Resource Abundance, Development, and Living Standards: Evidence from Oil Discoveries in Brazil

Francesco Caselli and Guy Michaels[1]

First draft: May 2008; This draft: January 2009

Introduction

Should communities that discover oil in their subsoil or off their coast, rejoice or mourn? After a long collection of seemingly miserable experiences in many countries, economists increasingly suspect that the simple answer (rejoice, of course) is too simple. Anecdotal evidence from Nigeria to Venezuela, from Congo to the Caucasus suggests that natural-resource abundance is associated with Dutch disease, extreme corruption and rent seeking, political instability and state failure, all of which lead to lower living standards, quite possibly even lower than they would have been without the resource endowment. The comparison between resource-rich African, Latin American and Middle-Eastern slow growers with resource-poor East Asian fast growers also contributes to justify a suspicious attitude towards resource endowments.

A number of authors have tried to move beyond these casual observations and provide more systematic statistical evidence. The near totality of this work focuses on inter-country comparisons. Using cross-sectional or panel data it typically presents regressions of per-capita growth on proxies for resource abundance. It is easy to see that this approach suffers from potentially severe problems. Different countries are characterized by wildly different institutional and cultural features, which may well correlate with resource abundance. For example, the Middle East is oil rich, but it differs from the rest of the World in many others ways as well (and did so even before oil was discovered), only some of which can be controlled for with available data. This makes inference very tricky. Another, perhaps even more severe problem with cross-country exercises is that they tend to measure resource abundance by flows of natural-resource revenues (often normalized by GDP, or total exports). But this is clearly an outcome variable, making inference once again difficult. Data quality is also clearly a serious problem in cross-country work, as is the fact that intrinsic limitations in the available variables prevent these studies from clearly identifying the specific mechanisms through which resources affect outcomes. By and large, these limitations are reflected in a certain lack of consensus coming out of this literature.[2] The only exception that we are aware of is Michaels (2008), who uses variation across US counties (therefore holding most institutional variables constant) and actual endowments (therefore an exogenous variable and not an outcome). No similar study has been

[1] LSE, CEPR, and NBER (Caselli, f.caselli@lse.ac.uk) and LSE and CEPR (Michaels, g.michaels@lse.ac.uk). We are grateful to Steve Pischke, Silvana Tenreyro, Adrian Wood and Alwyn Young and seminar participants at LSE, Oxford, Sussex, and Zurich and at the OxCarre 2008 Conference for helpful comments, and to Gabriela Domingues, Renata Narita, and Gunes Asik-Altintas for research assistance. Caselli also thanks ESRC for financial support.

[2] The "classic" cross-country study is Sachs and Warner (1997). Other contributions in this vein include Isham, Woolcock, Pritchett and Busby (2005), Kolstad (2007), Collier and Goderis (2007), and Brunnschweiler and Bulte (2008).

carried out for a developing country, where many of the more sinister political mechanisms that appear to affect resource-abundant communities are more likely to unfold.[3]

In this paper we develop a new dataset on oil abundance for Brazilian municipalities – local geographical units roughly corresponding to US counties, and somewhere between European cities and provinces.[4] Brazil of course has not, in the aggregate, historically been a major oil producer: oil accounts for little more than 2% 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 revenues per capita), oil royalties in 2000 accounted for almost a third (~31%) of municipal revenues (and royalties are not the only oil-related source of revenue for oil-rich municipalities).

Because the dataset exploits variation within a country, many of the institutional, cultural, and policy variables that confound the relationship between resources and macroeconomic outcomes at the country level are held constant, thereby greatly diminishing the potential for spurious correlations arising from the omission of hard-to-measure correlates, and enhancing our ability to make inference. In addition, we believe that, unlike much of the cross-country literature, we can make plausible claims of exogeneity for our measure of resource abundance, which is the output from oil fields located in, or off the coast of, each municipality. First, we can 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.[5] 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.

---

[3] A few papers come quite close, though. 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 a region's suitability for coffee production. Vicente (2008) compares changes in perceived corruption in Sao Tome (which recently found oil) with Cape Verde (which didn't). He argues (fairly plausibly) that the two island countries share similar histories, culture, and political institutions. He finds large increases in corruption following the oil discovery.

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

[5] As mentioned above, this is unlikely to be true at the national level. 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.

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. We find that 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 exploitation. However, for onshore oil, we also uncover some (modest) composition effects. While non-oil GDP in the industrial sector shrinks somewhat, services GDP expands. This is qualitatively consistent with (certain versions of) the Dutch disease mechanism, though quantitatively the effect is modest. For offshore oil, we do not observe any 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 municipality governance.

Next, we regress municipal revenues on contemporaneous oil output. We find that 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. The effects of oil on municipal revenues are similar for onshore and offshore oil. These large fiscal windfalls make it possible for us to investigate the effects of oil operating through the government.

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 welfare payments.[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 spending, one would expect sizable improvements in welfare-relevant outcomes for the local population. They main puzzle we uncover is that, to the contrary, the local population seems to experience very small, if any, benefit. In particular, despite the fact that much of the oil windfall is apparently spent on construction and infrastructure, the size and quality of housing for the general population are at best unchanged, and for some outcomes they actually get worse. Perhaps more damnably, we find that despite the claim of government to be spending more on welfare, households in oil-abundant municipalities report that if anything they

---

[6] Here, and in most of the exercises discussed subsequently, we run regressions both in levels, e.g. spending on contemporaneous revenues, 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 hundreds of a Real.

are receiving less welfare income. We do see some real increases in purchases of education and health inputs, although it can be argued that the observed increases are small compared with the reported increase in spending in these areas.

Having failed to find the expected sizable increase in the real provision of goods and services to match the reported increases in spending, we look at the possibility that the fiscal windfall exerts other beneficial effects not captured by our measures of public service and infrastructure. In particular, we look at effects on household income. We find no effects of increased spending by the local government on household income. 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.

Overall, then, it seems that the local population has little to rejoice with when oil is discovered in its subsoil. There are few if any positive economic spillovers from oil extraction to the non-oil economy. And the substantial increases in government revenues through the payment of royalties do not seem to translate into significant improvements in welfare. On the other hand, the evidence also does not conform with the most extreme “horror stories” suggested by the anecdotal evidence. In particular, there is no evidence that living standards fall. This does not mean that all is well. That increased government revenues (and increased reported government spending) does not translate into measurable improvements in the quantity and quality of services received by the population, even in those areas where the budget has expanded the most, points at best to severe waste in the local government, and at worst to rampant corruption.

A full assessment of which combination of the two is driving the results is left for future research. However, we do find both statistical and anecdotal evidence that points to at least some role for corruption and rent seeking. One such piece of evidence is that, unlike the general population, municipal employees do seem to be experiencing significant increases in the size of their houses following oil discoveries. In addition, we conducted a broad search for news stories on oil rich municipalities and we found many instances of alleged corruption, as we detail below.

As discussed above, we believe that relative to other studies our paper benefits through better identification of causal effects, better identification of channels, and richer and higher quality data on outcomes. But these benefits come at some cost. In particular, while we can identify the differential effects for oil-abundant municipalities relative to oil-scarce ones, we cannot identify the overall impact of oil on Brazil as a whole. However, recalling that oil accounts for roughly 2% of national GDP, we may perhaps conjecture that any aggregate effect would be rather small compared to the localized effect uncovered by our analysis. More importantly, our results may not readily generalize to other countries with different institutional structures. Indeed, Michaels (2008) finds substantial local benefits from oil in the US South, at least for a considerable period of time. If we take the US to have a good institutional set-up, and Brazil to be about average in the quality of its institutions, this suggests that in economies with even worse-performing institutions locals may see even fewer benefits, or perhaps even negative outcomes.

In the paper we discuss at some length the implications of our findings for (selected) theories of the effect of resource abundance. Among the market-based mechanisms in the literature, we find little support for the more standard ones, such as those where the resource sector drains productive factors from the non-resource economy, or those where a wealth effect reduces overall labor supply. Instead, at least in the case of onshore oil, we find some evidence for a "localized" version of Dutch disease where the services sector expands at the expense of the manufacturing sector in response to the increased demand for local services by oil workers and operations.

Among the political-economy hypotheses offered by the literature we find no evidence for rent-seeking mechanisms where effort is reallocated from productive endeavors to competition for the oil windfall, or for models where the planning horizon of the local political elite is significantly affected by the oil riches. Models of patronage also find little support. Most of all, the data seems consistent with an old conjecture in the political-science literature, according to which resource income is easier to divert for the personal uses of the political elite than incomes from taxation.

