

Do Oil Windfalls Improve Living Standards? Evidence from Brazil

Francesco Caselli and Guy Michaels¹

First draft: May 2008; This draft: June 2009

Abstract

We use variation in oil output among Brazilian municipalities to investigate the effects of resource windfalls. We find muted effects of oil through market channels: offshore oil has no effect on municipal non-oil GDP or its composition, while onshore oil has only modest effects on non-oil GDP composition. However, oil abundance causes municipal revenues and reported spending on a range of budgetary items to increase, mainly as a result of royalties paid by Petrobras. Nevertheless, survey-based measures of social transfers, public good provision, infrastructure, and household income increase less (if at all) than one might expect given the increase in reported spending. To explain why oil windfalls contribute little to local living standards, we use data from the Brazilian media and federal police to document that very large oil output increases alleged instances of illegal activities associated with mayors.

I. Introduction

Should communities that discover oil in their subsoil, or off their coast, rejoice or mourn? After a long collection of seemingly miserable experiences, from Nigeria to Venezuela and from Congo to the Caucasus, economists increasingly suspect that the simple answer (rejoice, of course) is too simple. The comparison between resource-rich African, Latin American and Middle-Eastern slow growers and resource-poor East Asian fast growers also contributes to a suspicious attitude towards resource endowments. Indeed, many social scientists have gone as far as describing natural-resource abundance as a “curse.”

There is a growing empirical literature attempting to go beyond casual observation and provide more systematic statistical evidence. The near totality of this work focuses on inter-country comparisons. Using cross-sectional or panel techniques it typically presents regressions of per-capita growth on proxies for resource abundance. While useful, these exercises have some limitations. Different countries are characterized by very different institutional and cultural features, which may well correlate with resource abundance. Only some of these can be controlled for with available data. Another possible problem is that resource abundance tends to be measured by flows of natural-resource exports (often normalized by GDP, or total exports), which could be argued to be an outcome variable in its own right – and hence not exogenous. Data quality is also a perennial concern in cross-country work, as is the fact that intrinsic

¹ LSE, CEPR, and NBER (Caselli, f.caselli@lse.ac.uk) and LSE and CEPR (Michaels, g.michaels@lse.ac.uk). We acknowledge comments from Marianne Bertrand, Doug Gollin, Todd Gormley, Seema Jayachandran, Andrei Levchenko, Marco Manacorda, Halvor Mehlum, Steve Pischke, Steve Redding, Silvana Tenreyro, Adrian Wood, and Alwyn Young; seminar participants at Bruegel, Columbia, LSE, Sussex, and Zurich; conference participants at OxCarre 2008, NBER growth conference in San Francisco, development conference at NYU, and ESSIM 2009. We thank Gabriela Domingues, Renata Narita, and Gunes Asik-Altintas for research assistance. Caselli also thanks ESRC for financial support.

limitations in the available variables prevent many of these studies from identifying the specific mechanisms through which resource abundance affects outcomes.²

In this paper we develop a new dataset on oil abundance for Brazilian municipalities – local geographical units roughly corresponding to US counties.³ Brazil has not historically been a major oil producer, and even today oil accounts for little more than 2 percent of national GDP. However, as we document below, oil fields are concentrated in certain locations. Hence, for municipalities that have it, oil could potentially have a major impact. For example, we estimate that in the 20 most oil abundant municipalities (as measured by oil output per capita), oil royalties in 2000 accounted for almost a third of municipal revenues (and royalties are not the only oil-related source of revenue for oil-rich municipalities).

Because we use variation within Brazilian states, many of the institutional, cultural, and policy variables that confound the relationship between resources and macroeconomic outcomes at the country level are held constant, enhancing our ability to make inference. In addition, we can make plausible claims of exogeneity for our main measure of resource abundance, which is per capita output from oil fields located in, or off the coast of, each municipality. First, we show that, conditional on a few geographical observables, municipalities with large oil endowments did not look statistically different, along a number of socio-economic dimensions, from non-oil municipalities *before the oil was discovered*. Second, conditional on the distribution of oil deposits across municipalities, we argue that prospecting and extraction decisions are taken by the national oil company, Petrobras, independently of local conditions.⁴ This is especially uncontroversial for offshore oilfields, so we focus our empirical work particularly on municipalities with oil off their coast.

Another advantage of our study is that we can say more about the channels of causation. In particular, we can distinguish between the effects of oil abundance operating through the market, and those operating through the (local) government. Furthermore, thanks to the richness of the set of variables we observe, we can shed an unusual amount of light on the way municipal governments spend their oil revenues and on the effects of such spending on welfare-relevant outcomes.

There are of course disadvantages to moving from the cross-country to a within-country setting. What may be true for Brazil may not generalize to other countries. In addition, we only capture the differential effect of oil between oil-rich and oil-poor municipalities, but any aggregate effect oil may have in the country as a whole would not be detected by our analysis.⁵ Most importantly, there are a number of prominent explanations for the “resource curse” that only operate at the

² The “classic” cross-country study on the effect of natural resources is Sachs and Warner (1997). Other contributions in this vein include Leite, C. and J. Weidmann (1999), Isham, Woolcock, Pritchett and Busby (2005), Kolstad (2007), Collier and Goderis (2007), and Brunnschweiler and Bulte (2008).

³ 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 oil and gas sector.

⁴ This is unlikely to be true when comparing across countries. For example, multinational oil companies may be more or less inclined to prospect and extract in a particular country depending on the level of corruption of that country’s government.

⁵ Given the relatively small size of the oil sector for Brazil as a whole, however, we do not expect this loss to be first order.

national level. For example, some “Dutch-disease” mechanisms work through (or are amplified by) the nominal exchange rate, and would therefore not show up across municipalities. Similarly, our analysis cannot address the hypothesis that resource abundance is a cause of political violence and civil war. We emphasize, therefore, that our work is a complement, and not a substitute, for cross-country analysis.

We begin by investigating the market effects of oil. We do this by regressing measures of municipal GDP, both aggregate and for subsectors, on contemporaneous municipal oil output. The market effects of oil abundance are fairly muted: to a first approximation, the GDP of the non-oil sector remains roughly unchanged in response to oil discovery and production. However, for onshore oil, we also uncover some (modest) composition effects. While non-oil GDP in the industrial sector shrinks somewhat, services GDP expands. For offshore oil there is no change in the composition of GDP, suggesting that offshore oil operations are completely segregated from economic activity on the mainland. This finding is of independent interest, but we also build on it later to identify the effects of oil abundance on municipal governance.

Next, we regress municipal revenues on contemporaneous oil output. Oil abundance generates a significant fiscal windfall for the local government, mostly – but not exclusively – in the form of royalties paid by Petrobras to communities on whose territory (or off whose coast) the oil is located. Evidently, royalty transfers are not undone by offsetting changes in state or federal governments. These large fiscal windfalls imply that oil abundance may have significant effects on municipal outcomes even if the market linkages are muted.

In order to identify the effects of the oil-related fiscal windfall, we begin by regressing municipal expenditures on various budgetary items on municipal revenues, using oil output as an instrument.⁶ The point of the instrumenting procedure is to isolate the effects on spending of the marginal oil-generated dollar. The instrument is valid as long as oil production affects spending (and other municipal-level outcomes) only through the municipal government budget. As discussed above, this is a highly plausible assumption, particularly for offshore oil.

It turns out that oil income results into increases in a wide range of reported budgetary items. About 20-25 cents of the marginal Real of oil-generated income are reported to be spent on “housing and urban development” (which normally accounts for about one-tenth of the budget); about 15 cents go to education and about 10 cents go to health; other items make up for the rest, including about five cents on social security and social assistance.⁷ Almost all of the revenue is accounted for in the form of increased spending. In particular, there is no Alaska-style lump-sum rebate to the population.

Given the significant expansion in reported spending, one would expect sizable improvements in welfare-relevant outcomes for the local population. The main result of the paper is that, to the contrary, the local population seems to experience very small, if any, benefits. In particular, despite the fact that much of the oil windfall is apparently spent on construction and

⁶ Here, and in most of the exercises discussed subsequently, we run regressions both in levels, e.g. spending on contemporaneous revenues in a cross-section of municipalities, and in differences. The latter exercise approximates a diff-in-diff approach to estimation. The results are almost always virtually indistinguishable.

⁷ Throughout the paper we use “cent” for “Centavos,” or one hundredths of a Real.

infrastructure, the size and quality of housing for the general population are at best unchanged, and for some outcomes they actually get worse. Other measures of infrastructure are similarly unaffected by oil income. The evidence on real purchases of education and health inputs is mixed, but even the most positive results are very small compared with the reported increase in spending in these areas. Also, despite the claim of local governments to be spending more on social security and social assistance, we find virtually no effects of increased spending by the local government on household income, or welfare receipts. We also show that oil-rich municipalities did not experience an 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 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 raise an important question: where are the oil revenues going? To partly address this question we put together a few pieces of tentative 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 corruption and the mayor in municipalities with very high levels of oil output (on an absolute, though not per capita, basis). Third, federal police operations are more likely to occur in municipalities with very high levels of oil output (again in absolute terms). And finally, we document anecdotal evidence of scandals allegedly involving mayors in several of the largest oil producing municipalities, some involving large sums of money. To partly explain why senior municipal workers may have thought that they could "get away" with large-scale alleged theft in a country where local elections are held regularly, we note that a survey in the largest oil producing municipality found considerable ignorance among residents regarding the scale of the municipal oil windfall.

After presenting our empirical results we relate our findings to the large theoretical literature on the consequences of resource abundance. As mentioned above, there are several proposed mechanisms that plausibly operate only at the national level, and are therefore invisible to our analysis. Still, we identify as many as seven theoretical conjectures that do have potential applicability at the sub-national level, all of which have distinctive predictions that are implicitly tested in our empirical work. The predictions are rejected in five of these cases. The remaining two, that appear quite consistent with our results, are that: (i) resource abundance changes the composition of non-resource economic activities, in response to demand for business and personal services by oil operations (an effect operating through the market); and: (ii) resource revenue is easier to embezzle than non-resource revenue, leading to an increase in corruption following resource windfalls (a political-economy effect).

A few recent studies have tried, like ours, to move beyond the cross-country correlations and examine resource discoveries in developing countries. Michaels (2008) finds substantial local benefits from oil in the US South, at least for a considerable period of time. This underscores the limitations of extrapolating our Brazilian results, as discussed above. Naritomi, Soares, and Assuncao (2007) single out Brazilian municipalities which were historically associated with sugar-cane production or gold extraction during the colonial period, and find that today they are worse governed, more unequal, and poorer. Bobonis (2008) studies elite behavior in respect to

labor practices and education policies in 19th century Puerto Rico as a function of various regions' suitability for coffee production. Vicente (2008) 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. Also closely related is Litschig's (2008) study of fiscal windfalls to Brazilian municipalities in the form of federal transfers, exploiting discontinuities in government rules. He finds that these windfalls translate into increased educational spending and gains in schooling.⁸

II. Oil in Brazil: A Brief Overview

Figure 1 presents a summary of the pace and timing of oil discoveries in Brazil. Onshore oil was first discovered in Brazil in 1939, and the number of finds reached a peak in the 1980s. Onshore prospecting activity has since dwindled. Offshore oil prospecting is a much more recent story, with finds growing very rapidly from almost nothing in the 1960s, 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 Tupi and Carioca occurred very recently. For our purposes, the important thing to 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 will show that in 1970 (subsequently) oil-rich municipalities looked indistinguishable from municipalities that will not later discover oil (conditional on appropriate controls).

As of 2005, the Brazilian oil sector (i.e. Petrobras) accounted for approximately 2% of world oil production, 1% of world oil reserves, and 2% of Brazilian GDP. (All of these figures will rise significantly when Tupi and Carioca begin production.) Importantly, offshore oil accounts for the vast majority of this production. For example in 2002 offshore oil output was 1,200 million barrels per day on average, while onshore output was about 200 million barrels per day. The relative importance of offshore oil continues to rise steadily.

Oil in Brazil is inextricably linked to Petrobras, the oil multinational. From 1953 to 1997 Petrobras was a fully state-owned monopolist both in oil extraction and refining. Since 1997 the oil industry has been liberalized, and Petrobras partially privatized, though the federal government retains a minority but controlling stake. Despite the liberalization and the appearance of some small new players, Petrobras still completely dominates the industry. Given the essentially monopolistic structure of the industry, the oil sector is heavily regulated. Since 1997 the industry regulator is Agência Nacional do Petróleo, Gás Natural e Biocombustíveis (ANP). One of the many important functions of ANP is to oversee the calculation and distribution of royalties from oil production. In the Appendix we give a detailed description of the (very complicated) rules for the allocation of royalties. Here we summarize the main points.

⁸ There is also a small literature in Portuguese investigating some differences in outcomes between municipalities according to the oil royalties they receive. Perhaps the most ambitious such study is by Postali (2008) who regresses the change in the municipality growth rate of per capita GDP over the two sub-periods 1996-1999 and 2001-2004 on average total municipality revenue over the period 2000-2004. He finds a negative coefficient. Other studies look at the correlation between royalty income and selected items of the spending budget or social indicators in selected sub-regions of the country [e.g. Leal and Serra (2002), Costa Nova (2005)].

Federal law mandates that Petrobras distribute close to 10% of the value of the gross output from its oilfields in the form of royalties. The recipients of royalties include: the ministry of the navy, the ministry of science and technology, 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. The shares of royalties going to these sets of recipients differ between onshore and offshore oil. As a rough order of magnitude, however, in both cases municipal governments are the ultimate beneficiaries of about 30% of the royalty pie, i.e. roughly 3% of the value of oil output. This can result in substantial royalty revenues for some municipalities: in our sample, in the top 25 municipalities by royalty revenue royalties accounted for about 30% of total revenues.

The rules for the allocation among municipalities of the municipal share of royalties also differ between onshore and offshore oil. In both cases, however, a municipality’s participation in the royalties depends on several factors. Some of these factors are purely geographic, and will be discussed 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, as well as for the landing of offshore oil, or 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. Finally, the allocation of the “special fund” is not based on geographic criteria. For these reasons, royalty income is not a credible exogenous measure of the windfall received by municipalities thanks to oil. This consideration will play a very important role in our identification strategy, which we now discuss.

III. Specification, Data, and Identification

III.A Specification.

