Discussion Paper No. 2001/147

Explaining Leakage of Public Funds

Ritva Reinikka[1] and Jakob Svensson[2]

December 2001

Abstract

Using panel data from a unique survey of public primary schools in Uganda we assess the degree of leakage of public funds in education. The survey data reveal that on average, during the period 1991–5, schools received only 13 percent of what the central government contributed to the schools' non-wage expenditures. The bulk of the allocated spending was either used by public officials for purposes unrelated to education or captured for private gain (leakage). Moreover we find that resource flows and leakages are endogenous to school characteristics. Rather than being passive recipients of flows from government, schools use their bargaining power vis-à-vis other parts of government to secure greater shares of funding. Resources are therefore not necessarily allocated according to the rules underlying government budget decisions, with potential equity and efficiency implications.

Keywords: public services, education, private gain, leakage

JEL classification: H52, H72, I22, I29, N37

[1] Development Research Group, The World Bank, 1818 H Street NW, Washington DC 20433, USA. E-mail: rreinikka@worldbank.org.

[2] Institute for International Economic Studies, Stockholm University, 106 91 Stockholm, Sweden. E-mail: jakob.svensson@iies.su.se.

This is a revised version of the paper presented at the UNU/WIDER development conference on Debt Relief, Helsinki, 17–18 August 2001.

UNU/WIDER gratefully acknowledges the financial contribution from the governments of Denmark, Finland and Norway to the 2000–2001 Research Programme.

Acknowledgements

We are grateful for many useful comments and suggestions by Jan Dehn, Jerey Hammer, Phillip Keefer, Michael Kremer, Edward Miguel, Oliver Morrissey, Abel Ojoo, Rohini Pahnde, Ashok Rai, David Strömberg, and Waly Wane, as well as seminar participants at the World Bank, conference participants at WIDER (Helsinki), and ISPE conference participants at Cornell University. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author(s) and do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent. Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

UNU World Institute for Development Economics Research (UNU/WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Camera-ready typescript prepared by the authors and Adam Swallow at UNU/WIDER
Printed at UNU/WIDER, Helsinki

ISSN 1609-5774
ISBN 92-9190-118-0 (printed publication)
ISBN 92-9190-119-9 (internet publication)

Explaining Leakage of Public Funds*

Ritva Reinikka† and Jakob Svensson‡

Abstract

Using panel data from an unique survey of public primary schools in Uganda we assess the degree of leakage of public funds in education. The survey data reveal that on average, during the period 1991-95, schools received only 13 percent of what the central government contributed to the schools' non-wage expenditures. The bulk of the allocated spending was either used by public officials for purposes unrelated to education or captured for private gain (leakage). Moreover we find that resource flows and leakages are endogenous to school characteristics. Rather than being passive recipients of flows from government, schools use their bargaining power vis-à-vis other parts of government to secure greater shares of funding. Resources are therefore not necessarily allocated according to the rules underlying government budget decisions, with potential equity and efficiency implications.

*We are grateful for many useful comments and suggestions by Jan Dehn, Jeffrey Hammer, Phillip Keefer, Michael Kremer, Edward Miguel, Oliver Morrissey, Abel Ojoo, Rohini Pahnde, Ashok Rai, David Strömberg, and Waly Wane, as well as seminar participants at the World Bank, conference participants at WIDER (Helsinki), and ISPE conference participants at Cornell University. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author(s) and do not necessarily represent the views of the World Bank, its Executive Directors, or the countries they represent. Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate. September 2001.

†Development Research Group, The World Bank, 1818 H Street NW, Washington DC 20433, USA. E-mail: rreinikka@worldbank.org.

‡Institute for International Economic Studies, Stockholm University, 106 91 Stockholm, Sweden. E-mail: jakob.svensson@iies.su.se.

1. Introduction

Data on official budget allocations are typically the only source of information on public spending in low-income countries. Unfortunately, such information poorly predicts what the intended beneficiaries actually receive in terms of resources and services. This is particularly so in countries with weak institutions. Surveying the supply side of service delivery can provide a useful reality check. In this paper we describe and analyze the results of an innovative survey tool implemented in Uganda to gauge the extent to which public resources actually filter down to the intended facilities. The survey compared disbursed flows from the central government (intended resources) with the resources actually received by 250 primary schools over a five-year period (1991-95). This unique panel data set let us study the level and determinants of leakage.

The results of the survey are striking. On average, schools received only 13 percent of central government allocations toward their non-wage expenditures. The bulk of the allocated spending did not reach the intended beneficiaries and was either used by local government officials for purposes unrelated to education or captured for private gain (defined as leakage).

The survey data also reveal large variations in leakage across schools. We develop a simple bargaining model to explain these differences. In the model, resource flows—and leakage—are endogenous to school characteristics, as schools use their bargaining power vis-à-vis other parts of government to secure greater shares of funding. These resources are therefore not allocated according to the rules underlying the government's budget decisions, with obvious equity and efficiency implications.

The model's predictions are confirmed by the data. Specifically, we find that larger schools receive a larger share of the intended funds (per student). Schools with children of wealthier parents also experience a lower degree of leakage, while schools with a higher share of unqualified teachers experience less leakage. After addressing potential selection and measurement issues, we find that these school characteristics have a quantitatively large impact on the degree of leakage.

The survey findings prompted a strong response from the central government: it began publishing the monthly transfers of public funds to the districts in newspapers, broadcasting the transfers on radio, and requiring primary schools to post information on inflows of funds for all to see. This not only lowered the information costs to parents and schools, but also signaled local government. An initial assessment of these reforms a few years later shows that the flow of funds improved dramatically, from 13 percent (on average) reaching schools in 1991-95 to over 95 percent of intended capitation grants reaching schools in 1999. The findings of the paper extend the emerging empirical literature on school funding (or more generally public goods provision) in developing countries.[1] Miguel (2000) shows that higher

[1]The literature on school funding in developed countries is large, particularly in the United States (see for instance Fernandez and Rogerson, 1996, and references given therein). Most of this literature explains the actual educational expenditures per student financed at the local level. This paper explores the extent

ethnic diversity is associated with sharply lower primary school funding (school fees collected from parents and local fund-raisers), and worse school-level facilities in western Kenya, suggesting that collective action problems may be more severe in the presence of greater cultural and linguistic differences. We focus instead on central government funding for schools and the influence of local political and socioeconomic factors on the actual outcomes. Although our study does not have information on ethnicity, it suggests that adverse effects of ethnic diversity on private school funding could be magnified by lower public funding through the reduced bargaining power of schools in ethnically diverse areas.

To the extent diverted funds are used for private gain (by district officials), this paper also provides, to our knowledge, the first quantitative attempt to systematically measure corruption in basic service delivery systems. Our findings provide new insight into an area almost exclusively studied using cross-country data.[2] We show that a large part of the variation in corruption at the local level can be explained by studying the interaction between the local officials and the end-users (schools) as a bargaining game. From an analytical point of view our approach differs from much of the existing literature on corruption, since we focus on the principal's (the school's) rather than the agent's (the district officials') incentives and constraints. Our results suggest that a systematic effort to increase the ability of citizens to monitor and challenge abuses of the system, and inform them about their rights and entitlements, are important aspects in controlling corruption.

The results of the paper also have implications for the large cross-country literature on public spending and growth in developing countries, as well as the literature on the macroeconomic impact of foreign aid. In particular, our findings highlight the identification problem in attempting to evaluate the efficacy of public capital or services with public spending data.[3] Given the extent of and variation in leakage, using central government budget allocation data to assess the impact of public spending on growth and social outcomes will severely underestimate any potential positive effect that the public capital or services actually created by public funds can have. Based on the existing cross-country work, the effect of government spending on growth and social development outcomes is ambiguous.[4] Our results suggest that increased spending does not necessarily translate into increased output and services.[5]

---

to which centrally financed educational expenditures are diverted at the local government level.

[2]For effects of corruption on investment and growth see Mauro (1995). On the determinants of corruption, see Ades and Di Tella (1997, 1999), Persson, Tabellini, and Trebbi (2000), Svensson (2000a), and Treisman (2000). A common theme in this literature is the use of subjective measures of corruption in a cross-country setting. Fisman and Svensson (2000), Svensson (2000b), and Di Tella and Schargrodsky (2000) are exceptions. They use quantitative micro-level data on corruption.

[3]Pritchett (1996), Reinikka and Svensson (2001) make a similar argument.

[4]Ram (1986) and Kormendi and Mequire (1985) find higher government expenditures associated with high growth, while Landau (1986), Barro (1991), Dowrick (1992), and Alesina (1997), find higher government expenditures associated with lower growth. Levine and Renelt (1992) show that government expenditures is not a robust (partial) correlate of growth.

[5]The empirical growth literature is abundant with explicit (and implicit) attempts to separate productive spending from expenditures that have no direct effect on productivity (for example by *ex ante* determining

In a similar vein, the recent literature on the macroeconomic impact of aid finds no statistical relationship between aid and growth or social development outcomes (Boone, 1995, 1996).[6]Our results provide a possible explanation for this finding. Since foreign aid is typically intermediated through the recipient's public sector, it is bound to suffer from similar deficiencies. If a large fraction of foreign aid does not result in actual public assets and services, a low correlation between aid and outcomes is to be expected.

The rest of this paper is structured as follows. The next section briefly reviews the institutional setting for school finance and decisionmaking in Uganda. Section 3 discusses the survey and the measurement of leakage. Section 4 sets out a simple bargaining model to investigate the relationship between school-related characteristics and expenditure leakage, whereas section 5 explores extensions to the model. Section 6 describes the empirical specification of the model that we use to examine leakage across schools and discusses the data. The results are presented in section 7. Section 8 concludes.

2. Institutional setting

It is commonly held that Uganda had a well-functioning public service delivery system in the 1960s. The government response to the political and military turmoil of the 1970s and early 1980s was to de facto retreat from funding and providing public services. In primary education parents gradually took over running the public schools. The survey data indicate that by 1991 this situation had not changed much. Parent-teacher associations (PTA) were the primary decisionmakers at the school level, and funding by parents was on average the most important source of income.

While the subsequent economic recovery increased public spending relatively rapidly, institutional reforms were much slower to come. In particular, the central government exercised weak oversight over the local governments (districts), which channeled public funding to the social sectors. District officials thus had discretion over how to use public funds supposedly earmarked for the schools.

