

**A Partial Identification
Approach to Identifying the
Determinants of Human
Capital Accumulation:
An Application to Teachers**

Nirav Mehta

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

A Partial Identification Approach to Identifying the Determinants of Human Capital Accumulation: An Application to Teachers

Abstract

This paper views teacher quality through the human capital perspective. Teacher quality exhibits substantial growth over teachers' careers, but why it improves is not well understood. I use a human capital production function nesting On-the-Job-Training (OJT) and Learning-by-Doing (LBD) and experimental variation from [Glewwe et al. \(2010\)](#), a teacher incentive pay experiment in Kenya, to discern the presence and relative importance of these forces. The identified set for the OJT and LBD components has a closed-form solution, which depends on experimentally estimated average treatment effects. The results provide evidence of an LBD component, as well as an informative upper bound on the OJT component.

JEL-Codes: I200, I280, J200, J240, J450, C100.

Keywords: human capital, teacher quality, on-the-job training, learning-by-doing, partial identification.

Nirav Mehta
Department of Economics
University of Western Ontario
London / ON / Canada
nirav.mehta@uwo.ca
web: <https://niravmehtax.github.io>

March 30, 2022

I thank Enrique Martin Luccioni and Tian Liu for research assistance, and Lance Lochner, Rory McGee, Magne Mogstad, Seth Richards-Shubik, and Camillie Terrier for their comments. This paper also benefited from comments from the 2019 meeting of the CESifo Economics of Education Group. I acknowledge funding from the SSHRC Insight Development Grant Program.

1 Introduction

Teacher quality, typically measured by a teacher’s value added to student achievement, is an important determinant of student achievement and economic growth (Rivkin et al., 2005; Hanushek, 2011). While the research and policy focus has typically been on cross-sectional variation in teacher quality, researchers have also documented substantial growth in quality over teachers’ careers (Harris and Sass, 2011; Wiswall, 2013; Papay and Kraft, 2015). For example, using data from North Carolina, Wiswall (2013) finds that the average growth in teacher quality over 35 years of experience is equal to almost one standard deviation of the cross-sectional teacher quality among novice teachers.

While the growth in quality over teachers’ careers is clearly important, little is known about why it occurs.¹ Economists have developed main two theories to explain how workers accumulate human capital, which is then used to produce output: *On-the-Job Training* (OJT) and *Learning-by-Doing* (LBD). In the “pure OJT” model, workers (in the current application, teachers) allocate work time away from production (here, classroom teaching) to invest in their human capital (here, e.g., teachers’ professional development) (Becker, 1964; Ben-Porath, 1967).² This model implies a tradeoff between current and future production because of the multitask problem (Hölmstrom and Milgrom, 1991), which emerges because OJT investments are unobserved and, therefore, noncontractible. This tradeoff is not present under the “pure LBD” model, wherein workers accumulate human capital via the act of production (Rosen, 1972; Blinder and Weiss, 1976).³ The strength of these forces has implications for the design of effective education policy. For example, if OJT was the main driver of teacher human capital accumulation then an incentive pay scheme that increased current achievement could reduce long-run teacher quality, by diverting teachers from making human capital investments. This would not be the case if instead LBD were the dominant force.

Unfortunately, we do not know about the importance of OJT versus LBD forces among teachers, or even workers more generally. Killingsworth (1982) shows there are no clear general testable implications of the OJT versus LBD theories of human capital accumulation when using standard observational data on workers. Additionally, as Heckman et al. (2002) discuss, identification in such contexts is further hampered by equilibrium re-adjustments that occur in the private sector.⁴ This has left researchers and policymakers lacking even basic, qualitative information about even the presence of OJT versus LBD forces for workers of any type. Indeed, prior empirical research has adopted “pure” specifications, only allowing for either OJT or LBD as the force generating

¹I discuss below the two economics papers studying how teacher experience might affect quality.

²For an example of an OJT-intensive job, consider being a student in school.

³For an example of an LBD-intensive job, consider “chicken sexing”, where workers gain an ability to determine the sex of baby chickens through their experience sorting chicks (McWilliams, 2018).

⁴Researchers using the OJT framework (see, e.g., Haley, 1976; Heckman, 1976) typically circumvent the fact that OJT investments are unobserved by solving for the optimal investment path in a worker’s dynamic program, which yields estimates of the importance of OJT in the technology of human capital accumulation. See Kuruscu (2006) for an example using a different approach.

human capital.⁵

To fill this gap, this paper develops a framework that yields new findings about the roles played by OJT and LBD forces in the technology governing teacher human capital accumulation. In light of the prominent identification difficulties, I adopt a partial identification approach to make the source of identification as clear as possible. The identified set contains all of the values of the OJT and LBD components of teacher human capital accumulation consistent with the data; this can then be projected onto the marginal identified set for either component. I estimate the identified set using data from an experimental intervention, [Glewwe et al. \(2010\)](#), which studies a teacher performance pay scheme enacted across many sites across Kenya. Labor prices are often fixed in the education sector because teachers are typically paid according to public salary schedules ([Podgursky and Springer, 2011](#)) and worker-specific output is observed, even if noisily so, making teacher quality trajectories a particularly attractive context for discriminating between OJT and LBD.

The starting point is workhorse models that have become ubiquitous in empirical research over the last five decades. The value-added (VA) model ([Hanushek, 1971](#); [Murnane, 1975](#)), which measures the contribution of different educational inputs to the production of student achievement, is used in the overwhelming mass of research on teacher quality (for more recent examples, see [Kane et al., 2013](#); [Chetty et al., 2014](#)).⁶ The literature studying human capital accumulation has from its inception typically used log-linear specifications when considering either OJT ([Brown, 1976](#); [Haley, 1976](#); [Heckman, 1976](#)) or LBD ([Blinder and Weiss, 1976](#)), and these specifications also continue to dominate in more recent research (see, e.g., [Heckman et al., 1998](#); [Fan et al., 2015](#); [Blandin, 2018](#)).⁷

I extend the modal specifications in the literature to nest both OJT and LBD forces. The resulting specification leverages experimentally estimated average treatment effects (ATEs) to report key information about average tendencies and guarantees that identification is not driven by assumed nonlinearities in functional forms, while not being very restrictive compared to specifications typically used in the literature. As in the “pure OJT” model, OJT investments are also unobserved in my framework. This closely matches the teacher quality application, as a large body of research has found that variables that seem like natural measures of training available to researchers, such as formal professional development, teacher certification, or additional education (e.g., Master’s degrees), do not help predict teacher quality ([Hanushek, 2003](#); [Hanushek and Rivkin, 2006](#); [Harris and Sass, 2011](#); [Podgursky and Springer, 2011](#); [Jackson et al., 2014](#)), and the general lack of avail-

⁵This is also true of recent research studying heterogeneity in the returns to experience for teachers ([Kraft and Papay, 2014](#)) or other types of workers (see, e.g., [Shaw and Lazear, 2008](#); [Haggag et al., 2017](#)). Here, too, the distinction between OJT versus LBD has implications for policy, as heterogeneous impacts of interventions would be exacerbated or attenuated, depending on how human capital was generated.

⁶See [McCaffrey et al. \(2003\)](#) and [Hanushek and Rivkin \(2012\)](#) for detailed discussions.

⁷As discussed by [Willis \(1985\)](#), these specifications for human capital production had already had a long history of use in labor economics as of three decades ago. These specifications are also consistent with those in the statistical literature spawned by [Abowd et al. \(1999\)](#). See [Shaw \(1989\)](#) for an empirical model of LBD using a different specification and see [Fu et al. \(2019\)](#) for an empirical model of OJT using a different specification.

ability of data on how teachers allocate their work time (Hanushek and Rivkin, 2012). This means my framework generates the classic identification problem of separating OJT and LBD forces.

Identification (i.e., informative parameter bounds) is obtained in the current paper because Glewwe et al. (2010) includes a follow-up measurement of the effects of the program. The intuition can be outlined in the following example: suppose that achievement increased in the treatment group while the incentive scheme was in place (as was the case in Glewwe et al. 2010) and that the random assignment of the intervention was balanced (as was also the case). Then the increase in achievement during the intervention must have come from an increase in treatment-group teachers' human capital allocated to production. Further, a positive treatment effect on achievement after the intervention ended, beyond that accounted for by the persistence of the during-intervention increase in achievement, would imply that teacher quality increased, pointing to the presence of the LBD component. On the other hand, a post-intervention treatment effect lower than that which could be explained by the persistence of the prior achievement effect would imply that teacher quality decreased, pointing to the presence of the OJT component. Without the follow-up data the identified set would be uninformative, meaning we would remain in the typical case in which we could not separate OJT and LBD forces.

I find that the 95% confidence set for the identified set for each parameter is informative. Specifically, the lower bound on the LBD component is greater than zero and the upper bound on the OJT component is lower than its uninformative level. I reject the “pure OJT” specification (in which LBD plays no role) at the 5% significance level, but I cannot reject the “pure LBD” specification, as the confidence set for the OJT component contains zero. When further imposing returns-to-scale type assumptions (i.e., restrictions on the sum of the OJT and LBD components in human capital production), bounds are tighter. Under the strongest assumption, constant returns to scale, the 95% confidence set for the LBD component lies strictly above that for the OJT component, meaning in this case we can also say with appreciable certainty that the LBD component is larger than the OJT component.