Our paper is also clearly of specific interest to students of local governance in Brazil. Particularly strongly related to ours is a study of the effects of windfall transfers from the central government by Litschig (2008). Using a regression-discontinuity design, he finds that windfall federal transfers translate one-for-one in increased spending, with no crowding-out of other sources of revenue. This is very similar to our findings for windfall oil revenues. He also finds that this increased in spending leads to substantial measurable gains in schooling and literacy. Recall that we also find that oil windfalls lead to some gains in spending (reported and real) on education, but not in other areas of the budget, where increases in spending do not appear to be matched by a real expansion ion the provision of services.[8]

Our findings may imply that oil-rich Brazilian municipalities should be given special consideration in the current trend towards greater decentralization [Lipscomb and Mobarak 2007)] and in the design of audit schemes aimed at curbing corruption [Ferraz and Finan (2008a, 2008b)]. This special focus may become even more important as the size of oil revenues and royalties flowing to oil-rich municipalities is bound to increase dramatically following the recent discovery of huge new offshore fields.[9]

The paper may also be relevant to the growing emphasis placed on *transparency* by international donors in their dealings with poor, resource-abundant countries. In particular it is increasingly common for conditionality-based programs to feature stringent reporting requirements, both on the part of multinational oil companies and recipient governments. Our results may be

---

[8] There is 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)].

[9] Indeed, the issue is clearly of political relevance. De Oliveira Cruz and Ribeiro (2008) list 9 major pending Federal legislative proposals to reform the royalty system, all submitted in 2008. Interestingly, most proposals tend to reduce the share of royalties going to local governments, as well as to reduce discretionality in the use of royalty revenues. See also Afonso and Gobetti (2008).

interpreted as suggestive that *accounting* transparency per se may be insufficient, and that reporting schemes should document the actual efficient disbursement of sums, and not merely their recording on balance sheets.

Theoretical Mechanisms

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. Here we present a brief review. We also discuss the extent to which these theories, usually developed to explain country-wide developments, can apply to the local context that is our focus. In some of the cases the applicability is indeed quite limited, which is of course a limitation of our analysis.

Market-based mechanisms are broadly referred to as instances of Dutch disease. There are a number of versions of the Dutch disease model, but a frequent ingredient is the reallocation of productive factors across sectors induced by the resource windfall. One of the reasons for this reallocation is the direct effect of increased demand for capital and labor by the resource sector, which drains the non-resource part of the economy. This direct factor-market channel is unlikely to operate at the level of Brazilian municipalities. Investment in oilfield development in Brazil is carried out almost exclusively by Petrobras, a global hydrocarbon giant with access to a global capital market: it is entirely implausible that its demand for capital will affect the AMC-level supply of capital. Also, oil production is inherently extremely capital intensive, and the relatively few workers required to operate oilfields tend to be highly specialized and, again, participate in a market that is at the very minimum national, if not worldwide. Hence, we also don't expect oilfields to directly draw significantly from the local labor pool.

This however leaves other, indirect mechanisms that operate through the goods market. The classic mechanism 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). This mechanism is sometimes fueled by (and contributes itself to fueling) an exchange rate appreciation that further causes the non-resource export sector to shrink. A further effect 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 (and other taxes) 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.

Another economic mechanism that may change the composition of non-oil GDP in our setting is the direct impact on relative demand by oil firms and oil workers. 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. In the absence of migration flows to fulfil this demand (and we will see below that such migration has not materialized), this would lead us to expect onshore oil to shift the composition of non-oil GDP away from industry and towards services.[10] 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.[11]

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 or between services to oil workers and firms and other activities. However, our data allow us to distinguish between the effect of onshore oil-field 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. Offshore operations are miles off the coast and are less directly linked to 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. The wealth effect should also be associated with a decline in labour supply.

The theoretical literature has also identified several political-economy mechanisms that may be triggered by resource riches. In the classic mechanism royalty windfall 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 and their allies a larger share of the income flowing into the municipality's budget [e.g. Tornell and Lane (1999) , Halvor, Moene and Torvik (2006a, 2006b)]. The rent-seeking effect interacts and complements a possible *patronage* effect: flushed with oil royalties local politicians may be both better able and more motivated to create semi-fictitious government jobs to reward political supporters [e. g. Robinson, Torvik, and Verdier (2002)]. Both these mechanisms imply an effective decline in labor supply and hence value added in the non-oil sector, and particularly in the private economy (while government employment may expand), in that human energy is tuned away from production and into rent seeking and/or disguised unemployment. We may also observe heightened political competition, especially in ethnically heterogeneous communities [Caselli and Coleman (2008)].[12]

---

[10] The net effect on aggregate municipality 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. We discussed this in greater detail below.

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

[12] Again, some of the mechanisms that have been suggested at the national level would not apply here. For example, various authors have linked resource abundance to civil war, which is obviously unlikely to arise at the municipality level.

More recently it has been noted that resource windfalls also have the potential to reduce the supply of government-provided infrastructure and other productive public services. Since with larger resource windfalls controlling the government becomes more attractive, political incumbents face more aggressive *challenges* for their position. This may reduce their planning horizon. If government services and infrastructure provide public and/or private benefits only with some lag, incumbents may be induced to reduce them [Caselli (2007)]. Also, faced with more aggressive political competition, incumbents may be forced to devote more time and energy to politics, and less to policy, again leading to a decline in publicly-provided productive inputs [Caselli and Cunningham (2008]. These mechanisms may therefore result in a fall in non-oil GDP.

Another issue especially emphasized in the political-science literature is increased *corruption*. In this view royalties on natural resources are more easily stolen by public officers 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.

The mechanisms reviewed so far tend to be adverse, or at best neutral, in terms of outcomes and welfare. This reflects the emphasis in the existing literature on finding possible causes for a conjectured (though never fully convincingly established empirically) "resource curse," according to which resource-abundant countries are outperformed by resource-scarce ones. Needless to say, it is also easy to conceive of channels through which greater resource abundance may lead to higher non-oil GDP. In theory, there may be productive spillovers, in which the oil sector transfers technology and ideas to the non-oil sector. Also the increased government revenues may be turned into productive purposes. For starters, they of course relax the government budget constraint, so if the government is benevolent they ought to increase, not reduce, the provision of productive government services. And even if the government is self-interested the increased need to respond politically to more aggressive challenges for power may induce the government to seek popularity by using public money more productively [see, again, Caselli (2007)].

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.

Another important fact is that offshore oil accounts today for the vast majority of oil 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 monopolists 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. 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.

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.

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 one third 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 or 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.

Because oil royalties are an important source of oil-related revenues for local governments, it would be tempting to use oil royalties as a right-hand-side variable when testing for some of the political-economy mechanisms associated with oil abundance. However, as we have just seen, some of the factors determining a municipality's share in the royalties are not purely geographic, implying that royalty income is potential endogenous to other municipality-level outcomes we are interested in. 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. Instead, to gauge the effect of oil-related revenue on municipality-level outcomes, we instrument total municipality revenue with municipal-level oil production, as we detailed below.

One last consideration on the royalties is in order. While the total amount of royalties received by oil municipalities may be endogenous to some municipality characteristic, it is important to keep in mind that there is no scope for municipalities to influence the amount of royalties they receive through lobbying or negotiation. As we have seen, royalties are allocated through a complicated but rigid formula set in federal law.[13]

Data

Municipality-level Oil Output

A key variable in our empirical work is a municipality-level measure of the value of oil extracted in that municipality. 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.[14]

[13] One 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).

[14] The reference price is 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

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 BDEP website provides Geographic Information System (GIS) maps showing the location of 339 oil and gas fields. From this database we selected 273 fields that had passed the stages of development and were already producing.[15] We then combined the map of field locations with another GIS maps showing the 1997 boundaries of Brazilian municipalities from the Instituto Brazileiro de Geografia e Estatistica (IBGE).[16] This allowed us to establish the geographical relationship between the various oil fields and the various municipalities.

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 that oilfield, 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 even more challenging in the case of offshore oilfields, which lie entirely outside the boundaries of any specific municipality. Fortunately for us, the authorities had to solve the same problem. As mentioned, Petrobras pays royalties 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" (ANP 2008) 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.[17] We refer again to the Appendix for a more detailed discussion.

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. Of the 5507 municipalities that existed in Brazil in 1997, 124 have a stake in at least one oilfield. Using the BDEP data we also determine when oil

---

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.

[15] The BDEP map database is at http://maps.bdep.gov.br/website/maps/viewer.htm.

[16] The IBGE map database is at http://mapas.ibge.gov.br/divisao/viewer.htm.

[17] We have done some unsystematic checks to make sure that the Percentuais Médios de Confrontação from the ANP document 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.

was first discovered in each oil-endowed municipality, and whether it has onshore or offshore oilfields (or both).

Other Municipality-Level Variables

The amount of oil royalties received by each municipality in each year is readily available from ANP.

In addition to calculating oil-related variables for each municipality, we use the Global GIS DVD (USGS 2003) to create an indicator for whether each municipality is adjacent to the coast.

From Municipalities to AMCs

Having constructed a dataset on oil production and royalties for municipalities, we use a crosswalk from **Instituto de Pesquisa Econômica Aplicada (IPEA)** to collapse the data to 1997-1970 áreas mínimas comparáveis (AMCs). AMCs are statistical units created by IPEA to allow a consistent analysis of municipality data over time. Each AMC contains one municipality (or more) such that the area of each AMC remains quite stable even when municipality boundaries change. Altogether, the 5507 municipalities that existed in 1997 are pooled into 3659 AMCs for 1970-1997.[18] We designate an AMC as endowed with oil if at least one of the municipalities it is made up of contains oil. The revenues of oil production in each AMC are the sum of those of its municipalities. And finally, an AMC is designated as coastal if at least one of its municipalities is coastal.

