantage. English Language Learner assignment does not change significantly.

I control for these observable student characteristics in all analyses. The event study design in Section 5 tests for parallel pre-trends, which would detect changes in unobservable student characteristics.

Table 2: Difference-in-Differences on Student Assignment by Co-Teaching Status

	Share Special Ed			Share FRPL		
	All	Gen Ed	Spec Ed	All	Gen Ed	Spec Ed
Ever Co-taught	0.008	-0.007^{**}	-0.090^{***}	0.025^{***}	0.022^{***}	0.002
	(0.005)	(0.002)	(0.011)	(0.004)	(0.004)	(0.009)
Ever Co-taught × Post	-0.006	0.020^{***}	-0.312^{***}	0.024^{***}	0.024^{***}	-0.018^{*}
	(0.005)	(0.002)	(0.012)	(0.004)	(0.004)	(0.009)
Observations	210,263	181,736	28,411	210,263	181,736	28,411
Dep. Var. Mean	0.232	0.159	0.695	0.507	0.489	0.622

Notes: This table uses difference-in-differences estimation to compare student characteristics assigned to teachers who ever co-teach versus those who never co-teach, before and after co-teaching begins. Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.3 Partner Assignment Shows Minimal Strategic Pairing

Administrators do not systematically pair experienced teachers with novices. I compare observed pairing frequencies to what would occur under random assignment (Figures A.5 and 3).

Novice-novice pairings are slightly less common than random assignment would predict, by 0.9 percentage points. But this deviation is small. Two inexperienced teachers form the most common pairing type, reflecting the mechanical fact that inexperienced teachers are most likely to co-teach. Current practice does not prioritize skill transfer in partnership

formation.

Figure 3: Difference From Random Pairing of Teachers

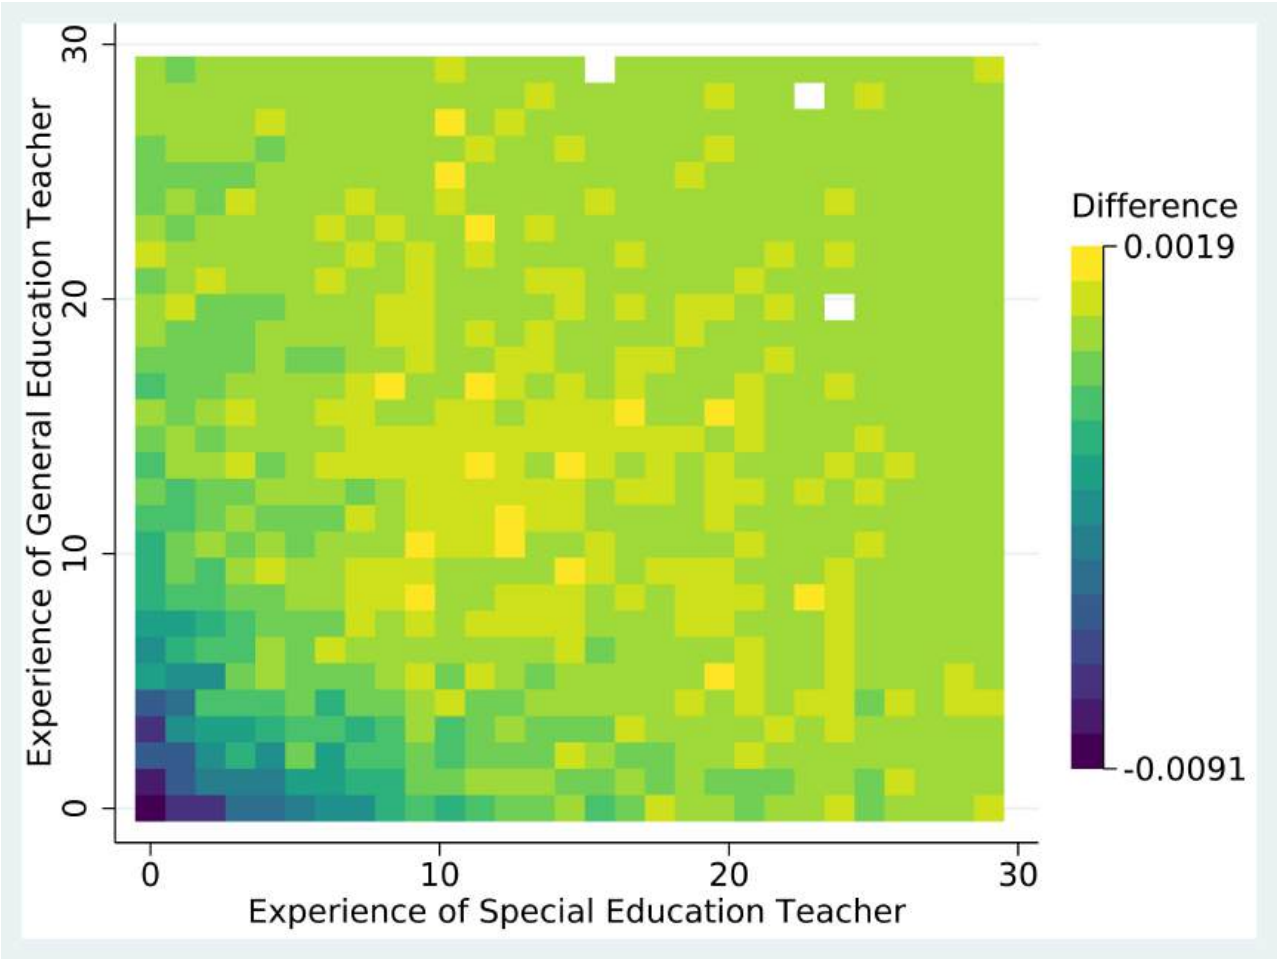

This heatmap displays the difference between observed co-teaching pairing frequencies by experience level and the frequencies that would be expected under random assignment of available teachers to partnerships. Figure A.5 shows the raw pairing frequencies.

4 Do Pairs Develop Match-Specific Skills?

Co-teaching could build match-specific capital or transferable individual skills. Match-specific capital makes a particular partnership productive but does not transfer when partners separate. Transferable skills persist after collaboration ends. I test for match-specific capital before examining individual skill development.

4.1 Longer-Lasting Pairs Are Not Better Matches

Partnership duration does not predict student achievement after controlling for teacher quality. I regress student test scores on partnership duration, with and without teacher fixed effects (Table 3).

Without fixed effects, longer-lasting pairs associate with higher test scores. This relationship disappears when I control for the individual teachers in the pair. Better teachers remain in partnerships longer. Partnerships do not improve through duration.

Table 3: Regression of Duration on Student Test Scores

	Student Test Scores		
	Combined	Math	English
Panel A: Without Teacher Fixed Effects			
Years Spent as Pair (Total)	0.00809^{*}	0.0125^{**}	0.00649
	(0.00365)	(0.00421)	(0.00470)
Observations	142,805	61,646	81,143
Panel B: With Teacher Fixed Effects			
Years Spent as Pair (Total)	−0.00455	-0.0313^{*}	0.000386
	(0.00688)	(0.0121)	(0.0105)
Observations	142,525	61,408	80,903

Notes: This table regresses student test scores on the total number of years a co-teaching pair worked together, with and without teacher fixed effects, to test whether longer-lasting pairs produce better student outcomes. Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.2 Pairs Do Not Improve Together Over Time

Teachers do not learn to work together more effectively over time. I test this using a two-step approach from Papay and Kraft (2015) that separates pair-specific learning from general experience effects.[2]

Without controlling for general teaching experience, student test scores increase with

[2]The single-step approach produces similar results (Figure A.6).

years together as a pair (Figure 4). This effect disappears entirely with experience controls. Teachers improve over time, but they improve equally whether that time is spent in the co-teaching pair or teaching solo.

Figure 4: Return To Experience in a Pair

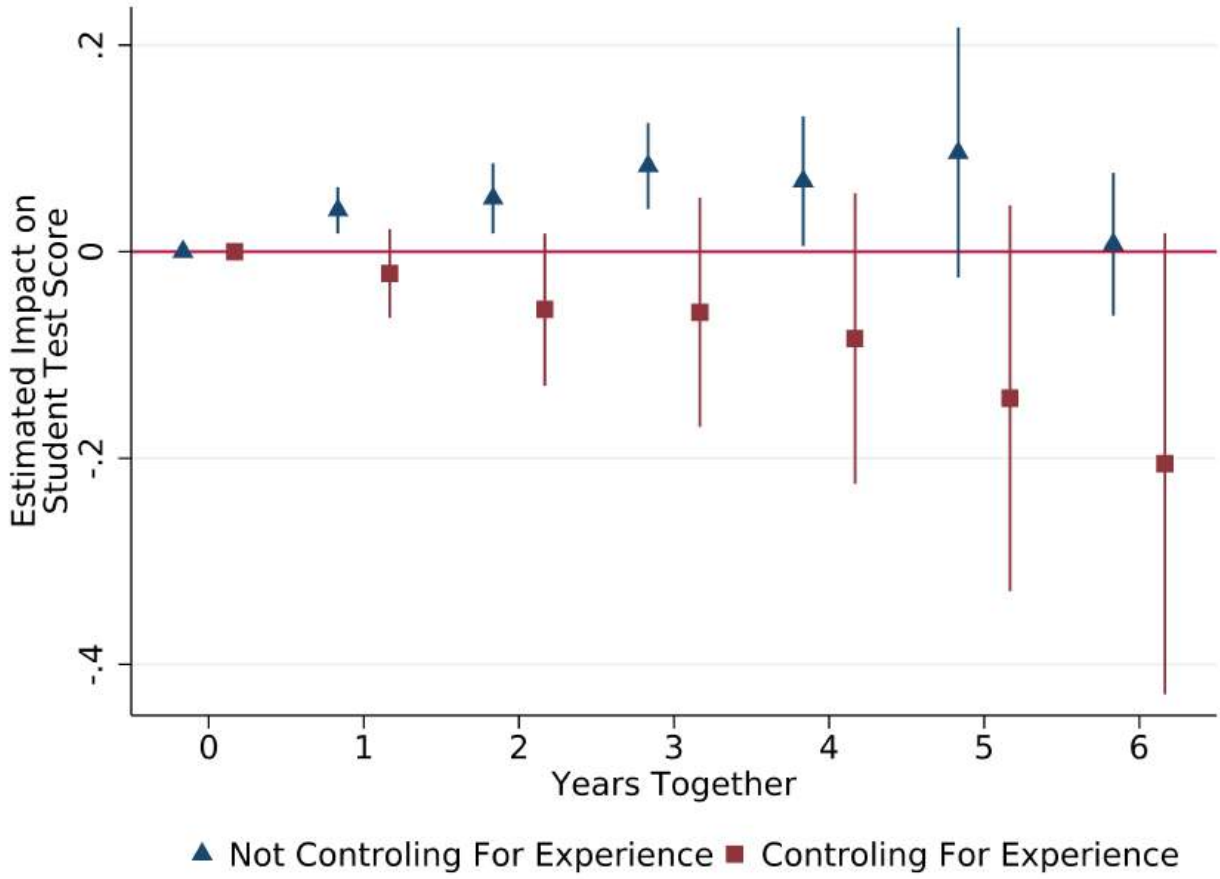

This figure plots student test scores against years of co-teaching partnership duration using a two-step estimation method to separate pair-specific learning from general teaching experience effects.

Co-teaching does not build match-specific capital. The next section tests whether it instead accelerates individual skill development that persists after partners separate.

5 Solo Teaching After Co-Teaching

5.1 Estimation Strategy

I compare teachers' student achievement before co-teaching to their performance after returning to solo instruction. I use a stacked event study design from Cengiz et al. (2019) that addresses staggered treatment timing. Traditional two-way fixed effects suffer from well-known biases when treatment effects vary over time (Goodman-Bacon, 2021). The stacked approach avoids comparing newly-treated units to already-treated units. A remaining concern is non-convex weighting, though methods addressing this typically cannot accommodate the concurrent controls necessary when student assignment varies (Wing et al., 2024). I examine sensitivity to alternative weighting schemes in Section 6.

