

How Did It Get This Way? Disentangling the Sources of Teacher Quality Gaps through Agent-Based Modeling

Dan Goldhaber

Matt Kasman

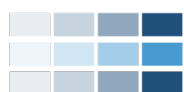
Vanessa Quince

Roddy Theobald

Malcolm Wolff

February 2023

WORKING PAPER No. 259-0223



CALDER

National Center for Analysis of
Longitudinal Data in Education Research



How Did It Get This Way? Disentangling the Sources of Teacher Quality Gaps through Agent-Based Modeling

Dan Goldhaber

*American Institutes for Research / CALDER
University of Washington*

Matt Kasman

Brookings Institution

Vanessa Quince

Mathematica Policy Research

Roddy Theobald

American Institutes for Research / CALDER

Malcolm Wolff

University of Washington

Contents

Contents.....	i
Acknowledgments	ii
Abstract	iii
1. Introduction	1
2. Framework and Literature Review	4
3. Methods	6
4. Results	14
5. Conclusions	17
References	21
Tables and Figures.....	26

Acknowledgments

This work is supported by the William T. Grant Foundation (grant #184925) and by the National Center for Analysis of Longitudinal Data in Education Research (CALDER), which is funded by a consortium of foundations. For more information about CALDER funders, see www.caldercenter.org/about-calder. We wish to thank Joe Walch for research assistance, as well as Chad Aldeman, James Cowan, Aliza Husain, and Jason Grissom for comments that improved the manuscript. The views expressed in this paper do not necessarily reflect those of the American Institutes for Research, the University of Washington, the project sponsors, or data providers.

CALDER working papers have not undergone final formal review and should be cited as working papers. They are intended to encourage discussion and suggestions for revision before final publication. Any opinions, findings, and conclusions expressed in these papers are those of the authors and do not necessarily reflect the views of our funders or the institutions to which the authors are affiliated. Any and all errors are solely the responsibility of the authors.

CALDER • American Institutes for Research
1400 Crystal Drive 10th Floor, Arlington, VA 22202
202-403-5796 • www.caldercenter.org

How Did It Get This Way? Disentangling the Sources of Teacher Quality Gaps through Agent-Based Modeling

Dan Goldhaber, Matt Kasman, Vanessa Quince, Roddy Theobald and Malcolm Wolff

CALDER Working Paper No. 259-0223

February 2023

Abstract

We use publicly available, longitudinal data from Washington state to study the extent to which three interrelated processes—teacher attrition from the state teaching workforce, teacher mobility between teaching positions, and teacher hiring for open positions—contribute to “teacher quality gaps” (TQGs) between students of color and other students in K–12 public schools. Specifically, we develop and implement an agent-based model simulation of decisions about attrition, mobility, and hiring to assess the extent to which each process contributes to observed TQGs. We find that eliminating inequities in teacher mobility and hiring across different schools would close TQGs within 5 years, while just eliminating inequities in teacher hiring would close gaps within 10 years. On the other hand, eliminating inequities in teacher attrition without addressing mobility and hiring does little to close gaps.

1. Introduction

There is abundant evidence of “teacher quality gaps” (TQGs) between advantaged and disadvantaged students in U.S. public schools. These TQGs are evident whether teacher quality is measured by experience, degrees, or other advanced credentials (e.g., Bruno et al., 2020; Clotfelter et al., 2005; Kalogrides and Loeb, 2013; Lankford et al., 2002; Rodriguez et al., 2023) or by “value added” measures of teacher effectiveness (e.g., Goldhaber et al., 2015; Isenberg et al., 2021; James and Wyckoff, 2022; Sass et al., 2012). For example, longitudinal data on public schools from North Carolina and Washington (Goldhaber et al., 2018) demonstrate that TQGs existed in every available year of data in these states for observable measures of student disadvantage (poverty level and being a student of color) and teacher quality (experience, licensure test scores, and value added). Importantly, there is also evidence that TQGs affect student achievement and explain ethnoracial gaps in student outcomes (Chetty et al., 2014; Goldhaber et al., 2022). In short, these gaps appear to be a meaningful driver of inequity in educational opportunity (e.g., Education Trust, 2020; Mehrotra et al., 2021).

What explains inequities in teacher quality across students? In this paper, we develop a novel methodology to explore the sources of one specific type of TQG that has been extensively documented in this prior literature (e.g., Clotfelter et al., 2005; Goldhaber et al., 2015; Lankford et al., 2022)—the relative exposure of students of color and white students to novice teachers (those with less than 5 years of experience). Specifically, we use data from Washington state to assess the extent to which TQGs arise and persist because of three distinct, but interrelated, processes: attrition from the state workforce, teacher mobility between public school teaching positions, and the hiring of teachers for open positions.

Consistent with prior literature, we document inequitable patterns in each of these processes. Our primary contribution, however, is to simulate how each of these processes

contributes to TQGs over time to better understand *how* they act in concert to create TQGs. A key challenge is that the teacher labor market exhibits substantial amounts of interdependence, adaptation, and heterogeneity; these are hallmarks of a *complex adaptive system* (Kasman et al, 2021; Miller & Page, 2009). Teachers' and schools' behaviors and outcomes are interconnected. For example, instances of mobility from one school to another can change the composition of the teaching body at both. In addition, the three processes that we focus on are fundamentally intertwined: teachers who leave a school, district, or the profession in one year generally must be replaced with new hires in the next. Teacher hiring and mobility decisions are a result of a “two-sided match” between employers and employees (Boyd et al., 2013). As such, teacher and school decision-making can both adapt in response to available jobs and applicant pools. And there is also good evidence that teacher hiring, mobility, and attrition are experienced differently across schools and districts (e.g., Boyd et al., 2011), leading to heterogeneity that may have important implications for TQGs.

To capture multiple, interacting and adaptive processes as well as the potential importance of heterogeneity in the attributes or behaviors of teachers and school administrators in attrition, mobility, and hiring decisions, we develop and deploy an agent-based modeling approach to characterize these aspects of the teacher labor market. Our first research goal is etiological in nature: we want to create a parsimonious computational simulation model that quantitatively characterizes teachers, schools, and the dynamics that drive attrition, mobility, and hiring over time. We assess the model's ability to reproduce observed patterns of TQGs given appropriate model input.

After determining whether our model has satisfactory explanatory power, we use it to identify key leverage points within the teacher labor market for addressing and reducing TQGs.

That is, we use this model not to simulate the impact of any specific policy but, rather, to identify important policy levers in states' efforts to close TQGs. Specifically, we use this model to simulate different hypothetical scenarios that allow us to examine the extent to which each of the three processes *independently* or *synergistically* contribute to TQGs between advantaged and disadvantaged students.

We argue that understanding the sources of TQGs is fundamental to closing them, a policy goal that has received a good amount of attention over the last decade. For instance, a federal directive from the U.S. Department of Education mandated state plans to reduce inequity in the distribution of teacher quality across public schools (Rich, 2014), though only 13 states ultimately provided timelines for eliminating TQGs as part of ESSA plans submitted in 2016 (Ross, 2019). The results from this analysis directly inform these plans. For example, if most of the inequity between advantaged and disadvantaged schools is the result of differential teacher attrition, states and districts might focus on policies to *keep* high-quality teachers in disadvantaged schools. If, on the other hand, discrepancies in teacher hiring explain most of the observed inequities, states and districts might instead prioritize policies to *attract* high-quality teachers to disadvantaged schools.¹

The paper proceeds as follows. In Section 2, we discuss our framework for investigating each of the labor market processes discussed above and the prior literature related to each process. We describe our methods in Section 3, discuss results in Section 4, and conclude with implications for policy and directions for future research in Section 5.

¹ It is, of course, possible to address TQGs by creating inequity in one process to offset an inequity in another process. For instance, disadvantaged schools might reduce the attrition of high-quality teachers below the level of attrition of high-quality teachers in advantaged schools to offset gaps in the ability of the two types of school to hire high-quality teachers. Assessing the costs of addressing TQGs is outside the scope of this paper, but we also believe that the politics of implementing solutions to TQGs are likely connected to the sources of the gaps, suggesting the need to better understand these.

2. Framework and Literature Review

A growing literature documents the existence and magnitude of TQGs in districts and states across the country (see Goldhaber et al. [2018] for a review). And as discussed above, under the Every Student Succeeds Act (ESSA) states are mandated to identify and address these uneven distributions of teaching quality and qualifications.² But to address them, it is important to understand how inequities formed in the first place. Following prior analyses of TQGs in public schools (e.g., Clotfelter et al., 2005; Goldhaber et al., 2015; James & Wyckoff, 2022; Lankford et al., 2002; Rodriguez et al., 2023), we focus on the relative exposure of students of color and white students to novice teachers (those with less than 5 years of experience). We use students of color as a proxy for disadvantage given historical and persistent inequities in their educational outcomes (e.g., Betts et al., 2003; Clotfelter et al., 2009). Our consideration of teacher experience as a proxy for teacher quality is motivated by the fact that it is observable for all teachers, it is called out explicitly in ESSA, and evidence from many different states and time periods that it predicts teacher effectiveness, particularly in the first five years of a teacher’s career (e.g., Clotfelter et al., 2005; Rivkin et al., 2005).

Several processes—the attrition of teachers from different types of schools, the movement of teachers between schools and districts, and the hiring of new teachers into their first jobs—may contribute to TQGs.

Teacher hiring

² Specifically, ESSA requires that states receiving a Title I, Part A grant to have plans that “describe how low-income and minority children enrolled in schools assisted under this part [of ESSA] are not served at disproportionate rates by ineffective, out-of-field, or inexperienced teachers, and the measures the State educational agency will use to evaluate and publicly report the progress of the State educational agency with respect to such description...” (ESSA, 1111 (g)(1)(B), 2015).