We focus on the period since 1970 for three sets of reasons. First, going back before 1970 would further reduce the number of AMCs due to boundary changes during the 1960s. Second, most of our outcome variables are available from 1970 onwards. Third, most oil discoveries (including those of most of the largest oilfields) were made after 1970, so not much is lost by not presenting results for the pre-1970 period.

Of the 3659 AMCs, 103 are oil-endowed according to the formula outlined above. In 44 of these oil endowed AMCs, oil was first discovered in 1970 or earlier, while in 59 oil was first discovered after 1970.[19]

In our analysis we use a range of geographic, demographic, and economic variables calculated at the AMC level. Many of the economic variables were taken directly from IPEA data and are calculated in R$2000. Other variables we downloaded were denominated in nominal R$, and we converted them to R$2000 using a CPI index from IPEA data.[20] We normalized many of the variables of interest by population. Up to the year 2000 population data comes from the Brazilian Censuses. To calculate population for years after 2000 we inflated the 2000 population from the

[18] Municipality boundaries appear to have been quite stable even after 1997, and the IPEA data website provides data for these AMCs from 1970-2005.
[19] Figure 2 shows the distribution of decades of discovery of onshore and offshore oilfields; only about a quarter of the oilfields were discovered in 1970 or earlier.
[20] The index we used is Índice Nacional de Preços ao Consumidor (INPC).

Census by IPEA's estimate of the percentage change in population residing in each AMC on 1 July of each year.[21]

The AMC data from IPEA also include geographic controls which we use throughout our empirical analysis (unless otherwise specified). These geographic controls include longitude, latitude, distance to the federal capital, and distance to the state capital. To these controls we add the coast indicator mentioned above and state fixed effects.

Summary Statistics

Figure 2 shows the geographic distribution of oil fields in Brazil, as well as its distribution among AMCs. It is apparent that oil is concentrated on the coast or offshore. It is also clear that there is a lot of variation among municipalities in whether or not they sit on (or inland from) oilfields.

Table 1 present some summary statistics from our dataset. In the first two columns, we divide our AMCs into two groups: those without oil deposits (whether onshore or offshore), of which there are 3556, and those with currently operating oilfields which were discovered after 1970, of which there are 59. Our reason for highlighting results for AMCs with operating oilfields discovered after 1970 is that, as we will see, our outcome data allow us to rule out that municipalities where oil was found after 1970 were systematically different in any outcome variables (after controlling for geography) from non-oil AMCs before discovery, but we cannot perform a similar check for pre-1970 oil AMCs. Nevertheless, we may as well note from now that all our results are robust to including the pre-1970 oil AMCs. In the third column we show data from all oil-abundant AMCs, which look very similar to the post-1970 oil AMCs.

The average size of oil and non-oil AMCs is very similar. Oil AMCs are on average 3 degrees of latitude (roughly 330 Km) to the North of non-oil ones. They are also five degrees of longitude (roughly 500 Km at Brazil's average longitude) to the East. This reflects the fact that many oil fields are offshore, as also indicated by the fact that 66% of oil AMCs are coastal (v. only 4% for the non-oil ones). Controls for latitude, longitude, and being on the coast are therefore crucial in our empirical analysis. There are also substantive differences in distances from federal and state capitals.

Oil AMCs were on average considerably more populous in 1970, probably because of the coastal location of many of them. Population growth after 1970 has also been slightly more rapid in oil AMCs (an average annual rate of 2.2 percent v. 2.0 in non-oil AMCs). We will see later that this difference disappears when geographic controls are included.

In 1970 oil AMCs were somewhat less productive than non-oil ones, with GDP per capita about 20% lower. However, as one would obviously expect, GDP growth per-capita in the oil-AMCs, at an annual average of 5.6 percent, has far outstripped growth in the non-oil ones, which was 3 percent, so that by the end of the period oil AMCs generated roughly 80 percent more in total

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

GDP per capita.[22] (To convert R$2000 in 2008 US dollars the appropriate conversion factor is roughly 1).

The next line reports our constructed measure of oil output for oil and non-oil AMCs, the latter being trivially 0. Next we show aggregate oil royalties. Recall from the previous discussion that municipalities may receive royalties based on criteria other than geographical proximity, and this may explain why some non-oil municipalities receive royalties. Nevertheless, almost all of the royalties go to oil AMCs.

Ultimately we are most interested in outcomes that are more directly welfare relevant, so in the next line we report household income per capita in 2000 in oil and non-oil AMCs. The most striking finding here is that, despite their 80 percent advantage in aggregate GDP per capita, oil AMCs in 2000 had 7 percent *lower* household income per capita than non-oil AMCs.

A substantial difference between aggregate output and household income at the AMC level is of course to be expected. That the difference is much larger for oil AMCs is also not surprising: oil output is clearly subject to much larger rates of taxation by the federal government, and a much larger share of the ownership of factors of production, both labour and capital, is likely to reside outside the oil-producing AMCs – particularly for offshore ones. Indeed the figures suggest that the vast majority of the extra output produced by the oil fields has been removed from the local economy.[23]

Loosely speaking, the municipalities in the “No oil” column can be thought of as our control group, while the municipalities in the two subsequent columns (municipalities with oil found after 1970, or all municipalities with oil) represent two alternative “treatment groups.” As we already discussed above, and as we further discuss below, there are reasons to believe that for some of the questions we are interested in one can achieve a particular clean identification by further distinguishing between two different “treatments:” onshore and offshore oil. Summary statistics for the two corresponding treatment groups are presented in the last two columns of Table 1 (there are 9 oil AMCs that have both, so they belong to neither control group). Offshore AMCs are richer, experience faster growth in per-capita output, and receive more royalties per capita. All our empirical exercises include the AMCs in the first column (control group) and the AMCs in one of the remaining four columns (treatment group).

Specification and Identification

[22] The reason for reporting GDP in 2002, instead of 2000 as for the other variables, will become apparent shortly.

[23] 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. One possible source of the omission of capital income is topcoding of income in the micro data, coupled with very high inequality in Brazilian society.

We present results from two sets of empirical models. The first set of results is generated by OLS estimation of the specification

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

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 AMC-level oil output, X_m is a set of AMC-level geographic controls, 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 are aggregate GDP, sectoral GDP, household income, poverty rates, 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 have to argue that Q is uncorrelated with the residual determinants in e.

The second set of results is from two-stage least square estimation of the following model

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

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. Note 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. We wish to interpret β as capturing the causal effect of a marginal oil-generated Real of revenue on the outcome W. For this interpretation to be legitimate, we need to argue that Q affects W only through its effect on R. This, of course, supposes that Q actually affects R, or in other words that we have a first stage regression for estimating (2). That this is indeed the case should come as little surprise given our previous discussion of royalties; we formally show that we have a strong first stage in the next section. But before we do that, we now discuss the plausibility of our assumption that oil affects outcomes only through its effect on revenues.

We begin by arguing that Q_{mt} is uncorrelated with e_{mt} in the estimates based on specification (1) above. The first step is to show that our outcomes of interest did not differ in oil-rich and oil-poor AMCs *before oil was discovered.* In other words, oil abundance is randomly assigned. It is clear from Table 1 that 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 oil is spuriously correlated with other covariates. But our claim is that oil is randomly 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. In the first column we run a panel regression of the following model

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

where Y is log-GDP per capita, and I_m is an indicator function that takes the value of 0 if the AMC does not have oil (i.e. it belongs to the first column of Table 1) and 1 if the AMC had positive oil output in 2000 but the oil was discovered before 1970 (i.e. the AMC belongs to the second column of Table 1). AMCs with oil discovered before 1970 are excluded. The time coverage is given by various dates from 1970 (earliest year for which we observe GDP) to 2005. 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 the oil-rich-in-2000 indicator is allowed to vary over time.[24]

It is quite clear that sizable systematic differences in log-GDP between oil and non-oil AMCs 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. Column 2, where the dependent variable is the level of GDP, tells essentially the same story.

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 2), 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.

As a further set of robustness checks, the remaining columns use oil revenues per capita in 2000 as a measure of oil abundance (i.e. we replace I_m by Q_m), 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.

The fact that our estimates for 1970 are almost universally not significant is an important element in our identification strategy. However there is one (marginally) significant estimate which may cause a slight discomfort. We will deal with this by also reporting the effects of oil on *changes* in outcomes relative to a base year, and not only the effects on the cross section of recent level outcomes. (This would be an appropriate approach even in a fully randomized experiment: controlling for relevant baseline characteristics improves efficiency).[25]

Further support for our claims of random assignment is provided in Appendix Table A1, where we repeat the specification in (3) (for t=1970 only and using Q_m order as independent variable to save space) for other dependent variables for which we have data from 1970 and on which we focus below: housing quality and education. Conditional on geography none of the outcomes

[24] This is essentially equivalent to running a series of cross-sections, one for each year.
[25] In some cases however we have to use 1991 as a baseline because of data limitations.

differed significantly between oil and non-oil AMCs before discovery (The only exception is electric lighting in 1970, for which the correlation with oil output in 2000 is negative and marginally significant).