I construct events by grouping teachers who share the same subject, year started teaching, year started co-teaching, and year left teaching. Within each event, treated teachers (who co-taught) are compared to control teachers with matching profiles who never co-taught. I restrict to teachers with 1-3 years of experience before co-teaching to ensure comparable pre-treatment periods.[3]

I estimate:

$$\pi_{dtpsy} = \sum_{i=-3}^{2} \omega_i YTC_{tyi} + \alpha X_{psy} + \delta_{dt} + \gamma_{dy} + \epsilon_{dtpsy} \tag{1}$$

where π_{dtpsy} is student p's test score in event d for teacher t in subject s in year y; YTC_{tyi} indicates years relative to first co-teaching; X_{psy} includes student demographics and a cubic polynomial for the previous two years' test scores; and δ_{dt} and γ_{dy} are event-specific teacher and year fixed effects.

The teacher fixed effects compare each teacher to themselves over time. The event-specific structure compares treated teachers only to never-treated teachers with similar career profiles.

[3]Teachers who first co-teach after 4+ years show different pre-trends, suggesting different selection into co-teaching (Figure A.7).

5.2 Average Effects

Teachers show higher student achievement after co-teaching than before (Figure 5). Prior to co-teaching, treated and control teachers follow parallel trends. The estimated differences in years -3 through -1 are small and statistically insignificant. After returning to solo instruction, teachers who co-taught show gains averaging 0.04σ in year 1. Rescaling to the teacher value-added distribution using middle school estimates from Chetty et al. (2014), this effect is equivalent to moving a teacher from the 50th to the 63rd percentile, implying approximately $20,000 in undiscounted lifetime earnings per student (in 2025 dollars). Year 2 estimates are similar but noisier.[4]

Figure 5: Difference in Student Test Scores By Time Since First Co-taught

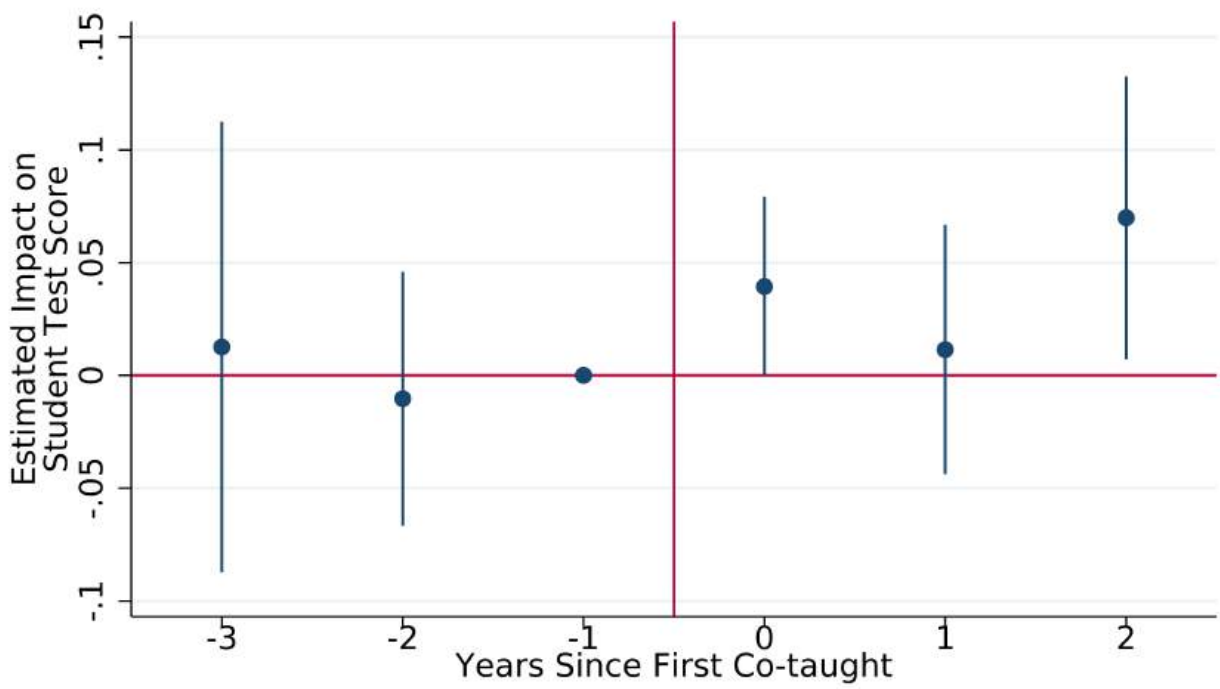

This figure presents the main event study results, showing the difference in student test scores for teachers who co-taught relative to those who did not, at various points before and after first co-teaching. The omitted category is one year before first co-teaching. Error bars represent 95% confidence intervals with standard errors clustered at the teacher level.

The parallel pre-trends support a causal interpretation. If co-teachers were on steeper improvement trajectories, we would see divergence before co-teaching began.

[4]Results are similar when controlling for current co-teaching status (Figure A.8).

5.3 Partner Experience

Partner experience explains most of the average effect. Figure 6 disaggregates results by co-teacher experience. Teachers paired with highly experienced partners (16+ years) show gains of approximately 0.10σ, equivalent to moving a teacher from the 50th to the 81st percentile of the value-added distribution and implying approximately $50,000 in undiscounted lifetime earnings per student (Chetty et al., 2014). Teachers paired with moderately experienced partners (8-15 years) show smaller, marginally significant gains. Teachers paired with inexperienced partners (0-7 years) show no benefit.

Figure 6: Difference in Student Test Scores By Time Since First Co-taught and Co-Teacher Experience

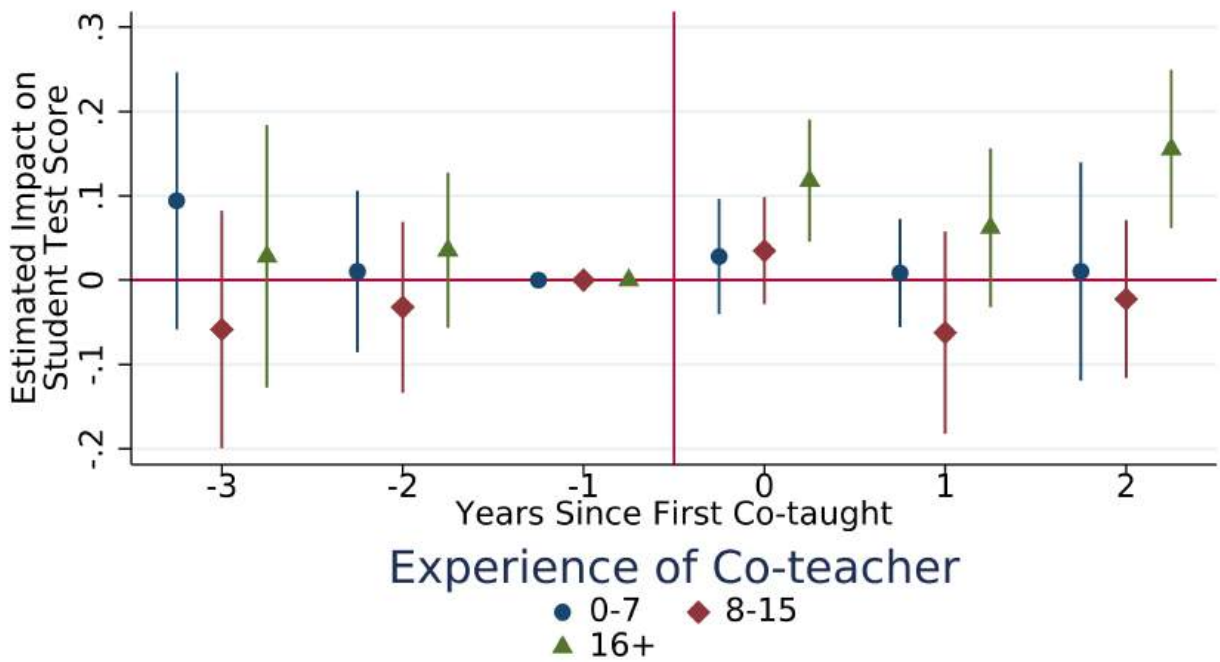

This figure disaggregates the main event study results by the experience level of the co-teaching partner, comparing effects across different partner experience categories.

This gradient is consistent with experienced teachers having accumulated knowledge that transfers through close collaboration. Partners must have something to offer. Pairing two novices produces no gains beyond what solo experience would provide.

I estimate separate difference-in-differences for finer experience bins to identify where

benefits are largest (Figure 7). Benefits increase with partner experience up to approximately 15-20 years, then level off. This pattern is robust to alternative functional form assumptions (Figure A.9).

Figure 7: Difference in Student Test Scores By Co-Teacher Experience

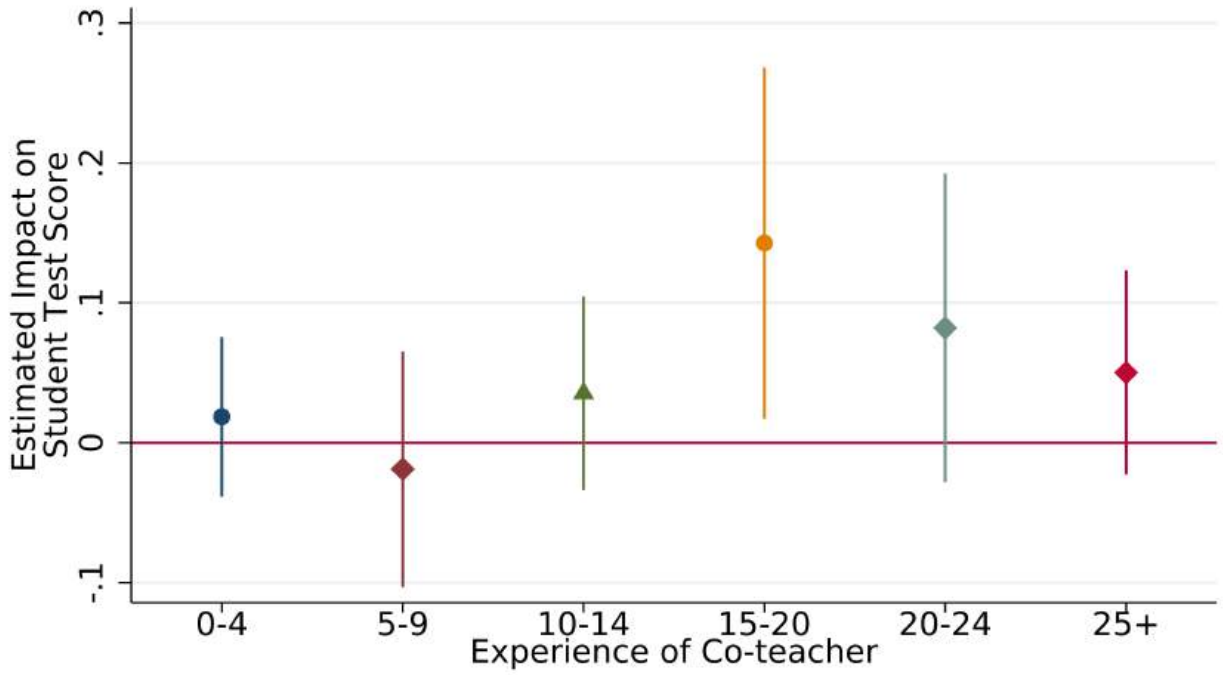

This figure shows the estimated impact of co-teaching on student test scores by the experience level of the co-teaching partner, using finer experience bins (5-year intervals). Each point represents a separate stacked difference-in-differences estimate. Error bars represent 95% confidence intervals with standard errors clustered at the teacher level.

The data do not reveal why benefits plateau after 15-20 years. This may be an artifact of statistical imprecision or a true phenomenon. One possibility is that pedagogical practices evolve, and teachers who developed their approaches decades ago may transfer less relevant techniques. Another is that very experienced teachers have automated their practice to the point where they struggle to articulate what they do.

