Corporate Interests and Campaign Donations in Brazilian Municipalities*

James Mao & Galileu Kim
Princeton University

May 04, 2020

Abstract

In this paper, we test the effects of Brazil's municipal anti-corrupton program on corporate campaign contributions to mayoral elections. We first leverage the random municipal audits to identify a clear and consistent effect – audited municipalities attract more corporate attention in the form of campaign contributions in the following election. We then use the framework of modern portfolio theory to explore the implications of these findings for debates on whether corporations behave as risk-reward optimizing investors in making campaign contributions.

1 Introduction

Investing in political candidates is risky: non-trivial sums of money are spent and potentially gained. A central concern for the health of democracies across the developing and developed world is the potentially destabilizing role played by private companies on electoral campaigns. Voters and political scientists alike voice concerns over the crowding out of individual donations by corporations whose interests may misalign with that of the median voter. Operation Car Wash (Operacao Lava-Jatos), an ongoing federal investigation by the Brazilian federal police, has

*Corresponding address: jsmao@princeton.edu. A previous version of this paper was presented in a junior workshop at Midwest Political Science Association 2019 under the title "Thumb on the Scale":" Corporate Interests and Campaign Donations to Dynastic Candidates in Brazil."

uncovered evidence of donations by private construction companies to politicians amounting to billions of dollars since 2014. These companies sought to secure and benefited from valuable, multimillion contracts to expand Brazil's infrastructure.[1] While illicit donations have been the primary focus of these federal investigations, legal donations have also been an important vehicle for companies to tilt the political playing field in their favor. Private legal donations by companies incriminated by the investigations have amounted to over 40% of legal donations to candidates in the 2014 electoral cycle.[2] According to The Intercept Brazil, the multinational, Rio-based construction company Odebrecht donated to over one thousand political candidates since 2002, for a total of $70 million.[3] Interestingly, for the 2014 presidential campaign, Odebrecht donated $4 million to PT candidate Dilma Roussef and $3 million to the runner up, Aécio Neves (PSDB).

While the study of campaign finance is well-developed in the U.S. context, we join the growing ranks of scholars who seek to extend its insights to more weakly institutionalized democracies (Nölke and Vliegenthart 2009; Samuels 2001a; Schneider and others 2004; Treisman 1998). We study corporate campaign finance at the local level in Brazil over three elections from 2004 to 2012. In each mayoral electoral cycle, there are over 15,000 potential mayoral candidates in which to invest. Over 5,000 potential municipal budgets are up for grabs, a valuable prize under direct purview of the winner. In our setting, corporations' contribution decisions interact with a unique anti-corruption initiative instituted by the Brazilian federal government in 2003. At the local level in Brazilian politics, party affiliations are ephemeral and campaign finance sources are limited (Samuels 2001b). Campaign financing by corporations makes up a large majority of mayoral candidates' funding and warrants further examination. We drill down into disparities in corporate responses to municipalities that are targeted by these corruption audits before an election. In doing so, we provide some support to the notion, advanced by Bezerra (1999) and Samuels (2001b), that corporate contributions are long-term investments in political assets.

The first step in our analysis is to establish that corporations' political involvement at the local level is affected by the random municipal audits. We show that this is indeed the case. Our preferred reduced-form specification reveals that the total corporate money dedicated to one mayoral race is five times greater in municipalities audited prior to an election compared to a control group. In addition, the number of unique corporations involved in the race is 1.3 times greater and the proportion of candidates backed by corporate funding in the race is 11 percentage points greater than in the control group races.

[1] See here for an excellent journalistic summary.
[2] For the Folha coverage, see here. In Estadão, here.
[3] See article here.

We next explore the implications of the corporate response to the audits with respect to the motivations for corporation contributions. Through a series of comparisons of corporate involvement in audited municipalities compared to non-audited ones, we suggest that corporations may indeed follow the long-term investment strategy suggested in extant literature on Brazilian politics. In this section, we make extensive use of insights from financial portfolio theory to describe patterns corporations' contribution behavior. We rule out the consumption rationale for political contributions by examining a subset of firms that contribute in multiple electoral cycles.

Our paper begins with a description of the local political environment in Brazil. We explain how the audits provide a unique opportunity to distinguish competing explanations for corporate giving, before diving into the analysis laid out above. Taken together, the evidence we adduce in this paper provide tentative support to the portrayal of corporate interests as forward-looking, rational political investors. A surprising finding we uncover is that Brazilian firms appear to use the shock induced by the audits as an opportunity to invest in candidates with less short-term returns on investment, but with the possibility of acceptable long-term returns. We recognize that these conclusions are suggestive at best and view this paper as an important contribution to understanding corporations' decision-making, given that the Brazilian audit program continues to be implemented and that corporations are now prohibited from contributing to political candidates, following a 2014 federal ruling. We discuss the implications of our findings for these recent political developments in the conclusion.

2 Background

2.1 Brazilian Local Politics

The motivation for and effects of political campaign donations, by both individuals and corporations, remains an unsettled question. Stratmann (2005) captures the central puzzle succinctly: "Do contributors consider contributions as pure consumption, as investment in policy, as a means of gaining access to the legislator, or as a way of influencing elections?" The first alternative sees contributions as a purely expressive instrument by actors who derive utility from the act of giving to a candidate for whom they feel an affinity. The next two possibilities both suggest that political candidates understand contributions to be a quid pro quo. The last potential explanation requires corporations to wield some degree of pivotality in affecting elections, which may be a strong requirement.

Scholars of Brazilian politics tend to favor quid pro quo explanations. However, an enormous

challenge to this interpretation is the commitment problem: politicians, once elected, can renege on implicit or explicit quid pro quo arrangements in the absence of full contracting (McCarty and Rothenberg 1996). Politicians can attempt to create reputations for reliability (Kroszner and Stratmann 2000, 2005), but this may be easier for some types of political candidates than others. Incumbent candidates may be better equipped to leverage reputational effects, which may explain the strong regularity in U.S. politics of the incumbency advantage in receiving contributions (Fouirnaies and Hall 2014; Poole and Romer 1985). But in general, candidates should find it difficult to promise to uphold commitments. Treisman (1998) argues that candidates in weakly institutionalized democracies face even greater difficulties in establishing reputations for delivering on investments.

On the other hand, the Brazilian context may be different. Bezerra (1999) argues that in a clientelistic society such as Brazil, politicians can acquire the label of being reliable investments. This happens especially when politicians and firms have repeated interactions as the politicians travel across different levels of government. Samuels (2001b) suggests that corporations "with budgetary interests take a long-term view and seek to invest in enduring relationships with politicians who will exert influence over government resources." Building on this insight, Boas, Hidalgo, and Richardson (2014) find that campaign contributions to Workers' Party legislators at the national level explain variation in public works contracts awarded to firms as recently as in 2010.

We argue that the long-term rationale should apply even more strongly to mayoral races. In Brazil, mayoral offices are highly sought after, even more so than federal legislative positions, though perhaps a notch below governors' offices (Samuels 2003; Santos and Pegurier 2011). Brazil is a presidential democracy with a highly decentralized, federal system of government (Falleti 2010). The 1988 Constitution, which ended military junta rule, established municipalities as autonomous local governments, meaning that city-level mayors wield inordinate power relative to politicians with similar titles in other democracies (Baiocchi 2006). Brazilian mayors have executive powers and are empowered, in conjunction with elected city councilors, to make budgetary and bureaucratic personnel decisions. Another way their political autonomy is apparent is in their electoral cycles, which do not coincide with that of governors and the president. Municipal elections are held in four-year cycles, with the latest one conducted in 2016.

Corporations that contribute to political candidates because of investment motivations must take into account both the immediate returns the candidate can return in the next mayoral term, conditional on election, as well as the long-term expected return on the candidate. We should expect, for instance, that corporations will find candidates with more political experience more

attractive. Santos and Pegurier (2011) show that candidates for federal legislative office who have previously held legislative positions at other levels of government are generally more likely to be elected and are more productive, while Meireles (2019) finds that mayoral candidates without previous federal experience are less likely to be elected. We build on these empirical patterns to explain corporate contribution allocations.

2.2 Mayoral Candidates as Political Investments

If mayoral candidates are investable assets, then corporations must form beliefs about the expected returns on each candidate, as well as the expected risk associated with that level of return. Posed this way, corporations' contribution decision is analogous to a portfolio allocation problem. Despite the prevalence of the metaphor of political contributions as investments, we are not aware of many studies that explicitly attempt to model the contribution decision as a financial investment decision. The closest is Snyder Jr (1990), who considers the risk-reward calculus of donating to U.S. House of Representative members and finds empirical support for his stylization of the investment problem. In the Brazilian local context, we use the insights of modern portfolio theory in the study of finance (Markowitz 1952) to achieve two ends. First, the framework of financial portfolio theory allows us to specify firms' preferences and describe firms' most preferred candidates. Second, taking seriously the metaphor of contributions as investments drives insights into firms' reactions to certain political institutions; in this paper, we study investments in the shadow of the anti-corruption audits as one example.

There are two determinants of expected return on a political candidate: their probability of winning the election and the material goods they can deliver once elected. Expected returns on financial assets are typically formulated based on a time series of past changes in the asset's price, but each individual mayoral candidate usually does not have an analogous history of data. In our dataset, for example, almost three-fifths of mayoral candidates across four election year have previously held political office one time or fewer.[4] We therefore assume that corporations make contribution allocations based on heuristic *types* of candidates, a reasonable assumption to make in an environment with thousands of candidates each cycle and no strong party labels (Mainwaring 1998). Candidate types are determined by four individual attributes. The first is political experience, as discussed above. The next two relate to incumbency status. Unlike in many other settings, incumbency advantage is not an established regularity in Brazilian local politics. Several studies have instead found that, at the mayoral level, incumbency can be a disadvantage rather than an advantage in Brazil (Brambor and Ceneviva 2011; Klavsnja and

[4]We introduce our data in greater depth in the next section.

Titiunik 2017; Titiunik 2009).

We confirm the primacy of political experience and the incumbency disadvantage in local elections using a simple quantitative exercise. Without making claim to causality, we assess which candidate attributes are the strongest predictors of electoral success using a least absolute shrinkage and selection operator (lasso) regression analysis. The advantage of the lasso regression is that it selects covariates that contribute most to the outcome variable, suppressing all other covariates to zero.[5] Across the 57,056 mayoral candidates in our data, we regress a dummy variable taking on the value of one if the candidate wins their electoral race and zero if not on a battery of individual-level attributes: incumbency status, presence of an incumbent in the election, whether the candidate shares a family name with a previous mayor in the same municipality (which we call progeny candidates),[6] candidate gender, previous political experience, whether the candidate is running under a party in the president's political coalition, the candidate's education level, the candidate's professional occupation prior to the election, the amount of corporate campaign donations to the candidate, and the amount of self-financing contributed by the candidate to their campaign. Our model is interpretable as a linear probability model where only the most relevant covariates are kept in the model.