The results so far establish that, conditional on geography, there were no systematic differences between oil- and non-oil AMCs before the beginning of oil discovery and exploitation. We believe that this in itself goes a long way provides sufficient support for our identifying assumption for model (1). However, in our empirical analysis, where the main explanatory variable is gross oil output, we also exploit variation 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 in that year. 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 ς_{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 ς_{mt}.

This is implausible. As already mentioned, 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. 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.[26] 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). In particular, when we estimate below the effect of oil output on AMC non-oil GDP we find essentially no effect. 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 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

[26] 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.

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.[27]
One last 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.

Oil Abundance and GDP

In Table 3 we look at the effects of oil abundance on the productive side of the local economy. We present results from panel regressions specified as in (1), where the main right-hand-side variable is the AMC's oil output, interacted with a year dummy to allow for time-varying coefficients, and the left-hand side is a measure of AMC GDP (aggregate or industry specific). The usual set of geographic dummy-year interactions is included. Because the results in the previous section indicate that the local effects of oil abundance only emerged in the late 90s and 2000s, we begin our sample period in 2000, or in the first year for which we have reliable data. 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 were the extraction took place. This mistake resulted in a vast overestimate of oil GDP at the AMC level, because it essentially amounted to using gross oil output as measure of oil GDP. Needless to say, the overestimate of oil GDP carried over to aggregate AMC GDP, which was thus also grossly overestimated. The year 2002 is the first year for which this mistake was removed.[28]

In interpreting the coefficients in Table 3 it is important to bear in mind 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 revenue 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 it turns out that at the *national* level the share of oil GDP in gross oil output is also fairly stable and around 0.4.[29] Under fairly standard assumptions average and marginal shares of GDP in gross output are the same, so to the extent that the

---

[27] Nevertheless, we include the onshore oil AMCs as a robustness check.
[28] This mismeasurement does not invalidate the falsification exercise we conducted in the previous section. The point of that exercise was to show that there were no systematic differences between treatment and control 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.
[29] 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/

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.[30]

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 aggregate GDP or industrial GDP. In columns 4 and 5 we show that these results are robust when AMCs 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 of offshore. While neither onshore nor offshore oil production are likely to draw directly from local factor markets,[31] 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 case of Dutch disease (though not necessarily malign).

Support for this hypothesis is found in the last four columns of Table 3. In offshore-only oil AMCs we found 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

---

[30] 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].

[31] Investment in oilfield development in Brazil over the sample period is carried out by a global hydrocarbon giant with access to a global capital market: it is entirely implausible that its demand for capital will affect the AMC-level supply of capital. Also, oil production is inherently extremely capital intensive, and the relatively few workers required to operate oilfields tend to be highly specialized and, again, participate in a market that it is at the very minimum national, if not worldwide. Hence, we also don't expect oilfields to directly draw significantly from the local labor pool.

offshore oil having no market impact on the local economy. On the other hand, in onshore-only oil AMCs the effect of oil on industrial value added is less than one-for-one, as the coefficient on gross oil output falls 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.

Oil Abundance and the Local Government Budget

Revenue side

Many of the possible channels through which oil abundance influences GDP (and other social and economic outcomes) flow through the government budget and the allocation of government spending. In this section we investigate the effect of fiscal oil windfalls both on the revenue and spending 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. Also, because of data limitations we focus on single cross-sections. The first two columns show the effect of oil production on current municipality revenue in 1991 and in 2000 (results using total revenues are very similar). The results in the first column show that oil-rich AMCs did not have appreciably larger government revenues in 1991. 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%. These reforms clearly resulted in a massive increase in royalties received by local governments after 1998.

In 2000, as a result, we find that one Real of gross oil output increases 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 AMCs. The effect is muted in the sample including onshore-only oil AMCs, where one Real of oil produced leads to just a 1.5 cent increase in government revenues. One shortcoming of the results in column 2 is that there are many missing values for municipality revenue in 2000. In column 3 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. This result is confirmed in column 4, where the dependent variable is the change in revenues from 1991 to 2000.[32]

[32] In column 4 we continue predicting 2000 municipal revenues for municipalities that did not report revenues that year using 2001 data. Similarly, we predict missing 1991 data on municipal revenues using 1992 data. This allows us to overcome most of the potential selection problem into reporting/non-reporting of municipal revenues.

In the last 2 columns 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. Evidently, oil production generates sources of income for municipalities over and above the royalties they receive from Petrobras.[33]

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.[34]

Spending side

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 various items on overall current revenues. The most important items are Education and Culture, on which municipalities report spending about a quarter 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. Overall, total reported spending accounts for about 95 cents of every Real of revenue.[35]

Clearly the OLS results are merely accounting – they describe the allocation of the average Real of revenues. But the direction of causality is somewhat unclear: municipalities that spend more (in general, or on specific items) may raise more revenue (through taxes or, more likely given the limited tax autonomy of municipalities, through more effective lobbying of the central government). Furthermore, the allocation of the average Real may not be the same as the allocation of the marginal Real arising from oil abundance. In order to both circumvent possible endogeneity issues, and to focus on the utilization of oil-related revenues, in the second and third

[33] 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 "partecipacao especial" is similar to the overall value of royalties. For example, in 2004 royalties amounted to R$5735, while the partecipacao 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)].

[34] This does not rule out the possibility of offsetting reductions in direct Federal transfers to families. For example, increases in welfare payments from the municipal government could conceivably be offset by reductions in payments from "Bolsa Escola" and "Bolsa Familia," that transferred Federal cash directly to the population. These programs, however, were launched in 2001 and 2003, so our estimates should be unaffected by them. As best as we can tell there were no significant direct Federal transfer programs to households in 1995-2000.

[35] 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.), urban services (garbage collection, etc.), and urban transport (buses, etc.). Legislation is mostly spending supporting the activities of the municipal assembly. Welfare is the sum of "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.

row of Table 5 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 (third row) because –as discussed above – the case for the validity of our instrument is particularly compelling in this case. But results using all oil AMCs as the treatment group give mostly identical results (second row).

Our IV results show that the largest beneficiary of the increase in government revenues is Housing and Urban Development, with about a quarter of the marginal "oil Real." Education falls into second place, with about 15 cents, and Health continues to receive about 10 cents. The overall effect on spending is also still about 95 cents for every Real.

Regressions in levels are easiest to interpret, but they do not necessarily fully control for initial municipality characteristics. In order to better control for initial characteristics, throughout the paper we also report regressions for differences over time in our outcome variables. For municipal spending, this is done in the bottom panel of Table 4, where we look at ten-year changes in spending as left-hand-side variable. Here, as nearly everywhere else in the paper, results using differenced outcomes are nearly identical to those using levels.[36]

The fact that increases in spending match virtually one-for-one increases in revenues implies that oil municipalities do not seem to engage in Alaska-style programs of rebating oil revenue to citizens by mailing them checks. Alternatively, politicians have private incentives to channel the money through the expenditure side of the budget rather than rebating it.[37]

Taken literally the results from Table 5 lend little support to theoretical constructs where government investments in productive public goods decline as a result of fiscal windfalls. The question is whether these increases in reported spending are genuine, and also whether they are productive rather than forms of disguised diversion or patronage.

While we have argued that the IV results in Table 5 uncover causal effect from extra revenue from oil to various spending items, an important limitation is that these results do not necessarily generalize to other exogenous sources of revenue. Since there are theoretical treatments of the effects of resource abundance where citizens exercise more or less monitoring on spending depending on the source of revenue, we should stress that non-oil income may be used differently from oil income. Indeed the OLS results strongly suggest that this is the case.

That oil revenues increase 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

[36] 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, recall that oil output, prices, and royalties increased considerably between the two periods, so using the 2000 oil output levels is probably a reasonable proxy for the change.

[37] Another potential use of oil revenues is to cut taxes. But, as already noted, there is very little taxation at the municipal level, so this is probably not a relevant option.

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, 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 an over-representation of oil AMCs among the biggest spenders on administration.

Oil Abundance and Public-Service Provision

Table 5 indicates that the biggest reported winner among budget items from oil-related government income is spending on housing and on urban services and infrastructure. We would naturally expect, therefore, to see significant improvements in this area in oil-rich AMCs. Table 6 looks at a variety of housing, urban service and infrastructure outcomes: housing quantity (rooms per person) and quality (fraction of population living in sub-standard housing), electricity, connection to water and sewage networks, and garbage collection. While the OLS results tend unsurprisingly to show a positive association between government revenue and these outcomes, the IV result 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!

Other “winners” from Table 5 were Education, Health, and Welfare. Table 7 looks for evidence that this reported spending increases did in fact materialize. Depending on the specification, variables associated with the provision of education services 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 with a 5-year lag.[38] It also leads to the eventual construction of two new classrooms. Unfortunately, these results are not robust to first-differencing.

To check on the reported expansion of the health budget we have looked at the supply of health infrastructure. As for education, the results are mixed, though in this case it is the first-difference regressions that look more favorable.