Our units of observations are all Brazilian áreas mínimas comparáveis (AMCs), statistical constructs slightly larger than the municipality for which we have detailed outcome variables (we’ll explain this shortly). For this population, we present results from two sets of empirical models. The first set of results is generated by OLS estimation of the specification

$$Y_{mt} = \alpha_t + \beta_t Q_{mt} + \gamma_t X_m + e_{mt}, \quad (1)$$

where m indexes AMCs and t indicates year, Y_{mt} is an AMC-level outcome in year t (e.g. AMC GDP), Q_{mt} is a measure of AMC-level oil output, 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, and a state fixed effect. The Greek letters are parameters to be estimated, and e_{mt} collects the effect on Y of the unobservables. Note that we allow the coefficients to be estimated to vary over time, though in practice we do not uncover particularly significant time-series variation. The outcome variables Y_{mt} that we consider for specification (1) are aggregate GDP, sectoral GDP, and municipal revenues. The time coverage is typically 2000-2005. To interpret this exercise as uncovering the causal effect of oil production on Y we will have to argue that Q is uncorrelated with the residual determinants in e .

The second set of results is from instrumental variable (IV) estimation of the following model

$$W_m = \alpha + \beta R_m + \gamma X_m + e_m, \quad (2)$$

where the set of instruments is $[Q_m X_m]$. In these specifications 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, employment; R_m is municipal-government revenue; X_m and Q_m are, as before, AMC-level geographic controls and oil output, respectively; the Greek letters are parameters to be estimated and e_m collects other determinants of the outcomes. Essentially, specification (2) uses specification (1) (with Y being municipal revenues) as its first-stage regression. Note however that variables and coefficients here are not time varying, i.e. this specification is for a single cross-section, typically for the year 2000. This is because of limitations in the time coverage of the data on the outcomes W .

In specification (2) we wish to interpret β as capturing the causal effect of a marginal oil-generated Real of municipal revenue on the outcome W . For this interpretation to be legitimate, we will need to argue that Q affects W only through its effect on R [and, of course, we need to show that we have a strong first stage, which will be seen to be the case when we present the results for (1)].

Specification (2) is most transparently interpreted in levels. However, as a robustness check, and to fully control for baseline characteristics, we also report results where (2) is estimated in first differences. The exact period over which we take first differences depends on availability of data on outcomes, but in most cases it is 1991-2000.⁹

As an alternative to specification (2), 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 AMC, which are observable. However, as explained above, some of the factors determining a municipality's share in the royalties are not purely geographic, implying that royalty income is potentially endogenous to other municipality-level outcomes. In particular, local conditions correlated with our outcomes of interest may also affect whether a municipality hosts oil-transportation infrastructure, the allocation of the special fund, and the size of the population.

In addition, oil royalties are not the only source of oil-related income for municipalities. For example, states within whose (land or maritime) borders are oil fields, also receive some royalties, and by constitutional law they must rebate some of these royalties to their municipalities. Relying on royalties alone may therefore distort the estimation of the budgetary effects of oil abundance. For both these sets of reasons, we use the royalty measure sparingly and, when we do, as an outcome rather than as a determinant.¹⁰

⁹ In the differenced regressions, we use 2000 oil output levels as instruments for changes in municipal revenues. The reason of course is that we do not have 1991 oil output levels. While this is clearly not ideal, we show below that royalties increased considerably between the two periods, so using the 2000 oil output levels is probably a reasonable proxy for the change.

¹⁰ One last possible concern with using royalties as the right-hand-side variables is that some municipalities may be able to manipulate ANP officials into distorting royalty payments in their favor. The ability to perform this manipulation may vary systematically among municipalities.

III.B Data

AMCs. Over the decades the number of Brazilian municipalities has steadily increased, as many of them have split into two or more – presumably as a consequence of population growth. This fragmentation makes it tricky to analyze panel data on municipalities, as municipalities that exist today did not exist twenty or more years ago. To deal with this problem, Instituto de Pesquisa Econômica Aplicada (IPEA), the official statistical agency of the Brazilian presidency, has created the AMCs. Each AMC contains one municipality (or more) such that the area of each AMC remains relatively stable even when municipality boundaries change.

Our empirical work is conducted at the AMC level. The main reason for this choice is that we wish to test for random assignment of oil. In order to do so, we need to compare outcomes at dates before (most of the) oil was discovered. This requires panel data. Altogether, more than 5500 municipalities that exist today are pooled into 3659 AMCs.¹¹

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 municipal level. Among these there is our key variable, oil output. In these cases we collapse the municipal-level data to the AMC level using a cross-walk from IPEA.¹²

Oil Output. A key variable in our empirical work is an AMC-level measure of the value of oil extracted in its municipalities. Here we give a detailed description of how we constructed this measure. This involves essentially three steps: (i) build a dataset of oil output for each oilfield; (ii) find the geographical position of each oilfield relative to each municipality; (iii) 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. ANP reports for each producing oilfield the reference price used to calculate royalties from oil and gas for every month since August 1998. For the same period, ANP also lists the quantity of oil and gas produced in each oilfield. Using these two datasets, we calculate the value of oil and gas produced each year in each oilfield from 2000 to 2005.¹³

Step (ii) is also easy. In 2000, ANP created Banco de Dados de Exploração e Produção (BDEP), a data base that contains information on exploration and production of oil and natural gas. The

¹¹ IPEA has multiple AMC partitions, for different spans of time (e.g. 1920-2000, 1940-2000, etc.). Because of the progressive fragmentation of municipalities, the longer the time span the coarser the disaggregation (i.e. the fewer the AMCs). The earliest data we use in this paper are for 1970, so we use 1970-2000 AMCs.

¹² The crosswalk we obtained maps municipalities existing in 1997 into 1970-2000 AMCs. One slight complication is that the municipal-level variables we use are for 2000-2005, so there are a few municipalities that did not exist in 1997. In these few cases before applying the crosswalk we first assigned the new municipalities to 1997 municipalities “by hand.”

¹³ The reference price is the maximum between the actual sale price of the oil extracted in a particular field and an imputed sale price (for oil delivered to Petrobras-owned refineries) based on prevailing world-market prices for oil with similar chemical composition. In practice, the reference price is essentially indistinguishable from the market price, so our measure of field-specific oil revenues should be very accurate. For details on the reference price see ANP (2001). The reference prices by month and field are at http://www.anp.gov.br/participacao_gov/precos_referencia.asp, and the quantities extracted are at http://www.anp.gov.br/participacao_gov/prod_petro_gas.asp.

BDEP website provides Geographic Information System (GIS) maps showing the location and status (i.e. producing, in development, etc.) of known oil and gas fields. From this database we selected 273 fields that had passed the stages of development and were already producing in December 2007.¹⁴ We then combined the map of field locations with another GIS map showing the boundaries of Brazilian municipalities from the Instituto Brasileiro de Geografia e Estatística (IBGE).¹⁵ This allowed us to establish the geographical relationship between the various oil fields and the various municipalities. A similar map of political boundaries and oilfields is reproduced in Figure 2, except that in the figure we show AMC boundaries rather than municipal ones.

Step (iii) is somewhat more challenging. While some onshore oilfields lie entirely within the boundaries of a single municipality, in which case it is natural to assign to it the entire output of those oilfields, many onshore oilfields straddle multiple municipalities, so a criterion has to be devised to apportion the oil output among them. Our solution is to simply share equally the oil from a certain field among the municipalities that lie above it.

The problem is a bit more involved in the case of offshore oilfields, which lie entirely outside the (terrestrial) boundaries of any specific municipality. Fortunately for us, the authorities had to solve the same problem. As discussed, Petrobras pays royalties (through the ANP) for oil extraction to municipal governments, and one component of the royalty allocation formula is geographic. Specifically, a certain percentage of the value of the output of each offshore oilfield must be paid to the “municipalities facing the oilfields,” so a mechanism had to be devised to determine for each oilfield which are the “facing” municipalities. The principle that has been followed has been to apportion the royalties based on the fraction of the oilfield that lies within each municipality’s borders’ extension on the continental shelf. The application of this principle, however, is complicated by the fact that there exist two sets of municipality maritime borders: one based on extending the land borders through parallel lines, and one based on perpendicular lines. This complication is finessed by distributing 50% of the royalties (due to facing municipalities) according to one set of borders, and the other 50% according to the other. The resulting percentage allocation is contained 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. We refer again to the Appendix for a more detailed discussion.¹⁶

Besides municipality-level oil output we also create an indicator for having a positive share of at least one oilfield based on the same criteria (i.e. for onshore oil the municipality lies above the

¹⁴ The BDEP map database is at <http://maps.bdep.gov.br/website/maps/viewer.htm>. The date of our download was March 3, 2008.

¹⁵ The IBGE map database is at <http://mapas.ibge.gov.br/divisao/viewer.htm>. For reasons discussed above we use the 1997 boundaries.

¹⁶ The Percentuais Médios de Confrontação we have used are those for February 2008, and can be found at http://www.anp.gov.br/doc/participacoes_governamentais/calculo/2008/fevereiro/Confrontacao_fev08.pdf. We have done some unsystematic checks to make sure that these shares do indeed reflect the stated geographical principles. In most cases, they seemed fairly consistent. However, there were a few smaller oilfields for which the allocation of percentages did not seem consistent with the stated criteria. We have been unable to establish what alternative criteria had been used in these cases. We have also checked the shares published for other months (e.g. March 2009) and, as expected, found little variation over time.

oilfield, and for offshore oil it has a positive Confrontação share. There are 124 municipalities with a stake in at least one (onshore or offshore) oilfield. Using the BDEP data we also determine when oil was first discovered in each oil-endowed municipality, and whether it has onshore or offshore oilfields (or both).

GDP. AMC-level GDP numbers, both aggregate and broken down into some sub-aggregates, are reported by IPEA (<http://www.ipeadata.gov.br/>). The construction of these GDP numbers 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).

Municipal budgetary variables. Detailed data on AMC-level revenues and expenditures, broken down by nature (current, capital, etc.) and function (education, health, etc.) are reported by IPEA. The amount of oil royalties received by each municipality in each year is readily available from ANP.

Other socio-economic outcomes. We investigate the effect of oil-induced municipal revenues on a battery of socio-economic outcomes. We give details of definitions and sources on an outcome-by-outcome basis when we present the results. Most of the data come from IPEA or from our own calculations using the Brazilian census, but in some cases we use more specialized sources.

Population. In estimating (1) and (2) we typically normalized both left- and right-side variables by population. Up to the year 2000 population data comes from IPEA, based on the Brazilian Censuses. To calculate population for years after 2000 we inflated the 2000 population from the Census by IPEA's estimate of the percentage change in population residing in each AMC on 1 July of each year.¹⁷

Geographic controls. Most of the geographic controls (latitude, longitude, distance from state and federal capital, and area) come from IPEA. In addition to calculating oil-related variables for each municipality, we use the political map of Brazil to create an indicator for whether a municipality is adjacent to the coast (an AMC is coastal if at least one of its municipalities is coastal).

Deflation. Many of the economic variables we used were obtained directly in R\$2000. Other variables were denominated in nominal R\$, and we converted them to R\$2000 using a CPI index from IPEA.¹⁸ A list of all the variables we use and their sources is in the Appendix.

Summary statistics and subsamples. Table 1 present some summary statistics from various subsamples in our dataset. The first column reports figures calculated from the subsample formed by the 3556 AMCs that do not share any oilfield. The second column is based on the subsample of 59 oil-endowed AMCs where the oilfield was discovered after 1970. The reason for highlighting AMCs with oilfields discovered after 1970 is that, as we will see, we can rule out that municipalities where oil was found after 1970 were systematically different in key

¹⁷ Similarly, there is one instance where we need population data for 1992, and, again, we used a similar interpolation from the 1991 Census.

¹⁸ The index we used is Índice Nacional de Preços ao Consumidor (INPC).

outcome variables (after controlling for geography) from non-oil AMCs before discovery. In the third column we show data from all 103 oil-abundant AMCs. In the fourth and fifth column we report data from the subsample of 31 AMCs that only have offshore oil, and the 63 that only have onshore oil, respectively. Loosely speaking, the municipalities in the “No oil” column can be thought of as our “control group,” while the municipalities in the four subsequent columns represent alternative “treatment groups” that we use throughout the paper.

There are clearly sizable differences in average GDP per capita, municipal revenues, and population between oil-rich and oil-poor AMCs, the oil-rich ones generally being richer and larger, and enjoying greater revenues (except for onshore ones, which get less revenues).¹⁹ These differences, however, can clearly not be treated as 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. More importantly, oil-rich AMCs are disproportionately coastal (the offshore ones by construction, but the onshore ones are also much more likely to be on the coast than the no-oil ones). There are also substantive differences in distances from federal and state capitals. To identify the causal effect of oil it is therefore essential to control for these geographic characteristics. We also control for state fixed effects.

The table also reports some statistics from the distribution of our constructed measure of oil output per capita, these being trivially 0 for the no-oil subsample.²⁰ 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 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, depending on the subsample, oil accounts for between 15 and 20% of GDP in 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 the fact that our identification of the effects of oil comes as much from within oil-rich variation as from between the no-oil and the oil-rich samples (hence the sense in which our use of the words “control” and “treatment” above should be taken very loosely.).

III.C Identification

We begin the discussion of our identification assumptions by arguing that Q_{mt} is uncorrelated with e_{mt} in specification (1). The first step is to show that key outcomes of interest 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, particularly with regards to their positions relative to the coast and their distance from federal and state capitals. This means that

assigned *conditional on geographic covariates* (state fixed effects, longitude, latitude, distance to federal capital, distance to state capital, and coastal dummies). In other words, once we compare oil and non-oil AMCs with similar geographic characteristics, oil-abundance status is random.

The main test for the validity of the conditional random-assignment assumption is reported in Table 2, which shows results from a panel regression of the following model

$$Y_{mt} = \delta_t + \eta_t Q_{m,2000} + \theta_t X_m + w_{mt}, \quad (3)$$

Where Y_{mt} is log-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. In the first column the sample is constituted by the “no oil” and the “post-1970 oil” subsamples, i.e. the AMCs in the first two columns of Table 1. The time coverage is given by various dates from 1970 to 2005. In the intervening years before 2000 we include all years for which per-capita GDP at the AMC level is available. After 2000 we have annual data and pick as “representative” dates 2002 and 2005, with the significance of 2002 still to be further explained below. Crucially, the coefficient on oil output in 2000 is allowed to vary over time.

The main reason why we focus 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.²¹ In addition, as we mention shortly, we performed falsification test on other outcomes as well, 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.

It is quite clear that sizable systematic effects from oil do not appear until well into the 1990s, and indeed we must wait for the 2000s to observe a clear relation between oil and GDP. Since the oil AMCs are only those where oil was discovered after 1970, this is strongly indicative that, conditional on our covariates, oil abundant and oil-poor AMCs were similar before the oil discoveries. As robustness checks, the remaining columns include all oil AMCs, and break down oil AMCs into onshore and offshore categories. The conclusion that oil-AMCs differ from non-oil ones only after the period of oil discovery (conditional on geography) seems extremely robust.²²

Further support for our claims of quasi-random assignment (conditional on the above mentioned covariates) is provided in Appendix Table A1, where we repeat the specification in (3) (for $t=1970$ only in order to save space) for other dependent variables for which we have data from 1970 and on which we focus below: housing quality and education (again, we discuss these variables in more detail below). Once again, conditional on geography none of the outcomes in

²¹ In going from 1970-2000 to 1960-2000 AMCs we would lose about a third of the sample size. Going back to the 1940-2000 AMCs would cause a further one-third loss.