As Figure 1 shows, among this set of covariates, the strongest predictors of electoral status are political experience, as measured by the number of previous offices held. Political veterans, defined as candidates who have previously held at least two political offices, are associated with a 49 percentage point increase in the probability of election. The next two influential predictors by absolute magnitude are candidate incumbency status and the presence of an incumbent. Being the incumbent candidate is associated with an 15 percentage point increase in the candidate's probability of winning the election. At the same time, being in the same race as an incumbent is associated with a 13 point decrease in the probability of winning.[7]

Notably, the amount of corporate campaign contributions is not considered by the regularization model to be a large determinant of electoral success. This result aligns with the interpretation of corporate contributions as investments, rather than as a way of influencing elections, the fourth possible contribution motivation offered by Stratmann (2005). An alternative specification of the lasso regression where corporate money is measured as the difference between the candidate's

[5] See the Appendix for a full explanation of our cross-fold validation approach to this prediction problem.

[6] We identify progeny candidates using a procedure similar to Braganca, Ferraz, and Rios (2015). For each elected mayor, we match any succeeding mayoral candidates in the same municipality through their family name. Note that in Brazil it is common practice for children to inherit the paternal last name from both parents. We do not restrict progeny to immediate children.

[7] The lasso estimator cannot guarantee unique optimal coefficient values when covariates are collinear, which the predictors incumbency status and incumbent in the race clearly are. However, running multiple iterations of the regularization through cross-fold validation alleviates this concern (see Appendix).

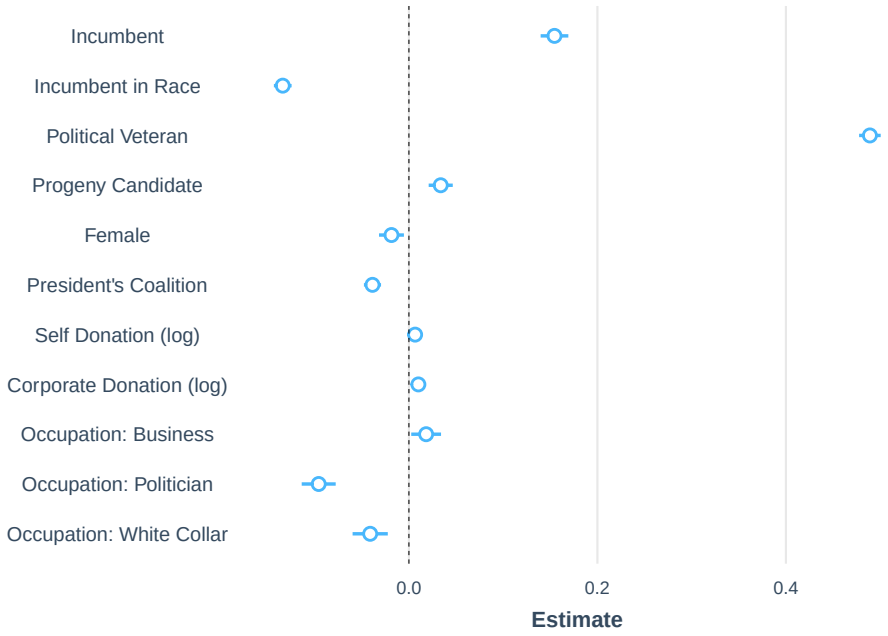

Figure 1: Predictors of electoral success in linear probability model, with predictors selected by lasso regularization. Political experience is by far the strongest predictor, followed by incumbent status and presence of incumbent in the race. Coefficients are percentage point change in probability of winning the election given a one unit change in predictor, shown with 95% confidence intervals. See main text for definitions of specific predictors. Sample is 42,777 mayoral candidates across 5,346 municipalities and three elections (2004, 2008, and 2012). $R^2 = 0.22$.

corporate receipts and the other candidates' receipts does not prevent the model from once again suppressing the coefficient on the variable to zero. Therefore, we are comfortable with proceeding under the assumption that campaign contributions are investments in political assets, with the caveat that we will address the possibility that mayoral candidates are consumption goods later in the paper.

We see that candidates who vary by political experience and incumbency are associated with varying probabilities of winning their races. Later, we will show that there is variation in the second component of expected returns, the material goods an elected official can return to their contributors. For now, we focus on the fourth individual attribute that defines candidate types. This attribute stems from Brazil's anti-corruption program. Building on previous studies of the effects of the anti-corruption program, we consider candidates that are running in municipalities that were audited by the federal anti-corruption office prior to the election as fundamentally different from candidates that are running in municipalities that were not audited before the election. In the next subsection, we first provide background on the audit program.

2.3 Brazil's Random Audit Program

Though we have stressed the autonomy of Brazilian municipalities, one way in which municipalities remain very much subordinate to the federal government is in their funding. A large proportion of the city budget is sourced from federal, intergovernmental transfers, especially in smaller municipalities (Arretche 1999; Brollo et al. 2013; Prado 2001). In Brazil, still a young democracy, accusations of corruption and public waste abound.

In May 2003, the federal government launched an ambitious anti-corruption program. Since then, the Corregedoria Geral da Uniao (CGU) has conducted several lotteries that select anywhere from 25 to 60 municipalities for random audits. Multiple lotteries are conducted each year, in conjunction with the national (gambling) lottery in the capital. All municipalities that are not state capitals and that have a population of 500,000 or fewer are eligible for selection, which is over 99% of the country. Audits are undertaken on site by the CGU over approximately ten days, and in teams of 10-15. Previous studies have established the near-unimpeachability of the auditors tasked with rooting out corruption (Bersch, Praca, and Taylor 2017; Ferraz and Finan 2008, 2011). The auditors examine documents relating to federal transfers to the municipal governments in the last three years, as well as interview key stakeholders about malfeasance directly. When irregularities are found, the auditors add them to a report that is swiftly publicized by the central CGU office, including the posting of the report on the CGU's online repository. The reports document the exact nature of each irregularity (Colonnelli and

Prem 2017). Media coverage of the audit reports is high, and the reports are also directly sent to the members of the judicial branch and the municipal legislative branch. Municipal councils and opposition candidates have used the reports to hurt sitting mayors' political prospects (Avis, Ferraz, and Finan 2018).

By now, many municipalities have been subjected to the audits. The sampling procedure for each lottery is stratified by state and is with replacement. Each municipality's implied audit probability is less than 3% a year (Colonnelli and Prem 2017). In our sample of 5,346 eligible municipalities, 1,108 were audited for the first time between 2006 and 2015. The average audit in our data covers 58 service items and $6.7 million in 2010 USD. Thus, Brazil's anti-corruption program is a true crackdown.

The random audits have been shown to produce a range of normatively desirable outcomes in Brazil's municipalities. Researchers have been able to exploit the randomness of municipal selection, which is key to causal inference, to study the effects of this novel anti-corruption mechanism. The program has been shown to reduce future corruption in the audited municipal government (Avis, Ferraz, and Finan 2018); increase firm activity and efficiency in the same municipality (Colonnelli and Prem 2017); reduce incumbent re-election probability as the amount of corruption uncovered by the auditors increases (Ferraz and Finan 2008); and discipline incumbent-eligible mayors to refrain from corruption (Ferraz and Finan 2011).[8] Most studies support the notion that the audits are working as intended: they have had a cleansing effect on municipalities.

We infer from these studies that municipalities that are audited before an election are different from municipalities that are not in a way that is pertinent to political investments. Given the timing of the lotteries, audits can shape the course of a mayoral campaign. We thus introduce audit status as the fourth characteristic of candidate types. These four characteristics define the set of political assets available to corporate contributors in each election.

To verify that audits do affect corporate contribution behavior, we exploit the randomness in the municipalities selected and assess the difference in corporate involvement in audited ("treated") mayoral races compared to a relevant control group of un-audited mayoral races. Because the audits are an unanticipated shock to the election, we are also able to make inferences about corporations' motivations for giving based on how they respond to the audits. We first estimate the effect of the audits on corporate behavior.

[8] In other work, one of us has demonstrated that the audits do not drive bureaucratic turnover, though they do seem to have an effect on improved management practices within municipal bureaucracies (Kim and Ferrali 2019).

3 The Beckoning Effect of Anti-Corruption Audits

In this section, we show that the random audits attract greater corporate involvement into audited municipalities. Firms, as a whole, increase campaign contributions to races that were audited before the election.

We use a similar empirical strategy as Ferraz and Finan (2008), where outcomes in municipalities audited prior to an election are compared to outcomes in municipalities audited after the election. The latter group is considered the control group since they are not affected by the treatment during the period when receipt of corporate contributions is permitted, which is restricted to the same year as the election. Unlike Ferraz and Finan (2008), who only consider the 2004 election, we pool our observations across the following two elections, 2008 and 2012.

Our data on elections comes from Brazilian electoral court's (TSE) extensive dataset of all political candidates at the local level. Each candidate is associated with a unique national identifier (CPF), which allows us to follow them over time and associate each candidate in each election with the number of political offices previously held. The candidate dataset also contains individual characteristics such as their complete birth name, age, education level, and professional occupation.[9] Campaign contribution data is also collected by the TSE. For every electoral cycle from 2002 to 2018, the electoral court mandates political candidates to disclose their campaign contributions. We are therefore able to map corporate contributors via their unique identifier (CNPJ) to the candidates they support via those CPF identifiers. This is how we build our dataset of company donations. While contributions data are self-reported, candidates who are found to have misreported their campaign finances are threatened with and have been effectively banned from running for office. Additionally, while this dataset may not capture illegal campaign contributions, it does provide us with a rich picture of campaign donations in Brazil and has been extensively used in other studies.[10] We deflate all contribution amounts to constant 2010 Brazilian reals, abbreviated R$.[11]

At the municipal level, wecollect a variety of municipal demographic from the National Institute of Geography (IBGE). Municipal covariates are used as controls in most of our estimations, following previous studies on the effects of the audits, and include municipal median wages, literacy rates, populations, rurality, and public goods provision in the form of garbage cleanup. All of these municipal covariates are measured from the 2000 Brazilian census, so they are prior to any of our electoral outcomes. We also collect data on municipal budgets from the Brazilian

[9] One caveat is that all of these characteristics are self-declared.
[10] See Boas, Hidalgo, and Richardson (2014) and Colonnelli, Prem, and Teso (2018).
[11] Cite deflate package in R.

Table 1: Total corporate involvement increased over elections, but the average amount per municipality remained roughly cosntant. All amounts in constant 2010 BR.

Year	Municipalities	Total Contributions	Amount per Municipality
2004	2,347	115,221,048	49,093
2008	3,549	191,524,400	53,966
2012	3,824	201,784,819	52,768

Finances (FINBRA). As we will discuss in the next section, municipal budget data is integral to our definition of political asset returns.

Finally, we collect data on the municipal audits from the reports themselves. We are able to extract data from 21 lotteries (lottery 20 through lottery 40), which encompass the years 2006 through 2015. This means that we can estimate the effect of the audits on two elections, 2008 and 2012. However, we still make use of candidate and contribution data pertaining to the 2004 election. Our final sample is therefore 16,038 mayoral races from 2004 to 2012, across 5,346 municipalities. The final sample is a function of eligibility for the random audits and the availability of budget data for at least one year in each of the electoral terms 2005-2008, 2009-2012, and 2013-2016.

