Appendices

A Inferring Firm Effectiveness from Electoral Outcomes

To organize and interpret the empirical results, we develop a model in this section of learning about firms' quality from electoral performance. We assume that candidates and parties cannot directly determine firms' effectiveness; they must attempt to learn the ability of a firm by observing the performance of the consultant's clients in elections.

We model this learning process as a sequence of Bayesian updating steps. The observer begins with a prior belief about the firm's quality and then observes a sequence of elections in which the consulting firm participates. Following each election, the observer updates its belief to incorporate the new information. The observer believes that the process governing realizations of the consultant's clients' election outcomes is as follows:

$$s_j^D = \Lambda(\alpha_j + q + \epsilon_j) \tag{1}$$

where s_j^D is the vote share of the Democratic candidate in election j; Λ is the logistic function: $\Lambda(x) = \frac{e^x}{1+e^x}$; α_j is an election-specific fixed effect, which is known to the observer;[1] q is the consultant's effectiveness, or quality; and ϵ_j is an independent and identically distributed error term drawn from a normal distribution. Electoral performance in this world is determined by a combination of the match between candidate and electoral district, an additive "boost" generated by the consulting firm's participation, and an unpredictable component. The challenge for the observer is to separate the consultant's contribution from the noise.

The observer forms beliefs about the consultant's quality by combining observations of election outcomes with a prior belief, which we will take to be:

$$g_0(q) \sim N(0, \sigma_q^2) \tag{2}$$

The choices of parametric and functional forms here are made for simplicity and ease of simulation. The specific functional forms we selected are not essential for the qualitative predictions discussed below to hold; what is important is that election results have some unpredictable component unrelated to consultant quality.

Learning from vote shares Suppose the observer observes vote share outcomes in each election. Then, given Bayesian updating, the observer's posterior belief follows the process defined by:

$$g_j = \frac{\phi\left(\log\left(\frac{s_j}{1-s_j}\right) - \alpha_j - q\right) g_{j-1}}{\int \phi\left(\log\left(\frac{s_j}{1-s_j}\right) - \alpha_j - q\right) g_{j-1}\, \mathrm{d}q} \tag{3}$$

where ϕ is the normal probability density function. Beginning with a prior belief, the observer sequentially updates to incorporate new information gleaned from election results. This generates a sequence of beliefs which, as the number of elections n grows large, converges onto the true value of q.

Learning from wins Suppose instead that the observer discards the vote share observation, and instead uses only the information contained in whether the firm's client won or lost. Letting $w_j = 1$ if the client wins and $w_j = 0$ otherwise, the posterior beliefs are:

$$g_j = \frac{\left(1 - \Phi\left(-\alpha_j - q\right)\right)^{w_i} \Phi\left(-\alpha_j - q\right)^{1-w_i} g_{j-1}}{\int \left(1 - \Phi\left(-\alpha_j - q\right)\right)^{w_i} \Phi\left(-\alpha_j - q\right)^{1-w_i} g_{j-1}\, \mathrm{d}q} \tag{4}$$

where Φ is the normal cumulative distribution function.

Simulation To illustrate how censoring the observation to a binary outcome reduces the informativeness of election results, we conducted a simulation of the processes defined by Equations 3 and 4. We generated 100 observations of election results for several values of α_j, representing more or less predictable elections: $\alpha_j \in \{0, 0.25, 1\}$. These values generate expected Democratic vote shares of 50 percent, 56 percent, and 73 percent, respectively. We then computed posterior beliefs following each realization using either the vote share or the binary win-loss outcome.[2]

Figure 1 clearly shows the consequences of censoring on the quality of information revelation. When elections are very competitive ($\alpha = 0$), learning from wins is almost as good as using the full vote share information - both processes fairly quickly approach the true quality value of 0.1. But once elections become even slightly tilted in favor of one candidate

($\alpha = 0.25$) the speed of learning in the censored process falls dramatically. In very uncompetitive races ($\alpha = 1$), there is essentially no information to be gained from observing wins, and the wins-only process never departs from the prior.