During the survey period (1991-95) the central government's financial contribution to primary education was threefold. First, the Ministry of Education paid salaries of primary school teachers either directly, if the teacher had a bank account, or most often through the district education officer and/or the headmaster. Second, the Ministry of Local Government transferred a capitation grant per enrolled student to the district administrations for non-wage expenditures like textbooks, instructional materials, and the costs of running schools. Capitation grants were not politicized in the sense that districts did not receive varying amounts by the central ministry, but a nationally set allocation per student per year. These

what types of spending are likely to be productive, see Barro, 1991). Unfortunately, partitioning expenditure categories does not address the core problem—that public funds may not reach the intended end-user.

[6]Burnside and Dollar (2000) and Svensson (1998), find similar unconditional results, but a positive relationship between aid and growth conditional on the recipient's policies (institutions). Hansen and Tarp (2001), using different methods, find a weakly positive relationship between aid and growth.

grants were managed by the districts on behalf of the central government. Third, the central government provided funding for capital expenditure also through the Ministry of Local Government. This funding was limited almost entirely to rehabilitation. In fact, since the 1970s the central government had virtually abandoned its responsibility for classroom construction. In principle, the provision of classrooms became the responsibility of local governments which passed it on to parents.[7]

The central government's total contribution to the funding of primary schools more than doubled between 1991 and 1995 in real terms, albeit from a negligible base (Table 2.1). In practice the entire increase was used to raise teachers' salaries, which had eroded to extremely low levels (equivalent to a few U.S. dollars a month) during the institutional and economic collapse of the 1970s and 1980s. The capitation grant was retained at the same nominal level throughout the survey period, therefore, its real value actually declined. There was an increase in rehabilitation and school construction spending toward the end of the survey period.

The central government's stated policy was to disburse capitation grants in full to the schools through the local government (districts). The grant was set in 1991 at the nominal rate of Ush 2,500 per child enrolled in grades one to four and Ush 4,000 per child enrolled in grades five to seven. These nominal rates remained the same until 1997. According to anecdotal evidence, the grants often ended up in the chief administrative officer's account in the district, and that the latter did not necessarily transfer the funds onward to the district education officer as was expected. There are also many anecdotes about highly inflated prices for school supplies procured by district officials, meaning little (either in monetary terms or in-kind) was received by schools.

Uganda implemented cash budgeting in 1992, which in many cases produced volatile monthly releases of funds from the Treasury. However, as part of the World Bank's structural adjustment programs non-wage recurrent expenditures for primary education were given priority program status, which protected schools from within-year budget cuts. Capitation grants were fully released by the center to the districts on a monthly basis. In the Ugandan treasury system, central ministries were unlikely to capture releases to local governments because they were subjected to relatively elaborate pre-audit procedures. Hence, from the spending program or agency perspective, the uncertainty of funding was greatest prior to the release. Released funds were very likely to arrive at their intended destination, which in the case of capitation grants was the district.

The central government policy regarding the capitation grant was not well-known to parents, particularly outside the capital city. Even if parents knew about the policy in principle, many similar policy statements were not implemented in practice at that time. Little information was available to the public, for example, on the spending items protected within the cash budget system. This worked well for the districts taking advantage of the

[7] In addition, central government is responsible for a share of the cost of donor-financed development projects (about 10 percent of the total project cost). It also incurs expenditure on teacher training, examinations, and school inspection. The latter was almost non-existent during the survey period.

asymmetric information about school funding; these districts could reduce disbursements or procure little for non-wage items to schools because they knew such action would not attract political attention. By contrast, failure to pay teachers' salaries attracted much more attention as, not surprisingly, teachers knew what their salaries were.

As Table 2.1 shows, parental contributions toward primary education consisted of PTA levies for investment and recurrent costs, top-ups to teachers' salaries, and tuition fees. The PTA fees and top-ups to teachers' salaries were entirely school-specific and set by each school's PTA, depending on the parents' ability to pay and the needs of the school. Parental contributions were clearly the mainstay of finance in government-aided primary schools. On average, during the sample period, parental contributions accounted for over 60 percent of total school income. In per-student terms, parents' average contribution increased by 35 percent in real terms during the sample period. Interviews at primary schools indicated that the parents who were not able to pay the agreed PTA fees were often alienated and even forced to take their children out of the school.

In theory, the tuition fee per student was set by the central government at the same level as the (matching) capitation grant. It was left to each district to determine how the funds raised through tuition fees should be redistributed among the schools. In some districts the schools were allowed to retain a certain percentage or a fixed amount of the tuition fee collected per student, with the balance transferred to the district education officer. In other districts all tuition fees collected were remitted to the district headquarters and subsequent onward disbursements to schools, either in cash or in-kind, may or may not have taken place. The efficiency of tuition fee collection was very low in 1991, but improved somewhat in subsequent years. Interviews at the schools suggested that low collection efficiency was due to adverse incentive: most schools were neither allowed to keep the collected funds, nor benefited from them in any other way.

Teacher recruitment was carried out by district education service committees on behalf of the national teacher service commission. Recruitment was supply driven, as all new teachers graduating from the primary teacher collages were usually hired. Although teachers were hired by the districts, their payroll was maintained by the central government. As a result, and contrary to non-wage spending, the central government provided some oversight for teacher recruitment and salaries through the maintenance of the national payroll. Once recruited, the district education officer posted the teacher to a specific school. Hence teachers had little opportunity to choose the school where they taught. If the demand for teachers exceeded the supply of training colleges, district education service committees recruited additional "licensed" teachers, who were often unqualified.

The PTA derived its authority from parents. The influence of the PTA over district officials depended on their competence to articulate their case. A typical PTA was run by an executive committee that had about six members elected during a general meeting, and the headmaster. According to anecdotal evidence it was common for influential and better-off parents who were close to the school establishment to serve on the executive committee.

Most students had few schools within a walking distance, particularly in rural areas. This lack of choice can be traced back to the tumult of the 1970s and early 1980s and the central government's gradual abandonment of school construction which the local governments were not able to pick up. School choice was also limited by Uganda's preference for "complete schools" (one school offering all seven grades) dating back to the colonial times.

3. Can leakage be measured?

Ideally, the public accounting system provides timely information about actual spending on various budget items and programs, and the reports accurately capture what the intended users receive. This is not often the case in low-income countries. Typically the accounting system functions poorly, institutions enhancing local accountability are weak, and there are few (if any) incentives to maintain adequate records at different levels of government. Consequently, little is known about the process of transforming budget allocations into services within most sectors.

These observations formed the basis for designing a new survey tool—a quantitative service delivery survey[8]—to gauge the extent to which public resources actually filtered down to the intended facilities. A survey covering 250 government primary schools was implemented in 1996, covering the period 1991-95 (see Reinikka, 2001, for details on survey design). At the time of the survey, about 8,500 government primary schools were supposed to receive a large proportion of their funding from the central government via district administrations.[9]

The objective of the survey was twofold. First, to measure the difference between intended resources (from central government) and resources actually received (by the school). Second, to collect quantitative data on service delivery at the frontline (i.e., the schools).

The initial intention to track all main spending categories through the entire delivery system, that is, the central government, districts, and schools, was not possible due to several deficiencies in the system. First, at the central government level, data were not available on salaries paid to primary school teachers either by district or by school. The only data available at the time of the survey were the aggregate salary payments, lumping together payments to teachers in primary, secondary, and tertiary levels, as well as to non-teaching staff. This made systematic comparison between budget allocations and actual spending at the school level impossible with respect to teachers' salaries. Because salary data was lacking or incomplete, we used systematic spending data on per-student capitation grants for non-wage spending available at the central government level as our core variable on intended funds. Second, the district-level records (for both non-wage and wage spending) were much worse than those at the central government level. The quality of available information both on transfers from the center and disbursements to schools was so poor that districts simply had to be excluded from the expenditure tracking exercise. Unlike primary schools,

[8] For a conceptual discussion on Quantitative Service Delivery Surveys (QSDS) and reference to ongoing survey work, see Dehn, Reinikka and Svensson (2001).

[9] The 1,500 private or community schools were not included in the survey.

some districts were also quite uncooperative during the survey exercise. School records, on the other hand, were relatively comprehensive. Thus, a detailed comparison of budgetary allocations and actual spending could be made between the central government outlays for non-wage spending on instructional materials and other running costs and the equivalent school income.

We believe the capitation grant data at the school level adequately capture what the schools receive for several reasons. First, the data collected directly from the school records were kept for the school's own needs. The school records were not submitted to any district or central authorities and were not the basis for current or future funding. Thus, there were no obvious incentives to misrecord the data. The concern that headmasters might have underreported school income in order to extract resources for themselves was allayed after interviews during the survey work, which did not support this possibility. This is not surprising since the PTA was typically the principal decisionmaker (and responsible for raising most of the income) at the school. Furthermore, parents who contributed the majority of school income presumably demanded financial information and accountability from the school (or PTA).

The central government simply assumed the funds reached schools. Audit systems at the center and local governments were weak, and there was little interest in ascertaining how the funds were actually used. The school survey brought out issues about the relationship between school authorities and district education officers for the first time.

Our school specific measure of degree of leakage is,

$$\frac{\text{capitation grants received}}{\text{intended capitation grants from the center}} \tag{3.1}$$

where a low value indicates a large leakage.

In theory, the denominator in (3.1), the intended capitation grants from the center, should be the product of the number of students in the school and the per-student capitation grant. A closer examination of records at the Ministry of EducatiAppendixon, however, revealed two sources of discrepancy from this formula. First, the growth in enrollment at the school level differed considerably from the central government statistics (see Reinikka, 2001, for a detailed discussion). Second, for the entire survey period (1991-95) the capitation grant was determined on the basis of the 1991 enrollment. Thus, the growth in enrollment observed at the school level over the period did not result in increased "intended capitation grants from the center" for the schools. For these reasons, we derive the denominator in (3.1) using 1991 enrollment data.

In order to bring out regional differences in the sample more clearly, the traditional four regions (North, East, West and Central) were reconfigured into seven regions (Northwest, North, Northeast, East, Central, Southwest and West). For each region, two or three districts were randomly chosen, together with the capital city, Kampala, to yield a sample of 18 districts, as illustrated in the appendix map.[10]

[10] The following 18 districts were selected: Arua, Moyo (Northwest); Apac, Gulu (North); Soroti, Moroto,

In the districts selected the number of schools visited ranged from 10 to 20. Bushenyi had the largest number of primary schools (399 in 1994), while Bundibugyo had the smallest number of schools (59). In the districts with less than 100 government schools the enumerators visited 10 randomly chosen schools. Where the number of schools was between 100 and 200, 15 schools were randomly selected for visits, and in the districts with more than 200 schools, 20 schools were randomly chosen for visits.