This paper's framework generates several contributions. First, it works with a human capital production function that nests both the OJT and LBD mechanisms, as opposed to the “pure” specifications used in prior empirical work. This allows for standard hypothesis testing about important qualitative features, in particular, the presence of either force in the human capital production function. Second, it uses a partial identification approach that features a closed-form solution for the identified set: it is a locus characterized by a line. This characterization sharply contrasts with the literature's lack of a proof showing how to separately identify OJT and LBD forces, yet obtains quite naturally here. Moreover, the primary input to the identified set is an experimentally estimated ATE, commonly viewed as the “gold standard” in causal inference. Thus, it is also easy to estimate the identified set and to do inference on it, which is rare in the partial identification literature (Imbens and Manski, 2004; Stoye, 2009; Tamer, 2010). Finally, the approach does not have to make assumptions regarding the optimality of input choices or outcomes observed

in the data.⁸

This paper seeks to understand the evolution of teacher quality from the perspective of the dominant theories of human capital accumulation. Somewhat separately from the aforementioned human capital literature, a small but substantively important set of papers examines other channels underlying teachers’ improvement as they gain more experience. [Ost \(2014\)](#), which measures the returns to teachers’ general and grade-specific experience, finds both to be important determinants of teacher quality growth. [Cook and Mansfield \(2016\)](#) extends this work to also allow for general and context-specific permanent components to teacher quality. The current paper nicely complements these papers by explicitly viewing teacher quality through the lens of the main conceptual frameworks for human capital, and by identifying and separating the OJT and LBD channels of human capital accumulation, which are not necessary to distinguish given these other papers’ goals.⁹ This paper also complements the extensive literature studying teacher quality more generally, recently discussed in [Hanushek and Rivkin \(2006\)](#); [Jackson et al. \(2014\)](#); [Strøm and Falch \(2020\)](#), and the sub-literature on teacher incentives (see, e.g., [Hanushek and Raymond, 2005](#); [Muralidharan and Sundararaman, 2011](#); [Imberman and Lovenheim, 2015](#); [Petronijevic, 2016](#)).¹⁰

2 Data and Variables

Study Design [Glewwe et al. \(2010\)](#) implemented a teacher incentive pay scheme in Kenya, which provided bonuses to teachers at schools where students did well on standardized exams administered as part of the standard curriculum in Kenya. The treatment group was exposed to the incentive scheme in 1998 and 1999, which I refer to as the “active-treatment” years. In addition to being observed during active-treatment years, outcomes for both the treatment and control groups were also observed for both a pre-treatment year (1997), and a post-active-treatment year (2000). [Table 1](#) summarizes the study design.

Variables I report on the variables that are needed for this paper’s analysis; readers are referred to [Glewwe et al. \(2010\)](#) for additional details. I measure output using the incentivized standardized achievement test presented in the main results of [Glewwe et al. \(2010\)](#), which is an average of seven standardized subject-specific achievement tests, where standardization was performed with respect to the means and standard deviations of the control group. It is important to note that the

⁸Researchers with different goals, say, of examining heterogeneous impacts or understanding how teacher quality would change under counterfactual incentive pay schemes, would have to make stronger statistical assumptions and/or impose behavioral assumptions. This would naturally represent fruitful and complementary avenues for future research.

⁹These papers may be viewed as part of the substantial empirical literature studying task-specific human capital for more general worker contexts. For recent examples, see [Poletaev and Robinson \(2008\)](#); [Sanders \(2010\)](#); [Yamaguchi \(2012\)](#); [Robinson \(2018\)](#). [Sanders and Taber \(2012\)](#) contains a detailed discussion of this and other extensions of the one-dimensional human capital model.

¹⁰Although less tightly related, this paper also relates to the literature studying teacher labor markets (see, e.g., [Dolton and Klaauw, 1999](#); [Stinebrickner, 2001](#); [Behrman et al., 2016](#); [Tincani, 2021](#); [Biasi, 2021](#); [Bobba et al., 2021](#); [Biasi et al., 2021](#)).

Table 1: Design of the [Glewwe et al. \(2010\)](#) Teacher Incentive Pay Experiment

Group	Teacher incentive pay?			
	1997	1998	1999	2000
Treatment	✗	✓	✓	✗
Control	✗	✗	✗	✗

Note: “✓” indicates that members of the group (e.g., “Treatment”) received incentive pay in that year and “✗” indicates that the group did not receive incentive pay that year.

experiment was balanced; of particular importance, the mean test scores in the pre-active-treatment year were balanced between the control and treatment groups.¹¹

The data also include information about whether teachers were present at schools and in classrooms during site visits, which are used as measures of teacher labor supply. While labor supply is the (observable) input to human capital accumulation in the LBD framework, the classic OJT framework treats OJT investments as unobserved. I do the same, for two reasons. First, [Kuruscu \(2006\)](#), who considers general worker contexts, argues that it is more appropriate to treat on-the-job training as unobserved because some training may not be observed by researchers; while this complicates identification in the typical setting, the current paper’s approach allows us to learn about the importance of OJT and LBD forces despite this fundamental data limitation. Second, the literature on teacher quality, which has largely focused on the United States, has consistently found that observable training measures like Master’s degrees, professional development work, or additional certificates are poor predictors of teacher quality ([Hanushek, 2003](#); [Hanushek and Rivkin, 2006](#); [Harris and Sass, 2011](#); [Podgursky and Springer, 2011](#)). Conceptually, OJT investments result in more effective teaching the next period, for example, a teacher might reflect on how her class went and then use this to improve her craft. If the teachers did not have any scope for making OJT investments, say, because they were inundated with classroom teaching duties, this would show up in the estimates because I allow there to be no role played by this force. Treating OJT investments as unobserved is therefore not only consistent with how the literature typically models OJT investments, but also makes the findings more applicable to contexts where measures of training have been found to have little bearing on teacher quality.

3 Model

The model provides the foundation for the partial identification of the components governing human capital accumulation. It describes how teachers allocate their human capital between on-the-job skill investments, production of output (i.e., contemporaneous student achievement), and leisure. The specifications for the production of human capital and output enable a closed-form constructive identification proof and calculation of bounds using ATE estimates, meaning the analysis captures

¹¹As in [Glewwe et al. \(2010\)](#), I exclude one of the districts (Teso), which did not offer the achievement test in the first year.

average behavior.¹² They are also not very restrictive either in terms of implications for behavior,¹³ or compared with specifications commonly used in the literature.¹⁴

3.1 Environment

Let the periods be indexed by $t = 0, 1, 2$, where $t = 0$ is the pre-treatment period, $t = 1$ is the “active-treatment” period, in which teachers in the treatment group are offered output-based incentives, and $t = 2$ is the post-active-treatment period, where teachers in either treatment or control group no longer are offered the incentives. I do not refer to $t = 1$ as the “treatment period” because, post active treatment, teachers in the treatment group may have been affected by their decisions made during active treatment.

The human capital of teacher i in period t , θ_{it} , is produced according to

$$\theta_{it} = \delta_{\theta}\theta_{it-1} + \delta_I I_{it-1} + \delta_h h_{it-1}, \quad (1)$$

where θ_{it-1} is teacher i ’s human capital last period, I_{it-1} is teacher i ’s on-the-job investment in human capital last period, and h_{it-1} is the time teacher i spent on production (e.g., actively teaching, engaging in teaching preparation, providing feedback to students, etc.) last period. Output produced by teacher i teaching student j in period t , y_{ijt} , which is measured by performance on a standardized achievement test, follows a standard value-added specification (Hanushek, 1971; Murnane, 1975; Hanushek, 1979):¹⁵

$$y_{ijt} = \beta_{\theta}\theta_{it} + \beta_h h_{it} + \beta_y y_{ijt-1} + \epsilon_{ijt}, \quad (2)$$

where y_{ijt-1} is student j ’s prior test achievement (meaning β_y measures the persistence of student knowledge), $\beta_{\theta}\theta_{it} + \beta_h h_{it}$ is teacher i ’s value added to student j ’s achievement in period t (i.e., *quality*), and ϵ_{ijt} is an ex post IID productivity shock.¹⁶

Time spent producing output, h_{it} , and on-the-job investment, I_{it} are rival, as captured by the

¹²While averages are not the only conceivable objects of policy interest, they certainly have garnered a vast amount of interest by researchers and policymakers, as they may average out unobserved heterogeneity and also may be useful for certain normative considerations (e.g., maximizing output).

¹³The human capital production function does not assume that growth in human capital is constant over time. Indeed, as the model does not assume optimality of teacher behavior, it allows for, e.g., decreasing OJT investments over the course of a teacher’s career, which would result in the concave value added profiles documented in the literature.

¹⁴The model developed in this section could be viewed as a log-linear approximation to a specification considering the behavior of a representative teacher and student in each of the control and treatment groups (see Appendix A), and is also consistent with (specifically, nested by) the commonly used translog specification used in the burgeoning literature using dynamic factor models to understand skill growth (see, e.g., Agostinelli and Wiswall, 2020; Bono et al., 2020; Freyberger, 2021). It is important to note that, while the link developed in the appendix might provide useful context, it is in no way necessary to this paper’s analysis.

¹⁵Recent research supports the view that controlling for prior achievement, as is done in a value-added model, does a reasonably good job of controlling for unobserved prior inputs (see, e.g., Kinsler, 2012; Chetty et al., 2014).

¹⁶I abstract from student characteristics here; as will become clear, what matters for the analysis is obtaining a consistent estimate of β_y .

constraint

$$h_{it} + I_{it} + l_{it} = H, \tag{3}$$

where *total work time* is $h_{it} + I_{it}$, and l_{it} is teacher i 's leisure time in period t . The variables h , I , and l could be viewed as shares of a teacher's total time, or of a teacher's overall focus or potential effort respectively allocated to production, OJT investment, and leisure.¹⁷ Under the latter interpretation, a teacher has a fixed "budget" of focus/potential effort, which can be allocated between h , I , and l ; naturally, a teacher must be working to engage in production or investment. I use the term "time" hereafter, but it might be useful to keep these alternative interpretations in mind. Although, as discussed below, there was no evidence of an average labor supply response to the incentive pay scheme, I explicitly model labor supply for the potential application of this approach to other contexts, and also allow for a labor supply response in the empirical application.