3.1 Corporate Contributions Data

Our data on corporate contributions covers 27,223 corporations. Since different branches of the same corporation can be assigned different CNPJ identifiers, we have matched parent corporations to subsidiaries to the best of our ability, based on their names, sector category, and legal registration status. Each subsequent election in our data has seen more corporate donations to more municipalities, although the average amount contributed per race remains constant, at about R$50,000 (see Table 1). Our data excludes service industry companies, as those were significantly more difficult to match parents to subsidiaries given the paucity of the TSE data. We have data on six sectors; their total contributions over time are shown in Figure 2. Sectors are categorized via the National Classification of Economic Activities (CNAE) code assigned by IBGE to each company. We map the CNAE code to the industry classifications in Dix-Carneiro (2014).

Corporate contributions have increased over time for each sector. We will examine whether sectors respond to the audits differently. Figure 2 shows that construction firms also account for a large share of corporate contributions, despite accounting for only 20% of total firms by count. Although we do not perceive much effect of the audits by sector, it is notable that

Colonnelli and Prem (2017) classify construction as a government-dependent sector based on their share of municipal-level procurement contracts. This is suggestive evidence of a quid pro quo relationship, from contributions to contracts. We examine the relationship in more detail when we discuss expected returns on political investments, below.

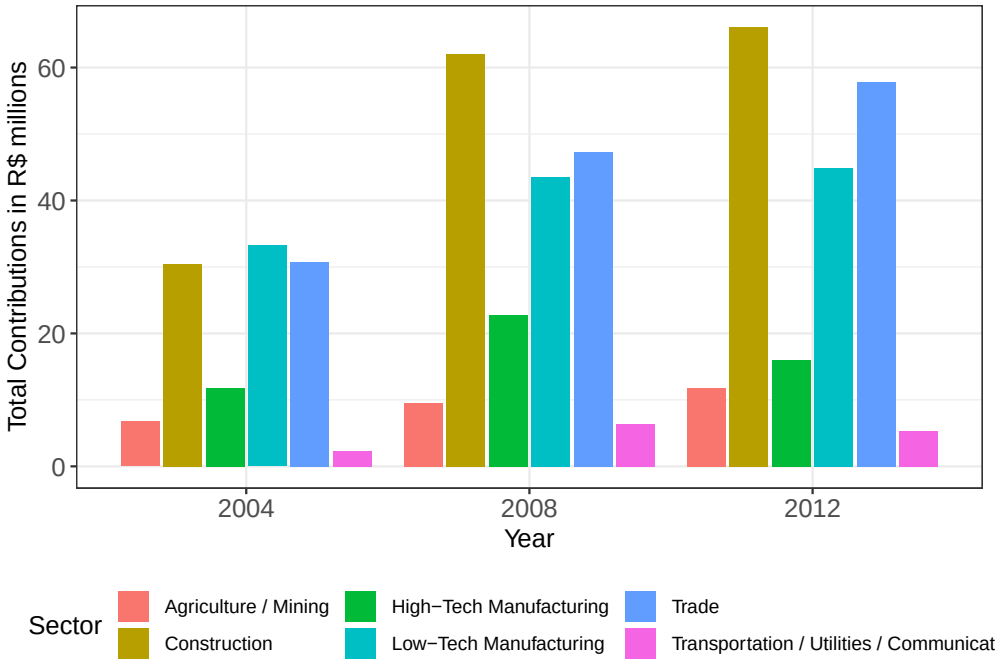

Figure 2: Sector-level corporate contributions over time. Of the 27,223 firms in our data, trade firms account for 44%, low-tech manufacturing accounts for 27%, construction accounts for 20%, high-tech manufacturing accounts for 4%, agriculture/mining accounts for 4%, and transportation/utility/communications accounts for 1%.

3.2 Audit Data

We consider both the baseline effect of the audits as well as heterogeneous effects based on the degree of municipal corruption. To that end, we construct a measure of corruption based on the information in the audit reports. As we described previously, the CGU auditors are sent to a set of randomly selected municipalities after each lottery to meticulously record every instance of misconduct by the local government. Each audit, called a "service order," reports irregularities that fall into one of three categories: notices, intermediate faults, and serious faults. The latter two indicate misconduct. Based on our reading of a random sample of 30 audit reports, we see that "intermediate faults" are generally acts of mismanagement, while "serious faults" capture clear corruption; that is, instances of abuse of public office for private gain that were clearly

intentional, with cases of over invoicing, rigging of auctions to favor companies, or the use of "ghost" employees for embezzlement. However, this delineation is not always clear-cut. Thus, following other studies of the Brazilian audits such as Avis, Ferraz, and Finan (2018), we use the sum of intermediate and serious faults ("total faults") as a measure of corruption discovered by the auditors. In our sample of 1,108 audits, the mean number of total faults discovered is 58.[12]

To construct a measure of corruption revealed by the audits, we bucketed municipalities into terciles of corruption. As Figure 3.A shows, the mean total faults per service order varies widely by lottery, so we first de-meaned the total faults per service order by lottery to control for this variation. Next, we separated each service order into terciles, with the lowest tercile being the least corruption exposed and the top tercile being the most corruption exposed. Figure 3.B shows the resulting distribution of corruption across audited municipalities.

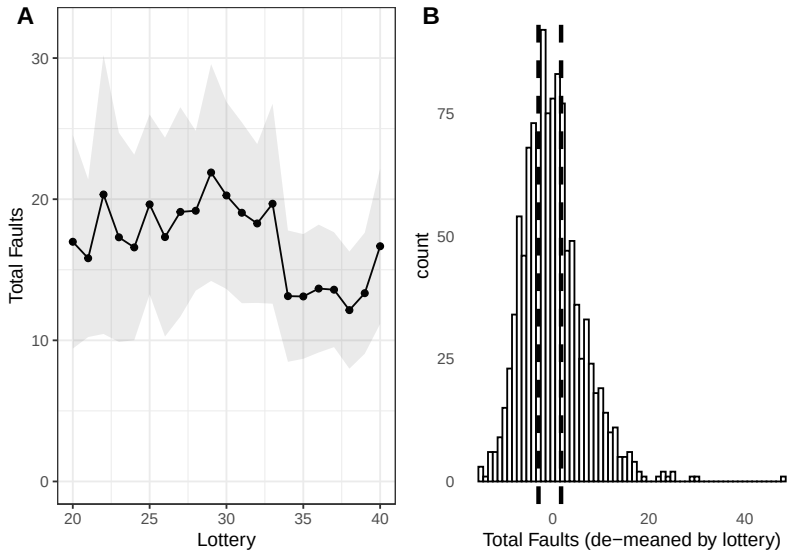

Figure 3: Distribution of total faults per service order. In A, the mean total fault number by lottery is plotted, with one standard error shown. In B, the dashed lines indicate the first and second terciles of the distribution.

[12]Note that our sample of 1,108 audits is less than the total number of audits that were undertaken in 2006-2015, for two reasons. The first is that we dropped audits for municipalities with missing data in our other measures, such as the budget. The second is that we only include the first audit for each municipality, which we discuss later in this section. Recall that the municipality selection process is done with replacement.

3.3 Audit Treatment Effect on Corporate Contributions

Our goal is to ascertain the effect of the random audits on corporate involvement in municipal elections. Given the random nature of the municipal selection for the audits, we consider municipalities that are audited for the first time *prior* to an election as the treatment group. We consider two sets of control groups. The first is all eligible municipalities that are still untreated in an election year. The second is our preferred specification, following Ferraz and Finan (2008), and comprises all municipalities that are treated for the first time between 2006 and 2015, but have not yet been audited in a given election year.[13] In this way, we compare the effect of the audits on municipalities that are treated *before* an election to the control, those that are treated *after* an election. When a municipality is treated for the first time, we drop them from comparison in the next election. Our sample includes 5,346 municipalities by the first definition of the control group and 1,087 municipalities by the second definition; this covers two elections, in 2008 and 2012.

We begin with a reduced-form model that estimates the average effect of a random audit on corporate contributions. We are interested in corporate contributions to a given mayoral race for the moment; in the next section, we delve into corporate behavior with respect to specific candidate types. The three outcomes of interest each capture a different aspect of firm engagement with a mayoral campaign:

1. **Total campaign donations to an electoral race.** This figure is summed across all firms who donate to a mayoral election and captures the total amount of money involved in the race. In all models run, we log this amount (cf. Boas, Hidalgo, and Richardson (2014)).
2. **Number of unique corporations donating to an electoral race.** This figure captures corporate competition within a mayoral campaign. Even if two firms donate to the same candidate, we would still like to consider that competition since the firms are now vying for influence over the same person. In all models run, we log this amount.
3. **Proportion of candidates who have received corporate donations in an electoral race.** In short, what fraction of a municipal election is backed by corporate money? This measure reflects the breadth of corporate involvement in a single mayoral election. Note that by definition this outcome is bounded [0,1].

Letting D_{msy} be the outcome of interest in municipality m, state s, and election year y, we estimate

$$D_{msy} = \alpha + \alpha_s + \alpha_\ell + \beta A_{msy} + X_{msy}\gamma + \epsilon_{msy}, \tag{1}$$

[13] Note that the second definition of the control group is a subset of the first definition.

where we control for state- and lottery-fixed effects (α_s and α_ℓ).[14] Our quantity of interest is β, since it captures the average effect of A_{msy}, which is an indicator variable that equals one when municipality m in state s is audited before the election in year y. Additionally, in some specifications we control for a vector of covariates X_{msy}, including whether the municipal-year race is incumbent-eligible; the log of the municipal population measured; an index of the municipality's rurality measured; and the municipality's literacy rate, following Ferraz and Finan (2008). The last three controls are time-invariant and measured in the 2000 census.

We must also control for the scope of the service order. Larger municipalities, for example, have larger budgets and therefore are usually found to have more irregularities. Consequently, the extent of auditing is always included in X_{msy}, across all model specifications. There are two alternatives to controlling for the extent of auditing; the first is to include the number of items audited in the regression and the second is to include the total value of items audited. In the models shown in the main text, we control for the total value of items audited, standardized to constant 2010 USD. However, we show in the Appendix that our main results hold when controlling non-parametrically for the number of items audited.

Table 2 shows the results of this first baseline check, across a variety of specifications. We find that the unconditional average effect of the audits is an increase in firm engagement with the ensuing election, as evidenced by the positive coefficients across all models shown. Corporations contribute about 3.9 times more in monetary value to mayoral races that are audited before the election compared to those that are not.[15] In the control group including only municipalities that will be treated in the future, but are not yet treated in the current election (i.e. models (3), (6), and (9) in Table 1), the average contribution is R$24,418. An increase of 4 times is therefore substantively relevant. In addition, model (6) shows that the number of unique corporations entering the race is 1.3 times greater in audited municipalities compared to not-yet-audited municipalities, resulting in a proportion of candidates that is about 11 percentage points higher.

