

Do Oil Windfalls Improve Living Standards? Evidence from Brazil

Francesco Caselli and Guy Michaels¹

First version: May 2008. This version: January 2011

Abstract

We use variation in oil output among Brazilian municipalities to investigate the effects of resource windfalls on government behavior. Oil-rich municipalities receive higher revenues (mostly from oil royalties) and report correspondingly higher spending on public goods and services. But survey data and administrative records indicate that this higher reported spending has led to little (if any) increase in social transfers, public good provision, infrastructure, and household income. To explain why some of the oil windfall has apparently gone missing, we show that large oil output increases instances of alleged illegal activities associated with mayors.

¹ LSE, CREI, CEPR, and NBER (Caselli, f.caselli@lse.ac.uk) and LSE and CEPR (Michaels, g.michaels@lse.ac.uk). We are grateful to Facundo Alvaredo, Igor Barenboim, Gadi Barlevy, Marianne Bertrand, Irineu de Carvalho, Claudio Ferraz, Fred Finan, Doug Gollin, Todd Gormley, Steve Haber, Andrea Ichino, Seema Jayachandran, Martin Koppensteiner, Andrei Levchenko, Marco Manacorda, Alan Manning, Halvor Mehlum, Marcos Mendes, Benoit Mojon, Steve Pischke, Steve Redding, Eustachio Reis, Silvana Tenreyro, Adrian Wood, Alwyn Young, and three anonymous referees for comments, data, or both; useful comments were also received from seminar participants at Bologna, Brown, Bruegel, Columbia, CREI, INSEAD, LSE, Oxford, Royal Holloway, Sussex, Toronto, Universita Nova Lisbon, Yale, and Zurich, and from conference participants at AMID/BREAD/CEPR 2009, Development Conference at NYU, ESSIM 2009, OxCarre 2008, NBER Growth Conference in San Francisco, NBER Summer Institute Political Economy Workshop, and ASSA Meetings 2010. We thank Gabriela Domingues, Renata Narita, and Gunes Asik-Altintas for research assistance. Caselli gratefully acknowledges the support of CEP, ESRC and Banco de España, the latter through the Banco de España Professorship.

I. Introduction

Should communities that discover oil in their subsoil or off their coast, rejoice or mourn? Should citizens be thrilled or worried when their governments receive fiscal windfalls? Ample anecdotal evidence on both questions has, over the years, shaken many economists' confidence that the answer is as obvious as it might at first seem. Oil, like other natural resources, has been variously described as a source of "disease" or even, by some, as a "curse," bringing adverse changes in relative prices, corruption, rent seeking, and other ills that result in the dissipation of most possible benefits – if not in extreme cases in an outright decline in living standards. Other types of fiscal windfalls, such as international aid or – in the case of local government – transfers from the central government, also often stand accused of creating similar problems.

To provide systematic empirical evidence on some of these issues we study the effects of an oil-induced fiscal windfall among Brazilian municipalities. Oil endowments, and hence production, vary widely within Brazil, and we argue that conditional on a few geographic controls this variation is exogenous to municipal characteristics. Furthermore, oil-producing municipalities are entitled to royalties, so we can investigate the consequences of oil-related revenues for the public provision of goods and services and living standards.

We begin by documenting that oil production significantly increases municipal revenues, and that the bulk of this increase is accounted for by oil royalties. Evidently, royalty payments are not undone by offsetting changes in other transfers from the state or federal governments (or by tax cuts). The increase in municipal revenues is matched by a similar increase in reported expenditures. Municipalities that receive oil windfalls report significant increases in spending on a variety of public goods and services, such as housing and urban infrastructure, education, health, and transportation, and transfers to households.

Given the significant expansion in reported spending, one would expect sizable improvements in welfare-relevant outcomes for the local population. We therefore look at measures of housing quality and quantity, supply of educational and health inputs, road infrastructure, and welfare receipts. The results paint a complex picture, with no apparent changes in some areas, small improvements in some others, and small worsening in others yet. On balance, however, the data appear to suggest that the actual flow of goods, services, and transfers to the population is not commensurate with the reported spending increases stemming from the windfall, a shortfall that we dub "missing money." To confirm that the windfall does

not trickle down to the population through other channels we look at household income, and find only minimal improvements. We also show that oil-rich municipalities did not experience a differential increase in population. This implies that our results are not driven by a dilution of the benefits of oil abundance. Furthermore, the fact that people do not flock to oil-abundant communities reinforces our message that oil abundance has not been seen as particularly beneficial by the population.

Our finding that oil windfalls translate into little improvement in the provision of public goods or the population's living standards raises an important question: where are the oil revenues going? To partly address this question we put together a few pieces of evidence. First, oil revenues increase the size of municipal workers' houses (but not the size of other residents' houses). Second, Brazil's news agency is more likely to carry news items mentioning alleged corruption and the mayor in municipalities with very high levels of oil output. Third, federal police operations are more likely to occur in municipalities with very high levels of oil output. And finally, we document anecdotal evidence of alleged scandals involving mayors in several of the largest oil producing municipalities, some of which involve large sums of money.

There is a growing empirical literature attempting to provide systematic statistical evidence on the effects of resource abundance and other fiscal windfalls, particularly international aid. The near totality of this work focuses on comparisons between countries.² Using variation within Brazil allows us to circumvent some of the well-known limitations of this approach. First, many of the institutional, cultural, and policy variables that potentially confound the cross-country relationship between resources (and aid) and outcomes are held constant, enhancing our ability to make inference. In addition, as we detail below, we can make plausible claims of exogeneity for our measure of resource abundance. This is rarely the case in cross-country work where typically resource abundance is measured in terms of resource exports (an outcome variable). Finally, we are able to focus on one specific channel of causation: the one operating through the change in the amount of public goods and services brought about by the fiscal windfall. Of course, this benefit also implies a limitation: our analysis is silent on other possible channels through which resource abundance may affect local socio-economic outcomes. Nevertheless, as we review below, most recent work on natural resources is centered on the way resource windfalls

² The "classic" cross-country study on the effect of natural resources is Sachs and Warner (1997), which has spawned a large literature. On foreign aid see, e.g., Rajan and Subramanian (2008) and the papers it references.

reverberates through the political process, so our focus is germane to an area of great concern.

A few recent studies have tried, like ours, to move beyond the cross-country correlations and examine resource discoveries using within-country regional variation [Aragon and Rud (2009), Michaels (2011), Naritomi, Soares, and Assunção (2007), Bobonis and Morrow (2010)]. None of these focus on the fiscal windfall associated with the resource boom. Vicente (2010) compares changes in perceived corruption in Sao Tome (which recently found oil) with Cape Verde (which didn't), and finds large increases in corruption following the oil discovery. On fiscal windfalls, the closest contribution is Litschig's (2008) study of federal transfers to Brazilian municipalities, exploiting discontinuities in the transfer-allocation rule. He finds that these windfalls translate into increased educational spending and gains in schooling.³

II. Oil in Brazil: A Brief Overview

Figure 1 presents a summary of the pace and timing of oil discoveries in Brazil.⁴ Meaningful onshore oil discovery began in the 1940s, and the number of finds peaked in the 1980s. Successful onshore prospecting activity has since dwindled. Offshore oil prospecting is a much more recent phenomenon, with finds growing rapidly from almost nothing in the early 1970s, to a peak in the 1980s. Subsequently, there has been a marked decline in the 1990s, and a significant pick up in the 2000s – the latter not reflected in the figure because the big finds at have occurred very recently. For our purposes, the important take away from the figure is that offshore oil is for all practical purposes a post-1970 development. This is important because later on we show that in 1970 (subsequently) oil-rich municipalities looked indistinguishable from municipalities that did not discover oil later in the century (conditional on appropriate controls).

As of 2005, the Brazilian oil sector accounted for approximately 2% of world oil production,

² Very recently there have been new contributions by Brollo et al. (2010) exploiting the same discontinuity as Litschig to look at corruption outcomes, and Monteiro and Ferraz (2010) who study the effects of oil royalties on political outcomes (campaign contributions, electoral competition, quality of politicians) in Brazilian municipalities. Other related work is the discovery by Reinikka and Svensson (2004) that the vast majority of public funds due to Ugandan primary schools never reach the intended recipients, which is reminiscent of our “missing money” result, and the “missing imports” finding of Fisman and Wei's (2004), which is similar to our “missing money.” The literature on the effects of transfers from central to local governments is, of course, very large, and to the extent that such transfers represent fiscal windfalls our paper relates to this entire line of research. Much of this literature focuses on the possibility of a “flypaper effect,” whereby local public expenditure appears more elastic to federal transfers than to (local) tax revenues [e.g. Hines and Thaler (1995) for a review.]

⁴ Throughout the paper we use “oil” as a shorthand for “oil and (natural) gas.” Oil accounts for about 90% of the value of output of the Brazilian oil and gas sector.

1% of world oil reserves, and 2% of Brazilian GDP (these figures will rise in years to come, as newly discovered oilfields begin to bear fruit). Offshore oil accounts for the vast majority of output. Oil in Brazil is inextricably linked to Petrobras, the oil multinational controlled by the Federal Government, which completely dominates the industry. The industry regulator is Agência Nacional do Petróleo, Gás Natural e Biocombustíveis (ANP). One of ANP's functions is to oversee the calculation of royalties due on each oilfield, collect the payments, and distribute them to recipients. In the unpublished Appendix we give a detailed description of the (very complicated) rules for the allocation of royalties. Here we summarize the main points.

Federal law mandates that Petrobras pay close to 10% of the value of the gross output from its oilfields in the form of royalties. The recipients of royalties include: some federal entities, state governments, and municipal governments, the latter two both directly, and indirectly through the division of a "special fund" into which some of the royalties are paid. Municipal governments are the ultimate beneficiaries of about 30% of the royalty pie, i.e. roughly 3% of the value of gross oil output. This can result in substantial royalty revenues for some municipalities: in the top 25 municipalities by per capita oil output, royalties accounted for about 30% of municipal revenues in 2000.

A municipality's royalty income depends on several factors. Some of these factors are purely geographic, and we discuss them in greater detail below. Other determinants of royalty participation, however, are not geographic. For example, municipalities on whose territory is located infrastructure for the storage and transportation of oil and gas or for the landing of offshore oil, or are even only "affected" by such operations, are also entitled to some. Furthermore, some components of the royalty allocation scheme depend on the size of the municipality's population. For these reasons, royalty income is not a credible exogenous measure of the windfall received by municipalities due to oil. This consideration plays an important role in our identification strategy, which we discuss below.

Another source of "Petro-Reais" for oil-producing municipalities is the "Special Participation" (not to be confused with the "Special Fund" mentioned above), a tax on oilfield output – a royalty in all but name – part of which, once again, is given out to municipalities bearing a close geographic relationship with the corresponding oilfields. The overall value of the "Participação Especial" is similar to the overall value of royalties. For example, in 2004 royalties amounted to R\$5735 Millions, while the Participação was R\$5995 Millions. However, royalties

are more important to municipalities, which receive between 20 and 30% of the royalties while producing/facing municipalities are only entitled to 10% of the Participação [de Oliveira Cruz and Ribeiro (2008), Mendes et al. (2008)].

III. Conceptual Framework

Our focus is the impact of oil-related revenues on the provision of public good and services. In order to organize ideas we adapt a simple framework from Caselli and Cunningham (2009). The key decision maker is an agent, or group of agents, with executive control over the windfall revenue. In Brazilian municipalities this tends to be the group formed by the Mayor and his or her close associates (including some members of the local legislature) so we simply refer to this (group of) agent(s) as the Mayor. The Mayor's problem is how to allocate this revenue.

We classify the Mayor's alternative uses of the revenues into three broad categories. The first category is the provision of public goods and services. The second category may be termed "unproductive self-preservation," and it covers activities the Mayor engages in with the sole goal of improving his chances of re-election. These may include: creation of fictitious "patronage" jobs whose "workers" receive wages without delivering the corresponding (meaningful) amount of labor services in exchange for votes at a future election; use of funds to buy votes; or use of funds to buy favorable media coverage. The third category includes revenues the mayor embezzles to enrich himself (and his close associates).⁵

These options are summarized in the following "Mayor's private budget constraint":

$$A_t = G_t + B_t + C_t, \quad (1)$$

where A_t is the value of oil-related revenues, G_t is spending on real public goods and services, B_t is self-preservation spending, and C_t is the value of funds stolen for the private benefits of the Mayor and his circle. To be sure, the boundaries between G_t , B_t and C_t in reality are somewhat blurred, but the abstraction adopted here is still useful.