I assume that $\delta_k \in D_k$, where $D_k = [0, \bar{\delta}]$, for $k = \theta, I, h$, and that $\beta_k \in B_k$, where $B_k = [0, \bar{\beta}]$, for $k = \theta, h, y$. The lower bound of zero for each parameter captures the natural assumption means that inputs cannot have negative effects, *ceteris paribus*. The upper bound for each parameter (i.e., $\bar{\delta}$ or $\bar{\beta}$) is taken to be large; I discuss below how the specific values of $\bar{\delta}$ and $\bar{\beta}$ do not affect this paper's main findings.

I define the "pure OJT" specification for the human capital production function as $\delta_I > 0$ and $\delta_h = 0$. Analogously, in the "pure LBD" specification, we have $\delta_h > 0$ and $\delta_I = 0$. Some scenarios representing different possibilities for true combinations of (δ_h, δ_I) are illustrated in Figure 1 (bounds for the *identified set* for (δ_h, δ_I) , which is denoted $d_h \times d_I$, are derived in Section 3.2). For example, the point labeled "pure OJT", on the horizontal axis, features a positive OJT component, with no LBD component, in contrast to the point labeled "pure LBD", on the vertical axis. The interior point, labeled "OJT and LBD both present", represents the possibility that teachers accumulate human capital via both OJT and LBD components; in contrast, the "pure" versions of the OJT and LBD human capital production functions are mutually exclusive.


3.2 Derivation of Bounds

This section develops bounds for (δ_h, δ_I) that only depend on period-specific estimates of the average treatment effects of the intervention on student achievement and on labor supply, and an estimate of the persistence of student knowledge (which could be estimated using the same dataset, or obtained from another source). The bounds are sharp, meaning they contain only the values of (δ_h, δ_I) that cannot be rejected given the data. For example, the sharp bound on δ_I would increase the lower bound for δ_I as much as possible while still being consistent with the data; if the lower bound were greater than zero one would reject the pure LBD specification.

Define the mean difference between the treatment and control groups for variable z in period t

¹⁷Because the data contain measures of leisure time, we can also identify total work time. However, the individual components of total work time are not separately observed, which means the scale of $h_{it} + I_{it}$ is fixed by the measure of l_{it} (which is discussed in detail in Section 4.1).

Figure 1: Examples of OJT and LBD Specifications



as $\Delta z_t := \overline{z_t^T} - \overline{z_t^C}$, where $\overline{z_t^T}$ and $\overline{z_t^C}$ respectively denote the treatment and control group means of z in t .

First, note that in any period the constant time endowment implies that

$$\Delta h_t + \Delta I_t + \Delta l_t = 0. \quad (4)$$

For the pre-treatment period, $t = 0$, the mean difference in achievement between the treatment and control groups is

$$\Delta y_0 = \beta_\theta \underbrace{\Delta \theta_0}_{=0} + \beta_h \underbrace{\Delta h_0}_{=0} + \beta_y \underbrace{\Delta y_{-1}}_{=0} + \underbrace{\Delta \epsilon_0}_{=0} = 0, \quad (5)$$

i.e., a balanced experimental design implies there will be no average difference in the pre-treatment average scores. Consistent with [Glewwe et al. \(2010\)](#), I assume the experiment was implemented with high fidelity, and consequently omit $\Delta \epsilon_t$ hereafter.

For the active-treatment period, $t = 1$, we have

$$\Delta y_1 = \beta_\theta \underbrace{\Delta \theta_1}_{=0} + \beta_h \Delta h_1 + \beta_y \underbrace{\Delta y_0}_{=0} = \beta_h \Delta h_1. \quad (6)$$

Equation (6) shows that the difference between treatment and control achievement in the active-treatment period can only come from the change in mean production time Δh_1 , as the fact that θ_{it} depends on lagged inputs and balance between the treatment and control groups implies that $\Delta \theta_0 = 0$, while, as shown in eq. (5), balance between the treatment and control groups implies that

$\Delta y_0 = 0$.

The mean difference in achievement between the treatment and control groups for the post-active-treatment period, $t = 2$, is

$$\begin{aligned}
\Delta y_2 &= \beta_\theta \Delta \theta_2 + \beta_h \Delta h_2 + \beta_y \Delta y_1 \\
&= \beta_\theta \delta_I \Delta I_1 + \beta_\theta \delta_h \Delta h_1 + \beta_h \Delta h_2 + \beta_y \Delta y_1 \\
&= \beta_\theta \delta_I \underbrace{[-\Delta l_1 - \Delta h_1]}_{=\Delta I_1, \text{ from eq. (4)}} + [\beta_\theta \delta_h + \beta_h \beta_y] \Delta h_1 + \beta_h \Delta h_2 \\
&= -\beta_\theta \delta_I \Delta l_1 + \left[\frac{\beta_\theta \delta_h - \beta_\theta \delta_I + \beta_h \beta_y}{\beta_h} \right] \underbrace{\Delta y_1}_{\beta_h \Delta h_1, \text{ from eq. (6)}} + \beta_h \Delta h_2, \tag{7}
\end{aligned}$$

which uses $\Delta \theta_1 = 0$ to go from the first to the second line and eq. (6) to go from the second to the third line. Equation (7) shows that, in general, only a locus of (δ_h, δ_I) will be identified, and, further, that many other variables appear in the same equation: the achievement production function parameters $(\beta_\theta, \beta_h, \beta_y)$ and, even after using eq. (6) to eliminate Δh_1 , the quantities $(\Delta y_1, \Delta y_2, \Delta l_1, \Delta h_2)$.¹⁸ The upper bound on the parameter spaces for δ_h and δ_I , $\bar{\delta}$, is also unknown. I estimate $(\beta_y, \Delta y_1, \Delta y_2, \Delta l_1)$ in Section 4.1. However, Δh_2 is unobserved and it cannot be eliminated, as was Δh_1 via eq. (6).

I now discuss how I arrive at values for $(\beta_\theta, \beta_h, \bar{\delta}, \Delta h_2)$. First note that β_θ can be normalized to 1 because it is attached to the latent variable θ_t . I show below that Δy_1 is significantly greater than zero (i.e., there was a positive effect of the intervention during the active-treatment period). In light of eq. (6) and the fact that $\beta_h \geq 0$, it is then reasonable to treat β_h as strictly positive, as the positive effect on achievement during the active-treatment period could only be rationalized by an increase in time allocated to production (i.e., $\Delta h_1 \geq 0$). Because the scale of β_h is not identified separately from $\bar{\delta}$, I fix $\beta_h = 1$ and $\bar{\delta} = 1$ hereafter.¹⁹ It is important to note that the specific values of β_h and $\bar{\delta}$ do not affect the main findings, such as whether bounds are informative or whether I can reject either the “pure OJT” or “pure LBD” specifications.²⁰

Assumption 1 summarizes the parameter values discussed thus far. Assumption 1(i) is maintained hereafter. Assumption 1(ii) corresponds to making no returns to scale assumption on the human capital production function. Section 3.4 explores how stronger assumptions about the returns to scale would tighten the identified set.

¹⁸Note that δ_θ does not appear in eq. (7). Intuitively, this parameter measures how differences in teacher human capital emanating from differences in inputs from two periods ago affect production today; the balanced experimental design means these differences are all zero, causing δ_θ to drop out.

¹⁹If the researcher viewed the model as a (log-linearized) approximation to a nonlinear model (see Appendix A for discussion), it would be natural to have $\beta_h = 1$ (i.e., the same as β_θ , which is consistent with the interpretation that a teacher’s output equals her share of human capital allocated to production), and also to not allow δ_h or δ_I to exceed 1 (i.e., $\bar{\delta} = 1$).

²⁰There are possible parameter values where the results would be affected by the choice of β_h relative to $\bar{\delta}$, but the estimated parameters are far from this region.

Assumption 1 (Baseline, no returns to scale assumption).

(i) $\beta_\theta = 1, \beta_h = 1$

(ii) $(\delta_h, \delta_I) \in [0, 1] \times [0, 1]$.

Finally, if there do not exist income effects, wherein the level of θ_t or prior compensation substantially affect how teachers allocate their time, then the increase in output-based pay for the treatment group has two useful implications: (i) time allocated to production will not decrease in the active-treatment period, i.e., $\Delta h_1 \geq 0$ and (ii) in the post-active-treatment period, even though there may be a mean difference in θ_2 between the treatment and control groups, h_2 will be similar in both the treatment and control groups. As discussed above, the data are consistent with $\Delta h_1 > 0$, which is consistent with the absence of such income effects. Arguably, this makes it reasonable to operate under the assumption that $\Delta h_2 = 0$, as the data do not reject the assumption upon which this was predicated. It is important to note that the assumption that $\Delta h_2 = 0$ does not rule out the intervention increasing a teacher's productivity by way of, e.g., preparing materials that could also be used in subsequent years. This increased productivity would be captured by a higher level of θ_2 , which could be generated by either OJT or LBD forces.

Using the values obtained thus far, eq. (7) becomes

$$\Delta y_2 = [\delta_h - \delta_I + \beta_y] \Delta y_1 - \delta_I \Delta l_1,$$

which we can rearrange to get the following expression for the LBD component δ_h as a function of the OJT component δ_I and remaining parameters:

$$\delta_h = \left[\frac{\Delta y_2}{\Delta y_1} - \beta_y \right] + \delta_I \left[1 + \frac{\Delta l_1}{\Delta y_1} \right], \tag{8}$$

i.e., eq. (8) characterizes the identified set, $d_h \times d_I$. We can find the identified set for either parameter by projecting the locus characterized by eq. (8) onto the relevant axis, e.g., d_h can be obtained by projecting $d_h \times d_I$ onto the δ_h -axis. It will be convenient to rewrite eq. (8) as

$$\delta_h = \pi_{\text{intercept}} + \pi_{\text{slope}} \delta_I, \tag{9}$$

where the intercept of the locus of permissible combinations of (δ_h, δ_I) is $\pi_{\text{intercept}} := \frac{\Delta y_2}{\Delta y_1} - \beta_y$ and the slope of the locus is $\pi_{\text{slope}} := 1 + \frac{\Delta l_1}{\Delta y_1}$.

