



The Education Policy Center
AT MICHIGAN STATE UNIVERSITY

WORKING PAPER #38

Evaluating Specification Tests in the Context of Value-Added Estimation

Cassandra M. Guarino
Mark D. Reckase
Brian W. Stacy
Jeffrey M. Wooldridge

Michigan State University

February 6, 2014

The content of this paper does not necessarily reflect the views of The Education Policy Center or Michigan State University

Evaluating Specification Tests in the Context of Value-Added Estimation

Author Information

Cassandra M. Guarino, Indiana University

Mark D. Reckase

Brian W. Stacy

Jeffrey W. Wooldridge

Michigan State University

This work was supported by grant number R305D10002 from the Institute for Education Sciences in the U.S. Department of Education.

Abstract

We study the properties of two specification tests that have been applied to a variety of estimators in the context of value-added measures (VAMs) of teacher and school quality: the Hausman test for choosing between random and fixed effects and a test for feedback (sometimes called a “falsification test”). We discuss theoretical properties of the tests to serve as background. An extensive simulation study provides important further insight to the VAM setting. Unfortunately, while both the Hausman and feedback tests have good power for detecting the kinds of nonrandom assignment that can invalidate VAM estimates, they also reject in situations where estimated VAMs perform very well. Consequently, the tests must be used with caution when student tracking is used to form classrooms.

Evaluating Specification Tests in the Context of Value-Added Estimation

Cassandra M. Guarino, Mark D. Reckase, Brian W. Stacy, Jeffrey M. Wooldridge

February 6, 2104

This work was supported by grant number R305D10002 from the Institute for Education Sciences in the
U.S. Department of Education.

Abstract: We study the properties of two specification tests that have been applied to a variety of estimators in the context of value-added measures (VAMs) of teacher and school quality: the Hausman test for choosing between random and fixed effects and a test for feedback (sometimes called a “falsification test”). We discuss theoretical properties of the tests to serve as background. An extensive simulation study provides important further insight to the VAM setting. Unfortunately, while both the Hausman and feedback tests have good power for detecting the kinds of nonrandom assignment that can invalidate VAM estimates, they also reject in situations where estimated VAMs perform very well. Consequently, the tests must be used with caution when student tracking is used to form classrooms.

1. Introduction

Measures of teacher and school quality based on value-added models (or VAMs) of student achievement are gaining increasing acceptance among policymakers as a tool for evaluating teaching and school effectiveness. Therefore, it is important for researchers and policy makers to understand the statistical properties of the estimates derived from VAMs, and to have some knowledge of when they can be expected to perform well – and when they do not perform well. One way to proceed is to apply statistical tests of the assumptions underlying VAMs to see if they appear justified. Rothstein (2010) and Harris, Sass, and Semykina (2010) are two examples of studies that develop and apply statistical tests of assumptions to VAMs designed to produce measures of teacher effectiveness.

In applying statistical diagnostics in VAM settings, it is imperative to be clear about the goal of the analysis. Is it to determine whether all of the assumptions underlying consistent estimation of a structural production function hold? Or is the main goal to get good estimates of value added – say of teachers or schools? In the literature and popular press, the main focus appears to be on getting good estimates of the value-added measures. Structural models are often used to motivate the estimation procedure, but the performance of the value added estimates is of primary interest. If we assume that providing relatively accurate performance estimates is the primary purpose of the VAM literature, then it is critical to understand the difference between rejecting assumptions of an underlying structural model and concluding that a particular procedure likely produces poor estimates of value added.

In earlier work (Guarino, Reckase, Wooldridge, forthcoming – hereafter, GRW), we provide a summary of the known theoretical properties of various approaches to estimating VAMs designed to yield teacher performance estimates. More importantly, we provide extensive

simulation evidence showing how six of the most commonly used estimators behave under different mechanisms used to match teachers and students. One of the key findings is that certain estimators that are not technically consistent can perform well in estimating teacher effects for the purpose of ranking. One of the estimators, in particular – ordinary least squares (OLS) applied to a dynamic gain-score equation, which we dubbed “dynamic OLS,” or “DOLS” – performs best across many scenarios, although other estimators are slightly better under some specific assignment mechanisms.

A plethora of models and estimation strategies have been applied to the task of estimating teacher and school effects, often producing quite different estimates. Given the several choices among estimators in the VAM context it would be helpful to have tools for choosing among different estimators, especially when they produce very different estimated effects or, more fundamentally, whether any estimator does a good enough job of estimating effects to enable the use of these estimates for policy purposes, such as rewarding or sanctioning teachers based on estimated performance. It is logical to turn to the large array of statistical tests that currently exist to help diagnose whether or not underlying assumptions are met; many have been developed in the statistical and econometric literature in reference to other topics and issues, and some recent tests have been proposed by education researchers for the express purpose of evaluating VAMs (for example, those in Rothstein (2010) and Harris, Sass, and Semykina (2010)).

The main purpose of this paper is to determine the usefulness of available tests in determining how well a VAM accomplishes its task; in this study the task is estimating teacher effects. To do so, we use simulations in which we originate student test score data based on known teacher effects but then act as if we do not know the true effects and estimate them. We then assess both the degree to which a model and estimation approach yield accurate teacher

effect estimates and the behavior of statistical tests of assumptions. Our goal is to determine the usefulness of statistical tests in revealing the quality of specific models and estimators for estimating teacher effects.

The two tests we focus on are those aimed at ferreting out conditions that might create bias in the estimated teacher effects. Both tests are designed to detect nonrandom teacher assignment – that is, teachers are assigned at least partly on the basis of observed or unobserved student characteristics. The first test – a robust version of the Hausman (1978) test comparing the random and fixed effects estimators – is primarily intended to uncover situations where teacher assignment is based on unobserved, time-constant student heterogeneity. The test has power for detecting other kinds of nonrandom assignment mechanisms but its main purpose is to determine whether teacher assignment is correlated with student heterogeneity.

The second test is perhaps better described as a class of tests, whose purpose is to detect dynamic teacher assignment mechanisms. One test was popularized in the VAM context by Rothstein (2010), who called it a “falsification” test. A version of the falsification test is known in the panel data literature as a test of the “strict exogeneity” assumption in the context of fixed effects estimation. The test essentially looks for feedback from shocks to student performance today into future teacher assignment. Such feedback effects cause inconsistency in the FE estimator.

Ideally, the diagnostic tests would reject various estimation methods when they produce poor value-added measures. Unfortunately, our findings are not very positive from a practitioner’s perspective. While in many cases the tests properly reject when the estimation methods produces poor value-added estimates, in other cases the tests strongly reject when the underlying estimation method is actually capable of producing very good estimates of the

VAMs. The particular situation that causes problems for both the Hausman and Rothstein-type falsification tests is when students are tracked based on some observable or unobservable factor, but the classrooms are randomly assigned to teachers. In such cases a variety of estimation methods produce reliable VAMs – depending on the nature of the tracking. Yet, as we show in Section 6, the specification tests strongly reject many of the best estimators.

An important consequence of our findings is that criticisms of VAMs on the basis of evidence provided by Hausman or feedback tests are likely to be unjustified. The bottom line is that, applied in the VAM context, the tests have power for detecting nonrandom assignment schemes that have nothing to do with whether popular estimators are doing their main job: provided good VAM estimates.

The rest of the paper is organized as follows. We discuss the value-added modeling framework in Section 2. Section 3 describes the statistical tests that we study – both those that have been applied to VAMs by other researcher and some that have not—and discusses their theoretical properties. We discuss the kinds of nonrandom grouping and assignment mechanisms that seem particularly relevant in Section 4. In Section 5 we discuss our simulation design, and Section 6 discusses the simulation results. We provide some concluding remarks in Section 7.

2. Conceptual Framework for Testing Value-Added Models

It is helpful to begin with a fairly general value-added equation and a brief discussion of the assumptions embedded in it. Assume that the achievement score, A_{it} , is generated as

$$A_{it} = \tau + \lambda A_{i,t-1} + E_{it}\beta_0 + c_i + u_{it} - \lambda u_{i,t-1} \quad (1)$$

$$u_{it} = \rho u_{i,t-1} + r_{it}, \quad t = 1, 2, \dots, T, \quad (2)$$

where A_{it} is a measure of achievement for student i in grade (or year) t and E_{it} is the (row) vector of educational inputs whose coefficients, β_0 , are of greatest interest. Generally, E_{it} can include inputs at the school, classroom, or even individual level. In the present paper, E_{it} is a vector of teacher assignment indicators. We assume $\{r_{it}\}$ is a sequence of independent, identically distributed normal random variables with mean zero so that $\{u_{it}\}$ follows an AR(1) model.

