The Potential Output Gains from Using Optimal Teacher Incentives: An Illustrative Calibration of a Hidden Action Model

by

Nirav Mehta

Western

Centre for Human Capital and Productivity (CHCP)

Working Paper Series

Department of Economics
Social Science Centre
Western University
London, Ontario, N6A 5C2
Canada

The Potential Output Gains from Using Optimal Teacher Incentives:

An Illustrative Calibration of a Hidden Action Model

Nirav Mehta

University of Western Ontario [*]

July 31, 2017

Abstract

This note examines the potential output gains from the implementation of optimal teacher incentive pay schemes, by calibrating the Hölmstrom and Milgrom (1987) hidden action model using data from Muralidharan and Sundararaman (2011), a teacher incentive pay experiment implemented in Andhra Pradesh, India. Findings suggest that the introduction of optimal individual incentive-pay schemes could result in very large increases in output, about six times the size of the (significant) results obtained in the experiment.

Keywords: hidden action, empirical contracts, teacher incentive pay

1 Introduction

Evidence that teacher quality is an important determinant of human capital that is hard to measure (Goldhaber and Brewer (1997); Rivkin et al. (2005); Hanushek (2011)) has generated substantial policy interest in output-based teacher incentive schemes and motivated a research agenda using randomized controlled trials to estimate whether such schemes affect teacher inputs (see, e.g., Glewwe et al. (2010); Fryer Jr (2013)).

The importance of adopting teacher incentive pay depends on the *potential gain* from their implementation, which can only be computed using *optimal incentive schemes*. Despite the

---

[*]nirav.mehta@uwo.ca

availability of high-quality experimental findings and a well-developed theoretical literature, the gains from using *optimal* incentives are unknown for two reasons. First, characterizing optimal contracts is technically very demanding in most economic environments. Further, quantifying the effects of implementing such a contract would also require knowledge of the model's underlying parameters.

If output is a noisy measure of teacher effort then output-based incentives could be suboptimally strong if, as incentive pay opponents argue, they expose teachers to too much risk. Alternatively, they could be too weak if they do not appreciably change teacher effort inputs, making it hard to discern significant effects. Theoretical work such as Barlevy and Neal (2012), which takes the first step in a multitask model of teacher effort provision, cannot quantify the gains from optimal contracts without taking the second step. At the same time, cleanly identified and precisely estimated causal effects from RCTs cannot speak to the gains from optimal incentives without the additional structure provided by theoretical work. Moreover, since RCT implementation is expensive, experiments typically study only a small number of different levels of incentive strength (i.e., treatment groups). Searching for optimal incentive strength via pure experimentation would be prohibitively expensive, motivating the use of additional structure to maximally leverage findings from RCTs for use in education policy.

This paper takes a step towards filling this gap by using the framework of Hölmstrom and Milgrom (1987), a hidden action, or "moral hazard" model of effort choice, to interpret findings from Muralidharan and Sundararaman (2011), an experimental study of teacher incentive pay implemented in Andhra Pradesh, a state in India. In the model, teachers choose an unobserved effort level, which determines their quality. The main advantage of this model is its closed-form solution of the optimal contract: the optimal incentive scheme is linear in output, which depends on teacher effort and a shock. A larger error variance or higher teacher risk aversion would reduce optimal incentive strength, or slope of remuneration in output. Hölmstrom and Milgrom (1987) is equivalent to the widely used CARA-Normal model, making it a particularly salient example environment.

Muralidharan and Sundararaman (2011) is particularly good for calibrating model parameters because the experiment introduced an output-based incentive scheme that, like the optimal contract in Hölmstrom and Milgrom (1987), is linear in the output of individual teachers. Additionally, the authors estimate a significant effect of individual incentive-pay schemes. Because this paper's goal is to assess potential gains from optimal contracts, it is most natural to focus on a well-designed incentive pay experiment reporting a statistically significant effect.