Note that, even with the parameter values obtained so far, the relationship between (δ_h, δ_I) still depends on $(\beta_y, \Delta y_1, \Delta y_2, \Delta l_1)$. Section 4.1 discusses estimation of these parameters. Briefly, all but β_y are estimated using the experimental variation. A statistical technique (treatment-year fixed effects in the production function, coupled with the statistical assumption that the experiment was appropriately balanced) yields consistent estimates of β_y when using just control group data.²¹

²¹This is because, for teacher-student pair ij in group g (i.e., treatment or control), we have $E_{ij \in g} [y_{ijt}] =$

3.3 Comparative Statics

We can understand the restrictions afforded by eq. (8) by considering some cases, where for the sake of illustration assume $\Delta y_1 > 0$. We start by maintaining Assumption 1. First suppose that $\Delta l_1 = 0$ and suppose that $\frac{\Delta y_2}{\Delta y_1} - \beta_y = \pi_{\text{iccept}} = 0$; this corresponds to case (a) in Table 2. Intuitively, the increase in test score coming via the positive effect of Δh_1 on Δy_1 entirely accounts for Δy_2 . The first line of eq. (7) then implies that $\Delta \theta_2 = 0$, i.e., teacher human capital post-active-treatment is on average the same in the treatment and control groups. Further, $\Delta l_1 = 0$ (i.e., $\pi_{\text{slope}} = 1$) implies that $\Delta I_1 = -\Delta h_1$, meaning the only way to satisfy eq. (8) is for $\delta_I = \delta_h$. That is, any increase in δ_I can satisfy the condition by a concomitant increase in δ_h . Without further information on either of these parameters, we cannot shrink the identified set. This scenario corresponds to the dashed, 45-degree, line in the left panel of Figure 2. Projecting the identified set onto each axis, we can see that the marginal identified set for either δ_I or δ_h (depicted by the dashed lines just outside that parameter's axis) has not shrunk at all, because any feasible value of, e.g., δ_I , can be rationalized by the same value for the coefficient on the other input to teacher human capital (in this example, δ_h). That is, the bounds in this case are uninformative.

Next consider case (b), which differs from case (a) in that $\frac{\Delta y_2}{\Delta y_1} - \beta_y < 0$, i.e., teacher human capital is lower in the treatment group, post active treatment ($\Delta \theta_2 < 0$). Because we still have $\Delta l_1 = -\Delta h_1$, the fact that $\Delta \theta_2 < 0$ means that we can rule out very low values of δ_I —the OJT component—and very high values of δ_h —the LBD component. Intuitively, if the net effect of increasing time spent on production on teacher human capital is negative, knowing that $\Delta I_1 = -\Delta h_1$ implies that the OJT parameter must be larger than the LBD one. At the same time, the slope of eq. (8) is unaffected because the one-to-one tradeoff between different time uses in the budget constraint (due to there being no average change in leisure time) implies a one-to-one tradeoff between δ_I and δ_h . This case is depicted by the dotted lines in the left panel of Figure 2.

In cases (a) and (b), there was no average difference in leisure time between the treatment and control groups during the active-treatment period (i.e., $\Delta l_1 = 0$, or $\pi_{\text{slope}} = 1$). Consider now case (c), where we start from case (b) but now assume that $\Delta l_1 < 0$ (here, $\pi_{\text{slope}} < 1$). Here, we know that $\Delta I_1 > -\Delta h_1$, i.e., the absolute difference in investment is smaller than the absolute difference in production time. As shown in the dash-dotted lines in the left panel, this rotates the locus eq. (8) downward from the intercept (which was already negative, as the starting point was case (b)), increasing the lower bound on δ_I and decreasing the upper bound on δ_h . Intuitively, all else equal, a smaller change in investment must be coupled with a relatively bigger technological effect of investment (δ_I) to rationalize the same data.

The cases discussed above are not exhaustive. For example, the signs of $\frac{\Delta y_2}{\Delta y_1} - \beta_y$ or Δl_1 could be opposite to those considered in cases (b) or (c), in which case the sharp bounds would be different.

$\overline{E_{ij \in g} [\beta_\theta \theta_{ijt} + \beta_h h_{it}] + \beta_y E_{ij \in g} [y_{ijt-1}] + \underbrace{E_{ij \in g} [\epsilon_{ijt+1}]}_{=0}}$ and $E_{ij \in g} [\beta_\theta \theta_{ijt} + \beta_h h_{it}]$ can be measured by the coefficient on a group-period indicator variable.

An example of this is case (d), which is depicted in the right panel of Figure 2, which illustrates how strengthening returns-to-scale-type assumptions yields tighter sharp bounds for the identified set. The next section explores this.

3.4 Bounds with Increasing Assumption Strength

A researcher might further find it natural to restrict the returns to scale in the human capital production function, by assuming they are nonincreasing. This is Assumption 2 below. Even stronger, the researcher might believe it reasonable to assume constant returns to scale (Assumption 3). This exploration of how assumptions about $\delta_h + \delta_I$ affect the identified sets is in the spirit of the ““worst-case”” approach of Horowitz and Manski (2000), which examines the sensitivity of findings to stronger sets of assumptions, some of which are made in the literature. It does not constitute an endorsement of making these stronger assumptions.

Assumption 2 (Nonincreasing returns to scale (NIRS)). *Assumption 1(i) and $\delta_h + \delta_I \leq 1$.*

Assumption 3 (Constant returns to scale (CRS)). *Assumption 1(i) and $\delta_h + \delta_I = 1$.*

The right panel of Figure 2 shows the additional information embedded in assumptions about the returns to scale for (δ_h, δ_I) , starting with Assumption 1 (case (d)); these cases are summarized in the lower part of Table 2. Case (d2) further imposes the restriction that $\delta_h + \delta_I \leq 1$ (Assumption 2), which means permissible combinations of (δ_h, δ_I) lie in the south/west right triangle.²² By looking at the dotted lines in the right panel of Figure 2, corresponding to this case, we can see that the marginal identified sets are tighter than those in case (d). Intuitively, the tighter upper bound for δ_h , combined with the nonincreasing returns to scale assumed in case (d2), yields a tighter upper bound for δ_I . Case (d3) further imposes the restriction that $\delta_h + \delta_I = 1$ (Assumption 3), which affords point identification (interior solid point in the right panel of Figure 2).

3.5 Expressions for Marginal Identified Sets

This section ends by characterizing the marginal identified sets for δ_h and δ_I , which are used in estimating the parameters’ bounds and confidence sets. Under Assumption 1, i.e., when making no additional assumptions about returns to scale, we can derive the identified set for δ_h , d_h , by varying δ_I over its domain, resulting in

$$\delta_h \in [\max\{\pi_{\text{icpt}}, 0\}, \min\{\pi_{\text{icpt}} + \pi_{\text{slope}}, 1\}]. \quad (10)$$

²²Note that this restriction in itself would not yield any information; without data (e.g., information about π_{icpt}), the bounds on either parameter, obtained from projecting the south/west right triangle onto either axis, would be uninformative.

Table 2: Bounds: Comparative Statics and Effects of Returns-to-Scale Assumptions

Case	$(\delta_h, \delta_I) \in$	Δl_1	$\frac{\Delta y_2}{\Delta y_1} - \beta_y$	Notes on bounds
(a)	$\delta_h \leq 1, \delta_I \leq 1$	0	0	Uninformative
(b)	$\delta_h \leq 1, \delta_I \leq 1$	0	< 0	Informative
(c)	$\delta_h \leq 1, \delta_I \leq 1$	< 0	< 0	Tighter than (b)
(d)	$\delta_h \leq 1, \delta_I \leq 1$	0	> 0	Informative
(d2)	$\delta_h + \delta_I \leq 1$	0	> 0	Tighter than (d)
(d3)	$\delta_h + \delta_I = 1$	0	> 0	Point identification

Note: $\Delta y_1 > 0$. Cases (a)-(c) correspond to comparative statics, and cases (d)-(d3) correspond to different returns-to-scale assumptions.