Needless to say, the "Mayor's private budget constraint" is very different from the "official" budget constraint, i.e. the budget constraint associated with the officially-published revenues and expenditures by the municipality. The official budget constraint reads simply

⁵ In principle there are other possible categories, such as lowering taxes and increasing municipal saving, but our results below show that neither is very relevant in the Brazilian municipal context, so we omit these for simplicity here. Explicitly considering these options would not change the conclusions of this section.

$$A_t^* = G_t^*,$$

where A_t^* is reported revenues and G_t^* is reported spending on goods and services. In corrupt municipalities we would typically expect $A_t^* < A_t$, or under-reporting of revenues, and $G_t^* > G_t$, i.e. over-reporting of spending on (real) public goods and services.

Note that in the two budget constraints we have implicitly normalized other sources of revenue to zero. This is for expositional simplicity and without loss of generality, since we are interested here in comparative statics with respect to changes in A_t . We are unable in this paper to say very much on the uses of non-resource revenues.

We will see later that A_t^* is roughly equal to A_t in the Brazilian context, possibly because oil royalties are allocated by a federal agency, which publishes the data widely and openly. Hence, the difference between G_t^* and G_t , or between reported spending and real public good provision, represents money used for illegitimate goals by the Mayor. In order to achieve this diversion of funds the Mayor could use various techniques. To create patronage jobs the Mayor can increase “official” public employment by giving jobs to political supporters. And to recycle public money into private funds for vote buying or self-enrichment, the Mayor can award fictitious public-procurement contracts to close associates who then rebate (most of) the money back to him.

We assume that the Mayor is infinitely lived and derives utility from the funds he is able to divert to himself.⁶ Hence, his flow utility is $U(C_t)$, where we assume $U'(C_t) > 0$ and $U''(C_t) < 0$. In every period, however, the Mayor is up for re-election, and he is re-elected with probability $\pi(A_t, G_t, B_t)$. The probability of winning re-election will typically be the outcome of a complex political game involving the Mayor, potential political challengers, and the voters. The probability of surviving in office may depend *directly* on the magnitude of the resource-income stream, A_t , as a larger resource pie may attract more and more talented challengers into the political fray, making it harder for the incumbent to win re-election. Hence, we conjecture that π is decreasing in its first argument.⁷ The probability of political survival may also depend on the amount of services provided to the community, G_t , if voters reward politicians who perform effectively in the common good. So we assume that π is increasing in G_t . Finally, to the extent

⁶ There is a two-term limit on re-election in Brazil and our results would be qualitatively similar if we gave mayors a two-period horizon. However as discussed above “Mayor” is a stand-in for a group of people, and this group will seek to find ways around the formal term limit by running another one of its member when the incumbent’s term expires. Hence, our infinite-horizon formulation may be empirically more accurate.

⁷ Contributions emphasizing increased power struggles following resource windfalls include Tornell and Lane (1999), Ross (2001a, 2001b, 2006), Mehlum, Moene and Torvik (2006a, 2006b), and Caselli and Coleman (2008).

that recipients of patronage jobs and vote-buying transfers reciprocate by supporting the re-election of the incumbent (or readers of media that has “sold out” to the government allow themselves to be influenced), the probability of staying in office is increasing in B_t .⁸

Our assumptions result in the following maximization problem for the Mayor

$$\text{Max } \sum_t \pi(A_t, G_t, B_t)^t U(C_t),$$

subject to the flow budget constraint (1). The choice variables are the infinite sequences of C_t , B_t , and G_t . Note that for simplicity we are abstracting from time discounting and we are normalizing utility outside of office to zero.

To tremendously simplify the model we assume $A_t=A$ for all t .⁹ With this, it is straightforward that the solution to the Mayor’s problem involves constant values for G , B , and C , and that these values maximize the object $U(C)/[1- \pi(A,G,B)]$ (always subject to (1)). To further simplify the exposition, we assume away cross-derivatives in π and rewrite this function as $\pi(A,G,B) = \pi_A(A) + \pi_G(G) + \pi_B(B)$. Then, assuming for simplicity that the π s are differentiable and concave, the first order conditions for an interior optimum are

$$U'(C) = \pi_G'(G) U(C) / [1 - \pi_A(A) - \pi_G(G) - \pi_B(B)] \quad (2)$$

and $\pi_G'(G) = \pi_B'(B)$. Equation (2) says that in an optimum the utility cost of an extra unit of foregone consumption must equal the benefits of an extra unit of provided public goods and services. The latter equals the improvement in the probability of surviving an extra period in office multiplied by the value of staying in office one more period. The condition $\pi_G'(G) = \pi_B'(B)$ says that the same must be true when the foregone consumption is used to buy votes (and of course it also says that the marginal benefit of extra public goods and services equals the marginal benefit of extra vote-buying).

The condition $\pi_G'(G) = \pi_B'(B)$ implicitly defines a function $B(G)$. Plugging this in (2), and using the implicit-function theorem we can find the effect of additional natural-resource revenues on the provision of goods and services – the main focus of our empirical work. We get

$$dG/dA = [-U''(1-\pi) + U'\pi_A' + U'\pi_G'] / [-U''(1+B')(1-\pi) - U\pi_G''].$$

⁸ See Robinson, Torvik, and Verdier (2002), and Acemoglu, Robinson, and Verdier (2004), for models emphasizing patronage as a source of votes for the incumbent. Needless to say if voters were perfectly informed (and fully rational) such practices would actually be detrimental to an incumbent’s political survival. However their ubiquity suggests incumbents actually find them useful.

⁹ In reality there is considerable growth of A_t over the years, so this assumption is certainly not realistic. However our empirical evidence focuses on differences across municipalities, and municipality differences in oil abundance are very persistent over time. Hence, our simplification comes at fairly modest cost in terms of guiding the interpretation of our estimates.

Assuming that both π_G'' and π_B'' are negative (the latter assures that B' is positive), the denominator is unambiguously positive. Hence, the sign of the response of G to A is the sign of the numerator. The first term in the numerator is positive. It captures the effect of diminishing marginal utility, which induces a desire to smooth consumption and increase one's chances of re-election by providing more goods and services. The second term is negative. Resource windfalls increase political competition, reducing the chances of political survival. This reduces the effective planning horizon of the current Mayor leading to more current consumption and less investment in public goods and services. The third term is positive. The resource windfall increases the value of staying in office (captured by U') and, if spending on public goods and services has a positive effect on the likelihood of political survival, then public spending increases. The sign of the response of public-service provision is thus a priori ambiguous.

Many other mechanisms that reinforce this ambiguity could be easily added to the analysis. If the Mayor also puts some weight on social welfare, we could write the utility function as $U(C,G)$. In this case, an increase in A expands the Mayor's budget constraint, and to the extent that his problem is homothetic this should contribute to an increase in G . In particular, a public-minded Mayor can increase his current utility by providing more goods and services (without reducing his consumption). Other possibilities would be opened by allowing for significant non-convexities in the function π . For example, there may be fixed costs in vote buying, as buying votes requires setting up an ad hoc operation. If such non-convexities are present, it is not difficult to generate examples where the Mayor switches from a G -based strategy to a B -based strategy over a certain threshold for A , thereby imparting a negative relation between G and A . Also, we could extend the range of possible outcomes. A richer model would have three possible outcomes instead of two: re-election, failure to be re-elected and return to private life (or politics as a party official), and loss of office and punishment following prosecution for corruption. If the punishment for corrupt activities is independent of the amount stolen, one could observe a shift from G to C over a certain threshold of oil revenues, again with a decline in G when the threshold is crossed.

Since the effects of an increase in resource income on real public-good provision are ambiguous, they need to be assessed empirically. This is the goal of the present paper.

IV. Specification, Data, and Identification

IV.A Specification

Our units of observation are Brazilian áreas mínimas comparáveis (AMCs), statistical constructs slightly larger than municipalities, for which we have detailed outcome variables (see data section). The main results of the paper are from instrumental variable (IV) estimation of the following model:

$$\Delta W_m = \delta + \theta \Delta R_m + \rho X_m + u_m, \quad (3)$$

where W_m is a set of AMC outcomes, including reported spending on various municipal-budget outcomes, real provision of public goods and services, transfers, household income and poverty rates, etc.; R_m , is an AMC-level measure of municipal revenues; and X_m is a set of the following AMC-level geographic controls: latitude, longitude, an indicator for whether the AMC is on the coast, distance from federal and state capital, a state capital dummy, and state fixed effects.

The set of instruments is $[Q_m X_m]$, where Q_m is a measure of AMC-level oil output. In other words we instrument changes in municipal revenues by municipal oil output. The idea behind the instrumental-variable approach is to isolate the average effect of a Real of municipal revenue due to oil (a “Petro-Real”). For this interpretation to be legitimate, we need (i) Q to affect ΔR , and (ii) Q to be (conditionally) random and to affect ΔW only through its effect on R . That (i) is the case will be shown in Section V.A, while the case of (ii) is made in Section IV.C.

The exact period over which we take first differences in (3) depends on availability of data on outcomes and municipal revenues, but in most cases it is 1991-2000 (i.e. between the last two censuses). Oil output is observed in 2000. In instrumenting the change in revenues between two dates by oil output in the final date we are proxying the change in oil output by the final level. This is not a bad approximation as oil output in the early 1990s was much lower than in 2000, and so were oil prices. This approximation introduces measurement error in the instrument, but that in itself is not a concern for identification, as long as the first stage is strong.

Specification (3), being in first differences, has the advantage of controlling for baseline characteristics. However, there are a few outcome variables for which we do not have baseline outcomes. In order to use these variables, we also report results from a level specification:

$$W_m = \delta + \theta R_m + \rho X_m + u_m, \quad (4)$$

where the instruments are always $[Q_m X_m]$. While we report the results from (4) for all our outcome variables, in the text we focus on the results for (3) as we feel they are better identified.

We only explicitly discuss the results for (4) when we do not have the first-differenced outcome. For (4) again the year varies but in the majority of cases it is 2000.

As an alternative to specifications (3) and (4), in order to gauge the effects of oil-related revenues, we could have simply regressed the socio-economic outcomes we are interested in on the oil royalties received by AMCs, which we observe. As explained above, however, some of the factors that determine municipalities' royalties are not purely geographic, so royalty income is potentially endogenous to other municipality-level outcomes. For example, more populous municipalities or those that host oil-transportation infrastructure receive more royalties, but those same factors may be correlated with outcomes of interest for other reasons.

IV.B Data

Over the decades the number of Brazilian municipalities has increased, as many of them have split into two or more – largely as a consequence of perverse incentives in the mechanism that assigns federal transfers to municipalities (transfers per capita are strongly decreasing in population size) [Brandt (2002)]. This fragmentation complicates the analysis of panel data on municipalities, as some of today's municipalities did not exist twenty or more years ago. To deal with this problem, Instituto de Pesquisa Econômica Aplicada (IPEA), which is the direct source for much of our data, makes data available at the AMC level. Each AMC contains one municipality (or more), and the area of each AMC remains relatively stable even when municipality boundaries change.

Our empirical work is conducted almost entirely at the AMC level. The main reason for this choice is that we wish to test for random assignment of oil, conditional on a set of geographic controls. In order to do so, we need to compare outcomes before (most of the) oil was discovered. This requires panel data. Altogether, more than 5500 municipalities that exist today are pooled into 3659 AMCs.

Many of the variables we use in the paper are directly available from IPEA at the AMC level. Others are available – or must be first constructed – at the municipality level. In these cases we collapse the municipality-level data to the AMC level using a cross-walk from IPEA. Most of the data we use is self-explanatory, so we relegate a detailed discussion of definitions, methods, and sources to an unpublished appendix (except for a few notes in the main text). Here we only give some of the details our key instrument.