2 Methods

2.1 Model

This section presents the workhorse CARA-Normal model of moral hazard, as developed in Bolton and Dewatripont (2005). Although this model assumes a linear contract, which need not be optimal, the solution is the same as that in Hölmstrom and Milgrom (1987), which studies a static one-period model split into a number of sub-periods, where in each sub-period an agent (i.e., teacher) controls the probability of success for a binomial random variable. In particular, Hölmstrom and Milgrom (1987) show that the optimal contract features an end-of-period payment that is a linear function of aggregated signals. The interpretation for our education context would be that, in each infinitesimal unit of time, the teacher could exert more or less effort to increase the probability a student obtains a sub-period-specific "bit" of human capital measured by an end-of-year exam.

The administrator has utility $q - w$, where q is output and w is the wage paid to the teacher. The teacher has constant absolute risk aversion (CARA) utility $-e^{-\xi(w-\psi(a))}$, where ξ is their coefficient of absolute risk-aversion and the cost of exerting effort a is $\psi(a) = \gamma a^2/2$. The teacher requires an expected utility of $\underline{u}$ to participate. Output from teacher i depends on their effort according to $q_i = a_i + \eta_i$, where the IID ex-post shock $\eta_i \sim \text{N}\left(0, \sigma_\eta^2\right)$ renders output a noisy measure of teacher effort.

Hölmstrom and Milgrom (1987) show that it is optimal for the administrator to pay the teacher using the linear contract $w = \beta_0 + \beta_1 q$, where β_1 is the share of output paid to the teacher. Therefore, the administrator solves

$$\max_{\beta_0,\beta_1} \ \text{E}_\eta\left[a + \eta - w(a,\eta)\right] \tag{1}$$

$$\text{s.t.} \ \ w(a,\eta) = \beta_0 + \beta_1(a + \eta)$$

$$\text{E}_\eta\left[-e^{-\xi(w(a,\eta)-\psi(a))}\right] \geq \underline{u} \tag{IR}$$

$$a \in \arg\max \ \text{E}_\eta\left[-e^{-\xi(w(a,\eta)-\psi(a))}\right]. \tag{IC}$$

The teacher problem yields a unique optimal effort level $a^* = \beta_1/\gamma$ by differentiating (IC) with respect to effort, and the optimal linear contract sets $\beta_1^* = 1/(1 + \xi\gamma\sigma_\eta^2)$. Therefore, expected output is $\text{E}\left[q^*\right] = \text{E}_\eta\left[a^* + \eta\right] = a^* = 1/\left(\gamma(1 + \xi\gamma\sigma_\eta^2)\right)$. Intuitively, as the signal quality worsens (i.e., σ_η^2 increases) the contract becomes lower powered (i.e., β_1^* decreases), resulting in lower effort a^* and expected output $\text{E}\left[q^*\right]$.

If noise increased, the resulting optimal contract would partially protect a risk-averse teacher by making incentives weaker in output, by reducing the slope of the linear contract

β_1^*. The more risk-averse the teacher, the more protected they would be from fluctuations in η.

2.2 Calibration

I calibrate the model parameters $(\gamma, \xi, \sigma_\eta^2)$ using a "sophisticated" back-of-the-envelope method, which is "sophisticated" because I calibrate using equilibrium implications of the hidden action model. As I show below, values for ξ and σ_η^2 can be obtained either directly from external sources or by transforming external data. However, to calibrate the effort cost parameter γ, we need to know how much teachers respond to incentive pay. Note that the "causal" or composite effect of teacher incentive pay reported in the experimental results could, in theory, also include changes in student and/or family inputs. However, assigning the total effect to changes in teacher effort is consistent with the theoretical model used to interpret these results. Note that effort and output are compared to their baseline levels, i.e., that obtained absent output-based incentives.

Muralidharan and Sundararaman (2011) estimate the effect of an output-based incentive scheme for teachers in the Indian state of Andhra Pradesh, in which teachers were paid according to a linear schedule, 500 rupees per percent increase in mean test scores, for test score gains above 5%. The study covered two years. As with any mapping between theory and data, assumptions have to be made. The benefit of using Hölmstrom and Milgrom (1987) to interpret Muralidharan and Sundararaman (2011) is that the linear scheme employed in the latter affords a clean mapping between their findings and the hidden action model. The same is true of their experimental research design, which obviates having to account for mean differences in output between treatment and control groups being based on selection on hidden types, allowing the calibration to proceed for a representative (average) teacher.[1] I convert currency into U.S. dollars for convenience. While this might raise concerns about external validity, CARA utility implies that risk aversion is independent of wealth, meaning the large wealth differences between teachers in India and the U.S. would only affect the intercept, not optimal teacher effort and output.

