Supporting Information

Do Elections Improve Constituency Responsiveness?
Evidence from U.S. Cities

Contents

SI 1 Information about Service Protocol and Time Line

Table SI.1 provides an overview of how 3-1-1 service requests are handled from when they are opened until they are closed.

Table SI.1 Service Request Protocol and Timeline

1. A request is called in to a 3-1-1 response center or submitted online.
2. A 3-1-1 representative (or online system) will determine the type of request, and an actioning agency will be assigned. Based on a pre-assigned workflow, the service representative will then do the intake of required information (as requested by the assigned agency) and submit the request. The request has been *opened*, and the date is recorded in the 3-1-1 database.
3. The request appears in the acting agency's queue for action, and will have a service level agreement (SLA) specifying a due date based on the request type. The agency may prioritize among different requests as they see fit.
4. Agency staff physically inspect the reported issue, with one of three potential outcomes, each of which results in the request being *closed*:
   a. Agency staff resolve the issue (may require revisits). Some issues are easily verified as resolved (e.g., graffiti removed), while others may require following up with the constituent.
   b. The issue is rerouted to a different agency. This may happen if the agency determines the issue falls outside its jurisdiction. For example, NYC Department of Transportation may reroute a highway issue to NY State, or a misreported issue may be rerouted to a different city agency. In this case the issue will be marked as closed once it has been resolved by the new agency.
   c. Agency staff determine that no work is warranted (e.g., trash was already picked up, prank call) and mark it as closed.

SI 2 Changes in Request Volume Following Term Extension (NYC 2008)

We explore whether changes in response times are driven by differential changes in the demand for constituency services. Concretely, we estimate the difference-in-differences specified in equation 1 using the number of requests as the dependent variable (omitting the fixed effects for request type). Figure SI.1 and table SI.2 do not suggest differential changes in request volume: newly eligible (treated) incumbents saw an increase of three to eight requests in their districts — small fluctuations relative to the average number of requests that cannot be distinguished from zero in two of three specifications.

Figure SI.1 Average Number of Requests around the Term Extension
No differential change in request volume around the term extension.

