Three Essays on Public Finance and Development

by

Traviss Cassidy

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2018

Doctoral Committee:

Professor James R. Hines, Jr., Chair
Professor Dean Yang
Associate Professor C. Hoyt Bleakley, Jr.
Professor Joel Slemrod

Traviss Cassidy

tmcassid@umich.edu

ORCID iD: 0000-0002-1943-4613

Dedication

This dissertation is dedicated to Alecia Cassidy.

Acknowledgments

I am grateful for support and helpful comments from my dissertation committee, Jim Hines, Dean Yang, Joel Slemrod, and Hoyt Bleakley; and from Alecia Cassidy, Mark Dincecco, Ugo Troiano, and many others who are thanked in the working paper versions of these three chapters.

Contents

List of Tables

List of Figures

List of Appendices

Abstract

This dissertation examines the determinants and consequences of tax and expenditure policy, with an emphasis on the link between public finance and development.

Chapter 1 analyzes local government responses to permanent and transitory changes in grant revenue in Indonesia. Exploiting national policy reforms and variation in resource endowments, I estimate the causal effects of two unconditional grants: a general grant that experienced a permanent change, and shared oil and gas revenue that exhibited significant transitory variation. I find that the general grant induced a front-loaded expenditure response and increased the provision of lumpy public goods and services, such as public schools, health facilities, and health personnel. In contrast, transitory fluctuations in oil and gas revenue produced a balanced fiscal response and had little impact on lumpy public goods and services. The two grants had similar effects on road quality, which depends on less lumpy maintenance expenditure. Together the results suggest that local governments respond to changes in permanent public income over a time horizon of three to five years. The results imply that the permanence of a grant reform could matter for both efficiency and countercyclical fiscal policy in a federation. Furthermore, they highlight the challenge to quantifying the accountability effects of local taxation when taxes and grants are subject to different types of shocks.

Chapter 2 estimates the long-run effects of oil wealth on development by exploiting spatial variation in sedimentary basins—areas where petroleum can potentially form. Instrumental variables estimates indicate that oil production impedes democracy and fiscal capacity development, increases corruption, and raises GDP per capita. Oil production may also increase internal armed conflict, however this point estimate is less precise. In many specifications failure to account for endogeneity leads to substantial underestimation of the adverse effects of oil, suggesting that countries with higher-quality political institutions and greater fiscal capacity disproportionately select into oil production. Countries that had weak executive constraints from 1950–1965 experienced the largest adverse effects of oil on democracy and fiscal capacity, yet they benefited the most in terms of income. The results confirm the existence of a political resource curse, while rejecting the economic resource curse hypothesis.

Chapter 3, coauthored with Alberto Alesina and Ugo Troiano, considers the role of a politician's age in Italian municipal governments. When the term limit is not binding, younger mayors engage more often in political budget cycles than older mayors. Thus younger politicians behave more strategically in response to electoral incentives, probably because they expect to have a longer political career and stronger career concerns. We discuss and rule out several alternative interpretations.

Chapter 1

Local Government Responses to Permanent vs. Transitory Grant Revenue: Evidence from Indonesia

1.1 Introduction

Intergovernmental grants are an important policy tool of federal systems that can address cross-jurisdictional externalities and promote horizontal equity (Oates, 1972). Understanding how subnational governments respond to intergovernmental grants informs the optimal design of fiscal policy in a federation and illuminates aspects of local government decision-making. However, little is known about how the dynamic properties of intergovernmental grants influence local public-good responses. This is a policy-relevant concern, because some of the most popular types of grants around the world vary markedly in their time-series properties. Many unconditional grants, such as equalizing grants, are allocated according to formulas which depend on slow-moving variables, such as population, land area, poverty and cost indices, and local tax bases. In the absence of major changes to the formula, current revenue approximates the long-run, or "permanent," revenue expected from the grant. Other popular unconditional grants, such as shared natural resource revenue, are functions of volatile variables, producing significant transitory variation in grant revenue. Do local governments respond differently, on a per-dollar basis, to permanent and transitory grant revenue?

I address this question in the context of Indonesia. A simple theoretical model guides the analysis. Consider a local government that provides durable and nondurable public goods financed by intergovernmental grants. Investment in durable public goods, such as schools, is "lumpy": additions to the current stock must exceed a minimum size threshold. Starting from an initial, optimal provision of public goods, the local government responds to two different types of grant-revenue shocks: (1) a permanent increase of one dollar and (2) a transitory (one-time) increase of one dollar. A myopic government, for which spending is a function of current-year revenue only, responds identically to the two shocks on a per-dollar basis. This is because it perceives the permanent shock to be a series of transitory shocks. In contrast, a non-myopic government conditions spending on current and future revenue over some time horizon. Because the present discounted value of the permanent revenue shock exceeds that of the transitory revenue shock, the former is more likely to lead a non-myopic government to overcome the lumpiness constraint and increase both the durable and nondurable public goods. A binding lumpiness constraint causes the government to adjust only the nondurable public good in response to the transitory revenue shock. Because durable goods require an upfront investment, the permanent increase in grant revenue induces a spending response that is front-loaded—spending increases by more than one dollar initially and less than one dollar subsequently. The spending response to the transitory increase in grant revenue is of equal size, and less than one dollar, across the years of the time horizon. Thus the permanence of a shock to grant revenue influences both the timing and composition of the fiscal response.

I test the predictions of the model using data on district governments in Indonesia following fiscal decentralization in 2001. I evaluate the effects of the country's two largest intergovernmental grants. The first is a general grant designed to equalize fiscal resources across regions. This grant has stable disbursements over time, with the exception of a permanent change in 2006, the magnitude of which varied according to district land area per capita. The second is shared revenue from local oil and gas production, which exhibits significant transitory variation as resource prices fluctuate. Both grants are unconditional, non-matching, and subject to the same level of central-government oversight. Hence, their primary distinguishing feature is their time-series variation. I exploit national policy reforms and variation in resource endowments to estimate the causal effects of the two grants on local public spending and provision of public goods and services. Consistent with the theory, the permanent reform to the general grant produced a front-loaded expenditure response and led to greater provision of lumpy public goods and services, such as public schools, health facilities, and health personnel. In contrast, transitory fluctuations in oil and gas revenue produced more balanced fiscal responses and had only modest impacts on lumpy public goods and services. The two grants had similar effects on road quality, which depends on less lumpy maintenance expenditure, suggesting that lumpiness, rather than graft, drives the differential responses of structures and personnel to the two grants.

The empirical results imply that local governments respond to changes in permanent public income over a time horizon of three to five years. This conclusion contrasts sharply with the commonly held view that local governments are myopic. A recent policy report expresses the concern that volatile grant revenue perpetuates a local resource curse:

> *Large and unpredictable transfers of natural resource revenues can destabilize a local economy. Cycles of boom and bust also harm economic growth, as governments are likely to spend on ostentatious projects during booms and not plan appropriately for downturns* (Natural Resource Governance Institute, 2016, p. 11).

The results of this paper show that local revenue volatility need not cause local expenditure volatility, contrary to conventional wisdom and standard assumptions made in the resource-curse literature (Cust and Viale, 2016; van der Ploeg, 2011).

Indonesia provides an ideal policy environment to study these questions. First, there are a large number of local governments—over 300—with spending authority similar to that of a U.S. state government. Second, there is very little local taxation due to national regulations. This eliminates an important margin of response to grants, making it possible to infer the government's time horizon using the pattern of the spending response to grant shocks. Finally, unique policy reforms resulted in the two largest grant programs exhibiting very different variation over time.

Despite Indonesia's distinctive features, the results hold lessons for many other countries, for two reasons. First, many countries have intergovernmental grants that are similar to the ones studied in this paper. Equalization grants are a popular policy tool used in countries from all parts of the income distribution, including Canada, China, Germany, India, and the United Kingdom. Natural resource revenue sharing likewise is very popular and can be found in over 30 countries (Natural Resource Governance Institute, 2016, p. 7). Second, across the developing world, central governments have devolved greater spending responsibilities to local governments in recent decades without devolving revenue-collection responsibilities to a similar degree, making intergovernmental grants especially important in lower-income countries (Gadenne and Singhal, 2014). The results show that inexperienced local governments can use their new decision-making power well by pursuing forward-looking fiscal policy.

This paper contributes to several related literatures in public finance and development. First, it is related to the literature on the so-called flypaper effect, the empirical regularity that local governments have a greater propensity to spend out of non-matching grants than out of local private income (Hines and Thaler, 1995; Inman, 2008). When testing for a flypaper effect, researchers ask how much additional local spending results from increasing grants by one dollar. Rarely do they distinguish between permanent and transitory increases of one dollar.[1] My paper builds on this literature by theoretically and empirically distinguishing between permanent and transitory changes in grant revenue. I show that the permanence of a grant increase determines how the additional spending is distributed across fiscal years. Knowing the timing of fiscal responses to grants is important for conducting countercyclical fiscal policy in a federation.

Second, this paper builds on the literature that exploits exogenous variation in grant allocations to estimate the causal effects of grants. A key concern with the older literature on intergovernmental grants is that the distribution of funds likely depends on local preferences for public good provision, so that the observed relationship between spending and grant

[1] I am aware of only two exceptions. Zou (1994) develops a theoretical model to explore government responses to permanent and temporary changes in various aspects of grant policy. Using a vector error-correction model, Buettner and Wildasin (2006) estimate budgetary responses of U.S. municipalities to temporary and permanent changes in total grant revenue.

revenue may not reflect the causal effect of increasing grants by one dollar.[2] To avoid this problem, recent research has identified features of intergovernmental grant policy that induce exogenous variation in grant allocations.[3] Compared to most of these studies, this paper exploits relatively large exogenous shocks to grants: the general grant reform permanently increased the grant by 46 percent on average and more than doubled the grant for 16 percent of districts.

Finally, this research is related to the development economics literature that examines the extent to which intergovernmental transfers improve conditions for their target populations. In some cases significant portions of targeted transfers have been captured by local politicians (Reinikka and Svensson, 2004), while in others general-purpose transfers have led to wasteful spending, corruption, and at best modest improvements in local public goods and services (Monteiro and Ferraz, 2012; Caselli and Michaels, 2013; Brollo, Nannicini, Perotti, and Tabellini, 2013). Perhaps motivated by these results, other researchers have examined whether increases in tax revenue lead to better local outcomes than increases in transfers (Borge, Parmer, and Torvik, 2015; Gadenne, 2017; Martínez, 2017). The idea is that voters may hold politicians more accountable for how they spend tax revenue as opposed to transfer revenue, perhaps due to endowment effects or because voters have more information about tax revenue than about transfer revenue. The development literature emphasizes the use of survey data on public goods and services as a check against reported spending, which local governments could falsify. Often the outcomes that are easiest to measure, such as the number of schools, require lumpy investment. This paper shows that panel-regression estimates of the impact of government revenue on lumpy public goods can depend on the nature of the identifying variation in revenue. Thus, the results underscore an important methodological challenge to quantifying the accountability effects of taxation: the estimated effects of grants and taxes on public goods can differ even when the two revenue streams are subject to the same degree of accountability, as long as they have different time-series properties.

The closest predecessor to this paper is Olsson and Valsecchi (2015), who examine the effects of shared oil and gas revenue on a variety of district outcomes in Indonesia.[4] Their paper differs from mine in several important ways. First, they ask whether a single intergovernmental grant has any positive effects on economic or social indicators. My paper considers both the general grant and oil and gas revenue and asks how responses to the two grants differ. Second, they do not distinguish theoretically or empirically between permanent and transitory changes in grant revenue. I highlight this distinction in a theoretical model, and I interpret the econometric results in light of the theory. Third, they do not account for the possible endogeneity of oil and gas revenue stemming from the district government's influence on oil and

[2] See Knight (2002) for a formalization of this argument.

[3] Recent examples include Gordon (2004), Baicker (2005), Dahlberg, Mörk, Rattsø, and Ågren (2008), Lutz (2010), Litschig and Morrison (2013), Gennari and Messina (2014), Lundqvist (2015), and Liu and Ma (2016).

[4] Another related paper, Cust and Rusli (2015), disentangles the direct production effect from the fiscal windfall effect of oil and gas production on district GDP in Indonesia.

gas production following decentralization.[5] I combine cross-district variation in predetermined resource endowments with time-series variation in aggregate grant disbursements by the central government to avoid the contaminating influence of the quality of local governance.

The paper proceeds as follows. Section 1.2 presents a model of public expenditure, Section 1.3 provides institutional background on Indonesia, Section 1.4 describes the data, Section 1.5 explains the identification strategy, Section 1.6 presents the empirical results, and Section 1.7 concludes.

1.2 Model

This section develops a simple model of public expenditure on durable and nondurable goods. It is similar to the model of Obstfeld and Rogoff (1996, pp. 96–98), with three main modifications. First, the government's time horizon is finite. This assumption could be justified on behavioral grounds—district heads simply fail to consider outcomes in the distant future—or political grounds—district heads are limited to two five-year terms and care only about citizen welfare while in office. Second, investment in durables is lumpy. This assumption captures the fact that publicly provided durable goods, such as schools or other structures, often must satisfy a minimum size requirement. The second assumption carries important implications for the effects of permanent and transitory changes in revenue on public good provision. Third, public spending is financed solely by intergovernmental grants.

Several testable implications emerge from the model. First, permanent and transitory shocks to grant revenue produce different responses in total expenditure. In particular, a one-time, permanent increase in grant revenue by one dollar causes total expenditure to increase by more than one dollar in the year of the shock, and less than one dollar in subsequent years. By contrast, a one-time, transitory increase in grant revenue by one dollar increases total expenditure by an equal amount, less than one dollar, in the year of the shock and subsequent years. Second, the size of the spending response to the transitory revenue shock is decreasing in the length of the government's time horizon. Third, permanent and transitory shocks to grant revenue have different effects on the composition of public expenditure. While permanent shocks are likely to increase consumption of both durable and nondurable public goods, transitory shocks are more likely to increase consumption of nondurable goods only. The above results hold when durables investment is "lumpy enough." When investment is not lumpy, permanent and transitory revenue shocks have identical per-dollar effects.

[5] The authors limit their sample to districts near the border between an oil-and-gas-producing province and a non-producing province. This strategy controls for cross-sectional differences in geography and culture without addressing differences in the quality of governance, which varies over time and could impact production.

1.2.1 The Environment

Suppose the local government provides a nondurable good, C, and a durable good, D. The durable good evolves according to the equation of motion, $D_t = (1-\delta)D_{t-1} + I_t$, where I_t is current investment in the durable good and δ is the depreciation rate. Let p_t denote the price of durable-good investment in units of the nondurable good in period t. Investment is "lumpy" in the following sense: the government may choose any level of "maintenance" investment up to the point of maintaining the entire durables stock from the previous period, but any increase in the durables stock must exceed a minimum size threshold, $\underline{I}$.[6] Total government spending in period t, G_t, is the sum of C_t and $p_t I_t$. The local government has access to a risk-free bond with exogenous rate of return r. Federal grants are the local government's only source of revenue.[7] Let A_t denote the government's stock of net assets, and let F_t denote federal grant revenue in period t. Assets evolve according to the equation of motion, $A_{t+1} = (1+r)(A_t + F_t - C_t - p_t I_t)$.

1.2.2 The Government's Problem

The local government acts as if it faces a finite time horizon of T periods, starting in period 0 and ending in period $T-1$. The intertemporal budget constraint in starting period t is therefore

$$\sum_{t=0}^{T-1} \frac{C_t + p_t I_t}{(1+r)^t} \leq A_0 + \sum_{t=0}^{T-1} \frac{F_t}{(1+r)^t}.$$

Let $\beta \in (0,1)$ denote the representative citizen's discount factor. The government has perfect foresight and maximizes the representative citizen's utility over the finite time horizon,

$$\sum_{t=0}^{T-1} \beta^t \Big(\gamma \log C_t + (1-\gamma) \log D_t\Big),$$

subject to the intertemporal budget constraint and the investment constraint.[8] Assuming that the initial stock of durables, D_{-1}, and the investment threshold, $\underline{I}$, are both small enough that the investment constraint does not bind, the necessary conditions for an optimum yield the

[6] Formally, the investment constraint is $I_t \in [0, \delta D_{t-1}] \cup [\delta D_{t-1} + \underline{I}, \infty)$, where $\underline{I}$ is the minimum size of new structures. Note that the investment constraint rules out selling any portion of the durables stock. This assumption is inconsequential if the initial stock of durables is small enough.

[7] District governments in Indonesia generate some revenue locally through business license fees, hotel and restaurant taxes, and utility fees. Locally sourced revenue accounts for only 9 percent of the district government budget on average, and districts are prohibited from introducing income or property taxes (World Bank, 2007).

[8] I abstract from private consumption in order to focus attention on the government's optimal expenditure plan. As there is no taxation in the model, adding private consumption would not change any of the results below as long as citizen preferences for private consumption and public consumption were separable.

two Euler equations

$$C_{t+1} = \beta(1+r)C_t,$$
$$\frac{\gamma p_t}{C_t} = \frac{1-\gamma}{D_t} + \beta\frac{\gamma p_{t+1}}{C_{t+1}}.$$

Combining the Euler equations yields the condition

$$\frac{(1-\gamma)C_t}{\gamma D_t} = p_t - \frac{1-\delta}{1+r}p_{t+1} \equiv \iota_t,$$

which states that the marginal rate of substitution between nondurables consumption and durables consumption equals the user cost of durables.

In order to simplify the dynamics of the solution, which will aid in the comparative statics exercise below, I make a number of parametric assumptions. First, assume that the citizen's discount rate equals the interest rate ($\beta = 1/(1+r)$), so that nondurables consumption is constant over time. Next, assume that the price of investment is constant over time ($p_t = p$, hence $\iota_t = pr/(1+r) \equiv \iota$). The citizen thus will want to consume the durable good and the nondurable good in constant proportion over time. Combining the first two assumptions, the citizen will desire a constant level of durables consumption over time. Finally, assume that the depreciation rate is zero ($\delta = 0$).[9] Together, the assumptions imply that all durables investment will occur in the first period—investment in subsequent periods is unnecessary to maintain a constant stock of durables, because there is no depreciation.

Define permanent public income to be

$$Y^P = \frac{r}{1+r-(1+r)^{1-T}}\left(A_0 + pD_{-1} + \sum_{t=0}^{T-1}\frac{F_t}{(1+r)^t}\right),$$

which is the constant resource flow that can be sustained over the government's time horizon (Flavin, 1981). In this model permanent public income is a function of initial financial wealth, the resale value of the initial stock of durables, and the present discounted value of grant revenue. The Euler equations, combined with the simplifying assumptions, imply that $C_{t+1} = C_t$ and $D_t = (\gamma\iota)^{-1}(1-\gamma)C_t$. Substituting the preceding equations into the intertemporal budget constraint yields

$$C_t = \gamma \cdot \frac{1+r-(1+r)^{1-T}}{1+r-\gamma(1+r)^{1-T}} \cdot Y^P, \quad D_t = \frac{1-\gamma}{\iota} \cdot \frac{1+r-(1+r)^{1-T}}{1+r-\gamma(1+r)^{1-T}} \cdot Y^P$$

for $t \in \{0, \ldots, T-1\}$. Thus durables and nondurables consumption are both constant fractions of permanent public income over the time horizon.

[9] Positive depreciation could be incorporated into the model without changing the qualitative nature of the results, as long as δ were small enough.

1.2.3 Response to a Permanent Revenue Shock

Starting at the interior optimum described above, consider a permanent increase in grant revenue by one unit: $\mathrm{d}F_t = 1$ for $t \in \{0, \dots, T-1\}$. Permanent public income increases by one unit: $\mathrm{d}Y^P = 1$. Assuming the revenue increase is large enough to push the government to a new interior optimum with positive investment in period 0, the consumption response to the permanent revenue shock is given by

$$\mathrm{d}C_t^{Perm} = \gamma \cdot \frac{1 + r - (1+r)^{1-T}}{1 + r - \gamma(1+r)^{1-T}}, \quad \mathrm{d}D_t^{Perm} = \frac{1-\gamma}{\iota} \cdot \frac{1 + r - (1+r)^{1-T}}{1 + r - \gamma(1+r)^{1-T}}$$

for $t \in \{0, \dots, T-1\}$. Durables and nondurables consumption immediately increase in period 0 and remain fixed at their new levels for the remainder of the time horizon. Because the initial stock of durables, D_{-1}, is predetermined and $I_t = D_t - D_{t-1}$, period 0 investment rises by $\mathrm{d}I_0{}^{Perm} = \mathrm{d}D_0^{Perm}$. Therefore the response of total public expenditure in period 0 is

$$\mathrm{d}G_0^{Perm} = \mathrm{d}C_0^{Perm} + p\mathrm{d}I_0^{Perm} = \frac{1 + r - \gamma}{r} \cdot \frac{1 + r - (1+r)^{1-T}}{1 + r - \gamma(1+r)^{1-T}}.$$

Thus when there are at least two time periods, total expenditure increases in the first period by more than the increase in permanent income ($\mathrm{d}G_0^{Perm} > 1$) due to the increase in upfront investment in durables. In addition, the increase in period-0 expenditure is smaller the higher the marginal utility of nondurables consumption relative to durables consumption. Because the revenue increase leaves investment unchanged in the ensuing periods,

$$\mathrm{d}G_t^{Perm} = \mathrm{d}C_t^{Perm} = \mathrm{d}C_0^{Perm} < 1$$

for $t \in \{1, \dots, T-1\}$.

To summarize, an increase in permanent grant revenue by one unit in period 0 increases public expenditure by more than one unit in period 0 and less than one unit in periods 1 through $T-1$. The expenditure response is more "front loaded" the stronger the consumer's preferences for durables consumption relative to nondurables consumption. Both durables and nondurables consumption increase in response to the permanent one-unit increase in grant revenue.

1.2.4 Response to a Transitory Revenue Shock

Next consider a temporary increase in grant revenue by one unit in period 0: $\mathrm{d}F_0 = 1$, and $\mathrm{d}F_t = 0$ for $t \in \{1, \dots, T-1\}$. Then assuming there are at least two time periods, the increase in permanent revenue, $\mathrm{d}Y^P = r/\left(1 + r - (1+r)^{1-T}\right)$, is less than one and is decreasing in the length of the time horizon, T. Three responses are possible depending on the parameter values.

Case 1. $\underline{I} \leq \frac{1-\gamma}{\iota} \cdot \frac{r}{1+r-\gamma(1+r)^{1-T}}$.

In the first case the investment constraint does not bind, because the increase in durables consumption that the government would choose in the absence of the investment constraint, $dD_t^{Temp} = dD_t^{Perm} \cdot dY^P$, exceeds $\underline{I}$. The government adjusts durables and nondurables consumption to a new interior solution, so $dC_t^{Temp} = dC_t^{Perm} \cdot dY^P$ and $dG_t^{Temp} = dG_t^{Perm} \cdot dY^P$ for $t \in \{0, \ldots, T-1\}$. The spending response per unit of revenue increase is identical to the case of the permanent revenue shock.

Case 2. $\frac{1-\gamma}{\iota} \cdot \frac{r}{1+r-\gamma(1+r)^{1-T}} < \underline{I} \leq \tilde{I}$, *where* $\tilde{I}$ *satisfies*

$$\gamma \log\left(C_{-1} + dY^P(1 - p\tilde{I})\right) + (1-\gamma)\log(D_{-1} + \tilde{I}) = \gamma \log(C_{-1} + dY^P) + (1-\gamma)\log D_{-1},$$

and C_{-1} *and* D_{-1} *are the pre-shock levels of nondurables and durables consumption, respectively.*

In the second case the investment constraint binds, yet $\underline{I}$ is small enough that the government chooses strictly positive investment in durables, setting $I_0 = \underline{I}$. The representative citizen is indifferent between (i) investing $\tilde{I}$ in durables and adjusting nondurables consumption to satisfy the lifetime budget constraint, and (ii) setting investment equal to zero and spending the entire revenue increase on nondurables consumption. For $\underline{I}$ less than $\tilde{I}$, the citizen would prefer a positive level of investment in durables.

Case 3. $\underline{I} > \tilde{I}$, *where* $\tilde{I}$ *is defined as in Case 2.*

In the final case the investment constraint binds, and investment is zero. The minimum size requirement, $\underline{I}$, is high enough that the citizen would rather spend the entire revenue increase on nondurables rather than invest at least $\underline{I}$ in durables. In this case $dG_t^{Temp} = dC_t^{Temp} = dY_t^P$ and $dD_t^{Temp} = 0$ for $t \in \{0, \ldots, T-1\}$.

To summarize, a transitory increase in grant revenue by one unit could lead to three possible outcomes. If the minimum investment size, $\underline{I}$, is small enough, the government will increase consumption of both durables and nondurables in proportion with the change in permanent income. In this case the per-dollar effects of permanent and transitory revenue shocks will be the same. For slightly larger values of $\underline{I}$, the government will invest the minimum required amount and adjust nondurables consumption to balance the budget. If the minimum size is large enough, the government will spend the entire revenue increase on nondurables consumption. Compared to a permanent revenue shock, a transitory revenue shock will produce a response skewed toward nondurables consumption if the minimum size of nondurables investment is large enough. In this case the total spending response evenly spreads the extra revenue across the time horizon.

1.2.5 Extensions

The model makes several simplifying assumptions for the purpose of tractability. The environment in which local governments operate in Indonesia may deviate from these assumptions in important ways. First, district governments may be liquidity constrained. Indeed, since decentralization was enacted, lending to district governments has been minimal (World Bank, 2007, p. 128). Liquidity constraints would lead to lower government spending in all periods—both when the constraints bind and when they do not. This is because the prospect of liquidity constraints binding in the future lowers current consumption (Zeldes, 1989). In theory, liquidity-constrained governments might respond more modestly to revenue increases—particularly permanent increases—than the model predicts. In practice, however, many district governments accumulated substantial reserves in the years following decentralization, suggesting that liquidity constraints were not a significant issue for a large number of districts.[10]

Second, districts may face uncertainty about future grant revenue. This would create a demand for precautionary saving, lowering current consumption relative to expected future consumption (Leland, 1968).[11] Whether the precautionary-saving motive influences how the government responds to a grant-revenue shock depends on how the shock affects the overall risk faced by the government. In a model in which the government can tax private income at any rate, Vegh and Vuletin (2015) show that the government's spending response to a permanent positive shock to grant revenue is larger, the weaker is the correlation between grant revenue and private income. The reason is that the shock increases the grant share of total income, which is assumed to be less than one half, diversifying the government's "portfolio."[12] The diversification effect is probably less relevant for Indonesia, where district governments cannot set tax rates on income and property. The central government sets and administers these taxes and rebates a portion back to the district. On average shared tax revenue accounts for only 11 percent of the district budget, and own-source revenue from business license fees, hotel and restaurant taxes, and utility fees accounts for nine percent of the budget. In contrast, grant revenue accounts for at least 71 percent of the district budget on average (World Bank, 2007, p. 120). In the Indonesian context a permanent increase in uncertain grant revenue may very well increase the total risk of public revenue, reducing the marginal propensity to spend out of public resources.

Third, bureaucratic delay may prevent local governments from immediately responding to grant revenue shocks. District governments in Indonesia often receive grant funds late in the fiscal year, face significant delays in the process of getting budgets approved by the provincial authorities, and have difficulty procuring goods and services in a timely manner. Bureaucratic delay in the budgeting process has been cited as the reason why districts accumulated substantial reserves in the wake of decentralization (World Bank, 2007, p. 128). Administrative bottlenecks imply that the fiscal responses predicted by the model may occur with a lag.

[10]Reserves were especially high in resource-rich regions, which will provide the bulk of the variation needed to estimate responses to transitory revenue shocks (World Bank, 2007, p. 127).

[11]That is, assuming the utility function has strictly positive third derivatives.

[12]The authors do not consider transitory shocks, though they claim that their main results would not change if shocks were assumed to be temporary.

1.2.6 Econometric Predictions

In order to facilitate interpretation of the regression estimates in Section 1.6, it is helpful to map the predictions of the theoretical model to the parameters of a linear econometric model. Consider the estimating equation

$$G_{dt} = \sum_{k=0}^{K} \beta_k F_{d,t-k} + \sum_{k=0}^{K} \delta_k H_{d,t-k} + U_{dt},$$

where G_{dt} is total expenditure per capita, F_{dt} is grant revenue per capita subject to permanent shocks, H_{dt} is grant revenue per capita subject to transitory shocks, and U_{dt} is the error term of district d in year t. Suppose that $K \geq T - 1$, so that the lag structure of the regression captures the district government's response over the entire time horizon. A permanent increase in F by one unit raises spending ℓ years later by $\sum_{k=0}^{\ell} \beta_k$, while a transitory increase in H by one unit raises spending ℓ years later by δ_ℓ. The sum $\sum_{k=0}^{\ell} \delta_k$ gives how much of the one-unit transitory increase in H is spent over the course of $\ell + 1$ years.

First consider a permanent shock to F. According to the baseline theoretical model, $\beta_0 = \mathrm{d}G_0^{Perm} > 1$ and $\beta_0 + \beta_1 = \mathrm{d}G_t^{Perm} < 1$ for $t \in \{1, \ldots, K\}$. It follows that $\beta_1 < 0$ and $\beta_k = 0$ for $k \in \{2, \ldots, K\}$. As mentioned above, bureaucratic delay can create a one- or two-year lag between budget proposal and budget implementation. How do the predictions for the econometric model change in the presence of bureaucratic delay? Suppose that a portion $\pi^0 \in (0, 1)$ of the desired budget adjustment is implemented in the current year, a portion $\pi^1 \in (0, 1 - \pi^0]$ is implemented in the following year, and the remaining portion $1 - \pi^0 - \pi^1$ is implemented with a two-year lag. Suppose further that $K = 3$. Then the fiscal response to a permanent one-unit increase in F is characterized by the equations

$$\begin{aligned}
\beta_0 &= \pi^0 \mathrm{d}G_0^{Perm}, \\
\beta_0 + \beta_1 &= \pi^1 \mathrm{d}G_0^{Perm} + \pi^0 \mathrm{d}G_1^{Perm}, \\
\beta_0 + \beta_1 + \beta_2 &= (1 - \pi^0 - \pi^1)\mathrm{d}G_0^{Perm} + (\pi^0 + \pi^1)\mathrm{d}G_1^{Perm}, \\
\beta_0 + \beta_1 + \beta_2 + \beta_3 &= \mathrm{d}G_1^{Perm},
\end{aligned}$$

using the fact that $\mathrm{d}G_t^{Perm} = \mathrm{d}G_1^{Perm}$ for $t \in \{1, \ldots, K\}$. Theory provides several testable predictions about the coefficient values. First, if π^1 is large, as suggested by policy reports, then $\beta_0 + \beta_1 > 1$. Second, if $\pi^0 + \pi^1 < 1$, then $\beta_3 < 0$. Finally, if on the other hand, $\pi^0 + \pi^1 = 1$, then $\beta_2 < 0$ and $\beta_3 = 0$.[13] More generally, the front-loaded nature of the response implies that the main prediction to test is $\beta_0 + \beta_1 > \beta_2 + \beta_3$.

Now consider a transitory shock to H under the assumption that Case 3 holds, i.e., the government only adjusts spending on nondurables. The baseline theoretical model predicts that

[13] Note that all three predictions follow from the fact that $\mathrm{d}G_0^{Perm} > \mathrm{d}G_1^{Perm}$.

$\delta_k = \delta_\ell = \mathrm{d}G_0^{Temp}$ for $k, \ell \in \{0, \ldots, K\}$. If instead there are bureaucratic delays, then a portion $\phi^0 \in (0,1)$ of the desired budget adjustment is implemented in the current year, a portion $\phi^1 \in (0, 1-\phi^0]$ is implemented in the following year, and the remaining portion $1-\phi^0-\phi^1$ is implemented with a two-year lag. In this case the response to a transitory one-unit increase in H is characterized by the equations

$$\begin{aligned}
\delta_0 &= \phi^0 \mathrm{d}G_0^{Temp}, \\
\delta_1 &= (\phi^0 + \phi^1)\mathrm{d}G_0^{Temp}, \\
\delta_2 &= \delta_3 = \mathrm{d}G_0^{Temp},
\end{aligned}$$

which use the fact that $\mathrm{d}G_t^{Temp} = \mathrm{d}G_0^{Temp}$ for $t \in \{0, \ldots, K\}$. Because the response to a transitory shock to H will be skewed toward spending on nondurables, which involves less red tape than capital projects (World Bank, 2007, pp. 98–103), the associated bureaucratic delays will be lower than in the case of durables. As a result, it is reasonable to assume that $\pi^0 \leq \phi^0$ and $\pi^0 + \pi^1 \leq \phi^0 + \phi^1$. Furthermore, spending on nondurables is unlikely to face a two-year delay, so $\phi^0 + \phi^1 = 1$. Mild bureaucratic delay therefore leads to the modified predictions that $\delta_0 < \delta_1 = \delta_2 = \delta_3 < 1$. If ϕ_0 is large (nondurables can be adjusted relatively quickly), then it is approximately true that $\delta_0 + \delta_1 = \delta_2 + \delta_3$, providing a prediction that is analogous to the main prediction for permanent shocks.

Finally, the theory predicts that nondurables consumption will increase more in response to a permanent increase in F than it will to a transitory increase in H. The differential response to permanent and transitory revenue shocks is especially likely to appear in measures of long-lived structures, such as schools and health facilities.

1.3 Institutional Background

1.3.1 Decentralization in Indonesia

The resignation of Suharto as president of Indonesia in 1998 marked the end of three decades of highly centralized authoritarian rule and paved the way for dramatic political and economic reforms. Indonesia now ranks as one of the most decentralized countries in the developing world (Shah, Qibthiyyah, and Dita, 2012). There are three levels of subnational government in Indonesia: province, district, and village. Indonesia currently has 34 provinces. The number of districts has grown from 336 in 2001 to 514 in 2014.[14] Districts are categorized as either rural districts (*kabupaten*) or municipalities (*kota*). Expenditure responsibilities are the same for both types of district.

The "Big Bang" fiscal decentralization reforms of 2001 devolved vast expenditure authority

[14] The increase is due to the widespread phenomenon of district splitting. See, for example, Fitrani, Hofman, and Kaiser (2005), Burgess, Hansen, Olken, Potapov, and Sieber (2012), and Bazzi and Gudgeon (2016) for details.

to provincial and (especially) district governments (World Bank, 2003). The share of total expenditures managed by subnational governments rose from 24 percent in the mid-1990s to 36 percent in 2011, with district governments accounting for most subnational spending.[15] However, own-source revenue accounts for only nine percent of total subnational revenue, necessitating large fiscal transfers from the central government (World Bank, 2007; Shah et al., 2012).

1.3.2 General Grant

The largest source of financing for most district governments is a federal grant known as the General Allocation Fund (*Dana Alokasi Umum*), or "general grant" for short, which accounts for around 56 percent of district revenue. The general grant is an equalization grant, intended to equalize the capacity to provide local public goods across regions. Equalization grants have the potential to promote equity by targeting areas populated by households with low earning potential. In real-world contexts, such as in Canada, such grants often distort household location decisions and fall short of equity goals (Albouy, 2012). Researchers and policymakers have argued that Indonesia's general grant is insufficiently equalizing and promotes inefficient spending on the civil service wage bill (Hofman, Kadjatmiko, Kaiser, and Sjahrir, 2006; World Bank, 2007).

Districts have complete discretion over how to spend the general grant. The total budget for the grant depends on long-term forecasts of factors determining the central government's budget health, such as the price of oil (World Bank, 2007). The allocation formula has two components: the basic allocation and the fiscal gap. The basic allocation consists of a lump-sum portion and a portion that is a function of the civil service wage bill. The fiscal gap is calculated as the difference between expenditure needs and fiscal capacity. The formula for the general grant is

$$\text{General Grant} = \text{Basic Allocation} + \text{Expenditure Needs} - \text{Fiscal Capacity}.$$

Expenditure needs are calculated as a weighted sum of indices related to population, land area, poverty, and cost of construction. Section 1.5 discusses the expenditure-needs formula in greater detail. Since 2002, fiscal capacity has been defined as the weighted sum of imputed own-source revenue, shared tax revenue, and shared natural resource revenue:[16]

$$\begin{aligned}\text{Fiscal Capacity} = a \cdot (\text{Imputed Own-Source Revenue}) + b \cdot (\text{Shared Tax Revenue}) \\ + c \cdot (\text{Shared Natural Resource Revenue}).\end{aligned}$$

[15] Indonesia's level of expenditure decentralization exceeds the OECD average and is higher than every East Asian country except China (World Bank, 2007).

[16] Own-source revenue is imputed based on a regression of actual own-source revenue on regional GDP (World Bank, 2007).

From 2002–2011 the value of a has varied between 0.5 and 1, b has varied between 0.73 and 1, and c has varied between 0.5 and 1. The fiscal gap component accounts for around half of the general grant budget. The general grant allocation is determined on a yearly basis. For the first two-thirds of the sample period, general grant disbursements followed a "hold-harmless" rule which ensured that general grant receipts would not fall below the previous year's receipts. The hold-harmless rule froze the general grant amount for many resource-rich districts which otherwise would have received much lower disbursements according to the formula (World Bank, 2007, p. 121).

1.3.3 Shared Oil and Gas Revenue

Districts rich in natural resources receive Shared Natural Resource Revenue (*Sumber Daya Alam*), which is allocated in proportion to resource production that occurs in the district and province. Oil and natural gas are by far the largest sources of natural resource revenue in Indonesia. Fifteen percent of oil revenue collected within a district is redistributed to subnational governments: 3 percent goes to the provincial government, 6 percent goes to the producing district, and the remaining 6 percent is evenly divided among the other districts located in the same province. The sharing rule for natural gas is more generous to subnational governments: 6 percent goes to the provincial government, 12 percent goes to the producing district, and another 12 percent is divided equally among the other districts in the province. Despite the less generous sharing rule, shared oil revenue on average exceeds shared gas revenue due to the higher value of oil production. Disbursements are to be made on a quarterly basis (Law No. 33/2004), though in practice the transfers often arrive late in the fiscal year (World Bank, 2007, p. 128). Districts have complete discretion over how to spend the shared oil and gas revenue.[17]

1.3.4 Political Institutions

The post-Suharto decentralization reforms also included considerable political decentralization. Starting in 1999, local parliaments were democratically elected through a proportional representation system. The district heads ("mayors") previously appointed by Suharto were allowed to finish their five-year terms, after which time each local parliament appointed a new district head. The political system was reformed yet again with the introduction of direct elections for district heads starting in 2005. Incumbent mayors were allowed to finish their terms before direct elections were held. For idiosyncratic reasons, terms of Suharto mayors expired in different years, so that the terms of indirectly appointed mayors also expired at

[17] In 2009 the central government slightly increased the amount of oil and gas revenue shared with subnational governments, earmarking this additional revenue for education (Law No. 33/2004). As a result, after 2009 around three percent of the district's oil grant, and two percent of the district's gas grant, was earmarked. This earmarking is unlikely to play any role in district spending decisions, as earmarked funds are extremely small relative to total education spending, which represents one third of the district budget on average.

different times. As a result, direct elections were introduced in a staggered manner with exogenously determined timing.[18]

1.4 Data

The data on federal grants and district revenue and expenditure come from reports by the Ministry of Finance (*Kementerian Keuangan*).[19] Data on district expenditure disaggregated by function come from the Ministry of Finance and the World Bank's Indonesia Database for Policy and Economic Research (INDO-DAPOER).[20] All fiscal variables are measured in constant 2010 IDR 10 million (approximately USD 1,000) per capita.[21] Data on grant revenue, other sources of revenue, and expenditure are available for the years 2001–2014.[22] Data on expenditure broken down by function are available for the years 2001–2012. INDO-DAPOER also provides information on district characteristics, such as land area and population.[23] Data on public good provision come from the Village Potential Statistics (*Pendataan Potensi Desa*, or PODES) survey waves of 2000, 2003, 2005, 2008, 2011, and 2014. The surveys act as a village census and thus are meant to cover every village.[24] Data on oil and gas reserves come from the proprietary UCube database maintained by Rystad Energy, an international oil and gas consulting company.[25]

To ensure that all districts in the sample operate under comparable institutional settings, I omit provinces that have a special administrative or fiscal arrangement with the central government.[26] Of the remaining districts, only those that existed as of 2005 or earlier are included in the sample.[27] The final sample for the analysis of public finance outcomes includes 372 districts from 29 provinces. The analysis of village-level public good outcomes is based on

[18]Skoufias, Narayan, Dasgupta, and Kaiser (2014) establish that the timing of the introduction of direct elections was unrelated to district characteristics.

[19]Each year district mayors are required to report on the district's finances to the Ministry of Finance. See http://www.djpk.depkeu.go.id/.

[20]The World Bank team in Jakarta spent considerable time and effort creating expenditure categories by economic classification and function that are consistent over time (Skoufias et al., 2014). The data are available at http://databank.worldbank.org/data/reports.aspx?source=1266.

[21]The abbreviation "IDR" stands for "Indonesian Rupiah."

[22]INDO-DAPOER provides data on revenue and expenditure broken down by economic classification up to either 2012 or 2013 depending on the variable. I add data from 2013–2014 using the budget report from the Ministry of Finance. I also replace missing or obviously incorrect values in INDO-DAPOER using the Ministry of Finance data.

[23]District population is available until 2013. I impute 2014 population using 2013 population and the median annual growth rate of approximately 1.4 percent.

[24]Due to a massive tsunami in 2004, the 2005 wave lacks data on districts on the islands of Nias (Nias, Nias Utara, Nias Barat, Nias Selatan, and Gunung Sitoli).

[25]For details on the UCube database, see https://www.rystadenergy.com/Products/EnP-Solutions/UCube.

[26]These provinces are DI Yogyakarta, which has special autonomy status; DKI Jakarta, whose districts are managed by the province; Nanggroe Aceh Darussalam, which has special autonomy status and receives special autonomy funds; and Papua and Papua Barat, which both receive special autonomy funds.

[27]I impose this restriction because the identification strategy exploits a policy reform in 2006. Therefore it is necessary to observe district outcomes before and after 2006.

a balanced panel of over 41,000 villages matched across the PODES waves and located within the same 372 districts and 29 provinces included in the public finance sample.[28] A significant number of villages split into multiple villages over the period 2000–2014. To maintain a consistent unit of observation in the public goods sample, I aggregate village outcomes up to 2000 borders.s[29]

1.5 Identification Strategy

I first consider fiscal responses to the general grant and shared oil and gas revenue. The structural equation for fiscal outcome Y is

$$Y_{dit} = \sum_{k=0}^{K} \beta_k GenGrant_{di,t-k} + \sum_{k=0}^{K} \delta_k OilGasRev_{di,t-k} + \alpha_d + \lambda_{it} + \varepsilon_{dit}, \tag{1.1}$$

where d indexes districts, i indexes island groups, and t indexes years. The model accounts for district fixed effects, α_d, as well as arbitrary island-specific time trends, λ_{it}.[30] I report standard errors that are robust to heteroskedasticity and two-way clustering at the district and province $\times$ year levels to account for within-district serial correlation and cross-district correlation within the same province and year (Cameron, Gelbach, and Miller, 2011).[31] As described in Section 1.2, the main hypotheses to test are $\beta_0 + \beta_1 = \beta_2 + \beta_3$ and $\delta_0 + \delta_1 = \delta_2 + \delta_3$, which describe the nature of the fiscal response over time. Also of interest are the sums of lagged effects. As already mentioned, $\sum_{k=0}^{K} \beta_k$ represents the effect of a permanent increase in the general grant by one unit on the level of Y K years later. By contrast, $\sum_{k=0}^{K} \delta_k$ represents the amount of a one-unit transitory increase in oil and gas revenue that is spent over the course of $K + 1$ years.

Next I consider the effects of the two grants on local public goods and services. The structural equation for village outcome Y is

$$Y_{vdis} = \beta\, \overline{GenGrant_{dis}} + \delta\, \overline{OilGasRev_{dis}} + \alpha_d + \lambda_{is} + \varepsilon_{vdis}, \tag{1.2}$$

where v indexes villages and s indexes time periods spanned by the PODES survey years.

[28] I drop villages with data that appear miscoded or indicate an incorrect merge. First, I drop villages with reported annual population growth of more than 25 percent or less than −25 percent in any time period. Second, I drop villages with reported population growth of at least 10 percent followed immediately by a population decline of at least 10 percent, or vice versa. Finally, I drop villages with implausibly large changes in public goods from one survey year to the next. To minimize the influence of outliers, I also drop villages with population below the 2nd percentile or above the 98th percentile in any year.

[29] Around a quarter of villages split from 2000–2014. I drop any village that was involved in an amalgamation during the sample period (roughly three percent of villages).

[30] Following the Indonesian Statistical Bureau, I code seven island groups: Sumatra, Java, Nusa Tenggara, Kalimantan, Sulawesi, Maluku, and Papua.

[31] The latter correlation arises from the fact that, in any given year, non-producing districts located in the same province receive the same level of oil and gas revenue.

The variable $\overline{GenGrant_{dis}}$ is the average annual general grant revenue during period s, and $\overline{OilGasRev_{dis}}$ is defined similarly. The outcome Y_{vdis} is a flow variable, such as the number of doctors per capita employed at the end of period s, or the average annual change in the stock of health clinics per capita over period s. For period s starting in year t_0 and ending in year t_1, the average annual general grant revenue and the average annual change in the stock of health clinics per capita are calculated as

$$\overline{GenGrant_{dis}} = \frac{1}{t_1 - t_0} \sum_{t=t_0+1}^{t_1} GenGrant_{dit}, \quad Y_{vdis} = \frac{1}{t_1 - t_0}(H_{vdit_1} - H_{vdit_0}),$$

where H_{vdit} is the stock of health clinics per capita in year t.[32] The regressions use as many time periods as possible, subject to the availability of data on the outcomes. The main hypothesis to test is $\beta = \delta$.

Both fiscal transfers could be endogenous in equations (1.1) and (1.2). The general grant is likely endogenous because it is a function of the civil service wage bill and fiscal need. Shared oil and gas revenue is potentially less problematic, but it could be endogenous if oil and gas production is affected by the local business environment, local economic shocks, conflict, or other unobservables that also affect district public-good outcomes. Furthermore, deviations of the two grants from the allocations prescribed by their respective formulas could reflect the relative political bargaining power of the district, introducing another source of endogeneity.[33] In order to consistently estimate the coefficients of interest, I exploit sources of exogenous variation in the grants, explained below.

1.5.1 General Grant

To estimate the effect of the general grant, I exploit variation induced by a large policy reform. The central government of Indonesia considers forecasts of its long-run budget health in determining how much money to allocate to the general grant. A key parameter in these forecasts is the assumption about the future price of oil. In 2006 the total general grant budget increased by 44 percent after the central government adjusted the oil price assumption from USD 30 per barrel to USD 60 per barrel (Agustina, Ahmad, Nugroho, and Siagian, 2012). The central government also adjusted the formula for expenditure needs in 2006, resulting in a larger share of the general grant budget going to less densely populated districts. Thus while most districts saw an increase in general grant revenue in 2006, the least densely populated districts saw the largest increases. Districts rich in oil and gas resources should have experienced a decline in general grant funds according to the formula. However, a hold-harmless provision froze the general grant allocation in place for these resource-abundant districts. Therefore the

[32] Note that $\frac{1}{t_1-t_0}(H_{vdit_1} - H_{vdit_0}) = \frac{1}{t_1-t_0}\sum_{t=t_0+1}^{t_1}(H_{vdi,t} - H_{vdi,t-1})$, the average annual change in H_{vdit} from t_0 to t_1.

[33] See the discussion in Dahlberg et al. (2008).

change in general grant revenue per capita received by district d from 2005 to 2006 is given by

$$GenGrant_{di,2006} - GenGrant_{di,2005} \approx \theta + \pi AreaPC06_{di} \times NonOilGas_{di} + Remainder_{di},$$

where $\pi > 0$, $AreaPC06$ is land area per capita in 2006, and $Remainder_{di}$ is much smaller than $\pi AreaPC06_{di}$ in absolute magnitude (World Bank, 2007). See the appendix for more details on the general grant formula and a derivation of the above approximation. Changes to the general grant in the pre-reform period (2001–2005) and post-reform period (2007–2014) were modest. As a result, general grant revenue per capita in district d and year t can be approximated as

$$GenGrant_{dit} \approx \theta + \pi AreaPC06_{di} \times NonOilGas_{di} \times 1(t \geq 2006) + Remainder_{dit},$$

where $1(t \geq 2006)$ equals one in years 2006 and later and zero in years prior to 2006. Both the increase in the total budget for the general grant and the change in the allocation formula were announced in 2004 (Law No. 33/2004).

Figure 1.1 graphs average general grant revenue per capita over time separately for three groups of districts divided according to area per capita in 2006. Panel (a) includes districts located in provinces with an insignificant oil and gas endowment per capita, while Panel (b) includes districts located in the six provinces with highly significant oil and gas endowments per capita.[34] In each figure, average general grant revenue per capita for districts exceeding the 75th percentile in land area per capita (among all districts) is shown with a solid blue line. The green long dashes apply to districts between the 50th and 75th percentiles in land area per capita, while yellow short dashes indicate districts below the 50th percentile and land area per capita. From 2001–2005, districts with greater area per capita received a larger general grant allocation in per capita terms. Over this period, both the level of general grant per capita in each group as well as the differences in general grant allocations between groups remained approximately constant over time. Starting in 2006, districts in resource-poor provinces with below-50th percentile land area per capita experienced only a small increase in general grant per capita. In contrast, districts in the third quarter of the distribution saw a moderate increase in general grant per capita, and districts in the top quarter experienced a massive increase in general grant per capita. The relative distribution of general grant revenue per capita by land area did not change much over time in provinces rich in oil and gas resources. The policy reform of 2006 thus provides significant cross-district variation in the size of a permanent shock to the general grant within provinces that lack significant oil and gas resources.

Figure 1.1 establishes that, in provinces with insignificant oil and gas resources, districts with greater land area per capita experienced a larger increase in general grant revenue per

[34]The six oil and gas rich provinces are Riau, Kepulauan Riau, Jambi, Sumatera Selatan, Kalimantan Timur, and Kalimantan Utara. See the appendix for the distribution of oil and gas revenue and oil and gas endowment across provinces. Note that while Kalimantan Utara was officially formed in 2012, as of 2014 its districts received shared oil and gas revenue as if they were still part of their former province of Kalimantan Timur. The appendix figures therefore combine the two provinces.

capita starting in 2006 than more densely populated districts. Consequently, the interaction term $AreaPC06_{di} \times NonOilGas_{di} \times 1(t \geq 2006)$ is a relevant instrument for general grant revenue per capita that summarizes the variation due to the reform in the most parsimonious way possible. Because district population growth is relatively slow, area per capita is approximately time invariant.[35] Therefore, the fixed-effects model allows for the possibility that the level of outcomes, such as spending or public goods, could depend on land area per capita. The exclusion restriction requires only that the direct effect of land area per capita on district outcomes be the same on average in the periods 2001–2005 and 2006–2014. Intuitively, this means that outcomes in districts with different levels of population density would have followed parallel paths over time in the absence of the general grant reform.

Even if the relationship between land area per capita and local preferences for spending were time-invariant, one may worry that the timing or overall size of the reform could be endogenous to the political or economic demands of less densely populated districts in resource-poor provinces. For example, members of the national legislature representing less densely populated districts may have pushed for the reform in order to help their own reelection prospects or the prospects of incumbents in the district legislatures. The timing of the reform is inconsistent with this story, however, as elections for both the national and district legislatures took place in 1999, 2004, 2009, and 2014. That is, the reform took effect three years prior to the next legislative elections, casting doubt on the claim that the timing of the reform was the result of political calculus. Alternatively, members of the national legislature may have wanted to improve the reelection prospects of incumbent mayors in less densely populated districts. If this were the case, then one would expect to see a disproportionate number of mayoral elections taking place in less densely populated districts in resource-poor provinces in 2006. In reality, among resource-poor provinces, the average land area per capita of districts with mayoral elections in 2005 is statistically indistinguishable from the average land area per capita of districts with mayoral elections in 2006, 2007, or 2008.[36] Thus, there is little reason to believe that the timing or overall size of the general grant reform were motivated by political considerations.

1.5.2 Shared Oil and Gas Revenue

For the purpose of natural resource revenue sharing, district territory includes sea territory that extends up to four nautical miles from the coastal shoreline (Law 22/1999). Government revenue collected from oil production within a district is divided as follows: 85 percent goes to the central government, 3 percent goes to the provincial government, 6 percent goes to the producing district, and the remaining 6 percent is divided equally among the non-producing districts located in the same province as the producing districts. Government revenue collected

[35] Median annual population growth is 1.4 percent.

[36] Results available upon request.

from gas production within a district is divided as follows: 70 percent goes to the central government, 6 percent goes to the provincial government, 12 percent goes to the producing district, and the remaining 12 percent is divided equally among the non-producing districts located in the same province as the producing districts. Let O_{dpt} and G_{dpt} denote oil and gas revenues produced in district d, located in province p, in year t. Shared oil and gas revenue per capita is

$$OilGasRev_{dpt} = \frac{1}{Pop_{dpt}} \left(0.06 \cdot O_{dpt} + 0.12 \cdot G_{dpt} + \frac{0.06}{N_{pt} - 1} \sum_{j \neq d} O_{jpt} + \frac{0.12}{N_{pt} - 1} \sum_{j \neq d} G_{jpt} \right),$$

where Pop_{dpt} is the population of district d in year t, and N_{pt} is the number of districts in province p in year t. Using the Rystad UCube database, I calculate the total amount of economically recoverable oil and gas resources as of 2000 (and known in 2000), prior to fiscal decentralization. I denote these measures as $Endow_{dt}^{Oil}$ and $Endow_{dt}^{Gas}$. The only reason the endowment measures could vary over time is because district borders sometimes change.[37] Using the sharing rule, I define the variable

$$\begin{aligned} EndowPC_{dpt} = \frac{1}{Pop_{dpt}} \Bigg(& 0.06 \cdot Endow_{dpt}^{Oil} + 0.12 \cdot Endow_{dpt}^{Gas} + \\ & \frac{0.06}{N_{pt} - 1} \sum_{j \neq d} Endow_{jpt}^{Oil} + \frac{0.12}{N_{pt} - 1} \sum_{j \neq d} Endow_{jpt}^{Gas} \Bigg), \end{aligned}$$

which represents oil and gas endowment per capita to which district d has a claim for revenue-sharing purposes.

Despite the formula established by law, the time variation in central government disbursements of shared oil and gas revenue does not perfectly match the time variation in the value of resource production. Panel (a) of Figure 1.2 graphs total oil and gas revenue shared with the districts against the weighted value of oil and gas production, where the value of oil production is given a weight of 0.06 and the value of gas production is given a weight of 0.12 as per the sharing formula. The weighted value of production should be roughly proportional to the central government's revenue base from oil and gas production. The two time series do not closely track each other, indicating that the central government frequently deviates from the revenue sharing rule on a discretionary basis.[38] The distribution of oil and gas revenue is highly skewed, with only six provinces—Jambi, Sumatera Selatan, Riau, Kepulauan Riau, Kalimantan Timur, and Kalimantan Utara—receiving significant amounts of revenue.[39] Panel (b) of Figure 1.2 graphs average oil and gas revenue separately for districts in the top five

[37] Fitrani et al. (2005) find no consistent relationship between natural resources and the likelihood of a district split from 1998–2004.

[38] Lags of weighted oil and gas production also do not closely track total shared oil and gas revenue.

[39] See the appendix.

percent in terms of oil and gas endowment, districts between the 90th and 95th percentiles, and districts in the bottom 90 percent. Oil and gas revenue is significant only for the top 10 percent of districts in terms of endowment. Furthermore, districts in the top five percent in terms of endowment experience very sharp increases and decreases in oil and gas revenue from one year to the next. Cross-district variation in resource endowments, combined with variation in aggregate oil and gas revenue disbursements over time, provide exogenous variation in the size and timing of *transitory* shocks to oil and gas revenue.

I instrument for oil and gas revenue with the interaction between aggregate disbursements of shared oil and gas revenue and predetermined oil and gas endowment per capita in 2000, $AggOilGasRev_t \times EndowPC_{dt}$. The validity of the instrument rests on two identifying assumptions. The first identifying assumption is that variation in aggregate oil and gas revenue over time is unrelated to district-level economic conditions. This assumption is likely to be satisfied, as each district accounts for less than 10 percent of total shared oil and gas revenue. Furthermore, while the time profile of aggregate disbursements does not perfectly match that of the value of oil and gas production, clearly the two are related, and the latter is driven more by fluctuations in resource prices than fluctuations in production.[40] In addition, the large increase in shared oil and gas revenue in 2006 is likely due in part to the central government's revised oil price forecast, which is exogenous from the standpoint of district governments. The second identifying assumption is that district oil and gas endowment is unrelated to unobservables driving local economic trends. One concern would be that better-managed districts attract more oil and gas exploration, which in turn increases known endowment. The instrument avoids contamination along these lines by measuring endowment known as of 2000, prior to fiscal decentralization. Before the decentralization reforms, the central government was the sole actor in negotiating with oil and gas companies. As a result, incentives to explore for oil and gas were roughly uniform across the archipelago prior to 2001.[41]

1.6 Results

1.6.1 Fiscal Responses

This section presents the estimates of district fiscal responses to the two grants. Panel A of Table 1.1 provides summary statistics for the district-level variables. District population averages around 560,000 and ranges from about 30,000 to 5.3 million. Fiscal variables are measured in constant 2010 IDR 1 million per capita. On average district revenue is IDR 2.1 million (approximately USD 210) per capita. Total revenue varies significantly across districts, ranging from IDR 130,000 (USD 13) to IDR 24 million (USD 2,400) per capita.

[40] Figures available upon request.

[41] Violent separatist movements in Aceh and Papua have disrupted resource extraction in the past. These regions are excluded from the sample.

Panel (a) of Figure 1.3 plots total spending over time for districts in resource-poor provinces and divided into three groups according to land area per capita in 2006. Trends in total spending were similar across the three groups prior to the general grant reform. Starting in 2006, total spending increases significantly more for districts with a high land area per capita than districts with a low land area per capita. Panel (b) of Figure 1.3 plots total spending over time for districts divided into three groups according to their oil and gas endowment in 2000. The most richly endowed districts spend considerably more than resource-poor districts over the sample period, and the spending gap between resource-rich and resource-poor districts grows over time—likely due to the increase in oil and gas revenue starting in 2006. Spending in the most richly endowed districts has a hump shape from 2001–2004, reflecting a response to the large increase in shared resource revenue at the start of decentralization followed by a delayed response to the ensuing decline in oil and gas revenue. Spending in these districts grows again, particularly sharply in 2008 following the large oil and gas revenue increase in 2006, before falling sharply in 2010. Overall the time path of spending in resource-rich districts is much smoother than the time path of total shared oil and gas revenue (dotted line).

Table 1.2 presents the first-stage results. To make the first-stage estimates readable, land area per capita is measured in tens of square kilometers per capita, and aggregate oil and gas revenue is measured in 2010 IDR trillions. The first instrument, $AreaPC06 \times NonOilGas \times 1(Year \geq 2006)$, has a positive effect on general grant revenue per capita that is significant at the one-percent level. The magnitude and statistical significance of this first-stage effect is insensitive to the inclusion of the second instrument, $AggOilGasRev \times EndowPC$, which has an insignificant effect on general grant revenue per capita. The second instrument has a positive effect on oil and gas revenue per capita that is significant at the one-percent level. Similarly, this first-stage effect is insensitive to the inclusion of the first instrument, which has an insignificant effect on oil and gas revenue per capita. In the second-stage regressions, the Sanderson and Windmeijer (2016) F statistic, which tests for weak identification of individual coefficients on the endogenous variables, is typically 17 or greater for the general grant and 50 or greater for the oil and gas revenue, indicating that the structural parameters are strongly identified.

Table 1.3 presents estimates of the effects of the two grants on the main alternative sources of revenue. Panel A presents the ordinary least squares estimates, and Panel B presents the two-stage least squares estimates. In column one, both the OLS and 2SLS estimates suggest the grants have little effect on own-source revenue, which mostly comes from business license fees, hotel and restaurant taxes, and utility fees. In column two, the OLS results suggest that increasing the general grant by one dollar per capita raises special allocation grants (*Dana Alokasi Khusus*, or DAK)—earmarked transfers given by the central government on a discretionary basis—by 11 cents per capita. The point estimate is significant at the one-percent level. However, the 2SLS estimate is one-fifth the size of the OLS estimate and is insignificant. One reason for the discrepancy could be that both the general and special grants are targeted towards poorer districts, and the OLS estimate reflects this source of endogeneity.

The estimated impact of oil and gas revenue on special allocation grants is modest at best. Finally, in column 3 the OLS estimate suggests that increasing oil and gas revenue by one dollar per capita raises shared tax revenue by 36 cents per capita. This point estimate is significant at the 10-percent level. However, the corresponding 2SLS estimate is half the magnitude of the OLS estimate and is insignificant. Both the OLS and 2SLS results indicate that general grant has little effect on shared tax revenue. Overall there is little indication that either grant significantly crowds out or crowds in other types of revenue.

Tables 1.4 and 1.5 present OLS and 2SLS estimates, respectively, of the expenditure responses to the two grants broken down by economic classification. The tables present estimates of the individual coefficients on current grant revenue and three lags, the sum of the lag coefficients, and tests for whether the spending responses to the grants are "front-loaded" in the sense that the sum of the coefficients on lags zero and one exceeds the sum of the coefficients on lags two and three. The latter test arises from the theoretical model, which predicts that, in the presence of administrative delay, the sum of the coefficients on lags zero and one will exceed the sum of the coefficients on lags two and three for the general grant but not for oil and gas revenue. For every spending category and both the OLS and 2SLS estimates, we cannot reject the hypothesis that the sum of the coefficients is equal for the two grants. This means that the overall propensity to spend out of the two grants is about the same. However, the timing of the spending responses differ for the two grants in the manner predicted by the theoretical model. The OLS results in Table 1.4 indicate that the spending response to the general grant is front-loaded in every category. For total spending, the sum of the first two coefficients is 0.96 whereas the sum of the last two coefficients is −0.02. In contrast, the spending response to oil and gas revenue is front-loaded only for goods and services. Even in that case, the sum of first two coefficients (0.24) is not significantly higher than the sum of the last two (0.16).

The 2SLS results of Table 1.5 are less precise, but they generally tell the same story. The spending response to the general grant is clearly front-loaded in three out of four categories. The test fails to reject the hypothesis that $\beta_0 + \beta_1 = \beta_2 + \beta_3$ in the case of capital expenditure, likely due to the imprecision of the point estimates. Nonetheless, the pattern of the point estimates suggests a front-loaded capital expenditure response. In the case of oil and gas revenue, the test fails to reject $\delta_0 + \delta_1 = \delta_2 + \delta_3$ in three out of four cases. Figure 1.4 displays the 2SLS estimates in graph form. While the response to the general grant is hump-shaped, the response to oil and gas revenue is fairly flat. In response to a permanent increase in the general grant by \$1, total spending increases by \$1.45 one year later, \$1.85 two years later, and \$1.24 three years later. Out of a \$1 transitory increase in oil and gas revenue, \$0.81 is spent in the first two years, and \$0.41 is spent in the next two years. Interestingly, practically the entire response of capital expenditure to the general grant occurs with a one-year delay, while there is a larger contemporaneous response of goods and services and personnel. The reason could be that the latter two categories may be associated with fewer bureaucratic delays compared to capital expenditure.

Table 1.6 presents OLS estimates of the expenditure responses broken down by function, and Table 1.7 presents the corresponding 2SLS estimates. These estimates are based on a smaller sample, because spending by category is only available until 2012. For all five categories the OLS estimates suggest the expenditure response to the general grant is front-loaded, whereas we can never reject the hypothesis that the response to the oil and gas revenue is evenly distributed over time. Once again, the 2SLS estimates are less precise. In all five cases $\beta_0 + \beta_1$ exceeds $\beta_2 + \beta_3$, but we can only reject the hypothesis that they are equal for health expenditure. In four out of five spending categories, the response to the oil and gas revenue is fairly evenly distributed, and we cannot reject the hypothesis that $\delta_0 + \delta_1 = \delta_2 + \delta_3$. We can only reject the hypothesis in the case of the administration spending response to oil and gas revenue, which surprisingly appears back-loaded.

The preponderance of the evidence on fiscal responses is consistent with the theoretical model, suggesting that the permanent increase in the general grant induced a front-loaded spending response in order to overcome an investment threshold. The spending increase in response to a transitory increase in oil and gas revenue is instead spread fairly evenly across four years.

Flypaper Effect

A voluminous literature finds that local governments increase spending more in response to an increase in unconditional grant revenue than to an equally sized increase in local private income (Inman, 2008). The result is an anomaly under the assumption that policy reflects the preferences of the median voter in the locality (Hines and Thaler, 1995). To address this literature, the appendix reports results from flypaper-style regressions which add oil and gas GDP and non-oil and gas GDP to the regression. The flypaper effect is extremely large—the marginal propensity to spend out of non-oil and gas GDP is only around 0.01, or one cent for every dollar of income, while the marginal propensity to spend out of each grant is at least one.[42] This is unsurprising given district governments' limited ability to tax local income. The propensity to spend out of non-oil and gas GDP exceeds the propensity to spend out of oil and gas GDP. The latter even appears to be negative, albeit statistically insignificant. This suggests that districts endowed with oil and gas engage in countercyclical fiscal policy to some extent.