5.4 Timing of Collaboration

Teachers benefit more from co-teaching if they have some prior solo experience. Figure 8 shows that teachers with 2-3 years of prior solo experience gain most from co-teaching.

First-year teachers gain less.

Figure 8: Difference in Student Test Scores By Experience When First Co-taught

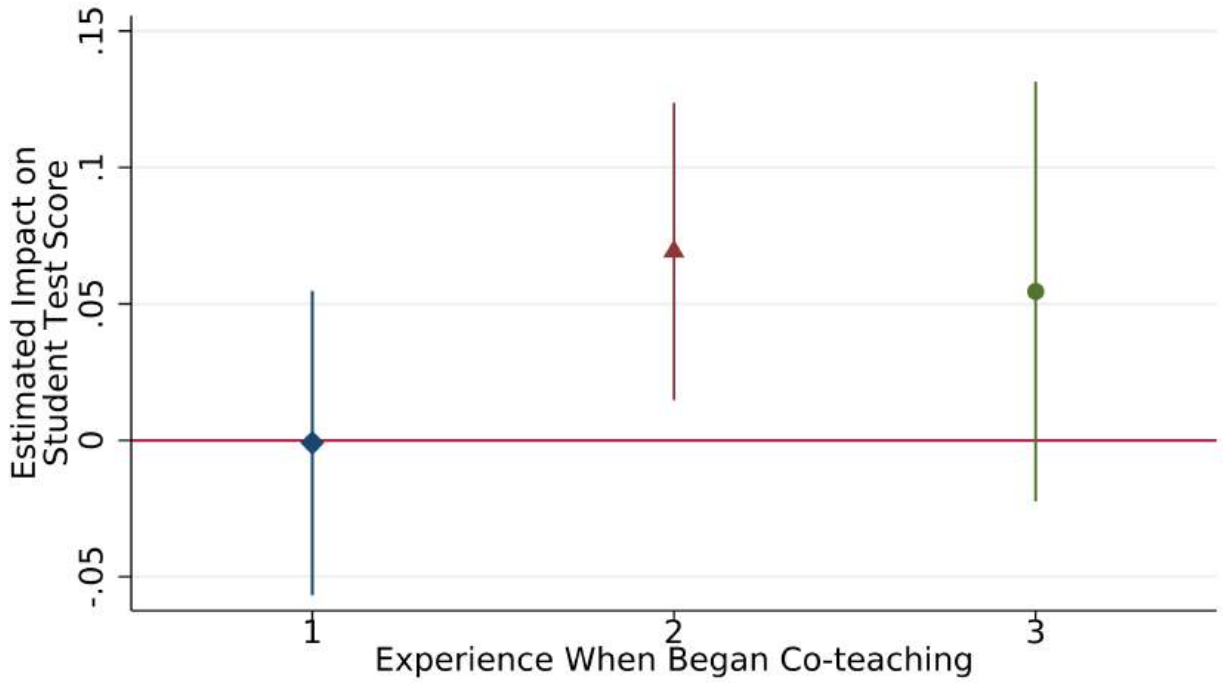

This figure examines heterogeneity in co-teaching effects based on when in a teacher's career they first experience co-teaching, comparing teachers with different amounts of prior solo teaching experience.

The data do not identify why first-year teachers benefit less. Perhaps some baseline competence helps teachers recognize and absorb what an experienced partner does. Complete novices may be too overwhelmed with basic classroom management to attend to subtle pedagogical techniques.

Partner education level does not predict gains (Figure A.10). Classroom experience, not formal credentials, correlates with what a partner can offer.

5.5 Duration of Collaboration

There is no evidence that longer collaborations produce larger benefits. I test whether skill transfer accumulates over time by estimating the marginal effect of each additional year of co-teaching (Figure 9).

Figure 9: Marginal Effect of Additional Years of Co-teaching

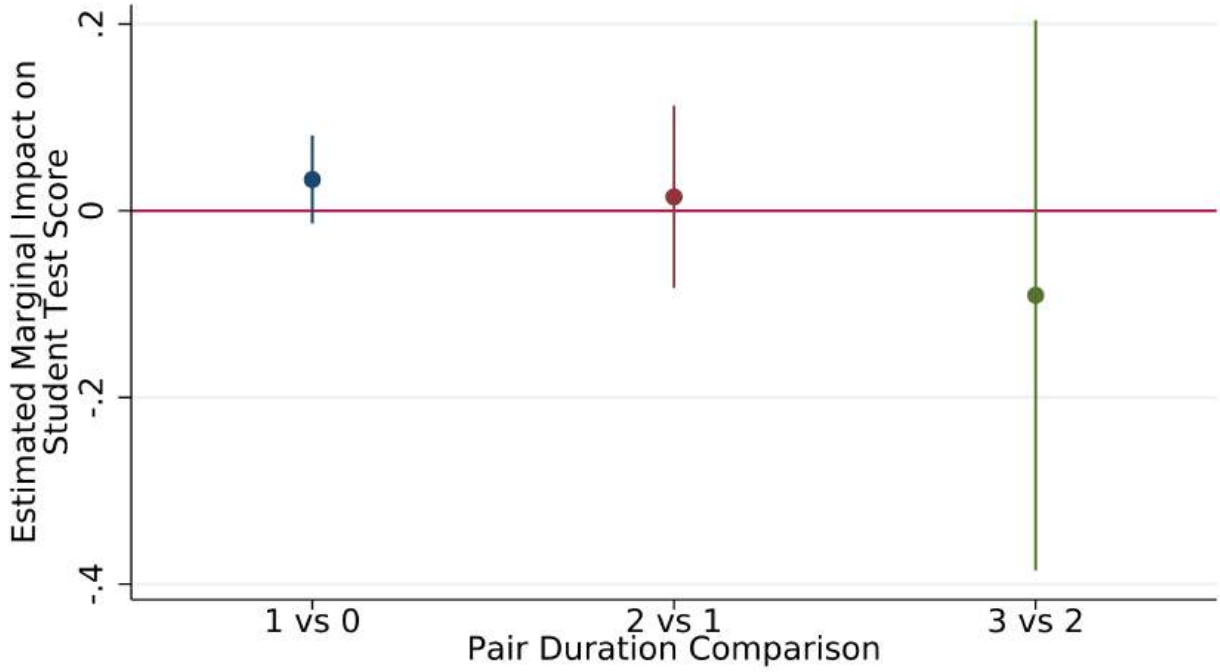

This figure presents stacked difference-in-differences estimates of the marginal effect of additional years of co-teaching on student test scores. Each point compares teachers who co-taught for n years to teachers who co-taught for $n-1$ years. Error bars represent 95% confidence intervals with standard errors clustered at the teacher level.

Additional years of co-teaching provide no observable incremental benefit. All estimates are statistically insignificant, but suggestively, the first year shows the largest effect. The second year adds nothing. The third year point estimate is negative, though confidence intervals are wide due to limited power.

This pattern is consistent with the finding in Section 4 that pairs do not improve together over time once general experience is controlled. Extended co-teaching may substitute for solo teaching experience. A teacher who spends three years co-teaching gains three years of general experience, but so does a teacher who co-teaches for one year and then teaches solo for two years. If the benefits of collaboration come from observing a partner's techniques rather than from cumulative practice together, the teacher with more solo experience may develop faster.

6 Robustness

6.1 Prior Student Achievement

Controlling for students' prior test scores is necessary to detect the positive effects of co-teaching. Without these controls, there is no significant effect (Figure A.11). This pattern reflects changes in student assignment. Teachers receive more students with disabilities and from disadvantaged backgrounds after co-teaching begins (Table 2). These students have lower expected achievement growth on average, so failing to control for prior scores biases estimates downward. Accurate productivity measurement requires conditioning on student inputs (Chetty et al., 2014). Using only one year of lagged scores rather than two produces similar results (Figure A.12).

6.2 Sample Restrictions

I restrict the main analysis to teachers with 1-3 years of experience before co-teaching. Teachers with 4+ years before co-teaching show different patterns, including negative pre-trends (Figure A.7). General education teachers are also less likely to begin co-teaching after many years (Figure 1), so administrators may observe characteristics of these late-entry teachers that I cannot measure in the administrative data. Including all teachers with 1+ years of prior experience slightly attenuates the estimated benefits but the qualitative pattern remains (Figure A.13).

6.3 Event Weighting

The stacked difference-in-differences implicitly weights events, which could affect estimates if treatment effects vary. Table A.1 compares three approaches: implicit weights from the stacked methodology (0.031σ), equal weights across events (0.017σ), and weights proportional to observations per event (0.055σ). This variation reflects heterogeneity in event-specific effects, but results are qualitatively similar regardless of approach.

7 Policy Simulation

Pairing early-career teachers with highly experienced colleagues produces lasting productivity gains, but implementing such a policy involves costs. The experienced teacher must leave their solo classroom, and a less effective instructor may take their place. I simulate this tradeoff to determine when strategic pairing yields net positive effects on student achievement.

7.1 Setup and Assumptions

Each year, teachers with at least two years of experience pair with highly experienced teachers (16+ years) for a single year of co-teaching. The number of pairings is constrained by the supply of experienced mentors. After co-teaching ends, both teachers return to solo instruction.

The development effect ($\delta = 0.10\sigma$) comes from my main estimates. The contemporaneous effect ($\beta = 0.02\sigma$) captures the additional benefit to students from being in a classroom with two teachers present. This estimate comes from Jones and Winters (2024). The opportunity cost ($\gamma = 0.14\sigma$) reflects the productivity gap between novice and experienced teachers estimated by Papay and Kraft (2015).

I assume 20% annual turnover, which implies approximately 2.9% of teachers have 16+ years of experience in steady state. My analysis suggests that co-teaching reduces turnover, which I account for among treated teachers in the baseline scenario (See Appendix Table A.2). Appendix B provides additional detail on all assumptions.

7.2 Simulation Logic

The simulation tracks how strategic pairing affects the teaching workforce over time. Each year, the policy removes experienced teachers from solo classrooms to co-teach with early-career colleagues. This creates an immediate opportunity cost because the experienced teacher's former classroom must be staffed by a novice replacement.

Benefits accumulate gradually while costs occur immediately. Treated teachers gain lasting skills and work for many years with enhanced productivity. Each mentor reassignment, by contrast, creates only a one-year cost. The simulation measures when total accumulated benefits exceed total accumulated costs.

The retention benefit creates a compounding effect. As treated teachers stay longer, more eventually reach 16+ years of experience and become mentors themselves. This expands the mentor pool and allows more pairings.

The simulation also calculates a break-even opportunity cost (γ^*). This is the threshold above which the policy would be harmful. Comparing γ^* to existing estimates indicates how much margin for error the policy has.

These estimates are conservative in one important respect. The simulation only credits the policy with the direct treatment effect and the expanded mentor pool. It does not account for the additional productivity gains that occur when treated teachers survive to higher experience levels and benefit from natural returns to experience (Papay and Kraft, 2015). Including this channel would make the policy estimates more favorable.

7.3 Results

Table 4 presents results for four scenarios that progressively remove optimistic assumptions. The baseline includes all estimated benefits. Columns (2) through (4) remove the turnover reduction, contemporaneous effect, and add skill decay, respectively.