Enumerators were trained and closely supervised by a local research team and survey experts from the World Bank to ensure quality and uniformity of data collection and standards for assessing recordkeeping at the schools. Standardized forms were used. In addition, interviewers made qualitative observations to supplement the quantitative data.

Do public resources reach the intended schools? How large is the leakage of public funds in education? Answering these questions was one of the key challenges when setting-up the data collection effort. Table 3.1 depicts information on our leakage variable, share of intended capitation grants received. Strikingly, on average only 13 percent of the total yearly capitation grant from the central government reached the school. Eighty-seven percent either disappeared or was used for purposes unrelated to education.[11]Most schools received very little or nothing (roughly 70 percent of the schools). In fact, based on yearly data 73 percent of the schools received less than 5 percent, while only 10 percent of the schools received more than 50 percent of the intended funds.

The picture looks slightly better when constraining the sample to the last year of the sample period. Still, only 22 percent of the total capitation grant from the central government reached the school in 1995. Thus, in 1995, for every dollar spent on nonwage education items by the central government, roughly 80 cents got diverted!

As illustrated in Table 3.1, there is variation across regions, although the bulk of the variation is within the regions. The standard deviation of leakage (share of intended capitation grants received) across regions is roughly one-third (9.2) of the average standard deviation within regions.[12] The Central region (including the capital) appears to be the only region with significantly lower leakage.[13] In the next two sections we attempt to account for this variation within (and across) regions.

Kapchorwa (Northeast); Jinja, Kamuli, Pallisa (East); Kampala, Mukono, Mubende (Central); Bushenyi, Kabale (Southwest); and Kabarole, Hoima, Bundibugyo (West).

[11]The classic argument of fiscal federalism is that local governments can better match public goods and services to preferences. Azfar and others (2000) analyzed whether district and sub-county government officials in Uganda are aware of household preferences in their jurisdictions, and whether they adjust resource allocations to respond to household preferences. Their results show that government officials at the national and sub-county levels, but not at the district level, are aware of household preferences. Actual resource allocations, however, reflect local preferences only weakly.

[12]The results are similar when comparing within and across districts rather than regions.

[13]This result is confirmed when running regression of leakage on the seven regional dummies. Only the null hypotheses of equal regional effects between the Central region and the other six regions can be consistently rejected.

4. A bargaining model of school expenditures

Below we set out a simple bargaining model to guide the empirical specification. The objective is to show that sociopolitical features of the school (parents and teachers) have implications for the equilibrium amount of leakage. The model assumes that the extent to which public funds reach the primary schools depends on the bargaining strength of the school vis-à-vis the district bureaucracy.

4.1. Basics

Consider a school i, $i \in I$, with n_i students. For simplicity we assume each student (child) belongs to a separate household h. Each (identical) household supplies inelastically one unit of labor and earns income y_i. Income is used to finance a private consumption good c_i, and educational services, e_i.

A household h with a child in school i has the following separable quasi-linear preference function:

$$U_{hi} = u(c_{hi}) + e_{hi} \ , \tag{4.1}$$

where e_{hi} is the (quantity and quality of) educational services provided to a student h in school i, and $u(.)$ is a standard utility function with $u' > 0$, $u'' < 0$.

We assume that e_i depends on both the amount of government-provided financial support, s_i, and the parents' own contribution, $\sum_h t_{hi}$. Thus, $e_{hi} = e_i = s_i + \frac{1}{n_i} \sum_h t_{hi}$. As an example, e_i could be text books or improved school facilities.

The I schools belong to a district which receives a grant g (per student) from the central government. The grants are intended for the schools, but are handled by the district bureaucracy (or a district official—we will use both terms interchangeably). The district official has discretion over the use of the funds and will disburse $s_i = g - x_i \geq 0$ per student to school i, where x_i is leakage. We assume the district official is an expected profit (rent) maximizer, thus he attempts to extract (in expected terms) as much of the public funds as possible. Formally, the district official maximizes,

$$EU^o = E \sum_{i=1}^{I} n_i x_i \ . \tag{4.2}$$

The n_i households and teachers associated with school i form a parent-teacher association $(\text{PTA})_i$ that is the effective decisionmaker at the school.[14] The PTA determines the

[14] By assuming that the PTA is the effective decisionmaker and can enforce its decisions (section 2 provides motivations for this assumption), we assume away free-riding problems. Given that most schools are fairly large (median school has 429 students), we believe that while free-riding may be a problem in reality, it will not be an important variable in explaining differences in leakage across schools. The reason being that the additional free-riding problem caused by increasing school size from say 300 to 400 (or 500) students is not likely to be large. The free-riding problem may be important when comparing very small school with large schools. In the empirical work we control for school size.

contribution schedules $t_i = (t_{1i}, t_{2i}, ..t_{n_i i})$, and bargains for resources from the district official to maximize joint (household) welfare. Specifically, at the beginning of the game the PTA receives an offer s_i. If it accepts, the game ends and educational services $e_i = s_i + \frac{1}{n_i}\sum_h t_{hi}$ (per student) is produced.

The problem for the PTA is that ex ante g is not known; i.e., g is private information to the district official. Alternatively we could assume that g is known, but that the school cannot determine (without a costly effort) if their district has actually received the funds from the central government. The PTA only knows that g is distributed on the interval $[0, \bar{g}]$ according to the distribution function $F(g)$.

The PTA can obtain information about both g, for example by contacting the central government, but this is costly. Let θ be the school-specific cost of finding out the true g.

In case the PTA does not accept the offer, it can exercise its voice option.[15] Voice can take many forms (see Hirschman, 1970), including individual or collective petition and/or appeal to a higher authority, including local chiefs, or through various types of actions and protests. There is a cost κ, defined in per-student terms, to launch a protest. We can conceptualize κ in a variety of ways. In order to initiate a (successful) protest the PTA (most likely) must disseminate the information about g to the parents; it must (most likely) build a coalition for action within the school, it might need to formulate an appeal to the Ministry of Education, and provide political contributions. All these actions are costly.

A protest is successful with probability π, in which case all intended funds (g) are disbursed to the school. With probability $1 - \pi$ the protest is not successful and the PTA will end up with s_i. π is assumed to be exogenously given $\pi \in (0, 1)$.

The timing of events are as follows. First, the PTA receives an offer s_i from the district official and sets PTA fees t_i. The PTA can either accept the offer or reject it. In case it rejects the offer, it can invest θ to find out the true g, and, if optimal, exercise its voice option (launch a protest). The order of events is as follows.

Timing of events

	(a) set school fees
PTA:	(b) obtain information of entitlement [yes, no]
	(c) if yes, form coalition and exercise voice option [yes, no] $\longrightarrow$
District:	(a) provide funds to schools
Nature:	(d) nature draws π

4.2. Equilibrium leakage

How much of the intended funds will the district official transfer to the school, and what factors make leakage more likely? The problem can be solved by working backwards. Consider

[15] In reality, parents may also use their exit option; i.e., move their children to another school/district. As discussed below, for our sample of public schools, this option is likely to be less relevant due to limited residential mobility in the presence of poorly functioning land markets and the scarce supply of schools in the rural areas.

a PTA who has invested θ. Clearly, it will find it optimal to launch a protest if the expected gain, $\pi g + (1-\pi)s_i - \kappa$ is larger than the certain payoff s_i. That is if,

$$g \geq \hat{g} \equiv s_i + \frac{\kappa}{\pi} \ . \tag{4.3}$$

Condition (4.3) can be re-stated as,

$$\pi (g - s_i) - \kappa \geq 0. \tag{4.4}$$

If the expected gain per student of a protest, the first term in (4.4), is larger than the expected cost per student, the PTA will launch a protest.

In the first stage, the PTA will decide to incur the information cost if its expected net benefit of doing so is nonnegative, that is,

$$\int_0^{\hat{g}} s_i f(g) dg + \int_{\hat{g}}^{\bar{g}} [\pi g + (1-\pi)s_i - \kappa] f(g) dg - \theta / n_i \geq s_i \ . \tag{4.5}$$

The left-hand side (LHS) of (4.5) represents the expected income when θ is incurred, while the right-hand side is the (certain) level of funding per student in the absence of the information investment. Equation (4.5) can be rewritten as,

$$\int_{\hat{g}}^{\bar{g}} [\pi (g - s_i) - \kappa] f(g) dg \geq \theta / n_i \ , \tag{4.6}$$

which clearly illustrates the consequence of an unknown g. Only if the expected net gain per student of a protest is sufficiently large, (LHS) of (4.6), will the PTA incur the cost of acquiring and disseminating information about public funding.

Equation (4.6) is a necessary condition for incurring the information cost. In addition, there is a liquidity constraint. The PTA must be able to afford the information investment and protest cost. That is,

$$n_i \kappa + \theta \leq \sum_h t_{hi} \tag{4.7}$$

At the beginning of the game, the PTA chooses the contribution schedule t_i. All individuals in a school district are identical. With quasi-linear utility, equilibrium school funding is simply

$$t_{hi} = t_i = y_i - u_c^{-1}(1) \ . \tag{4.8}$$

Consider next the district official's problem. By choosing a s_i such that (4.6) binds, the official can ensure that no protest will be voiced by the PTA. This will be an optimal response provided that the upper bound on the expected grant g ($\bar{g}$) is not too large. Specifically, if $\bar{g} = g$, which is the case if g is known but that the school cannot determine (without a costly effort) if the district has received all funds from the central government, it is optimal to choose s_i so (4.6) binds. Extracting more resources will lead the PTA to invest θ and protest, which yields strictly lower expected utility for the district official. By extracting less, the official simply gives up rents to the school.

