

When Incentives Matter Too Much: Explaining Significant Responses to Irrelevant Information

Tom Ahn
Naval Postgraduate School

Jacob L. Vigdor
University of Washington

When agents observe a continuous variable and a discrete signal based on that variable, theory suggests that the signal should not impact behavior conditional on the variable. Numerous empirical studies, many based on regression discontinuity design, contradict this basic prediction. We propose two rationalizations with testable implications. One is based on information acquisition costs and the other on learning and imperfect information. Using education data from North Carolina and exploiting a pay-for-performance system, we find support for the model of learning. This implies that rational responses to policy interventions may emerge gradually, and evaluations with short-term data may understate treatment effects.

I. Introduction

Traditional economic theory suggests that rational agents should exhibit no reaction to the introduction of irrelevant or redundant information. In scenarios where agents have access to a quantitative information set, supplementing the data with discrete signals derived directly from the same data should have no impact. For example, when a publicly available, continuous rating of product quality exists, the introduction of discrete categories based solely on the continuous measure should not alter rational consumer decisions.

A growing number of empirical analyses identify scenarios where this basic principle fails in practice. Numerous studies document economic

We are grateful to John Holbein for research assistance. Any opinions expressed herein are those of the authors and not any affiliated organization.

Electronically Published October 25, 2021.

Journal of Human Capital, volume 15, number 4, winter 2021.
© 2021 The University of Chicago. All rights reserved. Published by The University of Chicago Press.
<https://doi.org/10.1086/716785>

agents reacting significantly to discrete categorizations based on an underlying continuous variable even when the variable is in their information set.¹ School performance is also commonly measured using continuous variables, such as proficiency rates, but summarized with letter grades or binary designations.² In many cases, the failure of some agents to react rationally to irrelevant information creates profit opportunities for others. Previous literature has identified numerous cases of overreaction to irrelevant information, but explanations for this behavior and empirical tests of these explanations are still uncommon.

The phenomenon is particularly germane in studies employing regression discontinuity (RD) methodology, applied in scenarios where subjects are assigned to discrete categories from an underlying continuous measure. These studies often document significant reactions to discrete categorizations when the underlying difference in a continuous measure is trivial, such as when a car's odometer rolls from 9,999 to 10,000 miles. The mounting empirical evidence motivates a question: Is there a simple model of behavior that can explain significant responses to signals that should be irrelevant?

There are many important theoretical models of seemingly irrational behavior (Mullainathan 2002; Sarver 2008; for a review of the psychology literature, see Gigerenzer and Goldstein 1996). Informed by this literature, this paper presents two candidate explanations for significant responses to irrelevant information in a principal-agent style production framework with output incentives. It then uses data on school responses to an accountability program in North Carolina to evaluate their empirical implications.

The first model focuses on information costs, presuming that agents are "sparse maximizers" in the sense of Gabaix (2014), using the signal as a proxy for the underlying continuous variable. Intuitively, agents pay no attention to the underlying continuous variable until a discrete indicator informs them that something is wrong, much as a driver might refuel their vehicle only when the low fuel indicator illuminates.

The second posits that agents are Bayesian updaters who act rationally when fully informed. With imperfect information, however, incomplete understanding of the production process may lead them to react on the basis of the signal (Sims 2003). Here, agents pay attention to the continuous measure but are not quite sure whether actions they take will have a beneficial or detrimental impact. The discrete signal, in certain circumstances, motivates them to take a risky action. These two models yield divergent empirical predictions. Most importantly, incomplete information

¹ Examples include diners' response to sanitation grades, homeowners' electricity use based on grades from utility companies, and car buyers paying higher prices for small differences in odometer readings (see Jin and Leslie 2003; Card, Mas, and Rothstein 2008; Allcott 2011; Anderson and Magruder 2011; Berger and Pope 2011; Lacetera, Pope, and Sydnor 2012).

² Evidence from Figlio and Lucas (2004), Figlio and Kenny (2009), Martinez (2010), Papay, Murnane, and Willett (2011), and Ahn (2014) shows that teachers, parents, and students all respond to discrete grades in the education context.

problems should dissipate over time as agents learn the relationship between their actions and consequences. Sparse maximizers, by contrast, should exhibit persistent behavior.

In the empirical case examined here, personnel at public schools in North Carolina were provided with annual reports of school effectiveness, with effectiveness measured as a continuous variable. When this continuous variable exceeded certain predefined thresholds, the personnel were awarded salary bonuses. From a rational perspective, presuming that the effectiveness measure is stationary or subject to a smooth drift process, receipt of the bonus yields no additional information regarding the likelihood of being awarded a bonus in the future. We would expect rational, maximizing school personnel to behave similarly regardless of whether they barely missed or barely made the standard for receiving a bonus.

In contrast to this prediction, we show that schools below the threshold exhibited significant improvements in performance in following years relative to schools just above the threshold.³ The discontinuity is greatest among schools with relatively inexperienced principals and schools with a record of inconsistent performance. These patterns are more consistent with the learning model than with the information cost model.

In addition to offering insight into the nature of seemingly irrational responses to irrelevant information, this paper contributes to the large literature on educational accountability. Accountability systems are designed to improve student performance by introducing information, incentives, or market pressures that may be absent in public education systems.

Robust evidence exists from both US and international implementations of merit pay.⁴ Studies demonstrate small but statistically significant gains in math standardized test scores in elementary schools when merit pay systems induce teachers to higher effort. School-level incentives yield smaller effects compared with individual bonuses. Unsurprisingly, larger rewards are associated with larger effects. Studies find effects as large as 4% of a standard deviation increase in math scores when the incentive is equal to 10% of per capita income.⁵

Numerous empirical studies assessing the impact of accountability policies take advantage of sharp thresholds for qualifications in many systems, where the treatment is unanticipated.⁶ The implicit or explicit

³ One other study has confirmed these basic results (see Jinnai 2016).

⁴ In the United States, at least 40 states or localities have or had some type of incentive program. Internationally, well-known programs existed in China, England, India, Israel, Kenya, and Portugal.

⁵ For the United States, see Koretz and Barron (1998) for Kentucky, Fryer et al. (2012) for Chicago, Ahn (2013, 2017) for North Carolina, Ladd (1999) and Imberman and Lovenheim (2015) for Texas, Dee and Keys (2004) for Tennessee, and Winters et al. (2008) for Arkansas. Internationally, see Lavy (2002, 2009), Atkinson et al. (2009), Glewwe, Ilias, and Kremer (2010), Muralidharan and Sundararaman (2011), and Loyalka et al. (2019). See Neal (2011) and Pham, Nguyen, and Springer (2020) for a comprehensive review.

⁶ See Hanushek and Raymond (2005), West and Peterson (2006), Chiang (2009), Rockoff and Turner (2010), Chakrabarti (2011), and Craig, Imberman, and Perdue (2015) for RD implementations.

models of behavior underlying empirical analyses are usually one-shot games with fully informed agents, allowing little scope for effects of incentive programs to evolve as agents learn more about them.⁷ While some research points to principals and teachers learning through trial and error and altering behavior, in many cases these are considered as threats to the validity of estimated treatment effects and controlled for or minimized using econometric techniques.⁸

This study adds to the evidence that accountability systems can be useful in inducing performance gains while emphasizing that agents may exhibit dynamic learning and respond rationally using the limited (and sometimes erroneous) information at hand.

The evidence presented carries implications for the design of incentive systems in education and other domains. Principal-agent models usually presume that agents fully understand the incentives and the process by which their own efforts translate into outcomes. Our results suggest that agents may require several iterations to attain this mastery. As such, a one-size-fits-all incentive will have different effects on agents with varying amounts of experience, implying that short-term evaluations of incentive programs may misstate the long-run implications. Our results also suggest that efforts to tweak incentive programs to better align incentives need to be weighed against the costs of disrupting the learning process.

II. Conceptual Framework

As noted above, there have been numerous empirical exercises refuting traditional rational models in making decisions based on information when both a continuous measure and a more aggregated grade based on this continuous measure (and thus extraneous information) are available. Nonetheless, to frame our discussion of alternative models, we begin by outlining a basic rational framework based on a principal-agent model. In this scenario, agents who find themselves on either side of a discontinuous incentive threshold face near-equivalent incentives to adjust their behavior.

A. *Traditional Model*

Suppose that output y_{it} of employee i in period t , which in the context of educational production might be measured by improvements in student test scores, is a known function of a vector of inputs chosen by the employee,

⁷ Hemelt (2011), Ahn (2016), Macartney (2016), and Ahn and Vigdor (2020) are some exceptions.

⁸ Evidence on efficacy of principals is sparse. Some studies find that experience and tenure at the school and managerial practices of principals matter (see Clark, Martorell, and Rockoff 2009; Béteille, Kalogrides, and Loeb 2012; Miller 2013; Bloom et al. 2015; Grissom and Bartanen 2019; Ahn and Vigdor 2020).

x_{it} , which might generalize to effort in a simple model but might be more realistically thought of as an array of possible uses of time. A set of parameters governs the translation of these input choices into output, but the final value incorporates an idiosyncratic shock ϵ_{it} . In the most basic formulation, production function parameters are fixed and known; the shock may exhibit serial correlation. The employee's utility is a function of their wage, w_{it} , and a cost function based on input choices, $c_i(x_{it})$, which we take to be increasing and convex. We also allow for the possibility that $c_i(x_{it})$ may be less than zero for certain values of x_{it} , which would be the case if employees received some satisfaction or pride from turning in a certain level of effort even in the absence of monetary reward. The subscript also indicates that there may be permanent differences across teachers in the valuation of effort. The employee observes y_{it} and can thus determine the value of ϵ_{it} ex post.

To incentivize effort, the employer links compensation to the observed indicator of output, $w_{it}(y_{it})$. In this scenario, the employee's optimal choice of effort equates the expected marginal cost and benefit. The anticipated effect of the incentive scheme on effort thus depends on the strength of the relationship between output and effort and on the strength of the relationship between output and the wage.⁹

Consider the special case when the incentive payment is binary: w_{it} is incremented by some positive amount when output rises above a critical threshold. This case corresponds to many incentive pay programs for teachers, including the North Carolina program studied here. The expected marginal benefit to effort then reduces to the marginal impact of effort on the probability of pushing the output indicator above the critical value.

Now, consider a pool of identical employees who have optimally chosen effort according to the same rules. Any variation in compensation across these teachers reflects variation in ϵ_{it} . Under a variety of assumptions regarding the evolution of ϵ_{it} , we should expect no discontinuous change in the optimal effort choice as a function of incentive receipt. Were ϵ_{it} entirely uncorrelated across years, rational employees would behave exactly the same in the subsequent period. If $E(\epsilon_{it})$ evolves as a random walk, agents on either side of an arbitrary threshold face nearly equal incentives to adjust in response to a realization. Even under nonrandom evolution of ϵ_{it} , so long as $E(\epsilon_{it+1}|\epsilon_{it})$ does not exhibit a discontinuity precisely at the threshold distinguishing incentive recipients from nonrecipients, there is no reason to expect discontinuous behavior changes conditional on the value of ϵ_{it} . The most plausible scenario involving a discontinuous behavior change at the point of discontinuity would occur if there were no

⁹ In the case of teaching, a less stylized model would relax the assumption of a single-dimensional effort input; the actions taken to educate a student most can in fact vary along many dimensions.

year-to-year variation whatsoever in ϵ_{it} . In that scenario, the idiosyncratic determinant of output would be better described as an element of an employee's ability that is uncertain until the first period of employment.¹⁰ In the context of educational production, it is unrealistic to think of an outcome such as a class-level average change in standardized test scores as perfectly predictable conditional on information regarding agent effort choices. As such, we dismiss this scenario.

