The Public-Private Sector Wage Gap: Evidence from Kenyan Teachers *
Preliminary Draft

Nicholas Barton[†], Tessa Bold[‡], and Justin Sandefur[§]

October 28, 2016

Abstract

A large literature demonstrates a wage premium for public employment, with particularly high premiums in many developing economies. The observed wage gap may be efficient – due to compensating differentials, efficiency wages, or selection of candidates into public sector jobs based on unobserved human capital – or an inefficient rent causing labor misallocation. To differentiate between these explanations, we use a regression discontinuity design that exploits the Kenyan government's decision to hire roughly eighteen-thousand new teachers in 2010 based on a simple algorithm ranking applicants within each district. Many unsuccessful candidates found employment as teachers in the private sector. Fuzzy regression discontinuity estimates suggest a causal return to civil service employment of roughly 90%. In combination with data on the relatively low productivity and effort levels of civil service teachers, these results cast doubt on compensating differentials, efficiency wages, or human capital as explanations for Kenya's high public sector wages.

JEL classification:
Keywords: civil servants, teachers, public sector wages, wage gap

*We are grateful to Claudia Lagat and Mukhtar Ogle for helping to organize the survey data collection, and to David Lagakos and Andrew Zeitlin for helpful comments. All remaining errors are ours.

[†]Goethe University Frankfurt: barton@econ.uni-frankfurt.de
[‡]IIES Stockholm: tessa.bold@iies.su.se
[§]Center for Global Development: jsandefur@cgdev.org

1 Introduction

The government is important in providing essential goods and services: rule of law, infrastructure, health and education services. As well as providing goods and services, the government is also a major employer for many economies: in Kenya, roughly 30% of wage workers were employed by the public sector in 2012 (Kenya National Bureau of Statistics, 2014). In many countries, in particular those less developed, the public sector pays higher wages than the private sector. While this may not hold for the whole wage distribution, the average raw differential is found to be positive in favour of the public sector in most developing countries including Zambia at 38%-45% (Skyt-Nielsen and Rosholm, 2001), a range of countries in Latin America, from 40% in Chile to 111% in Colombia (Mizala et al., 2011), and Pakistan at roughly 50% (Aslam and Kingdon, 2009; Hyder and Reilly, 2005).

If this wage differential is due to the fact that the two sectors use different technologies leading to sorting of high-skilled workers into the public sector and low skilled workers into the private sector, then this is not necessarily a concern. In this instance, it is plausible that both sectors pay the marginal product of labor and wage differentials are simply a product of the different human capital levels in the two sectors.

Given the bulk of our knowledge about the public sector in developing countries, this is unlikely to be the case however, with teachers in the public sector being noticeably more absent than those in the private sector as well as being less likely to meet minimum knowledge requirements (Bold et al., 2016). In fact, it is much more likely that the pay gap is the symptom of an inefficiently run public sector that artificially inflates wages - with potentially serious consequences for total factor productivity: Cavalcanti and Rodrigues dos Santos (2015) estimate halving the public sector pay premium in Brazil would increase aggregate output by 10%.

Establishing whether the public pay premium is 'causal' in the sense that otherwise identical workers are rewarded more in the public than the private sector is therefore an important question of allocative efficiency. In this paper, we use a fuzzy regression discontinuity design (RD) to identify this 'causal' public-private pay gap. Specifically, we exploit information about applicants for 18,000 teacher vacancies in Kenyan public primary schools in 2010 and 2011. The Kenyan government organized this hiring exercise by assigning 66 vacancies to each of 250 constituencies and then listing applicants according to a simple algorithm that calculated a score based on time since graduation and grades in teacher training qualifications and secondary school examinations. The hiring rule prescribed that applicants in each constituency were ranked according to their score and those with a rank better than the cut-off, namely 66, should be offered a job. The fixed number of vacancies in each constituency therefore creates a discontinuity in the probability of being hired at the cut-off.[1] The fact that candidates are ranked according to a simple algorithm generating scores that change

[1]Rather than exhibiting a jump from zero to one in the treatment (hiring) status as would be the case in a sharp RD design, the probability of treatment exhibits a significant jump at the cut-off.

continuously at the cut-off, make it plausible that all unobservables related to differences in wages also change continuously. The latter is of course our identifying assumption that allows us to interpret the difference in wages of candidates just above and below the cut-off as causal.

We find a positive public sector premium in raw differences across the whole sample: wages are 50.9% higher for successful applicants working for the Teacher Service Commission (TSC) than for unsuccessful applicants, with the estimate of the local average treatment effect (LATE), i.e. the causal effect of being a civil servant, being higher at 91.8%. The large majority of the unsuccessful applicants work as teachers either in private school or in public schools on contracts similar to those faced by teachers in private schools: identical individuals doing identical jobs are paid different wages. That this is a real inefficiency in the public sector becomes even more obvious knowing that the 'private sector' teachers are at least as effective as the public sector ones (Duflo et al., 2015).

Our work is related to the large literature estimating the public sector premium (including Adamchik and Bedi (2000), Casero and Seshan (2006), Christofides and Pashardes (2002), Depalo et al. (2015), Giordano et al. (2011), Gunderson (1979), Mizala et al. (2011), Postel-Vinay and Turon (2007), Skyt-Nielsen and Rosholm (2001), Terrell (1993), van der Gaag and Vijverberg (1988)). Most of these papers rely on trying to control for as many worker (and firm) observables as possible. To the best of our knowledge, our paper is the first to exploit a fuzzy RD design in the case of the public-private sector wage gap, which we believe makes our interpretation of the estimated public sector premium as causal more convincing. This comes however at the cost of only being able to identify the *local* average treatment effect: namely the premium for newly hired civil service teachers who are marginally hired. We think that this drawback is more than made up for by the cleaner identification we can offer and complements the existing literature.

2 Wage Differentials and Misallocation

There is a wide and varying literature examining wage differentials between industries, genders and sectors. We focus on the difference in wages between the public and the private sectors. While this is mostly found to be positive in favor of the public sector, evidence is not unanimous across all countries. Since Ehrenberg and Schwarz's (1986) review of the literature considering public sector labor markets in the USA, interest in wage determination in the public sector, in particular the differential to the private sector, has grown appreciably. In general there is more work for developed countries, mostly due to data availability. One recent example is Depalo et al. (2015) which reviews data for 10 Eurozone countries finding a premium in favor of the public sector. In light of the prominence of the public sector as a wage employer in developing countries it is not surprising that researchers of the wage gap have increasingly exploited data from developing countries, where the issue of sectoral wage differentials is highly relevant.

The question of why wage differentials exist is an important one and could have a number of theoretical reasons as discussed in Gibbons and Katz (1992). The main distinction they draw is whether the differential is due to differences in workers' productive ability or if identical workers experience a true wage differential. If the latter occurs when the two workers are also carrying out the same tasks, we consider this to be reflective of an inefficiency in the market.

An observed wage differential may or may not be an indication of misallocation depending on the reason for it. One possible explanation would be that workers sort across employers such that those who are more productive earn more. Several early studies[2] on the public-private wage gap estimate wage setting equations, decomposing the wage an employee receives into observed characteristics of the employee (and sometimes the employer) as well as an unobserved component making use of an Oaxaca-Blinder decomposition (Oaxaca, 1973; Blinder, 1973). The part of the wage gap explained by worker characteristics is clearly not a misallocation of resources but rather a logical conclusion, and may result from complementarities between firm and worker characteristics. A second potential reason for a wage gap would be that it is compensating for other undesirable characteristics of a job, such as risk undertaken or unpleasant working conditions. These explanations rely on the employer's characteristics differing between jobs. If workers were carrying out the same tasks and working in the same place of employment, as is the case for teachers in our sample, it is difficult to believe that these would explain a wage gap. It could also be the case that there are non-pecuniary benefits which compensate for the difference in wages. In the case of teachers, civil servants normally receive better pension plans and other benefits which teachers on other contracts do not receive. This would mean that in the case of teachers whom we examine, non-pecuniary benefits (which remain unobserved in the data) would likely strengthen any observed compensation gap in favor of the public sector.