Table 4: Policy Simulation Results Across Scenarios

	(1) Full Model (Baseline)	(2) No Turnover Reduction	(3) No Contemp. Effect	(4) Most Conservative
Panel A: Assumptions				
Turnover reduction (6pp)	Yes	No	No	No
Contemporaneous effect	Yes	Yes	No	No
No skill decay	Yes	Yes	Yes	No
Panel B: Steady-State Outcomes				
Mentor pool (% of teachers)	7.8%	2.9%	2.9%	2.9%
Share of teachers treated	49.8%	14.4%	14.4%	14.4%
Panel C: Annual Effects (average school wide)				
Treatment effect on VA	$+0.042\sigma$	$+0.012\sigma$	$+0.012\sigma$	$+0.008\sigma$
Contemporaneous benefit	$+0.002\sigma$	$+0.001\sigma$	—	—
Opportunity cost	-0.011σ	-0.004σ	-0.004σ	-0.004σ
Net annual effect	$+0.033\sigma$	$+0.008\sigma$	$+0.007\sigma$	$+0.004\sigma$
Lifetime earnings per student	$+\$16,400$	$+\$4,000$	$+\$3,400$	$+\$1,900$
Panel D: Cost-Benefit Analysis (per classroom)				
Cost	$4,680	$1,740	$1,740	$1,740
Societal benefit	$328,000	$80,000	$68,000	$38,000
Government benefit (15%)	$49,200	$12,000	$10,200	$5,700

Notes: This table presents steady-state effects of strategic teacher pairing under four scenarios. Column (1) includes all estimated benefits. Columns (2) through (4) progressively remove the turnover reduction, contemporaneous effect, and add skill decay. Skill decay (column 4 only) means the 0.10σ development effect diminishes by 10% each year post-treatment. Annual effects are measured per average classroom. Lifetime earnings per student are undiscounted gains in 2025 dollars per year of exposure, calculated by converting the net annual effect to teacher value-added units and applying earnings estimates from Chetty et al. (2014), inflation-adjusted from 2010. Cost per classroom equals the mentor pool share times $60,000 (annual teacher salary). Societal benefit equals lifetime earnings per student times 20 students. Government benefit assumes 15% of societal benefit is captured through taxation.

The retention benefit drives the largest difference between scenarios. When treated teachers stay longer, the mentor pool expands from 2.9% to 7.8% of teachers, and nearly half of all teachers eventually receive treatment. Without this compounding effect, only 14% of teachers are ever treated. The net annual effect ranges from 0.004σ in the most conservative scenario to 0.033σ in the baseline.

Cumulative effects turn positive within six to eight years across all scenarios. This happens because a treated teacher benefits every subsequent cohort of students, while each

pairing creates only a one-year opportunity cost. The asymmetry favors the policy when the development effect is durable.

The policy has substantial margin for error. The break-even opportunity cost ranges from 0.29σ to 0.56σ across scenarios. These thresholds are two to four times the Papay and Kraft (2015) estimate of 0.14σ. The true opportunity cost in a given context could differ from existing estimates, but the policy would remain beneficial unless experienced teachers are far more productive than current research suggests.

The earnings implications are substantial. The policy generates between \$1,900 and \$16,400 in undiscounted lifetime earnings per student per year of exposure. These estimates convert the net annual effect to teacher value-added units and apply earnings estimates from Chetty et al. (2014), inflation-adjusted to 2025 dollars.

7.4 Cost-Benefit Analysis

The policy requires hiring additional teachers to cover classrooms vacated by experienced mentors. Assuming an annual salary of \$60,000, the cost per classroom equals the mentor pool share times \$60,000. This yields \$4,680 in the baseline scenario and \$1,740 in the conservative scenarios.

The benefits far exceed the costs. The societal benefit per classroom equals lifetime earnings gains times 20 students, ranging from \$38,000 to \$328,000. Assuming the government captures 15% through taxation, the government benefit ranges from \$5,700 to \$49,200 per classroom. In all scenarios, government benefits substantially exceed costs. The baseline implies a benefit-cost ratio of approximately 10:1.

8 Conclusion

Workplace collaboration can produce lasting skill development. Teachers who co-taught with experienced colleagues show higher student achievement after returning to solo instruction. The effect averages 0.04σ overall and reaches 0.10σ when paired with highly

experienced partners. Rescaling to the teacher value-added distribution (Chetty et al., 2014), these effects correspond to moving a teacher from the 50th to the 63rd percentile (average effect) or from the 50th to the 81st percentile (highly experienced partners), implying \$20,000 to \$50,000 in undiscounted lifetime earnings per student. This pattern is consistent with knowledge transfer from experienced to inexperienced teachers. Formal credentials do not predict learning gains. Only experience does.

Three findings inform collaboration design. First, partner selection matters. Pairing with experienced colleagues generates lasting benefits. Pairing two novices together produces none. Second, collaboration can be brief. One year appears sufficient, and extended partnerships may reduce development by substituting for solo experience. Third, timing matters. Teachers with 2-3 years of independent experience benefit most. Some baseline competence may be necessary to absorb what experienced partners offer.

These findings suggest that organizations underinvest in strategic collaboration. Current practice allows scheduling constraints rather than developmental objectives to shape partnership formation. The policy simulation indicates that pairing experienced workers with junior colleagues can generate annual gains of 0.033σ per classroom under the baseline assumptions. Benefits remain positive even under conservative assumptions about skill decay and auxiliary effects.

Several limitations warrant caution. The estimates carry substantial uncertainty, and confidence intervals are wide for some heterogeneity analyses. Co-teaching represents a specific form of collaboration that may not generalize to other settings. I cannot observe mechanisms directly. The knowledge transfer interpretation is consistent with the evidence but remains unproven. Results also depend on controlling for prior student achievement. This is standard in the teacher productivity literature but nonetheless a specification choice.

Collaboration's value may lie not just in what partners produce together, but in how working together changes what each can produce alone. Organizations that pair the right workers at the right time for the right duration may be able to accelerate workforce development. Current practice appears to leave many of these potential gains unrealized.

References

Andersen, S. C., Beuchert, L., Nielsen, H. S., and Thomsen, M. K. (2020). The effect of teacher's aides in the classroom: Evidence from a randomized trial. *Journal of the European Economic Association*, 18(1):469–505.

Arrow, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies*, 29(3):155–173.

Austin, V. L. (2001). Teachers' beliefs about co-teaching. *Remedial and Special Education*, 22(4):245–255.

Azoulay, P., Graff Zivin, J. S., and Wang, J. (2010). Superstar extinction. *The Quarterly Journal of Economics*, 125(2):549–589.

Ballis, B. and Heath, K. (2021). The long-run impacts of special education. *American Economic Journal: Economic Policy*, 13(4):72–111.

Battu, H., Belfield, C. R., and Sloane, P. J. (2003). Human capital spillovers within the workplace: Evidence for great britain. *Oxford Bulletin of Economics and Statistics*, 65(5):575–594.

Bessette, H. J. (2008). Using students' drawings to elicit general and special educators' perceptions of co-teaching. *Teaching and Teacher Education*, 24(5):1376–1396.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Chen, Y. (2021). Team-specific human capital and team performance: evidence from doctors. *American Economic Review*, 111(12):3923–3962.

Chetty, R., Friedman, J. N., and Rockoff, J. E. (2014). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American economic review*, 104(9):2593–2632.

Conderman, G. and Hedin, L. R. (2014). Co-teaching with strategy instruction. *Intervention in School and Clinic*, 49(3):156–163.

Cook, S. C., McDuffie-Landrum, K. A., Oshita, L., and Cook, B. G. (2017). Co-teaching for students with disabilities: A critical and updated analysis of the empirical literature. In *Handbook of Special Education*, pages 233–248. Routledge.

De Backer, L., Schelfhout, W., Simons, M., and Vandervieren, E. (2023). Student teachers' peer team teaching experiences from a quantitative perspective: Perceptions, profiles and transition probabilities. *Teaching and Teacher Education*, 135:104361.

De Grip, A. and Sauermann, J. (2012). The effects of training on own and co-worker productivity: Evidence from a field experiment. *The Economic Journal*, 122(560):376–399.

Gershenson, S., Lindsay, C. A., Papageorge, N. W., Campbell, R. A., and Rendon, J. H. (2023). Spillover effects at school: How black teachers affect their white peers' racial competency. Technical report, National Bureau of Economic Research.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Guryan, J., Kroft, K., and Notowidigdo, M. J. (2009). Peer effects in the workplace: Evidence from random groupings in professional golf tournaments. *American Economic Journal: Applied Economics*, 1(4):34–68.

Haggag, K., McManus, B., and Paci, G. (2017). Learning by driving: Productivity improvements by new york city taxi drivers. *American Economic Journal: Applied Economics*, 9(1):70–95.

Ichino, A. and Maggi, G. (2000). Work environment and individual background: Explaining regional shirking differentials in a large italian firm. *The Quarterly Journal of Economics*, 115(3):1057–1090.

Jackson, C. K. and Bruegmann, E. (2009). Teaching students and teaching each other: The importance of peer learning for teachers. *American Economic Journal: Applied Economics*, 1(4):85–108.

Jones, N., Kaler, L., Markham, J., Senese, J., and Winters, M. A. (2025). Service delivery models: Impacts for students with and without disabilities. *Educational Researcher*, 54(3):141–152.

Jones, N. and Winters, M. A. (2024). Are two teachers better than one?: The effect of coteaching on students with and without disabilities. *Journal of Human Resources*, 59(4):1180–1206.

Kraft, M. A., Blazar, D., and Hogan, D. (2018). The effect of teacher coaching on instruction and achievement: A meta-analysis of the causal evidence. *Review of Educational Research*, 88(4):547–588.

Mas, A. and Moretti, E. (2009). Peers at work. *American Economic Review*, 99(1):112–145.

Nix, E. (2020). Learning spillovers in the firm. Technical report, Working paper.

Ost, B. (2014). How do teachers improve? the relative importance of specific and general human capital. *American Economic Journal: Applied Economics*, 6(2):127–151.

Papay, J. P. and Kraft, M. A. (2015). Productivity returns to experience in the teacher labor market: Methodological challenges and new evidence on long-term career improvement. *Journal of Public Economics*, 130:105–119.

Papay, J. P., Taylor, E. S., Tyler, J. H., and Laski, M. E. (2020). Learning job skills from colleagues at work: Evidence from a field experiment using teacher performance data. *American Economic Journal: Economic Policy*, 12(1):359–388.

Shaw, K. and Lazear, E. P. (2008). Tenure and output. *Labour Economics*, 15(4):704–723.

Thompson, P. (2010). Learning by doing. *Handbook of the Economics of Innovation*, 1:429–476.

Thompson, P. (2012). The relationship between unit cost and cumulative quantity and the evidence for organizational learning-by-doing. *Journal of Economic Perspectives*, 26(3):203–224.

Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked difference-in-differences. Technical report, National Bureau of Economic Research.

A Appendix

Figure A.1: Hazard Rate for Continuing Co-Teaching

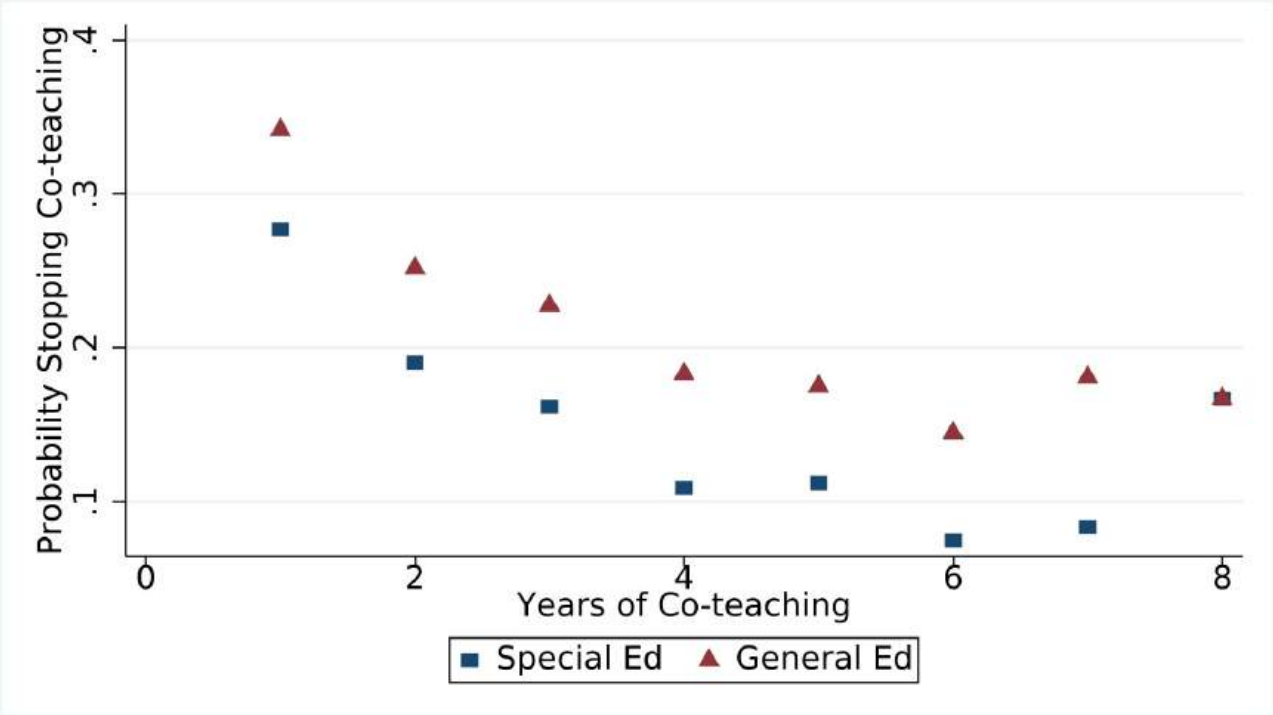

This figure displays the probability that a teacher stops co-teaching in the following year, conditional on co-teaching in the current year and continuing to teach, by teacher type and co-teaching experience.