Evidence on the robustness of teacher applicant pools is limited (Bleiberg and Kraft, 2022), but the quantitative work that does exist suggests that schools serving more disadvantaged students tend to have less robust applicant pools when hiring for an open position (Boyd et al., 2011; Bruno & Strunk, 2019; James et al., 2022). Existing evidence also finds that novice teachers tend to have more challenging school placements (e.g., Bruno et al., 2019), while emerging evidence suggests that disadvantaged schools are more likely to hire novice teachers for open positions in part because they tend to post positions later in the hiring cycle (e.g., James et al., 2022).³

Teacher mobility

The empirical evidence on teacher mobility suggests that within-teaching mobility (“transfers”) likely contributes to TQGs in two ways. First, teachers who teach in schools with higher proportions of disadvantaged students are more likely to transfer to another school in the same district than those who teach in schools with lower proportions of disadvantaged students (e.g., Goldhaber et al., 2011; Isenberg et al. 2021; Hanushek et al. 2004; Sass et al. 2012; Scafidi et al., 2007), likely due to differences in working conditions (e.g., Ladd, 2011). For example, Isenberg et al. (2021) find that on average, eleven percent of teachers in high-poverty schools transfer to another school in the district compared to five percent in low-poverty schools. This can contribute to TQGs because disadvantaged schools must hire more new teachers each year who, on average, are less experienced and less effective than the average teacher (e.g., Isenberg et al., 2021).

³ Some studies find that there are small to no statistical differences between the value added of teachers in high-poverty schools relative to low-poverty schools (Isenberg et al. 2021; Sass et al. 2012), while others (Xu et al., 2015) find that novice teachers in high-poverty schools were 0.02 standard deviations lower in terms of value added than teachers at low-poverty schools, and Loeb et al. (2012) find that effective schools are more likely to hire teachers with higher prior estimates of value added. The aforementioned papers do not focus specifically on teacher hiring, but simply compare the effectiveness of novice teachers in different school settings.

Second, teacher mobility can also contribute to TQGs if experienced teachers are disproportionately more likely to leave disadvantaged schools (or, alternatively, if inexperienced teachers are disproportionately likely to stay in disadvantaged schools). There is some prior evidence of this (e.g., Boyd et al. 2008; Goldhaber et al., 2011; Isenberg et al., 2021); for example, Isenberg et al. (2021) find that experienced teachers are more likely than inexperienced teachers to leave high-poverty schools for low-poverty schools in the same district.

Teacher attrition

A number of studies demonstrate that teachers in disadvantaged schools are more likely to leave the workforce, again likely because of differences in working conditions across schools (e.g., Geiger & Pivovarova, 2018). For instance, in a study of 26 school districts, Isenberg et al. (2021) find that ten percent of teachers leave high-poverty schools in a typical year compared to seven percent of teachers who leave low-poverty schools. Studies that focus on specific schools and districts tend to come to similar conclusions; for example, Barnes et al. (2007) find a six percent difference in turnover between high-poverty and low-poverty schools in Chicago and Illinois Public Schools. These observed patterns are consistent with findings from various individual states like Florida (Feng & Sass, 2017), Georgia (Scafidi et al, 2007), New York (Boyd et al., 2008), Texas (Hanushek et al. 2004), Wisconsin (Imazeki, 2005), North Carolina (Clotfelter et al., 2008; Goldhaber et al., 2011), and in research spanning three decades in our focal state of Washington (Krieg, 2006; Goldhaber et al., 2016; Gritz and Theobald, 1996).

3. Methods

3.1 Agent-Based Model Design

We use an agent-based model (ABM) model design to study the extent to which three mutually exclusive processes—teacher attrition from the state workforce, teacher mobility between public

school teaching positions (“transfers”), and teacher hiring for open positions—contribute to TQGs between advantaged and disadvantaged schools. ABM is a computational modeling approach in which individual entities (e.g., teachers or schools) as well as their interactions with one another over time are explicitly simulated. It is thus well suited for exploring broad questions about educational processes that are driven by the dynamic intersection of individual and organizational decision making (Maroulis et al., 2010).

Our model is designed to represent known features of the teacher labor market that present a challenge to study with traditional analytic techniques. Specifically, it can explicate the ways in which teacher decisions as well as the three dynamic processes that we want to characterize are interdependent. That is, teachers’ decisions about whether to leave a school incorporate decisions made by others, and these intrinsically link hiring, mobility, and attrition. When deciding whether to leave a school, teachers consider a school’s current teacher workforce. As we show in Section 3.3, teachers are more likely to leave their school as the proportion of novice teachers in the school increases; this may be because teachers specifically care about their peers’ presence, or because the proportion of novice teachers in the school provides a proxy for unobserved (to the researcher) working conditions in the school. Either way, as experienced teachers leave a school and are disproportionately replaced with novice teachers, our model can capture that remaining teachers become more likely to leave as well. Thus, we can observe the endogenous formation of dynamic feedback loops that might play a key role in TQGs. In the following subsections, we discuss the model design in more detail.

Agents (Teachers) and Environment (Schools)

The agents in this ABM represent teachers in K–12 public schools in Washington state. Following best practices in the design of complex systems models, we strive for a parsimonious

model that has substantial explanatory power for relevant phenomena, capable of representing both relevant agent-level behaviors and the system-level patterns in the outcomes of interest that they produce (Hammond, 2015). Agents in this model have a single attribute: years of experience, with agents having less than 5 years of teaching experience designated as “novice teachers.” The environment in which decisions are made that result in hiring, mobility, and attrition is the teacher labor market. We decompose this into two parts: K–12 schools, which are explicitly represented within our simulations, and exogenous, time-dependent factors (e.g., employment opportunities in other professions), which we represent implicitly. For the purposes of model parsimony, schools also have a single attribute: the percentage of students of color that they serve.⁴

Initialization and Dynamics

During the first simulated year of our model, we place a set of teachers in schools; we can initialize the model either “equitably” by doing so at random or “inequitably” by doing so in accordance with observed relationships between teacher experience and school composition in this initial year. All results in this paper are taken from runs using the “inequitable” initialization of the model (i.e., starting from the “real-world” data), but model output from ones with equitable initialization are available upon request. In subsequent years in our model, teacher and school interactions follow a serial, three-stage process:

1. *Attrition*. Teachers decide whether to stop working as public school K–12 teachers.
2. *Mobility*. Teachers who remain after the previous step decide whether they wish to transfer to an available position in another school.

⁴ Of course teachers also may have other characteristics correlated with being novices, and schools have other characteristics correlated with the percentage of students of color in the school, all of which are also related to the various processes we investigate. But we select these two characteristics for model parsimony and to connect this simulation to prior documented TQGs that are a function of these two characteristics.

3. *Hiring.* Finally, schools fill positions from two distinct pools of available applicants, movers and new hires, the latter of whom may have previous teaching experience (e.g., be returning to the profession or entering from another state).

We allow for each of these stages to be characterized either “equitably”—with behavior occurring without any consideration of school composition—or “inequitably,” allowing the process to occur differentially across schools serving different types of students. Specifically, when attrition and mobility occur equitably, the probability that teachers leave their current position is a function solely of their experience, while in the inequitable scenario attrition and mobility are also functions of other school-level factors (e.g., the average teacher experience in the school and the percentage of students of color in the school) and interactions with teacher characteristics (see Section 3.2).⁵ In the case of teacher hiring, when hiring is equitable, open positions are filled randomly, while under inequitable hiring, open positions are filled based on the level of disadvantage of the school—i.e., schools are probabilistically matched to mover and new hire applicant pools based on levels of school disadvantage.

3.2 Data Sources and Summary Statistics

We use Washington state’s S-275 database, which contains information from OSPI’s personnel-reporting process and includes school assignment of all certified employees in the state in addition to a measure of teaching experience in the state.⁶ S-275 data are available from the 1983–84 school year through the 2021–22 school year, although our analysis in Washington spans 13 years (2001–02 through 2013–14) to be consistent with a prior analysis of TQGs (Goldhaber et al., 2018). We consider teachers with a full-time teaching appointment in a single

⁵ We remove average teacher experience in the school in the equitable scenario to prevent path dependent preservation of past inequity (i.e., reflecting historical artifacts rather than the process itself).

⁶ The S-275 contains the experience that teachers are credited with for pay purposes, which do not include out-of-state teaching, teaching in a private school, or substitute teaching.

school within a given year. We use the S-275 dataset to quantify each teacher’s experience, and specifically whether a teacher has less than 5 years of experience which is where returns to teacher experience are greatest (e.g., Clotfelter et al., 2005; Rivkin et al., 2005). We then connect these data to publicly available data on the percentage of students of color in each teacher’s school from the Washington State Report Card.⁷ The final longitudinal dataset in Washington includes 564,296 teacher-year observations (86,241 unique teachers) spanning 13 school years.

In **Figure 1** we provide summary statistics from this dataset that replicate and extend prior work discussed in the previous section about inequities in teacher attrition, mobility, and hiring. Panels A and B illustrate that teacher attrition and mobility are more prevalent in disadvantaged schools in Washington. There is also some evidence of *differential* mobility and attrition that reinforces prior findings; for example, Panel A of Figure 1 illustrates that non-novice teachers are disproportionately likely to leave advantaged schools relative to novice teachers. These findings largely replicate prior trends discussed in Section 2, but more novel are the findings for teacher hiring; in particular, the hiring proportions (Panel C in Figure 1) illustrate that, conditional on an open position, disadvantaged schools are more likely to hire an inexperienced teacher than an advantaged school.