Sometimes it is useful to subtract $A_{i,t-1}$ from both sides of (1) to obtain an equation for the gain score, ΔA_{it} :

$$\Delta A_{it} = \tau_t + \alpha A_{i,t-1} + E_{it}\beta_0 + c_i + u_{it} - \lambda u_{i,t-1}, \quad (3)$$

where $\alpha = \lambda - 1$.

Equation (1) can be derived from a general cumulative effects model (CEM) under various assumptions, and (2) adds the assumption that the errors $\{u_{it}\}$ have a particular pattern of serial correlation. The parameter λ is the decay parameter in the CEM. Even this restricted version of the CEM is never estimated, as accounting for the combined issues of heterogeneity, ρ different from λ , and the lagged dependent variable is a challenging econometric problem.

There are several types of misspecifications that can cause difficulty for standard estimators of β_0 . One is failure of the so-called ‘‘common factor restriction’’ (CFR), $\lambda = \rho$. Under the CFR the errors $u_{it} - \lambda u_{i,t-1}$ in (1) have no serial correlation, but if $\lambda \neq \rho$ then (1)

contains serial correlation. In general, serial correlation in the presence of a lagged dependent variable causes inconsistent estimation for many estimation procedures, including OLS applied to (1) (where we ignore both the presence of c_i and serial correlation in $u_{it} - \lambda u_{i,t-1}$). The Arellano and Bond (1991) instrumental variables procedure, which removes c_i , relies on no serial correlation in the errors in (1).

McClain and Wooldridge (1995) propose a simple test of the null hypothesis that the CFR holds in the context of time series regression models. The test is easily adapted to the panel data case when there is no heterogeneity, that is, when c_i is not in (1). Unfortunately, it is not clear how to extend the test to allow for heterogeneity. Any neglected serial correlation, due to violation of the CFR, higher order autoregressive properties, or the presence of c_i will cause a rejection of the CFR restriction. In GRW we found that violation of the CFR did not appreciably affect the dynamic OLS estimator in the sense that DOLS still provided estimates of the teacher effects that produced reliable rankings among teachers. For these reasons we do not study the CFR test further in this paper.

A second kind of misspecification arises in setting λ in equation (1) equal to unity [which is the same as dropping $A_{i,t-1}$ in equation (3)]. Including the lagged achievement in equation (3) is a simple, effective way to detect dynamic misspecification. Harris, Sass, and Semykina (2010) apply tests for dynamic misspecification by including lagged teacher assignment using data from Florida. We do not study the properties of dynamic misspecification tests in the current paper as they are standard tests of omitted variables. Here we are mainly interested in the behavior of exogeneity tests when λ is incorrectly set to unity. Incidentally, in the current testing setting, omitting $A_{i,t-1}$ from (1) and failure of the CFR restriction can be expected to have similar consequences because, in effect, a variable that can predict the gain score is omitted from the

equation – and teacher assignment might be correlated with that variable. From here on we focus on misspecified dynamics due to setting λ to an incorrect value (unity in our case).

Finally – and most importantly for this paper – we are interested in what happens when teachers were assigned to students in such a way to make E_{it} endogenous in an estimating equation. In this third kind of “misspecification” it is not necessarily true that (1) is an incorrect equation; it is that inputs have been chosen in a way to violate certain exogeneity requirements, resulting in inconsistent estimators.

3. A Discussion of the Tests in the VAM Setting

The general purpose of the specification tests we study is to detect nonrandom assignment of students to teachers. We consider tests of both static and dynamic assignment. Static assignment occurs when teachers are (partly) assigned on the basis of unobserved student heterogeneity – that is, students with fixed but unobserved characteristics are matched with particular teacher effectiveness levels. Dynamic assignment occurs when the prior test scores of students are matched to particular teacher effectiveness levels.

In what follows it is important to distinguish between two mechanisms that can be used for generating classrooms of students taught by particular teachers. Students may be first grouped on the basis of unobserved or observed characteristics – a process often referred to as “tracking.” This kind of grouping might be done even if teachers are randomly assigned to classrooms. Nonrandom teacher assignment occurs when classrooms with different average levels of ability or achievement are systematically assigned to teachers with different levels of competence.

Although little research exists on evaluating specification tests within the VAM setting, the literature on the most commonly applied tests is vast in the sense that the tests are fairly

standard in the panel data literature. For example, Wooldridge (2010, Section 10.7.3) discusses different versions of the Hausman test used to compare the RE and FE estimators. In the VAM context, the Hausman test is primarily a test of static assignment mechanisms because it is intended mainly to pick up correlation between student heterogeneity and the observed inputs – teacher assignment in this case. (Nevertheless, the test generally has power against other misspecifications that cause systematic deviation between the RE and FE estimators.) As discussed in Wooldridge (2010), it is generally important to use a version of the Hausman test that is robust to violations of assumptions that are not required for consistently estimating the parameters – in this case, the teacher value-added measures. In particular, we prefer tests that are robust to serial correlation and heteroskedasticity in the students’ idiosyncratic shocks. A regression-based test, which we review in Section 3.2, provides a straightforward method for obtaining a fully robust Hausman test.

In the VAM context, Rothstein (2010) proposes a different test, which he calls a “falsification test,” to detect nonrandom assignment, and he provided a small simulation study in addition to applying the test to data from North Carolina. In the panel data nomenclature, Rothstein’s test can be viewed as detecting violation of strict exogeneity of the explanatory variables. In the VAM context, violations of strict exogeneity can occur under dynamic assignment – that is, when students are assigned to teachers at least partly based on realization of past test scores (or shocks to those scores). Another way to think of violation of strict exogeneity is that shocks to the test score today feed into future teacher assignment. To test this assumption, Rothstein includes future teacher assignments in a current gain-score equation. Importantly, he does not allow for student fixed effects. Without fixed effects, standard estimators, such as OLS, do not require strict exogeneity for consistent estimation. Thus, it is not clear why one wants to

test for the presence of dynamic assignment in such cases. By contrast, because failure of strict exogeneity does result in inconsistency when student-level fixed effects are included, Wooldridge (2010, Section 10.7.1) shows how feedback effects are easily tested in the context of fixed effects estimation. The most straightforward way to test for feedback effects is to include future values of the explanatory variables – usually one-period ahead – and test their significance using a robust Wald test after FE estimation. We discuss variants of the falsification test in Section 3.3.

In addition to Rothstein (2010), some other recent empirical papers have applied falsification tests in the VAM context. Koedel and Betts (2009) applied falsification tests in the context of VAM estimation using data from the San Diego School District. Harris, Sass, and Semykina (2010) (HSS) apply a battery of tests (several of which we do not study here) in an empirical context. In their application to data from Florida, HSS generally find evidence against random assignment of students to teachers and find estimated VAMs that vary widely across procedures.

Other papers have independently studied Rothstein’s falsification test using both theoretical calculations and simulation arguments. Goldhaber and Chaplin (2012) provide an evaluation of Rothstein’s test, first by studying whether Rothstein’s statistic properly detects omitted variable bias. They conclude that it is possible to have data generating mechanisms that do not produce biased VAMs but where the Rothstein test will reject the specification. We come to a similar conclusion but using a different route. In particular, we generate test score data using different tracking and teacher assignment mechanism in order to mimic how principals might actually match students and teachers. Also, we consider the Hausman test for comparing the random effects and fixed effects estimators and focus on a more common panel data test for

“feedback” in the context of fixed effects estimation (and also study how it works for dynamic regression).

Kinsler (2012) studies a particular version of Rothstein’s falsification test in the presence of student heterogeneity based on Chamberlain’s (1984) correlated random effects approach. His conclusions are broadly similar to those found by Goldhaber and Chaplin and what we find here. We instead use the regression based version of the Hausman test that is computationally simple and can be applied to unbalanced panels. We discuss these tests further in the next section.

3.1. A Basic Gain-Score Equation and Exogeneity Assumptions

In describing tests for endogeneity of teacher assignment we start with a gain-score equation because both the Hausman test and the falsification test (test of strict exogeneity) can be reasonably applied. It is rare to see such tests applied when the dependent variable is a level-score rather than a gain-score, but the following discussion applies when the level is used.

A standard gain-score equation is

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + X_{it}\gamma_0 + c_i + e_{it}, \quad t = 1, \dots, T. \quad (4)$$

The vector X_{it} includes controls, many of which may be constant, that are often included in empirical VAM studies, such as gender, race/ethnicity, disability status, free-and-reduced lunch eligibility. In this paper we do not include extra controls in our simulation study, but for a general discussion of how to apply tests, it is useful to explicitly include X_{it} .