Constructing the value of oil extracted in each AMC involves two steps: (i) build a dataset of

oil output for each oilfield; (ii) allocate the oil output of each oilfield among municipalities according to an appropriate rule based on their mutual geographical relationship. Step (i) is relatively easy, as since 1999 ANP reports detailed price and production data for each oilfield. This gives us the value of oil produced each year in each oilfield from 1999 to 2005.

Step (ii) differs for onshore and offshore oil. For offshore oil we take advantage of the geographic component of the royalty-allocation formula. As discussed, Petrobras pays royalties (through ANP) for oil extraction to municipal governments, and one component of the royalty allocation formula is based on the principle that a certain percentage of the value of the output of each offshore oilfield must be paid to the “municipalities facing the oilfields.” To implement this principle a mechanism had to be devised to determine for each oilfield which are the “facing” municipalities. The principle that has been followed according to Brazilian law apportions the royalties based on the fraction of the oilfield that lies within each municipality’s borders’ extension on the continental shelf. The resulting percentage allocation is collected in a document called “Percentuais Médios de Confrontação” or average shares of “facing,” i.e. shares of each municipality in an offshore oilfield based on the “facing” criterion. We use these shares to allocate oil output from each field to the various municipalities.

For onshore oil we were able to use a simpler algorithm. We combined GIS data on the (terrestrial) boundaries of municipalities with similar data on the boundaries of onshore oilfields. We then shared equally the oil from a certain field among the municipalities that lie above it.

There are 124 municipalities (or 103 AMCs) with a stake in at least one (onshore or offshore) oilfield. Figure 2 shows a map of Brazil with AMC boundaries and oilfields.

Table 1 present summary statistics from various subsamples in our dataset. Column (1) reports figures calculated from the subsample formed by the 3556 AMCs that do not share any oilfield. Column (2) is based on the subsample of all 103 AMCs endowed with oil, whether onshore or offshore. In column (3) we show data from the subsample of 31 AMCs that only have offshore oil (in other words, they have some offshore oil and no onshore oil). As we discuss below, in this paper we use the municipalities in the “No oil” column as our “control group,” while the municipalities in the “offshore only” subsample are our “treatment group.” All our results from now on, therefore, exclude AMCs with onshore oil. The reason for this choice of control and treatment groups is explained in detail at the end of Section IV.C. Detailed results using alternative treatment groups (e.g. “all oil AMCs,” “all oil AMCs with oil discovered after

1970,” and “AMCs with onshore oil only”) are reported in the working paper version [Caselli and Michaels (2009)], and generally paint a very similar picture of the effects of oil windfalls on the provision of public goods and services as that described in the current paper.

Table 1. Summary Statistics for Brazilian AMCs

	(1)	(2)	(3)
	No oil	All with oil	Offshore only
GDP per capita in 2002 (Brazilian R\$2000)	4,313	6,813	8,058
Municipal revenues per capita in 2000 (Brazilian R\$2000)	508	539	699
Population in 2000	45,116	90,944	93,372
Latitude	-16.6	-12.7	-17.8
Longitude	45.0	39.3	42.1
Coast dummy	0.04	0.55	1.00
Distance to the federal capital (kilometers)	1,016	1,271	1,180
State capital dummy	0.006	0.039	0.032
Distance to the state capital (kilometers)	244.6	99.4	99.2
Mean oil output per capita in 2000 (Brazilian R\$2000)	0	2,824	3,894
90th percentile of oil output per capita in 2000 (Brazilian R\$2000)	0	6,949	6,949
95th percentile of oil output per capita in 2000 (Brazilian R\$2000)	0	10,249	21,319
Maximum oil output per capita in 2000 (Brazilian R\$2000)	0	45,221	45,221
Observations (AMCs)	3,556	103	31

Notes: Each AMC includes one or more municipalities. All values reported are means, unless otherwise specified. Municipal revenues in 2000 are only available for 3,242 out of a total of 3,659 AMCs.

There are sizable differences in average GDP per capita, municipal revenues and population between oil-rich and oil-poor AMCs.¹⁰ These differences, however, are clearly not causal. As is clear from the geographic variables also reported in the table, the distribution of oil is far from uniform throughout Brazil. Oil-rich AMCs tend to be systematically to the North and to the East of non-oil ones. To identify the causal effect of oil we therefore control for these geographic characteristics and for state fixed effects.

The table also reports statistics from the distribution of our constructed measure of oil output per capita.¹¹ It is important to keep in mind that our oil output measure corresponds to a gross output concept, so it is not directly comparable to the GDP numbers in the table. Nevertheless the following back-of-the-envelope calculation can be used to get a sense of the importance of

¹⁰ To convert R\$2000 to 2008 US dollars the appropriate conversion factor is roughly 1. Our reason for reporting GDP for 2002 (instead of 2000 as for the other variables) is discussed later.

¹¹ We should point out that there are a few cases of zero oil output even among the “oil AMCs.” This is because the oil-AMC dummy is constructed based on having a positive share in an oilfield that was operating in 2007 (see Appendix 2). Some of these fields were still in the development stage (or still undiscovered) in 2000.

oil in oil-rich municipalities. In the national accounts, value added in the oil sector is about 40% of gross output. Applying that percent to the average gross output number in Table 1 we find that oil accounts for almost 20% of GDP in the offshore-only oil AMCs (17% in the full sample of oil AMCs). Another important message from Table 1 is that there is massive variation in oil output within oil-rich subsamples, with the 90th, 95th, and 100th percentiles all being large multiples of the mean. This underscores that our identification of oil's effects comes as much from variation within oil-rich municipalities as from variation between the no-oil and the oil-rich subsamples.

IV.C Identification

We begin the discussion of our identification assumptions by arguing that Q_m is exogenous to local characteristics and local shocks. We start this argument by showing that a number of municipal-level outcomes did not differ in oil-rich and oil-poor AMCs *before oil was discovered*. As we have just seen, oil and non-oil AMCs differ in a number of geographical characteristics. This means that oil is spuriously correlated with other covariates. But our claim is that oil is as good as randomly assigned *conditional on geographic covariates* (state fixed effects, longitude, latitude, distance to federal capital, distance to state capital, state-capital dummies, and coastal dummies). In other words, once we compare oil and non-oil AMCs with similar geographic characteristics, oil-abundance is essentially random.

To establish conditional random assignment we relate pre-discovery socio-economic outcomes to *subsequent* oil abundance, as measured by oil output in 2000. The outcomes we check are municipal outcomes in 1970, as explained below. Table 2 shows results from a sequence of cross-sectional regressions of the following form

$$Y_{mt} = \delta_t + \eta_t Q_{m,2000} + \theta_t X_m + w_{mt}, \quad (5)$$

where Y_{mt} is GDP per capita in AMC m and year t , and $Q_{m,2000}$ is oil output per capita in AMC m in the year 2000.¹²

¹² The construction of the GDP numbers (both aggregate and sectoral) appears to be based mainly on firm- and consumer surveys as well as on tax returns. A description of the principles underlying the construction of these numbers can be found in IBGE (2008).

Table 2. Tests of Conditional Random Assignment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Outcome in 1970						
	GDP per capita	Average years of schooling among people aged 25 and over	Fraction illiterate among people aged 15 and over	Residential capital per capita	Percent of households with electric lighting	Percent of households with toilets linked to main network	Percent of households with water linked to main network
Oil output per capita in 2000	-0.032 (0.021)	-0.005 (0.005)	0.116 (0.122)	-0.023 (0.015)	-0.350 (0.181)	-0.096 (0.184)	-0.142 (0.131)
Observations (AMCs)	3,587	3,587	3,587	3,587	3,587	3,587	3,587

Notes: Each cell reports coefficients from a regression using a cross-section of AMCs (each AMC includes one municipality or more). The sample includes all AMCs without oil and AMCs with offshore oil only. All values are in Brazilian R\$2000. All regressions control for latitude, longitude, coast dummy, state capital dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

Our main reason for focusing on the period since 1970 for our falsification test is that going back before 1970 would significantly reduce the number of AMCs, due to boundary changes during and before the 1960s. In addition, as we mention shortly, we also perform falsification tests on outcomes other than GDP, and some of these (particularly related to housing – an important variable for us) are not available before 1970, irrespective of the level of AMC aggregation. On the other hand, most oil discoveries (and nearly all of the offshore discoveries) were made after 1970, so not much is lost by not presenting results for the pre-1970 period.

Table 2 shows that conditional on the geographic covariates, oil output per capita in 2000 is uncorrelated with GDP per capita, and various measures of education and health in 1970. This in itself goes a long way in providing support for the identification of models (3) and (4). However, even if initial conditions were invariant to the oil abundance (and to 2000 oil output), one could in principle still be concerned that among oil AMCs the quantity of oil extracted, say, in 2000 is endogenous to other AMC-level shocks occurring after discovery. Similarly, one could be concerned that prospecting decisions and discovery events after 1970 could have been influenced by shocks occurring after 1970.

We argue that this is implausible. Oilfield operations in Brazil over the sample period were carried out by Petrobras, a global hydrocarbon giant with access to global factor and product markets. Neither its highly specialized equipment, nor its equally-specialized labor force could realistically be expected to be drawn locally, so local factor prices should not be a consideration.

Other than the physical presence of the oil, and the morphological characteristics of the oilfield, we think it utterly unlikely that Petrobras is influenced by temporary local conditions in deciding how much oil to extract from a given oilfield, and even less that it is swayed by local economic outcomes in its prospecting plans.

Another possible concern is that municipalities compete to lobby and/or bribe Petrobras to drill near them, or to influence the amount of oil extracted in a given location. This is exceedingly unlikely. First, municipalities are tiny and it is nearly inconceivable that they will have the political heft and financial resources to sway the decisions of Petrobras, one of the world's biggest companies. It is interesting in this respect that in a regression predicting oil output in 2000 with longitude, latitude, coast dummy, distances from state and federal capital, and a state-capital dummy (and state fixed effects) distance from state capital has a significant positive coefficient (p-value 0.04) while the state-capital dummy has a (borderline significant) *negative* coefficient (p-value 0.11) – a result that is hardly consistent with the view that oil output is affected by political influence (distance from federal capital is insignificant - results available on request). Second, unlike many Brazilian institutions, Petrobras actually has a strong record and reputation for integrity – at least in recent years. This record has been explicitly recognized by international NGOs operating in the natural-resource area, e.g. Transparency International (2008).

Even if oil deposits are randomly assigned, and oil production is exogenous to local considerations, we still need to argue that oil output affects outcomes of interest (mainly spending by the local government, provision of public goods, services, and transfers, and household income) only through the revenues it generates for the municipality. Support for this claim is based on showing that oil production has little or no effect on non-oil economic activity, as measured by non-oil GDP. This is seen in Table 3, where the specification is

$$Y_{m2002} = \delta + \eta Q_{m,2002} + \theta X_m + w_m, \quad (6)$$

i.e. the same as in Table 2 except that now we regress GDP on the *contemporaneous* value of oil output instead of a constant, post-discovery value.

At this point we need to explain why we use 2002 data (rather than 2000 data) for analyzing local GDP. It turns out that measured GDP in oil-abundant municipalities experiences a dramatic discrete drop between 2001 and 2002. An investigation of the data-construction measures behind the IPEA figures reveals that up to 2001 inputs into oil extraction were misattributed to the AMC

where operations headquarters were located, rather than – correctly – to the AMC were the extraction took place. This mistake resulted in a vast overestimate of oil GDP at the AMC level, because it essentially amounted to using gross oil output to measure oil GDP. Needless to say, the overestimate of oil GDP carried over to aggregate AMC GDP, which was thus also grossly overestimated. The year 2002 is the first year for which this mistake was removed.¹³

Table 3. The Effect of Oil Output on GDP, by Sector (2002)

	(1)	(2)	(3)
Dependent variable: GDP per Capita in	All sectors	Industry	Non-Industry
Oil output per capita	0.390 (0.027)	0.376 (0.008)	0.014 (0.019)
Observations	3,587	3,587	3,587

Notes: Same notes as Table 2. Additional notes: Industry includes manufacturing, mineral extraction, civilian construction, and public utilities. The calculation of GDP in industry (and total GDP) from oil changed in 2002.