Notably, the real-world electoral environment in Congressional elections is much closer to the $\alpha = 1$ case than to the more competitive cases, particularly in elections involving an incumbent candidate. This observation implies that, if observers pay attention only to wins and not vote shares, elections involving Congressional incumbents are likely to be almost completely uninformative about the quality of the consulting firms involved. A consulting firm with sufficiently established reputation to be capable of selectivity in its choice of clients could take advantage of this fact by choosing to work only for incumbents. Such a strategy would essentially halt the learning process, locking in the firm's reputation indefinitely and preventing clients from learning any new, potentially unfavorable, information about the firm's ability.

B Robustness

We estimated several alternative specifications of our productivity regressions to check their robustness. We present each variant and describe the results in this section.

Linear probability model To guard against the possibility that the distinction we observe between residuals in terms of vote share and in terms of win probability is due to estimating the former using a linear model and the latter using a generalized linear model, we estimated linear probability models of client wins, and re-ran the exercise reported in Tables 7 and 8 of the article. The results, shown below in Tables 1 and 2, are substantively the same as the GLM version presented in the main results section.

Stratification on competitiveness We constructed our vote share residuals separately within relatively competitive races (those rated "Tossup," "Lean Republican," and "Lean Democrat" by Cook) and relatively safe races (the remainder), and then estimated the market response model of Table 7 of the article separately for each. The purpose of this exercise was to allow for different market responses to unexpected vote share in close versus uncompetitive