Finally, there is also no indication whatsoever of an increase in welfare income in the census. This last result is particularly troubling. While it is conceivable that our measures of public-service provision in the areas of housing, education, and health do not fully capture the range of possible uses of public funds devoted to these sectors; or that we have not allowed for sufficient lags between spending and outcomes; neither of these concerns can be relevant for welfare. Here we have a situation where the municipal government reports a direct transfer to households, and households report having received no such transfer. At least for welfare, therefore, the

[38] A potentially useful way to think about this number is to begin by noting that according to Table 5 a 1 million increase in municipal revenues leads to about 150 thousand in reported spending on education. We don’t have yet data on primary teacher salaries, but assuming that they are roughly in the order of twice per-capita GDP would suggest that the government could hire 15 extra teachers with the money it says it spends on education.

conclusion that there is “missing money” seems inescapable. (We don’t have base-year welfare-income numbers, so this result is only available for the cross-section).

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 actually materialize in services to the population. Particularly worrisome is that there is absolutely no evidence of improvement in the quality and quantity of the housing infrastructure, despite the fact that this item wins the lion’s share of the spending increases, and that households do not report receiving extra welfare payments, despite claims by the government that welfare transfers go up. 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.[39]

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, while the OLS coefficients are unsurprisingly positive, 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.

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 two quintiles of the income distribution. In the full sample the benefits extend to the third quintile. Since, as we already have seen, welfare receipts by households do not increase, the most likely explanation is that public employment is drawn predominately from the lower segments of the income distribution, and that there is an expansion in public employment (drawn from the unemployment pool) and/or an increase in public wages. 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

[39] We do not have a specific mechanism in mind in this section. We are merely asking whether oil windfalls raise living standards in ways not captured by the outcomes we have looked at so far. Since oil does not seem to affect aggregate non-oil GDP, and since oil GDP is produced by factors owned mostly by residents outside the local economy, we would not expect first-order effects on household income through increased payments to local productive factors. On the other hand, oil does significantly expand government reported spending, particularly on public services in areas such as housing, education, and health, as well as increased welfare payments. The greater reported spending on government services, if true, should translate into a combination of expanded employment in government jobs and higher hourly wages. The latter should unambiguously increase average household income. The former has ambiguous effects: if the expanded government employment is entirely drawn from the private (non-oil) economy overall household income need not increase. The expansion of welfare payments also has ambiguous effects: if the greater availability of welfare support discourages market employment once again household income needs not raise.

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 we find evidence that oil may have reduced poverty. The effect is minuscule, however. The coefficient estimate (for the full sample) 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). Taken together, these estimates suggest that poverty reductions due to oil were modest, at best.

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.

We can actually be somewhat more specific about the sources of the (small) increase in household income. As we have argued above, the possible candidates are increased employment, through an expansion of the government sector, as long as it is not accompanied by a corresponding shrinkage of employment in the private sector; increased public-sector wages; and increased welfare payments. We have already seen in Table 7 that welfare payments are unchanged. In Table 9 we look at the employment channel. In the first column we find that oil-generated government revenues cause an increase in municipal employment as a share of the population. In the second column we show that at least some of this expansion in municipal employment results in a net increase in total labor supply, so this is clearly a source of increased household income. Furthermore, in results that we do not report, we found that municipal wages do not increase so we tend to attribute the entire increase in household income to the expansion in public employment. Since this expansion appears to have been largely unproductive, this lends some support to theories in which resource-abundance fuels patronage.

The slight increase in labor supply reported in Table 9 is also important from another point of view. Recall that some of the Dutch-disease mechanisms studied in the literature involve a wealth effect that reduces labor supply. This does not seem to be at work in the Brazilian context.

The more important result of this section, however, is that the increase in household income is extremely small, as we have seen in Table 8. Hence, the evidence points to a large amount of "missing" money.

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. We do have one outcome variable that may speak to this issue, albeit very indirectly. Census data allow us to identify separately the quality of housing enjoyed by municipal employees. Table 10 reports the results. It seems clear that oil-related revenue increases the quality of housing for municipal workers – but not, as we already know, for everyone else. Whatever the mechanism, municipal workers seem to be able to obtain for themselves more spacious accommodations.

Anecdotal evidence gleaned from the Brazilian media also suggests a frequent link between abundant oil revenues and corruption. One of the highest profile corruption cases at the local level in the last few years is "Operação Telhado De Vidro" ("Operation Glass Ceiling") in the municipality of Campos de Goytacazes, one of the biggest recipients of oil royalties. In March 2008, a large number of local-government officers at the highest level have recently been accused of diverting up to R$250m. More generally, in an informal search using Google, we found stories about ongoing corruption allegations in 13 of the top 25 municipalities ranked by oil output.

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 almost one-for-one 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.

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 money does not come from tax income, or whether they have less accurate information on the amounts flowing to the government in the form of royalties, we cannot say with the available data.

References

Afonso, J. R. Rodrigues and S. Wulff Gobetti (2008): "Renda do Petroleo no Brasil: Alguns Aspectos Fiscais e Federativos," *Revista do Bindes*, 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): *Analise do impacto social de receitas provenientes de royalties do petroleo em municipios do estado da Bahia*. Master's thesis, Universidade de Brasilia.

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.

Halvor, M., Moene, K.O. and Torvik, R. (2006a). 'Institutions and the Resource Curse', *The Economic Journal* 116(508): pp.1–20

Halvor, M., Moene, K.O. and Torvik, R. (2006b). 'Cursed by Resources or Institutions?' The World Economy, Blackwell Publishing, vol. 29(8), pp. 1117-1131, 08

IBGE (2008): "Produto Interno Bruto dos Municipios," 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)*.

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.

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 dod 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.

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.

Transparecy 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 see 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 e Sao 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 municipal ones. 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.[40] 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.

---

[40] 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 a1(m,f)% of field f lies inside the maritime borders of municipality m by the perpendicular-line criterion, and a2(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

$$(1/2) \times [a1(m,f) + a2(m,f)] \times 0.225 \times q(f),$$

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

$$(1/2) \times [a1(m,f) + a2(m,f)] \times q(f).$$

Table 1. Summary Statistics for Brazilian AMCs

	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
Area (square kilometers)	2,332	2,360	1,531	948	1,745
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
Population in 1970	24,960	39,086	42,466	47,869	36,492
Population in 2000	45,116	74,031	90,944	93,372	82,175
GDP per capita in 1970 (Brazilian R$2000)	1,756	1,478	1,889	2,091	1,862
GDP per capita in 2002 (Brazilian R$2000)	4,313	6,288	6,813	8,058	6,305
Oil output in 2000 (Brazilian R$2000)	0	3,388	2,832	3,894	2,332
Oil royalties per capita in 2000 (Brazilian R$2000)	1	84	70	83	51
Household income per capita in 2000 (Brazilian R$2000)	2,111	1,983	1,753	2,658	1,274
Observations (AMCs)	3,556	59	103	31	63

Notes: The table reports summary statistics for Brazilian AMCs (each AMC includes one municipality or multiple municipalities). Throughout the table "oil" denotes both oil and natural gas.

Table 2. Effect of Oil Output Per Capita on GDP Per Capita

	Dependent variable:					
	ln(GDP per capita)	GDP per capita	GDP per capita	GDP per capita	GDP per capita	GDP per capita
	AMCs where oil was first discovered after 1970 (or no oil was discovered)			All AMCs	AMCs with offshore oil (or no oil) only	AMCs with onshore oil (or no oil) only
(Oil discovered after 1970) x (year = 1970)	0.030	141				
	(0.080)	(270)				
(Oil discovered after 1970) x (year = 1980)	0.033	186				
	(0.082)	(395)				
(Oil discovered after 1970) x (year = 1996)	0.122	641				
	(0.087)	(466)				
(Oil discovered after 1970) x (year = 2002)	0.502	2689				
	(0.084)	(659)				
(Oil discovered after 1970) x (year = 2005)	0.537	3316				
	(0.087)	(876)				
(Oil output per capita in 2000) x (year = 1970)			-0.025	-0.028	-0.034	0.003
			(0.018)	(0.021)	(0.020)	(0.041)
(Oil output per capita in 2000) x (year = 1980)			-0.011	-0.010	-0.015	0.037
			(0.025)	(0.026)	(0.026)	(0.064)
(Oil output per capita in 2000) x (year = 1996)			0.060	0.093	0.075	0.177
			(0.106)	(0.108)	(0.131)	(0.161)
(Oil output per capita in 2000) x (year = 2002)			0.548	0.518	0.552	0.454
			(0.035)	(0.048)	(0.026)	(0.149)
(Oil output per capita in 2000) x (year = 2005)			0.803	0.779	0.858	0.608
			(0.044)	(0.070)	(0.032)	(0.214)
AMCs	3,615	3,615	3,615	3,659	3,587	3,619
Observations	18,066	18,075	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). We use data for 1970, 1980, 1996, 2002, and 2005. 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. The calculation of GDP due to 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. Robust standard errors clustered by AMC are in parentheses.