Figure A.2: Hazard Rate for Permanent Co-teaching Pair Separation

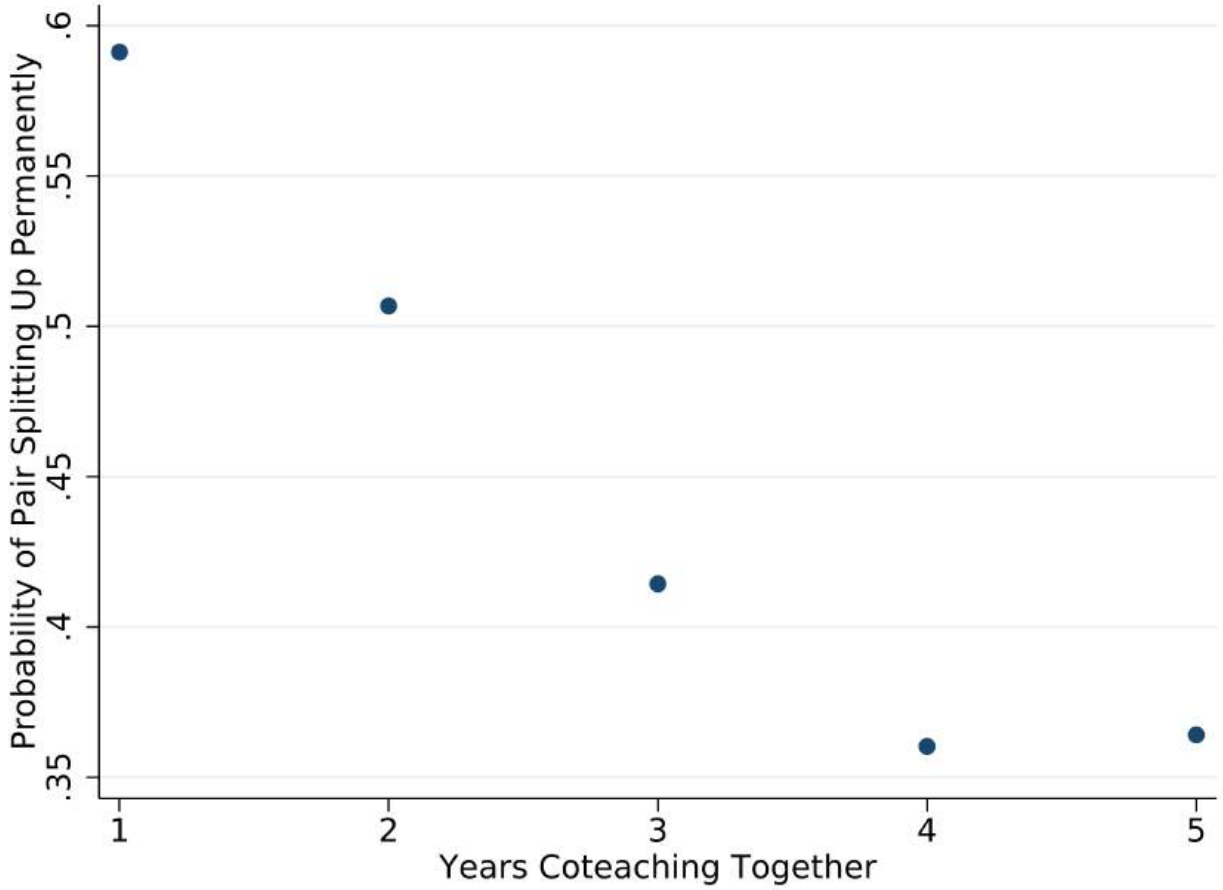

This figure shows the probability that co-teaching pairs never work together again (permanent separation) rather than temporary separation followed by potential reunion.

Figure A.3: What General Education Teachers Do After Co-Teaching

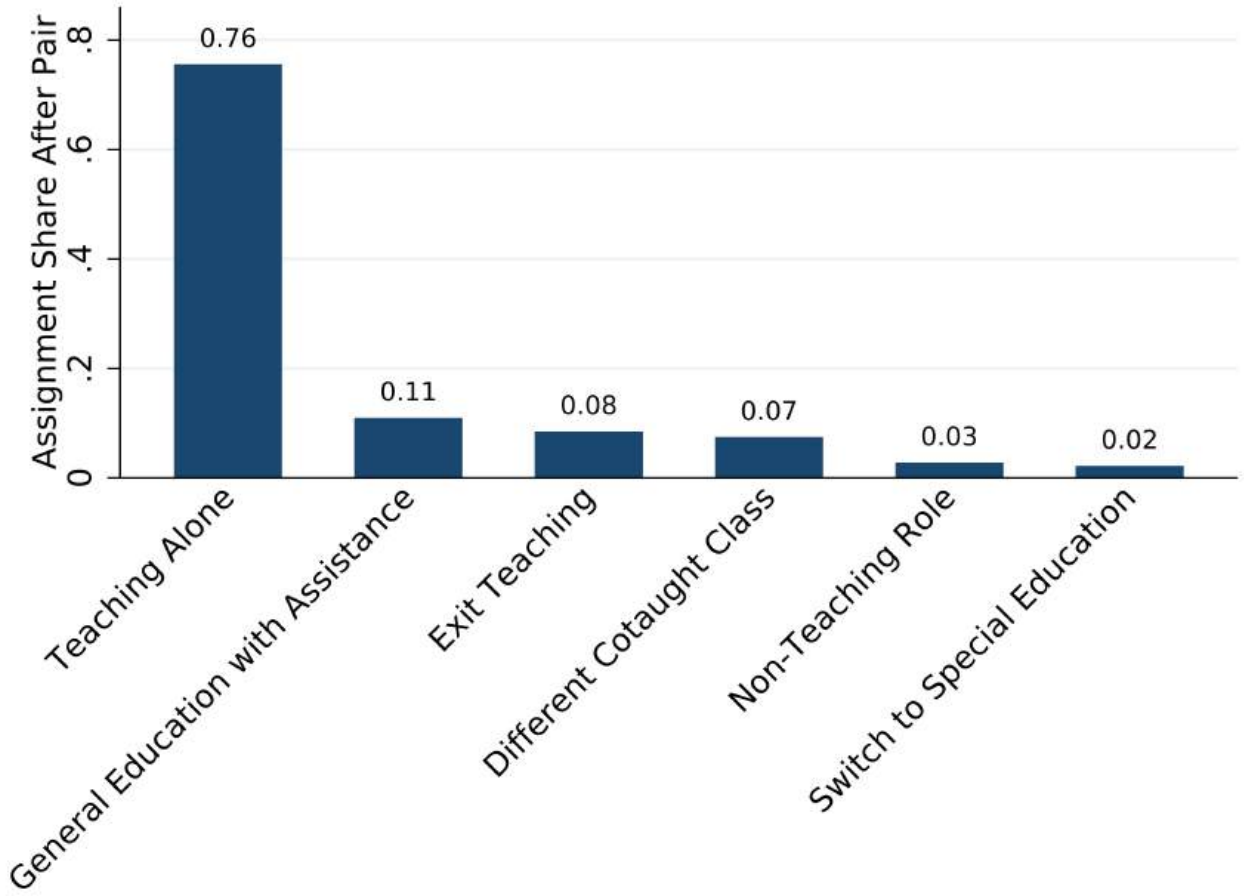

This figure shows the distribution of classroom assignments for general education teachers in the year immediately following the end of a co-teaching partnership.

Figure A.4: What Special Education Teachers Do After Co-Teaching

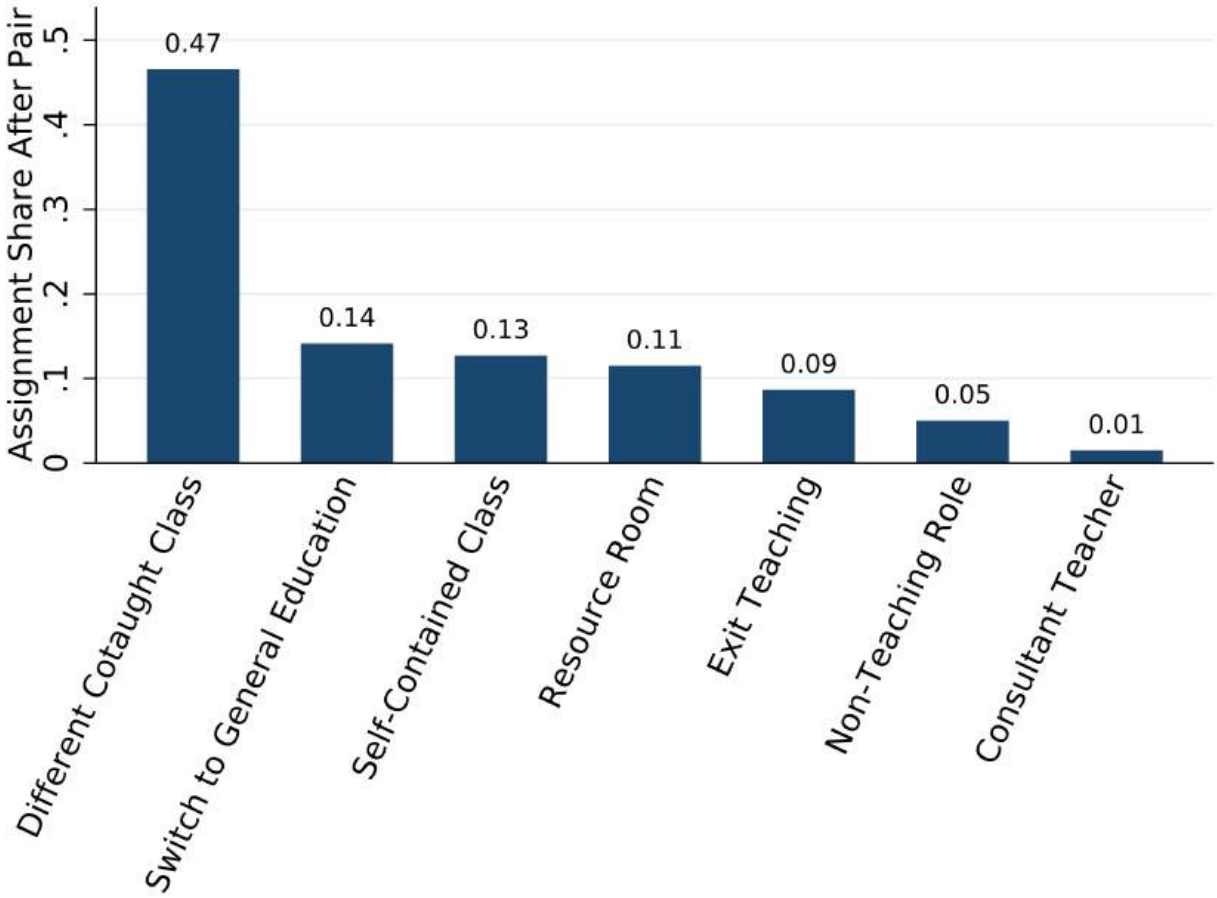

This figure displays the distribution of classroom assignments for special education teachers in the year immediately following the end of a co-teaching partnership.