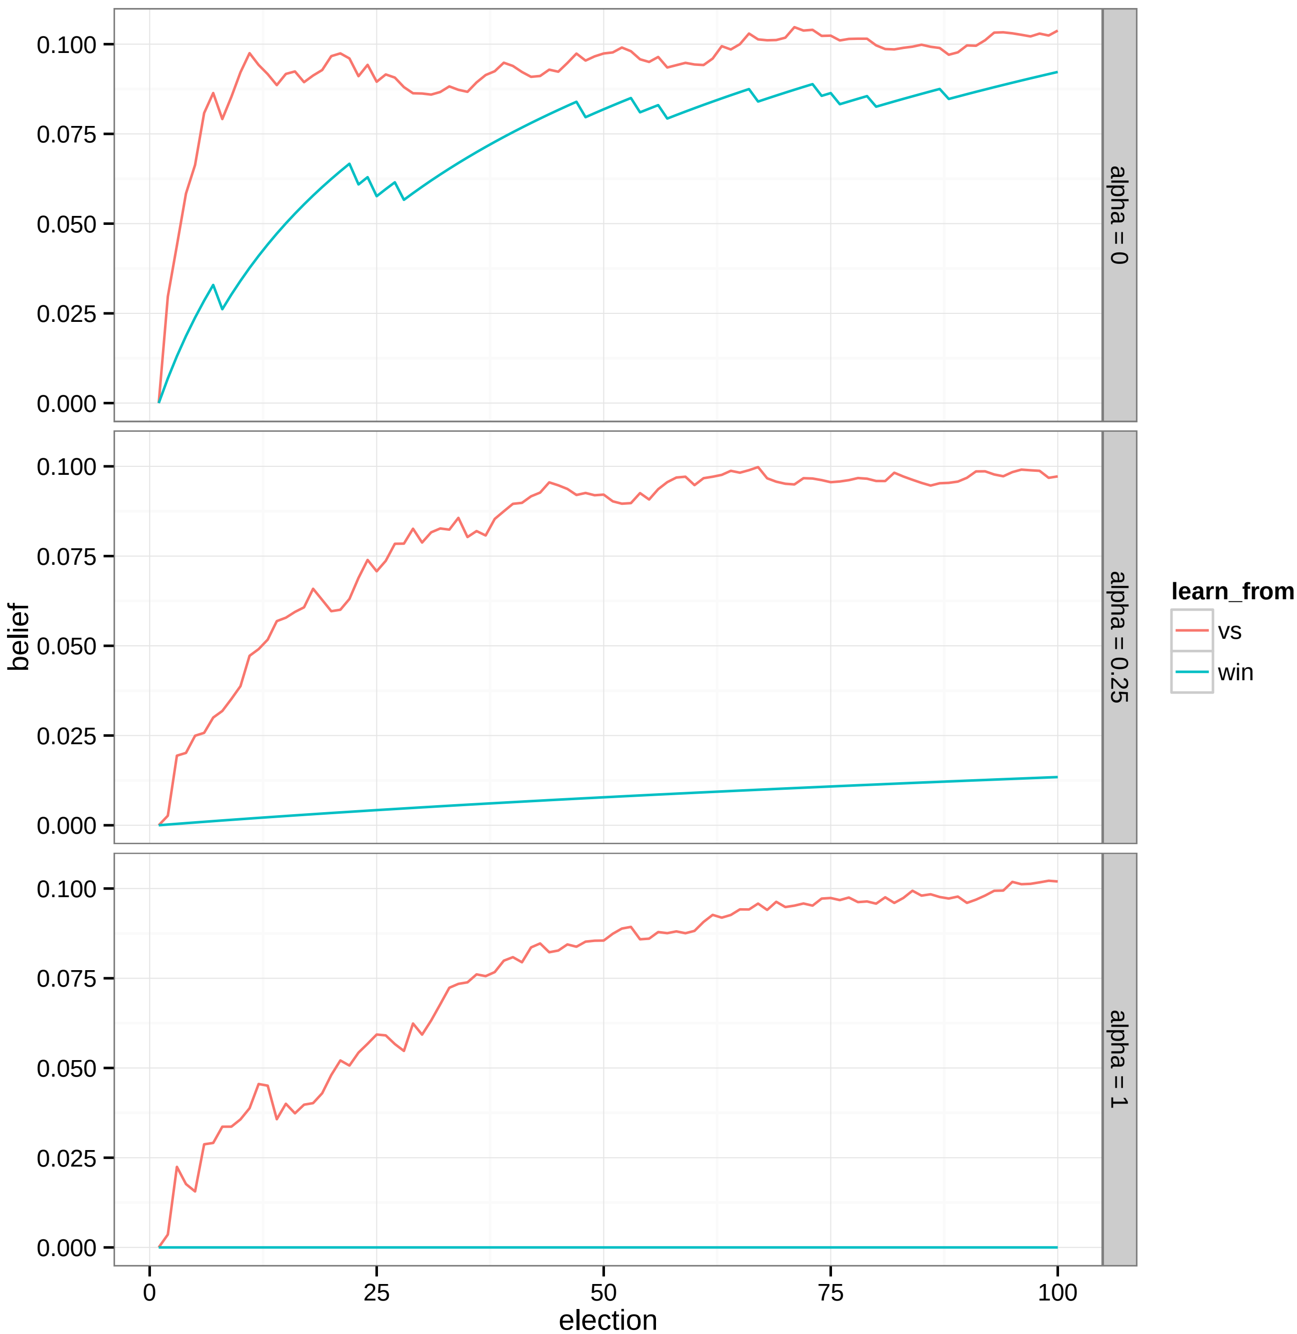

Figure 1: Simulation of learning processes using vote share vs. win information.

Coefficient	(1)	(2)	(3)	(4)	(5)
Prev. Cycle Productivity (Win)	17,242		17,854	17,823	11,366
	(5,501)		(5,457)	(5,531)	(11,010)
Prev. Cycle Average Residual (Win)		0.181			
		(0.029)			
Prev. Cycle Productivity (VS)			-16.595	-27.736	8.408
			(118.686)	(126.529)	(191.380)
Prev. Cycle Market Share	0.865	0.862	0.865	0.865	
	(0.018)	(0.018)	(0.019)	(0.019)	
2006 Cycle				-0.011	-0.010
				(0.007)	(0.007)
2008 Cycle				-0.008	-0.013
				(0.006)	(0.006)
2010 Cycle				-0.007	-0.018
				(0.006)	(0.006)
(Intercept)	0.023	0.023	0.024	0.030	
	(0.002)	(0.002)	(0.002)	(0.005)	
Firm Fixed Effects?	N	N	N	N	Y