The constants τ_t allow for different intercepts for different grades (or, with many cohorts, allows for cohort effects). Because we have few grades (time periods), these can be estimated precisely with a large number of students. The c_i are the time-constant student unobserved effects, sometimes called “student heterogeneity.” The presence of c_i causes the composite error,

$v_{it} = c_i + e_{it}$, to be serially correlated. More importantly, if c_i is correlated with the inputs E_{it} , leaving c_i in the error term can cause inconsistency in estimating β_0 .

The idiosyncratic errors, $\{e_{it}\}$, are time-varying unobserved factors that affect gain scores. Generally, these can be serially correlated or heteroskedastic, or both. In the context of an underlying cumulative effects model, e_{it} is a linear combination of the errors appearing in the structural production function; see, for example, GRW for a more extensive discussion of the structural model.

An important assumption required of the most common panel data estimators that recognize the presence of c_i – RE and FE – is strict exogeneity of the inputs conditional on the student heterogeneity, namely,

$$E(e_{it} | E_{iT}, E_{i,T-1}, \dots, E_{i1}, X_{iT}, X_{i,T-1}, \dots, X_{i1}, c_i) = 0, \quad t = 1, \dots, T. \quad (5)$$

Note that in (5) the expectation of the error term, conditional on heterogeneity and all current, past, and *future* inputs is zero.

To interpret the strict exogeneity assumption, drop the $\{X_{is}: s = 1, \dots, T\}$ for simplicity. Then, when we combine (5) with equation (1), we have

$$E(\Delta A_{it} | E_{iT}, \dots, E_{i1}, c_i) = \tau_t + E_{it}\beta_0 + c_i. \quad (6)$$

Equation (6) means that, once we control for student heterogeneity, only inputs at time t , E_{it} , appear in the gain-score equation at time t . This restriction implies that past inputs – in this case, previous teachers – have no effect on the current gain score, once current teacher and student heterogeneity have been accounted for. As discussed in HSS, it is simple to test such an assumption: simply include, say, $E_{i,t-1}$ and test for joint significance (using whatever estimation method one settles on). Of course, including lagged inputs costs us a year of data, but such tests typically can be carried out for the kinds of data sets available.

For our purposes assumption (5) has another important implication: *future* values, such as $E_{i,t+1}$, do not appear on the right-hand side of (6). Generally, if teacher assignment at time $t + 1$ depends on ΔA_{it} or A_{it} , $E_{i,t+1}$ will be (partially) correlated with e_{it} . Shortly, we use this observation to obtain a test of (5).

In addition to strict exogeneity of the inputs conditional on c_i , another important assumption in the panel data literature is

$$E(c_i | E_{iT}, E_{i,T-1}, \dots, E_{i1}) = E(c_i) = 0, \quad (7)$$

where we have again dropped the $\{X_{is}; s = 1, \dots, T\}$. The assumption that $E(c_i) = 0$ is without loss of generality when the gain-score equation has an intercept (or a full set of time intercepts). When we combine assumptions (6) and (7), the inputs $\{E_{it}\}$ are strictly exogenous with respect to the composite error:

$$E(v_{it} | E_{iT}, E_{i,T-1}, \dots, E_{i1}) = 0, \quad t = 1, \dots, T. \quad (8)$$

Assumption (8) is important, as it justifies generalized least squares estimation (GLS) – including the popular RE estimator – applied to

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + v_{it}, \quad t = 1, \dots, T. \quad (9)$$

A special case of GLS is pooled OLS (POLS), where any serial correlation in v_{it} due to the presence of c_i – in fact, any serial correlation – is ignored. Inference is handled by using a robust variance matrix estimator (which is also robust to heteroskedasticity of arbitrary form). From (9) it is easily seen that consistency of POLS only requires that v_{it} and E_{it} are uncorrelated; the strict exogeneity assumption in (8) is not needed. However, because v_{it} includes c_i , POLS requires that the inputs are uncorrelated with the student-specific heterogeneity.

3.2. The Hausman Test Comparing RE (or POLS) to FE

In many empirical panel data applications, including VAM estimation, one often estimates the gain-score equation (4) by both RE and FE. Both estimators require the strict exogeneity assumption stated in (5) for consistency. In addition, RE uses the heterogeneity exogeneity condition in (7). Therefore, it is common to compare the RE and FE estimates as a test of (7). However, it is important to remember that RE and FE – and, for that matter, POLS – will typically have different probability limits if the strict exogeneity assumption (1) is violated. Therefore, any test that explicitly or implicitly compares the RE and FE estimators (or the POLS and FE estimators) generally has power against violation of (4) or (5), and one cannot use the outcome of the Hausman test to conclude which assumption fails, or whether both fail.

The traditional form of the Hausman (1978) statistic uses a quadratic form based on the differences between the RE and FE estimators. A critical point in applying the traditional form is that it assumes that the RE estimator is (asymptotically) efficient: the variance-covariance matrix appearing in the quadratic form is valid only when RE is asymptotically efficient. As discussed by Wooldridge (2010, Section 10.7.3), the relative efficiency of the RE estimator holds only when the idiosyncratic errors $\{e_{it}\}$ are serially uncorrelated and homoskedastic – both conditional on the covariates and c_i . (See Wooldridge, 2010, 10.7.3 for a formal statement of the assumptions.) Yet the Hausman test has no power for detecting serial correlation or heteroskedasticity in $\{e_{it}\}$ because these problems do not cause inconsistency in either the RE or FE estimator. In the language of Wooldridge (1990), the traditional form of the Hausman test adds “auxiliary assumptions,” which are used to get a standard null distribution even though the test has no power for detecting failure of the assumptions.

It is possible but computationally cumbersome to modify the usual Hausman statistic to be robust to arbitrary serial correlation and heteroskedasticity in $\{e_{it}\}$. One problem is that the variance-covariance matrix is singular when the estimated equation includes time effects, which is very common. A much more straightforward approach is to use a robust, regression-based test.

The regression-based Hausman test is based on the correlated random effects specification

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + \bar{E}_i\xi + a_i + e_{it} \quad (10)$$

where $\bar{E}_i = T^{-1} \sum_{r=1}^T E_{ir}$ is the time average and $c_i = \bar{E}_i\xi + a_i$. In this formulation, we explicitly model the heterogeneity c_i as a linear function of the time average of the inputs (which is where the name “correlated random effects” comes from). Equation (7) still contains unobserved heterogeneity, a_i , but it is uncorrelated with the entire history of inputs, $\{E_{it}\}$. If we maintain strict exogeneity conditional on c_i then strict exogeneity holds conditional on a_i in (10). Therefore, equation (10) can be estimated by POLS or random effects.

A well-known algebraic result (for example, Wooldridge, 2010, Section 10.7.3) is that when POLS or RE is applied to (10), the resulting estimate of β_0 is the fixed effects estimator that uses deviations from time averages to remove c_i from the equation

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + c_i + e_{it} \quad (11)$$

Therefore, equation (10) is very useful for presenting a unified setting for RE and FE estimation. In particular, if (10) is estimated by random effects (that is, feasible GLS using the RE structure), it is straightforward to construct a robust Wald test of $H_0: \xi = 0$, which has as many degrees of freedom as there are inputs E_{it} ¹. Obtaining a Wald test that is robust to arbitrary serial

¹ The POLS and RE estimates of β_0 are equal to the FE estimator. Further, with a balanced panel the POLS and RE estimates of ξ are the same. (They generally differ with an unbalanced panel, in which case RE will be more efficient under the standard RE assumptions.) Whether POLS or RE is used, the test should be made fully robust to serial correlation and heteroskedasticity in $\{e_{it}\}$.

correlation or heteroskedasticity in $\{e_{it}\}$, while remaining asymptotically efficient under the traditional RE assumptions, is straightforward using popular packages that support RE and FE estimation². The regression-based test is asymptotically equivalent (against local alternatives) to the traditional Hausman test when the $\{e_{it}\}$ are serially uncorrelated and homoskedastic.

A rejection of $H_0: \xi = 0$ is typically taken to mean that c_i is correlated with \bar{E}_i but, as mentioned earlier, this interpretation is based on maintaining assumption (5). If we reject $H_0: \xi = 0$ then we have found that \bar{E}_i is correlated with the composite error, $c_i + e_{it}$, which warrants a statistical rejection of the RE estimator. However, as we will see in our simulations in Section 6, in the context of estimating VAMs one must be cautious in using the Hausman test in this way. It could be that RE is statistically rejected but provides better estimates of the VAM coefficients than its natural alternative, FE. Even though the RE estimates of VAMs might be systematically biased, they typically have less sampling variation – sometimes much less – and the bias may be such that the estimated VAMs do a good job of ranking teachers. We will have more to say on this in Section 6.