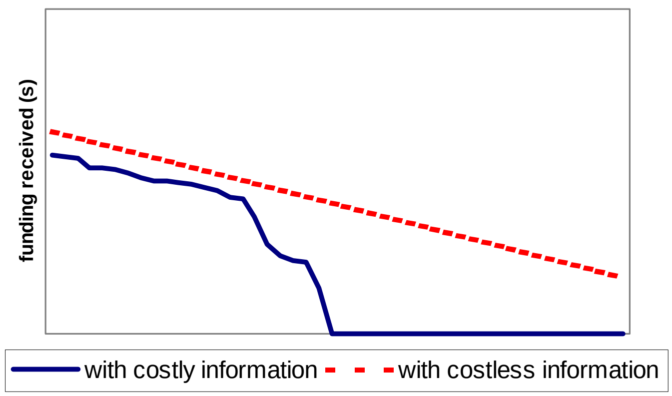

Figure 4.1: s_i^* as a function of κ_i.

Proposition 1. *If π is sufficiently large, there exists an equilibrium without protest in which funding to the school (leakage) s_i is a non-increasing function of the information cost θ and the protest cost κ, and a non-decreasing function of average income y_i and the size of the school, n_i.*

Proof. See appendix. ■

The intuition for the results summarized in proposition 1 is straightforward. The cost of acquiring information and the cost of exercising voice have direct bearing on the cost-benefit decisions in (4.4) and (4.6). As θ is fixed, a larger school implies lower per-student costs of acquiring information. Lower per student costs in turn implies that the district official must disburse more funds ex ante to avoid a protest. Parental income influences equilibrium leakage through the liquidity constraint. Higher parental income implies that (4.4) is less likely to bind. The school can then threaten to initiate a protest.

An implication of proposition 1 is that if income is too low or the cost of acquiring information is too high, the school may end up with no funding. In this case, condition (4.4) may still hold with strict inequality, implying that a well informed school would initiate a protest (net gain of protest > 0). However, because the cost of acquiring information may be too high, the school chooses not to invest θ. As a result, the district official can divert all funds. An example of such an outcome is illustrated in Figure 4.1.

This simple model illustrates two crucial points. First, information on central government spending allocations can be misleading in explaining outcomes, in particular when institutions and oversight in the public sector are weak. That is, g and s may differ substantially. Second, the equilibrium amount of leakage x_i^* is a function of the school's relative bargaining strength vis-à-vis the district bureaucracy.

5. Extensions

Before proceeding to specify an empirical model, it is useful to consider relaxing some of the simplifying assumptions in the model. This is important to better understand the empirical findings presented below and our choice of empirical strategy.

The stylized model set up in section 4 identifies a set of cost factors as important determinants of leakage. These cost factors in turn are determined by various school-specific factors, such as the quality of the school leadership and the social cohesion in the school/community. It is plausible that a school with skilled leadership will require less resources to acquire information and initiate a protest. The social network determines the cost of agreeing, coordinating, and minimizing free-riding problems in the case of a protest.[16] Other (partly unobservable) factors that may influence the school's bargaining strength vis-a-vis the district are distance to district headquarters, whether or not the school is located in an area that supported the (local) government, and access to media. In the empirical work we attempt to measure some of these underlying determinants of κ and θ. As several of them are time invariant (at least in a 5-year perspective), we can deal with the potential omitted variable problem by using school-specific fixed effects.[17]

In the model, students cannot choose which school to attend (or not attend at all). Allowing multiple school choices may result in local sorting that would influence the observed relationship between n and x. While this is a serious concern in principle, as discussed in section 2, we believe it to be less of a concern in reality for the survey period. For most parents and students there was little choice with respect to primary school.[18] Only in some urban areas did parents have a choice where to send their children, so the sorting bias is likely to be small.

The model takes the location of schools as given. In reality school construction is endogenous. In education systems like Uganda's where local financing is important, more affluent communities can afford to build more schools (or support private schools), suggesting fewer students per school in richer communities (cf. Duflo, 2000). Parents in these communities are likely to be more educated, have better political and bureaucratic access, and thus have

[16] Studies on the role of social networks in overcoming coordination problems and reducing transaction costs in developing countries include Narayan and Pritchett, 1999 and Wade, 1988. To the extent that ethnic ties proxy for social networks, Miguel (2000) argues that ethnically diverse communities are less able to ensure enough social pressure for sustaining primary school contributions in rural western Kenya. In related work Gugerty and Miguel (2000) show that higher ethnic diversity is associated with lower community participation in school meetings. Anecdotal evidence suggests that similar mechanisms apply to most parts of Uganda.

[17] It seems implausible to assume that the omitted variables are orthogonal to our set of regressors. See Miguel (2000) on the relationship between ethnicity and private contributions.

[18] An important explanation for this is simply the limited number of schools in most rural areas, which in turn can be traced back to the tumultuous period in the 1970s and early 1980s.

a better chance to capture its share of local funding. This non-randomness in school construction would bias the coefficient on school size downward. Empirically, we deal with the sorting problem and the non-randomness in school construction by instrumenting for school size.

We have not allowed any heterogeneity across districts. The focus on school and community characteristics seems relevant given that the bulk of the variation in leakage is within districts (region). However, it is feasible that for instance high-income districts are better run (lower leakage) and that processes at the district level, rather than at the school level, make it harder for officials to divert funds. At the extreme, all variation in leakage could stem from district characteristics. This alternative hypothesis is tested below.

6. Specification

The model identifies four explanatory variables n_i, κ_i, θ_i, and y_i. Generically, our empirical model can thus be stated as

$$x^*_{ijt} = X(n_{ijt}, \theta_{ijt}, \kappa_{ijt}, y_{ijt}) + \varepsilon_{ijt} \,, \tag{6.1}$$

where subscripts i, j, t refer to school, district, and year, respectively, and ε_{ijt} is an error term. Below we discuss how we attempt to measure the variables in (6.1).

Only one of the explanatory variables in (6.1) is directly observable, namely the number of students (n_{it}). Thus, n_{it} is measured as the number of students in primary school (P1-P7) i at time t, denoted by students. We do not have data on parental income. However, we do have information on parents' financial involvement in the school. PTA income is the average (per student) contribution by parents to the school. In the simple model of section 4 there is a one-to-one relationship between y and t. Thus, increased income implies larger contributions to the school at the margin. The cost variables κ_i and θ_i are proxied by two variables. The first proxy is a time-variant measure of the quality of the school/PTA leadership, defined as the number of qualified teachers to the total number of teachers in the school (share of qualified teachers). This is a suitable proxy if formal education signals competence and competence determines the amount of resources that must be invested to acquire information and voice a complaint. The second proxy is a time-invariant, school-specific effect η_i. As discussed in section 5, many of the underlying determinants of κ_i and θ_i, such as degree of social cohesion, political access, and distance to district headquarters, can (in the short run) be treated as fixed. A detailed description of all variables are provided in appendix 2.

Obviously, when estimating the determinants of x_{it}, it is necessary somehow to scale the level of leakage. As indicated earlier, the most natural approach is to define leakage as share of grants received by school i at time t to what the school should have received $(s/g)_{ijt}$.

Log-linearizing (6.1), our empirical model is then,

$$\log \left(\frac{s}{g}\right)_{ijt} = \beta_0 + \beta_1 \log \text{qualified teachers}_{it} + \beta_2 \log \text{PTA income}_{it}$$
$$+\beta_3 \log \text{students}_{it} + \mathbf{w}_{it}\gamma + \eta_i + \mu_t + \varepsilon_{ijt}\,, \tag{6.2}$$

where $\mathbf{w}$ is a vector of other controls, μ_t is a time-specific effect, and ε_{ijt} is an error term. We allow for district and year-specific random effects. Thus, $\varepsilon_{ijt} = \bar{\varepsilon}_{it} + \hat{\varepsilon}_{jt}$, where $\bar{\varepsilon}_{it}$ is an idiosyncratic error term and $\hat{\varepsilon}_{jt}$ is a random district (j) and year (t) effect. The model suggest that $\beta_1,\ \beta_2,\ \beta_3 > 0$. Note that $\left(\frac{s}{g}\right)_{ijt}$ is censored from below; i.e., $s \geq 0$. For further references, let $\mathbf{z}_{it} = [\log \text{qualified teachers}_{it}, \log \text{PTA income}_{it}, \log \text{students}_{it}]$.

7. Results

Before proceeding, it is useful to take an initial look at the sample. Some schools did not report data for all five years, either due to missing records for these years or because the school was not operational in the earlier years. Excluding a handful of misrecorded observations, we ended up with roughly 950 observations for 239 school.

Descriptive statistics are reported in Table 7.1 and Figure 7.1. In the sample, average school size is 492 students. There are large variations, however, with the smallest school having 35 students and the largest one having roughly 100 times as many. The distribution is illustrated in Figure 7.1. The average student/teacher ratio is 32 students per teacher, with 68 percent of the teachers being qualified. Thirty-four schools (14 percent), reported that they did not have any qualified teachers for at least one year during the sample period. Only 13 schools had only qualified teachers at least one year during the sample period, and only one school had only qualified teachers during the whole sample period.

Parents contributed on average US$10 (in 1990 prices) to school expenditures. The data, however, again reveal large variations. Twenty-seven schools (11 percent) reported no supplementary income from the parents in any year in which data were reported, while there are 44 school-year observations (5 percent) with PTA income per student above US$50. The median yearly contribution per student is US$1.60. As with the leakage variable, the variation in PTA income per student, share of qualified teachers, and school size is mainly within the districts (regions).

7.1. Basic findings

We start by looking at the simple relationship between leakage and the school characteristics, recognizing that there are several econometric issues, including censoring, sorting, endogeneity, and measurement problems, that have not yet been addressed. We deal with these concerns in the following subsections.