Table 1: Second stage regression of firms' market share on productivity measures, using linear productivity model for wins.

races. The results here are illuminating: unexpected vote share performance in close races does appear to have a significant effect on firms' market shares in the next cycle; unexpected vote share performance in uncompetitive races does not. This result is consistent with the distinction observed in the main specification between win- and vote-share-based residual measures, as changes in vote share in close races are very highly correlated with changes in win probability.

As candidates are very likely to care about maximizing their chance of winning an election, but may or may not care about maximizing their vote share (Milyo, 2001), this result may seem unsurprising. However, the problem of deciding what actions to take in one's own campaign is quite distinct from the problem of evaluating the usefulness of consultants based on their performance in other campaigns. Focusing entirely on wins (or, equivalently, on vote share in close races only) is a sensible approach in the former problem but not in the latter. Because most Congressional races are not competitive, candidates are effectively throwing away the majority of available information on consultant performance.

The result is an unusual set of incentives for consulting firms: firms with established reputations can lock in their past success by working in safe races; working in close contests is likely to be unattractive because of the risk of damage to the firm's reputation.[3] Upstart firms with no track record (and no existing relationships to worry about losing) are likely to

Coefficient	(1)	(2)	(3)	(4)	(5)
Productivity (Win)	6.517E+10	7.098E+10	6.683E+10	6.623E+10	1.981E+10
	(1.269E+10)	(1.397E+10)	(1.336E+10)	(1.358E+10)	(1.561E+10)
Productivity (VS)		-2.461E+08	-2.981E+08	-2.756E+08	1.088E+08
		(1.292E+08)	(1.410E+08)	(1.410E+08)	(1.559E+08)
Prev. Cycle Payment	0.370	0.369	0.369	0.369	
	(0.022)	(0.022)	(0.022)	(0.022)	
Fundraising Total	0.061	0.061	0.062	0.062	0.084
	(0.005)	(0.005)	(0.005)	(0.005)	(0.011)
Cook Forecast - D1	-44,416	-44,300	-38,388	-38,282	-26,600
	(8,488)	(8,307)	(7,936)	(8,187)	(13,298)
Cook Forecast - D2	-26,972	-26,907	-21,257	-20,290	-16,507
	(8,759)	(8,530)	(8,437)	(8,319)	(13,266)
Cook Forecast - D3	-52,935	-52,419	-48,652	-46,318	-38,372
	(8,564)	(8,551)	(8,643)	(8,619)	(17,443)
Cook Forecast - R1	-27,101	-27,058	-24,647	-24,190	-37,704
	(9,258)	(8,831)	(8,960)	(8,997)	(16,383)
Cook Forecast - R2	-51,818	-51,629	-47,841	-47,078	6,617
	(9,691)	(9,648)	(9,650)	(9,670)	(13,064)
Cook Forecast - R3	-48,505	-48,377	-46,131	-45,273	-15,199
	(8,303)	(8,307)	(8,469)	(8,347)	(14,184)
2006 Cycle			6,544	6,810	-412.508
			(4,976)	(5,030)	(3,728)
2008 Cycle			-21,222	-20,586	-17,792
			(3,830)	(3,790)	(3,695)
2010 Cycle			-9,519	-9,357	-4,729
			(3,400)	(3,398)	(3,513)
TV Advertising Price Index				-7.746	43.141
				(1.965)	(12.103)
(Intercept)	42,308	42,563	45,679	46,399	
	(9,606)	(9,787)	(10,159)	(9,850)	
Firm-Candidate Fixed Effects?	N	N	N	N	Y

Table 2: Second stage regression of candidate-to-firm payments on productivity measures, using linear probability model for wins.

prefer to take the gamble and work in the closest contests.