3.3 Statistical Estimation

As described in Section 3.1, parameterization of the ABM such that it reflects inequitable patterns observed in the real world requires predicted probabilities of teacher attrition, mobility, and hiring behaviors and opportunities given teacher and school characteristics (i.e., the percentage of students of color and teacher experience). To obtain these values, we estimate a series of Bayesian models predicting the attrition, mobility, and hiring of teachers between 2002

⁷ Following Goldhaber et al. (2018), we define students of color as American Indian, Black, or Hispanic.

and 2014. Teacher attrition and mobility are jointly estimated using the mixed-effects multinomial logistic regression model described by equations 1a-1c:

$$\beta_{\cdot,} \sim N(0, \mathbf{I}), \quad \gamma_{\cdot,} \sim N(0, .1\mathbf{I}), \quad \sigma_{\tau} \sim N(0, \mathbf{I}), \quad \sigma_{\zeta} \sim N(0, \mathbf{I}), \quad (1a)$$

$$\ln \frac{P(Y_{jst} = \text{Attrit})}{P(Y_{jst} = \text{Stay})} = \beta_{A,0} + T_{jt}^{\top} \beta_{A,1} + S_{st}^{\top} \beta_{A,2} + E_{jst}^{\top} \beta_{A,3} + \tau_{A,j} + \zeta_{A,s} + \gamma_{A,t} + \varepsilon_{A,jst}, \quad (1b)$$

$$\ln \frac{P(Y_{jst} = \text{Move})}{P(Y_{jst} = \text{Stay})} = \beta_{M,0} + T_{jt}^{\top} \beta_{M,1} + S_{st}^{\top} \beta_{M,2} + E_{jst}^{\top} \beta_{M,3} + \tau_{M,j} + \zeta_{M,s} + \gamma_{M,t} + \varepsilon_{M,jst}. \quad (1c)$$

T_{jt} is a 4×1 vector of teacher-by-year characteristics including novice status, a novice/year interaction term, and a linear and quadratic term of teaching experience for teacher j and year t . S_{st} is a 3×1 vector of school-by-year characteristics including average experience, proportion of novice teachers, and proportion of students of color in school s and year t . E_{jst} is a 4×1 vector of interactions between the proportion of students of color in school s and individual novice status, individual experience, average experience, and proportion of novice teachers. $\tau_{\cdot,j}$ is a teacher random effect, $\zeta_{\cdot,s}$ is a school random effect, $\gamma_{\cdot,t}$ is a year fixed effect, and $\varepsilon_{\cdot,jst}$ is a mean-zero error term. To obtain ABM parameters for counterfactual scenarios where these processes are equitable, we use estimates from an alternative model that restricts equations 1a-1c so that only coefficients associated with the linear experience terms and the intercepts are nonzero.

Teacher hiring practices in year t are estimated using the two-stage mixed-effects model described by equations 2a-2c:

$$\beta_{\cdot,} \sim N(0, \mathbf{I}), \quad \gamma_{\cdot,} \sim N(0, .1\mathbf{I}), \quad \sigma_{\tau} \sim N(0, \mathbf{1}), \quad \sigma_{\zeta} \sim N(0, \mathbf{I}), \quad (2a)$$

$$\ln \frac{P(M_{st} = 0)}{P(M_{st} > 0)} = \beta_{N,0} + S_{st}^{\top} \beta_{N,1} + E_{st}^{\top} \beta_{N,2} + \zeta_{N,s} + \gamma_{N,t} + \varepsilon_{N,st}, \quad (2b)$$

$$M_{st} = \beta_{H,0} + S_{st}^{\top} \beta_{H,1} + E_{st}^{\top} \beta_{H,2} + \zeta_{H,s} + \gamma_{H,t} + \varepsilon_{H,st}. \quad (\text{for } M_{st} > 0) \quad (2c)$$

M_{st} is the proportion of new hires in year t teaching in a different Washington state school in year $t - 1$. S_{st} is a 4×1 vector of school-year characteristics including the proportion of novice teachers, average teacher experience, the proportion of zero experience teachers, and the proportion of students of color. E_{st} is a 3×1 vector of interaction terms between the proportion of students of color in school s and year t and other school-by-year characteristics. $\zeta_{.s}$ is a school random effect, $\gamma_{.t}$ is a year fixed effect, and $\varepsilon_{.st}$ is a mean-zero error term. To obtain ABM parameters for counterfactual scenarios in which the hiring process is equitable, we again estimate an alternative equitable model by restricting equations 2a-2c so that only coefficients associated with the linear experience terms and the intercepts are nonzero.

Prior distributions on each of the coefficients and standard deviation of random effects are chosen to be independently distributed according to a standard normal to remain uninformative. We set the standard deviation of the year fixed effect prior distribution to be 0.1 to reduce the occurrence of unrealistically large year effects having substantial impact on outcome patterns from ABM runs.⁸ Importantly, the parameter values used for each ABM run that determine the probabilities of subsequent teacher placements are sampled from the conjugate normal posterior distributions of these estimates; this means that subsequent simulation estimates incorporate both estimation error (i.e., reflected in the standard errors in Table 1) as well as simulation error reflected in differences across the different simulation runs described below.

The estimates from the models in equations 1a–2c (provided in Table 1) strongly comport with the summary statistics discussed in Section 3.2 and the prior literature discussed in Section 2; teachers in schools that serve more students of color are more likely to leave public K–12

⁸ In expectation, this selection should not affect our key findings, which rest upon comparisons between dynamic TQG trends under baseline and counterfactual conditions. However, by choosing to limit stochasticity with this decision, we can meaningfully conduct these comparisons with a smaller set of repeated simulation runs. Given the large computational time required for each run, this is logistically beneficial.

schools than teachers in schools with fewer students of color, and schools that serve a high proportion of students of color are more likely to hire a novice teacher when they have an opening than schools that serve fewer students of color. As in the summary statistics, there is also evidence of differential inequity, as experienced teachers are particularly likely to leave schools with higher proportions of students of color.

We use the coefficients from Table 1 to generate predicted probabilities of attrition, mobility, and hiring for open positions (as described in Section 3.1). Our primary model outputs are distributions of TQG derived from proportions of novice teachers serving students of color and white students in each year across 100 simulation runs in each of eight “inequity” configurations (i.e., with each potential combination of inequality present or not for attrition, mobility, and hiring). As discussed above, this allows us to simultaneously account for both estimation error (i.e., through sampling of parameter values) and model stochasticity, and thus to explore a substantial amount of model sensitivity using our primary analysis data.

That the regressions in Table 1 account for the average experiences and proportion of novice teachers (and the interactions between these variables and the percent of students of color in the school) means that the simulations allow for *indirect* effects of attrition, mobility, and hiring on future attrition and mobility decisions through the changing characteristics of teachers in a school (e.g., if a school is losing experienced teachers they tend to lose even more experienced teachers in future years because of the positive relationship between school percent novice teachers and the probability of teacher mobility in Table 1). As a robustness check, though, we run a separate set of simulations that do not allow for these indirect effects (i.e., setting the relationships between school-level teacher characteristics and attrition and mobility

from Table 1 to zero) to test how important these indirect effects are in explaining our simulation results.

4. Results

4.1 Model use and experimentation

In **Figures 2 and 3**, we explore the “generative sufficiency” of our model in two key respects (Epstein, 2012). In Figure 2, we show that rates of attrition, mobility, hiring, and overall proportions of novice teachers in these simulations track the true rates quite closely. More importantly, when each of our three inequity pathways is active, we expect to generate processes and TQGs similar to those observed in the real world during the years that correspond to our simulations. The solid lines in the first two panels of Figure 3 plot the simulated probability (i.e., the probability from the median simulation run) that students are assigned to a novice teacher over simulation years for students of color (Panel A) and other students (Panel B) alongside the true assignment rates (dotted lines) and 90 percent confidence intervals for these estimates (dashed lines) that, as described in the previous section, account for both estimation error from the models in equations 1a–2c in Section 3 *and* simulation error across 100 iterations of each simulation. As shown in these figures, the simulated probability that both groups of students are assigned to a novice teacher tracks very closely with the observed probabilities, and well within the margin of error for these estimates across simulation years. As a result, and as shown in Panel C of Figure 3, the simulated TQG tracks the observed TQG (documented for these same years of data in Goldhaber et al. [2018]) very closely across simulations years and, again, well within the margin of error for these estimates.

Finally, when each of the inequity pathways is *inactive*, we expect that TQGs will be functionally non-existent. This will establish the degree to which our characterizations of teacher

labor market processes—and only these characterizations—are capable of broadly explaining TQGs. In Panel D of Figure 3, we show that when all of the inequity pathways are set to be equitable across schools the TQG goes to zero within about 8 simulation years in the average simulation run. So, the figures in Panels C and D showing the simulated TQG when all processes are inequitable (Panel C) and equitable (Panel D) represent bounds on the simulated TQGs discussed in the next section that turn on and off these inequity pathways alone or in combination.

4.2 *Simulation Results*

We begin by providing an overview of the most straightforward simulations that turn off (i.e., set to “equitable”) one inequity pathway at a time. The intuition behind these simulations is that, if only one process in the teacher labor market becomes equitable while the others remain inequitable, then any subsequent reductions in the overall TQG (relative to the TQG when all processes are inequitable) provides an estimate of how much that specific process contributes to TQGs. Alternatively, these exercises provide an estimate of how much we might expect TQGs to decrease in response to hypothetical interventions that manage to remove all inequity in one of the three processes we investigate.

In **Figure 4**, we first repeat in Panel A the figure showing the evolution of the simulated TQG across simulation years when all inequity pathways are on. We then present analogous figures for simulations in which the attrition (Panel B), mobility (Panel C), and hiring (Panel D) inequity pathways are turned off. We find that turning off inequity in terms of teacher mobility and hiring (Panels C and D of Figure 4) are both simulated to substantially reduce TQGs. In the case of mobility, the TQG reduces to about half of the observed TQG by the fifth year of the simulation and stays at about half of the observed TQG for the remainder of the simulation years. The simulated reduction in TQGs resulting from turning off the hiring inequity pathway is less

rapid but larger over time; TQGs are projected to be reduced to near zero by year 8. These simulation results suggests that, even under observed rates of attrition and mobility from different schools by experienced and inexperienced teachers, TQGs could be eliminated in less than 10 years if the probability of *replacing* these teachers with an experienced teacher was the same across advantaged and disadvantaged schools.