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

	AMCs where oil was first discovered after 1970 (or no oil was discovered)			All AMCs		AMCs with offshore oil (or no oil) only		AMCs with onshore oil (or no oil) only	
	Dependent variable:			Dependent variable:		Dependent variable:		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	-	0.012	-	-0.009	-	0.119
			(0.016)		(0.020)		(0.013)		(0.049)
(Oil output per capita in 2001) x (year = 2001)	-	-	0.025	-	0.041	-	0.005	-	0.161
			(0.030)		(0.036)		(0.021)		(0.069)
(Oil output per capita in 2002) x (year = 2002)	0.400	0.381	0.019	0.356	0.025	0.375	0.012	0.297	0.091
	(0.032)	(0.012)	(0.022)	(0.025)	(0.024)	(0.008)	(0.019)	(0.117)	(0.046)
(Oil output per capita in 2003) x (year = 2003)	0.453	0.434	0.019	0.410	0.025	0.439	0.009	0.324	0.100
	(0.051)	(0.042)	(0.023)	(0.044)	(0.024)	(0.046)	(0.020)	(0.135)	(0.051)
(Oil output per capita in 2004) x (year = 2004)	0.364	0.354	0.010	0.335	0.020	0.354	0.002	0.284	0.098
	(0.033)	(0.015)	(0.021)	(0.032)	(0.022)	(0.015)	(0.021)	(0.127)	(0.045)
(Oil output per capita in 2005) x (year = 2005)	0.449	0.447	0.002	0.423	0.009	0.451	-0.002	0.311	0.084
	(0.026)	(0.016)	(0.015)	(0.037)	(0.016)	(0.015)	(0.015)	(0.144)	(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). Throughout the table "oil" denotes both oil and natural gas. 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. Robust standard errors clustered by AMC are in parentheses.

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

	Dependent variable: per capita municipality expenditures on					
	(1)	(2)	(3)	(4)	(5)	(6)
	Current municipality revenues per capita in 1991	Current municipality revenues per capita in 2000	Current municipality revenues per capita in 2000 (see footnote)	Change in current municipality revenues per capita from 1991-2000 (see footnote)	Royalties from oil in 2000 (only AMCs in column 3)	Royalties from oil in 2000 (no restriction on AMCs)
A. Only AMCs where oil was first discovered after 1970						
Oil outputs per capita in 2000	0.0016	0.0277	0.0280	0.0278	0.0195	0.0196
	(0.0008)	(0.0060)	(0.0058)	(0.0048)	(0.0024)	(0.0024)
Observations (AMCs)	3,502	3,208	3,512	3,493	3,512	3,615
B. All AMCs						
Oil output per capita in 2000	0.0024	0.0279	0.0308	0.0299	0.0189	0.0189
	(0.0008)	(0.0057)	(0.0048)	(0.0041)	(0.0021)	(0.0021)
Observations (AMCs)	3,546	3,242	3,553	3,534	3,553	3,659
C. Same as B., but using AMCs with offshore oil (or no oil) only						
Oil output per capita in 2000	0.0019	0.0292	0.0295	0.0277	0.0181	0.0181
	(0.0008)	(0.0059)	(0.0058)	(0.0051)	(0.0020)	(0.0020)
Observations (AMCs)	3,475	3,183	3,484	3,466	3,484	3,587
D. Same as B., but using AMCs with onshore oil (or no oil) only						
Oil output per capita in 2000	0.0047	0.0160	0.0325	0.0340	0.0168	0.0168
	(0.0017)	(0.0073)	(0.0082)	(0.0065)	(0.0028)	(0.0028)
Observations (AMCs)	3,507	3,205	3,514	3,495	3,514	3,619

Notes: Each column reports coefficients from a regression using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). Throughout the table "oil" denotes both oil and natural gas. All values are in Brazilian R$2000. Since we only have municipal revenues for about 90 percent of the AMCs in 2000, columns 3 predict 2000 municipal revenues from 2001 municipal revenues using a linear regression. The dependent variable in column 4 is the change in municipal revenues from 1991-2000, where we use 2001 municipal revenues to predict missing 2000 values and 1992 municipal revenues to predict missing 1991 values. Column 5 restricts the sample to the AMCs in column 3. 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

	Dependent variable: per capita municipality current expenditures in 2000, by category					
	Education and culture	Health and sanitation	Housing and urban develop-ment	Legislation	Welfare	Total
A. OLS using all AMCs						
Per capita municipal revenues in 2000	0.284 (0.010)	0.127 (0.006)	0.109 (0.008)	0.042 (0.002)	0.068 (0.005)	0.937 (0.024)
Observations (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553
B. IV using all AMCs						
Per capita municipal revenues in 2000	0.158 (0.029)	0.096 (0.015)	0.256 (0.062)	0.042 (0.009)	0.046 (0.013)	0.957 (0.114)
Observations (AMCs)	3,553	3,553	3,553	3,553	3,553	3,553
C. IV, using AMCs with offshore oil (or no oil) only						
Per capita municipal revenues in 2000	0.139 (0.022)	0.106 (0.016)	0.194 (0.027)	0.049 (0.007)	0.050 (0.008)	0.866 (0.093)
Observations (AMCs)	3,484	3,484	3,484	3,484	3,484	3,484

	Dependent variable: change in per capita municipality current expenditures from 1991-2000, by category					
	Education and culture	Health and sanitation	Housing and urban develop-ment	Legislation	Welfare	Total
D. OLS using all AMCs						
Change in per capita municipal revenues from 1991-2000	0.281 (0.013)	0.122 (0.008)	0.098 (0.011)	0.037 (0.002)	0.071 (0.006)	0.890 (0.030)
Observations (AMCs)	3,423	3,423	3,423	3,423	3,423	3,423
E. IV using all AMCs						
Change in per capita municipal revenues from 1991-2000	0.161 (0.029)	0.085 (0.016)	0.246 (0.061)	0.039 (0.009)	0.043 (0.015)	0.938 (0.105)
Observations (AMCs)	3,423	3,423	3,423	3,423	3,423	3,423
F. IV, using AMCs with offshore oil (or no oil) only						
Change in per capita municipal revenues from 1991-2000	0.150 (0.024)	0.108 (0.018)	0.194 (0.033)	0.047 (0.007)	0.050 (0.008)	0.883 (0.106)
Observations (AMCs)	3,355	3,355	3,355	3,355	3,355	3,355

Notes: Each column reports coefficients from a regression using a cross section of AMCs (each AMC includes one municipality or more). Throughout the table "oil" denotes both oil and natural gas. 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. In IV regressions, the instrument is oil output per capita in 2000. 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 Quality

	Dependent variable:						
	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
A. Outcomes in 2000							
OLS using all AMCs							
Per capita municipal	0.159	-0.0004	0.0028	0.0029	0.0032	-0.0013	-0.0022
revenues in 2000	(0.035)	(0.0003)	(0.0005)	(0.0008)	(0.0008)	(0.0012)	(0.0016)
Obs.	3,553	3,553	3,553	3,553	3,553	3,553	3,553
IV using all AMCs							
Per capita municipal	0.079	-0.0097	0.0095	0.0097	0.0045	-0.0054	-0.0074
revenues in 2000	(0.123)	(0.0033)	(0.0059)	(0.0041)	(0.0063)	(0.0141)	(0.0090)
Obs.	3,553	3,553	3,553	3,553	3,553	3,553	3,553
IV, using AMCs with offshore oil (or no oil) only							
Per capita municipal	0.198	-0.0146	0.0010	0.0071	-0.0065	-0.0221	-0.0134
revenues in 2000	(0.171)	(0.0018)	(0.0017)	(0.0034)	(0.0039)	(0.0111)	(0.0094)
Obs.	3,484	3,484	3,484	3,484	3,484	3,484	3,484
B. Changes in outcomes from 1991-2000							
OLS using all AMCs							
ΔMunicipal revenues /	0.015	-0.00027	0.0012	0.0085	0.0026	0.0044	0.0017
capita from 1991-2000	(0.030)	(0.00014)	(0.0007)	(0.0019)	(0.0008)	(0.0011)	(0.0014)
Obs.	3,534	3,534	3,534	3,534	3,534	3,534	3,534
IV using all AMCs							
ΔMunicipal revenues /	-0.114	-0.0052	-0.0108	-0.0005	0.0005	-0.0068	-0.0174
capita from 1991-2000	(0.085)	(0.0012)	(0.0066)	(0.0093)	(0.0042)	(0.0085)	(0.0059)
Obs.	3,534	3,534	3,534	3,534	3,534	3,534	3,534
IV, using AMCs with offshore oil (or no oil) only							
ΔMunicipal revenues /	-0.131	-0.0069	0.0010	-0.0048	-0.0014	-0.0158	-0.0230
capita from 1991-2000	(0.093)	(0.0007)	(0.0021)	(0.0043)	(0.0022)	(0.0054)	(0.0063)
Obs.	3,466	3,466	3,466	3,466	3,466	3,466	3,466