In interpreting the coefficients in Table 3 it is important to recall that the right-hand-side variable, oil output, is a measure of *gross output*, while the left-hand-side, municipal GDP, is a measure of *value-added*. Consider what this implies, for example, for the regression in column (1), where the dependent variable is aggregate AMC GDP and the coefficient on oil output is approximately 0.4. Because aggregate GDP is the sum of oil and non-oil GDP, this 0.4 is the sum of the direct effect of \$1 worth of oil extracted on oil GDP and its indirect (or spillover) effect on non-oil GDP. As already mentioned, at the *national* level the share of oil GDP in gross oil output is also fairly stable and around 0.4. Under fairly standard assumptions average and marginal shares of GDP in gross output are the same, so if national numbers are representative of local production relations the results in column (1) are *prima facie* evidence that oil production has little (if any) positive or negative spillovers on non-oil economic activity.¹⁴

¹³ This mis-measurement does not invalidate the falsification exercise using GDP presented in Table 2. The point of that exercise was to show that differences among municipalities were not systematically related to oil abundance before (virtually all of) the offshore oil discoveries. Inflation in oil GDP numbers in oil-rich municipalities would only work against our case, by tending to make the effect of oil to seem to “kick-in” earlier than it did.

¹⁴ In formulas, begin by the identity $GDP = \text{NON-OIL GDP} + \text{OIL GDP}$. From column 1 we have $d(\text{NON-OIL GDP})/d(\text{Gross oil output}) + d(\text{OIL GDP})/d(\text{Gross oil output}) \approx 0.4$. From data at the national level $d(\text{OIL GDP})/d(\text{Gross oil output}) \approx 0.4$ So that $d(\text{NON-OIL GDP})/d(\text{Gross oil output}) \approx 0$. Needless to say, it would have been cleaner to simply obtain a measure of non-oil GDP and regress it on oil output. Regrettably, despite numerous attempts, we have been unable to obtain the figures used by IBGE for oil GDP, so we cannot net it out of aggregate GDP to obtain non-oil GDP. We do know that oil GDP at the municipal level is computed by distributing Petrobras value added according to a geographical formula similar to the one used by ANP to allocate (the geographical component of) royalties to municipalities [IBGE (2008) and e-mail exchanges with IBGE staff].

Brazilian AMC GDP can also be disaggregated into industrial (manufacturing, construction, mining, and utility services) and non-industrial (agriculture, government, and services) GDP. In columns (2) and (3) we look at the effects of gross oil extraction on these two subaggregates. Since oil GDP is part of industrial GDP, column (2) has much the same interpretation as column (1), and since coefficients are still stable and close to 0.4 they suggest that in the typical oil-rich AMC oil production has little if any spillovers on other industrial subsectors. Similarly, column (3) shows essentially no spillovers from oil to the service sector. This last result is important in establishing that the no-spillover conclusion does not rely on an (admittedly uncertain) estimate of the share of oil GDP in gross oil output, as is the case for aggregate GDP or industrial GDP.

There is reason to expect that the extent of spillovers from oil production to the rest of the economy may differ depending on whether the oil is located onshore or offshore. As we discussed above neither onshore nor offshore oil production are likely to draw directly from local factor markets. However, onshore oil production could affect the composition of demand on non-oil product markets. In particular, it could increase the relative demand for personal services to the oilfield workers and business services to the oilfield operations. In the absence of migration flows to fulfill this demand (and we show below that such migration has not materialized), this would lead us to expect onshore oil to shift the composition of non-oil GDP away from industry and towards services, a particular (though not necessarily malign) case of Dutch Disease.

In results reported in the working-paper version of this paper, we find some support for this hypothesis. In onshore-only oil AMCs the effect of oil on industrial value added is less than one-for-one, as the coefficient on gross oil output falls approximately to 0.3. Continuing to use 0.4 as the rule of thumb for the share of value added in gross oil output this implies that a one Real increase in onshore oil GDP causes a 25-cent decline in non-oil industrial output.¹⁵ At the same time, however, we find a symmetric positive effect on non-industrial output: the coefficient of about 0.1 implies that one extra Real of oil GDP increases non-industrial GDP by 25 cents. It seems, then, that onshore oil causes some minor reallocation of local productive factors from industrial to non-industrial activities.

The results for onshore oil are the main reason why in this paper we focus on the “offshore-only” subsample as our control group. For AMCs with offshore oil only, oil seems to have no *market* effects on economic activity. We infer that any effect from oil likely arises from the

¹⁵ Throughout the paper we use “cent” for “Centavos,” or one hundredth of a Real.

revenues it brings to the municipal government, making this the ideal control group to study the effects of oil-related fiscal windfalls on the provision of public goods and services.

Admittedly, once we focus on the offshore group only we rely on a small sample of “treated” AMCs, and the results are bound to depend on variation among a handful of AMCs. For example, the top two AMCs ranked by oil output per capita are critical to identify the effect of oil abundance on revenues.¹⁶ While this is obviously not a problem conceptually (that’s where the variation is!), it is perhaps useful to note that all the main conclusions of the paper go through when using the full sample with all municipalities (where the results are extremely robust to taking out subsets as large as the 10 top AMCs, again as measured by oil output per capita). As mentioned, these results are reported in detail in the working-paper version.

Another issue relevant to identification is the role of population flows. Since our outcome variables are per capita, and since for many of the outcomes we tend to find little if any positive welfare effect from oil abundance, one possible concern is that oil discoveries in a certain locale attract migratory flows which dilute the benefits on a per-capita basis. Appendix Table A1 shows that there is no significant effect of oil on population, so our conclusions below are probably not driven by changes in the denominator.

As a final robustness check on our identification strategy, we re-estimated the regressions in our paper using only the AMCs that have offshore oil (and no onshore oil) and the adjacent AMCs. The benefit of this alternative strategy is that it uses AMCs that were likely more similar to those that produce oil before oil discoveries took place. The cost is that this alternative strategy reduces sample size, and that nearby AMCs might be indirectly affected by oil. Hence, in general, we prefer the full-sample results and do not report these estimates in the paper. We note, however, that the results using this alternative strategy were generally very similar (both in magnitudes and precision) to the results that we do report (we flag at the appropriate points the few occurrences where these results differed).

¹⁶ The top two AMCs by oil output include the municipalities of Rio das Ostras, Casimiro de Abreu, Macaé, Quissamã, and Carapebus, all of which are well-known large royalty recipients. When they are simultaneously omitted from the offshore-only sample, the point estimates do not change very much but the standard errors increase massively. When AMCs are dropped one at a time all the results in the paper are robust.

V. Results

V.A Oil Abundance and the Local Government Budget - Revenues

We begin the results section by confirming that oil increases municipal revenues. Column (1) of Table 4 shows estimates from a cross-section regression for 2000. The coefficient implies that one Real of gross oil output increases total local-government revenues by 3 cents. One shortcoming of this estimate is that municipal revenues in 2000 are missing for about 11 percent of the AMCs. In column (2) we use 2001 values to impute the missing observations for 2000 (so we are now missing only about 3 percent of AMCs), with almost no change in the results.

In column (3) we investigate the sources of the increase in revenues, by looking at the effect of oil production on royalties. The increase in royalty income accounts for almost two-thirds of the overall increase in municipality income due to oil production.¹⁷ The bulk of the remaining one-third is almost certainly accounted for by the Participação Especial discussed in Section II.

Table 4. Effect of Oil Output Per Capita on Municipal Revenues Per Capita

	(1)	(2)	(3)	(4)
	Total revenues in 2000	Total revenues in 2000 (see footnote)	Royalties from oil in 2000	Change in total revenues 1991-2000
Oil output per capita in 2000	0.0308	0.0311	0.0181	0.0293
	-(0.0039)	-(0.0039)	-(0.0020)	-(0.0032)
Observations (AMCs)	3,183	3,484	3,484	3,466

Notes: Same notes as Table 2. Additional notes: in columns (2), (3), and (4) missing 2000 revenues are predicted from 2001 revenues using a linear regression. In column (4) missing 1991 values are predicted from 2000.

Column (4) shows a specification in semi-differences, i.e. with the 1991-2000 change in revenues regressed on 2000 oil output – which, for reasons discussed above, is also a reasonable proxy for the change in oil output. The coefficient is very close to the coefficient in the level regression for 2000. Note that column (4) is essentially the first stage for our main IV estimates of equation (3) [and column (2) is the first stage for equation (4)].

One implication of Table 4 is that the money received from oil operations is not offset by a reduction in federal government transfers to the local government. Indeed, the fact that the increase in revenues is larger than the royalties suggests that there is not even a partial offset. Similarly, since revenues increase substantially, it does not seem that municipal governments

¹⁷ Column 3 only includes AMCs included in column 2. When we include all municipalities for which we have royalty income the coefficients are the same up to the fourth decimal figure.

take advantage of royalty income to cut local taxes.¹⁸

V.B Oil Abundance and the Local Government Budget - Spending

So oil brings money to the local government. What does the local government do with it? To be answering this question, Table 5 shows the effect of oil on *reported* spending. To establish a baseline, the first row of the top panel shows simple OLS regressions of spending on some of the functions that account for the largest shares of the average municipality budget on total revenues. The most important items are Education and Culture, on which municipalities report spending about 27 cents of each Real that comes into their coffers, and Health and Sanitation and Housing and Urban Development, each of which receives about 10 cents on the Real. Transportation and Transfers to Households also receive significant shares of spending by function.¹⁹ Overall, total reported spending accounts for about 90 cents of every Real of revenue, consistent with the fact that Brazilian aggregate municipal statistics show a surplus for 2000.

It is difficult to interpret the OLS estimates as causal, due to multiple sources of endogeneity and measurement-error bias. But even if such causal interpretation were possible, these estimates would describe the mean reported allocation of revenues independent, of their source. In order to causally identify the utilization of oil-related revenues, in Panels B and D we turn to our empirical models (1) and (4), where municipal revenues are instrumented for by oil output. In other words, we treat the regressions in Table 4 as first-stage regressions in a two-stage least-square estimation of the effect of increases in revenues on spending.

¹⁸ Municipal taxation in Brazil is relatively unimportant, so there are little taxation that could be cut. Not surprisingly, regressions of changes in local tax revenues on changes in municipal revenues, instrumented as usual by oil output, with our usual controls, yield a coefficient of -.011 with a standard error of .015.

¹⁹ Education spending by municipal governments is mostly allocated to primary schooling. Health spending includes local clinics and hospitals. Housing comprises the planning, development and construction of housing in both rural and urban areas. Urban Development includes urban infrastructure. Transfers to households include “Social Assistance” (to the aged, to the handicapped, to children and communities) and “Social Security.” We do not have the year 2000 breakdown of these two items but in 2004 (and subsequently) the latter accounted for about 2/3 municipal social transfers. Nevertheless, social security is probably fairly tightly linked to retirement patterns, and hence to the demographic structure of the AMC’s population. Hence, we conjecture that social assistance is more discretionary and hence the relevant component at the margin.

Table 5. Effect of Municipal Revenues from Oil on Municipal Expenditures

	(1)	(2)	(3)	(4)	(5)	(6)
	Education and Culture	Health and Sanitation	Housing and urban development	Transportation	Social Transfers	Total
Changes (1991-2000)						
A: OLS	0.274 (0.012)	0.121 (0.008)	0.085 (0.007)	0.047 (0.006)	0.070 (0.005)	0.871 (0.028)
B: IV	0.142 (0.017)	0.106 (0.011)	0.185 (0.021)	0.126 (0.019)	0.049 (0.006)	0.841 (0.038)
Obs. (AMCs)	3,355	3,355	3,355	3,355	3,355	3,355
Levels (2000)						
C: OLS	0.274 (0.010)	0.123 (0.007)	0.095 (0.005)	0.067 (0.006)	0.067 (0.005)	0.902 (0.021)
D: IV	0.131 (0.017)	0.105 (0.010)	0.186 (0.017)	0.132 (0.019)	0.049 (0.006)	0.829 (0.029)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484

Notes: Dependent variable is changes in municipal revenues per capita from 1991-2000 in Panels A and B and municipal revenues per capita in 2000 in Panels C and D. Cross sectional regressions on a sample of no oil and offshore-oil-only AMCs. Panels A and C: OLS. Panels B and D instrument municipal revenues per capita with oil output per capita. Brazilian R\$2000. Missing 2000 (1991) revenues predicted using 2001 (1992) values and a linear regression. All regressions control for latitude, longitude, coast dummy, state capital dummy, distances to the state and federal capitals, and state dummies. Robust standard errors are in parentheses.