B. Introducing Information Costs and Selective Reoptimization

To rationalize the existence of discontinuous responses to signals based on available information, we must introduce some additional complication to the basic model. The most obvious extensions would involve departures from complete information. A simple extension would introduce costs to the agent of observing the realization of ϵ_{it} and thereby obtaining information on the distribution of future realizations. An alternate extension would be to introduce uncertainty and time variation in the parameters translating input choices into output, with the agent required to incur some cost in order to update information on parameter values. The notion of information costs underlies other theoretical models of boundedly rational behavior (e.g., Gabaix 2014).

To generate a prediction of asymmetric response at the incentive threshold, the model would need to yield a decision rule that involved paying the cost to realize the information only conditional on receiving a signal. Plausibly, if there is significant serial correlation in the value of ϵ_{it} or the parameters linking input choices to output, agents might find it optimal to pay the information cost only in the event of an unexpected shock, such as the failure to receive a discrete bonus payment after a steady period of receiving it, as the investment in the information cost would yield an expected return in the form of resolving uncertainty regarding the potential return to altering their vector of input choices.¹¹ If the agent were to learn that their performance lay significantly close to the threshold, they might rationally choose to increase it.

In this scenario, one would predict a significant improvement in next-period performance among agents who barely missed the performance threshold relative to those who barely made it, because only the former

¹⁰ This basic scenario can be straightforwardly translated to simple models of consumer choice. When selecting a product, such as a used car, the quality of the product may vary monotonically as a function of a continuous measure, such as mileage, but in most cases, there is no reason for quality inferences to vary discontinuously at any arbitrary point in the distribution. The exception would be in a lemons-type scenario, where potential suppliers of a good selectively choose to sell because of a known proclivity among buyers. For example, potential sellers of cars with an odometer reading of 10,000 miles might rationally withhold cars from the market, given consumer discounting.

¹¹ The intuition here is similar to that of the finite adjustment cost or $[S, s]$ model (Bertola and Caballero 1990).

group would engage in the reoptimization made possible by incurring the information cost.

Importantly, as this is a scenario where agents face real costs to obtain information, the model would not predict that more experienced agents exhibit any kind of differential response pattern. Conditional on paying the cost, all agents have full information. As such, while school-level resources or distribution of student or teacher ability may impact test scores, the school principal's knowledge gained through prior education or experience should not impact the school's ability to generate test score gains. One might expect that agents might gain better intuition for what is happening with state variables over time, which could be modeled as gradual reductions in information costs. Learning is the focus of our second candidate explanation.

C. Introducing Incomplete Information regarding the Production Process

A second variation on the imperfect information theme regards uncertainty in the nature of the production process itself. Here we consider a model of uncertainty and Bayesian updating closely related to a category of problems in the bandit framework (e.g., Rothschild 1974; Berry and Fristedt 1985; Easley and Kiefer 1988; El-Gamal and Sundaram 1993; Bala and Goyal 1995). Generally, the problem involves a long-lived agent with a set of prior beliefs regarding the state of the world whose actions result in some observed outcome that is used to derive a new posterior set of beliefs, which in turn becomes the new prior guiding the next round of action. We consider agents that begin their careers or their experiences under the incentive regime, uncertain of exactly how their input choices translate into output (parameters) and uncertain of the role of exogenous factors in production (variance and other properties of ϵ_{it}). This arguably characterizes the status of new entrants to a profession such as teaching (or managing teachers).

After the first production period, agents receive information on a continuous measure of their output. This single data point is insufficient to identify the unknown parameters of the education production function and variance of the error term. They also observe a binary indicator of whether they met a production threshold and receive the cash bonus if they crossed the threshold. Agents receiving different signals—particularly inexperienced agents whose prior beliefs lean toward stationary and deterministic rather than stochastic production functions—face differing incentives to alter their effort allocation in the next period. Such agents would infer that there are minimal incentives to alter their input choices if they barely qualify for the bonus, given the probability that altering input would reduce rather than increase output. On the other hand, the impetus to alter effort allocations is strongest among those who receive the initial signal that their output was insufficient, as they are inclined to believe that they will continue to fail to cross the threshold

in subsequent periods without adjusting their inputs. As agents gain a better appreciation for the role of stochastic factors in the production process, these discontinuous responses should dissipate.

In subsequent periods, agents who maintain the same input choices gain information on the variance and other properties (e.g., serial correlation) of the error term but cannot update their prior beliefs on parameter values. Agents who alter their input choices have some capacity to update their prior beliefs on parameter values, limited by continued draws from the error distribution. For any equilibrium where the first-period decision rule is to alter input choices only following the receipt of a negative discrete signal, decisions at intermediate experience levels may depend on the agent's history. Agents receiving their first negative signal after a string of positive signals have relatively accurate beliefs regarding the properties of the error term but limited information on how adjustments to their inputs might raise output. We might expect agents on either side of the signal threshold to be similarly motivated to make adjustments to their input choices but similarly limited in their efficacy.

Agents who receive multiple negative signals will gain information on the relationship between input choices and output while maintaining fairly uninformed beliefs on the role of error. If their prior beliefs ascribe a negligible role to stochastic determinants of output, negative signals in the present period may continue to motivate them to make further adjustments to input choices even when output falls arbitrarily close to the threshold. In certain scenarios, we then expect agents with a history of past negative signals to respond more strongly and more effectively to a negative signal in the current period.

In the long run, agents converge to full information regarding the nature of the production process, at which point the model simplifies to the rational fully informed agent model above. This presents a contrast with the information cost story, as presented above. Discontinuous responses are expected to be concentrated among relatively inexperienced agents.

The learning model implies that the impact of an incentivization policy may take time to emerge. Policy evaluations based on short-term implementation may not reveal the long-run impact obtained when agents have more accurate beliefs regarding how their input choices translate into output.

Additionally, if the goal of an incentivization policy is to increase output, the learning model illustrates a clear trade-off between tweaking how the model operates and allowing agents to learn how to respond optimally to it. A tweak in this circumstance might involve an action as simple as changing the threshold for receiving a positive signal, or it might involve a more complicated reform, such as altering how output is measured. In the context of education policy, reforms could include basing incentives on a different standardized test, a different school- or classroom-level summary statistic, or moving away from incentives based on standardized tests entirely.

III. Empirical Application: The North Carolina ABC Program

Beginning in the 1996–97 school year, the state of North Carolina implemented the ABCs of Public Education accountability plan, which introduced a system of cash bonuses awarded to all teachers in schools meeting test score–based performance goals. Initially, the bonus amount was set to \$1,000 per teacher, but after 1 year the state switched to a two-tiered bonus structure, with payment amounts of \$750 and \$1,500. The performance measure used to assess schools, the *composite growth index*, was based on year-over-year changes in test scores for enrolled students, which makes the program distinct from the federal No Child Left Behind (NCLB) program or other incentive schemes based purely on proficiency rates. The formula for computing the performance measure changed after the 2004–5 school year; our analysis below focuses on the measure in place during the more recent period.

Details regarding the computation of the performance measure can be found in Vigdor (2009). Importantly, a bonus of \$750 per teacher was awarded if the school’s measure exceeded a predetermined threshold, and a \$1,500 bonus was awarded in schools where the measure exceeded a second, higher threshold. This implies that the effect of being awarded a bonus (or of failing to receive a bonus) can be estimated with an RD design.

Figure 1, reprinted from Vigdor (2009), shows the proportion of schools eligible for bonus payments from the inception of the program through 2006–7. Between 50% and 90% of schools were eligible for at least some bonus payment in every year, while the proportion eligible for the full \$1,500 bonus varied between 10% and 70%.

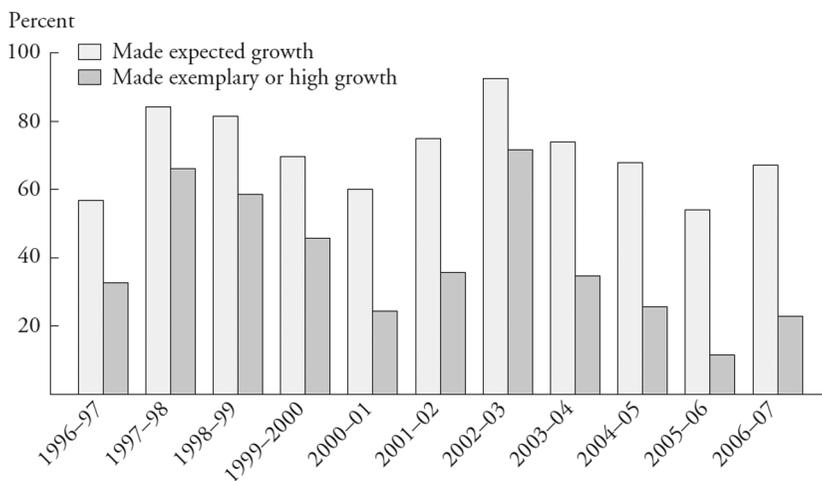


Figure 1.—Proportion of schools qualifying for North Carolina bonus. Expected-growth reward is \$750 per teacher, and high growth reward is \$1,500 (from Vigdor 2009).

TABLE 1
 AYP AND ABC STATUS

AYP	ABC	
	Yes	No
Yes	956	284
No	423	635

Note.—The table shows AYP for NCLB and expected-growth bonus receipt for the ABC program. Yes = qualified under accountability regime; no = failed to qualify under accountability regime.

From the 2002–3 school year forward, the NCLB program imposed a simultaneous but distinct set of requirements and sanctions upon public schools in North Carolina. Because these sanctions were based on student proficiency rates and not test score growth, the correlation between qualifying for positive sanctions—bonus receipt in the state system, adequate yearly progress (AYP) in NCLB—is modest. Table 1 shows a cross tabulation of AYP status and bonus receipt for school years 2005–6 and 2006–7. Over 40% of schools qualify for some bonus payment even though they have failed to make AYP, and about 30% receive no bonus in spite of the fact they have made AYP.

It is important to emphasize that there is no direct connection between a school's performance in year $t - 1$ and the stakes for making or missing the bonus threshold in year t . The substantial fluctuation in the proportion of schools receiving bonus payments from year to year underscores the importance of noise in the evolution of school performance measures over time. Schools on either side of the bonus threshold in year $t - 1$ should have derived little or no information regarding their prospects for receiving a bonus in year t , particularly after conditioning on the composite growth index.

IV. Data and Methods

A. Data

We use individual-level test score data provided by the North Carolina Education Research Data Center (NCERDC) to analyze differences in student performance on either side of the bonus discontinuity. The NCERDC data provide longitudinal links for students in grades 3–8 based on standardized test score records. We use these records to compute individual-level gain scores, which when aggregated to the school level yield our measure of output. We also observe a range of demographic and socioeconomic indicators at the individual level, including race, gender, limited English proficiency, and free/reduced-price lunch participation. Table 2 presents summary statistics for our sample, which consists of students enrolled in schools serving grades 3–5 in the 2005–6 and

TABLE 2
SUMMARY STATISTICS

Variable	Mean (SD)
Δ math score	.1385 (.6007)
Δ reading score	-.0544 (.6205)
Math proficiency level	2.8763 (.8399)
Reading proficiency level	3.326 (.7676)
Minority	.3803 (.4855)
Eligible for free/reduced-price lunch	.4677 (.4990)
Female	.4948 (.5000)
Limited English proficient	.0642 (.2450)
Years since last bonus	.6663 (.9780)
Number of no-bonus years in past 5 years	1.1577 (1.1944)
Years since AYP made	.6692 (1.0012)
Number of AYP failed since 2002–3	1.2075 (1.1467)
School size	211.5 (113.2)
Principal years of experience	6.1873 (3.8569)
Principal years of tenure at current school	4.4226 (3.1135)
Observations	338,240

Note.—The table shows NCERDC data for elementary schools and students from 2005–6 and 2006–7. Math and reading scores are c-scores. A student is proficient in a subject with a level 3 or 4. Minority students are black, Hispanic, or American Indian.