This effect may be driven by a variety of factors. Firms may believe the audits are a signal about future economic conditions. Firms that are attracted to highly corrupt places may view the audits as free information about where to find those places. Or, firms may prefer not to engage in corrupt activity, and an audit portends clean transactions. We show that these possibilities are not likely to be driving the result. Corporate contributions in response to the audits are, somewhat surprisingly, invariant to the degree of corruption in the municipality.

[14] Lottery fixed effects are important because the program has changed over time. Some lotteries, especially after 2010, were targeted in specific areas. Furthermore, within a given lottery a municipality has the same probability of having been audited as other municipalities in the same state (Avis, Ferraz, and Finan 2018).

[15] For a coefficient of 1.35 in a log-linear model, the effect of a one-unit increase in the independent variable is an increase of $\exp(1.23) = 4$ times in the un-logged outcome variable.

We account for heterogeneous audit effects by using our measure of corruption based on the terciles of total faults. We operationalize the corruption tercile using a dummy variable, $C_{msy} \in \{1, 2, 3\}$, and include both its main effect as well as its interaction with the audit. In this way, we can separate the baseline effect of the audit from the additional effect of the degree of corruption in the municipality. In Table 3, we estimate for each of the corporate outcomes the regression

$$D_{msy} = \alpha + \alpha_s + \alpha_\ell + \beta_0 C_{msy} + \beta_1 A_{msy} + \beta_2 (C_{msy} \times \beta_1 A_{msy}) + X_{msy}\gamma + \epsilon_{msy}, \tag{2}$$

where β_2 now captures the conditional causal effect of the audit. Here, we focus on the sample that directly compares across the set of municipalities that are audited for the first time in 2006 to 2015, not the broader set of eligible municipalities, in line with Ferraz and Finan (2008).

Interestingly, there is no additional effect of higher levels of revealed corruption revealed on campaign contributions to audited municipalities. For the outcomes about total value of contributions and number of corporations, the baseline effect of audits persists and remains similar in magnitude to the estimates reported in Table 2. These effects are attenuated when incumbents are eligible for the race, as evidenced by the negative sign on the relevant coefficients across all three models, likely because incumbents do actually tend to enter the race: within the audited municipalities in our data, 73% of eligible incumbent mayors choose to run again in the next election, compared to 56% of eligible mayors in municipalities that are not audited before an election.

We address two potential concerns with our estimates. First, there is considerable heterogeneity in how different scholars choose to control for the variation in service order scope. In our main specifications, we have chosen to include the raw value of items audited in the municipality, expressed in constant 2010 USD. We believe that this way on controlling for service order scope is less biased than other choices, such as taking logarithms of the count of total items audited. We also choose a monetary value version of service order scope because it aligns well with our first outcome of interest, the monetary contributions made by corporations. To check that our estimates are robust to alternative measures of service order scope, we use a non-parametric measure of total items audited, as in Avis, Ferraz, and Finan (2018). Here, we create a dummy variable for every count of total faults. Using this alternative measurement, our results survive and in fact some of our estimates are even larger in magnitude. In Table 7, we show the results for the first two outcomes of interest.

Table 2: Average audit effect on corporate contributions to mayoral races. The unconditional average effect of audits is increased corporate contributions to the municipality, but these models do not show whether different levels of corruption revealed drive the result. Columns (3) and (6) are based on the sample of municipalities that were audited at least once in 2006-2011; other columns are based on all audit-eligible municipalities. All models include lottery and state fixed effects and control for the total items audited in the service order. Standard errors are clustered at the municipality level. See main text for definition of corruption and details about controls.

	Dependent Variable (by mayoral race):					
	Corporate Money in Race			N Corporations		
	(1)	(2)	(3)	(4)	(5)	(6)
Pre-election Audit	1.093*	1.232**	1.383**	0.197*	0.227**	0.263**
	(0.624)	(0.59)	(0.64)	(0.118)	(0.109)	(0.122)
Controls	N	Y	N	N	Y	N
Observations	15578	15578	1854	15578	15578	1854
R^2	0.059	0.227	0.1	0.06	0.284	0.11

*p<0.1, **p<0.05, ***p<0.01

Table 3: Main results for audit effect on corporate contributions to mayoral races. The baseline effect of audits is to attract corporations to the ensuing mayoral election, as measured across three different outcomes. More corporate campaign contributions (measured on a log scale in constant 2010 reals) and more unique corporations (measured in logs as well) flow to the electoral race. F-test p-values are derived from a joint significance test with null hypothesis that the sum of Pre-audit effect and the relevant interaction effect is zero. These tests suggest higher levels of corruption revealed by audits result in fewer corporations competing in a post-audit race, but it is unclear whehter that results in money in the race (note the conflicting signs on the interactions in Models (1) and (2)). All models include lottery and state fixed effects and control for the total amount audited. Standard errors are clustered at the municipality level. See main text for definition of corruption and details about controls.

	Dependent Variable (by mayoral race):			
	Corporate Money in Race		N Corporations	
	(1)	(2)	(3)	(4)
Pre-election Audit	1.69*	1.534*	0.337**	0.316**
	(0.692)	(0.661)	(0.128)	(0.119)
Moderate corruption	0.333	0.266	0.068	0.059
	(0.394)	(0.362)	(0.064)	(0.057)
High corruption	1.015**	0.361	0.199**	0.07
	(0.411)	(0.371)	(0.072)	(0.062)
Pre-election audit x moderate corruption	-0.08	-0.009	-0.042	-0.027
	(0.48)	(0.466)	(0.08)	(0.077)
Pre-election audit x high corruption	-0.083	0.024	-0.074	-0.05
	(0.488)	(0.466)	(0.082)	(0.076)
Incumbent-eligible		-0.861***		-0.134***
		(0.207)		(0.036)
F-test (Moderate corruption p-value)	0.027	0.021	0.019	0.009
F-test (High corruption p-value)	0.035	0.024	0.045	0.022
Controls	N	Y	N	Y
Observations	1854	1854	1854	1854
R^2	0.106	0.264	0.115	0.322

*p<0.1, **p<0.05, ***p<0.01

Second, there is a possibility that firms may not learn about the audits in time to change their contributions behavior. The effect of this possibility may be overstate our estimated coefficients, if we assume that firms would otherwise have been deterred from contributing to an audited municipality. This possibility is most likely to manifest with respect to audits undertaken and completed in the same year as the election. We therefore re-run our tests with municipalities that were audited in 2008 and 2012 excluded. Our results once again survive, as shown in Table 7.

3.4 Discussion of Audit Effects

We have causally identified a striking result: pre-election audits attract more corporate contributions to mayoral races. What is driving this effect? Since the effect we identify is aggregated by mayoral race, we explore two sources of variation, by the sector of the corporate donor and by the candidate type as discussed above.

We first address sectoral variation by measuring the proportion of municipal-level intake of money attributable to corporations from each sector in our data. If any particular sector is driving the audit effect, we might expect to see a large shift in that sector's share of total contributions to a municipality following an audit. Continuing to work within the context of our sample of 1,108 municipalities that are audited for the first time between 2006 and 2015, we compare the sectoral composition of corporate contributions to mayoral races in the election prior to an audit and the election after an audit. On the left hand side of Table 4, we report that there is not much change in which sectors account for the money flowing to these races. The columns denoted "Before t" refer to the set of races where the audit occurred before an election at time t. We compare sectoral inflows to the race in the election at $t-1$ to the election at t. There is not much change in sectoral composition. The only notable increase in the share of the pie is from construction firms, but this shift – almost the exact same percentage point value – also occurs in our placebo set of municipalities.

Our placebo set of municipalities is on the right hand side of Table 4. We examine the set of races where the audit occurs *after* the election, so that the audit should have no effect. We measure the change in sectoral composition of contributions *two* elections prior, $t-2$, to the election just prior to an audit, t. The similarity in the "Change" columns suggests that the sector-level interest in political investments is constant across audit status. On the one hand, this is fairly obvious, since sector interest in a mayoral race should be determined by municipal-level characteristics that are time invariance; for example, a municipality rich in natural resources will be attractive to a mining company with or without a pre-election audit. On the other hand, we confirm that the greater flows of corporate contributions to audited municipalities is not driven by any particular sector.

We repeat the exercise based on the composition of contributions towards candidate types. Based on the earlier discussion of political experience and incumbency, we focus on six types of candidates: political novices running in open seats and against incumbents, political veterans running in open seats and against incumbents, and incumbents who are novices and veterans. We define a political novice as a candidate who has held less than two previous political offices. A change in contribution composition towards or away from a specific candidate type may

	Audit Timing					
	Before t			After t		
Sector	t-1	t	Change	t-2	t-1	Change
Ag/Min	0.07	0.08	0.01	0.09	0.08	-0.01
Const.	0.29	0.33	0.04	0.26	0.30	0.04
HT Man.	0.12	0.09	-0.03	0.06	0.09	0.03
LT Man	0.20	0.20	0.00	0.25	0.19	-0.06
Trade	0.30	0.25	-0.06	0.32	0.31	-0.00
TUC	0.02	0.05	0.03	0.02	0.02	0.00

Table 4: Sector composition of corporate contributions to mayoral races. In the 'Before t' columns, we compare the proportion of contributions from each sector to municipalities audited prior to an election in the two elections around an audit. In the 'After t' columns, we compare the proportion for municipalities audited after an election in the two elections before the audit, as a placebo test for the effect of the audits on sectoral composition of corporate inflows.

indicate what is driving the audit effect, though because we are still looking at aggregate, composite figures, we will not yet be able to fully identify the cause. Table 5 shows a striking disparity. In municipalities that are audited before an election, proportionally more corporate contributions are dispensed to novice candidates running in open seats in the ensuing election, relative to the proportion in the election prior. Novices in open seats increase from accounting for 30% to accounting for 40% of corporate contributions to mayoral races following a random audit, even though two-thirds of such races are incumbent-eligible and, of those, incumbents run for re-election 73% of the time. Meanwhile, the money that flows to veterans challenging incumbents drops by 13 percentage points. A chi-squared goodness of fit test comparing the post-audit contribution amounts by candidate type to the proportions in the pre-audit election confirms the difference.

The right hand side of Table 5 shows the same candidate composition of corporate contributions for municipalities that are audited after an election, comparing across two elections prior and the election immediately prior. Interestingly, the pattern in this placebo set of municipalities is the opposite of the left hand side of the table. Corporations here reallocate from challengers running in open seats to challenging incumbents. A tentative inference to draw from this comparison is that audits shift corporations away from challenging incumbents, possibly because the incumbents that choose to run anyway following an audit are particularly high quality or because the audits showed that the incumbent was not too corrupt (Ferraz and Finan (2008)). Meanwhile, corporations might have disproportionately competed for the remaining open seats,

driving up the price of those seats.