1.6.2 Provision of Public Goods and Services

This section presents estimates of the effects of the two grants on the provision of public schools, health facilities, and health personnel. I focus on public goods related to education and

[42]For the OLS results, we can reject the hypotheses that each grant produces the same total spending response as each type of GDP at the one-percent level. The 2SLS estimates are less precise—the smaller sample size due to GDP data being available only until 2013 exacerbates the problem. The spending response to oil and gas revenue is statistically distinguishable from the response to GDP at the 10-percent level, though the difference between the responses to the general grant and GDP just misses statistical significance.

health, because these are the areas that district governments have the primary responsibility to address.[43] The analysis is conducted at the village level.[44] Panel B of Table 1.1 provides summary statistics for the village-level variables. The average village population is around 3,410. On average there is 0.94 public schools per 1,000 villagers, and the number of public primary schools is over six times the number of public secondary schools.[45] Villages average 0.11 doctors, 0.54 midwives, and 0.17 primary health care centers (known as *puskesmas*) per 1,000 villagers. The main village road is made of asphalt—as opposed to gravel, dirt, or other materials—in 70 percent of villages. The annual change in public schools and health care centers is 0.01 of either sign on average. A typical district in Indonesia contains hundreds of villages. In the sample of villages successfully merged across all waves of the village census, the average number of villages per district is 204.

I measure each grant variable in terms of average annual revenue per capita over the inter-survey period. Both grants are in units of constant 2010 IDR 1 million (approximately USD 100) per capita. USD 100 per capita is approximately the average increase in the general grant experienced by districts above the 75th percentile in land area per capita, and it would also be a "typical" yearly fluctuation in oil and gas revenue for a district above the 95th percentile in terms of endowment. (See Figures 1.1 and 1.2.) USD 100 per capita is large relative to district income, representing 7 percent of average non-oil/gas GDP per capita (USD 1,453) and 35 percent of average oil/gas GDP per capita (USD 289).

Table 1.8 displays estimates of the effects of the two grants on public goods and services (β and δ in equation (1.2)). Panel A presents the OLS estimates, and Panel B presents the 2SLS estimates. The first column contains the results for the sum of public primary and secondary schools, while columns 2 and 3 present the results for public primary schools and public secondary schools, respectively. The OLS results indicate that increasing the general grant by USD 100 per capita raises the annual change in public schools per 1,000 villagers by 0.008. This means that permanently raising the general grant by USD 100 per capita in the nine-year,

[43]Village governments play a lead role in the upgrading and maintenance of local infrastructure, such as roads, bridges, and piped water systems. Districts contribute to the financing of village infrastructure projects, but in most cases village governments initiate and implement the projects. Furthermore, villages incur the vast majority of local infrastructure costs (World Bank, 2010, pp. 21–22).

[44]In principle, one could aggregate the outcome variables to the district level using one of two methods. First, one could sum the public goods across villages within the district and divide by district population. Second, one could take the average of village public goods per villager within the district. The first method is inappropriate, because the sample of villages is non-comprehensive due to an imperfect merge across survey waves. Furthermore, the percentage of villages that successfully merged varies by district. Consequently, the aggregate district measure would underestimate the true public good provision per capita of the district to a degree that varies according to the merge rate of villages in the district, introducing bias of an unknown form. The second method does not cause bias, but it leads to inefficient estimators. The method would produce numerically identical estimates as the village-level regressions if each district contained the same number of villages. However, because the number of villages varies by district, the aggregate regression loses efficiency by giving the same weight to each district. By contrast, the village-level regression implicitly weights each district by $\sqrt{M_d}$, where M_d is the number of villages in district d. This is the efficient weighting scheme when the village-level disturbance term is conditionally homoskedastic.

[45]Here I define "public schools" as the sum of primary and secondary public schools.

post-reform period (2006–2014) has the cumulative effect of increasing the stock of public schools per 1,000 villagers by 0.072, almost eight percent of the sample mean. The estimate is significant at the 10-percent level. Columns 2 and 3 show that the effect of the general grant on public schools is driven by the effect on secondary schools, which are much less numerous than primary schools. The effect of oil and gas revenue on public schools is −0.002 and is statistically insignificant and statistically distinguishable from the effect of the general grant at the 10-percent level. The 2SLS estimates are larger in magnitude. According to these estimates, an increase in the general grant of USD 100 per capita raises the annual change in schools per 1,000 villagers by 0.029, while the corresponding effect of oil and gas revenue is 0.008. The former estimate is significant at the one-percent level, and the latter is insignificant. Once again, the estimates suggest that the impact of the general grant on public schools is concentrated on secondary schools. The effect of the general grant is economically quite large. Permanently raising the general grant by USD 100 per capita in the post-reform period increases the stock of public schools per 1,000 villagers by 0.261, which is 28 percent of the sample mean. The 2SLS estimates for the two grants are statistically different from each other at the one-percent level for all public schools and public secondary schools.

Columns 4–6 of Table 1.8 present estimates of the impact of the grants on health facilities and personnel. Similar to the results on public schools, the 2SLS estimates are all larger than the OLS results by a factor of roughly two or more. The general grant has a sizable impact on health facilities. According to the 2SLS estimates, increasing the general grant by USD 100 per capita raises the annual change in health care centers per 1,000 villagers by 0.087, which implies an economically large effect. Permanently raising the general grant by USD 100 per capita in the post-reform period increases the stock of health care centers per 1,000 villagers by 0.783, which is 461 percent of the sample mean. Note, however, that the 95-percent confidence interval is wide, ranging from 0.003 to 0.171. The corresponding point estimate for oil and gas revenue is 0.017 and is statistically insignificant. The effects of the two grants on health care centers are statistically different from one another at the five-percent level. The general grant also significantly increased the number of health personnel: increasing the general grant by USD 100 per capita raises the number of doctors and midwives per 1,000 villagers by 0.038 and 0.206, respectively, representing a 35 percent increase in doctors and a 38 percent increase in midwives relative to their respective sample means. Both effects are significant at the five-percent level. The effects of oil and gas revenue on health personnel are an order of magnitude smaller and are statistically indistinguishable from zero. For both doctors and midwives, we can reject the hypothesis that the two grants have the same effect on health personnel at the one-percent level.

Column 7 of Table 1.8 presents estimates of the impact of the grants on road quality. The outcome equals one if the main village road is made of asphalt, and zero otherwise.[46] The outcome measures the quality of an existing road, not the construction of a new road.

[46] The other road-material categories are "gravel," "dirt," and "other."

According to the OLS estimates, increasing the general grant by USD 100 per capita raises the probability of the main road being made of asphalt by 0.015. This effect is statistically insignificant. The corresponding estimate for oil and gas revenue is 0.042 and is significant at the 10-percent level. We fail to reject the hypothesis that the two OLS estimates are equal. The 2SLS estimate of the effect of the general grant is much larger at 0.066 and is statistically significant at the 10-percent level. This effect represents a nine-percent increase relative to the sample mean. The 2SLS estimate for oil and gas revenue is of a similar magnitude—0.051—and is significant at the five-percent level. This effect is seven percent of the sample mean. Once again, we fail to reject the hypothesis that the two grants have the same effect on road quality.

Two lessons emerge from this subsection. First, the general grant induced larger increases in lumpy public goods and services, per dollar of revenue, than the oil and gas revenue. The general grant caused increases in every lumpy category of public goods and services—durable structures and personnel—and these increases were both economically large and statistically distinguishable from the effect of oil and gas revenue. By contrast, the oil and gas revenue had small and statistically insignificant effects on lumpy public goods and services. Thus, the results confirm the theoretical prediction that permanent grant revenue shocks will have a greater impact on public goods that require lumpy investment than transitory grant revenue shocks. Of course, health personnel do not require such an investment. However, adding an additional doctor or midwife requires incurring an upfront fixed cost associated with committing funds toward paying a salary for a year or more. Such a transaction is lumpy compared to, say, raising doctor wages by a small amount. The lumpiness of the personnel hiring decision is magnified by the fact that public workers in Indonesia enjoy significant job security.[47] In contrast to the results for lumpy public goods, the two grants had similar effects on road quality. Road maintenance in one year does not commit the government to maintenance in future years and thus represents a less lumpy outcome. The fact that both grants increase road quality to a similar degree indicates that lumpiness, rather than graft, drives the differential responses of structures and personnel to the two grants.

Second, the discrepancy between the OLS and 2SLS estimates is typically much larger for the general grant than for oil and gas revenue, suggesting that endogeneity concerns are more important for the general grant. This is important for policymakers to keep in mind when evaluating the effects of the general grant, which is the most important source of funding for district governments in Indonesia.

1.6.3 Threats to Validity

One potential concern is that the results for the general grant could simply reflect catch-up growth by more remote, less developed regions. The instrument captures variation in the

[47] In field interviews, public-sector midwives in Yogyakarta said that they could earn significantly more in the private sector but stayed in the public sector due to job security (UNFPA Indonesia, 2014, p. 47).

general grant driven by the increased importance of land area per capita in the allocation formula in the years 2006 and later. If more densely populated districts were experiencing different time trends in outcomes than less densely populated districts for reasons other than the general grant reform, the 2SLS estimates would be asymptotically biased. While districts with different levels of land area per capita may differ in their level of public goods and services, the identifying assumption is that they would have followed parallel trends over time in the absence of the reform. While this assumption is untestable, it produces a corollary which is testable: the partial effect of land area per capita on outcomes should be constant over time prior to the reform. To test this assumption, I run the diagnostic regressions

$$\begin{aligned} Y_{vdis} = &\sum_{j \in \mathcal{J}} \theta_j AreaPC06_{di} \times NonOilGas_{di} \times 1(s = j) \\ &+ \sum_{j \in \mathcal{J}} \gamma_j EndowPC_{di} \times 1(s = j) + \phi_d + \rho_{is} + u_{vdis}, \end{aligned}$$

where $\mathcal{J}$ is the set of time periods s for which data on Y exist.[48] The (omitted) reference period is 2004–2005. The parameters $\{\theta_j\}_{j \in \mathcal{J}}$ capture how outcomes vary over time according to the district's exposure to the general grant reform. Likewise, the parameters $\{\gamma_j\}_{j \in \mathcal{J}}$ reflect how exposure to fluctuations in oil and gas revenue affects outcomes over time. Each parameter represents the partial effect of grant exposure in a time period relative to its effect in 2004–2005, the time period prior to the general grant reform. Panel (a) of Figure 1.5 plots estimates of $\{\theta_j\}_{j \in \mathcal{J}}$, and Panel (b) plots estimates of $\{\gamma_j\}_{j \in \mathcal{J}}$. As shown in the figure, we fail to reject the hypothesis that exposure to the general grant reform had the same impact on outcomes in the period 2001–2003 as it did in the period 2004–2005 at the 95-percent level. This result implies that districts with varying levels of exposure to the general grant reform were on parallel trends prior to the reform, lending credence to the causal interpretation of the main results for public goods and services.

1.7 Conclusion

Indonesia's fiscal decentralization reforms produced large increases in unconditional intergovernmental grants to district governments. The manner in which these grants were delivered depended on district characteristics: districts with greater land area per capita and few natural resources saw a larger permanent increase in general grant revenue starting in 2006. Districts that were richly endowed with oil and natural gas saw little variation in the general grant over this period, but they experienced large swings in the oil and gas revenue that factored heavily into their budgets. Thus, while one set of districts experienced a large shift in permanent grant revenue, another faced frequent transitory shocks to grant revenue. How did the nature of

[48] $EndowPC_{di}$ is average endowment over the sample period.

these fluctuations impact the behavior of local governments?

This paper's theoretical model provides two predictions: both the timing and the composition of the government's expenditure response to a grant revenue shock depend on whether the shock was permanent or transitory. A permanent increase in revenue by one dollar is more likely to allow the government to overcome a minimum size requirement for investment in new durable goods, such as schools, leading to a front-loaded spending response on both durables and nondurables. A transitory increase in revenue by one dollar, on the other hand, has a smaller impact on permanent public income and is less likely to allow the government to overcome the up-front investment constraint, skewing spending towards nondurables which can be varied continuously. Using large-scale policy reforms following decentralization in Indonesia, this paper tests the main predictions of the model. The results confirm that local governments respond to a permanent increase in grant revenue by front-loading expenditure and increasing the provision of lumpy public goods. Transitory shocks to grant revenue elicit flatter fiscal responses and have more modest effects on lumpy public goods.

The results of this paper are informative for policymakers in central governments, which have the option of changing intergovernmental grant allocations on a temporary or permanent basis. The efficiency of a grant reform could depend on its permanence. This is because the permanence of the reform influences the composition of the local fiscal response, and externalities vary by the type of spending. The permanence of a grant reform can also influence aggregate output responses, for at least two reasons. First, it influences the timing of the fiscal response, and output multipliers may vary according to labor market slackness (Michaillat, 2014). Second, it influences the composition of the fiscal response, and government spending on durables and nondurables may have different output multipliers (Boehm, 2016).

This paper also makes an important methodological point: panel-regression estimates of the impact of grants on lumpy public goods will be sensitive to the nature of the identifying variation in grants. This fact poses a challenge to research that compares government responses to different sources of revenue—particularly the recent literature comparing the effects of tax revenue to those of intergovernmental grants (Borge et al., 2015; Gadenne, 2017; Martínez, 2017). In developing countries with low fiscal capacity, local governments are more likely to increase tax revenue via base broadening rather than raising rates. Changes in tax revenue driven by base broadening arguably represent permanent shocks to public revenue, while shocks to intergovernmental grants may be either permanent or transitory. One must determine the nature of the identifying variation in revenue in order to properly interpret estimates of local government responses.

To what extent are the results of this paper informative for other countries? The relative unimportance of local taxation in Indonesia contrasts sharply with the federal systems in many high-income countries, such as the United States. The results may therefore be more applicable to developing countries, where local taxation is less important (Gadenne and Singhal, 2014). Certainly, the *absolute* level of the expenditure response to grant revenue should be lower when there is scope for cutting local taxes, and indeed this is the case (Hines and Thaler, 1995; Inman, 2008). Nonetheless, the results of this paper may be predictive of the *relative* responses to permanent and transitory shocks to local government revenue. National reforms often produce both types of shocks to local taxes. In the United States, the Tax Reform Act of 1986 broadened the definition of taxable income and increased the tax rate on capital gains, among (many) other things. The former reform permanently increased state tax revenue in states that used the federal definition of taxable income (Ladd, 1993), while the latter induced a transitory increase in state tax revenue as capital-gains realizations spiked right before the higher tax rate was to take effect (Auten, 1999). Future research should examine how the time-series properties of grant revenue influence local government responses in contexts with significant local taxation.

1.8 Figures

Figure 1.1: General Grant Revenue per Capita by Land Area per Capita

(a) Non-Oil/Gas Provinces

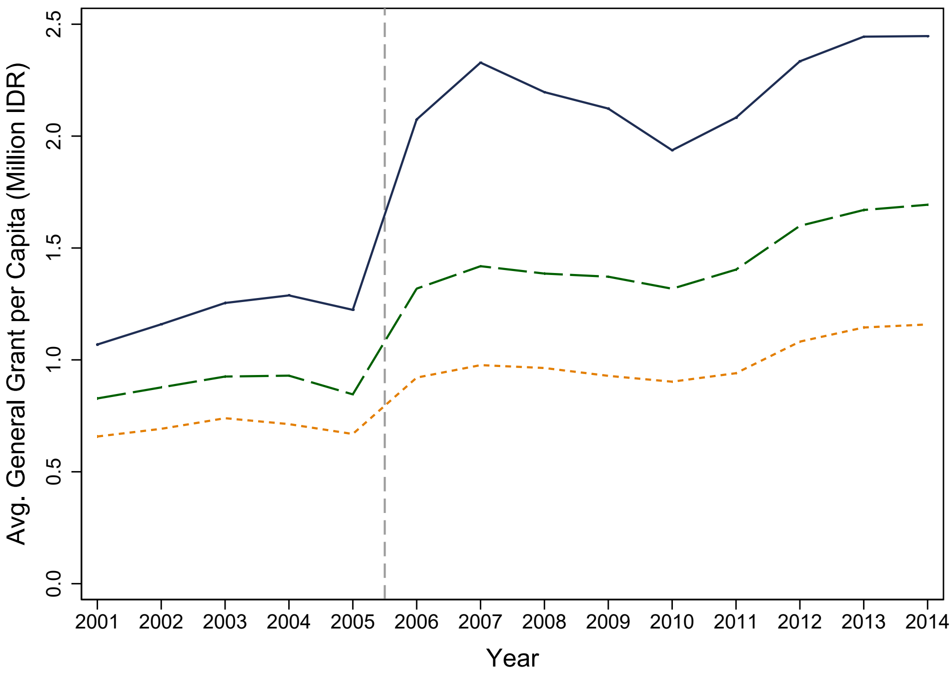

(b) Oil/Gas Provinces

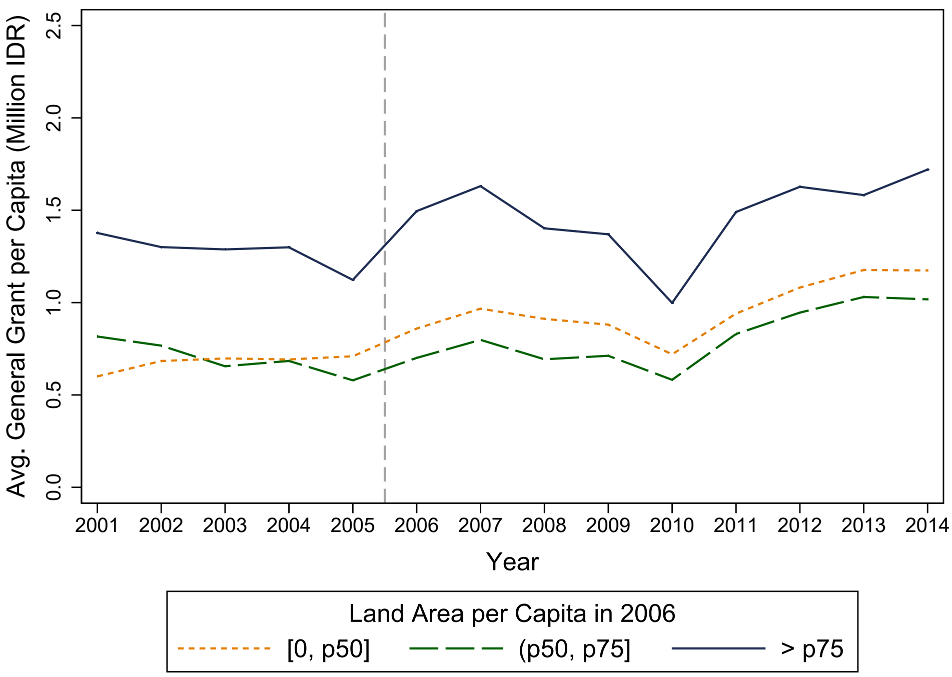

Notes. This figure plots average general grant revenue in constant 2010 rupiah per capita (millions) for districts divided into three groups according to land area per capita in 2006. Panel (a) uses districts located in non-oil-and-gas provinces, and Panel (b) uses districts located in oil and gas provinces. Oil and gas provinces are those that receive a non-trivial amount of oil and gas revenue per capita: Riau, Kepulauan Riau, Jambi, Sumatera Selatan, Kalimantan Timur, and Kalimantan Utara. The gray dashed line indicates the timing of the general grant reform.

Figure 1.2: Shared Oil and Gas Revenue and Value of Production

(a) Aggregate Shared Oil and Gas Revenue and Weighted Value of Production

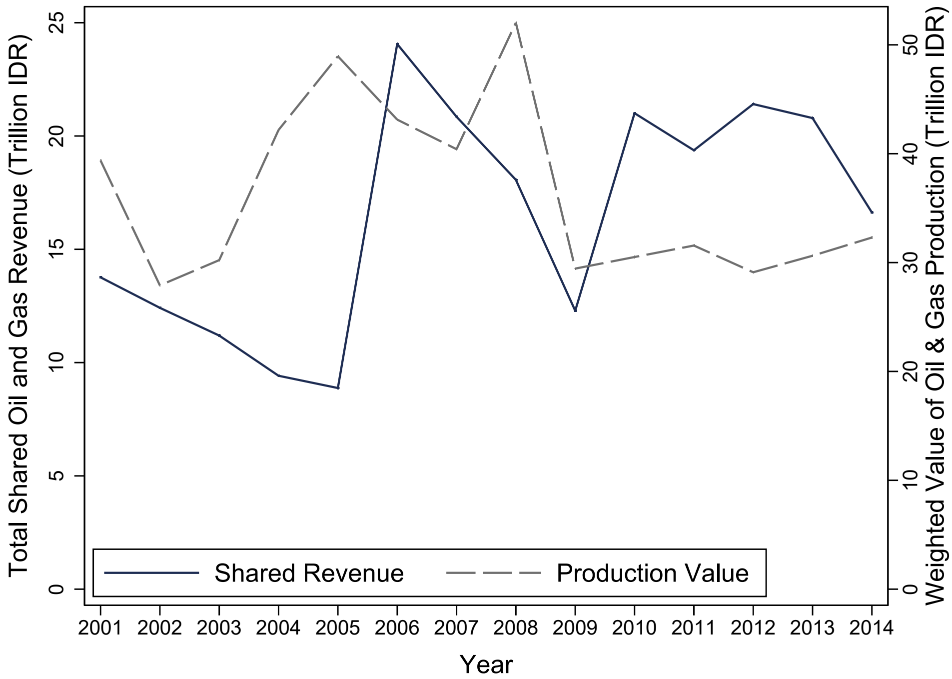

(b) Average Shared Oil and Gas Revenue per Capita by Endowment

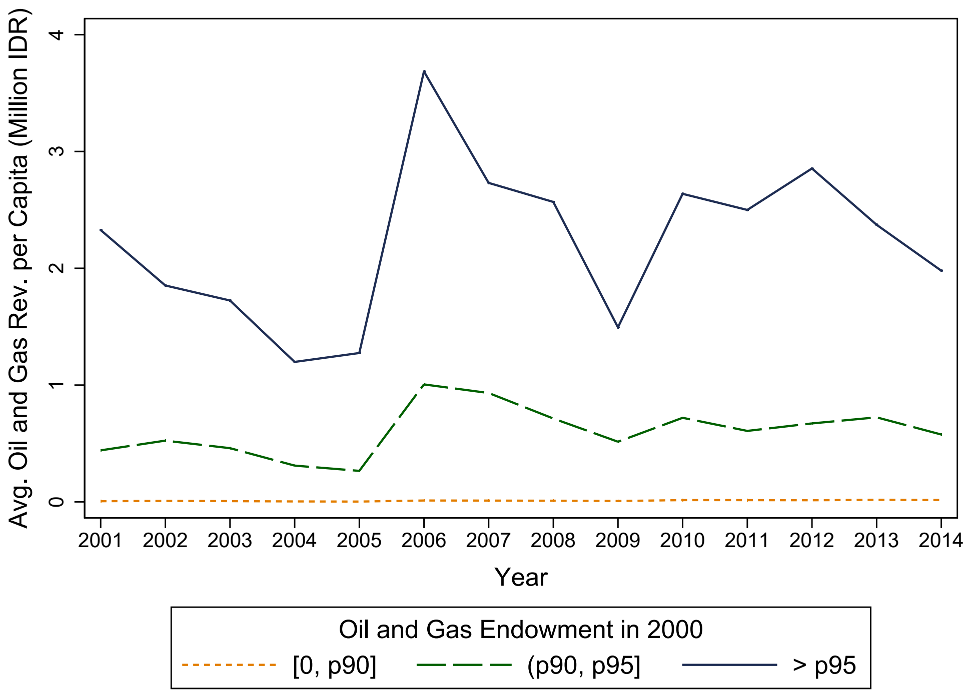

Notes. Panel (a) plots total oil and gas revenue shared with districts (thick line) and the weighted value of oil and gas production (dashed line), defined using the weights from the central government's revenue sharing rule: $0.06 \cdot P_t^{oil} \cdot Q_t^{oil} + 0.12 \cdot P_t^{gas} \cdot Q_t^{gas}$. Both series are expressed in constant 2010 rupiah (trillions). Panel (b) plots average oil and gas revenue in constant 2010 rupiah per capita (millions) for districts divided into three groups according to oil and gas endowment.

Figure 1.3: Total Expenditure Responses to the Two Grants

(a) Total Expenditure per Capita by Area per Capita, Non-Oil/Gas Provinces

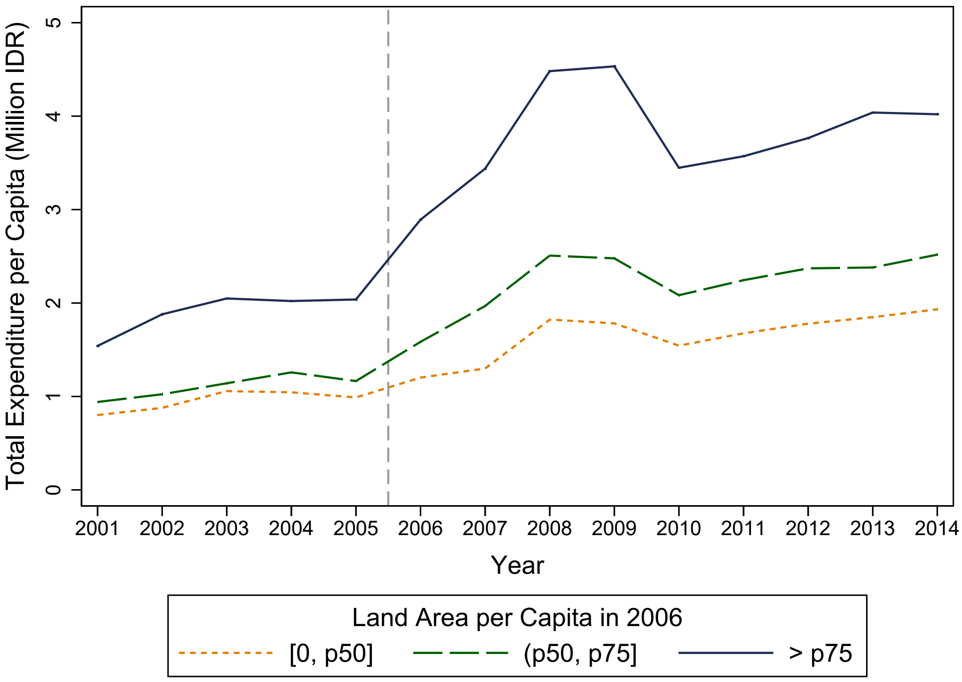

(b) Total Expenditure per Capita by Oil and Gas Endowment

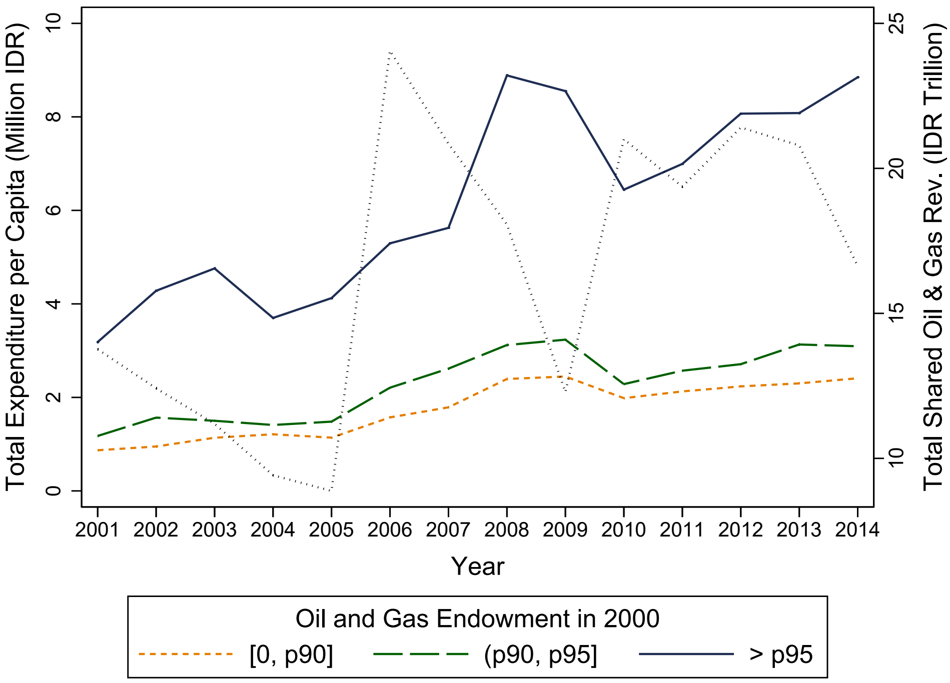

Notes. Panel (a) plots average total expenditure in constant 2010 rupiah per capita (millions) for districts divided into three groups according to land area per capita in 2006. The sample excludes the six provinces with significant oil and gas revenue per capita—Riau, Kepulauan Riau, Jambi, Sumatera Selatan, Kalimantan Timur, and Kalimantan Utara. The gray dashed line indicates the timing of the general grant reform. Panel (b) plots average total expenditure in constant 2010 rupiah per capita (millions) for districts divided into three groups according to oil and gas endowment in 2000.

Figure 1.4: Expenditure Responses to the Grants over Time

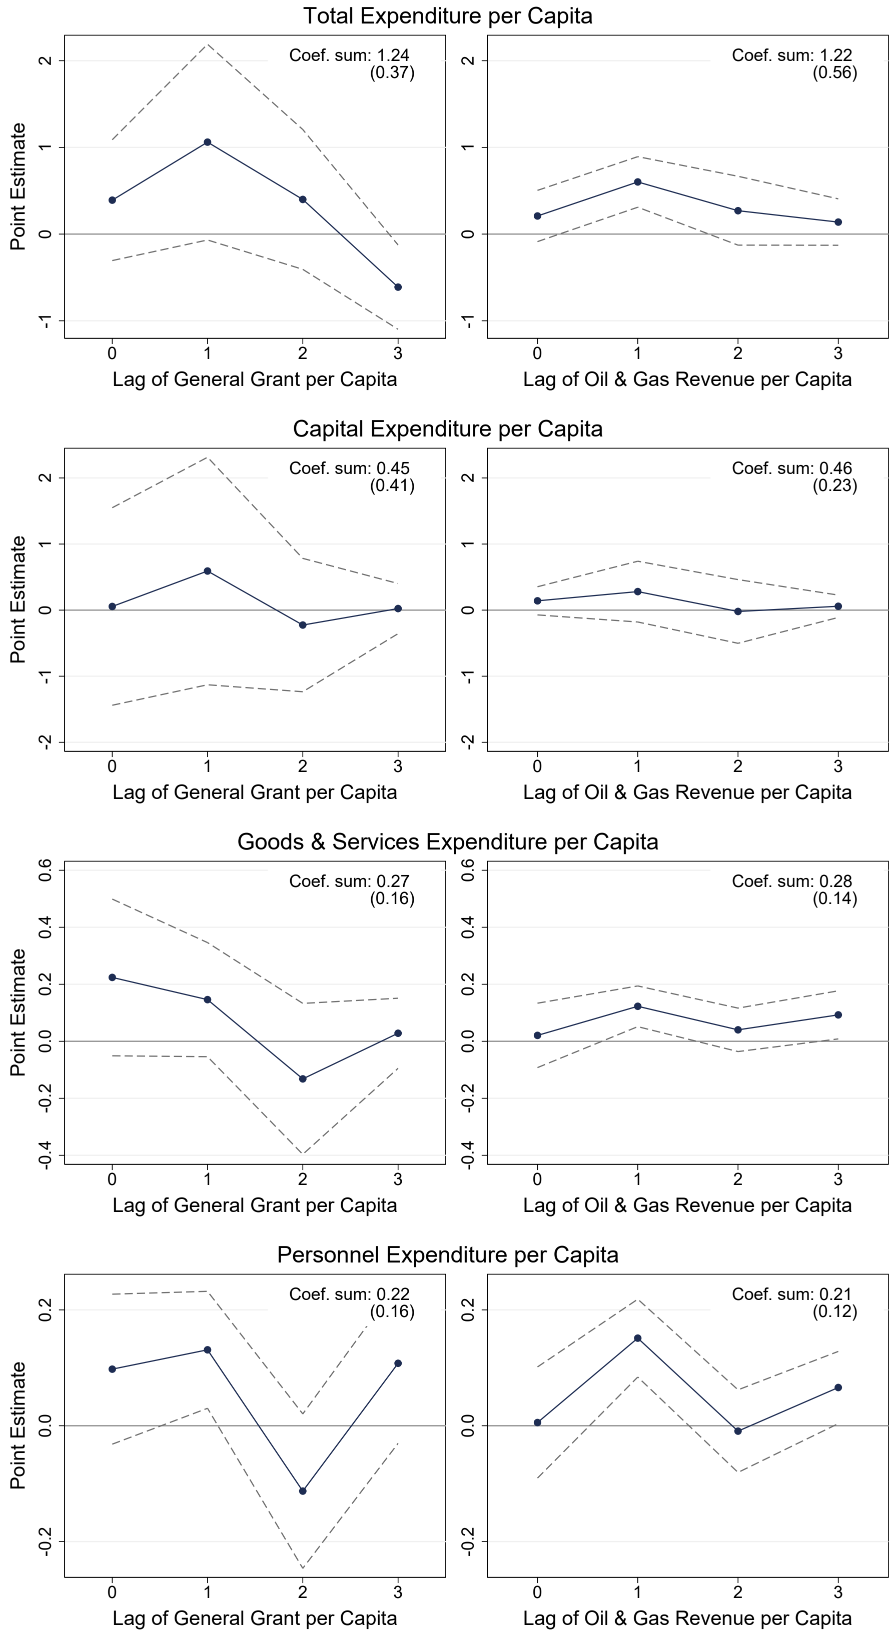

Notes. This figure plots 2SLS estimates and 90-percent confidence intervals of the effects of lags of the general grant (first column) and oil and gas revenue (second column) on expenditure per capita.

Figure 1.5: The Effect of Grant Exposure on Public Goods and Services over Time

(a) The Effect of Land Area per Capita by Year

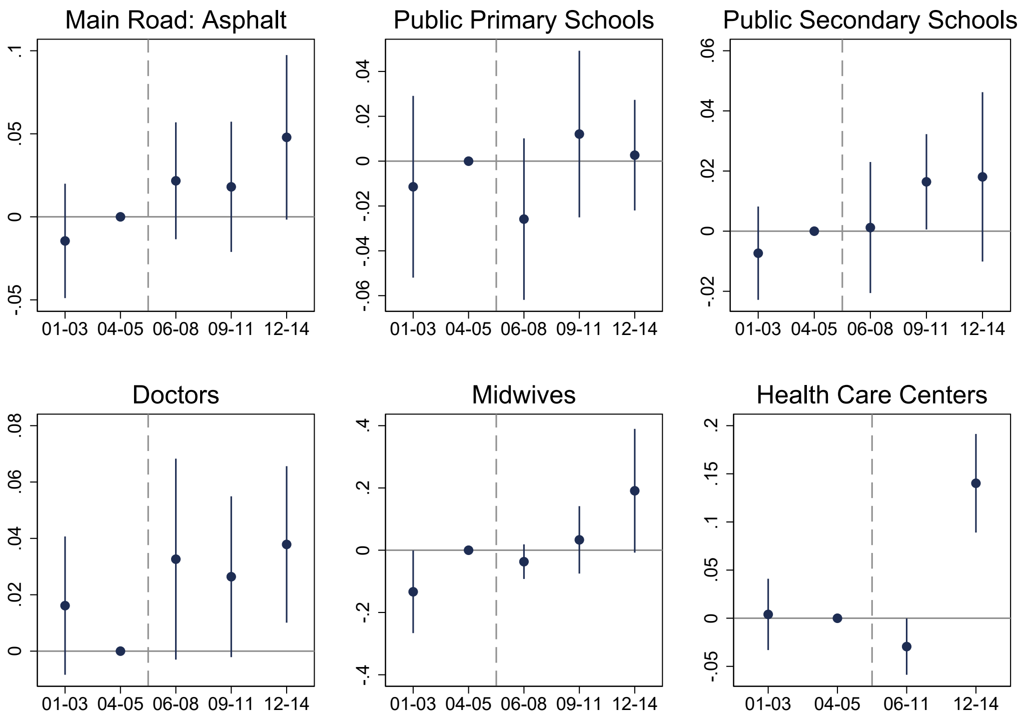

(b) The Effect of Oil and Gas Endowment by Year

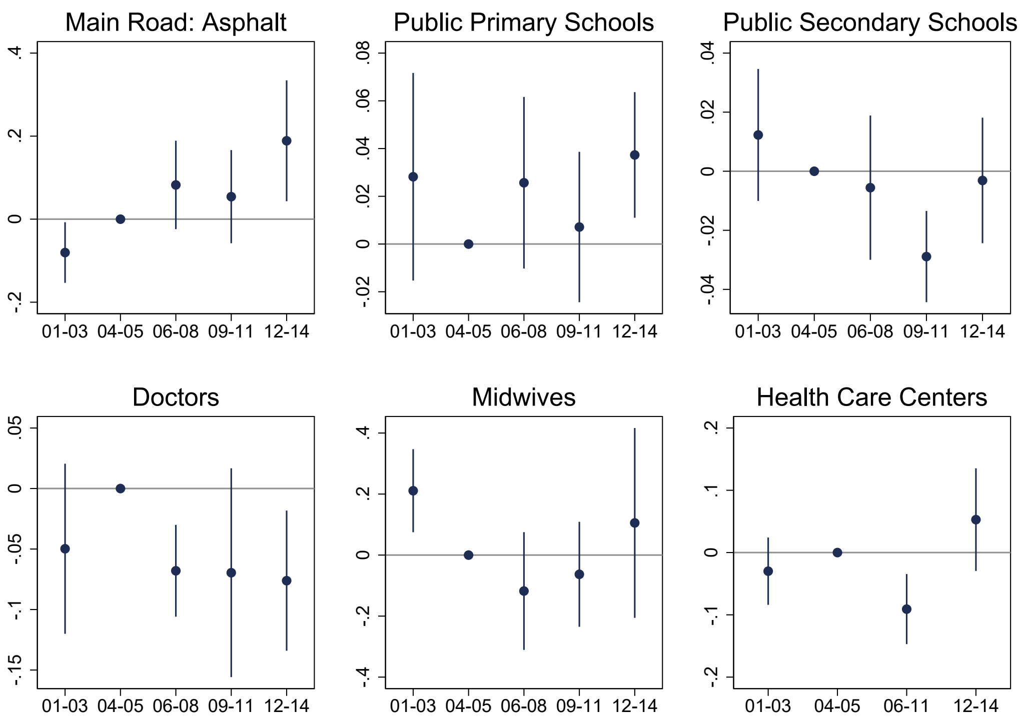

Notes. This figure displays point estimates and 95-percent confidence intervals from the regression $Y_{vdis} = \sum_{j\in\mathcal{J}} \theta_j AreaPC06_{di} \times NonOilGas_{di} \times 1(s=j) + \sum_{j\in\mathcal{J}} \gamma_j EndowPC_{di} \times 1(s=j) + \phi_d + \rho_{is} + u_{vdis}$, where $\mathcal{J}$ is the set of years for which data on Y exist, and $EndowPC_d$ is average endowment over the sample period. The (omitted) reference year is $s = 2005$. Panel (a) plots $\{\theta_j\}_{j\in\mathcal{J}}$, and Panel (b) plots $\{\gamma_j\}_{j\in\mathcal{J}}$.

1.9 Tables

Table 1.1: Summary Statistics

	Mean	Std. Dev.	Min.	Max.	Obs.
Panel A: District-Level Variables					
Total Revenue per Capita	2.11	1.89	0.13	23.71	4,838
Own-Source Revenue per Capita	0.14	0.19	0.00	3.66	4,827
Special Allocation Revenue per Capita	0.13	0.17	0.00	3.30	4,735
Shared Tax Revenue per Capita	0.17	0.42	0.00	15.55	4,696
Total Expenditure per Capita	2.04	1.83	0.01	21.55	4,607
Capital Expenditure per Capita	0.56	0.78	0.00	11.05	4,635
Goods & Services Expenditure per Capita	0.39	0.43	0.00	7.45	4,642
Personnel Expenditure per Capita	0.90	0.57	0.01	6.69	4,655
Education Expenditure per Capita	0.53	0.33	0.00	3.10	3,910
Administration Expenditure per Capita	0.64	0.76	0.01	11.18	3,991
Infrastructure Expenditure per Capita	0.35	0.59	0.00	10.76	3,906
Health Expenditure per Capita	0.16	0.15	0.00	1.80	3,909
Agriculture Expenditure per Capita	0.09	0.11	0.00	1.12	3,892
General Grant Revenue per Capita	1.19	0.87	0.00	7.95	5,003
Oil & Gas Revenue per Capita	0.17	0.66	0.00	10.17	5,003
AreaPC06 × Non-Oil/Gas × Year ≥ 2006	0.10	0.27	0.00	2.72	5,025
Agg. Oil & Gas Rev. × Endow. per Capita	0.46	1.79	0.00	27.54	5,025
Non-Oil/Gas GDP per Capita	14.53	16.18	0.63	262.21	4,652
Oil/Gas GDP per Capita	2.89	24.59	0.00	563.51	4,652
Population (Millions)	0.56	0.59	0.03	5.28	5,025
Panel B: Village-Level Variables					
Public Schools per 1,000 People	0.94	0.71	0.00	11.86	204,485
Public Primary Schools per 1,000 People	0.81	0.57	0.00	8.82	204,485
Public Secondary Schools per 1,000 People	0.13	0.32	0.00	7.91	204,485
Primary Health Care Centers per 1,000 People	0.17	0.36	0.00	6.98	163,070
Doctors per 1,000 People	0.11	0.32	0.00	8.86	204,485
Midwives per 1,000 People	0.54	0.69	0.00	10.71	204,485
Main Road Made of Asphalt	0.70	0.46	0.00	1.00	201,950
Avg. Δ Public Schools per 1,000 People	−0.01	0.14	−3.50	3.14	204,485
Avg. Δ Public Primary Schools per 1,000 People	−0.01	0.12	−2.35	2.77	204,485
Avg. Δ Public Secondary Schools per 1,000 People	0.01	0.08	−2.31	2.04	204,485
Avg. Δ Primary Health Care Centers per 1,000 People	0.01	0.12	−2.45	2.33	163,070
Population (Thousands)	3.41	2.62	0.22	17.83	204,485
Villages per District (Hundreds)	2.04	1.13	0.01	4.57	204,485

Notes. All fiscal and GDP variables are measured in constant 2010 IDR 1 million (≈ USD 100) per capita. Village-level variables are measured per 1,000 villagers.

Table 1.2: First Stage: General Grant and Oil and Gas Revenue

	General Grant p.c.		Oil & Gas Revenue p.c.	
	(1)	(2)	(3)	(4)
AreaPC06 × Non-Oil/Gas × Year ≥ 2006	0.65***	0.65***		−0.00
	(0.17)	(0.17)		(0.02)
Agg. Oil & Gas Rev. × Endow. per Capita		−0.03	0.37***	0.37***
		(0.02)	(0.04)	(0.04)
Observations	5,003	5,003	5,003	5,003
District clusters	372	372	372	372
Prov. × year clusters	384	384	384	384

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.3: Effects of the Grants on Alternative Revenue Sources

Panel A: OLS	Revenue per Capita		
	(1) Own-Source	(2) Special Allocation	(3) Shared Taxes
General Grant p.c.	0.02*	0.11***	−0.01
	(0.01)	(0.02)	(0.04)
Oil & Gas Revenue p.c.	0.02	0.03*	0.36*
	(0.02)	(0.02)	(0.20)
Observations	4,708	4,708	4,577
District clusters	372	372	372
Prov. × year clusters	384	384	384
Test: coefs equal	0.941	0.004	0.027

Panel B: 2SLS	Revenue per Capita		
	(1) Own-Source	(2) Special Allocation	(3) Shared Taxes
General Grant p.c.	0.01	0.02	0.05
	(0.03)	(0.06)	(0.10)
Oil & Gas Revenue p.c.	0.01	0.01	0.16
	(0.02)	(0.02)	(0.10)
Observations	4,708	4,708	4,577
District clusters	372	372	372
Prov. × year clusters	384	384	384
F-stat. Gen. Grant	17.9	17.9	17.5
F-stat. Oil & Gas Rev.	91.9	91.9	92.2
Test: coefs equal	0.830	0.857	0.557

Notes. Panel A presents OLS estimates, and Panel B presents 2SLS estimates. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. Sanderson and Windmeijer (2016) first-stage F-statistics are reported for each endogeneous variable. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.4: Expenditure Responses by Economic Classification (OLS)

	Expenditure per Capita			
	(1) Total	(2) Capital	(3) Goods & Services	(4) Personnel
General Grant p.c.	0.89***	0.39***	0.12***	0.31***
	(0.11)	(0.11)	(0.04)	(0.06)
Lag 1	0.07	0.00	0.05*	0.08**
	(0.15)	(0.10)	(0.03)	(0.04)
Lag 2	0.17**	0.20***	−0.06	0.01
	(0.09)	(0.07)	(0.04)	(0.03)
Lag 3	−0.19*	−0.19***	−0.00	0.07*
	(0.11)	(0.04)	(0.03)	(0.04)
Oil & Gas Revenue p.c.	0.48**	0.19***	0.10	0.08
	(0.21)	(0.05)	(0.09)	(0.07)
Lag 1	0.45***	0.04	0.14***	0.14***
	(0.09)	(0.06)	(0.03)	(0.03)
Lag 2	0.51***	0.15***	0.07*	0.09**
	(0.09)	(0.03)	(0.04)	(0.04)
Lag 3	0.18*	0.02	0.09*	0.12***
	(0.11)	(0.04)	(0.05)	(0.03)
Coef. sum: General Grant p.c.	0.93	0.40	0.11	0.47
	(0.20)	(0.11)	(0.07)	(0.08)
Coef. sum: Oil & Gas Revenue p.c.	1.63	0.41	0.39	0.42
	(0.37)	(0.08)	(0.15)	(0.09)
Test: Coef. sums equal	0.181	0.911	0.124	0.757
Test: Gen. Grant L0 + L1 = L2 + L3	0.000	0.001	0.000	0.000
Test: Oil & Gas L0 + L1 = L2 + L3	0.291	0.501	0.021	0.900
Observations	3,573	3,612	3,614	3,623
District clusters	372	372	372	372
Prov. × year clusters	306	306	306	306

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.5: Expenditure Responses by Economic Classification (2SLS)

	Expenditure per Capita			
	(1) Total	(2) Capital	(3) Goods & Services	(4) Personnel
General Grant p.c.	0.39	0.06	0.22	0.10
	(0.42)	(0.91)	(0.17)	(0.08)
Lag 1	1.06	0.59	0.15	0.13^{**}
	(0.69)	(1.05)	(0.12)	(0.06)
Lag 2	0.40	−0.23	−0.13	−0.11
	(0.49)	(0.61)	(0.16)	(0.08)
Lag 3	-0.61^{**}	0.02	0.03	0.11
	(0.30)	(0.23)	(0.07)	(0.08)
Oil & Gas Revenue p.c.	0.21	0.14	0.02	0.01
	(0.18)	(0.13)	(0.07)	(0.06)
Lag 1	0.60^{***}	0.28	0.12^{***}	0.15^{***}
	(0.18)	(0.28)	(0.04)	(0.04)
Lag 2	0.27	−0.02	0.04	−0.01
	(0.24)	(0.29)	(0.05)	(0.04)
Lag 3	0.14	0.06	0.09^{*}	0.07^{*}
	(0.16)	(0.10)	(0.05)	(0.04)
Coef. sum: General Grant p.c.	1.24	0.45	0.27	0.22
	(0.37)	(0.41)	(0.16)	(0.16)
Coef. sum: Oil & Gas Revenue p.c.	1.22	0.46	0.28	0.21
	(0.56)	(0.23)	(0.14)	(0.12)
Test: Coef. sums equal	0.969	0.958	0.965	0.939
Test: Gen. Grant L0 + L1 = L2 + L3	0.019	0.232	0.057	0.068
Test: Oil & Gas L0 + L1 = L2 + L3	0.134	0.415	0.901	0.032
Observations	3,573	3,612	3,614	3,623
District clusters	372	372	372	372
Prov. × year clusters	306	306	306	306

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.6: Expenditure Responses by Function (OLS)

	Expenditure per Capita				
	(1) Education	(2) Administration	(3) Infrastructure	(4) Health	(5) Agriculture
General Grant p.c.	0.17***	0.18	0.36***	0.06***	0.04***
	(0.05)	(0.11)	(0.12)	(0.02)	(0.01)
Lag 1	0.08***	0.06	0.14***	0.04***	0.03***
	(0.02)	(0.09)	(0.05)	(0.01)	(0.01)
Lag 2	0.01	0.01	0.02	−0.02	0.00
	(0.02)	(0.05)	(0.04)	(0.02)	(0.01)
Lag 3	0.02	−0.05	−0.05**	−0.00	−0.01
	(0.02)	(0.04)	(0.02)	(0.01)	(0.01)
Oil & Gas Revenue p.c.	0.08**	0.26*	0.21**	0.03	0.02*
	(0.04)	(0.15)	(0.10)	(0.03)	(0.01)
Lag 1	0.07*	0.07	0.23***	0.03	0.02***
	(0.04)	(0.15)	(0.05)	(0.03)	(0.01)
Lag 2	0.13***	0.13	0.21**	0.04	0.01
	(0.03)	(0.11)	(0.10)	(0.02)	(0.01)
Lag 3	0.10***	0.35***	0.29***	0.03	0.03***
	(0.03)	(0.12)	(0.07)	(0.03)	(0.01)
Coef. sum: General Grant p.c.	0.29	0.20	0.47	0.09	0.07
	(0.05)	(0.12)	(0.13)	(0.03)	(0.02)
Coef. sum: Oil & Gas Revenue p.c.	0.39	0.81	0.94	0.13	0.09
	(0.09)	(0.45)	(0.15)	(0.10)	(0.03)
Test: Coef. sums equal	0.453	0.248	0.000	0.696	0.653
Test: Gen. Grant L0 + L1 = L2 + L3	0.000	0.025	0.000	0.000	0.000
Test: Oil & Gas L0 + L1 = L2 + L3	0.326	0.243	0.767	0.709	0.259
Observations	2,893	2,948	2,893	2,892	2,879
District clusters	372	372	372	372	372
Prov. × year clusters	249	249	249	249	249

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.7: Expenditure Responses by Function (2SLS)

	Expenditure per Capita				
	(1) Education	(2) Administration	(3) Infrastructure	(4) Health	(5) Agriculture
General Grant p.c.	0.09	1.12***	−0.38	−0.02	0.02
	(0.11)	(0.37)	(0.67)	(0.03)	(0.04)
Lag 1	0.09	−0.92	0.92	0.13***	0.04
	(0.16)	(0.64)	(0.87)	(0.03)	(0.06)
Lag 2	−0.30*	0.16	−0.22	−0.06	0.05
	(0.17)	(0.52)	(0.47)	(0.04)	(0.06)
Lag 3	0.19*	−0.05	0.02	0.03	−0.04
	(0.11)	(0.30)	(0.26)	(0.04)	(0.06)
Oil & Gas Revenue p.c.	0.04	0.04	0.21	0.01	0.01
	(0.03)	(0.18)	(0.20)	(0.02)	(0.01)
Lag 1	0.07*	−0.35	0.36**	0.02	0.02
	(0.04)	(0.23)	(0.17)	(0.03)	(0.01)
Lag 2	0.13**	0.12	0.00	0.02	0.00
	(0.07)	(0.22)	(0.18)	(0.02)	(0.01)
Lag 3	0.09**	0.07	0.32**	0.03	0.02
	(0.04)	(0.23)	(0.15)	(0.04)	(0.01)
Coef. sum: General Grant p.c.	0.07	0.31	0.33	0.08	0.08
	(0.12)	(0.36)	(0.24)	(0.04)	(0.03)
Coef. sum: Oil & Gas Revenue p.c.	0.32	−0.12	0.90	0.08	0.06
	(0.06)	(0.66)	(0.32)	(0.09)	(0.04)
Test: Coef. sums equal	0.112	0.504	0.193	0.981	0.676
Test: Gen. Grant L0 + L1 = L2 + L3	0.209	0.909	0.226	0.004	0.507
Test: Oil & Gas L0 + L1 = L2 + L3	0.275	0.079	0.518	0.494	0.791
Observations	2,893	2,948	2,893	2,892	2,879
District clusters	372	372	372	372	372
Prov. × year clusters	249	249	249	249	249

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.8: Second Stage: Public Goods and Services

Panel A: OLS

	Average Annual Change in Stock per Capita				Personnel per Capita		Main Road
	(1) All Public Schools	(2) Public Primary Schools	(3) Public Secondary Schools	(4) Health Care Centers	(5) Doctors	(6) Midwives	(7) Asphalt
Avg. General Grant p.c.	0.008*	−0.001	0.009***	0.023*	0.019***	0.109***	0.015
	(0.004)	(0.002)	(0.003)	(0.012)	(0.006)	(0.033)	(0.010)
Avg. Oil & Gas Revenue p.c.	−0.002	−0.003	0.000	−0.006	−0.003	−0.033	0.042*
	(0.004)	(0.004)	(0.003)	(0.020)	(0.007)	(0.034)	(0.023)
Observations	204,485	204,485	204,485	163,070	204,485	204,485	201,950
District clusters	372	372	372	372	372	372	372
Prov. × year clusters	137	137	137	109	137	137	137
Test: Coefficients equal	0.052	0.724	0.042	0.190	0.011	0.001	0.233

Panel B: 2SLS

	Average Annual Change in Stock per Capita				Personnel per Capita		Main Road
	(1) All Public Schools	(2) Public Primary Schools	(3) Public Secondary Schools	(4) Health Care Centers	(5) Doctors	(6) Midwives	(7) Asphalt
Avg. General Grant p.c.	0.029***	0.005	0.024**	0.087**	0.038**	0.206**	0.066*
	(0.009)	(0.013)	(0.010)	(0.043)	(0.015)	(0.085)	(0.034)
Avg. Oil & Gas Revenue p.c.	0.008	0.003	0.005	0.017	−0.007	0.001	0.051**
	(0.006)	(0.007)	(0.006)	(0.027)	(0.007)	(0.043)	(0.025)
Observations	204,485	204,485	204,485	163,070	204,485	204,485	201,950
District clusters	372	372	372	372	372	372	372
Prov. × year clusters	137	137	137	109	137	137	137
F-stat. Gen. Grant	24.1	24.1	24.1	26.9	24.1	24.1	20.5
F-stat. Oil & Gas Rev.	69.2	69.2	69.2	49.5	69.2	69.2	66.0
Test: Coefficients equal	0.006	0.825	0.009	0.018	0.000	0.007	0.666

Notes. Panel A presents OLS estimates, and Panel B presents 2SLS estimates. In columns 1–3 and 5–7, the sample includes the time periods 2001–03, 2004–05, 2006–08, 2009–11, and 2012–14. In column 4 the time periods are 2001–03, 2004–05, 2006–11, and 2012–14, due to missing data on health care centers in 2008. The outcomes in columns 1–4 are measured as the average annual change in the stock of the public good per 1,000 villagers over the time period. The outcomes in columns 5–6 are measured as the number of health personnel per 1,000 villagers at the end of the time period. The outcome in column 7 is an indicator variable equal to one if the village main road is made of asphalt, and zero otherwise. The grant variables are measured as the average annual revenue in constant 2010 IDR 1 million (≈ USD 100) per capita at the district level. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. Sanderson and Windmeijer (2016) first-stage F-statistics are reported for each endogeneous variable. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 2

The Long-Run Effects of Oil Wealth on Development: Evidence from Petroleum Geology

2.1 Introduction

Does natural resource abundance promote or hinder economic and political development? Despite decades of research, the question remains largely unresolved.[1] Much of the disagreement owes to the difficulty of identifying exogenous variation in resource wealth.[2] Country-level resource exploration and extraction are endogenous to political, institutional, and economic conditions.[3] Recent contributions to the literature have exploited subnational data and short-term fluctuations in world resource prices in order to identify short-run causal effects of

[1]Early studies argued that resource wealth lowered economic growth via the Dutch Disease (Corden and Neary, 1982; Sachs and Warner, 1995, 1999, 2001), but recent studies call the Dutch Disease hypothesis into question, showing that oil discovery and production can cause positive spillovers for manufacturing and boost aggregate investment and employment (Michaels, 2011; Allcott and Keniston, 2015; Arezki, Ramey, and Sheng, 2016a). Arezki et al. (2016a) and Smith (2015) both present evidence that oil wealth raises GDP, using cross-country panel data. Influential early studies in the political science literature claimed that resource rents promoted authoritarianism (Ross, 2001; Jensen and Wantchekon, 2004). However, Herb (2005) and Haber and Menaldo (2011) argue that there is no robust relationship between oil rents and democracy. See, however, responses to the latter study by Andersen and Ross (2014) and Wiens, Poast, and Clark (2014). Alexeev and Conrad (2009) argue that the negative cross-sectional association between oil and the quality of institutions disappears after controlling for (instrumented) GDP. Brückner, Ciccone, and Tesei (2012) present evidence that oil exports improve democratic institutions. For recent surveys of the resource curse literature, see Ross (2014), van der Ploeg (2011), Frankel (2010), and Torvik (2009). See Cust and Poelhekke (2015) for a survey of the subnational evidence for the resource curse.

[2]See, for example, the discussions in Brunnschweiler and Bulte (2008) and van der Ploeg and Poelhekke (2010).

[3]David and Wright (1997) argue that the United States became the world's premier mineral producer from 1870–1910 not because of a fortuitous mineral endowment relative to other countries, but because its superior technology and institutions allowed it to more efficiently extract resources. Bohn and Deacon (2000) find that democratic institutions and political stability positively affect investment in oil exploration. Cust and Harding (2015) show that when oil is potentially located on a national border, 95 percent more exploratory drilling occurs in the country with relatively better institutions.

resource income.[4] However, several important outcomes, such as the political regime and fiscal capacity of the central government, require analysis at the national level. Furthermore, the interaction between resource wealth and economic and political variables may develop over long periods of time. Political and fiscal institutions develop and consolidate over many years—as do their effects.[5] In addition, both the "greed" and "grievance" motives for conflict (Collier and Hoeffler, 2004) can be deeply rooted in the presence of resource endowments. Therefore the long-run effects of natural resources are of great interest. Understanding how natural resource wealth affects long-run development will inform not only domestic resource policy (e.g., royalties and drilling rights), but also federal transfer policy and foreign aid, as natural resource revenue and other forms of windfall income are believed to have similar effects (e.g., Djankov, Montalvo, and Reynal-Querol, 2008; Brollo et al., 2013).

This paper examines the long-run effects of oil wealth on development using a new identification strategy that exploits the geological characteristics of countries. Hydrocarbons—notably crude oil and natural gas—are produced by the heating and compression of organic matter buried within sedimentary basins. Our instrumental variables approach uses new data on the spatial distribution of sedimentary basins to isolate exogenous cross-country variation in oil wealth.[6]

Addressing endogeneity is crucial in this context, because the sign of the bias of ordinary least squares is *a priori* ambiguous. If wealthier or more democratic countries attract greater private investment in resource exploration and production, perhaps due to their stronger property-rights protections, then the estimated effect of resource wealth on development will be biased upwards (Cust and Harding, 2015). On the other hand, if low-income or less democratic countries have more lax regulation of the resource sector or are governed by politicians who personally benefit from rapid extraction rates, then the estimate will be biased downwards (Robinson, Torvik, and Verdier, 2006).

Other studies have used instrumental variables (Brunnschweiler and Bulte, 2008; van der Ploeg and Poelhekke, 2010; Tsui, 2011; Borge, Parmer, and Torvik, 2016), price shocks (Brückner et al., 2012; Dube and Vargas, 2013; Andersen et al., 2014; Caselli and Tesei, 2016; Carreri and Dube, 2015), giant oil field discoveries (Lei and Michaels, 2014; Smith, 2015; Arezki et al., 2016a), and quasi-experimental policy variation (Litschig, 2012; Brollo et al., 2013) to estimate the causal effects of natural resource abundance. Studies based on panel variation have the advantage of controlling for unit fixed effects, though the quality of institutions could influence which

[4]Subnational studies include Vicente (2010), Michaels (2011), Litschig (2012), Monteiro and Ferraz (2012), Aragón and Rud (2013), Brollo et al. (2013), Caselli and Michaels (2013), Dube and Vargas (2013), Allcott and Keniston (2015), Aragón and Rud (2015), Aragón, Rud, and Toews (2015), and Carreri and Dube (2015). For empirical strategies that exploit price shocks, see Brückner et al. (2012), Dube and Vargas (2013), Andersen, Johannesen, Lassen, and Paltseva (2014), Caselli and Tesei (2016), and Carreri and Dube (2015).

[5]See Besley and Persson (2011) for a model of fiscal capacity as a stock variable, and see Persson and Tabellini (2009) on the implications of democratic capital.

[6]Bartik, Currie, Greenstone, and Knittel (2017) use an index of geological suitability for hydrocarbons accessible by fracking to predict the prevalence of fracking at the U.S. county level.

countries discover resources—and when. In Section 2.4 we argue that this paper provides the most credible strategy thus far for estimating long-run effects of oil in cross-country samples. Previous cross-country studies have used initial subsoil natural resources as an instrument for resource wealth (see, e.g., van der Ploeg and Poelhekke, 2010; Tsui, 2011). However, these measures of known resource endowment depend on exploration effort, which endogenously responds to economic and political conditions. In contrast, our instrumental variable measures geophysical properties that are not influenced by economic or political factors. Furthermore, conditional on an appropriate set of geographic controls, sedimentary basins affect development only through the channel of oil wealth.

The instrumental variables estimates indicate that an increase in average annual oil production from 1966–2008 significantly reduces the level of democracy in 2008 as well as the average level of democracy from 1966–2008. Increasing oil production also leads to more purges of political rivals and reduces average tax revenue as a share of GDP from 2000–2008. The corresponding OLS estimates understate the negative effects of oil, suggesting that countries with better political institutions and greater state capacity disproportionately select into oil production. The evidence on corruption, internal armed conflict, and coup attempts is less conclusive. Finally, we find evidence that oil production raises GDP. The results are consistent with recent research showing that oil negatively impacts political institutions without leading to noticeably worse economic outcomes on average (Ross, 2012). The results are robust to controlling for a wide variety of geographic covariates and initial population and changing the construction of the instrument and oil wealth variable.

Several studies have argued that natural resources have heterogeneous effects which depend on country-specific factors, such as institutions.[7] Following this literature, we test for heterogeneous effects, finding that the negative long-run effects of oil wealth on democracy and tax revenue are concentrated in the subsample of countries with weak institutional constraints on executive decision-making from 1950–1965. Interestingly, countries with weak executive constraints from 1950–1965 benefited the most from oil in terms of income, probably reflecting the fact that lower-income countries have the highest potential GDP gains from oil (Smith, 2015). We view the evidence on heterogeneous effects as suggestive rather than causal, because initial institutions may be correlated with unobserved country characteristics which affect modern-day outcomes.

The results on heterogeneous effects of oil on democracy are most similar to those of Tsui (2011) and Caselli and Tesei (2016), who find that resource wealth causes non-democracies to become less democratic but has no effect on the political regime in democracies. Unlike those studies, however, we condition on initial rather than contemporary political institutions to

[7] For the argument that the effect of natural resources on income depends on the quality of institutions, see, e.g., Lane and Tornell (1996), Tornell and Lane (1999), Mehlum, Moene, and Torvik (2006), Robinson et al. (2006), and Boschini, Pettersson, and Roine (2007). Other studies emphasize that resource rents influence politician behavior in different ways depending on preexisting political institutions; see, e.g., Aslaksen and Torvik (2006), Bhattacharyya and Hodler (2010), Tsui (2011), Andersen and Aslaksen (2013), and Caselli and Tesei (2016).

(partially) alleviate concerns about the endogeneity of political institutions. Theory predicts that natural resource wealth will have heterogeneous effects on corruption and conflict depending on the quality of institutions (Bhattacharyya and Hodler, 2010; Besley and Persson, 2011). However, our empirical results provide little support for these predictions. Our finding that oil wealth reduces fiscal capacity is related to the theoretical predictions of Besley and Persson (2009a, 2010); Besley and Persson (2011) and is consistent with previous empirical studies (Jensen, 2011; Cárdenas, Ramírez, and Tuzemen, 2011). To our knowledge this is the first paper to empirically test how the effect of oil on tax revenue depends on initial institutions. Recent research on fiscal capacity and natural resources emphasizes the role of the marginal value of public funds,[8] however our results are more consistent with a "rentier state" model (Mahdavy, 1970; Ross, 2001) which focuses on an autocrat's ability to use public finance to produce a quiescent population.[9]

The paper proceeds as follows. Section 2.2 provides background information on petroleum geology and describes the construction of the instrumental variable. Section 2.3 describes the data, Section 2.4 describes the identification strategy, Section 2.5 presents the main results, Section 2.6 discusses the evidence of heterogeneous effects, and Section 2.7 concludes.

2.2 Petroleum Geology and Instrumental Variables

2.2.1 Formation of Hydrocarbons

This section provides a brief overview of petroleum geology and defines the instrumental variable. There are five geological prerequisites for oil reservoir formation. First, there must be a *source rock*, a sedimentary rock rich in organic material deposited by algae and zooplankton millions of years ago. Source rocks form within a sedimentary basin—a region of the Earth's crust characterized by prolonged subsidence, in which tectonic movements cause the surface area to sink and sediments from surrounding regions to fill in the depressed area (Southard, 2007). Extreme heat and pressure convert the buried organic material into hydrocarbons, notably natural gas and crude oil (Kvenvolden, 2006). Second, a *migration pathway* must connect the source rock to an area where the reservoir will form. For example, this migration pathway may be a fracture caused by seismic activity. Third, a *reservoir rock* must be located along the migration pathway. This highly porous and permeable rock, usually a sandstone or carbonate, collects and absorbs the migrating hydrocarbons (Chen, 2009). Fourth, a highly impermeable *caprock* must seal the hydrocarbons within the reservoir rock, preventing the hydrocarbons from leaking to the surface and dissipating. The final requirement is the presence of what is known as a *trap*, which concentrates the hydrocarbons in specific locations where

[8] See Besley and Persson (2011) and Jensen (2011).
[9] A formal version of this model is in the appendix.

they can be exploited (Allen and Allen, 2005).[10] See the appendix for illustrations.