2006–7 school years.¹² North Carolina is a racially and socioeconomically heterogeneous state, with a rapidly growing immigrant population and a mix of prosperous metropolitan areas and poorer rural and inner city regions. The data set contains roughly 340,000 student/year observations in 2,248 elementary schools/years.

The math and reading gain scores are computed by subtracting a student's prior year standardized math or reading score from his or her prior year's standardized score in the same subject. Besides a continuous score, elementary school students in North Carolina are placed into one of four proficiency levels for reading and mathematics, with level 3 indicating sufficient mastery of the subject, which equates to grade-level proficiency. On average, students in North Carolina are slightly below grade-level proficiency for math and slightly above for reading.

To proxy for the number of opportunities available to the school leadership to update its priors on noise in the education production process and/or the efficacy of inputs, we use years of experience of the principal. We use payroll data for teachers and administrators to track the experience level of the principals. Payroll data are available from 1992, which allows us to count up to 13 years of experience for a principal in the 2005–6 academic year. Therefore, years of experience for principals suffers from right censoring. However, principals with 13 or more years of experience comprise less than 10% of the sample.

¹² The set of schools that are considered are schools capped at grade 5. Schools that contain both middle school grades (grades 6, 7, and/or 8) and elementary school grades are excluded from the analysis. Because students in these upper grades may move across classes and teachers, the teacher utility maximization problem is complicated.

We couple these individual-level data with official school-by-year records from the state's Department of Public Instruction. These record the official value of the composite growth index along with a few other school-level summary statistics. The composite growth index ranges from -0.45 to 0.66 ; schools with values above zero qualify for the \$750 bonus.¹³

B. Methodology

Our basic goal is to examine the impact of bonus receipt on student performance in the next academic year, using RD analysis. RD can be performed either parametrically or nonparametrically. In both varieties, the outcome is modeled as a smooth function of the assignment variable, with the possibility of a discrete jump at the threshold. We use the Hahn, Todd, and van der Klaauw (2001) nonparametric specification in this study by estimating a local linear regression to fit a smooth function to either side of the discontinuity. While it is not necessary to specify a functional form using this method, a bandwidth—effectively, the number of data points incorporated into the local linear regression at any point—must be selected. As the bandwidth increases, the local linear regression approaches a simple linear model; small bandwidths permit a greater number of inflection points in model fit. We report results for a variety of bandwidths centered around the optimal bandwidth, as defined by Imbens and Kalyanaraman (2009), as well as the estimate for the robust bandwidth, as defined by Calonico, Cattaneo, and Titiunik (2014).

Our data contain student-level records, yet assignment to the treatment is at the school level. To estimate the impact of the treatment correctly, we collapse the data to the school-level averages. We then weight observations by the number of student observations used in the school-specific means of the dependent and predetermined student sociodemographic variables.

C. Standard Tests of RD Validity

To attach a causal interpretation to RD estimates of the difference in test score growth on either side of the bonus discontinuity, we must verify a series of assumptions that underlie the method (Imbens and Lemieux 2008; Lee and Lemieux 2010).

First, we check for evidence of manipulation: that schools employ strategies to place their composite growth index score just above the critical value. Figure 2 shows the distribution of average growth performance measure across all school-year observations, with a normal density overlaid. If schools manipulated their growth performance to affect the bonus

¹³ Lending credence to our assertion that it takes effort to understand the incentive scheme and construct a best response, we were unable to perfectly duplicate the state's growth scores using individual-level data. In addition, while the state has been making statistical information available on the web since before the ABC program was in place, the growth scores were made public for only the 2005–6 and 2006–7 school years.

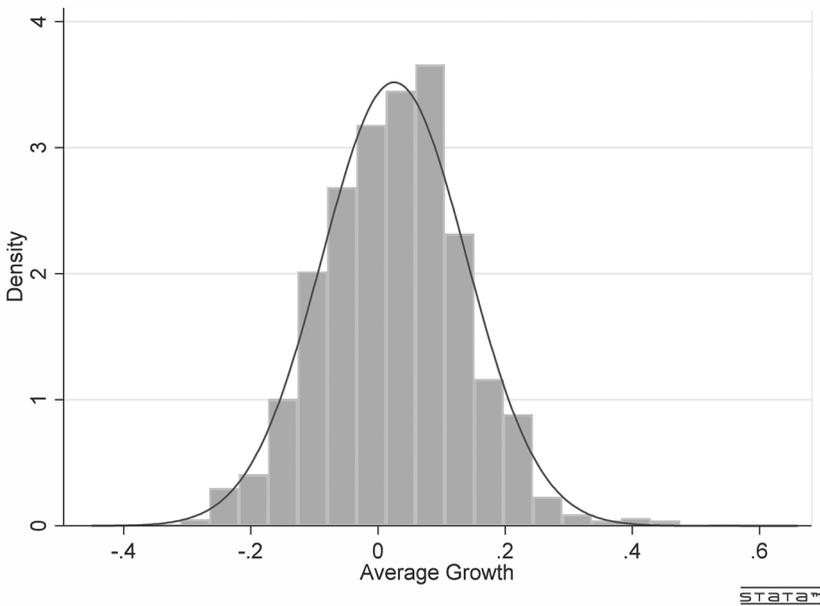


Figure 2.—Density of observations across assignment variable of full sample of schools, which shows little evidence of manipulation at threshold cutoff (normal density overlaid). Threshold cutoff is centered at zero.

outcome, such manipulation would be expected to create a discontinuity in the density of school-year observations at the bonus threshold (McCrary 2008). Table 3 shows that using the entire sample as well as all subsamples used to explore the deeper implications of the conceptual frameworks, the McCrary test statistic is uniformly insignificant.¹⁴ Intuitively, as the composite growth index is a function of the performance of dozens to hundreds of students, calculated using a complex formula and revealed *ex post*, it would seem difficult to engage in *ex ante* manipulation.

We next check for balance in covariates on both sides of the bonus threshold. Table 4 looks for discontinuities at the threshold for predetermined student sociodemographic variables: percent of school population that is female, minority, eligible for free/reduced-price lunch, and limited English proficient as well as enrollment size of the school. Because of the large number of subsample analyses, we present only the optimal Imbens-Kalyanaraman bandwidth result.¹⁵ Alternative bandwidth results are presented in tables A1–A3. Most of the RD estimates are indistinguishable from zero, which implies that the observed discontinuity in test scores is less likely to arise coincidentally from sharp changes in demographic characteristics. The one exception is that there is a discernible pattern

¹⁴ Standard McCrary test density figures for relevant subsamples are presented in figs. A1–A4 (figs. A1–A13 are available online).

¹⁵ A subset of RD figures showing the lack of discontinuity for these placebo runs is presented in figs. A5–A8.

TABLE 3
 MCCRARY TEST FOR CONTINUITY OF ASSIGNMENT
 VARIABLE AT THRESHOLD

Sample	Table	Discontinuity	SE
All schools	5	.0855	(.0910)
No fails	6	.0851	(.1717)
One or more fails		.0627	(.1346)
Years of experience:	7		
<5		.0530	(.1717)
5–10		–.0576	(.1403)
>10		–.0025	(.1969)
No fails:	8		
<5 years of experience		.0785	(.2767)
5–10 years of experience		.1987	(.3935)
>10 years of experience		–.2446	(.4822)
One or more fails:	8		
<5 years of experience		.0681	(.2147)
5–10 years of experience		.0470	(.2318)
>10 years of experience		.0976	(.3127)
Achievement level:	9		
1		.0593	(.1151)
2		.0808	(.1132)
3		.0828	(.1137)
4		.0881	(.1143)

Note.—McCrary statistics are calculated for every subsample. Relevant subsamples are analyzed in tables 5–9.

in student enrollment size, especially for the subsample of schools that have experienced at least one bonus failure over the past 5 years. In these, conditional on the principal being inexperienced, schools that fail to get the bonus this year is substantively larger compared with schools that just qualify. If these schools are led by a principal with 5–10 years of experience, the opposite holds, with schools that just fail being smaller than schools that just receive the bonus. Although odd, these results are not consistently replicated in other bandwidth estimates.

Next, we verify that there is in fact a discontinuity—that schools on either side of the eligibility threshold were in fact differentially likely to receive a bonus. Figure 3 shows teachers' bonus receipt as a function of the composite growth index. There is a sharp discontinuity in probability of bonus receipt (from 0 to 1) at the assignment variable value of 0. Teachers to the right of the discontinuity receive a bonus of at least \$750. There is an additional fuzzy discontinuity around 0.1–0.2 in average growth, above which teachers receive \$1,500.¹⁶

These results support the validity of the North Carolina ABCs program for causal estimation of a local average treatment effect (LATE). As noted above, a traditional rational behavioral model would predict a LATE of zero as the bonus payments are applied *ex post*, and schools on either

¹⁶ The fuzziness reflects the fact that eligibility for the \$1,500 bonus was determined by additional factors beyond the composite growth index. For example, in certain years, test score gains needed to be distributed sufficiently broadly across students in order for schools to receive the \$1,500 bonus. See the appendix for further discussion.

TABLE 4
 PLACEBO RD WITH PREDETERMINED VARIABLES FOR ALL SUBSAMPLES
 DEPENDENT VARIABLE: SCHOOL-LEVEL DEMOGRAPHIC POPULATION MEASURES

	All Principals	Years of Experience		
		>10	5–10	<5
A. All Schools				
Percent female	.0008 (.0041)	-.0098 (.0085)	-.0005 (.0064)	.0084 (.0066)
Percent minority	-.0101 (.0261)	.0454 (.0500)	-.0330 (.0430)	-.0102 (.0467)
Percent poverty	.0099 (.0213)	.0728 (.0512)	-.0298 (.0303)	.0288 (.0344)
Percent LEP	-.0059 (.0073)	.0042 (.0154)	-.0087 (.0115)	-.0041 (.0108)
School size	-10.2669 (15.6733)	-52.4741 (42.4510)	41.7285* (23.4730)	-42.7783** (19.5751)
B. Qualified for Bonus Every Year				
Percent female	-.0003 (.0043)	-.0184 (.0127)	-.0018 (.0068)	.0165** (.0081)
Percent minority	-.0040 (.0379)	.0810 (.0675)	-.0379 (.0593)	-.0147 (.0606)
Percent poverty	.0026 (.0314)	.1184* (.0698)	-.0782* (.0466)	.0366 (.0460)
Percent LEP	-.0072 (.0109)	.0016 (.0184)	-.0214 (.0182)	.0000 (.0161)
School size	2.0791 (20.3082)	-52.4069 (41.4863)	58.7390* (33.9442)	-28.9612 (22.5820)
C. Failed to Qualify for Bonus in at Least 1 Year				
Percent female	-.0023 (.0092)	.0102 (.0164)	-.0073 (.0125)	-.0086 (.0150)
Percent minority	-.0072 (.0398)	-.0166 (.0674)	-.0174 (.0697)	.0268 (.0737)
Percent poverty	.0218 (.0325)	-.0142 (.0746)	.0499 (.0558)	.0521 (.0535)
Percent LEP	-.0007 (.0125)	-.0053 (.0210)	.0179 (.0139)	.0011 (.0232)
School size	-16.7534 (21.6384)	-10.0840 (44.9934)	44.1183** (19.4519)	-81.1473** (37.1202)

Note.—Standard errors are in parentheses. Results for optimal bandwidth from the Imbens-Kalyanaraman algorithm are presented. Alternative bandwidth results, including estimator with robust bias-corrected confidence interval from Calonico, Cattaneo, and Titiunik (2014), are available in tables A1–A3. LEP = limited English proficient.