Interestingly, and by contrast, turning off *only* the attrition inequity pathway (Panel B)—i.e., ensuring that the probability of teacher attrition is the same for teachers across all school settings in the state—does little to reduce TQGs. In fact, the TQG from the “all inequitable” simulation is within the confidence interval for the estimated TQG under this simulation in each of the 12 simulation years. In other words, eliminating inequities in teacher attrition without addressing mobility and hiring does little in the short run to close these gaps.

As discussed in the introduction, these simulations represent a complex adaptive system (Kasman et al., 2021; Miller & Page, 2009) in which all three processes are dynamically interconnected. It is therefore important to consider how these different processes might *interact* to better understand how policies that affect more than one process—e.g., incentives to keep teachers in their current school that might affect both attrition and mobility, or interventions to recruit teachers to disadvantaged schools that might affect both mobility and hiring—might affect TQGs.

In **Figure 5** the estimated TQG and associated confidence intervals after 5 simulation years (Panel A) and after 10 simulation years (Panel B) for the simulation associated with *each combination* of inequity pathways. The estimates for “1 Attrition,” “2 Mobility,” and “3 Hiring” correspond exactly with the corresponding estimates in Panels B–D of Figure 4, but the simulations that turn off multiple inequity pathways lead to additional conclusions. Of particular

note, the simulation that turns off both the mobility and hiring inequity pathways results in a TQG after just 5 years that is indistinguishable from zero. Since the primary mechanism through which teacher mobility affects TQGs is through the movement of teachers from less advantaged to more advantaged schools (Goldhaber et al., 2016), this suggests that interventions that target teacher recruitment to disadvantaged schools—whether through new teacher hiring *or transfers from other schools*—could close TQGs even faster than interventions that target just one of these processes.

Finally, as noted above, the simulations allow for both direct effects of attrition, mobility, and hiring (i.e., specific teachers leaving and entering positions) as well as indirect effects that operate through the average characteristics of teachers in a school (e.g., teachers are more likely to switch schools as the proportion of novice teachers in the school increases). A natural question is how important these indirect effects are relative to the direct effects of these processes. To explore this, we run one more set of simulations that “turn off” the indirect pathways by setting all coefficients and interactions involving school-level teacher characteristics to zero. The results from these simulations are shown in Figure 6, which mirrors the format of Figure 5. The simulated effects of removing inequity pathways alone or in combination are very slightly more modest, but nearly indistinguishable from the estimates in Figure 5. This suggests that the overwhelming majority of teacher quality gaps are related to the experience levels of the specific teachers that leave, transfer, or are hired as opposed to how each of these processes are influenced by the average experience levels of teachers in schools.

5. Conclusions

We believe that the results from these simulations can inform both recommendations for policymakers seeking to close TQGs and provide a framework for considering the extent to which specific interventions might be expected to affect them. Of particular interest is the

finding that teacher mobility and hiring play outsized roles in contributing to TQGs. In contrast, teacher attrition *from the workforce* contributes relatively little to observed TQGs, at least when other processes (mobility and hiring) are still inequitable. That teacher attrition appears to play a relatively minor role in TQGs is notable given that it has been the focus of many efforts to close TQGs (e.g., Carver-Thomas & Darling-Hammond, 2017; Mehrotra et al., 2021).⁹ That said, a secondary but important finding from this paper is that very little of the impacts of attrition, mobility, and hiring on TQGs operate indirectly through the average teacher characteristics in different schools, meaning that teachers seem to respond most directly to school conditions rather than just follow one another's lead. In other words, retention of a few key senior teachers in a disadvantaged school is unlikely to stave off exodus of experienced teachers to other schools at a greater rate than is seen in an advantaged school, nor substantially attract experienced teachers to that school.

These findings suggest that we need to learn more about what affects recruitment to and transfers from disadvantaged schools as there is only a small literature on these issues. For instance, there is evidence that transfer provisions in collective bargaining agreements increase the inequity in the distribution of teacher experience by providing more senior teachers advantages within districts when schools downsize or new slots open in a district (Goldhaber et

⁹ For instance, an oft-cited example of an intervention to make teacher retention more equitable is a bonus policy implemented in North Carolina in the early 2000s in which certified math, science, and special education teachers working in disadvantaged schools received an annual bonus of \$1800. Clotfelter et al. (2008) found strong evidence that the policy reduced the combined attrition and mobility of teachers from disadvantaged schools, and importantly, found that experienced teachers demonstrated the greatest response to the policy. Our simulations suggest that *even such a successful intervention* will only have a substantial impact on TQGs if: a) overall patterns of mobility (not just rates of retention) are also impacted; and b) patterns of hiring to replace departing teachers are made more equitable as well.

al., 2016). But there is very little quantitative evidence about the factors that affect the equity of teacher quality related to recruitment and hiring processes.¹⁰

In developing the ABM used in this paper, we make a number of abstractions and simplifying assumptions. Foremost among these is our representation of the labor market as a probabilistic flow of individual teachers into and out of schools and applicant pools. There are ways in which the models of the processes that influence teacher quality gaps could become more sophisticated. We do not represent specifics such as how current or prospective teachers decide where to apply, how schools adjudicate between applicants to make job offers, or how teachers might respond when they receive multiple offers. One implication of this abstraction is that our model does not incorporate the possibility of adaptive responses across processes. For example, it is possible that a school experiencing high rates of attrition might respond in a way that changes its hiring process (e.g., recruitment of more experienced teachers in an effort to stem future attrition).

The assumptions used in our ABM design were driven by two primary factors. The first is that our research goal was to explain system-level patterns of TQGs to provide general insight into where policymakers and researchers should focus their attention to close them; as described above, we believe that we have been successful in this. The second is that we do not have theory, evidence, or data necessary for a finer-grained characterization of this system. Thus, our work can best be viewed as a first step along an iterative research process.

What we have done highlights several future research directions where investment of resources into data collection and analysis might be especially fruitful. First, it would be useful

¹⁰ The cost of any intervention is also an important consideration. Indeed, it is conceivable that the most cost-effective means of closing TQGs is to make a process inequitable (in favor of disadvantaged schools) to offset inequities in another process that may be harder to address.

to explore the extent to which our findings generalize to other measures of teacher quality and school disadvantage (e.g., socioeconomic status of students attending schools). Second, it is important to understand whether these findings are consistent across contexts and, if not, how they differ. Data collection and model application can help identify how context-specific factors such as differences in teacher supply (James et al., 2022), the regulatory environment determining employment eligibility (Redding & Smith, 2016), and the level of school segregation (James and Wyckoff, 2022), all might play a role in the relative and synergistic impact of hiring, mobility, and attrition on inequity.

Additional information about teacher labor markets would also permit researchers to add sophistication to the simulation framework. This would allow for modeling to incorporate nuance in how we represent teacher hiring and retention. Relatedly, this would allow us to make statements about how interventions targeted at specific schools or teachers might affect equity in those schools and the district as a whole over time. We view this as an important, if likely long-term, line of future research that can be used as a powerful tool to proactively inform efforts to close TQGs in K-12 public schools..

References

- Barnes, G., Crowe, E., & Schaefer, B. (2007). The cost of teacher turnover in five school districts: A pilot study. National Commission on Teaching and America's Future.
- Betts, J. R., Zau, A., & Rice, L. 2003. Determinants of student achievement: New evidence from San Diego. San Francisco: Public Policy Institute of California.
- Bleiberg, J., & Kraft, M. A. (2022). What happened to the k-12 education labor market during COVID? The acute need for better data systems. EdWorkingPaper: 22-544. Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/2xw0-v642>
- Boyd, D., Lankford, H., Loeb, S., Ronfeldt, M., & Wyckoff, J. (2011). The role of teacher quality in retention and hiring: Using applications to transfer to uncover preferences of teachers and schools. *Journal of Policy Analysis and Management*, 30(1), 88-110.
- Boyd, D., Lankford, H., Loeb, S., & Wyckoff, J. (2013). Analyzing the determinants of the matching of public school teachers to jobs: Disentangling the preferences of teachers and employers. *Journal of Labor Economics*, 31(1), 83-117.
- Boyd, D., Grossman, P., Lankford, H., Loeb, S., & Wyckoff, J. (2008). Who leaves? Teacher attrition and student achievement (No. w14022). National Bureau of Economic Research.
- Bruno, P., Rabovsky, S. J., & Strunk, K. O. (2020). Taking their first steps: The distribution of new teachers in school and classroom contexts and implications for teacher effectiveness. *American Educational Research Journal*, 57(4), 1688-1729.
- Bruno, P., & Strunk, K. O. (2019). Making the cut: The effectiveness of teacher screening and hiring in the Los Angeles Unified School District. *Educational Evaluation and Policy Analysis*, 41(4), 426-460.
- Carver-Thomas, D. & Darling-Hammond, L. (2017). Teacher turnover: Why it matters and what we can do about it. Learning Policy Institute, August 2017.

- Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *American Economic Review*, 104(9), 2633-79.
- Clotfelter, C., Glennie, E., Ladd, H., & Vigdor, J. (2008). Would higher salaries keep teachers in high-poverty schools? Evidence from a policy intervention in North Carolina. *Journal of Public Economics*, 92(5), 1352-1370.
- Clotfelter, C. T., Ladd, H. F., & Vigdor, J. L. (2009). The academic achievement gap in grades 3 to 8. *The Review of Economics and Statistics*, 91(2), 398-419.
- Clotfelter, C. T., Ladd, H. F., & Vigdor, J. (2005). Who teaches whom? Race and the distribution of novice teachers. *Economics of Education review*, 24(4), 377-392.
- The Education Trust. (2020). Is your state prioritizing teacher diversity & equity? <https://edtrust.org/educator-diversity/#US> .
- Epstein, J. M. (2012). Generative social science. Princeton University Press.
- Feng, L., & Sass, T. R. (2017). Teacher quality and teacher mobility. *Education Finance and Policy*.
- Geiger, T., & Pivovarova, M. (2018). The effects of working conditions on teacher retention. *Teachers and Teaching*, 24(6), 604-625.
- Goldhaber, D., Gross, B., & Player, D. (2011). Teacher career paths, teacher quality, and persistence in the classroom: Are public schools keeping their best?. *Journal of Policy Analysis and Management*, 30(1), 57-87.
- Goldhaber, D., Lavery, L., & Theobald, R. (2015). Uneven playing field? Assessing the teacher quality gap between advantaged and disadvantaged students. *Educational Researcher*, 44(5), 293-307.