An alternative third explanation for a wage gap would be some form of efficiency wage argument, in which increasing wages increases profit for firms over some range of possible wages (Krueger and Summers, 1988), or schooling outcomes increase in the case of teachers. Such a situation may arise due to a lack of perfect information potentially leading to moral hazard or adverse selection. In the moral hazard case an employer may offer higher wages to stop workers shirking, to attract or retain more productive workers. To prevent shirking, efficiency wages mean the consequence of being fired when not exerting effort would cost an employee more. In the case of civil servants who are virtually unsackable and still shirk to a large extent, this seems a little far fetched. Duflo et al. (2011) find that civil service teachers in their Kenyan sample are only in the classroom 45% of the time they should be while Glewwe et al. (2010) finds that across all teachers there is a 20% absence rate. Bold et al. (2016) find that public sector teachers are absent more often so a difference in wages does not seem to be having a positive effect on absenteeism. The alternative idea from an efficiency wage argument is that of overcoming adverse selection, namely attracting

[2]Ehrenberg and Schwarz (1986) provides a good summary of early work considering the US, and Bender (1998) summarizes much of the early literature focussing on countries beyond the USA

good workers who may otherwise not apply for the job. This appears to be supported in a randomized experiment in Dal Bo et al. (2013), which finds that higher wages seem to attract more able applicants and increase acceptance of job offers. Where informational frictions are removed, such 'efficiency wages' would in fact be inefficiently high and wages would depend only on a worker's characteristics for otherwise identical jobs. A fourth, and particularly attractive explanation in the case of the public sector (Ehrenberg and Schwarz, 1986), is that union power allows one group of workers to push wages to an artificially high level, accruing rent.

Paying rents may also be politically desirable if they allow politicians to give public sector jobs, which pay relatively higher wages, to their own potential supporters in order to secure more votes. In situations where direct transfers as a means of redistribution would face severe political opposition, it may be more politically feasible to use less obvious methods of patronage/clientelism to achieve political aims. According to Coate and Morris (1995), redistribution to interest groups may be undertaken using inefficient methods where voters are not perfectly informed about the effects of policy and the predispositions of politicians. Robinson and Verdier (2013) develop a model in which jobs allow the giver and receiver of patronage to enter in to a bilateral commitment, where employment is credible, selective and reversible. They find that such a system of clientelism leads to inefficiencies in the supply of public goods, and that such a system is more attractive if inequality is high or productivity is low. Empirical evidence that politicians use the allocation of public sector jobs for political reasons is abundant. Alesina et al. (2001) estimate that half the public wage bill in the South of Italy can be interpreted as subsidy from the richer North to the poorer South, both because of an overinflated size of the public sector and through a public wage premium compared with outside alternatives. Calvo and Murillo (2004) consider the use of public service jobs in Argentina and argue that the poorer and less educated people are more susceptible to clientelism, as they need to be paid a lower wage premium compared to their outside option. They find that in provinces governed by the Peronist party, whose supporters are traditionally less educated, there is a higher wage premium paid to low skilled workers, and that 21% more public sector workers are employed than in provinces governed by the Union Civica Radical, a party more associated with the middle class.

If wages are set optimally for a firm motivated by profit, wages will equal marginal productivity. The logical conclusion of this fact would mean that two identical workers who are carrying out identical tasks should produce the same value of output. If they are earning different amounts, this may be indicative that they are not free to move in such a way as to equalize their marginal productivities thus reducing total factor productivity of the economy. This misallocation of workers will lead to a loss of output for the whole economy, reducing total welfare. The issue of inputs not being efficiently allocated are discussed in Cavalcanti and Rodrigues dos Santos (2015)[3]. They include a literature review covering possible reasons that marginal productivities are not equalized, including credit market imperfections and

[3]Restuccia and Rogerson (2013) provide a more in depth overview of the literature on misallocation and productivity

frictions, taxes and regulations, and trade policies among others. Their model includes an endogenous choice of sectors and analyzes the effect of a wage premium on aggregate output, finding a strong negative effect when estimating the model with Brazilian data.

Why should wage gaps indicate or cause inefficiency? The fact that wages are not equal across sectors may be indicative of restricted movement between sectors preventing workers being able to work where they are most efficient. This is borne out in the observed job queueing which occurs in many countries. Job queueing could in theory be a result of workers equalizing expected public sector wages with private sector wages (Upadhyay, 1994) and so remaining unemployed while waiting (often called wait-unemployment). While individually rational, this is inefficient for the economy as workers waiting for a better job are not producing anything. In fact, the larger the gap, the longer potential employees are willing to wait (Boudarbat, 2004). Queues could also cause problems if the incentives of the bureaucrats are not aligned with those of the nation as a whole (for whom it would be optimal to maximise output under a given resource constraint). The bureaucrats may have incentives to hire people from the queue who command a lower wage (Mengistae, 1999), and are then likely to be less productive. This would lead to more well qualified people working in the private sector and lower qualified people enjoying the wage premium of the public sector.

The political motivation of government wage setting exacerbates this problem in that governments often want to be seen as a fair employer and so are prepared to pay over the odds for less well qualified workers, while they don't want to be seen to be contributing towards inequality and so are unwilling to pay extremely high wages. This leads to the compressed wage distributions observed in the data, which we later return to, and likely leads to worse quality civil servants being hired. Alternatively, the higher wages may attract people who have very little public service motivation and whose values are not aligned with the aims of their employer. Handy and Katz (1998) propose a model for non-profit organizations which offer lower wages with the aim of positive self-selection. Teachers often seem unmotivated, in particular civil servants earning high wages who we have seen to be so often absent. This would fit with the argument of Handy and Katz (1998). While mission alignment may lead to better outcomes, this is far from clear cut in the existing evidence and so it is not obvious that offering high wages does harm through motivation rather than attracting better qualified applicants.

Misallocation of resources can be problematic for productivity, as shown in Hsieh and Klenow (2009) who find that by reallocating capital and labor to equalize marginal products, there would be total factor productivity gains of 30%-50% in China and 40%-60% in India. Hsieh and Klenow (2010) suggest that TFP may explain 50%-70% of differences in income between countries and that misallocation may play an important role in the differences of TFP, though they remain quiet on the sources of this misallocation. More specific to the labor market is Hsieh et al. (2013), in which the authors use a Roy model (Roy, 1951) to study the inefficiency of misallocation of productive black and female labor. They find that frictions in the model account for 15% to 20% of growth in aggregate output since 1960. In

an alternative model for the size of the public service, Jaimovich and Rud (2014) show that when bureaucrats with low public service motivation are attracted to the public sector by high wages, they may in turn hire many low-skilled workers thus causing productivity to fall.

A number of papers have investigated the wage differential between the public and private sectors, using various empirical techniques. The most simple is to run an OLS regression for wages including an indicator variable for working in the public sector (as described in Ehrenberg and Schwarz (1986)). This imposes the restriction that other included characteristics have the same impact on wages in both sectors and also fails to take selection of employees into different sectors into account. Many other papers attempt to break down the raw difference into observable and unobservable characteristics using an Oaxaca-Blinder decomposition. Some then extend this with a two stage estimation to allow for selection into the worker's chosen sector. This can be accounted for either in a Heckman two step selection model (Heckman, 1979), in essence instrumenting for the sectoral selection, or with a structural switching model which simultaneously models sector selection with the wage equations. Sample selection correction was used in Terrell (1993) for data from Haiti and Casero and Seshan (2006) in Djibouti, who both discovered positive rents accruing to public sector workers. In the wider literature, Abowd et al. (1999) exploited longitudinal data with matched firms and workers, to not only find the effects of worker characteristics on wages, but also firm effects. They found that individual effects and not firm effects were driving wage differences between firms. Unfortunately, their dataset does not include the public sector, so given the belief that the public sector is somehow structurally different, their results are perhaps not directly transferable to our setting, though the results do contradict the early findings of Krueger and Summers (1988) and Gibbons and Katz (1992). Due to a paucity of data, their technique is not often applicable to our setting. As an alternative to a two step selection model, van der Gaag and Vijverberg (1988) were among the first to use a switching model[4] for sector specific wage determination models, and find that simple OLS regressions lead to upwardly biased results for the public sector in their sample from the Côte d'Ivoire. These different techniques are often combined with quantile regressions in order to investigate the importance of working in the public sector at different points of the earning distribution. Papers using this technique, such as Skyt-Nielsen and Rosholm (2001) in Zambia and Hyder and Reilly (2005) in Pakistan, often find that the public sector wage distribution is more compressed, with those at the lower end earning a higher public sector premium and those at the upper end of the distribution earning a lower (or even negative) premium.