We then used the vote-share residuals calculated from competitive races only to repeat the exercise in Figure 2 of the article. The results, shown in Figure 2, continue to show numerous races in the "safe" categories - corresponding to established incumbents - employing high-quality effective consultants.

Media firms We conducted the exercise in Table 7 of the article for firms in the "media" category only, due to evidence that these firms may have different compensation plans than do other types of firms (Grossmann, 2009). It does appear that there is more cycle-to-cycle variability in the compensation of these firms, as evidenced by the substantially lower coefficient on previous-cycle market share and the fact that models using our per-dollar-of-revenue

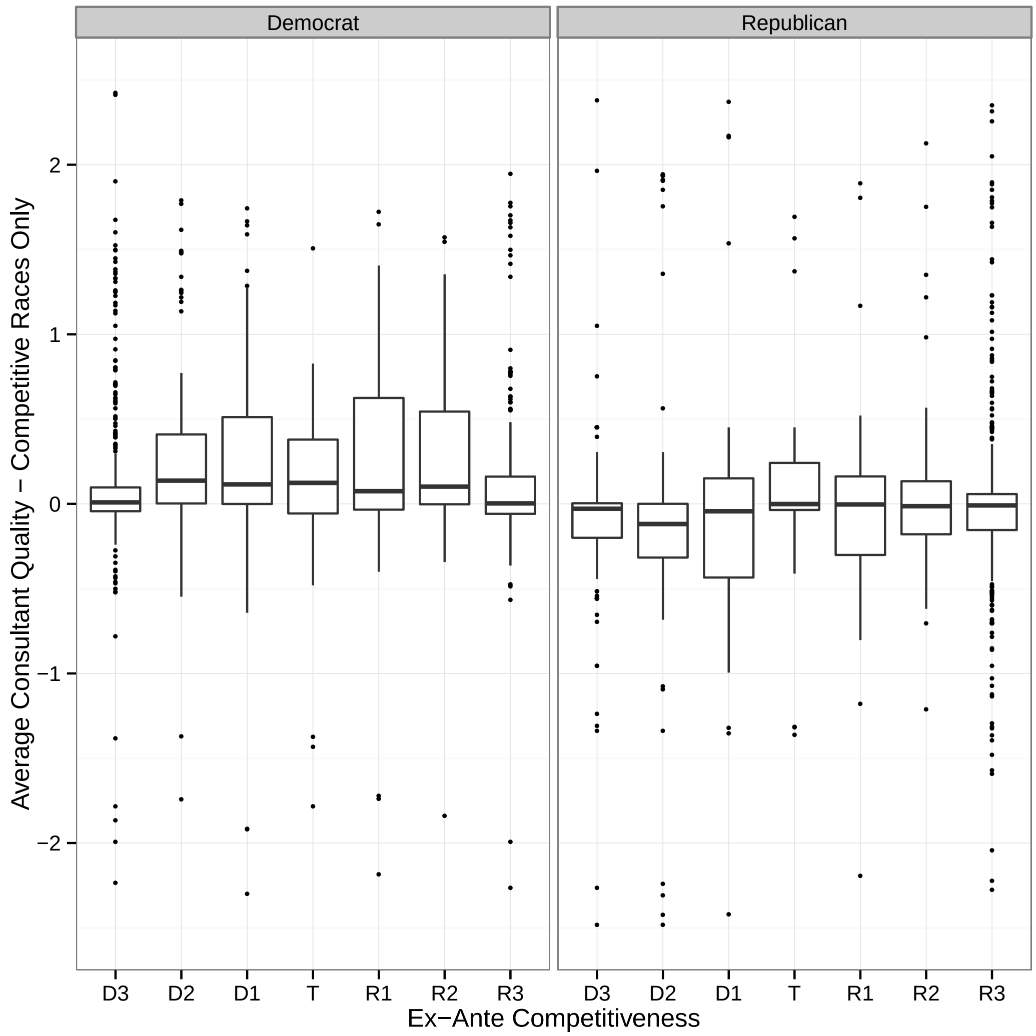

Figure 2: Box plot of ex-ante race competitiveness against the previous election cycle vote share residual of firms assigned to that race in the current cycle, using alternative close-race-only vote share residual measure.