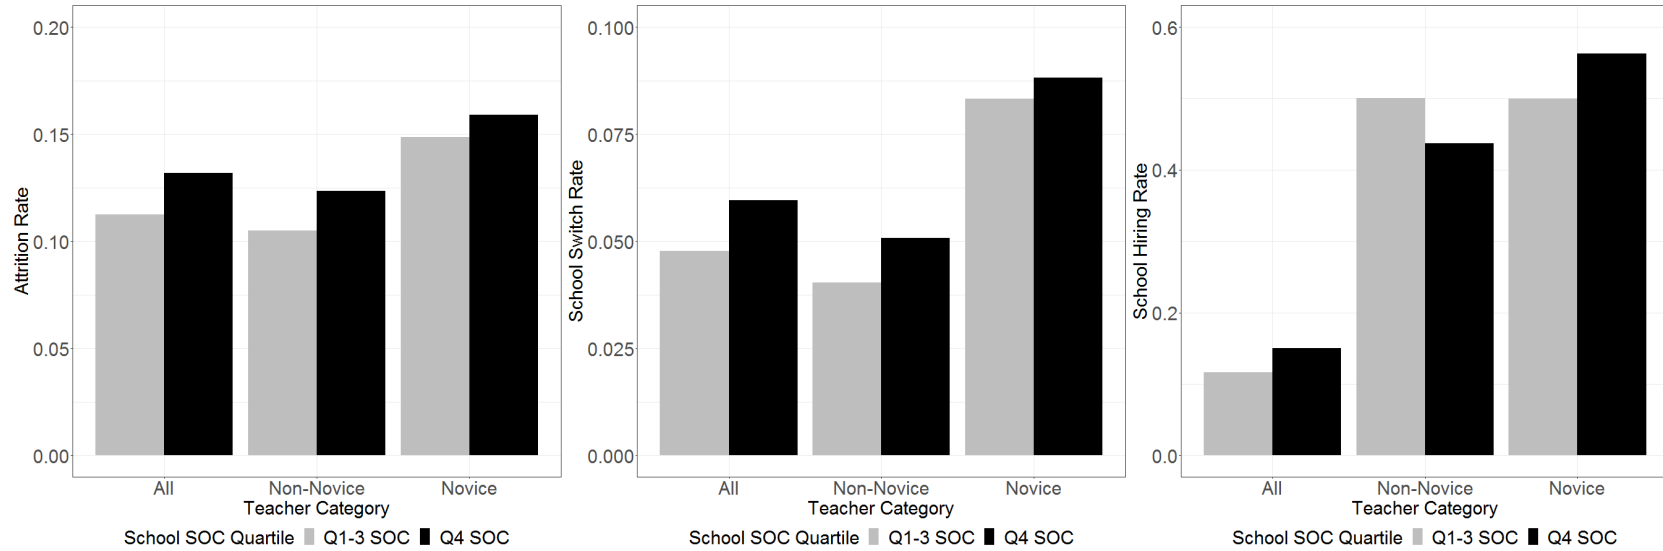
- Goldhaber, D., Lavery, L., & Theobald, R. (2016). Inconvenient truth? Do collective bargaining agreements help explain the mobility of teachers within school districts?. *Journal of Policy Analysis and Management*, 35(4), 848-880.
- Goldhaber, D., Theobald, R., & Fumia, D. (2022). The importance of teacher assignments in explaining STEM outcome gaps. *Social Science Research*.
- Goldhaber, D., Quince, V., & Theobald, R. (2018). Has it always been this way? Tracing the evolution of teacher quality gaps in U.S. public schools. *American Educational Research Journal*, doi 10.3102/0002831217733445.
- Gritz, M., and Theobald, N. (1996). The effects of school district spending priorities on length of stay in teaching. *The Journal of Human Resources*, 31(3) 477-512.
- Hammond, R. A. (2015). Considerations and best practices in agent-based modeling to inform policy. In *Assessing the use of agent-based models for tobacco regulation*. National Academies Press (US).
- Hanushek, E., Kain, J., and Rivkin, S. (2004). Why public schools lose teachers. *Journal of Human Resources*, 39(2), 326–354.
- Imazeki, J. (2005). Teacher salaries and teacher attrition. *Economics of Education Review*, 24(4), 431-449.
- Isenberg, E., Max, J., Gleason, P., Deutsch, J. (2021). Do low-income students have equal access to effective teachers? Evidence from 26 districts. *Educational Evaluation and Policy Analysis*.
- James, J., and Wyckoff, J. (2022). School segregation, teacher sorting, and the distribution of teachers. (EdWorkingPaper: 22-659). Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/tv5x-6t2>.

- James, J., Kraft, M. A., & Papay, J. (2022). Local Supply, Temporal Dynamics, and Unrealized Potential in Teacher Hiring. EdWorkingPaper: 22-518. Retrieved from Annenberg Institute at Brown University: <https://doi.org/10.26300/1yfe-gs84>
- Kalogrides, D., and Loeb, S. (2013). Different teachers, different peers: The magnitude and effects of student sorting within schools.” *Educational Researcher*, 42(6).
- Kasman, M., Breen, N., & Hammond, R. A. (2021). Complex systems science. *The Science of Health Disparities Research*, 243-256.
- Krieg, J. M. (2006). Teacher quality and attrition. *Economics of Education review*, 25(1), 13-27.
- Ladd, H. F. (2011). Teachers’ perceptions of their working conditions: How predictive of planned and actual teacher movement?. *Educational Evaluation and Policy Analysis*, 33(2), 235-261.
- Lankford, H., Loeb, S., & Wyckoff, J. (2002). Teacher sorting and the plight of urban schools: A descriptive analysis. *Educational evaluation and policy analysis*, 24(1), 37-62.
- Loeb, S., Kalogrides, D., & Bêteille, T. (2012). Effective schools: Teacher hiring, assignment, development, and retention. *Education Finance and Policy*, 7(3), 269-304.
- Maroulis, S., Guimera, R., Petry, H., Stringer, M. J., Gomez, L. M., Amaral, L. A. N., & Wilensky, U. (2010). Complex systems view of educational policy research. *Science*, 330(6000), 38-39.
- Mehrotra, S., Morgan, I., & Socol, A. (2021). Getting Latino students better access to non-novice teachers. Education Trust, <https://edtrust.org/wp-content/uploads/2014/09/Getting-Latino-Students-Better-Access-to-Non-Novice-Teachers-December-2021.pdf>.
- Miller, J. H. & Page, S. E. (2009). Complex adaptive systems. Princeton university press.

- Redding, C., & Smith, T. M. (2016). Easy in, easy out: Are alternatively certified teachers turning over at increased rates?. *American Educational Research Journal*, 53(4), 1086-1125.
- Rich, M. (2014). U.S. to focus on equity in assigning of teachers. *The New York Times*.
- Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. *Econometrica*, 73(2), 417-458.
- Rodriguez, L. A., Nguyen, T. D., & Springer, M. G. (2023). Revisiting Teaching Quality Gaps: Urbanicity and Disparities in Access to High-Quality Teachers Across Tennessee. *Urban Education*, 00420859231153409.
- Ross, E. (2019). Ensuring equitable access to great teachers: State policy priorities. *Phi Delta Kappan*. <https://kappanonline.org/ensuring-equitable-access-great-teachers-state-policy-priorities-ross/>.
- Sass, T. R., Hannaway, J., Xu, Z., Figlio, D. N., & Feng, L. (2012). Value added of teachers in high-poverty schools and lower poverty schools. *Journal of Urban Economics*, 72(2-3), 104-122.
- Scafidi, B., Sjoquist, D. L., & Stinebrickner, T. R. (2007). Race, poverty, and teacher mobility. *Economics of Education Review*, 26(2), 145-159.
- Xu, Z., Özek, U., & Hansen, M. (2015). Teacher performance trajectories in high-and lower-poverty schools. *Educational Evaluation and Policy Analysis*, 37(4), 458-477

Tables and Figures

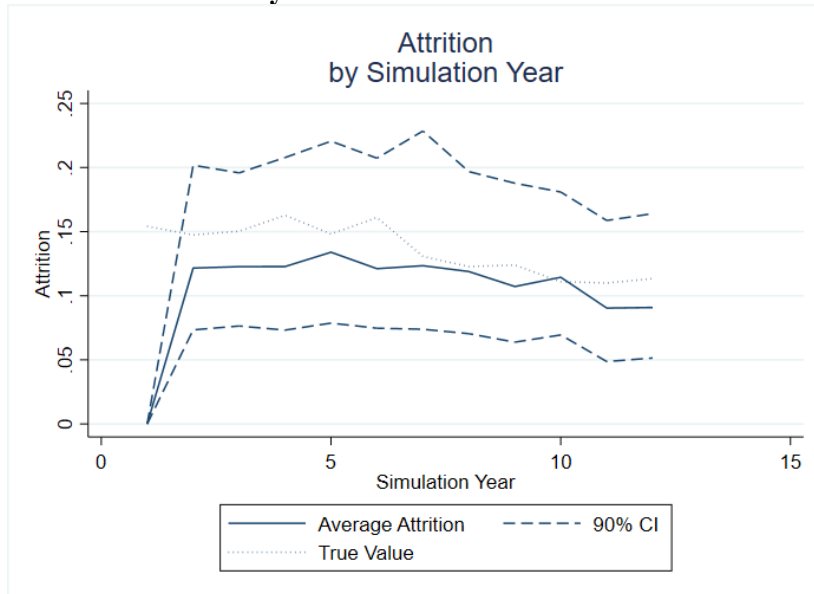
Figure 1. Teacher attrition, mobility, and hiring by quartile of school percent students of color (SOC) and teacher experience