While these previous papers attempt to control for selection into each of the two sectors, either by the worker or the employer, this is difficult to control for completely. We exploit a shock in the demand for civil service teachers to find a LATE of working in the public sector as a teacher. We find what we argue to be a causal effect of simply being employed by the public sector for a marginally hired teacher.

[4]As first used by Lee (1978) in his investigation of the role of unions in wage setting

3 Background

In January 2003 the Free Primary Education (FPE) act was introduced in Kenya in line with one of the more prominent aims of the Millennium Development Goals; to achieve universal primary education. UNESCO (2011) estimate that net enrolment in Kenya increased from 62% in 1999 to 76% in 2010. This occurred alongside a stop placed on expanding the civil service teacher workforce between 1998 and 2010. These factors combined led to increasing student-teacher ratios, which is believed to be detrimental to student outcomes (Angrist and Lavy, 1999). In 2009, the primary student-teacher ratio was at 47:1, which is significantly above the government target of 40:1. In fact in 2011 there was perceived to be a shortfall of 61,000 teachers (Bold et al., 2013).

In the face of this problem and lacking the funds to hire enough civil service teachers, Parent Teacher Associations (PTAs) as well as the government hired some teachers on fixed term contracts, paying them considerably less than the civil servants as well as giving them worse contractual conditions and non-pecuniary benefits. According to Duflo et al. (2015) PTA teachers received roughly $25 per month while civil service teachers received about $120 per month as well as other benefits including a housing allowance, medical insurance and a pension. In 2010 the decision was made to hire additional civil service teachers, rather than simply replacing the existing number of teachers. 18,000 new teachers were hired with 66 vacancies allotted to each district regardless of size and existing pupil-teacher ratios. In some cases districts were split into sub-districts, with the sum still adding up to 66.

Applicants were required to be fully qualified to teach at primary school level and were ranked according to a well defined algorithm, based on a combination of how long ago the applicant had qualified and the grade obtained in the teaching qualification. Where these were equal, KCSE (Kenyan Certificate of Secondary Education) grades were used to separate candidates. We use the ranking (normalized by subtracting the cut-off) as our running variable in a regression discontinuity set-up. Given that applicants have to present their qualification certificate upon application, which includes the information for the ranking, an applicant cannot affect their own rank after the ranking system has been announced. This should make the hiring of civil service teachers a suitable setting in which to use RD.

From the list of applicants, survey enumerators were instructed to interview 40 applicants in each district, 20 above and 20 below. Where an individual was not reached, the next applicant away from the cut-off on the list was interviewed until at least 40 applicants had been interviewed in each district. In our regressions we use all the data collected, which is naturally bunched around the official cut-off (see Figure 2).

An additional challenge to our set-up was the hiring of more teachers in 2011. Some districts only hired in 2011 and some hired additional teachers at this point. Clearly, teachers hired in 2011 also enjoyed the benefits of civil service status when interviewed in 2012. Thus, whether they were hired as a civil servant at all is the more relevant question for the issue of a wage gap between those employed as civil servants and those not. We follow Urquiola (2006) in dealing with the question of what cut-off to use, as explained in section 5.

4 Descriptive Statistics

Our data include a total of 1597 applicants, some of whom were hired in 2010, some in 2011 and some were not employed by the TSC at the time of the survey. We do not know whether 448 were hired by the TSC from their survey responses. 810 respondents were either hired in 2010 (704 observations) or hired in 2011 if they had not been in 2010 (106 observations). Some of the 810 who were hired by the TSC in the full sample reported that they are no longer working for the TSC, thus the final number working for the TSC is slightly lower at 803. The fact that so few leave their position working for the TSC is indicative that these civil service jobs are highly sought after, potentially due to the accrual of rents by civil service teachers.

To consider the question of the wage premium received from working as a civil servant, we must restrict our sample to wage earners. While there is clearly a value to non-wage work, we have no means by which to assess this and so choose to compare the income of TSC teachers with others who hold jobs which pay a regular wage. When we additionally restrict the sample to those with completed responses for all variables used in our analysis, we have 764 observations remaining, whose descriptive statistics are shown in Table 1. While the non-response may not be random, there is no reason to believe that this would change discontinuously at the cut-off and so our identifying assumption should remain unharmed.

Panel A displays the entire sample we use, which is then subsequently divided into subsets according to occupation stated in the survey. Of our sample, 584 were hired by the TSC. Of those 180 not hired by the TSC, 156 are working as a teacher under some other form of contract (i.e. in the private sector, whether at a private school or locally hired in state run schools). Of the 24 not working as a teacher, 19 are undertaking wage work, 2 are self-employed and 3 are farming.

Before considering the causal impact of being hired as a civil servant it is useful to compare the mean values of income (column 2) and other variables for those working as a civil servant hired by the TSC and those not working as a civil servant. Table 1 shows that the mean of monthly income is higher for those with a TSC job (13829 KSh) than those without (9972 KSh), while interestingly those not teaching (Panel C) have a higher mean income (11804 KSh) than those teaching without civil service status (9690 KSh). While only suggestive due to the small sample of those outside the TSC, it appears that alternative employment may be more lucrative in the short run for those failing to receive a post from the TSC. However, those who failed in their initial application may be working as a teacher to remain in the queue to become a TSC teacher, thus maximising their lifetime income. The civil service teachers are earning roughly $130 on average, which is in line with the $120 per month as reported in Duflo et al. (2015). There do not appear to be many major differences in the other variables. Individuals hired by the TSC are on average slightly older, which should be expected given that years waited plays such an important role in the hiring procedure. Among those hired by the TSC there is a higher proportion of women than in the subset not hired. A similar proportion of applicants is married and the household size is also similar in

the panels A-C where the sample size has not dropped to such low levels. It appears that the applicants are similar in most aspects other than their income. Any differences that do exist do not pose a threat to our identification in the RD strategy, explained later, as long as they are continuously changing at the cut-off.

5 Empirical strategy

Where a treatment is assigned randomly, it is possible to use OLS to estimate the effect of a treatment, in our case being hired by the TSC. Deciding randomly who receives a job as a teacher is clearly neither feasible nor morally desirable, given the high stakes at play in educating the next generation. An OLS regression of the following form will then indicate the raw differences in wages between sectors, without any causal interpretation. We first estimate the following regression with OLS to gain an accurate idea of the raw wage gap between the public and private sectors.

$$Y_i = \beta_0 + \beta_1 \text{TSC}_i + \varepsilon_i \tag{1}$$

Where Y is the outcome of interest (i.e. some measure of income) for applicant i, TSC is indicator variable for working as a civil service teacher. Treatment is likely to be correlated with unobserved variables leading to a bias in $\hat{\beta}_1$ if trying to find a causal effect. As previously discussed, if there is positive selection of the public sector, ability of the teachers will influence their income along with other characteristics. As we have seen in section 2, previous work investigating the public private wage gap has attempted to control for many different characteristics and investigate how these contribute to a wage setting equation in each sector. We feel that controlling for all selection between the sectors is difficult and so believe that exploiting a fuzzy RD set-up to find a LATE on the income of newly hired civil service teachers is a useful contribution.