2.2.2 Sedimentary Basin Classification

The Fugro Robertson, Ltd. (2013) Tellus GIS database provides the name, location, description, and geological classification of every onshore and offshore sedimentary basin. See Figure 2.1 for a map of the basins. Geologists rely on three general techniques to collect data on sedimentary basins: (i) surface mapping, (ii) core sampling, and (iii) subsurface geophysics such as seismic profiling (Southard, 2007). Aerial photographs provide a base map of the surface, and survey work on the ground complements the photographs in the construction of surface maps (Marjoribanks, 2010, ch. 2). Core sampling involves the removal of a cylindrical piece of subsurface material using a drill. Geologists use seismic air guns to initiate seismic waves underground. They use seismic detectors to record the arrival of the waves at different points under the surface. Geologists then use the data collected by the seismic detectors to draw seismic profiles (Britannica, 2015).

Fugro Robertson, Ltd. (2013) divides sedimentary basins into 24 classification groups according to their plate-tectonic environment, primary mechanism of subsidence, and other details regarding the nature of faulting and subsidence and the relative location of the basin on the tectonic plate. Each basin forms in one of three general plate-tectonic environments. The first is a divergent environment, in which adjacent tectonic plates pull away from each other. The second is a convergent environment, in which tectonic plates collide head on, causing one plate to pass underneath the other in a process known as subduction. Convergent environments are further divided according to whether they feature continental plates, oceanic plates, or both. The third is a wrench environment, in which adjacent tectonic plates move in opposite, parallel directions, rubbing alongside each other. The mechanism of subsidence is mechanical (a.k.a., "tectonic"), thermal, or thermo-mechanical. Mechanical subsidence is caused by the movement of tectonic plates due to faulting. Thermal subsidence is caused by the thickening of the Earth's crust due to cooling of the underlying mantle, which causes the crust to become denser than its surroundings. Thermo-mechanical subsidence is caused by some combination of the aforementioned mechanical and thermal processes.

Table 2.1 lists the name, classification code, and plate-tectonic environment ("sub-regime") of each of the 24 Fugro Robertson basin types. The classification code consists of two or three elements. The first element indicates the general plate-tectonic environment. It takes the value of "D" for "Divergent," "C" for "Convergent," and "W" for "Wrench." For codes consisting of three elements, the second element indicates the involvement of continental tectonic plates, oceanic tectonic plates, or both. A second-element value of 1 indicates the presence of two continental plates, 2 indicates the presence of one continental and one oceanic plate, and 3 indicates the presence of two oceanic plates. For example, a basin with code starting with "C.1"

[10] I am indebted to Mike Waite, a former geophysicist at Chevron, for explaining this process to me.

exists in an environment in which two continental plates are converging, while a basin with code starting with "C.2" exists in an environment in which a continental plate and an oceanic plate are converging. For codes consisting of three elements, the third element indicates the location of the basin relative to the plates and areas of faulting. For example, codes ending in "F" indicate foreland basins, which are formed adjacent to a mountain range caused by the subduction of two plates. The code "C.1.F" corresponds to a foreland basin formed in the context of two continental plates colliding, while "C.3.F" corresponds to a foreland basin formed from the collision of two oceanic plates. To give another example, codes ending in "E" indicate extensional basins, which are formed in areas characterized by the stretching of the crust or lithosphere. For codes consisting of only two elements, the second element indicates the location of the basin relative to the plates and areas of faulting. In sum, the final element of the code indicates local characteristics of the basin formation, while the preceding elements of the code indicate global characteristics of the plate-tectonic environment.

Figure B.3.6 in the appendix displays diagrams for two common basin types. The first basin type, C.1.F or "peripheral foreland basin," exists in a convergent plate-tectonic environment and is characterized by a mechanical subsidence mechanism. Peripheral foreland basins are found adjacent to mountain ranges formed by the subduction of two continental plates. Large peripheral foreland basins exist in the Persian Gulf and Arabian Peninsula, adjacent to the Zagros mountains in Iran. The second basin type, D.4 or "passive margin basin," forms within a divergent plate-tectonic environment and features a thermal subsidence mechanism. Passive margins occupy areas where an oceanic plate and a continental plate have diverged, such as the eastern coastlines of the Americas and all coastlines of Africa, among other places.

2.2.3 Instrument Construction

The next task is to specify the candidate instrument sets. The composition of each instrument set depends on two choices. The first choice is how to aggregate the 24 Fugro Robertson basin categories into a smaller number of exhaustive and mutually exclusive basin categories. Aggregating the basin categories is reasonable *a priori* as many of the disaggregated categories account for a very small fraction of the earth's surface area and thus are unlikely to have much predictive power. The second choice is which aggregate basin categories to include in the set of instruments. Section 2.4 describes the instrument selection procedure.

We pursue two approaches to basin aggregation. The first is based on the global characteristics of the basin environment—the general plate-tectonic environment and primary mechanism of subsidence. Fugro Robertson, Ltd. (2013) provides a grouping that assigns each basin type to one of five plate-tectonic environments—divergent, convergent continent-continent, convergent ocean-continent, convergent ocean-ocean, and wrench—and one of three subsidence categories—mechanical, thermo-mechanical, and thermal. This method results in eight groups of basin types that actually exist, as shown in Table 2.2.[11] The second approach is based on the local characteristics of the basin as indicated by the final element of the Fugro Robertson, Ltd. (2013) code. As already mentioned, the local characteristics involve the location of the basin relative to the plates and areas of faulting. This approach produces ten basin groups, as shown in Table 2.3. The appendix provides maps of the aggregated basin categories.

We assign values of each aggregate basin type to countries by calculating the log of the sovereign area (in square kilometers) per 1000 inhabitants in 1960 covered by the basin.[12] Sovereign territory is inclusive of maritime boundaries. Data on country land borders are from Erle and Gilles (2013), and data on maritime borders are from the Flanders Marine Institute (2013).[13]

2.3 Other Data Sources

This section describes the other data sources used in the empirical analysis. The sample period is 1966–2008.[14] Data on oil production, our primary measure of oil wealth, come from Ross (2013), who cleaned and compiled data from the U.S. Geological Survey, the U.S. Energy Information Administration's International Energy Statistics, the World Bank, and the BP Statistical Review. This dataset covers 172 countries, of which 96 have produced oil, from 1932–2011.[15] Oil production is measured as the log of average annual metric tons per 1000 inhabitants from 1966–2008. We also use data on oil discoveries compiled by the Association for the Study of Peak Oil (ASPO) and provided by Cotet and Tsui (2013). This dataset, first used by Tsui (2011), includes the amount of oil discovered each year in the 64 top oil-producing countries over the period 1930–2003. It also contains oil production data for the years 1930–2008. Discovery by countries not included in the ASPO dataset is imputed as the sum of cumulative production and remaining publicly reported reserves. Oil wealth in these countries is modest compared to that of the 64 top oil-producing countries. Oil discovery is measured as the log of average annual millions of barrels per 1000 inhabitants from 1966–2003. Discovery is measured until 2003, instead of 2008, due to data availability. However, the measures are comparable in the sense

[11] Basins with convergent ocean-ocean tectonics and thermal subsidence covered only 1,331 square kilometers of sovereign area among countries in the sample, which is several orders of magnitude less than any other basin group defined by the tectonic environment and subsidence mechanism. These basins exists in essentially just one country included in the sample. (St. Kitts and Nevis contains 1,329 square kilometers of this basin type, while Venezuela contains two square kilometers.) We therefore combine these basins with those with convergent ocean-ocean tectonics and mechanical subsidence.

[12] All geographic variables are normalized by population in 1960, prior to the sample period, because population may be endogenous to oil production through changes in migration (Michaels, 2011) or fertility (Ross, 2008).

[13] All geographic calculations use the Cylindrical Equal Area projected coordinate system, which preserves area measure.

[14] The sample ends in 2008 to avoid the depths of the Great Recession.

[15] An advantage of this dataset is that it also includes information on oil exports as well as natural gas production and exports. Natural gas often accumulates near crude oil reservoirs, so the sedimentary basin instrument also predicts natural gas endowment. The empirical analysis focuses on oil production to facilitate comparison to past studies, however the results are very similar when the explanatory variable is oil and gas production.

that the average lag between discovery and the start of production is roughly five years.[16]

Data on oil quality come from the World Oil and Gas Review (ENI, 2015), which defines ten ordered categories of oil, ranging from lowest quality ("heavy and sour") to highest quality ("ultra light"). The report provides the volume of production of each category of oil for 54 top oil-producing countries in the years 2000, 2005, and 2010–2014. To construct a variable representing oil quality, we assign numbers to each category, ranging from 1 (no oil) to 11 (ultra light). We then take the average oil quality weighted by production volume for each country and year. For most countries, our final measure of oil quality is based on production in 2000, though we use data from 2005 or 2010 for a few countries with missing data for 2000. The oil quality variable is available for 54 large oil producers and non-producers; it is missing for small oil producers.

To ensure that the basin instrument satisfies the exclusion restriction, we include controls for geographic features that are possibly correlated with both sedimentary basins and economic and political outcomes. The basin variable will naturally be correlated with the physical size of the country, so we include a control for total land area calculated from GIS data. Gallup, Sachs, and Mellinger (1998) show that countries with more land in the tropics and less access to waterways tended to grow more slowly over their sample period. We use their data to construct a measure of land area in the tropics. Data on country coastline are obtained from the CIA World Factbook (CIA, 2015). We also use data on the area of mountainous land from Fearon and Laitin (2003), who argue that mountainous terrain is associated with higher levels of insurgency and civil war. Finally, we control for soil quality, which could influence development directly through its effect on agricultural productivity, or indirectly through the division of labor and the evolution of gender norms (Alesina, Giuliano, and Nunn, 2013). We use the FAO's Global Agro-Ecological Zones (GAEZ) database (Fischer, van Velthuizen, Shah, and Nachtergaele, 2002) to calculate each country's land area containing "good" soil.[17] Soil quality depends on nutrient availability, nutrient retention capacity, rooting conditions, oxygen availability, presence of excess salts, toxicity, and workability.

As with the *Basin* variables, all geographic controls measuring surface area are expressed as the log of the surface area (in square kilometers) per 1000 inhabitants in 1960. The coastline variable is expressed as the log of the coastline (in kilometers) per 1000 inhabitants in 1960.[18] Data on population come from Maddison (2013).

We measure democracy using the standard POLITY2 index from the Polity IV database (Marshall and Gurr, 2014), which depends on qualities of executive recruitment, constraints

[16] Arezki et al. (2016a) calculate an average delay of 6.4 years for offshore discoveries and 4.6 years for onshore discoveries. The average delay for all discoveries is 5.4 years.

[17] The GAEZ database divides zones according to the moisture regime (dry, moist, sub-humid, or humid) and soil quality (good, moderate, or poor). We define "good soil" as soil with "good" quality falling in any of the moisture regimes. We use the most recent version of the database available, version 3.0.

[18] Due to the presence of zero values, each "log" transformation in the empirical analysis is in fact a differentiable and monotonic transformation $h(w) = \log(w)$ for $w > w_0$ and $h(w) = \log(w_0) - 1 + w/w_0$ for $w \leq w_0$. In practice w_0 is set equal to the minimum positive value of the random variable observed in the sample.

on executive authority, and political competition. The index takes integer values from −10 to 10. POLITY2 codes cases of foreign "interruption" as missing and cases of "interregnum," or anarchy, as zero. Furthermore, the POLITY2 score is prorated starting from zero during periods of transition following interruption or interregnum. This can give the false impression that, say, a period of anarchy in an autocratic country represents a movement towards democracy. We follow the recent literature (Brückner and Ciccone, 2011; Caselli and Tesei, 2016) and code periods of interregnum as missing. Furthermore, we prorate the score during periods of transition starting from the most recent non-missing POLITY2 score. We normalize POLITY2 to take values between zero and one, with one being the most democratic. Two different democracy outcomes are used: (1) democracy in 2008 and (2) average democracy from 1966–2008 in years in which the country was independent. The measure of executive constraints is the XCONST variable from the Polity IV database, also normalized to take values between zero and one. This variable measures the "extent of institutionalized constraints on the decision-making powers of chief executives," where the constraints can be imposed by any accountability group (Marshall and Gurr, 2014).

Data on corruption and conflict come from several sources. Our corruption measure comes from the Political Risk Services and focuses on corruption within the political system.[19] The index ranges from zero to six, with higher numbers indicating less corruption. We measure corruption in 2008. Three variables capture different aspects of political conflict. First, we use the UCDP/PRIO dataset (Gleditsch, Wallensteen, Eriksson, Sollenberg, and Strand, 2002) to calculate the number of internal or internationalized internal armed conflicts per year in which the country was independent from 1966–2008. The dataset counts only conflicts in which the government is a party and which involve at least 25 battle-related deaths. Second, we use the Polity IV database (Marshall and Marshall, 2016) to count the number of (failed or successful) coup attempts per year in which the country was independent from 1966–2008.[20] Finally, we use the dataset by Banks and Wilson (2016) to calculate the number of purges per year in which the country was independent from 1966–2008.[21]

Revenue data come from the ICTD Government Revenue Dataset, compiled by Prichard, Cobham, and Goodall (2014) on behalf of the International Centre for Tax and Development (ICTD). The series covers the period 1980–2013 for 204 countries, although a nontrivial amount of data are missing, particularly in earlier years. Previously available cross-country tax and revenue datasets were plagued by many missing observations, inconsistent accounting definitions, and inadequate decomposition of tax and revenue by source, among other problems.

[19] According to the Political Risk Services, the measure accounts for excessive patronage, nepotism, job reservations, 'favor-for-favors,' secret party funding, and suspiciously close ties between politics and business.

[20] A coup is defined as a "forceful seizure of executive authority and office by a dissident/opposition faction within the countryfis ruling or political elites that results in a substantial change in the executive leadership and the policies of the prior regime (although not necessarily in the nature of regime authority or mode of governance)" (Marshall and Marshall, 2016).

[21] A purge is defined as "any systematic elimination by jailing or execution of political opposition within the ranks of the regime or the opposition" (Banks and Wilson, 2016).

In particular, accounting treatment of natural resource revenue is notoriously variable across countries, making cross-country analysis difficult. The authors of the ICTD dataset combined and manually cleaned data from several international databases, improving data coverage and consistency. For the purposes of this paper, the ICTD dataset is particularly valuable because it is based on a standardized approach to revenue from natural resources.[22] We focus on two government revenue outcomes: total revenue and tax revenue. All revenue variables exclude social contributions. Total revenue is the sum of all tax and non-tax revenue. Crucially, total revenue includes both resource tax revenue (e.g., corporate taxes paid by private natural-resource firms) and non-tax resource revenue (e.g., royalties paid by private companies and profits from state-owned natural-resource companies). Following the ICTD classification, tax revenue is defined as the sum of all non-resource tax revenue.[23] To maximize sample size and smooth out fluctuations due to business cycles, revenue variables are measured as the log of their average share of GDP from 2000–2008.

Fiscal capacity—the state's maximum administrative ability to collect tax revenue—is unobservable. Following the empirical fiscal-capacity literature (Besley and Persson, 2011; Jensen, 2011; Cárdenas et al., 2011), we use tax revenue as a proxy for fiscal capacity. Tax revenue collection requires investment in tax administration and entails higher information and enforcement costs than other forms of revenue, such natural-resource royalties (Besley and Persson, 2011). We thus expect variation in tax revenue to largely reflect variation in the state's administrative capacity to collect taxes.

We measure the log of GDP per capita in 2008 (constant 2011 international dollars) using the World Bank's World Development Indicators. When we examine the effect of oil on GDP in different years starting in 1960, we use the dataset of Maddison (2013), which has greater data coverage for early years.[24]

[22]Despite the extensive efforts made to construct a reliable dataset, some problems remain due to the limitations of primary sources. In some cases the data appear not credible, and in other cases it is impossible to isolate natural resource revenue from other types of revenue. These problematic observations are flagged in the dataset and are excluded from the empirical analysis.

[23]This definition is conceptually appealing, as we are interested in how resource wealth affects investments in state capacity. Taxing a few large resource firms requires much less administrative capacity than, say, enforcing a personal income tax.

[24]Maddison (2013) measures GDP per capita in constant 1990 U.S. dollars.

2.4 Identification Strategy

2.4.1 Estimating Equations

This section describes the identification strategy. We estimate the effect of oil wealth on country outcomes using sedimentary basin areas as instruments. The estimating equations are

$$y_{cr} = \beta Oil_{cr} + \boldsymbol{\delta}'\boldsymbol{x}_{cr} + \eta_r + \varepsilon_{cr}$$

$$Oil_{cr} = \boldsymbol{\pi}'\boldsymbol{Basin}_{cr} + \boldsymbol{\phi}'\boldsymbol{x}_{cr} + \nu_r + \xi_{cr},$$

where c indexes countries and r indexes regions. The variable y represents a country-level outcome, such as level of democracy or tax revenue. *Oil* is a measure of average annual per capita oil wealth over the period of interest. $\boldsymbol{Basin}$ is a possibly multidimensional vector of sedimentary basin measures. The main threat to identification is the possibility that some geographic features omitted from the model are correlated with elements of $\boldsymbol{Basin}$ and development outcomes. We address this concern by controlling for several geographic characteristics that have been shown to be correlated with economic and political development.[25] The vector $\boldsymbol{x}$ comprises total land area, mountainous area, tropical area, good-soil area, and length of coastline. The parameter η_r represents an unobserved region-specific determinant of development.[26] We eliminate the potential bias produced by η_r by including region indicator variables.

The first identifying assumption of the model is that, conditional on the set of geographic covariates, $\boldsymbol{Basin}$ is independent of potential development outcomes and potential selection into oil discovery. Informally, this assumption says that $\boldsymbol{Basin}$ does not have a direct effect on development outside the channel of oil discovery, and that basin prevalence is not systematically related to country exploration technology or any other propensity for discovery, after controlling for geographic covariates. Given that we control for geographic features that are both plausibly correlated with $\boldsymbol{Basin}$ and may affect development outcomes, the first assumption is likely to hold. The second identifying assumption is that increasing the prevalence of sedimentary basins would never cause a country to produce less oil, for example because of lower exploration effort. This is the familiar monotonicity assumption.[27] It is likely to hold in all but the most implausible scenarios. The final identifying assumption is that $\boldsymbol{Basin}$ and *Oil* have non-zero correlation. If these assumptions hold, then the two-stage least squares estimand identifies the average causal effect of *Oil* on y in countries where a marginal change in basin area induces a change in *Oil* (Imbens and Angrist, 1994; Angrist and Imbens, 1995).

[25] See Gallup et al. (1998) for geographic correlates of economic development, and see Fearon and Laitin (2003) for the correlation between mountainous terrain and insurgency.

[26] The regions are Africa, Europe/Northern America/Oceania, Asia, and Latin America/Caribbean.

[27] See Imbens and Angrist (1994) and Angrist and Imbens (1995).

Our identification strategy is related to studies which use a measure of the initial resource endowment as an instrument for resource wealth over a specific time period (van der Ploeg and Poelhekke, 2010; Tsui, 2011). The resource endowment of a country is typically measured as the sum of cumulative resource discoveries and a geological estimate of undiscovered subsoil resources. The disadvantage of this measure is that known resource endowments represent a non-random sample of true resource endowments. Resource discovery depends on exploration effort, which is likely to be correlated with country characteristics such as property rights institutions (Bohn and Deacon, 2000; Cust and Harding, 2015; Arezki, van der Ploeg, and Toscani, 2016b). Hence the difference between true endowment and known endowment is a function of country characteristics that influence development. In contrast, sedimentary basins are unaffected by country-level outcomes. The next section will discuss robustness checks comparing estimates using the basin instrument to estimates using the oil endowment instrument from Tsui (2011).

In contrast to the empirical strategy presented here, researchers commonly use commodity price shocks, either directly (Caselli and Tesei, 2016) or interacted with a time-invariant measure of resource abundance (Brückner et al., 2012; Dube and Vargas, 2013; Andersen et al., 2014; Carreri and Dube, 2015), as a source of exogenous variation in resource wealth. The strategy appears very credible when applied to subnational data. However, in cross-country studies, the approach raises two concerns. First, the commodity price may not be exogenous to all countries. Producers with significant market share, such as members of OPEC, may adjust production to manipulate prices in response to changing global or domestic economic conditions. This concern is alleviated by dropping large producers from the sample, but at the expense of external validity. Second, the time-invariant measure of resource abundance, usually calculated in an initial period or averaged over several periods, is endogenous in cross-country regressions for reasons already mentioned. Identification issues aside, the price-shock strategy is suited for estimating short-run effects. However, the long-run effects of natural resource wealth on democracy and fiscal capacity are also of great interest, especially because both outcomes can be viewed as the result of many years of investment.[28]

2.4.2 Instrument Selection

No definitive ranking of sedimentary basin types by hydrocarbon potential exists in the petroleum geology literature.[29] Therefore, we pursue a data-driven procedure for instrument selection. In selecting a set of valid instrumental variables, the researcher generally faces a trade-off between bias and efficiency. Starting from a baseline set of valid instruments, adding additional valid instruments potentially improves asymptotic efficiency (Wooldridge, 2010,

[28] Persson and Tabellini (2009) develop the idea of "democratic capital," and Besley and Persson (2009a, 2010); Besley and Persson (2011) model fiscal capacity as a state variable.

[29] Kingston, Dishroon, and Williams (1983) admit that "there is no magic formula which can separate sedimentary basins into oil-and-gas-prone versus barren."

pp. 229–230). However, the finite-sample bias of 2SLS generally grows with the number of instruments used (Donald and Newey, 2001), posing a particularly severe problem when the added instruments are weak (Bound, Jaeger, and Baker, 1995). Furthermore, the presence of weak instruments can render inference based on the standard normal approximations invalid (Staiger and Stock, 1997; Stock and Yogo, 2003). In light of these concerns, we search for the (possibly singleton) set of instruments that maximizes the first-stage F statistic, rather than including all possible instruments. In this way we prioritize minimizing bias and making valid inferences over maximizing efficiency. Specifically, for each of the two basin aggregation methods described in Section 2.2, we estimate a first-stage regression for every possible subset of basins. For each regression, we calculate the Kleibergen and Paap (2006) robust rk Wald F statistic for the excluded instruments.

The main results will be based on the set of instruments that maximizes this F statistic, though we will also report results using the F statistic-maximizing instrument set for each set size. It is important to note that the instrument selection procedure does not invalidate second-stage inference. The reason is that model selection is performed at the service of predicting oil production, not second-stage outcomes. Nonetheless,

2.5 Empirical Results

2.5.1 Descriptive Statistics

Table 2.4 provides general summary statistics. Average democracy in 2008 (0.69) greatly exceeds average democracy in 1966 (0.44), reflecting a general trend toward democratization. Table 2.5 summarizes variables separately according to whether the country produced any oil from 1966–2008. In the sample period 96 countries had positive oil production, and 76 had zero production. In 1966 average democracy in non-oil countries was three percentage points higher than average democracy in oil countries. By 2008 this difference had increased to seven percentage points, though neither difference is statistically significant ($p = 0.677$, $p = 0.182$). Corruption levels and the number of coup attempts were similar in the two groups, however oil countries had more internal conflict and purges ($p = 0.068$, $p = 0.067$). While oil countries had greater total revenue as a proportion of GDP from 2000–2008 compared to non-oil countries ($p < 0.001$), total non-resource tax revenue was lower in oil countries than in non-oil countries ($p = 0.179$). Oil countries tended to be richer than non-oil countries, both in 1966 ($p < 0.001$) and in 2008 ($p < 0.001$). Average executive constraints from 1950–1965 were slightly stronger in oil countries, although the difference is statistically insignificant ($p = 0.594$). Unsurprisingly, all sedimentary basin measures are higher for oil countries, with the exception of the relatively rare convergent ocean-ocean basins, though the difference in average values is statistically insignificant. Average land area, coastline, mountainous area, and good-soil area are statistically indistinguishable in the two groups, although oil countries contain less tropical area on average ($p = 0.029$). The average year of independence is 1771 among oil countries and 1896 among non-oil countries ($p = 0.044$). It is important to note that the categories mask considerable heterogeneity in production levels, as the distribution of oil production is highly skewed.

2.5.2 First Stage

We focus on the effect of oil production on the outcomes of interest, because the results are very similar when other measures of oil wealth are used.[30] Tables 2.6 and 2.7 present the first-stage results. We report the Kleibergen and Paap (2006) robust *rk* Wald *F* statistic, which tests for weak identification and is robust to heteroskedasticity. In each table, column *N* reports the results using the instrument set of size *N* that maximizes the first-stage *F* statistic.

Table 2.6 presents first-stage results based on basin categories defined according to the general plate-tectonic environment and primary mechanism of subsidence. The singleton instrument set that maximizes this *F* statistic is the basin type with convergent continent-continent tectonics and mechanical subsidence, which achieves an *F* statistic of 25.3. The aforementioned basin type, together with the basin type with convergent ocean-continent tectonics and thermal subsidence, constitute the two-instrument set that maximizes the *F* statistic, achieving an *F* statistic of 17.6. Inspection of the remaining columns reveals that, with one exception, adding an additional instrument reduces the *F* statistic. When every instrument is included, the *F* statistic equals 9.4

Table 2.7 presents first-stage results based on basin categories defined according to the final element of the Fugro Tellus code, which indicates local properties of the depositional environment. The singleton instrument set that maximizes the first-stage *F* statistic is the foreland basin type, which achieves an *F* statistic of 16.4. Foreland basins and intracratonic sag basins constitute the two-instrument set that maximizes the first-stage *F* statistic, achieving an *F* statistic of 17. From column 2 to column 10, the *F* statistic declines monotonically in the number of instruments included, equaling only 6.6 when every instrument is included.

Comparing the results across Tables 2.6 and 2.7, the instrument set that maximizes the *F* statistic is the singleton basin type with convergent continent-continent tectonics and mechanical subsidence. The baseline second-stage results will be based on this instrument set, though we will report results using the other instrument sets as well. The optimal instrument set's *F* statistic of 25.3 indicates that strong-instrument asymptotic theory applies. Nonetheless, to be conservative we also report 95-percent Anderson and Rubin (1949) confidence intervals for the coefficient on oil. Unlike the usual Wald test, the Anderson-Rubin test has correct size even in the presence of weak instruments.

[30] Results using oil discovery, oil reserves, and oil endowment are reported in the appendix, and results using oil and gas production, oil exports, and oil and gas exports are available upon request.

2.5.3 Second Stage

Tables 2.8 and 2.9 present the second-stage results. In each table, Panel A presents the OLS estimates, and Panel B presents the IV estimates. Below the IV estimates in Panel B, we report the p-value to a test of whether oil production is endogenous. The endogeneity test is the Hansen (1982) overidentification test of the null hypothesis that oil production is exogenous.[31]

Political Resource Curse

Table 2.8 presents tests of the political resource curse hypothesis. The regressions presented in the first two columns provide strong evidence that oil wealth impedes democracy. The IV estimates indicate that a one-percent increase in average annual oil production per capita from 1966–2008 reduces the level of democracy in 2008 by 0.038. The same increase in oil production reduces average democracy during 1966–2008 by 0.039. The effects are statistically significant at the five- and one-percent levels, respectively, and appear to be large in political-economic terms. An increase in oil production by one standard deviation (4.24 log points) reduces the 2008 democracy score by 0.16, or half a standard deviation. This is roughly equal to the difference between the scores of Colombia or Kenya (0.85) and the United States (1.0). In both democracy specifications, the OLS estimates are smaller in absolute magnitude than the IV estimates; in the second specification we can statistically reject the exogeneity of oil production ($p = 0.063$), although in the first we cannot.

The results in column 3 suggest that oil wealth increases corruption, consistent with conventional wisdom and previous empirical evidence (e.g., Bhattacharyya and Hodler, 2010). (Recall that the corruption index ranges from 0 to 6, where higher numbers indicate less corruption.) An increase in oil production by one standard deviation reduces the corruption index by 0.58 points, or half a standard deviation. The OLS estimates are much smaller in absolute magnitude and are statistically insignificant. The discrepancy between the OLS and IV results is consistent with more corrupt countries attracting less oil exploration and production, perhaps due to a poor business environment. In this specification we can statistically reject the exogeneity of oil production ($p = 0.053$).

The results in columns 4 through 6 provide little evidence that oil wealth increases conflict—contrary to conventional wisdom, though consistent with previous research (Cotet and Tsui, 2013). The OLS results suggest that oil wealth has a positive and significant effect on internal armed conflict, though the corresponding IV estimate is half the size of OLS and is statistically insignificant. Both the OLS and IV regressions find that the effect of oil wealth on coup attempts and purges is statistically insignificant.

Columns 7 and 8 examine the effect of oil production on government revenue. The IV estimate of the effect of oil production on total government revenue is positive but statistically insignificant. In contrast, the IV estimate of the effect of oil production on tax revenue

[31] The test is valid under the assumption that *Basin 2* is exogenous.

is negative and statistically significant. A one-percent increase in oil production causes a 0.16-percent reduction in tax revenue as a share of GDP from 2000–2008. The effect on tax revenue is significant at the one-percent level. An increase in oil production by one standard deviation causes a decline in tax revenue by 0.69 log points, or one standard deviations. This is roughly the difference between Burundi (−2.01) and France (−1.32). The corresponding OLS estimates are much smaller in absolute magnitude. The Hansen (1982) test decisively rejects the exogeneity of oil production in the tax revenue specification ($p < 0.001$) but not the total revenue specification.

Economic Resource Curse

Table 2.9 presents tests of the economic resource curse hypothesis. Column 1 presents results for (log) GDP per capita, while columns 2 through 5 present results for disaggregated measures of (log) GDP per capita. Both the OLS and IV estimates indicate that oil wealth raises GDP. According to the IV estimate, a one-percent increase in average oil production per capita raises GDP per capita in 2008 by 0.07 percent. The effect is statistically significant at the ten-percent level. Raising oil production by one standard deviation causes an increase in GDP by 0.31 log points, or 0.25 standard deviations. This is roughly the difference between Norway (11.09) and Ireland (10.78) or between Algeria (9.45) and Ecuador (9.14). Thus oil production appears to be economically less significant for GDP than for democracy or fiscal capacity.

The results in column 1 could be consistent with oil wealth harming the non-resource sectors of the economy, as long as the positive effects on the resource sector outweigh the negative effects on the non-resource sectors. The OLS results in columns 2 through 5 indicate that oil wealth actually raises non-resource GDP and manufacturing GDP. The IV estimates for non-resource GDP are similar to the OLS estimates, though less precise. Together they suggest that a one-percent increase on oil production raises non-resource GDP by 0.05 to 0.07 percent. The OLS and IV estimates of the effect of oil wealth on manufacturing significantly diverge. The OLS estimate indicates that a one-percent increase on oil production raises manufacturing GDP by almost 0.08 percent, and this estimate is significant at the one-percent level. On the other hand, the IV estimate is negative and statistically insignificant. We reject the exogeneity of oil production in the manufacturing GDP equation ($p = 0.079$).

In four of the 13 specifications, the Hansen (1982) test rejects the exogeneity of oil production at the 10-percent level. This outcome is unlikely to be due simply to chance or multiple hypothesis testing. For example, if oil production were in fact exogenous in each of the 13 regressions, the probability of rejecting the null hypothesis of exogeneity at the 10-percent level in four or more of the specifications is 0.034 (assuming the tests are independent).[32] Furthermore, whenever the OLS and IV estimates diverge considerably, OLS understates the negative effects of oil relative to IV. Thus the results are consistent with the possibility

[32] Under the stated assumptions, the number of rejections, W, has a binomial distribution with $n = 13$ and $p = 0.1$. Therefore $P(W \geq 4) = 0.034$.

that countries with stronger political and fiscal institutions disproportionately select into oil discovery and production.

2.5.4 Varying the Instrument Set

The results discussed so far are based on the optimal (singleton) instrument set which maximizes the first-stage F statistic. We now consider how the results change when the instrument set changes. Figures 2.2 and 2.3 plot the second-stage results for the political and economic outcomes, respectively, using instrument sets categorized according to the general plate-tectonic environment and primary mechanism of subsidence. The results based on N instruments use the instrument set of size N that achieves the highest first-stage F statistic. For each outcome, the dashed gray line indicates the value of the corresponding OLS estimate. As Table 2.6 showed, each of the eight instrument sets is at least moderately strong, however the first (singleton) instrument set is significantly stronger than the others with a first-stage F statistic of 25.3. Because of this, along with the fact that the bias of 2SLS generally increases with the number of instruments (Donald and Newey, 2001), we would expect results based on the first instrument set to have lower bias, but also lower precision, compared to results based on the other instrument sets. Consistent with this prediction, the estimated effects of oil production on democracy, average democracy, corruption index, and tax revenue are substantially more negative and less precise when using one instrument—or even two instruments—compared to estimates based on larger instrument sets. Adding additional, weaker instruments appears to push the 2SLS estimates toward the OLS estimate, which we expect to be biased upwards for these four outcomes. The estimates of the effect of oil production on internal conflict and purges show a somewhat different pattern: estimates based on small instrument sets imply effects roughly equal to zero, while estimates based on larger instrument sets imply positive and marginally significant effects. The point estimates for coup attempts and total revenue do not change much as the instrument set varies. For every measure of GDP, the point estimates based on smaller instrument sets are smaller than the point estimates based on larger instrument sets. This pattern is especially apparent for non-resource GDP and manufacturing GDP. Similar to the results for democracy, corruption, and tax revenue, the GDP results are consistent with the fact that richer countries with stronger institutions engage in more resource exploration and production.

Figures 2.4 and 2.5 plot the second-stage results for the political and economic outcomes, respectively, using instrument sets categorized according to the local properties of the depositional environment. Once again, results based on N instruments use the instrument set of size N that achieves the highest first-stage F statistic. The coefficient patterns are qualitatively similar to those in Figures 2.2 and 2.3. The main difference is that the estimates based on different instrument sets diverge less from each other, perhaps because the smaller instrument sets are weaker than in the case of the tectonic-subsidence grouping. Another difference is that the sign and statistical significance of the estimated effect of oil production is less sensitive to the instrument set—at least for the political outcomes—than when instrument sets based on the tectonic-subsidence grouping are used. In fact, nearly every instrument set implies that oil production has a negative and significant effect on democracy, average democracy, the corruption index, and tax revenue; a positive and significant effect on internal conflict, purges, and total revenue; and an insignificant effect on coup attempts. The preponderance of the evidence suggests that OLS understates the adverse political effects of oil production, though the OLS and 2SLS estimates often are not statistically different from one another.

2.5.5 Validity of the Instrument

Measurement

We now consider several potential objections to the validity of the $\boldsymbol{Basin}$ instrument. The first relates to measurement. Two of the three methods used to map sedimentary basins—core sampling and seismic profiling—require the use of advanced technology and physical access to the area under investigation. One might therefore worry that the precision or reliability of the basin data is increasing in "good" institutions like property rights protections. In that case the *variance* of the basin measurement error would be decreasing in the quality of institutions. However, it does not follow that the measurement error is *correlated* with the quality of institutions, so the above form of measurement error need not produce asymptotic bias.

Another version of the measurement argument supposes that basin area is systematically underestimated in countries with poor institutions, invalidating the instrument. This argument is unconvincing for two reasons. First, it is inconsistent with the pattern of basin coverage by region. Table B.2.2 in the appendix summarizes the portion of sovereign area covered by sedimentary basins separately for seven regions defined by common geographical location and history. Basin coverage is actually higher on average in Eastern Europe and Central Asia (0.67) and the Middle East and North Africa (0.86)—areas associated with relatively weak property-rights protections—than in the extensively prospected areas of Northern, Central, Western, and Southern Europe and Neo-Europes (0.57).[33] This pattern is visually confirmed in Figure 2.1. Second, even if basin area were underestimated in countries with poor institutions, the vast majority of the conclusions drawn in this paper would hold up. This type of non-classical measurement error would cause the IV estimates to understate the negative effects of oil on democracy, corruption, conflict, and fiscal capacity, so that the estimated coefficients would often provide informative (absolute) lower bounds on the true effect.[34]

[33] The "Neo-Europes" are Australia, Canada, New Zealand, and the United States.

[34] Let Z, Z^*, and X be the measured *Basin*, the true *Basin*, and *Oil*, respectively, after netting out the control variables using population projections. If the measurement error in *Basin*, e, is uncorrelated with the control variables, then $Z = Z^* + e$ (Wooldridge, 2010, p. 29). Then the probability limit of $\hat{\beta}_{IV}$ is $\beta + \text{Cov}(\varepsilon, e) / \text{Cov}(X, Z)$. Because $\text{Cov}(X, Z)$ is positive, the sign of the bias depends on the sign of $\text{Cov}(\varepsilon, e)$. For "good" outcome variables like democracy, the example in the text implies that the bias is positive, whereas for "bad" outcome variables like

Reverse-Engineering of the Basin Classification

The next potential objection is that sedimentary basin classification could be reverse-engineered: the known presence or absence of hydrocarbons may influence how geologists categorize a basin, based on their knowledge of other hydrocarbon-rich or hydrocarbon-poor basins. Therefore some of the correlation between hydrocarbons and particular basin types may be spurious rather than based on true geological features.

This issue is unlikely to invalid our results, for two reasons. First, reverse-engineering of basin categories would bias the 2SLS estimates towards the OLS estimates. The intuition is simple: in the extreme case of reverse-engineering, a few basin types would have 100-percent hydrocarbon success rates and would jointly predict oil production almost perfectly. To the extent that the 2SLS and OLS estimates differ, the 2SLS estimates would still provide useful bounds on the true effects of oil production.

Second, as already discussed, Figures 2.2, 2.3, 2.4, and 2.5 show that the results are broadly similar whether instruments are constructed based on global characteristics of basins or local characteristics. It is unlikely that both the global and local categorizations of basins could be reverse-engineered.

Predetermined Confounders

Another potential objection is that $\boldsymbol{Basin}$ could be correlated with omitted determinants of development, causing an asymptotic bias of unknown sign. To explore this possibility, Table 2.10 reports the results from regressing several predetermined variables on the basin instrument, controls, and region effects. The first outcome is the urbanization rate in 1850, which is the last year in the series provided by Chandler (1987). The next outcome is an indicator for have a British legal origin, taken from William Easterly's Global Development Network Growth Database (Easterly, 2001). The third outcome is an indicator for having a legacy as a communist country, taken from the list of communist countries in Kornai (1992). The next three outcomes measure the percentage of the population that was Christian, Muslim, or Hindu in 1950. These data come from the World Religion Database (Johnson and Grim, 2017). The final three outcomes are measures of ethnic, religious, and linguistic fractionalization produced by Alesina, Devleeschauwer, Easterly, Kurlat, and Wacziarg (2003). Seven of the nine estimated coefficients on $\boldsymbol{Basin}$ are statistically insignificant, suggesting that the instrument is uncorrelated with historical determinants of long-run economic development, legal origin, communist legacy, the presence of Christians or Hindus, or religious or linguistic fractionalization. The basin instrument has a strong, positive correlation with the percentage of the population that was Muslim in 1950. A large portion of this correlation is driven by the religious composition and presence of basins in the Middle East and North Africa; adding a dummy variable for

conflict, the bias is negative.

this region causes the coefficient on $\boldsymbol{Basin}$ to fall by half.[35] The basin instrument also has a positive correlation with ethnic fractionalization that is significant at the ten-percent level. It is therefore important to examine how the main results change when we control for these two variables.

Table 2.11 reports the main results for the political outcomes using the optimal instrument while controlling for the percentage of the population that was Muslim in 1950. The OLS estimates of the effect of oil production on the political outcomes generally move slightly closer to zero while maintaining the same pattern of signs and similar levels of statistical significance: oil production still has a negative and significant effect on democracy, average democracy, and tax revenue, while having a positive and significant effect on internal conflict and total revenue. Controlling for the Muslim population causes the 2SLS estimates to become more imprecise, due to a weakened first stage. The 2SlS estimates for the effect of oil production on democracy, average democracy, and corruption all move towards zero while remaining below the OLS point estimates, once again suggesting that OLS may understate the adverse effects of oil wealth on institutions. These three point estimates are now statistically insignificant. Given that OLS likely provides an upper bound on the effect of oil production on democracy, we are still able to conclude that oil impedes democracy. Controlling for Muslim population pushes the 2SLS estimate of the effect of oil production on tax revenue slightly closer to zero, however this estimate remains sizable and highly significant.

Table 2.12 reports the main results for the economic outcomes using the optimal instrument while controlling for the percentage of the population that was Muslim in 1950. Both the OLS and 2SLS estimates are broadly similar to those in the baseline specification, in terms of both magnitude and significance. Controlling for Muslim population leads to slightly larger positive estimated effects of oil production on GDP.

Overall, Tables 2.11 and 2.12 suggest that the baseline 2SLS estimates may have slightly overstated the adverse effects of oil wealth on democracy and taxation while still providing strong evidence that such adverse effects exist. The results weaken the original claim that the OLS results for average democracy and corruption were substantially biased, while confirming the claim that the OLS results for taxation were substantially biased. The robustness check confirms the baseline OLS and 2SLS estimates for the GDP regressions.

Are the above results limited to the optimal instrument, or do they apply to all instrument sets? Figures 2.6, 2.7, 2.8, and 2.9 replicate the main results using optimal instrument sets of different sizes while controlling for the percentage of the population that was Muslim in 1950. The pattern of coefficient estimates based on different instrument sets is very similar to the pattern in the original figures. The two main differences are that some point estimates move slightly closer to zero, and the confidence intervals of all point estimates grow.

Tables 2.13 and 2.14 report the main results using the optimal instrument while controlling for ethnic fractionalization. The results are remarkably similar to the baseline results, which

[35] Result not shown but available upon request.

is perhaps unsurprising given that the partial correlation between $\boldsymbol{Basin}$ and ethnic fractionalization is weak. Figures 2.10, 2.11, 2.12, and 2.13 confirm that the results using different instrument sets hardly change when we control for ethnic fractionalization.

Placebo Tests

While it is reassuring that our conclusions do not change significantly when accounting for the influence of potential confounders described above, there may be other determinants of political and economic development that are correlated with the basin instrument. We address this possibility with two placebo tests. If $\boldsymbol{Basin}$ impacts development only through the channel of oil wealth, then it should have no impact on economic and political outcomes in years when oil was not a commercially valuable commodity or when world oil production was minimal. Before 1859 the value of oil was modest. The year 1859 saw both the first modern oil well (by Edwin Drake) and the first commercially successful internal-combustion engine (by Étienne Lenoir) (Britannica, 2015). Prior to 1920 no country produced a significant amount of oil, defined as $100 per capita (in constant 2007 USD) (Andersen and Ross, 2014). In 1940 there were three significant oil producers, and by 1950 there were 10. For context, 56 countries were significant oil producers in 2008 (Ross, 2013).

Figure 2.14 plots estimates of the reduced-form effect of $\boldsymbol{Basin}$ on the polity index in different years. To examine the changing influence of the basin instrument over time, we fix the sample of countries. The four graphs are based on fixed country samples starting in 1900, 1910, 1920, and 1930. All four graphs tell the same story: prior to 1940, the effect of $\boldsymbol{Basin}$ on democracy was statistically indistinguishable from zero. Starting in 1940, $\boldsymbol{Basin}$ had a negative and statistically significant effect, and this negative effect persisted to 2008. Figure 2.15 is similar, presenting four graphs of the reduced-form effect of $\boldsymbol{Basin}$ on log population density over time. We focus on population density, because GDP data prior to 1950 are available only for a small number of countries. Prior to 1950, the availability of population data from Maddison (2013) varies considerably from year to year. We choose to measure log population density in the years 1820, 1870, and 1913, because these are the only years prior to 1950 for which population data are available for more than 65 countries. The graphs suggest that $\boldsymbol{Basin}$ had no influence on log population density prior to 1970; the effect of $\boldsymbol{Basin}$ on log population density becomes positive and statistically significant starting in 1990. Together Figures 2.14 and 2.15 suggest that $\boldsymbol{Basin}$ did not influence political and economic development in periods in which the value of oil production was insignificant. These results strengthen the claim that the baseline results are not simply driven by variables that are correlated with both $\boldsymbol{Basin}$ and long-run development.

Predetermined Borders

The validity of the basin instrument rests on the assumption that national borders were drawn without consideration for the locations of sedimentary basins. The most plausible violation of this assumption would occur in geographic regions where modern borders were established after the discovery of oil. If oil-field acquisition (hence basin acquisition) via border changes were systematically related to potential outcomes—e.g., if more economically or militarily powerful countries acquired more oil fields through territorial conquest or delimiting colonial dependencies—then the IV estimator would be inconsistent for the treatment effect of interest.

To address this concern we replicate the main analysis on the subsample that excludes any country whose land borders could have plausibly been influenced by the location of oil fields.[36] We first record the year of the earliest known oil discovery for each country, according to Thieme, Lujala, and Rød (2007). We then record the year of the earliest establishment of modern borders, using the information in Strang (1991), Britannica (2015), and CIA (2015). It is important to note that the modern borders of most former colonies and former satellite states were drawn decades before independence. Finally, we record the dates of all changes to homeland territory (as opposed to dependency territory) since 1816, according to Tir, Schafer, Diehl, and Goertz (1998). Using data on country contiguity from Correlates of War Project (2007) (described in Stinnett, Tir, Schafer, Diehl, and Gochman (2002)) to identify neighboring countries, we implement the following procedure:

1. Exclude country A if country A first discovered oil before its modern borders were set.
2. Exclude country A if country A's neighbor, country B, first discovered oil before country A's modern borders were set, *and* country B's modern borders were not set prior to the discovery.
3. To minimize unnecessary exclusions, include countries that were to be excluded according to Rule 1 or Rule 2 if either (a) there are no known onshore oil fields within 200 kilometers of the border in question, or (b) there are no land basins within 200 kilometers of the border in question.
4. Include countries with borders set prior to 1859, even if they qualify for exclusion according to Rule 1 or Rule 2.[37]

The procedure results in the exclusion of 61 countries from the baseline sample of 157 countries. Tables 2.15 and 2.16 report regression results based on the sample of countries with borders that were not plausibly influenced by the location of oil fields or basins. The results

[36] We focus this robustness check on land borders for tractability, as maritime borders are often ambiguous or disputed.

[37] Before 1859 petroleum was arguably not a very valuable commodity and thus would not have influenced border formation. The year 1859 saw both the first modern oil well (by Edwin Drake) and the first commercially successful internal-combustion engine (by Étienne Lenoir) (Britannica, 2015).

are remarkably similar to the results from the full sample, both qualitatively and quantitatively. The broad similarity of the results to the main results suggests that countries with borders drawn after the discovery of oil are not systematically different than countries with borders drawn before the discovery of oil.

2.5.6 Comparison to Endowment Instrument

The closest predecessor to the identification strategy in this paper is Tsui (2011), who uses oil endowment as an instrument for oil discovery. To facilitate comparison between Tsui (2011) and the current paper, we normalize the oil endowment variable in the same manner that we normalize the basin variables: *Endowment* is the (log of) total oil endowment in millions of barrels divided by 1960 population.[38] As mentioned in the introduction, there are *a priori* reasons to worry that known oil endowment is endogenous. We find suggestive statistical evidence that this is indeed the case. Tables 2.17 and 2.18 compare the OLS results, 2SLS results based on *Endowment*, and 2SLS results based on *Basin*. The first-stage F statistic on *Endowment* is extremely large—410 in the full sample—and IV estimates using *Endowment* are almost always closer than IV estimates using *Basin* to the OLS estimates. In addition, the Hansen (1982) overidentification test rejects the exogeneity of *Endowment* in the average democracy, corruption, tax revenue, and manufacturing GDP specifications, though it fails to reject exogeneity in the other specifications.[39] Nonetheless, the *Endowment* and *Basin* produce the same qualitative conclusions, so the results of this paper may be interpreted as confirming the findings of Tsui (2011).

2.5.7 Discussion

The estimated negative effects of oil production on democracy and tax revenue indicate that oil wealth has a tendency to degrade—or retard the development of—democratic institutions and fiscal capacity over the long run. Oil wealth increases the value of holding political power, which in theory could make a coups d'état more attractive in the eyes of potential usurpers. However, the resource revenue also strengthens the government's hand, potentially funding investment in defense.[40] The results of this section suggest that oil wealth increases government repression in the form of purges.[41] However, in equilibrium oil wealth does not lead to more coup attempts. This result is consistent with the model of Tsui (2010), which predicts that the number of political insurgents will be independent of the size of resource

[38] The data on oil endowment is shared online by Cotet and Tsui (2013).

[39] The overidentification test evaluates the exogeneity of *Endowment* under the assumption that *Basin* is exogenous.

[40] Cotet and Tsui (2013) find that oil discoveries increase military spending in nondemocratic countries.

[41] Note that we find a positive, significant effect of oil production on purges using many instrument sets of size greater than one. The evidence on internal armed conflict is inconclusive.

wealth.[42] The reason is that an increase in resource wealth induces the ruler to invest in political entry barriers which deter potential insurgents.

The fact that OLS underestimates the pernicious effects of oil on democracy and tax revenue suggests that countries with better political institutions and greater state capacity have a greater propensity to select into oil production.[43] The results are consistent with recent evidence that the drilling decisions of international oil companies are highly sensitive to the quality of national institutions (Cust and Harding, 2015).

2.6 Heterogeneous Effects by Executive Constraints

2.6.1 Theory

Several political economy models predict that the political and economic effects of natural resource wealth will depend on the quality of institutions. In some models institutions determine the extent to which incumbents can spend resources to increase their likelihood of staying in power. The degree to which resource booms promote autocracy or resource misallocation within the economy thus depends on institutions (Robinson et al., 2006; Caselli and Tesei, 2016). In a similar vein, democratic institutions determine the degree to which popular support (or lack thereof) affects the incumbent's chances of staying in power. While resource booms increase the scope of corruption, incumbents are less likely to embezzle state funds when democratic institutions are strong (Bhattacharyya and Hodler, 2010). In addition, resource rents are more likely to promote repression and civil war when political checks and balances are weak (Besley and Persson, 2009b, 2011). Finally, resource abundance can reduce economic growth when institutions favor rent-seeking over productive activities (Mehlum et al., 2006).[44]

In the appendix we present a theoretical model that predicts that institutions will determine the effect of resource revenue on the incumbent's *joint* decision over the political regime and tax policy. An autocrat faces the threat of a popular uprising and must decide whether to allow a transition to democracy or suppress the movement using bribes. Under democracy the median voter, who is poor, chooses a positive tax rate. When the autocrat chooses to suppress democracy, his optimal strategy involves bribing the rich citizens and setting taxes equal to zero. Both the autocrat's ability and willingness to suppress democracy increase in the amount of resource rents accruing to the autocrat. However, executive constraints create transaction costs associated with stealing resource rents from government coffers and making bribes. As a result, a resource boom increases the likelihood that autocracy and low taxation persist *if and*

[42] This result depends on the counterinsurgent technology having constant returns to scale.

[43] Prospecting intensity probably accounts for most of the differential selection into oil production. While known subsoil assets in the OECD countries are valued at around US$265,000 per square kilometer, in sub-Saharan Africa known subsoil assets are valued at only US$45,000 per square kilometer (Collier and Laroche, 2015).

[44] See Tsui (2010) for a model that combines the economic and political dimensions of the resource curse while modeling institutions as the deadweight costs associated with rent appropriate and political entry deterrence.

only if executive constraints are sufficiently weak. See the appendix for details.

2.6.2 Evidence

To test the implications of the theoretical models described above, we estimate the effects of oil production, allowing for heterogeneity in the response according to the strength of executive constraints. We construct a measure of initial executive constraints by averaging each country's XCONST score (Polity IV) from 1950–1965.[45] The variable XCONST is measured on a scale of one to seven, with one indicating "unlimited authority," three indicating "slight to moderate limitation on executive authority," five indicating "substantial limitations on executive authority," and seven indicating "executive parity or subordination." Numbers two, four, and six denote intermediate categories. We construct an indicator variable, *weak constraints*, which equals one for countries that averaged a score of three or lower from 1950–1965. In our sample the median score for average XCONST over this period is three.

We split the sample into two subsamples—countries with relatively strong executive constraints and those with relatively weak constraints—and estimate the structural equation separately for each subsample. We then compare the IV estimates obtained in each subsample. While we have data on democracy in 2008 for 157 countries, we observe *weak constraints* for only 116 countries. This is because countries that gained independence after 1965 have missing values for XCONST for all years from 1950–1965.

Tables 2.19 and 2.20 present the results of the heterogeneity analysis. The validity of the exercise relies on the assumption that *weak constraints* is uncorrelated with unobserved determinants of development. In the appendix we discuss some supporting evidence for the exogeneity of *weak constraints*. However, exogeneity is a strong assumption, and the results in this section should be interpreted with caution. The optimal basin instrument in the full-sample analysis leads to excessively small first-stage F statistics in the subsample analysis. We therefore report results based on the instrument set {Foreland, Intracratonic Sag}, which produces modestly sized first-stage F statistics in the subsamples. We checked the results using the seven strongest instrument sets according to Tables 2.6 and 2.7, and the pattern of second-stage coefficients is very similar using different instrument sets.

Political Resource Curse

Table 2.19 presents the results of the heterogeneity analysis for the political outcomes. As shown in Panel A, in the sample of strong-constraints countries, oil production has a statistically insignificant effect on each political variable, with the exception of total revenue. In contrast, Panel B shows that, in the sample of weak-constraints countries, oil production has a statistically significant effect on six of the eight outcomes, reducing democracy in 2008, average democracy

[45]Naturally, the sample is restricted to countries with at least one observation of XCONST from 1950–1965. We use a 16-year average to reduce noise and maximize sample size.

from 1966–2008, and tax revenue; and increasing internal conflict, purges, and total revenue. The effects of oil production on corruption and coup attempts are statistically insignificant in the sample of weak-constraints countries.

A one-percent increase in oil production reduces the level of democracy in 2008 by 0.044. The effect is significant at the five-percent level. In the weak-constraints sample, an increase in oil production of one standard deviation (4.24 log points) reduces 2008 democracy by 0.19, or 0.59 standard deviations.[46] This is roughly equal to the difference between the scores of Algeria (0.6) and Malawi (0.8).[47] The negative effect of oil on democracy in 2008 is smaller in magnitude (−0.027) and statistically insignificant in the strong-constraints sample.

Oil production also has a large effect on tax revenue in the sample of weak-constraints countries. A one-percent increase in oil production reduces the tax-revenue-to-GDP ratio by 0.108 percent in the sample of countries with weak constraints. The estimate is significant at the five-percent level. Among countries with weak constraints, increasing oil production by one standard deviation (4.24 log points) reduces the tax revenue share of GDP by 0.46, or 0.67 standard deviations, which is roughly equal to the difference in tax revenue between Nicaragua (−1.83) and Mexico (−2.33).[48] The negative effect of oil on tax revenue is smaller in magnitude (−0.012) and statistically insignificant in the strong-constraints sample.

Overall, the results suggest that the adverse political consequences of oil wealth are concentrated in the sample of countries with weak initial constraints on the executive. While the point estimates in the two subsamples often differ substantially, Panel C shows that we are unable to reject the hypothesis that the point estimates are equation in any of the equations, perhaps owing to the small sample sizes.

Economic Resource Curse

Table 2.20 presents the results of the heterogeneity analysis for the economic outcomes. As shown in Panel A, in the sample of strong-constraints countries, oil production has a positive effect on each economic variable, though each coefficient is statistically insignificant. The point estimates for the sample of weak-constraints countries, reported in Panel B, are slightly larger than those in Panel A, and they are all significant at least at the ten-percent level. In the weak-constraints sample, a one-percent increase in oil production raises GDP per capita by 0.15 percent, and the effect is almost identical for manufacturing GDP—contrary to the Dutch Disease hypothesis.

[46] The standard deviation of 2008 democracy in the weak-constraints sample is 0.31.

[47] While both Algeria and Malawi had weak executive constraints from 1950–1965, Algeria produced a significant amount of oil from 1966–2008, and Malawi produced no oil.

[48] Both Nicaragua and Mexico had weak executive constraints from 1950–1965. From 1966–2008 oil production was substantial in Mexico and nil in Nicaragua.

Weak Constraints

The heterogeneity analysis would be invalid if, for example, pre-1966 oil production affected both *weak constraints* and post-1966 democracy and tax revenue. However, the $\boldsymbol{Basin}$ measures have virtually no statistical association with *weak constraints*, as shown in the appendix. The only basin types that have a statistically significant association with *weak constraints* are convergent ocean-ocean basins, convergent wrench basins, and fore-arc basins. In all three cases, the association with *weak constraints* is negative, which contradicts the claim that pre-1966 oil production adversely affected pre-1966 institutions.

2.6.3 Discussion

We find evidence that the long-run effects of oil wealth on development may be heterogeneous. In particular, the adverse effects of oil on democracy and fiscal capacity are concentrated in the subsample of countries that had weak executive constraints from 1950–1965. This result is consistent with other recent findings. Tsui (2011) finds that the discovery of oil impeded democratization only in countries that were non-democratic at the time of discovery. Similarly, Caselli and Tesei (2016) show that resource windfalls cause autocratic countries to become even more autocratic, whereas they have no effect on the regime in democratic countries or in deeply entrenched autocracies. Finally, Andersen and Aslaksen (2013) show that oil wealth positively affects political survival (measured as the leader's duration in office) in intermediate and autocratic regimes, but not in democracies. In contrast to the results of Bhattacharyya and Hodler (2010), we do not observe heterogeneous effects of oil on corruption. Neither does the effect of oil on conflict seem to differ according to institutional quality. The finding that oil has a larger positive effect on GDP in weak-constraints countries is consistent with other evidence that less developed countries have the largest GDP gains from oil production (Alexeev and Conrad, 2009; Smith, 2015).

In order to identify a heterogeneous effect of oil, the dimension of heterogeneity (e.g., political institutions) must be uncorrelated with unobserved determinants of future political outcomes and oil wealth. Of course, this assumption is unlikely to hold. However, unlike the studies mentioned above, we condition on initial rather than contemporaneous political institutions to (partially) alleviate concerns about the simultaneity of political institutions and resource production.

The heterogeneity results are interesting in light of the recent literature on the determinants of fiscal capacity. Previous empirical studies find that resource wealth tends to negatively impact tax revenue (Cárdenas et al., 2011; Jensen, 2011; Crivelli and Gupta, 2014). However, these studies do not test for heterogeneous effects. The fiscal capacity model of Besley and Persson (2009a, 2010); Besley and Persson (2011) predicts that a "common-interest" state emerges when institutions are "cohesive" enough. In their model institutional cohesion depends on the ability of the group in power to redistribute resources away from the group not in power. In a common-interest state, politicians invest in fiscal capacity, because they know that future capacity to tax will be used to raise funds for common-interest public goods rather than for redistributing income away from the group not in power. Because the marginal utility from public goods is assumed to be declining, a relaxation of the government's budget constraint due to a resource windfall causes the group in power to invest less in fiscal capacity. When institutions are not cohesive, no group invests in fiscal capacity, regardless of the level of resource revenue. Therefore the model predicts that resource rents lower future tax revenue only in countries with cohesive institutions. In contrast, we find that the negative effect of oil production on future tax revenue is strongest in countries that lack cohesive institutions. Our results are not wholly inconsistent with the fiscal capacity model, however they do suggest the importance of low taxation as a means of political survival.

2.7 Conclusion

Using a new instrumental variables approach, we show that oil wealth impedes democracy, increases purges, reduces tax revenue, and raises GDP. We find only weak evidence in support of the proposition that oil wealth increases corruption and internal conflict. In several specifications OLS substantially underestimates the detrimental effects of oil, suggesting that there is positive selection into oil discovery and production. For outcomes such as democracy and fiscal capacity, the initial strength of executive constraints appears to determine whether subsequent oil production is a curse or a blessing. However, initial institutions seem to matter less for how oil affects corruption, conflict, and purges. Despite suffering a political resource curse, countries with weak initial institutions saw the greatest economic gains from oil wealth, at least in aggregate terms.

This paper's identification strategy is useful to researchers studying the long-run impact of oil wealth on any outcome in cross-country data. The strategy can also be applied to subnational analyses, because the spatial distribution of sedimentary basins generates exogenous within-country variation in oil endowment. Furthermore, the general idea of the strategy—that we should exploit variation in the geophysical processes that produce subsoil resources—may prove useful for studying the effects of mineral resources other than oil.

2.8 Tables

Table 2.1: Fugro Robertson Global Basin Classification Codes

Sub-Regime Group	Code	Sub-Regime Name
Convergent (Continent-Continent)	C.1.F	Peripheral Foreland (Continent-Continent)
	C.1.F(p)	Peripheral Foreland with Piggyback (Continent-Continent)
	C.1.POE	Late to Post-Orogenic Extension (Continent-Continent)
	C.1.SOE	Syn-Orogenic Extensional (Continent-Continent)
	C.1.TOC	Trapped Oceanic Crustal Sag (Continent-Continent)
	C.1.W	Intramontane Wrench (Continent-Continent)
Convergent (Ocean-Continent)	C.2.E	Retro-Arc Extensional (Ocean-Continent)
	C.2.F	Retro-Arc Foreland (Ocean-Continent)
	C.2.FA	Fore-Arc (Ocean-Continent)
	C.2.S	Retro-Arc Post-Extensional Sag (Ocean-Continent)
	C.2.W	Arc-Related Wrench (Ocean-Continent)
Convergent (Ocean-Ocean)	C.3.E	Retro-Arc Extensional (Ocean-Ocean)
	C.3.F	Retro-Arc Foreland (Ocean-Ocean)
	C.3.FA	Fore-Arc (Ocean-Ocean)
	C.3.S	Retro-Arc Post-Extensional Sag (Ocean-Ocean)
	C.3.W	Arc-Related Wrench (Ocean-Ocean)
Divergent	D.1	Rift
	D.2	Intracratonic Sag
	D.3	Post-Rift Sag
	D.3(i)	Post-Rift Sag with Inversion
	D.4	Passive Margin
	D.4(i)	Passive Margin with Inversion
Wrench	W.1	Intracratonic Wrench
	W.2	Wrench (Ocean-Continent)

Source. Fugro Robertson, Ltd. (2013).

Table 2.2: Grouping by Plate-Tectonic Environment and Primary Subsidence Mechanism

Number	Tectonics	Subsidence	Basin Aggregation in Group
1	Convergent C-C	Mechanical	C.1.F + C.1.F(p) + C.1.SOE + C.1.W
2	Convergent C-C	Thermo-Mechanical	C.1.POE + C.1.TOC
3	Convergent O-C	Mechanical	C.2.E + C.2.F + C.2.FA + C.2.W
4	Convergent O-C	Thermal	C.2.S
5	Convergent O-O	Mechanical or Thermal	C.3.E + C.3.F + C.3.FA + C.3.S + C.3.W
6	Divergent	Mechanical	D.1
7	Divergent	Thermal	D.2 + D.3 + D.3(i) + D.4 + D.4(i)
8	Wrench	Mechanical	W.1 + W.2

Notes. The categorization is from Fugro Robertson, Ltd. (2013). See Table 2.1 for the basin types associated with each code. In "C-C," "O-C," and "O-O," "C" stands for continent, and "O" stands for "Ocean."

Table 2.3: Grouping by Final Component of Fugro Tellus Code

Number	Group Name	Basin Aggregation in Group
1	Foreland	C.1.F + C.1.F(p) + C.2.F + C.3.F
2	Fore-Arc	C.2.FA + C.3.FA
3	Extensional	C.1.POE + C.1.SOE + C.2.E + C.3.E
4	Convergent Sag	C.1.TOC + C.2.S + C.3.S
5	Convergent Wrench	C.1.W + C.2.W + C.3.W
6	Rift	D.1
7	Intracratonic Sag	D.2
8	Post-Rift Sag	D.3 + D.3(i)
9	Passive Margin	D.4 + D.4(i)
10	Wrench	W.1 + W.2

Notes. The categorization is from Fugro Robertson, Ltd. (2013). See Table 2.1 for the basin types associated with each code.