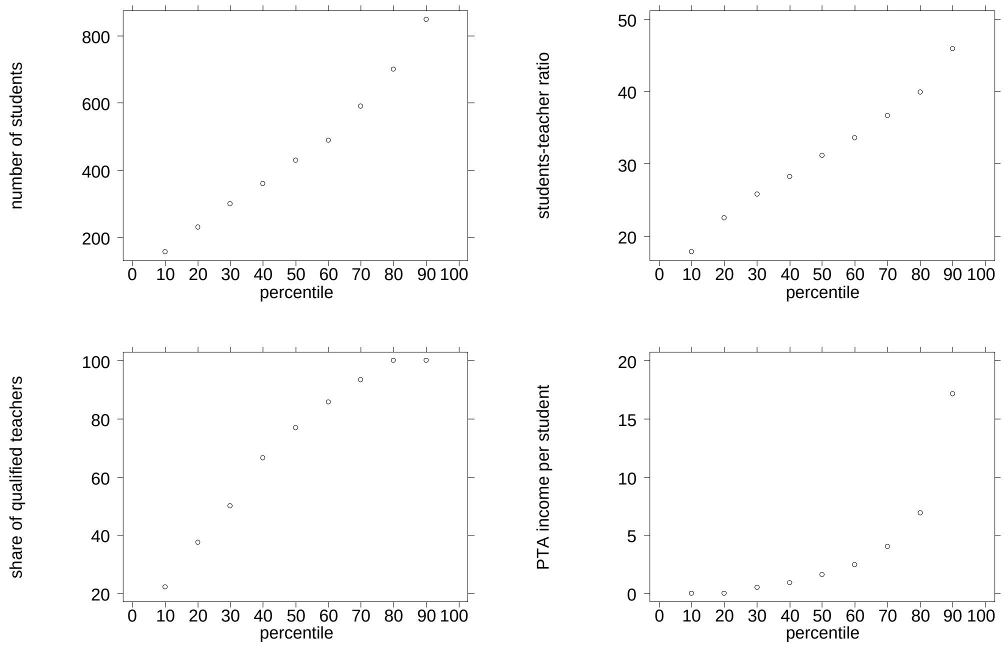

Figure 7.1: Cumulative distribution of explanatory variables

As a reference point, Table 7.2, column 1, reports a cross-section regression; i.e., equation (6.2) without school-specific fixed-effects. The share of qualified teachers enters significantly and with the predicted sign. PTA income per student also enters with the right sign, but is not significant at standard significance levels. The variable school size, however, enters with a negative sign. These results provide some weak support for the bargaining hypothesis. However, given that we do not control for any (time invariant) school/community characteristics, the results should be viewed accordingly. If the school-specific effects are correlated with the vector $\mathbf{z}_{it}$, the coefficients suffer from omitted variable biases.

Columns 2-5 report the results of estimating (6.2) with fixed-effects least squares. The first three columns show the partial effect of PTA income per student, share of qualified teachers, and school size, on the share of intended capitation grant received, controlling for other (time invariant) community characteristics. All three variables enter with predicted signs and are highly significant, suggesting that local sociopolitical factors influence the schools' bargaining powers, and thus the degree of leakage of public funds. The base regression is depicted in column (4). As evident, the variables are both individually and jointly highly significant, and the estimated effects are quantitatively important. A 1-percent increase in school size reduces leakage by 0.8 percent. Similarly a 1-percent increase in PTA support (higher parental income) increases the amount of public funding that reaches the school by 0.3 percent, and a 1-percent increase in the share of qualified teachers raises the amount of public funding that reaches the school by 0.4 percent.

Table 7.2 also reports two specification tests. F is the F ratio for the null hypothesis that all school-effects (η_i) are equal. H is the Hausman (1978) test statistic for testing the hypothesis that η_i and $\mathbf{z}_{it}$ are uncorrelated; that is, a test for fixed or random effects. As evident, both hypotheses can be soundly rejected, thus providing support for our choice of a fixed effect estimator.

The preliminary findings reported in Table 7.2 support the main hypothesis of the paper: the equilibrium amount of leakage is a function of the schools' relative bargaining strength. The bargaining power, in turn, is a function of (average) parental income, school size, the quality of the school leadership, and a set of (time invariant) community/school characteristics. In section 4 we provide a plausible explanation for why these variables should matter. Acquiring information and initiating a protest are costly actions. Schools with students of relatively wealthy parents are more likely to be able to afford these costs. The skill-level of the school leadership determines the investment costs (θ, κ), and to the extent that the costs are partly fixed, the per-student cost is also inversely related to school size. The school-specific effects capture fixed factors such as degree of social cohesion, political access, and distance to district headquarters. The data shows that these fixed factors are also important.

In the following subsections we show that these qualitative results are robust.

7.2. Censoring and time effects

Table 7.3 reports the same set of regressions estimated by maximum likelihood (MLE). With censored data fixed-effects least squares is inconsistent. All coefficients remain highly significant. As expected, the MLE estimates are larger than the fixed-effects least squares estimates. A simple comparison, however, is misleading since the unscaled coefficient vector ($\boldsymbol{\beta}^{MLE}$) only captures $d\left(\frac{s}{g}\right)_{it}/dz_{it} \mid \left(\frac{s}{g}\right)_{it} > 0$. The left column of Figure 7.2 plots $dE\left(\frac{s}{g}\right)_{it}/dz_{it}$; that is, the expected marginal effect on grants received to what should have been received of an increase in the explanatory variables. The right column of Figure 7.2 plots the same derivatives $dE\left(\frac{s}{g}\right)_{it}/dz_{it}$ for all but the top 10-percentile observations. All derivatives are evaluated at the mean of the explanatory variables. For most schools, that is for smaller schools, schools in poorer communities, and schools with relatively few qualified teachers, the marginal impact is small.

With school-specific fixed effects, $\boldsymbol{\beta}_z$ is identified from the deviation from school means. This identification strategy may be problematic if all variables have a common time trend. On the other hand, the data is noisy and including time effects places a strong restriction on the data. As shown in column (5), the effects remain intact when adding time effects. The coefficients are jointly highly significant, although the coefficient estimates on PTA income per student and share of qualified teachers are smaller. With time effects, PTA income per student becomes marginally insignificant (at the 10-percent level).

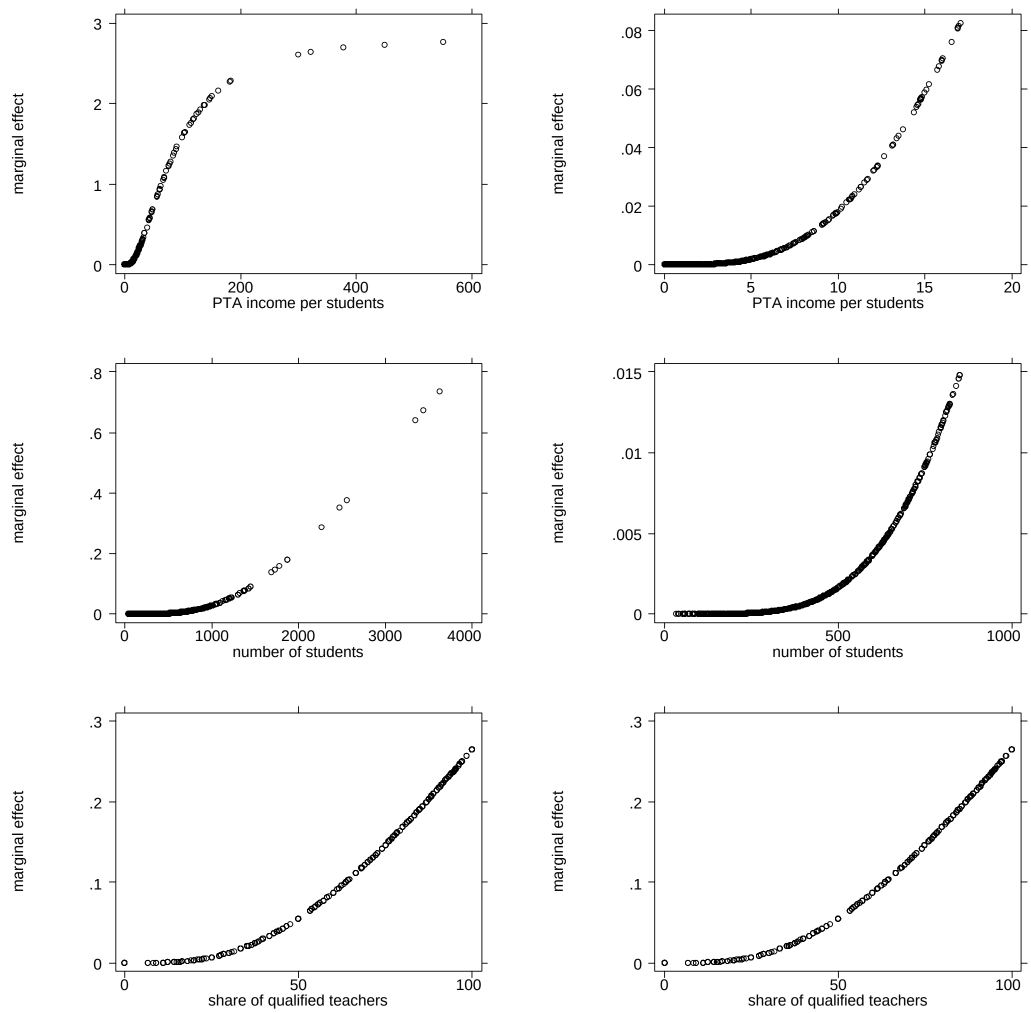

Figure 7.2: Marginal effects (in %) of changes in explanatory variables, $dE\left(\frac{s}{g}\right)_{it}/dz_{it}$. Left column: all sample. Right column: sample excluding top 10-percentile observations [Table 7.3, column (4)].

7.3. Self-selection, endogeneity and measurement errors

Until now we have relied on the restrictions of the model to estimate the relationship between leakage and the schools' bargaining strength. However, as argued above, the model is restricted in that the agents' action space is reduced to exercise "voice". In reality, parents may also use their "exit option". Specifically, poorly financed schools (schools suffering from extensive leakage) may not be able to attract many students. That is, students may self-select into well financed schools and/or may choose not to attend poorly funded schools. If this is the case, the estimated relationship between school size and the share of intended capitation grant received suffers from a sorting bias that would bias the coefficient on school size upward. On the other hand, in education systems relying on local financing, more affluent communities can afford to build more schools (or support private schools), suggesting fewer students per school in richer communities. Parents in these communities are likely to be more educated, have better political and bureaucratic access, and thus have a better chance to capture its share of local funding. This non-randomness in school construction would bias the coefficient on school size downward. We deal with the potential sorting and non-randomness biases by instrumenting for school size using district population data (denoted by district population). While there might be some sorting within given districts, there is very limited mobility across districts. Likewise, to the extent that the variation in school construction intensity is mostly local, district population mitigates the potential non-randomness bias.

Using instrument techniques also addresses another significant estimation issue, the impact of "noisy" data. The problem with non-sampling measurement errors is a general concern when using micro-level data. While there are no strong incentives for the school to misreport the number of students in its own records, measurements or recording errors can still be expected.[19] The district-level population data should serve to mitigate the effects of measurement error, since we generally think of these errors as being largely idiosyncratic to the school.