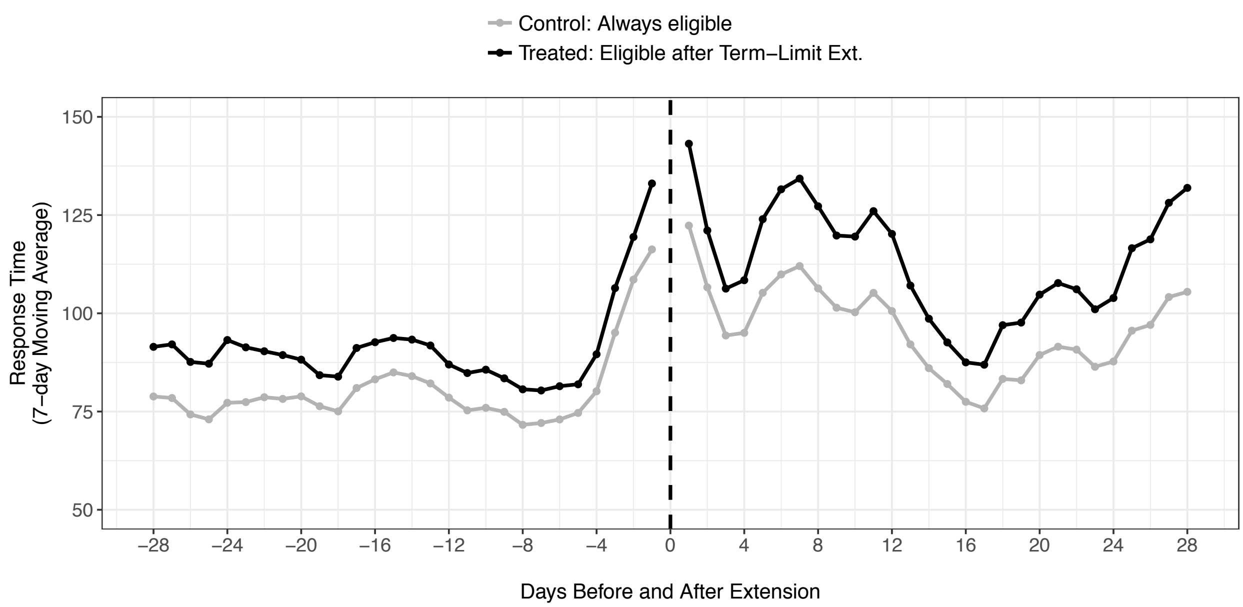

Table SI.2 Effect of Term Extension of Request Volume

	Dependent variable:		
	Number of Requests		
	2	3	4
$\hat{\beta}^{\ddagger}$	3.424	4.451	7.611
	(2.908)	(2.973)	(4.331)
	p = 0.240	p = 0.135	p = 0.080
Observations	1,160	1,720	2,280

‡ Difference-in-differences estimator (see Eq. 1)
Standard errors clustered on districts in parentheses.

SI 3 Changes in Request Types Following Term Extension (NYC 2008)

We also consider whether the composition of service requests changes differentially in districts where incumbents' election eligibility changes. For the top twenty service requests (which account for 94 percent of observations), we count the frequency of each type of request in every day and district. We then regress our treatment indicator (D_{dt}) on these frequencies, including district and day fixed effects. (We use two weeks on either side of the policy change for this analysis.) This is comparable to a joint test of orthogonality when evaluating balance and allows us to evaluate whether changes in the frequency of certain requests predicts treatment.

Looking at Table SI.3, the coefficients are uniformly small. We focus, however, on the joint hypothesis that the coefficients on these frequencies all equal zero: our F-statistic is 1.417, $p = 0.17$. We cannot, thus, reject the null that that the coefficients on all of these categories equal zero.

Table SI.3 Joint Test that Request Composition does not Predict Treatment

	Dependent variable:
	D_{dt}
Animal Care	−0.003 (0.006)
Appliance	−0.005 (0.008)
Broken Meter	−0.001 (0.004)
Building/Use	0.003 (0.003)
Complaint	0.009 (0.006)
Construction/Plumbing	0.002 (0.001)
Damaged Tree	−0.004 (0.006)
Derelict Vehicle	0.004 (0.004)
Electric	0.0002 (0.003)
Heating	0.0002 (0.0003)
Noise	−0.004 (0.003)
Nonconst	−0.005 (0.003)
Paint/Graffiti	−0.002 (0.001)
Sanitation/Cleaning	0.005 (0.002)
Sidewalk/Sewer	0.002 (0.001)
Street Light Condition	0.002 (0.002)
Street/Highway/Bridge Issues	−0.004 (0.002)
Taxi	0.003 (0.005)
Traffic Signal Condition	−0.001 (0.004)
Water	−0.002 (0.003)
Observations	1,120
F-statistic	1.417
p-value	0.17
Note:	Standard errors clustered on districts in parentheses

SI 4 Robustness to Including Incumbents Not Seeking Third Term (NYC 2008)

The treatment group for our main analysis consists of incumbents who were termed out before the October 23, 2008 decision but ran for a third term after the decision. Not all of the newly eligible councilors took advantage of the extension; eight incumbents left politics or ran for different positions (e.g., four ran for comptroller). In Table SI.4, we show that our results hold but decline in magnitude when we include all term-limited incumbents in our treatment group. The results attenuate as expected: we would not expect councilors intending to leave office to exert more effort following the term limit extension.

Table SI.4 Effect of Term Extension on Constituency Services, All Term-Limited Incumbents

	Dependent variable:			
	Response Time			
Time frame†	2	3	4	12
$\hat{\beta}^{\ddagger}$	−1.001	−1.202	−1.089	−0.441
	(0.703)	(0.627)	(0.496)	(0.808)
	p = 0.155	p = 0.056	p = 0.029	p = 0.586
	$\mathbb{1}$(Response Time < 5 Days)			
Time frame†	2	3	4	12
$\hat{\beta}^{\ddagger}$	0.026	0.013	0.015	0.001
	(0.010)	(0.008)	(0.009)	(0.009)
	p = 0.013	p = 0.125	p = 0.095	p = 0.868
Observations	145,229	211,947	284,927	822,431

†Weeks on either side of the extension used to estimate Eq. 1

‡Difference-in-differences estimator (see Eq. 1)

Standard errors clustered on districts in parentheses

SI 5 Heterogeneous Effects by Vote on the Extension (NYC 2008)

Table SI.5 indicates councilors' votes on the term extension based on their treatment status and decision to run for a third-term in office.