Table 2.4: Summary Statistics

	Mean	Std. Dev.	Min.	Max.	Obs.
Democracy, 2008	0.69	0.32	0.00	1.00	157
Democracy, 1966	0.44	0.38	0.00	1.00	117
Avg. democracy, 1966–2008	0.53	0.31	0.00	1.00	160
Corruption, 2008	2.56	1.18	0.00	6.00	136
Internal conflicts per year, 1966–2008	0.21	0.48	0.00	3.86	172
Coup attempts per year, 1966–2008	0.06	0.08	0.00	0.35	160
Purges per year, 1966–2008	0.06	0.13	0.00	1.12	172
Total revenue, 2000–2008 (log avg.)	−1.53	0.45	−3.05	−0.54	165
Total tax revenue, 2000–2008 (log avg.)	−1.97	0.69	−5.03	−0.77	167
GDP, 2008 (log p.c.)	9.06	1.26	6.36	11.71	166
GDP, 1966 (log p.c.)	7.69	1.00	6.05	10.37	136
Non-Oil GDP, 2008 (log p.c.)	9.27	1.16	6.33	11.46	132
Non-Oil/Gas GDP, 2008 (log p.c.)	9.28	1.15	6.33	11.45	129
Non-Resource GDP, 2008 (log p.c.)	8.93	1.30	5.92	11.45	166
Executive constraints, 1950–1965	0.47	0.37	0.00	1.00	116
Oil production, 1966–2008 (log avg. p.c.)	−4.88	4.24	−9.03	4.45	172
Oil discovery, 1966–2003 (log avg. p.c.)	−9.03	3.24	−11.14	1.73	172
Oil reserves, 1966–2003 (log avg. p.c.)	−5.26	3.14	−7.30	4.69	172
Oil endowment (log p.c.)	−10.17	2.76	−11.93	−0.31	172
Oil quality	3.44	3.28	1.00	10.44	127
Convergent C-C mechanical area (log p.c.)	−8.36	2.97	−10.34	0.20	172
Convergent O-C thermal area (log p.c.)	−8.85	0.70	−8.99	−4.22	172
Convergent O-C mechanical area (log p.c.)	−9.62	3.69	−11.92	−1.06	172
Convergent ocean-ocean area (log p.c.)	−9.80	2.38	−10.66	−0.31	172
Divergent thermal area (log p.c.)	−6.51	4.91	−13.75	1.27	172
Wrench mechanical area (log p.c.)	−13.25	3.66	−14.94	−0.73	172
Divergent mechanical area (log p.c.)	−10.64	2.68	−12.10	−2.02	172
Convergent C-C thermo-mechanical area (log p.c.)	−8.52	0.92	−8.75	−2.91	172
Foreland area (log p.c.)	−7.07	2.93	−9.60	0.20	172
Intracratonic sag area (log p.c.)	−11.00	5.21	−14.70	−0.03	172
Passive margin area (log p.c.)	−8.78	5.19	−13.75	1.27	172
Convergent sag area (log p.c.)	−15.09	3.20	−16.14	−2.91	172
Post-rift sag area (log p.c.)	−10.18	3.56	−12.52	−0.73	172
Wrench area (log p.c.)	−13.25	3.66	−14.94	−0.73	172
Extensional area (log p.c.)	−9.90	2.07	−10.68	−1.60	172
Convergent wrench area (log p.c.)	−11.67	3.20	−13.12	−0.32	172
Fore-arc area (log p.c.)	−10.35	3.16	−11.92	−0.80	172
Rift area (log p.c.)	−10.64	2.68	−12.10	−2.02	172
Land area (log p.c.)	−3.23	1.58	−7.79	0.49	172
Coastline (log p.c.)	−9.35	2.99	−14.01	−3.17	172
Mountainous area (log p.c.)	−6.54	2.99	−11.21	−0.53	172
Tropical area (log p.c.)	−6.21	3.85	−10.47	−0.08	172
Good soil area (log p.c.)	−6.79	2.64	−11.49	−0.37	172
Year of independence	1826.21	408.38	−660.00	1993.00	172

Notes. See the appendix for variable definitions. Due to the presence of zero values, the "log" transformation of the oil and geographic variables is actually a differentiable and monotonic transformation $h(w) = \log(w)$ for $w > w_0$ and $h(w) = \log(w_0) - 1 + w/w_0$ for $w \leq w_0$. This function was suggested by James Hamilton of UC San Diego. In practice w_0 is chosen for each variable as the minimum positive value observed in the sample.

Table 2.5: Summary Statistics by Oil Presence

	Oil Countries	Non-Oil Countries	Difference	p-value
Democracy, 2008	0.66	0.73	−0.07	0.182
Democracy, 1966	0.43	0.46	−0.03	0.677
Avg. democracy, 1966–2008	0.52	0.54	−0.02	0.683
Corruption, 2008	2.62	2.48	0.14	0.508
Internal conflicts per year, 1966–2008	0.27	0.13	0.13^{*}	0.068
Coup attempts per year, 1966–2008	0.05	0.07	−0.01	0.330
Purges per year, 1966–2008	0.08	0.04	0.04^{*}	0.067
Total revenue, 2000–2008 (log avg.)	−1.41	−1.68	0.27^{***}	0.000
Total tax revenue, 2000–2008 (log avg.)	−2.04	−1.89	−0.14	0.179
GDP, 2008 (log p.c.)	9.46	8.58	0.88^{***}	0.000
GDP, 1966 (log p.c.)	8.01	7.24	0.77^{***}	0.000
Non-Oil GDP, 2008 (log p.c.)	9.34	9.11	0.23	0.300
Non-Oil/Gas GDP, 2008 (log p.c.)	9.36	9.11	0.25	0.256
Non-Resource GDP, 2008 (log p.c.)	9.27	8.51	0.76^{***}	0.000
Executive constraints, 1950–1965	0.48	0.45	0.04	0.594
Oil production, 1966–2008 (log avg. p.c.)	−1.60	−9.03	7.42^{***}	0.000
Oil discovery, 1966–2003 (log avg. p.c.)	−7.36	−11.14	3.79^{***}	0.000
Oil reserves, 1966–2003 (log avg. p.c.)	−3.64	−7.30	3.67^{***}	0.000
Oil endowment (log p.c.)	−8.78	−11.93	3.15^{***}	0.000
Oil quality	7.07	1.00	6.07^{***}	0.000
Convergent C-C mechanical area (log p.c.)	−7.37	−9.61	2.23^{***}	0.000
Convergent O-C thermal area (log p.c.)	−8.74	−8.99	0.25^{**}	0.021
Convergent O-C mechanical area (log p.c.)	−8.81	−10.64	1.83^{***}	0.001
Convergent ocean-ocean area (log p.c.)	−9.92	−9.63	−0.29	0.428
Divergent thermal area (log p.c.)	−5.90	−7.29	1.38^{*}	0.066
Wrench mechanical area (log p.c.)	−12.82	−13.80	0.97^{*}	0.084
Divergent mechanical area (log p.c.)	−10.43	−10.89	0.46	0.264
Convergent C-C thermo-mechanical area (log p.c.)	−8.40	−8.68	0.28^{**}	0.047
Foreland area (log p.c.)	−5.96	−8.47	2.51^{***}	0.000
Intracratonic sag area (log p.c.)	−10.05	−12.21	2.16^{***}	0.007
Passive margin area (log p.c.)	−8.26	−9.45	1.19	0.136
Convergent sag area (log p.c.)	−14.39	−15.98	1.59^{***}	0.001
Post-rift sag area (log p.c.)	−9.57	−10.97	1.40^{**}	0.010
Wrench area (log p.c.)	−12.82	−13.80	0.97^{*}	0.084
Extensional area (log p.c.)	−9.82	−9.99	0.17	0.589
Convergent wrench area (log p.c.)	−11.47	−11.93	0.46	0.355
Fore-arc area (log p.c.)	−10.03	−10.75	0.72	0.141
Rift area (log p.c.)	−10.43	−10.89	0.46	0.264
Land area (log p.c.)	−3.09	−3.41	0.32	0.185
Coastline (log p.c.)	−9.22	−9.51	0.29	0.528
Mountainous area (log p.c.)	−6.23	−6.92	0.69	0.132
Tropical area (log p.c.)	−6.77	−5.49	-1.28^{**}	0.030
Good soil area (log p.c.)	−6.70	−6.91	0.21	0.599
Year of independence	1770.61	1896.43	-125.82^{**}	0.044
Observations	96	76		

Notes. This table defines oil countries as those countries that had positive oil production at any time from 1966–2008. Averages are reported in the first two columns. The third column reports the difference of the averages, and the fourth column reports the p-value corresponding to the two-sided test of equality of the averages. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.6: First-Stage Estimates for Optimal Sets of Basin Measures by Plate-Tectonic Environment and Primary Mechanism of Subsidence

	Log Avg. Oil Production per capita, 1966–2008							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Convergent C-C mechanical	0.599***	0.592***	0.589***	0.584***	0.607***	0.610***	0.609***	0.604***
	(0.119)	(0.119)	(0.124)	(0.124)	(0.125)	(0.124)	(0.126)	(0.127)
Convergent O-C thermal		0.589***		0.371**	0.340*	0.285	0.267	0.271
		(0.175)		(0.168)	(0.176)	(0.185)	(0.194)	(0.194)
Convergent O-C mechanical			0.359***	0.337***	0.320***	0.329***	0.327***	0.329***
			(0.084)	(0.087)	(0.088)	(0.091)	(0.091)	(0.091)
Convergent O-O			−0.362**	−0.377***	−0.344**	−0.345**	−0.341**	−0.341**
			(0.139)	(0.142)	(0.149)	(0.151)	(0.152)	(0.153)
Divergent thermal					0.058	0.062	0.060	0.059
					(0.069)	(0.069)	(0.071)	(0.071)
Wrench mechanical						0.070	0.073	0.075
						(0.072)	(0.074)	(0.075)
Divergent mechanical							0.026	0.026
							(0.116)	(0.117)
Convergent C-C thermo-mechanical								0.059
								(0.327)
Observations	157	157	157	157	157	157	157	157
R^2	0.318	0.327	0.394	0.398	0.400	0.403	0.404	0.404
F statistic	25.3	17.6	18.9	16.5	14.0	12.6	10.8	9.4

Notes. See Tables 2.1 and 2.2 for basin variable definitions. See the appendix for other variable definitions. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.7: First-Stage Estimates for Optimal Sets of Basin Measures by Final Component of Fugro Tellus Code

	Log Avg. Oil Production per capita, 1966–2008									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Foreland	0.576***	0.608***	0.613***	0.585***	0.677***	0.677***	0.701***	0.710***	0.700***	0.701***
	(0.142)	(0.143)	(0.139)	(0.140)	(0.144)	(0.142)	(0.141)	(0.139)	(0.142)	(0.142)
Intracratonic sag		0.213***	0.209***	0.201***	0.155**	0.153**	0.164**	0.159**	0.159**	0.164**
		(0.069)	(0.068)	(0.068)	(0.073)	(0.073)	(0.074)	(0.075)	(0.075)	(0.077)
Passive margin			0.091	0.102		0.080	0.086	0.085	0.086	0.080
			(0.076)	(0.076)		(0.076)	(0.077)	(0.076)	(0.077)	(0.077)
Convergent sag				0.199***	0.164***	0.174***	0.183***	0.190***	0.189***	0.184***
				(0.063)	(0.061)	(0.064)	(0.063)	(0.066)	(0.066)	(0.066)
Post-rift sag					0.231**	0.218**	0.233**	0.227**	0.228**	0.213**
					(0.101)	(0.100)	(0.101)	(0.102)	(0.103)	(0.106)
Wrench					0.119	0.119	0.122	0.122	0.122	0.129
					(0.079)	(0.078)	(0.078)	(0.078)	(0.079)	(0.079)
Extensional							−0.196	−0.182	−0.199	−0.224
							(0.173)	(0.172)	(0.180)	(0.188)
Convergent wrench								−0.049	−0.069	−0.061
								(0.118)	(0.110)	(0.109)
Fore-arc									0.077	0.077
									(0.107)	(0.105)
Rift										0.084
										(0.107)
Observations	157	157	157	157	157	157	157	157	157	157
R^2	0.315	0.357	0.364	0.383	0.409	0.415	0.420	0.421	0.422	0.424
F statistic	16.4	17.0	14.4	12.6	11.7	10.3	9.5	8.4	7.4	6.6

Notes. See Tables 2.1 and 2.3 for basin variable definitions. See the appendix for other variable definitions. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.8: Testing for a Political Resource Curse

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Ordinary Least Squares							
Oil production	−0.019***	−0.014***	−0.032	0.012**	0.001	0.003	0.032***	−0.044***
	(0.005)	(0.004)	(0.023)	(0.006)	(0.002)	(0.002)	(0.007)	(0.012)
Observations	157	160	136	172	160	172	165	167
R^2	0.441	0.536	0.334	0.204	0.203	0.093	0.463	0.471
	Panel B: Two-Stage Least Squares							
Oil production	−0.038**	−0.039***	−0.136**	0.007	0.002	−0.001	0.021	−0.163***
	(0.017)	(0.015)	(0.060)	(0.018)	(0.003)	(0.005)	(0.016)	(0.039)
Observations	157	160	136	172	160	172	165	167
F statistic	25.3	26.7	23.2	31.4	26.7	31.4	29.3	27.3
A-R 95% CI	[−0.081, −0.008]	[−0.077, −0.014]	[−0.280, −0.027]	[−0.030, 0.045]	[−0.004, 0.009]	[−0.011, 0.010]	[−0.015, 0.053]	[−0.268, −0.100]
Oil exog.	0.209	0.063	0.053	0.780	0.578	0.445	0.446	0.000

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.9: Testing for an Economic Resource Course

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Ordinary Least Squares				
Oil production	0.092***	0.043**	0.040**	0.071***	0.076***
	(0.015)	(0.017)	(0.017)	(0.016)	(0.021)
Observations	166	132	129	166	145
R^2	0.661	0.623	0.608	0.657	0.599
	Panel B: Two-Stage Least Squares				
Oil production	0.074*	0.045	0.039	0.054	−0.037
	(0.040)	(0.044)	(0.044)	(0.041)	(0.075)
Observations	166	132	129	166	145
F statistic	29.3	20.4	20.1	29.3	14.3
A-R 95% CI	[−0.010, 0.157]	[−0.047, 0.139]	[−0.055, 0.132]	[−0.033, 0.137]	[−0.246, 0.097]
Oil exog.	0.632	0.963	0.967	0.646	0.079

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.10: Partial Correlation between *Basin* and Predetermined Variables

	Urbanization, 1850	British Legal Origin	Communist Legacy	Percentage Christian, 1950	Percentage Muslim, 1950	Percentage Hindu, 1950	Fractionalization: Ethnic	Fractionalization: Religious	Fractionalization: Linguistic
Convergent C-C mechanical	−0.050	0.003	−0.014	−0.357	5.759***	0.015	0.015*	0.002	0.008
	(0.277)	(0.015)	(0.012)	(0.597)	(0.865)	(0.214)	(0.008)	(0.008)	(0.007)
Observations	84	163	172	172	172	172	171	172	165
R^2	0.449	0.086	0.060	0.800	0.495	0.078	0.407	0.093	0.402

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

Table 2.11: Testing for a Political Resource Curse (Controlling for Percentage Muslim in 1950)

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Ordinary Least Squares							
Oil production	-0.014^{***}	-0.011^{**}	-0.018	0.012^{*}	-0.000	0.003	0.034^{***}	-0.033^{***}
	(0.005)	(0.004)	(0.022)	(0.006)	(0.002)	(0.002)	(0.007)	(0.010)
Observations	157	160	136	172	160	172	165	167
R^2	0.477	0.567	0.359	0.204	0.231	0.093	0.464	0.530
	Panel B: Two-Stage Least Squares							
Oil production	-0.016	-0.025	-0.093	0.003	-0.003	-0.001	0.023	-0.138^{***}
	(0.019)	(0.015)	(0.077)	(0.022)	(0.005)	(0.006)	(0.020)	(0.044)
Observations	157	160	136	172	160	172	165	167
F statistic	12.4	15.3	10.6	18.9	15.3	18.9	16.3	14.8
A-R 95% CI	[−0.062, 0.025]	[−0.067, 0.004]	[−0.298, 0.073]	[−0.047, 0.048]	[−0.014, 0.006]	[−0.015, 0.012]	[−0.025, 0.064]	[−0.274, −0.068]
Oil exog.	0.893	0.339	0.318	0.675	0.470	0.488	0.585	0.006

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

Table 2.12: Testing for an Economic Resource Course (Controlling for Percentage Muslim in 1950)

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Ordinary Least Squares				
Oil production	0.099***	0.049***	0.045**	0.078***	0.090***
	(0.015)	(0.017)	(0.017)	(0.015)	(0.021)
Observations	166	132	129	166	145
R^2	0.668	0.628	0.611	0.665	0.616
	Panel B: Two-Stage Least Squares				
Oil production	0.112**	0.087	0.074	0.096*	0.011
	(0.051)	(0.060)	(0.061)	(0.052)	(0.097)
Observations	166	132	129	166	145
F statistic	17.1	11.7	10.8	17.1	6.8
A-R 95% CI	$[0.005, 0.228]$	$[-0.033, 0.245]$	$[-0.058, 0.234]$	$[-0.012, 0.216]$	$[-0.344, 0.224]$
Oil exog.	0.793	0.499	0.631	0.735	0.401

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.13: Testing for a Political Resource Curse (Controlling for Ethnic Fractionalization)

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Ordinary Least Squares							
Oil production	−0.019***	−0.014***	−0.030	0.011*	0.001	0.003	0.035***	−0.042***
	(0.005)	(0.004)	(0.023)	(0.006)	(0.002)	(0.002)	(0.007)	(0.012)
Observations	156	159	135	171	159	171	164	166
R^2	0.439	0.537	0.336	0.216	0.207	0.093	0.510	0.479
	Panel B: Two-Stage Least Squares							
Oil production	−0.040**	−0.041***	−0.133**	0.003	0.002	−0.000	0.027	−0.165***
	(0.018)	(0.015)	(0.060)	(0.019)	(0.003)	(0.005)	(0.017)	(0.041)
Observations	156	159	135	171	159	171	164	166
F statistic	22.4	24.6	21.3	28.0	24.6	28.0	26.2	24.4
A-R 95% CI	[−0.086, −0.009]	[−0.080, −0.014]	[−0.282, −0.023]	[−0.038, 0.041]	[−0.004, 0.010]	[−0.011, 0.011]	[−0.009, 0.061]	[−0.279, −0.099]
Oil exog.	0.181	0.059	0.058	0.664	0.579	0.476	0.613	0.001

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.14: Testing for an Economic Resource Course (Controlling for Ethnic Fractionalization)

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Ordinary Least Squares				
Oil production	0.097***	0.047***	0.044***	0.076***	0.083***
	(0.015)	(0.017)	(0.017)	(0.015)	(0.021)
Observations	165	131	128	165	144
R^2	0.674	0.637	0.621	0.667	0.618
	Panel B: Two-Stage Least Squares				
Oil production	0.091**	0.059	0.052	0.070*	−0.005
	(0.040)	(0.043)	(0.043)	(0.041)	(0.075)
Observations	165	131	128	165	144
F statistic	26.1	18.8	18.4	26.1	11.9
A-R 95% CI	[0.008, 0.177]	[−0.030, 0.156]	[−0.038, 0.148]	[−0.015, 0.157]	[−0.220, 0.139]
Oil exog.	0.878	0.773	0.846	0.883	0.186

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.15: Testing for a Political Resource Curse: Subsample with Predetermined Borders

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Ordinary Least Squares							
Oil production	−0.017***	−0.018***	−0.044	0.014**	0.002	0.002	0.020**	−0.043***
	(0.006)	(0.006)	(0.030)	(0.006)	(0.002)	(0.002)	(0.009)	(0.014)
Observations	96	97	80	108	97	108	104	105
R^2	0.395	0.580	0.376	0.219	0.201	0.119	0.441	0.467
	Panel B: Two-Stage Least Squares							
Oil production	−0.035**	−0.042***	−0.151**	0.009	0.001	−0.000	0.006	−0.138***
	(0.014)	(0.012)	(0.068)	(0.012)	(0.003)	(0.008)	(0.019)	(0.047)
Observations	96	97	80	108	97	108	104	105
F statistic	21.7	21.7	20.4	27.8	21.7	27.8	30.3	25.4
A-R 95% CI	[−0.067, −0.008]	[−0.071, −0.021]	[−0.311, −0.018]	[−0.016, 0.035]	[−0.006, 0.009]	[−0.017, 0.017]	[−0.035, 0.044]	[−0.260, −0.058]
Oil exog.	0.157	0.038	0.086	0.637	0.698	0.809	0.387	0.033

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.16: Testing for an Economic Resource Course: Subsample with Predetermined Borders

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Ordinary Least Squares				
Oil production	0.081***	0.036	0.036	0.062***	0.049
	(0.020)	(0.023)	(0.023)	(0.021)	(0.031)
Observations	105	73	70	105	89
R^2	0.686	0.647	0.632	0.678	0.596
	Panel B: Two-Stage Least Squares				
Oil production	0.112**	0.099*	0.096	0.107**	0.006
	(0.050)	(0.059)	(0.060)	(0.051)	(0.084)
Observations	105	73	70	105	89
F statistic	30.4	19.4	18.5	30.4	13.9
A-R 95% CI	$[0.020, 0.229]$	$[-0.005, 0.254]$	$[-0.008, 0.252]$	$[0.013, 0.227]$	$[-0.181, 0.196]$
Oil exog.	0.484	0.225	0.239	0.331	0.577

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. The p-value of the test of the endogeneity of oil wealth is from the Hansen (1982) overidentification test of the null hypothesis that oil wealth is exogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.17: Testing for a Political Resource Curse: Basin vs. Endowment

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Ordinary Least Squares							
Oil production	−0.019***	−0.014***	−0.032	0.012**	0.001	0.003	0.032***	−0.044***
	(0.005)	(0.004)	(0.023)	(0.006)	(0.002)	(0.002)	(0.007)	(0.012)
Observations	157	160	136	172	160	172	165	167
R^2	0.441	0.536	0.334	0.204	0.203	0.093	0.463	0.471
	Panel B: 2SLS using Endowment							
Oil production	−0.026***	−0.018***	−0.034	0.003	0.001	0.001	0.043***	−0.080***
	(0.005)	(0.005)	(0.024)	(0.009)	(0.002)	(0.002)	(0.008)	(0.017)
Observations	157	160	136	172	160	172	165	167
F statistic	326.8	350.4	286.1	409.6	350.4	409.6	401.2	395.2
	Panel C: 2SLS using Basin							
Oil production	−0.038**	−0.039***	−0.136**	0.007	0.002	−0.001	0.021	−0.163***
	(0.017)	(0.015)	(0.060)	(0.018)	(0.003)	(0.005)	(0.016)	(0.039)
Observations	157	160	136	172	160	172	165	167
F statistic	25.3	26.7	23.2	31.4	26.7	31.4	29.3	27.3
Overident. p-value	0.418	0.084	0.063	0.752	0.547	0.708	0.129	0.003

Notes. See the appendix for variable definitions. Panel A presents OLS estimates for comparison. Panel B presents IV estimates using initial oil endowment as an instrument for oil production. Panel C presents IV estimates using Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The Hansen (1982) overidentification test p-value corresponds to the null hypothesis that both Endowment and Basin 2 are valid instruments. Assuming that Basin 2 is a valid instrument, rejection implies that Endowment is endogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.18: Testing for an Economic Resource Course: Basin vs. Endowment

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Ordinary Least Squares				
Oil production	0.092***	0.043**	0.040**	0.071***	0.076***
	(0.015)	(0.017)	(0.017)	(0.016)	(0.021)
Observations	166	132	129	166	145
R^2	0.661	0.623	0.608	0.657	0.599
	Panel B: 2SLS using Endowment				
Oil production	0.114***	0.066***	0.061***	0.085***	0.070***
	(0.016)	(0.017)	(0.017)	(0.016)	(0.024)
Observations	166	132	129	166	145
F statistic	407.9	224.3	424.4	407.9	322.3
	Panel C: 2SLS using Basin				
Oil production	0.074*	0.045	0.039	0.054	−0.037
	(0.040)	(0.044)	(0.044)	(0.041)	(0.075)
Observations	166	132	129	166	145
F statistic	29.3	20.4	20.1	29.3	14.3
Overident. p-value	0.273	0.587	0.563	0.373	0.093

Notes. See the appendix for variable definitions. Panel A presents OLS estimates for comparison. Panel B presents IV estimates using initial oil endowment as an instrument for oil production. Panel C presents IV estimates using Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The Hansen (1982) overidentification test p-value corresponds to the null hypothesis that both Endowment and Basin 2 are valid instruments. Assuming that Basin 2 is a valid instrument, rejection implies that Endowment is endogenous. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.19: Political Resource Curse: Heterogeneous Effects by Initial Institutional Quality

	Democracy, 2008	Avg. Democracy, 1966–2008	Corruption, 2008	Internal Conflict, 1966–2008	Coup Attempts, 1966–2008	Purges, 1966–2008	Total Revenue, 2000–2008	Tax Revenue, 2000–2008
	Panel A: Countries with Relatively Strong Executive Constraints from 1950–1965							
Oil production	−0.027	−0.009	0.054	0.010	−0.000	0.018	0.057**	−0.012
	(0.019)	(0.018)	(0.070)	(0.037)	(0.005)	(0.012)	(0.025)	(0.030)
Observations	53	54	51	54	54	54	52	53
F statistic	4.2	5.5	6.0	5.5	5.5	5.5	4.1	4.2
	Panel B: Countries with Relatively Weak Executive Constraints from 1950–1965							
Oil production	−0.044**	−0.022**	−0.033	0.035**	0.007	0.013*	0.071***	−0.108**
	(0.017)	(0.010)	(0.045)	(0.017)	(0.006)	(0.007)	(0.022)	(0.047)
Observations	60	62	54	62	62	62	58	60
F statistic	10.2	12.7	10.6	12.7	12.7	12.7	11.1	12.4
	Panel C: Difference between Estimates							
Difference	0.017	0.013	0.087	−0.025	−0.007	0.005	−0.014	0.096
	(0.054)	(0.029)	(0.134)	(0.063)	(0.010)	(0.025)	(0.072)	(0.079)
p-value	0.754	0.658	0.518	0.689	0.491	0.836	0.846	0.227

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. Column titles refer to the sample of countries used in the regression. Countries in the "Strong" subsample averaged strictly greater than three points out of seven on the executive constraints index, XCONST (Polity IV), from 1950–1965. Countries in the "Weak" subsample averaged three points or fewer out of seven on XCONST from 1950–1965. A score of three points for XCONST indicates "slight to moderate limitation on executive authority" (Polity IV). In practice "Weak" indicates having an average XCONST score equal to or below the median average XCONST score from 1950–1965. The standard errors and p-values in Panel C are calculated by a bootstrap procedure based on 200 repetitions. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.20: Economic Resource Curse: Heterogeneous Effects by Initial Institutional Quality

	GDP, 2008	Non-Oil GDP, 2008	Non-Oil/Gas GDP, 2008	Non-Resource GDP, 2008	Manufacturing GDP, 2008
	Panel A: Countries with Relatively Strong Executive Constraints from 1950–1965				
Oil production	0.117	0.061	0.059	0.104	0.122
	(0.075)	(0.066)	(0.068)	(0.078)	(0.118)
Observations	51	46	45	51	50
F statistic	3.4	2.2	2.1	3.4	3.3
	Panel B: Countries with Relatively Weak Executive Constraints from 1950–1965				
Oil production	0.152***	0.109**	0.104*	0.132***	0.149***
	(0.040)	(0.055)	(0.053)	(0.045)	(0.057)
Observations	59	48	47	59	50
F statistic	12.0	6.9	7.0	12.0	6.8
	Panel C: Difference between Estimates				
Difference	−0.035	−0.048	−0.045	−0.029	−0.027
	(0.112)	(0.168)	(0.137)	(0.114)	(0.284)
p-value	0.753	0.774	0.745	0.800	0.925

Notes. See the appendix for variable definitions. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. The IV specifications use Basin 2 as an instrument for oil production. The F statistic is the Kleibergen and Paap (2006) rk statistic, which tests for weak identification and is robust to heteroskedasticity. The A-R 95% confidence interval is based on the Anderson and Rubin (1949) χ^2 test of the null hypothesis that the coefficients on the endogenous variables in the structural equation are jointly equal to zero. The A-R test is robust to the presence of weak instruments. Column titles refer to the sample of countries used in the regression. Countries in the "Strong" subsample averaged strictly greater than three points out of seven on the executive constraints index, XCONST (Polity IV), from 1950–1965. Countries in the "Weak" subsample averaged three points or fewer out of seven on XCONST from 1950–1965. A score of three points for XCONST indicates "slight to moderate limitation on executive authority" (Polity IV). In practice "Weak" indicates having an average XCONST score equal to or below the median average XCONST score from 1950–1965. The standard errors and p-values in Panel C are calculated by a bootstrap procedure based on 200 repetitions. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.9 Figures

Figure 2.1: Sedimentary Basins and Giant Oil and Gas Fields

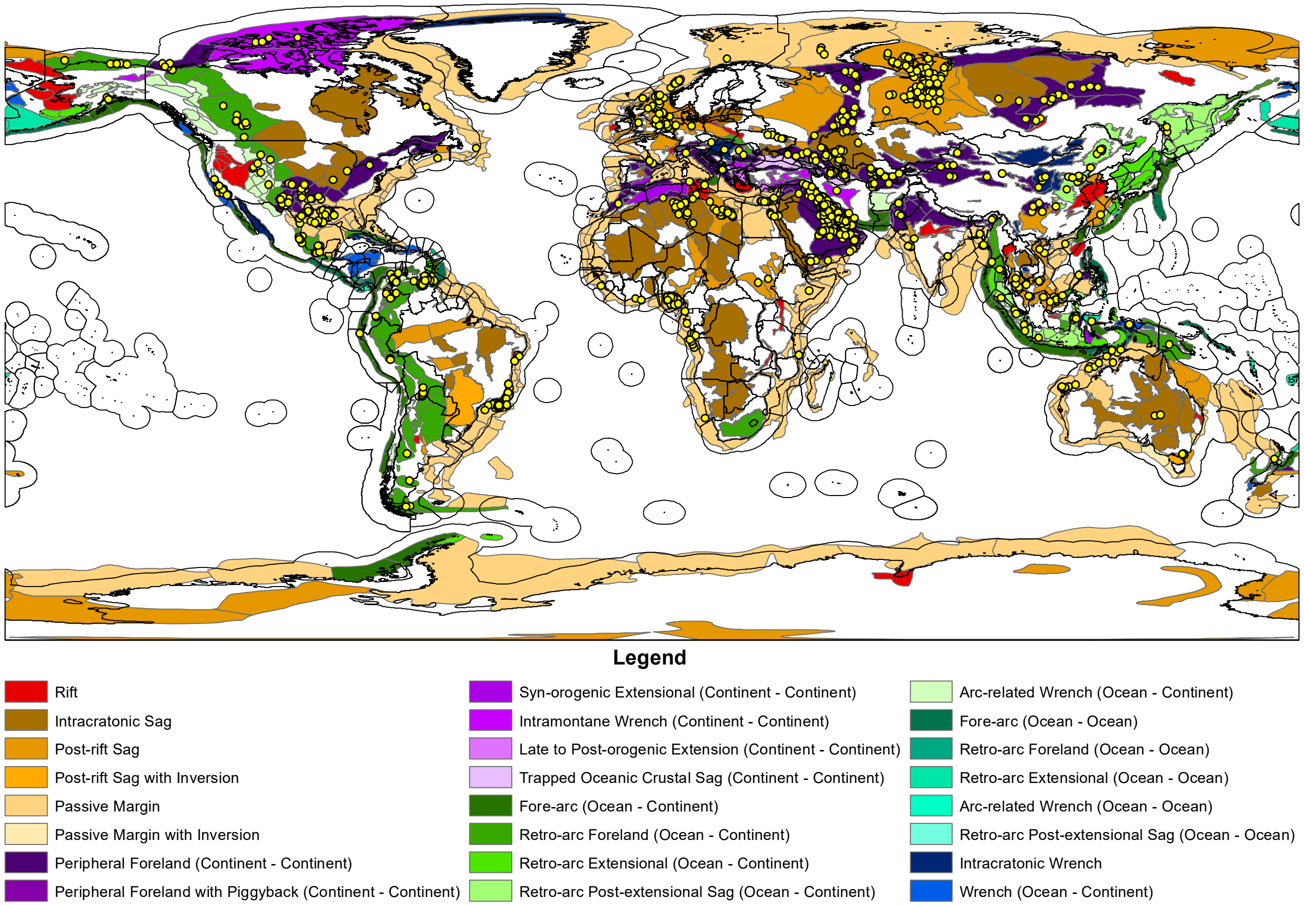

Notes. Colored areas represent sedimentary basins, and yellow dots represent giant oil and gas fields. The GIS data on sedimentary basins come from Fugro Robertson, Ltd. (2013), and the GIS data on giant oil and gas fields come from Horn (2004).

Figure 2.2: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping

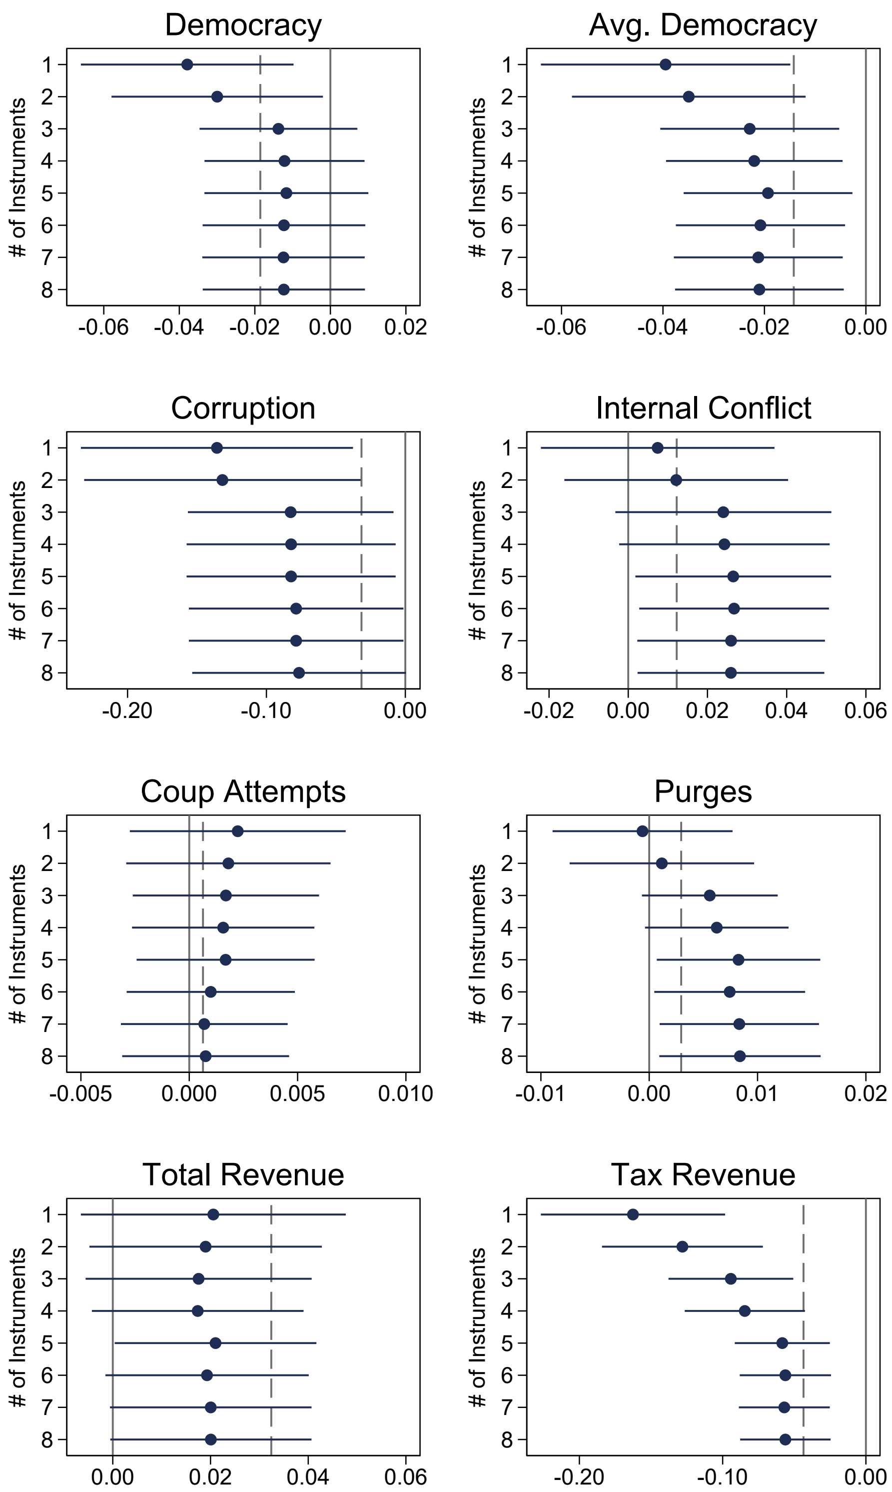

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.3: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping

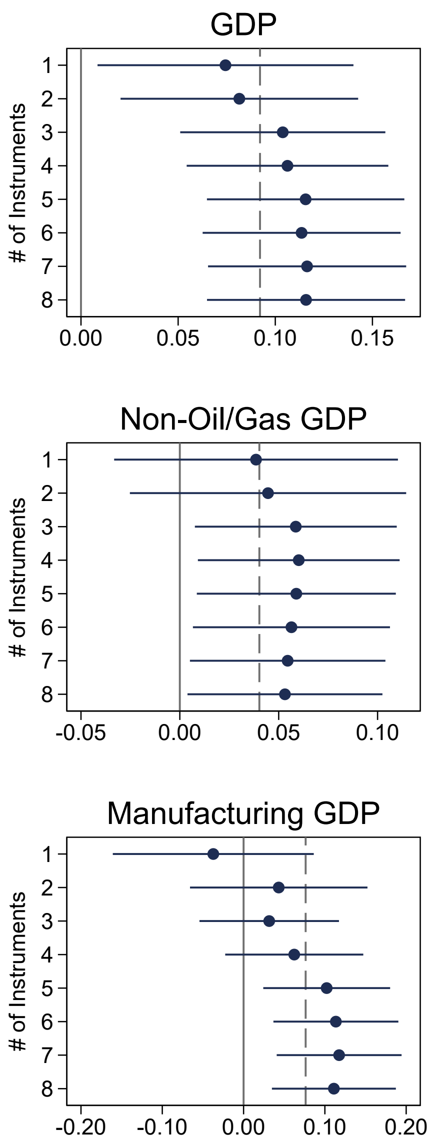

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.4: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping

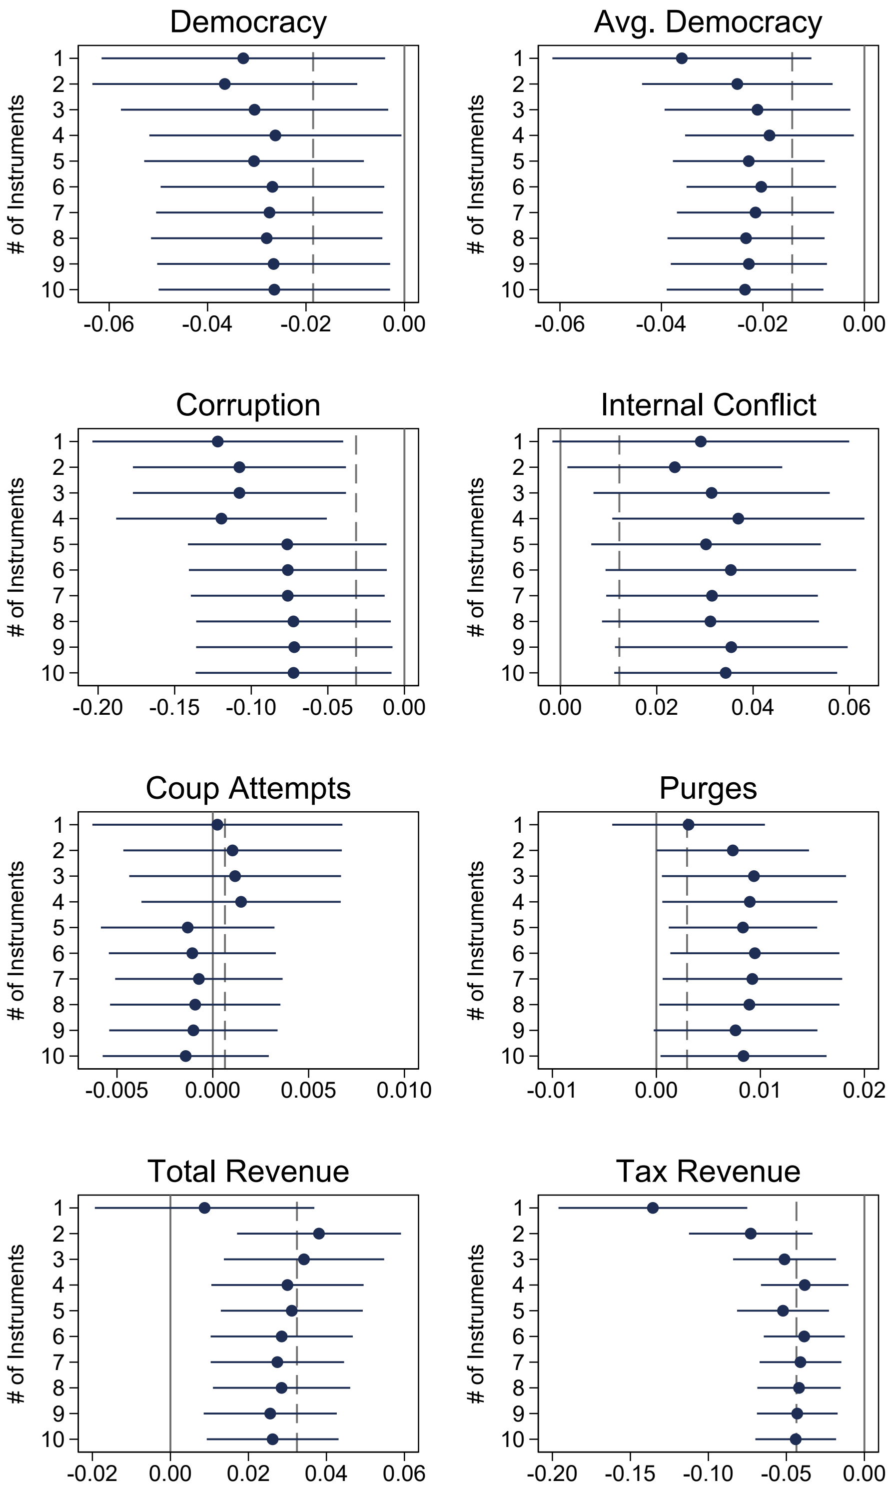

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.5: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping

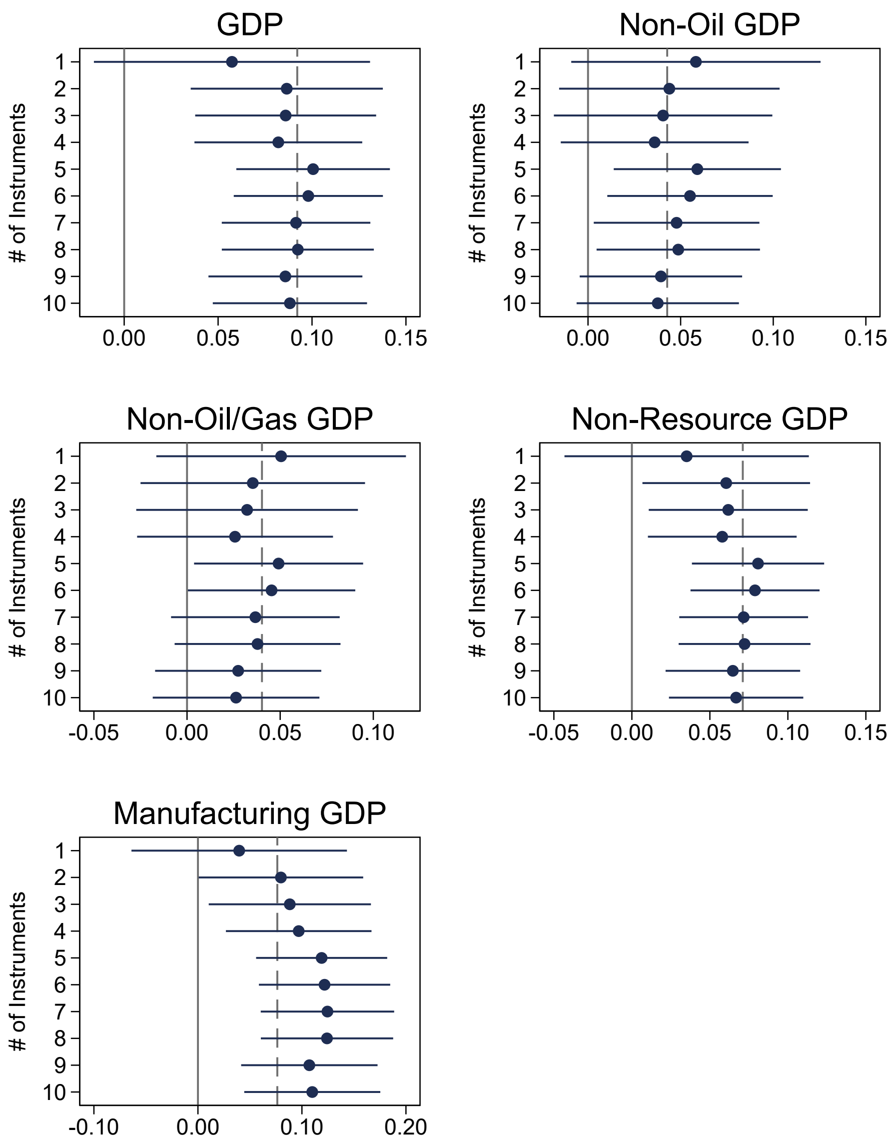

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.6: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping (Controlling for Percentage Muslim in 1950)

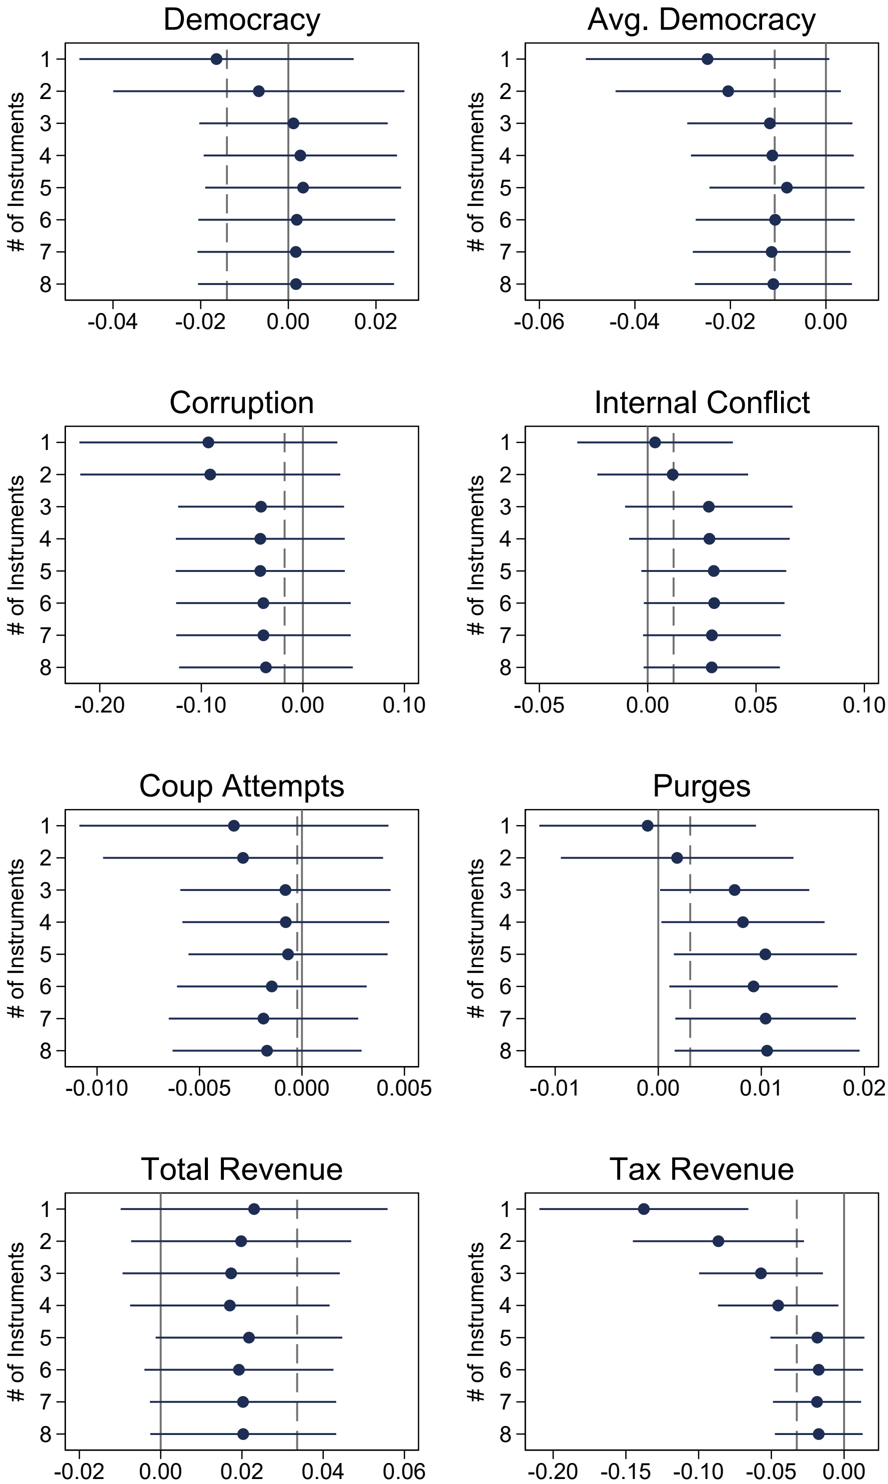

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.7: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping (Controlling for Percentage Muslim in 1950)

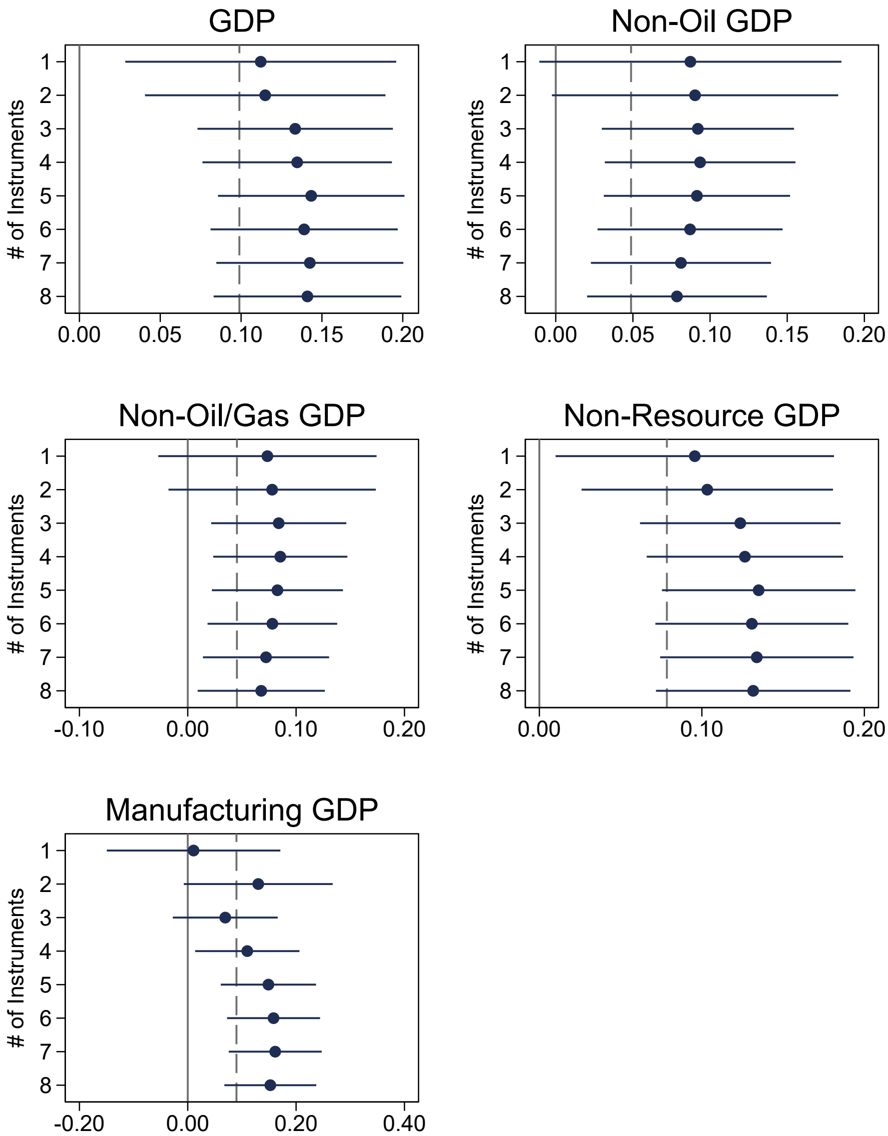

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.8: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping (Controlling for Percentage Muslim in 1950)

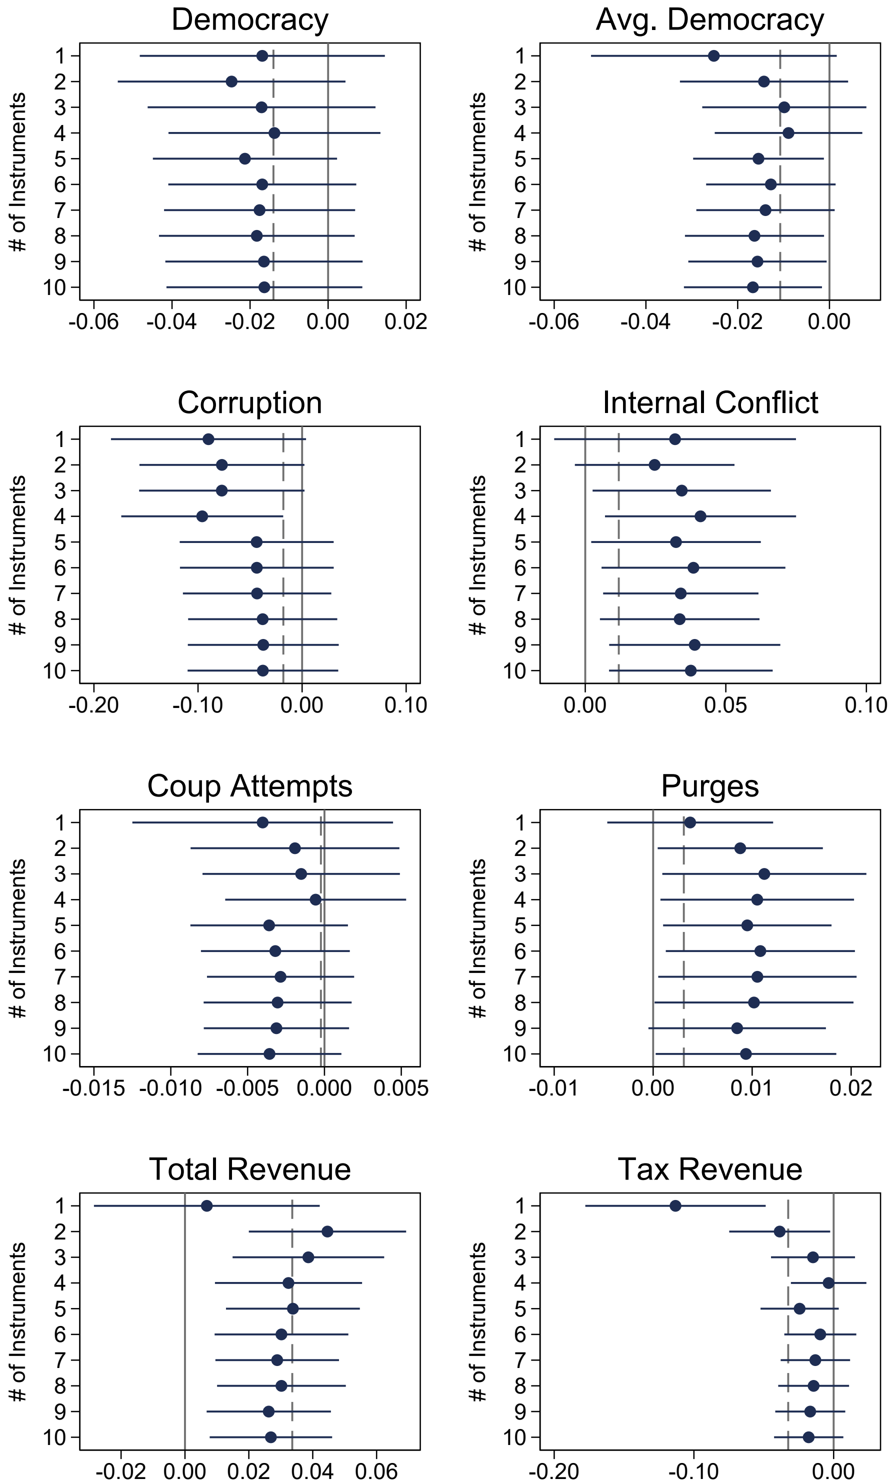

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.9: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping (Controlling for Percentage Muslim in 1950)

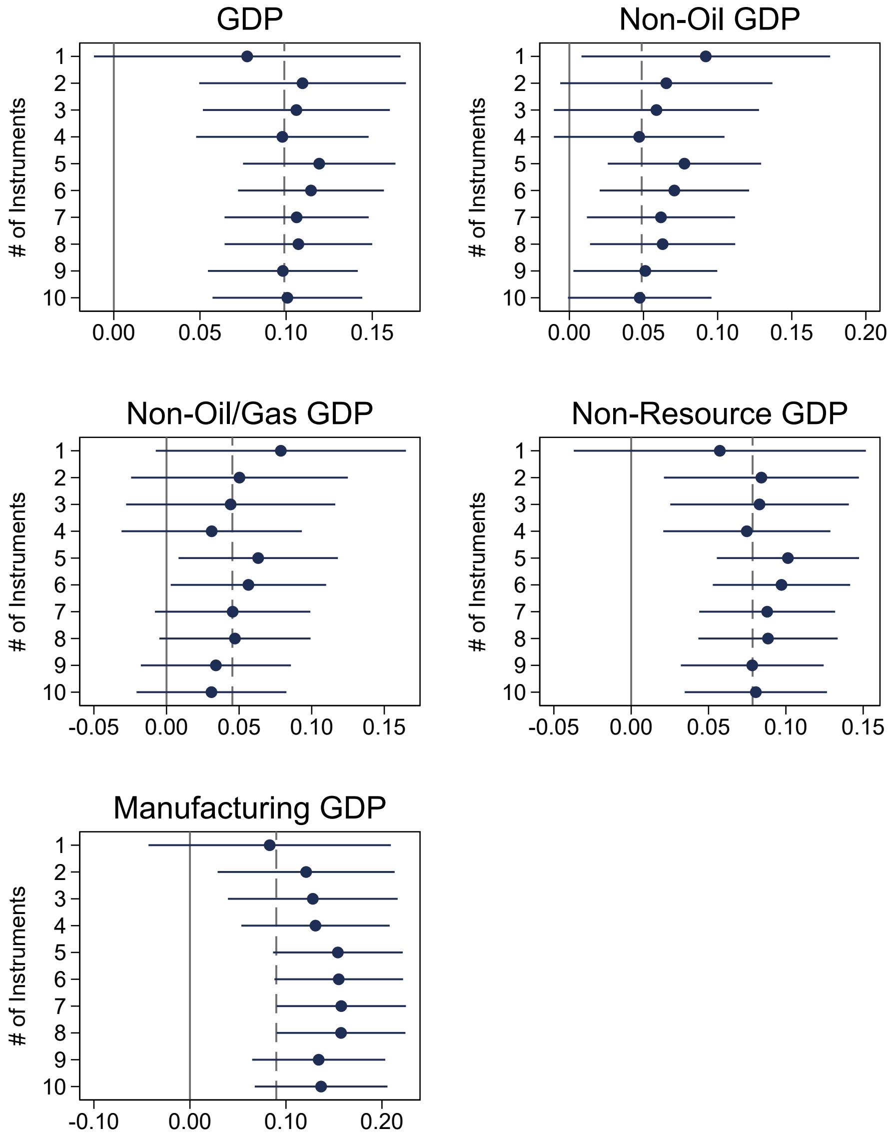

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.10: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping (Controlling for Ethnic Fractionalization)

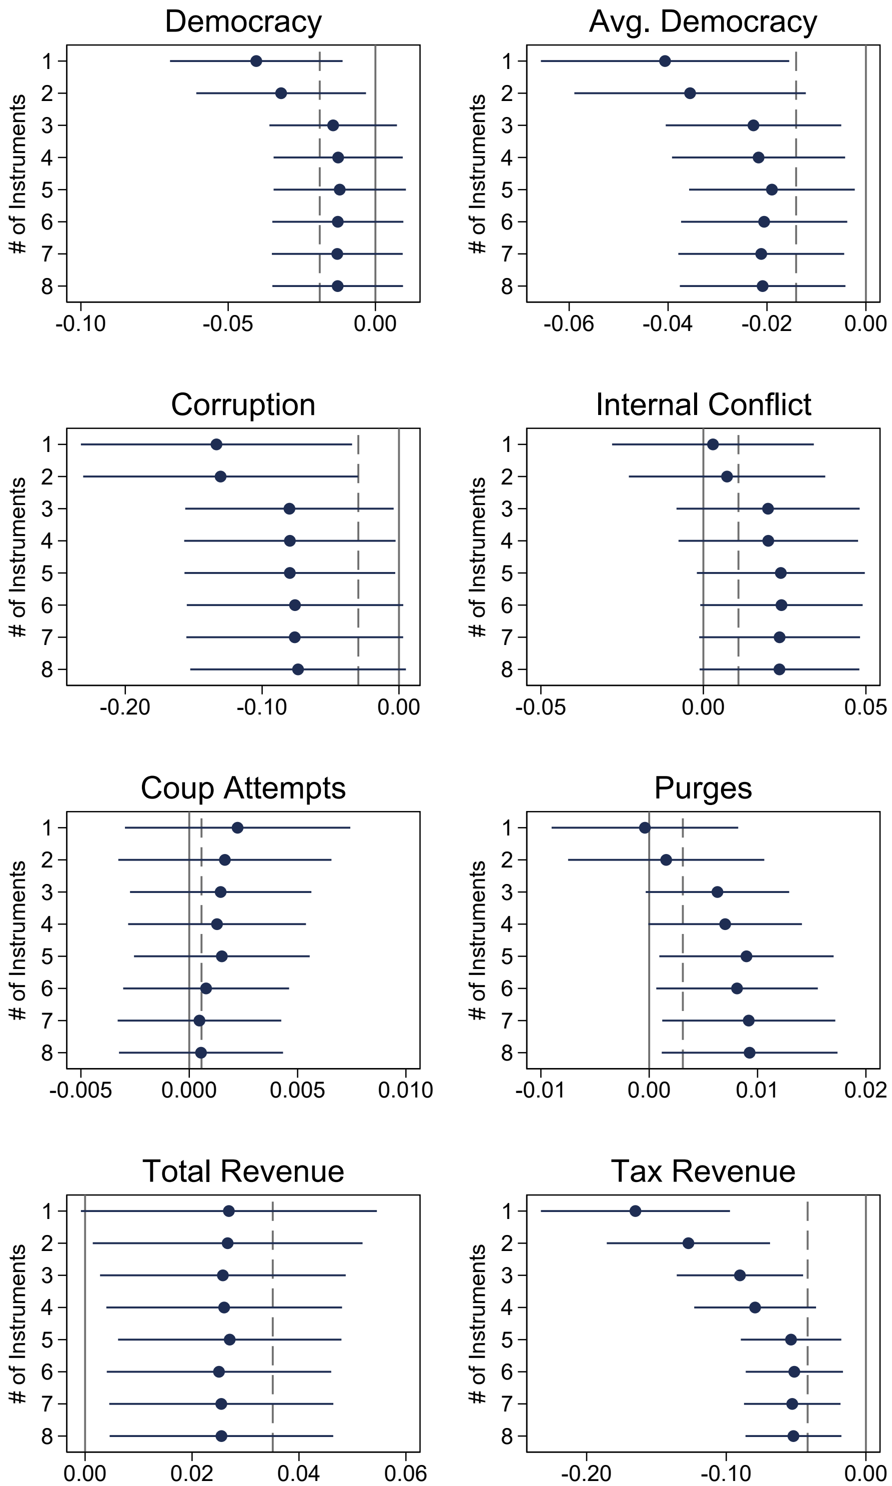

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.11: 2SLS Estimates by Size of Instrument Set, Tectonic-Subsidence Grouping (Controlling for Ethnic Fractionalization)

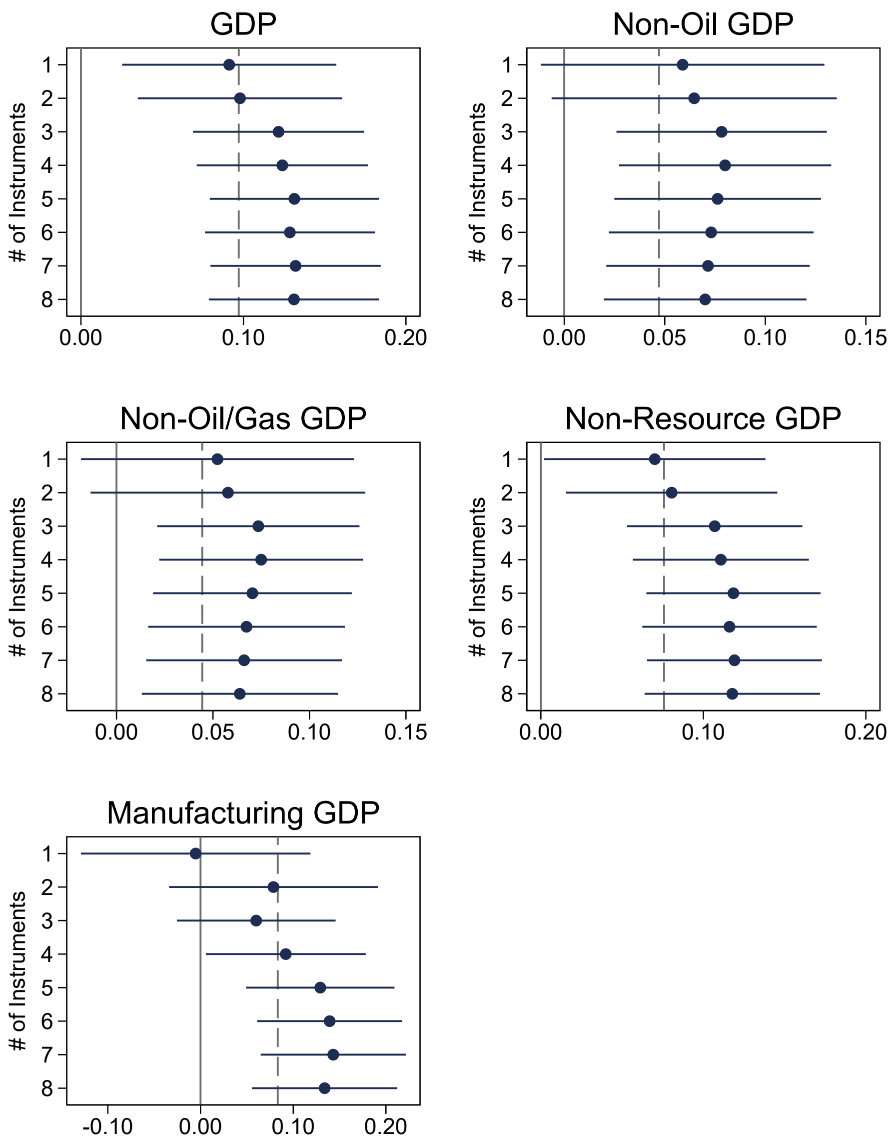

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.12: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping (Controlling for Ethnic Fractionalization)

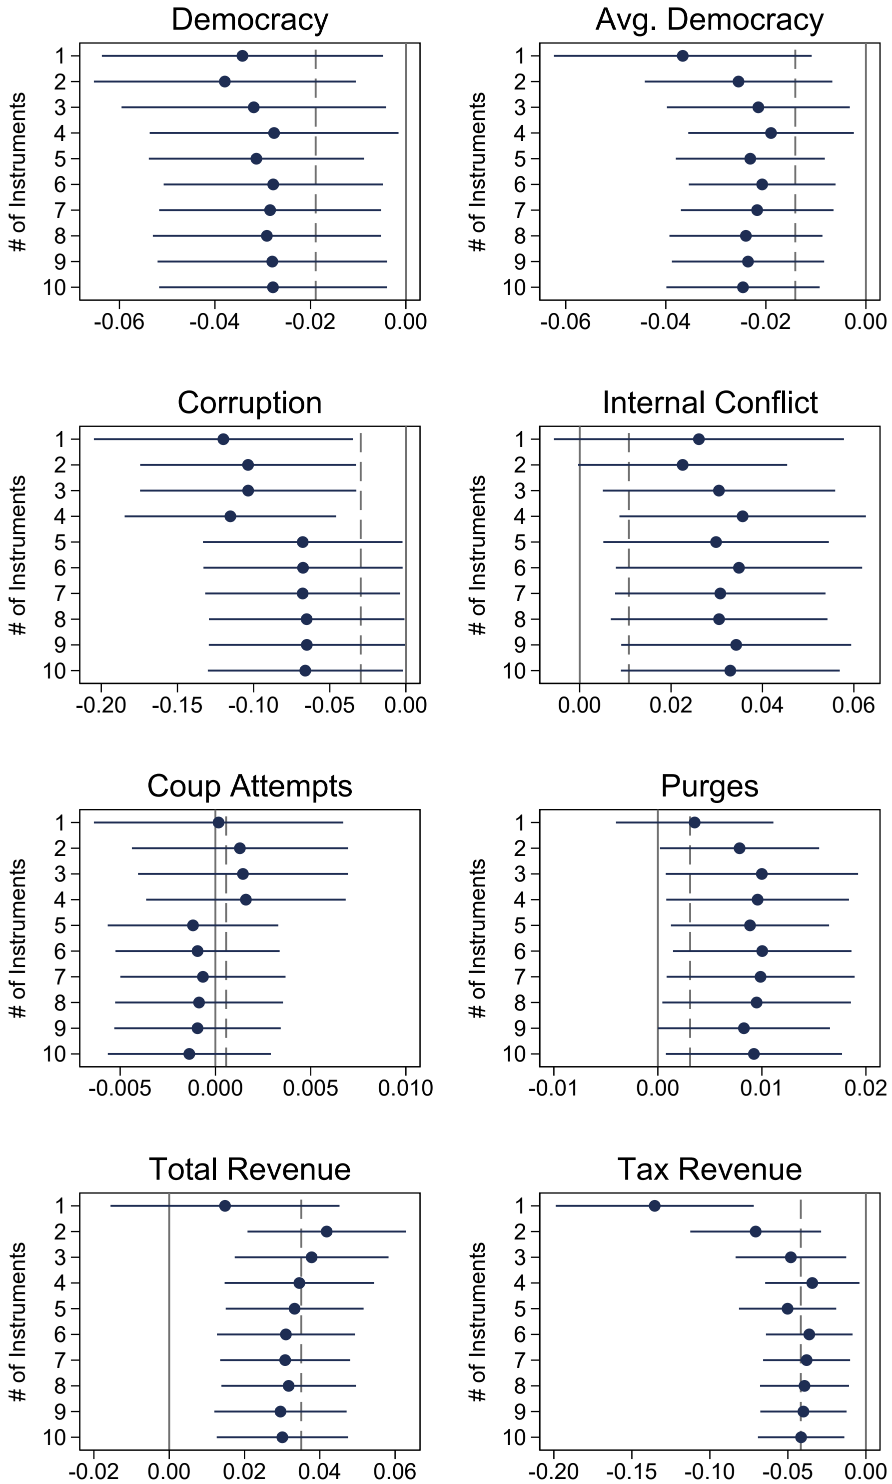

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.13: 2SLS Estimates by Size of Instrument Set, Final Component of Code Grouping (Controlling for Ethnic Fractionalization)

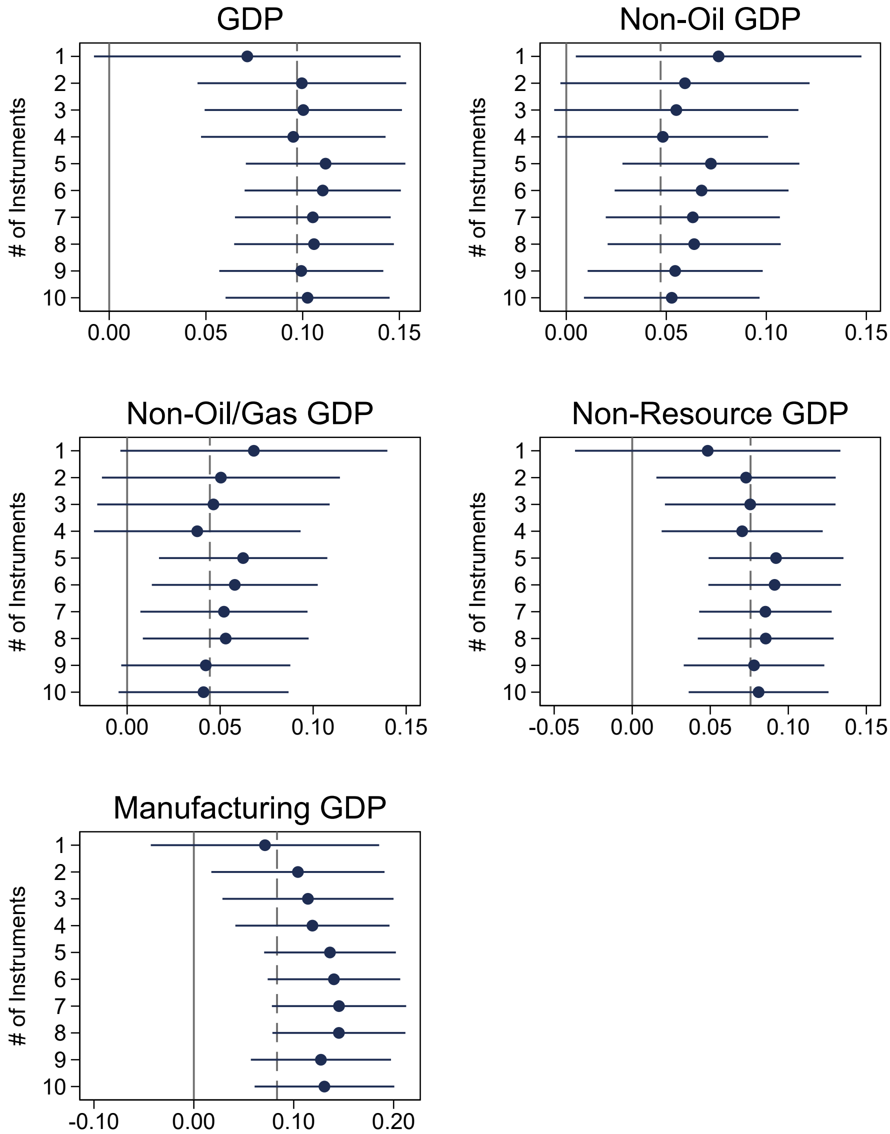

Notes. This figure plots point estimates and 90-percent confidence intervals for the coefficient on oil production, using optimal instrument sets of varying sizes. The gray, dashed line marks the value of the OLS estimate.

