

Squeaky Wheels and Inequality in Bureaucratic Service Provision

Tara Slough*

August 8, 2021

Abstract

Bureaucrats produce and distribute public goods and services with wide scope to influence “who gets what.” Under what conditions do bureaucrats’ actions create inequalities in access to services? I contend that citizens’ principal mechanism of control over bureaucrats is the complaint to a politician. When politicians respond to complaints by tightening oversight of bureaucrats, differences in the likelihood of complaint induce bureaucrats to devote more effort to the loudest citizens. I probe this mechanism using a national-scale factorial audit experiment of two large Colombian programs to measure bureaucratic effort behaviorally. Bureaucrats devoted less effort in responding to marginalized citizens. Consistent with the mechanism, this bias manifests in municipalities with greater inequalities in citizens’ access to the state and on tasks where oversight from politicians is most likely. This paper shows that inequality in access to public services can emerge even when budget allocations to public goods are equitable.

* Assistant Professor, New York University, tara.slough@nyu.edu. This project is supported in part by an NSF Graduate Research Fellowship, DGE-11-44155; NSF Doctoral Dissertation Research Improvement Grant 1729334; and grants from J-PAL, Evidence in Governance and Politics (EGAP), the Center for the Study of Development Strategies, and the Columbia Department of Political Science. The experiment was conducted with IRB approval from Columbia University (#AAAR3617) and Innovations for Poverty Action (#14442). The administrative data was collected with IRB approval from Columbia University (#AAAR0174). I am indebted to Juan Sebastian Leiva, Sofia Jaramillo, and Yuliana Silva Contreras for outstanding research assistance. I thank Innovations for Poverty Action, the Departamento Administrativo de la Función Pública, Prosperidad Social, and Departamento Nacional de Planeación for their contributions to the development and execution of this project. Thanks to Lindsay Dolan, Natalia Garbiras-Díaz, Ben Goodrich, Don Green, Kolby Hanson, Shigeo Hirano, John Huber, Macartan Humphreys, Navin Kartik, Kimuli Kasara, Horacio Larreguy, Evan Lieberman, Yotam Margalit, John Marshall, Vicky Murillo, Suresh Naidu, Gareth Nellis, Carlo Prato, Pia Raffler, Tesalia Rizzo, Cyrus Samii, Jake Shapiro, Jack Snyder, Mike Ting, and Joonseok Yang; seminar participants at American University, Columbia, Rochester, Universidad de los Andes, Universidad del Rosario, the University of Pennsylvania, Washington University in St. Louis, and the WZB Social Science Center; policy audiences at Prosperidad Social, Departamento Nacional de Planeación, and Departamento Administrativo de la Función Pública; and conference attendees at NEWEPS-11, the Harvard Experiments Working Group Conference, APSA 2018, and MPSA 2018 for comments. This project was pre-registered with EGAP prior to the realization of outcome data. This paper was previously circulated under the title: “Bureaucrats Driving Inequality in Access: Experimental Evidence from Colombia.”

1 Introduction

Inequalities in access to public goods and services have long challenged developing democracies, inhibiting efforts to reduce socioeconomic inequality and promote economic development. Some inequality results from budget constraints or politicians' targeting of resources, but inequalities may also emerge in the production of these goods and services. As "producers of public goods," bureaucrats map politicians' budget allocations into outputs. I argue that bureaucrats are prone to exert differential effort in providing service to different groups of citizens. Variation in bureaucratic effort, in turn, creates inequality in citizen access to public services. These disparities can emerge even when budget allocations are equitable.

I focus on bureaucrats' role as service providers, emphasizing interactions between street-level bureaucrats and citizens.¹ Across contexts, citizens engage the bureaucracy to gain access to benefits and services. Bureaucrats, in turn, distribute these services because political principals delegate tasks of program administration. With this delegation comes oversight. Politicians monitor the work of bureaucrats, doling out punishments or rewards upon observation of bureaucrats' performance.

I consider the role of citizen complaints in directing politicians' oversight of bureaucrats. Such complaints, conceptualized broadly as communication from a citizen (or client) about the actions of an agent to the agent's principal, represent the primary form of citizen control over bureaucrats (Prendergast, 2003, 2007). Complaints incentivize the politician to target monitoring to specific decisions of the bureaucrat. Empirically, laws regarding citizen complaints and responses are quite common in developing democracies and enshrine this form of citizen control over bureaucrats as a right. Such citizen appeals (complaints) and politicians' remedies animate recent discussions of constituency service in developing democracies (Bussell, 2020). This paper instead focuses on how differences in citizen propensity to complain generate disparate treatment by bureaucrats, promulgating inequality in citizen access to services.