²² Behind the gradual increase in coefficients over time there are two factors. First, the distribution over time of oil discoveries post-1970 in our AMCs is fairly uniform (see Figure 1), so in the earlier years only a fraction of the “oil AMCs” is producing oil. Second, even for the early starters, there are inevitable lags between the time of discovery and the time where the oilfield is being exploited at its full capacity.

1970 is correlated with oil output in 2000. (The only exception is electric lighting in 1970, for which the correlation with oil output in 2000 is negative and marginally significant).

Hence, conditional random assignment seems to hold. This in itself goes a long way in providing support for the identification of model (1). However, in our empirical analysis, we also use variation in the quantity of oil output among oil AMCs, and not only between oil- and non-oil ones. In principle, then, one could 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.

Formally, one could think of the error term e_{mt} in equation (1) as being the sum of a time-invariant AMC characteristic ω_m and a time varying shock ζ_{mt} . Our results from Tables 2 and A1 can be interpreted as showing that Q_{mt} and ω_m are uncorrelated. But we still have to deal with the possibility that Q_{mt} is correlated with ζ_{mt} .

We argue that this is implausible. Oilfield operations in Brazil over the sample period were carried out by a global hydrocarbon giant that has full access to global factor and product markets. Neither its highly specialized equipment, nor its equally-specialized labor force could realistically be 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 will be influenced by temporary local conditions in deciding how much oil to extract from a given oilfield, and even less that it will be swayed by local economic outcomes in its prospecting plans. This is likely true everywhere, but particularly so for the case of offshore oilfields.²³ For this reason, for each set of results we present separate regressions where we include only non-oil AMCs and AMCs where the oil is offshore (i.e., exclude AMCs with onshore oilfields). We feel that these results should be particularly sharply identified.

We now briefly turn to identification of model (2). As mentioned, here the key assumption is that our instrument, oil output Q_m , affects outcomes of interest at the municipality level (mainly spending by the local government, provision of public goods, and household income) only through the revenues R_m it generates for the municipal budget (the bulk of which is represented by oil royalties). Our main defense of this identifying assumption is given by an anticipation of the results from estimating (1). As we show below, the effect of oil output on AMC non-oil GDP is essentially zero. For offshore oil, we also find no effects on the composition of non-oil GDP (onshore oil has a minor effect). This strongly suggests that oil has little *market* effects on economic activity at the AMC level (and for offshore oil the effect is nil). Given our previous

²³ The same applies to the possible concern that demonstrations or lobbying by local groups or tribes and/or by environmental organizations may affect the cost of drilling or complementary investments (e.g. building oil pipelines and related infrastructure). While this may occasionally be an issue onshore, it is exceedingly unlikely to arise in the case of offshore activities. Another possible concern is that municipalities compete to lobby and/or bribe Petrobras to drill near them. This is exceedingly unlikely. First, municipalities are tiny (see below) and it is nearly unconceivable that they will have the political heft and financial resources to sway the decisions of Petrobras, one of the World's biggest companies. Second, unlike many Brazilian institutions, Petrobras actually has a strong record and reputation for integrity. This record has been explicitly recognized by international NGOs operating in the natural-resource area, e.g. Transparency International (2008)

discussion, this is not surprising: Petrobras operations tend to be fairly isolated from the local economy, especially offshore. This suggests that any effect from oil likely arises from the revenues it brings to the municipal government. Since this reasoning appears to be particularly robust for offshore oil, the most cleanly identified results are those pertaining to the subsample where the treatment group is composed of municipalities that derive their oil only from offshore fields.²⁴

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 A2 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 also 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 there is a concern that nearby AMCs might be indirectly affected by oil. But while we do not report these estimates in the paper they were generally very similar (both in magnitudes and precision) to the results that we do report.

IV. Results

IV.A Oil Abundance and GDP

In Table 3 we look at the effects of oil abundance on the productive side of the local economy. The specification is (1), with AMC GDP on the left-hand side, and AMC oil output, interacted with a year dummy to allow for time-varying coefficients, as the main right-hand-side variable. A full set of interactions between each of our usual geographic variables and year dummies are also included.

We begin our sample period in 2000 - or in the first year for which we have reliable data.²⁵ This brings us to the repeatedly promised discussion of the significance of 2002. It turns out that IPEA GDP data 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 where 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

²⁴ Nevertheless, we include the onshore oil AMCs as a robustness check.

²⁵ The first year for which we have both GDP and oil output numbers is 1999, but since in the rest of the paper data limitations force us to focus on the year 2000, we decided to begin with that year here as well.

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.²⁶

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, 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 fairly stable over time and hovers around 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. Now as already mentioned at the *national* level the share of oil GDP in gross oil output is also fairly stable and around 0.4.²⁷ Under fairly standard assumptions average and marginal shares of GDP in gross output are the same, so to the extent that the 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.²⁸

We also have AMC-level GDP numbers 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 it suggests 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 because in this case the no-spillover conclusion does not rest on an (admittedly uncertain) estimate of the share of oil GDP in gross oil output, as is the case for

²⁶ This mismeasurement does not invalidate the falsification exercise we conducted in the previous section. The point of that exercise was to show that differences among municipalities were not systematically related to oil abundance before (and for several years after) the 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.

²⁷ Here is the annual time series of the ratio of GDP to gross output in the oil sector in the national accounts between 2000 and 2005: 0.49, 0.40, 0.35, 0.36, 0.35, 0.42. Source: ftp://ftp.ibge.gov.br/Contas_Nacionais/Sistema_de_Contas_Nacionais/Referencia_2000/2004_2005_novembro2007/Tabelas_de_Recursos_e_Usos/

²⁸ Begin with the identity

$$\text{GDP} = \text{NON-OIL GDP} + \text{OIL GDP}$$

From the results in 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 we also infer

$$d(\text{OIL GDP})/d(\text{Gross oil output}) \approx 0.4$$

which then implies

$$d(\text{NON-OIL GDP})/d(\text{Gross oil output}) \approx 0.$$

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 email exchanges with IBGE staff].

aggregate GDP or industrial GDP. In columns 4 and 5 we show that these results are robust when AMC's where oil was discovered before 1970 are included in the analysis.

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 in Section III.C, 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 fulfil this demand (and we have seen above 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.

Support for this hypothesis is found in the last four columns of Table 3. In offshore-only oil AMC's we find the usual one-for-one increase in industrial GDP with oil GDP (i.e. roughly 0.4 coefficient on gross oil output), and no change in non-industrial GDP. This is consistent with offshore oil having no market impact on the local economy. On the other hand, in onshore-only oil AMC's the effect of oil on industrial value added is less than one-for-one, as the coefficient on gross oil output falls to approximately 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. At the same time, however, we find a symmetric positive effect on non-industrial output: the coefficient of about 0.10 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.

VI.B Oil Abundance and the Local Government Budget - Revenues

Having failed to find significant market effects of oil abundance, we turn to possible effects flowing through the government budget. We begin by investigating the effect of oil on the revenue side of Brazilian municipalities' budgets. Table 4 confirms that oil riches flow in part into local-government budgets. The specification is still the same as (1), only with various measures of municipal revenues as the outcome variable. To save on space we focus on a single cross-section, and since for several other variables we analyze below we only have data for 2000 we chose 2000 here as well.

Column 1 shows significant increases in revenues received by local governments from oil in 2000. One Real of gross oil output increases total local-government revenues by almost 3 cents. This is true for the subset of oilfields discovered after 1970 as well as for the full sample, and for the subsample including only offshore-only oil AMC's. The effect is muted in the sample including onshore-only oil AMC's, where one Real of oil produced leads to just a 1.5 cent increase in government revenues. One shortcoming of the results in column 1 is that there are many missing values for municipality revenue in 2000. In column 2 we use 2001 values to impute the missing observations for 2000, and the anomaly for the offshore-only subsample disappears: one Real of oil output increases revenues by about 3 cents in all subsamples.

In column 3 of Table 4 we investigate the sources of the increase in revenues. In particular, we look at the effect of oil production on royalty income. The increase in royalty income accounts for almost two-thirds of the overall increase in municipality income due to oil production.²⁹ Evidently, oil production generates sources of income for municipalities over and above the royalties they receive from Petrobras.³⁰

One very important implication of Table 4, and, in particular, of the fact that oil municipalities have larger revenues, 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 take advantage of royalty income to cut local taxes.³¹

As mentioned in Section III.A, in the remainder of the paper we estimate specification (2), in which municipal revenue is the right-hand-side variable, both in levels and first differences (whenever we can). The earliest date for which we have good coverage of both municipal revenues and outcomes of interest is 1991, so many of the first-differenced regressions will be for 1991-2000. In order to establish the validity of 1991 as a base year for the differenced regressions column 4 of Table 4, shows the coefficient of a regression of 1991 revenues on oil output in 2000. The results show that municipal revenues in 1991 were much more weakly related to oil-abundance than in 2000. Some of the reasons for this are already familiar: many of the oilfields were discovered late in the century, and development lags further delay the impact of the discoveries on municipality budgets. Furthermore, the local-government “take” in local oil output increased dramatically after a wide ranging reform enacted in 1998 that, among other things, radically increased the “reference price” used to evaluate output for the purposes of computing royalties (basically making it the market price), and increased the typical overall tax – to be distributed in the form of royalties – from 5% to 10%.

The validity of 1991 as a baseline is reinforced by column 5, where we regress the change in revenues between 1991 and 2000 on 2000 oil output. The coefficient is essentially the same as in the level regression for 2000. Note that column 2 is essentially the first stage for our subsequent IV regressions in levels and column 5 for our subsequent IV regressions in differences.³²

VI.C Oil Abundance and the Local Government Budget - Spending

²⁹ 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.

³⁰ The bulk of these additional sources are contributions from “participacao especial.” This is an ad hoc tax levied by the Federal government on each oilfield, and depends on a variety of field characteristics. The overall value of the “participacao especial” is similar to the overall value of royalties. For example, in 2004 royalties amounted to R\$5735, while the participacao was R\$5995. 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 “participacao” [de Oliveira Cruz and Ribeiro (2008)].

³¹ Taxes, however, account on average for less than 5% of municipal revenues, so the scope for tax cuts would have been fairly limited anyway.

³² In columns 4 and 5 we predict missing 1991 data on municipal revenues using 1992 data, and in column 5 we continue predicting 2000 municipal revenues for municipalities that did not report revenues that year using 2001 data.

So oil brings money to the local government. What does the local government do with it? We begin in Table 5 with what the government says it does, i.e. we look at 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. The most important items are Education and Culture, on which municipalities report spending about 27 cents of the average 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, about 7 cents on the Real.³³ 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.

The OLS results describe the allocation of the *average* Real of revenues, independent of its source. In order to identify the utilization of oil-related revenues, in Panels B and C we turn to our empirical model (2), 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. We emphasize results for the sample in which the “treatment group” is composed of the AMCs with only offshore oil (Panel C) because – as discussed above – the case for the validity of this procedure is particularly compelling in this case. But results using all oil AMCs as the treatment group give mostly identical results (Panel B).

Our IV results show that the largest beneficiary of the increase in government revenues from oil is Housing and Urban Development, with about a fifth of the marginal “oil Real.” Education falls into second place, with close to 15 cents, Transportation receives about 13 cents, Health continues to receive about 10 cents, and Transfers to households 5 cents. Hence, we find significant differences between the allocations of the marginal and average oil Real. The overall effect on spending is also different, as it drops (for offshore oil) to about 83 cent per Real, indicating that the saving rate out of oil-related revenue is higher than for general revenue.

The bottom panel of Table 5 reports the results in ten-year differences. Results using differenced outcomes are nearly identical to those using levels.

That oil output increases the size of municipal-government budgets is confirmed by simple summary statistics on the size of administration and personnel costs of municipal governments. We computed administration costs in 2000 (as usual using 2001 values where 2000 information is missing), for the top 25 AMCs ranked by oil output per capita. More specifically, we measured municipal administration costs after controlling for our usual geographic covariates (this is just the residual from a regression of municipal administration costs per capita on state dummies,

³³ Education spending by municipal governments is mostly in the area of 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 (roads etc.) and transportation (buses, etc.). 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 of the total.

coastal dummy, etc.). We found that of the 25 top oil AMCs in 2000, 10 were in the top decile and an additional 5 were in the second decile of the administrative-cost distribution. Moreover, 4 (out of 25) were in the top 1 percent! (Results using unadjusted administrative costs are slightly less dramatic, but still show a large over-representation of oil AMCs among the biggest spenders on administration.

IV.D. 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 a variety of measures of real outcomes in all of these areas, to see to what extent the increased reported spending leads to material improvements in living standards.

Table 6 looks at a variety of housing, urban service and infrastructure outcomes: overall value of the residential housing stock, a proxy for housing quantity (rooms per person) and measures of quality of housing and infrastructure: fraction of population living in sub-standard housing, electricity, connection to water and sewage networks, garbage collection, and extent of roads under municipal jurisdiction. The number of rooms at home is a variable we computed ourselves from the Brazilian census. The length of roads under municipal supervision is from administrative records from the Ministry of Transport, and we aggregate at the AMC level as described above. The remaining variables are directly available from IPEA at the AMC level.³⁴

While the OLS results tend to show a positive association between government revenue and housing and transportation outcomes, the IV results are almost uniformly indistinguishable from 0. The exceptions are percent of population *not* living in sub-standard housing, and the percentages linked to water and sewage, but in all three cases the coefficient has the “wrong” sign: oil-related government income leads to a worsening of housing quality and infrastructure.³⁵ It is also worth noting that while the IV estimates using offshore oil are (not surprisingly) less precisely estimated than the OLS, their confidence intervals do not overlap in 7 of the 18 regressions, and in each of these 7 cases the IV estimate is smaller. Panel B shows that these results are robust to first differencing.

In Table 7 we look at real outcomes in education, health, and certain transfers received by households. The education and health data come from administrative records from the Ministries of Education and Health, respectively. Transfers received by households are our own computations from the census, and mainly include transfers for the alleviation of poverty, unemployment benefits, and incentives for schooling.

In the specification in levels, variables associated with the provision of education services (columns 1-4 of top panel) do indeed increase significantly with oil-generated revenues. For example, the coefficients in the regressions for the offshore oil sample imply that a million Reales of extra revenue leads to the hiring of 2 new teachers contemporaneously and almost 5

³⁴ Residential-capital values are also based on Census data on housing characteristics and locations, that are then converted into Reals through an hedonic model.