There were on average 3.14 teachers and 37.5 pupils per teacher in the incentive schools. Student achievement increased by an average of 0.15 sd, per year. Students' annual wages increased by an average of 2,156 rupees per student[2]; the average cost of the incentive scheme was 20,000 rupees.[3] With a conversion rate of 45 rupees per dollar, this corresponds

---

[1] The linearity of the administrator's objective implies that she can solve a separate problem for each teacher.

[2] See footnote 34 on page 72 of Muralidharan and Sundararaman (2011).

[3] The incentive scheme cost an average of 10,000 rupees for each of two years.

to \$1,796.67 (=\$47.91×37.5) in total output produced by the average teacher and \$141.54 (=\$444.44/3.14) paid to the average teacher. Then, the slope of the contract is the per-teacher income increase (\$141.54) divided by the increase in output (\$1,796.67), or 0.0788; i.e., teachers were paid a piece rate of 7.88% of output.

We can exploit the teacher's optimal choice of action, which solves (IC) in (1) but does not rely on optimality of the slope β_1, to map (β_1, a) to the effort cost γ. The value of γ which rationalizes this increase is then $\gamma = \beta_1/a = 0.0788/1,796.67 = 4.385 \times 10^{-5}$. Teacher risk aversion matters for how incentives are structured (Nadler and Wiswall (2011)). I set the CARA parameter to $\xi = 6.7 \times 10^{-3}$, the mean estimated CARA from the benchmark model of Cohen and Einav (2007), Table 5. This is likely conservative, as Dohmen and Falk (2010) document that teachers are more risk-averse than other workers.

Assuming mean test scores[4] $\overline{y}$ are converted to output via $q = \beta_q \overline{y}$, the conversion factor β_q can be calibrated by noting that the scheme increased mean test scores by 0.15 sd and output per teacher by \$1,796.67, resulting in a conversion factor β_q =\$11,977.78 (=\$1,796.67/0.15). Student j's test score depends on their teacher i's effort and a student-specific shock distributed IID according to $\epsilon_{ji} \sim \mathrm{N}\left(0, \sigma_\epsilon^2\right)$, which captures idiosyncratic factors affecting student achievement on the administered test instrument. The variance of the mean test score then can be computed by dividing the variance of test score error σ_ϵ^2 by the average number of students per teacher in the data, i.e., $\sigma_{\overline{y}}^2$ =0.953/(37.5).[5] To obtain the variance of output σ_η^2 we then square the test-score-to-income parameter and multiply by the variance of mean test score, i.e., σ_η^2 =6,076,631\$^{2}(=\$11,977.78$^2 \times$ 0.953/(37.5)).[6]

3   Results

Using the calibrated parameter values, we can solve for the optimal slope of $\beta_1^* = 0.483$, which is over six times steeper than in the experiment. This results in an optimal effort level/output gain of $a^* = \$11,011.34$, which corresponds to an average increase in student achievement of 0.919 sd. Accordingly, these increases are also more than six times larger than the estimated increases stemming from the much weaker incentives provided under the experiment.

---

[4]Note that everywhere, I refer to test score gains.

[5]Schochet and Chiang (2012) compile estimates of the variances from a large number of studies in their study of error rates in value-added models, providing a good source for typical values for σ_ϵ^2 of $\sigma_\epsilon^2 = 0.953$. Results available upon request.

[6]This is because the variance of q, i.e., σ_η^2, is $\beta_q^2 \sigma_{\overline{y}}^2$.