Figure 2.14: Placebo Test: Effect of *Basin* on Democracy over Time

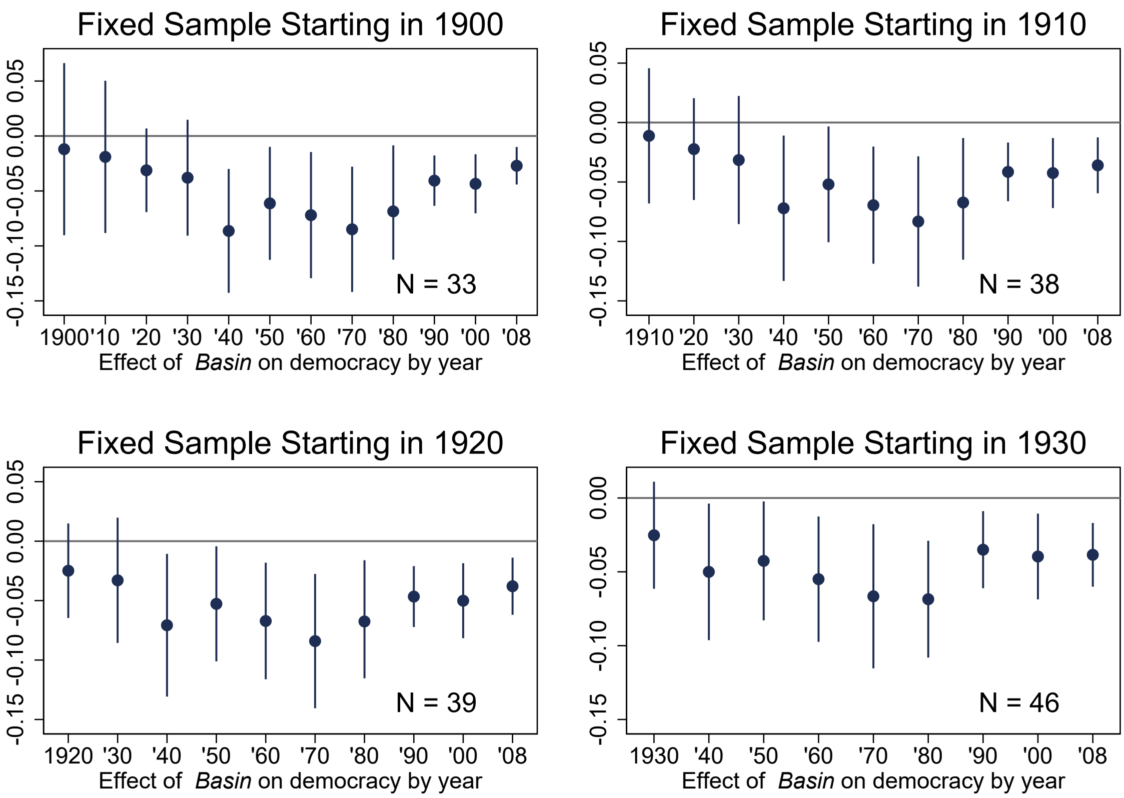

Notes. This figure plots point estimates and 95-percent confidence intervals for the reduced-form effect of the optimally chosen *Basin* variable on democracy over time. In each graph, the sample of countries is fixed.

Figure 2.15: Placebo Test: Effect of *Basin* on Log Population Density over Time

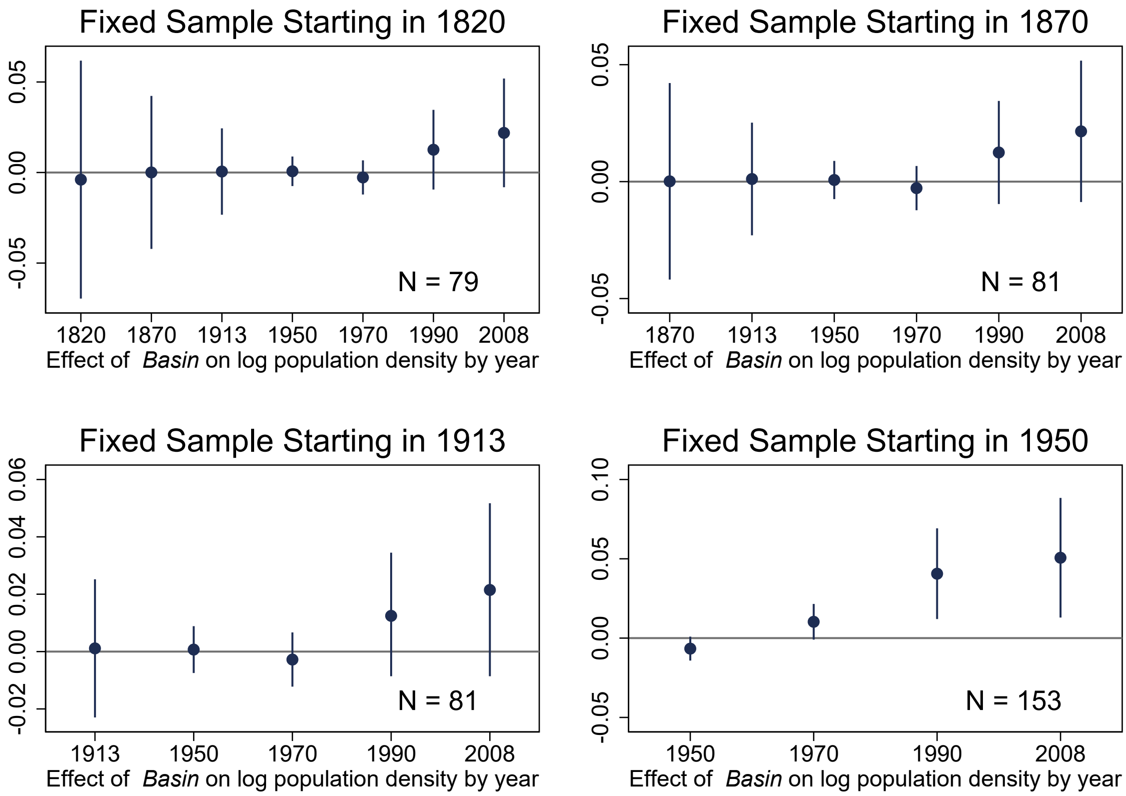

Notes. This figure plots point estimates and 95-percent confidence intervals for the reduced-form effect of the optimally chosen *Basin* variable on log population density over time. In each graph, the sample of countries is fixed.

Chapter 3

Old and Young Politicians

Coauthored with Alberto Alesina and Ugo Troiano

3.1 Introduction

When discussing whom to vote for, citizens commonly mention the age of the candidate, saying that he or she is too young (too inexperienced or too aggressive and eager) or too old (unmotivated, not energetic enough). In fact, concerns over a politician's age often extend beyond the casual worries of voters, and there are many examples of laws which limit the eligibility of candidates based on their age. Some of these laws date back at least two thousand years: the *lex Villia annalis*, established in Rome in 180 B.C., set minimum ages for senatorial magistrates.[1] The law was approved shortly after Publius Cornelius Scipio Africanus moved swiftly through the *cursus honorum* (the levels of the political career in ancient Rome) and became consul at the age of 30. The main rationale for the law was that the established elite believed that the policy choices made by young political leaders were driven too much by personal ambition ("career concerns," in modern terminology) and the desire to emulate Scipio's example (Kuiper, 2010).[2] Today many countries—including Australia, Austria, Brazil, Chile, Germany, Indonesia, Italy, Nigeria, Norway, the Philippines, the United States, and Venezuela—impose age restrictions for holding public office. The age threshold is often greater than the minimum age to vote and can be as high as 40 (Hong Kong, Italy, and Nigeria) or 45 (the Philippines). Another issue which might become relevant in the policy debate is whether politicians should have a retirement age. A prominent Indian congressman has advocated for a retirement age for politicians on the grounds that the politicians in office are often older than the voters that they represent (Pilot, 2004).

[1] Under the law, the *consulship* could be held from age 42, the *praetorship* from age 39, and the *curule adediliship* from age 36 (Kuiper, 2010).

[2] For instance, Cicero stated that: "For it is in old men that reason and good judgment are found, and had it not been for old men no state would have existed at all."

Policies which place age restrictions on politicians and the concern of voters must be based on the view that the age of a policymaker matters for public policy. However, to the best of our knowledge, politician's age on policy has been largely neglected by the literature. The exceptions are Altindag and Mocan (2015), who find, among members of the Turkish Parliament, a negative correlation between a politician's age and the probability of strategically switching parties when facing electoral uncertainty. While suggestive, the latter result does not imply that age causally affects public policies, because many omitted geographic variables are correlated with the age of the elected politicians. Additionally, Bertrand, Burgess, Chawla, and Xu (2015) study the effects of a bureaucrat's age in India. Younger entrants in the Administrative Service display stronger career concerns, even though, unlike politicians, their careers are decided by their superiors, not voters. Previous research has examined the effects of other politician characteristics on policy. For instance, gender (Chattopadhyay and Duflo, 2004; Gagliarducci and Paserman, 2012; Brollo and Troiano, 2016), religion (Meyersson, 2014), tenure (Coviello and Gagliarducci, 2013), salary (Gagliarducci and Nannicini, 2013), education (Besley, Montalvo, and Reynal-Querol, 2011), and race (Vogl, 2014) have all been shown to matter in a variety of ways for policies. In this paper we focus on age.

Younger politicians may differ from older ones for at least five reasons. One is that they have a potentially longer political career ahead of them and therefore have stronger career concerns. The second is simply that, as younger citizens, they have a longer horizon and therefore they may have an incentive to adopt more long-term policies.[3] The third and more mundane reason is that younger politicians may be more energetic and productive at work. The fourth reason is that there could be different self-selection patterns by age: because people of different ages have different opportunities in the labor market, this may affect the decision of becoming a politician. The fifth reason is that politicians of different ages may have different political connections, innate or accumulated doing the course of their previous work.

We use data on Italian municipalities which are in charge of a vast array of public goods, such as education, transportation, and waste management. On the revenue side, Italian mayors have the power to propose changes in property tax rates and municipal income tax rates.[4]

The age of a mayor may be correlated with many city and individual characteristics. For instance, cities with more favorable attitudes toward young people may be more likely to have young mayors and would likely have different policy preferences. In order to identify our effect, we use a fixed-effects model to control for time-invariant municipal characteristics. The effect of age is identified by using changes in the mayor's age across mayoral terms within a municipality. Furthermore, we account for time-varying municipal confounders by focusing on mayoral terms following close elections.

[3] See Bisin, Lizzeri, and Yariv (2015), Prato (2015), and Alesina and Passarelli (2015) for theoretical models on the effects of discount rates and time horizons on policies.

[4] Factors that have been shown to matter for urban public finance include decentralization (Besley and Coate, 2003), ethnic composition (Alesina, Baqir, and Hoxby, 2004), and geographical distribution (Ades and Glaeser, 1995).

We first find that younger mayors are more likely to be reelected and are also more likely to move to higher levels of elected government, both provincial and regional. We verify that the reelection effect is not merely explained by the fact that younger mayors are also more likely to run again. If we take reelection as a proxy of good government, these results may be consistent with the view that younger politicians implement better policies. Perhaps younger politicians implement policies of higher quality, because they are better selected or exert more effort, and this is why they are more often reelected. Measuring the quality of public policies is difficult, and we cannot decisively reject these hypotheses. However, capitalization models imply that better policies or better governors should be incorporated into house prices (Oates, 1969; Yinger, 1982). With this in mind, we test whether the age of the mayor affects house prices. We find no effect. A second possibility is that younger mayors may respond more quickly to the needs of their constituents. To test this hypothesis, we check whether the mayor's age affects the speed of public good provision.[5] We find that younger mayors are not faster in actually providing to voters the public goods that were budgeted at the beginning of the year, suggesting that this channel is unlikely to explain why younger mayors are more likely to get reelected.

In the second part of the paper, we then examine the mayor's policies implemented while in office. Younger and older mayors choose similar levels of expenditure and revenue on average during the term. However, the timing of expenditure differs by age: younger mayors are more likely to increase capital expenditure right before the upcoming election. The fact that budget cycles occur on capital expenditure is consistent with the result of Cioffi, Messina, and Tommasino (2012), who argue that capital expenditure is highly visible and easily targeted to specific groups of voters.[6] Alesina and Paradisi (2017) provide evidence of political budget cycles on real estate taxes in Italian cities. Other papers dealing with cycles in Italian municipalities include Bartolini and Santolini (2009) and Bonfatti and Forni (2017) Surveys of the literature on political cycles include Alesina, Roubini, and Cohen (1997) Drazen (2000) Cioffi et al. (2012), and Alesina and Passalacqua (2017).[7] This literature shows that political budget cycles occur in some circumstances and not in others. In this context we make the novel point that political budget cycles are more likely to occur when the career concerns of politicians is stronger.

Do younger mayors have a higher probability of reelection *because* they strategically increase spending right before the election, thus fooling voters? We do find a positive correlation between the cycle in capital expenditure and reelection. This correlation is nearly significant

[5] Casaburi and Troiano (2015) show that the political returns of implementing tax enforcement policies are higher for mayors who are faster in providing public goods.

[6] Repetto (2016) finds that a 2008 reform requiring all Italian municipalities to disclose their balance sheets before elections significantly reduced budget cycles in capital expenditure. Consistent with this result, we find that the effect of the mayor's age on capital expenditure cycles is strong prior to 2008 but non-existent after the reform. These results are available upon request. Other papers dealing with cycles in Italian municipalities include Bartolini and Santolini (2009) and Bonfatti and Forni (2017).

[7] The formalization of political business cycle models was pioneered by Nordhaus (1975), and subsequently by Rogoff and Sibert (1988) and Rogoff (1990).

at the ten-percent level when mayor controls and city and election-year effects are excluded, however the estimate becomes less precise when either mayor controls or city and election-year effects are added to the regression, even if the size of the coefficient remains virtually unchanged. Thus these results are only suggestive.

Song, Storesletten, and Zilibotti (2012) predict that the age of voters may matter for fiscal policy: young voters may have a disciplining effect on the implemented fiscal policy, because old voters do not internalize the future costs of a present loose fiscal policy. Thus, reading our results in light of theirs, young politicians would like to engage in short term fiscal policies to be reelected while young voters would want to discipline them, which is an interesting contrast. As a young citizen, a young politician to some extent would prefer long-term, non-strategic policies. As an ambitious politician, he might prefer the opposite. Our results show that for the average young politician the second effect dominates.

The paper is organized as follows. Section 3.2 describes our data and the institutional setting. Section 3.3 describes our methodology. Section 3.4 presents our results, and the final section concludes.

3.2 Data and Institutional Framework

3.2.1 Institutional Information

The Italian municipal government (*Comune*) is composed of a mayor (*Sindaco*), an executive committee (*Giunta*) appointed by the mayor, and an elected city council (*Consiglio Comunale*) responsible for authorizing the annual budget proposed by the mayor. The mayor and the executive committee propose policies, such as changes in the tax rates or expenditure. Subsequently, the city council votes on the proposals. Municipalities manage around 10 percent of total public expenditure in Italy and are in charge of many public services, such as preschools, waste management, municipal roads, and municipal public housing. Expenditures are divided into two types: capital expenditure, which relates to multi-year production factors, where amortization does take place, and current expenditure, which relates only to the current fiscal year.

A law in 1993 changed the mayoral electoral rule from party to individual ballot and introduced a two-term limit. In 2000 the duration of the mayoral term was extended from four to five years. Municipalities with more than 15,000 inhabitants adopt a runoff system to elect mayors, while a single-round system is in place in cities with a population below this threshold.[8] The number of city councilors depends on the size of the municipality.[9] We

[8]Bordignon, Nannicini, and Tabellini (2016) study the effect of the electoral rules on policies, finding that under runoff elections, the number of political candidates is larger, but the influence of extremist voters on equilibrium policy is smaller.

[9]The electoral rule for the city councilors also depends on the size of the municipality: in cities with fewer than 15,000 inhabitants, two thirds of the seats are assigned to councilors in the mayoral coalition, while the rest

calculate the margin of victory for municipalities with runoff as the margin of victory in the second election, while we use the margin of victory of the first (and only) election for the other municipalities.

3.2.2 Data and Descriptive Statistics

Our main database includes administrative data on municipal elections and politicians from 1998–2014, provided by the Italian Department of the Interior (*Ministero degli Interni*). The data contain information on every municipal election and every appointed administrator in the municipal, provincial, regional, and national administrations. For every election, we have data on the number of candidates, and the vote, the party affiliation, and demographic information on each candidate who is appointed to any position in the administration. We complement this dataset with socio-economic and demographic information on Italian municipalities from the National Statistical Office and administrative data on financial reports from the Italian Department of the Interior (*Ministero degli Interni*), covering the years 1998–2013. The financial-reports data contain yearly information on revenues and expenditures. We also obtained access to administrative data on house prices from the Italian Agency of the Territory (*Agenzia del Territorio*) for the years 2002–2011.

Table 3.1 presents summary statistics for the entire sample: mayor characteristics, municipality characteristics, political outcomes, public good and housing outcomes, and budget outcomes. We include in the sample all observations with non-missing data on the vote counts and ages of first- and second-place mayoral candidates. This creates a sample of around 23,000 mayoral terms, 16,000 of which have data on public finance outcomes.[10] In order to limit the potential impact of outliers, we winsorize the public finance variables at the 99-percent and one-percent levels. The results are very similar without winsorizing those variables.

As shown in Table 3.1, the older candidate wins the election in around 50 percent of the cases. On average Italian mayors are 49 years old and have 6.4 years of experience in elected municipal office and 0.8 years of experience in unelected municipal office at the time of election. Roughly 11 percent of mayors are women, and on average the mayor faces two rivals on the electoral ballot. Fifty-five percent of the municipalities we consider are in the North, 31 percent are in the South, and 15 percent are in the Center. Amongst all non-term-limited mayors, 45 percent ran for reelection, and 35 percent were reelected to a second term, where the latter number refer to all non-term-limited mayors, not just those who ran for reelection. Eight percent of mayors were elected to the provincial administration within 10 years of their first

of the seats are assigned proportionally to the vote shares. In cities with more than 15,000 inhabitants, if the winning mayor's coalition won at least 60 percent of the vote in the first round, every seat of the city council is assigned according to the proportional rule. If the winning mayor's coalition won less than 60 percent but more than 40 percent of the votes in the first round, they are granted 60 percent of the seats, and the remaining seats are assigned proportionally.

[10] To be precise, we have data on election outcomes and mayor characteristics for 22,789 out of 25,950 elections (88 percent) held during the sample period.

election to the mayor office, one percent of mayors were later elected to a position in the regional administration, and one percent of mayors were later elected to a position in the national administration.

The age of a mayor may be correlated with many city and individual characteristics. For instance, cities with more favorable attitudes toward young people may be more likely to have young mayors and would likely have different policy preferences. In order to identify our effect we use a fixed-effects model to control for time-invariant municipal characteristics. The effect of age is identified by using changes in the mayor's age across mayoral terms within a municipality.

3.3 Empirical Strategy

Simply comparing the outcomes of municipalities governed by mayors of different ages would not allow us to identify our effect of interest, because many other variables, not only observable but also unobservable, are correlated with the age of the mayor. The summary statistics provided in Table 3.2 show that many individual and municipality characteristics differ according to whether the older or younger mayoral candidate won the last election. For instance, older mayors have more elected political experience, less unelected political experience, and less education. Older mayors are also less likely to be a woman, more likely to have been born locally, and more likely to govern a lower-income municipality. Therefore, we pursue a fixed-effects strategy to alleviate the aforementioned concerns, and we explain why a regression discontinuity strategy is inappropriate in our setting.

3.3.1 Fixed-Effects Model

We begin our analysis by using this model:

$$Y_{mt} = \beta Age_{mt} + \boldsymbol{\delta}' \boldsymbol{Z}_{mt} + \eta_m + \gamma_t + \varepsilon_{mt}, \tag{3.1}$$

where Y is an outcome, Age is the age of the mayor as of the date of the most recent election, and $\boldsymbol{Z}$ is a vector of mayor characteristics. The letter m indexes municipalities, and t indexes election years. The parameters η_m and γ_t represent municipality fixed effects and election-year effects. The outcome Y_{mt} is measured over the term in office of the winner in municipality m and election year t.

The effect of age is identified using within-municipality variation in the mayor's age across mayoral terms. Our identifying assumption is that the mayor's age is exogenous to the time-varying municipality and mayor unobservables, ε_{mt}. This assumption will hold if $\boldsymbol{Z}_{mt}$ contains all outcome-relevant mayor characteristics correlated with age, and the municipality and election-year fixed effects absorb all unobserved municipal and temporal heterogeneity that is correlated with both outcomes and the age of the mayor. If the identification assumptions are satisfied, then β represents the expected change in Y from increasing the mayor's age by one year, holding the controls constant.

The main threat to our identifying assumption is that changes in voter preferences for public goods may be correlated with changes in the mayor's age. We address this possibility by also reporting results for the subsample of mayoral terms following elections determined by a vote margin of five percentage points or less. By focusing on outcomes following close elections, we isolate changes in the mayor's age due to idiosyncratic electoral outcomes rather than changes in voter preferences.

Equation (3.1) assumes a linear relationship between mayor age and outcomes. To test whether this is a good approximation, we also estimate the flexible equation

$$Y_{mt} = \sum_{k=1}^{8} \beta_k 1(Age_{mt} \in Bin_k) + \boldsymbol{\delta}' \boldsymbol{Z}_{mt} + \eta_m + \gamma_t + \varepsilon_{mt}. \tag{3.2}$$

Each age bin is five years long, with the exception of the (omitted) reference age bin (18–30 years) and the final bin (66+ years), both of which are aggregated to avoid very small bin sizes.

Another important parametric assumption of equation (3.1) is that the non-age characteristics enter the equation in a linear fashion. In the appendix we relax this assumption by using propensity-score-matching and inverse-probability-weighting estimators. To implement these estimators, we specify treatment assignment as an indicator variable equal to one if the older candidate won the election, and the propensity score as a logistic function of $\boldsymbol{Z}$. The estimators condition on individual covariates without requiring a linear relationship between $\boldsymbol{Z}$ and the outcomes. The results using these estimators are very similar to the baseline results.

3.3.2 Dynamic Panel Model

To examine how political budget cycles vary according to the age of the mayor, we follow (e.g., Brender and Drazen, 2005) and estimate two dynamic panel models,

$$Y_{m,t} = \sum_{k=1}^{K} \phi_k Y_{m,t-k} + \beta_1 Elec_{m,t} + \beta_2 Old_{m,t} + \beta_3 Elec_{m,t} \times Old_{m,t} + \alpha_m + \gamma_t + \varepsilon_{m,t}, \tag{3.3}$$

$$Y_{m,t} = \sum_{k=1}^{K} \rho_k Y_{m,t-k} + \delta_1 Elec_{m,t} + \delta_2 Age_{m,t} + \delta_3 Elec_{m,t} \times Age_{m,t} + \eta_m + \zeta_t + \nu_{m,t}, \tag{3.4}$$

where m indexes municipalities, and t indexes calendar years. The variable $Y_{m,t}$ represents a fiscal outcome, such as capital expenditure. The variable $Elec_{m,t}$ equals one when t is the year prior to an election in municipality m, and zero otherwise. The variable $Old_{m,t}$ equals one if the mayor of municipality m in year t was the older of the top two candidates in the most recent election, and zero otherwise. The variable $Age_{m,t}$ is the age of the current mayor as of the date of the most recent election.

The model includes K lagged values of Y to capture the persistence of fiscal outcomes. We report results for $K = 1$ to maximize sample size, however the estimates are very similar in magnitude and significance when $K = 2$ or $K = 3$. The model also includes municipality fixed effects and year fixed effects. We measure political budget cycles as the difference in fiscal outcomes in the year prior to an election compared to all other years of the mayor's term. The model allows both the level of spending and the size of the political budget cycle to vary according to the mayor's age. In equation (3.3), β_1 captures the political budget cycle under younger mayors, and $\beta_1 + \beta_3$ captures the political budget cycle under older mayors. Therefore, β_3 captures the difference in political budget cycles under older and younger mayors. In equation (3.4), $\delta_1 + \delta_3 Age$ captures the political budget cycle under a mayor of age *Age*. Therefore, δ_3 represents the marginal effect on the political budget cycle of increasing the mayor's age by one year.

There are two reasons why standard fixed-effects estimators (i.e., "within" estimators) of equations (3.3) and (3.4) would be asymptotically biased. The first is the well-known Nickell (1981) bias due to the presence of lagged dependent variables. This bias is likely to be substantial, given that we have a short panel. To address this issue, we estimate equations (3.3) and (3.4) using system GMM, which exploits moment conditions for the equation, in both first-differences and levels, identified by Arellano and Bond (1991), Arellano and Bover (1995), and Blundell and Bond (1998).[11] Because instrument proliferation can result from exploiting all moment conditions in system GMM, we follow the guidelines of Roodman (2009) and Bazzi and Clemens (2013) to avoid the problem of many weak instruments. Specifically, we "collapse" the instrument matrix and use only twice-lagged instruments for y and contemporaneous instruments for the other variables.

The second potential source of bias is the mayor's age. In the full sample of municipalities, both *Old* and *Age* could be endogenous in the above equations, even after accounting for municipality fixed effects. The reason is that changes in voter preferences for spending could be correlated with either changes in the electoral performance of older or younger candidates—rendering both *Old* and *Age* endogenous—or changes in the age profile of the pool of candidates—rendering *Age* endogenous. We address this possibility by limiting the sample to municipality-years following a close election, where a "close" election is decided by a vote margin of five percentage points or less.[12] We thus focus on time variation in *Old* and *Age* driven by arguably idiosyncratic election outcomes rather than shifts in voter preferences. We estimate equation (3.4) in two ways. The first treats *Age* and its interaction as strictly exogenous (conditional on the aforementioned sample selection), and the second treats "Age" and its interaction as endogenous and uses "Old" and its interaction as instruments.

[11] Results obtained by difference GMM, which relies on weaker assumptions by only exploiting the moment conditions of Arellano and Bond (1991), are very similar in terms of magnitude and statistical significance.

[12] The results are similar using much smaller or larger vote-margin thresholds, such as two or ten.

Because the incentive to manipulate fiscal outcomes for electoral advantage differs depending on whether the mayor's term limit is binding, we present results separately for mayoral terms when the term limit was binding and those when it was not. In Appendix we discuss why a regression discontinuity design would not be fully appropriate in our context.

Equation (3.4) assumes a linear relationship between the mayor's age and the size of the political budget cycle. To test whether this is a good approximation, we also estimate the flexible equation

$$\begin{aligned} Y_{m,t} = \rho Y_{m,t-1} + \alpha Elec_{m,t} + \sum_{k=1}^{8} \beta_k 1(Age_{m,t} \in Bin_k) \\ + \sum_{k=1}^{8} \delta_k Elec_{m,t} \times 1(Age_{m,t} \in Bin_k) + \eta_m + \zeta_t + \nu_{m,t}, \end{aligned} \tag{3.5}$$

which uses the same age bins as in (3.2).

3.4 Results

3.4.1 The Effect of Age on Reelection

First we estimate the effect of the mayor's age on reappointment. We limit the sample to elections in which the incumbent mayor does not face a binding term limit. In Table 3.3 we present estimates of the coefficient on age in the fixed-effects model, both with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city he or she governs. The results in Panel A are based on the full sample of mayoral terms, and the results in Panel B are based on mayoral terms following close elections. In the full sample of mayoral terms, a one-year increase in age reduces the mayor's probability of running again for mayor by 0.65 percentage points and reduces the probability of reelection to another mayoral term by 0.71 percentage points. This means that a one-standard-deviation increase in the mayor's age (10 years) reduces the probability of running again by 14 percent and reduces the probability of reelection by 20 percent relative to their respective means. We also find that a one-year increase in age reduces the probability of being appointed in the provincial administration (within five years after the end of the political term) by about 0.17 percentage points and reduces the probability of being appointed in the regional administration by about 0.02 percentage points. This means that a one-standard-deviation increase in the mayor's age reduces the probability of joining the provincial and regional administrations by 21 percent and 20 percent relative to their respective means. The effect of the mayor's age on reaching the national administration is small and statistically insignificant. In the sample of mayoral terms following close elections, the effect of age on running again and reelection is larger in absolute magnitude than the baseline estimates, and the effect of age on reaching higher levels of government is similar to the baseline estimates. Overall the estimates seem not to be sensitive to the inclusion of controls for other mayor characteristics. Note that for the latter four outcomes we do not condition on the mayor's decision to run again, which is endogenous. Results reported in the appendix show that for second-term (i.e., term-limited) mayors, the negative effect of age on reaching higher political office is more statistically significant and larger in absolute magnitude compared to the baseline results using first-term mayors

Is the lower reelection rate for older mayors entirely explained by the fact that older mayors are less likely to run for reelection? We use the dichotomous age measure, *Old*, which indicates whether the older candidate won the election. Table C.1.14 in the appendix show that the probability of running for reelection is 9.1 percentage points lower for the older candidate, and the probability of being reelected is 13.6 percentage points lower for the older candidate. The results imply that a share of the reelection effect can be explained by the (endogenous) decision to run again. Since the choice to run for reelection is endogenous and depends on the expected probability of winning, the causal effect of age on reelection conditional on running again is not easily recovered. If all mayors faced the same probability of reelection conditional on running again,which at the sample mean is 78 percent, (the reelection rate (0.35) divided by the rate of running again (0.45), then the effect of age on running again would result in older mayors having a 7-percentage-point lower reelection probability compared to younger mayors. In fact, older mayors have a reelection probability that is 13.6 percentage points lower than that of younger mayors. Thus the effect of age on running again by itself explains roughly half of the estimated reelection effect. The third panel of Table C.1.14 shows that the difference of the reelection effect and the effect on running again is statistically significant at the five-percent level.

The baseline regression specifies a linear relationship between mayor age and outcomes. Is this a good approximation? Figure 3.1 plots point estimates and 95-percent confidence intervals for the coefficients $\{\beta_k\}$ in (3.2), where β_ℓ represents the difference in average outcomes for mayors in bin ℓ and mayors in the 18–30 year bin, accounting for individual controls, municipality fixed effects, and election-year effects. The effect of age on the probability of running again, being reelected, and moving to the provincial administration appears quite linear. The results for moving to the regional administration appear less linear, though the estimates are noisy due to the fact that this is a very rare event.

In Table 3.4 we report the effect of age on moving to higher office for term-limited mayors. We only report results for the full sample of second-term mayors, as there is not sufficient variation in the outcomes in the subsample of second terms following close elections. The full sample results show that age continues to have a negative effect on reaching higher office, though the results are more statistically significant and larger in absolute magnitude than the baseline results using first-term mayors.

3.4.2 The Effect of Age on Average Revenue and Expenditure

Now we explore why younger candidates tend to be elected by exploring what they do when they are in office. We first examine the effect of the mayor's age on average revenue and expenditure over the term, both measured in euro per capita. The first panel of the Table 3.5 shows estimates of the coefficient on age in the fixed-effects model, both with and without controls for mayor characteristics, using the full sample of mayoral terms. The second panel of Table 3.5 shows the corresponding estimates based on the subsample of terms following close elections. Of the 12 point estimates presented in the table, none is significant at the 10-percent level. The magnitudes of the estimates are economically small. There is no evidence that the age of the mayor matters for revenue or expenditure on average. Figure 3.2 presents the estimates of equation (3.2) for the average public finance outcomes. According to the first two graphs, mayors in the age range of 18–30 have higher average revenue and average capital expenditure than mayors in the other age bins. Beyond the 30-year mark, however, the effect of age on these two outcomes is roughly zero—hence linear—across the age distribution. Thus for average revenue and average capital expenditure, the effect of age is nonlinear only in the age range 18–35. For these outcomes, we view the linear specification as an approximation. The third graph shows that the effect of age on average current expenditure is roughly zero across the entire age distribution. The results are very similar when we remove municipality-specific linear trends from the revenue and expenditure outcomes. See the appendix for details.

3.4.3 The Effect of Age on Measures of Quality of Governance

A career-concerns model, in which politicians care about both citizen welfare and reelection and have heterogeneous values of reelection, would predict that politicians that value reelection more. In this section we test whether the quality of their policies is superior because they invest more effort.[13] Younger mayors may improve their chances at reelection by means,of a variety of policies, and their quality which we may not observe in our dataset.

If the capitalization model is correct (Oates, 1969; Yinger, 1982) one would expect that better policies translate into higher house prices. As shown in Table 3.6, there is little evidence that house prices respond to the age of the mayor in office. While the mayor's age has a statistically significant, negative effect on house prices in the specification without controls using the subsample of terms following a close election, adding controls results in a dramatically reduced point estimate which becomes statistically insignificant. The results are very similar when we remove municipality-specific linear trends from the outcome. (See the appendix for details.)

A second interpretation is that while younger and older mayors implement similar policies, both quantitatively and qualitatively, the younger ones are faster in responding to the citizens' needs. For instance, Casaburi and Troiano (2015) find that the returns to otherwise identical

[13] See Persson and Tabellini (2000) for a discussion of the career-concerns model.

tax enforcement policies are stronger for mayors who are faster in translating these policies to actual public goods. To investigate this possibility, we adopt the same measure of speed of public good provision adopted in Casaburi and Troiano (2015) and test whether this measure responds to the age of the mayor. The measure is the ratio of actually provided public goods over the public goods that were promised to voters in the provisional budget at the end of the year. The estimates in Table 3.6 indicate that the speed of public good provision does not depend on the age of the mayor. The results are very similar when we remove municipality-specific linear trends from the outcome. (See the appendix for details.)

Figure 3.3 displays the estimates of (3.2) for the speed of public good provision and house prices. The results confirm that these outcomes are typically similar for older and younger mayors.

3.4.4 Age and Political Budget Cycles

Political budget cycles lower the welfare of citizens by distorting the path of fiscal policy. However, for younger politicians this social cost could be outweighed by the benefit of increasing the reelection probability, because younger politicians expect to have a long political career. If voters are gullible or rational but poorly informed about the negative consequences of budget cycles, this would in turn imply that younger voters have a higher reelection probability. We therefore test whether the mayor's age affects the timing of revenue and expenditure relative to the electoral cycle.

We do this by estimating equations (3.3) and (3.4) using system GMM. For the sake of brevity, we report estimates only for specifications that exclude controls for mayor characteristics. The results, available upon request, are very similar when we add these controls.

Panel A in Table 3.7 presents the results for the full sample of mayoral terms, while Panel B presents the results for terms following close elections. Columns 1, 4, and 7 present estimates of equation (3.3). Columns 2, 5, and 8 present estimates of equation (3.4) treating *Age* and its interaction as strictly exogenous, while in columns 3, 6, and 9, *Age* and its interaction are treated as endogenous and instrumented with *Old* and its interaction. Column 1 of Panel A shows that younger mayors increase revenue by 12.54 euro per capita in the year before an election relative to other years of the mayoral term. This effect is significant at the five percent level. In contrast, older mayors increase revenue by only 5.41 euro per capita in the year before an election, however the difference in revenue cycles for older and younger mayors is statistically insignificant. Similarly, the results in columns 2 and 3, which model a linear relationship between the mayor's age and the size of the budget cycle, show that revenue cycles are smaller for older mayors, but the marginal effect of age is statistically insignificant. The results in Panel B based on terms following close elections confirm the conclusion that the mayor's age does not have a statistically significant effect on cycles in total revenue.

The specification in column 4 of Table 3.7 tests whether political budget cycles in capital expenditure are different for older and younger mayors. In the sample of all mayoral terms (Panel A), younger mayors increase capital expenditure by 7.25 euro per capita in the year before the election, and in the sample of terms following close elections (Panel B), younger mayors increase capital expenditure by 9.48 euro per capita. These two estimates are significant at the one-percent and five-percent levels, respectively. The corresponding estimates of capital expenditure cycles for older mayors, on the other hand, are very close to zero—28 cents per capita and −55 cents per capita. In both samples, the difference in cycles for older and younger mayors is statistically significant at the five-percent level. Columns 5 and 6 test for an effect of the mayor's age on capital expenditure, using a continuous measure of age. Considering the results in both Panels A and B, the point estimates suggest that political cycles in capital expenditure are smaller for older mayors, and the marginal effect of age is statistically significant at the five-percent level in three out of four specifications. The results suggest that increasing the mayor's age by one year reduces the size of the cycle in capital expenditure by 0.26 to 0.98 euro per capita. To compare the results to the specification using a dichotomous measure of age, consider column 6 of Panel B. The average age of a younger mayor who wins the election is 43 years, while the average age of an older mayor who wins the election is 55 years. Therefore the results suggest that a typical younger mayor will increase capital expenditure before the election by 9.66 euro per capita, while a typical older mayor will actually decrease capital expenditure before the election by 2.1 euro per capita. Therefore the specifications using the dichotomous and continuous measures of age produce broadly similar results.

The baseline size and heterogeneity of political cycles in capital expenditure are both larger in the sample of mayoral terms following close elections. Younger mayors with career concerns are more likely to engage in political budget cycles in more politically competitive municipalities where the stakes are higher. However, since political competition is not randomly distributed across municipalities, and could be correlated with other municipality characteristics.

The results in columns 7 through 9 of Table 3.7 indicate that there is little evidence that younger mayors engage in political cycles in current expenditure or that cycles in current expenditure vary by the mayor's age. This is consistent with Cioffi et al. (2012) who find that budget cycles in Italian municipalities are driven by changes in capital expenditure. This is because capital expenditure is highly visible, easily targeted to specific groups of voters, and largely exempted from balanced-budget rules. All of the political-budget-cycle results are very similar when we remove municipality-specific linear trends from the revenue and expenditure outcomes. (See the appendix for details.)

Is the relationship between the mayor's age and political budget cycles nonlinear? Figure 3.4 plots point estimates and 95-percent confidence intervals for the coefficients $\{\delta_k\}$ in (3.5), where δ_ℓ represents the difference in average political budget cycles for mayors in bin ℓ and mayors in the 18–30 year bin, accounting for municipality fixed effects, year effects, and the persistence of public finance outcomes. The first and third graphs show that cycles in total revenue and current expenditure are similar across different age bins. For these two outcomes, the size of the cycle for mayors aged 18–30 is statistically indistinguishable from the size of the cycle for mayors in any of the older age bins. In contrast, the second graph shows that cycles in capital expenditure are smaller for mayors aged 46–50, 61–65, or over 66 compared to mayors aged 18–30, and this difference is statistically significant. Focusing on the point estimates, we document four patterns. First, increasing mayor age from the 18–30 bin to the 31–35 bin has virtually no effect on capital expenditure cycles. Second, the marginal effect of mayor age on capital expenditure cycles is negative moving from the 31–35 bin to the 36–40 bin. Third, the marginal effect of age is roughly zero over the 36–60 age range. Finally, the marginal effect of age again becomes negative for ages 61 and greater. Though one should avoid over-interpreting these patterns in the face of substantial statistical uncertainty, they provide suggestive evidence of nonlinearities in the relationship between mayor age and cycles in capital expenditure. The linear specification should therefore be interpreted as an approximation.

In regressions reported in the appendix, we examine whether the effect of the mayor's age on political budget cycles varies according to the mayor's wage, the presence of a balanced-budget rule, or the electoral system. We exploit the fact that, for Italian municipalities, these three variables are each a step function of population, changing discontinuously at different population thresholds.[14] We find that the effect of age on political budget cycles is smaller when the mayor's wage is higher, the municipality is subject to a balanced-budget rule, or the municipality uses runoff elections as opposed to single-round plurality elections.[15]

3.4.5 Political Budget Cycles and Reelection

Two of our main findings are that younger mayors increase investment spending right before an election by a greater amount than older mayors, and younger mayors are more likely to be reelected than older mayors. Do younger mayors have a higher probability of reelection *because* they strategically increase spending right before the election, thus fooling voters? This is a difficult question to answer, because spending decisions are endogenous and may be correlated with, among many other things, expectations of reelection, which are unobservable. For instance, mayors who feel certain about their reelection may be less prone to political budget cycles, and mayors who feel electorally vulnerable may be more prone to cycles.

As shown in the last table, in the sample of terms following close elections, there is a positive correlation between the cycle in capital expenditure and reelection. This correlation

[14]The mayor's wage increases at the following population thresholds: 1,000, 3,000, 5,000, 10,000, 30,000, and several higher thresholds that are relevant to only a tiny fraction of municipalities. (See Gagliarducci and Nannicini, 2013, for details.) A balanced-budget rule was introduced in 1999 and relaxed in 2001 only for municipalities with populations below 5,000. (See Grembi, Nannicini, and Troiano, 2016, for details.) Municipalities with populations below 15,000 elect the mayor according to a single-round plurality system, while municipalities with populations of 15,000 or greater elect the mayor according to a runoff system. (See Bordignon et al., 2016, for details.)

[15]The evidence is only suggestive because the size of the municipalities could influence the effect of age on political budget cycles through many channels besides the wage, balanced-budget rule, or electoral system.

is statistically insignificant (but close to the ten-percent level) when mayor controls and city and election-year effects are excluded, however it becomes less precise when either mayor controls or city and election-year effects are added to the regression. It is however re-assuring that the magnitude of the coefficient, in close elections, does not change with the inclusion of the controls. While suggestive, these estimates may be biased and should not be interpreted as causal. Mayors that strategically increase spending before the election may exert more effort towards reelection in other dimensions, for instance in political speeches, media appearances, and so on. Furthermore, spending decisions may be correlated with expectations of reelection, which are unobservable.

In a previous version of the paper we have shown that the cycle in residual expenditure (the difference between the promised and the delivered expenditures) is correlated with reelection in a statistically significant way (the result is available upon request). This may be consistent with the view that voters may pay more attention to the promises of politicians rather than to the implemented policies.

3.5 Conclusion

We study the role of a politician's age in determining their policy choices. We find that younger politicians engage in political budget cycles on public investment more than older ones. This result is consistent with a career-concerns model in which younger politicians with a longer horizon in their career are more willing to engage in strategic policies to ensure reelection. Interestingly, this is more likely to occur when elections are close and therefore even a small effect on voting behavior could be critical. Younger politicians are more likely to be reelected, and there is suggestive evidence of a positive correlation between reelection and political budget cycles, albeit statistically insignificant.

3.6 Tables

Table 3.1: Summary Statistics

	Mean	Std. Dev.	Min.	Max.	Obs.
Mayor					
Older candidate won	0.49	0.50	0.00	1.00	22,789
Older candidate's margin of victory	-0.28	28.01	-98.66	98.08	22,789
Age	49.12	9.93	18.64	86.31	22,789
Previously held any city office	0.80	0.40	0.00	1.00	22,789
City political experience (elected)	6.40	5.24	0.00	24.00	22,789
City political experience (unelected)	0.78	1.99	0.00	16.00	22,789
Woman	0.11	0.31	0.00	1.00	22,789
Born locally	0.47	0.50	0.00	1.00	22,789
High school degree	0.89	0.31	0.00	1.00	22,789
College degree	0.45	0.49	0.00	1.00	22,789
Center-right party	0.08	0.26	0.00	1.00	22,789
Number of rivals on ballot	1.87	1.25	1.00	18.00	22,789
Term limit binding	0.35	0.48	0.00	1.00	22,789
Municipality					
Population (2001)	7,481.38	43,095.43	33.00	2,546,804.00	22,789
Population per square km (2005)	309.41	656.51	1.00	12,624.00	22,789
Active pop. / total pop. (2005)	0.41	0.06	0.16	0.60	22,789
Elderly index (2005)	1.89	1.53	0.00	35.00	22,789
Family size (2005)	2.46	0.31	1.20	4.20	22,789
Production units per capita (2005)	0.08	0.03	0.02	0.34	22,789
Employed / total pop. (2005)	0.26	0.18	0.02	3.03	22,789
Income per capita (2005 €)	13,506.76	3,056.29	5,013.00	44,949.00	22,789
Altitude (meters)	488.15	434.27	0.00	2,851.00	22,789
North	0.55	0.50	0.00	1.00	22,789
Central	0.15	0.35	0.00	1.00	22,789
South	0.31	0.46	0.00	1.00	22,789
Political Outcomes					
Ran for reelection (not term-limited)	0.45	0.50	0.00	1.00	14,897
Reelected (not term-limited)	0.35	0.48	0.00	1.00	14,897
Provincial administration after term	0.08	0.28	0.00	1.00	22,789
Regional administration after term	0.01	0.12	0.00	1.00	22,789
National administration after term	0.01	0.09	0.00	1.00	22,789
Public Good and Housing Outcomes					
Speed of public good provision	0.77	0.07	0.56	0.90	15,937
Avg. house price	12.66	44.64	-370.71	669.00	9,000
Budget Outcomes					
Revenue per capita	719.72	503.87	180.11	3,449.92	127,171
Capital expenditure per capita	78.03	152.93	0.00	1,046.25	127,525
Current expenditure per capita	539.04	269.68	246.90	1,914.84	127,521

Notes. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. The outcomes "Ran for reelection" and "Reelection" are set to missing for mayoral terms in which the incumbent faces a binding term limit. Budget outcomes are yearly, while the other outcomes are measured on the basis of mayoral terms.

Table 3.2: Summary Statistics by Winning Candidate

	Mean by Winning Candidate				
	Older	Younger	Difference	p-value	Obs.
Mayor					
Age	54.31	44.17	10.14***	0.000	22,789
Previously held any city office	0.80	0.79	0.02***	0.004	22,789
City political experience (elected)	6.84	5.98	0.86***	0.000	22,789
City political experience (unelected)	0.73	0.82	−0.09***	0.001	22,789
Woman	0.09	0.12	−0.03***	0.000	22,789
Born locally	0.54	0.40	0.14***	0.000	22,789
High school degree	0.86	0.92	−0.06***	0.000	22,789
College degree	0.41	0.48	−0.08***	0.000	22,789
Center-right party	0.08	0.08	−0.00	0.451	22,789
Number of rivals on ballot	1.86	1.89	−0.03*	0.051	22,789
Term limit binding	0.37	0.32	0.05***	0.000	22,789
Municipality					
Population (2001)	7011.87	7930.13	−918.26	0.108	22,789
Population per square km (2005)	307.46	311.28	−3.82	0.661	22,789
Active pop. / total pop. (2005)	0.41	0.41	−0.00	0.844	22,789
Elderly index (2005)	1.90	1.89	0.00	0.915	22,789
Family size (2005)	2.46	2.46	0.00	0.970	22,789
Production units per capita (2005)	0.08	0.08	−0.00***	0.001	22,789
Employed / total pop. (2005)	0.26	0.27	−0.00**	0.034	22,789
Income per capita (2005 €)	13447.44	13563.45	−116.01***	0.004	22,789
Altitude (meters)	486.71	489.52	−2.80	0.626	22,789
North	0.55	0.54	0.01	0.340	22,789
Central	0.14	0.16	−0.02***	0.001	22,789
South	0.31	0.30	0.01	0.150	22,789
Political Outcomes					
Ran for reelection (not term-limited)	0.41	0.48	−0.07***	0.000	14,897
Reelected (not term-limited)	0.30	0.39	−0.09***	0.000	14,897
Provincial administration after term	0.07	0.10	−0.03***	0.000	22,789
Regional administration after term	0.01	0.02	−0.00***	0.003	22,789
National administration after term	0.01	0.01	−0.00**	0.010	22,789
Public Good and Housing Outcomes					
Speed of public good provision	0.77	0.77	0.00	0.143	15,937
Avg. house price	11.30	13.86	−2.56**	0.010	8,170
Budget Outcomes					
Revenue per capita	675.82	683.06	−7.24***	0.008	110,183
Capital expenditure per capita	65.36	65.74	−0.38	0.634	110,417
Current expenditure per capita	516.04	518.60	−2.57*	0.086	110,413

Notes: See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. The outcomes "Ran for reelection" and "Reelection" are set to missing for mayoral terms in which the incumbent faces a binding term limit. Budget outcomes are yearly, while the other outcomes are measured on the basis of mayoral terms. The first and second columns report averages in years following the victory of the older and younger mayoral candidate, respectively. The third column reports the difference of the averages, the fourth column reports the p-value corresponding to the two-sided test of equality of the averages, and the fifth column reports the sample size. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.3: The Effect of the Mayor's Age on Political Outcomes

	Ran again		Reelected		Provincial admin.		Regional admin.		National admin.	
Panel A: Fixed-Effects Estimates										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Age	-0.0057^{***}	-0.0065^{***}	-0.0061^{***}	-0.0072^{***}	-0.0016^{***}	-0.0018^{***}	-0.0003^{**}	-0.0002^{*}	-0.0001	-0.0001
	(0.0005)	(0.0005)	(0.0005)	(0.0005)	(0.0003)	(0.0003)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	6813	6813	6813	6813	6813	6813	6813	6813	6813	6813
Observations	14897	14897	14897	14897	14897	14897	14897	14897	14897	14897
Panel B: Fixed-Effects Estimates, Close Elections										
	Ran again		Reelected		Provincial admin.		Regional admin.		National admin.	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Age	-0.0073^{***}	-0.0091^{***}	-0.0093^{***}	-0.0101^{***}	-0.0016^{*}	-0.0014	-0.0001	-0.0002	0.0001	0.0000
	(0.0018)	(0.0019)	(0.0016)	(0.0018)	(0.0009)	(0.0011)	(0.0002)	(0.0003)	(0.0004)	(0.0003)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	2550	2550	2550	2550	2550	2550	2550	2550	2550	2550
Observations	3108	3108	3108	3108	3108	3108	3108	3108	3108	3108

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. All regressions include municipality fixed effects and year-of-election effects. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors in the fixed-effects specification are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.4: The Effect of the Mayor's Age on Political Outcomes (Binding Term Limit)

Panel A: Fixed-Effects Estimates

	Provincial admin.		Regional admin.		National admin.	
	(1)	(2)	(3)	(4)	(5)	(6)
Age	-0.0044^{***}	-0.0046^{***}	-0.0009^{***}	-0.0011^{***}	-0.0005^{**}	-0.0002
	(0.0007)	(0.0007)	(0.0003)	(0.0004)	(0.0003)	(0.0003)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	5533	5533	5533	5533	5533	5533
Observations	7892	7892	7892	7892	7892	7892

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is binding. All regressions include municipality fixed effects and year-of-election effects. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors in the fixed-effects specification are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.5: The Effect of the Mayor's Age on Average Budget Outcomes

Panel A: Fixed-Effects Estimates

	Avg. revenue		Avg. capital expend.		Avg. current expend.	
	(1)	(2)	(3)	(4)	(5)	(6)
Age	0.10	0.23	−0.08	−0.03	0.16	0.22
	(0.38)	(0.41)	(0.13)	(0.13)	(0.16)	(0.18)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	6584	6584	6584	6584	6584	6584
Observations	10244	10244	10244	10244	10244	10244

Panel B: Fixed-Effects Estimates, Close Elections

	Avg. revenue		Avg. capital expend.		Avg. current expend.	
	(1)	(2)	(3)	(4)	(5)	(6)
Age	−0.65	−0.94	−0.22	−0.20	−0.36	−0.35
	(0.92)	(1.05)	(0.35)	(0.42)	(0.37)	(0.44)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	1902	1902	1902	1902	1902	1902
Observations	2122	2122	2122	2122	2122	2122

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. All regressions include municipality fixed effects and year-of-election effects. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors in the fixed-effects specification are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.6: The Effect of the Mayor's Age on House Prices and the Speed of Public Good Provision

Panel A: Fixed-Effects Estimates

	Avg. house price		Speed of public good provision	
	(1)	(2)	(3)	(4)
Age	−0.48	−0.57	0.000057	0.000074
	(0.57)	(0.61)	(0.000084)	(0.000089)
Controls	*No*	*Yes*	*No*	*Yes*
Municipalities	5234	5234	6584	6584
Observations	5996	5996	10244	10244

Panel B: Fixed-Effects Estimates, Close Elections

	Avg. house price		Speed of public good provision	
	(1)	(2)	(3)	(4)
Age	-2.73^{**}	−0.11	0.00030	0.00036
	(1.37)	(1.40)	(0.00024)	(0.00026)
Controls	*No*	*Yes*	*No*	*Yes*
Municipalities	1173	1173	1902	1902
Observations	1228	1228	2122	2122

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.7: The Effect of the Mayor's Age on Political Budget Cycles

Panel A: Dynamic Panel Estimates

	Revenue per capita			Capital expenditure per capita			Current expenditure per capita		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Year before election	13.71**	22.43	54.76	7.73***	16.16**	38.75***	−2.81	−10.42*	−15.99
	(5.43)	(18.68)	(36.13)	(2.13)	(7.80)	(14.97)	(1.89)	(6.18)	(13.36)
Old	0.34			0.89			0.37		
	(2.69)			(1.04)			(1.00)		
Old × Year before election	−7.31			−6.60**			2.34		
	(7.31)			(2.97)			(2.72)		
Age		−0.09	0.05		−0.05	0.10		−0.01	0.03
		(0.14)	(0.28)		(0.05)	(0.11)		(0.05)	(0.10)
Age × Year before election		−0.26	−0.93		−0.24	−0.71**		0.19	0.30
		(0.38)	(0.75)		(0.16)	(0.31)		(0.13)	(0.28)
AR(2) test p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Municipalities	7064	7229	7064	7064	7229	7064	7064	7229	7064
Observations	60721	63018	60721	60838	63149	60838	60837	63148	60837

Panel B: Dynamic Panel Estimates, Close Elections

	Revenue per capita			Capital expenditure per capita			Current expenditure per capita		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Year before election	6.60	27.24	4.43	9.85**	34.22**	50.80**	−0.86	−8.95	9.05
	(10.90)	(32.86)	(62.03)	(3.83)	(14.47)	(23.07)	(3.56)	(11.02)	(23.26)
Old	−2.72			−0.26			3.78		
	(6.78)			(2.74)			(2.64)		
Old × Year before election	0.20			−9.73**			−2.08		
	(12.73)			(4.83)			(4.84)		
Age		0.01	−0.32		−0.04	−0.01		0.16	0.38
		(0.34)	(0.69)		(0.13)	(0.28)		(0.12)	(0.27)
Age × Year before election		−0.42	0.04		−0.60**	−0.94**		0.15	−0.22
		(0.63)	(1.24)		(0.28)	(0.47)		(0.21)	(0.48)
AR(2) test p-value	0.660	0.661	0.660	0.890	0.901	0.905	0.514	0.518	0.516
Municipalities	2470	2470	2470	2472	2472	2472	2472	2472	2472
Observations	11464	11464	11464	11481	11481	11481	11481	11481	11481

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. Both panels include observations following all mayoral elections. The results in Panel A are from the sample for which the mayor is not term limited. The results in Panel B are from the sample for which the mayor is not term limited and the most recent election was decided by a vote margin of five percentage points or less. In columns 1, 4, and 7, "Old" is an indicator for the mayor being older than the electoral opponent coming in second place. "Old" is treated as strictly exogenous. In the remaining columns, "Age" is the age of the mayor. In columns 2, 5, and 8, "Age" is treated as strictly exogenous. In columns 3, 6, and 9, "Age" and its interaction are treated as endogenous and instrumented with "Old" and its interaction. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.8: The Correlation between Political Budget Cycles and Reelection

Panel A: Cycles in Capital Expenditure

	Reelected			
	(1)	(2)	(3)	(4)
Δ Capital expend.	0.00039	−0.000090	−0.0076	−0.010
	(0.0051)	(0.0051)	(0.0085)	(0.0084)
Mayor controls	*No*	*Yes*	*No*	*Yes*
City and election-year effects	*No*	*No*	*Yes*	*Yes*
R^2	0.000	0.022	0.074	0.092
Municipalities	6581	6581	6581	6581
Observations	10239	10239	10239	10239

Panel B: Cycles in Capital Expenditure, Close Elections

	Reelected			
	(1)	(2)	(3)	(4)
Δ Capital expend.	0.019	0.016	0.017	0.019
	(0.012)	(0.012)	(0.035)	(0.035)
Mayor controls	*No*	*Yes*	*No*	*Yes*
City and election-year effects	*No*	*No*	*Yes*	*Yes*
R^2	0.001	0.021	0.186	0.199
Municipalities	1902	1902	1902	1902
Observations	2123	2123	2123	2123

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. The equations in the first two columns are estimated by OLS, and the equations in the last two columns are estimated using municipality fixed effects and year-of-election effects. For the sake of readability, Δ Capital expend. is measured in hundreds of euro per capita. Standard errors are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.7 Figures

Figure 3.1: Mayor Age and Political Outcomes

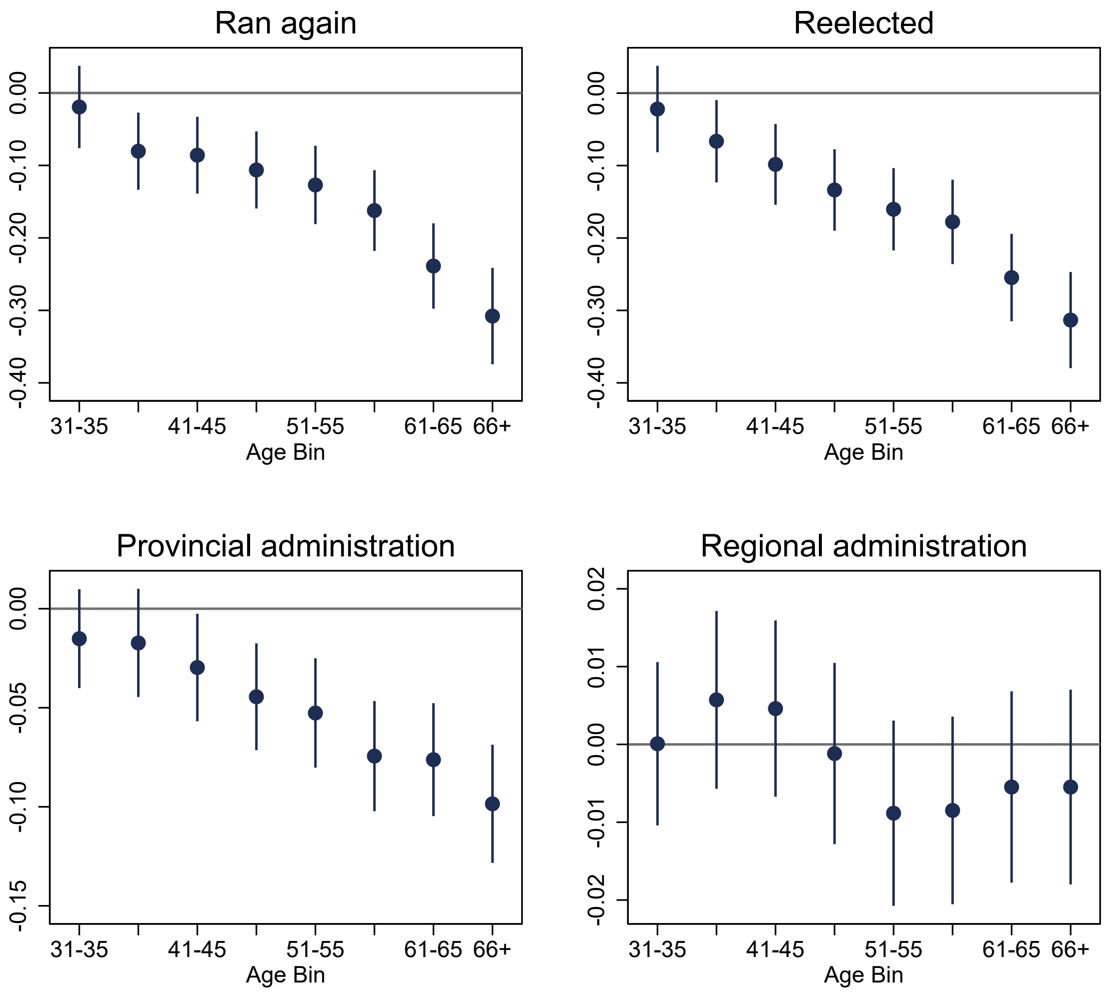

Notes: This figure plots point estimates and 95-percent confidence intervals for the coefficients $\{\beta_k\}$ in the regression $Y_{mt} = \sum_{k=1}^{8} \beta_k 1(Age_{mt} \in Bin_k) + \boldsymbol{\delta}' \boldsymbol{Z}_{mt} + \eta_m + \gamma_t + \varepsilon_{mt}$. The (omitted) reference age category is 18–30. Each included age bin is five years long, with the exception of the final bin, which is aggregated to avoid very small bin sizes. The vector $\boldsymbol{Z}$ contains controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city he or she governs.