Table SI.5 Voting for Term Extension by Treatment Status

	Group	Freq.	Prop. Voting Yes
1	Control	12	0.330
2	Second-Term, Contesting	28	0.790
3	Second-Term, Retiring	11	0.270

In Table SI.6, we estimate whether treated incumbents that supported the term extension on October 23, 2008 saw larger reductions in response time. This analysis employs a four-week window around the policy change and does *not* indicate heterogeneous effects based on councilors' support for the reform; councilors supporting the reform do not achieve larger reductions in response times.

Table SI.6

	Dependent variable:	
	Response Time	
	Constesting Incumbents	All Incumbents
D_{dt}	-1.641	-1.188
	(0.967)	(0.632)
	p = 0.090	p = 0.061
$D_{dt} \times \mathbb{1}$(Voted Yes)	0.487	0.141
	(1.033)	(0.688)
	p = 0.637	p = 0.838
Observations	230,067	284,927

Standard errors clustered on districts in parentheses.

SI 6 Post-Extension Trends as 2009 NYC Election Approaches

In Table SI.7, we explore post-extension trends in response times leading up to the November 3, 2009 election in NYC. Since all councilors were eligible to run for reelection after the term limit extension, we do not expect different response time *trends* in districts represented by previously term-limited councilors (treated) and districts represented by councilors who were always eligible for reelection (control).

We test this expectation using Equation 2. We use two time periods: a year and six months before the election. If there were no differential changes to response times in treated and control districts, then $\hat{\beta}$ from Equation 2 should be indistinguishable from 0. Table SI.7 provides strong evidence for this: the estimates of differential time trends are very close to zero.

We emphasize that this analysis is different from our tests of the impact of the term-limit extension, for which we use Equation 1. The extension analysis shows that response times dropped in districts represented by previously term-limited councilors in the weeks after the extension. However, as we also show, this effect attenuates with time. This makes sense given the analysis presented here: as the 2009 election approaches, response time trends in treated and control districts become indistinguishable.

Table SI.7 Tests for Differential Trends before the 2009 NYC Election

	Dependent variable:	
	Response Time	
Days until election:	365	182
$\hat{\beta}$	0.0001	−0.0002
	(0.0001)	(0.0004)
	p = 0.639	p = 0.635
Observations	1,294,925	581,386

Note: Standard errors clustered on districts in parentheses

SI 7 Robustness to Trimming

The tables below demonstrate the robustness of our results to different decisions about how to trim the dependent variable, which has a long right tail.

Table SI.8 Robustness of Election Eligibility Results to Trimming Decisions

	Dependent variable:		
	Response Time		
$\max\{y\}$	1 year	2 years	3 years
$\hat{\beta}^{\ddagger}$	-0.801	-0.709	-1.241
	(0.525)	(0.521)	(0.562)
	p = 0.127	p = 0.174	p = 0.028
Observations	229,349	229,570	230,067

Uses 4 weeks on either side of the extension to estimate Eq. 1.
‡Difference-in-differences estimator (see Eq. 1)
Standard errors clustered on districts in parentheses.

Table SI.9 Robustness of Differential Time Trends to Trimming Decisions (β in Equation 2)
The linear trend in responsiveness falls significantly faster where incumbents can run for reelection.

	Dependent variable:		
	Response Time		
$\max\{y\}$	1 year	2 years	3 years
	Days to Primary (NYC) or Filing Date (SF)		
$\hat{\beta}$	-0.020	-0.020	-0.020
	(0.004)	(0.005)	(0.005)
	$p < 0.001$	$p < 0.001$	$p < 0.001$
Observations	2,293,059	2,306,984	2,307,747
	Days to General Election		
$\hat{\beta}$	-0.014	-0.015	-0.015
	(0.005)	(0.005)	(0.005)
	p = 0.003	p = 0.007	p = 0.006
Observations	2,434,722	2,450,274	2,451,003

Uses 547 days before election date to estimate Eq. 2.
Standard errors clustered on districts in parentheses.