¹Following Lipsky (2010), I define street-level bureaucrats as those individuals that interface directly with citizens to implement policies that they do not create.

This paper makes three innovations. First, I advance a theory that variation in citizen ability to check the bureaucracy via complaint induces bias in bureaucratic effort. The theory emphasizes how such biases map onto inequality in service provision across groups in society. Second, I develop novel behavioral measures of bureaucratic effort. Finally, I derive tests of the mechanisms to separate complaint-driven bias from taste-based biases.

I introduce a simple model of service provision that seeks to understand a bureaucrat's decision to exert effort to deliver a service to a citizen. In equilibrium, a bureaucrat's effort is increasing in both their taste for serving the citizen and the anticipated intensity of oversight by a politician. A politician's intensity of oversight increases in their taste for recovering service for a citizen and in response to citizen complaints about the bureaucrat's outputs. Thus, the bureaucrat works harder to serve a citizen if they anticipate that lackluster service provision will draw a complaint. Building from these predictions, I define bureaucratic bias in effort as the difference in average effort rendered to citizens from different groups. The three factors that drive bureaucratic effort therefore introduce three mechanisms underpinning bureaucratic bias: the bureaucrat's tastes; anticipation of citizen complaints; and the politician's tastes. *Complaint-driven bias*, a central focus of this paper, emerges when bureaucrats anticipate that citizens from different groups will complain at different rates and therefore devote more effort to "squeaky wheels," or the group most likely to complain. I use comparative statics on total bureaucratic bias in effort to generate testable implications to disentangle complaint-driven bias from the taste-driven biases.

I develop a research design to measure bureaucrats' behavior through a pre-registered national-scale phone audit experiment in Colombia. The audits pertain to two national social welfare programs and were conducted in consultation with three national agencies overseeing the bureaucracy and these programs. The audited programs, a conditional cash transfer program (CCT) and a means testing service represent two of the largest social programs in Colombia by enrollment. While both programs are funded and directed by the national government, bureaucrats within each local government (*alcaldía*) assume responsibility for enrolling and actualizing the status of ser-

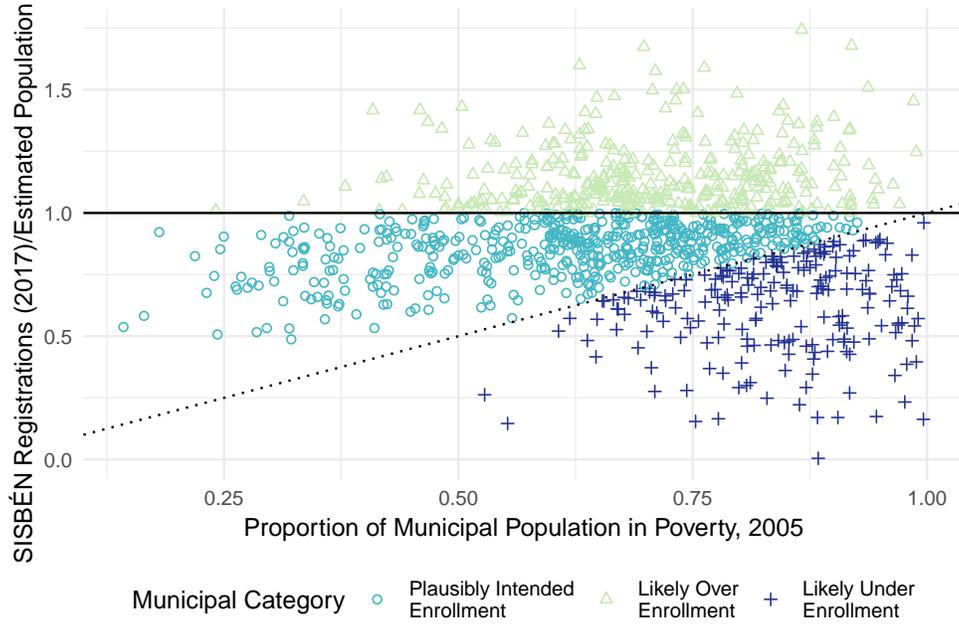


Figure 1: The relationship between municipal poverty rates and SISBÉN (a means-testing service) enrollment as a proportion of the population.