A practical problem with using equation (10) as the basis for the Hausman test is that, with many teachers, (10) contains many regressors: the original teacher dummies and then the proportion of times the student sees that teacher over a student’s entire observed history (that is, the time average). Computationally, many regressors are not too difficult to handle with modern computers and statistical packages. A more pressing concern is potential finite-sample distortions in using large-sample critical values (which is what the Hausman approach necessarily uses). The proper asymptotics relies on the number of students per teacher getting “large.” In practice, we may not have many student outcomes associated with some teachers. In such cases, a one

² In Stata, a fully robust Wald test is easily obtained using the “cluster,” option, which is how we carry out the test in our simulations.

degree-of-freedom test may have better size properties. Rather than including the entire vector \bar{E}_i and testing joint significance, which necessitates a separate variable for each teacher in the dataset, we propose a new test. This consists of substituting, for each student, the estimated *average teacher effect* across all years for the vector \bar{E}_i . This is identical to estimating the equation

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + \alpha(\bar{E}_i\hat{\beta}_0) + error_{it} \quad (12)$$

and performing a t test of $\alpha = 0$. The estimate $\hat{\beta}_0$ is the RE estimate from the equation (11) obtained in a first stage. Like equation (10), equation (12) can be estimated by RE, preferably with a fully robust t statistic.

A test using (12) rather than (10) conserves on degrees of freedom, but it may not detect certain kinds of teacher assignment mechanisms. In our simulation we study the properties of both tests and find that the one-degree-of-freedom test has substantial power against nonrandom assignment alternatives.

In many applications of RE estimation in the VAM context, other explanatory variables are included as controls. Often such controls are student characteristics, such as family background, socioeconomic status, or baseline test scores that do not vary over time. (Test scores lagged one or more period are not allowed in RE estimation because lagged dependent variables always violate the strict exogeneity assumption.) When available, it is important to include such controls in equation (10), leading to an equation such as

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + \bar{E}_i\xi + Z_i\gamma + a_i + e_{it}, \quad (13)$$

where Z_i is the vector of time-constant controls. With good controls, it is more plausible that the (remaining) unobserved heterogeneity is uncorrelated with $\{E_{it}\}$. One can also include time-varying, strictly exogenous controls, say $\{X_{it}\}$, and then (10) becomes

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + \bar{E}_i\xi + Z_i\gamma + X_{it}\eta + \bar{X}_i\lambda + a_i + e_{it}, \quad (14)$$

where we also include the time averages of $\{X_{it}\}$. To test whether the inputs are partially correlated with heterogeneity we would still test $H_0: \xi = 0$; failing to reject means we can drop \bar{E}_i from (11) and estimate the equation by RE, typically obtaining a more precise estimator of β_0 .³ The test described in equation (12) can also be applied when additional covariates are included in the model. In our simulations, we only consider an equation with teacher dummies and no other inputs.

Wooldridge (2009) shows that equation (13) can be used as the basis for a Hausman test even in the case of unbalanced panels – provided the reason the panel is unbalanced is appropriately exogenous. One subtle point is that a time period should be used in constructing the time averages only when observations on all variables are available. In the simulations later we only consider balanced panels but most panel data sets are, at least initially, unbalanced.

3.3. A Test of Strict Exogeneity Using Fixed Effects

If the RE estimator is rejected using the regression-based Hausman statistic from Section 3.3, a natural step is to use the FE estimator so that arbitrary correlation is allowed between c_i and $\{E_{it}\}$. Because consistency of the FE estimator relies on strict exogeneity, it is (potentially) important to test that assumption. Here we are interested in testing for feedback, assuming under the null that only current inputs appear in the gain-score equation at time t .

An auxiliary equation that leads to a simple test is

$$\Delta A_{it} = \tau_t + E_{it}\beta_0 + E_{i,t+1}\delta + c_i + r_{it}, \quad t = 1, \dots, T - 1, \quad (15)$$

³ Guggenberger (2010) warns of the problems of using the Hausman test as a pre-test for choosing between RE and FE. The regression-based version of the Hausman test makes it clear that the Hausman pre-testing problem is essentially the same as the problem of pre-testing whether a set of regressors belongs in an equation and then using an F or Wald test to determine whether those regressors appear in the final model.

where we lose the last time period (grade) by putting the future inputs into the equation at time t . Equation (15) should be estimated by fixed effects in order to allow the heterogeneity and inputs E_{it} to be correlated under the null, making a test of $H_0: \delta = 0$ a pure test of strict exogeneity (feedback in this case). Naturally, the test should be made robust to arbitrary serial correlation and heteroskedasticity in $\{r_{it}\}$ (through what are commonly called “cluster robust” test statistics).

We can add additional time-varying covariates to equation (15) and we may or may not include their lead values. As in any testing context, including a lot of irrelevant variables (lead values in this case) tends to reduce the power of the test. In our simulation study we do not have extra covariates.

Rothstein (2010) uses a version of the test from equation (15) but he applies the test one grade at a time. By using deviations from school means, Rothstein allows school fixed effects, but he does not allow unobserved student effects that are correlated with teacher assignment. In effect, Rothstein imposes the restriction $c_i = 0$, something that is important to recognize in interpreting the outcome of the test. Rothstein effectively applies the test to a cross-sectional regression with school fixed effects. Importantly, strict exogeneity of teacher assignment is *not* required for OLS with school dummies to consistently estimate teacher effects – provided there are many children per school (which is true in Rothstein’s setting and reasonable in general). In other words, if Rothstein thinks it is sufficient to control for school but not student effects, then he is testing an assumption that is not needed for consistent estimation of teacher effects⁴. It is only when student fixed effects are allowed in panel data that feedback necessarily causes inconsistent estimation of the teacher effects.

⁴ Rothstein (2010) proposes two versions of the test, one that excludes current teacher assignment and one that includes it. In practice, one should include current teacher assignment because it may be correlated with the next grade’s teacher assignment.

Rothstein also applies a version of the feedback test that is equivalent to testing the coefficients on $E_{i,t+1}$ in the following equation (which does not include student heterogeneity):⁵

$$A_{i,t-1} = \tau_t + \alpha A_{it} + E_{i,t+1} \delta + r_{it}, \quad t = 1, \dots, T - 1; \quad (16)$$

see also Goldhaber and Chaplin (2012), who focus on this particular test among those proposed by Rothstein. This test can be interpreted as checking whether future teacher assignment is related to the test score two years prior after controlling for previous year's score. For example, it checks for whether fifth grade teacher assignment depends on the third grade score once the fourth grade score has been partialled out. Consequently, this test should have power for detecting dynamic assignment mechanisms that depend on multiple lagged test scores, but it has little to do with whether standard VAM estimators are consistent. Essentially, the Rothstein test is irrelevant for evaluating VAM estimators provided we are willing to use dynamic regression with multiple lags of student achievement to control for nonrandom assignment. Goldhaber and Chaplin (2012) make a similar argument and obtain bias formulas for the estimated teacher effects under some simple scenarios. But it is easier, and more general, to simply understand that the regression in (16) is just one way of testing whether $E_{i,t+1}$ and $A_{i,t-1}$ are correlated after partialling out A_{it} . The absence of partial correlation is neither necessary nor sufficient for dynamic VAM estimators to consistently estimate teacher VAMS, let alone provide good rankings of teachers⁶.

In the context of dynamic regression where, for simplicity, we include only a single lagged test score, a more natural test comes from the equation

⁵ The test based on (16) is the same if $A_{i,t-1}$ is replaced with the gain score, ΔA_{it} because A_{it} is included as a regressor. Therefore, the coefficients on $E_{i,t+1}$ are the same whether $A_{i,t-1}$ or the gain score is used.

⁶ Rothstein (2010) rejects almost all specifications that use either no lags or a single lagged test score, which is consistent with teacher assignment that may depend on more than just the most recent test score. This testing outcome likely explains why Rothstein finds that the VAM estimates differ when more flexible dynamic models are used. But this simply means one should use the dynamic models that include several lags

$$A_{it} = \tau + \lambda A_{i,t-1} + E_{it}\beta_0 + E_{i,t+1}\delta + r_{it}. \quad (17)$$

The test of strict exogeneity of teacher assignment is that all elements of δ are zero. Unlike the Rothstein approach, (17) properly controls for current teacher assignment, and answers the question: Is future assignment correlated with current test scores after we partial out lagged test scores *and* current teacher assignment? This is the relevant test of strict exogeneity.