Sensitivity Analysis Figure 1 presents contour maps of model outcomes for a grid of points covering a wide range of alternative values of σ_η^2 and ξ, ranging from one half to ten times the calibrated value of each parameter.[7] Note that, because γ was recovered using the teacher's effort action choice and can be recovered by using the slope of incentives in the experiment and increase in output, it does not depend on (σ_η^2, ξ). Figure 1a is a contour map of the optimal output share, or β_1^*. Figure 1b is a contour map of optimal output, i.e., $\mathrm{E}\left[q^*\right]$. In both figures, the value corresponding to the calibrated values of σ_η^2 and ξ is indicated by a red dot. We can see that as teachers become more risk averse (increasing ξ) or output becomes noisier (increasing σ_η^2), both incentive strength (Figure 1a) and output gains decrease (Figure 1b). For example, the increase in output ranges from about 1.5 sd in student achievement (at the bottom-left) to around 0.5 sd when teachers are ten times more risk averse than their calibrated value of $\xi = 6.7e - 3$. This latter figure is only about three times the estimated effect of the incentive scheme, but still of considerable magnitude when compared with the effects of other educational interventions, while not being implausibly large.[8] Put another way, teachers would have to be extremely risk averse and/or output would have to be far noisier than typical test instruments to have optimal incentives be even close to as flat as those implemented in Muralidharan and Sundararaman (2011).

Finally, Figure 1c presents a contour map of the expected share of teacher income comprised by variable compensation, i.e., $\mathrm{E}\left[\beta_1^* q^*\right] / \mathrm{E}\left[\beta_0^* + \beta_1^* q^*\right]$. As with the slope and output, this share declines as the output shock variance and degree of risk aversion increase.[9] The optimal expected share of income that is variable pay under the calibrated parameter values would be around 7%.

4 Discussion and Conclusion

This paper produces the first assessment of the potential gains to implementing optimal teacher incentive pay. The findings point to large potential gains to implementing optimal contracts, which are six times steeper than those in the experiment. This finding suggests that the estimated null effect found in many implemented studies of incentive pay (see, e.g., Glewwe et al. (2010) and Fryer Jr (2013)) could potentially be attributed to weaker-than-optimal incentive strength.

The simplicity of this paper's approach allows me to study an environment for which the optimal contract has already been characterized and then use a well-implemented empirical

---

[7]Table 2 in Babcock et al. (1993) shows that a higher-end estimate of ξ is about 0.35, well above the range considered in the parameter grid here.
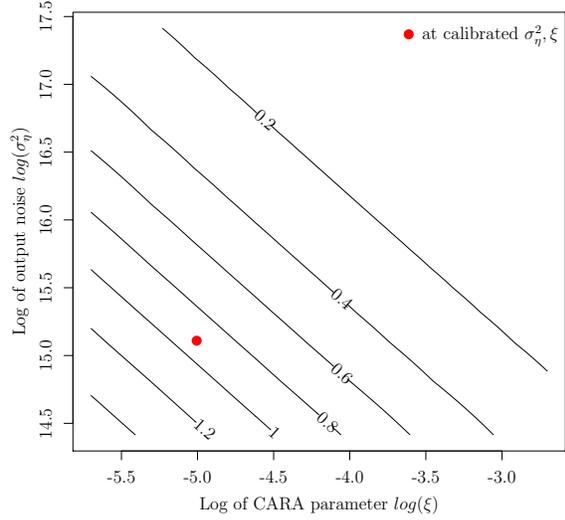
[8]Cohen (1988) classifies gains of 0.80sd and higher as "large".

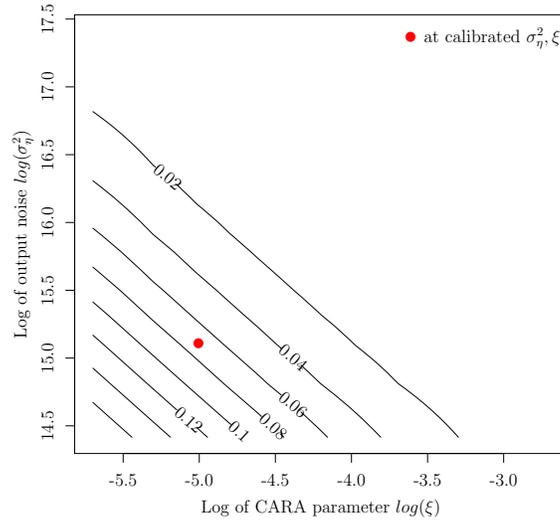[9]This was computed using a certainty equivalent value of $70,000 (Himes (2015)).