* Significant at the 10% level.

** Significant at the 5% level.

side of the threshold should have nearly identical expectations regarding the potential returns to effort in the following year.

V. Results

We next present our findings from the RD analysis. Throughout this section, we present results that are in line with or contradict each of our

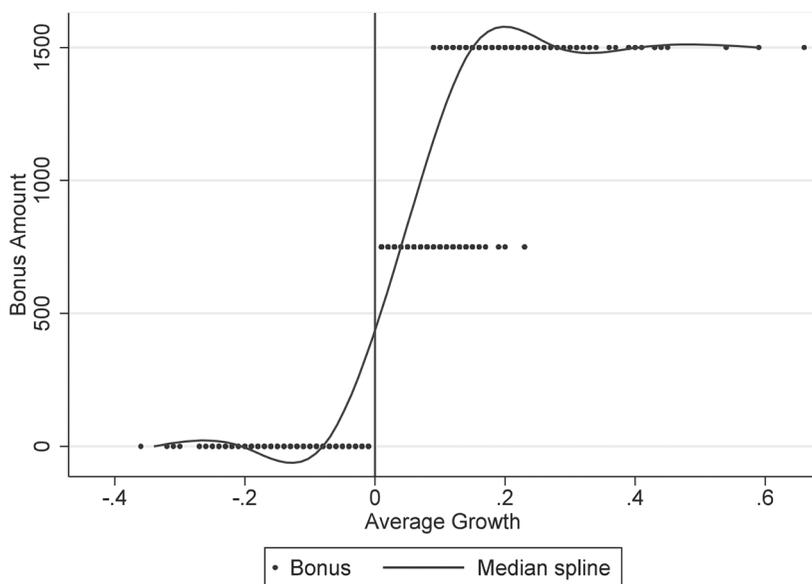


Figure 3.—Existence of discontinuity in probability of bonus receipt at policy change, with sharp discontinuity at \$750 for expected growth and fuzzy discontinuity at \$1,500 for greater than expected growth.

three conceptual frameworks, in order. Therefore, in section V.A, we explore regression results that evaluate whether principals operate with complete information. Section V.B examines whether the RD results corroborate information costs, and section V.C investigates learning. Section V.D presents evidence of effect heterogeneity suggesting that school leaders may harbor confusion about the summary statistic incentivized in the North Carolina bonus program.

A. Documenting the Basic Effect

Figure 4 presents a graphical representation of our most basic RD estimates, and table 5 reports the associated effects and standard errors. In the case of math scores, our estimates indicate—as promised—that schools just below the bonus eligibility threshold exhibit higher test score gains relative to barely eligible schools. The estimated effect is fairly robust to bandwidth choice, ranging from 0.0260 to 0.0564, with higher point estimates in models with narrower bandwidths (across the Imbens-Kalyanaraman estimates).¹⁷ The estimated discontinuity is inconsistent with the simple, traditional complete information model.

These are substantial local average treatments effects, about one and a half times larger in magnitude compared with discontinuities estimated with the same data set for the impact of failing to make adequate yearly

¹⁷ Figures with wider and narrower bandwidths are presented in figs. A9–A11.

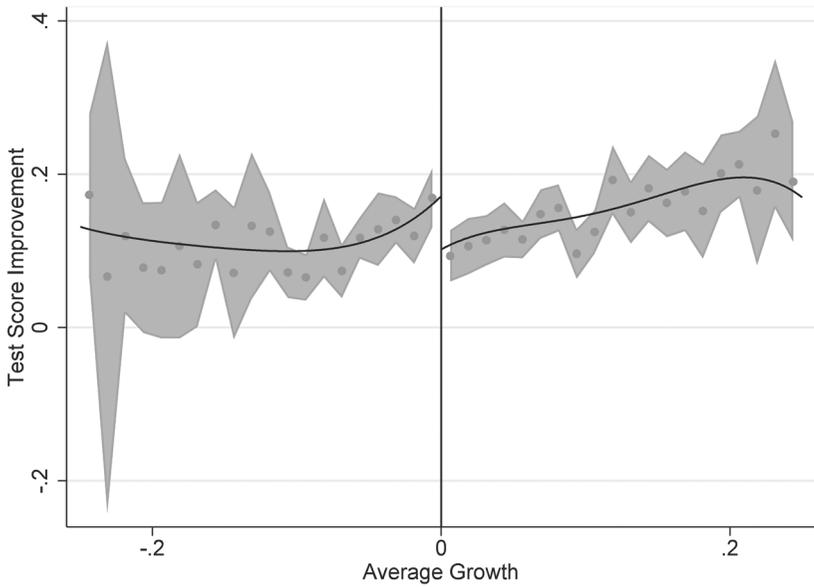


Figure 4.—Simple RD illustration of math score improvement in year $t + 1$ conditional on just being below qualification for bonus in year t . Line is generated with a polynomial of degree 4. Local averages are presented with 40 bins. Observations for assignment variable less than -0.25 or greater than 0.25 are dropped for presentation purposes. Shading indicates the 95% confidence interval.

progress under NCLB (Ahn and Vigdor 2019). In addition, figure 4 shows that this improvement is quite meaningful for schools in close proximity to the bonus threshold. The association of larger effects with narrower bandwidth—and hence more flexible functional form—is consistent with an incentivization effect that is highly localized to the area immediately adjacent to the discontinuity. This localization suggests that school personnel pay attention to more than just the discrete signal, as conditional

TABLE 5
RD ESTIMATES OF IMPACT OF FAILING TO RECEIVE BONUS
DEPENDENT VARIABLE: READING AND MATH STANDARDIZED SCORE GROWTH

	Reading ($N = 2,276$)	Math ($N = 2,276$)
IK optimal bandwidth	-.0098 (.0111)	-.0351** (.0140)
IK half optimal bandwidth	-.0226 (.0155)	-.0458** (.0192)
IK twice optimal bandwidth	-.0057 (.0096)	-.0260** (.0116)
CCT bandwidth (robust SE)	-.0462 (.0295)	-.0564** (.0287)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).
** Significant at the 5% level.

on receiving a negative signal the response appears to vary as a function of distance to the threshold.¹⁸ This could be consistent with either an information cost or a learning framework.

For reading scores, the pattern of RD estimates across bandwidths is similar to math, with larger estimates at narrower bandwidths. However, most are statistically insignificant.¹⁹ Point estimates are reported in table 5. From this point, our discussion will center on math results, as reading results are mostly statistically insignificant in all specifications.

The specification for table 5 is the standard for almost every RD study of merit pay (or any accountability policy). Therefore, while finding a substantive effect for math in the context of the North Carolina program is interesting, it is certainly not novel. However, our point is that the standard results are in tension with a story of principals and teachers operating with complete information.

We confirm the robustness of our initial findings by conducting a set of analyses to check for possible confounding factors. Detailed descriptions are in the appendix. To summarize, we (1) examine whether the estimated math treatment effect is robust to bandwidth choice (figs. A9–A12), (2) run the analysis with artificial thresholds that do not relate to the policy in reality (fig. A13), and (3) estimate a set of parametric RD specifications that mirror our nonparametric model (table A4, specification [A1]). We find that (1) the treatment effect is robust to bandwidth choice, (2) estimates away from the true threshold yield imprecise estimates, and (3) the parametric RD yield estimates for the discontinuity at about 0.048, which is bracketed by the Imbens-Kalyanaraman estimate with the original bandwidth and the Calonico, Cattaneo, and Titiunik estimate.

In summary, the strong estimated reaction to bonus receipt appears robust. It does not, however, make sense in the context of a traditional rational model. From this point forward, we assess alternate explanations for the effect.

B. Testing the First Alternate Model: Bonus as Signal

The model outlined above suggests that schools act to assess and potentially reoptimize their behavior only in the presence of a signal that such activity may yield dividends. Results to this point suggest that failure to receive a bonus might serve as such a signal and that schools within a narrow band short of the bonus threshold believe that reoptimization is necessary

¹⁸ A possible reason for the gains may be teachers moving in response to the bonus outcome. If less effective teachers in failing schools are more likely to transfer, this may generate the discontinuity. However, the modest bonus amount is most likely not enough to induce a large-scale exodus. Similarly, while poorly performing students transferring might produce the discontinuity, student transfer rates in North Carolina are low (see Ahn and Vigdor 2019).

¹⁹ This is in line with most of the literature that finds teachers and schools less able to impact reading scores compared with math scores (e.g., see Jacob 2005; Reback 2008; Rouse et al. 2013). A few studies have found the opposite (see Muralidharan and Sundararaman 2011).

to push them into the eligible category. This is, so far, in line with the information cost framework. As discussed further in section II.B, this explanation suggests that the signal value of bonus receipt (or nonreceipt) is strongest when a school experiences a change in status after a period of relative stability. For this reason, we now study how reactions to bonus receipt vary across schools with differing histories, and therefore differing expectations, regarding the bonus.

Table 6 shows RD estimates for subsets of schools divided by their past performance in the bonus system. Schools are divided into those that have continuously qualified for the bonus and those that have at least one failure in the past 5 years. If the failure to qualify for the bonus serves as an easy-to-interpret signal, we would expect to see strong reaction from high-performing schools upon first failure. Additionally, we may expect lower-performing schools exerting maximal effort to stay above the bar once they qualify.²⁰ That is to say, we might expect an opposite-signed effect among schools with a history of infrequent bonus attainment.

Results from table 6 contradict predictions of a stronger response by schools upon their first failure to qualify for the bonus. In fact, schools that have never failed do not exhibit any test score response to their first failure. The discontinuity is not observed across any bandwidth when schools have qualified for the bonus in all years prior. Schools that have had previous failures in the recent past register substantial extra gains after a near miss relative to a near make. For math scores, academic performance increases by approximately 0.05 of a standard deviation after the next failure (for confirmation of these results using parametric RD, see table A4, specification [A2]). As we will show in section V.C, another prediction from the information cost framework fails to hold as well: principal experience impacts a school's ability to respond effectively to incentives. Clearly, the results fail to support a simple story of bonus receipt as a cheap-to-acquire signal.

C. Testing the Second Alternate Model: Uncertain Production Technology and Learning

School administrators' asymmetric responses at the point of bonus discontinuity may reflect incomplete knowledge about the nature of incentivization and the production process more generally. Whereas the bonus-as-signal model predicts more significant responses to the bonus over time—because it has very little signal value at the beginning of time—the learning model suggests that agents will adopt rational behavior in the long run while suggesting that behavioral asymmetries might not translate effectively into output asymmetries in the very short run. Indeed, our results above identifying stronger asymmetries among schools with a track record of poor performance is entirely consistent with the predictions in

²⁰ One may argue that these schools may incrementally increase academic growth, instituting more costly reforms as required. However, one would assume that schools this savvy would not repeatedly fail to qualify for the bonus.