5.1 Selection

A fuzzy RD is estimated as a two stage least square regression (Angrist and Lavy (1999); Lee and Lemieux (2010)) predicting the treatment in the first stage. In our case, the first stage alludes to the process of hiring which is also a topic of interest. District education officials seem to deviate widely from the public hiring algorithm described above. If it were perfectly observed, all those better than the cut-off would be hired and employed as civil servants, while all those below would not be. This may not be the case for several potential reasons. People above the cut-off may reject the job (or move on having already been employed), or they may never be offered the job in the first place. The former would not be a problem as far as following the rule is concerned, while the latter may be suggestive of ulterior motives of those deciding who receives a job offer. Unfortunately, we have no way of telling which of these occurs. Those beyond the cut-off may be employed to fill the spaces unfilled by those

above legitimately not taking the job, or they may have been employed in spite of their rank due to some other unknown reasons.

An additional issue in our data is that we only have the ranking information for 2010, but several of those in our sample were employed in 2011. Some districts hired only later, and so the ranking for 2010 represents the applicants for those positions filled in 2011, while some other districts hired both in 2010 and 2011, when there was likely another round of applications. Given the weight placed on years waited in the ranking algorithm, it stands to reason that those just beyond the cut-off in the initial round would now be near the top of a later ranking and so are highly likely to be employed. We want to consider the effect of currently being a civil servant as this is the factor that will affect the wages of subjects in our sample, and so we follow Urquiola (2006) in using the cut-off found in the data (rather than the officially prescribed cut-off) for our main regressions in the following section.[5] The official cut-off, includes all those ranked 66 or below (or the sub-district relevant value), and can be seen as region A in Figure 1. The 2010 "de facto" (2010 DF) cut-off includes up to the obvious cut for those hired in 2010 only and so includes regions A and B. This then takes into account that some applicants who were offered a job as a civil service teacher may have rejected the offer, and so some individuals slightly beyond the original official cut-off then took their place. This means that for the RD regressions using the 2010 DF cut-off, we are essentially finding a LATE at the boundary between B and C. For the 2011 DF cut-off, all applicants hired in 2010 or 2011 are predicted in the first stage, which includes regions A, B and C in Figure 1 and so the LATE is found at the boundary between C and D. The use of 2011 DF therefore includes all applicants who could be in TSC employment when the survey was carried out. The use of the 2011 de facto cut-off is positively correlated with using the official cut-offs with a correlation coefficient of 0.59.

Figure 1: Applicants and cut-offs

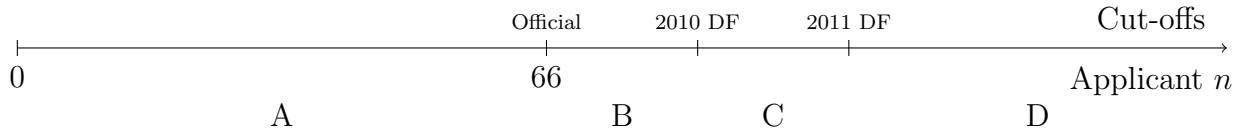

5.2 Wage Gap

Our primary interest is in the impact of holding a civil service job on wage income. To estimate this impact, we regress income (Y_i) on an indicator variable for civil service employment ($\text{TSC}_i \in \{0, 1\}$) using two stage least squares, controlling for each candidate's rank

[5] We investigate for each district which cut-off best predicts a change in the probability of being hired, and thus the most noticeable cut-off.

as described above, and the interaction of civil service employment and rank. From this we can see the relationship between working as a civil servant and income.

$$Y_i = \beta_0 + \beta_1 \text{TSC}_i + \beta_2 \widetilde{\text{Rank}}_i + \beta_3 \widetilde{\text{Rank}}_i \times \text{TSC}_i \tag{2}$$

Here the tilde denotes that the rank variable has been recentered at the hiring cutoff, ($\widetilde{\text{Rank}}_i \equiv \text{Rank}_i - \text{Cut-off}$). Thus the coefficient β_1 can be read as the difference between civil servants and non civil servants.

An intention to treat or sharp regression discontinuity would simply use a dummy for being above the cut-off as the main regressor of interest. However, due to imperfect compliance with the hiring algorithm, we have to use fuzzy regression discontinuity. Treatment is potentially endogenous here, as district officials may deviate from the hiring rule to select candidates based on characteristics correlated with our outcome of interest. To overcome this endogeneity, we use the cut-off as an instrument for observed hiring and estimate instrumental variables (IV) by two stage least squares. This first stage regression is:

$$\text{TSC}_i = \beta_0 + \beta_1 \text{I}_{\text{rank} \leq \text{ cut-off}} + \beta_2 \widetilde{\text{Rank}}_i + \beta_3 \widetilde{\text{Rank}}_i \times \text{I}_{\text{rank} \leq \text{cut-off}} \tag{3}$$

We instrument for TSC as well as for the interaction term with the normalized rank. The predicted values are then used in thr second stage regression, Equation 2.

In order to use RD we rule out that any treatment effect is being driven by discontinuities in other variables at the cut-off. It can be seen in Figure 3 that there are no clear discontinuities in the other variables we observe. While this does not rule out discontinuities in unobserved variables potentially driving any treatment effect, we feel that the lack of any observed discontinuities is an indication that other potential influences are not varying at the cut-off we use. Of the observed variables it is particularly reassuring that the parent teacher association points are not discontinuous at the cut-off. If the PTA had a strong influence in the hiring decision, despite not being part of the official ranking, then this could give support to the possibility of manipulation at the cut-off. We do not actively restrict our bandwidth to a narrow interval around the cut-off (though the sampling methodology will lead to most of the information being taken from around the cut-off) and so include additional specifications with demographic variables as controls.

6 Results

To take into account the second round of hiring in 2011, our preferred specification is to use the cut-off most apparent from the data; the "de facto" cut-off (DF). We start by showing the results from the first stage regression where the discontinuity of hiring probability underlines our assertion that fuzzy RD is a suitable method to analyze this data. We then estimate the effect of working as a civil service teacher on income and employment for the marginally hired teacher.

6.1 Selection

Using our preferred specification of the de facto cut-off from 2011, it can be seen in Figure 4 that there is a jump in the probability of currently working as a TSC teacher at the cut-off. This result is upheld in the results for the first stage regression shown in Table 2. Having a rank lower than the cut-off (i.e. being better than the cut-off with a negative normalized rank) increases the probability of being hired by around 20% (24% for the 2011 de facto cut-off and 17% for the 2011 de facto). If the rule were being strongly adhered to, we would expect this to take a higher value but it is nonetheless positive and highly significant, allowing us to continue with our fuzzy RD analysis. Additionally, this first stage regression for 2011 DF has an F-statistic of 15.77, which allows us to reject the null that the first stage is jointly insignificant and lay concerns of weak instruments aside. The coefficient for the official cut-off is positive as expected but insignificant. This is probably due to a number of people just beyond the official cut-off being hired, after those initially offered jobs turned the job offer down.

6.2 Wage Gap

The contracts offered to civil service teachers are often vastly superior to those received by other teachers, or even those who are unsuccessful and no longer teach. Income is measured on a monthly basis in Kenyan Shillings. We run regressions with the income in Kenyan Shillings (columns (1)-(5)) and the log of this monthly income (columns (6)-(10)). Table 3 shows the result from five specifications for each of income and the logarithm of income: a simple OLS as in Equation 1 in columns 1 and 6, an IV using the de facto cut-off for hiring in 2010 both with (column 2 and 7) and without (column 3 and 8) additional control variables, as well as with the de facto cut-off for hiring in 2011 with (column 4 and 9) and without (column 5 and 10) additional control variables. The control variables are standardized with mean zero and standard deviation equal to one, such that the constant value reported is still the average (log) income for those (predicted to be) not hired by the TSC. The IV regressions are estimated using two stage least squares from Equation 3 and Equation 2. As well as instrumenting for the TSC status of a teacher, we also instrument for the interaction with the adjusted rank variable.

From the OLS with only TSC status as a right hand side variable income is found to be roughly 4215 KSh higher for public sector employees. This is perfectly in line with the result in the descriptive statistics (see Table 1) and as can be seen in the log income regression corresponds to a roughly 50% increase in wages. The result of this very basic OLS with only TSC status enables us to estimate the raw wage gap alongside our income equations.