Solving eq. (9) instead for δ_I , we can analogously obtain d_I by varying δ_h over its domain:

$$\delta_I = \frac{\delta_h - \pi_{\text{icept}}}{\pi_{\text{slope}}} \Rightarrow \delta_I \in \left[\max \left\{ \frac{-\pi_{\text{icept}}}{\pi_{\text{slope}}}, 0 \right\}, \min \left\{ \frac{1 - \pi_{\text{icept}}}{\pi_{\text{slope}}}, 1 \right\} \right]. \quad (11)$$

Invoking non-increasing returns to scale (Assumption 2) does not affect the lower bound of d_h , but does tighten its upper bound:

$$\begin{aligned} \delta_h &= \pi_{\text{icept}} + \pi_{\text{slope}} \delta_I \Rightarrow \delta_h \leq \pi_{\text{icept}} + \pi_{\text{slope}} [1 - \delta_h] \Rightarrow \\ \delta_h &\in \left[\max \{ \pi_{\text{icept}}, 0 \}, \min \left\{ \frac{\pi_{\text{icept}} + \pi_{\text{slope}}}{1 + \pi_{\text{slope}}}, 1 \right\} \right]. \end{aligned} \quad (12)$$

We can analogously tighten the upper bound on d_I by applying Assumption 2:

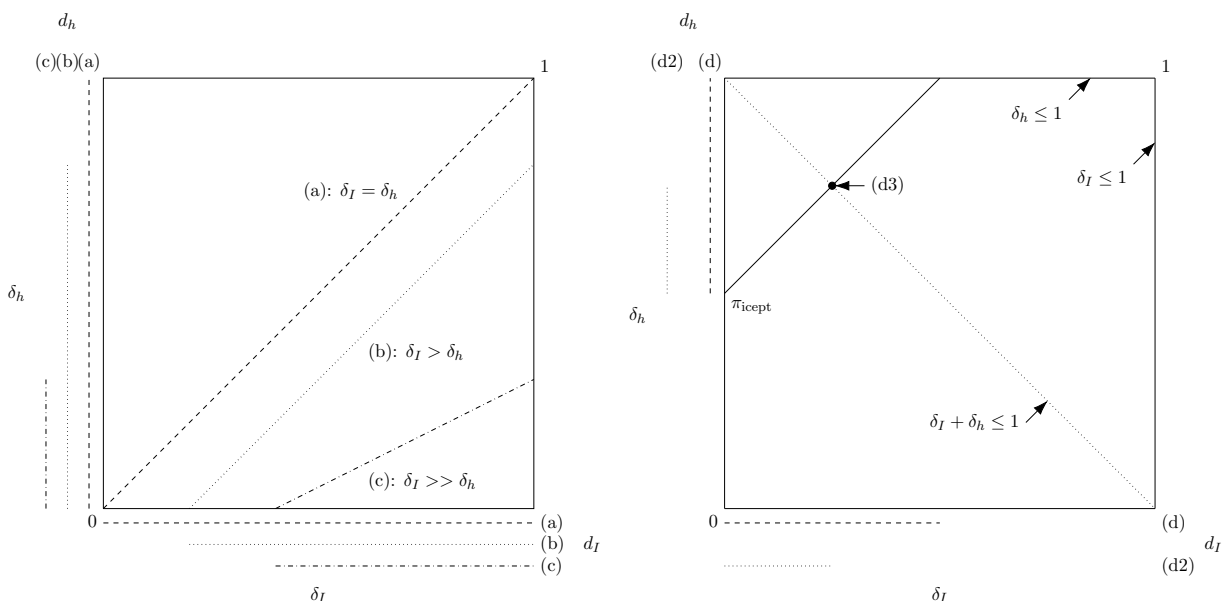
$$\begin{aligned} \delta_h &= \pi_{\text{icept}} + \pi_{\text{slope}} \delta_I \Rightarrow [1 - \delta_I] \geq \pi_{\text{icept}} + \pi_{\text{slope}} \delta_I \Rightarrow [1 + \pi_{\text{slope}}] \delta_I \leq 1 - \pi_{\text{icept}} \\ \delta_I &\in \left[\max \left\{ \frac{-\pi_{\text{icept}}}{\pi_{\text{slope}}}, 0 \right\}, \min \left\{ \frac{1 - \pi_{\text{icept}}}{1 + \pi_{\text{slope}}}, 1 \right\} \right]. \end{aligned} \quad (13)$$

Finally, imposing Assumption 3, we achieve point identification for both parameters:

$$\delta_h = \pi_{\text{icept}} + \pi_{\text{slope}} [1 - \delta_h] \Rightarrow \delta_h = \frac{\pi_{\text{icept}} + \pi_{\text{slope}}}{1 + \pi_{\text{slope}}} \quad (14)$$

$$[1 - \delta_I] = \pi_{\text{icept}} + \pi_{\text{slope}} \delta_I \Rightarrow \delta_I = \frac{1 - \pi_{\text{icept}}}{1 + \pi_{\text{slope}}}. \quad (15)$$

Figure 2: Illustration of Example Identified Set Cases



Note: Marginal identified sets for parameters are indicated by the lines outside their respective axes. Cases depicted in the figure are summarized in Table 2. The left panel corresponds to comparative statics, cases (a)-(c). The right panel corresponds to different returns-to-scale assumptions, cases (d)-(d3).

4 Empirical Results

4.1 Estimation of Parameters Determining Identified Set for (δ_h, δ_I)

From eq. (8), the identified set for (δ_h, δ_I) depends on $(\Delta y_1, \Delta y_2, \beta_y, \Delta l_1)$. This section discusses how these parameters are estimated, and Section 4.2 shows how the resulting marginal identified sets for (δ_h, δ_I) are estimated.

I pool the two active-treatment years to map the empirical application, which spans four years of data, to the three-period structure of the model. The pre-treatment year (1997) corresponds to model period $t = 0$. Both active-treatment years (1998 and 1999) are pooled into one period, corresponding to $t = 1$ in the model, and the post-active-treatment year (2000) corresponds to $t = 2$. While the estimates presented here are based on the pooled data, I have examined their sensitivity to using unpooled data when feasible (i.e., for β_y , which can be estimated without the experimental variation), and the results were essentially unaffected.

Table 3 presents the achievement ATE results for the active-treatment period (active * treated) and the post-active treatment period (post * treated). I also include the treatment year as a regressor (active-treatment, post-active-treatment) to control for secular trends. The results for the treatment group in the active-treatment period indicate that the average treatment effect pooled over both active-treatment years, Δy_1 , is positive and significantly different than zero (0.089). The next row indicates that student achievement in the post-active-treatment year, Δy_2 , remained significantly higher (0.098) in the treatment group. Through the lens of the model, this positive

Table 3: Estimates of ATEs on Achievement

	<i>Dependent variable:</i>
	Test score y_{ijt}
active-treatment ($t = 1$)	0.001 (0.019)
post-active-treatment ($t = 2$)	-0.001 (0.017)
active * treated (Δy_1)	0.089*** (0.013)
post * treated (Δy_2)	0.098*** (0.025)
Constant	0.007 (0.008)
Observations	26,537

Note: *p<0.1; **p<0.05; ***p<0.01. The t corresponds to a pooled period.

ATE could be due to persistence of higher achievement from the active-treatment period and (potentially) higher teacher human capital for the treatment group in the post-active-treatment period. Establishing the relative importance of these effects is the goal of the next section.

Table 4 presents an estimate of the persistence component β_y using control group data, where the lagged score is that from the previous pooled period. I estimate the persistence component to be 0.544, and statistically greater than zero. The estimate of persistence is not driven by the use of pooled data: when instead using unpooled data the estimate is 0.618. Even more to the point, this estimate is in the range of those in [Andrabi et al. \(2011\)](#), which estimated the persistence of a variety of cognitive skills using a dynamic panel data model applied to Pakistani schoolchildren. Allowing for both measurement error and unobserved student heterogeneity (in contrast with the specifications researchers have typically used to estimate achievement production functions), they estimate the (annual) persistence of cognitive skills to range from 0.2 to 0.55, across a variety of subjects. Indeed, they argue that the upper part of their range is possibly too high. This matters because a smaller value of β_y will further tighten the estimated bounds, in light of the achievement ATE estimates, which already yield a positive, significant, estimate of π_{icpt} . Intuitively, the post-active treatment achievement ATE is higher than can be explained by the active-treatment achievement ATE persisting into post-active treatment, pointing to the presence of an LBD component; lowering the value of β_y would only make this effect more prominent.

I now discuss how I obtain a value for the effect of the intervention on leisure during the active-

Table 4: Estimate of the Persistence Component

	<i>Dependent variable:</i>
	Test score y_{ijt}
lagged score $y_{j,t-1}$ (β_y)	0.544*** (0.010)
Constant	-0.047*** (0.008)
Observations	5,007

Note: *p<0.1; **p<0.05; ***p<0.01. These estimates are obtained from a regression of current test score on the lagged test score, run on the pooled data for pooled periods 1 and 2, for the control group. Specifically, the (period) t subindex on the test score y_{ijt} refers to a pooled period, meaning student j 's lagged test score, $y_{j,t-1}$, is the one from the previous pooled period.

treatment period, Δl_1 . The literature estimating OJT models of human capital accumulation typically treats total work time (which I define to be the sum of time spent in production and investment) as observed (see, e.g., [Brown, 1976](#); [Heckman, 1976](#)). I follow this literature and use measures of total work time, which are sufficient for the current paper because the change in total work time ($\Delta I_t + \Delta h_t$) is the complement of the change in leisure (Δl_t).

The [Glewwe et al. \(2010\)](#) data contain two measures of teachers' total work time in a period: (i) the fraction of teachers in attendance at the school during a site visit by the research team that period and (ii) the fraction of teachers present in their classroom during the site visit.²³ [Glewwe et al. \(2010\)](#) finds no evidence that teachers in the treatment group on average altered their school attendance or classroom presence (see Table 5, Panels A and B, columns (2) and (3)). While either would be a reasonable measure of total work time, it seems more appropriate to use the share of teachers present at the school, as teachers could make on-the-job investments outside the classroom (e.g., to avoid being interrupted by students). Perhaps unsurprisingly, Table 5 shows the treatment effect on the share of teachers in attendance, -0.017, is not significantly different from zero.²⁴ This means the point estimate for the effect on leisure, which is the negative of the effect on total work time, is positive, at 0.017 (and, naturally, also insignificantly different from zero).²⁵

4.2 Estimates of Bounds

The parameters characterizing the identified set for (δ_h, δ_I) , π_{icept} and π_{slope} , can be estimated using $\widehat{\pi}_{\text{icept}} := \frac{\Delta y_2}{\Delta y_1} - \beta_y$ and $\widehat{\pi}_{\text{slope}} := 1 + \frac{\Delta l_1}{\Delta y_1}$, which are computed using plug-in estimators. To simulate the joint distribution of $(\widehat{\pi}_{\text{icept}}, \widehat{\pi}_{\text{slope}})$, taking into account the variability of the inputs

²³Teacher-level data on either measure were not available in the data, so I use school-period-level averages.