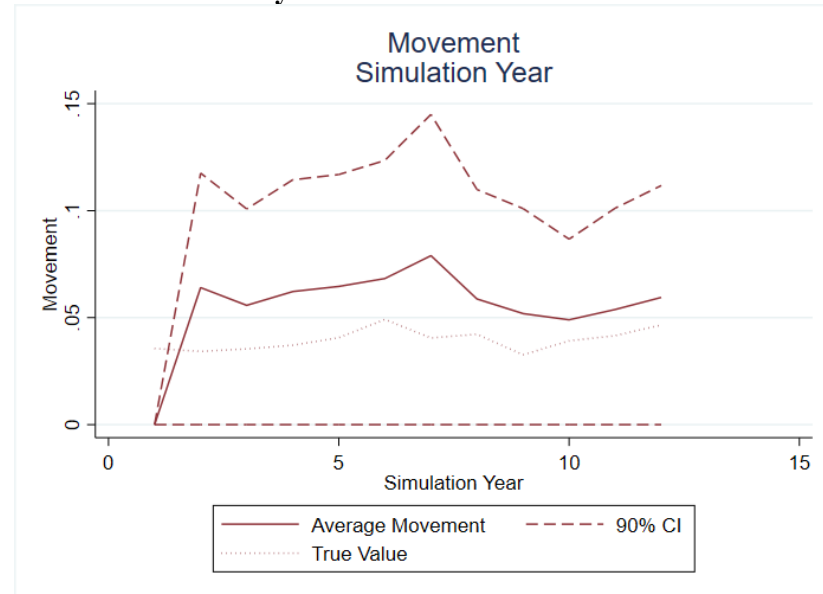
Note. SOC = students of color, Q1-3 = bottom three quartiles, Q4 = top quartile

Figure 2. Simulation processes at observed values (all processes inequitable)

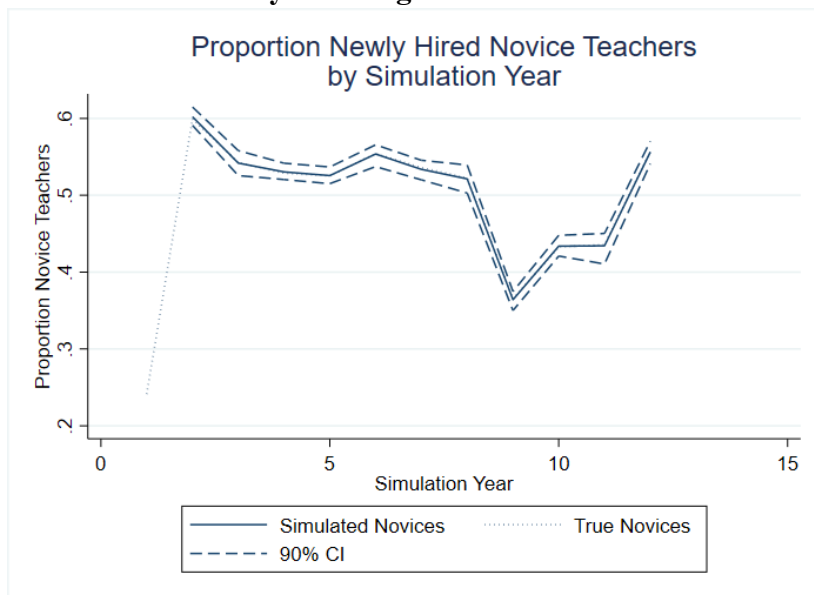
Panel A. Probability of attrition



Panel B. Probability of school move



Panel C. Probability of hiring novice teacher



Panel D. Proportion of novice teachers

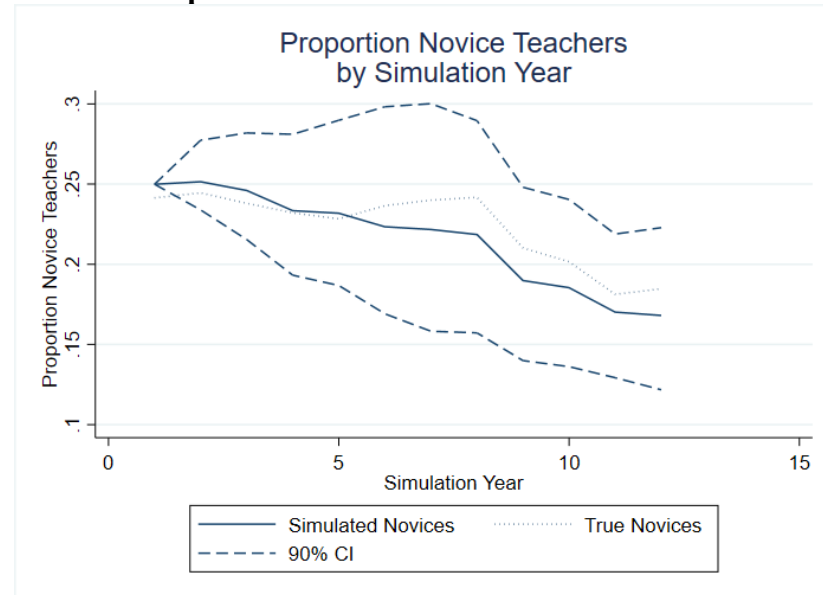
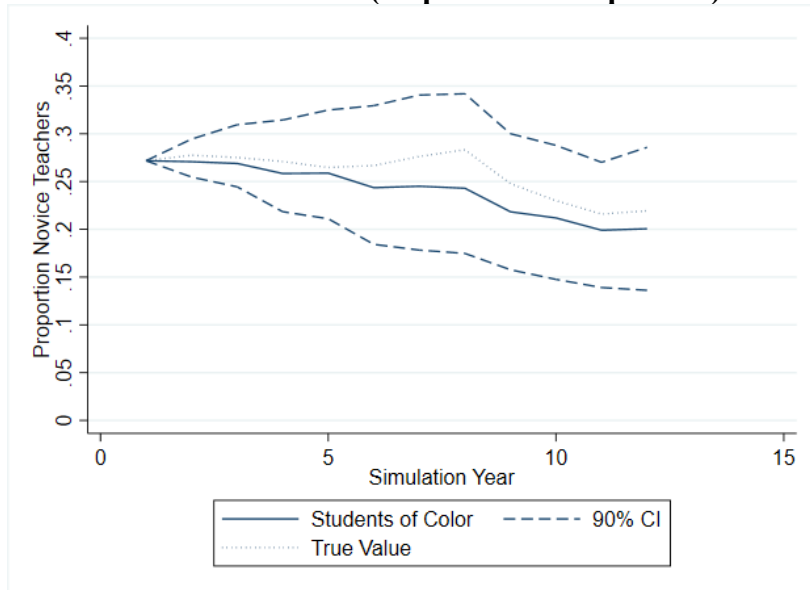
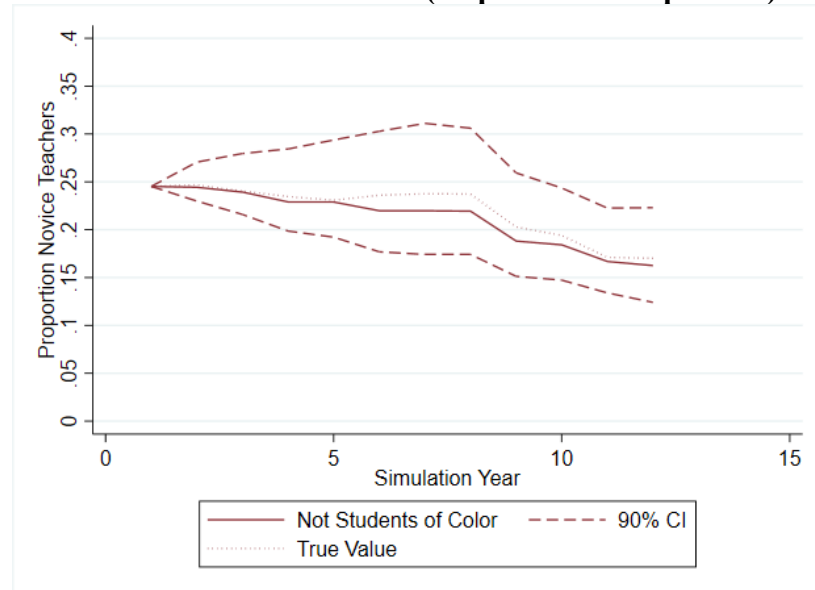


Figure 3. Simulated novice teacher exposure and teacher quality gaps, by simulation year

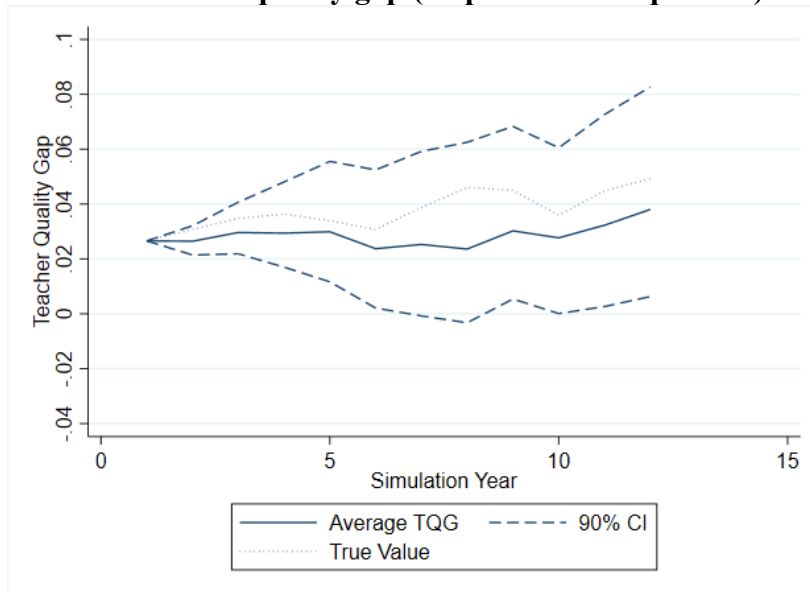
Panel A. Students of color (all processes inequitable)