Nevertheless, even though we prefer (17) to Rothstein's approach, we must emphasize again that dynamic OLS does *not* require strict exogeneity of teacher assignment to consistently estimate the teacher VAMs. Consequently, it is not clear what we can learn, in general, from such a test. Nevertheless, because Rothstein-type tests are popular, we will evaluate the tests in a simulation study in the chance that the tests provide useful information.

As in the case of the Hausman test, a one-degree-of-freedom test can be used as an alternative to conserve on degrees of freedom. Rather than include the full set of teacher indicators, one includes the estimated teacher effect for next year's teacher. If $\hat{\beta}_0$ denotes the estimated teacher effect – using whichever method under study – then the regressor is simply $E_{i,t+1}\hat{\beta}_0$, a single linear combination of the lead teacher dummies. As always, it is prudent to make the t statistic robust to arbitrary within-student serial correlation and heteroskedasticity.

As a more complicated test that uses the panel structure and allows for student effects to be correlated with teacher assignment, Rothstein (2010) uses Chamberlain's (1984) minimum distance approach to unobserved effects models. While this approach can deliver somewhat more power when idiosyncratic shocks are serially correlated, it has several drawbacks. First, it requires testing nonlinear restrictions and therefore requires special programming. Second, Chamberlain's approach is very difficult to adapt to unbalanced panels, something that is important in practice. Rothstein avoids this problem by dropping observations until he has a

balanced panel, possibly leading to a more severe sample selection problem than would be otherwise present. Third, under the traditional fixed effects assumptions – essentially, that the idiosyncratic errors have no serial correlation or heteroskedasticity – the simpler test based on equation (15) is asymptotically efficient. As we will see in the simulations, the simpler test from (15) has plenty of power against nonrandom assignment mechanisms. In previous applications of the falsification test, failure to reject has not been an issue. In settings with complicated forms of serial correlation, Chamberlain’s approach can yield more power, but that would be traded off against the complicated nature of the test and the drawbacks to artificially balancing the sample. Kinsler (2012) studies only the Chamberlain form of the test, focusing on small-sample properties. Kinsler find that the test rejects the null too often, which is what we find here even without small sample bias.

4. Student Grouping, Teacher Assignment, and Behavior of the Tests

In our previous study of the properties of various estimators (GRW, forthcoming), we used several mechanisms for grouping students into classrooms and assigning teachers to those classrooms. In this paper, we study the behavior of the tests described in Section 3 under the same scenarios.

As in GRW, we consider grouping students – to simulation the practice of “tracking” – in four different ways. The first method of grouping students is random grouping (RG), which means there is no tracking. We then consider grouping students on the basis of their most recent test score (dynamic grouping, or DG), on the basis of their by their base (second-grade) test score (base grouping, or BG), and on the basis of their unobserved student heterogeneity (heterogeneity grouping, or HG). In the latter three cases noise is added to the grouping of

students to reflect the reality that, even with tracking, not all of the top students will be assigned to the same class.

We consider three ways of assigning teachers to classes: random assignment (RA), assignment where good teachers – based on their teacher effects – are assigned to better classes (positive assignment, or PA), and assignment where good teachers are assigned to worse classes (negative assignment, or NA). With random grouping of students there is only random assignment of teachers, but all three kinds of teacher assignments can be applied to the three different ways of tracking students. Therefore, in total there are 10 different grouping/assignment scenarios.

It is important to keep separate the issues of tracking and teacher assignment. As discussed in GRW, tracking by itself does not cause problems for VAM estimates. Even when the dynamics are misspecified in the regression analysis, most of the common estimators perform well in terms of ranking teachers. By contrast, several of the estimators perform poorly when teachers are nonrandomly assigned to groups of students.

To study the tests under dynamic misspecification we consider $\lambda = .5$ along with the baseline $\lambda = 1$. We also considered a scenario where the student heterogeneity is uncorrelated with the base score, and this has implications for some of the tests in certain scenarios. But we present in our tables the more realistic case where c_i and A_{i2} are positively correlated.

Given the discussion of the various specification tests in Section 3, we can predict the outcomes of the tests across different scenarios. It is important to remember that specification tests are intended to detect inconsistent parameter estimation and not to determine when various procedures may or may not do well ranking teachers. For example, as discussed in GRW, in some scenarios the VAM estimates are amplified, and this actually makes it easier to rank

teachers even though comparing the magnitudes of the estimated teacher effects could be misleading. Unfortunately, we cannot expect specification tests to distinguish between biases that help with ranking and those that hurt. Currently available tests are devised to detect inconsistent estimation of parameters.

Table 1 shows the predicted outcomes for the Hausman test based on the RE estimator. We assume that the common factor restriction holds. In constructing the tables it is useful to show how the tests would behave if we had an infinite amount of data. In other words, we do not worry about sampling error and the fact that with any particular sample size we can always make Type I or Type II errors. Thus, the entries in the tables are “Reject” or “Accept.”

Consider first the case $\lambda = 1$ and random assignment. This is a clear-cut case where the Hausman test should not reject RE estimation in favor of FE estimation: assignment of teachers is exogenous with respect to the student heterogeneity c_i and strictly exogenous with respect to the idiosyncratic shocks e_{it} in the equation

$$\Delta A_{it} = \tau + E_{it}\beta_0 + c_i + e_{it}. \tag{18}$$

Therefore, no function of the history of teacher indicators should help to predict the gain score from grade $t - 1$ to t .

Table 1. Predicted Outcomes for Hausman Test

Grouping/Assignment Mechanism	Lambda = 1	Lambda < 1
RG/RA	ACCEPT	REJECT
DG/RA	REJECT	REJECT
DG/NRA	REJECT	REJECT
BG/RA	REJECT	REJECT
BG/NRA	REJECT	REJECT
HG/RA	REJECT	REJECT
HG/NRA	REJECT	REJECT

Unfortunately, the conclusion for RG/RA does not carry over to other scenarios with random assignment of teachers to classrooms. Consider the DG/RA case (still with $\lambda = 1$), where students are grouped together based on past test scores but the resulting classrooms are randomly assigned to teachers. As shown via simulation in GRW, the RE estimator works quite well in this case, producing rank correlations between the estimated and true teacher effects on the order of .90 across different parameter settings. Yet the Hausman test will reject because past teacher assignment contains information on the ability level of the student. For example, if students with above average previous test scores tend to be grouped together, having had a third-grade teacher with a high estimated VAM tells us that, on average, the student has higher ability. Therefore, in a fourth-grade gain score equation the third grade teacher assignment has some predictive power for the gain score because third-grade assignment is correlated with ability, c_i .

A similar mechanism comes into play in the other random assignment, nonrandom grouping mechanisms.

The situation is even worse when $\lambda < 1$. In this case, the Hausman test will reject even in the random grouping, random assignment case. Rejection occurs because when $\lambda < 1$ equation (18) effectively omits the lagged dependent variable. While it is true that under RA the assignment of a teacher in grade t does not depend on $A_{i,t-1}$, teacher assignment at time $t - 1$ is correlated with $A_{i,t-1}$ whenever teachers have an effect on achievement (which we assume here). Thus, $E_{i,t-1}$ is correlated with the error term at time t because the error term effectively includes a fraction of $A_{i,t-1}$. In other words, lagged teacher assignment helps to predict the gain score, conditional on current teacher assignment, because the lagged teacher assignment is correlated with $A_{i,t-1}$. It is important to understand that, unlike in the usual settings where the Hausman test is applied, in the current scenario neither the random or fixed effects estimator produces consistent estimates of the teacher VAMs. Rather, the RE and FE estimators have different (incorrect) probability limits.

Rejection in the RG/RA scenario with $\lambda < 1$ is unfortunate because, as shown in GRW, the RE estimator again fares well – in fact, it is scarcely worse than in the $\lambda = 1$ case.

Table 1 shows that dynamic misspecification ($\lambda < 1$) in any grouping/assignment combination results in rejection of the RE estimator. Again, such an outcome is unfortunate because the RE estimator does well in several (but not all) of these scenarios. It is not surprising that the Hausman test detects dynamic misspecification in the RE model, but when we couple this analysis with the findings in GRW we are left to conclude that the Hausman test for choosing between RE and FE is not very information for VAM applications.

Because the remaining $\lambda < 1$ scenarios combine various grouping/assignment mechanisms along with dynamic misspecification, we use simulations to obtain an idea of how often the Hausman test rejects.