In principle, the estimated relationship between share of qualified teachers and the share of intended capitation grant received suffers from a similar sorting bias: qualified teachers might self-select into well-financed schools. However, teachers could not shop around for jobs themselves because the appointments during the sample period were made by the districts. Teachers had limited choice about choosing which schools to work in within a district. Good (qualified) teachers could try to get into private primary schools, but since our sample consists of only public schools this selection problem is less of a concern. The allocation of (quality) teachers across schools within a district may be partly determined by the relative bargaining strength of the schools.[20] However, to the extent that our explanatory variables n, y, and

[19] It is plausible that the incentive to exaggerate the number of students is stronger for small schools, thus introducing a bias that would mask negative relationship between school size and leakage.

[20] It is worth noting that with respect to the hiring of teachers, the central government (Ministries of Education and Public Service and the Teacher Service Commission) clearly exercised some oversight over the district educational officers and district education service commissions.

the school-specific effects η capture the relative bargaining strength of schools this will not cause a problem. Only to the extent that there are time-variant school-specific effects that influence both the allocation of teachers across schools and the share of intended capitation grant received will the coefficient on share of qualified teachers be affected. We therefore choose to treat the share of qualified teachers as exogenous.

In the model PTA income per student is an endogenous variable, although in a one-to-one relationship with y. In a more general set-up, however, parents' contributions would depend on both income and amount of funds received from the district. For a given y_j, well-financed schools (low leakage) will receive less contributions from the parents (substitution effect). This endogenous response will tend to mask the positive relationship between PTA income per student and share of intended capitation grant received, and thus work against us. In addition, PTA income per student may also be measured with error. These problems may be mitigated by instrumenting for PTA income per student. Our instrument (denoted by mean consumption) is created in three steps using household expenditure data. The 1992 Integrated Household Survey data (IHS 1992) provide the basis for the instrument. First, the IHS 1992 were used to derive district mean consumption levels in 1992.[21] Second, since the survey data are not representative at the district rural-urban level, we use the statistically robust ratio between urban and rural consumption at the regional level (central, east, west, north) as a scale factor to decompose mean district consumption into mean district urban and mean district rural consumption. Finally, while subsequent household survey sample sizes were too small to be representative at the district level, they are large enough to robustly capture regional (central, east, west, north) differences. Thus, average annual growth rates from 1992-95 were calculated using data on regional-urban-rural mean consumption levels. The average annual growth rate over the period was then used to infer the urban-rural-district mean consumption levels in 1991. Combining the growth data for 1991-95 with the district mean consumption levels in 1992, we derived our instrument: mean consumption levels across district-urban-rural location in 1991-95.

Table 7.4 depicts the first-stage regressions. The instruments perform well. Mean consumption [district population] is a significant predictor of PTA income per student [school size]. In both regressions, the instruments pick up roughly 3 percent of the variation in the explanatory variables.

To deal with the censoring and the selection/measurement problems we estimate the model by conditional maximum likelihood.[22] The results are given in Table 7.5. The IV-estimates are significantly larger than the ML-results given in Table 7.3. The large coefficient on school size suggests that selection issues are of less concern, but that the ML-results suffer from a measurement error bias and possibly a bias due to non-randomness in school construction. Under plausible assumptions, both these types of biases push the estimates toward zero. Similarly, measurement and simultaneity problems mask the relationship be-

[21] We wish to thank Simon Appleton for providing some of these data.

[22] The conditional log-likelihood function for a simultaneous limited dependent variable model is given in Smith and Blundell (1986).

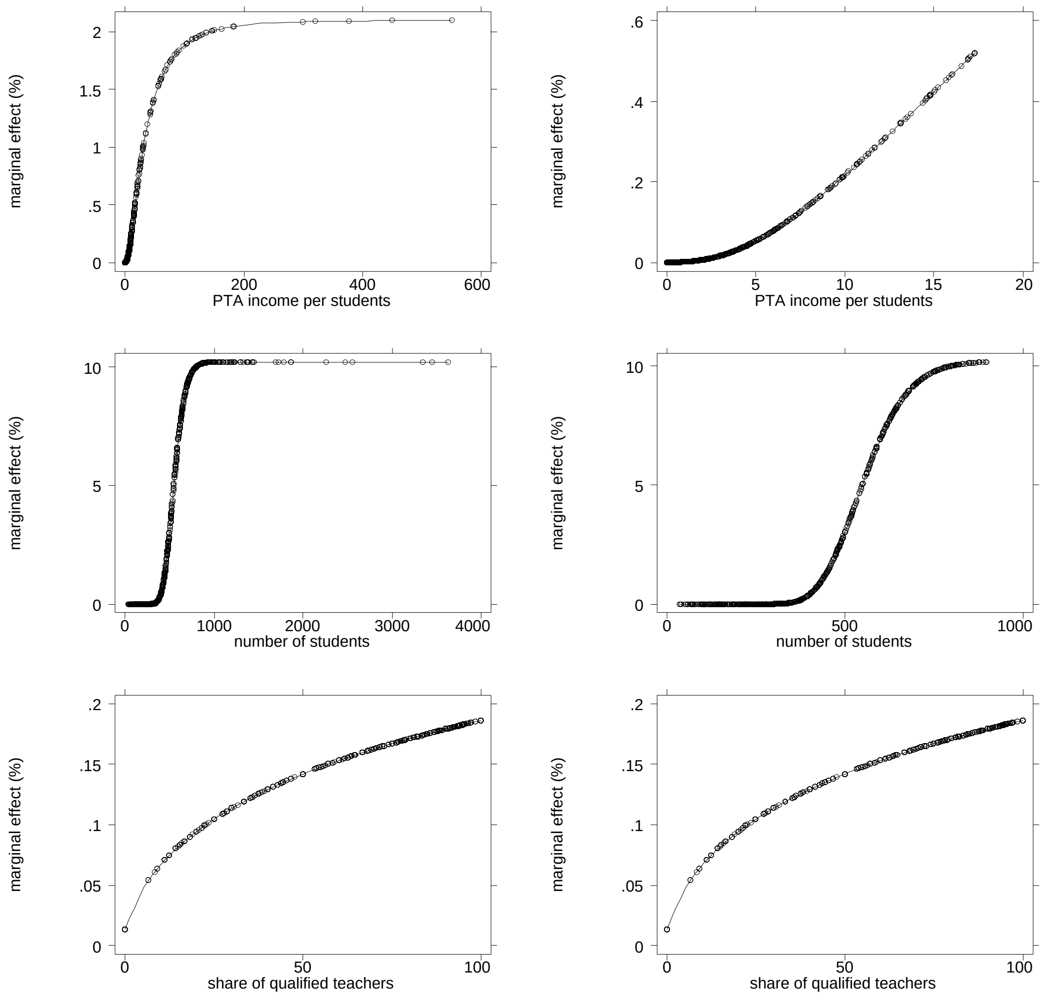

Figure 7.3: Marginal effects (in %) of changes in explanatory variables (IV-estimates), $dE\left(\frac{s}{g}\right)_{it}/dz_{it}$. Left column: all sample. Right column: sample excluding top 10-percentile observations [Table 7.5, column (2)].

tween income and share of intended capitation grant received in Table 7.3. These problems are mitigated when instrumenting for PTA income per student.[23] As evident from column (2), the results remain intact when including time-effects, although the coefficient estimates are smaller.[24]

The simultaneous limited dependent variable estimates are qualitatively large, also for smaller schools and schools with less wealthy parents. Figure 7.3 (left column) again plots $dE\left(\frac{s}{g}\right)_{it}/dz_{it}$, with the right column depicting the derivatives for all but the top 10-percentile observations. A 1-percent increase in school size (evaluated at the mean of all explanatory variables) reduces leakage by 2 percentage points. A 1-percent increase in PTA support increases the amount of public funding that reaches the school by 0.25 percentage points, and a 1-percent increase in the share of qualified teachers reduces leakage by 0.27 percent.

To summarize, once dealing with potential measurement, endogeneity, and selection problems, we find the identified school characteristics have a quantitatively large impact on the degree of leakage.

7.4. Additional robustness tests

We ran a number of additional robustness tests on the results reported above. One concern is outliers. Until now, we have taken an extremely conservative approach with respect to outliers: only a few observations, which seem quite clearly to be a result of misrecording, have been dropped. However, some fairly serious outliers remain. In particular, there are 17 [3] observations on PTA income per student [school size] taking values of more than 3 standard deviations above the mean. While there is no theoretical justification for deleting these observations, it would be of considerable concern if our results were completely driven by them. To examine this possibility, we dropped all observations on school size and PTA income per student with values larger than 3 standard deviations above the mean. The results are similar to those reported above.

We added additional controls, including the students-teacher ratio and the tuition fee per student. Adding these variables did not change the results. Only tuition fee per student had some explanatory power, as can be seen from Table 7.6, column (1). All other variables remain unchanged.

In column (2) we add the district-based instrument variables mean consumption and district population to the basic regression. When instrumenting, the parameters are identified solely based on variation across districts. One might expect that there are processes at the district level, rather than at the school level, which influence the degree of leakage and thus explain our results. Specifically, it is plausible that our instruments, district income

[23]We note a similar pattern by comparing the fixed-effects least squares results in Table 7.3, with the two-stage, fixed-effects least squares estimation (results available upon request).

[24]It is worth noting that if share of qualified teachers also is measured with error, the resulting attenuation bias pushes the estimate toward zero. Thus, the estimates in Table 7.5 are most likely to constitute a lower bound on the effects of a more qualified teaching staff.

and size, could directly influence the officials' possibilities to divert funds; that is, they have an independent effect on $\frac{s}{g}$. However, once controlling for the set of school-specific characteristics, the evidence suggests that these district characteristics are unimportant. The proxy for district income (district mean consumption level) even enters with a negative sign. The finding that the share of intended capitation grant received does not appear to be driven by these district specific variables is important and suggests that they are suitable as instruments. The result also supports the maintained assumption of the paper: to focus on school/community characteristics.