In our preferred RD specification found in column 9 of Table 3, we see a positive effect on income. While the size of the coefficients seems to have increased when considering the LATE for a marginally hired teacher, the noise has also increased considerably. While it is not possible to claim with any confidence that the causal effect on income is stronger than

the raw difference, it seems clear that there is a positive causal effect on the wage from working as a public sector teacher. This holds for the monthly income specification and the log income specification. Those with a public sector job appear to be earning about 92% more than those in a private sector job. Column 7 of Table 3 shows the same regressions with the cut-off apparent in the data only for those hired in 2010. The coefficients on being a public sector teacher are smaller here, but this is likely due to the fact that some applicants were hired later in 2011 and so are in fact now working in the public sector earning a higher wage. The results from the regressions without additional control variables are remarkably similar to those with control variables again supporting our argument that other variables are not influencing our locally causal results at the cut-off. It could be argued that we are in fact comparing the wage with a permanent contract with the wage of a temporary contract. Theoretically, if the same individual prefers to have a permanent contract, as would seem logical, then he must be compensated with a higher wage to take the temporary contract. Given that the wage gap is positive for those with the permanent contract, this suggests that such a compensating differential argument is not at play here.

7 Robustness

The aforementioned results were for all individuals observed in our data. We are interested in the inefficiency of the public-private sector wage gap. If the regressions from Table 3 are rerun with the restriction that all those not working for the TSC are still working as teachers, and so employed in as close to an identical job as could be hoped for, the previous results still hold and are in fact even stronger. This can be seen in Table 5, where the monthly causal difference in income is more strongly significant and larger at 12271KSh instead of 8978KSh in column 4 of Table 5 and Table 3 respectively. The fact that applicants stay working as a teacher despite seemingly earning less than their outside option is suggestive that they are either motivated to work as teachers, or believe that they stand a better chance of becoming a civil service teacher by staying in the teaching profession than if they leave the market. The latter would make sense in the context that applicants are scored with a heavy weight given to the number of years waiting since qualifying.

A typical issue when using RD is the question of what the bandwidth should be. For the OLS and IV specifications with controls in Table 3, we run sensitivity checks as to how the coefficients react to varying bandwidths. These can be found in Figure 5 and find that the coefficients do not change much when the bandwidth is changed. When the bandwidth closes around the cut-off, the noise dramatically increases and makes it increasingly difficult to reject the null hypothesis that the coefficient is equal to zero. In spite of this, the character of the result for income holds and it still appears that there is a positive effect of being employed by the civil service.

One potential criticism of our results would be that those with higher income are more likely to report their income in a survey and so those earning less will not report their income.

In this case, public sector workers should be less likely to withhold information about their income given that they on average earn more. If this were true and there were private sector workers with an income which they do not disclose, then it is likely to strengthen the causal difference between public and private sector workers. Alternatively, those with higher incomes may under report their income. Again this would mean that the wage gap is larger than that which we find. More problematic would be if people earning a reasonably high regular monetary income, from sources outside the TSC, do not report this. If there were more individuals earning a high enough income to partially close the income gap we observe, then them not reporting their income would lead to an overestimation of the income gap. There is no reason to believe this to be the case.

We also consider the question of whether coethnics of the president or minister of education at the time of hiring are more likely to be hired. Ethnicity is only one characteristic which could be relevant in trying to discern whether the wage gap arises due to patronage. Including a dummy for individuals equal to one when the applicant is the same ethnicity as an individual in power in the first stage hiring regression will give a first indication of whether those in power are receiving preferential treatment and that civil service teaching jobs are therefore being used to redistribute resources to their own ethnicity. Not all individuals answered the question about their ethnicity, so our sample size is slightly reduced. The first stage results with these two additional dummies included can be seen in Table 4. Being a coethnic of the minister of education does not seem to have any effect, but being a coethnic of the president appears to be negatively correlated with the probability of receiving a civil service job. This result is unexpected and suggests that coethnics of the president may even be discriminated against in the hiring process. However, why this is the case is a question we cannot answer. It may be the case that these individuals have better outside options due to their ethnicity and so are less likely to accept a job teaching. Nonetheless, it is at least indicative that patronage is not the main reason for the wage gap.

8 Conclusions

In this paper we exploit an unexpected policy change in the hiring of teachers in Kenya to investigate the causal impact on wages of working in the public sector. If identical workers who are carrying out the same tasks are earning vastly different incomes, this is likely indicative of or a consequence of an inefficiency in the labor market. By using a RD set-up, we are able to estimate locally the causal effect of being a civil servant on teacher wages, and so can rule out that the positive public sector wage gap, which we find, is due to sorting or adverse selection. The previous literature has attempted to control for selection either by use of an instrument in a Heckman selection model or with a two sector structural estimation, but we feel that by using a RD set-up, we feel that we more credibly control for any selection and find a causal effect. In an additional check we restrict our sample to only teachers. We can thereby ensure that the tasks undertaken are as close to identical as could be feasibly hoped for and so effectively rule out that the wage differential is due

to compensating differentials. Given the problems of absenteeism observed in the Kenyan teaching force (Bold et al., 2016), it also seems unlikely that the higher wages are being used to avoid shirking as is true under a moral hazard argument. The possible explanations which survive these successive exclusions are based on the extraction of rents, either due to union powers in the public sector or due to clientelism and other political reasons. These surviving explanations are those which indicate that the difference in wages is likely to be inefficient. We believe that we have found strong evidence for a causal difference in wages between the public and private sector for those marginally hired which is indicative of inefficiency in the public sector in Kenya.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

Adamchik, V. A. and A. S. Bedi (2000). Wage differentials between the public and the private sectors: Evidence from an economy in transition. *Labour Economics 7*(2), 203–224.

Alesina, A., S. Danninger, and M. Rostagno (2001). Redistribution through public employment: The case of italy. *IMF Staff Papers 48*(3), 447–473.

Angrist, J. D. and V. Lavy (1999). Using maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics 114*(2), 533–575.

Aslam, M. and G. Kingdon (2009). Public-private sector segmentation in the pakistani labour market. *Journal of Asian Economics 20*, 34–49.

Bender, K. A. (1998). The central government-private sector wage differential. *Journal of Economic Surveys 12*(2), 177–220.

Blinder, A. S. (1973). Wage discrimination: Reduced form and structural estimates. *The Journal of Human Resources 8*(4), 436–455.

Bold, T., D. Filmer, G. Martin, E. Molina, B. Stacy, C. Rockmore, J. Svensson, and W. Wane (2016). What do teachers know and do in primary schools in sub-saharan africa? *Preliminary Draft*.

Bold, T., M. Kimenyi, G. Mwabu, A. Ng'ang'a, and J. Sandefur (2013). Scaling up what works: Experimental evidence on external validity in kenyan education. *Center for Global Development Working Paper 321*.

Boudarbat, B. (2004). Employment sector choice in a developing labor market. *Centre of Labour and Empirical Economics Research Working Paper*.

Calvo, E. and M. V. Murillo (2004). Who delivers? partisan clients in the argentine electoral market. *American Journal of Political Science 48*(4), 742–757.

Casero, P. A. and G. Seshan (2006). Public-private sector wage differentials and returns to education in djibouti. *World Bank Policy Research Working Paper* (3923).

Cavalcanti, T. and M. Rodrigues dos Santos (2015). (mis)allocation effects of an overpaid public sector. *REAP - Rede de Economia Aplicada Working Paper 82*.

Christofides, L. N. and P. Pashardes (2002). Self/paid-employment, public/private sector selection, and wage differentials. *Labour Economics 9*(6), 737–762.

Coate, S. and S. Morris (1995). On the form of transfers to special interests. *Journal of Political Economy 103*(6), 1210–1235.

Dal Bo, E., F. Finan, and M. A. Rossi (2013). Strengthening state capabilities: The role of financial incentives in the call to public service. *The Quarterly Journal of Economics 128*(3), 1169–1218.

Depalo, D., R. Giordano, and E. Papapetrou (2015). Public–private wage differentials in euro-area countries: Evidence from quantile decomposition analysis. *Empirical Economics 49*(3), 985–1015.