TABLE 6
RD ESTIMATES BY ABC BONUS HISTORY OF SUCCESS AND FAILURES
DEPENDENT VARIABLE: READING AND MATH STANDARDIZED SCORE GROWTH

	Reading ($N = 817$)	Math ($N = 817$)
A. Qualified for Bonus Every Year		
IK optimal bandwidth	-.0077 (.0194)	-.0381 (.0361)
IK half optimal bandwidth	-.0246 (.0252)	-.0384 (.0535)
IK twice optimal bandwidth	-.0045 (.0179)	-.0440 (.0304)
CCT bandwidth (robust SE)	-.0452 (.0511)	-.0243 (.0768)
B. Failed to Qualify for Bonus in at Least 1 Year		
	Reading ($N = 1,431$)	Math ($N = 1,431$)
IK optimal bandwidth	-.0139 (.0137)	-.0415*** (.0160)
IK half optimal bandwidth	-.0324 (.0204)	-.0485** (.0221)
IK twice optimal bandwidth	-.0118 (.0117)	-.0292** (.0132)
CCT bandwidth (robust SE)	-.0523 (.0339)	-.0699** (.0354)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

** Significant at the 5% level.

*** Significant at the 1% level.

section II.C. In this section, we present further tests based on principal experience levels.

Table 7 reproduces the basic results from table 5 by splitting schools into those headed by principals with low (less than 5 years as a principal), medium (5–10 years), and high (more than 10 years) experience.²¹ Estimates indicate that school with principals of midlevel experience just below the bonus eligibility threshold exhibit higher test score gains relative to barely eligible schools. The effect ranges from 0.0442 to 0.0999. These are large improvements, comparable to the impact of replacing an ineffective principal from NCLB sanctions (Ahn and Vigdor 2019). Similar effects are not observed for principals with low or high levels of experience. In fact, not only are the discontinuities statistically insignificant, but also the estimated magnitudes are 30%–70% smaller than those estimated for principals of midlevel experience. This pattern of discontinuity estimates across principal experience is consistent with imperfect information and learning.

²¹ The principal experience coding is based on the total number of years holding the title of principal rather than tenure at an individual school. Splitting the sample by tenure at a given school yields similar results, consistent with the notion that information about the production process at one school may not translate directly to another.

TABLE 7
RD ESTIMATES BY EXPERIENCE LEVEL OF PRINCIPAL (Math Scores Only)
DEPENDENT VARIABLE: MATH STANDARDIZED SCORE GROWTH

	Years of Experience		
	>10 (N = 456)	5–10 (N = 824)	<5 (N = 968)
IK optimal bandwidth	–.0302 (.0310)	–.0456** (.0213)	–.0209 (.217)
IK half optimal bandwidth	–.0373 (.0417)	–.0536** (.0274)	–.0507 (.0302)
IK twice optimal bandwidth	–.0161 (.0253)	–.0442** (.0200)	–.0157 (.0175)
CCT bandwidth (robust SE)	.0261 (.0889)	–.0999** (.0424)	–.0682 (.0421)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

** Significant at the 5% level.

Even with imperfect information and learning, we find discontinuities in table 7 for a subset of principals because incentive effects are observed only after failing to get the bonus.

The statistically *insignificant* effects for principals with more than 10 years of experience is consistent with the prediction that highly experienced principals will have had periods of sustained success, allowing them to learn about the stochastic shocks (luck) in the bonus outcome. These principals converge to full information regarding the nature of education production, understanding that the impetus to adjust input choices in period t is not a discontinuous function of outcomes in period $t - 1$ and possessing the expertise necessary to make these adjustments effectively.

On the other hand, the statistically significant discontinuity observed for principals with 5–10 years of experience aligns with the behavior of principals who have learned about the links between effort and output due to experimentation but as yet are relatively uninformed about the role of exogenous factors.

The lack of an observed discontinuity for principals with less than 5 years of experience could be attributed to the lack of understanding about how to improve. Without the experience of learning about the production process, these principals may indeed be trying to improve but with mixed success, leading to insignificant estimates.

Splitting schools by accountability history and principal experience further supports our learning framework. Table 8 presents these results. While schools with spotless records do not have statistically significant discontinuities in response to their first failure regardless of principal experience, the exceptionally small estimates for schools with highly experienced principals indicate that these schools do not respond at all to the first failure, supporting our hypothesis that they rationally attribute the aberrant result to chance and do not implement wholesale changes in response.

TABLE 8
RD ESTIMATES BY ABC BONUS HISTORY \times EXPERIENCE LEVEL (Math Scores Only)
DEPENDENT VARIABLE: MATH STANDARDIZED SCORE GROWTH

	Years of Experience		
	>10 (<i>N</i> = 180)	5–10 (<i>N</i> = 285)	<5 (<i>N</i> = 352)
A. Qualified for Bonus Every Year			
IK optimal bandwidth	.0179 (.0665)	−.0823 (.0544)	−.0510 (.0557)
IK half optimal bandwidth	−.0095 (.0897)	−.0746 (.0685)	−.0618 (.0883)
IK twice optimal bandwidth	.0050 (.0588)	−.0908* (.0522)	−.0549 (.0453)
CCT bandwidth (robust SE)	.0853 (.2151)	.0062 (.1270)	−.1049 (.0852)
B. Failed to Qualify for Bonus at Least Once			
	>10 (<i>N</i> = 276)	5–10 (<i>N</i> = 539)	<5 (<i>N</i> = 616)
IK optimal bandwidth	−.0420 (.0334)	−.0597** (.0273)	−.0109 (.0215)
IK half optimal bandwidth	−.0550 (.0441)	−.0660* (.0376)	−.0317 (.0278)
IK twice optimal bandwidth	−.0314 (.0295)	−.0438** (.0229)	−.0168 (.0184)
CCT bandwidth (robust SE)	.0065 (.0807)	−.1067** (.0538)	−.0624 (.0437)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

* Significant at the 10% level.

** Significant at the 5% level.

At intermediate experience levels, the learning model predicts significantly different patterns of discontinuous response to the discrete signal. Those with a strong track record correctly impute that their likelihood of retaining the incentive payment in the next period is effectively identical on either side of the discrete threshold—and in any event, they know little about how to improve their performance were it to be indicated. Principals with a weaker track record are more empowered to take actions to improve performance when they are asked to do so and may be more likely to underestimate the role of chance in determining whether they stay above or below the threshold in the following period. The overreaction around the bonus threshold is concentrated among schools with histories of poor performance headed by principals with midlevel experience, with a response of about 0.0438–0.1067 of a standard deviation.²²

²² It is interesting to note that although there seems to be some response by highly experienced principals upon additional failures, once principals with very short tenures (less than 3 years) are eliminated, the (still insignificant) discontinuity drops to similar magnitude as schools with no failures. We hypothesize that some of the short-tenure, highly experienced principals may be tasked with resuscitating underachieving schools under the NCLB regime.

This result is also consistent with the principal's prior on the variance of shock (luck) evolving from a small value (where production is assumed to be mostly deterministic) to a larger value (where production is at least partially dependent on luck) with experience: principals initially learn on the job to systematically overemphasize the role of effort relative to luck in scenarios where they have been repeatedly exhorted to exert more effort or find new ways to increase test scores. Upon finding sustained success, the experienced principal learns not to respond as he or she learns about the stochastic component of the education production function. Parametric RD results in table A4 (specifications [A3]–[A5]) replicate these findings.

D. Effect Heterogeneity and the Impact of Dueling Accountability Systems

As described above, North Carolina's accountability program paid bonuses on the basis of year-over-year test score gains. There is no reason to believe that gains are easier to produce among students in close proximity to the proficiency threshold. Indeed, depending on the test, large gains may be easiest to produce in the tails of the distribution. By contrast, the federal NCLB system incentivizes proficiency rates, which quite clearly gives schools an incentive to target instructional resources on those students in close proximity to the proficiency threshold (Neal and Schanzenbach 2010). Given the low degree of correlation between bonus receipt and NCLB sanction shown in table 1, evidence that principals focus on students near the threshold when the state system gives them strong reasons to focus on generating gains would be consistent with a fundamental misunderstanding about the nature of the incentive system.

Table 9 presents an analysis of math score improvements for students stratified by initial achievement level. The unit of observation continues to be school/year, but only data on students in a given performance level are used to compute the average test score growth statistic. We see that statistically significant discontinuities exist for students at achievement levels 2 and 3. The border between these levels is the bar for proficiency. While it is clear that schools that just failed to qualify for the bonus respond substantively, the apparent focus on students near the proficiency level suggests that they may not fully understand the bonus program.

These results suggest one specific dimension where there is some scope for learning among school leaders. In a system where there are dueling accountability systems, it may take time for principals to understand that the summary statistic used to assess a school in one system (mean year-over-year test score gain) is not necessarily maximized by a strategy that focuses on a different metric (percent proficient).

VI. Conclusion

This paper identifies a scenario in the education context where, relative to a rational maximization model, agents exhibit too much sensitivity to

TABLE 9
RD ESTIMATES BY MATH PROFICIENCY LEVEL

Level	RD Estimate
1. Insufficient mastery ($N = 2,078$):	
IK optimal bandwidth	.0121 (.0348)
IK half optimal bandwidth	.0092 (.0502)
IK twice optimal bandwidth	.0046 (.0283)
CCT bandwidth (robust SE)	-.0015 (.0573)
2. Inconsistent mastery ($N = 2,212$):	
IK optimal bandwidth	-.0342* (.0182)
IK half optimal bandwidth	-.0324 (.0218)
IK twice optimal bandwidth	-.0336* (.0175)
CCT bandwidth (robust SE)	-.0310 (.0319)
3. Sufficient mastery ($N = 2,195$):	
IK optimal bandwidth	-.0399** (.0185)
IK half optimal bandwidth	-.0498* (.0263)
IK twice optimal bandwidth	-.0317** (.0152)
CCT bandwidth (robust SE)	-.0633** (.0306)
4. Superior mastery ($N = 2,161$):	
IK optimal bandwidth	-.0261 (.0168)
IK half optimal bandwidth	-.0255 (.0221)
IK twice optimal bandwidth	-.0236 (.0152)
CCT bandwidth (robust SE)	-.0324 (.0370)

Note.—Standard errors are in parentheses. Dependent variables are defined in the text. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

* Significant at the 10% level.

** Significant at the 5% level.

discrete signals based on continuous variables available in their information sets. While we do not claim to have identified the only possible explanation for this phenomenon, we show evidence that this behavior is consistent with a Bayesian learning model, in which imperfectly informed agents focus on the discrete signal before they fully appreciate the continuous measure that underlies it. In our case, principals obtain information about their school's performance and learn to pay attention to only the underlying continuous information.

This paper also sheds light on issues in evaluating and operating incentive systems. Until agents learn, through repetition, how input choices map into output, incentive programs may appear to have no impact. This may explain why evaluations of one-shot educational incentive schemes sometimes find no effects, while schemes implemented over time show important effects. Our findings suggest that it may be inappropriate to evaluate incentive systems on the basis of short-term implementation experiments. Administrators should also resist the temptation to calibrate the system to maximize effort, as this may hamper learning by principals.