Figure A.5: Pairing Frequency by Experience

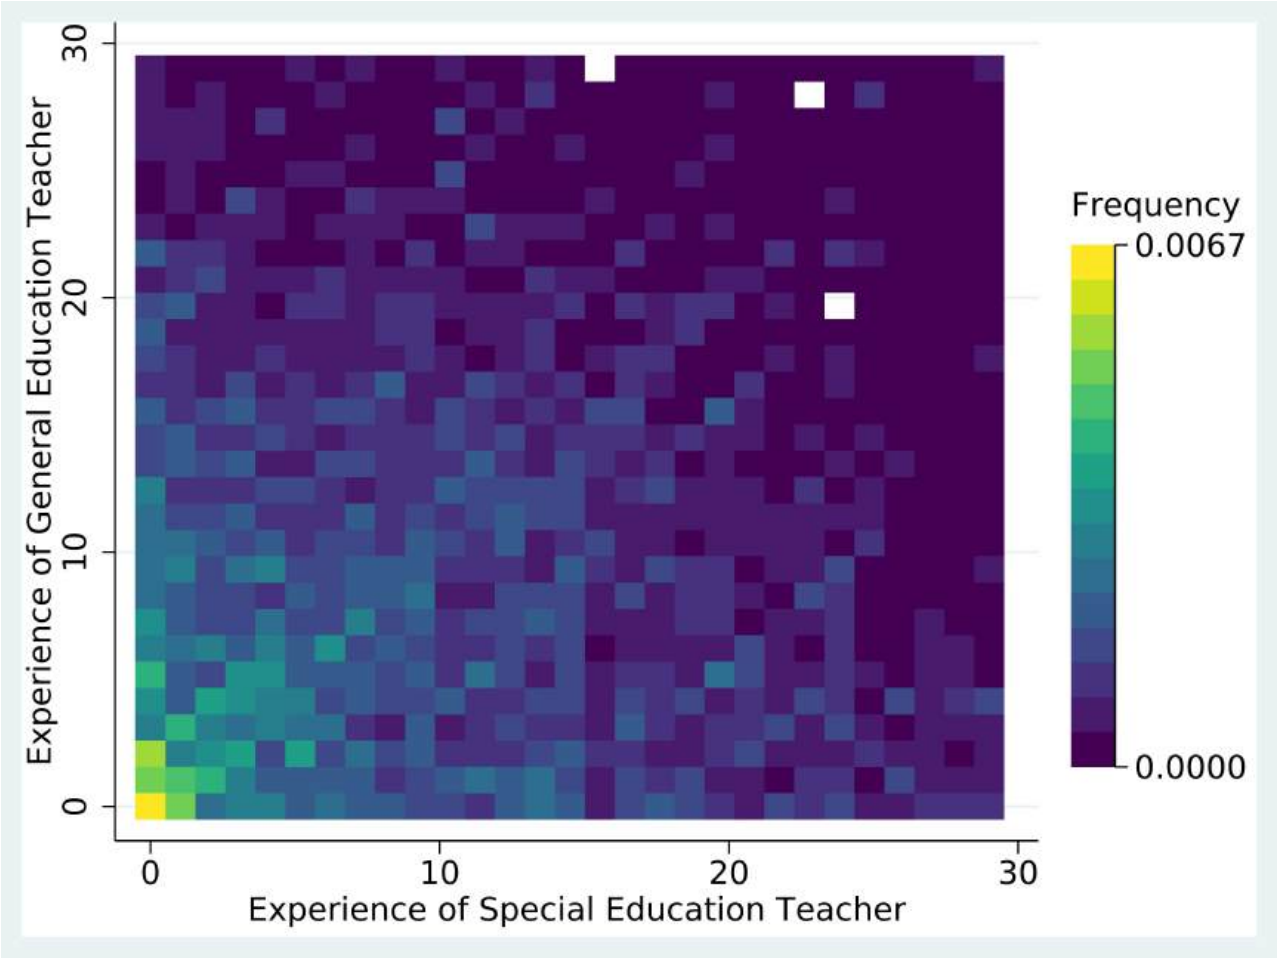

This heatmap shows the observed frequency of co-teaching pairs based on the experience levels (in years) of both the special education and general education teachers in the partnership. Figure 3 shows the difference between these observed frequencies and what would be expected under random assignment.

Figure A.6: Return To Experience in a Pair (Estimated in One Step)

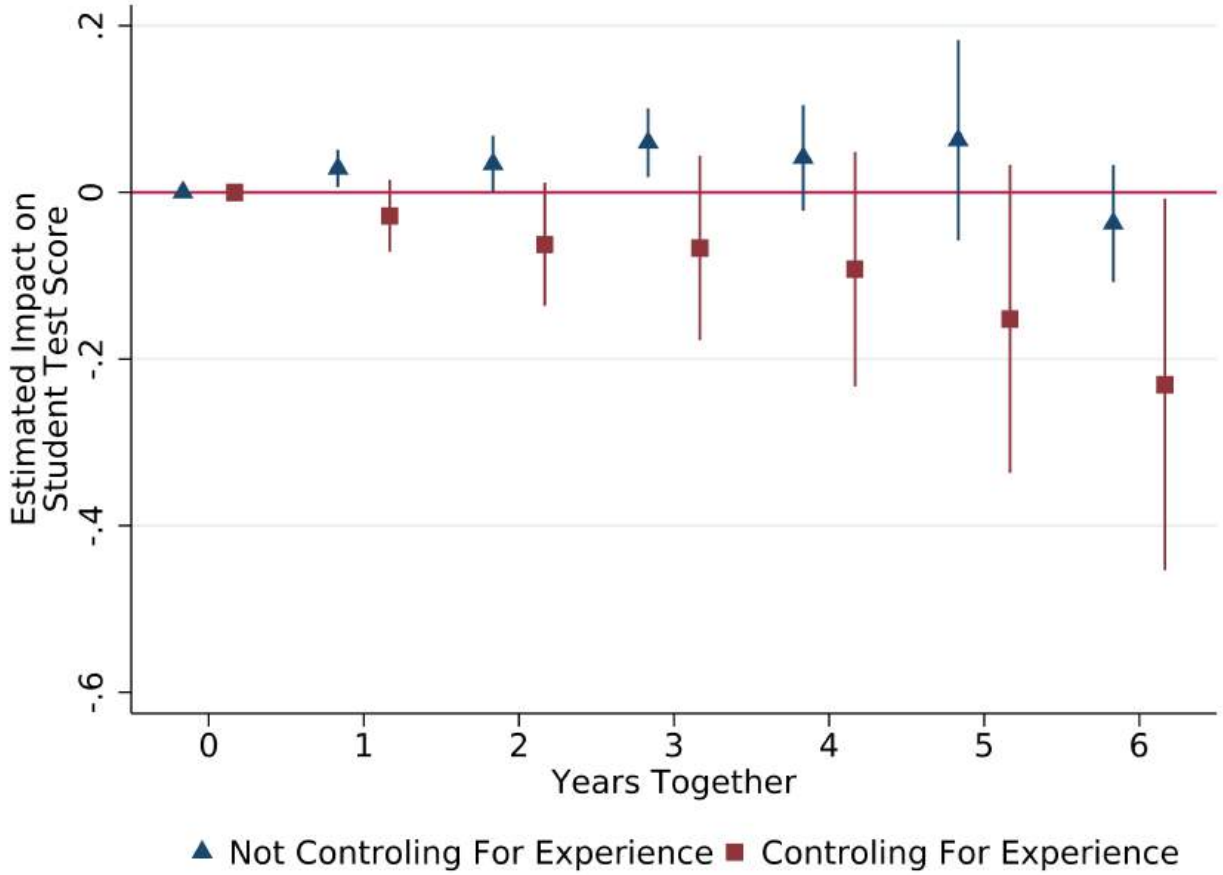

This figure presents the same analysis as Figure 4 but using a single-step estimation method rather than the two-step approach, as a robustness check on the methodology.

Figure A.7: Difference in Student Test Scores By Time Since First Co-taught (Only Using More Than Three Years of Prior Experience)

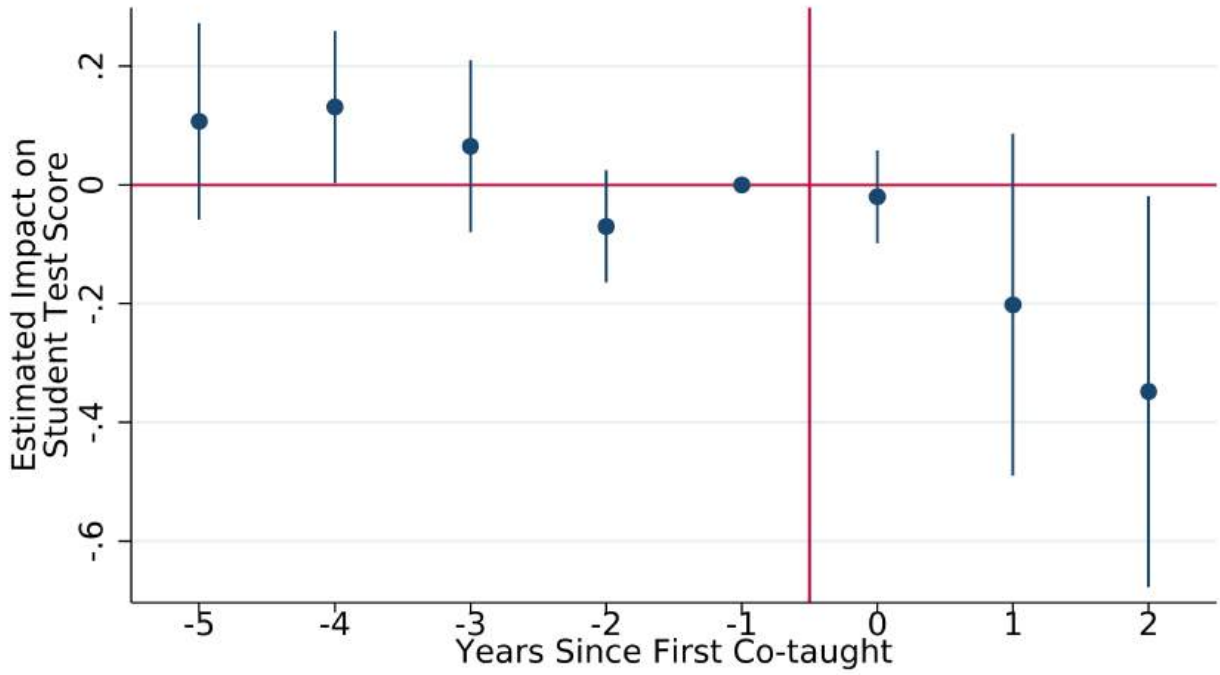

This figure restricts the sample to teachers who had more than three years of solo teaching experience before first co-teaching, testing the sensitivity of results to sample composition.