Our IV results show that the largest reported beneficiary of the increase in government revenues from oil is Housing and Urban Development, with almost a fifth of the marginal “oil Real.” Education and Transportation share second place, with, respectively, 14 and 13 cents. Health receives about 10 cents, and Transfers to households 5 cents. Note that we find significant differences between the IV and OLS results, though we cannot tell whether these differences are due to biases in the OLS coefficients (which could be either upward or downward), or to the fact that municipalities treat oil-related revenues differently. The overall effect on spending using the IV estimates is about 83 cent per Real.

V.C. Oil Abundance and Public Service Provision

Table 5 shows that oil-related revenues feed increased reported spending on housing and urban services, transportation, education, health, sanitation, and transfers to households. The purpose of this section is to look at various measures of related real outcomes, to see to what extent the increased reported spending actually improves living standards.

Table 6. Effect of Municipal Revenues from Oil on Housing and Infrastructure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Residential capital per capita	Rooms at home per 1000 people aged 16-64	Pop. not living in favelas (%)	Pop. with electricity (%)	Pop. with garbage collection (%)	Pop. with piped water (%)	H.holds with mains water (%)	H.holds with toilets linked to network (%)	Km of municipal paved roads per million people in 2005
Changes (1991-2000)									
A: OLS	0.00043 (0.00008)	0.018 (0.029)	-0.0004 (0.0001)	0.0021 (0.0007)	0.0092 (0.0019)	0.0029 (0.0008)	0.0043 (0.0011)	0.0016 (0.0014)	
B: IV	-0.00106 (0.00054)	-0.128 (0.083)	-0.0064 (0.0008)	0.0004 (0.0019)	-0.0055 (0.0040)	-0.0016 (0.0021)	-0.0153 (0.0050)	-0.0214 (0.0045)	
Obs. (AMCs)	3,466	3,466	3,466	3,466	3,466	3,466	3,466	3,466	
Levels (2000, except for column 9)									
C: OLS	0.00018 (0.00016)	0.160 (0.033)	-0.0003 (0.0002)	0.0020 (0.0004)	0.0023 (0.0007)	0.0025 (0.0007)	-0.0028 (0.0012)	-0.0037 (0.0016)	0.174 (0.051)
D: IV	-0.00112 (0.00059)	0.225 (0.157)	-0.0142 (0.0016)	0.0015 (0.0016)	0.0078 (0.0030)	-0.0051 (0.0031)	-0.0196 (0.0108)	-0.0112 (0.0089)	-0.022 (0.013)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484	3,484	3,484	3,484

Notes: Same notes as Table 5.

Table 6 looks at a variety of housing, urban service and infrastructure outcomes: the overall value of the residential housing stock, a proxy for housing quantity (rooms per person) and measures of quality of housing and infrastructure, namely the fraction of the population living in favelas, connection to electricity, water and sewage networks, piping, garbage collection, and extent of roads under municipal jurisdiction. All these variables bar roads are constructed from the micro-data of the Brazilian household census (Censo Demográfico)²⁰, except the length of roads under municipal supervision, which we constructed using administrative records.

²⁰ Residential capital values are based on Census data on housing characteristics and location, which are then converted into Reais through a hedonic model. The number of rooms is the total number of rooms, not just bedrooms. We also note for readers unfamiliar with Brazilian data that the Brazilian “census” is really a representative sample covering approximately 12% of the population.

Table 7. Effect of Municipal Revenues from Oil on Education, Health & Transfers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Municipal teachers	Municipal classrooms	Municipal teachers	Municipal classrooms	Municipal hospitals	Municipal clinics	Social transfers
	Changes from:						
	1996-2000	1996-2000	1996-2005	1996-2005	1992-2002	1992-2002	
A: OLS	3.2 (0.4)	1.4 (0.2)	3.5 (0.5)	1.8 (0.2)	0.001 (0.012)	0.155 (0.031)	
B: IV	0.4 (1.3)	-0.2 (0.5)	3.2 (1.6)	0.7 (0.6)	0.011 (0.006)	0.166 (0.062)	
Obs. (AMCs)	3377	3377	3378	3378	3466	3466	
	Levels in:						
	2000	2000	2005	2005	2002	2002	2000
C: OLS	4.9 (0.4)	3.0 (0.2)	5.0 (0.4)	3.3 (0.3)	0.051 (0.009)	0.299 (0.034)	-0.008 (0.009)
D: IV	1.6 (0.8)	0.7 (0.5)	4.3 (0.9)	1.5 (0.6)	0.011 (0.005)	-0.034 (0.055)	-0.005 (0.002)
Obs. (AMCs)	3481	3481	3484	3484	3484	3484	3484

Notes: Same as Table 5. Dep. Var. in columns (1)-(6) are per million people; column (7) per capita.

Our preferred IV results in first differences are statistically indistinguishable from 0 in four cases, and statistically significantly *negative* (at the 5% level) in the other four cases. Specifically, oil-induced municipal revenues appear to reduce the value of residential capital, increase the percent of people living in favelas, and reduce the percent of households receiving water from the main network and having toilets connected to the municipal sewage network. No single measure of housing and infrastructure quality appears to improve with oil-related revenues. We do not have a base-year measure for the extension of municipal paved roads so for this variable we have to rely on the (somewhat less satisfactory) level regression (model 2). This shows no association with oil-related revenues in the IV regressions.

In Table 7 we look at actual inputs into education (number of teachers and of classrooms, both in 2000 and 2005) and health (hospitals and clinics in 2002), and certain transfers received by households in 2000 (these include transfers for the alleviation of poverty, unemployment benefits, and incentives for schooling for poor families).²¹

The results on education and health are slightly more encouraging than those for housing and road networks. While there is no statistically significant effect on classrooms (columns (2) and

²¹ For the first-difference specification we use 1992 and 1996 as our base years for health and education variables since those are the earliest years for which we have outcomes comparable to those in our later year of data.

(4)), in our preferred IV specification in differences larger oil-related revenues appear to be associated with increases in municipal teachers over the period 1996-2005 (column (3)). The point estimate implies that a million Reais of extra revenue leads to the hiring of 3 teachers. This hiring also appears to be lagged, as we do not find a corresponding increase for the 1996-2000 period (column (1)). Even so, these numbers seem disappointing. According to Table 5, a 1 million increase in municipal revenues leads to about a R\$140,000 increase in reported spending on education, so if all spending on education was on teachers and classrooms, they would imply that hiring a new teacher costs $140,000/3 \approx R\$47,000$ (interpreting our results as indicating that there are no new classrooms built). Given that per capita income is roughly R\$4,000, this implies that either primary teachers are paid in the order of 12 times the average income, or there are substantial amounts of missing money in the education budget.²²

In columns (5) and (6) we look at municipal health infrastructure. In our preferred differenced-IV specification there is no statistically significant effect on hospitals (column (5)), but we find a significant increase in clinics (column (6)). The coefficients can be interpreted as saying that a R\$10M increase in oil-related revenues leads to the construction of approximately 1.7 extra clinics. While R\$10M is an enormous amount in the context of Brazilian municipalities, it is difficult to say with confidence whether this number is too large or too small, or just right, as health spending is probably targeted at other items as well, so we can't infer the "effective price" of a health-care facility for an oil-rich municipality from this figure alone.

Finally, in column (7) we look at the effect of oil-related revenues on poverty- and unemployment-related social transfers from the population census. Since we have no baseline we have to look at the level effect. There is no indication whatsoever that these welfare-like payments increase with oil revenues. To the contrary, the coefficient is significantly negative!

Taken together, the results from Tables 5, 6 and 7 are potentially troubling. Reported spending on housing, transportation, education, health, and social transfers all respond strongly to revenues from oil, but when we look at indicators of real outcomes in these areas we find effects that seem extremely small compared to the reported budget items. The only possible exception is in the health-infrastructure area, where we are unable to benchmark our estimated real outcomes to assess their magnitude vis-à-vis what one should plausibly expect.

V.D Limitations of the Results on Public Service Provision

²² There are many items in the education budget, but teachers and classrooms should make up the bulk of it.

We now discuss some important limitations to the conclusions we have just drawn.

(i) *Wrong Outcomes*. It is possible, in principle, that we are looking at the wrong outcomes. Namely, it could be that increased spending on housing, transportation, etc. show up in variables other than the ones we used to identify outcomes in these areas. Alternatively, it is conceivable that much socially-productive spending is misclassified under the headings of Table 5, and shows up in areas entirely different from the ones we have drawn our outcome variables from. This is a real concern because the allocation of municipal spending across various functions is notoriously inaccurate, so conceivably the “true spending” is elsewhere.

In principle, we could try to address this concern directly by adding further outcome variables to the left hand sides of our regressions. However, a truly exhaustive search over all the possible socio-economic outcomes would quickly become unmanageable, and we must draw the line somewhere. The variables we have selected and used in Tables 6 and 7 reflect our – admittedly subjective and informal – ex-ante assessment of their relevance to households’ welfare *and* the likelihood that municipal governments will be able to influence these outcomes. Our failure to find convincing evidence of productive use of oil revenues is spread over a sufficiently wide range of variables that it seems implausible that we have systematically oversampled from the subset of outcomes with no positive effects. However, as a partial way of addressing these concerns, in the next section we further look at the effects of oil revenues of household income, as an alternative (and summary) measure of living standards.

(ii) *Time to build*. Another possible concern is that we fail to identify positive effects because spending produces benefits only with some lag. For example some of the spending is directed at infrastructure projects and these may take a few years to complete. But the comparison between columns (1)-(2) and (3)-(4) of Table 7 (where outcomes seem to respond more strongly when a lag) lends only partial credence to this view when one considers that oil output (and the corresponding transfers to oil-rich municipalities) was much higher in 2005 than in 2000. In further assessing the importance of this concern a number of considerations are relevant. First, only some of our outcome variables are plausibly subject to “time to build.” In particular, there is no great delay needed to hire teachers or mail transfer payments to households. Indeed, the fact that the lag effect appears for teachers but not for classrooms seem inconsistent with a time to build interpretation. Second, there is no effect whatever of spending in 2000 on municipal roads in 2005, even though transportation is one of the significant winners from oil revenues in the

reported spending. And even the positive education outcomes in 2005 are suggestive of inefficient spending, as argued above.

Third, it is important to keep in mind that municipal revenues from oil are persistent over time. AMCs with relatively large revenues in 2000 tend to have had relatively large revenues for several years, so our coefficients should not necessarily be interpreted as measuring the impact effect of contemporaneous revenues. Rather, they should be thought of as capturing the cumulated effects of several-year worth of Petro-Reais. To make this point more concrete, in Appendix 1 we present an attempt to estimate cumulative municipal revenues since 1991 associated with one Real worth of oil production in 2000. In performing this calculation we make conservative assumptions, and our estimates are almost certainly downward biased. Still, we find this lower bound to be in the order of at least 6 cents, namely double the “flow” coefficient for 2000 we reported in Table 4. This confirms that municipalities receiving oil-related revenues in 2000 had been doing so in previous years as well.

When one takes into account that not all our outcome variables are plausibly subject to time to build constraints; that even when observing outcomes with a lag we find either no effect or effects that still seem too small; and that flow municipal revenues understate significantly the cumulative revenues received from 1991 to 2000, there is probably no strong reason to worry about time to build.