³⁵ The fraction of population with garbage collection increases with oil revenues in the cross section, though not in the time-differenced specification.

with a 5-year lag. It also leads to the eventual construction of two new classrooms. Unfortunately, when we take first-differencing these results disappear (except, perhaps for teachers), so there seems to be no robust effect from oil-related revenues on any education-related variable (bottom panel).³⁶

In column 5 and 6 we look at municipal health infrastructure. There is a positive and statistically significant effect on the number of municipal health establishments with impatient care, and this holds both in levels and in first differences. The coefficients can be interpreted as saying that a R\$100M increase in oil-related revenues leads to the construction of one extra impatient care facility. While R\$100M 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. In the first-difference specification there is also a positive effect on facilities without impatient care, but this vanishes in the level regressions.

Finally, in column 7 we look at the effect of oil-related revenues on poverty- and unemployment-related social transfers from the population census. There is no indication whatsoever that these welfare-like payments increase with oil revenues. (We don't have base-year welfare-income numbers, so this result is only available in levels).³⁷

Taken together, the results from Tables 5, 6 and 7 are fairly 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 almost universally find either no effect or 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 real outcomes to assess their magnitude vis-à-vis what one should plausibly expect.

We should acknowledge some important limitations to the conclusions we have just drawn. First, of course, we might be 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. While we can't rule out either of these

³⁶ Note that 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 that are comparable to those in our later year of data. But even ignoring the first-differenced results, and taking the level coefficients at face value, the implied outcomes would probably be disappointing relative to the reported spending. According to Table 5 a 1 million increase in municipal revenues leads to about a R\$130,000 increase in reported spending on education. We don't have yet data on primary teacher salaries, but assuming (generously) that they are roughly in the order of twice per-capita GDP would suggest that after hiring two extra teachers the government would still be left with more than R\$100,000 to spend on infrastructure. Let's say (again, generously) that the two classrooms we observe in 2005 are the results of these investments. That's R\$50,000 per classroom – which is enormous given that R\$50,000 will buy a large independent family house in most locations in Brazil. Of course there are other items on which the education budget may be spent, but teachers and classrooms should by far be the largest.

³⁷ In Table 7 the confidence intervals of the OLS and IV estimates for offshore oil do not overlap in 5 cases (out of 13) and in each of these 5 cases the IV estimate is lower.

possibilities, we don't think they can be seriously entertained. Our failure to find convincing evidence of productive use of oil revenues is spread over a very wide range of different outcomes and it is implausible that we have systematically oversampled from the subset of outcomes with no positive effects. In any event, 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 measure of living standards.

A second 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. The comparison between columns 1-2 and 3-4 of Table 7 lends some credence to this view, where 2005 outcomes seem to respond more strongly to 2000 revenues than 2000 outcomes. In assessing the importance of this concern a number of considerations are relevant. First, not all of our outcome variables are plausibly subject to "time to build." In particular, there is no reason to great delay needed to hire teachers or mail transfer payments to households. 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. And even the positive education outcomes in 2005 are suggestive of very inefficient spending, as argued above. Third, and perhaps most crucially, it is important to keep in mind that municipal revenues from oil are very persistent over time. AMCs with relatively large revenues in 2000 tend to have had relatively large revenue for many years before, so our coefficients should not necessarily be interpreted as measuring the impact effect of 2000 revenues. Rather, they should be thought of as capturing the effect of "permanent income" from oil. In this respect, there is perhaps no strong reason to worry about time to build.

Another potential issue is crowding out of state and federal spending. The items in the state and federal functional spending budget are essentially the same as in the municipal budget, and the "division of labor" between different levels of government is not always clearly defined. It is therefore conceivable that state and federal bodies will withdraw funding in areas where they are aware of increased spending by municipal governments. For some of our outcomes we can rule out this crowding out. In particular, our road, teacher, classroom, and health establishment data all strictly relate to provision by municipal governments, so they are net of any state and federal contributions. Furthermore, for these same variables we also have data on state and federal provision (e.g. teachers in schools run by the states/federal government), and we found no effect of oil-revenues on these. Nevertheless, we acknowledge that the possibility of crowding out remains, at least in theory, for the housing variables and social transfers.³⁸

IV.E. Oil Abundance and Household Income

So far we have looked at the effects of oil abundance on the productive structure of the non-oil economy, on the government budget, and on the supply of a number of public services. The latter exercise generates some questions as to the extent to which the reported spending increases

³⁸ As already mentioned our measure of social transfers includes subsidies to the poor, unemployment insurance, and a program to incentivize poor families to send their children to school, called Bolsa Escola. The last of these was still a municipal-level program in 2000 (it went federal in 2001) so it is not subject to the crowding out problem. But the other programs conceivably might be. The other famous federal transfer program, "Bolsa Familia," was launched in 2003, so it is not an issue for us.

actually materialize in services to the population. Nevertheless, it is still quite possible 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 this section we study the effect of oil-induced government revenue on a summary measure of living standards: namely household income, which we compute from the Brazilian census.

Table 8 reports estimates of the effects of oil-generated municipal revenue on household income per capita. The specification is always as in (2). Column 1 shows that there is no effect on average household income whatever in the IV regressions. This result is robust to choice of sub-sample and to specification in levels or first differences. 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.³⁹

In the next five columns we look at the effect of oil on household income by quintile. This gives a somewhat more nuanced view than looking at the average effect. In particular, we do find fairly robust evidence that household income increases in the bottom quintile of the income distribution. In the full sample the benefits extend to the second quintile (and, in the first-difference specification, to the third quintile, but this is not the case in the levels regression). Nevertheless, it is important to notice that these increases are extremely small: for every per-capita dollar of increased revenue (and spending), the increase in income is in the order of ten cents. Clearly, much of the reported public spending is either "missing," or it is offset by reduced market income.

In the last column of Table 8 we look at the effect of oil on poverty rates. In the full sample, using the specification in levels, we find evidence that oil may have reduced poverty, but this result vanishes in first differences and, more importantly, in our better-identified regressions using offshore oil only as the treatment. Taken together, these estimates suggest that poverty reductions due to oil were modest, at best.⁴⁰

³⁹ It is useful to relate this finding to our previous results. We already know that oil affects local GDP only through its own direct contribution, and has no spillovers on non-oil GDP. Still, household income could increase if there were local under-utilized factors of production that became mobilized in oil production and participated in the distribution of oil GDP. If this was the case, our instrument would be invalid – but the IV coefficient would be positive, because it would reflect the positive reduced-form relation between oil GDP and incomes. The fact that this is not so therefore reinforces our previous contention that oil production does not use local factors. (To double check this we estimated directly the reduced-form relation finding insignificant coefficients.) The fact that non-oil GDP does not change in response to oil production also rules out many channels through which the fiscal windfall could ultimately increase household income. If (real) demand for goods and services by the government increases, and is satisfied at least in part by locally-owned factors of production, particularly labor, household income would likely increase, but so would GDP (if the increased demand by the government completely crowds out production targeted to the market then neither household income nor GDP will increase). This leaves increased transfers to the population as the only mechanism through which oil revenues could increase household income that is not implicitly ruled out by our previous finding that non-oil GDP is unaffected. To some extent, therefore, the result in Column 1 may be interpreted as a robustness check on our GDP results.

⁴⁰ Even the significant coefficient for the level regression in the full sample implies minuscule effects.. The coefficient estimate implies that municipal revenues due to oil need to increase by 100 Reales *per capita* to see a reduction in poverty of one percentage point. This is more than the average revenue from oil royalties in the oil abundant sample (see Table 1).

Given the foregoing considerations, the most plausible reading of Table 8 is that oil production does very little for household income. Individual citizens seem to be almost isolated from the oil windfall, at least as measured by their income. Results from the reduced-form regression of household income on oil abundance (available on request) tell a very similar story. Hence, the evidence points to a large amount of “missing” money.⁴¹

V. Where is the missing money going?

Where is the missing money going? It is difficult to resist the suspicion that much of it is diverted to private use by government officers. One outcome variable that may speak to this issue, albeit very indirectly, is the relative size of houses enjoyed by municipal employees, which can be easily identified in Census data from the already-used quality of housing variables cross-checked against respondent’s sector of employment. Table 9 reports the results. It seems clear that oil-related revenue increases the quality of housing for municipal workers – but, as we already know, not for everyone else. Whatever the mechanism, municipal workers seem to be able to obtain for themselves relatively more spacious accommodations in oil-rich municipalities.

To shed more light on this question, we have constructed new data on the frequency with which different municipalities are cited in the news media in connection with corruption. Specifically, for each municipality in Brazil we have performed a search of all the news items in the archives of Brazil’s News Agency (Agência Brasil). Each time, the search keywords were the name of the municipality, the (Brazilian) Portuguese word for mayor (*prefeito*), and the Portuguese word for fraud (*fraude*).⁴² We then constructed a municipality-level dummy variable that took the value of 1 if the search delivered at least one hit, and 0 otherwise.⁴³ When we regressed this variable on per capita oil output in 2000, with our usual set of controls, we found insignificant estimates. However, when we used the absolute level of oil output in 2000 our estimates were statistically significant for all AMCs, and also separately for offshore and onshore oil separately (see Table 10). We conclude that high oil output *per capita* does not make a municipality more likely to be cited for fraud in the news, but more oil output *per mayor* does.

⁴¹ A somewhat worrisome aspect of our data is that there is a seeming discrepancy with the national accounts. In our data household income is in the order of 60% of GDP. In the national accounts national income, which should be roughly similar to household income, is in the order of 95% of GDP. The discrepancy does not seem to arise from GDP: when we aggregate our AMC-level GDP numbers we recover a figure very close to the national accounts, so the discrepancy seems to be between household income and national income. To check our household income data, which we obtained from IPEA, we constructed an independent measure of household income from the 2000 census. The resulting IPEA-based and census-based aggregate household incomes are very close. The bulk of household income in the census data is due to wages, and when we aggregate data on wages we get something very close to the figures for remuneration in the national accounts. This suggests that most of the income we are missing in our household income data is capital income.

⁴² In an English-speaking country the best word to use for this search would have been the more specific “embezzlement,” rather than “fraud.” In Portuguese, however, the phrase for embezzlement is “desvio de fundos,” which is often shortened into “desvio” when the context allows no ambiguity. Unfortunately “desvio” by itself also means “detour.” Hence, a search using “desvio de fundos” would have led us to miss many articles dealing with embezzlement (when simply desvio had been used) and one using desvio would have led to many false positives.

⁴³ We used a simple 0-1 indicator instead of the actual number of hits because many municipality names are identical to state names or have other meanings in Portuguese. This means that there are some huge outliers in the distribution of hits. To prevent these outliers from dominating the results we have chosen the dummy-variable approach. 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.

To further examine whether oil-rich municipalities are likely to be involved in illegal activities requiring federal involvement, we examined a list of Federal Police operations for 2003-2008.⁴⁴ For each of the 713 federal police operations we searched the (very short) abstract for mention of each of the municipality names. We then repeated the analysis above, where this time the dependent variable is an indicator for a mention of (at least) one of the municipalities in the AMC in a federal police operation. The results (not reported) are very similar to the news articles: per capita oil output in 2000 does not increase the odds of a federal police operation, but the absolute level of oil output in 2000 did increase those odds. As before this result holds for all AMCs and also when we use offshore and onshore oil AMCs separately. Finally, we restricted ourselves to federal police operations that involved mayors, and again the results (in Table 10) show that the absolute (though not per capita) level of oil output increased the odds of having at least one such operation in the AMC – for offshore oil, for onshore oil, and for both types together.

As an informal check on the nature of the allegations behind news stories and police raids, we made a broader search of the news media on the 10 most oil-rich municipalities (as measured by oil output in 2000). Some of the stories we picked up are summarized in Table 11, and they appear to support the view that corruption in oil-rich municipalities is a serious concern. One of the stories we found, related to “Operação Telhado De Vidro” (“Operation Glass Ceiling”) in the municipality of Campos de 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 up to R\$250 million (which were worth at the time about 140 million dollars). This is a very large sum, much larger than the average municipality’s budget, though still only a fraction of the royalties received by Campos de Goytacazes over the past decade. Of course, some of the allegations may turn out to be unfounded. But our results about the very limited improvements in standards of living and the recent nature of some of the cases suggest that these stories may only be “the tip of the iceberg”, and that many other cases may have as yet gone undetected.

There is also some circumstantial evidence that the general public is ill informed about the magnitude of the oil-related fiscal windfall. In Campos de Goytacazes, the largest oil producer of Brazil, the local news bulletin “Petróleo, Royalties & Região” reports results from a survey aimed at assessing the information existing in the population concerning the amount of royalties flowing to the government (issue of September 2003). In the results 58% of local residents ignored the existence of oil royalties; of those who are aware of the royalties, 54% did not know how they are used.⁴⁵ This is remarkable as Campos has received R\$ billions in oil revenues over the last few years, and the surveys were conducted well after royalty money started to flow. Also, the level of education in Campos de Goytacazes is higher than in most of Brazil (an average over 6 years in 2000, compared with about 4.2 for the average AMC), and in the survey better educated people displayed relatively more knowledge regarding royalties. Finally, being in

⁴⁴ These are official administrative data available from the Federal Police’s web page: http://www.dpf.gov.br/DCS/Resumo_OP_200X.html, where X=5, 6, 7, 8 or 3-2004 (for example, for 2008 the link is here: http://www.dpf.gov.br/DCS/Resumo_OP_2008.html; 2003 and 2004 are reported together).

⁴⁵ These results were reported in the bulletin’s first issue. In subsequent years awareness of the royalties seems to increase substantially, perhaps a sign of success for the Bulletin’s itself, whose mission is indeed to diffuse knowledge about the fiscal consequences of the royalties. For a list of issues see: <http://www.royaltiesdopetroleo.ucam-campos.br/index.php?cod=1>.

the state of Rio De Janerio, Campos is probably better covered by the media than most of Brazil. Thus, we expect that public ignorance regarding the extent of oil revenues may be even higher in more rural and remote parts of Brazil.