SI 8 Placebo Tests for Election Timing Results

We randomly draw 20 dates from the period following the 2005 election in NYC and the 2010 election in SF. We then re-estimate β from Equation 2. We use 730 days to estimate these placebo regressions. Thus, our "placebo" election dates have to be drawn from an interval of 200 days that is at least two years after the actual election date. This explains why the placebo dates fall well to the right of the estimate associated with the actual election dates (the left-most, black points). These tests suggest that the differential time trends that we discover in our analysis do not persist after the election.

Figure SI.2 Placebo Results for Election Timing

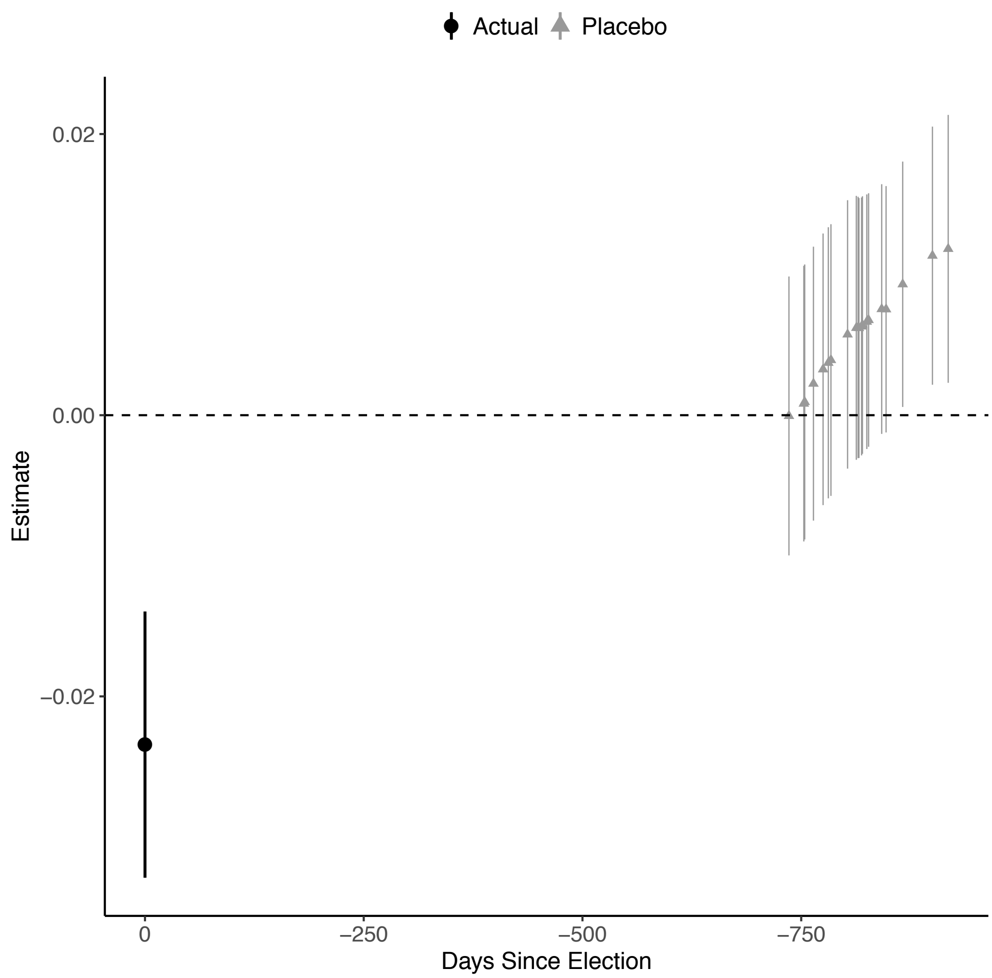

Displayed above are the estimates and 95% confidence intervals for $\hat{\beta}$ from Equation 2 using two years of data before each placebo date. The estimate from the actual term-limit extension is the left-most, black point.

SI 9 Separate Results for NYC 2005 and SF 2010

Table SI.10 Separate Estimates of Differential Time-Trends for NYC 2005 and SF 2010 (β in Eq. 2)
The linear trend in responsiveness falls significantly faster where incumbents can run for reelection.