(iii) *Crowding out of state and federal spending.* There is considerable overlap between the items in the municipal, state and federal functional spending budgets, and the “division of labor” between different levels of government in Brazil is often blurred. It is therefore conceivable that state and federal bodies withdraw funding in areas where they are aware of increased spending by municipal governments. But we can rule out this concern in the areas of education, health, and transportation, because our road, teacher, classroom, and health establishment variables all come with the qualifier “municipal,” as they refer to provision by municipal governments. They are therefore net of state and federal contributions, and as such not subject to crowding out.²³

Still, it is interesting to see if federal and state provision in these areas responds to municipal oil revenues. We investigate this question in appendix Table A2. We find very few state- and

²³ There are federal and state transfers to municipalities earmarked for education, health, and road construction, but crucially these transfers show up as revenues in the municipal budget, as well as spending items in the corresponding functional categories (we have confirmed this in private communications with a Brazilian fiscal expert). But now recall that we observe both revenues and reported spending increasing with oil, so crowding out of these transfers cannot be the explanation for our results.

federally-provided goods and services are significantly affected by oil-related municipal revenues. Only in the health field we find a robust negative coefficient. Note that this deepens the gloom of the picture that is emerging. In Table 7 the only plausible piece of good news was an increase in *municipal* clinics, and now we have just seen that oil revenues reduce *state and federal* clinics (and also hospitals).

Another variable that speaks to the issue of crowding out is the value of federal contracts per capita. Brazilian mayors can individually negotiate deals with the federal government to finance specific projects. It is possible that mayors who are awash in oil revenues exert less effort to secure such funding.²⁴ In column (8) of Table A2 we test this conjecture and find weak support for it: the coefficient is not quite significant, but close. However, the crowding out is minimal: one Real of oil-related revenues only displaces about 1.5 cents of federal contracts, so the overall impact of oil on the funds available to the municipality is virtually unaffected. In conclusion, we don't feel that crowding out is a likely first order driver of our results.²⁵

(iv) *Large confidence intervals*. Perhaps the most serious concern is that – as is frequent for IV estimates – the standard errors are quite large. Our failure to find significant improvements in the provision of public goods and services may therefore stem as much from imprecision of the estimates as from a true lack of effectiveness of public spending. There is a handful of cases where we can rule this out convincingly, namely when the upper bound of the 95% confidence interval is *negative*. This obviously occurs in Table 6 in Columns (1), (3), (6), and (7), as well as in Table 7 in column (7). At a minimum, we can conclude with considerable confidence that oil-related fiscal windfalls *destroys* the provision of at least some public goods and services.

For the outcome variables with positive 95% upper bounds on the coefficient estimates we present in Appendix 2 back-of-the-envelope calculations to infer possible implied “minimum prices” of public goods (or bundles of public goods). This exercise leads to a fairly clear conclusion in the case of the housing and urbanization outcomes in Table 6: even if *all* first-difference coefficients in columns (1)-(8) simultaneously took their 95% upper-bound values

²⁴ We are grateful to Fred Finan for pointing out this issue and for giving us the data. If federal contracts were included in municipal budgets, we would already know that total municipal revenues *net* of any crowding out of mayoral effort to secure these contracts increase with oil, and all our inferences would be unchanged. The problem only arises if these contracts are off-budget, so that one may worry that the observed increase in budgetary revenues is offset by a decline in off-budget contract financing. Since we were unable to establish conclusively whether these contracts are on- or off-budget, we report the results in Column (8).

²⁵ Litschig (2008) also looks at whether fiscal windfalls due to discontinuities in the allocation formula for federal transfers crowd out other sources of revenues, and finds no evidence of this.

(and those for which the 95% upper bound is negative are set to 0), it would cost more than R\$200,000 to house and connect to services 5 families – while reasonable independent figures for the same bundle would suggest that the cost should be in the neighborhood of R\$100,000. Similarly, the cost of roads seems excessive, even in the very optimistic scenario that the coefficient in column (9) takes its upper bound value.

Unfortunately, in the case of the education and health variables, the results of the bounding exercise are not as conclusive: if true coefficients *all* take their 95% upper bound values, then the corresponding bundles of goods and services are not necessarily over-priced. Coupled with the fact that our bundles are “incomplete,” in the sense that we do not observe all of the goods and services the oil money is buying, we cannot completely finesse the large-standard-error issue.²⁶ Hence, we can rule out a positive effect on the infrastructure, transportation, and welfare payments measures we observe, but not on education and health. Given that the first three categories are at least as important as the last two in terms of reported spending out of a Petro-Real (about 35 cents compared to about 25 cents), we still suspect that a substantial amount of the money may have gone missing.

V.E. Oil Abundance and Household Income

The results presented so far raise questions regarding the extent to which reported spending increases in oil-rich municipalities materialize in services to the population. Nevertheless, it is still conceivable that the population benefits from the government’s expansion of the budget in ways that are not directly captured by our indicators of public-good provision. Hence, in Table 8 we study the effect of oil-induced government revenue on a summary measure of living standards – household income per capita, which we compute from the Brazilian census.

Column (1) shows that municipal revenues from oil have no effect on average household income in the IV regressions. These results suggest that the reported expansion in the government budget has not lead to aggregate increases in living standards that we have somehow missed in the previous section. In the next five columns we look at the effect of oil on average income within each quintile of the household income distribution. This gives a somewhat more nuanced view than looking at the average effect. In particular, we do find that household income

²⁶ Clearly the probability that all the coefficients take simultaneously their upper-bound value is nearly 0. Another dimension in which our calculations are very aggressive is that they ignore accumulated revenues from previous years that could have translated into a better stock in the cross-section. See the discussion above on time to build.

increases in the bottom two quintiles of the income distribution.²⁷ Nevertheless, it is important to notice that these increases are small: for every per-capita Real of increased revenue (and spending), the increase in income is in the order of ten cents. To benchmark this number, consider this: suppose the government mechanically rebated each Real of oil-related income to all households in the municipality equally (and there were no additional general equilibrium effects). Then *all five* of the coefficients in columns (2)-(6) should *simultaneously* be 1. In the last column of Table 8 we look at the effect of oil on poverty rates. There is no statistical evidence that oil-related municipal revenues reduce poverty.²⁸ Taken together, our findings about oil's effect on local household income reinforce our view that some of the money is missing.

Table 8. Effect of Municipal Revenues from Oil on Household Income Per Capita

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Average	1st quintile	2nd quintile	3rd quintile	4th quintile	5th quintile	% poor
Changes (1991-2000)							
A: OLS	0.17 (0.08)	0.05 (0.02)	0.07 (0.02)	0.13 (0.04)	0.18 (0.07)	0.44 (0.30)	-0.0036 (0.0009)
B: IV	0.13 (0.23)	0.08 (0.03)	0.10 (0.05)	0.11 (0.10)	0.19 (0.20)	0.12 (0.83)	-0.0036 (0.0031)
Obs. (AMCs)	3,376	3,376	3,375	3,375	3,375	3,376	3,376
Levels (2000)							
C: OLS	0.17 (0.12)	0.06 (0.01)	0.09 (0.03)	0.15 (0.06)	0.24 (0.12)	0.27 (0.41)	-0.0020 (0.0006)
D: IV	0.23 (0.27)	0.10 (0.03)	0.11 (0.08)	0.13 (0.13)	0.19 (0.27)	0.62 (0.87)	-0.0045 (0.0027)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484	3,484

Notes: Same as Table 5

VI. Where is the missing money going?

VIA Oil Abundance and Corruption: Suggestive Evidence

Where is the missing money going? One outcome variable that may speak to this issue, albeit very indirectly, is the size of houses enjoyed by municipal employees, which can be identified in

²⁷ We note that the effects of municipal revenues from oil on the top of the income distribution are imprecisely estimated, making it difficult to tell whether (some) wealthy individuals benefitted disproportionately from it.

²⁸ Even abstracting from statistical significance, the point estimate implies minuscule effects. The coefficient estimate implies that municipal revenues due to oil need to increase by 100 Reais *per capita* to see a reduction in poverty of 0.36 of one percentage point. This is more than the average revenue from oil royalties in the oil abundant sample (see Table 1).

census data when the previously-mentioned quality of housing variables are combined with respondents' sector of employment. Table 9 shows that oil-related revenue increases the quality of housing for municipal workers – but, as we already know, not for everyone else.

Table 9. Effect of Municipal Revenues from Oil on Rooms at Home

	(1)	(2)	(3)
	municipal employees	Non-municipal employees	Difference
Changes (1991-2000)			
A: OLS	-0.05 (0.09)	0.02 (0.03)	-0.07 (0.09)
B: IV	0.38 (0.17)	-0.14 (0.08)	0.52 (0.14)
Obs. (AMCs)	3,457	3,457	3,457
Levels (2000)			
C: OLS	0.08 (0.05)	0.16 (0.03)	-0.07 (0.05)
D: IV	0.52 (0.18)	0.21 (0.16)	0.31 (0.13)
Obs. (AMCs)	3,481	3,481	3,481

Notes: Same as Table 5. In 1 (2) dep. var. is rooms per 1000 municipal employees (non-mun. emp.).

To shed more light on the “missing money,” we gather new data on the frequency with which municipalities are cited in the news media in connection with corruption. Specifically, for each municipality we search all the news items in the archives of Brazil's News Agency's (Agência Brasil) website. In a first search, we include the name of each municipality and the (Brazilian) Portuguese words for mayor (prefeito) and for embezzlement (desvio). We then construct a municipality-level dummy that takes the value of 1 if the search delivered at least one hit, and 0 otherwise.²⁹ When we regress this variable on per capita oil output in 2000, with our usual set of controls, we find insignificant estimates. However, when we use the absolute level of oil output in 2000, the estimates become statistically significant (top panel of Table 10). In a second search, we use the word for corruption (corrupção) instead of embezzlement, with very similar results (mid-panel of Table 10). Results using the word for fraud (fraude), which are not shown, are also

²⁹ We used an indicator instead of counts because many municipality names are identical to state names or have other meanings in Portuguese, giving rise to outliers in the distribution of counts. Using an indicator prevents these outliers from dominating the results. The search was over all 5,507 municipalities that existed in 2000. The corresponding AMC-level dummies took a value of 1 if at least one municipality in the AMC took a value of 1.

qualitatively identical. We conclude that high oil output *per capita* does not make a municipality more likely to be cited for irregularities in the news, but more oil output *per mayor* does.³⁰

Table 10. Effect of Oil on News About Fraud and on Federal Police Operations (2000)

	Effect of:	Oil output per capita	Oil output
<u>Dependent variable:</u>			
A: Indicator for news on alleged embezzlement possibly related to a mayor		7.3 (10.8)	0.116 (0.029)
B: Indicator for news on alleged corruption possibly related to a mayor		4.4 (10.0)	0.101 (0.028)
C: Indicator for federal police operation involving a mayor		13.1 (12.0)	0.149 (0.031)
Obs. (AMCs)		3,587	3,587

Notes: Coefficients from regressions using a cross section of AMCs. Coefficients in column (1) [(2)] are multiplied by 10^6 [10^9]. Sample: all AMCs without oil and AMCs with offshore oil only. Brazilian R\$2000. Dependent variables. Panel A: at least one webpage on Brazil's News Agency mentioning embezzlement, the mayor, and the name of one of the municipalities in the AMC. Panel B: same, except the keyword was corruption. Search conducted from 1-9 July 2009. Panel C at least one current or former mayor investigated by a federal police operation from 2003-2008. All regressions control for latitude, longitude, coast dummy, state capital dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

We also examined a list of 723 Federal Police operations for 2003-2008, available from the Federal Police web site. For each operation, we searched the abstract for an occurrence of at least one of the following 12 words: *benefícios, cidade, cofre, corrupção, desvio, fraude, fundo, irregular, licita, municipais, município, prefeito*. We found 241 operations that satisfy this requirement. For each of these operations we then read the top 10 items in a Google search for the name of the operation. If we found that a current or past mayor was involved, we recorded this mayor's municipality. As before, we used this to construct an indicator for whether at least one of the municipalities in each AMC had a mayor involved in a federal police investigation. The results are very similar to the news articles (bottom panel of Table 10).³¹

As an informal check on the nature of the allegations behind news stories and federal police operations, we conducted a broader search of news stories on the 10 most oil-rich municipalities (by oil output in 2000). The results (summarized in Unpublished Appendix A4) show evidence of alleged corruption in 6 of the top 10 oil producing municipalities, suggesting that corruption is

³⁰ The findings about oil per mayor are not mechanically driven by the size of oil producing municipalities, since adding the log of population (or the log of municipal revenues) as a control has little effect on the estimates.