Before concluding this section it may be appropriate to discuss how the “missing money” result and our hypothesis that most of it is appropriated by top local officials relates to national-account identities that, after all, need to hold. There can be no question that GNP in oil-rich communities increases. While local factors of production do not seem to share in the distribution of oil GDP, the government receives a significant fiscal windfall. Hence, $\Delta \text{GNP} \approx \text{Fiscal Oil Revenues}$. On the expenditure side, the change in net exports (value of the oil minus the royalties, which are kept locally) cancels out with the value of net factor payments from abroad (since all inputs are owned outside the municipalities, all exports net of the royalties are negative NFPs), leaving the change in GNP to be balanced by a combination of increased consumption (C), government spending (G), and investment (I) (the expenditure-side formula for GNP is $C+I+G+NX+NFP$). As we have seen, municipal governments essentially balance this by reporting $\Delta G \approx \text{Fiscal Oil Revenues}$. Since our results indicate that the ΔG does not quite materialize we think that what really happens is $\Delta C \approx \text{Fiscal Oil}$, where ΔC is extremely concentrated among the very top municipal officials, and therefore goes undetected in the household income numbers. Alternatively top government officials save the embezzled funds in bank accounts outside the municipalities, which would “show up” on the expenditure side as a further increase in net exports.

VI. Theoretical Implications

The theoretical literature has proposed both *market-based* and *political-economy* mechanisms through which natural-resource abundance may affect economic outcomes and the welfare of the populace. As mentioned in the Introduction, there are several theoretical channels that are only applicable at the national level, and a limitation of our analysis is that we cannot detect the operation of such aggregate forces. Still, this leaves a significant number of theoretical mechanisms that could potentially operate in the local context that is our focus. Here we review these mechanisms and assess them in light of our results. For completeness, we also briefly mention (in footnotes) some of the non-observable (to us) theories.

Market-based mechanisms are broadly referred to as instances of Dutch disease. A classic source of Dutch disease arises from a **wealth effect** (sometimes called a “spending effect” in the Dutch-disease literature). The mineral riches trigger a surge in demand for consumption goods. To satisfy the extra demand for non-tradables resources are reallocated from the (non-resource) tradable to the non-tradable sector, while the extra demand for tradables is accommodated through increased imports (financed by the resource exports and, in some cases, external debt). A further consequence of the wealth effect is that it potentially depresses overall labor supply, leading to a combination of higher wages and lower overall non-resource GDP. This mechanism is clearly potentially at work in the local Brazilian setting. In particular a wealth effect could arise from the royalties paid by Petrobras to the local government, to the extent that these royalties are rebated (directly or in the form of goods and services) to the local population. Evidence for such a mechanism would be represented by a decline of GDP in the (non-resource)

tradable sectors and an increase in the GDP of the non-tradable sector. Evidence of a wealth effect would also be found in a decline in overall labor supply and an increase in local wages.

Another market mechanism that may be present in our setting is due to the direct impact on **relative demand by oil firms and oil workers** for goods and services. In particular, there could be an increase in the relative demand for personal services to the oilfield workers, and of business services to the oilfield operations. This would lead us to expect oil to shift the composition of non-oil GDP away from industry and towards services.⁴⁶ This particular type of Dutch disease is not often studied in the theoretical literature, because it is not likely to be important at the national level. But it could be quite relevant at the level of the local economy.

Both the wealth effect mechanism and the direct relative demand shift from oil firms and workers imply a change in the composition of non-oil GDP, and our data do not allow us to establish whether the change is between tradables and non-tradables – as predicted by the wealth effect -- or between services to oil workers and firms and other activities – as predicted by the oil-operation as local demand-shifters effect. However, our data allow us to distinguish between the effect of onshore oilfield operations and offshore ones. This is extremely useful because both offshore and onshore operations pay royalties to the local government – so they are both potentially liable to create a wealth effect. Instead, only onshore operations are likely to significantly directly affect the composition of demand for local goods and services. As we have discussed, offshore operations are miles off the coast and are largely disconnected from local factor and goods markets. We can therefore conclude that changes in the composition of non-oil GDP observed both in onshore and offshore oil AMCs are more likely to be due to a wealth effect, while changes confined to onshore AMCs are more likely due to the direct demand impact of oil-firm operations on local product markets. Furthermore, the wealth effect, unlike the relative demand from oil operations, should also be associated with a decline in labour supply and an increase in wages (outside the oil sector).

As we have seen, onshore oil operations change (modestly) the composition of non-oil GDP, while offshore ones leave it unchanged; furthermore, in regressions not reported here there is no evidence of a decline in labor supply, nor of a significant increase in wages. This leads us to conclude that the wealth effect is probably not present in the local Brazilian context, while we do find some (modest) effect on relative demand for goods and services from onshore oilfield operations. One possible reason for the absence of a wealth effect, of course, is that the oil money is misappropriated by political leaders, so the population fails to experience an increase in wealth. Another possibility is that the oil income is used to provide public goods that are poor substitutes for private ones (and for leisure). Clearly our results are more consistent with the former explanation.⁴⁷

⁴⁶ The net effect on aggregate municipality non-resource GDP is ambiguous. If the shrinking and expanding sectors have the same capital intensity aggregate GDP should increase. However if the expanding sector is more labor intensive aggregate GDP could remain unchanged or even fall.

⁴⁷ Among the versions of the Dutch disease model that our data *cannot* speak to, are those that focus on the reallocation of productive factors away from the non-resource and towards the resource sector. As discussed in detail in Section III.C, neither capital nor labor inputs are drawn significantly from the local economy, so this mechanism is unlikely to operate in our setting. Similarly we obviously do not expect Dutch-disease symptoms to be exacerbated by exchange rate appreciation, as is often thought to happen at the national level. Oft-cited models of

The theoretical literature has also identified several *political-economy* mechanisms that may be triggered by resource riches. In the classic mechanism royalty windfalls may increase **rent seeking**. Instead of producing marketable goods and services, a larger fraction of the population may be drawn into competing for political power and influence to secure for themselves a larger share of the income flowing into the municipality's budget [e.g. Tornell and Lane (1999) , Mehlum, Moene and Torvik (2006a, 2006b)]. Due to the effective decline in (productive) labor supply rent seeking is consistent with a decline in value added in the non-oil sector, and particularly in the private sector. As we have seen, we do not detect any such decline, so there is little support in the data for rent-seeking (as defined above) to be a dominant mechanism in the local Brazilian context.⁴⁸

Other political-economy mechanisms focus on the incentives and constraints faced by incumbent political leaders. One force at work is that with larger resource windfalls controlling the local government becomes more attractive, so political incumbents face more aggressive challenges for their position. Faced with a greater likelihood of losing power, incumbents may adopt a **shorter planning horizon**, resulting in a decline in the supply of government-provided infrastructure and other productive public services. As a consequence, non-oil GDP may also decline. Other predictions of this mechanism are a decline in average duration in office for mayors (or, roughly equivalently, a decline in incumbency advantage), as well as an increase in embezzlement of public revenues. On the other hand mayors may respond to the potential threat from political challengers by improving the **opportunity cost** of such challengers. To this end, they may actually expand the provision of productive infrastructure and public services with a view to increase private-sector opportunities. Hence, the opportunity-cost argument has diametrically opposed prediction to the planning-horizon mechanism. Our results suggest that the provision of productive public services is essentially invariant to the amount of oil revenues received by the municipal government, and are therefore consistent with neither of these views.

Yet another response by political incumbents may be to increase unproductive spending aimed at shoring up the political power base of the incumbent mayor, particularly through an increase in **patronage**, i.e. in an expansion of semi-fictitious government jobs to reward political supporters [e. g. Robinson, Torvik, and Verdier (2002), Acemoglu, Robinson, and Verdier (2004)]. This mechanism will mostly show up in an expansion in municipal-government employment.⁴⁹ We do find a very modest increase in public employment, so qualitatively there is something to be said

Dutch-disease featuring various combinations of the mechanisms discussed above include Corden and Neary (1982), Corden (1984), Krugman (1987), Wijnbergen (1994), Younger (1992), and Torvik (2001).

⁴⁸ Of course a clearly important form of resource-triggered rent seeking takes the extreme form of civil war [e.g. Caselli and Coleman (2008), Ross (2006)]. This is another example of a mechanism that we cannot pick up with within-country data, as is the fact that in some cases rent seeking over natural resources triggers the collapse of national-level institutions [Ross (2001a, 2001b)].

⁴⁹ The expansion in public employment following a resource windfall is a very rough proxy for the patronage effect. To the extent that some of the increase in public employment reflects genuine needs to manage the expanded size of the government budget, it is an overestimate. To the extent that patronage works by creating extra fictitious jobs in private firms that benefit from contracts from the municipal government, it is an underestimate.

for this view. Quantitatively, however, the increase in patronage is extremely small, so this mechanism does not appear to be of first-order importance.⁵⁰

The shorter-horizon story is not the only one to predict increased embezzlement by public officials. Another view especially emphasized in the political-science literature is that royalties on natural resources are more **easily stolen** than revenue flowing from general taxation. The political science literature explains this difference on semi-behavioral grounds: citizens tend to monitor more closely the utilization of funds coming directly from their own pockets (taxes) than those not arising, so to speak, from their own efforts (royalties). An alternative explanation, which is perhaps more theoretically palatable to economists, is that royalty revenue is intrinsically less *transparent*, so citizens do not have a precise estimate of how much money the government has and cannot accurately assess the extent of diversion to private uses by government officials. In our view, the ease-of-stealing story appears very consistent with the empirical results. Oil money does appear to be uncorrelated with the provision of public goods and services, while other forms of revenue seem to be positively correlated with spending. In addition, circumstantial evidence indicates that oil-rich municipalities are more corrupt. Finally, anecdotal evidence reveals widespread ignorance among the electorate on the magnitude of the oil-related fiscal revenues.

VII. Conclusions

We summarize our findings as follows. Offshore oil has no appreciable *market* linkages with the local economy: it does not use local factors of production, either directly or indirectly through purchases of locally-produced goods and services. As a result, local non-oil GDP is unaffected by the existence of offshore oil operations. Onshore oil triggers some reallocation of local factors of production from manufacturing activities to services, presumably through the direct demand for services by oil workers and firms. The net effect on aggregate non-oil GDP, however, is once again nil. There is therefore little evidence for traditional models of Dutch disease, particularly those operating through a wealth effect.

Both onshore and offshore oil generate significant increases in local-government revenues, mostly in the form of royalties. In turn, these revenue windfalls are matched to a very large extent by reported spending increases, particularly in the areas of urban infrastructure and housing, education and health services. However, various socio-economic outcomes that would be expected to respond to the recorded spending increases are unchanged (if not worsened) by the oil-induced income. Furthermore, increases in household income associated with oil-induced government revenues are at best in the order of one-fifth of what might be expected given the reported spending increases. This is strongly suggestive that a large fraction of the government revenue generated by oil is wasted, if not stolen.

⁵⁰ Caselli and Cunningham (2009) present a stylized toy model that nests the shortened-horizon, the opportunity cost, and the patronage effect, as well as several others that are unlikely to operate at the municipality level (e.g. more spending on political repression). Caselli (2007) shows that the relative importance of these mechanisms can vary with the level of development and with the size of the resource windfall, leading to potentially very non-monotonic relations between resource revenues and outcomes such as government investment, political stability, and growth.

This evidence allows us to discriminate to some extent among the main political-economy models that deal with the effects of resource abundance. Rent seeking models tend to emphasize the reallocation of effort from productive to unproductive activities, i.e. from producing goods and services to competing over the resource windfall – in this case the royalties flowing to the government. Since non-oil GDP does not fall, however, this mechanism does not seem to play a first-order effect here.

Similarly, models that emphasize the effect of oil revenues on the intensity of the political challenges faces by the political elite, and hence on its planning horizon, predict either an increase in socially productive spending (if the oil windfall makes the elite more secure) or a decline (if it induces it to shorten its planning horizon). Since “real” spending (as opposed to reported) hardly changes, these mechanisms also do not seem to play a first-order effect here.

There is definitely some support for a model of patronage, in that municipal employment increases, without an attendant increase in the quality and quantity of services received by the general municipality. But the overall increase in patronage-like spending, as measured by the increase in employment, is small. Furthermore, municipal-worker hourly wages are unchanged. Hence, this does not appear to be the main political-economy story coming out of the data.

Perhaps surprisingly, the theory that seems most consistent with the data is the old political-science view that oil revenues are somehow more “stealable” than other types of revenues. The (indirect) evidence for this interpretation is that OLS results (roughly capturing the use of the average Real of revenue) are very different from the IV results (capturing the effect of oil-related revenue). Whether this is because citizens themselves are more tolerant of corruption when the

and that reporting schemes should document the actual effective disbursement of sums, and not merely their recording on balance sheets.

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.

Afonso, J. R. Rodrigues and S. Wulff Gobetti (2008): "Renda do Petróleo no Brasil: Alguns Aspectos Fiscais e Federativos," *Revista do Binde*, 15(30), 231-269.

ANP (2001): Guia dos Royalties do Petróleo e do Gás Natural 2001

Bobonis, G. J. (2008): 'Endowments, Coercion, and the Historical Containment of Education,' unpublished, University of Toronto

Brunnschweiler, C.N. and Bulte, E.(2008). 'Natural Resources and Violent Conflict: Resource Abundance, Dependence and the Onset of Civil Wars', Economics working paper series 08/78, CER-ETH - Center of Economic Research (CER-ETH), ETH Zurich

Caselli, F (2007): "Power Struggles and the Natural Resource Curse," unpublished, LSE.

Caselli, F. & T. Cunningham (2007) "Leader Behaviour and the Natural Resource Curse," unpublished, LSE.

Caselli and Coleman (2006). 'On the Theory of Ethnic Conflict,' unpublished, LSE

Collier, P. and Goderis, B. (2007). 'Commodity Prices, Growth, and the Natural Resource Curse: Reconciling a Conundrum,' CSAE Working Paper: CSAE WPS/2007-15.

Corden, W.M. and Neary, J.P., 1982. Booming sector and de-industrialisation in a small open economy. *Economic Journal* **92**, pp. 825–848

Corden, W.M. (1984), "Booming Sector and Dutch Disease Economics, Survey and Consolidation" *Oxford Economic Papers*, (36) 1984 - Oxford Univ Press.

Costa Nova, L. (2005): *Análise do impacto social de receitas provenientes de royalties do petróleo em municípios do estado da Bahia*. Master's thesis, Universidade de Brasília.

De Oliveira Cruz, B. and M. B. Ribeiro (2008): "Sobre Maledicoes e Bencaos: e possivel gerir recursos naturais de forma sustentavel? Uma analise sobre os royalties e as compensacoes financeiras no Brasil," unpublished, Diretoria de Estudos Regionais e Urbanos.

Ferraz, C. and F. Finan (2008a): "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes," forthcoming, *Quarterly Journal of Economics*.

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

IBGE (2008): “Produto Interno Bruto dos Municípios,” Serie Relatorios Metodologicos, 29.

Isham, J. Woolcock, M., Pritchett, L and Busby, G. (2005). ‘The Varieties of Resource Experience: Natural Resource Export Structures and the Political Economy of Economic Growth’, *World Bank Economic Review*, Oxford University Press, vol. 19(2), pp. 141-174.