Figure 1: Optimal output share and ratio of output for $(\sigma_\eta^2, \xi)-$grid
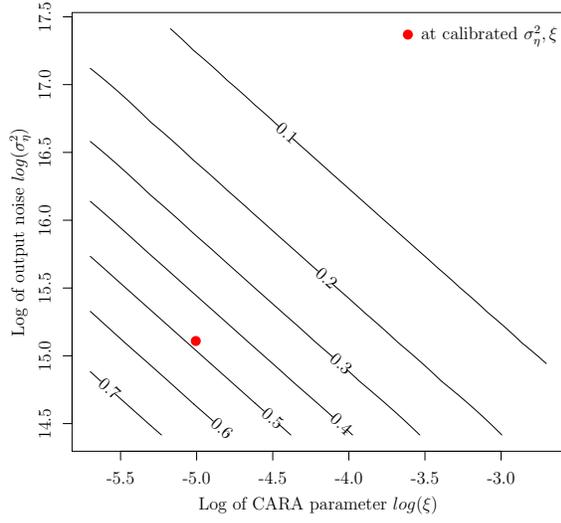
(a) Output share, β_1^*

(b) Output (sd), $\mathrm{E}\left[q^*\right]$

(c)   Variable   share   of   income, $\mathrm{E}\left[\beta_1^* q^*\right] / \mathrm{E}\left[\beta_0^* + \beta_1^* q^*\right]$

study to recover the relevant parameters. It provides an example of the potential gains to adopting optimal contracts in educational production. Calculating the potential output gains from moving to optimal contracts in other hidden action environments (e.g., Imberman and Lovenheim (2015)) constitutes an important avenue for future research.

References

Babcock, B. A., E. K. Choi and E. Feinerman, "Risk and Probability Premiums for CARA Utility Functions," *Journal of Agricultural and Resource Economics*, pp. 17–24, 1993.

Barlevy, G. and D. Neal, "Pay for Percentile," *American Economic Review*, 102(5):1805–31, 2012.

Bolton, P. and M. Dewatripont, *Contract Theory*, MIT Press, 2005.

Cohen, A. and L. Einav, "Estimating Risk Preferences From Deductible Choice," *American Economic Review*, pp. 745–788, 2007.

Cohen, J., *Statistical Power Analysis for the Behavioral Sciences*, vol. 2, Hillsdale, NJ: Lawrence Earlbaum Associates, 1988.

Dohmen, T. and A. Falk, "You Get What You Pay For: Incentives and Selection in the Education System," *Economic Journal*, 120(546):F256–F271, 2010.

Fryer Jr, R. G., "Teacher Incentives and Student Achievement: Evidence from New York City Public Schools," *Journal of Labor Economics*, 31(2):373–407, 2013.

Glewwe, P., N. Ilias and M. Kremer, "Teacher Incentives," *American Economic Journal: Applied Economics*, 2(3):205–227, 2010.

Goldhaber, D. D. and D. J. Brewer, "Why Don't Schools and Teachers Seem to Matter? Assessing the Impact of Unobservables On Educational Productivity." *Journal of Human Resources*, 32(3), 1997.

Hanushek, E. A., "The Economic Value of Higher Teacher Quality," *Economics of Education Review*, 30(3):466–479, 2011.

Himes, T., "LAUSD Educators Typically Earned \$75,504 Last Year," *Los Angeles Daily News*, 2015.

Hölmstrom, B. and P. Milgrom, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, pp. 303–328, 1987.

Imberman, S. A. and M. F. Lovenheim, "Incentive Strength and Teacher Productivity: Evidence from a Group-Based Teacher Incentive Pay System," *Review of Economics and Statistics*, 97(2):364–386, 2015.

Muralidharan, K. and V. Sundararaman, "Teacher Performance Pay: Experimental Evidence from India," *Journal of Political Economy*, 119(1):39–77, 2011.

Nadler, C. and M. Wiswall, "Risk Aversion and Support for Merit Pay: Theory and Evidence From Minnesota's Q Comp Program," *Education Finance and Policy*, 6(1):75–104, 2011.

Rivkin, S., E. Hanushek and J. Kain, "Teachers, Schools, and Academic Achievement," *Econometrica*, 73(2):417–458, 2005, ISSN 1468-0262.

Schochet, P. and H. Chiang, "What Are Error Rates for Classifying Teacher and School Performance Using Value-Added Models?" *Journal of Educational and Behavioral Statistics*, 2012.