	Audit Timing					
	Before t			After t		
Candidate Type	t-1	t	Change	t-2	t-1	Change
Incumbent Novice	0.04	0.07	0.03	0.09	0.06	-0.03
Incumbent Veteran	0.22	0.21	-0.01	0.16	0.28	0.12
Novice vs. OS	0.33	0.41	0.08	0.45	0.30	-0.15
Novice vs. Incum.	0.20	0.10	-0.10	0.13	0.12	-0.01
Veteran vs. OS	0.15	0.16	0.02	0.13	0.21	0.08
Veteran vs. Incum.	0.07	0.05	-0.02	0.04	0.04	-0.01

Table 5: Candidate type composition of corporate contributions to mayoral races. In the 'Before t' columns, we compare the proportion of contributions to candidate types in municipalities audited prior to an election in the two elections around an audit. In the 'After t' columns, we compare the proportion for municipalities audited after an election in the two elections before the audit, as a placebo test for the effect of the audits on sectoral composition of corporate inflows.

We can only speculate up to this point. Still, it is striking that the differences over time in candidate composition of corporate contributions are almost completely opposite for audited municipalities versus municipalities where there was no pre-election audit. The lack of difference in Table 4, on the other hand, suggests that the audit effect is driven by a pull effect (towards certain candidate types) and not a push effect (from certain sectors).

While the two composition comparisons we make shed some light on the drivers of the audit effect on corporate behavior, we are still dealing in aggregate behaviors. In order to more fully examine the question, we return to first principles: What are the corporations' utility functions? We ask whether the patterns observed in the changes in composition of corporate contributions induced by the random audits can tell us about the motivations for corporations' campaign contributions. To be consistent with the image of a future policy return-maximizing corporation, these shifts in allocations across political investments should represent an optimal tradeoff of risk and return.

We therefore require an explicit framework of risk and return to address the issue. This is the task of the next section.

4 Campaign Contributions Through the Lens of Portfolio Theory

We start from the premise that firms make campaign donations in order to maximize their return on investment while minimizing risk. Let the firm's portfolio of campaign contributions to n candidates be called p, and let r_p be defined as the total return on the portfolio. Formally, we define the corporation's utility in mean-variance terms:

$$U = \mathbb{E}(r_p) - \frac{\gamma}{2} Var(r_p) \tag{3}$$

where r_p is the return on the corporation's portfolio of donations across mayoral candidates in Brazil in a given year and γ is a risk tolerance parameter. Note that r_p is the weighted total of the return r on each asset i the corporation invests in. The corporate donor chooses a portfolio of campaign donations before every election, so the maximization problem is

$$max_{\{w_i\}} \mathbb{E}(r_p) - \frac{\gamma}{2} Var(r_p)$$
$$\text{s.t. } w_i \geq 0 \text{ and } \Sigma_{i=1}^{n} = 1$$

where w_i is the proportion of a corpopration's total donations during one election year allocated to candidate i (i.e. the weight). The expected return and variance on a portfolio of campaign donations is simply the weighted sum of the expected return and variance of its component candidates.

4.1 Defining Expected Returns

Following our discussion on corporate motivates to build long-term relationships with candidates, we can define expected return on a mayoral candidate as the corporation's expectation over the candidate's ability to return value in the immediate electoral term as well as in the future. Measuring the present value of the corporation's long-term projection about the candidate is difficult. We approach the problem by first measuring expected returns based on the projected return in the next mayoral term, then assess whether there is a systematic discrepancy that can be explained by the missing long-term component.

As previously discussed, rather than constructing expected returns on individual candidates, we define expected returns by candidate types. We posit that these candidate types, as defined

by the four key attributes introduced earlier in the paper (political experience, incumbency, presence of incumbent in the election, and pre-election audit), are differentially able to deliver value to their patrons. There are six basis candidate types, or "asset classes", which are then multiplied by two to take into account whether these candidates are participating in races that have just borne a municipal audit. We show the count of candidates per class in each election that we have pre-election audit data on in Table 6.

Asset Class	2008	2012
Incumbent Novice - Pre-Election Audit	61	89
Incumbent Novice - No Audit	712	735
Incumbent Veteran - Pre-Election Audit	182	125
Incumbent Veteran - No Audit	1844	1246
Novice for Open Seat - Pre-Election Audit	510	627
Novice for Open Seat - No Audit	5273	6142
Novice vs. Incumbent - Pre-Election Audit	320	264
Novice vs. Incumbent - No Audit	3528	2511
Veteran for Open Seat - Pre-Election Audit	79	141
Veteran for Open Seat - No Audit	867	1544
Veteran vs. Incumbent - Pre-Election Audit	51	86
Veteran vs. Incumbent - No Audit	457	725

Table 6: Number of candidates by asset class in each of the 2008 and 2012 mayoral elections.

To quantify returns, we use the municipal budget as our starting point. The budget is an attractive measure of returns for two reasons. First, it includes the federal transfers that firms covet and it is under the executive control of the mayor (Brollo et al. 2013). Second, it is relevant to the random audits because those audits check municipal budget expenditures for faults. We define the returns to a municipality as the total budget expenditures in a municipality during the term after the election. Formally, we define for municipality m in electoral term following the election year y the return $B_{m,y}$, where B is the municipal budget expenditures (in millions constant 2010 reals). The return on a candidate i of a type c is equal to $B_{m,y}$ multiplied by adummy indicating whether the candidate won. If there are n_y^c candidates in class c running in election year y, the conditional return on that class is

$$\mu_y^c = \frac{1}{n_y^c} B_{m,y} \mathbb{1}(Elected_{i,m})$$

We assume that the expected returns and risk on candidate classes is constant across the 2008

and 2012 elections. Suppressing the class-specific superscript, we define the expected return for a given asset class as

$$\mu = \frac{n_{08}}{n_{08}+n_{12}}\mu_{08} + \frac{n_{08}}{n_{08}+n_{12}}\mu_{12}. \tag{4}$$

The assumption that the first two moments of each asset class are fixed during the time period we consider means that we assume a corporation making investment decisions across the candidate types in 2012 does not learn from what transpired in 2008.[16] Instead, the corporation understands that candidate types are endowed with a certain ability to produce returns, and this is time invariant. We justify this assumption by making the observation that the two components of political expected returns are highly correlated over time. The probability of election by candidate type has a correlation across the two years of 99%. Similarly, the pairwise correlations of budget expenditures by municipality in 2004, 2008, and 2012 are all in excess of 97%. Furthermore, the correlation of budget returns by candidate class has a correlation of 96% across 2008 and 2012.

In Figure 4, we show the risk-return tradeoff on each asset class. Points more to the top-left are more attractive political assets, because the more to the top left a point is, the greater the expected return for a smaller amount of risk. Incumbent novices without a pre-election audit are the least attractive assets: although the incumbent novice with a pre-election audit has a lower return, the type without a pre-election audit is carries only slightly greater expected return but much more volatility.

Veteran candidate types are clustered at the top right hand corner of the figure, indicating that while they systematically generate higher expected returns than novice candidates, the variance on their returns is higher as well. The ratio of return to risk, however, is still in the veteran candidates' favor: while all novice-type candidates have a return to risk ratio less than 0.3, all veteran-type candidates clear this level. This is in line with Santos and Pegurier (2011)'s argument that successful federal deputies can bring their skills to bear in mayoral offices. Relatedly, Samuels (2002) and Leoni, Pereira, and Renno (2004) claim that running for mayor is an indicator of political ambition. The self-selection of politically experienced candidates into mayoral elections produces the sharp divergence seen in Figure 4.

We confirm these intuitions with a simple association between the budget return and candidate attributes. In Figure **??**, we regress the municipal-election-level budget return on attributes

[16] An explanation for how we constructed the covariance matrix of returns on asset classes can be found in the Appendix.

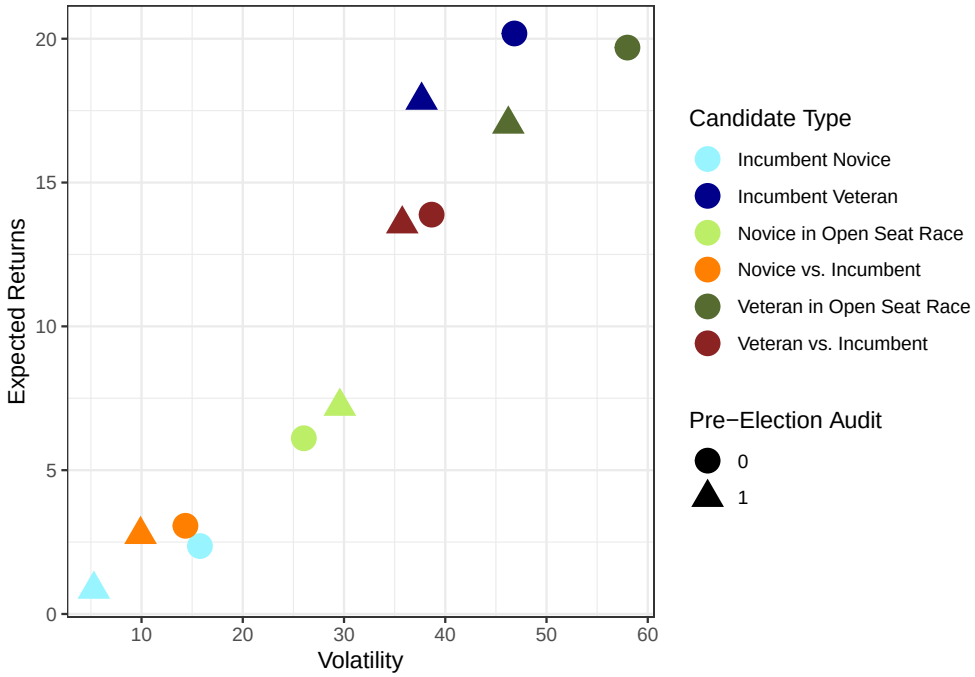

Figure 4: Risk-return tradeoff by asset class.

of the electoral winner, controlling for election and state fixed effects. We use the overall budget return, as well as specific subsets of the budget (agriculture and energy resources budget, housing and urban development budget, and the transportation and communications budget), as alternative outcome variables. Once again, we underscore the point that politically experienced candidates are the most attractive assets.

It is noticeable that incumbent candidates are associated with less in budget returns, because their expected returns (Figure 4) are so high. This apparent contradition is the result of two factors: incumbents are more likely to win, and incumbent veterans almost always win. Incumbent novices, however, almost never win. In 2008, no incumbent novices running after a pre-election audit won their race; in 2012, this figure improved to 18% of such candidates. Incumbency on its own, it seems, does not generate returns for corporate givers.

To bolster the validity of the association, we also show the relationship between progeny candidates and budget returns. Extant literature suggets that progeny candidates are less efficient in their expenditures, though in some categories, such as healthcare spending, Brazilian mayors who are progenies have been shown to spend more (Braganca, Ferraz, and Rios 2015; Dal Bó, Dal Bó, and Snyder 2009; Querubin and others 2016). We confirm the negative association between progeny types and expected returns.