²⁴I also include the treatment year as a regressor here to control for secular trends.

²⁵Given that the point estimate of the effect on leisure is small and insignificant, the estimated bounds and confidence sets are similar if I instead fix $\Delta l_1 = 0$.

Table 5: Estimate of ATE on Teacher Attendance

	<i>Dependent variable:</i>
	Present at school in period t
active-treatment ($t = 1$)	0.012 (0.025)
post-active-treatment ($t = 2$)	0.052** (0.025)
active * treated (Δl_1)	-0.017 (0.029)
post * treated	-0.011 (0.041)
Constant	0.833*** (0.014)
Observations	202

Note: *p<0.1; **p<0.05; ***p<0.01. Dependent variable is school-level average share of site visits during which teachers were in attendance. The t corresponds to a pooled period.

to the plug-in estimators, I bootstrap the joint distribution of $(\widehat{\Delta y}_1, \widehat{\Delta y}_2, \widehat{\beta}_y, \widehat{\Delta l}_1)$, where in each bootstrap replication, the individual elements are estimated as described just above.²⁶

I provide an overview here of how I estimate the bounds and confidence sets for the parameters; see Algorithm 1 in Appendix B for more detail. I first construct the identified set, $\hat{d}_h \times \hat{d}_I$, and then project this set onto each marginal dimension, and then construct confidence sets to contain each parameter (not the identified set for each parameter) at the pre-specified significance level (I use a confidence level of 95%). Estimated bounds are the average across the identified sets. By projecting the identified set onto the marginal dimensions, I avoid the problem of overly conservative confidence sets described by Kaido et al. (2019). That being said, the estimated confidence sets are virtually identical when containing parameter or identified set with particular probability, because each confidence set contains at least one end point of the parameter space.

Results Figure 3 illustrates the estimated bounds and 95% confidence sets for δ_h and δ_I under the different assumptions about $\delta_h + \delta_I$. The table below presents the corresponding estimates and also reports bound widths. Starting with the LBD parameter in panel (a), we can see that when we do not impose a returns-to-scale-type assumption (“none”), the estimated upper and lower bounds (thick, black, line) for δ_h are informative, and the lower bound is greater than zero at the 95%

²⁶I bootstrap using 100,000 replications of the data, stratified by treatment status.

confidence level (thin, grey, line), leading us to reject the pure OJT specification (in which $\delta_h = 0$). Similarly, in panel (b) we can see that the estimated upper and lower bounds for δ_I are informative, and that the 95% confidence set for δ_I does not contain the upper bound of 1 when no assumption about returns to scale is made (“none”).


Imposing non-increasing returns to scale (“NIRS” on the horizontal axis in each panel) tightens the estimated upper bound for δ_h , although it does not affect the 95% confidence set, because the estimated lower bound, which governs the lower bound of the confidence set, is unchanged. However, imposing NIRS does tighten the estimated upper bound and upper bound on the confidence sets for δ_I . Intuitively, high values of both δ_h and δ_I are no longer mutually feasible under NIRS, lowering the upper bounds for both parameters. Consequently, we can see in the accompanying table that the width of the estimated bounds falls by about one half for both parameters when imposing NIRS (e.g., from 0.433 to 0.233 for δ_h), and the same is true of the width of the 95% confidence set for δ_I . Finally, imposing constant returns to scale (“CRS” on the horizontal axis in each panel) yields point identification for both parameters, corresponding to the smallest width confidence sets. The estimated bounds are of course zero width in this case.

Overall, under all the assumptions about returns to scale, the achievement ATE in the post-active-treatment period is larger than would be accounted for by the positive ATE in the active-treatment period (caused by an increase in h_1 , an input to contemporaneous student achievement) and persistence of this increased student achievement. That is, the OJT component, operating through $\Delta I_1 \leq 0$, is dominated by the LBD component, operating through $\Delta h_1 > 0$. We can reject the “pure OJT” specification in which $\delta_h = 0$ across all returns-to-scale assumptions at the 95% confidence level. However, we cannot reject that $\delta_I = 0$ in any of the returns-to-scale assumptions at the 95% confidence level; this means we cannot reject the “pure LBD” specification. Further, if one were willing to assume CRS, then one could infer that δ_h was greater than δ_I , as the 95% confidence sets under CRS do not overlap; the same cannot be said under either of the weaker assumptions.

Discussion The finding that LBD, at least in part, explains growth in teacher quality means that, on average, teachers improve by teaching their students. Of course, teachers might also improve by making OJT investments, as the confidence set for the OJT component includes strictly positive values. While this might seem quite intuitive, economists to date had not been able to identify the force within human capital theory behind why teacher quality should increase with experience.

Coming back to the specific context underlying this paper’s estimates, [Glewwe et al. \(2010\)](#) assessed that the intervention may have increased “teaching to the test”, not more general student knowledge. They surmised this because, in contrast to the results for the incentivized exam (used in the current paper), student achievement for a non-incentivized exam that covered similar material did not significantly increase in the treatment group during active treatment. That being said, it is important to note that the associated change in teacher human capital was not limited to

Figure 3: Estimated Bounds and Confidence Sets



Notes: The left panel depicts estimated bounds (—) and 95% confidence sets (—) for δ_h , under Assumptions 1, 2, and 3 (denoted on the bottom axis via “none”, “NIRS”, and “CRS”, respectively). The right panel depicts the analogous results for δ_I .

Parameter	Ass. about $\delta_h + \delta_I$	Estimated bounds			95% confidence set		
		min	max	width	min	max	width
δ_h	1: none	0.565	0.998	0.433	0.067	1	0.933
	2: NIRS	0.565	0.798	0.233	0.067	1	0.933
	3: CRS	0.798	0.798	0	0.562	1	0.438
δ_I	1: none	0.003	0.378	0.375	0	0.853	0.853
	2: NIRS	0.003	0.202	0.199	0	0.438	0.438
	3: CRS	0.202	0.202	0	0	0.438	0.438

Notes: Assumption 1 corresponds to no assumption about the returns to scale for $\delta_h + \delta_I$, and Assumptions 2 and 3 respectively correspond to non-increasing returns to scale and constant returns to scale.

a one-time change contained only to the active treatment period, as it did show up in the post-active-treatment period, meaning that the current paper does identify a force generating a form of teacher human capital. However, just as the above assessment of [Glewwe et al. \(2010\)](#) warrants a modicum of caution when interpreting the findings of that paper, it does so for those of this paper too.

While I use the incentivized test to most clearly demonstrate the current methodology for separately identifying OJT and LBD, we can also use the developed framework to help think about the importance of LBD versus OJT in generating (desirable) teacher human capital, even in light of the above caveat. Somewhat loosely, imagine there were two forms of OJT investment, one for “incentivized teacher human capital”, and another for more general (or unincentivized) teacher human capital, and analogously that a teacher can allocate her time to production of either form of human capital in her students (which could manifest in different degrees in student achievement). Through the lens of the model it still must be the case that $\Delta h_1 > 0$, i.e., teacher human capital allocated to production increased in the treatment group during the active-treatment period. While all we know is that this increase in production time may have been directed to incentivized (and not more general) student knowledge, the fact that total work hours did not change means the OJT investment likely decreased for either type of human capital. Then, the fact that the post-active treatment effect was still greater than that which would be explained by the persistence of the active-treatment achievement effect is still consistent with OJT not playing the paramount role, insofar as achievement depended on both test-specific and general inputs.

5 Conclusion

I develop a framework nesting the OJT and LBD forces of human capital accumulation, and derive theoretical bounds for OJT and LBD components. The developed bounds are sharp, and yield novel information about the presence and relative importance of the forces generating human capital. The derived bounds only require information about ATEs and the persistence of student achievement. The estimated bounds are informative, and under even the weakest assumptions about the returns to scale allow one to reject the “pure OJT” model. That is, the data are consistent with the presence of an LBD component to teacher human capital accumulation. This suggests the dynamic multitasking problem inherent to the “pure OJT” model is at least tempered by the presence of an LBD component to human capital accumulation.

Overall, this paper constitutes an important step towards designing effective educational policy that targets teachers and also shows how a partial identification approach can exploit existing data, designed for another purpose, to answer an important policy relevant question. This paper’s framework could also be applied to other contexts, in education and otherwise. It is common to collect follow-up measures to gauge the longer-run effects of educational interventions. The framework developed in this paper provides a way to interpret such follow-up data on interventions

targeting teachers, through the lens of classic conceptual frameworks for human capital: longer-run effects stem from the persistence of student knowledge and changes in teacher human capital. Future teacher incentive pay experiments that collected follow-up data would be able to apply this paper's methodology to identify the forces underlying teacher human capital growth. The strategy developed in this paper could also be adapted to other applications in which there were outcome-based incentives and a followup measure of output, to quantify the importance of different channels underlying human capital development.²⁷

In light of the well known identification difficulties, it may be surprising that we can learn something new about human capital accumulation, even under the transparent and relatively simple approach taken here. While this paper focuses on the straightforward bounds attainable under the specifications most commonly used in the literatures on teacher quality and human capital, it is important to note that other assumptions and different estimation methods would also yield estimates, but these would likely be less transparent due to being cast in a less tractable framework. A very promising, complementary, tack would be the structural econometric approach, which would require different (some stronger) assumptions but could then also answer other important questions about the importance of OJT and LBD forces in teachers' human capital accumulation. Such an approach would also be well-suited to rationalize the observed patterns in the data, and would yield other benefits, such as allowing for heterogeneity in teacher human capital accumulation trajectories and, thus, heterogeneity in the growth of teacher quality. It would also permit simulation of behavior and outcomes under counterfactual incentive schemes. This is left for future research.

²⁷In this vein, note that I maintain the classic OJT assumption that training investments are unobserved, which increases its applicability to other contexts.