Kolstad, I. (2007). ‘The Resource Curse: Which Institutions Matter?’ CMI Working Papers wp200x-x, CMI (Chr. Michelsen Institute), Bergen, Norway

Krugman, P (1987). "The Narrow Moving Band, the Dutch Disease, and the Competitive Consequences of Mrs. Thatcher". *Journal of Development Economics* 27 (1-2), p. 41.

Leal, J.A.A. and R.V. Serra (2002): “Notas sobre os Fundamentos Economicos da Distribucao Espacial dos Royalties Petroliferos no Brasil,” *Anais do XXX Encontro Naciuonal de Economia (ANPEC)*.

Leite, C. and J. Weidmann (1999) ‘Does Mother Nature Corrupt? Natural Resources, Corruption and Economic Growth’, IMF Working Papers 99/85.

Lipscomb, M. and A. M. Mobarak (2007): “Decentralization and Water Pollution Spillovers: Evidence from the Re-drawing of County Boundaries in Brazil,” unpublished, University of Colorado.

Litschig, S. (2008): “Intergovernmental Transfers and Elementary Education: Quasi-Experimental Evidence from Brazil,” unpublished, Universitat Pompeu Fabra.

Manzano & Rigobon (2007) in: Lederman D. and Maloney W.F. (eds.), 2007, *Natural Resources: Neither Curse Nor Destiny*, Stanford University Press.

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

Michaels, G. (2008) “The Long Term Consequences of Resource Based Specialization,” unpublished, LSE

Naritomi, Joana, Rodrigo R. Soares, and Juliano J. Assuncao (2007): “Rent Seeking and the Unveiling of 'De Facto' Institutions: Development and Colonial Heritage within Brazil,” NBER Working Paper No. 13545.

Postali, F. A. Slaibe (2008): “Efeitos da distribuicao de royalties do petroleo sobre o crescimento dos Municipios no Brasil,” unpublished manuscript, Universidade de Sao Paulo.

Robinson, J. A., Torvik, R. and Verdier, T. (2002). ‘Political foundations of the resource curse’, CEPR Discussion Paper No. 3422.

Ross, Michael L. (2001a): *Timber Booms and Institutional Breakdown in Southeast Asia*, Cambridge University Press.

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

Ross, Michael 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.

Torvik (2001). "Learning by doing and the Dutch disease". *European economic review (0014-2921)*, 45 (2), p. 285.

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

Vicente, P. C. (2008): “Does Oil Corrupt? Evidence from a Natural Experiment in West Africa,” unpublished, Oxford University.

Wijnbergen, S.V. (1994), “The ‘Dutch Disease’: A Disease After All?” *The Economic Journal*, Vol. 94, No. 373 (Mar., 1984), pp. 41-55

Younger, Stephen D., Aid and the Dutch Disease: Macroeconomic Management when Everybody Loves You. *World Development*, Vol. 20, No. 11, pp. 1587-1597, November 1992.

Appendix: Rules governing the allocation of royalties from oil and gas in Brazil

This appendix is based on ANP (2001).

The current allocation of royalties is the result of a series of incremental legislative changes between 1953 and 1998. The incremental nature of the legislation has resulted in a rather complicated structure, which we now try to describe.

The total amount of royalty payments from each oilfield is the sum of two components. The first component is a fixed 5% of the value of the oil extracted. We call it the fixed quota. The second is a further percentage that must be between 0 and 5%. We call this the variable quota. However, even the variable quota is almost always set at the maximum of 5%. This is because the legislation authorizes ANP to assign a quota less than the maximum only in the case of lower-quality or higher-risk fields. As a result about 90% of the oilfields and all of the large oilfields pay between 9.1 and 10% in royalties. Another 9% pays between 8.1 and 9%. Only 1% pays less than 8%. The weighted average royalty is 9.8%.

This does not mean that the distinction between fixed quota and variable quota is irrelevant, though, because the sets of recipients of the two quotas, and the way the quotas are distributed among the various recipients, are very different. In particular, the fixed quota is divided as follows. For onshore fields, 70% to “producing” states, 20% to “producing” municipalities, and 10% to municipalities with significant offshore-oil related infrastructure (essentially, terminals for bringing offshore oil to land). For offshore fields, 30% to “facing” states, 30% to “facing” municipalities, 20% to the Navy, 10% to a “special fund” to be divided between all states and all municipalities, and 10% to municipalities with significant oil-related infrastructure. We come back to the definitions of “producing” and “facing” below.

The variable quota has even more recipients. For onshore oilfields, 52.5% goes to “producing states,” 25% to the Ministry of Science, 15% to “producing” municipalities, and 7.5% to municipalities “affected” by operations connected with the landing of offshore oil. For offshore oil, 25% to the Ministry of science, 22.5% to “facing” states, 22.5% to “facing” municipalities, 15% to the Navy, 7.5% to the “special fund,” and 10% to “affected” municipalities.

If one were to combine the percentages from the fixed and the variable quota, in the (typical) case of an oilfield paying the maximum royalty (i.e. 10%) then the total share going to *producing* municipalities in an onshore oilfield is 18%. For *offshore* fields, the percent of a 10% royalty going to *facing* municipalities would be 26%. Unfortunately, however, in the case of offshore oil things are not so simple, because the definition of “facing” is different in the case of the fixed quota and in the case of the variable quota. We thus now turn to the definitions of producing and facing.

For onshore oil, a state (municipality) is a “producing” state (municipality) vis-à-vis a certain oilfield if and only if there are wells tapping into that particular oilfield that are inside the state’s (municipality’s) borders. Each producing state (municipality) participates into the total royalties allocated to producing states (municipalities) from a certain oilfield in proportion to the *output*

share of the wells situated in that state (municipality) in the total oilfield output. This is true both for the fixed and for the variable quota.

For offshore oil, things get quite complicated. First of all, state and municipality maritime boundaries are needed. The relevant legislation assigns this task to the Brazilian Geographical Institute (BGE), and this has resulted into two distinct sets of borders. The first set is based on *perpendicular lines*. It begins by picking 25 points on the Brazilian coast, and connecting them by straight lines. This is necessary because the fractal nature of the coastline would otherwise make it impossible to draw perpendicular lines going out to sea. The 25 points include all the points at which two coastal state's boundaries reach the shore, but they also include a few extra points to accommodate extreme irregularities of the coastline inside a state.

Once the coastline has been 'linearized,' parallel lines going out to sea are drawn from the points where the state borders reach the coastline. These lines are deemed to be the continuation of the state border onto the continental shelf, and they "end" when they meet the outside boundary of the Brazilian continental shelf (i.e. the end of Brazil's territorial water).

Municipality boundaries based on perpendicular lines follow similar principles, with a few adjustments. First, in the states of Rio de Janeiro and São Paulo, a few more points are added to the 'linearization' of the coast so the linearized version used for municipal boundaries is a bit more jagged than the one used for state boundaries. Second, municipal borders end either when they reach the continental shelf boundaries (as was the case for state borders), or when they reach the state boundary.

The second set of boundaries is based on *parallel lines*. It amounts to identifying state (municipal) boundaries with the parallel passing by the point where state (municipal) boundaries reach the coast (and ending at the continental-shelf boundary).

Once state boundaries are available, a municipality is "facing" a certain oilfield if there are wells tapping into this oilfield that lie inside the municipality's maritime border. If the two sets of borders map the same well into two different municipalities they both have equal rights to the royalties.

This is far from the end of the story, though. For, each identified facing municipality must share the "facing" quota with a set of neighboring municipalities, called the geo-economic area. The construction of the geo-economic area begins by identifying the "geographic mesoregion" to which the facing municipality belongs. The geographic mesoregion is a purely geographic construct that exists independently of the royalty allocation mechanism. Each municipality belongs to one and only one geographic mesoregion.

Next, within the mesoregion, IBGE identifies a "main production zone." The facing municipality must always belong to the main production zone. In addition, this zone includes municipalities with at least three of the following: (i) infrastructure for processing, treating, storing, and shipping oil (excluding pipelines); (ii) infrastructure supporting exploration, extraction, and shipment of oil (ports, airports, manufacture and maintenance of oil-rig equipment, etc.). However, it turns out that very few (i.e. eight) municipalities fulfill this criterion, so the vast

majority of the municipalities in the main production zone are the municipalities facing the oil wells.

Once the main production zone is identified, the geo-economic area is the union of two sets: all the municipalities in the mesoregion, and all the municipalities which border the main production zone. In several cases, of course, the latter set is a subset of the former, but often that is not the case. Municipalities in a geo-economic area that are not in the main production zone are assigned to the “geographic zone contiguous to the main production zone.” Municipalities in the geo-economic area whose territory is crossed by pipelines transporting offshore oil/gas are assigned to the “secondary production zone.” There are only 8 of these. All municipalities in the secondary production zone will therefore also belong to the geographically contiguous zone, but if one municipality receives royalties by virtue of being in the secondary zone then it is excluded from division of royalties based on the contiguity criterion.

Given the total royalties from the fixed quota going to a certain geo-economic area based on the “facing” principle (i.e. 30% of 5%), there is a first round of allocation that works as follows: 60% to the main zone, 10% to the secondary zone, and 30% to the contiguous zone. Next, within each zone, each municipality’s share depends on its population size.⁵² Recall that in practice, the main zone tends to be constituted almost exclusively by the facing municipality, so in practice we should expect close to 60% of the 30% of the 5% to go to facing municipalities.

Note that, as we have seen, the same oil well may well be inside two municipality’s borders, depending on the perpendicular or parallel principle. If the two municipalities are in the same geo-economic area this fact has no implications whatsoever for the allocation of royalties, as the identity of the municipality in whose border the well lies does not affect the allocation within the geo-economic area. But, of course, if the two municipalities are in different geo-economic areas then the number of municipalities sharing in the royalties increase accordingly.

Finally, how is the total “facing” component of the royalty allocated between geo-economic areas? The principle is the same as for onshore oil, i.e. each geo-economic area receives royalties in proportion to the *output share* of the wells situated inside its maritime borders in the total oilfield output.

Things are very different, and much simpler, for the variable quota. First, here it is facing municipalities only: nothing goes to the geo-economic areas. (Although there is a 7.5% separate quota for municipalities with infrastructure, but this is outside the “facing” quota). Second, and more importantly, the identification of a facing municipality is no longer based on the location of wells, but on the location of fields. In particular, for each field, the set of facing municipalities is the set of municipalities whose borders’ extensions on the continental shelf (whether drawn with perpendiculars or parallels) contain any portion of the field.

⁵² However, if one municipality in the main zone has at least three pieces of land-based infrastructure for processing, treatment, storage, and disposal of offshore oil it must receive at least one-third of the overall royalties going to the zone in which it sits. Hence, if the allocation based on population implies that this municipality receives less than one-third of the total, a new allocation is made where it receives one-third and the remainder is divided among the others based on the population criterion. In practice, only two municipalities satisfy this criterion.

The allocation of the overall 22.5% of the variable quota among facing municipalities is pro-rated based on the simple average of the municipality's share in the total field area based on perpendiculars and based on parallels. Hence, if $a_1(m,f)\%$ of field f lies inside the maritime borders of municipality m by the perpendicular-line criterion, and $a_2(m,f)\%$ according to the parallel-line criterion, then the royalties of municipality m (based on the “facing” criterion alone, and only on the variable quota) are

$$(1/2) \times [a_1(m,f) + a_2(m,f)] \times 0.225 \times q(f),$$

where $q(f)$ is the value of the output of field f . Accordingly, the formula by which we seek to assign offshore output to municipalities is

$$(1/2) \times [a_1(m,f) + a_2(m,f)] \times q(f).$$

Table 1. Summary Statistics for Brazilian AMCs

	(1)	(2)	(3)	(4)	(5)
	No oil	Only AMCs where oil was first discovered after 1970	All AMCs with oil	AMCs with offshore oil only	AMCs with onshore oil only
GDP per capita in 2002 (Brazilian R\$2000)	4,313	6,288	6,813	8,058	6,305
Municipal revenues per capita in 2000 (Brazilian R\$2000)	508	566	539	699	440
Population in 2000	45,116	74,031	90,944	93,372	82,175
Latitude	-16.6	-13.5	-12.7	-17.8	-10.6
Longitude	45.0	40.5	39.3	42.1	38.2
Coast dummy	0.04	0.66	0.55	1.00	0.27
Distance to the federal capital (kilometers)	1,016	1,320	1,271	1,180	1,297
Distance to the state capital (kilometers)	245	134	99	99	96
Mean oil output per capita in 2000 (Brazilian R\$2000)	0	3,388	2,824	3,894	2,319
90th percentile of oil output per capita in 2000 (Brazilian R\$2000)	0	9,193	6,949	6,949	6,737
95th percentile of oil output per capita in 2000 (Brazilian R\$2000)	0	17,577	10,249	21,319	9,193
Maximum oil output per capita in 2000 (Brazilian R\$2000)	0	45,221	45,221	45,221	15,661
Observations (AMCs)	3,556	59	103	31	63

Notes: Each AMC includes one municipality or multiple municipalities. In all tables "Oil" denotes both oil and natural gas. 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.

Table 2. Test of Conditional Random Assignment

Dependent variable: GDP per capita	(1)	(2)	(3)	(4)
	Post-1970	All	offshore	onshore
(Oil output per capita in 2000) x (year = 1970)	-0.025 (0.018)	-0.028 (0.021)	-0.034 (0.020)	0.003 (0.041)
(Oil output per capita in 2000) x (year = 1980)	-0.011 (0.025)	-0.010 (0.026)	-0.015 (0.026)	0.037 (0.064)
(Oil output per capita in 2000) x (year = 1996)	0.060 (0.106)	0.093 (0.108)	0.075 (0.131)	0.177 (0.161)
(Oil output per capita in 2000) x (year = 2002)	0.548 (0.035)	0.518 (0.048)	0.552 (0.026)	0.454 (0.149)
(Oil output per capita in 2000) x (year = 2005)	0.803 (0.044)	0.779 (0.070)	0.858 (0.032)	0.608 (0.214)
AMCs	3,615	3,659	3,587	3,619
Observations	18,075	18,295	17,935	18,095

Notes: Each column reports coefficients from a regression using a panel of AMCs (each AMC includes one municipality or multiple municipalities). All panels include all AMCs without oil. Column 1 adds AMCs where oil was first discovered after 1970; Column 2 all AMCs with oil; Column 3 those AMCs with offshore oil only; and Column 4 those AMCs with onshore oil only. We use data for 1970, 1980, 1996, 2002, and 2005. All values are in Brazilian R\$2000. All regressions control for year dummies interacted with latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors clustered by AMC are in parentheses.