We can create a similar table for the feedback (or falsification) test for three commonly used estimators: random effects, fixed effects, and dynamic OLS. (The entries for pooled OLS on the gain-score equation (18) are identical to RE.) It includes only the $\lambda = 1$ case. We should emphasize that Wooldridge (2010) applies the leads test only to the fixed effects estimator; not to POLS, RE, or DOLS. The test can provide useful information for RE in the sense that it can detect correlation between E_{it} and the two sources of error in (18). But the Hausman test for choosing between RE and FE can too, and it is usually applied to the RE estimator to see whether one should use FE. If the RE estimator is rejected based on the Hausman test, the leads test is applied to the FE estimator because FE relies on the strict exogeneity assumption.

As discussed in Section 3, the case for applying the feedback test to DOLS is a priori weak. It can only tell us whether the random-grouping/random-assignment scenario holds – not whether DOLS does a good job estimating the teacher VAMs. Because DOLS works well in many scenarios, the feedback test is likely to be very misleading.

To isolate the key problem with the feedback test, suppose that equations (1) and (2) hold with the common factor restriction and no student heterogeneity, so we can write

$$A_{it} = \tau + \lambda A_{i,t-1} + E_{it}\beta_0 + r_{it}, \quad (19)$$

where $\{r_{it}\}$ is unpredictable given past test scores and current and past inputs. This is the ideal setup for DOLS estimation of β_0 (and λ): the estimated teacher effects will be consistent and, under a homoskedasticity assumption on $\{r_{it}\}$, asymptotically efficient. This is true regardless of whether E_{it} is correlated with $A_{i,t-1}$; in fact, the main reason for including the lagged test scores

is to allow this kind of nonrandom assignment. It is exactly because E_{it} and $A_{i,t-1}$ are correlated that the lead teacher assignments likely will be significant when added to (19). Other versions of the test, such as Rothstein's (2010), potentially reject when assignment is based on the past two years of test scores. Generally, though, such nonrandom assignment mechanisms are easily handled by including, if needed, additional lagged test scores in equation (19). We include the DOLS estimator in this study because a version of the leads test has been applied by Rothstein (2010) and Harris, Sass, and Semykina (2010).

Table 2 contains the predicted outcomes if the leads test is applied to the three estimators. With random grouping of students and random assignment of teachers, none of the tests should reject – this is the first row of Table 2. Any deviation from random group or random assignment causes the feedback test to reject for RE and DOLS. Again, rejection is not surprising when the teacher assignment is nonrandom. For example, if teacher assignment is based on past test score, then next grade's teacher will predict the current gain score regardless of the estimation method. As with the Hausman test, the reason for rejection with random assignment but nonrandom grouping is more subtle. If, say, the students are grouped based on their unobserved heterogeneity, and the better teachers get the better students, then the estimated lead teacher effect is, on average, higher for the better students – and so is the student's gain score. It is the opposite for the worse teachers and lower performing students. Therefore, the estimated lead teacher effects are positively correlated with the students' gain scores.

As with the Hausman test, it is unfortunate that the leads test rejects both RE and DOLS in cases where they produce very reliable teacher VAMs. In Section 6 we include RE and DOLS in the simulations to see how often the tests actually reject in reasonable scenarios.

The rejection scenarios for the FE estimator are more subtle. Because FE removes a time-constant student effect, grouping on the basis of time-constant variables does not cause a rejection using the leads test – provided the assignment of teachers to classrooms does not depend on time-varying factors, such as a lagged test score. Therefore, when grouping of students is done using the base score or student heterogeneity, the FE test will not reject. It is precisely assignment based on time-constant factors that FE is intended to be robust against. So the leads test is informative for FE: it tests whether grouping or assignment (or both) are based on an omitted factor that varies over time.

Table 2. Predicted Outcomes for Feedback Test: Lambda = 1

Grouping/Assignment Mechanism	Random Effects	Fixed Effects	DOLS
RG/RA	ACCEPT	ACCEPT	ACCEPT
DG/RA	REJECT	REJECT	REJECT
DG/NRA	REJECT	REJECT	REJECT
BG/RA	REJECT	ACCEPT	REJECT
BG/NRA	REJECT	ACCEPT	REJECT
HG/RA	REJECT	ACCEPT	REJECT
HG/NRA	REJECT	ACCEPT	REJECT

5. The Simulation Design

In the tables and surrounding discussion in Section 4 we effectively assumed that we have an infinite amount of data. In practice, we will not always reject with certainty for entries labeled “REJECT.” Some estimation methods will control for more of the factors causing nonrandom assignment, in which case rejection rates will be lower. Also, when a test should not reject in theory it might due to poor finite-sample performance. To learn about the size and power of the tests it is very helpful to simulate the statistics in plausible scenarios to see how they perform.

Our simulation design closely follows that in GRW, although we restrict our attention to the case where students and teachers are randomly assigned to schools. (In GRW, where we evaluated the ability of VAM estimates to track the true teacher effects, we considered mechanisms where students and teachers sorted into schools. Such sorting had little effect on the rankings of the different estimators.) An important reason for following the GRW design is the GRW findings show which estimators work well across a variety of situations. We can compare those findings with the properties of the test statistics to determine when the tests provide useful information – and when they do not. Along with the rejection frequencies for the tests reported in Section 6, we also computed the statistics in GRW measuring how well the estimated VAMs mimic the true teacher effects. For space reasons, we do not report these in tables. We will draw mainly on the rank correlations between estimated and true teacher effects; the findings are very similar to those in GRW, and the reader is referred to that paper for a complete set of simulation results.

In generating the data, we assume that test scores are perfect reflections of the sum total of a child's learning (that is, no measurement error) and that they are on an interval scale that remains constant across grades. We assume that teacher effects are constant over time and that unobserved child-specific heterogeneity has a constant effect in each time period. We allow for unobserved time-varying shocks to the test scores, but we do not allow other time-varying factors (such as family effects correlated with teacher assignment). Also, we omit school effects and peer effects, and we assume that teachers have the same effect on each student in a class (so no interactions between students and teachers). We also assume a constant decay parameter, and we assume the shocks in the gain score equation are serially uncorrelated.

Our data represent three elementary grades per student in a hypothetical district. We can think of these as grades 3 through 5 over the course of three years, where we observe an initial second-grade test score. We create data sets that contain students nested within teachers nested within schools, with students followed over time. Our simple baseline data generating process (DGP) as follows:

$$\begin{aligned}
 A_{i3} &= \lambda A_{i2} + \beta_{i3} + c_i + e_{i3} \\
 A_{i4} &= \lambda A_{i3} + \beta_{i4} + c_i + e_{i4} \\
 A_{i5} &= \lambda A_{i4} + \beta_{i5} + c_i + e_{i5}
 \end{aligned}
 \tag{20}$$

where A_{i2} is a baseline score reflecting the subject-specific knowledge of child i entering third grade, λ is a time constant decay parameter, β_{it} is the teacher-specific contribution (the true teacher value-added effect), c_i is a time-invariant child-specific effect, and e_{it} is a random deviation for each student. Because we assume independence of e_{it} over time, we are maintaining the common factor restriction in the underlying cumulative effects model.

The random variables A_{i2} , β_{it} , c_i , and e_{it} are drawn from normal distributions with mean zero, where we adjust the standard deviations to allow different relative contributions to the scores. We choose the same second moments as in GRW; we refer the reader to that paper for a survey of the literature underlying our choices. Specifically, the standard deviation of the teacher effect is .25, while that of the student fixed effect is .5, and that of the random noise component is 1, each representing approximately 5, 19, and 76 percent of the total variance in gain scores, respectively. Also, the correlation between the time-invariant child-specific heterogeneity c_i and the baseline score A_{i2} is about .5.

Our data structure has the following characteristics that do not vary across simulation scenarios:

- 10 schools
- 3 grades (3rd, 4th, and 5th) of scores and teacher assignments, with a base score in 2nd grade
- 4 teachers per grade (thus 120 teachers overall)
- 20 students per classroom
- 4 cohorts of students
- No crossover of students to other schools

To create different scenarios, we vary certain key features: the grouping of students into classes, the assignment of classes of students to teachers, and the amount of decay in prior learning from one period to the next.