8. Conclusion

In this paper we have provided, to our knowledge, the first quantitative assessment of leakage in a large public expenditure program in a developing country. Even though the institutional environment in Uganda is not identical to other low-income (Sub-Saharan African) countries, we believe our estimate of leakage can nevertheless be viewed as a first approximation of similar programs elsewhere. Furthermore, we have argued that resource flows (leakage) are endogenous to school characteristics. Rather than being passive recipients of flows from government, schools use their bargaining power vis-à-vis other parts of government to secure greater shares of funding. Resources are therefore not allocated according to the rules underlying government budget decisions, with substantial equity and efficiency implications. One implication of this finding is that estimates of the actual budget allocation across end-users (in this case schools), requires an understanding of the local political economy. In the case of school funding in Uganda we have argued that this involves studying the bargaining game between the intended user (school) and the provider of funds (the district officials). Three variables seem important in explaining the variation of leakage across schools: school size, income, and the extent to which teachers are qualified. Our results also indicate that a large part of the variation in leakage can be explained by (time invariant) school/community characteristics. Identifying what characteristics matter is an important area for future research.

As an example, anecdotal evidence indicates that the headmaster's relationship with district officials was an important factor in obtaining funding from the local government. Similarly, academically well-performing schools were often favored by district officials because they projected a positive image of them and the district as a whole. Well-performing schools attracted visitors from the center. Local officials, in turn, rewarded them by transferring more capitation grants. These anecdotes are consistent with the school survey data which show that, despite dismal spending outcomes overall, some schools were able to obtain most of their intended capitation grants.

The contribution of this paper is not only empirical. A methodological contribution is the design of a new survey tool—the quantitative service delivery survey—that can be used to gather data on government resource flow and frontline service delivery. In countries with poor accounting systems and in the absence of incentives to maintain adequate administrative records, such a survey can provide policymakers with valuable information both on inputs and outputs of the service delivery system. In addition, information disseminated directly to the public can play a critical role in improving spending outcomes. In fact, the Uganda survey findings prompted a strong response from the central government. It began to publish monthly transfers of public funds to districts in newspapers and broadcast them on radio. It also required primary schools to post notices on all inflows of funds. On the one hand, these measures aimed at empowering the user by lowering the cost of information θ, and strengthening the schools' bargaining position vis-à-vis the districts, whereas on the other hand, they aimed at changing the nature of the game by strengthening the oversight by the central government. Hence, instead of a bargaining game between the schools and district bureaucracies, the new situation could be described as a principal-agent game, with the central government as principal.

An initial assessment of these reforms suggests hugely improved outcomes (Republic of Uganda, 2000). Instead of about 20 percent in 1995 over 90 percent of the intended capitation grants reached the schools in 1999. These qualitative results are in accordance with the bargaining model presented. By lowering the cost of accruing information θ, the school's bargaining position improves, thus leading to lower leakage.

Similar quantitative service provider surveys are presently being implemented in Ghana, Honduras, Mozambique, and Tanzania, and several others are likely to follow suit. We have shown that the type of data collected with such tools on local public goods provision can be used to analyze problems in service delivery systems in developing countries and, in the end, improve policy and outcomes.

References

[1] Ades, A., and R. Di Tella. 1997. National champions and corruption: some unpleasant interventionist arithmetic. The Economic Journal 107(443):1023-42.

[2] Ades, A. and R. Di Tella. 1999. Rents, Competition, and Corruption. American Economic Review 89(4): 982-993.

[3] Alesina, A. 1997. The political economy of high and low growth. Paper presented at the 1997 ABCDE Conference, World Bank.

[4] Azfar, Omar, Satu Kähkönen, Jeffrey Livingston, Patrick Meagher, and Diana Rutherford. 2000. Making Decentralization Work: An Empirical Investigation of Governance and Public Services in Uganda. University of Maryland, College Park. Processed.

[5] Barro, R. 1991. Economic growth in a cross section of countries. Quarterly Journal of Economics 106 (2): 407-443.

[6] Boone, P. 1995. The impact of foreign aid on savings and growth. London School of Economics. Processed.

[7] Boone, P. 1996. Politics and the effectiveness of foreign aid. European Economic Review 40: 289-329.

[8] Burnside, C., and D. Dollar. 1997. Aid, policies and growth. Policy Research Working 1777. World Bank, Development Research Group, Washington, D.C.

[9] Dehn, J., R. Reinikka, and J. Svensson. 2001. Quantitative Service Delivery Surveys. Development Research Group, World Bank, Washington, D.C. Processed.

[10] Di Tella, R., and E. Schargrodsky. 2000. The Role of Wages and Auditing During a Crackdown on Corruption in the City of Buenos Aires. Harvard Business School. Processed.

[11] Dowrick, Steve. 1992. Estimating the impact of government consumption: growth accounting and optimising models. Working Paper 243. Department of Economics, Australian National University, Canberra.

[12] Duflo, E. 2000. Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. Massachusetts Institute of Technology, Cambridge, Mass. Processed.

[13] Fernandez, R., R. Rogerson. 1996. Income distribution, communities, and the quality of public eductaion. Quarterly Journal of Economics 111(1): 135-164.

[14] Fisman, R., and J. Svensson. 2000. Are corruption and taxation really harmful to growth? Firm level evidence. Policy Research Working Paper 2485. World Bank, Development Research Group, Washington, D.C.

[15] Gugerty, M.K., and E. Miguel 2000. Community participation and social sanctions in Kenyan schools. University of Berkeley. Processed.

[16] Hansen, Henrik, and Finn Tarp. Aid and growth regressions. Journal of Development Economics 64(2, April): 547-570.

[17] Hausman, Jerry A. 1978. Specification Tests in Econometrics. Econometrica 46(6, November): 1251-71.

[18] Hirschman, Albert O. 1970. Exit, voice, and loyalty: Response to decline in firms, organizations, and states. Cambridge, Mass.: Harvard University.

[19] Kormendi, R., and P. Meguire. 1985. Macroeconomic determinants of growth: cross-country evidence. Journal of Monetary Economics 16: 121-163.

[20] Landau, D. 1986. Government and economic growth in Less Developed Countries: An empirical study. Economic Development and Cultural Change 35: 35-75.

[21] Levine, R., and D. Renelt. 1992. A sensitivity analysis of cross-country growth regressions. American Economic Review 82 (4): 942-63.

[22] Mauro, P. 1995. Corruption and growth. Quarterly Journal of Economics 110: 681-712.

[23] Miguel, Edward. 2000. Ethnic diversity and school funding in Kenya. Department of Economics, University of California at Berkeley. Processed.

[24] Narayan, D., and L. Pritchett. 1999. Cents and sociability: household income and social capital in rural Tanzania. Economic Development and Cultural Change 47(4): 871-97.

[25] Persson, T., G. Tabellini, and F. Trebbi. 2000. Electoral rules and corruption. IIES, Stockholm University. Processed.

[26] Pritchett, L. 1996. Mind your P's and Q's: The cost of public investment is not the value of public capital. Policy Research Working Paper 1660. World Bank, Development Research Group, Washington, D.C.

[27] Ram, R. 1986. Government size and economic growth: A new framework and some evidence from cross-section and time-series data. American Economic Review 76(1): 191-203.

[28] Reinikka, R. 2001. Recovery in Service Delivery: Evidence from Schools and Health Centers. In R. Reinikka and P. Collier (eds.), Uganda's Recovery: The Role of Farms, Firms, and Government. Regional and Sectoral Studies. Washington, D.C.: The World Bank.

[29] Reinikka, R., and J. Svensson. 2001. Coping with Poor Public Capital. Journal of Development Economics (forthcoming).

[30] Republic of Uganda. 2000. Tracking the flow of and accountability for UPE funds. Report by International Development Consultants, Ltd, Ministry of Education and Sports, Kampala.

[31] Smith, R. J., amd R.W. Blundell. 1986. An exogeneity test for simultaneous equation Tobit model with an application to labor supply. Econometrica 54(3): 679-685.

[32] Svensson, J. 2000a. Foreign aid and rent-seeking. Journal of International Economics 51(2): 437-461.

[33] Svensson, J. 2000b. Who must pay bribes and how much? Policy Research Working Paper 2486. World Bank, Development Research Group, Washington, D.C.

[34] Svensson, Jakob. 1999. Aid, Growth and Democracy. Economics and Politics 11(3): 275-297.

[35] Treisman, D. 2000. The causes of corruption: a cross-national study. Journal of Public Economics 76(3): 399-457.

[36] Wade, R. 1988. Village Republics. Cambridge, U.K.: Cambridge University Press.

9. Appendix

9.1. Proof of proposition 1

Let $\tilde{s}_i^* = g - \frac{\kappa}{\pi}$: that is, $\tilde{s}_i^*$ is the s_i such that (4.4) binds. Consider the case when the school has made the information investment θ. If $s_i \leq \tilde{s}_i^*$ the PTA will choose to initiate a protest and the official's expected payoff is $(1-\pi)(g-s_i)$. If $s_i > \tilde{s}_i^*$ the PTA will not protest and the official's expected payoff is $g - s_i$. Clearly the official will then either choose $s_i = 0$ (in which case the school will protest), or $s_i = \tilde{s}_i^*$ thereby avoiding a protest. Ensuring no protest by providing funding $\tilde{s}_i^*$ is optimal if

$$E\left[x_i \mid no\ protest\right] - E\left[x_i \mid protest\right] = \pi g - \frac{\kappa}{\pi} \geq 0\ . \tag{9.1}$$

Condition (9.1) is most likely to hold when π is large. Thus, for sufficiently high π, the official will ensure enough funding so that no protest will be initiated.

Consider now the situation before the PTA makes its choice whether or not to acquire information about g. Let $\hat{s}_i^*$ be the cutoff value of s_i implicitly defined by (4.5). That is,

$$\int_{\hat{s}_i^*+\frac{\kappa}{\pi}}^{\bar{g}} \left[\pi \left(g - \hat{s}_i^*\right) - \kappa\right] f(g)dg - \theta/n = 0 \tag{9.2}$$

Comparing (9.2) and (4.5) it is obvious that $\hat{s}_i^* < \tilde{s}_i^*$. Thus, if the district official offers $s_i < \hat{s}_i^*$ the PTA will invest θ (per student) and once g is knows also initiate a protest. If π is sufficiently high this will result in expected payoff $(1-\pi)(g-s_i)$ which is strictly lower than $g - \hat{s}_i^*$. Thus, provided that the credit constraint (4.7) does not bind, equilibrium leakage is given by $x_i^* = g - \hat{s}_i^*$.