Table 3. The Effect of Oil Output Per Capita on GDP Per Capita, by Sector

	Post-1970			All		Offshore		Onshore	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	GDP Per Capita	GDP Per Capita in Industry	GDP Per Capita in Non-Industry	GDP Per Capita in Industry	GDP Per Capita in Non-Industry	GDP Per Capita in Industry	GDP Per Capita in Non-Industry	GDP Per Capita in Industry	GDP Per Capita in Non-Industry
(Oil output per capita in 2000) x (year = 2000)	-	-	0.004 (0.016)	-	0.012 (0.020)	-	-0.009 (0.013)	-	0.119 (0.049)
(Oil output per capita in 2001) x (year = 2001)	-	-	0.025 (0.030)	-	0.041 (0.036)	-	0.005 (0.021)	-	0.161 (0.069)
(Oil output per capita in 2002) x (year = 2002)	0.400 (0.032)	0.381 (0.012)	0.019 (0.022)	0.356 (0.025)	0.025 (0.024)	0.375 (0.008)	0.012 (0.019)	0.297 (0.117)	0.091 (0.046)
(Oil output per capita in 2003) x (year = 2003)	0.453 (0.051)	0.434 (0.042)	0.019 (0.023)	0.410 (0.044)	0.025 (0.024)	0.439 (0.046)	0.009 (0.020)	0.324 (0.135)	0.100 (0.051)
(Oil output per capita in 2004) x (year = 2004)	0.364 (0.033)	0.354 (0.015)	0.010 (0.021)	0.335 (0.032)	0.020 (0.022)	0.354 (0.015)	0.002 (0.021)	0.284 (0.127)	0.098 (0.045)
(Oil output per capita in 2005) x (year = 2005)	0.449 (0.026)	0.447 (0.016)	0.002 (0.015)	0.423 (0.037)	0.009 (0.016)	0.451 (0.015)	-0.002 (0.015)	0.311 (0.144)	0.084 (0.046)
Observations	14,460	14,460	21,690	14,636	21,954	14,348	21,522	14,476	21,714

Notes: Each column reports coefficients from a regression using a panel of AMCs (each AMC includes one municipality or multiple municipalities). All panels include all AMCs without oil. Columns 1-3 add AMCs where oil was first discovered after 1970; Columns 4-5 all AMCs with oil; Columns 6-7 those AMCs with offshore oil only; and Columns 7-8 those AMCs with onshore oil only. Industry includes manufacturing, mineral extraction, civilian construction, and public utilities. The calculation of GDP in industry (and total GDP) from oil changed in 2002 - see paper for details. All values are in Brazilian R\$2000. All regressions control for year dummies interacted with latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies.

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

	(1)	(2)	(3)	(4)	(5)
	Total municipality revenues per capita in 2000	Total municipality revenues per capita in 2000 (see footnote)	Royalties from oil in 2000 (only AMCs in column 2)	Total municipality revenues per capita in 1991 (see footnote)	Change in total municipality revenues per capita from 1991- 2000 (see footnote)
A: Post-1970	0.0293 (0.0042)	0.0296 (0.0041)	0.0195 (0.0024)	0.0017 (0.0007)	0.0288 (0.0034)
Observations (AMCs)	3,208	3,512	3,512	3,594	3,493
B: All	0.0296 (0.0040)	0.0327 (0.0037)	0.0189 (0.0021)	0.0028 (0.0009)	0.0308 (0.0031)
Observations (AMCs)	3,242	3,553	3,553	3,638	3,534
C: Offshore	0.0304 (0.0040)	0.0308 (0.0039)	0.0181 (0.0020)	0.0018 (0.0008)	0.0291 (0.0032)
Observations (AMCs)	3,183	3,484	3,484	3,567	3,466
D: Onshore	0.0212 (0.0067)	0.0363 (0.0075)	0.0168 (0.0028)	0.0069 (0.0017)	0.0341 (0.0062)
Observations (AMCs)	3,205	3,514	3,514	3,598	3,495

Notes: Each column reports coefficients on Oil Output per Capita in 2000 from a regression using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). All samples include all AMCs without oil. Panel A adds AMCs where oil was first discovered after 1970; Panel B all AMCs with oil; Panel C those AMCs with offshore oil only; and Panel D those AMCs with onshore oil only. All values are in Brazilian R\$2000. Since we only have municipal revenues for about 91 percent of the AMCs in 2000, columns (2)-(5) predict 2000 municipal revenues from 2001 municipal revenues using a linear regression to complete missing 2000 values. Similarly, in columns (4) and (5) we use 1992 municipal revenues to predict missing 1991 values. All regressions control for latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

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

	(1)	(2)	(3)	(4)	(5)	(6)
	Functional expenditures					
	Education and Culture	Health and Sanitation	Housing and urban development	n	Social Transfers	Total
	Levels (2000)					
A: All, OLS	0.271 (0.009)	0.122 (0.006)	0.105 (0.007)	0.065 (0.006)	0.066 (0.004)	0.904 (0.020)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553
B: All, IV	0.149 (0.023)	0.090 (0.013)	0.242 (0.050)	0.090 (0.027)	0.044 (0.013)	0.904 (0.060)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553
C: Offshore, IV	0.133 (0.017)	0.102 (0.009)	0.186 (0.017)	0.132 (0.020)	0.048 (0.006)	0.829 (0.029)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484
	Changes (1991-2000)					
D: All, OLS	0.270 (0.012)	0.120 (0.008)	0.097 (0.010)	0.045 (0.006)	0.068 (0.005)	0.874 (0.027)
Obs. (AMCs)	3,423	3,423	3,423	3,423	3,423	3,423
E: All, IV	0.155 (0.025)	0.082 (0.015)	0.239 (0.054)	0.089 (0.026)	0.042 (0.014)	0.910 (0.066)
Obs. (AMCs)	3,423	3,423	3,423	3,423	3,423	3,423
F: Offshore, IV	0.143 (0.018)	0.103 (0.011)	0.185 (0.021)	0.127 (0.019)	0.048 (0.006)	0.841 (0.038)
Obs. (AMCs)	3,355	3,355	3,355	3,355	3,355	3,355

Notes: Each cell reports the coefficient on municipal revenues per capita in 2000 (Panels A-C) or on its change between 1991-2000 from a regression using a cross section of AMCs (each AMC includes one municipality or more). All samples include all AMCs without oil. Panels A, B, D, and E add all AMCs with oil; Panels C and F only those AMCs with offshore oil. Panels A and D report simple OLS coefficients while Panels B, C, D, and E instrument revenues with oil output per capita. 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, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Per capita residential capital	Rooms at home per 1000 people aged 16-64	Percent of population living in "standard" (not sub- standard) housing	Percent of population living in housing with electricity	Percent of population living in housing with garbage collection	Percent of population living in housing with piped water	Percent of households with water linked to main network	Percent of households with toilets linked to main network	Per capita 1000km of paved roads under municipal jurisdiction in 2005
Levels (2000, except for column 9)									
A: All, OLS	0.00029 (0.00016)	0.163 (0.033)	-0.0003 (0.0002)	0.0026 (0.0005)	0.0027 (0.0007)	0.0029 (0.0007)	-0.0018 (0.0012)	-0.0027 (0.0015)	0.00008 (0.00004)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553	3,553	3,553	3,553
B: All, IV	-0.00072 (0.00047)	0.070 (0.114)	-0.0092 (0.0035)	0.0089 (0.0051)	0.0091 (0.0035)	0.0041 (0.0058)	-0.0053 (0.0134)	-0.0073 (0.0085)	-0.00002 (0.00002)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553	3,553	3,553	3,553
C: Offshore, IV	-0.00143 (0.00070)	0.190 (0.150)	-0.0140 (0.0016)	0.0010 (0.0016)	0.0068 (0.0028)	-0.0063 (0.0035)	-0.0212 (0.0108)	-0.0129 (0.0090)	-0.00002 (0.00001)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484	3,484	3,484	3,484
Changes (1991-2000)									
D: All, OLS	0.00044 (0.00008)	0.017 (0.028)	-0.00030 (0.00013)	0.0015 (0.0006)	0.0085 (0.0018)	0.0028 (0.0008)	0.0046 (0.0011)	0.0020 (0.0014)	
Obs. (AMCs)	3,534	3,534	3,534	3,534	3,534	3,534	3,534	3,534	
E: All, IV	-0.00071 (0.00029)	-0.112 (0.082)	-0.0050 (0.0013)	-0.0104 (0.0061)	-0.0004 (0.0090)	0.0004 (0.0040)	-0.0067 (0.0083)	-0.0171 (0.0054)	
Obs. (AMCs)	3,534	3,534	3,534	3,534	3,534	3,534	3,534	3,534	
F: Offshore, IV	-0.00110 (0.00056)	-0.125 (0.083)	-0.0066 (0.0008)	0.0009 (0.0020)	-0.0046 (0.0039)	-0.0013 (0.0021)	-0.0150 (0.0050)	-0.0219 (0.0046)	
Obs. (AMCs)	3,466	3,466	3,466	3,466	3,466	3,466	3,466	3,466	

Notes: Each cell reports the coefficient on municipal revenues per capita in 2000 (Panels A-C) or on its change between 1991-2000 from a regression using a cross section of AMCs (each AMC includes one municipality or more). All samples include all AMCs without oil. Panels A, B, D, and E add all AMCs with oil; Panels C and F only those AMCs with offshore oil. Panels A and D report simple OLS coefficients while Panels B, C, D, and E instrument revenues with oil output per capita. 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, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Municipal teachers per million people	Municipal classrooms per million people	Municipal teachers per million people	Municipal classrooms per million people	Municipal health establishments with inpatient care per million people	Municipal health establishments without inpatient care per million people	Social transfers per capita
Levels in:							
	2000	2000	2005	2005	2002	2002	2000
A: All, OLS	4.8 (0.4)	2.8 (0.2)	5.0 (0.4)	3.1 (0.3)	0.050 (0.009)	0.284 (0.032)	-0.008 (0.008)
Obs. (AMCs)	3550	3550	3553	3553	3553	3553	3553
B: All, IV	2.1 (0.9)	0.5 (0.5)	4.2 (0.9)	1.3 (0.5)	-0.006 (0.017)	-0.015 (0.073)	-0.002 (0.003)
Obs. (AMCs)	3550	3550	3553	3553	3553	3553	3553
C: Offshore, IV	1.9 (0.8)	0.8 (0.5)	4.6 (1.0)	1.7 (0.6)	0.012 (0.005)	-0.020 (0.057)	-0.005 (0.003)
Obs. (AMCs)	3481	3481	3484	3484	3484	3484	3484
Changes in:							
	1996-2000	1996-2000	1996-2005	1996-2005	1992-2002	1992-2002	
D: All, OLS	3.0 (0.4)	1.3 (0.2)	3.3 (0.5)	1.7 (0.2)	0.003 (0.011)	0.145 (0.030)	
Obs. (AMCs)	3445	3445	3446	3446	3534	3534	
E: All, IV	0.5 (1.1)	0.3 (0.7)	2.6 (1.3)	1.2 (1.0)	0.027 (0.017)	0.105 (0.095)	
Obs. (AMCs)	3445	3445	3446	3446	3534	3534	
F: Offshore, IV	0.4 (1.3)	-0.2 (0.5)	3.3 (1.7)	0.7 (0.6)	0.012 (0.006)	0.174 (0.063)	
Obs. (AMCs)	3377	3377	3378	3378	3466	3466	

Notes: Each cell reports the coefficient on municipal revenues per capita in 2000 (Panels A-C) or on its change between 1991-2000 from a regression using a cross section of AMCs (each AMC includes one municipality or more). All samples include all AMCs without oil. Panels A, B, D, and E add all AMCs with oil; Panels C and F only those AMCs with offshore oil. Panels A and D report simple OLS coefficients while Panels B, C, D, and E instrument revenues with oil output per capita. 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, 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 welfare income of households for 1991.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Per capita household income	Per capita household income: 1st (=bottom) quintile	Per capita household income: 2nd quintile	Per capita household income: 3rd quintile	Per capita household income: 4th quintile	Per capita household income: 5th quintile	Percent poor
Levels (2000)							
A: All, OLS	0.21 (0.12)	0.07 (0.01)	0.10 (0.03)	0.17 (0.06)	0.27 (0.11)	0.41 (0.39)	-0.0025 (0.0006)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553	3,553
B: All, IV	0.16 (0.22)	0.11 (0.04)	0.17 (0.08)	0.20 (0.11)	0.18 (0.21)	0.17 (0.74)	-0.0101 (0.0042)
Obs. (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553	3,553
C: Offshore, IV	0.08 (0.30)	0.09 (0.04)	0.09 (0.08)	0.08 (0.14)	0.06 (0.30)	0.09 (0.99)	-0.0035 (0.0029)
Obs. (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484	3,484
Changes (1991-2000)							
D: All, OLS	0.18 (0.08)	0.06 (0.02)	0.08 (0.02)	0.13 (0.04)	0.19 (0.07)	0.46 (0.29)	-0.0034 (0.0009)
Obs. (AMCs)	3,444	3,444	3,443	3,443	3,443	3,444	3,444
E: All, IV	0.11 (0.18)	0.14 (0.05)	0.18 (0.05)	0.17 (0.07)	0.17 (0.15)	-0.14 (0.72)	-0.0059 (0.0027)
Obs. (AMCs)	3,444	3,444	3,443	3,443	3,443	3,444	3,444
F: Offshore, IV	0.08 (0.25)	0.07 (0.03)	0.09 (0.05)	0.09 (0.11)	0.15 (0.21)	-0.07 (0.90)	-0.0038 (0.0031)
Obs. (AMCs)	3,376	3,376	3,375	3,375	3,375	3,376	3,376

Notes: Each cell reports the coefficient on municipal revenues per capita in 2000 (Panels A-C) or on its change between 1991-2000 from a regression using a cross section of AMCs (each AMC includes one municipality or more). All samples include all AMCs without oil. Panels A, B, D, and E add all AMCs with oil; Panels C and F only those AMCs with offshore oil. Panels A and D report simple OLS coefficients while Panels B, C, D, and E instrument revenues with oil output per capita. 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, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses

Table 9. Effect of Municipal Revenues from Oil on House Size

	(1)	(2)	(3)
	Rooms at home per 1000 municipal employees	Rooms at home per 1000 adults who are not municipal employees	Rooms at home per 1000 adults: municipal workers minus other adults
	Levels (2000)		
A: All, OLS	0.08	0.16	-0.08
	(0.05)	(0.03)	(0.05)
Obs. (AMCs)	3,550	3,550	3,550
 B: All, IV	 0.10	 0.06	 0.04
	(0.23)	(0.12)	(0.22)
Obs. (AMCs)	3,550	3,550	3,550
 C: Offshore, IV	 0.49	 0.17	 0.32
	(0.17)	(0.15)	(0.13)
Obs. (AMCs)	3,481	3,481	3,481
	Changes (1991-2000)		
D: All, OLS	-0.05	0.02	-0.06
	(0.09)	(0.03)	(0.09)
Obs. (AMCs)	3,525	3,525	3,525
 E: All, IV	 0.40	 -0.12	 0.52
	(0.29)	(0.08)	(0.30)
Obs. (AMCs)	3,525	3,525	3,525
 F: Offshore, IV	 0.38	 -0.14	 0.52
	(0.17)	(0.08)	(0.14)
Obs. (AMCs)	3,457	3,457	3,457

Notes: Each cell reports the coefficient on municipal revenues per capita in 2000 (Panels A-C) or on its change between 1991-2000 from a regression using a cross section of AMCs (each AMC includes one municipality or more). All samples include all AMCs without oil. Panels A, B, D, and E add all AMCs with oil; Panels C and F only those AMCs with offshore oil. Panels A and D report simple OLS coefficients while Panels B, C, D, and E instrument revenues with oil output per capita. 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, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

Table 10. Effect of Oil Output on News About Fraud and Federal Police Operations

	(1)	(2)	(3)	(4)	(5)	(6)
	Dependent variable: indicator for news on fraud possibly related to a mayor (coefficient multiplied by 10 ⁶)					
	All	Offshore	Onshore	All	Offshore	Onshore
Oil output per capita in 2000	5.3 (7.7) 3659	6.7 (10.6) 3587	3.9 (10.9) 3619			
Oil output in 2000				0.118 (0.030) 3659	0.113 (0.029) 3587	0.748 (0.257) 3619
	Dependent variable: indicator for federal police operation involving mayor (coefficient multiplied by 10 ⁶)					
Oil output per capita in 2000	12.4 (9.0) 3659	12.6 (11.9) 3587	13.0 (12.0) 3619			
Oil output in 2000				0.153 (0.034) 3659	0.147 (0.031) 3587	0.829 (0.212) 3619

Notes: Each cell reports coefficients from a regression using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). All samples include all AMCs without oil. Columns 1 and 4 add all AMCs with oil; Columns 2 and 5 only AMCs with offshore oil; and Columns 3 and 6 only AMCs with onshore oil. All values are in Brazilian R\$2000. The dependent variable in the top panel is an indicator for at least one webpage on Brazil's News Agency mentioning fraud, the mayor, and the name of at least one of the municipalities in the AMC. The search code was: fraude prefeito "<Municipality name in Portuguese>" site:www.agenciabrasil.gov.br/noticias. For example, for the municipality Belém do Brejo do Cruz, the search was: fraude prefeito "Belém do Brejo do Cruz" site:www.agenciabrasil.gov.br/noticias. The search was conducted from 21-26 March 2009. The dependent variable in the bottom panel is an indicator for having at least one current or former mayor investigated by a federal police operation from 2003-2008. All regressions control for latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

Table 11. News About Allegations of Corruption and Fraud in Oil-Rich Municipalities

Rank (Oil Output in 2000)	Rank (Per Capita Oil Output in 2000)	Municipality	Event	Authority involved	Amount involved	Source
1	4	Campos dos Goytacazes	Mayor removed from office and several associates arrested in a case involving fraud in bidding and illegal hiring in the city council.	Federal Police Operation "Telhado de Vidro"	Up to R\$240 million	O Globo, O Dia Online
3	12	Macaé	State prosecution took legal action against the mayor between 1988-2004, who is accused of misusing public funds and participating in a scheme directing biddings for contracts of municipal works.	Prosecution of the State of Rio de Janeiro	Not specified	Gazetta Mercantil
"	"	"	Re-elected mayor and ex-mayor accused of dishonest administration and cheating in bidding of school lunches.	Federal Public Prosecution	Cumulative R\$ 1.5 million	Agencia Brasil
5	20	Cabo Frio	Former mayor and current senior municipal employee exchange allegations regarding an investigation by the federal police that targeted the municipality (among other municipalities).	Federal Police Operation "Joao de Barra"	Estimated R\$700 million (over many municipalities)	Agencia Brasil, O Globo, Jornal do Brasil
6	10	Coari	Mayor and associates arrested in a federal police operation under allegations of corruption and criminal organization. This criminal organization was accused of directing bids, overbilling, and faking work to appropriate resources allocated by the federal government and Petrobras for the exploration of oil and gas. 24 hours later, the headquarters of the daily newspaper, Diario do Amazonas, which reportedly covered the case, was attacked by gunshots.	Federal Police Operation "Vorax"	Allegedly at least R\$50 million	Agencia Brasil, Ultimo Segundo, and a blog
8	5	Armação de Búzios	Mayor ordered to appear before the state court to explain irregularities in accounts.	Court of the State of Rio de Janeiro	Not specified	Amarribo (an NGO)
9	3	Carapebus	Two former mayors accused of corruption and fraud in bidding and were investigated as part of a federal police operation	Federal Police Operation "Pasárgada"	Not specified	O Globo

Notes: This table lists journalistic accounts of events involving alleged fraud, corruption, and other illegal activities associated with mayors in municipalities that ranked in the top 10 in Brazil in total oil output in 2000.

Table A1. Further Tests of Conditional Random Assignment

	Dependent variable: outcome in 1970					
	(1)	(2)	(3)	(4)	(5)	(6)
	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
Post-1970	-0.000002 (0.000006)	-0.000055 (0.000193)	-0.000022 (0.000014)	-0.000339 (0.000169)	-0.000084 (0.000162)	-0.000134 (0.000121)
Observations (AMCs)	3,615	3,615	3,615	3,615	3,615	3,615
All	0.000002 (0.000007)	-0.000174 (0.000233)	-0.000016 (0.000015)	-0.000146 (0.000230)	-0.000095 (0.000138)	-0.000068 (0.000137)
Observations (AMCs)	3,659	3,659	3,659	3,659	3,659	3,659
Offshore	-0.000008 (0.000005)	0.00014 (0.00012)	-0.00003 (0.00001)	-0.00041 (0.00018)	-0.00012 (0.00018)	-0.00019 (0.00013)
Observations (AMCs)	3,587	3,587	3,587	3,587	3,587	3,587
All, absolute value	0.000046 (0.000031)	-0.00120 (0.00058)	0.00001 (0.00008)	0.00054 (0.00109)	0.00071 (0.00077)	0.00073 (0.00065)
Observations (AMCs)	3,659	3,659	3,659	3,659	3,659	3,659

Notes: The table reports coefficients on oil output per capita in 2000 (Panels A-C) or in absolute value (Panel D) from regressions using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). Outcomes are measured for 1970. All values are in Brazilian R\$2000 (or in millions of R\$2000 in the bottom panel). All regressions control for latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors are in parentheses.

Table A2. No Significant Effect of Oil on Population

	All AMCs		AMCs with offshore oil (or no oil) only	
	(1)	(2)	(3)	(4)
	Dependent variable: ln(population)	Dependent variable: population	Dependent variable: ln(population)	Dependent variable: population
(Oil output per capita in 2000) x (year = 1970)	-0.0000211 (0.0000129)	-2.996 (2.326)	-0.0000148 (0.0000161)	-3.397 (3.117)
(Oil output per capita in 2000) x (year = 1980)	-0.0000216 (0.0000129)	-4.091 (2.750)	-0.0000147 (0.0000155)	-4.541 (3.687)
(Oil output per capita in 2000) x (year = 1991)	-0.0000185 (0.0000125)	-4.688 (2.999)	-0.0000119 (0.0000150)	-4.826 (4.010)
(Oil output per capita in 2000) x (year = 1996)	-0.0000135 (0.0000122)	-4.715 (3.041)	-0.0000071 (0.0000142)	-4.736 (4.063)
(Oil output per capita in 2000) x (year = 2000)	-0.0000130 (0.0000122)	-5.020 (3.199)	-0.0000063 (0.0000138)	-4.971 (4.269)
(Oil output per capita in 2000) x (year = 2005)	-0.0000126 (0.0000123)	-5.462 (3.447)	-0.0000054 (0.0000131)	-5.268 (4.589)
Years: 1970, 1980, 1991, 1996, 2000, 2005	X	X	X	X
AMCs	3,615	3,615	3,615	3,615
Observations	21,690	21,690	21,690	21,690

Notes: Each column reports coefficients from a regression using a panel of AMCs (each AMC includes one municipality or multiple municipalities). The 44 AMCs in which oil was first discovered before 1971 are excluded from the sample. Throughout the table "oil" denotes both oil and natural gas. All regressions control for year dummies interacted with latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and state dummies. Robust standard errors clustered by AMC are in parentheses.

Appendix 2: data sources (to be completed)

Variable	Source(s)	Geographic level at which the variable was obtained	URL or where we got the data	Primary source
Area (square kilometers)	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE
Latitude	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE
Longitude	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE
Coast dummy	IBGE map (our calculation)	Municipality (1997 boundaries, which were those used for the 2000 census)	http://maps.ibge.gov.br/divisao/viewer.htm	IBGE
Distance to the federal capital (kilometers)	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE
Distance to the state capital (kilometers)	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE
Population in 1970	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE (Household Census)
Population in 2000	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IBGE (Household Census)
GDP per capita in 1970 (Brazilian R\$2000)	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Sistema de Contas Regionais Referência 2002
GDP per capita in 2002 (Brazilian R\$2000)	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA
Oil output in 2000 (Brazilian R\$2000)	ANP	Municipality (1997 boundaries, which were those used for the 2000 census)	see text	ANP, IBGE
Oil royalties per capita in 2000 (Brazilian R\$2000)	ANP	Municipality (1997 boundaries, which were those used for the 2000 census)	see text	ANP, IBGE
Household income per capita in 2000 (Brazilian R\$2000)	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Date of discovery	ANP - BDEP GIS database	Oilfield - matched to municipality (1997 boundaries)	http://maps.bdep.gov.br/website/maps/viewer.htm	ANP - BDEP
Producing oilfield dummy	ANP - BDEP GIS database	Oilfield - matched to municipality (1997 boundaries)	http://maps.bdep.gov.br/website/maps/viewer.htm	ANP - BDEP
Offshore dummy	ANP - BDEP GIS database; IBGE map (our calculation)	Municipality (1997 boundaries, which were those used for the 2000 census)	http://maps.bdep.gov.br/website/maps/viewer.htm ; http://maps.ibge.gov.br/divisao/viewer.htm	ANP, ANP - BDEP
Onshore dummy	ANP - BDEP GIS database; IBGE map (our calculation)	Municipality (1997 boundaries, which were those used for the 2000 census)	http://maps.bdep.gov.br/website/maps/viewer.htm ; http://maps.ibge.gov.br/divisao/viewer.htm	ANP, ANP - BDEP
GDP Per Capita	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA
GDP Components Per Capita	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA
Per capita total municipality revenues	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita municipal functional expenditure on education and culture	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita municipal functional expenditure on health and sanitation	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita municipal functional expenditure on housing and urban development	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita municipal functional expenditure on legislation	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita municipal functional expenditure on welfare	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita total municipal expenditure	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Ministry of Finance (Fonte: Ministério da Fazenda - Secretaria do Tesouro Nacional)
Per capita residential capital	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Rooms at home per 1000 people aged 16-64	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Percent of population living in "standard" (not sub-standard) housing	IPEA Data	AMC0070	Obtained informally	Household Census
Percent of population living in housing with electricity	IPEA Data	AMC0070	Obtained informally	Household Census
Percent of population living in housing with garbage collection	IPEA Data	AMC0070	Obtained informally	Household Census
Percent of population living in housing with piped water	IPEA Data	AMC0070	Obtained informally	Household Census
Percent of households with water linked to main network	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Household Census
Percent of households with toilets linked to main network	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	Household Census
Per capita 1000km of paved roads under municipal jurisdiction in 2005	Ministry of Transportation: BASE DE DADOS GEORREFERENCIADA	Municipality (2005 boundaries) - matched to 1997 municipality boundaries	http://www.transportes.gov.br/PNLT/DVD_BD_2006/Index_BD.htm	Ministry of Transportation
Municipal teachers per million people in 2000	Ministério da Educação	Municipality (2000 boundaries) - matched to 1997 municipality boundaries	Obtained informally	
Municipal classrooms per million people in 2000	Ministério da Educação	Municipality (2000 boundaries) - matched to 1997 municipality boundaries	Obtained informally	
Municipal teachers per million people in 2005	Ministério da Educação	Municipality (2005 boundaries) - matched to 1997 municipality boundaries	Obtained informally	
Municipal classrooms per million people in 2005	Ministério da Educação	Municipality (2005 boundaries) - matched to 1997 municipality boundaries	Obtained informally	
Health establishments with inpatient care per million people in 2002	IBGE/DataSUS (?)	Municipality (2002 boundaries) - matched to 1997 municipality boundaries	http://w3.datasus.gov.br/datasus/datasus.php	
Health establishments without inpatient care per million people in 2002	IBGE/DataSUS (?)	Municipality (2002 boundaries) - matched to 1997 municipality boundaries	http://w3.datasus.gov.br/datasus/datasus.php	
Welfare income per capita	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		
Per capita household income	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Per capita household income: 1st (=bottom) quintile	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Per capita household income: 2nd quintile	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Per capita household income: 3rd quintile	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Per capita household income: 4th quintile	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Per capita household income: 5th quintile	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Percent poor	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)		Household Census
Municipal employees per 1000 people	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Annual hours worked per capita	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Rooms at home per 1000 municipal employees	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Rooms at home per 1000 adults who are not municipal employees	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
Rooms at home per 1000 adults: municipal workers minus other adults	Household Census	Municipality (1997 boundaries, which were those used for the 2000 census)	Obtained informally	Household Census
News on fraud possibly related to a mayor	Brazilian News Agency (Agência Brasil)	Municipality (1997 boundaries)	www.agenciabrasil.gov.br/noticias	Brazilian News Agency (Agência Brasil)
Federal police operation involving mayor	Federal Police Department (Departamento de Polícia Federal)	Municipality (1997 boundaries)	see text	Federal Police Department (Departamento de Polícia Federal)
News on fraud possibly related to a mayor	Various news websites (google search)	Municipality (1997 boundaries)		
Average years of schooling amon people aged 25 and over	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Fraction illiterate among people aged 15 and over	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Residential capital per capita	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Percent of households with electric lighting	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Percent of households with toilets linked to main network	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)
Percent of households with water linked to main network	IPEA Data	AMC0070	http://www.ipeadata.gov.br/ipeaweb.dtl/	IPEA (based on household census)

Figure 1. Number of Oilfields Discovered by Decade




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