Our findings of incentive effects dissipating as the RD bandwidths increase point to the limitation of one (or a few) threshold(s) for incentive qualification. In designing future policy, we may wish to consider a continuous payout that rewards gains anywhere along the growth domain or places the threshold at the point with the highest density of schools.²³

In addition, beyond models of competition among teachers to induce gains, information and learning may also have a role to play in education production. Ratcheting up pressure may lead to more effort by principals and teachers to produce gains but only to the extent that they have the capacity to make improvements. Those without experience or insight into the education process may be unable to effectively respond. Related to this insight, incentive thresholds should be based on the ability of the principal. While experienced principals may be held to a higher standard, demands on new school leaders should be tempered and experimentation encouraged.

Finally, sharing of knowledge and best practices across schools may be helpful in speeding up gains. More collaborative incentives, such as rewards based on joint production of a group of experienced and new principals, may hit the sweet spot of inducing both effort exertion and information sharing. Ultimately, the research shows that, as is the case with most well-meaning incentive systems, a one-size-fits-all program that expects all principals to make similar gains, if only they were appropriately motivated to do so, may be both inefficient and unfair.

Appendix

A1. Robustness Checks for RD

Figure A12 presents evidence of robustness of the math treatment effect estimate to bandwidth choice by representing effect estimates and 95% confidence intervals at 20 different bandwidths. Nineteen of the twenty estimates lie within a narrow band indicating an effect size between 0.04 and 0.06. Fifteen of the 20 estimates have

²³ Alternative systems may have their own issues. Continuous payout schemes may have resistant buy-in if it appears too opaque. Setting the threshold at the fattest part of the density may mean that it gets set too low (high), leading to the bonus being a foregone (impossible) outcome for many schools.

confidence intervals that exclude zero. The contradictory effect estimates correspond to extremely narrow bandwidth values, where the effect is estimated using a smaller set of data points, leading to progressively less precise estimates. The estimate with the narrowest bandwidth is the only one outside the 0.04–0.06 range; this estimate is so imprecise that we fail to reject effect sizes three times the value of other estimates in either direction.

Figure A13 shows the results of standard falsification tests using thresholds of no policy relevance. RD estimates using placebo composite growth index values from -0.2 to 0.3 yield imprecise estimates that are never significantly different from zero.

A2. Parametric RD Specifications

We run a set of parametric RD estimations. The benefit of this analysis is twofold: parametric RD allows us to use the entire sample of schools in one regression, increasing the power, and it allows us to replicate the nonlinear effects of experience, which we show in later sections using nonparametric RD. Our base specification (A1) is as follows:

$$\Delta\text{math}_{it} = \beta_0 + X_{it}\beta_1 + \beta_2 I_{it} + f(\Delta y_{i,t-1}) + f(I_{it} \cdot \Delta y_{i,t-1}) + \epsilon_{it}. \quad (\text{A1})$$

The dependent variable, Δmath_{it} , is the change in math score in year t . The assignment variable $\Delta y_{i,t-1}$ is school i 's normalized academic growth rate in year $t - 1$. I_{it} is an indicator variable that equals 1 if the school qualified for the bonus in year $t - 1$, with $\Delta y_{i,t-1} \geq 0$. X_{it} is a vector of control variables—such as percent of students who are female, minority, limited English proficient, and eligible for free and reduced-price lunch—as well as year and school dummy variables. The $f(\cdot)$ term is a flexible function. The inclusion of the second $f(\cdot)$ with $I_{it} \cdot \Delta y_{i,t-1}$ allows the conditional mean function on the other side of the discontinuity to have a different shape. The idiosyncratic error term is represented by ϵ_{it} . All regressions are weighted by the number of students in the school. Estimates of the parametric discontinuity are shown to be about 0.048 in the first set of results in table A4. This value is bracketed by the Imbens-Kalyanaraman estimate with the optimal bandwidth and the Calonico, Cattaneo, and Titiunik estimate.

In addition to the base specification (A1), we also run alternative models that are used to assess the other information models. Specification (A2) includes accountability history of the school, with a dummy variable that equals 1 for schools that have consistently qualified for the bonus, Z_{it} . Each argument inside the function $f(\cdot)$ has polynomial controls in its own and all possible interaction terms:

$$\Delta\text{math}_{it} = \beta_0 + X_{it}\beta_1 + \beta_2 I_{it} + \beta_3 Z_{it} + \beta_4 I_{it} \cdot Z_{it} + f(\Delta y_{i,t-1}, Z_{it}) + f(I_{it} \cdot \Delta y_{i,t-1}, Z_{it}) + \epsilon_{it}. \quad (\text{A2})$$

Results from parametric RD specification (A2), which separately estimates treatment effects for schools by their accountability history, show that the treatment effect on schools with consistent success in attaining the bonus is negligible from zero, while the treatment effect on schools with at least one failure is marginally statistically significant. The difference in estimated magnitude of the parameters is striking, as the latter estimate is over eight times larger than the former estimate.

We also run models that include measures of the principal's experience. Specification (A3) accounts for experience and learning in a basic way, as follows:²⁴

$$\begin{aligned} \Delta\text{math}_{it} = & \beta_0 + X_{it}\beta_1 + \beta_2 I_{it} + \beta_3 \text{LEV}_{it} + \beta_4 I_{it} \cdot \text{LEV}_{it} + f(\Delta y_{i,t-1}, \text{LEV}_{it}) \\ & + f(I_{it} \cdot \Delta y_{i,t-1}, \text{LEV}_{it}) + \epsilon_{it}. \end{aligned} \quad (\text{A3})$$

LEV_{it} is a dummy variable for the experience level for the principal in school i . Specification (A4) uses a continuous measure of principal experience and its squared term to confirm that results are not being driven by the cut-off values for experience:

$$\begin{aligned} \Delta\text{math}_{it} = & \beta_0 + X_{it}\beta_1 + \beta_2 I_{it} + \beta_3 \text{exp}_{it} + \beta_4 \text{exp}_{it}^2 + \beta_5 I_{it} \cdot \text{exp}_{it} + \beta_6 I_{it} \\ & \cdot \text{exp}_{it}^2 + f(\Delta y_{i,t-1}, \text{exp}_{it}, \text{exp}_{it}^2) + f(I_{it} \cdot \Delta y_{i,t-1}, \text{exp}_{it}, \text{exp}_{it}^2) + \epsilon_{it}. \end{aligned} \quad (\text{A4})$$

The final specification (A5) includes accountability history of the school interacted with the experience measure:

$$\begin{aligned} \Delta\text{math}_{it} = & \beta_0 + X_{it}\beta_1 + \beta_2 I_{it} + \beta_3 \text{exp}_{it} + \beta_4 \text{exp}_{it}^2 + \beta_5 I_{it} \cdot \text{exp}_{it} \cdot Z_{it} + \beta_6 I_{it} \\ & \cdot \text{exp}_{it}^2 \cdot Z_{it} + \beta_7 I_{it} \text{exp}_{it} (1 - Z_{it}) + \beta_8 I_{it} \text{exp}_{it}^2 (1 - Z_{it}) \\ & + f(\Delta y_{i,t-1}, \text{exp}_{it}, \text{exp}_{it}^2, Z_{it}) + f(I_{it} \cdot \Delta y_{i,t-1}, \text{exp}_{it}, \text{exp}_{it}^2, Z_{it}) + \epsilon_{it}. \end{aligned} \quad (\text{A5})$$

Parametric regression runs buttress our results from the nonparametric RD analysis. Continuing with specification (A3), which contains dummy variables for low, mid, and high experience level, the parameter of interest here is β_4 . Following the results in table 7, we would expect β_4 for midexperience principals to be negative.

When the continuous measure of experience is used as in specification (A4), the parameters of interest are β_4 and β_6 . If experience is nonlinear as portrayed in the nonparametric RD results, we would expect β_5 to be negative and β_6 to be positive.

Finally, to replicate table 8, we interact experience with school accountability history in specification (A5). Here, we are interested in β_5 , β_6 , β_7 , and β_8 . We would expect β_5 and β_6 to be statistically insignificant, β_7 to be negative, and β_8 to be positive. The results for three specifications above are presented in the third, fourth, and fifth sets of results in table A4.²⁵

As demonstrated, the relevant parameter estimates in all three specifications are consistent with the nonparametric RD results. In particular, in specification (A4), the shape of the experience response shows that the peak of response is between 6 and 7 years of experience.

The inclusion of experience as a control variable (in the parametric RD) is standard in many education production function estimations that include teacher or principal characteristics. Our unique setup—in splitting the school sample by accountability history and principal experience in the nonparametric RD and including the experience and history measures interacted with the indicator for threshold in the parametric RD—allows us to find support for the gradual learning model,

²⁴ For example, if there are only two experience levels, $f(y, \text{LEV}) = \sum_{k=1}^n \gamma_k y^k + \theta_k \text{LEV}^k + \delta_k (y \cdot \text{LEV})^k$.

²⁵ Further robustness checks with a dummy variable for the maximum observable value of experience (13 years) yielded no qualitative differences.

providing a consistent explanation for the seemingly irrational behavior of schools in response to a modest and straightforward merit pay system.

A3. *The Impact of the \$1,500 Threshold*

The North Carolina incentive system is relatively straightforward in its implementation, which lends itself to a simple RD framework. One possible issue is the inclusion of a second, higher threshold that can introduce complications to the estimation of the LATE.

This second threshold requires higher year-over-year test score gains than the expected-growth criterion. In addition, other criteria—such as ensuring that gains are sufficiently broadly distributed across the student population—must be met. This second threshold potentially creates threats to the validity of our estimate because of two factors:

1. The incentive impact of the second threshold may be *additive*, leading to an overestimate of the treatment effect. To the extent that the second threshold, which would double the monetary reward, would incentivize the principal and teachers to higher exertion to attain the bonus, this may lead to higher gains.
2. The incentive impact of the second threshold may be *subtractive*, leading to an underestimate of the treatment effect. Average test score growth at the school could decline if the additional conditions required for qualifying for the higher bonus pull effort and attention away from maximizing average gains.

The potential impact of these factors on the RD estimate depends on the following three factors:

1. The relative size difference between the additive and subtractive factors will determine how much contamination is possible.
2. The number of schools induced to alter behavior in response to the second threshold, especially schools close to our relevant threshold, will impact the estimate. The number of schools induced to change behavior would depend on the perceived probability of success in attaining the second bonus.
3. Related to point 2, bandwidth for RD estimation will change the number of potentially contaminated schools. Keeping in mind that the treated group is schools that just failed (schools to the left of the threshold), a wider bandwidth selection will contain more schools that are less likely to be impacted by incentives from the second threshold. It also increases the power of the estimator, since sample size increases. However, increasing the bandwidth also decreases the incentive effects from our relevant threshold. In fact, this tension underlies the selection of optimal bandwidth in all RD estimation.

As such, it is interesting to note that comparing the Imbens-Kalyanaraman estimates in table 5 at half, optimal, and twice optimal bandwidth, we find that the size of the estimated treatment effect decreases in magnitude from about -0.046 to -0.026 . This may imply that the overall incentive effect from the second threshold is additive. However, it is impossible to truly disentangle the relative sizes of these incentives.

It does appear that the second threshold is a difficult bar to cross, with approximately 15% of schools receiving the bonus in our sample. In contrast, over 60% manage to qualify for at least the lower bonus. This lends some cautious optimism that most schools would consider the second threshold as out of reach. Then, the impact on the first threshold may be modest.