Coefficient	(1)	(2)
Prev. Cycle VS Residual (Competitive Elections)	0.013	
	(0.004)	
Prev. Cycle VS Residual (Safe Elections)		0.001
		(0.001)
Prev. Cycle Market Share	0.863	0.865
	(0.019)	(0.019)
(Intercept)	0.024	0.024
	(0.002)	(0.002)

Table 3: Second stage regression of firms' market share on productivity measures, stratified by competitive / uncompetitive elections.

measures have lower precision. Nonetheless, the substantive conclusions are unchanged.

Coefficient	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Prev. Cycle Productivity (Win)	7,706				11,013	8,764	1.240E+05
	(8,061)				(9,868)	(8,903)	(28,646)
Prev. Cycle Productivity (VS)		-366.213			278.978	199.795	-21,224
		(474.096)			(535.880)	(587.255)	(2,341)
Prev. Cycle Average Residual (Win)			0.058				
			(0.012)				
Prev. Cycle Average Residual (VS)				0.003			
				(0.001)			
Prev. Cycle Market Share	0.579	0.574	0.633	0.581	0.586	0.597	
	(0.041)	(0.040)	(0.042)	(0.042)	(0.043)	(0.043)	
2006 Cycle						9.882E-04	-0.008
						(0.011)	(0.015)
2008 Cycle						-0.015	-0.042
						(0.008)	(0.014)
2010 Cycle						-0.011	-0.039
						(0.010)	(0.014)
(Intercept)	0.043	0.044	0.034	0.041	0.042	0.048	
	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)	(0.008)	
Firm Fixed Effects?	N	N	N	N	N	N	Y

Table 4: Second stage regression of firms' market share on productivity measures; media firms only.

Challengers only It is possible that some of the persistence we observe in matches between firms and candidates may be desirable from the point of view of the party, for instance if firms develop expertise specific to a particular district or candidate over the course of a campaign. Parties may therefore have an interest in having effective firms and incumbent candidates continue to work together over several cycles, even if those incumbents do not appear to be under serious threat in the current cycle. Logic of this kind may partly explain the null relationship between ex ante competitiveness and quality of firm assignment that we observed in Figure 2 of the article.

However, such candidate-specific expertise should not be a relevant factor for new (i.e. non-incumbent) candidates. We therefore repeated the exercise of Figure 2 of the article for non-incumbent candidates only. While the results, shown in Figure 3, display somewhat more of the expected inverted-U shape as compared to the version including incumbent candidates, the difference in average quality between competitive races and uncompetitive ones is still not significant.

House and Senate leadership As a final check that our results are not driven by omitted variable bias, we included dummy variables for House and Senate leadership roles[4] in our predictive regressions. The results with this additional predictive variable in the first stage regression are qualitatively very similar to the main results. We are also concerned about the possibility that party leaders, who are on average quite electorally secure, may be employing high-quality consulting firms. To rule out the possibility that our results on the allocation of consultant quality across races is spuriously driven by party leaders, we examined whether the result changes when we omit party leaders from the analysis. The relationship of firm quality to race competitiveness (Figure 2 of the article) is very nearly identical when members of the leadership are excluded from the sample.

C Descriptive Statistics

Figure 4 shows the distribution of per-client revenues in our sample. The median firm bills about $23,000 per client, and works with two campaigns per cycle. However, there is a substantial group of firms charging in the $100,000 - $1 million range per client.