	Dependent variable:			
	Response Time			
Time frame†	730	547	365	182
	NYC 2005			
$\hat{\beta}$	-0.012	-0.006	0.005	-0.022
	(0.004)	(0.004)	(0.008)	(0.019)
	p = 0.003	p = 0.097	p = 0.533	p = 0.249
Observations	2,376,717	2,238,742	1,578,977	774,375
	SF 2010 (Control: No Election)			
$\hat{\beta}$	-0.014	-0.025	-0.039	-0.002
	(0.004)	(0.006)	(0.009)	(0.022)
	p = 0.001	$p < 0.001$	$p < 0.001$	p = 0.943
Observations	160,678	120,549	81,801	42,571
	SF 2010 (Control: Term-Limited)			
$\hat{\beta}$	-0.022	-0.039	-0.071	-0.047
	(0.011)	(0.020)	(0.036)	(0.040)
	p = 0.050	p = 0.055	p = 0.049	p = 0.237
Observations	135,393	101,458	68,933	35,531

†Days before general election used to estimate Eq. 2
Standard errors clustered on districts in parentheses

Table SI.11 Separate Estimates of Changes in Legislative Activity for NYC 2005 and SF 2010
Eligible incumbents either maintain or increase their legislative effort, suggesting no reallocation.

	Dependent variable:					
	Total Legislative Actions Introduced					
	NYC 2005		SF 2010 (Control: No Election)		SF 2010 (Control: Term-Limited)	
Time frame†	365	182	365	182	365	182
$\hat{\beta}$	-0.0004	-0.0002	0.006	0.016	0.002	0.008
	(0.003)	(0.006)	(0.0004)	(0.002)	(0.002)	(0.001)
Observations	1,224	612	287	168	205	120

†Days before general election used to estimate model (excluding γ_{type}).
Standard errors clustered on districts in parentheses

SI 10 DiD Estimates by Racial Composition

Our findings suggest that public service responsiveness improves when public officials are eligible to seek reelection and as elections approach. We are also interested in whether some neighborhoods benefit more from these improvements than others. In particular, we are interested in whether the racial composition of neighborhoods is responsible for heterogeneous treatment effects in responsiveness. Public officials may, for example, favor coethnic constituents or constituents from a particular racial group.

To carry out this analysis we begin by matching the 3-1-1 database with census data on the racial composition of every block in NYC. We then code two characteristics that indicate each block's relationship with the city council member representing the electoral district in which the block is located: whether or not its largest group is coethnic with the city council member; and whether or not its majority group (if any) is coethnic with the city council member. We also create variables for each block's plurality group and (if applicable) majority group without regard for its relationship with the city council member. We then re-run our analyses on different subsets of the data based on these variables.

The results are displayed in Figures SI.3 and SI.4. They suggest that the heightened responsiveness that followed the term-limit extension in 2008 is not driven by ethnic favoritism: neighborhoods in which the plurality (or majority) group is coethnic with the city council member representing the district do not see larger drops in response times than other neighborhoods. We find some evidence, however, that neighborhoods populated primarily by Hispanics or Asians see larger drops in response times, though this may not reflect the ethnicity of these neighborhoods but rather some unobserved characteristic of the neighborhoods that we are not controlling for here (e.g., location).

Moving to the 2005 council elections, response times dropped quite uniformly across neighborhoods two years to a year and a half before the elections. In general, there are few interesting heterogeneous treatment effects to report, perhaps with the exception that Asian neighborhoods saw larger drops closer to the election.

Our failure to uncover consistent heterogeneous effects related to neighborhoods' ethnic composition should not be taken to imply that there are no ethnic disparities in public service delivery. Our empirical strategy leverages changes in responsiveness and, thus, does not address level differences in service delivery across neighborhoods of varying composition.