Figure A.8: Difference in Student Test Scores By Time Since First Co-taught (Controlling For Current Co-teaching)

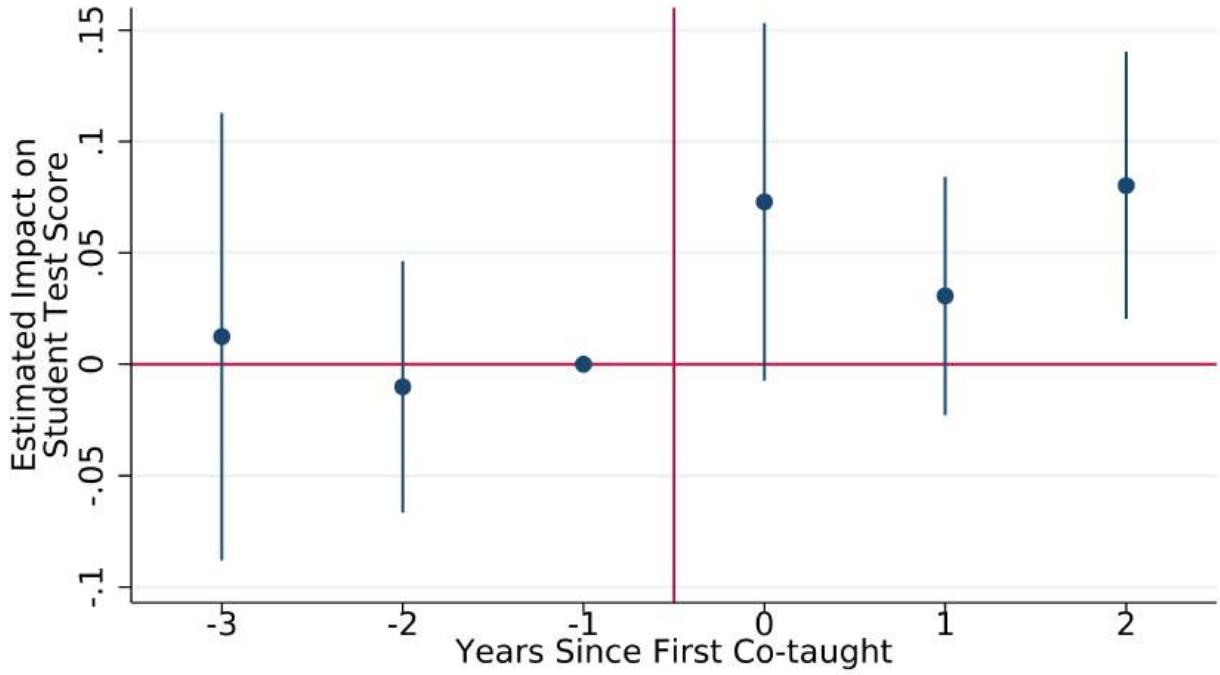

This specification adds indicator variables for whether teachers are currently co-teaching to separate persistent effects from contemporaneous collaboration benefits.

Figure A.9: Co-Teacher Experience Effects Under Alternative Functional Forms

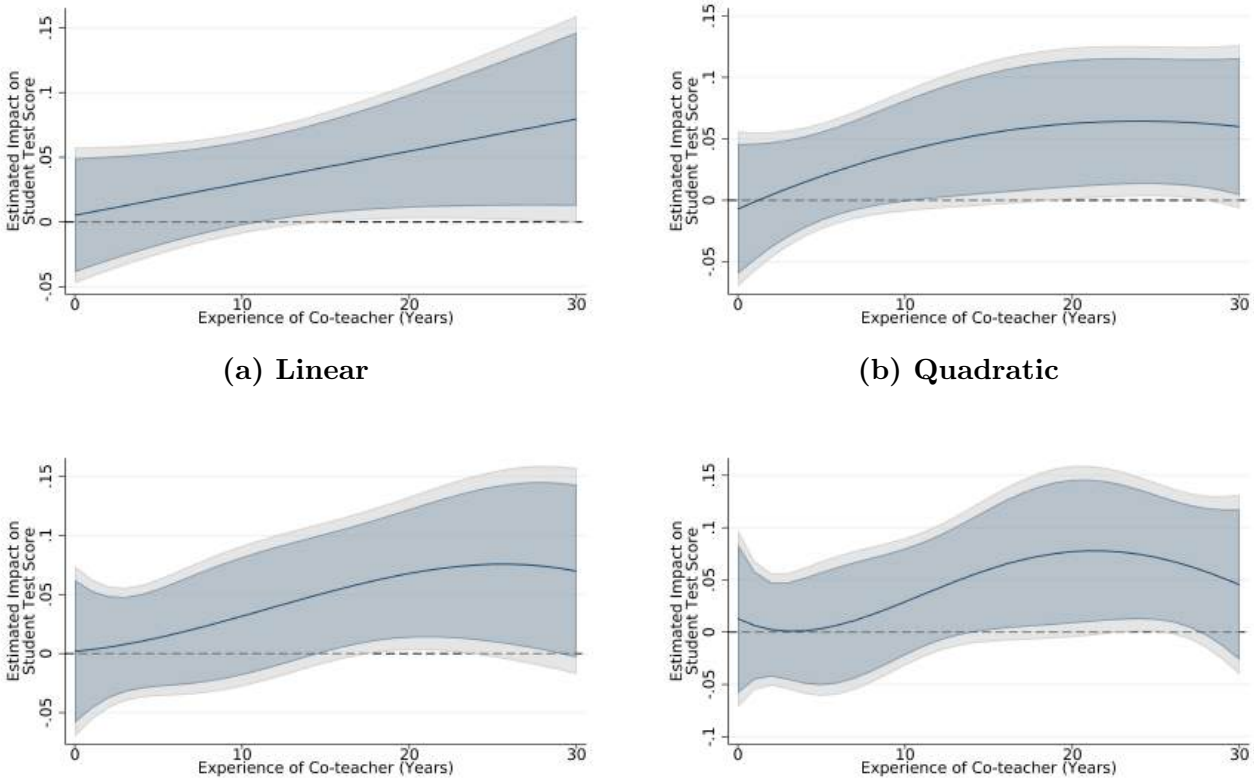

(a) Linear **(b) Quadratic**

(c) Cubic **(d) Quartic**

This figure presents the relationship between co-teacher experience and student test score improvements under alternative polynomial specifications. The polynomial specifications interact treatment status (post-co-teaching) with polynomials of co-teacher experience, with shaded regions representing 90% (darker) and 95% (lighter) confidence intervals. All specifications include the full set of student and classroom controls, event-specific teacher and year fixed effects, and cluster standard errors at the teacher level. The consistent pattern across specifications suggests this relationship is robust to functional form assumptions.

Figure A.10: Difference in Student Test Scores By Co-teacher Degree

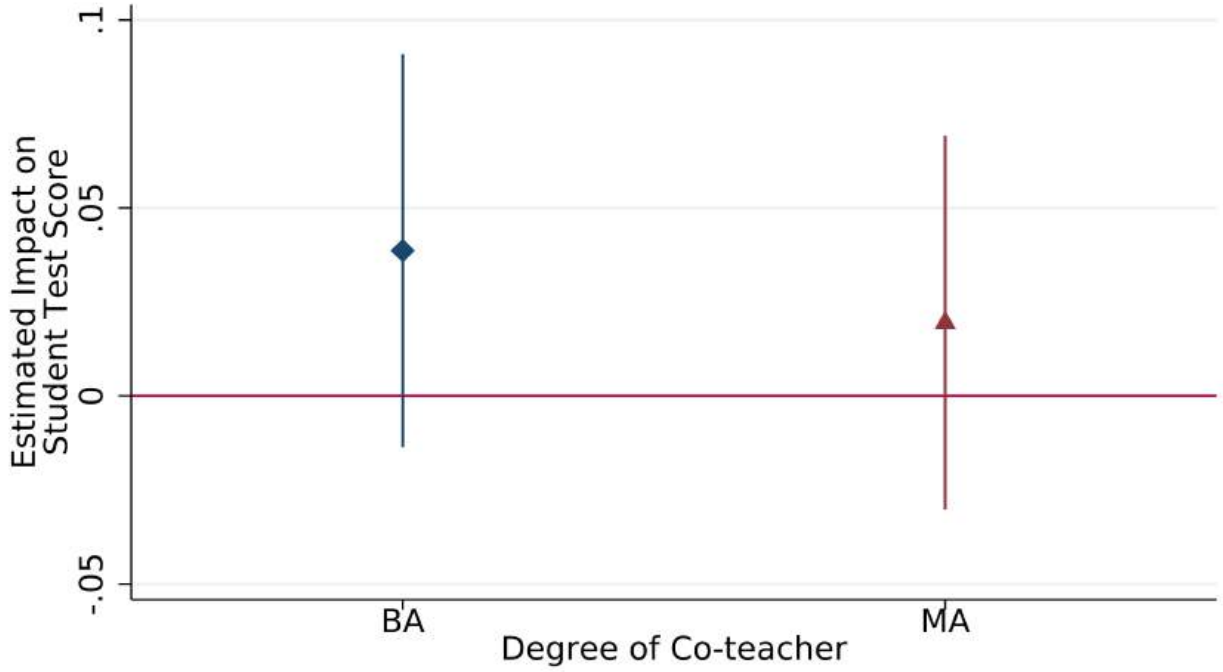

This figure examines whether the educational attainment (degree level) of the co-teaching partner affects the magnitude of developmental benefits from co-teaching.

Figure A.11: Difference in Student Test Scores By Time Since First Co-taught (No Lagged Test Scores)

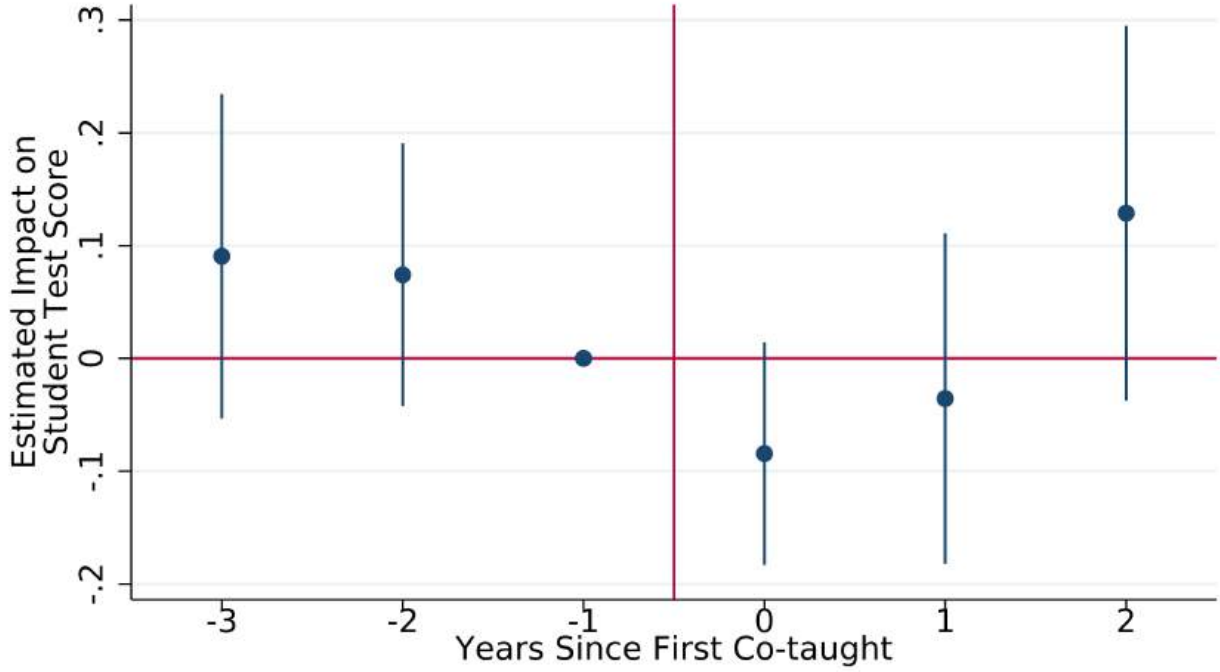

This robustness check presents the main event study results without including student prior test score controls in the regression specification.