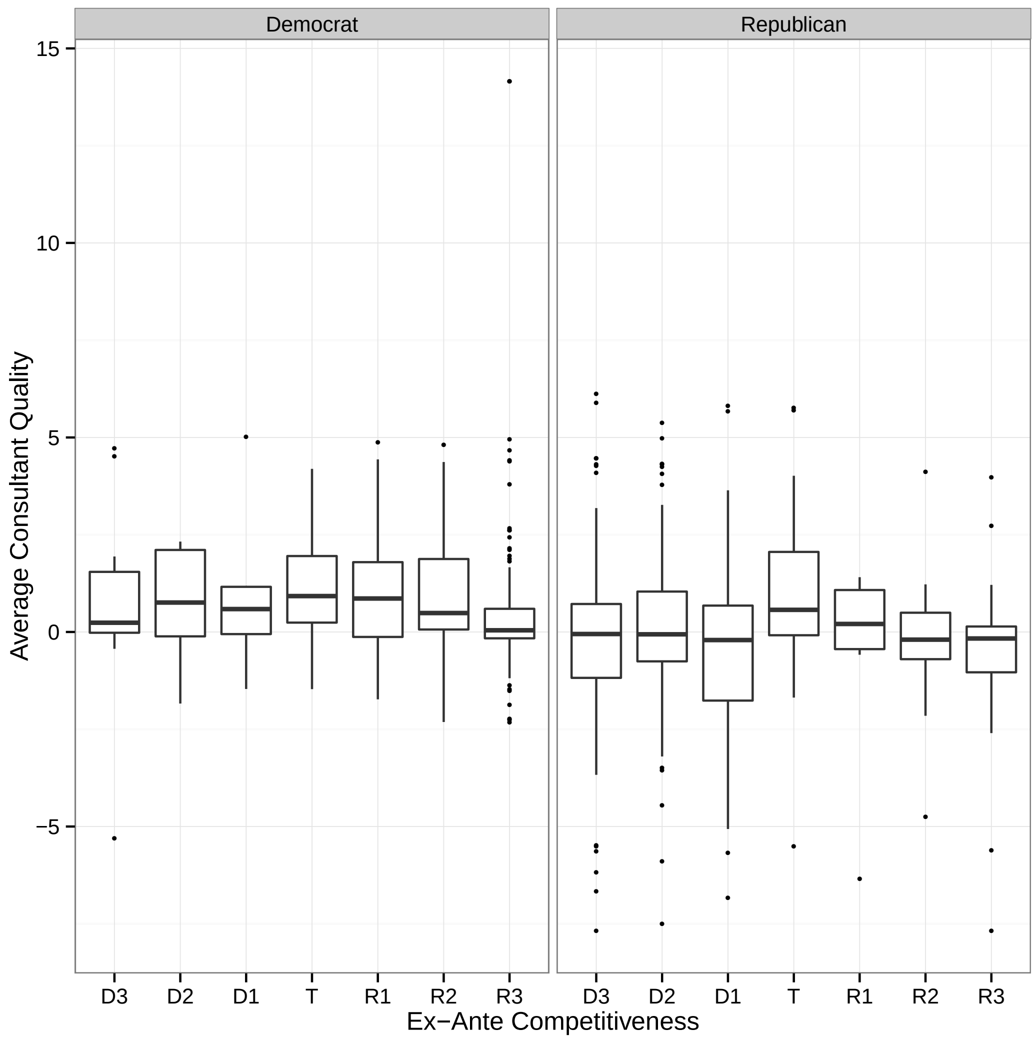

Figure 3: Box plot of ex-ante race competitiveness against the previous election cycle vote share residual of firms assigned to that race in the current cycle, for non-incumbent candidates only.