³¹ If we simply use as a dependent variable an indicator for whether any of the abstracts of the 723 Federal Police operations mention one of the AMC's municipalities the results are very similar.

a serious and common concern. Some of the news stories discuss “Operação Telhado De Vidro” (“Operation Glass Ceiling”) in the municipality of Campos dos Goytacazes, the largest oil producer in 2000. In March 2008, a large number of local-government officials at the highest level were accused of diverting the equivalent of 140 million dollars. This is larger than an average municipality’s annual budget, though still only a fraction of the royalties received by Campos dos Goytacazes over the past decade.

VI.B Back to the Conceptual Framework

In Section III we discussed the effects of resource windfalls in terms of a Mayor’s “official” and “private” budget constraints. The results show that Mayors make sure that the official budget constraint is satisfied, or $A^*=G^*$. However, true spending G appears to be less than reported spending G^* , suggesting a role for the other items in the Mayor’s private constraint: self-enrichment (C), and unproductive self-preservation (B).

To try to shed a little light on how the missing money is allocated between C and B we have looked at re-election rates for the Mayor’s party. Specifically, for the 2000 election we have constructed a dummy variable taking the value of one if the elected new Mayor belongs to the same party as the outgoing Mayor (this of course includes as a special case the incumbent Mayor being re-elected). We then regressed this variable on oil output in 2000.³² Neither per-capita oil-output nor total oil output significantly affect re-election probabilities. We also looked at the effect of 2000 oil output on 2004 re-election, and again found no effect.

Note that these regressions for re-election can be interpreted as estimates of the *total* derivative of the re-election probability function $\pi(A,G,B)$ in Section III. Taking our no-effect result to mean that the total derivative is zero, we then have

$$0 \approx \pi_A'(A) + \pi_G'(G) dG/dA + \pi_B'(B) dB/dA.$$

The main empirical result of the paper is that $dG/dA \approx 0$. Note that the Mayor’s private budget constraint $A = G + B + C$ then implies $d(C+B)/dA \approx 1$. We then have

$$0 \approx \pi_A'(A) + \pi_B'(B) dB/dA \quad (7)$$

As discussed in Section III, we would expect $\pi_A'(A)$ to be negative, as oil riches should strengthen political competition, making harder to stay in power. For (7) to hold, then, at least some of the diverted funds/missing money must have been used to illegally support the Mayor’s

³² To be precise, the dependent variable in these regressions is the average of the re-election dummy among the municipalities in the AMC.

(or the incumbent party's) re-election i.e. $dB/dA > 0$.

VII. Conclusions

Oil production generates significant increases in revenues for Brazilian municipalities. These windfalls are matched by reported spending increases, particularly in urban infrastructure and housing, education, health services, and transfers. However, we are unable to detect commensurate improvements in various outcomes that would be expected to respond to the recorded spending increases. Furthermore, increases in household income are at best very modest, and people do not seem to flock to the oil-rich municipalities. Finally, there is suggestive evidence that some of the “missing money” is accounted for by embezzlement.

In the context of a simple conceptual framework for the Mayor's objectives and constraints, these findings suggest that incumbent Mayors are able to divert much of the oil revenues that accrue to the municipality. The diverted funds are probably allocated to a combination of self-enrichment and vote buying, with the latter succeeding in neutralizing any additional political competition brought about by the oil riches (as well as overcoming any possible unpopularity stemming from illegal behavior).

A particularly important question for future research is whether and to what extent the implications of fiscal windfalls from natural resources are different from other types of windfalls, such as foreign aid or federal transfers, and – if so – why.³³ Absent problems of identification, our OLS coefficients capture the effect of the marginal Real of “general” revenues, while the IV coefficients capture the marginal effect of petro-Reais. The fact that OLS coefficients and IV coefficients are very often statistically different may hint at differences in their allocation and effects. Unfortunately, as is obvious, the OLS coefficients are hard to interpret causally.

If there is a systematic difference between the effects of natural-resource based revenues and other types of fiscal revenues, we think that this difference may stem from voter ignorance. Observing and estimating royalty income may be harder than forming an idea of the likely magnitude of other types of incomes. There is some circumstantial evidence that the general public is ill informed about the magnitude of the oil-related fiscal windfall. In Campos dos Goytacazes, mentioned above as the largest oil producer of Brazil, a local news bulletin “Petróleo, Royalties & Região” conducted surveys to assess the local population's knowledge

³³ Dalgaard and Olsson (2008) argue with cross-country data that resources corrupt more than aid.

about oil royalties. In May 2004 respondents chose a range of values for the monthly fiscal receipts due to oil. 66% of respondents underestimated, i.e. picked a range below the “correct” interval of \$R20m to \$R50m. Only 15% overestimated.^{34,35}

References

Acemoglu, D., J. Robinson and T. Verdier (2004) ‘Kleptocracy and Divide-and-Rule: A Model of Personal Rule’, *Journal of the European Economic Association*, 2 (2–3): 162–92.

Bobonis, G. J. and P. M. Morrow (2010). “Labor coercion and the Accumulation of Human Capital,” unpublished, University of Toronto.

Brandt, C. (2002). “The Effects of the Changes in the Revenue Sharing Fund of the Municipalities After the 1988 Federal Constitution,” unpublished, GWU.

Brollo, F., Nannicini, T., Perotti, R. and Tabellini, G. (2010). “The Political Resource Curse” NBER Working Paper 15705.

Caselli, F. and Coleman, J. (2006). ‘On the Theory of Ethnic Conflict,’ unpublished, LSE

Caselli, F. and Cunningham, T. (2009). “Leader Behaviour and the Natural Resource Curse,” *Oxford Economic Papers*, 2009.

Caselli, F. and Michaels, G. (2009). “Do oil windfalls improve living standards? Evidence from Brazil,” NBER Working Paper 15550.

Dalgaard, C-J. and Olsson, O. (2008). “Windfall gains, Political Economy and Economic Development,” *Journal of African Economies*, 17 (Supplement 1), pp. i72-i109.

Ferraz, C. and F. Finan (2008a). “Exposing Corrupt Politicians: The Effects of Brazil’s Publicly Released Audits on Electoral Outcomes,” *Quarterly Journal of Economics*. MIT Press, vol. 123(2), pages 703-745, 05.

Ferraz, C. and F. Finan (2008b): “Electoral Accountability and Corruption: Evidence from the Audits of Local Governments,” unpublished, UCLA.

³⁴<http://www.royaltiesdopetroleo.ucam-campos.br/index.php?cod=1>. Incidentally this quarterly bulletin is a very interesting window on the concerns raised locally by the oil windfall. Recurrent themes are the lack of transparency in the utilization of the royalties, and the perceived ineffectiveness of royalty money in promoting development and living standards. The September 2004 issue pointedly reports that one year of royalty money would be enough to build 18,890 social housing units (*casa populares*) or pay for a presidential election campaign.

³⁵ Specific to Brazil, our findings may suggest that oil-rich municipalities should be given special consideration in the trend towards greater decentralization (Lipscomb and Mobarak 2007) and in the design of audit schemes aimed at curbing corruption (Ferraz and Finan 2008a, 2008b). This focus may become even more important following the recent discovery of huge new offshore fields. Indeed, the issue is of massive current political relevance, as the government has recently proposed a radical reform of the royalty regime.

- Fisman, R. and Wei, S-J. (2004): "Tax Rates and Tax Evasion: Evidence from "Missing Imports" in China," *Journal of Political Economy*, vol. 112(2), pages 471-500, April.
- Hines, J.R. and Thaler, R.H. (1995): "Anomalies: The Flypaper Effect," *The Journal of Economic Perspectives*, Vol. 9, No. 4 (Autumn), pp. 217-226.
- IBGE (2008): "Produto Interno Bruto dos Municípios," Relatorios Metodologicos, 29.
- Lipscomb, M. and Mobarak, A.M. (2007). "Decentralization and Water Pollution Spillovers: Evidence from the Re-drawing of County Boundaries in Brazil," University of Colorado.
- Litschig, S. (2008). "Intergovernmental Transfers and Elementary Education: Quasi-Experimental Evidence from Brazil," unpublished, Universitat Pompeu Fabra.
- Mehlum, H., Moene, K.O., and Torvik, R. (2006a). "Institutions and the Resource Curse," *The Economic Journal* 116(508): pp.1–20
- Mehlum, H., Moene, K.O., and Torvik, R. (2006b). "Cursed by Resources or Institutions?" *The World Economy*, Blackwell Publishing, vol. 29(8), pp. 1117-1131, 08
- Mendes, M.J., Miranda, R.B., and Blanco, F.C. (2008): "Transfêrencias Intergovernamentais no Brasil: Diagnóstico e Proposta de Reforma," Textos Para Discussão 40, Consultoria Legislativa do Senado Federal.
- Michaels, G. (2011). "The Long Term Consequences Of Resource-Based Specialisation," *Economic Journal*, forthcoming.
- Monteiro, J. and Ferraz, C. (2010). "Does Oil Make Leaders Unaccountable? Evidence from Brazil's off shore oil boom," unpublished, PUC-Rio.
- Naritomi, J., Soares, R.J., and Assunção, J.J. (2007). "Rent Seeking and the Unveiling of 'De Facto' Institutions: Development and Colonial Heritage within Brazil," NBER WP 13545.
- Rajan, R.G. and Subramanian, A. (2008). "Aid and Growth: What Does the Cross-Country Evidence Really Show?" *Review of Economics and Statistics*, November, Vol. 90, No. 4, Pages 643-665: 643-665.
- Reinikka, R., and Svensson, J. (2004): "Local Capture: Evidence from a Central Government Transfer Program in Uganda," *Quarterly Journal of Economics*, 679-705.
- Robinson, J. A., Torvik, R. and Verdier, T. (2002). "Political foundations of the resource curse," CEPR Discussion Paper No. 3422.
- Ross, M.L. (2001a): *Timber Booms and Institutional Breakdown in Southeast Asia*, Cambridge University Press.

Ross, M.L. (2001b). “Does Oil Hinder Democracy?” *World Politics*, 53, 325–361.

Ross, M.L. (2006): “A Closer Look at Oil, Diamonds, and Civil War,” *Annual Review of Political Science*, 9, 265–300.

Sachs, J.D. and Warner, A.M. (1997a). “Natural resource abundance and economic growth – revised version,” Working Paper, Harvard University.

Tornell, A. and Lane, P.R. (1999). “The voracity effect,” *American Economic Review*, vol. 89, pp. 22–46.

Transparency International (2008). *Promoting Revenue Transparency: Report on Transparency of Oil and Gas Companies*.

Vicente, P.C. (2010). “Does Oil Corrupt? Evidence from a Natural Experiment in West Africa,” *Journal of Development Economics*, Elsevier, vol. 92(1), pages 28-38, May.

Appendix 1: A Rough estimate of cumulative oil-related revenues over the 1990s

In order to estimate the cumulative revenues associated with oil production in 2000 we estimate a set of regressions of municipal revenues in each preceding year on oil output in 2000. We have done this already in Table 4 for 2000, but we repeat this for every year from 1991-1999. The cumulative effect of one Real of oil in 2000 on revenues between 1991 and 2000 should be the sum of all these coefficients. The results are in the unpublished appendix. Unfortunately, the municipal-level data on total revenues suffer from unacceptably high numbers of missing values in 1998 and 1999. In each of these two years, 25% of the AMCs in the offshore oil sample fail to report municipal revenues, so we can't put any store on the coefficients from these two years. We do have, separately, municipal royalties in 1999 (with no missing values). Hence, as a partial remedy to the “hole” for that year we can regress 1999 royalties on 2000 oil output. We get a coefficient of 0.0100 (s.e. 0.0013). To construct a lower-bound estimate of cumulative revenues between 1991 and 2000 we do the following (i) sum over the *statistically significant* coefficients of revenues on 2000 output (these are 1991-1995, and 2000); (ii) add 0.01 for 1999 from the royalty regression. This gives 0.0638, i.e. roughly double the flow estimate for 2000.