Figure A.12: Difference in Student Test Scores By Time Since First Co-taught (One Year of Lagged Test Scores)

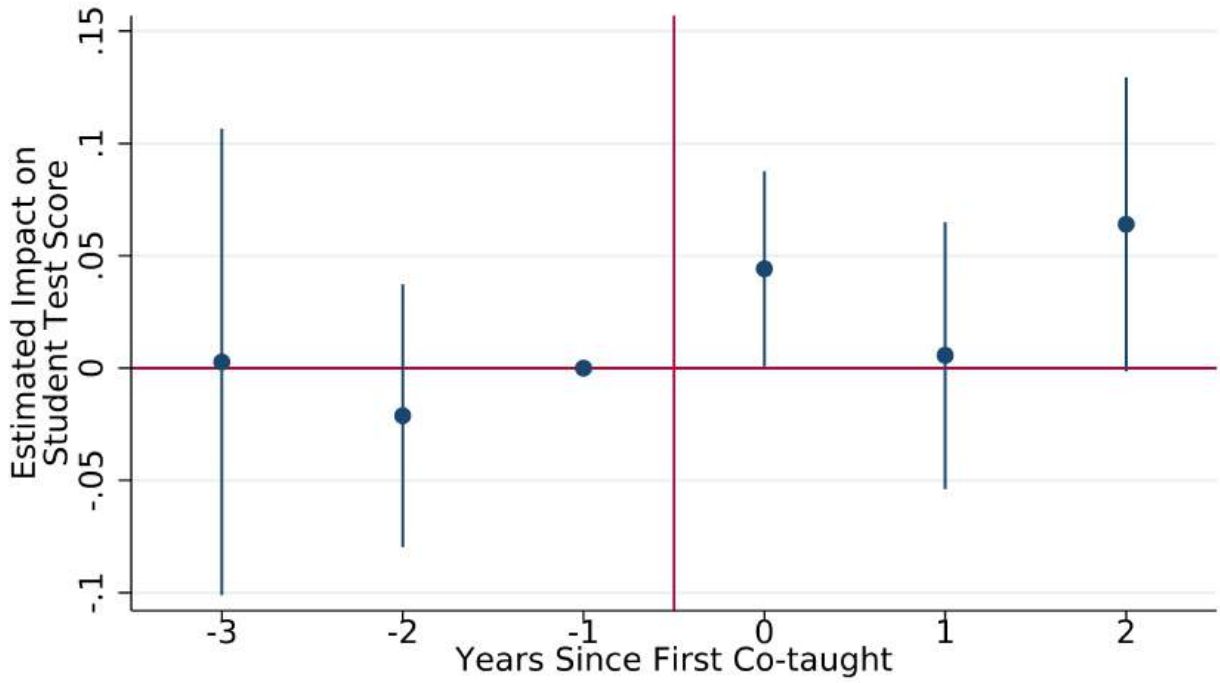

This robustness check uses only one year of prior student test scores in the control polynomial rather than the two years used in the main specification.

Figure A.13: Difference in Student Test Scores By Time Since First Co-taught (One or More Years of Prior Experience)

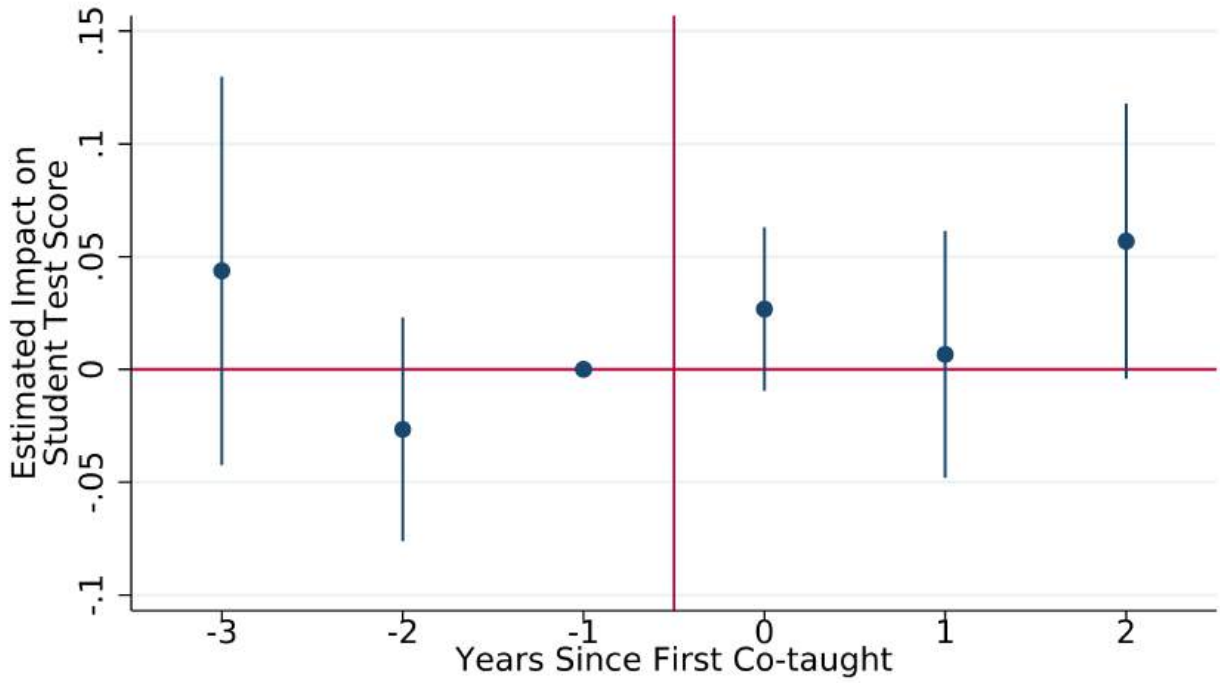

This figure expands the sample to include all teachers with at least one year of prior experience before co-teaching, rather than restricting to 1-3 years as in the main specification.

Table A.1: Difference-in-Differences Estimates Under Alternative Event Weighting

	Student Test Scores		
	Implicit Weights	Even Weights	Observation Weights
Estimated Difference	0.031	0.017	0.055

Notes: This table compares the main difference-in-differences estimate under three different weighting schemes: implicit weights from the stacked methodology, equal weights across events, and weights proportional to the number of observations per event.

Table A.2: Effect of Co-Teaching on Teacher Retention

	Full Controls		Classroom Controls		No Controls	
	Stay State	Stay School	Stay State	Stay School	Stay State	Stay School
Co-taught	0.087***	0.128***	0.067***	0.093***	0.062***	0.085***
	(0.015)	(0.018)	(0.010)	(0.012)	(0.010)	(0.012)
Observations	9,496	9,496	20,202	20,202	20,202	20,202
Baseline Mean	0.869	0.753	0.892	0.794	0.873	0.760

Notes: This table presents difference-in-differences estimates of the effect of co-teaching on teacher retention, measured as the probability of staying in the state or staying in the same school. Full controls include student test score lags, classroom demographics, and subject indicators. Co-teaching is associated with 6-9 percentage point higher retention rates, though this may reflect selection rather than a causal effect of co-teaching on retention. Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B Policy Simulation Details

This appendix provides additional detail on the assumptions underlying the policy simulation in Section 7.

B.1 Policy Structure

The simulation considers a policy where teachers with 2-3 years of solo teaching experience pair with highly experienced teachers (16+ years) for one year of co-teaching. After the co-teaching year ends, both teachers return to solo instruction. The early-career teacher carries any skill gains forward. The experienced teacher returns to their original role and remains available for future pairings.

The number of pairings each year is constrained by the supply of experienced mentors. Each mentor can support one pairing per year. I assume all available mentors participate and all eligible early-career teachers are paired if a mentor is available.

B.2 Parameter Values

Development effect. The early-career teacher gains $\delta = 0.10\sigma$ in effectiveness after co-teaching with a highly experienced partner. This estimate comes from Figure 6, which shows gains of approximately 0.10σ for teachers paired with partners who have 16+ years of experience.

Contemporaneous effect. Students in the co-taught classroom may experience direct benefits from having two teachers present. Jones and Winters (2024) estimate effects of approximately 0.02σ for students with disabilities and 0.01σ for students without. I use $\beta = 0.02\sigma$ in scenarios that include this benefit. Scenarios 3 and 4 set this to zero.

Opportunity cost. When the experienced teacher enters the co-taught classroom, their former solo classroom must be staffed by a replacement. I assume the replacement is a novice teacher. Based on Papay and Kraft (2015), the effectiveness gap between a novice and a highly experienced teacher is approximately $\gamma = 0.14\sigma$. This is the per-classroom opportunity cost of removing an experienced teacher from solo instruction. This estimate is conservative relative to the returns-to-experience function used in the simulation, which implies a gap of approximately 0.10σ between novice and 16-year teachers.

Teacher turnover. Teachers leave the profession with probability $\tau = 0.20$ per year. This rate is consistent with national estimates of teacher attrition. The turnover rate determines both the steady-state experience distribution and the size of the mentor pool. With 20% annual turnover and no retention benefit, approximately 2.9% of teachers have 16+ years of experience.

Retention benefit. In the baseline scenario, treated teachers have 6 percentage points lower turnover after their second year post-treatment. This estimate comes from Appendix Table A.2, though this may reflect selection rather than a causal effect of co-teaching on retention. Scenarios 2-4 set this benefit to zero to test sensitivity.

Skill decay. In Scenario 4, the development effect decays at 10% per year. A teacher who gained 0.10σ in year 1 retains 0.09σ in year 2, 0.081σ in year 3, and so on. This reflects

the possibility that skills acquired through collaboration may fade without continued practice or reinforcement. Scenarios 1-3 assume no decay.

Returns to experience. Teachers improve with experience even without co-teaching. I use estimates from Papay and Kraft (2015) Table 2, averaging across math and reading. Teachers gain approximately 0.064σ over their first five years (0.0128σ per year), 0.031σ over years 5-15 (0.0031σ per year), and minimal gains thereafter (0.0003σ per year). These experience effects apply to all teachers regardless of treatment status. The simulation uses this returns-to-experience function to calculate baseline teacher productivity but does not explicitly credit the policy for shifting the experience distribution. Including this channel would increase the estimated benefits.

B.3 Steady-State Calculation

I simulate the policy forward until the distribution of teacher experience and treatment status reaches steady state. In each year:

1. Teachers with 16+ years of experience are assigned as mentors.
2. The share of mentors determines how many pairings occur.
3. Early-career teachers (2+ years of experience) are paired with mentors.
4. Treated teachers carry their skill gains (δ) forward, potentially with decay.
5. Teachers turn over at rates that depend on treatment status and time since treatment.
6. The workforce composition evolves based on these flows.

The simulation runs for 150 years to ensure convergence to steady state. Results report steady-state values after this convergence.

B.4 Benefit and Cost Calculations

Annual VA benefit. The simulation tracks the share of classrooms staffed by treated teachers and multiplies by the treatment effect (δ), adjusted for any skill decay. In scenarios with contemporaneous effects, it also adds the benefit from having two teachers present in co-taught classrooms (β times the mentor pool share).

Opportunity cost. Each pairing removes one experienced teacher from a solo classroom. The opportunity cost equals the mentor pool share (fraction of classrooms with mentors) multiplied by $\gamma = 0.14\sigma$.

Net annual effect. This equals total annual benefits minus annual opportunity cost, measured in units of student achievement (σ).

Lifetime earnings translation. I convert achievement effects to earnings using estimates from Chetty et al. (2014). Their middle school estimates imply that a 0.01σ increase in test scores translates to approximately \$500 in undiscounted lifetime earnings (in 2010 dollars). I inflate this to 2025 dollars using the CPI. The lifetime earnings per student equal the net annual effect times this conversion factor.

Years to positive cumulative. The simulation tracks cumulative net effects year by year. The "years to positive cumulative" is the first year where the sum of all past net effects becomes positive.

Break-even opportunity cost (γ^*). This is the opportunity cost that would make net annual benefits equal to zero: $\gamma^* = (\text{total benefits})/(\text{mentor pool share})$. Comparing this to the assumed $\gamma = 0.14\sigma$ shows how much margin for error the policy has.