Figure 3.2: Mayor Age and Average Public Finance Outcomes

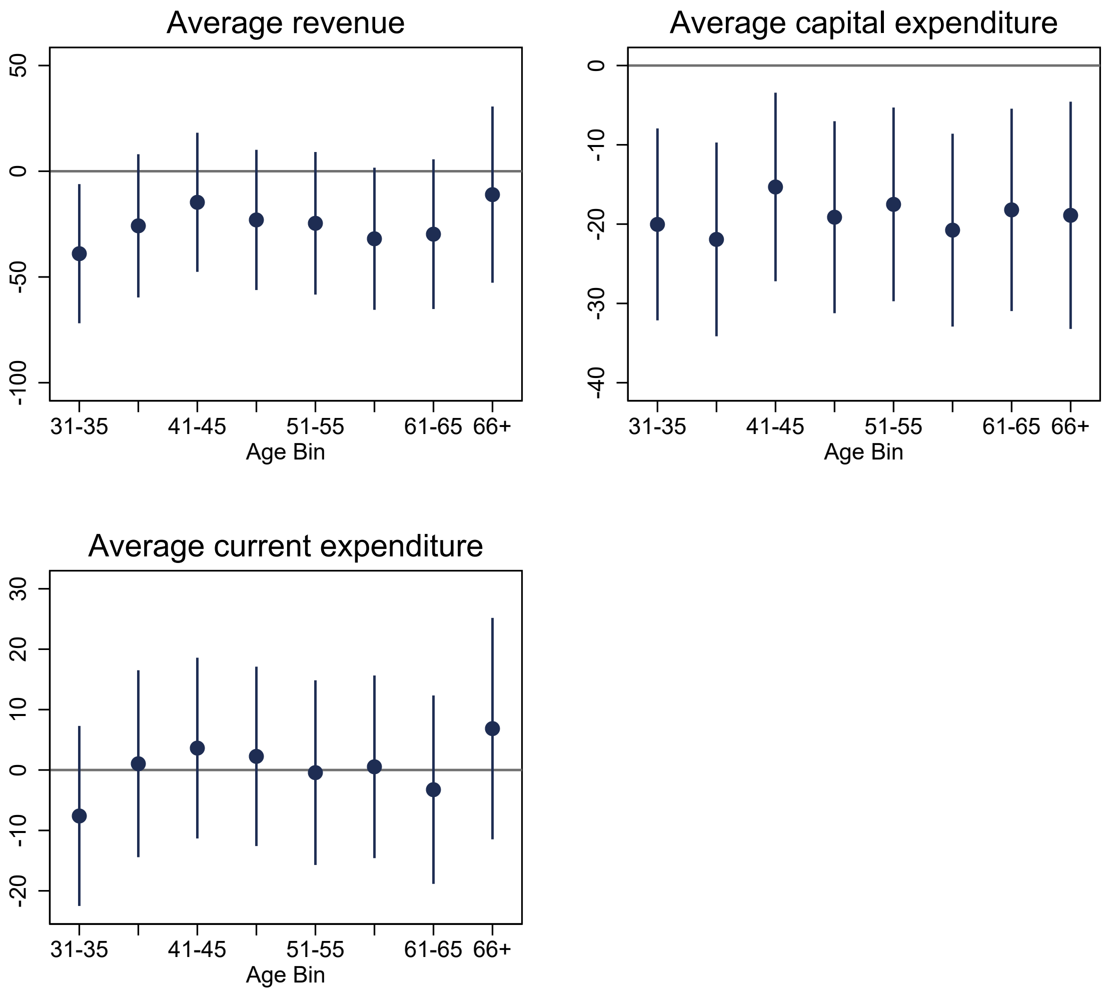

Notes: This figure plots point estimates and 95-percent confidence intervals for the coefficients $\{\beta_k\}$ in the regression $Y_{mt} = \sum_{k=1}^{8} \beta_k 1(Age_{mt} \in Bin_k) + \boldsymbol{\delta}' \boldsymbol{Z}_{mt} + \eta_m + \gamma_t + \varepsilon_{mt}$. The (omitted) reference age category is 18–30. Each included age bin is five years long, with the exception of the final bin, which is aggregated to avoid very small bin sizes. The vector $\boldsymbol{Z}$ contains controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city he or she governs.

Figure 3.3: Mayor Age and Other Municipal Outcomes

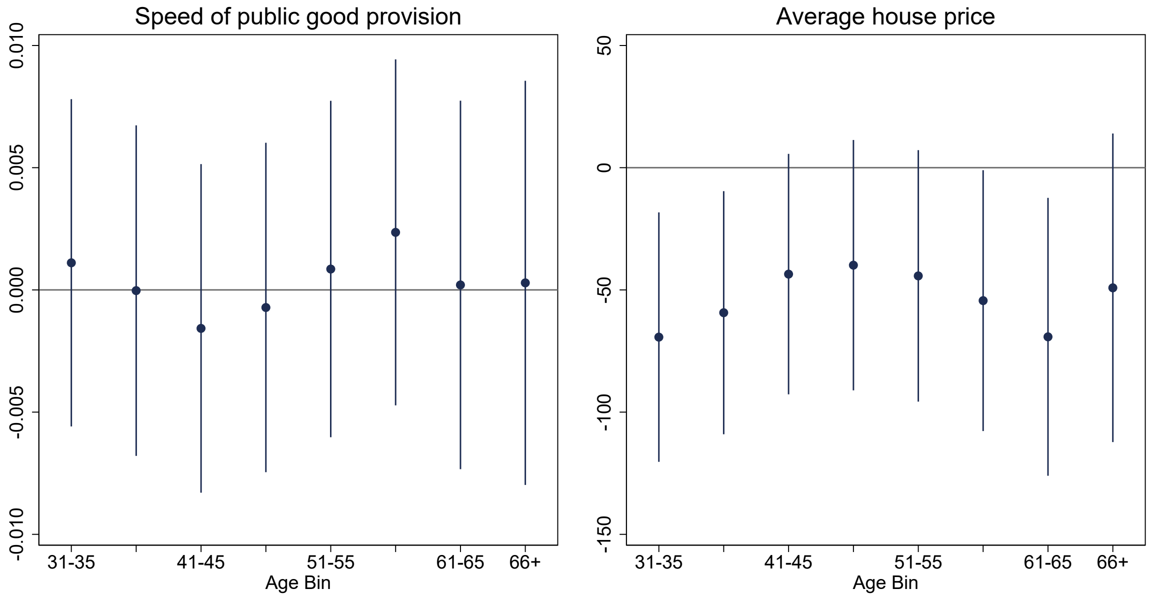

Notes: This figure plots point estimates and 95-percent confidence intervals for the coefficients $\{\beta_k\}$ in the regression $Y_{mt} = \sum_{k=1}^{8} \beta_k 1(Age_{mt} \in Bin_k) + \boldsymbol{\delta}' \boldsymbol{Z}_{mt} + \eta_m + \gamma_t + \varepsilon_{mt}$. The (omitted) reference age category is 18–30. Each included age bin is five years long, with the exception of the final bin, which is aggregated to avoid very small bin sizes. The vector $\boldsymbol{Z}$ contains controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city he or she governs.

Figure 3.4: Mayor Age and Political Budget Cycles

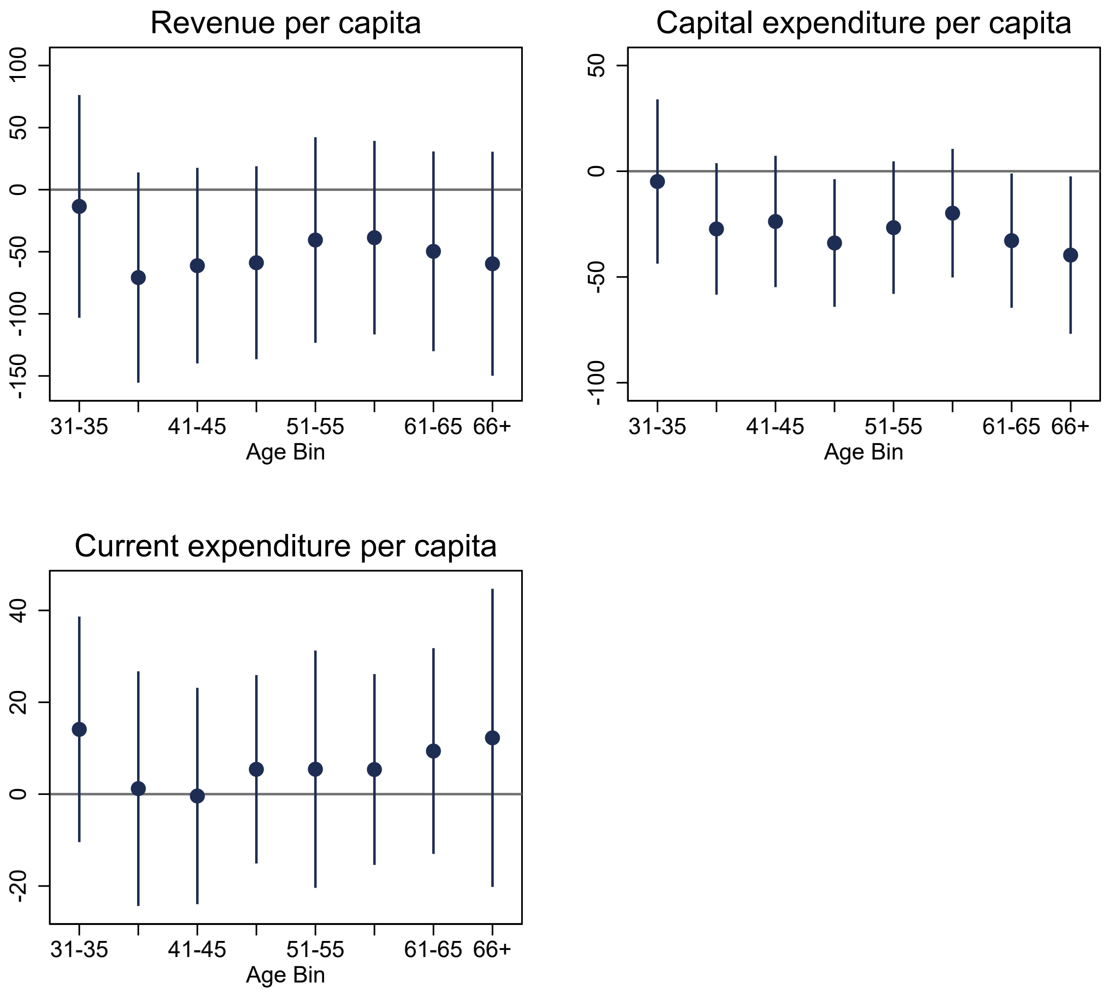

Notes: This figure plots point estimates and 95-percent confidence intervals for the coefficients $\{\delta_k\}$ in the regression $Y_{m,t} = \rho Y_{m,t-1} + \alpha Elec_{m,t} + \sum_{k=1}^{8} \beta_k 1(Age_{m,t} \in Bin_k) + \sum_{k=1}^{8} \delta_k Elec_{m,t} \times 1(Age_{m,t} \in Bin_k) + \eta_m + \zeta_t + \nu_{m,t}$. The sample is restricted to years following "close" elections decided by a vote margin of five percentage points or less. The (omitted) reference age category is 18–30. Each included age bin is five years long, with the exception of the final bin, which is aggregated to avoid very small bin sizes.

Appendix A

A.1 Details on the General Grant Reform

From 2002–2005 the expenditure-needs formula was

$$AvgExp \cdot (0.4 \cdot PopIndex_d + 0.1 \cdot PovGapIndex_d + 0.1 \cdot AreaIndex_d + 0.4 \cdot CostIndex_d),$$

where $AvgExp$ is average expenditure of all district governments, $PopIndex_d$ is the population of district d divided by average district population, and the other indices are defined analogously. Starting in 2006, the formula was

$$\begin{aligned} AvgExp \cdot (0.3 \cdot PopIndex_d + 0.1 \cdot 1/HDI_d + 0.15 \cdot GDPIndex_d \\ +0.15 \cdot AreaIndex_d + 0.3 \cdot CostIndex_d), \end{aligned}$$

where HDI stands for Human Development Index. The expenditure-needs formula changed in three ways. First, $AvgExp$ increased as a result of the budget expansion. Second, the poverty gap index was replaced by the (inverse of) the human development index and the regional GDP per capita index.[1] This change had little effect on equalization (World Bank, 2007). Third, the weights of the population, area, and cost indices changed. In particular, greater weight was giving to less densely populated districts. Rural districts tend to be poorer than urban districts in Indonesia. As a result, in 2006 the general grant increased for most districts, and the increase was much larger for poor, rural districts (World Bank, 2007). Furthermore, the policy change was persistent, as the expenditure-needs formula changed very little from 2006–2011 (Shah et al., 2012).[2] Holding fixed the Basic Allocation and Fiscal Capacity, the change in the

[1] The latter index is the regional GDP per capita relative to the average district GDP per capita.

[2] In 2010 and 2011 the weight on the area index changed to 0.1325 and 0.135, respectively, and the weights on the inverse HDI index and the GDP index increased slightly.

per capita general grant allocation to district d from 2005 to 2006 is given by

$$\begin{aligned}\frac{GenGrant_{d,06}}{Pop_{d,06}} - \frac{GenGrant_{d,05}}{Pop_{d,05}} &= \left(0.3 \cdot \frac{AvgExp_{06}}{AvgPop_{06}} - 0.4 \cdot \frac{AvgExp_{05}}{AvgPop_{05}}\right) \\ &+ \left(0.15 \cdot \frac{AvgExp_{06}}{AvgArea} \cdot \frac{Area_d}{Pop_{d,06}} - 0.1 \cdot \frac{AvgExp_{05}}{AvgArea} \cdot \frac{Area_d}{Pop_{d,05}}\right) \\ &+ \left(0.3 \cdot \frac{AvgExp_{06}}{Pop_{d,06}} \cdot \frac{Cost_{d,06}}{AvgCost_{06}} - 0.4 \cdot \frac{AvgExp_{05}}{Pop_{d,05}} \cdot \frac{Cost_{d,05}}{AvgCost_{05}}\right) \\ &+ \left(0.1 \cdot \frac{AvgExp_{06}}{Pop_{d,06}} \cdot \frac{1}{HDI_{d,06}} + 0.15 \cdot \frac{AvgExp_{06}}{Pop_{d,06}} \cdot \frac{GDP_{d,06}}{AvgGDP_{06}}\right. \\ &\quad \left. - 0.1 \cdot \frac{AvgExp_{05}}{Pop_{d,05}} \cdot \frac{PovGap_{d,05}}{AvgPovGap_{05}}\right).\end{aligned}$$

A useful approximation to the above expression obtains under the assumption of zero district population growth, zero change in the relative cost of construction across districts, and zero change in the relative poverty gap across districts.[3] Under these assumptions, the change in per capita general grant allocation can be expressed in terms of the total general grant budgets in 2005 and 2006 and district characteristics measured in 2006:

$$\begin{aligned}\frac{GenGrant_{d,06}}{Pop_{d,06}} - \frac{GenGrant_{d,05}}{Pop_{d,05}} &\approx \frac{(0.3 \cdot AvgExp_{06} - 0.4 \cdot AvgExp_{05})}{AvgPop_{06}} \\ &+ \frac{(0.15 \cdot AvgExp_{06} - 0.1 \cdot AvgExp_{05})}{AvgArea} \cdot \frac{Area_d}{Pop_{d,06}} \\ &+ \frac{(0.3 \cdot AvgExp_{06} - 0.4 \cdot AvgExp_{05})}{Pop_{d,06}} \cdot \frac{Cost_{d,06}}{AvgCost_{06}} \\ &+ \left(0.1 \cdot \frac{AvgExp_{06}}{Pop_{d,06}} \cdot \frac{1}{HDI_{d,06}} + 0.15 \cdot \frac{AvgExp_{06}}{Pop_{d,06}} \cdot \frac{GDP_{d,06}}{AvgGDP_{06}}\right. \\ &\quad \left. - 0.1 \cdot \frac{AvgExp_{05}}{Pop_{d,06}} \cdot \frac{PovGap_{d,06}}{AvgPovGap_{06}}\right).\end{aligned}$$

The second term on the right-hand side accounts for a large fraction of the cross-district variation in the general grant allocation change. The quantity $(0.15 \cdot AvgExp_{06} - 0.1 \cdot AvgExp_{05})$ is large and positive due to the overall general grant budget increase. This term is scaled by relative area per capita, $Area_d/(AvgArea \cdot Pop_{d,06})$. The change in general grant revenue received by district d from 2005 to 2006 can be approximated as

$$\frac{GenGrant_{d,06}}{Pop_{d,06}} - \frac{GenGrant_{d,05}}{Pop_{d,05}} \approx \theta + \pi \frac{Area_d}{Pop_{d,06}} + Remainder_d.$$

The above expression yields the approximate change in general grant revenue per capita for districts for which the reform to the expenditure-needs formula was binding. The formula dictated that districts rich in natural resources, which had substantial "fiscal capacity" according

[3] District annual population growth averaged 1.3 percent over the sample period, and median annual population growth was 1.4 percent.

to the formula, should have experienced a decline in general grant revenue over this period. Instead, a hold-harmless provision froze the general grant amount for such districts over this period.

A.2 Figures

Figure A.2.1: Shared Oil and Gas Revenue and Endowment per Capita by Province, 2014 Borders

(a) Shared Oil and Gas Revenue per Capita by Province, 2014

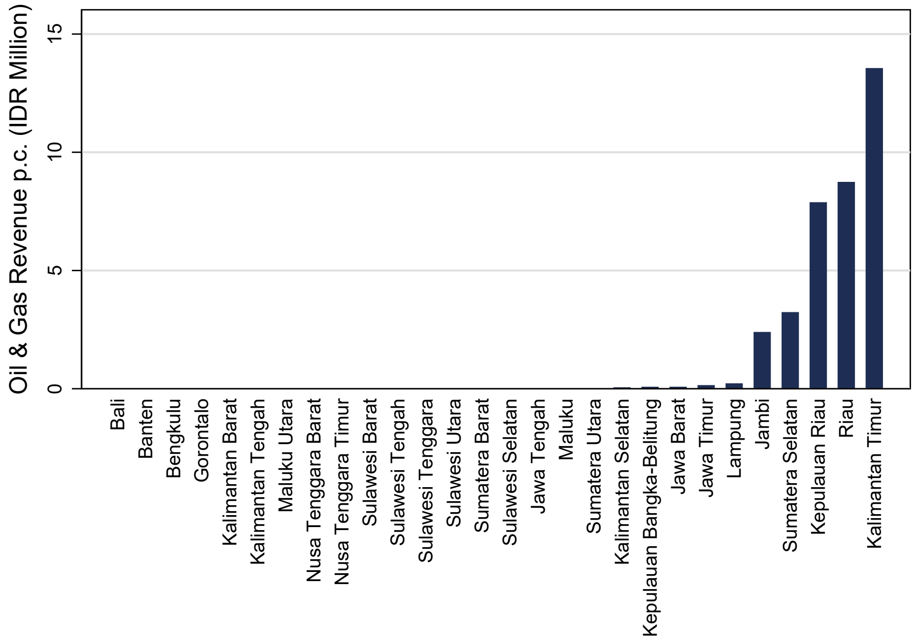

(b) Oil and Gas Endowment per Capita by Province

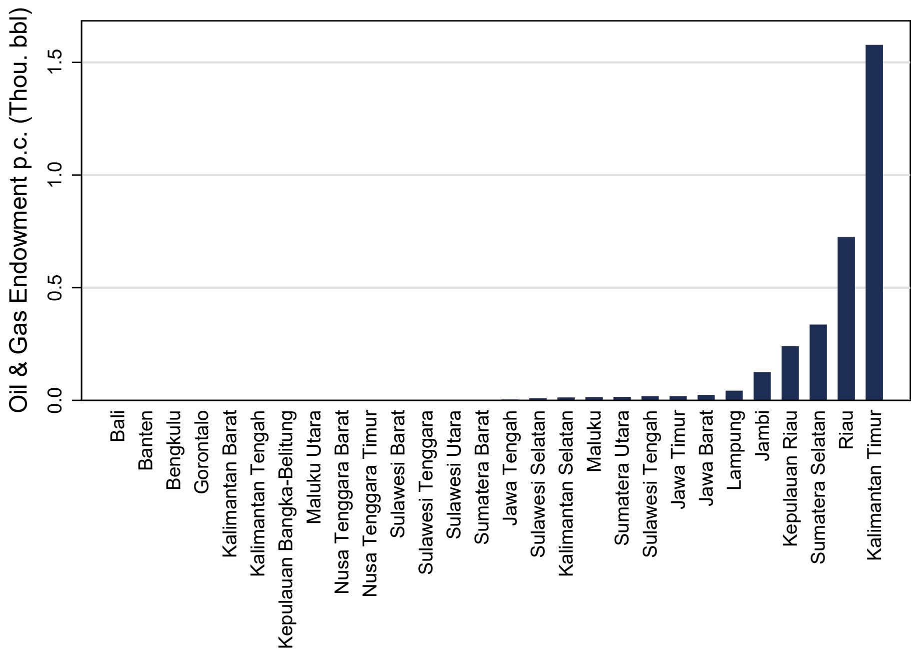

Notes. Panel (a) shows the total shared oil and gas revenue per capita by province in 2014. Revenue is expressed in constant 2010 rupiah (millions). Panel (b) shows the oil and gas endowment per capita for each province known as of the year 2000 based on 2014 borders and population. Oil and gas endowment per capita is expressed in thousands of barrels of oil equivalent per capita. In both panels Kalimantan Utara is combined with its parent province, Kalimantan Timur, consistent with the national government's revenue sharing policy through 2014.

Figure A.2.2: Oil and Gas Fields of Indonesia

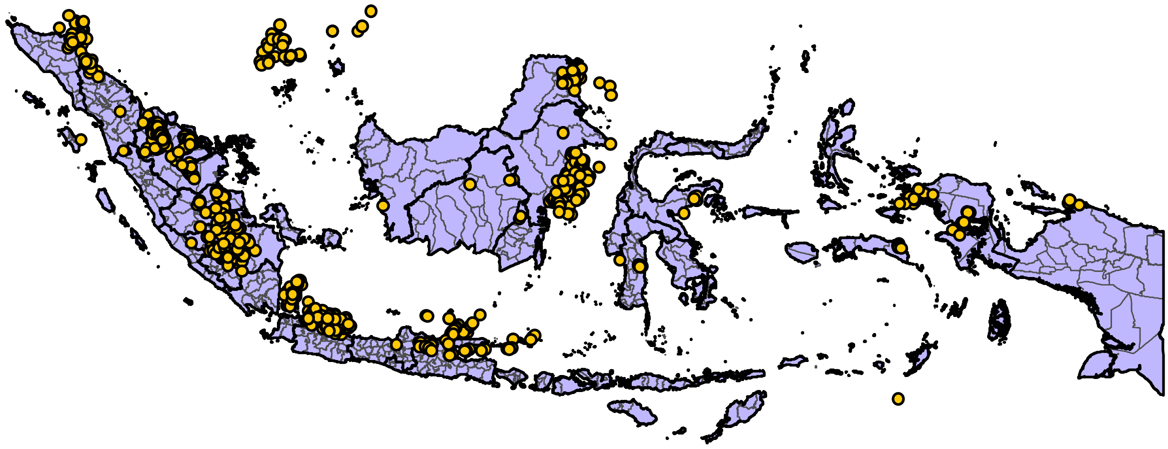

Source. Rystad Energy.

A.3 Tables

Table A.3.1: Flypaper Effect

	Total Expenditure per Capita	
	(1) OLS	(2) 2SLS
General Grant p.c.	0.808***	0.357
	(0.181)	(2.393)
Lag 1	0.288*	1.040
	(0.164)	(2.595)
Lag 2	0.080	0.394
	(0.169)	(1.404)
Lag 3	−0.158	−0.461
	(0.177)	(0.593)
Oil & Gas Revenue p.c.	0.633***	0.432*
	(0.203)	(0.255)
Lag 1	0.586***	0.637
	(0.177)	(0.817)
Lag 2	0.531***	0.321
	(0.189)	(0.548)
Lag 3	0.392***	0.377
	(0.109)	(0.346)
Non-Oil/Gas GDP p.c.	0.005	0.012
	(0.006)	(0.020)
Lag 1	0.002	−0.007
	(0.006)	(0.025)
Lag 2	0.002	−0.002
	(0.005)	(0.023)
Lag 3	0.001	0.005
	(0.008)	(0.013)
Oil/Gas GDP p.c.	0.010**	0.009**
	(0.004)	(0.004)
Lag 1	−0.018***	−0.018**
	(0.005)	(0.009)
Lag 2	0.001	0.003
	(0.002)	(0.003)
Lag 3	−0.002	−0.000
	(0.004)	(0.005)
Coef. sum: General Grant p.c.	1.018	1.330
	(0.189)	(0.835)
Coef. sum: Oil & Gas Revenue p.c.	2.142	1.766
	(0.513)	(1.016)
Coef. sum: Non-Oil/Gas GDP p.c.	0.009	0.008
	(0.008)	(0.009)
Coef. sum: Oil/Gas GDP p.c.	−0.008	−0.006
	(0.005)	(0.010)
Test: equal sums, Gen. Grant & Non-Oil/Gas GDP	0.000	0.115
Test: equal sums, Gen. Grant & Oil/Gas GDP	0.000	0.113
Test: equal sums, Oil & Gas Rev. & Non-Oil/Gas GDP	0.000	0.082
Test: equal sums, Oil & Gas Rev. & Oil/Gas GDP	0.000	0.080
Observations	3,214	3,214
District clusters	372	372
Prov. × year clusters	278	278

Notes. Each regression includes a full set of district and island × year dummies. Standard errors, reported in parentheses, are robust to heteroskedasticity and two-way clustering by district and province × year. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix B

B.1 Heterogeneous Effects: Theory

B.1.1 The Environment

Suppose the economy is populated by an autocrat and a continuum of citizens. There are two time periods, indexed by $t \in \{1, 2\}$. In period one the state of the world is autocracy, and in period two the state is either autocracy or democracy and is denoted by $S \in \{A, D\}$. There are two types of (exogenous) income in the economy: private income and natural resource rents. Each period citizens receive state-dependent private income, and the government receives natural resource rents in the amount of $R_t \geq 0$.[1] Following Acemoglu and Robinson (2006), we assume that there are two groups of citizens, the rich and the poor.[2] The individual private incomes of the rich and the poor in state S are y_S^r and y_S^p, respectively. The total population of citizens is normalized to unity, and a fraction δ are rich, where $\delta < 1/2$. Total private income coincides with average private income and is equal to $\bar{y}_S = \delta y_S^r + (1 - \delta) y_S^p$. Letting φ denote the fraction of total income held by the rich, the per capita incomes of the rich and poor can be written as

$$y_S^r = \frac{\varphi \bar{y}_S}{\delta} \qquad \text{and} \qquad y_S^p = \frac{(1 - \varphi)\bar{y}_S}{1 - \delta}, \tag{B.1.1}$$

where $\varphi > \delta$. All citizens are risk neutral.

Private income is potentially taxed under both autocracy and democracy. Under autocracy citizens receive group-specific transfers, or "bribes," from the autocrat, whereas under democracy all citizens receive a lump-sum transfer of equal size. Thus the indirect utilities of citizen

[1] For example, natural resource rents could arrive in the form of profits from state-owned resource firms or royalties paid by international resource firms.

[2] In contrast to Acemoglu and Robinson (2006), here the rich group is separate from the ruling elite and can potentially challenge the power of the elite.

i in states A and D, respectively, are

$$V_A^i = (1 - \tau_A)y_A^i + b^i \quad \text{and} \quad V_D^i = (1 - \tau_D)y_D^i + T,$$

where τ_S is the tax rate, b^i is the group-specific bribe, and T is the lump-sum transfer. There is an aggregate cost of taxation that is proportional to total income, $C(\tau_S)\bar{y}_S$. We assume that costs are low at low levels of taxation and are increasing and convex for strictly positive tax rates: $C(0) = 0$, $C'(\cdot) > 0$, and $C''(\cdot) > 0$. We also assume $C'(0) = 0$ and $C'(1) = 1$ to ensure an interior solution to the problem that follows. The capacity to tax is nil in period one, but τ_S may be positive in period two.

Under democracy tax revenue and resource rents are shared equally among the citizens. Thus period-two transfers satisfy the budget constraint,

$$T \leq (\tau_D - C(\tau_D))\bar{y}_D + R_2.$$

The (deposed) autocrat receives income normalized to zero.

Under autocracy the autocrat confiscates the tax revenue and resource rents.[3] However, there are transaction costs associated with stealing government revenue, so the autocrat receives only a fraction $(1 - \theta)$ of government revenue, where $\theta \in [0, 1]$. Transaction costs may stem from transparency of the budget or administrative procedures (Persson and Tabellini, 2000). More generally, transaction costs depend on the strength of accountability groups which constrain executive power.[4] The greater is the capacity of accountability groups to constrain the executive's ability to act unilaterally, the higher is θ. Let aggregate bribes be denoted by $b = \delta b^r + (1 - \delta)b^p$. When the autocrat makes aggregate bribes in the amount of b, he incurs a cost of $(1 + \gamma)b$ in period one and group i enjoys the benefits of b^i in period two.[5] Similar to θ, the parameter $\gamma > 0$ captures the marginal transaction cost of making bribes and depends on executive constraints. Assume that the autocrat is risk neutral and discounts future utility by the factor $\beta \in (0, 1)$, where $\beta > \varphi$. The autocrat's indirect utility in period t under autocracy is

[3]Using data on deposits to offshore bank accounts, Andersen et al. (2014) show that political elites appropriate oil rents in oil-rich autocracies but not in oil-rich democracies.

[4]A powerful legislature and an independent judiciary are archetypal accountability groups, but in nondemocracies executive accountability may derive from other sources. In a one-party government the executive may be constrained by senior officials in the ruling party. In a monarchy a council of nobles may provide a check on the king's power. The military may even provide a counterbalance in coup-prone polities (Marshall and Gurr, 2014). Finally, powerful producer groups, such as the cattle ranchers in Botswana, can restrain executive power (Acemoglu, Johnson, and Robinson, 2003). Strong accountability groups force the autocrat to use convoluted, opaque methods of stealing the rents, costing the autocrat θR_t. Alternatively, one can think of θ as the fraction of rents the autocrat must pay to accountability groups as bribes in exchange for keeping a fraction $1 - \theta$ of the rents. Interpreting the allocation of rents as the result of a Nash bargaining game, θ represents the bargaining power of accountability groups relative to the ruler.

[5]This timing assumption could capture the fact that many potential group-specific transfers—public employment, targeted public goods, or exclusive production rights—are enjoyed with a time lag. The autocrat's period-one cost of providing b could reflect an upfront investment cost or an opportunity cost of guaranteeing liquidity in period two.

equal to consumption, c_t, where

$$0 \leq c_1 \leq (1-\theta)R_1 - (1+\gamma)b$$

$$\text{and} \qquad 0 \leq c_2 \leq (1-\theta)\left[R_2 + (\tau_A - C(\tau_A))\bar{y}_A\right].$$

Note that we have assumed that the autocrat is credit-constrained. This is a reasonable assumption to a first approximation: the more unilateral authority the ruler has, the less likely he is to be compelled to repay a loan, making him a risky borrower.[6]

B.1.2 The Political Game

Timing. The timing of events is as follows. In the beginning of the first period, the autocrat receives $(1-\theta)R_1$ and announces period-two policies (τ_A, b^r, b^p). We assume that the autocrat can fully commit to period-two policies in period one.[7] Tax policy is set with a one-period delay, so the autocrat can only choose period-two taxes.[8] At the end of the first period, the citizens decide whether to stage a revolution to depose the autocrat. We assume that the revolution succeeds if and only if both groups of citizens participate. A group of citizens participate in the revolution if and only if their period-two payoff under democracy strictly exceeds their period-two payoff under autocracy, given the (binding) promises of the autocrat. We assume that citizens can commit to their period-two rebellion decision in period one. If the revolution succeeds, then the state transitions to democracy, the autocrat receives zero income, and rich and poor citizens vote on the tax rate and transfers and receive payoffs V_D^r and V_D^p. If the revolution fails, then the autocrat stays in power, implements policies (τ_A, b^r, b^p), and receives $(1-\theta)\left[R_2 + (\tau_A - C(\tau_A))\bar{y}_A\right]$; and rich and poor citizens receive payoffs V_A^r and V_A^p.

Period-two equilibrium. To characterize the subgame perfect Nash equilibrium, we work backwards and first consider the Nash equilibrium starting in period two. If the state is autocracy in the second period, then each player's strategy and payoff is determined by policy commitments made in the first period. Citizen i receives $(1-\tau_A)y_A^i + b^i$ and the autocrat receives $(1-\theta)\left[R_2 + (\tau_A - C(\tau_A))\bar{y}_A\right]$.

If the state is democracy in the second period, then citizens vote on the tax rate, τ_D, and the level of lump-sum transfers, T. Because utility is strictly increasing in transfers (all else equal), the budget constraint will always bind: $T = (\tau_D - C(\tau_D))\bar{y}_D + R_2$. For a given value of

[6] See, for example, North and Weingast (1989).

[7] Thus we abstract from the possibility that democratization could result from the elite's inability to commit to future policy (Acemoglu and Robinson, 2006).

[8] Taxation requires significant investments in the government's ability to monitor citizens and enforce the tax code (Besley and Persson, 2011). For simplicity we capture this fact by assuming that tax policy is implemented with a delay, abstracting from investment costs.

τ_D, the payoff of citizen i under democracy is

$$(1-\tau_D)y_D^i + (\tau_D - C(\tau_D))\bar{y}_D + R_2. \tag{B.1.2}$$

Let τ_D^i denote the most preferred tax rate of citizen i. Because there are no public goods in this economy, the sole function of the tax is redistribution. Therefore $\tau_D^r = 0$. Substituting (B.1.1) into (B.1.2), it is straightforward to show that the most preferred tax rate of a poor citizen satisfies

$$C'(\tau_D^p) = \frac{\varphi - \delta}{1-\delta}.$$

It follows from our assumptions that $\tau_D^p \in (0,1)$ and τ_D^p is increasing in the amount of inequality, φ. It is possible to show that both poor and rich citizens have single-peaked preferences over τ_D.[9] Suppose that under democracy τ_D is chosen by pairwise majority voting in an environment with no uncertainty. Then by the median-voter theorem, the most preferred policy of the median voter, τ_D^p, is selected (Black, 1948; Downs, 1957). The equilibrium payoff to citizen i under democracy is then

$$V_D^i = (1-\tau_D^p)y_D^i + (\tau_D^p - C(\tau_D^p))\bar{y}_D + R_2.$$

Period-one equilibrium. At the end of period one, each citizen chooses whether to participate in the revolution, given the period-two equilibrium policies under autocracy, (τ_A, b^r, b^p), and under democracy, (τ_D^p, T). Citizen i participates in the revolution if and only if $V_D^i > V_A^i$. Equivalently, for each value of τ_A, citizen i participates in the revolution if and only if $b^i < \tilde{b}^i(\tau_A)$, where

$$\tilde{b}^i(\tau_A) = (1-\tau_D^p)y_D^i + (\tau_D^p - C(\tau_D^p))\bar{y}_D + R_2 - (1-\tau_A)y_A^i.$$

Note that $\tilde{b}^i(\tau_A)$ is strictly increasing in τ_A: increasing the tax rate under autocracy causes citizen i to demand a larger reservation bribe in exchange for not rebelling. The following assumption ensures that democracy is sufficiently appealing relative to autocracy that $\tilde{b}^i(\tau_A) > 0$ for any R_2 and τ_A.

Assumption B.1.1. *$G^i \equiv (1-\tau_D^p)y_D^i + (\tau_D^p - C(\tau_D^p))\bar{y}_D - y_A^i > 0$ for $i \in \{r, p\}$.*

In the beginning of period one, the autocrat chooses period-two policies, (τ_A, b^r, b^p), to maximize his lifetime discounted utility, taking the strategies of citizens as given. Letting $(\tau_A, b^r, b^p) \in \mathcal{P}$, the function $\rho : \mathcal{P} \mapsto \{0,1\}$ indicates whether the revolution is prevented,

[9] The strict convexity of $C(\cdot)$ guarantees that the indirect utility function is strictly concave in τ. This is a sufficient condition for preferences to be single-peaked.

where $\rho(\tau_A, b^r, b^p) = 1$ indicates prevention. The autocrat's problem is

$$\max_{\tau_A, b^r, b^p} (1-\theta)R_1 - (1+\gamma)b + \rho(\tau_A, b^r, b^p)\beta(1-\theta)\left[R_2 + (\tau_A - C(\tau_A))\bar{y}_A\right]$$

$$\begin{aligned} \text{subject to} \qquad & b = \delta b^r + (1-\delta)b^p \\ & (1+\gamma)b \leq (1-\theta)R_1 \\ & \rho(\tau_A, b^r, b^p) = \begin{cases} 1 & \text{if } b^r \geq \tilde{b}^r(\tau_A) \text{ or } b^p \geq \tilde{b}^p(\tau_A) \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Strictly speaking, R_2 denotes expected period-two resource rents from the perspective of period one.

For each τ_A it is optimal for the autocrat to pay bribes (b^r, b^p), with $b^i > 0$ for some $i \in \{r, p\}$, if and only if three conditions are satisfied:

(i) Sufficiency: $\rho(\tau_A, b^r, b^p) = 1$

(ii) Feasibility: $(1+\gamma)b \leq (1-\theta)R_1$

(iii) Desirability: $(1+\gamma)b \leq \beta(1-\theta)\left[R_2 + (\tau_A - C(\tau_A))\bar{y}_A\right]$.

The bribes are sufficient if they prevent the revolution, they are feasible if the autocrat has enough income in period one to cover the cost of the bribes, and they are desirable if the autocrat's expected benefit from staying in power exceeds the cost of the bribes. If no set of bribes satisfy all three conditions, the autocrat sets $b^r = b^p = 0$ and the state transitions to democracy in period two.

If the autocrat chooses to pay bribes to avert a revolution, it is optimal to bribe only one group of citizens. To simplify the analysis, we assume that the rich are cheaper to bribe than the poor.

Assumption B.1.2. $\delta\tilde{b}^r(\tau_A) \leq (1-\delta)\tilde{b}^p(\tau_A)$ *for all* $\tau_A \in [0, 1]$.

This assumption is reasonable because the rich are less numerous and have more to lose from democracy than the poor.[10] Assumption B.1.2 is more likely to hold the smaller is δ and the larger are τ_D^p, R_2, and $\bar{y}_A$. When the autocrat chooses to pay bribes, he will pay each rich citizen exactly $\tilde{b}^r(\tau_A)$ so that $b = \delta\tilde{b}^r(\tau_A)$.

We make the following parametric assumptions for γ.

Assumption B.1.3. $\beta/\varphi - 1 < \gamma < \beta/\delta - 1$.

The first inequality rules out the situation in which the autocrat both taxes and bribes the rich citizens in order to prevent a revolution. To see this, note that when $b = \delta\tilde{b}^r(\tau_A)$,

[10] Note that the assumption is weaker than assuming that $\tilde{b}^r(\tau_A) \leq \tilde{b}^p(\tau_A)$ for all $\tau_A \in [0, 1]$, because $\delta < 1/2$.

the autocrat's marginal cost of increasing τ_A is $(1+\gamma)\varphi\bar{y}_A$, while his marginal benefit is $\beta(1-\theta)(1-C'(\tau_A))\bar{y}_A$. Assumption B.1.3 guarantees that the marginal cost of increasing τ_A exceeds the marginal benefit for all values of τ_A and θ. Thus the autocrat will set $\tau_A = 0$ whenever $b = \delta\tilde{b}^r(\tau_A)$. The second inequality guarantees that a threshold value $\theta^*(\gamma)$, which will be described below, is strictly positive.

Noting that $\rho(0, \tilde{b}^r(0), 0) = 1$ and $\tilde{b}^r(0) = G^r + R_2$, where G^r is defined in Assumption B.1.1, the autocrat will set $\tau_A = 0$ and $b = \delta\tilde{b}^r(0)$ if and only if the following conditions are satisfied:

(i) Feasibility: $\delta(1+\gamma)(G^r + R_2) \leq (1-\theta)R_1$

(ii) Desirability: $\delta(1+\gamma)(G^r + R_2) \leq \beta(1-\theta)R_2$.

The following definitions are useful for studying the comparative statics of the model.

Definition B.1.4. *A **resource boom** is an increase in both R_1 and R_2.*

Definition B.1.5. *In an economy with parameter values (δ, γ, θ), a **balanced resource boom** is a resource boom that satisfies*

$$\frac{\Delta R_2}{\Delta R_1} < \frac{1-\theta}{\delta(1+\gamma)}.$$

Because an increase in R_2 increases the attractiveness of democracy to the citizens, if the increase in R_2 far exceeds the increase in R_1, the autocrat will be unable to pay the reservation bribe of the rich. In contrast, a balanced resource boom increases the likelihood that the feasibility constraint is satisfied, because the increase in R_2 is not "too large" relative to the increase in R_1. Because Assumption B.1.3 implies that $\delta(1+\gamma) < \beta$, a balanced resource boom could involve $\Delta R_2 > \Delta R_1$. Note that a resource boom is more likely to be balanced the smaller are γ and θ.

B.1.3 Results

We are now ready to state the main results.

Proposition B.1.6. *For each γ there exists a threshold value $\theta^*(\gamma) \in (0,1)$ such that for $\theta < \theta^*(\gamma)$, a balanced resource boom makes the transition to democracy less likely, the lower is θ. For $\theta \geq \theta^*(\gamma)$ the state transitions to democracy for any $(R_1, R_2) \geq 0$.*

Proof. First note that the feasibility constraint is satisfied because the resource boom is balanced. Let $\theta^*(\gamma) = 1 - \delta(1+\gamma)/\beta$, which is in $(0,1)$ by Assumption B.1.3. When $\theta < \theta^*(\gamma)$, we have that $\beta(1-\theta) - \delta(1+\gamma)$ is positive and decreasing in θ. This means that an increase in R_2 increases the likelihood that the desirability constraint is satisfied, and the marginal effect of R_2 on desirability is decreasing in θ. When $\theta \geq \theta^*(\gamma)$, the desirability constraint is always violated. □

Proposition B.1.6 states that a balanced increase in resource rents will lower the chances of democratization when constraints on the ruler are sufficiently weak. However, when constraints are strong, no resource boom can impede democratization. The assumption that the autocrat is credit-constrained necessitates that the resource boom be balanced. Note that both types of marginal transaction costs induced by executive constraints, γ and θ, matter for the outcome. For example, lowering γ (subject to Assumption B.1.3 holding) increases $\theta^*(\gamma)$, raising the likelihood that a balanced resource boom impedes democratization for a given value of θ.

Corollary B.1.7. *There exists a threshold value $\theta^* \in (0,1)$ such that for $\theta < \theta^*$, a balanced resource boom is more likely to result in zero tax revenue, the lower is θ. For $\theta \geq \theta^*$ taxes are positive for any $(R_1, R_2) \geq 0$.*

Proof. The result follows immediately from Proposition B.1.6 by noting that under autocracy, $\tau_A = 0$, while under democracy, $\tau_D = \tau_D^p > 0$. □

The prediction of Corollary B.1.7 contrasts with that of the fiscal capacity model of Besley and Persson (2011). In their model political transitions are exogenous and taxation is used either to fund a public good or to redistribute income to the group in power. An increase in resource wealth leads to lower taxes only when institutions are "cohesive," i.e., θ is large. This is because in their model tax revenue is spent on the public good when institutions are cohesive, and the diminishing marginal utility of the public good implies that tax revenue is less valuable after a resource windfall that relaxes the budget constraint. For small values of θ, resource wealth does not affect equilibrium taxation in their model. In our model the mechanism determining the tax rate is quite different: the political transition is endogenous, and equilibrium taxation depends on the incumbent's ability and willingness to use patronage to remain in power. Figure B.3.4 graphically demonstrates how the effect of a resource boom on the suppression decision depends on the strength of executive constraints.

B.2 Tables

Table B.2.1: Variable Descriptions and Sources

Variable	Definition	Source
Democracy, 2008	POLITY2 index in 2008, normalized to take values between zero and one	Polity IV
Avg. Democracy, 1966–2008	Average normalized POLITY2 index from 1966–2008 in years in which the country was independent	Polity IV
Corruption, 2008	Corruption index in 2008 ranging from 0 to 6, with higher numbers indicating less corruption	PRS
Internal Conflict, 1966–2008	Internal or internationalized internal armed conflicts per year in which country was independent from 1966–2008	UCDP/PRIO
Coup Attempts, 1966–2008	(Failed or successful) coup attempts per year in which country was independent from 1966–2008	Polity IV
Purges, 1966–2008	Political purges per year in which country was independent from 1966–2008	CNTS
Total Revenue, 2000–2008	Log of average government revenue share of GDP from 2000–2008	ICTD
Tax Revenue, 2000–2008	Log of average tax revenue share of GDP from 2000–2008	ICTD
GDP, 2008	Log of per capita GDP in 2008 in constant 2011 international dollars	WDI
GDP, 1960, . . . , GDP, 2008	Log of per capita GDP in different years in constant 1990 U.S. dollars	Maddison
Executive Constraints, 1950–1966	Average XCONST index from 1950–1965 after normalizing XCONST to take values between zero and one	Polity IV
Weak Constraints, 1950–1966	Indicates having averaged three points or fewer out of seven on XCONST from 1950–1965	Polity IV
Oil Production, 1966–2008	Log of average annual metric tons of oil produced per 1000 inhabitants from 1966–2008	Ross
Oil Discovery, 1966–2003	Log of average annual millions of barrels discovered per 1000 inhabitants from 1966–2003	ASPO
Oil Reserves, 1966–2003	Log of average annual millions of barrels of reserves per 1000 inhabitants from 1966–2003	ASPO
Oil Endowment	Log of total oil endowment in millions of barrels per 1000 inhabitants in 1960	ASPO
Oil Quality	Average oil quality weighted by production volume; quality ranges from 1 (no oil) to 11 (ultra light)	WOGR
Basin Type Area	Log of sovereign area covered by a type of basin in square km per 1000 inhabitants in 1960; see Tables 2.1, 2.2, 2.3	Tellus
Land Area	Log of land area in square km per 1000 inhabitants in 1960	GIS
Coastline	Log of length of coastline in km per 1000 inhabitants in 1960	CIA
Mountainous Area	Log of mountainous land area in square km per 1000 inhabitants in 1960	FL
Tropical Area	Log of land area falling within tropics in square km per 1000 inhabitants in 1960	GSM
Good Soil Area	Log of land area containing "good" soil in square km per 1000 inhabitants in 1960	GAEZ

Notes. Polity IV stands for the Polity IV Project (Marshall and Gurr, 2014; Marshall and Marshall, 2016). PRS stands for Political Risk Services. UCDP/PRIO stands for the UCDP/PRIO Armed Conflict Dataset (Gleditsch et al., 2002). CNTS stands for Cross-National Time-Series Data Archive (Banks and Wilson, 2016). ICTD stands for International Centre for Tax and Development (Prichard et al., 2014). WDI stands for the World Bank World Development Indicators. Maddison stands for the Maddison Project (Maddison, 2013). Ross stands for Ross (2013). ASPO stands for Association for the Study of Peak Oil. WOGR stands for the World Oil and Gas Review published by ENI (ENI, 2015). Tellus stands for the Fugro Robertson, Ltd. (2013) Tellus GIS database. GIS stands for author's calculation using ArcGIS. CIA stands for CIA World Factbook (CIA, 2015). FL stands for Fearon and Laitin (2003). GSM stands for Gallup et al. (1998). GAEZ stands for the FAO's Global Agro-Ecological Zones database (version 3.0) (Fischer et al., 2002).

Table B.2.2: Total Basin Coverage of Sovereign Area by Region

	Mean	Std. Dev.	Min.	Max.	Obs.
East Asia and the Pacific	0.39	0.25	0.00	0.75	20
Eastern Europe and Central Asia	0.67	0.28	0.13	1.00	23
Rest of Europe and Neo-Europes	0.57	0.32	0.00	1.00	26
Latin America and the Caribbean	0.56	0.22	0.12	0.99	30
Middle East and North Africa	0.86	0.20	0.35	1.00	21
Sub-Saharan Africa	0.44	0.28	0.00	0.90	45
South Asia	0.55	0.32	0.03	1.00	7
Total	0.56	0.30	0.00	1.00	172

Notes. This table summarizes the portion of country sovereign area containing any type of sedimentary basin.

Table B.2.3: Weak Executive Constraints and Basins: Tectonic-Subsidence Grouping

	Weak Executive Constraints, 1950–1965							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Convergent C-C mechanical	0.028							
	(0.019)							
Convergent O-C thermal		−0.047						
		(0.062)						
Convergent O-C mechanical			−0.018					
			(0.015)					
Convergent O-O				−0.036*				
				(0.022)				
Divergent thermal					−0.003			
					(0.014)			
Wrench mechanical						0.016		
						(0.012)		
Divergent mechanical							−0.003	
							(0.017)	
Convergent C-C thermo-mechanical								−0.020
								(0.052)
Observations	116	116	116	116	116	116	116	116
R^2	0.184	0.175	0.183	0.187	0.172	0.184	0.172	0.172

Notes. See Table B.2.1 for variable definitions. The variable "weak constraints" is an indicator for having averaged three points or fewer out of seven on XCONST from 1950–1965. A score of three points for XCONST indicates "slight to moderate limitation on executive authority" (Polity IV). In practice "weak constraints" indicates having an average XCONST score equal to or below the median average XCONST score from 1950–1965. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.2.4: Weak Executive Constraints and Basins: Final Component of Code Grouping

	Weak Executive Constraints, 1950–1965									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Foreland	0.001 (0.020)									
Intracratonic sag		0.003 (0.011)								
Passive margin			−0.013 (0.012)							
Convergent sag				0.000 (0.019)						
Post-rift sag					−0.004 (0.013)					
Wrench						0.016 (0.012)				
Extensional							−0.010 (0.024)			
Convergent wrench								−0.032** (0.014)		
Fore-arc									−0.035** (0.017)	
Rift										−0.003 (0.017)
Observations	116	116	116	116	116	116	116	116	116	116
R^2	0.171	0.172	0.180	0.171	0.172	0.184	0.172	0.205	0.201	0.172

Notes. See Table B.2.1 for variable definitions. The variable "weak constraints" is an indicator for having averaged three points or fewer out of seven on XCONST from 1950–1965. A score of three points for XCONST indicates "slight to moderate limitation on executive authority" (Polity IV). In practice "weak constraints" indicates having an average XCONST score equal to or below the median average XCONST score from 1950–1965. All specifications include geographic controls (land area, coastline, and mountainous area), climatic controls (tropical area and good soil area), and region fixed effects. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.3 Figures

Figure B.3.1: First Stage

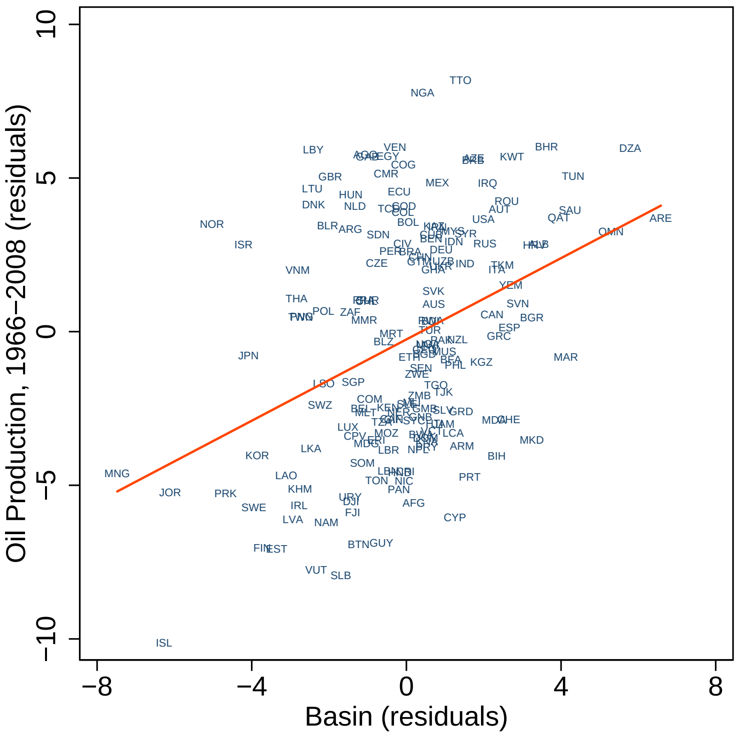

Notes. The figure plots oil production residuals against the residuals from *Basin*, where the residuals are obtained from separate regressions on the full set of geographic and climatic controls and region dummies.

Figure B.3.2: Second Stage

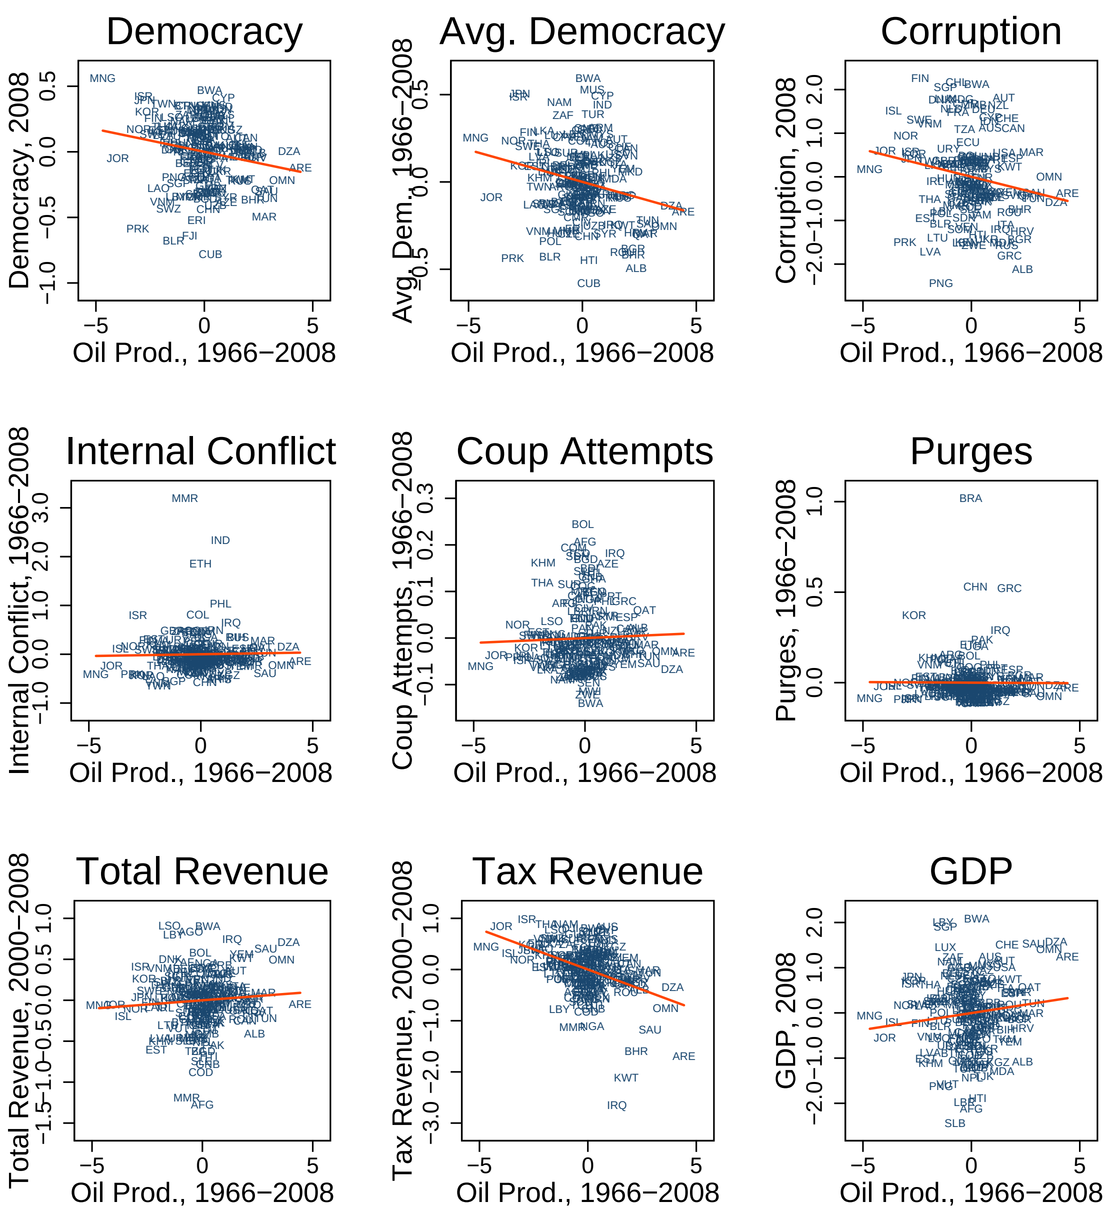

Notes. The figures plot outcome residuals against Oil Production predicted residuals. Each outcome residual is obtained by regressing the outcome variable on the full set of geographic and climatic controls and region dummies. The Oil Production predicted residuals are obtained by regressing the predicted values of Oil Production from the first stage on the full set of geographic and climatic controls and region dummies.

Figure B.3.3: Endowment and Basin

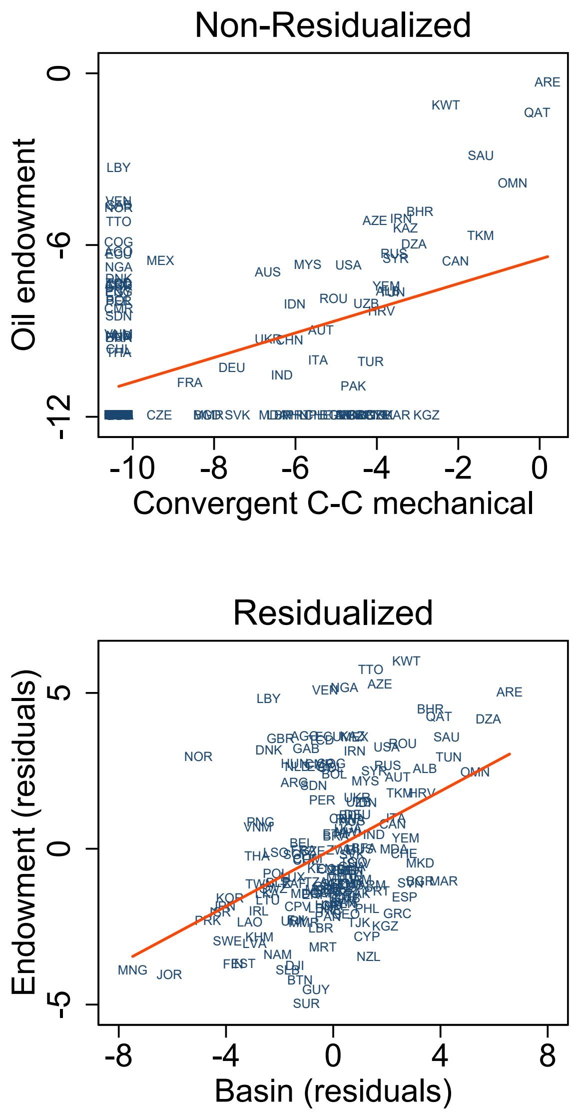

Notes. The figure plots oil endowment against *Basin*. The first graph is a raw scatterplot where the residuals are obtained from separate regressions on the full set of geographic and climatic controls and region dummies.

Figure B.3.4: Suppression Decision

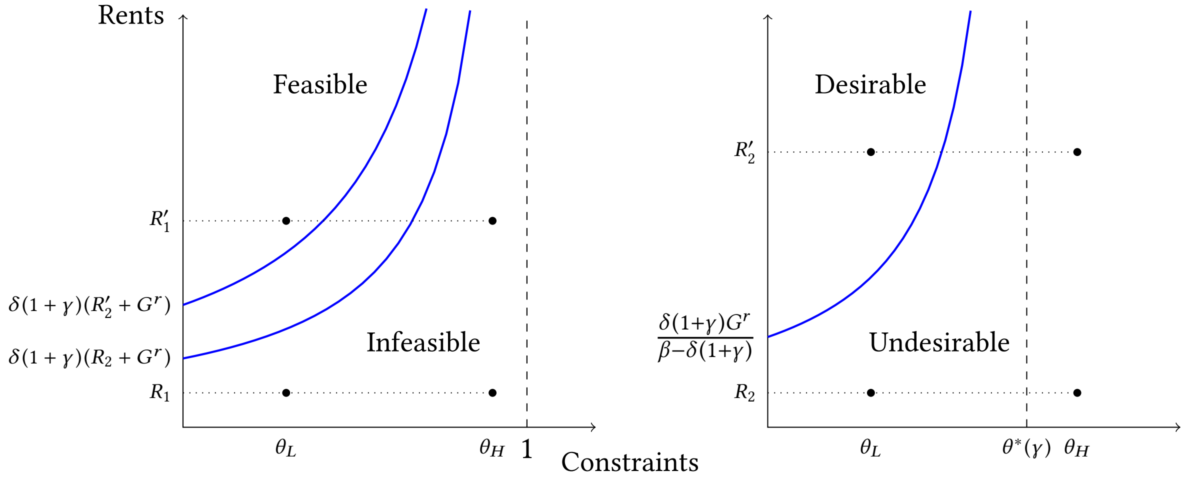

Notes. The figure shows the effect of a resource boom on the decision to suppress democracy in two different countries: one with weak executive constraints, θ_L, and the other with strong executive constraints, θ_H. In both countries period-one rents increase from R_1 to R_1', and period-two rents increase from R_2 to R_2'. The resource boom is balanced from the perspective of the country with weak constraints. Democracy is repressed if and only if (θ, R_1') lies above the blue line in the first graph (feasibility) and (θ, R_2') lies above the blue line in the second graph (desirability). Note that the increase in R_2 causes the blue line in the feasibility graph to shift upward, because it raises the reservation bribe of the rich group. In the country with weak constraints, the resource boom leads to repression, while the country with strong constraints transitions to democracy.

Figure B.3.5: Petroleum System

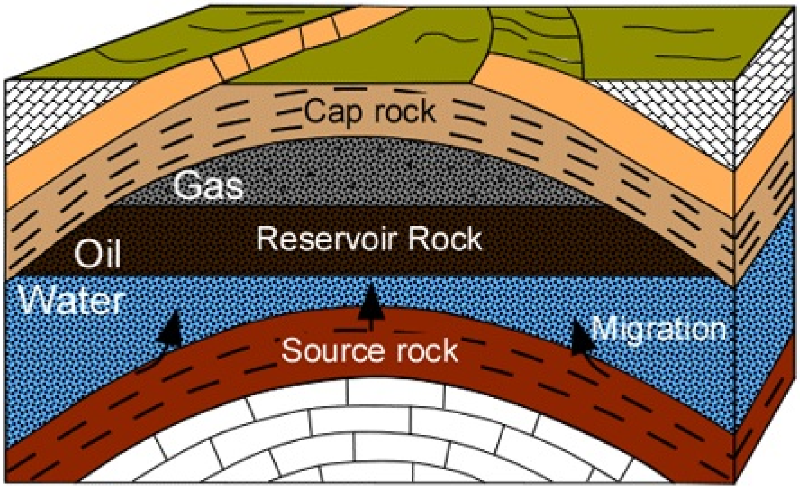

Source. Petrolia Haldimand Project.