There is room for some cautious optimism that whatever contamination exists, it would be relatively minor. For example, the Tennessee Project on Incentives in Teaching system examined in Springer et al. (2010) offered much more generous benefits (up to \$15,000) yet saw little incentive effects. This may have been because only 20% of teachers had a realistic chance to qualify for any bonus conditional on past performance, leading to an optimal nonresponse by most teachers. With the low fraction of schools that ended up qualifying for the second threshold, it is arguable that given the much smaller bonus amounts in North Carolina, the contamination effects would also have been much more modest.

TABLE A1
PLACEBO RD WITH PREDETERMINED VARIABLES FOR ALL SCHOOLS \times PRINCIPAL EXPERIENCE
DEPENDENT VARIABLE: SCHOOL-LEVEL POPULATION MEASURES

	All	Years of Experience		
		>10	5–10	<5
Percent female:				
IK optimal	.0008 (.0041)	–.0098 (.0085)	–.0005 (.0064)	.0084 (.0066)
IK 0.5 \times optimal	.0064 (.0056)	–.0084 (.0112)	.0049 (.0084)	.0139 (.0091)
IK 2 \times optimal	.0002 (.0034)	–.0083 (.0074)	.0004 (.0055)	.0048 (.0055)
CCT	.0130 (.0091)	–.0088 (.0205)	.0111 (.0119)	.0249* (.0142)
Percent minority:				
IK optimal	–.0101 (.0261)	.0454 (.0500)	–.0330 (.0430)	–.0102 (.0467)
IK 0.5 \times optimal	–.0040 (.0355)	.0397 (.0645)	–.0457 (.0575)	–.0363 (.0669)
IK 2 \times optimal	–.0200 (.0223)	.0177 (.0420)	–.0442 (.0383)	–.0216 (.0369)
CCT	–.0130 (.0518)	.0369 (.1082)	–.0090 (.1011)	–.0803 (.0901)
Percent poverty:				
IK optimal	.0099 (.0213)	.0728 (.0512)	–.0298 (.0303)	.0288 (.0344)
IK 0.5 \times optimal	.0080 (.0298)	.0658 (.0705)	–.0394 (.0367)	.0298 (.0518)
IK 2 \times optimal	.0010 (.0182)	.0475 (.0418)	–.0327 (.0291)	.0103 (.0276)
CCT	–.0272 (.0534)	.0687 (.1091)	–.0923 (.0955)	–.0564 (.0864)
Percent LEP:				
IK optimal	–.0059 (.0073)	.0042 (.0154)	–.0087 (.0115)	–.0041 (.0108)
IK 0.5 \times optimal	–.0067 (.0099)	.0092 (.0198)	–.0163 (.0157)	–.0006 (.0139)
IK 2 \times optimal	–.0062 (.0064)	–.0051 (.0126)	–.0059 (.0107)	–.0044 (.0103)
CCT	–.0296 (.0216)	.0019 (.0276)	–.0645 (.0401)	–.0219 (.0349)

TABLE A1 (Continued)

	All	Years of Experience		
		>10	5–10	<5
School size:				
IK optimal	-10.2669 (15.6733)	-52.4741 (42.4510)	41.7285* (23.4730)	-42.7783** (19.5751)
IK 0.5× optimal	-16.8675 (21.4177)	-60.0398 (62.1447)	48.4007 (33.6117)	-60.5118*** (22.6186)
IK 2× optimal	-6.4109 (13.0207)	-44.3210 (32.5272)	37.0807* (21.2442)	-29.6402* (16.6786)
CCT	-20.3536 (39.3275)	-67.6484 (103.3714)	44.6793 (65.4574)	-43.8053 (51.9305)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT). LEP = limited English proficient.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

TABLE A2
PLACEBO RD WITH PREDETERMINED VARIABLES FOR SCHOOLS WITH AT LEAST
ONE FAILURE (Failed to Qualify at Least 1 Year) × PRINCIPAL EXPERIENCE
DEPENDENT VARIABLE: SCHOOL-LEVEL POPULATION MEASURES

	All	Years of Experience		
		>10	5–10	<5
Percent female:				
IK optimal	-.0003 (.0043)	-.0184 (.0127)	-.0018 (.0068)	.0165** (.0081)
IK 0.5× optimal	.0038 (.0055)	-.0297* (.0180)	.0023 (.0083)	.0300*** (.0114)
IK 2× optimal	.0002 (.0040)	-.0157 (.0100)	-.0019 (.0064)	.0097 (.0065)
CCT	.0144 (.0106)	-.0540* (.0301)	.0184 (.0188)	.0397** (.0155)
Percent minority:				
IK optimal	-.0040 (.0379)	.0810 (.0675)	-.0379 (.0593)	-.0147 (.0606)
IK 0.5× optimal	-.0177 (.0560)	.0505 (.0837)	-.0291 (.0860)	-.0342 (.0899)
IK 2× optimal	-.0126 (.0298)	.0562 (.0562)	-.0491 (.0485)	-.0189 (.0475)
CCT	-.0115 (.0653)	.1003 (.1686)	-.0272 (.1072)	-.0243 (.1081)
Percent poverty:				
IK optimal	.0026 (.0314)	.1184* (.0698)	-.0782* (.0466)	.0366 (.0460)
IK 0.5× optimal	-.0178 (.0468)	.0925 (.0959)	-.1079 (.0668)	.0307 (.0713)
IK 2× optimal	.0060 (.0246)	.0928* (.0551)	-.0600 (.0385)	.0321 (.0359)
CCT	-.0410 (.0697)	.1270 (.1689)	-.1411 (.1183)	-.0029 (.1093)
Percent LEP:				
IK optimal	-.0072 (.0109)	.0016 (.0184)	-.0214 (.0182)	.0000 (.0161)

TABLE A2 (Continued)

	All	Years of Experience		
		>10	5–10	<5
IK 0.5× optimal	-.0163 (.0158)	.0123 (.0225)	-.0414 (.0261)	.0048 (.0228)
IK 2× optimal	-.0069 (.0084)	-.0059 (.0158)	-.0116 (.0139)	.0007 (.0134)
CCT	-.0358 (.0264)	-.0112 (.0312)	-.1054** (.0517)	-.0061 (.0457)
School size:				
IK optimal	2.0791 (20.3082)	-52.4069 (41.4863)	58.7390* (33.9442)	-28.9612 (22.5820)
IK 0.5× optimal	-8.1219 (28.1384)	-61.1657 (55.6294)	46.0301 (53.0092)	-43.0077* (25.6467)
IK 2× optimal	1.5589 (16.7618)	-44.5182 (36.1177)	48.8386* (28.5425)	-23.5783 (19.7112)
CCT	-17.7314 (50.6042)	-97.1022 (132.7873)	36.6106 (93.3691)	-35.0157 (65.0315)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

TABLE A3
PLACEBO RD WITH PREDETERMINED VARIABLES FOR SCHOOLS WITH NO FAILURES
(Qualified for Bonus Every Year) × PRINCIPAL EXPERIENCE
DEPENDENT VARIABLE: SCHOOL-LEVEL POPULATION MEASURES

	All	Years of Experience		
		>10	5–10	<5
Percent female:				
IK optimal	-.0023 (.0092)	.0102 (.0164)	-.0073 (.0125)	-.0086 (.0150)
IK 0.5× optimal	.0050 (.0130)	.0400* (.0228)	-.0090 (.0145)	-.0169 (.0204)
IK 2× optimal	-.0045 (.0079)	.0091 (.0141)	-.0066 (.0122)	-.0110 (.0128)
CCT	.0268 (.0218)	.0804 (.0550)	.0401 (.0832)	-.0091 (.0331)
Percent minority:				
IK optimal	-.0072 (.0398)	-.0166 (.0674)	-.0174 (.0697)	.0268 (.0737)
IK 0.5× optimal	.0071 (.0461)	-.0175 (.0895)	-.0248 (.0914)	-.0085 (.0925)
IK 2× optimal	-.0117 (.0378)	-.0222 (.0611)	-.0242 (.0667)	.0143 (.0643)
CCT	-.0203 (.0957)	-.0183 (.3521)	.2707 (.3583)	-.1688 (.1755)
Percent poverty:				
IK optimal	.0218 (.0325)	-.0142 (.0746)	.0499 (.0558)	.0521 (.0535)
IK 0.5× optimal	.0433 (.0425)	-.0546 (.1069)	.0287 (.0767)	.0612 (.0864)
IK 2× optimal	.0076 (.0304)	.0068 (.0675)	.0278 (.0531)	.0136 (.0453)

TABLE A3 (Continued)

	All	Years of Experience		
		>10	5–10	<5
CCT	.0484 (.0744)	.0689 (.3182)	.1017 (.1326)	-.0113 (.1338)
Percent LEP:				
IK optimal	-.0007 (.0125)	-.0053 (.0210)	.0179 (.0139)	.0011 (.0232)
IK 0.5× optimal	-.0026 (.0152)	-.0031 (.0295)	.0177 (.0158)	-.0104 (.0311)
IK 2× optimal	-.0006 (.0117)	-.0122 (.0186)	.0241* (.0134)	-.0024 (.0210)
CCT	-.0064 (.0369)	-.0070 (.0998)	.0593 (.0497)	-.0347 (.0549)
School size:				
IK optimal	-16.7534 (21.6384)	-10.0840 (44.9934)	44.1183** (19.4519)	-81.1473** (37.1202)
IK 0.5× optimal	-32.1159 (24.3204)	-32.0789 (55.8020)	35.4749 (23.0889)	-73.1363* (40.1655)
IK 2× optimal	-8.4285 (20.8356)	2.9969 (41.0666)	50.6010*** (18.8038)	-55.7719 (34.0292)
CCT	77.9242 (58.0762)	-15.0868 (97.8231)	92.6551 (82.2141)	129.7609 (98.6491)

Note.—Standard errors are in parentheses. Bandwidth determination is by the Imbens-Kalyanaraman (IK) algorithm. Estimator with robust bias-corrected confidence interval is from Calonico, Cattaneo, and Titiunik (2014; CCT).

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at the 1% level.

TABLE A4
PARAMETRIC ROBUSTNESS CHECKS
DEPENDENT VARIABLE: MATH STANDARDIZED SCORE GROWTH

	Coefficient
Specification (A1): base case equivalent to table 5:	
β : treatment at threshold	-.0480** (.0211)
Specification (A2): accountability history equivalent to table 6:	
β_2 : no fail schools	-.0133 (.0586)
β_4 : schools with at least one failure	-.1105* (.0670)
Specification (A3): by experience dummies equivalent to table 7:	
β_4 : low experience	-.0260 (.0216)
β_4 : midexperience	-.0467** (.0203)
β_4 : high experience	-.0134 (.0260)
Specification (A4): by years of experience equivalent to table 7:	
β_5 : linear experience	-.0155** (.0066)
β_6 : quadratic experience	.0010** (.0005)

TABLE A4 (Continued)

	Coefficient
Specification (A5): interaction of experience and accountability history equivalent to table 8:	
β_5 : linear experience (no fails)	-.0377 (.0400)
β_6 : quadratic experience (no fails)	.0040 (.0029)
β_7 : linear experience (at least one fail)	-.0671** (.0339)
β_8 : quadratic experience (at least one fail)	.0054** (.0024)
Observations	2,248

Note.—Standard errors are in parentheses. Specifications control for percent minority, limited English proficient, eligible for free/reduced-price lunch, and female as well as year and school dummies. Regression is weighted by school size. Cubic polynomial controls for assignment variables are presented. Different degrees of polynomial results and parameter estimates for all control variables are available at <http://sites.google.com/site/tomsyahn/>.