Notes: The table reports coefficients from regressions using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). The data are from the 2000 and 1991 Census micro data. Throughout the table "oil" denotes both oil and natural gas. The regressor of interest is municipal revenues per capita (or changes in that variable from 1991-2000). The instrument is oil output per capita in 2000. All values are in Brazilian R$2000. For municipalities that did not report current municipal revenues in 2000, we predicted 2000 revenues using 2001 revenues. 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 & Welfare

	Dependent variable:						
	Municipal teachers per million people in 2000	Municipal classrooms per million people in 2000	Municipal teachers per million people in 2005	Municipal classrooms per million people in 2005	Health establishments with inpatient care per million people in 2002	Health establishments without inpatient care per million people in 2002	Welfare income per capita
OLS using all AMCs							
Per capita municipal revenues in 2000	5.1	2.9	5.2	3.3	0.052	0.287	-0.006
	(0.4)	(0.2)	(0.4)	(0.3)	(0.009)	(0.035)	(0.007)
Obs.	3550	3550	3553	3553	3553	3553	3553
IV using all AMCs							
Per capita municipal revenues in 2000	2.2	0.5	4.5	1.4	-0.006	-0.015	-0.002
	(1.0)	(0.5)	(1.0)	(0.6)	(0.018)	(0.077)	(0.003)
Obs.	3550	3550	3553	3553	3553	3553	3553
IV, using AMCs with offshore oil (or no oil) only							
Per capita municipal revenues in 2000	1.9	0.9	4.8	1.8	0.012	-0.021	-0.006
	(0.8)	(0.6)	(1.2)	(0.7)	(0.005)	(0.059)	(0.003)
Obs.	3481	3481	3484	3484	3484	3484	3484

	Dependent variable:					
	Δmunicipal teachers per million people from 1996-2000	Δmunicipal classrooms per million people from 1996-2000	Δmunicipal teachers per million people from 1996-2005	Δmunicipal classrooms per million people from 1996-2005	Δhealth establishments with inpatient care per million people 1992-2002	Δhealth establishments without inpatient care per million people 1992-2002
OLS using all AMCs						
ΔMunicipal revenues / capita from 1991-2000	3.2	1.4	3.5	1.8	0.003	0.146
	(0.4)	(0.2)	(0.5)	(0.3)	(0.012)	(0.032)
Obs.	3445	3445	3446	3446	3534	3534
IV using all AMCs						
ΔMunicipal revenues / capita from 1991-2000	0.5	0.3	2.8	1.3	0.028	0.110
	(1.2)	(0.7)	(1.4)	(1.0)	(0.018)	(0.099)
Obs.	3445	3445	3446	3446	3534	3534
IV, using AMCs with offshore oil (or no oil) only						
ΔMunicipal revenues / capita from 1991-2000	0.4	-0.2	3.4	0.8	0.012	0.183
	(1.4)	(0.5)	(1.9)	(0.7)	(0.006)	(0.074)
Observations (AMCs)	3377	3377	3378	3378	3466	3466

Notes: Each column reports coefficients from a regression using a cross section of AMCs (each AMC includes one municipality or more). Throughout the table "oil" denotes both oil and natural gas. All values are in Brazilian R$2000. IV regressions use oil output per capita in 2000 as an instrument for municipal revenues per capita in 2000. 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. All regressions control for latitude, longitude, coast dummy, distance to the state capital, distance to the federal capital, and

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

	Dependent variable:						
	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
A. Outcomes in 2000							
OLS using all AMCs							
Per capita municipal	0.25	0.07	0.11	0.19	0.31	0.54	-0.0028
revenues in 2000	(0.12)	(0.01)	(0.03)	(0.06)	(0.11)	(0.40)	(0.0006)
Obs.	3,553	3,553	3,553	3,553	3,553	3,553	3,553
IV using all AMCs							
Per capita municipal	0.17	0.12	0.18	0.21	0.19	0.18	-0.0107
revenues in 2000	(0.23)	(0.04)	(0.08)	(0.12)	(0.22)	(0.78)	(0.0048)
Obs.	3,553	3,553	3,553	3,553	3,553	3,553	3,553
IV, using AMCs with offshore oil (or no oil) only							
Per capita municipal	0.08	0.09	0.09	0.09	0.07	0.09	-0.0036
revenues in 2000	(0.31)	(0.04)	(0.09)	(0.15)	(0.31)	(1.03)	(0.0030)
Obs.	3,484	3,484	3,484	3,484	3,484	3,484	3,484
B. Changes in outcomes from 1991-2000							
OLS using all AMCs							
Per capita municipal	0.19	0.06	0.08	0.14	0.20	0.51	-0.0034
revenues in 2000	(0.08)	(0.02)	(0.02)	(0.04)	(0.07)	(0.30)	(0.0009)
Obs.	3,444	3,444	3,443	3,443	3,443	3,444	3,444
IV using all AMCs							
Per capita municipal	0.12	0.15	0.19	0.18	0.18	-0.15	-0.0062
revenues in 2000	(0.18)	(0.05)	(0.06)	(0.07)	(0.15)	(0.75)	(0.0028)
Obs.	3,444	3,444	3,443	3,443	3,443	3,444	3,444
IV, using AMCs with offshore oil (or no oil) only							
Per capita municipal	0.08	0.08	0.10	0.09	0.15	-0.07	-0.0040
revenues in 2000	(0.26)	(0.03)	(0.05)	(0.11)	(0.21)	(0.95)	(0.0031)
Obs.	3,376	3,376	3,375	3,375	3,375	3,376	3,376

Notes: The table reports coefficients from regressions using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). The data are from the 2000 and 1991 Census micro data. Throughout the table "oil" denotes both oil and natural gas. The regressor of interest is municipal revenues per capita (or changes in that variable from 1991-2000). The instrument is oil output per capita in 2000. All values are in Brazilian R$2000. For municipalities that did not report current municipal revenues in 2000, we predicted 2000 revenues using 2001 revenues. 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 Employment

	First stage (for IV regs only)	Dependent variable: Municipal employees per 1000 people	Dependent variable: Annual hours worked per capita
A. Outcomes in 2000			
OLS using all AMCs			
Per capita municipal revenues in 2000		0.000021	0.00041
		(0.000002)	(0.00019)
Obs.		3,550	3,553
IV using all AMCs			
Per capita municipal revenues in 2000	0.0308	0.000007	0.00071
	(0.0048)	(0.000004)	(0.00082)
Obs.	3,550	3,550	3,553
IV, using AMCs with offshore oil (or no oil) only			
Per capita municipal revenues in 2000	0.0295	0.000004	0.00126
	(0.0058)	(0.000003)	(0.00065)
Obs.	3,481	3,481	3,484
B. Changes in outcomes from 1991-2000			
OLS using all AMCs			
Per capita municipal revenues in 2000		0.000011	-0.00013
		(0.000002)	(0.00023)
Obs.		3,525	3,534
IV using all AMCs			
Per capita municipal revenues in 2000	0.0299	0.000006	-0.00002
	(0.0041)	(0.000003)	(0.00109)
Obs.	3,525	3,525	3,534
IV, using AMCs with offshore oil (or no oil) only			
Per capita municipal revenues in 2000	0.0277	0.000006	0.00091
	(0.0051)	(0.000002)	(0.00047)
Obs.	3,457	3,457	3,466

Notes: The table reports coefficients from regressions using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). The data are from the 2000 and 1991 Census micro data. Throughout the table "oil" denotes both oil and natural gas. The regressor of interest is municipal revenues per capita (or changes in that variable from 1991-2000). The instrument is oil output per capita in 2000. All values are in Brazilian R$2000. For our sample of adults we use all people aged 14-64. For municipalities that did not report current municipal revenues in 2000, we predicted 2000 revenues using 2001 revenues. 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 Municipal Revenues from Oil on House Size

	Dependent variable:		
	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
A. Outcomes in 2000			
OLS using all AMCs			
Per capita municipal revenues in 2000	0.07	0.15	-0.08
	(0.05)	(0.03)	(0.05)
Obs.	3,550	3,550	3,550
IV using all AMCs			
Per capita municipal revenues in 2000	0.10	0.07	0.03
	(0.24)	(0.12)	(0.23)
Obs.	3,550	3,550	3,550
IV, using AMCs with offshore oil (or no oil) only			
Per capita municipal revenues in 2000	0.51	0.18	0.33
	(0.22)	(0.17)	(0.14)
Obs.	3,481	3,481	3,481
B. Changes in outcomes from 1991-2000			
OLS using all AMCs			
Per capita municipal revenues in 2000	-0.04	0.01	-0.06
	(0.09)	(0.03)	(0.09)
Obs.	3,525	3,525	3,525
IV using all AMCs			
Per capita municipal revenues in 2000	0.41	-0.12	0.54
	(0.29)	(0.09)	(0.31)
Obs.	3,525	3,525	3,525
IV, using AMCs with offshore oil (or no oil) only			
Per capita municipal revenues in 2000	0.40	-0.15	0.54
	(0.18)	(0.09)	(0.17)
Obs.	3,457	3,457	3,457