4.2 Corporations Appear Rational...

One hypothesis is that if corporations are investing as maximizing agents, they should be directing most of their resources towards the assets with the highest return-risk ratios (or the lowest risk-return ratios). In this market, that is incumbent veterans who have been audited, incumbent veterans without an audit, and political veterans running in municipalities that have a pre-electio audit – both against an incumbent and in open seats. Is that the case?

Figure 5 generall confirms that firms invest optimally. Here, we show the average donation per candidate within each candidate class. Candidate types are separated into two graphs based on whether the race has seen a pre-election audit. Veteran challengers contesting open seats receive the most on average, followed by veteran incumbents, which is roughly in line with the expected return versus volatility trade-off. Furthermore, this figure supports the pattern found in Table 5. There we showed that, in municipalities where there is a pre-election audit, the ensuing corporate campaign inflows disproportionately shift toward veteran challengers in open seats, while Figure 5 confirms that not only is the aggregate figure large, but the per-candidate

contribution dwarfs the contributions to other candidate classes in post-audit elections as well.

In fact, the patterns in Figure 5 support our earlier conjecture, based on Table 5, that the main driver behind the audit effect on corporate campaign contributions is that corporations invest relatively less in races involving incumbents, and relatively more in open seat races (compared to similar races in non-audited municipalities). Comparing across 5.A and 5.NA, we find that the only asset classes to which corporations give more per candidate when it is a post-audit setting rather than a no-audit setting are the two classes involving open seats: Novices running for an open seat (light green) and Veterans running for an open seat (dark green).

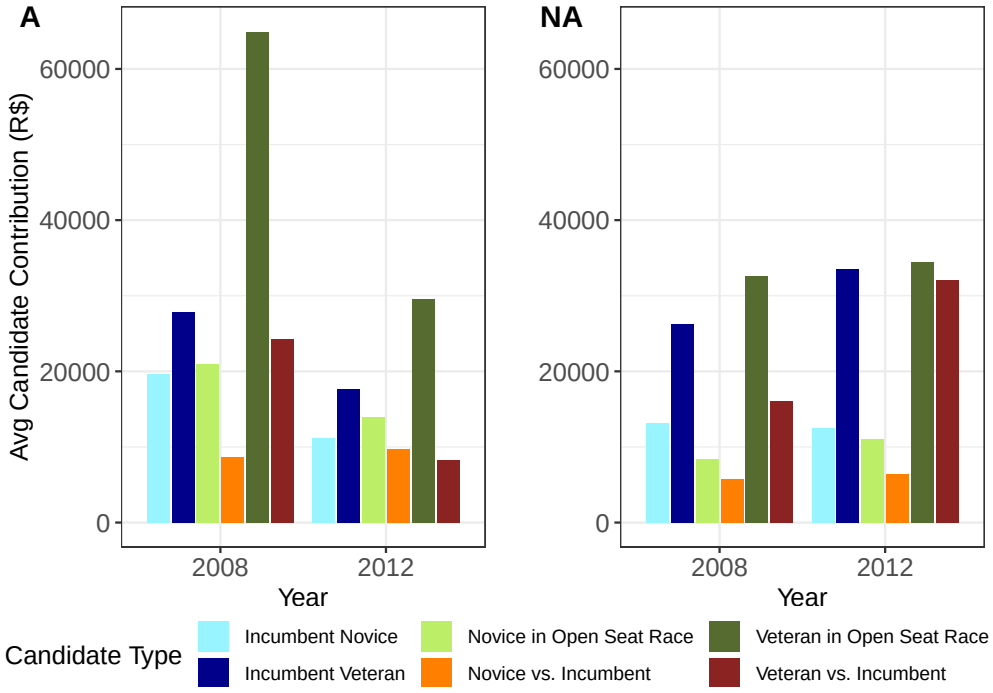

Figure 5: Within the set of municipalities that have been audited, we compare the average donation to a candidate of each asset class among municipalities that were audited pre-election (A) to the average among municipalities that were not audited pre-election (NA).

Corporations are not necessarily vying against each other by backing different assets in these races, either. The mean number of corporations involved in an open-seat race that has been audited prior to the election is 3.2, but the number of candidates backed by corporate money is on average just 1.2 in these races. In addition to consistency with our earlier discussion about the composition of contributions by candidate type, the fact that firms in post-audit races target novices running for open seats aligns with another pattern in the risk-return relationship across candidate types. In Figure 4, we see that the plotted triangles are almost all to the left of and a below their twins in circle. This means that a candidate class has less expected return in the

audited municipalities than in the un-audited variety. The only exception is the light green pair of triangle and circle: here, novice in open seats are better investments when their municipality has been audited rather than not.

Interestingly, the higher returns expected from novice in pre-audited, open seat races are not driven by an increased probability of winning. We measured the proportion of candidates within each asset class who won their mayoral race. The relevant comparison to make is between pairs of similar candidate types across audited and non-audited varieties. We report that 32% of novices in open-seat races won their race, no matter if there was a pre-election audit or not.

5 . . . But May Not Be Completely Maximizing Their Returns?

The puzzle of corporate contributors' affinty for the novice candidates running in open seats after an audit calls into question whether corporations are behaving optimally with respect to their campaign contributions. We show using standard portfolio theory techniques that, based on our definition of expected returns and risk, corporations do not seem to be fully optimizing over their utility functions. However, we do not consider this a death knell for the view of corporations as political investors. We discuss a few reasons corporations might not appear to be maximizing under our definition, the most important of which is that we might be systematically understating the expected returns.

Under modern portfolio theory, there is a set of optimal allocations across asset types that varies depending on the investor's risk tolerance. Each optimal allocation, which is found using quadratic programming to solve the problem in Equations (4)-(5), has a portfolio expected return and variance. Plotting these portfolios in return-risk space yields the efficient frontier. All portfolios to the right and downwards of the efficient frontier in resturn-risk space are, by definition, inefficient. This is because for all points within the efficient frontier, a lower volatility can be achieved for the same expected return (by traveling horizontally to the efficient frontier), or, equivalently, a higher epected return can be achieved for a lower expected return (by traveling vertically to the efficient frontier).

We plot the portfolios of only those investors who are involved in the asset classes "Veteran-Open Seat-Audit" and "Novice-Open Seat-Audit", since we are interested in assessing whether this response to the random audits still falls under the rubric of the optimizing political investor. Figures 6 and 7 seems to suggest not. There also does not seem to be a systematic pattern across sectors in terms of which sectors' allocations lie closest to the optimal frontier.

There are a variety of reasons that firms may appear to be suboptimal. The first possibility is that

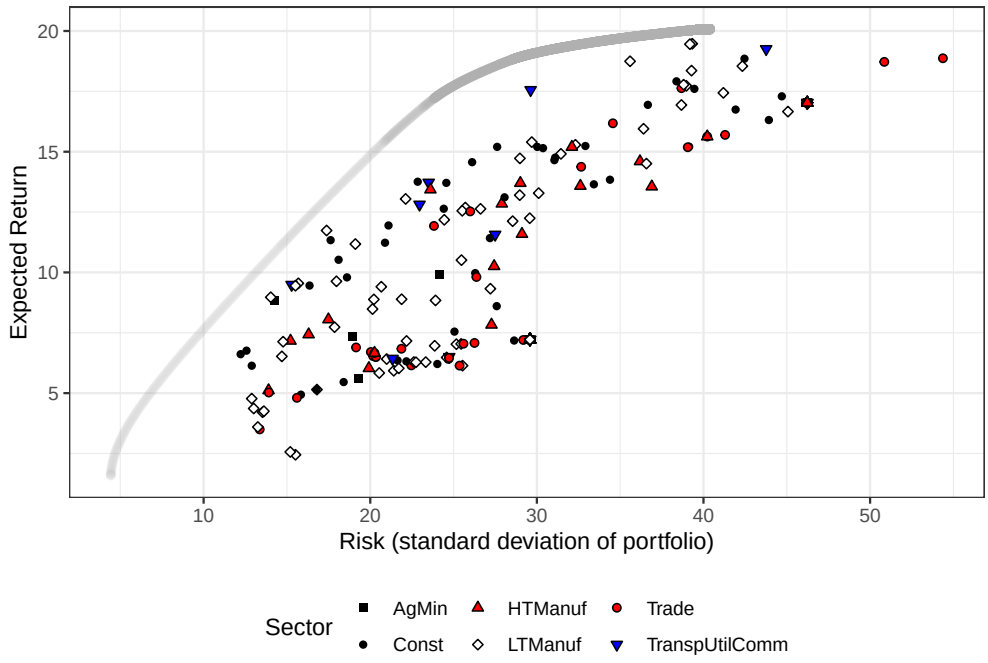

Figure 6: Actual portfolios of 608 corporate contributors who donated at least partially to open seat races in municipalities following an audit prior to the 2008 election. The efficient frontier is traced in gray. Because the observed portfolios all lie to the left and downward of the frontier, they are all considered inefficient.

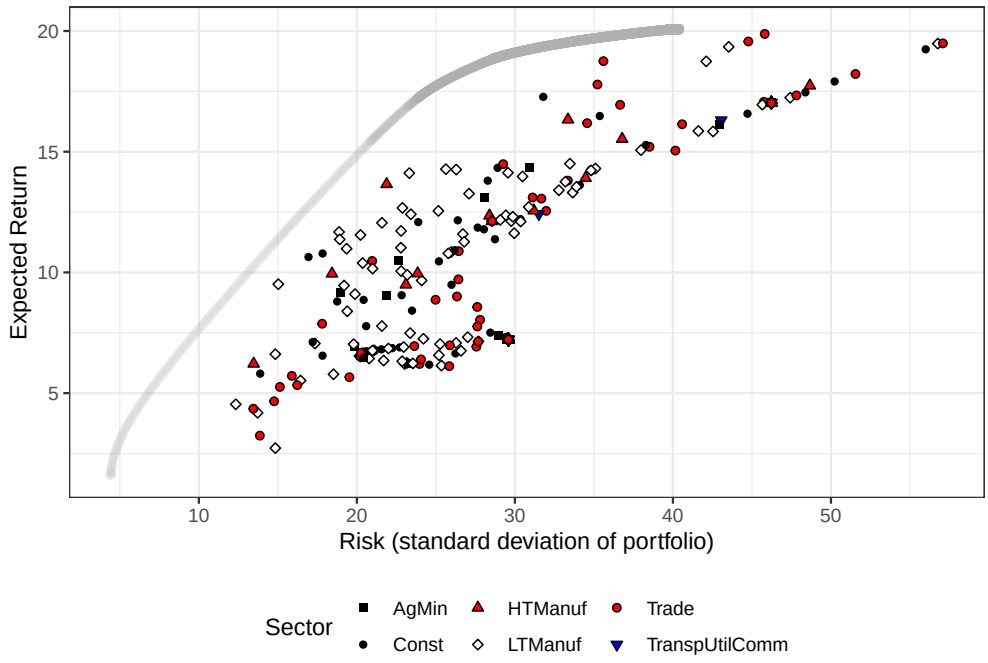

Figure 7: Actual portfolios of 896 corporate contributors who donated at least partially to open seat races in municipalities following an audit prior to the 2012 election. Again, the portfolios are all deemed inefficient. High-tech manufacturing and trade firms appear to get the closest to the frontier.

we may have simply measured the wrong type of returns. Firms may still be optimizing, just over some other good. For example, Besley and Mueller (2018) shows that firms in developing countries prioritize security of physical assets. Brazilian corporations may favor protection given that the crime rate in their country is extremely high.