Figure SI.3 The Term Limit Extension of 2008 and Heterogeneous Effects by Neighborhood Race

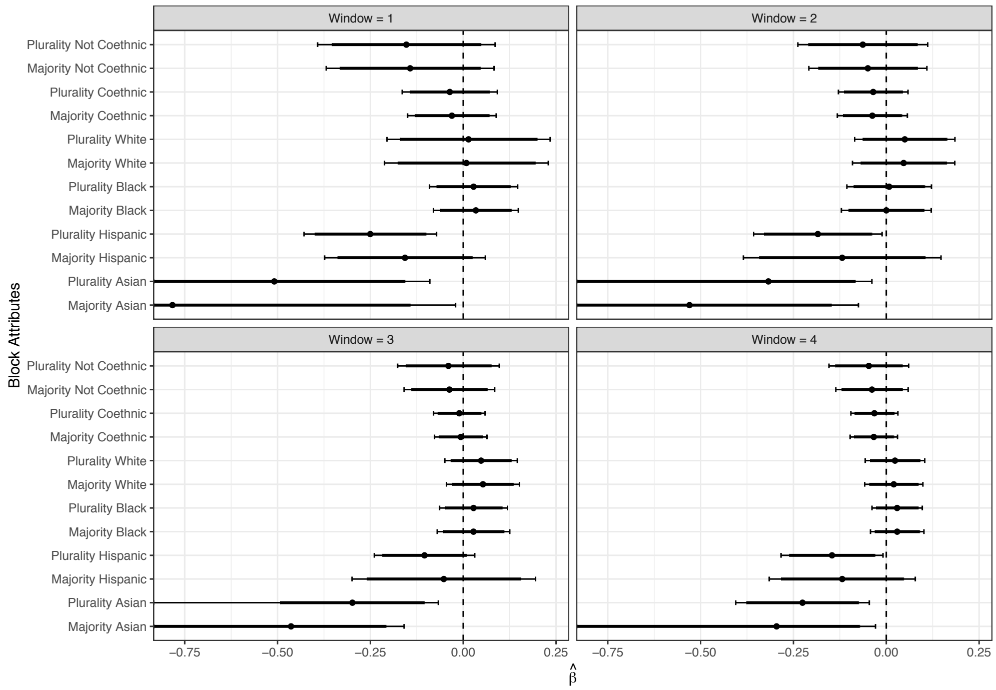

We run Equation 1 in subsets of neighborhoods defined by the attributes on the y-axis. "Window" gives the number of weeks on either side of the extension we use. The estimates of β (with 95% CIs) from these regressions are displayed above.

Figure SI.4 The 2005 NYC Election and Heterogeneous Effects by Neighborhood Race

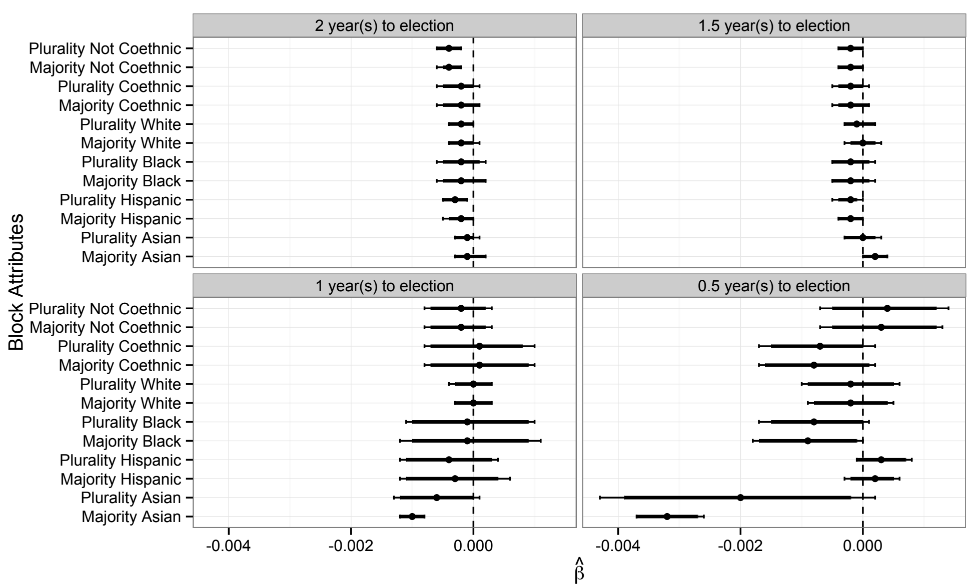

We run Equation 2 in subsets of neighborhoods defined by the attributes on the y-axis, for different time windows before the election. The estimates of β (with 95% CIs) from these regressions are displayed above.