Figure B.3.6: Peripheral Foreland and Passive Margin Basins

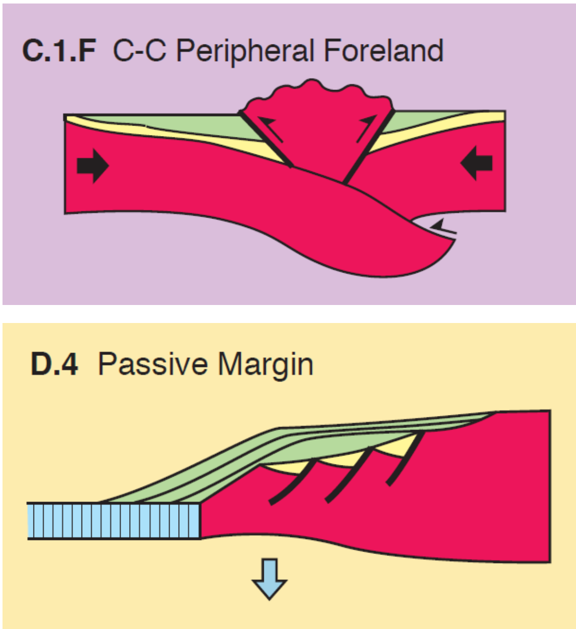

Source. Fugro Robertson, Ltd. (2013).
Notes. The tan region is old sediments, and the light blue-green region is newer sediments.

B.3.1 Basins grouped by plate-tectonic environment and primary subsidence mechanism

Figure B.3.7: Basins: Convergent Continent-Continent Tectonics, Mechanical Subsidence

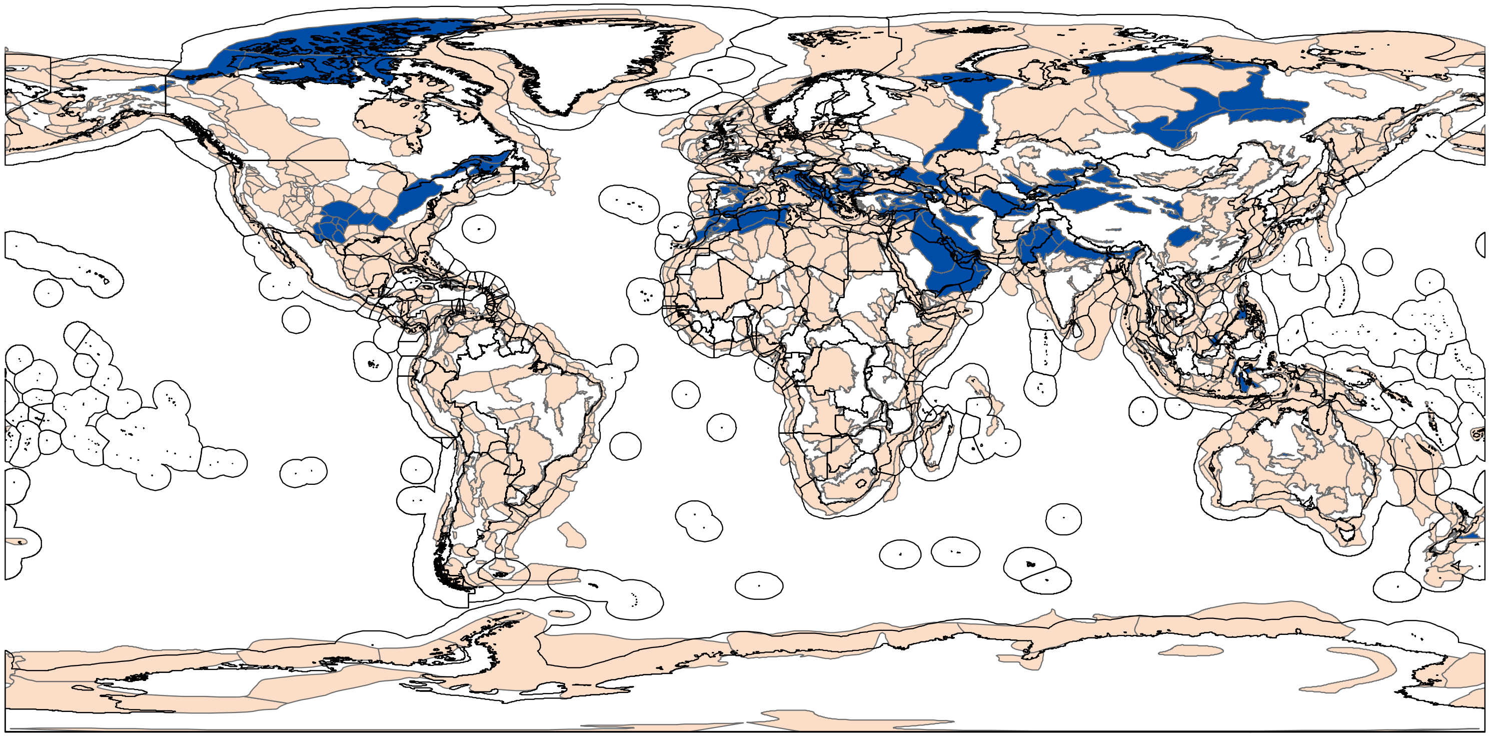

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.8: Basins: Convergent Continent-Continent Tectonics, Thermo-Mechanical Subsidence

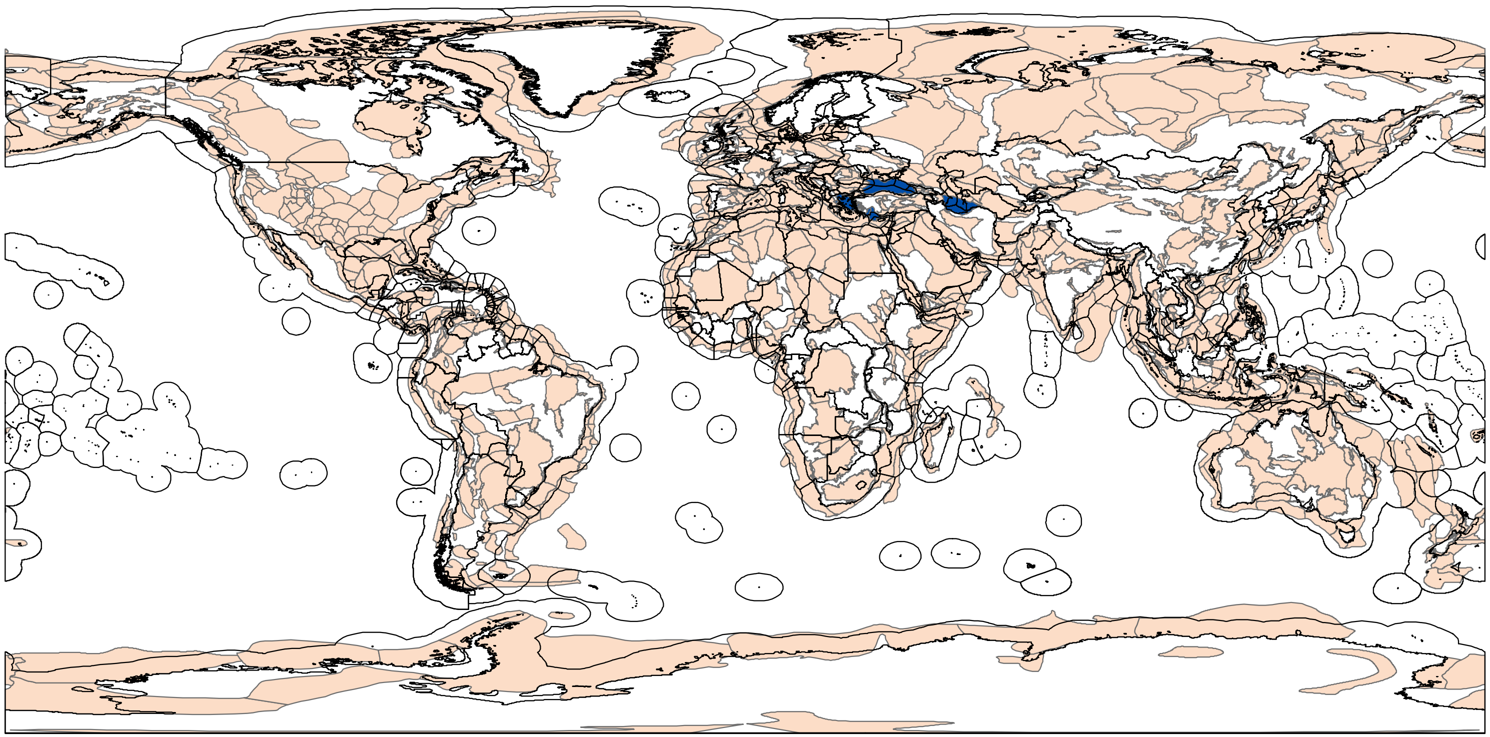

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.9: Basins: Convergent Ocean-Continent Tectonics, Mechanical Subsidence

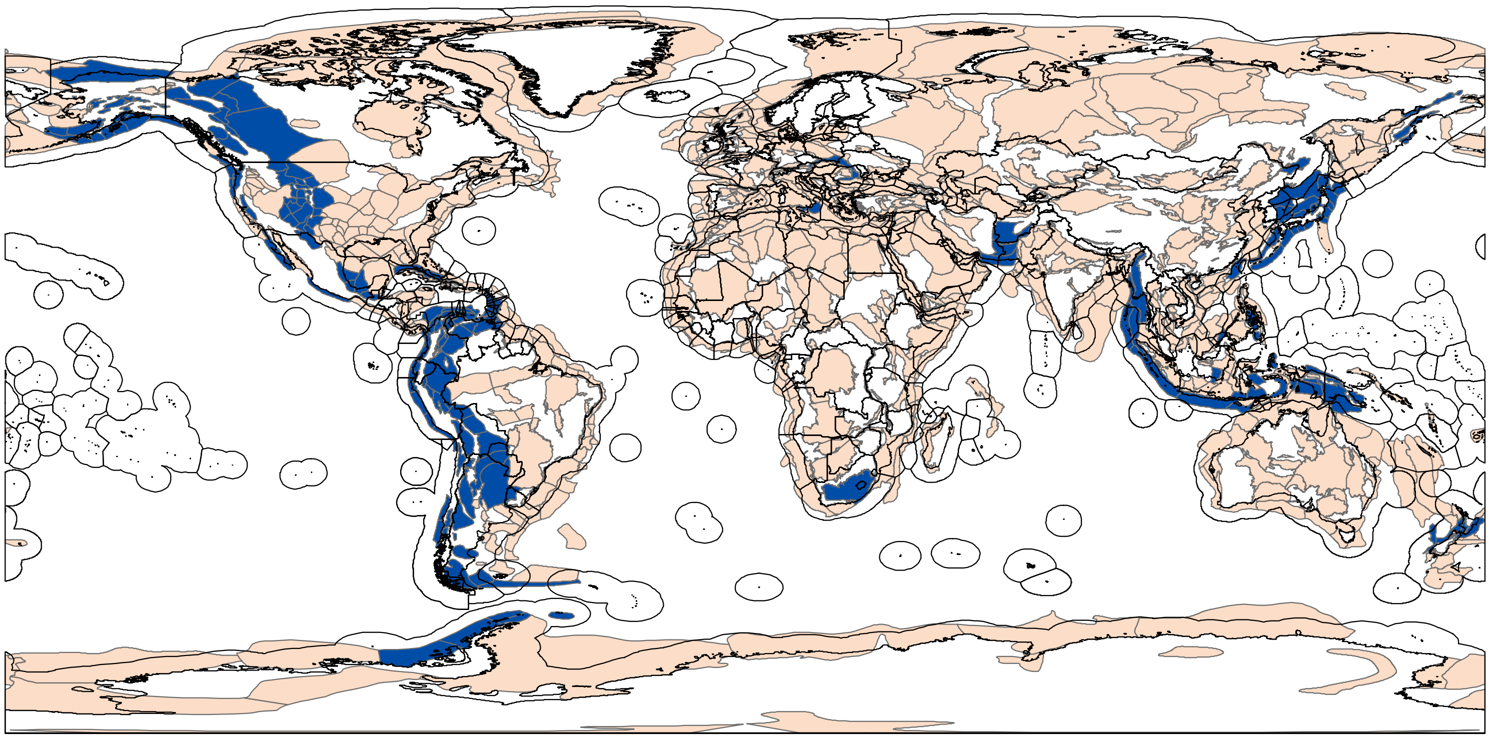

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.10: Basins: Convergent Ocean-Continent Tectonics, Thermal Subsidence

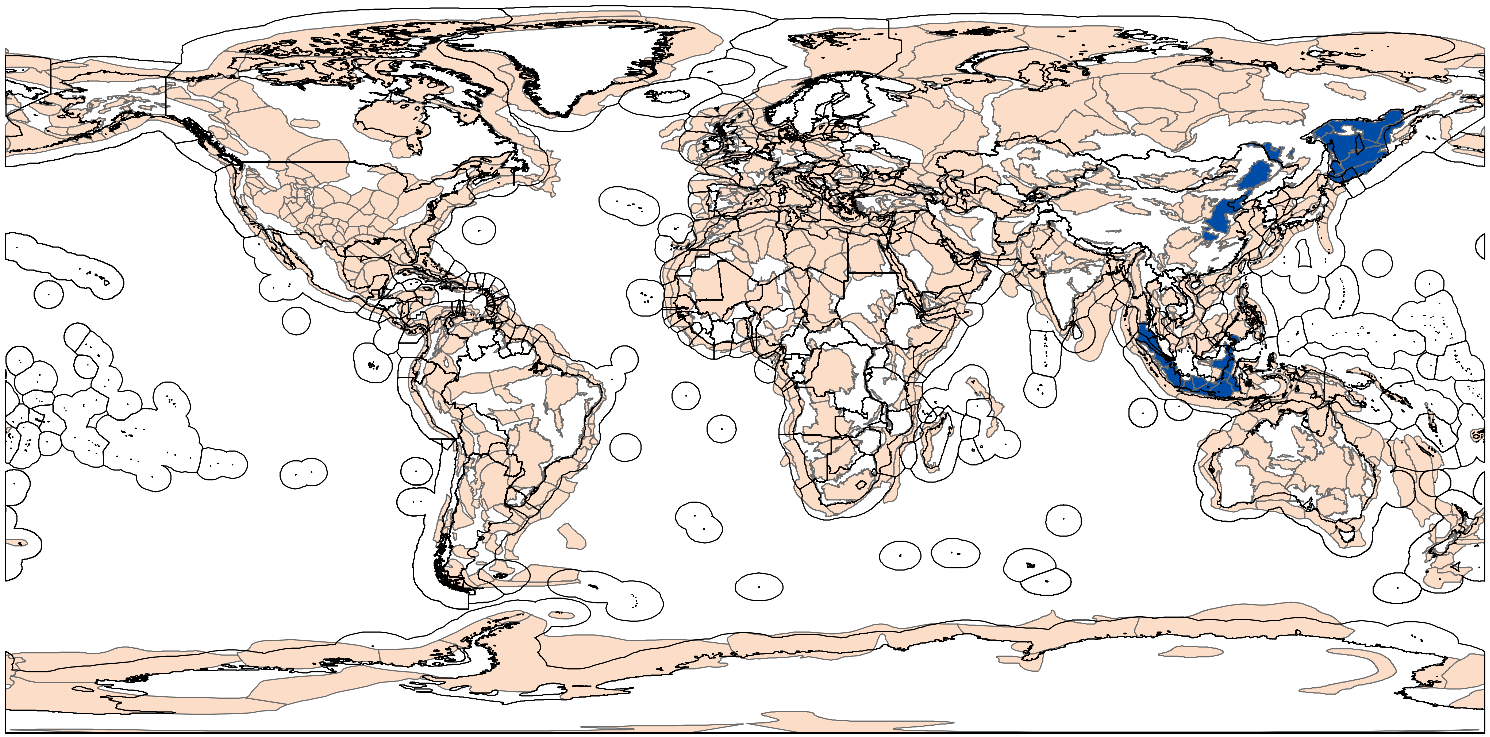

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.11: Basins: Convergent Ocean-Ocean Tectonics

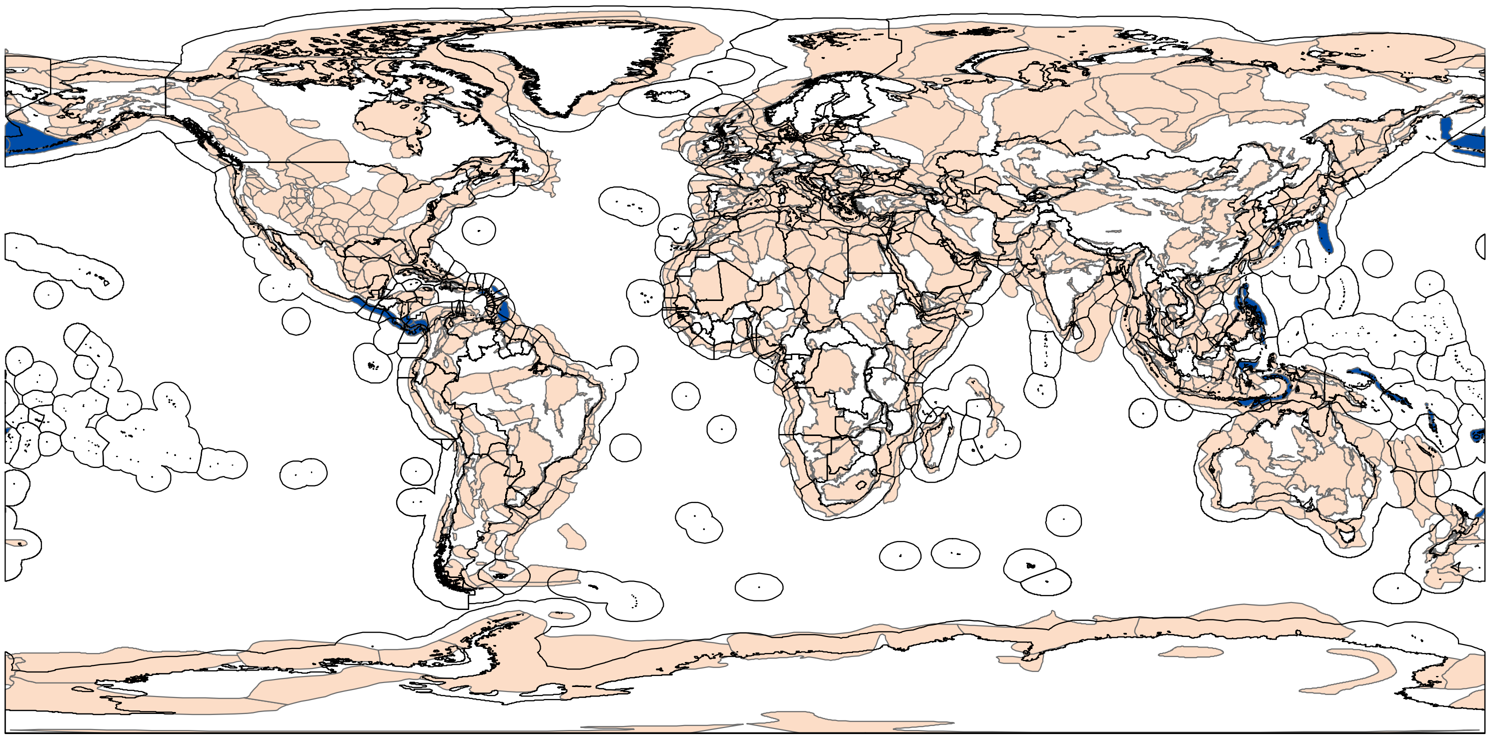

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.12: Basins: Divergent Tectonics, Mechanical Subsidence

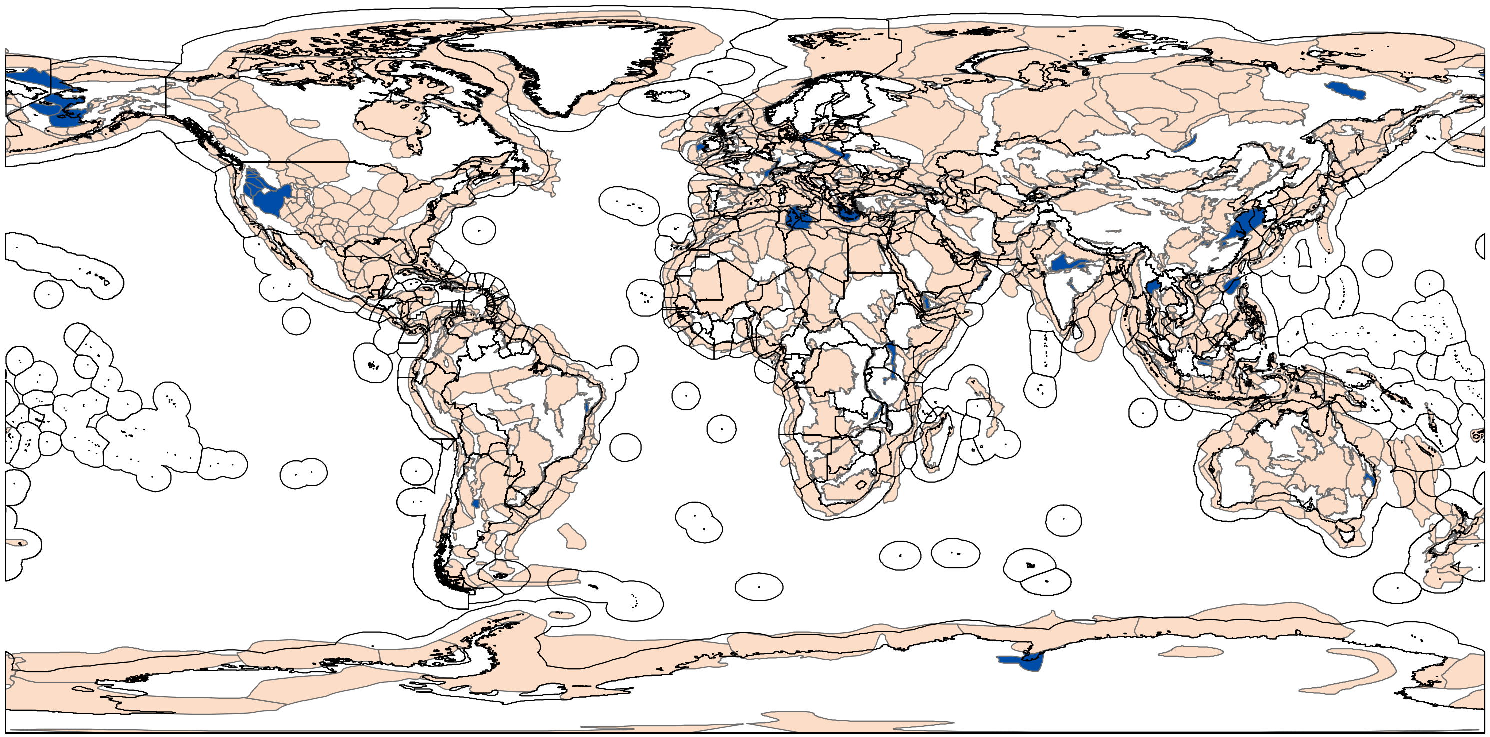

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.13: Basins: Divergent Tectonics, Thermal Subsidence

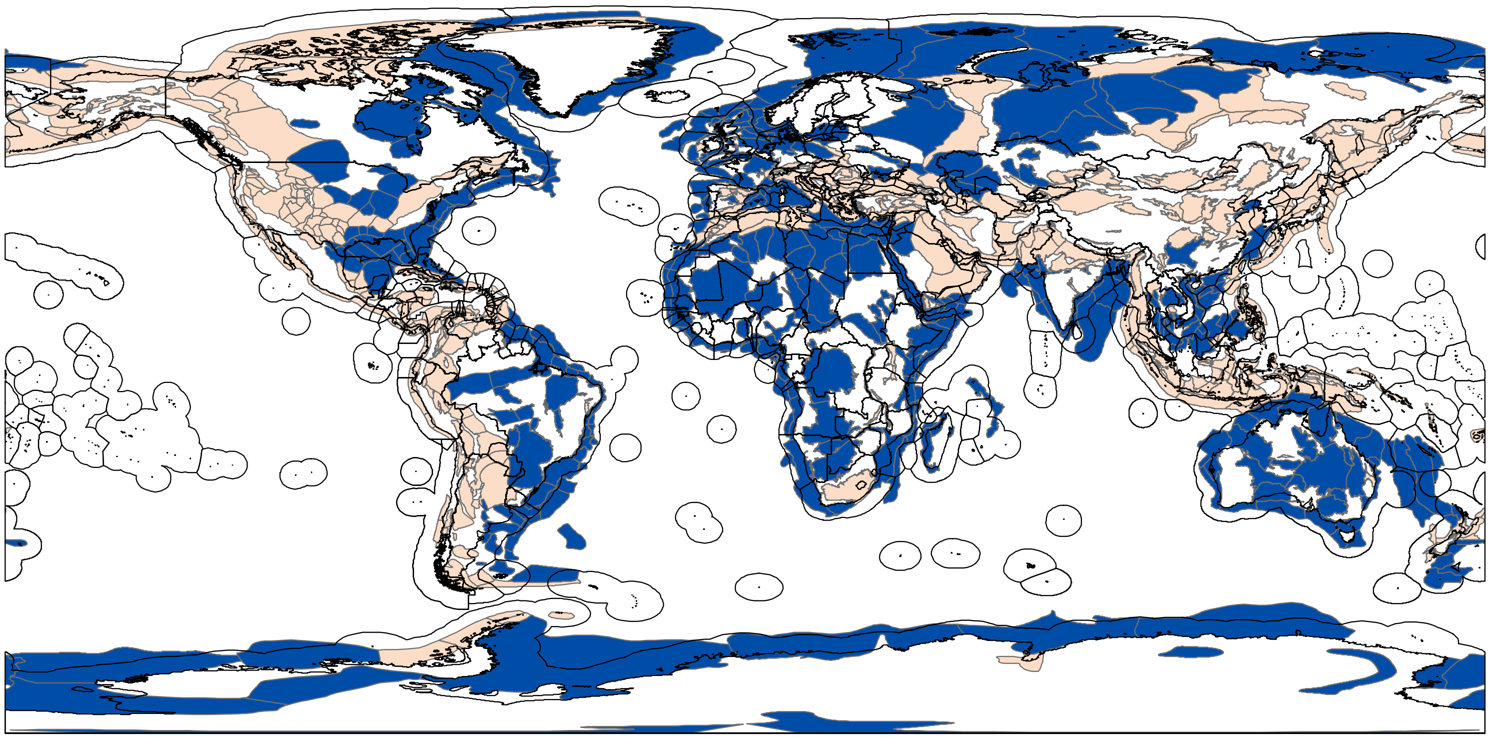

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.14: Basins: Wrench Tectonics, Mechanical Subsidence

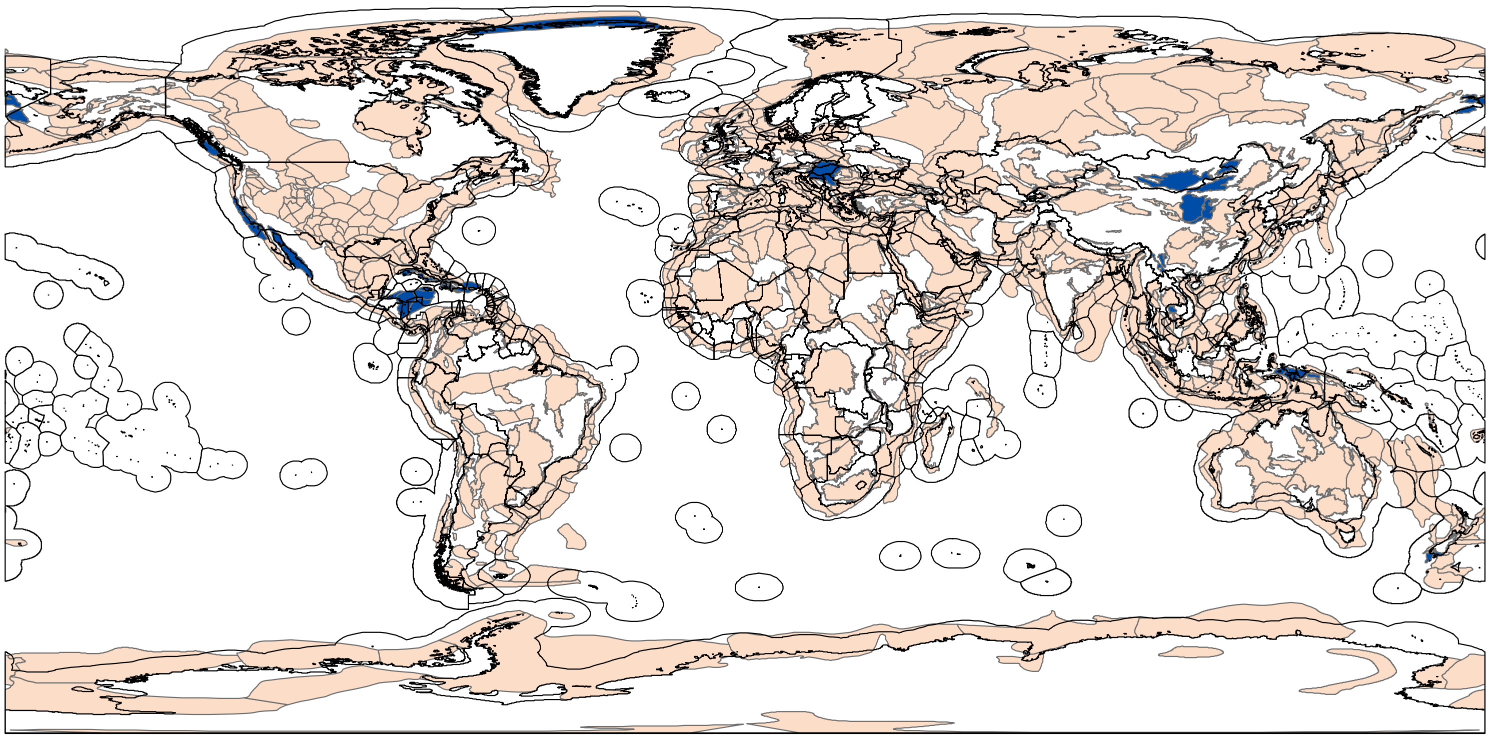

Source. Fugro Robertson, Ltd. (2013).

B.3.2 Basins grouped by final component of Fugro Tellus code

Figure B.3.15: Foreland Basins

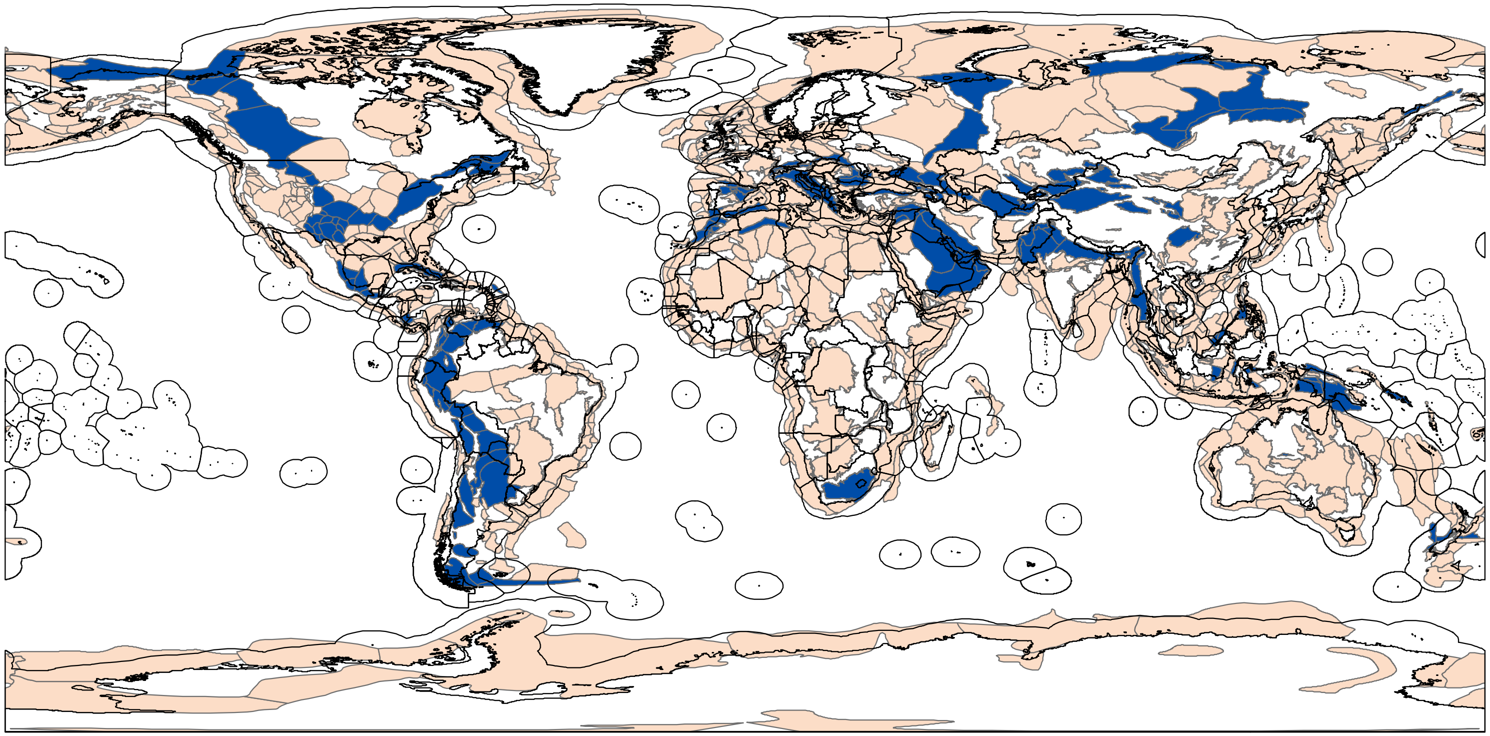

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.16: Fore-Arc Basins

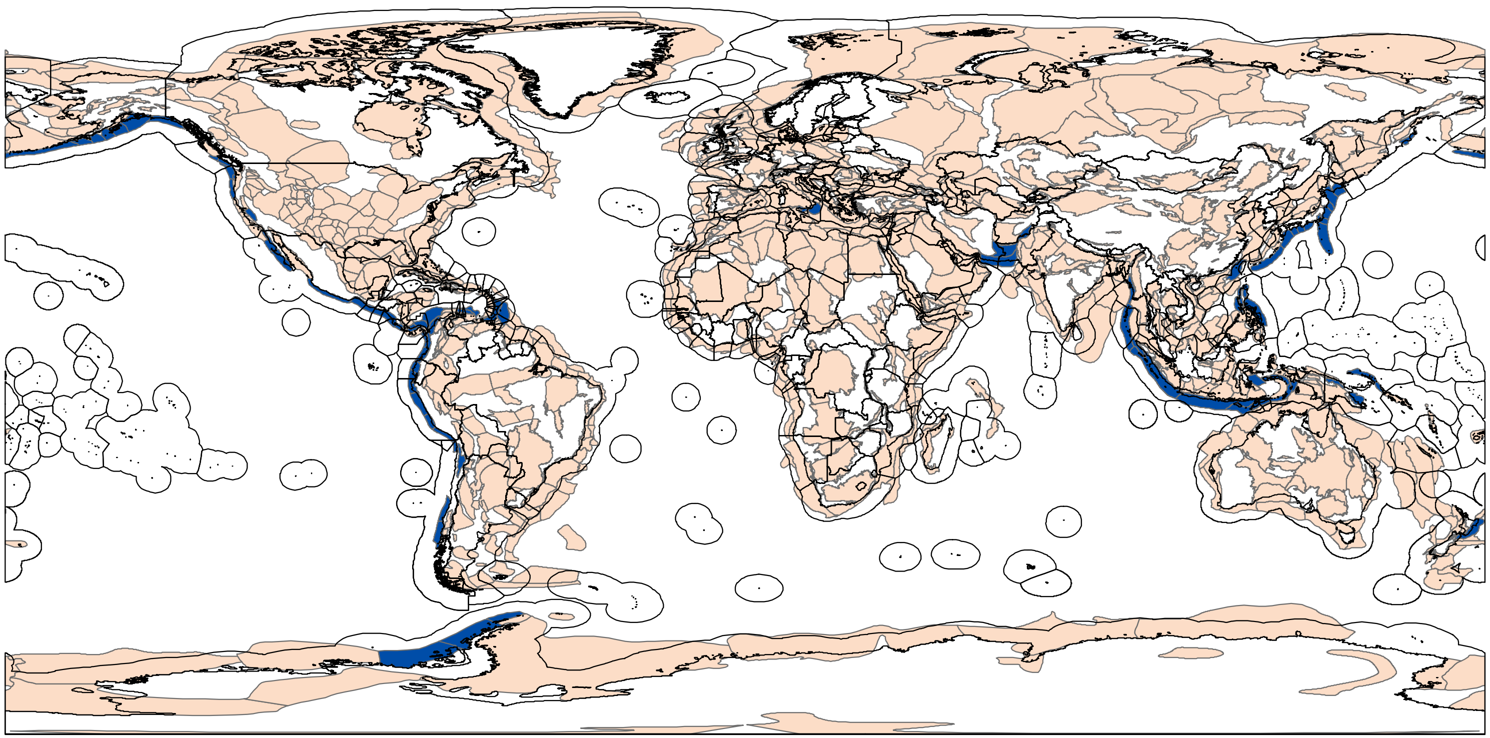

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.17: Extensional Basins

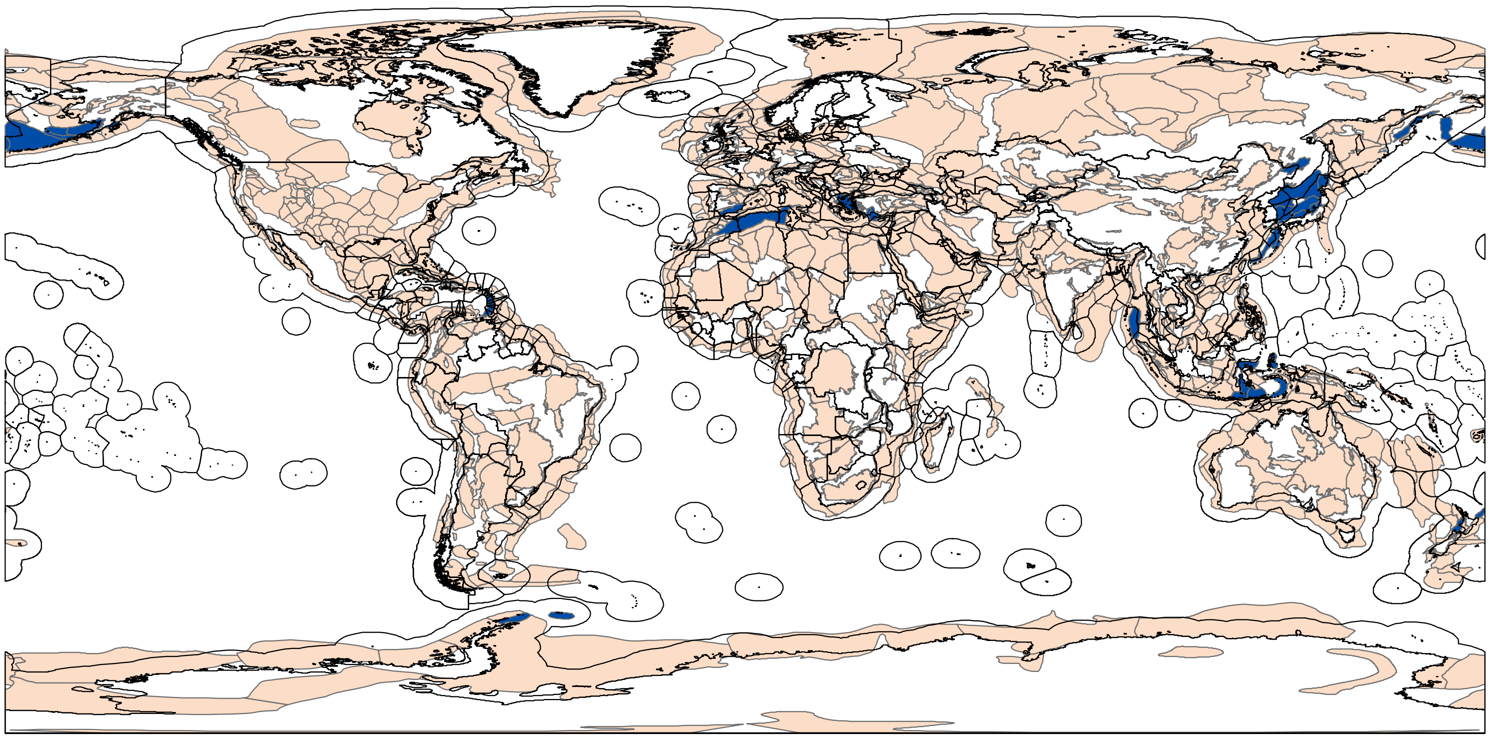

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.18: Convergent Sag Basins

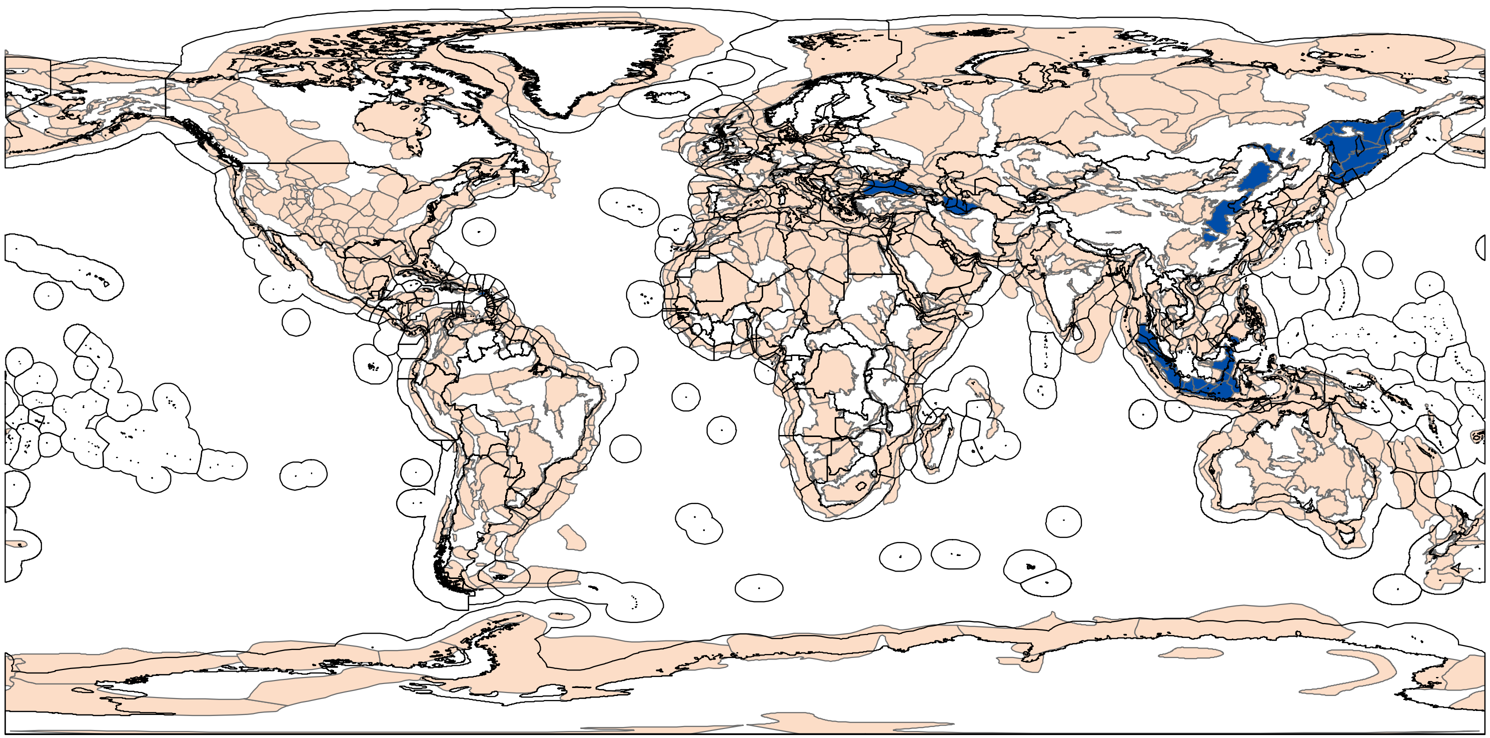

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.19: Convergent Wrench Basins

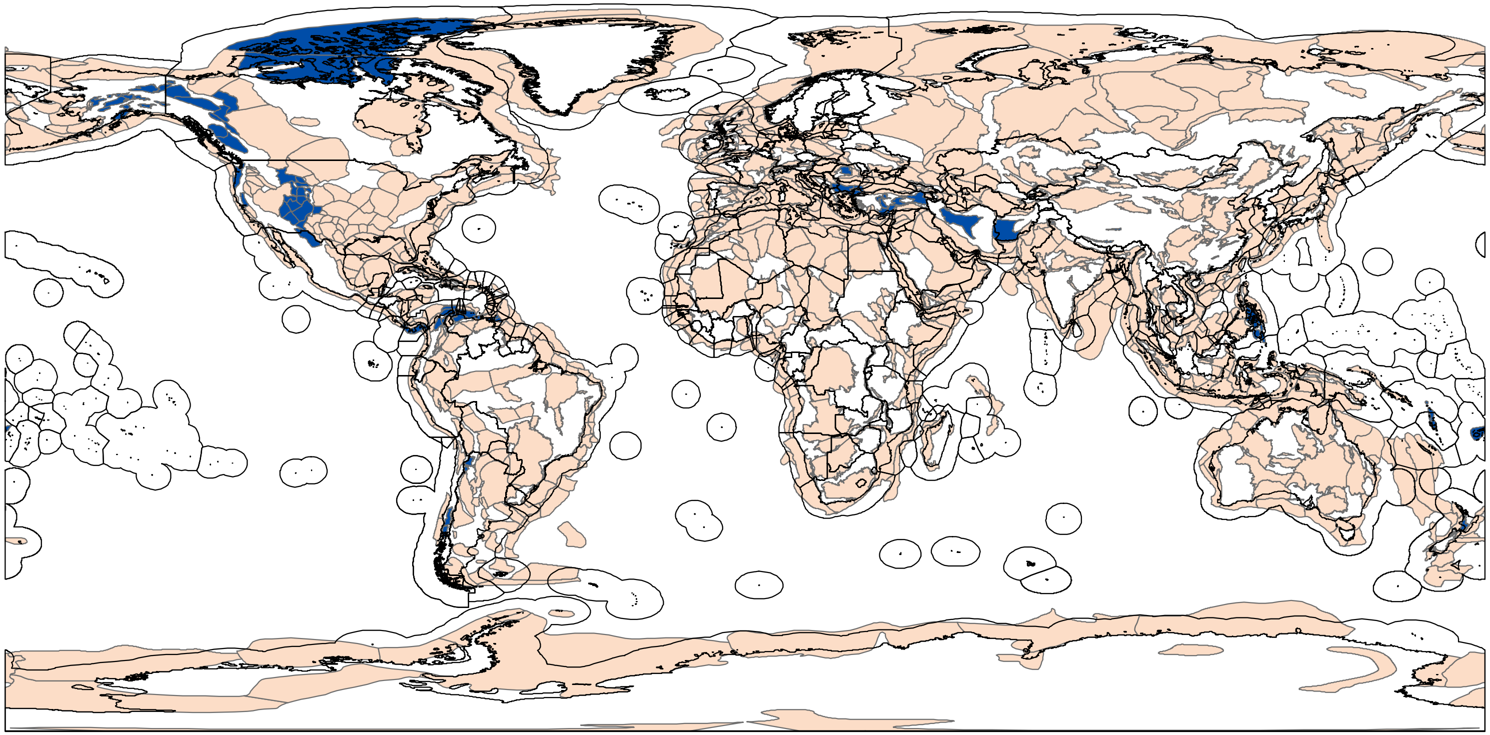

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.20: Rift Basins

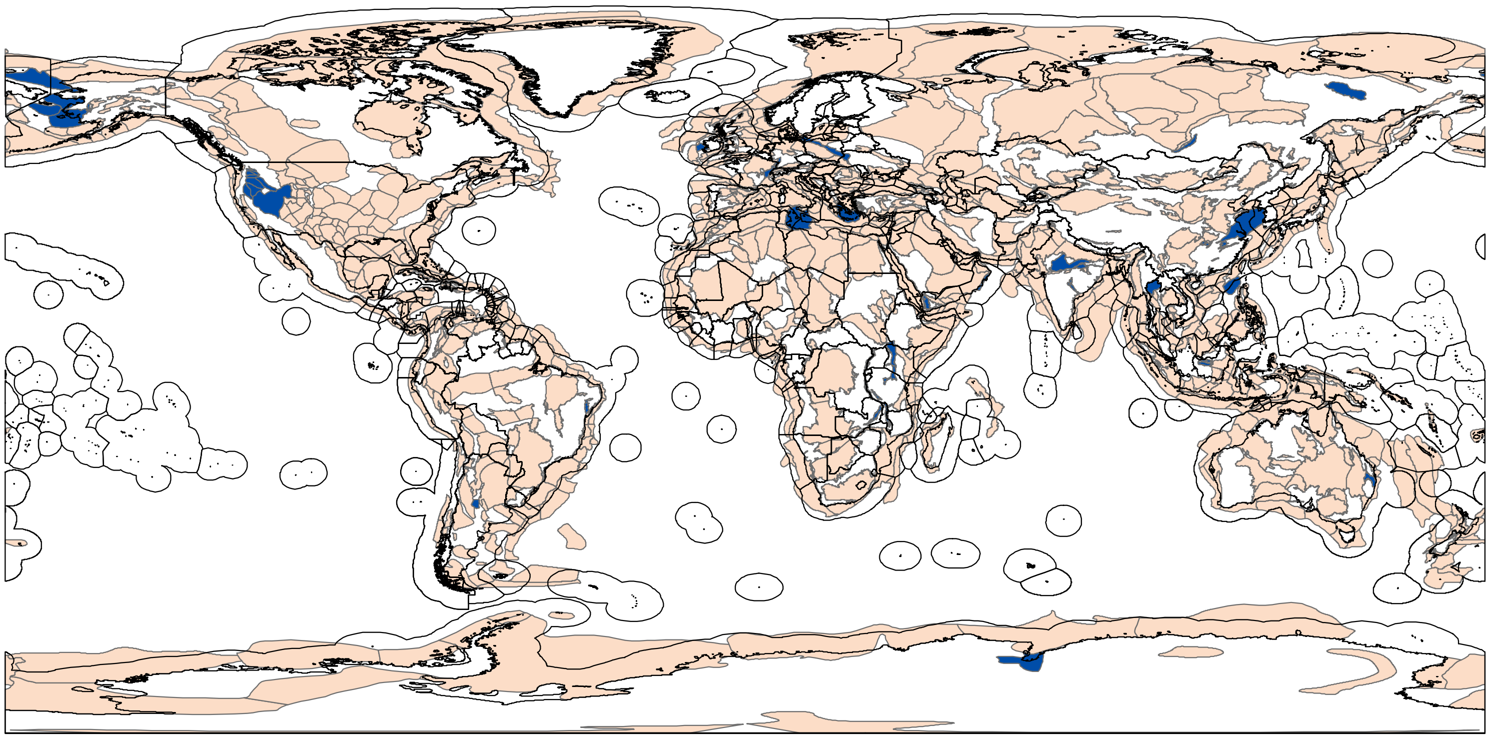

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.21: Intracratonic Sag Basins

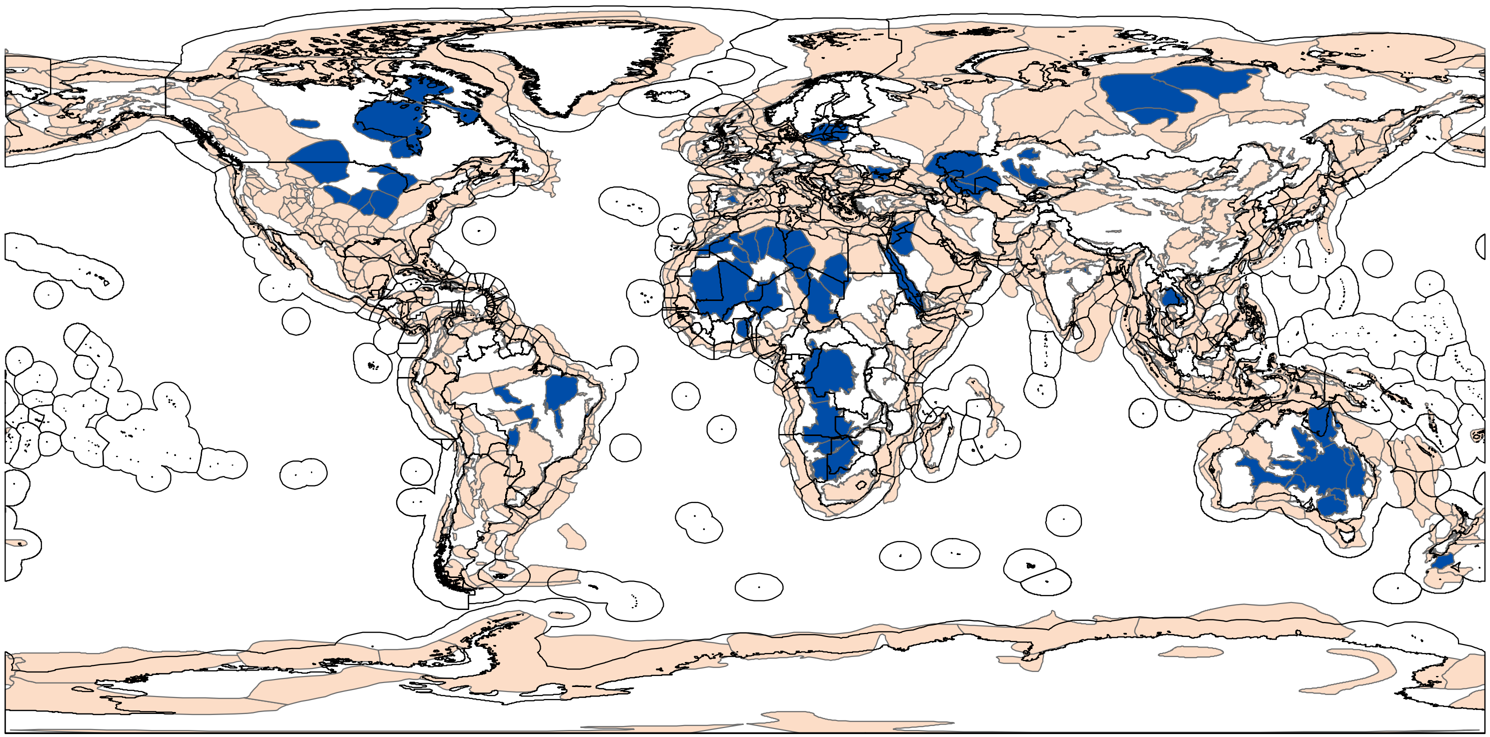

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.22: Post-Rift Sag Basins

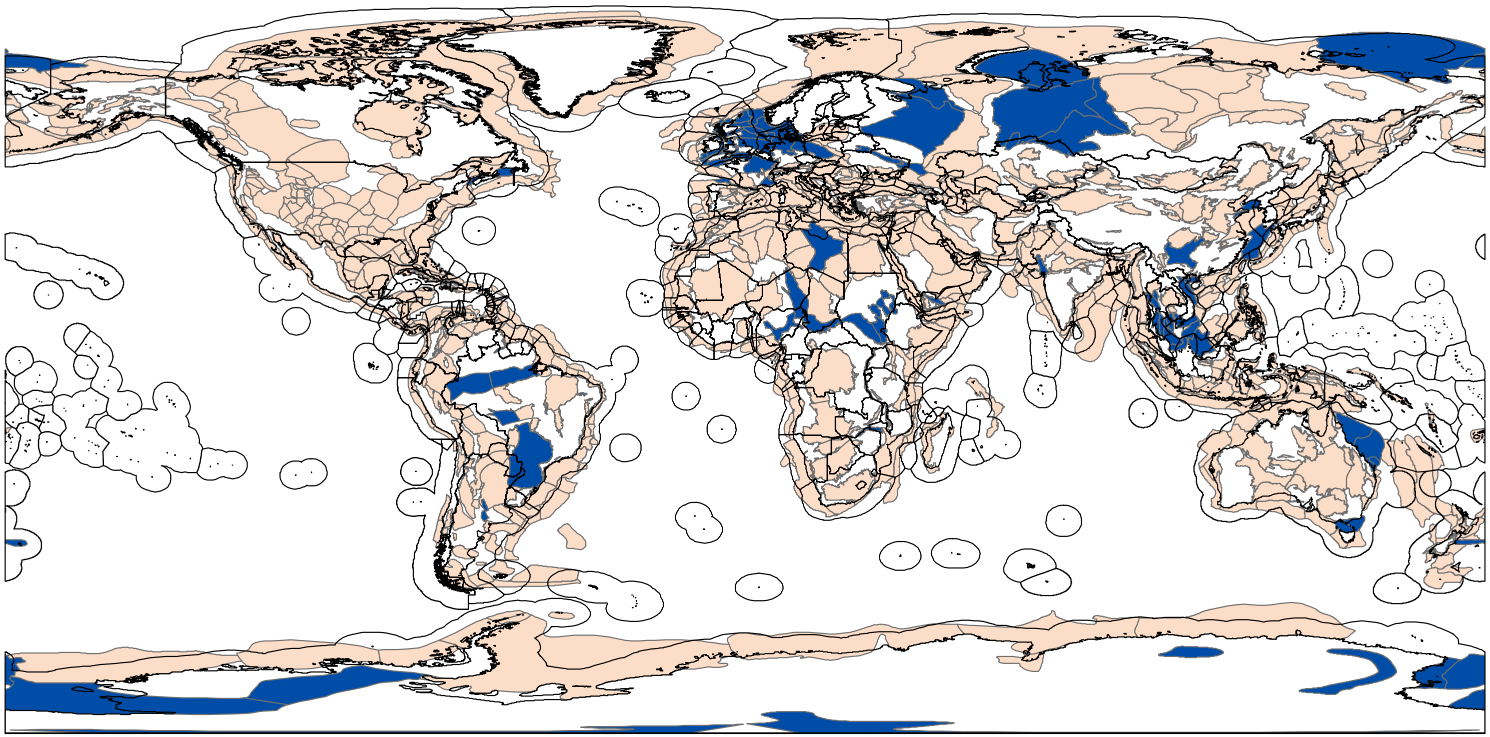

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.23: Passive Margin Basins

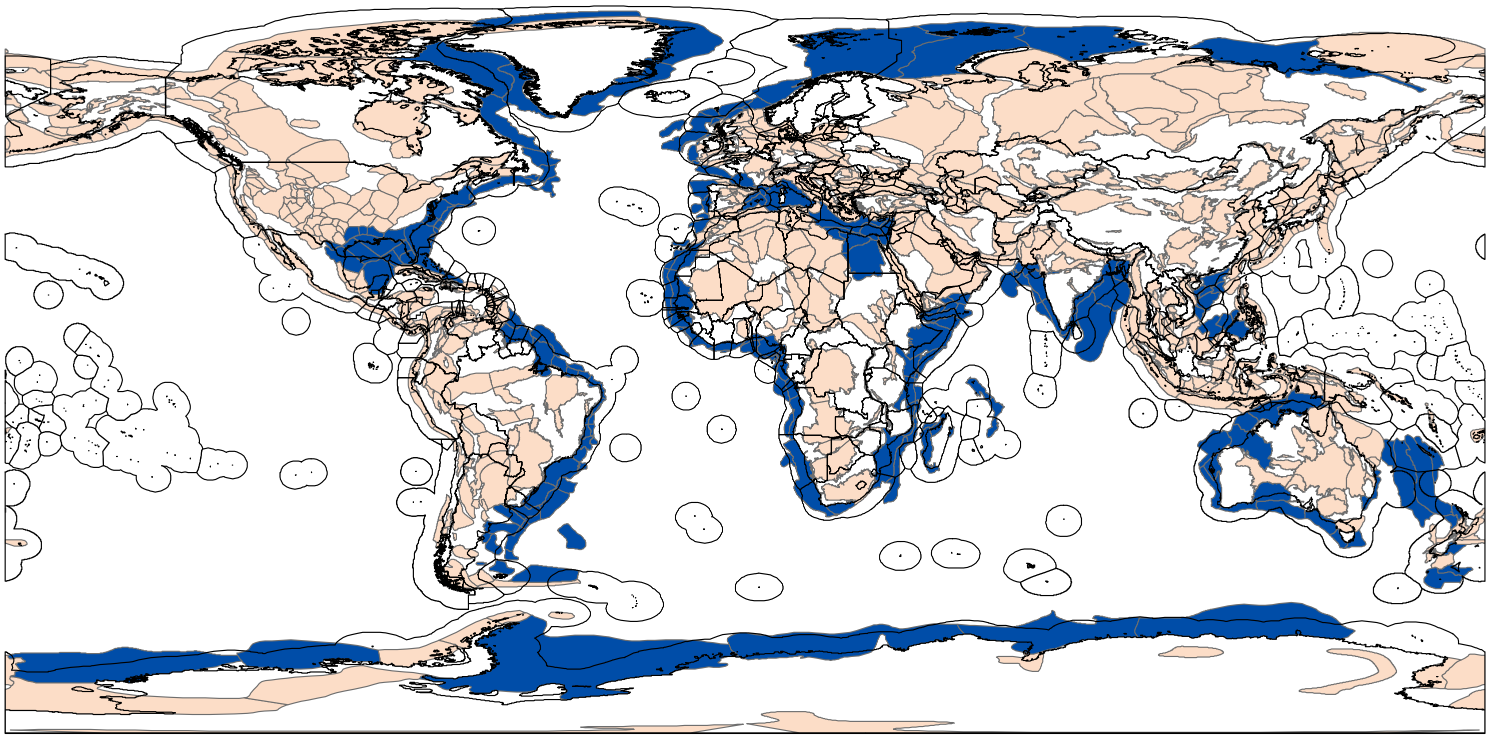

Source. Fugro Robertson, Ltd. (2013).

Figure B.3.24: Wrench Basins

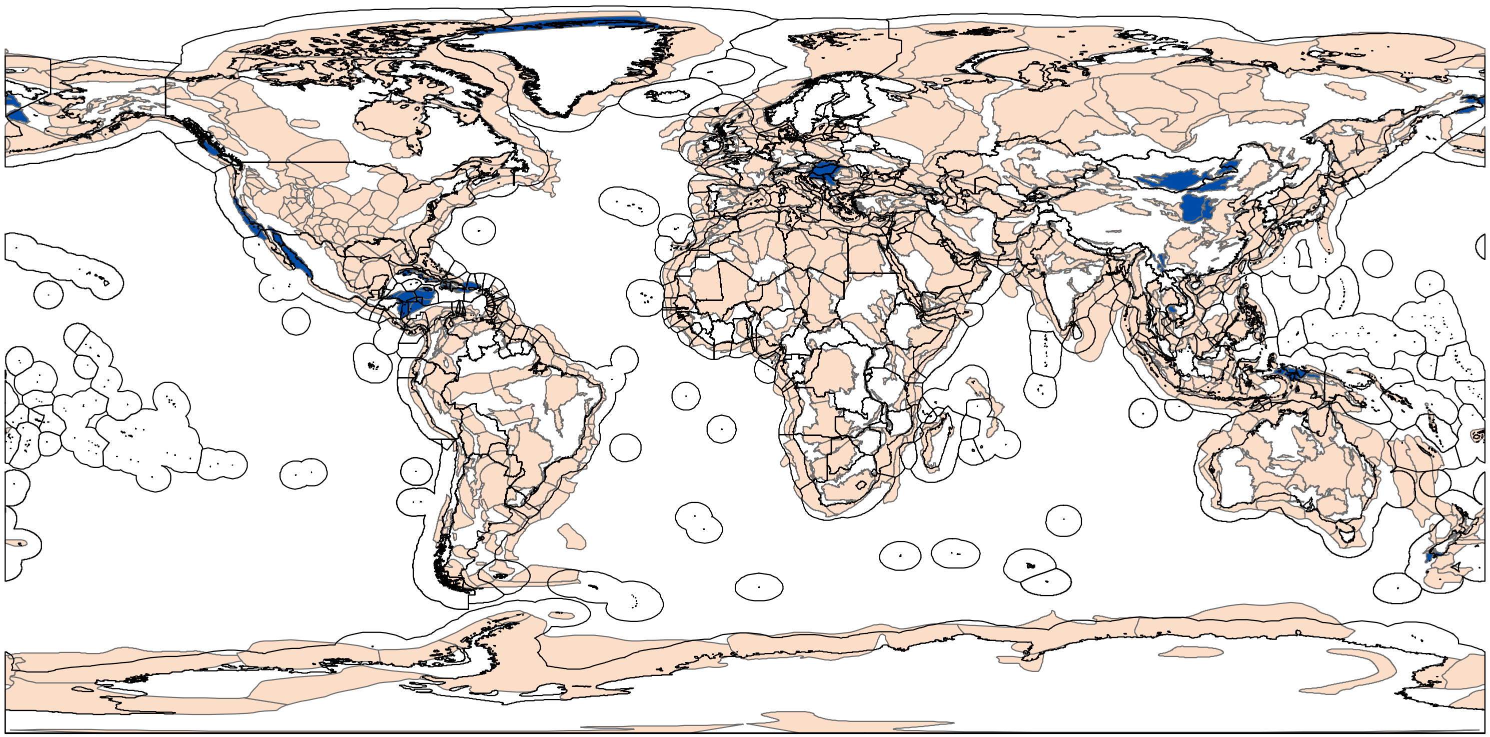

Source. Fugro Robertson, Ltd. (2013).