* Significant at the 10% level.

** Significant at the 5% level.

References

- Ahn, T. 2013. "The Missing Link: Estimating the Impact of Incentives on Effort and Effort on Production Using Teacher Accountability Legislation." *J. Human Capital* 7 (3): 230–73.
- . 2014. "A Regression Discontinuity Analysis of Graduation Standards and Their Impact on Students' Academic Trajectories." *Econ. Educ. Rev.* 37:64–75.
- . 2016. "A Theory of Dynamic Investment in Education in Response to Accountability Pressure." *Econ. Letters* 149:75–78.
- . 2017. "Strategic Matching of Teachers and Schools with (and without) Accountability Pressure." *Educ. Finance and Policy* 12 (4):516–35.
- Ahn, T., and J. L. Vigdor. 2019. "The Impact of No Child Left Behind's Accountability Sanctions on School Performance: Regression Discontinuity Evidence from North Carolina." Working Paper no. 20511, NBER, Cambridge, MA.
- . 2020. "Opening the Black Box: Behavioral Responses of Teachers and Principals to Pay-for-Performance Incentive Programs." Working paper.
- Allcott, H. 2011. "Social Norms and Energy Conservation." *J. Public Econ.* 95 (9–10): 1082–95.
- Anderson, M., and J. Magruder. 2011. "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database." *Econ. J.* 122 (563): 957–89.
- Atkinson, A., S. Burgess, B. Croxson, et al. 2009. "Evaluating the Impact of Performance-Related Pay for Teachers in England." *Labour Econ.* 16 (3): 251–61.
- Bala, V., and S. Goyal. 1995. "A Theory of Learning with Heterogeneous Agents." *Internat. Econ. Rev.* 36 (2): 303–23.
- Berger, J., and D. Pope. 2011. "Can Losing Lead to Winning?" *Management Sci.* 57 (5): 817–27.
- Berry, D., and B. Fristedt. 1985. *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman & Hall.
- Bertola, G., and R. Caballero. 1990. "Kinked Adjustment Costs and Aggregate Dynamics." In *NBER Macroeconomics Annual*, vol. 5, edited by O. J. Blanchard and S. Fischer, 237–96. Cambridge, MA: MIT Press.

- Béteille, T., D. Kalogrides, and S. Loeb. 2012. "Stepping Stones: Principal Career Paths and School Outcomes." *Soc. Sci. Res.* 41 (4): 904–19.
- Bloom, N., R. Lemos, R. Sadun, and J. Van Reenen. 2015. "Does Management Matter in Schools?" *Econ. J.* 125 (584): 647–74.
- Calonico, S., M. Cattaneo, and R. Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82 (6): 2295–326.
- Card, D., A. Mas, and J. Rothstein. 2008. "Tipping and the Dynamics of Segregation." *Q. J.E.* 123 (1): 178–218.
- Chakrabarti, R. 2011. "Incentives and Responses under No Child Left Behind: Credible Threats and the Role of Competition." *J. Public Econ.* 110:124–46.
- Chiang, H. 2009. "How Accountability Pressure on Failing Schools Affects Student Achievement." *J. Public Econ.* 93:1045–57.
- Clark, D., P. Martorell, and J. Rockoff. 2009. "School Principals and School Performance." CALDER Working Paper no. 38, Nat. Center Analysis Longitudinal Data Educ. Res., Washington, DC.
- Craig, S. G., S. A. Imberman, and A. Perdue. 2015. "Do Administrators Respond to Their Accountability Ratings? The Response of School Budgets to Accountability Grades." *Econ. Educ. Rev.* 49:55–68.
- Dee, T. S., and B. J. Keys. 2004. "Does Merit Pay Reward Good Teachers? Evidence from a Randomized Experiment." *J. Policy Analysis and Management* 23 (3): 471–88.
- Easley, D., and N. Kiefer. 1988. "Controlling a Stochastic Process with Unknown Parameters." *Econometrica* 56:1045–64.
- El-Gamal, M., and R. Sundaram. 1993. "Bayesian Economists, Bayesian Agents: An Alternative Approach to Optimal Learning." *J. Econ. Dynamics and Control* 17 (3): 355–83.
- Figlio, D. N., and L. W. Kenny. 2009. "Public Sector Performance Measurement and Stakeholder Support." *J. Public Econ.* 93 (9–10): 1069–77.
- Figlio, D. N., and M. E. Lucas. 2004. "What's in a Grade? School Report Cards and the Housing Market." *A.E.R.* 94 (3): 591–604.
- Fryer, R. G., S. D. Levitt, J. List, and S. Sadoff. 2012. "Enhancing the Efficacy of Teacher Incentives through Loss Aversion: A Field Experiment." Working Paper no. 18237, NBER, Cambridge, MA.
- Gabaix, X. 2014. "A Sparsity-Based Model of Bounded Rationality." Working paper.
- Gigerenzer, G., and D. Goldstein. 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Rev.* 103 (4): 650–69.
- Glewwe, P., N. Ilias, and M. Kremer. 2010. "Teacher Incentives." *American Econ. J. Appl. Econ.* 2 (3): 205–27.
- Grissom, J. A., and B. Bartanen. 2019. "Principal Effectiveness and Principal Turnover." *Educ. Finance and Policy* 14 (3): 355–82.
- Hahn, J., P. Todd, and W. van der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design." *Econometrica* 69 (1): 201–9.
- Hanushek, E. A., and M. E. Raymond. 2005. "Does School Accountability Lead to Improved Student Performance?" *J. Policy Analysis and Management* 24 (2): 297–327.
- Hemelt, S. W. 2011. "Performance Effects of Failure to Make Adequate Yearly Progress (AYP): Evidence from a Regression Discontinuity Framework." *Econ. Educ. Rev.* 30 (4): 702–23.
- Imbens, G. W., and K. Kalyanaraman. 2009. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." Working Paper no. 14726, NBER, Cambridge, MA.

- Imbens, G. W., and T. Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice." *J. Econometrics* 142 (2): 615–35.
- Imberman, S. A., and M. F. Lovenheim. 2015. "Incentive Strength and Teacher Productivity: Evidence from a Group-Based Teacher Incentive Pay System." *Rev. Econ. and Statis.* 97 (2): 364–86.
- Jacob, B. A. 2005. "Accountability, Incentives, and Behavior: The Impact of High-Stakes Testing in the Chicago Public Schools." *J. Public Econ.* 89:761–96.
- Jin, G., and P. Leslie. 2003. "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards." *Q.J.E.* 118 (2): 409–51.
- Jinnai, Y. 2016. "The Effects of a Teacher Performance-Pay Program on Student Achievement: A Regression Discontinuity Approach." *Econ. Bull.* 36 (2): 993–99.
- Koretz, D. M., and S. I. Barron. 1998. "The Validity of Gains on the Kentucky Instructional Results Information System (KIRIS)." Santa Monica, CA: RAND.
- Lacetera, N., D. G. Pope, and J. R. Sydnor. 2012. "Heuristic Thinking and Limited Attention in the Car Market." *A.E.R.* 102 (5): 2206–36.
- Ladd, H. F. 1999. "The Dallas School Accountability and Incentive Program: An Evaluation of Its Impacts on Student Outcomes." *Econ. Educ. Rev.* 18 (1): 1–16.
- Lavy, V. 2002. "Evaluating the Effect of Teachers' Group Performance Incentives on Pupil Achievement." *J.P.E.* 110 (6): 1286–317.
- . 2009. "Performance Pay and Teachers' Effort, Productivity and Grading Ethics." *A.E.R.* 99 (5): 1979–2021.
- Lee, D. S., and T. Lemieux. 2010. "Regression Discontinuity Designs in Economics." *J. Econ. Literature* 48 (2): 281–355.
- Loyalka, P., S. Sylvia, C. Liu, J. Chu, and Y. Shi. 2019. "Pay by Design: Teacher Performance Pay Design and the Distribution of Student Achievement." *J. Labor Econ.* 37 (3): 621–62.
- Macartney, H. 2016. "The Dynamic Effects of Educational Accountability." *J. Labor Econ.* 34 (1): 1–28.
- Martinez, E. 2010. "Do Housing Prices Account for School Accountability?" Working paper.
- McCrary, J. 2008. "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test." *J. Econometrics* 142 (2): 698–714.
- Miller, A. 2013. "Principal Turnover and Student Achievement." *Econ. Educ. Rev.* 36:60–72.
- Mullainathan, S. 2002. "A Memory-Based Model of Bounded Rationality." *Q.J.E.* 117 (3): 735–74.
- Muralidharan, K., and V. Sundararaman. 2011. "Teacher Performance Pay: Experimental Evidence from India." *J.P.E.* 119:39–77.
- Neal, D. 2011. "The Design of Performance Pay in Education." In *Handbook of the Economics of Education*, vol. 4, edited by E. A. Hanushek, S. Machin, and L. Woessmann, 495–550. Amsterdam: North-Holland.
- Neal, D., and D. Schanzenbach. 2010. "Left Behind by Design: Proficiency Counts and Test-Based Accountability." *Rev. Econ. and Statis.* 92 (2): 263–83.
- Papay, J., R. J. Murnane, and J. B. Willett. 2011. "How Performance Information Affects Human-Capital Investment Decisions: The Impact of Test-Score Labels on Educational Outcomes." Working Paper no. 17120, NBER, Cambridge, MA.
- Pham, L. D., T. D. Nguyen, and M. G. Springer. 2020. "Teacher Merit Pay: A Meta-Analysis." *American Educ. Res. J.* 58 (3): 527–66.
- Reback, R. 2008. "Teaching to the Rating: School Accountability and the Distribution of Student Achievement." *J. Public Econ.* 92 (5–6): 1394–415.
- Rockoff, J., and L. J. Turner. 2010. "Short-Run Impacts of Accountability on School Quality." *American Econ. J. Econ. Policy* 2 (4): 119–47.
- Rothschild, M. 1974. "A Two-Armed Bandit Theory of Market Pricing." *J. Econ. Theory* 9 (2): 185–202.

- Rouse, C. E., J. Hannaway, D. Goldhaber, and D. Figlio. 2013. "Feeling the Florida Heat? How Low-Performing Schools Respond to Voucher and Accountability Pressure." *American Econ. J. Econ. Policy* 5:251–81.
- Sarver, T. 2008. "Anticipating Regret: Why Fewer Options May Be Better." *Econometrica* 76:263–305.
- Sims, C. 2003. "Implications of Rational Inattention." *J. Monetary Econ.* 50:665–90.
- Springer, M. G., D. Ballou, L. Hamilton, et al. 2010. "Teacher Pay for Performance: Experimental Evidence from the Project on Incentives in Teaching." Nashville, TN: Nat. Center Performance Incentives.
- Vigdor, J. L. 2009. "Teacher Salary Bonuses in North Carolina." In *Performance Incentives: Their Growing Impact on American K–12 Education*, edited by M. G. Springer. Washington, DC: Brookings Inst. Press.
- West, M. R., and P. E. Peterson. 2006. "The Efficacy of Choice Threats within School Accountability Systems: Results from Legislatively Induced Experiments." *Econ. J.* 116:C46–C62.
- Winters, M., J. P. Greene, G. Ritter, and R. Marsh. 2008. "The Effect of Performance-Pay in Little Rock, Arkansas on Student Achievement." Working Paper no. 2008-02, Nat. Center Performance Incentives, Nashville, TN.

Copyright of Journal of Human Capital is the property of University of Chicago and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.