North Carolina Paper Summary

This paper asks whether achievement gains from student–teacher allocation primarily reflect comparative advantage—assigning students to teachers who are relatively more effective for their particular characteristics—or absolute advantage—exposing more students to teachers with higher average value added. Using North Carolina administrative records linking students, teachers, and math and reading achievement in grades 3–8 from 2006 through 2024, I evaluate both the predictive content of teacher heterogeneity and its implications for allocation.

I estimate overall teacher value added and type-specific value added by economic disadvantage, race, sex, and prior achievement. I evaluate whether comparative-advantage forecasts improve prediction beyond overall value added using chronological holdouts and a teacher-turnover design. I also use matrix completion—a low-rank model that uses historical student–teacher outcomes to estimate latent match quality—to test whether individual matches contain predictable information beyond observed student characteristics and average teacher effects.

The main findings are:

- **There is evidence of comparative advantage, but it does not add predictive value.** Corrected correlations between teacher effects across student types are high but below one, consistent with some small but persistent heterogeneity. However, adding estimated comparative advantage does not improve out-of-sample prediction beyond overall teacher value added. And they are consistently not fully forecast unbiased.
- **I find no evidence of predictable latent student–teacher match quality using matrix completion.** Such matches could reflect teaching style or other interactions not captured by observed student characteristics. The model fits historical outcomes well but fails to predict future outcomes. This does not rule out individual match effects; it shows that the estimated matches cannot reliably guide prospective assignment.
- **The apparent gains from matching depend on which dimension is optimized and evaluated.** With teacher capacities fixed, sorting on prior-achievement comparative advantage produces model-implied gains of approximately 0.0021 standard deviations in math and 0.0022 in reading. Using an equal-weight ensemble of all four type-specific predictions reduces these gains to approximately 0.0011 in each subject. This motivates a critique of evaluating assignment policies through a single selected dimension: if multiple dimensions capture relevant heterogeneity, gains measured along one dimension ignore other dimensions that also have some level of match quality, and seem to often move in opposite directions, which mitigate predicted benefits.
- **Reassigning teachers to fixed classroom (the policy used by the paper I am responding to) bundles can exploit both comparative advantage and differences in overall teacher value added**, because bundles differ in size as well as student composition. The remaining decomposition (the main thing I still need to do) will compare assignments based on overall value added alone with assignments that also

incorporate comparative advantage, holding the feasible assignments and evaluation criterion fixed. This will establish how much of the simulated gain comes from assigning higher-value-added teachers to larger classes, and how much comparative advantage adds.

- **Moving a moderate number of students from lower- to higher-value-added teachers within schools generates larger model-implied gains in both subjects.** These gains require shifting roughly 4–5 percent of enrolled seats. The headline estimates assume no dilution of teacher effectiveness as workloads increase; sensitivity analyses relax that assumption.

Taken together, the results favor expanding access to higher-value-added teachers over fine-tuning assignments using estimated comparative advantage. Comparative advantage adds little out-of-sample predictive information, and modest workload reallocation generates gains larger than student sorting. The outstanding decomposition will determine whether the larger gains from teacher-to-class reassignment likewise primarily reflect absolute advantage rather than comparative advantage.