Panel B. Not students of color (all processes inequitable)



Panel C. Teacher quality gap (all processes inequitable)



Panel D. Teacher quality gap (all processes equitable)

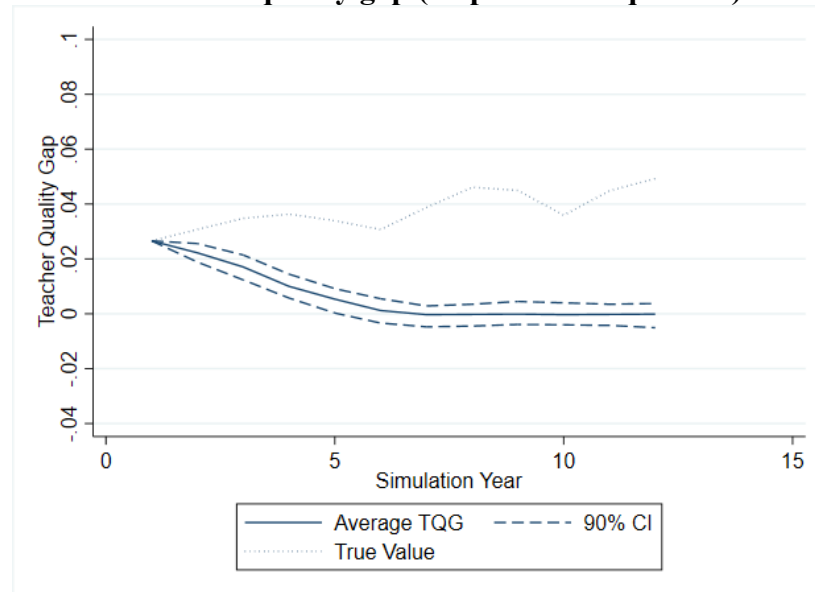
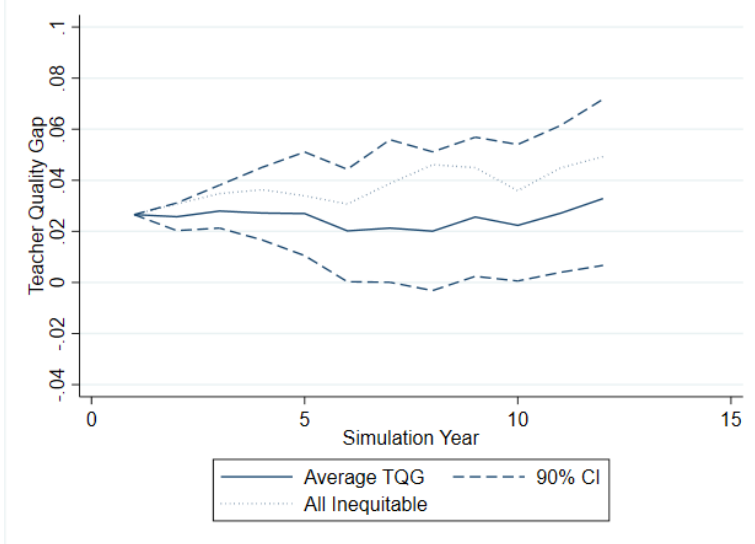


Figure 4. Individual simulation results over simulation years

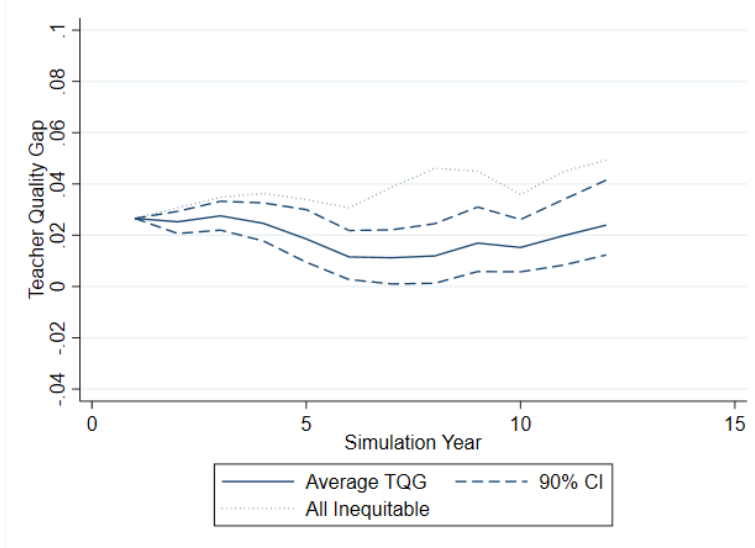
Panel A. Teacher quality gap (all processes inequitable)



Panel B. Teacher quality gap (attrition equitable)



Panel C. Teacher quality gap (mobility equitable)



Panel D. Teacher quality gap (hiring equitable)

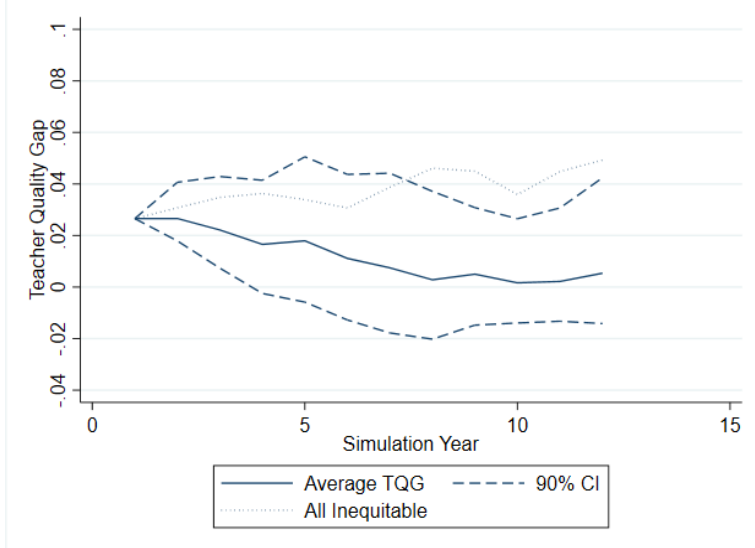
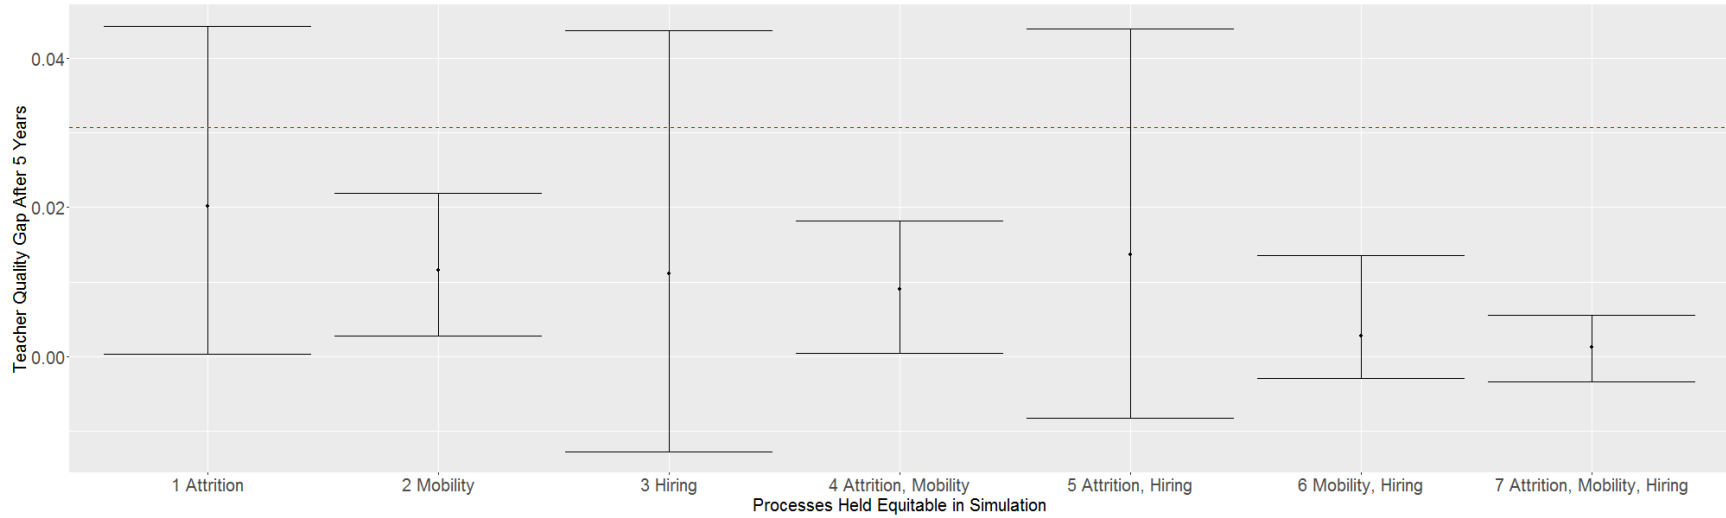


Figure 5. Results across simulations, selected simulation years

Panel A. Teacher quality gaps after 5 years (horizontal line = true gap)



Panel B. Teacher quality gaps after 10 years (horizontal line = true gap)