Duflo, E., P. Dupas, and M. Kremer (2011). Peer effects, teacher incentives, and the impact of tracking: Evidence from a randomized evaluation in kenya. *American Economic Review 101*(5), 1739–1774.

Duflo, E., P. Dupas, and M. Kremer (2015). School governance, teacher incentives, and pupil–teacher ratios: Experimental evidence from kenyan primary schools. *Journal of Public Economics 123*, 92–110.

Ehrenberg, R. G. and J. L. Schwarz (1986). Chapter 22 public-sector labor markets. In Ashenfelter, Orley. Ed and R. Layard (Eds.), *Handbook of labor economics. 2 vols*, Volume 2 of *Handbook of Labor Economics*, pp. 1219–1260. Amsterdam and Oxford: North-Holland.

Gibbons, R. and L. F. Katz (1992). Does unmeasured ability explain inter-industry wage differentials. *The Review of Economic Studies 59*(3), 515–535.

Giordano, R., D. Depalo, M. Coutinho Pereira, B. Eugène, E. Papapetrou, J. J. Perez, L. Reiss, and M. Roter (2011). The public sector pay gap in a selection of euro area countries. *ECB Working Paper Series* (1406).

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

Gunderson, M. (1979). Earnings differentials between the public and private sectors. *The Canadian Journal of Economics 12*(2), 228–242.

Handy, F. and E. Katz (1998). The wage differential between nonprofit institutions and corporations: Getting more by paying less? *Journal of Comparative Economics 26*, 246–261.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica 47*(1), 153–161.

Hsieh, C.-T., E. Hurst, C. Jones, and P. Klenow (2013). *The Allocation of Talent and U.S. Economic Growth*. Cambridge, MA: National Bureau of Economic Research.

Hsieh, C.-T. and P. Klenow (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Hsieh, C.-T. and P. J. Klenow (2010). Development accounting. *American Economic Journal: Macroeconomics 2*(1), 207–223.

Hyder, A. and B. Reilly (2005). The public and private sector pay gap in pakistan: A quantile regression analysis. *The Pakistan Development Review 44*(3), 271–306.

Jaimovich, E. and J. P. Rud (2014). Excessive public employment and rent-seeking traps. *Journal of Development Economics 106*, 144–155.

Kenya National Bureau of Statistics (2014). Kenya facts and figures 2014.

Krueger, A. B. and L. H. Summers (1988). Efficiency wages and the inter-industry wage structure. *Econometrica 56*(2), 259–293.

Lee, D. S. and T. Lemieux (2010). Regression discontinuity designs in economics. *Journal of Economic Literature 48*(2), 281–355.

Lee, L.-F. (1978). Unionism and wage rates: A simultaneous equations model with qualitative and limited dependent variables. *International Economic Review 19*(2), 415–433.

Mengistae, T. (1999). *Wage Rates and Job Queues: Does the Public Sector Overpay in Ethiopia?* The World Bank.

Mizala, A., P. Romaguera, and S. Gallegos (2011). Public–private wage gap in latin america (1992–2007): A matching approach. *Labour Economics 18*, S115–S131.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review 14*(3), 693–709.

Postel-Vinay, F. and H. T. Turon (2007). The public pay gap in britain: Small differences that (don't?) matter. *The Economic Journal 117*, 1460–1503.

Restuccia, D. and R. Rogerson (2013). Misallocation and productivity. *Review of Economic Dynamics 16*(1), 1–10.

Robinson, J. A. and T. Verdier (2013). The political economy of clientelism. *The Scandinavian Journal of Economics 115*(2), 260–291.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers 3*(2), 135–146.

Skyt-Nielsen, H. and M. Rosholm (2001). The public-private sector wage gap in zambia in the 1990s: A quantile regression approach. *Empirical Economics 26*.

Terrell, K. (1993). Public-private wage differentials in haiti. *Journal of Development Economics 42*, 293–314.

UNESCO (2011). UNESCO institute for statistics: Statistics in brief: Education(all levels) profile - kenya.

Upadhyay, M. P. (1994). Accumulation of human capital in ldcs in the presence of unemployment. *Economica 61*(243), 293–314.

Urquiola, M. (2006). Idntifying class size effects in developing countries: Evidence from rural bolivia. *The Review of Economics and Statistics 88*(1), 171–177.

van der Gaag, J. and W. Vijverberg (1988). A switching regression model for wage determinants in the public and private sectors of a developing country. *The Review of Economics and Statistics 70*(2), 244–252.

9 Appendix

Figure 2: Distribution of surveyed individuals by 2011 De Facto Cut-Off

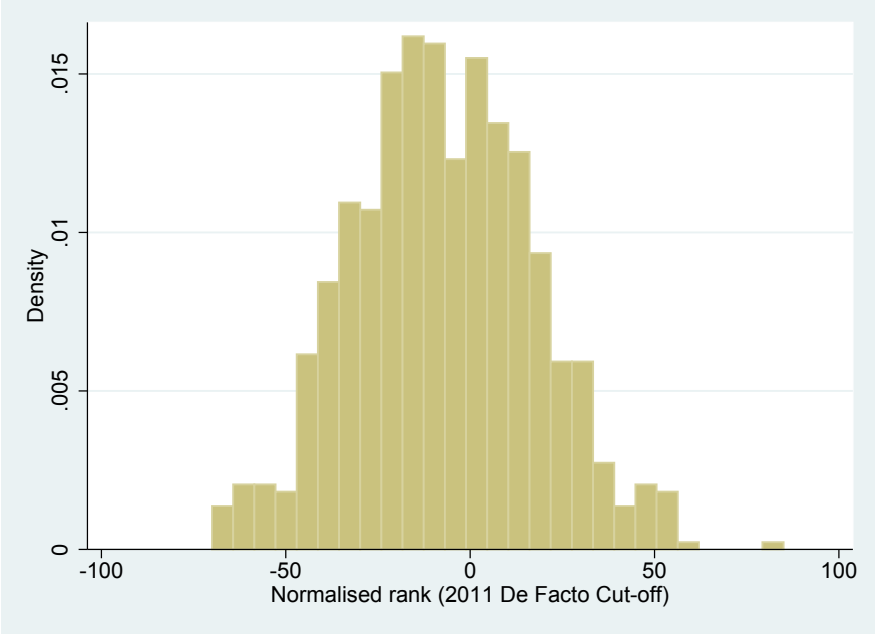

Table 1: Descriptive statistics

	(1) Observations	(2) Income (KSh)	(3) ln(Income)	(4) Age	(5) Male	(6) Married	(7) Number of Children
Panel A:							
All	764	12920	9.330	31.61	0.457	0.723	4.242
		(6978)	(0.559)	(4.160)	(0.498)	(0.448)	(2.497)
Panel B:							
TSC Status							
TSC	584	13829	9.435	31.85	0.437	0.719	4.228
		(5324)	(0.493)	(4.168)	(0.496)	(0.450)	(2.491)
Not TSC	180	9972	8.990	30.85	0.522	0.733	4.289
		(10188)	(0.623)	(4.053)	(0.501)	(0.443)	(2.525)
Panel C:							
If Not TSC: Teaching?							
Teacher	156	9690	8.966	30.65	0.513	0.737	4.327
		(10239)	(0.618)	(3.821)	(0.501)	(0.442)	(2.431)
Not Teacher	24	11804	9.151	32.13	0.583	0.708	4.042
		(9866)	(0.645)	(5.245)	(0.504)	(0.464)	(3.113)
Panel D:							
If Not Teacher: Job?							
Farming	3	5933	8.688	31	0.333	1	7.333
		(115.5)	(0.0196)	(7.937)	(0.577)	(0)	(5.132)
Self-employed	2	12500	9.346	36	0	1	3.5
		(7071)	(0.599)	(0)	(0)	(0)	(0.707)
Wage work	19	12657	9.203	31.89	0.684	0.632	3.579
		(10725)	(0.685)	(5.120)	(0.478)	(0.496)	(2.714)