Given the different ways of grouping students, assigning teachers, and specifying the amount of decay, we have 20 different scenarios. (In Section 6 we briefly consider the possibility that the base test score is uncorrelated with the student effect.) Because we have many different

scenarios and estimators that require nontrivial computational effort in some cases, we limit ourselves to 100 replications per simulation.⁷

6. Simulation Results

We begin with the Hausman test for comparing the RE estimator to the FE estimator. For completeness, we also include the POLS estimator. As discussed in Section 3, in practice one should allow for general serial correlation and heteroskedasticity in the composite error term, and so we report findings for the test robust to cluster correlation at the student level. (The findings are similar when we use the nonrobust tests provided the nonrobust tests are asymptotically valid.)

Because the grouping mechanisms generate within-classroom correlation, we also, when it is mechanically possible, cluster at the school level. We cannot cluster at the classroom level because the students change classrooms over time. Besides, grouping different students into different classrooms in different grades generally creates correlation across all students at a school within a grade. Thus, if one is to cluster at a level higher than the individual student then the school level is natural. We use school-level clustering even though we only have 10 schools, which makes applying the asymptotic theory where the number of schools is large suspect. Nevertheless, some simulation studies have shown clustering with as few as 10 clusters can work reasonably well, and much better than doing nothing. We do not have enough schools to cluster when using the test that includes the average of all teacher indicators.

Table 3 contains the results for the Hausman test with $\lambda = 1$. The first panel considers clustering at the student level. Each scenario and estimator has two entries. The first row contains the rejection frequencies of the one degree-of-freedom test that includes the estimated

⁷ However, we have tested the sensitivity of our results to much higher numbers of replications and found no substantive difference in the results.

teacher effect across all three grades. The second row has the rejection rates for the test that includes the full set of teacher dummies.

We focus on the RE results because POLS seems to have more small sample bias. Under the RG/RA scenario, the one-df test rejects about 6% of the time. The full test somewhat over rejects (13% rejection rate). Somewhat surprisingly, clustering at the school level leads to a test with pretty good size (8%).

The remaining predictions from Section 4 are born out as well. The test that includes the estimated teacher effect almost always rejects in every other kind of grouping/assignment scenario. To see the practical problems this causes, consider the HG-RA scenario in Table 3. The test clustered at the student level rejects 100% of the time, and the lowest rejection rate is 60% (RE with clustering at the school level). This means that we would traditionally reject the RE estimator in favor of FE. Yet in this simulation the rank correlation of the estimated VAMs for the RE estimator is .85 compared with .63 for FE. In fact, of the four estimators – POLS, DOLS, RE, and FE – RE works the best.

As mentioned in Section 4, the rejection of RE using the Hausman test in the HG-RA scenario is essentially mechanical due to the fact that good students are grouped with other good students. But this grouping has no effect on the quality of RE as an estimator of teacher VAMs. We are forced to conclude that the Hausman test is very misleading in this case.

In other scenarios the situation is even worse. For example, in the HG-PA setting – where we fully expect RE to be rejected, and it is – the RE estimator actually does even better in ranking teachers. The rank correlation jumps to .91; because FE removes the heterogeneity when estimating the teacher effects, we expect the FE rank correlation to remain about the same, and it does to two decimal places.

As if things were not bad enough, against the one alternative where some versions of the Hausman test do not detect nonrandom assignment, HG-NA, the FE estimator outperforms the RE estimator (with rank correlations of .62 and .55, respectively). In other words, when we want to reject RE in favor of FE the Hausman test has the lowest power. To be fair, the version of the test that uses the estimated teacher effects has unit power when we do not cluster, but this is not the standard form of the Hausman test.

Table 4 contains the simulation rejection rates when $\lambda = 1/2$. With a handful of exceptions, the test rejects the RE estimator 100% of the time.

The situation is somewhat improved for the leads test in the sense that it has roughly size 5% for the fixed effects estimator under static assignment mechanisms and it detects dynamic forms of teacher assignment. Table 5 contains the rejection frequencies.

In the RG-RA case, the test has size roughly 5% for all estimators with the exception of pooled OLS, where the rejection rates are somewhat high. For FE under base score and heterogeneity grouping, the feedback test using the estimated lead teacher effect rejects between 4% and 9% of the time – reasonably close to a 5% significance level. Clustering by school causes some distortions, but the rates are acceptable with only 10 schools.

The FE estimator is strongly rejected when dynamic grouping is coupled with nonrandom teacher assignment – either positive or negative: the rejection rates are all 100%. This shows that the test works as it is supposed to in detecting a failure of strict exogeneity when using FE estimation. Under dynamic grouping but random teacher assignment, the test rejects about 16% of the time. What appears to be happening is that removing a student fixed effect largely, but not entirely, accounts for the grouping by past test scores.

Table 5 also shows that POLS and RE are strongly rejected in most scenarios – even though in some of these RE is the best estimator for ranking the teacher effects. We already discussed a similar situation for the Hausman test. DOLS is rejected much less often with base score grouping than are POLS and RE. One way to understand why this happens is that DOLS is “almost” controlling for the right variable that determines teacher assignment: the most recent test score rather than the base score. The feedback test applied to DOLS is strongly rejected in the HG case: grouping is based directly on a factor affecting test scores (the student effect) and controlling for the lagged test score is not sufficient when assignment is based on c_i . To see why applying the leads test to DOLS is problematical, we again turn to the rank correlations between the DOLS VAM estimates and the true teacher effects. In the HG-PA scenario the leads test rejects 100% of the time, yet the rank correlation is about .87. Therefore, DOLS is doing a good job of ranking teachers even though the falsification test virtually always rejects.

The feedback test also rejects DOLS 100% of the time in the DG-PA and DG-NA cases even though these are cases where DOLS does well in ranking the teachers (rank correlation .76). Interestingly, both POLS and RE do notably better, with both having rank correlations of about .90. Of course, the leads test strong rejects when they are used, too. If we relied on the rejection by this “falsification test” to determine whether the estimates are doing a good job estimating the teacher effects, we would be led badly astray: we would conclude none of the estimates can be trusted. In effect, Rothstein’s (2010) conclusion was that the VAM estimates could not be trusted because his version of the falsification test always rejected. Our simulations show that this test actually has very little to say about when POLS, RE, and DOLS are working well or not.

7. Concluding Remarks

In this paper we have discussed two specification tests that are applied in the literature. The first test is a robust, regression-based version of the Hausman test that compares the random effects and fixed effects estimators. The second test is a feedback or leads test that was originally designed to test for violation of the strict exogeneity assumption in the context of fixed effects estimation. Versions of this test were used in an influential paper by Rothstein (2010) to detect nonrandom teacher assignment in the context of several regression equations, including dynamic equations.

The most important point of this paper is that neither the Hausman test nor the feedback test is very helpful for choosing among estimators or in determining whether a particular estimation method is providing good estimates of teacher VAMs. The Hausman test rejects RE in favor of FE in many cases where the RE estimator is clearly superior for ranking teachers based on estimated VAMs. The source of the problem with the Hausman test is that nonrandom grouping of students – often called tracking – leads to rejection even though teacher assignment to classrooms is random. Under random teacher assignment, RE does well for estimating teacher value added.

The feedback test is a little more successful, but only when it is applied to the fixed effects estimator – the original application of the test described in Wooldridge (2010). The test has good size properties under static assignment mechanisms and detects dynamic assignment – which has deleterious effects on the FE VAM estimates – with high probability. Nevertheless, we must emphasize that the falsification test applied to pooled OLS, random effects, and dynamic regression produces misleading results. Often the test rejects even though the estimation

method is working well. Conversely, sometimes the test fails to reject when the estimated VAMs are poor.

The findings in this paper can be combined with those in GRW to provide some practical advice to those wanting to estimate teacher VAMs. GRW found that, generally, dynamic regression methods provide the best and most robust estimates – although there are notable exceptions, such as RE estimation with random teacher assignment and correctly specified dynamics. The current paper shows that applying a falsification test to dynamic regression – whether it is the simple form studied here, with just a single lagged score, or more sophisticated methods with multiple lags – is a poor idea. A rejection has very little to do with whether dynamic regression produces good VAM estimates. A similar comment holds for RE, whether one applies the Hausman test or the falsification test: the outcome of the test is practically useless for the main aim of estimating VAMs.

The inappropriateness of applying the falsification test to dynamic regression methods for estimating VAMs can be further understood by viewing dynamic regression through the lens of estimating average treatment effects – where being assigned a particular teacher is the “treatment” – rather than thinking of dynamic regression as estimating a structural cumulative effects model. From the modern treatment effects perspective, controlling for lagged test scores, and perhaps other observables, is intended to make teacher assignment random conditional on the observables. This “unconfoundedness of treatment assignment” assumption is at the heart of regression, propensity score, and matching methods for estimating treatment effects; see, for example, Imbens and Wooldridge (2009). As is well known, the unconfoundedness assumption is not testable: it exactly identifies the teacher effects. Moreover, the treatment is assumed, or at least allowed, to be correlated with the conditioning variables – usually the past test scores.