vice recipients at the request of citizens.² Figure 1 compares municipal poverty rates to rates of SISBÉN program enrollment at the time of this study. A simple bounding exercise defines three regions of the plot and suggests dramatic variation in municipal administration SISBÉN across Colombia’s 1102 municipalities. Above the horizontal line, SISBÉN has more enrollees than the ostensible municipal population (32.8%). Below the 45-degree line, the program fails to cover the share of the municipal population in poverty – the minimal intended population of enrollees and targets of associated social programs (19.1%). This underenrollment is consequential because it impedes potential recipients’ access to at least 24 other means-tested social programs (Table A10). Finally, between the two lines, enrollment plausibly aligns with intended administration (48.1%).³

My empirical strategy directly parallels the theoretical model. I measure bureaucratic effort devoted to individual citizens in response to informational queries about locally-administered services associated with each program. The use of a phone audit offers rich measurement of bureau-

²For the rest of the paper, I use the Spanish word for local government, *alcaldía*, to refer to the government entities that I study.

³I pursue this bounding exercise because administrative data that facilitates comparison of individual enrollees and non-enrollees is not public.

crats' behavior, capturing access to officials and provision of information (the service). In order to identify bias in effort, I use a factorial experimental design that varies both characteristics of petitioners (socioeconomic class, internal migrant status, and regional accent) to estimate levels of bureaucratic bias in effort. Finally, I follow the tests of the mechanism implied by the model by leveraging both experimental manipulation of petition difficulty as well as observational variation in the program and municipal characteristics to probe the new mechanism that I propose – complaint-driven bias. Because complaint-driven bias refers to bureaucratic anticipation of oversight in response to complaints, I measure bias in bureaucrat's allocation of effort in response to informational queries; petitioners did not lodge complaints in response to poor service, nor do I analyze complaints in the experimental analysis.⁴

I find robust evidence of bias in effort: lower class individuals and internal migrants received substantially less information than their lower-middle class and resident counterparts, respectively. In the context of the audit experiment, I identify these biases on outcomes measuring bureaucratic effort. I then show that bias in the provision of information in the experiment correlates with municipal-level measures of SISBÉN underprovision, occurring only in the municipalities where SISBÉN (the means testing program) is administratively underprovided. This finding provides evidence consistent with the theoretical model's link between bureaucratic effort and the ultimate service provision outcomes depicted in Figure 1.

Drawing upon the testable implications of the model, I provide clear evidence that the results reflect complaint-driven bias. First, to separate bureaucratic tastes from oversight, I show that bias is attenuated to zero as oversight by politicians becomes less likely. This suggests bias is unlikely to be driven by the bureaucrats' tastes alone. Second, I separate complaint-based bias from politician tastes by showing that anti-poor bias emerges only in poorer places where the ostensible differences between experimental petitioners' abilities to complain is greatest.

Theoretically, this paper unites and extends two foundational approaches to the analysis of

⁴I provide observational data on complaints in the Colombian context. Because bureaucrats provide better service in anticipation of complaints, the model does not produce a falsifiable directional prediction on the incidence of complaints by group.

agency problems in public service provision. First, it brings the moral hazard problem facing bureaucrats into a large literature on distributive politics focused on relationships between politicians and citizens (voters) (Golden and Min, 2013). In so doing, I argue that the strategic behavior of intermediary bureaucrats may distort politicians' desired allocation of public services. My findings imply that even if politicians were to allocate funds equally across citizens or constituencies, variation in citizen control over bureaucrats may lead to an inequitable distribution of services.

Second, it brings citizens into models of politician's oversight of bureaucrats (e.g., Banks, 1989; Ting, 2017). In this regard, I endogenize the information that politicians receive to guide oversight by emphasizing the role of citizen complaints to a principal (Prendergast, 2003, 2007). While citizen complaints to a political principal increase overall bureaucratic effort, they generate inequality when some citizens are less able to engage politicians via complaint, inducing biased oversight by the principal. I find that oversight can either increase or reduce bias in bureaucratic effort, countering existing assumptions that oversight deters bureaucratic bias (Hemker and Rink, 2017; White, Nathan, and Faller, 2015).

The experiment represents one of the first audit experiments on local bureaucrats, generating new insights about the administration of social programs.⁵ I join a literature that identifies substantial variation in access to social programs in Latin America even among ostensible target populations (Holland, 2015, 2017; Niedzwiecki, 2018). In contrast to existing electoral and partisan explanations for this variation, I identify a complementary mechanism through which everyday service provision by bureaucrats generates disparities in access.