APPENDIX

A Relationship to a Log-Linear Specification

This section illustrates one way in which the linear technologies (1)-(2) relate to nonlinear specifications. The illustration considers a representative teacher, teaching a representative student, in each of the control and treatment groups; therefore I suppress the teacher and student subscripts in this section. I maintain the assumption of balance of the experimental design.

Consider the following production function for teacher human capital, denoted here as κ_{it} :

$$\kappa_t = \kappa_{t-1}^{\gamma_\kappa} [\kappa_{t-1} \iota_{t-1}]^{\gamma_\iota} [\kappa_{t-1} \zeta_{t-1}]^{\gamma_\zeta}, \quad (16)$$

where κ_{t-1} is the teacher's human capital last period (which may depreciate), ι_{t-1} is the share of the teacher's human capital last period spent on OJT investment, and ζ_{t-1} is the share of the teacher's human capital last period spent on production. The parameters of interest, respectively representing the OJT and LBD components of human capital accumulation in (16), are $(\gamma_\iota, \gamma_\zeta) \in [0, 1]^2$.²⁸

Using $\tilde{\cdot}$ to denote the logarithm of a variable, we can write the log-linearized version of the human capital production function, eq. (16):

$$\tilde{\kappa}_t = [\gamma_\kappa + \gamma_\iota + \gamma_\zeta] \tilde{\kappa}_{t-1} + \gamma_\iota \tilde{\iota}_{t-1} + \gamma_\zeta \tilde{\zeta}_{t-1}. \quad (17)$$

Let w_t measure the cognitive skill of the student in period t , which is produced according to²⁹

$$w_t = \kappa_t^{\lambda_\kappa} \zeta_t^{\lambda_\zeta} w_{t-1}^{\lambda_w}, \quad (18)$$

which, in logs, is

$$\tilde{w}_t = \lambda_\kappa \tilde{\kappa}_t + \lambda_\zeta \tilde{\zeta}_t + \lambda_w \tilde{w}_{t-1}. \quad (19)$$

This equation is a log-linearized value added specification for cognitive achievement, where the value added to log achievement is $\lambda_\kappa \tilde{\kappa}_t + \lambda_\zeta \tilde{\zeta}_t$.

As before, the bounds on the parameters of interest will depend on the ATEs for achievement and leisure. With a representative teacher and student, we have $\Delta \tilde{z}_t = \tilde{z}_t^T - \tilde{z}_t^C$ for $z = \kappa, \iota, \zeta, w$.

For the pre-treatment period, $t = 0$, the mean difference in achievement between the treatment and control groups is

$$\Delta \tilde{w}_0 = \lambda_\kappa \underbrace{\Delta \tilde{\kappa}_0}_{=0} + \lambda_\zeta \underbrace{\Delta \tilde{\zeta}_0}_{=0} + \lambda_w \underbrace{\Delta \tilde{w}_{-1}}_{=0} = 0, \quad (20)$$

²⁸Note that the unit interval is conservative, as it allows for parameter values that could yield explosive growth and would thus likely be ruled out a priori if estimating this model.

²⁹I have not included an error because there is a representative student in each group.

i.e., a balanced experimental design implies there will be no average difference in the pre-treatment average scores, as was also the case in Section 3.2.

For the active-treatment period, $t = 1$, we have

$$\Delta\tilde{w}_1 = \lambda_\kappa \underbrace{\Delta\tilde{\kappa}_1}_{=0} + \lambda_\zeta \Delta\tilde{\zeta}_1 + \lambda_w \underbrace{\Delta\tilde{w}_0}_{=0} = \lambda_\zeta \Delta\tilde{\zeta}_1. \quad (21)$$

Similar to eq. (6), eq. (21) shows that the difference between treatment and control achievement in the active-treatment period can only come from the change in mean working time, $\Delta\tilde{\zeta}_1$.

The mean difference in achievement between the treatment and control groups for the post-active-treatment period, $t = 2$, is

$$\begin{aligned} \Delta\tilde{w}_2 &= \lambda_\kappa \Delta\tilde{\kappa}_2 + \lambda_\zeta \Delta\tilde{\zeta}_2 + \lambda_w \Delta\tilde{w}_1 \\ &= \lambda_\kappa \gamma_l \Delta\tilde{l}_1 + \lambda_\kappa \gamma_\zeta \Delta\tilde{\zeta}_1 + \lambda_\zeta \Delta\tilde{\zeta}_2 + \lambda_w \Delta\tilde{w}_1, \end{aligned} \quad (22)$$

which uses $\Delta\theta_1 = 0$ to go from the first to the second line. Noting that the intervention had no effect on leisure, we have

$$\Delta\zeta_t + \Delta\iota_t = 0, \quad (23)$$

i.e., the effect on OJT investment is opposite that on production shares. Substituting using eqs. (21) and (23) and maintaining the assumption that post-active treatment production shares will not be different between the control and treatment groups (i.e., $\Delta\tilde{\zeta}_2 = 0$), eq. (22) becomes

$$\Delta\tilde{w}_2 = \left[\frac{\lambda_\kappa \gamma_\zeta - \lambda_\kappa \gamma_l + \lambda_\zeta \lambda_w}{\lambda_\zeta} \right] \underbrace{\Delta\tilde{w}_1}_{\lambda_\zeta \Delta\tilde{\zeta}_1, \text{ from eq. (21)}} \quad (24)$$

The last step is to obtain values for the relevant quantities in (24), λ_κ , λ_ζ , λ_w , $\Delta\tilde{w}_1$, $\Delta\tilde{w}_2$. Analogous to Section 4, I set $\lambda_\kappa = 1$ and $\lambda_\zeta = 1$; both of these are normalizations of sorts,³⁰ and these parameters having the same value is consistent with a teacher's value added being the share of her human capital allocated to production. Next consider the remaining parameters, λ_w , $\Delta\tilde{w}_1$, $\Delta\tilde{w}_2$. It is well known that test scores measuring, e.g., cognitive skill, have no inherent scale, meaning any monotonic (increasing) transformations (e.g., logarithms) are also valid measures (see, e.g., Cunha and Heckman, 2008). Given the earlier argument that qualitative differences in inputs drive the analysis, and the known sensitivity of achievement tests to monotonic transformations (Bond and Lang, 2013), it is reasonable to use the estimator of β_y as an approximation for λ_w . In a similar way, we can use Δy_t , which corresponds to the estimated difference in value added between the treatment and control groups, to measure $\Delta\tilde{w}_t$, the difference in achievement for the students re-

³⁰ λ_κ fixes the scale of pre-treatment teacher human capital, which is unobserved and yields no direct testable implications. Setting $\lambda_\zeta = 1$ may not strictly be a pure normalization, as a null ATE on achievement in the active-treatment period could stem from $\lambda_\zeta = 0$ and/or $\Delta\tilde{\zeta}_1 = 0$. However, both of these parameters must be nonzero to match the positive ATE in the data; given this, then, setting $\lambda_\zeta = 1$ is innocuous.

spectively representing the treatment and control groups. Putting all of this together, the bounds obtained for (δ_I, δ_h) would also apply to $(\gamma_\iota, \gamma_\zeta)$.

B Estimation of Confidence Sets

This appendix describes the algorithm used to estimate the confidence sets for δ_h and δ_I .

Algorithm 1 Bootstrap estimation of confidence sets

for $s = 1 \dots nSamp$ **do**

 Sample observations from treatment and control groups (stratified by treatment group)

 Estimate $(\hat{\beta}_y^s, \hat{\Delta}l_1^s, \hat{\Delta}y_1^s, \hat{\Delta}y_2^s)$

 Use simulated values to create locus defined by eq. (8)

 Project locus onto marginals, obtaining random intervals \hat{d}_h^s, \hat{d}_I^s

end for

for $k = h, I$ **do**

for $d_k^{try} \subseteq D_k$ **do**

 retain d_k^{try} iff $\frac{1}{nSamp} \sum_{s=1}^{nSamp} \mathbf{1}\{\hat{d}_k^s \subseteq d_k^{try}\} \geq 1 - \alpha$

end for

end for

$(1 - \alpha)\%$ confidence set for δ_k is $\arg \min_{d_k^{try} \text{ retained}} \{\max\{d_k^{try}\} - \min\{d_k^{try}\}\}$

References

- Abowd, J. M., F. Kramarz and D. N. Margolis, “High Wage Workers and High Wage Firms,” *Econometrica*, 67(2):251–333, 1999.
- Agostinelli, F. and M. Wiswall, “Estimating the Technology of Children’s Skill Formation,” Tech. rep., National Bureau of Economic Research, 2020.
- Andrabi, T., J. Das, A. I. Khwaja and T. Zajonc, “Do Value-Added Estimates Add Value? Accounting for Learning Dynamics,” *American Economic Journal: Applied Economics*, 3(3):29–54, 2011.
- Becker, G. S., *Human Capital Theory*, Columbia University Press, New York, 1964.
- Behrman, J. R., M. M. Tincani, P. E. Todd and K. I. Wolpin, “Teacher Quality in Public and Private Schools Under a Voucher System: The Case of Chile,” *Journal of Labor Economics*, 34(2):319–362, 2016.
- Ben-Porath, Y., “The Production of Human Capital and the Life Cycle of Earnings,” *Journal of Political Economy*, 75(4):352–365, 1967.
- Biasi, B., “The Labor Market for Teachers under Different Pay Schemes,” *American Economic Journal: Economic Policy*, 13(3):63–102, 2021.
- Biasi, B., C. Fu and J. Stromme, “Equilibrium in the Market for Public School Teachers: District Wage Strategies and Teacher Comparative Advantage,” Tech. rep., National Bureau of Economic Research, 2021.
- Blandin, A., “Learning by Doing and Ben-Porath: Life-cycle Predictions and Policy Implications,” *Journal of Economic Dynamics and Control*, 90:220–235, 2018.
- Blinder, A. S. and Y. Weiss, “Human Capital and Labor Supply: A Synthesis,” *Journal of Political Economy*, 84(3):449–472, 1976.
- Bobba, M., T. Ederer, G. Leon-Ciliotta, C. Neilson and M. G. Nieddu, “Teacher Compensation and Structural Inequality: Evidence from Centralized Teacher School Choice in Perú,” Tech. rep., National Bureau of Economic Research, 2021.
- Bond, T. N. and K. Lang, “The Evolution of the Black-White Test Score Gap in Grades K–3: The Fragility of Results,” *Review of Economics and Statistics*, 95(5):1468–1479, 2013.
- Bono, E. D., J. Kinsler and R. Pavan, “A Note on the Importance of Normalizations in Dynamic Latent Factor Models of Skill Formation,” 2020.
- Brown, C., “A Model of Optimal Human-Capital Accumulation and the Wages of Young High School Graduates,” *Journal of Political Economy*, 84(2):299–316, 1976.
- Chetty, R., J. N. Friedman and J. E. Rockoff, “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” *American Economic Review*, 104(9):2593–2632, 2014.
- Cook, J. B. and R. K. Mansfield, “Task-specific Experience and Task-specific Talent: Decomposing the Productivity of High School Teachers,” *Journal of Public Economics*, 140:51–72, 2016.