Figure 3: Possible discontinuities in control variables

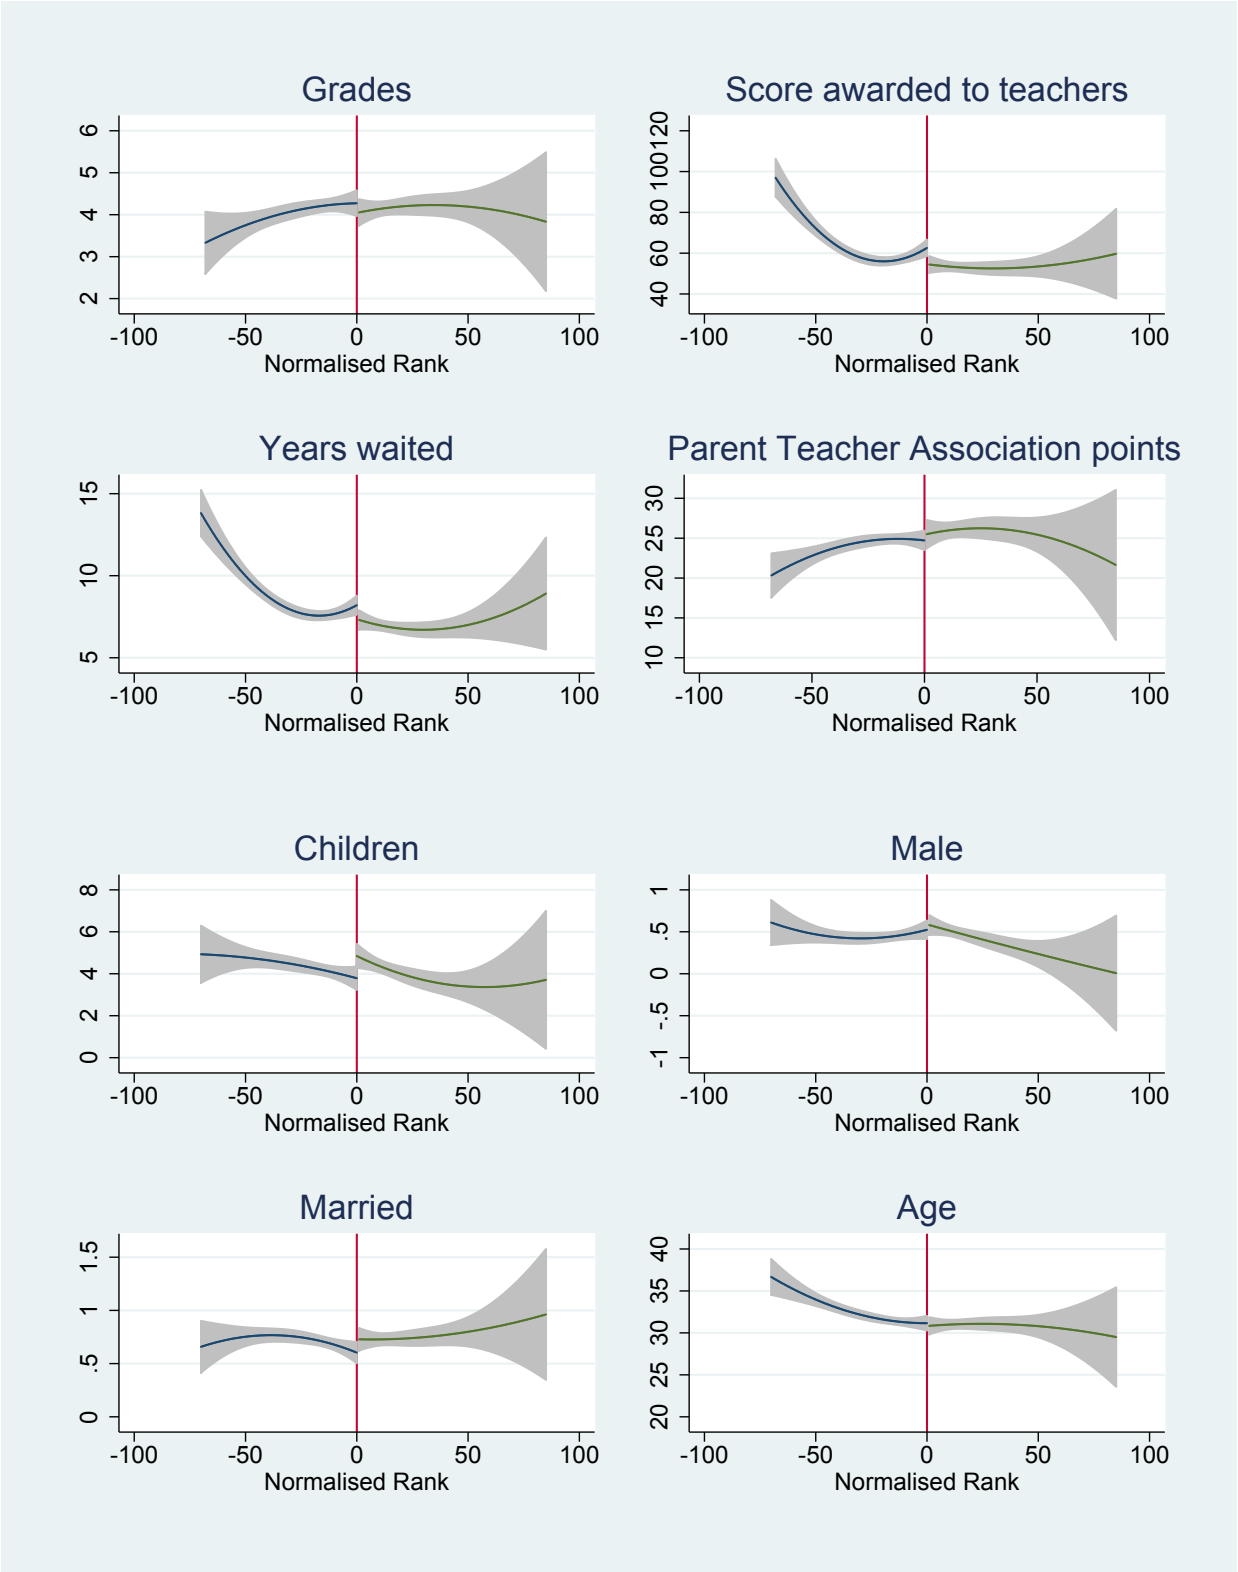

Figure 4: Discontinuity in treatment

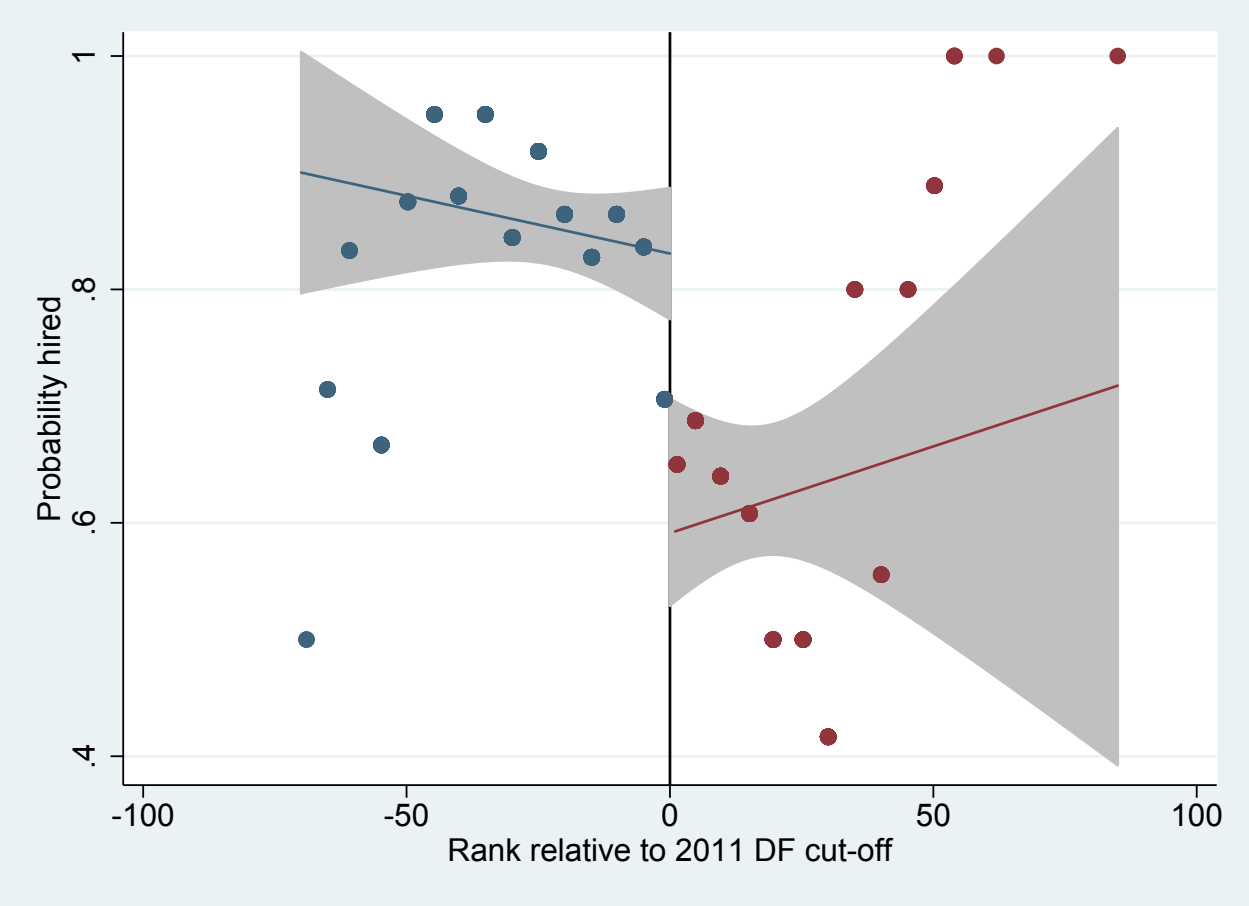

Table 2: First-stage regressions

	(1) Official 2010	(2) de facto 2010	(3) de facto 2011
Below cutoff	0.053	0.17***	0.24***
	(0.048)	(0.047)	(0.048)
Rank	-0.0078***	-0.0046***	-0.0015
	(0.0016)	(0.0014)	(0.0017)
Rank × Below cutoff	0.0060**	0.0031	-0.00060
	(0.0024)	(0.0020)	(0.0022)
Obs.	764	763	764
R-squared	0.13	0.14	0.15

Additional controls included: Age, Age^2, Married, Male, Number of children

Table 3: Full regression table for Income and ln(income)

	Income					Log income				
	(1) OLS	(2) IV (2010)	(3) IV (2010)	(4) IV (2011)	(5) IV (2011)	(6) OLS	(7) IV (2010)	(8) IV (2010)	(9) IV (2011)	(10) IV (2011)
TSC	4215.4***	8876.5*	8927.9*	8978.0*	8829.5*	0.509***	0.725**	0.688*	0.918**	0.866**
	(570.8)	(4941.5)	(4905.0)	(5244.6)	(5136.2)	(0.0404)	(0.360)	(0.351)	(0.442)	(0.420)
TSC*Rank		-220.1	-217.9	-406.2	-388.0		-0.0208	-0.0187	-0.0455	-0.0417
		(242.8)	(229.3)	(337.0)	(313.4)		(0.0177)	(0.0164)	(0.0284)	(0.0256)
Rank		176.9	174.4	329.6	313.4		0.0146	0.0128	0.0356	0.0324
		(182.8)	(173.7)	(275.8)	(257.5)		(0.0133)	(0.0124)	(0.0232)	(0.0210)
Age		1184.9		1372.6			0.0952		0.0850	
		(2796.2)		(2942.5)			(0.204)		(0.248)	
Age^2		-981.4		-1117.3			-0.0798		-0.0654	
		(2747.5)		(2893.0)			(0.200)		(0.244)	
Male		110.4		189.2			0.0267		0.0394	
		(281.5)		(335.9)			(0.0205)		(0.0283)	
Married		139.1		111.5			-0.0119		-0.0144	
		(271.0)		(303.3)			(0.0198)		(0.0255)	
Children		-136.2		-201.2			-0.00102		-0.0146	
		(334.8)		(393.1)			(0.0244)		(0.0331)	
Constant	9698.2***	5673.7	5638.6	5269.8	5408.5	8.941***	8.729***	8.763***	8.531***	8.577***
	(485.4)	(3965.9)	(3920.0)	(4337.3)	(4198.9)	(0.0343)	(0.289)	(0.281)	(0.365)	(0.343)
Observations	764	763	763	764	764	764	763	763	764	764

Standard errors in parentheses
Control variables (Age, Age^2, Male, Children) are standardized with mean=0 and standard error=1
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: First-stage regressions with Patronage controls