It is important to highlight the reasons why this is a lower bound. First, we have treated the two insignificant estimates in 1996 and 1997 as zeros, even though at least the 1996 figure is borderline significant. And second, we have assumed a zero coefficient for 1998, just because

the sample size is too small for revenues and we have no royalty data at the municipal level.

Appendix 2: A Bounding Exercise for the Cost of Public Goods and Services when the Financing comes from Oil-Related Revenues

To find lower bounds for the cost of goods and services we begin by choosing *upper bounds* for the coefficients in the tables. We begin with the upper bounds of the 95% confidence intervals for the offshore, differenced IV regressions for the housing variables in Table 6. When the upper bound is negative the estimates suggest that oil revenues “destroy” the corresponding services, but to focus on the rosier possible scenario in this exercise we treat negative upper bounds as zero. The upper bound for per-capita residential capital is 0 to the fifth decimal digit, so we ignore it as well. This leaves us with positive upper bounds on rooms at home per 1000 adults, and percentages of the population receiving electricity, garbage collection, and piped water.

To interpret these upper bounds begin with rooms at home per 1000 adults. The relationship we estimated is $\text{Rooms} \times 1000 / [P \times 2/3] = a + 0.0345 \times R/P$, where P is population, R denotes municipal revenues due to oil, 2/3 is a rough estimate of the adult population in the total population, a is a constant and 0.0345 is the upper bound on the estimated effect of oil revenue. This calculation suggests that a R\$1 million increase in oil-related revenues adds 23 new rooms to the housing stock. Similar calculations for the other positive upper bounds imply that a R\$1 million increase in oil-related revenues results in connecting 4.1 more people to the electricity network, and bring garbage collection and piped water to 2.4 and 2.5 more people, respectively.

Now recall that about 20% of the marginal-oil related Real goes to “Housing and Urban Development.” Assuming that the kind of house that gets built with oil-related revenue is made of approximately five rooms, and that a family comprises fewer than four people, we conclude from this bounding exercise that a bundle of (fewer than) 5 housing units and electricity, garbage collection, and water provision for (less than) one family costs *at least* R\$200,000.

This figure seems rather large. The marginal cost of extending electricity, garbage collection, and water to an extra family should be extremely low, so the bulk of these R\$200,000 should be accounted for by the extra five housing units. An official estimate of the cost of a social housing unit (casa popular) is approximately \$R25,000 in 2003 (*Petróleo, Royalties, & Região*, September 2004), equivalent to approximately \$R20,200 in 2000. Hence, 5 new housing units should cost approximately R\$100,000. This seems to suggest that at least some of the

R\$200,000 have gone missing. And recall that R\$200,000 is a lower bound on the cost of the bundle: the bulk of the probability mass places the cost at a much higher level.

There are clearly caveats to this exercise. On the one hand, we have ignored urban and sanitation services destroyed by oil, even according to the upper-bound coefficient. Also, we took the upper bound as the situation in which all 8 coefficients happen to simultaneously take their individual 95% upper-bound value. The probability of this happening is, of course, vastly less than 5%. Both these biases go in the direction of exaggerating the possible benefits of oil-income. On the other hand, there may be other services paid for by the Housing and Urban Development budget that we do not observe. Whether this omission leads to an overestimate or underestimate of the benefits of oil is ambiguous, as we have seen from Table 5 that oil revenues are as likely to destroy as to create new services.

For the municipal-road variable in column (9) we only have level regressions. The 95% upper bound implies that R\$1,000,000 of revenues translate *at best* into 3 meters of road. Table 5 implies that transportation gets about R\$130,000 out of this million. We do not have an independent estimate of the marginal cost of building 3m of municipal road but R\$130,000 seems like an enormous sum. Of course the same caveats as before apply.

Moving on to the education-related variables in columns (1)-(4) of Table 7, the upper bounds from the difference regressions are imply that \$R1 million in 2000 (of which approximately \$R140,000 goes to education according to Table 5) translates into fewer than 3 teachers and 1 classroom contemporaneously. Allowing for a five year lag, however, leads to an upper-bound of 6 teachers and 2 classrooms, which seems much more acceptable.

For municipal health establishments the upper bounds in the difference regressions (coupled with the coefficients in Table 5) imply that R\$100,000 of oil-related revenue spent on health and sanitation in 2000 lead at most to 0.02 hospitals and 0.28 clinics. One way to interpret these numbers is that R\$5M could buy a bundle of one hospital and 4 clinics. This does not seem an outrageous price tag so the bounding exercise is again inconclusive.

The upper bound for social transfers Column (7) is negative so clearly the “no beneficial effect” result is robust.

Figure 1. Number of Oilfields Discovered by Decade

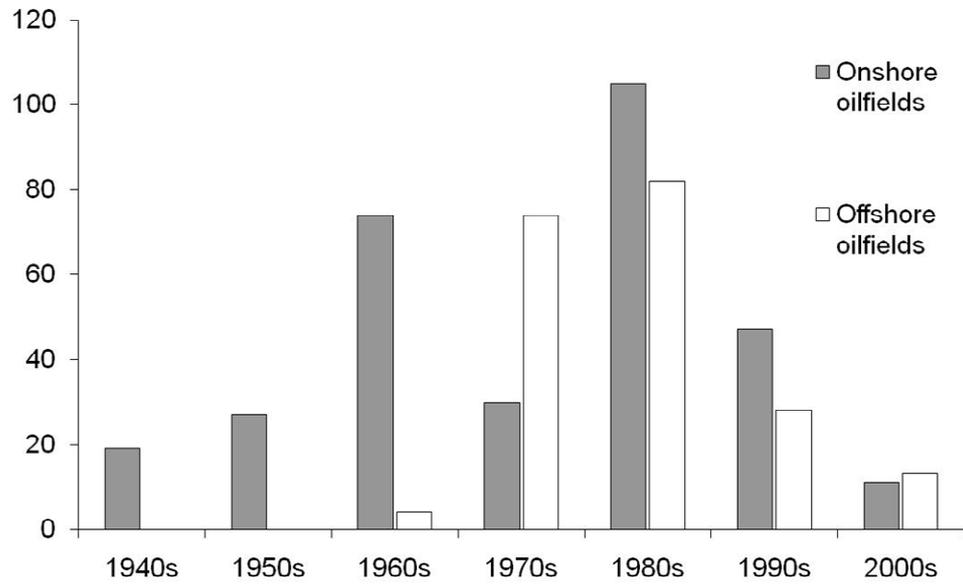


Figure 2. AMCs (from 1970) and Oilfields in Brazil

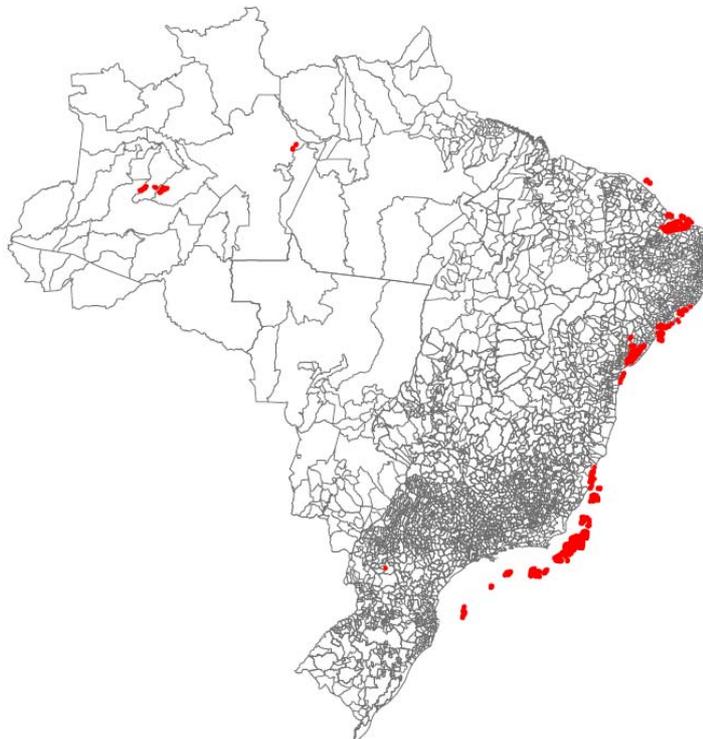


Table A1. No Significant Effect of Oil on Population

	(1)	(2)	(3)	(4)	(5)	(6)
A: Dependent variable: ln(population)	1970	1980	1991	1996	2000	2005
Oil output per capita in 2000	-0.0000091 (0.0000169)	-0.0000089 (0.0000161)	-0.0000061 (0.0000155)	-0.0000013 (0.0000147)	-0.0000005 (0.0000142)	0.0000004 (0.0000134)
Observations (AMCs)	3,587	3,587	3,587	3,587	3,587	3,587
B: Dependent variable: population	1970	1980	1991	1996	2000	2005
Oil output per capita in 2000	-1.8 (2.5)	-2.4 (2.9)	-2.4 (3.1)	-2.2 (3.2)	-2.3 (3.3)	-2.3 (3.5)
Observations (AMCs)	3,587	3,587	3,587	3,587	3,587	3,587

Notes. Each cell reports coefficients from a regression using a panel of AMCs (each AMC includes one municipality or more). All panels include all AMCs without oil. Columns 1 and 2 add all AMCs with oil; Columns 3 and 4 those AMCs with offshore oil only. All values are in Brazilian R\$2000. All regressions control for year dummies interacted with latitude, longitude, coast dummy, state capital dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors clustered by AMC are in parentheses.

Table A2. Test for Crowding Out of State and Federal Investments

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Fed. and State teachers per million people	Fed. and State classrooms per million people	Fed. and State teachers per million people	Fed. and State classrooms per million people	Fed. and State hospitals per million people	Fed. and State clinics per million people	Km. of paved roads under non-municipal jurisdiction per million people in 2005	Value of Fed. contracts per capita
	Changes from:							
	1996-2000	1996-2000	1996-2005	1996-2005	1992-2002	1992-2002		
A: OLS	-2.3 (0.3)	-0.7 (0.1)	-1.1 (0.4)	-0.2 (0.1)	-0.004 (0.003)	-0.132 (0.024)		
Obs. (AMCs)	3,377	3,377	3,378	3,378	3,466	3,466		
B: IV	0.3 (0.5)	-0.3 (0.4)	-0.7 (0.5)	0.0 (0.4)	-0.009 (0.004)	-0.100 (0.029)		
Obs. (AMCs)	3,377	3,377	3,378	3,378	3,466	3,466		
	Levels in:							
	2000	2000	2005	2005	2002	2002	2005	2000
C: OLS	-0.3 (0.2)	0.3 (0.1)	0.8 (0.3)	0.7 (0.1)	0.000 (0.001)	0.001 (0.002)	4.8 (0.6)	0.0247 (0.0041)
Obs. (AMCs)	3,481	3,481	3,484	3,484	3,484	3,484	3,484	3,484
D: IV	-0.4 (0.7)	-0.8 (0.3)	-1.4 (0.8)	-0.6 (0.3)	0.007 (0.003)	-0.008 (0.004)	-0.4 (0.6)	-0.0147 (0.0085)
Obs. (AMCs)	3,481	3,481	3,484	3,484	3,484	3,484	3,484	3,484

Notes: Each cell reports the coefficient on changes in municipal revenues per capita between 1991-2000 (Panels A and B) or on municipal revenues per capita in 2000 (Panels C and D) from a regression using a cross section of AMCs (each AMC includes one municipality or more). Panels A and C report OLS coefficients, while Panels B and D instrument municipal revenues per capita with oil output per capita. The sample includes all AMCs without oil and AMCs with offshore oil only. All values are in Brazilian R\$2000. For municipalities that did not report expenditures or revenues in 2000 (1991), we predicted these using 2001 (1992) values and a linear regression. All regressions control for latitude, longitude, coast dummy, state capital dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses. For municipalities that did not report health establishments, we assumed that there were no health establishments. We have no data on roads or federal contracts for 1991.