- Cunha, F. and J. J. Heckman, “Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation,” *Journal of Human Resources*, 43(4):738–782, 2008.
- Dolton, P. and W. v. d. Klaauw, “The Turnover of Teachers: A Competing Risks Explanation,” *Review of Economics and Statistics*, 81(3):543–550, 1999.
- Fan, X., A. Seshadri and C. Taber, “Estimation of a Life-Cycle Model with Human Capital, Labor Supply and Retirement,” 2015.
- Freyberger, J., “Normalizations and Misspecification in Skill Formation Models,” 2021.
- Fu, C., H.-C. L. Lin and A. Tanaka, “College Loans and Post-Schooling Skill Accumulation,” *Working Paper*, 2019.
- Glewwe, P., N. Ilias and M. Kremer, “Teacher Incentives,” *American Economic Journal: Applied Economics*, 2(3):205–227, 2010.
- Haggag, K., B. McManus and G. Paci, “Learning by Driving: Productivity Improvements by New York City Taxi Drivers,” *American Economic Journal: Applied Economics*, 9(1):70–95, 2017.
- Haley, W. J., “Estimation of the Earnings Profile from Optimal Human Capital Accumulation,” *Econometrica*, pp. 1223–1238, 1976.
- Hanushek, E. A., “Teacher Characteristics and Gains in Student Achievement: Estimation Using Micro Data,” *American Economic Review*, 61(2):280–288, 1971.
- , “Conceptual and Empirical Issues in the Estimation of Educational Production Functions,” *Journal of Human Resources*, pp. 351–388, 1979.
- , “The Failure of Input-Based Schooling Policies,” *Economic Journal*, 113(485):F64–F98, 2003, ISSN 1468-0297.
- , “The Economic Value of Higher Teacher Quality,” *Economics of Education Review*, 30(3):466–479, 2011.
- Hanushek, E. A. and M. E. Raymond, “Does School Accountability Lead to Improved Student Performance?” *Journal of Policy Analysis and Management*, 24(2):297–327, 2005.
- Hanushek, E. A. and S. G. Rivkin, “Teacher Quality,” *Handbook of the Economics of Education*, 2:1051–1078, 2006.
- , “The Distribution of Teacher Quality and Implications for Policy,” *Annual Review of Economics*, 4(1):131–157, 2012.
- Harris, D. N. and T. R. Sass, “Teacher Training, Teacher Quality and Student Achievement,” *Journal of Public Economics*, 95(7):798–812, 2011.
- Heckman, J. J., “A Life-Cycle Model of Earnings, Learning, and Consumption,” *Journal of Political Economy*, 84(4):S11–44, 1976.
- Heckman, J. J., L. Lochner and R. Cossa, “Learning-by-Doing vs. On-the-Job Training: Using Variation Induced by the EITC to Distinguish Between Models of Skill Formation,” *NBER Working Paper Series*, p. 9083, 2002.

- Heckman, J. J., L. Lochner and C. Taber, “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents,” *Review of Economic Dynamics*, 1(1):1–58, 1998.
- Hölmstrom, B. and P. Milgrom, “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics, and Organization*, 7:24–52, 1991.
- Horowitz, J. L. and C. F. Manski, “Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data,” *Journal of the American Statistical Association*, 95(449):77–84, 2000.
- Imbens, G. W. and C. F. Manski, “Confidence Intervals for Partially Identified Parameters,” *Econometrica*, 72(6):1845–1857, 2004.
- Imberman, S. A. and M. F. Lovenheim, “Incentive Strength and Teacher Productivity: Evidence from a Group-Based Teacher Incentive Pay System,” *Review of Economics and Statistics*, 97(2):364–386, 2015.
- Jackson, C. K., J. E. Rockoff and D. O. Staiger, “Teacher Effects and Teacher-related Policies,” *Annual Review of Economics*, 6(1):801–825, 2014.
- Kaido, H., F. Molinari and J. Stoye, “Confidence Intervals for Projections of Partially Identified Parameters,” *Econometrica*, 87(4):1397–1432, 2019.
- Kane, T. J., D. F. McCaffrey, T. Miller and D. O. Staiger, “Have We Identified Effective Teachers? Validating Measures of Effective Teaching Using Random Assignment,” *Research Paper. MET Project. Bill & Melinda Gates Foundation*, 2013.
- Killingsworth, M. R., “‘Learning by Doing’ and ‘Investment in Training’: A Synthesis of Two ‘Rival’ Models of the Life Cycle,” *Review of Economic Studies*, 49(2):263–271, 1982.
- Kinsler, J., “Assessing Rothstein’s Critique of Teacher Value-Added Models,” *Quantitative Economics*, 3(2):333–362, 2012.
- Kraft, M. A. and J. P. Papay, “Can Professional Environments in Schools Promote Teacher Development? Explaining Heterogeneity in Returns to Teaching Experience,” *Educational Evaluation and Policy Analysis*, 36(4):476–500, 2014.
- Kuruscu, B., “Training and Lifetime Income,” *American Economic Review*, 96(3):832–846, 2006.
- McCaffrey, D. F., J. Lockwood, D. M. Koretz and L. S. Hamilton, *Evaluating Value-Added Models for Teacher Accountability. Monograph.*, ERIC, 2003.
- McWilliams, J., “The Lucrative Art of Chicken Sexing,” *Pacific Standard*, 2018.
- Muralidharan, K. and V. Sundararaman, “Teacher Performance Pay: Experimental Evidence from India,” *Journal of Political Economy*, 119(1):39–77, 2011.
- Murnane, R. J., *The Impact of School Resources on the Learning of Inner City Children.*, Cambridge, MA: Ballinger, 1975.
- Ost, B., “How Do Teachers Improve? The Relative Importance of Specific and General Human Capital,” *American Economic Journal: Applied Economics*, 6(2):127, 2014.

- Papay, J. P. and M. A. Kraft, "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, 2015.
- Petronijevic, U., "Incentives and Productivity in the Economics of Education," Ph.D. thesis, University of Toronto (Canada), 2016.
- Podgursky, M. and M. Springer, "Teacher Compensation Systems in the United States K-12 Public School System," *National Tax Journal*, 64(1):165–192, 2011.
- Poletaev, M. and C. Robinson, "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000," *Journal of Labor Economics*, 26(3):387–420, 2008.
- Rivkin, S. G., E. A. Hanushek and J. F. Kain, "Teachers, Schools, and Academic Achievement," *Econometrica*, 73(2):417–458, 2005, ISSN 1468-0262.
- Robinson, C., "Occupational Mobility, Occupation Distance, and Specific Human Capital," *Journal of Human Resources*, 53(2):513–551, 2018.
- Rosen, S., "Learning and Experience in the Labor Market," *Journal of Human Resources*, pp. 326–342, 1972.
- Sanders, C., "Skill Uncertainty, Skill Accumulation, and Occupational Choice," *Working Paper*, 2010.
- Sanders, C. and C. Taber, "Life-cycle Wage Growth and Heterogeneous Human Capital," *Annu. Rev. Econ.*, 4(1):399–425, 2012.
- Shaw, K. and E. P. Lazear, "Tenure and Output," *Labour Economics*, 15(4):704–723, 2008.
- Shaw, K. L., "Life-cycle Labor Supply with Human Capital Accumulation," *International Economic Review*, pp. 431–456, 1989.
- Stinebrickner, T. R., "A Dynamic Model of Teacher Labor Supply," *Journal of Labor Economics*, 19(1):196–230, 2001.
- Stoye, J., "More on Confidence Intervals for Partially Identified Parameters," *Econometrica*, 77(4):1299–1315, 2009.
- Strøm, B. and T. Falch, "The Role of Teacher Quality in Education Production," in "The Economics of Education," pp. 307–319, Elsevier, 2020.
- Tamer, E., "Partial Identification in Econometrics," *Annual Review of Economics*, 2(1):167–195, 2010.
- Tincani, M. M., "Teacher Labor Markets, School Vouchers and Student Cognitive Achievement: Evidence from Chile," *Quantitative Economics*, 12:173–216, 2021.
- Willis, R. J., *Wage Determinants: A Survey and Reinterpretation of Human Capital Earnings Functions*, Economics Research Center/NORC, 1985.
- Wiswall, M., "The Dynamics of Teacher Quality," *Journal of Public Economics*, 100:61–78, 2013.
- Yamaguchi, S., "Tasks and Heterogeneous Human Capital," *Journal of Labor Economics*, 30(1):1–53, 2012.