Appendix C

C.1 Regression Discontinuity Design

A major concern with the fixed-effects model is that time-varying unobservables may be correlated with both the outcomes and the mayor's age. For example, a change in voter preferences may drive changes in outcomes and in the age of the winning mayoral candidate. Alternatively, the age distribution of the mayoral candidate pool may change as a result of changes in policy outcomes, if the candidate's decision to run for office as a function of current policy differs by age.

In order to address these concerns, we could use a regression discontinuity (RD) design and focus on close elections between mayoral candidates of different ages. The running variable, X, would be defined as the older candidate's margin of victory, measured in percentage points. The main intuition is that, for each close election involving two candidates of different ages, whether the older candidate wins would be as good as randomly assigned. Let $D = 1$ if the older candidate wins, and zero otherwise. Adopting the potential outcomes framework, let Y_j denote the policy outcome in the event that the winning politician is j years old, where $j \in \{1, 2, \ldots, J\}$. In addition, let $A_D \in \{1, 2, \ldots, J\}$ be the age of the winning politician when the election outcome is D. The usual monotonicity condition is satisfied, because by construction $A_1 > A_0$ for all elections. Importantly, to apply this strategy, one would need to assume that $\mathbb{E}\left[Y_{A_D} \mid X = x\right]$ is continuous at $x = 0$ for $D \in \{0, 1\}$. The monotonicity and continuity conditions jointly suffice to identify the causal effect of the politician's age on policy outcomes (Hahn, Todd, and Van der Klaauw, 2001). In practice we would estimate the equations

$$\begin{aligned} A_{mt} &= \alpha + \delta D_{mt} + P(X_{mt}) + \nu_{mt} \\ Y_{mt} &= \gamma + \rho A_{mt} + P(X_{mt}) + \xi_{mt} \end{aligned} \tag{C.1.1}$$

using data from elections decided by a relatively narrow margin of victory.

Unfortunately, in our setting one of the assumptions required for identification is likely not satisfied: the continuity of potential outcomes. This assumption is not testable, but a way to provide supporting evidence is to show that a broad range of municipal and individual covariates are balanced at the threshold. While the balancedness of municipal covariates is typically satisfied in close elections, the strategy does not guarantee that individual covariates will be balanced, as noted by Brollo and Troiano (2016). In Table C.1.4 we verify that observable municipal characteristics are indeed balanced in close elections (Panel B), but we also show that many individual covariates are not balanced (Panel A).

We find that older mayors that win a close election are on average 11 years older than younger mayors that win a close election. However, they differ on other dimensions as well. Older mayors have more elected political experience, are less likely to be a woman, are more likely to be born in the city where they become mayor, and are less educated. It should be noted that the continuity assumption is not satisfied in this setting because age in our case is a "compounded treatment" (age being inherently correlated with other potential treatments at the individual level). In contrast, all of the municipal characteristics are balanced in close elections, including population, geography, income, and demographic structure. The unbalancedness of individual covariates implies that we cannot apply a standard regression discontinuity design in this setting. In a previous version of the paper we showed that our results are robust if we implement a regression discontinuity combined with a matching strategy, to alleviate the concerns of unbalancedness, and following Keele and Titiunik (2015), and those results are available upon request.

Table C.1.1: Variable Descriptions and Sources

Variable	Definition and measure	Sample	Source
Age	Age of the mayor, in years	1998–2013	IMI
Marg. of victory	Margin of victory of older candidate, in percentage points	1998–2013	IMI
Prior office	Equal to 1 if mayor had prior appointment to any city office	1998–2013	IMI
Exper. (elected)	Mayor's experience in elected city political office, in years	1998–2013	IMI
Exper. (unelected)	Mayor's experience in unelected city political office, in years	1998–2013	IMI
Woman	Equal to 1 if the mayor is a woman	1998–2013	IMI
Born locally	Equal to 1 if mayor was born in the city that s/he governs	1998–2013	IMI
High school	Equal to 1 if mayor has a high-school degree	1998–2013	IMI
College	Equal to 1 if mayor has a college (bachelor's) degree	1998–2013	IMI
Center-right	Equal to 1 if the mayor is a member of a center-right party	1998–2013	IMI
# Rivals	Number of rival candidates	1998–2013	IMI
Term limited	Equal to 1 if mayor has a binding term limit	1998–2013	IMI
Pop.	Population	2001	Census
Pop. dens.	Population density, measured as the number of people per square kilometer	2005	SAIM
Prop. active pop.	Proportion of population that is over 15 years old and either has a job or is looking for one	2005	SAIM
Elderly	Elderly index, equal to (population over 65 years old)/(population under 14 years old)	2005	SAIM
Fam. size	Average family size	2005	SAIM
Prod. p.c.	Number of production units per capita	2005	SAIM
Employment	Proportion of population that is employed	2005	SAIM
Income p.c.	Disposable income per capita in euro	2005	SAIM
Altitude	Altitude of the city, in meters	2001	SAIM
North	Equals 1 if municipality is in the North	2001	SAIM
Central	Equals 1 if municipality is in the Center	2001	SAIM
South	Equals 1 if municipality is in the South (includes Sicily and Sardinia)	2001	SAIM

Notes: IMI stands for Italian Ministry of the Interior. SAIM stands for Statistical Atlas of Italian Municipalities. Census stands for the Italian census.

Table C.1.2: Variable Description and Sources

Variable	Definition and measure	Sample	Source
Ran again	Incumbent mayor decided to run for office again (non-term-limited mayors only)	1998–2013	IMI
Reelected	Incumbent mayor was reelected (non-term-limited mayors only)	1998–2013	IMI
Prov. admin.	Mayor held a position in provincial administration within 5 years of first mayoral mandate	1998–2014	IMI
Region. admin.	Mayor held a position in regional administration within 5 years of first mayoral term	1998–2014	IMI
Nation. admin.	Mayor held a position in national administration within 5 years of first mayoral term	1998–2014	IMI
Avg. rev.	Average total collected revenue over the years of the mayor's term (€ per capita)	1998–2013	IMI
Avg. cap. exp.	Average implemented capital expenditure over the years of the mayor's term (€ per capita)	1998–2013	IMI
Avg. curr. exp.	Average implemented current expenditure over the years of the mayor's term (€ per capita)	1998–2013	IMI
Δ *Revenue*	Total collected revenue in pre-election year minus average revenue in previous years of term (€ per capita)	1998–2013	IMI
Δ *Cap. exp.*	Implemented capital expenditure in pre-election year minus average cap. expend. in previous years of term (€ p.c.)	1998–2013	IMI
Δ *Curr. exp.*	Implemented current expenditure in pre-election year minus average curr. expend. in previous years of term (€ p.c.)	1998–2013	IMI
Speed of public good provision	Average ratio of paid to committed current expenditure over the years of the mayor's term	1998–2013	IMI
House price	Average price of residential housing over term (€ per square meter)	2002–2011	IRS

Notes: IMI stands for Italian Ministry of the Interior. SAIM stands for Statistical Atlas of Italian Municipalities. IRS stands for Italian Revenue Service. All budget variables exclude the election year.

Table C.1.3: Covariate Balance Following Close Elections

	Older Candidate Won	Younger Candidate Won	Difference	p-value
Mayor				
Age	55.34	43.44	11.91^{***}	0.000
Previously held any city office	0.66	0.63	0.03	0.245
City political experience (elected)	5.17	4.18	0.99^{***}	0.000
City political experience (unelected)	0.62	0.70	−0.08	0.450
Woman	0.11	0.15	-0.04^{**}	0.026
Born locally	0.52	0.39	0.13^{***}	0.000
High school degree	0.87	0.93	-0.06^{***}	0.000
College degree	0.43	0.56	-0.13^{***}	0.000
Center-right party	0.10	0.09	0.00	0.873
Number of rivals on ballot	2.05	1.93	0.12^{*}	0.089
Municipality				
Population (2001)	6685.76	5453.53	1232.23	0.182
Population per square km (2005)	321.09	286.62	34.47	0.238
Active pop. / total pop. (2005)	0.41	0.41	0.00	0.930
Elderly index (2005)	1.78	1.78	−0.00	0.965
Family size (2005)	2.50	2.50	0.01	0.594
Production units per capita (2005)	0.08	0.07	0.00^{*}	0.072
Employed / total pop. (2005)	0.27	0.25	0.02^{**}	0.024
Income per capita (2005 €)	13263.34	13217.52	45.82	0.783
Altitude (meters)	455.06	482.74	−27.68	0.230
North	0.52	0.50	0.02	0.390
Central	0.11	0.13	−0.02	0.295
South	0.36	0.37	−0.00	0.861
Observations	792	754		

Notes: The sample is restricted to non-term-limited mayoral terms following an election decided by a margin of two percentage points or fewer. The first two columns report averages from elections decided by a margin of victory of less than two percentage points. The third column reports the difference of the averages, and the fourth column reports the p-value corresponding to the two-sided test of equality of the averages. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.4: Discontinuities in Mayor and Municipality Characteristics

Panel A: Mayor Characteristics

	Age	Prior office	Exper. (elected)	Exper. (unelected)	Woman	Born locally
RD	11.2^{***}	0.042^{**}	0.98^{***}	-0.060	-0.036^{***}	0.15^{***}
	(0.32)	(0.017)	(0.23)	(0.074)	(0.012)	(0.019)
Bandwidth	15.4	14.6	11.6	12.1	15.3	14.2
Observations	10789	10371	8537	8814	10766	10123

	High school	College	Center-right	# Rivals	Term limited
RD	-0.069^{***}	-0.13^{***}	-0.0042	0.0022	0.027
	(0.012)	(0.019)	(0.011)	(0.058)	(0.018)
Bandwidth	12.7	13.3	14.3	9.94	10.4
Observations	9202	9571	10159	7511	7829

Panel B: Municipality Characteristics

	Pop.	Pop. dens.	Prop active pop.	Elderly	Fam. size	Prod. p.c.
RD	-413.1	18.7	0.0013	0.057	-0.012	0.00057
	(1466.8)	(24.4)	(0.0028)	(0.054)	(0.012)	(0.0010)
Bandwidth	11.5	13.4	8.39	12.6	12.4	11.5
Observations	8487	9603	6451	9137	9027	8448

	Employment	Income p.c.	Altitude	North	Central	South
RD	0.0063	96.2	-34.2^{*}	0.0045	0.0030	-0.0098
	(0.0086)	(149.3)	(19.5)	(0.020)	(0.013)	(0.021)
Bandwidth	9.82	8.23	9.83	12.8	13.8	11.2
Observations	7440	6348	7441	9294	9903	8313

Notes: Above are sharp regression discontinuity estimates using the older candidate's margin of victory as the running variable. Estimates are obtained by local linear regression using a uniform kernel and the Calonico, Cattaneo, and Titiunik (2014) bandwidth selector. Heteroskedasticity-robust standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.5: The Effect of the Mayor's Age on Average Budget Outcomes (Municipality-Specific Linear Trends Removed)

Panel A: Fixed-Effects Estimates

	Avg. revenue		Avg. capital expend.		Avg. current expend.	
	(1)	(2)	(3)	(4)	(5)	(6)
Age	4.91	10.45	−0.02	−0.11	4.55	9.51
	(8.53)	(10.92)	(0.44)	(0.54)	(7.59)	(9.57)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	6581	6581	6581	6581	6581	6581
Observations	10239	10239	10239	10239	10239	10239

Panel B: Fixed-Effects Estimates, Close Elections

	Avg. revenue		Avg. capital expend.		Avg. current expend.	
	(1)	(2)	(3)	(4)	(5)	(6)
Age	3.99	4.49	0.20	0.18	4.73	5.79
	(4.06)	(4.57)	(0.33)	(0.36)	(4.38)	(4.93)
Controls	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*
Municipalities	1902	1902	1902	1902	1902	1902
Observations	2123	2123	2123	2123	2123	2123

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. All regressions include municipality fixed effects and year-of-election effects. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors in the fixed-effects specification are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.6: The Effect of the Mayor's Age on House Prices and the Speed of Public Good Provision (Municipality-Specific Linear Trends Removed)

Panel A: Fixed-Effects Estimates

	Avg. house price		Speed of public good provision	
	(1)	(2)	(3)	(4)
Age	−0.19	−0.20	−0.000035	−0.000041
	(0.18)	(0.19)	(0.000082)	(0.000078)
Controls	*No*	*Yes*	*No*	*Yes*
Municipalities	5232	5232	6581	6581
Observations	5995	5995	10239	10239

Panel B: Fixed-Effects Estimates, Close Elections

	Avg. house price		Speed of public good provision	
	(1)	(2)	(3)	(4)
Age	−0.71	0.23	−0.000099	0.00026
	(0.44)	(0.49)	(0.00052)	(0.00032)
Controls	*No*	*Yes*	*No*	*Yes*
Municipalities	1173	1173	1902	1902
Observations	1228	1228	2123	2123

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. Results are reported with and without controls for prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and whether the mayor was born in the city s/he governs. Standard errors are robust to heteroskedasticity and clustering at the level of municipality. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.7: The Effect of the Mayor's Age on Political Budget Cycles (Municipality-Specific Linear Trends Removed)

Panel A: Dynamic Panel Estimates

	Revenue per capita			Capital expenditure per capita			Current expenditure per capita		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Year before election	13.68***	20.14	42.55*	7.65***	16.94**	35.24***	−0.15	−5.39	−7.80
	(3.58)	(12.55)	(24.26)	(1.94)	(7.24)	(13.46)	(1.04)	(3.45)	(7.12)
Old	2.07			1.37			−0.21		
	(1.87)			(0.87)			(0.73)		
Old × Year before election	−5.92			−6.23**			2.30		
	(4.89)			(2.73)			(1.46)		
Age		0.07	0.24		0.02	0.15		0.00	−0.01
		(0.09)	(0.19)		(0.05)	(0.09)		(0.04)	(0.08)
Age × Year before election		−0.19	−0.66		−0.25*	−0.63**		0.14*	0.18
		(0.25)	(0.50)		(0.15)	(0.28)		(0.07)	(0.15)
AR(2) test p-value	0.838	0.838	0.827	0.007	0.003	0.008	0.019	0.015	0.019
Municipalities	7064	7229	7064	7064	7229	7064	7064	7229	7064
Observations	60721	63018	60721	60838	63149	60838	60837	63148	60837

Panel B: Dynamic Panel Estimates, Close Elections

	Revenue per capita			Capital expenditure per capita			Current expenditure per capita		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Year before election	15.64**	33.19	37.97	11.04***	35.50***	57.61***	−1.07	−8.88	−26.76
	(6.87)	(22.38)	(42.52)	(3.31)	(13.17)	(20.81)	(2.39)	(6.99)	(17.96)
Old	−0.54			2.16			−1.27		
	(4.35)			(1.72)			(2.36)		
Old × Year before election	−4.33			−10.95**			6.24*		
	(8.70)			(4.28)			(3.76)		
Age		−0.07	−0.05		0.02	0.24		−0.09	−0.13
		(0.21)	(0.44)		(0.08)	(0.17)		(0.11)	(0.24)
Age × Year before election		−0.41	−0.50		−0.62**	−1.07**		0.22	0.59
		(0.44)	(0.86)		(0.25)	(0.42)		(0.14)	(0.37)
AR(2) test p-value	0.315	0.369	0.314	0.738	0.735	0.753	0.050	0.052	0.051
Municipalities	2497	2497	2497	2499	2499	2499	2499	2499	2499
Observations	12612	12612	12612	12636	12636	12636	12636	12636	12636

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. Both panels include observations following all mayoral elections. The results in Panel A are from the sample for which the mayor is not term limited. The results in Panel B are from the sample for which the mayor is not term limited and the most recent election was decided by a vote margin of five percentage points or less. In columns 1, 4, and 7, "Old" is an indicator for the mayor being older than the electoral opponent coming in second place. "Old" is treated as strictly exogenous. In the remaining columns, "Age" is the age of the mayor. In columns 2, 5, and 8, "Age" is treated as strictly exogenous. In columns 3, 6, and 9, "Age" and its interaction are treated as endogenous and instrumented with "Old" and its interaction. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.8: Mayor Age, Political Budget Cycles, and Population Thresholds for Mayor Wage

Dynamic Panel Estimates	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	78.96***	30.67***	12.87**
	(16.92)	(6.74)	(5.93)
Yr pre × Pop. ∈ [1K, 3K)	−68.02***	−23.92***	−19.08***
	(19.26)	(7.14)	(6.62)
Yr pre × Pop. ∈ [3K, 5K)	−95.60***	−38.30***	−21.05***
	(20.40)	(7.42)	(7.76)
Yr pre × Pop. ∈ [5K, 10K)	−86.36***	−29.37***	−20.63***
	(20.16)	(7.65)	(6.95)
Yr pre × Pop. ∈ [10K, 30K)	−100.63***	−27.91***	−19.40**
	(21.02)	(7.55)	(8.65)
Yr pre × Pop. ≥ 30K	−75.45***	−33.82***	−20.15***
	(21.06)	(8.30)	(7.26)
Old	123.94***	23.95***	58.52***
	(17.75)	(5.21)	(8.67)
Old × Pop. ∈ [1K, 3K)	−147.42***	−25.74***	−69.90***
	(20.72)	(5.75)	(10.06)
Old × Pop. ∈ [3K, 5K)	−160.00***	−29.91***	−70.56***
	(24.35)	(6.89)	(11.61)
Old × Pop. ∈ [5K, 10K)	−172.16***	−33.04***	−81.85***
	(25.57)	(7.37)	(12.31)
Old × Pop. ∈ [10K, 30K)	−169.18***	−35.04***	−82.38***
	(26.24)	(7.84)	(12.64)
Old × Pop. ≥ 30K	−177.16***	−32.85***	−84.10***
	(29.56)	(8.32)	(14.73)
Old × Yr pre	−103.12***	−25.25***	−30.15***
	(23.01)	(9.66)	(8.00)
Old × Yr pre × Pop. ∈ [1K, 3K)	116.82***	16.21	41.46***
	(26.29)	(10.82)	(9.16)
Old × Yr pre × Pop. ∈ [3K, 5K)	134.94***	30.81***	40.57***
	(27.56)	(11.14)	(10.68)
Old × Yr pre × Pop. ∈ [5K, 10K)	120.54***	27.32**	44.26***
	(26.88)	(10.77)	(9.56)
Old × Yr pre × Pop. ∈ [10K, 30K)	121.69***	24.20**	40.07***
	(28.60)	(10.99)	(11.50)
Old × Yr pre × Pop. ≥ 30K	131.70***	40.76***	46.87***
	(36.07)	(15.08)	(12.10)
AR(2) test p-value	0.000	0.000	0.001
Municipalities	7063	7063	7063
Observations	60687	60802	60801

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.9: Mayor Age, Political Budget Cycles, and Population Thresholds for Mayor Wage

Dynamic Panel Estimates, Close Elections

	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	127.17***	54.99***	24.43*
	(42.79)	(14.01)	(14.10)
Yr pre × Pop. ∈ [1K, 3K)	−130.05***	−50.71***	−23.63
	(43.95)	(14.26)	(14.97)
Yr pre × Pop. ∈ [3K, 5K)	−132.89***	−54.97***	−32.23**
	(46.94)	(14.89)	(16.29)
Yr pre × Pop. ∈ [5K, 10K)	−138.78***	−57.71***	−28.14
	(51.89)	(15.18)	(19.21)
Yr pre × Pop. ∈ [10K, 30K)	−166.09***	−49.45***	−37.70*
	(53.23)	(16.43)	(21.23)
Yr pre × Pop. ≥ 30K	−63.17	−45.97**	−7.64
	(60.44)	(18.32)	(21.92)
Old	96.59*	24.97**	65.58***
	(51.44)	(9.96)	(23.27)
Old × Pop. ∈ [1K, 3K)	−91.28*	−22.71**	−63.68***
	(55.01)	(10.34)	(24.02)
Old × Pop. ∈ [3K, 5K)	−114.15*	−28.97**	−71.63***
	(62.61)	(11.71)	(26.81)
Old × Pop. ∈ [5K, 10K)	−142.99**	−35.41***	−90.35***
	(66.48)	(12.70)	(30.72)
Old × Pop. ∈ [10K, 30K)	−126.00*	−37.18***	−83.76***
	(68.98)	(13.15)	(30.54)
Old × Pop. ≥ 30K	−113.35*	−31.42**	−65.82***
	(60.32)	(13.81)	(23.13)
Old × Yr pre	−191.48***	−58.89***	−63.55***
	(54.22)	(17.65)	(20.78)
Old × Yr pre × Pop. ∈ [1K, 3K)	211.52***	54.55***	65.90***
	(59.23)	(19.46)	(22.26)
Old × Yr pre × Pop. ∈ [3K, 5K)	209.65***	54.63***	83.13***
	(59.95)	(20.55)	(23.21)
Old × Yr pre × Pop. ∈ [5K, 10K)	221.97***	66.94***	77.39***
	(62.13)	(19.40)	(25.51)
Old × Yr pre × Pop. ∈ [10K, 30K)	243.65***	54.60***	87.34***
	(64.57)	(20.32)	(26.73)
Old × Yr pre × Pop. ≥ 30K	142.74*	59.77**	43.46
	(75.72)	(23.20)	(28.36)
AR(2) test p-value	0.196	0.449	0.243
Municipalities	2496	2498	2498
Observations	12603	12625	12625

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited and the most recent election was decided by a vote margin of five percentage points or less. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.10: Mayor Age, Political Budget Cycles, and the Balanced-Budget Rule

Dynamic Panel Estimates	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	25.94***	11.09***	−0.76
	(7.39)	(2.96)	(2.62)
Yr pre × Pop. ≥ 5K	−36.05	−4.69	−6.95
	(42.22)	(7.93)	(21.01)
Yr pre × Pop. ≥ 5K × Yr ≥ 2001	0.71	−5.56	1.12
	(42.91)	(7.23)	(22.08)
Old	20.08***	5.51***	10.59***
	(4.31)	(1.58)	(2.12)
Old × Pop. ≥ 5K	−82.27***	−13.42	−50.60***
	(29.15)	(8.27)	(14.79)
Old × Pop. ≥ 5K × Yr ≥ 2001	19.55	−1.86	19.03
	(28.23)	(7.77)	(13.71)
Old × Yr pre	−19.05**	−10.97***	−2.05
	(9.67)	(4.00)	(3.54)
Old × Yr pre × Pop. ≥ 5K	−0.65	−7.87	5.99
	(66.76)	(13.90)	(34.41)
Old × Yr pre × Pop. ≥ 5K × Yr ≥ 2001	39.10	23.09*	8.58
	(67.15)	(13.38)	(35.11)
AR(2) test p-value	0.000	0.000	0.001
Municipalities	7063	7063	7063
Observations	60687	60802	60801

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.11: Mayor Age, Political Budget Cycles, and the Balanced-Budget Rule

Dynamic Panel Estimates, Close Elections

	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	28.96*	16.06***	4.78
	(15.22)	(5.26)	(4.61)
Yr pre × Pop. ≥ 5K	−57.74	−5.19	−29.42
	(69.09)	(10.32)	(38.73)
Yr pre × Pop. ≥ 5K × Yr ≥ 2001	17.58	−10.72	22.09
	(66.59)	(7.78)	(38.36)
Old	20.38**	6.05*	14.80***
	(9.97)	(3.12)	(5.51)
Old × Pop. ≥ 5K	−17.00	−8.95	−32.46*
	(46.66)	(15.60)	(18.92)
Old × Pop. ≥ 5K × Yr ≥ 2001	−44.74	−8.57	−2.78
	(44.65)	(14.09)	(17.74)
Old × Yr pre	−32.46*	−18.20***	−9.41
	(18.26)	(6.49)	(6.59)
Old × Yr pre × Pop. ≥ 5K	83.56	−1.30	65.81
	(80.99)	(18.46)	(41.93)
Old × Yr pre × Pop. ≥ 5K × Yr ≥ 2001	−22.40	25.02	−46.88
	(77.98)	(16.52)	(41.48)
AR(2) test p-value	0.193	0.441	0.273
Municipalities	2496	2498	2498
Observations	12603	12625	12625

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited and the most recent election was decided by a vote margin of five percentage points or less. The AR(2) test p-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.12: Mayor Age, Political Budget Cycles, and Single-Round vs. Runoff Elections

Dynamic Panel Estimates	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	−12.95	1.74	−2.90
	(12.19)	(4.10)	(9.11)
Yr pre × Pop. ≥ 15K	19.85	7.98*	9.50
	(17.16)	(4.73)	(11.39)
Old	15.68*	3.69*	7.71
	(8.62)	(2.15)	(6.29)
Old × Pop. ≥ 15K	−20.65	−4.11	−12.81
	(13.10)	(3.06)	(9.34)
Old × Yr pre	5.99	−7.43	3.35
	(20.86)	(7.28)	(14.56)
Old × Yr pre × Pop. ≥ 15K	−31.68	−2.83	−12.99
	(27.52)	(8.68)	(20.18)
AR(2) test *p*-value	0.078	0.438	0.781
Municipalities	857	857	857
Observations	7016	7022	7022

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited. The sample includes only municipalities with populations between 10,000 and 29,000 in order to hold the mayor's wage constant. The AR(2) test *p*-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.13: Mayor Age, Political Budget Cycles, and Single-Round vs. Runoff Elections

Dynamic Panel Estimates, Close Elections

	Revenue p.c.	Capital expenditure p.c.	Current expenditure p.c.
	(1)	(2)	(3)
Yr pre	−38.37	7.55	−19.16
	(26.42)	(8.35)	(25.42)
Yr pre × Pop. ≥ 15K	30.97	−4.13	28.43
	(32.39)	(9.88)	(26.00)
Old	15.76	3.55	0.25
	(25.65)	(5.36)	(6.96)
Old × Pop. ≥ 15K	−23.85	−6.67	−2.55
	(27.77)	(5.17)	(6.94)
Old × Yr pre	28.30	−19.01	21.72
	(35.72)	(11.64)	(28.94)
Old × Yr pre × Pop. ≥ 15K	−8.19	18.84	−16.87
	(42.67)	(13.99)	(28.28)
AR(2) test *p*-value	0.706	0.762	0.240
Municipalities	296	296	296
Observations	1341	1341	1341

Notes: All regressions include one lag of the dependent variable, account for municipality fixed effects and year effects, and are estimated by system GMM. The results are from the sample for which the mayor is not term limited and the most recent election was decided by a vote margin of five percentage points or less. The sample includes only municipalities with populations between 10,000 and 29,000 in order to hold the mayor's wage constant. The AR(2) test *p*-value corresponds to a test of serial correlation in the error term. Standard errors, reported in parentheses, are robust to heteroskedasticity and clustering at the level of municipality. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.14: The Effect of the Mayor's Age on Political Outcomes (Propensity-Score Methods)

Panel A: IPW and Matching Estimates, Close Elections

	Ran again		Reelected		Provincial admin.		Regional admin.		National admin.	
	IPW	Match	IPW	Match	IPW	Match	IPW	Match	IPW	Match
ATE of Old	-0.099^{***}	-0.110^{***}	-0.143^{***}	-0.154^{***}	-0.019	-0.017	-0.001	-0.001	0.008^{**}	0.007^{**}
	(0.028)	(0.031)	(0.026)	(0.029)	(0.012)	(0.013)	(0.005)	(0.006)	(0.004)	(0.003)
Observations	1281	1281	1281	1281	1281	1281	1281	1281	1281	1281

Panel B: Difference of ATE on Reelected and ATE on Ran again

	IPW	Match
Difference	-0.044^{**}	-0.044^{**}
	(0.021)	(0.022)

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. Panel A reports semiparametric estimates of the average treatment effect of the older candidate winning the election, using the subsample of elections determined by a margin of victory of two percentage points or fewer. For each outcome, the first column reports the inverse-probability-weighted regression-adjustment estimate, and the second column reports the propensity-score matching estimate. Both estimates utilize a propensity score estimated using a logistic model and the following covariates: prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and born locally. Heteroskedasticity-robust standard errors for the propensity-score matching estimates account for the first-stage estimation of the propensity score and are calculated according to Abadie and Imbens (2016). Panel B reports the difference of the average treatment effects on Reelected and Ran again, with standard errors based on 100 bootstrap repetitions. Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.15: The Effect of the Mayor's Age on Average Budget Outcomes (Propensity-Score Methods)

IPW and Matching Estimates, Close Elections	Avg. revenue		Avg. capital expend.		Avg. current expend.	
	IPW	Match	IPW	Match	IPW	Match
ATE of Old	12.53	2.93	−2.21	−3.63	13.44	12.05
	(26.79)	(27.37)	(5.82)	(5.81)	(13.57)	(14.20)
Observations	859	859	859	859	859	859

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. The table reports semiparametric estimates of the average treatment effect of the older candidate winning the election, using the subsample of elections determined by a margin of victory of two percentage points or fewer. For each outcome, the first column reports the inverse-probability-weighted regression-adjustment estimate, and the second column reports the propensity-score matching estimate. Both estimates utilize a propensity score estimated using a logistic model and the following covariates: prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and born locally. Heteroskedasticity-robust standard errors for the propensity-score matching estimates account for the first-stage estimation of the propensity score and are calculated according to Abadie and Imbens (2016). Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.16: The Effect of Age on House Prices and the Speed of Public Good Provision (Propensity-Score Methods)

IPW and Matching Estimates, Close Elections				
	Avg. house price		Speed of public good provision	
	IPW	Match	IPW	Match
ATE of Old	57.3	68.7	−0.0015	−0.0019
	(38.6)	(45.5)	(0.0047)	(0.0049)
Observations	514	514	859	859

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. The table reports semiparametric estimates of the average treatment effect of the older candidate winning the election, using the subsample of elections determined by a margin of victory of two percentage points or fewer. For each outcome, the first column reports the inverse-probability-weighted regression-adjustment estimate, and the second column reports the propensity-score matching estimate. Both estimates utilize a propensity score estimated using a logistic model and the following covariates: prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and born locally. Heteroskedasticity-robust standard errors for the propensity-score matching estimates account for the first-stage estimation of the propensity score and are calculated according to Abadie and Imbens (2016). Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.1.17: The Effect of the Mayor's Age on Political Budget Cycles (Propensity-Score Methods)

IPW and Matching Estimates, Close Elections						
	Δ Revenue		Δ Capital expend.		Δ Current expend.	
	IPW	Match	IPW	Match	IPW	Match
ATE of Old	3.82	0.18	-11.41^{**}	-15.65^{**}	1.15	-1.86
	(14.92)	(16.23)	(5.75)	(6.13)	(4.21)	(4.41)
Observations	859	859	859	859	859	859

Notes: In all specifications the sample is restricted to mayoral terms in which the term limit is not binding. The table reports semiparametric estimates of the average treatment effect of the older candidate winning the election, using the subsample of elections determined by a margin of victory of two percentage points or fewer. For each outcome, the first column reports the inverse-probability-weighted regression-adjustment estimate, and the second column reports the propensity-score matching estimate. Both estimates utilize a propensity score estimated using a logistic model and the following covariates: prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and born locally. Heteroskedasticity-robust standard errors for the propensity-score matching estimates account for the first-stage estimation of the propensity score and are calculated according to Abadie and Imbens (2016). Standard errors are in parentheses. See Tables C.1.1 and C.1.2 in the appendix for variable definitions and data sources. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.1.1: Density of Estimated Propensity Scores

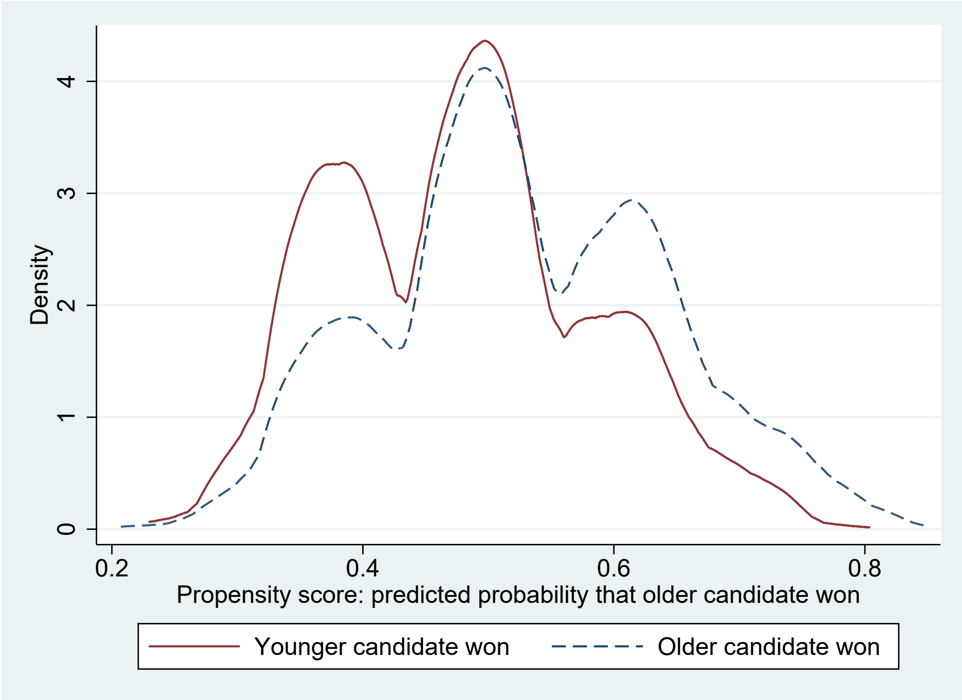

Notes: The graph plots the density of the estimated propensity score for elections in which the younger candidate won and elections in which the older candidate won. The propensity score for the older candidate winning is estimated using a logistic model and the following covariates: prior appointment to city office, elected city political experience, unelected city political experience, gender, high school education, college education, and born locally. Densities are estimating using the standard Epanechnikov kernel and the subsample of elections determined by a margin of victory of two percentage points or fewer.

Bibliography

ABADIE, A. AND G. W. IMBENS (2016): "Matching on the Estimated Propensity Score," *Econometrica*, 84(2), 781–807.

ACEMOGLU, D., S. JOHNSON, AND J. A. ROBINSON (2003): "An African Success Story: Botswana," in *In Search of Prosperity: Analytical Narratives on Economic Growth*, ed. by D. Rodrick, Princeton, NJ: Princeton University Press.

ACEMOGLU, D. AND J. A. ROBINSON (2006): *Economic Origins of Dictatorship and Democracy*, New York, NY: Cambridge University Press.

ADES, A. F. AND E. L. GLAESER (1995): "Trade and Circuses: Explaining Urban Giants," *Quarterly Journal of Economics*, 110(1), 195–227.

AGUSTINA, C. D., E. AHMAD, D. NUGROHO, AND H. SIAGIAN (2012): "Political Economy of Natural Resource Revenue Sharing in Indonesia," Working Paper 55, Asia Research Centre.

ALBOUY, D. (2012): "Evaluating the Efficiency and Equity of Federal Fiscal Equalization," *Journal of Public Economics*, 96, 824–839.

ALESINA, A., R. BAQIR, AND C. HOXBY (2004): "Political Jurisdictions in Heterogeneous Communities," *Journal of Political Economy*, 112(2), 348–396.

ALESINA, A., A. DEVLEESCHAUWER, W. EASTERLY, S. KURLAT, AND R. WACZIARG (2003): "Fractionalization," *Journal of Economic Growth*, 8 (2), 155–194.

ALESINA, A., P. GIULIANO, AND N. NUNN (2013): "On the Origins of Gender Roles: Women and the Plough," *Quarterly Journal of Economics*, 128 (2), 469–530.

ALESINA, A. AND M. PARADISI (2017): "Political Budget Cycles: Evidence from Italian Cities," *Economics and Politics*, 29, 157–177.

ALESINA, A. AND A. PASSALACQUA (2017): "The Political Economy of Government Debt," in *Handbook of Macroeconomics Volume 2*, ed. by J. Taylor and H. Uhlig, Amsterdam: North-Holland, forthcoming.

ALESINA, A. AND F. PASSARELLI (2015): "Loss Aversion in Politics," Working Paper 21077, National Bureau of Economic Research.

ALESINA, A., N. ROUBINI, AND G. COHEN (1997): *Political Cycles and the Macroeconomy*, Cambridge, MA: MIT Press.

ALEXEEV, M. AND R. CONRAD (2009): "The Elusive Curse of Oil," *Review of Economics and Statistics*, 91 (3), 586–598.

Allcott, H. and D. Keniston (2015): "Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America," Working Paper 20508, NBER.

Allen, P. A. and J. R. Allen (2005): *Basin Analysis: Principles and Applications*, Malden, MA: Blackwell Publishing, second ed.

Altindag, D. T. and N. Mocan (2015): "Mobile Politicians: Opportunistic Career Moves and Moral Hazard," Working paper.

Andersen, J. J. and S. Aslaksen (2013): "Oil and Political Survival," *Journal of Development Economics*, 100, 89–106.

Andersen, J. J., N. Johannesen, D. D. Lassen, and E. Paltseva (2014): "Petro Rents, Political Institutions, and Hidden Wealth: Evidence from Offshore Bank Accounts," Working paper.

Andersen, J. J. and M. L. Ross (2014): "The Big Oil Change: A Closer Look at the Haber-Menaldo Analysis," *Comparative Political Studies*, 47 (7), 993–1021.

Anderson, T. W. and H. Rubin (1949): "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations," *Annals of Mathematical Statistics*, 20, 46–63.

Angrist, J. D. and G. W. Imbens (1995): "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association*, 90 (430), 431–442.

Aragón, F. M. and J. P. Rud (2013): "Natural Resources and Local Communities: Evidence from a Peruvian Gold Mine," *American Economic Journal: Economic Policy*, 5 (2), 1–25.

——— (2015): "Polluting Industries and Agricultural Productivity: Evidence from Mining in Ghana," *Economic Journal*, forthcoming.

Aragón, F. M., J. P. Rud, and G. Toews (2015): "Mining Closure, Gender and Employment Reallocations: The Case of UK Coal Mines," Working paper.

Arellano, M. and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *Review of Economic Studies*, 58, 277–297.

Arellano, M. and O. Bover (1995): "Another Look at the Instrumental Variable Estimation of Error-Components Models," *Journal of Econometrics*, 68, 29–51.

Arezki, R., V. A. Ramey, and L. Sheng (2016a): "News Shocks in Open Economies: Evidence from Giant Oil Discoveries," *Quarterly Journal of Economics*, 132 (1), 103–155.

Arezki, R., F. van der Ploeg, and F. Toscani (2016b): "Shifting Frontiers in Global Resource Wealth: The Role of Policies and Institutions," Working paper.

Aslaksen, S. and R. Torvik (2006): "A Theory of Civil Conflict and Democracy in Rentier States," *The Scandinavian Journal of Economics*, 108 (4), 571–585.

Auten, G. (1999): "Capital Gains Taxation," in *The Encyclopedia of Taxation and Tax Policy*, ed. by J. J. Cordes, R. D. Ebel, and J. Gravelle, Washington, D.C.: Urban Institute Press, 58–61, first ed.

Baicker, K. (2005): "Extensive or Intensive Generosity? The Price and Income Effects of Federal Grants," *Review of Economics and Statistics*, 87 (2), 371–384.

Banks, A. S. and K. A. Wilson (2016): "Cross-National Time-Series Data Archive," http://www.cntsdata.com/.

Bartik, A. W., J. Currie, M. Greenstone, and C. R. Knittel (2017): "The Local Economic and Welfare Consequences of Hydraulic Fracturing," Working paper.

Bartolini, D. and R. Santolini (2009): "Fiscal Rules and the Opportunistic Behaviour of the Incumbent Politician: Evidence from Italian Municipalities," Cesifo wp. n. 2605.

Bazzi, S. and M. A. Clemens (2013): "Blunt Instruments: Avoiding Common Pitfalls in Identifying the Causes of Economic Growth," *American Economic Journal: Macroeconomics*, 5 (2), 152–186.

Bazzi, S. and M. Gudgeon (2016): "Local Government Proliferation, Diversity, and Conflict," Working paper, Boston University.

Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2015): "Determinants and Consequences of Bureaucratic Effectiveness: Evidence from the Indian Administrative Service," Working paper.

Besley, T. and S. Coate (2003): "Centralized Versus Decentralized Provision of Local Public Goods: A Political Economy Approach," *Journal of Public Economics*, 87(12), 2611–2637.

Besley, T., J. G. Montalvo, and M. Reynal-Querol (2011): "Do Educated Leaders Matter?" *The Economic Journal*, 121 (554), F205fi?!227.

Besley, T. and T. Persson (2009a): "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, 99 (4), 1218–1244.

——— (2009b): "Repression or Civil War?" *American Economic Review: Papers and Proceedings*, 99 (2), 292–297.

——— (2010): "State Capacity, Conflict, and Development," *Econometrica*, 78 (1), 1–34.

——— (2011): "The Logic of Political Violence," *Quarterly Journal of Economics*, 126 (3), 1411–1445.

Besley, T. and T. Persson (2011): *Pillars of Prosperity: The Political Economics of Development Clusters*, Princeton, NJ: Princeton University Press.

Bhattacharyya, S. and R. Hodler (2010): "Natural Resources, Democracy, and Corruption," *European Economic Review*, 54, 608–621.

Bisin, A., A. Lizzeri, and L. Yariv (2015): "Government Policy with Time Inconsistent Voters," *American Economic Review*, forthcoming.

Black, D. (1948): "On the Rationale of Group Decision-Making," *Journal of Political Economy*, 56 (1), 23–34.

Blundell, R. and S. Bond (1998): "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, 87, 115–143.

Boehm, C. E. (2016): "Government Spending and Durable Goods," Working paper, U.T. Austin.

Bohn, H. and R. T. Deacon (2000): "Ownership Risk, Investment, and the Use of Natural Resources," *American Economic Review*, 90 (3), 526–549.

Bonfatti, A. and L. Forni (2017): "Fiscal Rules to Tame the Political Budget Cycle: Evidence from Italian Municipalities," Imf working paper.

Bordignon, M., T. Nannicini, and G. Tabellini (2016): "Moderating Political Extremism: Single Round vs. Runoff Elections under Plurality Rule," *American Economic Review*, 106(8), 2349–2370.

Borge, L.-E., P. Parmer, and R. Torvik (2015): "Local Natural Resource Curse?" *Journal of Public Economics*, 131, 101–114.

——— (2016): "Local Natural Resource Curse?" *Journal of Public Economics*, forthcoming.

Boschini, A. D., J. Pettersson, and J. Roine (2007): "Resource Curse or Not: A Question of Appropriability," *Scandinavian Journal of Economics*, 109 (3), 593fi?!–617.

Bound, J., D. A. Jaeger, and R. M. Baker (1995): "Problems with Instrumental Variables Estimation When the Correlation between the Instruments and the Endogenous Explanatory Variable Is Weak," *Journal of the American Statistical Association*, 90 (430), 443–450.

Brender, A. and A. Drazen (2005): "Political Budget Cycles in New Versus Established Democracies," *Journal of Monetary Economics*, 52, 1271–1295.

Britannica, E. (2015): "Britannica Academic," http://www.britannica.com/, accessed June 2015.

Brollo, F., T. Nannicini, R. Perotti, and G. Tabellini (2013): "The Political Resource Curse," *American Economic Review*, 103 (5), 1759–1796.

Brollo, F. and U. Troiano (2016): "What Happens When a Woman Wins an Election? Evidence from Close Races in Brazil," *Journal of Development Economics*, 122, 28–45.

Brückner, M. and A. Ciccone (2011): "Rain and the Democratic Window of Opportunity," *Econometrica*, 79 (3), 923–947.

Brückner, M., A. Ciccone, and A. Tesei (2012): "Oil Price Shocks, Income, and Democracy," *Review of Economics and Statistics*, 94 (2), 389–399.

Brunnschweiler, C. N. and E. H. Bulte (2008): "The Resource Curse Revisited and Revised: A Tale of Paradoxes and Red Herrings," *Journal of Environmental Economics and Management*, 55 (3), 248–264.

Buettner, T. and D. E. Wildasin (2006): "The Dynamics of Municipal Fiscal Adjustment," *Journal of Public Economics*, 90, 1115–1132.

Burgess, R., M. Hansen, B. A. Olken, P. Potapov, and S. Sieber (2012): "The Political Economy of Deforestation in the Tropics," *Quarterly Journal of Economics*, 127 (4), 1707–1754.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82 (6), 2295–2326.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011): "Robust Inference With Multiway Clustering," *Journal of Business and Economic Statistics*, 29 (2), 238–249.

Cárdenas, M., S. Ramírez, and D. Tuzemen (2011): "Commodity Dependence and Fiscal Capacity," Working paper, Brookings Institution.

Carreri, M. and O. Dube (2015): "Do Natural Resources Influence Who Comes to Power, and How?" Working paper.

Casaburi, L. and U. Troiano (2015): "Ghost-House Busters: The Electoral Response to a Large Anti Tax Evasion Program," *Quarterly Journal of Economics*, forthcoming.

Caselli, F. and G. Michaels (2013): "Do Oil Windfalls Improve Living Standards? Evidence from Brazil," *American Economic Journal: Applied Economics*, 5 (1), 208–238.

Caselli, F. and A. Tesei (2016): "Resource Windfalls, Political Regimes, and Political Stability," *Review of Economics and Statistics*, 98 (3), 573–590.

Chandler, T. (1987): *Four Thousand Years of Urban Growth: An Historical Census*, Lewiston, NY: St. David's University Press.

Chattopadhyay, R. and E. Duflo (2004): "Women as Policy Makers: Evidence from a Randomized Policy Experiment in India," *Econometrica*, 72(5), 1409–1443.

Chen, H.-Y. (2009): "Reservoir Rock and Source Rock Types: Classification, Properties and Symbols," mimeo, New Mexico Tech.

CIA (2015): "The World Factbook," https://www.cia.gov/library/publications/the-world-factbook/, accessed June 2015.

Cioffi, M., G. Messina, and P. Tommasino (2012): "Parties, institutions and political budget cycles at municipal level: evidence from Italy," Working Paper 885, Bank of Italy.

Collier, P. and A. Hoeffler (2004): "Greed and Grievance in Civil War," *Oxford Economic Papers*, 56 (4), 563–595.

Collier, P. and C. Laroche (2015): "Harnessing Natural Resources for Inclusive Growth," Growth brief, International Growth Centre.

Corden, W. M. and J. P. Neary (1982): "Booming Sector and De-Industrialisation in a Small Open Economy," *The Economic Journal*, 92 (368), 825–848.

Correlates of War Project (2007): "Direct Contiguity Data, 1816–2006, Version 3.1," http://correlatesofwar.org, accessed June 2015.

Cotet, A. and K. Tsui (2013): "Oil and Conflict: What Does the Cross-Country Evidence Really Show?" *American Economic Journal: Macroeconomics*, 5 (1), 49–80.

Coviello, D. and S. Gagliarducci (2013): "Tenure in Office and Public Procurement," Working paper.

Crivelli, E. and S. Gupta (2014): "Resource Blessing, Revenue Curse? Domestic Revenue Effort in Resource-Rich Countries," Working paper, IMF.

Cust, J. and T. Harding (2015): "Institutions and the Location of Oil Exploration," Working paper.

Cust, J. and S. Poelhekke (2015): "The Local Economic Impacts of Natural Resource Extraction," Working paper.

Cust, J. and R. D. Rusli (2015): "Disentangling the Effects of Resource Extraction: Local Government and Investment Multipliers," Working paper, University of Oxford.

Cust, J. and C. Viale (2016): "Is There Evidence for a Subnational Resource Curse?" Policy paper, Natural Resource Governance Institute.

Dahlberg, M., E. Mörk, J. Rattsø, and H. Ågren (2008): "Using a Discontinuous Grant Rule to Identify the Effect of Grants on Local Taxes and Spending," *Journal of Public Economics*, 92, 2320–2335.

David, P. A. and G. Wright (1997): "Increasing Returns and the Genesis of American Resource Abundance," *Industrial and Corporate Change*, 6 (2), 203–245.

Djankov, S., J. G. Montalvo, and M. Reynal-Querol (2008): "The Curse of Aid," *Journal of Economic Growth*, 13 (3), 169–194.

Donald, S. G. and W. K. Newey (2001): "Choosing the Number of Instruments," *Econometrica*, 69 (5), 1161–1191.

Downs, A. (1957): *An Economic Theory of Democracy*, New York, NY: Harper & Row.

Drazen, A. (2000): "Political Business Cycle After 25 Years," *NBER Macroeconomics Annuals, MIT Press, Cambridge, MA.*

Dube, O. and J. F. Vargas (2013): "Commodity Price Shocks and Civil Conflict: Evidence from Colombia," *Review of Economic Studies*, 80, 1384–1421.

Easterly, W. R. (2001): "Global Development Network Growth Database," http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/0,,contentMDK:20701055~pagePK:64214825~piPK:64214943~theSitePK:469382,00.html.

ENI (2015): *World Oil and Gas Review*, Italy: ENI SPA.

Erle, S. and S. Gilles (2013): "World Borders Dataset, version 0.3," http://thematicmapping.org, accessed October 2013.

Fearon, J. D. and D. D. Laitin (2003): "Ethnicity, Insurgency, and Civil War," *The American Political Science Review*, 97 (1), 75–90.

Fischer, G., H. van Velthuizen, M. Shah, and F. Nachtergaele (2002): *Global Agro-Ecological Assessment for Agriculture in the 21st Century: Methodology and Results*, Rome: Food and Agriculture Organization of the United Nations.

Fitrani, F., B. Hofman, and K. Kaiser (2005): "Unity in Diversity? The Creation of New Local Governments in a Decentralising Indonesia," *Bulletin of Indonesian Economic Studies*, 41 (1), 57–79.

Flanders Marine Institute (2013): "Maritime Boundaries Geodatabase, version 7," http://www.marineregions.org, accessed October 2013.

Flavin, M. A. (1981): "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy*, 89 (5), 974–1009.

Frankel, J. A. (2010): "The Natural Resource Curse: A Survey," Working Paper 15836, NBER.

Fugro Robertson, Ltd. (2013): "Fugro Tellus Sedimentary Basins of the World Map," http://www.datapages.com/gis-map-publishing-program/gis-open-files/global-framework/robertson-tellus-sedimentary-basins-of-the-world-map.

Gadenne, L. (2017): "Tax Me, But Spend Wisely? Sources of Public Finance and Government Accountability," *American Economic Journal: Applied Economics*, 9 (1), 274–314.

Gadenne, L. and M. Singhal (2014): "Decentralization in Developing Economies," *Annual Review of Economics*, 6, 581–604.

Gagliarducci, S. and T. Nannicini (2013): "Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection," *Journal of the European Economic Association*, 11(2), 369fi?!398.

Gagliarducci, S. and M. D. Paserman (2012): "Gender Interactions with Hierarchies: Evidence from the Political Arena," *Review of Economic Studies*, 79, 1021–1052.

Gallup, J. L., J. D. Sachs, and A. D. Mellinger (1998): "Geography and Economic Development," Working Paper 6849, NBER.

Gennari, E. and G. Messina (2014): "How Sticky are Local Expenditures in Italy? Assessing the Relevance of the Flypaper Effect through Municipal Data," *International Tax and Public Finance*, 21, 324–344.

Gleditsch, N. P., P. Wallensteen, M. Eriksson, M. Sollenberg, and H. Strand (2002): "Armed Conflict 1946–2001: A New Dataset," *Journal of Peace Research*, 39 (5).

Gordon, N. (2004): "Do Federal Grants Boost School Spending? Evidence from Title I," *Journal of Public Economics*, 88, 1771–1792.

Grembi, V., T. Nannicini, and U. Troiano (2016): "Do Fiscal Rules Matter?" *American Economic Journal: Applied Economics*, 8 (3), 1–30.

Haber, S. and V. Menaldo (2011): "Do Natural Resources Fuel Authoritarianism? A Reappraisal of the Resource Curse," *American Political Science Review*, 105 (1), 1–26.

Hahn, J., P. Todd, and W. Van der Klaauw (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69 (1), 201–209.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50 (4), 1029–1054.

Herb, M. (2005): "No Representation without Taxation? Rents, Development, and Democracy," *Comparative Politics*, 37 (3), 297–316.

Hines, J. R. and R. H. Thaler (1995): "The Flypaper Effect," *Journal of Economic Perspectives*, 9 (4), 217–226.

Hofman, B., Kadjatmiko, K. Kaiser, and B. S. Sjahrir (2006): "Evaluating Fiscal Equalization in Indonesia," Policy Research Working Paper 3911, World Bank.

Horn, M. (2004): "Giant Oil and Gas Fields of the World," http://www.datapages.com/AssociatedWebsites/GISOpenFiles/HornGiantFields.aspx.

Imbens, G. W. and J. D. Angrist (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62 (2), 467–475.

Inman, R. P. (2008): "The Flypaper Effect," Working Paper 14579, NBER.

Jensen, A. (2011): "State-Building in Resource-Rich Economies," *Atlantic Economic Journal*, 39, 171–193.

Jensen, N. and L. Wantchekon (2004): "Resource Wealth and Political Regimes in Africa," *Comparative Political Studies*, 37 (7), 816–841.

Johnson, T. M. and B. J. Grim (2017): "World Religion Database," http://www.worldreligiondatabase.org.

Keele, L. J. and R. Titiunik (2015): "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 23, 127–155.

Kingston, D., C. Dishroon, and P. Williams (1983): "Global Basin Classification System," *AAPG Bulletin*, 67, 2175–2193.

Kleibergen, F. and R. Paap (2006): "Generalized Reduced Rank Tests Using the Singular Value Decomposition," *Journal of Econometrics*, 133 (1), 97–126.

Knight, B. (2002): "Endogenous Federal Grants and Crowd-Out of State Government Spending: Theory and Evidence from the Federal Highway Aid Program," *American Economic Review*, 92 (1), 71–92.

Kornai, J. (1992): *The Socialist System: The Political Economy of Communism*, Princeton, NJ: Princeton University Press.

Kuiper, K. (2010): *Ancient Rome: From Romulus and Remus to the Visigoth Invasion*, New York, NY: Rosen Education Service.

Kvenvolden, K. A. (2006): "Organic geochemistry—A retrospective of its first 70 years," *Organic Geochemistry*, 37 (1).

Ladd, H. F. (1993): "State Responses to the TRA86 Revenue Windfalls: A New Test of the Flypaper Effect," *Journal of Policy Analysis and Management*, 12 (1), 82–103.

Lane, P. R. and A. Tornell (1996): "Power, Growth, and the Voracity Effect," *Journal of Economic Growth*, 1, 213–241.

Lei, Y.-H. and G. Michaels (2014): "Do Giant Oilfield Discoveries Fuel Internal Armed Conflicts?" *Journal of Development Economics*, 110, 139–157.

Leland, H. E. (1968): "Saving and Uncertainty: The Precautionary Demand for Saving," *Quarterly Journal of Economics*, 82 (3), 465–473.

Litschig, S. (2012): "Are Rules-based Government Programs Shielded from Special-interest Politics? Evidence from Revenue-sharing Transfers in Brazil," *Journal of Public Economics*, 96, 1047–1060.

Litschig, S. and K. M. Morrison (2013): "The Impact of Intergovernmental Transfers on Education Outcomes and Poverty Reduction," *American Economic Journal: Applied Economics*, 5 (4), 206–240.

Liu, C. and G. Ma (2016): "Taxation without Representation: Local Fiscal Response to Intergovernmental Transfers in China," *International Tax and Public Finance*, 23, 854–874.

Lundqvist, H. (2015): "Granting Public or Private Consumption? Effects of Grants on Local Public Spending and Income Taxes," *International Tax and Public Finance*, 22, 41–72.

Lutz, B. (2010): "Taxation with Representation: Intergovernmental Grants in a Plebiscite Democracy," *Review of Economics and Statistics*, 92 (2), 316–332.

Maddison, A. (2013): "Statistics on World Population, GDP and Per Capita GDP, 1–2006 AD," http://www.ggdc.net/maddison/.

Mahdavy, H. (1970): "The Patterns and Problems of Economic Development in Rentier States: The Case of Iran," in *Studies in Economic History of the Middle East*, ed. by M. A. Cook, London: Oxford University Press, 428–467.

Marjoribanks, R. (2010): *Geological Methods in Mineral Exploration and Mining*, Springer-Verlag Berlin Heidelberg, 2nd ed.

Marshall, M. G. and T. R. Gurr (2014): "Polity IV Project: Political Regime Characteristics and Transitions, 1800–2013," http://www.systemicpeace.org/polity/polity4.htm.

Marshall, M. G. and D. R. Marshall (2016): "Polity IV Project: Coups d'État, 1946–2015," http://www.systemicpeace.org/inscrdata.html.

Martínez, L. R. (2017): "Sources of Revenue and Government Performance: Evidence from Colombia," Working paper, University of Chicago.

Mehlum, H., K. Moene, and R. Torvik (2006): "Institutions and the Resource Curse," *Economic Journal*, 116 (508), 1–20.

Meyersson, E. (2014): "Islamic Rule and the Empowerment of the Poor and Pious," *Econometrica*, 82(1), 229–269.

Michaels, G. (2011): "The Long Term Consequences of Resource-Based Specialisation," *The Economic Journal*, 121, 31–57.

Michaillat, P. (2014): "A Theory of Countercyclical Government Multiplier," *American Economic Journal: Macroeconomics*, 6 (1), 190–217.

Monteiro, J. and C. Ferraz (2012): "Does Oil Make Leaders Unaccountable? Evidence from Brazil's Offshore Oil Boom," Working paper, PUC-Rio.

Natural Resource Governance Institute (2016): "Natural Resource Revenue Sharing," Technical report, Natural Resource Governance Institute.

Nickell, S. (1981): "Biases in Dynamic Models with Fixed Effects," *Econometrica*, 49 (6), 1417–1426.

Nordhaus, W. (1975): "The political business cycle," *Review of Economic Studies*, 42, 169–190.

North, D. C. and B. R. Weingast (1989): "Constitutions and Commitment: The Evolution of Institutions Governing Public Choice in Seventeenth-Century England," *Journal of Economic History*, 49 (4), 803–832.

Oates, W. (1969): "The Effects of Property Taxes and Local Public Spending on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis," *Journal of Political Economy*, 77(6), 957–971.

——— (1972): *Fiscal Federalism*, New York, NY: Harcourt Brace Jovanovich.

Obstfeld, M. and K. Rogoff (1996): *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press, first ed.

Olsson, O. and M. Valsecchi (2015): "Resource Windfalls and Local Government Behavior: Evidence from a Policy Reform in Indonesia," Working Paper 635, SWOPEC.

Persson, T. and G. Tabellini (2000): *Political Economics: Explaining Economic Policy*, Cambridge, MA: MIT Press.

——— (2009): "Democratic Capital: The Nexus of Political and Economic Change," *American Economic Journal: Macroeconomics*, 1 (2), 88–126.

Pilot, S. (2004): "Should Politicians Have a Retirement age?" *Times of India*.

Prato, C. (2015): "Electoral Competition and Policy Feedback Effects," Working paper.

Prichard, W., A. Cobham, and A. Goodall (2014): "The ICTD Government Revenue Dataset," Working Paper 19, International Centre for Tax and Development.

Reinikka, R. and J. Svensson (2004): "Local Capture: Evidence from a Central Government Transfer Program in Uganda," *Quarterly Journal of Economics*, 119 (2), 679–705.

Repetto, L. (2016): "Political Budget Cycles with Informed Voters: Evidence from Italy," Working paper, Uppsala University.

Robinson, J. A., R. Torvik, and T. Verdier (2006): "Political Foundations of the Resource Curse," *Journal of Development Economics*, 79, 447–468.

Rogoff, K. (1990): "Equilibrium Political Budget Cycles," *American Economic Review*, 80, 21–36.

Rogoff, K. and A. Sibert (1988): "Elections and Macroeconomic Policy Cycles," *The Review of Economic Studies*, 55, 1–16.

Roodman, D. (2009): "A Note on the Theme of Too Many Instruments," *Oxford Bulletin of Economics and Statistics*, 71 (1), 135–158.

Ross, M. L. (2001): "Does Oil Hinder Democracy?" *World Politics*, 53 (3), 325–fi?!61.

——— (2008): "Oil, Islam, and Women," *American Political Science Review*, 102 (1), 107–123.

——— (2012): *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations*, Princeton, NJ: Princeton University Press.

——— (2013): "Oil and Gas Data, 1932–2011," http://hdl.handle.net/1902.1/20369.

——— (2014): "What Have We Learned about the Resource Curse?" Working paper.

Sachs, J. D. and A. M. Warner (1995): "Natural Resource Abundance and Economic Growth," Working Paper 5398, NBER.

——— (1999): "The Big Push, Natural Resource Booms And Growth," *Journal of Development Economics*, 59 (1), 43–76.

——— (2001): "Natural Resources and Economic Development: The Curse of Natural Resources," *European Economic Review*, 45, 827–838.

Sanderson, E. and F. Windmeijer (2016): "A Weak Instrument F-test in Linear IV Models with Multiple Endogenous Variables," *Journal of Econometrics*, 190, 212–221.

Shah, A., R. Qibthiyyah, and A. Dita (2012): "General Purpose Central-Provincial-Local Transfers (DAU) in Indonesia: From Gap Filling to Ensuring Fair Access to Essential Public Services for All," Working paper, World Bank.

Skoufias, E., A. Narayan, B. Dasgupta, and K. Kaiser (2014): "Electoral Accountability and Local Government Spending in Indonesia," Policy Research Working Paper 6782, World Bank.

Smith, B. (2015): "The Resource Curse Exorcised: Evidence from a Panel of Countries," *Journal of Development Economics*, 116, 57–73.

Song, Z., K. Storesletten, and F. Zilibotti (2012): "Rotten Parents and Disciplined Children: A Politico-Economic Theory of Public Expenditure and Debt," *Econometrica*, 80(6), 2785–2803.

Southard, J. (2007): "Sedimentary Basins," Lecture notes, MIT OpenCourseWare.

Staiger, D. and J. H. Stock (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65 (3), 557–586.

Stinnett, D. M., J. Tir, P. Schafer, P. F. Diehl, and C. Gochman (2002): "The Correlates of War Project Direct Contiguity Data, Version 3," *Conflict Management and Peace Science*, 19 (2), 58–66.

Stock, J. H. and M. Yogo (2003): "Testing for Weak Instruments in Linear IV Regression," Working Paper 284, NBER.

Strang, D. (1991): "Global Patterns of Decolonization, 1500–1987," *International Studies Quarterly*, 35 (4), 429–454.

Thieme, N., P. Lujala, and J. K. Rød (2007): "The Petroleum Dataset: Country Profiles," Working paper.

Tir, J., P. Schafer, P. F. Diehl, and G. Goertz (1998): "Territorial Changes, 1816–1996: Procedures and Data," *Conflict Management and Peace Science*, 16 (1), 89–97.

Tornell, A. and P. R. Lane (1999): "The Voracity Effect," *The American Economic Review*, 89 (1), 22–46.

Torvik, R. (2009): "Why Do Some Resource-abundant Countries Succeed While Others Do Not?" *Oxford Review of Economic Policy*, 241–256.

Tsui, K. (2010): "Resource Curse, Political Entry, and Deadweight Costs," *Economics and Politics*, 22 (3), 471–497.

——— (2011): "More Oil, Less Democracy: Evidence from Worldwide Crude Oil Discoveries," *Economic Journal*, 121, 89–115.

UNFPA Indonesia (2014): "Midwifery Consultancy Report," Technical report, United Nations Population Fund.

van der Ploeg, F. (2011): "Natural Resources: Curse or Blessing?" *Journal of Economic Literature*, 49 (2), 366–420.

van der Ploeg, F. and S. Poelhekke (2010): "The Pungent Smell of "Red Herrings": Subsoil Assets, Rents, Volatility and the Resource Curse," *Journal of Environmental Economics and Management*, 60 (1), 44–55.

Vegh, C. A. and G. Vuletin (2015): "Unsticking the Flypaper Effect in an Uncertain World," *Journal of Public Economics*, 131, 142–155.

Vicente, P. C. (2010): "Does Oil Corrupt? Evidence from a Natural Experiment in West Africa," *Journal of Development Economics*, 92, 28–38.

Vogl, T. S. (2014): "Race and the politics of close elections," *Journal of Public Economics*, 109, 101–113.

Wiens, D., P. Poast, and W. R. Clark (2014): "The Political Resource Curse: An Empirical Re-evaluation," *Political Research Quarterly*, 67 (4), 783–794.

Wooldridge, J. M. (2010): *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT Press, 2nd ed.

World Bank (2003): "Decentralizing Indonesia: A Regional Public Expenditure Review Overview Report," Technical report, World Bank, Jakarta.

——— (2007): "Spending for Development: Making the Most of Indonesia's Opportunities," Technical report, World Bank, Jakarta.

——— (2010): "Village Capacity in Maintaining Infrastructure: Evidence from Rural Indonesia," Technical report, World Bank, Jakarta.

Yinger, J. (1982): "Capitalization and the Theory of Local Public Finance," *Journal of Political Economy*, 90(5), 917–943.

Zeldes, S. P. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97 (2), 305–346.

Zou, H. (1994): "Dynamic Effects of Federal Grants on Local Spending," *Journal of Urban Economics*, 36, 98–115.