The second possibility is that the information environment in Brazilian local politics is not sufficient to apply financial portfolio theory. The property of efficient asset prices means mandates that prices must reflect all available information by all market participants. In a weakly institutionalized democracy, it is not clear to what extent this should hold in the context of elections. However, we do measure the response to a common shock in the form of the random audits. Studying corporate contributions in the context of the anti-corruption program, which is known to be well-publicized whenever audit reports are released, should hopefully overcome the information concern. The information problem may also bite in the other direction, where some firms learn about electorally relevant data before others do.

We recognize that we cannot fully overcome these concerns. Nevertheless, we run one final exercise to adjudicate between the possibilities that firms are or are not investment optimizers. Instead of taking our metric for budget returns as given and assessing the portfolios of firms with respect to the efficient frontier, we invert the question. Suppose that all firms lie on the efficient frontier. What would expected returns per asset class need to be for this to be true?

We posit that if our measure of expected returns is completely, utterly off-base, then the required expected returns in order to make all firms' portfolio allocations optimal will be unrecognizable compared to Figure 4. The ordinal position of asset classes will be jumbled compared to our estimates. Furthermore, the cardinal values of the expected returns should be significantly different.

Following the methodology of Best and Grauer (1985), we estimate the returns required to generate universal optimality across firms' actual portfolios. We do this for the portfolios observed in 2012 only. The full methodology is explained in the Appendix.

Figure 8 shows the result of the simulation. While the magnitudes of asset class returns are greater than our estimates across the board, we also note that the relative magnitudes across asset classes, as well as the rank orderings of asset classes, is largely preserved. The two notable exceptions are the novice candidates competing in open seat races and against incumbents in municipalities without a pre-eleciton audit. The slight deviations from what would be consistent with optimal behavior have all centered on corporations' contributions to novice-type candidates in open seat races. It is clear that a better understanding of how political strangers come to enter Brazilian politics, and how corporations decide to support them.

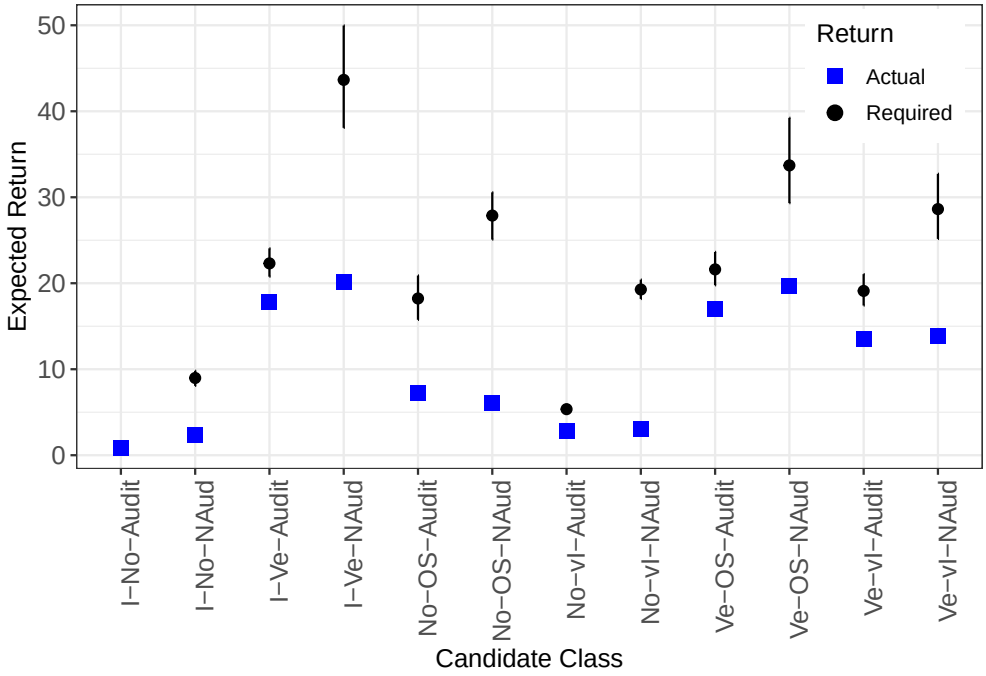

Figure 8: Required expected returns given observed portfolio allocations in 2012 election, such that all corporations are investing optimally across candidate types. Abbreviations are I for Incumbent, Ve for Veteran, No for Novice, OS for Open Seat race, vI for versus Incumbent, Audit for pre-election Audit, and NAud for No pre-election Audit. The required returns are very different from our estimated returns in magnitude, but their rank orderings and relative sizes to each other are not far off from that of our estimates.

What is missing for certain from our measure of expectd returns is the long-term component of returns that firms seek when they give to campaigns. The gaps between our estimated returns and the required returns for all firms to be simultaneously optimizing are systematically greater for those candidates in un-audited races, which may be capturing some of that missing component. Further investigation would be required to tease out these relationships.

6 Conclusion

We have shown that political campaign contributions from corporations are responsive to Brazil's anti-corruption audit program. The audits have a beckoning effect on corporate giving. This effect mainly mainfests in increased contributions to a few specific types of candidates, especially relatively politically inexperienced candidates who run in open seat races in the wake of audits.

Our evidence is suggestive of an investment-motivated corporate contributor. However, our measure of expected returns suffers from its inability to capture the long-term relationship that Brazil scholars such as Bezerra (1999) and Samuels (2001b) have indicated is integral to the corporation's utility. We have proposed on way to capture this missing component of expected returns, by imputing the required returns to make all firms appear as mean-variance utility maximizers, but we recognize that this is just a first step towards understanding how corporate interests interact with the Brazilian audit program. A more complete accounting of the factors that drive the audit effect we see will also help scholars think about how corporate influence might be wielded in the future, given the 2014 prohibition on corporate campaign contributions.

7 Appendix

7.1 Robustness Check on Main Audit Effect

	Dependent Variable (by mayoral race):			
	Corporate Money in Race		N Corporations	
	(1)	(2)	(3)	(4)
Pre-election Audit	4.778***	6.128***	3.315***	4.628***
	(1.142)	(1.191)	(1.121)	(1.229)
Moderate corruption		0.43		0.43
		(0.375)		(0.38)
High corruption		0.522		0.522
		(0.397)		(0.398)
Pre-election audit x moderate corruption		-1.044		-1.044
		(0.536)		(0.565)
Pre-election audit x high corruption		-1.455		-1.455
		(0.637)		(0.681)
Incumbent-eligible		-0.578**		-0.891***
		(0.231)		(0.239)
F-test (Moderate corruption p-value)	0.026	0.228	0.026	0.228
F-test (High corruption p-value)	0.017	0.266	0.017	0.266
Controls	N	Y	N	Y
Observations	1854	1854	1671	1671
R^2	0.189	0.343	0.182	0.339

*p<0.1, **p<0.05, ***p<0.01

Table 7: Robustness check for main results of audit effect on corporate contributions to mayoral races. In Models (1) and (2), the extent of the service order is controlled for non-parametrically by count of total items audited rather than a continuous value audited, and election year fixed effects are added. In Models (3) and (4), we drop all observations from lotteries in 2008 and 2012 to mitigate information effects. Here, we continun to use the non-paramteric control for extent of service order. Results from the main text still hold with these alternative specifications, although the joint significance tests are no longer statistically significant. All models still include lottery and state fixed effects. Standard errors are clustered at the municipality level. See main text for definition of corruption and details about controls.

7.2 Selecting the Predictors of Electoral Success

We use the least absolute shrinkage and selection operator (lasso) to reduce our laundry list of potential predictors of electoral success. Letting the outcome $Y = 1$ be where candidate i is elected, we have

$$P(Y=1|x_i) = \frac{\exp(\beta_0 + \beta_1 + x_{i1} + \cdots + \beta_k x_{ik})}{1 + \exp(\beta_0 + \beta_1 + x_{i1} + \cdots + \beta_k x_{ik})},$$

where $x_i = (x_{i1}, x_{i2}, \ldots, x_{ik})$ are the potential predictors and β_j is the coefficient corresponding to each of the k predictors. We run a logistic lasso estimator that finds $\hat{\beta}_0, \hat{\beta}_1, \ldots, \hat{\beta}_k$ that miniizes the negative log likelihood:

$$\Sigma_{i=1}^{n}[-y_i(\beta_0 + \beta_1 x_{i1} + \cdots + \beta_k x_{ik}) + \log(1 + \exp(\beta_0 + \beta_1 x_{i1} + \cdots + \beta_k x_{ik}))] \tag{5}$$

subject to $\Sigma_{j=1}^{k}||\beta_j||_1 \leq \lambda$. Here, $\lambda > 0$ is a parameter that controls how sparse the estimator will be. We select λ by using ten-fold cross-validation. We found that $\lambda = 0.002$ yielded the minimum mean cross-validated error. Taking the predictors with $\beta_j > 0$ at that value of λ, we then ran a linear probability model on this sparse set of covariates to yield the estimates seen in Figure 1.

7.3 Estimating the Covariance Matrix for Mayoral Expected Returns

Here we detail the derivation of the expected returns covariance matrix.

First, we define the variance of an asset class. Let n_y be the number of candidates who fall into an asset class in year y and let μ_y be the mean return in year y as defined in the main text. Let σ_y^2 be the within-class variance in returns for the election year y.

Recall that we defined the expected return of an asset class in the case with two time periods a, b as $\mu = \frac{n_a}{n_a+n_b}\mu_a + \frac{n_b}{n_a+n_b}\mu_b$. Now let $y \in \{08, 12\}$, our two time periods. By the law of total variance we can proceed:

$$\sigma^2 = \mathbb{E}[Var(\mu|y)] + Var(\mathbb{E}[\mu|y]) \tag{6}$$

$$= \left[\frac{n_{08}}{n_{08}+n_{12}}\sigma_{08}^2 + \frac{n_{12}}{n_{08}+n_{12}}\sigma_{12}^2\right] + \left[\frac{n_{08}}{n_{08}+n_{12}}(\mu_{08}-\mu)^2 + \frac{n_{12}}{n_{08}+n_{12}}(\mu_{12}-\mu)^2\right] \tag{7}$$

$$= \frac{n_{08}\sigma_{08}^2 + n_{12}\sigma_{12}^2}{n_{08}+n_{12}} + \frac{n_{08}n_{12}}{(n_{08}+n_{12})^2}(\mu_{08}-\mu_{12})^2 \tag{8}$$

where the final expression takes some algebra.