Differentiating (9.2) yields,

$$\frac{ds}{dn} = \frac{\theta}{n^2\Lambda} \geq 0$$

$$\frac{ds}{d\kappa} = \frac{-\left[\pi \left(g - \hat{s}_i^*\right) - \kappa/n\right] f(\hat{g})\pi^{-1} - \int_{\hat{g}}^{\bar{g}} f(g)dg}{\Lambda} \leq 0$$

$$\frac{ds}{d\theta} = -\frac{1}{n\Lambda} \leq 0$$

where

$$\Lambda = \left[\pi \left(g - \hat{s}_i^*\right) - \kappa\right] f(\hat{g}) + \int_{\hat{g}}^{\bar{g}} \pi f(g)dg > 0$$

Substituting (4.8) into the credit constraint (4.7), yields

$$\kappa + \theta/n_i \leq y_i - u_c^{-1}(1)\ . \tag{9.3}$$

Clearly (9.3) holds for a wider range of parameter values κ and θ the larger average income y_i.

9.2. Data description

- average share of teachers = average share of qualified teachers to total number of teachers in the district-urban-rural location.
- district population = district population (source: Bureau of Statistics, Republic of Uganda).
- mean consumption = mean consumption level in the district-urban-rural location (source: constructed using the 1992-1995 Uganda Household Surveys data).
- PTA income per student = real PTA total income in US 1990 dollars/number of student (adjusted for inflation using end of year calendar data from the Department of Statistics).

- school size = number of students in P1-P7.
- share of intended capitation grant received = capitation grant received as share of what should have been received. The amount that should have been provided is based on the number of students in 1991 (or first year it was recorded), scaled by the ratio between number of students in the school according to the survey and the number of students in the school according to official statistics in 1991.
- share of qualified teachers = share of qualified teachers to total number of teachers.
- students-teacher ratio = students-teacher ratio.

Table 2.1. School income data, 1991–95
(1991 prices in millions of U Sh)

	1991	*1992*	*1993*	*1994*	*1995*
Government	539.8	437.1	606.9	1,017.7	1,202.9
Teacher salaries	213.9	214.7	381.3	748.6	914.6
Capitation grants*	252.1	159.9	152.0	150.4	141.2
Rehabilitation and other	73.8	62.5	73.6	118.7	147.1
Parents (PTA)	772.3	840.5	1,087.8	1,371.8	1,649.9
PTA levies	591.1	609.6	775.2	934.9	1,032.7
Teacher salaries	125.8	134.1	196.0	300.7	475.9
Tuition fees	55.4	96.8	116.6	136.2	141.3
Total	1,312.1	1,277.61	1,694.7	2,389.5	2,852.8
(percent)					
Government	100	100	100	100	100
Teacher salaries	40	49	63	74	76
Capitation grants*	47	37	25	15	12
Rehabilitation and other	13	14	12	11	12
Parents (PTA)	100	100	100	100	100
PTA levies	77	73	71	68	63
Teacher salaries	16	16	18	22	29
Tuition fees	7	11	11	10	8
Total	100	100	100	100	100
Government	41	34	36	43	42
Parents (PTA)	59	66	64	57	58

*****Capitation grants based on what schools should have received; tuition fees are those actually collected from parents; other items are actual receipts by the schools.

Table 3.1. Share of intended capitation grant received
(in percent)

	Mean	*Median*	*St. dev.*	*Maximum*	*Minimum*	*Obs.*
All schools						
1991–95	12.6	0	26.7	115.9	0	944
1995	21.9	0	33.7	108.9	0	208
Regions						
North	11.5	0	22.8	104.4	0	136
West	11.8	0	25.4	109.8	0	143
Southwest	8.1	0	23.7	101.6	0	131
Northwest	7.6	0	22.8	105.9	0	101
East	11.4	0	25.6	107.2	0	137
Northeast	17.5	0	27.2	108.9	0	146
Central	18.3	0	34.3	115.9	0	150
Region-year average	11.8	0	9.2	36.8	0	35

Table 7.1. Descriptive statistics

Variable	*Mean*	*Med.*	*St. dev.*	*Max.*	*Min.*	*Obs.*
Number of students	492	429	350	3,828	35	942
Student-teacher ratio	32.0	31.2	12.3	110	6	942
Percent qualified teachers	68.4	76.9	29.9	100	0	938
PTA income per student [real 1990 US$]	10.1	1.6	36.4	550.7	0	942

Table 7.2. Explaining leakage across schools

Equation	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*
Time	*1991-95*	*1991-95*	*1991-95*	*1991-95*	*1991-95*
Method	*OLS*	*FE-LS*	*FE-LS*	*FE-LS*	*FE-LS*
PTA income per student	0.133	0.421			0.336
	(.107)	(.096)			(.092)
	[.216]	[.000]			[.000]
School size	–0.332		0.828		0.827
	(.114)		(344)		(.324)
	[.005]		[.018]		[.012]
Share of qualified teachers	0.093			0.449	0.397
	(.054)			(.118)	(.124)
	[.088]			[.000]	[.002]
Wald					8.87
					[.000]
F		3.59	3.49	3.54	3.68
		[.000]	[.000]	[.000]	[.000]
H		12.35	21.72	16.74	43.65
		[.000]	[.000]	[.000]	[.000]
No. schools	239	239	239	239	239
No. obs.	938	942	942	938	938
Adj. R^2	.02	.39	.39	.39	.42

Note: Estimation by OLS (column 1) and fixed-effects least squares (cols. 2-5) with random district and year effects. Dependent variable is the share of intended capitation grant received. Standard errors in parenthesis and p-values in brackets. Wald is the test statistic for the null hypothesis that the coefficients on *PTA income per student*, *school size*, and *share of unqualified teachers* are zero, with p-values reported in brackets. F is the *F*-ratio for the null hypothesis that all fixed effects are equal, with p-values reported in brackets. H is the Hausman (1978) test statistic for the null hypothesis that the fixed effects are uncorrelated with the explanatory variables (**z**), with p-values reported in brackets.

Table 7.3. Explaining leakage across schools: Limited dependent variable estimation

Equation	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*
Time	*1991–95*	*1991–95*	*1991–95*	*1995–95*	*1995–95*
Method	*MLE*	*MLE*	*MLE*	*MLE*	*MLE*
PTA income per student	3.061			2.756	0.932
	(.423)			(.423)	(.356)
	[.000]			[.000]	[.009]
School size		3.421		3.043	2.754
		(.780)		(.713)	(.607)
		[.000]		[.000]	[.000]
Share of qualified teachers			3.387	2.559	0.559
			(.550)	(.511)	(.361)
			[.000]	[.000]	[.122]
σ	2.515	2.648	2.551	2.343	1.840
Proportion y > 0	0.26	0.26	0.25	0.25	0.25
LR				120.8	41.23
				[.000]	[.000]
Time effects	No	No	No	No	Yes
No. schools	239	239	239	239	239
No. obs.	942	942	938	938	938

Note: Estimation by maximum likelihood. Dependent variable is the share of intended capitation grant received. Standard errors in parenthesis and p-values in brackets. LR is the likelihood ratio test statistic for the null hypothesis that the coefficients on *PTA income per student*, *school size*, and *share of unqualified teachers* are zero, with p-values reported in brackets.

Table 7.4. First-stage regressions

Equation	*(1)*	*(2)*	*(3)*	*(4)*
Time	*1991–95*	*1991–95*	*1991–95*	*1991–95*
Dep. Variable	*PTA income per student*	*School size*	*PTA income per student*	*School size*
Method	*FE-LS*	*FE-LS*	*FE-LS*	*FE-LS*
Mean consumption (district)	1.889	–0.355	1.753	–0.345
	(.632)	(.262)	(.622)	(.263)
	[.003]	[.176]	[.005]	[.191]
Population (district)	2.6E-5	2.1E-6	-4.4E-6	1.9E-6
	(1.4E-6)	(5.7E-7)	(2.0E-6)	(8.4E-7)
	[.053]	[.000]	[.025]	[.023]
Time effects	No	No	Yes	Yes
No. schools	239	239	239	239
No. obs.	942	942	942	942
Adj. R^2	0.82	0.90	0.83	0.90

Note: Estimation by fixed-effects least squares. Standard errors in parenthesis and p-values in brackets.

Table 7.5. Explaining leakage across schools: Instrument techniques

Equation	*(1)*	*(2)*
Time	*1991–95*	*1991–95*
Method	*Conditional MLE*	*Conditional MLE*
PTA income per student	5.320	2.055
	(1.432)	(1.239)
	[.000]	[.098]
School size	24.76	10.15
	(6.213)	(4.586)
	[.000]	[.027]
Share of qualified teachers	0.971	0.577
	(.373)	(.351)
	[.009]	[.101]
Share of qualified teachers (squared)		
σ	2.027	1.839
Proportion y > 0	0.25	0.25
LR	207.2	35.71
	[.000]	[.000]
Time effects	No	Yes
No. schools	239	239
No. obs.	938	938

Note: Estimation by conditional maximum likelihood (Smith and Blundell, 1986). Dependent variable is the share of intended capitation grant received. Standard errors in parenthesis and p-values in brackets. LR is the likelihood ratio test statistic for the null hypothesis that the coefficients on *PTA income per student*, *school size*, and *share of unqualified teachers* are zero, with p-values reported in brackets.

Table 7.6. Explaining leakage across schools: Additional robustness tests

Equation	*(1)*	*(2)*
Time	*1991–95*	*1991–95*
Method	*MLE*	*MLE*
PTA income per student	2.351	0.989
	(.425)	(.360)
	[.000]	[.006]
School size	3.186	2.754
	(.704)	(.610)
	[.000]	[.000]
Share of qualified teachers	2.386	0.575
	(.507)	(.365)
	[.000]	[.116]
Tuition fee per student	1.676	
	(.412)	
	[.000]	
Mean consumption (district)		–6.916
		(6.318)
		[.274]
Population (district)		1.6E-5
		(1.4E-5)
		[.272]
LR1	99.42	41.38
	[.000]	[.000]
LR2		1.45
		[.484]
Time effects	No	Yes
No. schools	239	239
No. obs.	938	938

Note: Estimation by maximum likelihood. Dependent variable is the share of intended capitation grant received. Standard errors in parenthesis and p-values in brackets. LR1 is the likelihood ratio test statistic for the null hypothesis that the coefficients on *PTA income per student*, *school size*, and *share of unqualified teachers* are zero, with p-values reported in brackets. LR2 is the likelihood ratio test statistic for the null hypothesis that the coefficients on *mean consumption* and *population* are zero, with p-values reported in brackets.

Districts Included in the Primary School Survey (1996)