Generally, testing whether teacher assignment is correlated with past test scores is not a test of the unconfoundedness assumption unless some strong assumptions are imposed about the nature of any nonrandom assignment. Imbens and Wooldridge (2009) discuss a set of sufficient conditions, but the spirit of them can be easily described. In effect, in order to construct a falsification test one must assume that unconfoundedness holds conditional on a short history of test scores, with more lags excluded from the conditioning set. We see no reason to think such assumptions are plausible when trying to estimate teacher effectiveness.

Given the frailty of the cumulative effects model as a description of educational production, viewing dynamic regression methods as flexible ways to estimate VAMs – without worrying about “structural” parameters – appears to be the most promising way forward.

References

- Arellano, M., Bond, S. (1991). Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. *The Review of Economic Studies* 58(2), 277-297.
- Chamberlain, G., "Panel Data," in *Handbook of Econometrics*, Vol. II, Z. Griliches and M. D. Intriligator, eds., Amsterdam: Elsevier North-Holland (1984)
- Guarino, C., Reckase, M., & Wooldridge, J.M. (forthcoming). Can Value-Added Measures of Teacher Performance be Trusted?
- Guggenberger, P. (2010). The Impact of a Hausman Pretest on the Size of a Hypothesis Test: The Panel Data Case. *Journal of Econometrics* 156(2), 337-343
- Harris, D., Sass, T., Semykina, A. (2010). Value-Added Models and the Measurement of Teacher Productivity. *Working Paper*. Downloadable at <http://www.urban.org/url.cfm?ID=1001508>
- Imbens, G.W. and J.M. Wooldridge (2009), Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature* 47, 5-86.
- Kinsler, J. (2012), Assessing Rothstein's Critique of Teacher Value-Added Models. *Quantitative Economics* 3, 333-362.
- Koedel, C., Betts, J. (2009). Does Student Sorting Invalidate Value-Added Models of Teacher Effectiveness? An Extended Analysis of the Rothstein Critique. *Working Paper*. Downloadable at <http://econpapers.repec.org/paper/umcwpaper/0902.htm>.
- McClain, K.T., Wooldridge J.M., (1995) A Simple Test of the Consistency of Dynamic Linear Regression in Rational Distributed Lag Models. *Economic Letters* 48(3), 235-240.
- Rothstein, J., (2010). Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement. *Quarterly Journal of Economics* 125(1), 175-214.
- Todd, P., Wolpin, K. (2003). On the Specification and Estimation of the Production Function for Cognitive Achievement. *The Economic Journal*, 113 (2003), F3-F33.
- Wooldridge, J.M. (1990). A Unified Approach to Robust, Regression-Based Specification Tests, *Econometric Theory* 6, 17-43
- Wooldridge, J.M. (2009), "Correlated Random Effects Models with Unbalanced Panels," draft, Michigan State University Department of Economics.
- Wooldridge, J.M. (2010), *Econometric Analysis of Cross Section and Panel Data*, second edition. Cambridge, MA: MIT Press.

Table 3: Hausman Test Rejection Rates. Results from 100 replications. Vertically scaled test scores. $\lambda=1$. Correlation of student fixed effect with scorebase is .5. Row 1: Rejection rate of test with estimated mean teacher effect. Row 2: Rejection rate of test with mean teacher indicators.

Hausman Test $\lambda=1$	Cluster at Student Level		Cluster at School Level	
	POLS	RE	POLS	RE
Assignment Mechanism				
RG-RA	0.19 0.12	0.06 0.13	0.18	0.08
DG-RA	1 1	0.95 1	0.99	0.67
DG-PA	1 1	1 1	1	1
DG-NA	1 0.91	1 0.6	0.75	0.83
BG-RA	0.99 1	0.99 1	0.71	0.45
BG-PA	1 0.25	1 0.89	1	1
BG-NA	1 0.99	1 0.83	0.07	0.78
HG-RA	1 1	1 1	0.86	0.6
HG-PA	1 0.76	1 0.67	1	1
HG-NA	1 0.19	1 0.06	0.24	0.27

Table 4: Hausman Test Rejection Rates. Results from 100 replications. Vertically scaled test scores. $\lambda=.5$. Correlation of student fixed effect with scorebase is .5. Row 1: Rejection rate of test with estimated mean teacher effect. Row 2: Rejection rate of test with mean teacher indicators.

Hausman Test $\lambda=.5$	Cluster at Student Level		Cluster at School Level	
	POLS	RE	POLS	RE
Assignment Mechanism				
RG-RA	1 0.94	1 0.94	1	1
DG-RA	1 1	1 1	1	1
DG-PA	1 1	1 1	1	1
DG-NA	1 1	1 1	1	1
BG-RA	1 1	1 1	1	1
BG-PA	1 1	1 1	1	1
BG-NA	0.53 1	0.53 1	0.32	0.41
HG-RA	1 0.98	1 0.98	1	1
HG-PA	0.63 0.97	0.63 0.97	0.47	0.53
HG-NA	1 1	1 1	1	1

Table 5: Leads Test Rejection Rates. Results from 100 replications. Vertically scaled test scores. $\lambda=1$. Correlation of student fixed effect with scorebase is .5. Row 1: Rejection rate of test with estimated lead teacher effect. Row 2: Rejection rate of test with future teacher indicators.

Leads Test $\lambda=1$	Cluster at Student Level				Cluster at School Level				
	Estimator	POLS	DOLS	RE	FE	POLS	DOLS	RE	FE
Assignment Mechanism									
RG-RA	0.1	0.03	0.01	0.07	0.15	0.03	0.05	0.07	
	0.05	0.04	0.05	0.03					
DG-RA	1	0.36	0.95	0.16	0.97	0.02	0.57	0.16	
	1	1	1	0.18					
DG-PA	1	1	1	1	1	1	1	1	
	1	1	1	1					
DG-NA	1	1	1	1	0.74	1	0.8	1	
	1	1	1	1					
BG-RA	0.65	0.08	0.35	0.09	0.55	0.03	0.31	0.06	
	0.68	0.09	0.61	0.09					
BG-PA	1	0.14	1	0.06	1	0.15	1	0.04	
	0.97	0.11	0.97	0.1					
BG-NA	0.13	0.22	0.6	0.06	0.06	0.22	0.49	0.13	
	0.97	0.1	0.97	0.1					
HG-RA	0.86	0.36	0.54	0.04	0.7	0.25	0.43	0.1	
	0.94	0.73	0.88	0.07					
HG-PA	1	1	1	0.06	1	1	1	0.06	
	1	1	1	0.09					
HG-NA	0.59	0.76	0.56	0.08	0.19	0.51	0.2	0.06	
	1	1	1	0.09					

Table 6: Leads Test Rejection Rates. Results from 100 replications. Vertically scaled test scores. $\lambda=.5$. Correlation of student fixed effect with scorebase is .5. Row 1: Rejection rate of test with estimated lead teacher effect. Row 2: Rejection rate of test with future teacher indicators.

Leads Test $\lambda=.5$	Cluster at Student Level				Cluster at School Level				
	Estimator	POLS	DOLS	RE	FE	POLS	DOLS	RE	FE
Assignment Mechanism									
RG-RA	0.13	0.03	0.13	0.06	0.11	0.02	0.16	0.09	
	0.02	0.04	0.02	0.03					
DG-RA	1	0.38	1	0.62	0.99	0.02	0.99	0.55	
	1	1	1	0.38					
DG-PA	1	1	1	1	1	1	1	1	
	1	1	1	1					
DG-NA	1	1	1	1	1	1	1	1	
	1	1	1	1					
BG-RA	0.1	0.09	0.1	0.09	0.11	0.03	0.14	0.05	
	0.52	0.08	0.52	0.13					
BG-PA	0.98	0.28	0.98	0.07	0.88	0.26	0.95	0.07	
	0.9	0.15	0.91	0.14					
BG-NA	0.78	0.32	0.78	0.17	0.65	0.25	0.71	0.14	
	0.78	0.23	0.78	0.14					
HG-RA	0.1	0.35	0.1	0.04	0.06	0.25	0.1	0.06	
	0.23	0.73	0.23	0.05					
HG-PA	0.71	1	0.71	0.12	0.56	1	0.65	0.15	
	0.7	1	0.71	0.12					
HG-NA	0.99	0.63	0.99	0.07	0.97	0.32	0.97	0.12	
	0.76	1	0.77	0.12					