Next, we define the off-diagonals of the covariance matrix. Let μ_{cy} be the return on asset class c in election year y. Following Litterman and Winkelmann (1998) (pp.12-13), we define

the off-diagonals of the covariance matrix between asset classes as $\sigma_{ij} = [(w_{12}\mu_{i,12}\mu_{j,12}) + (w_{08}\mu_{i,08}\mu_{j,08})]/(w_{08} + w_{12})$, where w_y denotes the total number of candidates in the classes i and j. We loop through i and j for all ten asset classes (but use the variance as defined above to fill in the diagonals).

7.4 Simulating Returns Required for Universal Optimality

We estimate here the required expected returns on each "asset type" (i.e., candidate type) such that all corporate contributors are behaving optimally. The goal of this exercise is to place some bounds on how plausible municipal budgets are as the good that firms maximize over when they make campaign contributions. If, for example, the estimated required returns are orders of magnitude higher than the estimated counterparts in the main text, we would question whether such a high level of returns is reasonable. On the other hand, if the required returns we simulate are actually very similar to our budget-estimated figures, then it makes no sense to speak of suboptimal behavior as measured by our estimates, since the observed portfolio weights are compatible with an equally plausible world where firms are universally investing optimally.

The estimation approach takes the observed portfolio allocations for the set of corporations in an electoral cycle as given and backs out what the expected returns must have been in order to observe those portfolio allocations. We follow the approach of Best and Grauer (1985) and make two simplifying assumptions in order to make the problem tractable. First, we maintain the assumption that all firms are optimizing over one dimension, as in the main text of the paper. This is undoubtedly a strong assumption. Heterogeneity by industry sector, for example, could be a reasonable first step in constructing multiple objective functions. Nevertheless, we believe that this assumption does not do too much violence to reality, as different objects of optimization could conceivably be decomposed into a single dimension. Our estimation here considers only a single set of two moments (mean and variance) for each asset class, commonly known to all firms.

Second, we take the covariance matrix estimated in the paper as given. This assumption means that we hold fixed the correlations in expected returns across candidate types. If we have mis-specified the covariance matrix, then this exercise is doomed from the start. However, this assumption is indispensable given the structure of mean-variance portfolio analysis. Maximizing expected returns is meaningless without understanding the costs incurred by return variance; minimizing return variance is meaningless without understanding the tradeoff with potential gains.

The method of Best and Grauer (1985) is suitable only for estimating expected returns for one

time period, so we use observables from the 2012 municipal elections. Let the observed set of weights across the $m = 10$ candidate types discussed in the paper be denoted ($\mathbf{x}_{mt}$, where $t = 2012$. The covariance matrix of returns is Σ. Let θ_{1t} be the return on a "zero-beta" asset, which in this context means the return on an asset that is perfectly uncorrelated with the municipal mayors political market in 2012. Finally, define θ_{2t} as the investor's risk tolerance. Best and Grauer (1985) establish that, given the two knowns ($\mathbf{x}_{mt}$ and Σ) and two unknowns (θ_{1t} and θ_{2t}), the required expected returns that render the observed weights simultaneously and individually optimal are

$$\mu_t = \theta_{1t}\iota + \theta_{2t}\Sigma\mathbf{x}_{mt}, \tag{9}$$

where ι is a vector of ones the same length as the covariance matrix.

Now all that remains is to estimate the zero-beta rate θ_{1t} and the investor risk tolerance θ_{1t}. Best and Grauer (1985) (p. 92) find that the market risk-free rate is a suitable proxy for the zero-beta rate, so we use the 5-year U.S. Treasury Bond yield around October 2012, the month of the municipal elections. For the risk tolerance parameter, we simulate required returns based on a range of risk tolerance levels compatible with the efficient frontier estimated from budget returns in the main text. We also simulate required returns on a range of risk-free rates, randomly sampled from days in 2012. Finally, we adopt a bootstrap approach to estimate the equation above: we randomly sample with replacement, 500 times, 1000 rows from the matrix of observed portfolio weights and apply each iteration to the full range of risk-free and risk tolerance parameters.

The result of the bootstrap estimation is a 500 by 10 matrix of required expected returns. We generate a point estimate and 95% confidence intervals by using the distribution of returns on each candidate class. This produces Figure 8.

References

Arretche, Marta TS. 1999. "Social policies in brazil: Decentralization in a federal state." *Brazilian journal of social sciences* 14(40): 111–141.

Avis, Eric, Claudio Ferraz, and Frederico Finan. 2018. "Do government audits reduce corruption? Estimating the impacts of exposing corrupt politicians." *Journal of Political Economy* 126(5): 1912–1964.

Baiocchi, Gianpaolo. 2006. "Inequality and innovation: Decentralization as an opportunity structure in brazil." *Decentralization and local governance in developing countries* 2006: 53–80.

Bersch, Katherine, Sergio Praca, and Matthew M Taylor. 2017. "State capacity, bureaucratic politicization, and corruption in the brazilian state." *Governance* 30(1): 105–124.

Besley, Timothy, and Hannes Mueller. 2018. "Predation, protection, and productivity: A firm-level perspective." *American Economic Journal: Macroeconomics* 10(2): 184–221.

Best, Michael J, and Robert R Grauer. 1985. "Capital asset pricing compatible with observed market value weights." *The Journal of Finance* 40(1): 85–103.

Bezerra, Marcos Ot 'a vio. 1999. *In the name of the "bases": Politics, favor and personal dependence*. Relume Dumar 'a.

Boas, Taylor C, F Daniel Hidalgo, and Neal P Richardson. 2014. "The spoils of victory: Campaign donations and government contracts in brazil." *The Journal of Politics* 76(2): 415–429.

Braganca, Arthur, Claudio Ferraz, and Juan Rios. 2015. "Political dynasties and the quality of government." *Unpublished manuscript*.

Brambor, Thomas, and Ricardo Ceneviva. 2011. "Incumbency advantage in brazilian mayoral elections." In *APSA 2011 annual meeting paper*,

Brollo, Fernanda et al. 2013. "The political resource curse." *American Economic Review* 103(5): 1759–96.

Colonnelli, Emanuele, and Mounu Prem. 2017. "Corruption and firms: Evidence from randomized audits in brazil." *Available at SSRN 2931602*.

Colonnelli, Emanuele, Mounu Prem, and Edoardo Teso. 2018. "Patronage and selection in public sector organizations." *Available at SSRN 2942495*.

Dal Bó, Ernesto, Pedro Dal Bó, and Jason Snyder. 2009. "Political dynasties." *The Review of Economic Studies* 76(1): 115–142.

Dix-Carneiro, Rafael. 2014. "Trade liberalization and labor market dynamics." *Econometrica*

82(3): 825–885.

Falleti, Tulia G. 2010. *Decentralization and subnational politics in latin america*. Cambridge University Press.

Ferraz, Claudio, and Frederico Finan. 2011. "Electoral accountability and corruption: Evidence from the audits of local governments." *American Economic Review* 101(4): 1274–1311.

Ferraz, Claudio, and Frederico Finan. 2008. "Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes." *The Quarterly journal of economics* 123(2): 703–745.

Fouirnaies, Alexander, and Andrew B Hall. 2014. "The financial incumbency advantage: Causes and consequences." *The Journal of Politics* 76(3): 711–724.

Kim, Galileu, and Romain Ferrali. 2019. "Are anti-corruption audits effective? Evidence from brazilian municipal bureaucracies." *Evidence from Brazilian municipal bureaucracies (December 5, 2019)*.

Klavsnja, Marko, and Rocio Titiunik. 2017. "The incumbency curse: Weak parties, term limits, and unfulfilled accountability." *American Political Science Review* 111(1): 129–148.

Kroszner, Randall S, and Thomas Stratmann. 2005. "Corporate campaign contributions, repeat giving, and the rewards to legislator reputation." *The Journal of law and Economics* 48(1): 41–71.

Kroszner, Randall S, and Thomas Stratmann. 2000. *Does political ambiguity pay? Corporate campaign contributions and the rewards to legislator reputation*. National bureau of economic research.

Leoni, Eduardo, Carlos Pereira, and Lucio Renno. 2004. "Political survival strategies: Political career decisions in the brazilian chamber of deputies." *Journal of Latin American Studies* 36(1): 109–130.

Litterman, Robert, and Kurt Winkelmann. 1998. "Risk management series."

Mainwaring, Scott. 1998. "Party systems in the third wave." *Journal of democracy* 9(3): 67–81.

Markowitz, Harry. 1952. "The utility of wealth." *Journal of political Economy* 60(2): 151–158.

McCarty, Nolan, and Lawrence S Rothenberg. 1996. "Commitment and the campaign contribution contract." *American Journal of Political Science*: 872–904.

Meireles, Fernando. 2019. "Political careers in the house of representatives: A quasi-experimental analysis." *DADOS: Revista de Ciencias Sociais* 62(4).

Nölke, Andreas, and Arjan Vliegenthart. 2009. "Enlarging the varieties of capitalism: The emergence of dependent market economies in east central europe." *World politics* 61(4): 670–702.

Poole, Keith T, and Thomas Romer. 1985. "Patterns of political action committee contributions to the 1980 campaigns for the united states house of representatives." *Public Choice* 47(1): 63–111.

Prado, Sergio. 2001. "Transfer tax rates and municipal financing in brazil." *Research Report Decentralizes Fiscal and Cooperation Intergovernmental Financial. EBAP / Funda Konrad Adenauer*.

Querubin, Pablo, and others. 2016. "Family and politics: Dynastic persistence in the philippines." *Quarterly Journal of Political Science* 11(2): 151–181.

Samuels, David. 2003. *Ambition, federalism, and legislative politics in brazil*. Cambridge University Press.

Samuels, David. 2001a. "Money, elections, and democracy in brazil." *Latin American Politics and Society* 43(2): 27–48.

Samuels, David. 2001b. "Money, elections, and democracy in brazil." *Latin American Politics and Society* 43(2): 27–48.

Samuels, David. 2002. "Progressive ambition, federalism, and pork-barreling in brazil." *Legislative Politics in Latin America* 315: 340.

Santos, Fabiano GM, and Fabiano JH Pegurier. 2011. "Political careers in brazil: Long-term trends and cross-sectional variation." *Regional and Federal Studies* 21(2): 165–183.

Schneider, Ben Ross, and others. 2004. *Business politics and the state in twentieth-century latin america*. Cambridge University Press.

Snyder Jr, James M. 1990. "Campaign contributions as investments: The us house of representatives, 1980-1986." *Journal of Political Economy* 98(6): 1195–1227.

Stratmann, Thomas. 2005. "Some talk: Money in politics. A (partial) review of the literature." In *Policy challenges and political responses*, Springer, p. 135–156.

Titiunik, Rocio. 2009. "Incumbency advantage in brazil: Evidence from municipal mayor elections." *University of California-Berkeley, mimeo*.

Treisman, Daniel. 1998. "Dollars and democratization: The role and power of money in russia's transitional elections." *Comparative Politics*: 1–21.