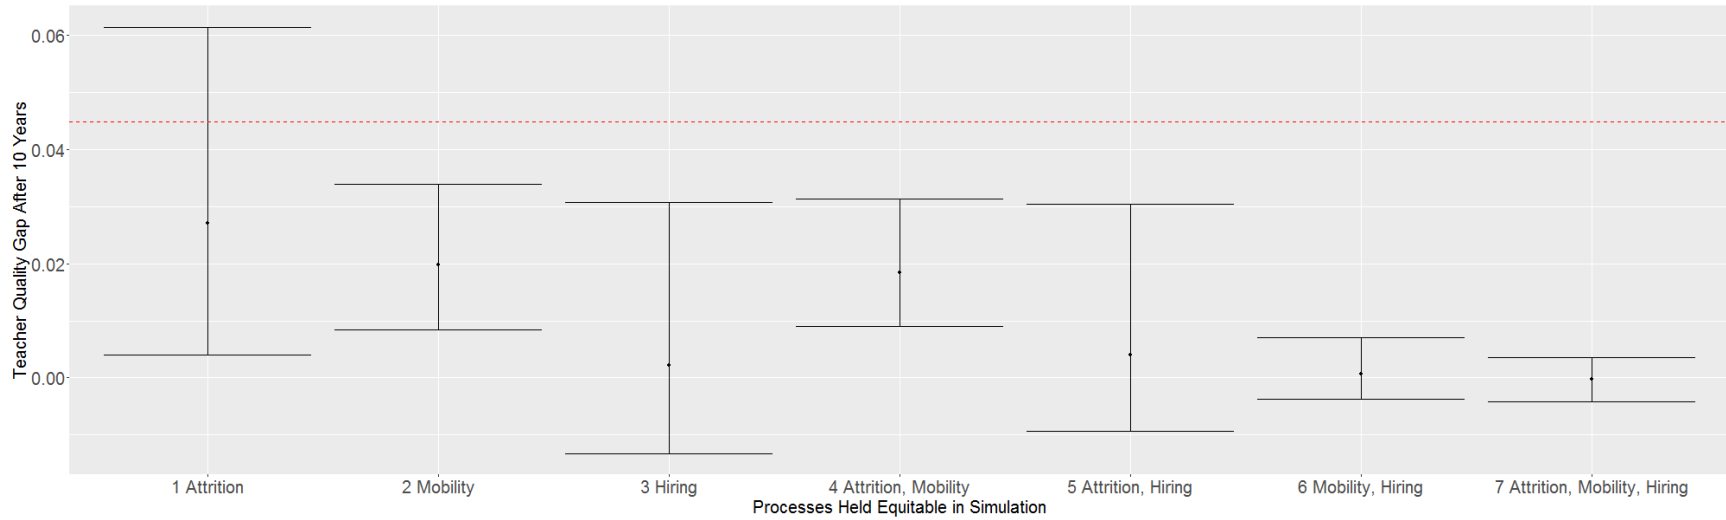
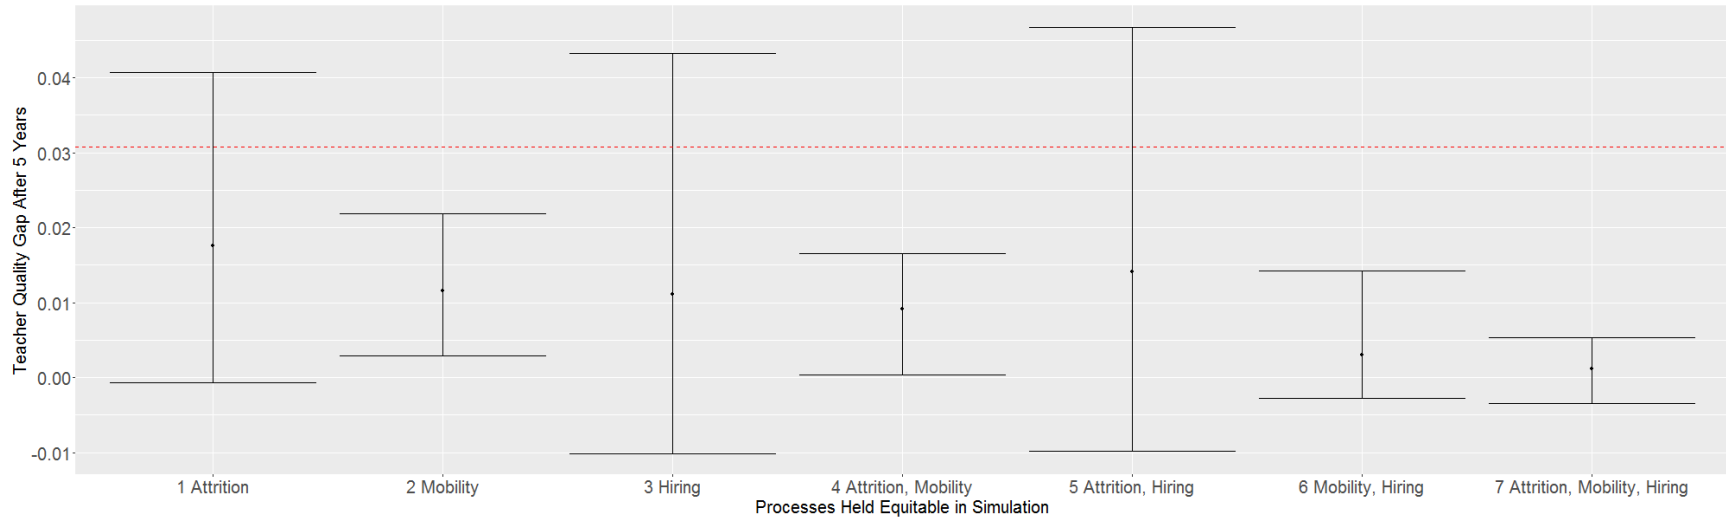


Figure 6. Results across simulations, selected simulation years, not allowing for indirect effects

Panel A. Teacher quality gaps after 5 years (horizontal line = true gap)



Panel B. Teacher quality gaps after 10 years (horizontal line = true gap)

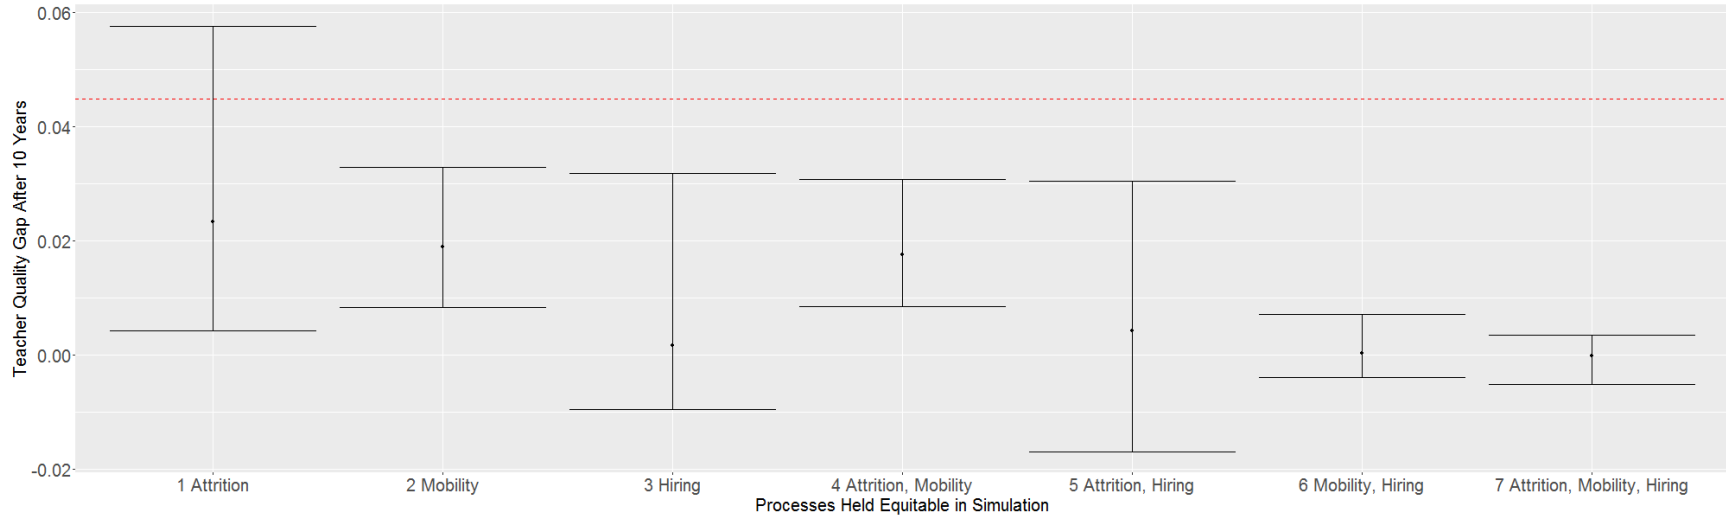


Table 1. Selected coefficients from first-stage Bayesian regression models

Column	(1)	(2)	(3)
Model	Multinomial Logit for Attrition and Mobility		Bernoulli New Hiring Model
Outcome	Teacher Attrition	Teacher Mobility	School Hires New Teacher
Novice Teacher	0.56*** (0.03)	0.14 (0.05)	
Teacher Experience	0.02*** (0.00)	-0.04 (0.00)	
School Proportion of Novice Teachers	0.15 (0.13)	0.43** (0.20)	-0.52 (0.27)
Average Experience in School	0.01 (0.01)	-0.03*** (0.01)	0.05*** (0.01)
Percent SOC in School	0.77** (0.32)	-0.67 (0.50)	0.98* (0.55)
Novice Teacher * Percent SOC in School	-0.01 (0.09)	-0.10 (0.15)	
Teacher Experience * Percent SOC in School	0.02*** (0.00)	0.03*** (0.01)	
School Proportion of Novice Teachers * Percent SOC in School	-0.26 (0.37)	0.09 (0.55)	0.14 (0.64)
Average Experience in School * Percent SOC in School	-0.05** (0.02)	0.00 (0.03)	-0.05* (0.03)
Observations	269,725		12,367

Note. *P*-values from two-sided *t*-test: **p*<.10; ***p*<.05; ****p*<.01. SOC = students of color.