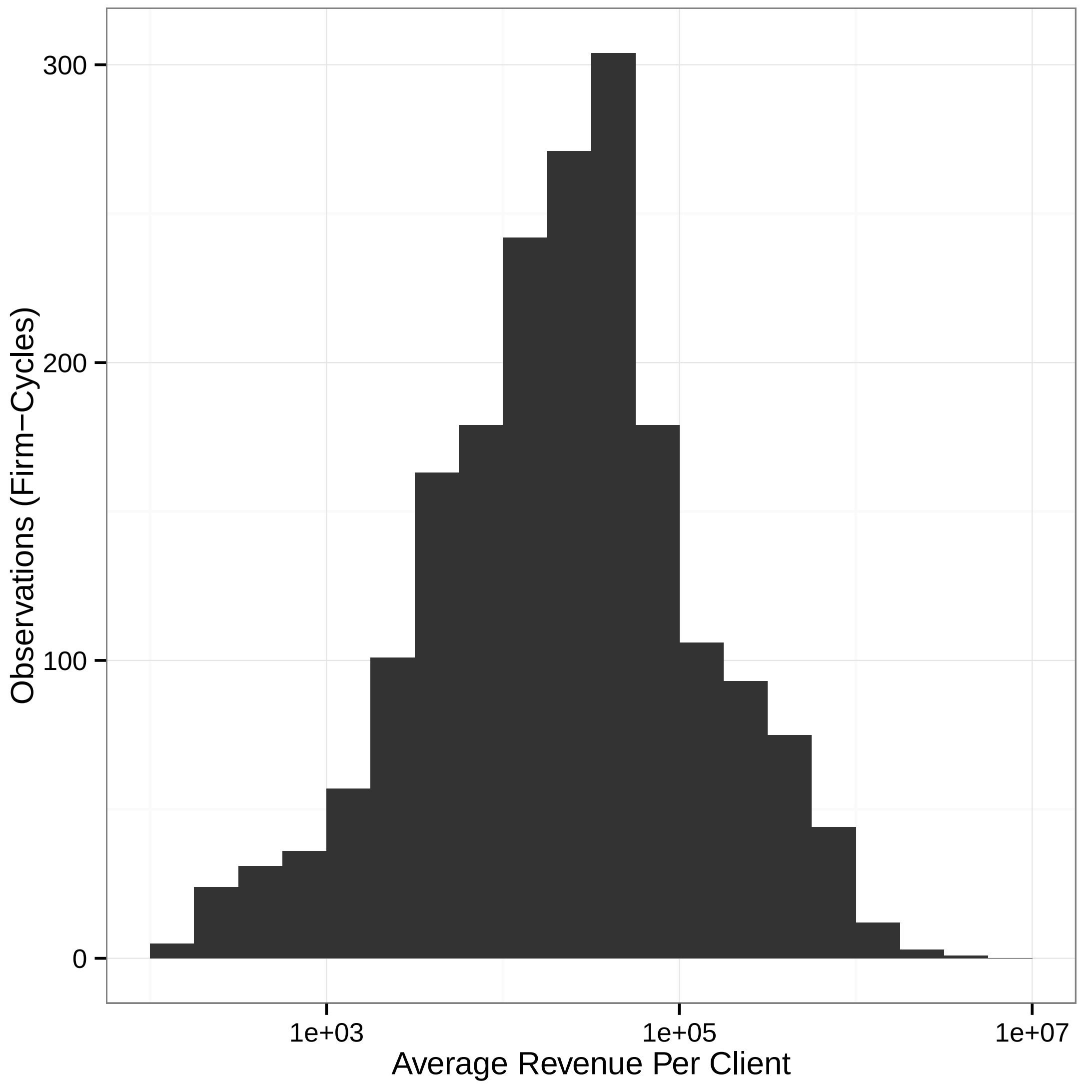

Figure 4: Histogram of firms' average revenue per campaign (log scale)

Coefficient	(1)	(2)
Prev. Cycle Average Residual (Win)	0.044	
	(0.010)	
Prev. Cycle Average Residual (VS)		0.002
		(7.53E-04)
Prev. Cycle Market Share	0.877	0.878
	(0.019)	(0.019)
(Intercept)	0.021	0.022
	(0.002)	(0.002)

Table 5: Second stage regression of firms' market share on residual measures, using leadership status as predictors.

References

Grossmann, Matt. 2009. "Campaigning as an Industry: Consulting Business Models and Intra-Party Competition." *Business and Politics* 11(1):1–19.

Milyo, Jeffrey. 2001. "What Do Candidates Maximize (and Why Should Anyone Care)?" *Public Choice* 109(1-2):119139.