	(1) Official 2010	(2) de facto 2010	(3) de facto 2011
Below cutoff	0.042	0.18∗∗∗	0.24∗∗∗
	(0.048)	(0.047)	(0.049)
Rank	-0.0084∗∗∗	-0.0049∗∗∗	-0.0021
	(0.0016)	(0.0014)	(0.0018)
Rank × Below cutoff	0.0066∗∗∗	0.0037∗	0.00024
	(0.0024)	(0.0020)	(0.0022)
Same ethniciy as president	-0.091∗	-0.100∗∗	-0.10∗∗
	(0.047)	(0.046)	(0.046)
Same ethniciy as education minister	-0.065	-0.048	-0.053
	(0.067)	(0.066)	(0.066)
Obs.	754	753	754
R-squared	0.13	0.15	0.15

Additional controls included: Age, Age^2, Married, Male, Number of children

Table 5: Teachers only regression table for Income and ln(income)

	Income					Log income				
	(1) OLS	(2) IV (2010)	(3) IV (2010)	(4) IV (2011)	(5) IV (2011)	(6) OLS	(7) IV (2010)	(8) IV (2010)	(9) IV (2011)	(10) IV (2011)
TSC	4861.2***	10708.0**	11005.7**	12271.3**	12197.5**	0.566***	0.929**	0.900**	1.216**	1.174**
	(580.4)	(5431.8)	(5283.7)	(6040.4)	(5791.2)	(0.0410)	(0.409)	(0.388)	(0.500)	(0.466)
TSC*Rank		-338.6	-292.8	-390.7	-336.0		-0.0307	-0.0259	-0.0427*	-0.0368*
		(298.9)	(262.6)	(302.3)	(263.0)		(0.0225)	(0.0193)	(0.0250)	(0.0212)
Rank		295.6	262.6	363.9	318.9		0.0247	0.0209	0.0378*	0.0327*
		(245.2)	(215.8)	(271.5)	(237.4)		(0.0185)	(0.0158)	(0.0225)	(0.0191)
Age		4275.3		4155.0			0.317		0.346	
		(3742.2)		(3501.5)			(0.282)		(0.290)	
Age^2		-4141.7		-3999.6			-0.306		-0.331	
		(3623.0)		(3395.8)			(0.273)		(0.281)	
Male		105.8		129.3			0.0269		0.0332	
		(299.2)		(312.9)			(0.0225)		(0.0259)	
Married		66.10		-4.700			-0.0156		-0.0242	
		(279.4)		(307.6)			(0.0210)		(0.0255)	
Children		-232.8		-292.5			-0.00977		-0.0204	
		(337.2)		(367.6)			(0.0254)		(0.0304)	
Constant	9120.3***	3877.1	3739.8	2816.5	2981.9	8.890***	8.543***	8.575***	8.312***	8.355***
	(502.9)	(4604.4)	(4420.8)	(4993.9)	(4724.7)	(0.0355)	(0.346)	(0.324)	(0.414)	(0.380)
Observations	740	739	739	740	740	740	739	739	740	740

Standard errors in parentheses
Control variables (Age, Age^2, Male, Children) are standardized with mean=0 and standard error=1
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 5: Alternative bandwidths