Notes: The table reports coefficients from regressions using a cross section of AMCs (each AMC includes one municipality or multiple municipalities). The data are from the 2000 and 1991 Census micro data. Throughout the table "oil" denotes both oil and natural gas. The regressor of interest is municipal revenues per capita (or changes in that variable from 1991-2000). The instrument is oil output per capita in 2000. All values are in Brazilian R$2000. For our sample of adults we use all people aged 14-64. For municipalities that did not report current municipal revenues in 2000, we predicted 2000 revenues using 2001 revenues. 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 A1. Falsification: the Effect of Oil Output Per Capita on Social and Economic Outcomes

	Dependent variable:					
	Residential capital per capita	Fraction of households with electric lighting	Fraction of households with sanitary installations	Fraction of households with canalized water	Average years of schooling (see footnote)	Fraction illiterate among those aged 15 and over (see footnote)
Estimates for 1970						
Oil output per capita in 2000	-0.0000222	-0.0000034	-0.0000008	-0.0000013	-0.0000022	-0.0000545
	(0.0000140)	(0.0000017)	(0.0000016)	(0.0000012)	(0.0000060)	0.0002
t-stat	*-1.59*	*-2.01*	*-0.52*	*-1.11*	*-0.37*	*-0.28*
Observations (AMCs)	3,615	3,615	3,615	3,615	3,615	3,615

Notes: The table reports coefficients from regressions using a cross section 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 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

	Dependent variable: ln(population)	Dependent variable: population	Dependent variable: ln(population)	Dependent variable: population
(Oil discovered after 1970) x (year = 1970)	0.106	-20993		
	(0.148)	(20288)		
(Oil discovered after 1970) x (year = 1980)	0.144	-27542		
	(0.154)	(24677)		
(Oil discovered after 1970) x (year = 1991)	0.182	-32453		
	(0.153)	(27380)		
(Oil discovered after 1970) x (year = 1996)	0.216	-33977		
	(0.151)	(28077)		
(Oil discovered after 1970) x (year = 2000)	0.216	-37128		
	(0.150)	(29764)		
(Oil discovered after 1970) x (year = 2005)	0.209	-42197		
	(0.150)	(32458)		
(Oil output per capita in 2000) x (year = 1970)			-0.0000081	-2.887
			(0.0000162)	(2.652)
(Oil output per capita in 2000) x (year = 1980)			-0.0000074	-3.894
			(0.0000159)	(3.129)
(Oil output per capita in 2000) x (year = 1991)			-0.0000047	-4.296
			(0.0000150)	(3.409)
(Oil output per capita in 2000) x (year = 1996)			0.0000007	-4.232
			(0.0000145)	(3.457)
(Oil output per capita in 2000) x (year = 2000)			0.0000012	-4.470
			(0.0000140)	(3.632)
(Oil output per capita in 2000) x (year = 2005)			0.0000016	-4.804
			(0.0000131)	(3.905)
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.

Figure 1. Number of Oilfields Discovered by Decade

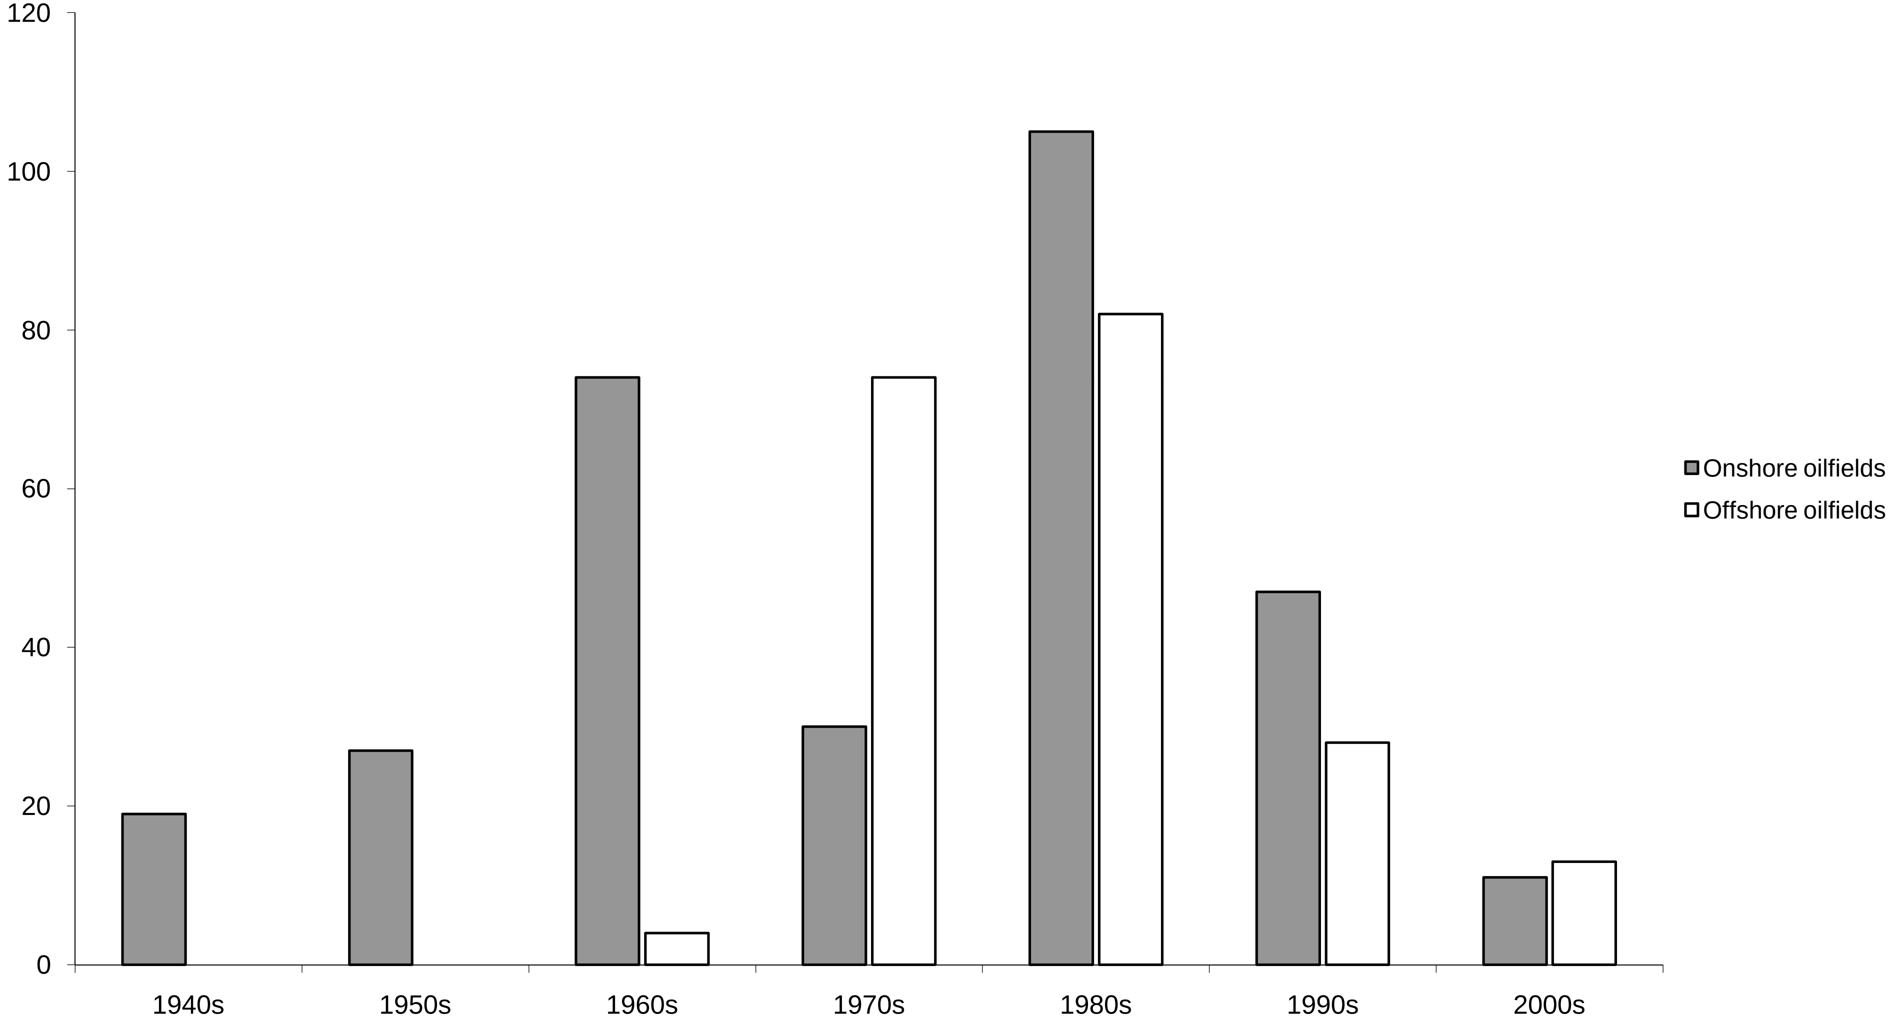

Figure 2. AMCs (from 1970) and Oilfields in Brazil