Taken together, I show that socioeconomic inequalities generate political inequalities in citizens' ability to extract oversight over bureaucrats responsible for service provision. These inequalities in voice engender inequality in access to poverty reduction programs intended to mitigate existing disparities. This analysis thus reveals a new mechanism for understanding how inequalities in political voice map onto inequalities in policy outcomes, rooted in quotidian interactions between citizens and bureaucrats. By showing that the administration of social programs by bu-

⁵For seminal audit experiments on politicians conducted in low- and middle-income settings, see McClendon (2016) and Distelhorst and Hou (2017).

reaucrants can reinforce inequality traps in developing contexts, I highlight the magnitude of the challenge in the design of large scale programs to effectively reduce poverty in light of recent proliferation of these programs in Latin America (De la O, 2015; Garay, 2016).

2 Theory

2.1 Model

The model consists of three actors: a citizen (C), a street-level bureaucrat (B), and a politician (P). The citizen identifies with one of two groups, $g \in \{x, y\}$. The two groups of citizens vary in perceived costs of accessing the state. These costs capture some combination of physical distance, familiarity with bureaucratic procedures, and education (Kruks-Wisner, 2018; Rizzo, 2018). Costs, $c_C \in \mathbb{R}_+$, are common knowledge and are distributed according to group-specific densities $f_g(\cdot)$ with cdfs denoted $F_g(\cdot)$. Without loss of generality, assume that $F_y(c_C) \leq F_x(c_C)$, or that F_y first-order stochastically dominates F_x . This implies that in the aggregate, group y faces higher costs of accessing the state than group x . Both the bureaucrat and the politician, indexed by i , may prefer to provide the citizen of one group with the service than the other group. These tastes are represented as $\gamma_i^g \in [0, 1]$ drawn from densities Γ_i^g . This bias is strictly taste-based (Becker, 1957). Alternatively, one could consider an altruistic bureaucrat and politician that internalize the benefits when the citizen receives service, with weights proportional to γ_i^g .

A bureaucrat responds to an exogenous citizen request for service by allocating effort, $e \in [0, 1]$. The service is provided with probability e . The bureaucrat's cost of exerting effort are proportional to the difficulty of the task, $c_B \geq 1$. When the service is not delivered, the citizen decides whether to complain to the politician, denoted $q \in \{0, 1\}$, at cost c_C . These citizen complaints are akin to the citizen appeals to politicians for constituency service described by Bussell (2019).

The politician observes whether or not a citizen lodged a complaint and subsequently chooses a level of effort to invest in auditing the work of the bureaucrat, $a \in [0, 1]$. Politicians therefore audit underprovision as opposed to misallocation of services. This setting characterizes many service provision settings where all citizens have a right to request service, referring to the majority of

tasks or programs in which equal treatment is a goal, at least through the initial application or enrollment stage.

With probability a , the politician is able to recover the service for the citizen. The politician benefits reputationally and thus electorally from the increase in service provision when she detects underprovision, parameterized as $S \in \mathbb{R}_+$.⁶ A biased politician will also gain utility from providing the service to a favored citizen. Failing to remedy a complaint induces a separate reputational cost normalized to 1. Finally, audits are costly, which constrains the intensity of auditing; the marginal cost of an audit on a given task is c_P . To avoid corner solutions, I assume that $c_P > S + 2$. In addition, $c_P > c_B$ implies that it is costlier for politicians to recover the service than for bureaucrats to provide it in the first place. This assumption is consistent with standard arguments that bureaucrats have greater expertise than their political principals (Weber, 2009; Gailmard and Patty, 2007). The politician's expected utility can thus be expressed as:

$$E[U_P(a)] = a(S + \gamma_P^g) - (1 - a)(q) - \frac{c_P a^2}{2} \quad (1)$$

The citizen receives a utility of $b > 0$ if she receives the service. The citizen's expected utility conditional on not having received the service from the bureaucrat is a function of the probability that the oversight process will recover the service and her decision to complain (q):

$$E[U_C(q)] = ba - qc_C. \quad (2)$$

Finally, consider the bureaucrat's utility. He gains utility proportional to γ_B^g by (directly) providing a favored citizen with service. If a decision is reversed during the course of an audit, bureaucrats incur a penalty normalized to 1. In practice, these costs range from a reprimand, to transfer, or even termination. The bureaucrat's expected utility is thus:

$$E[U_B(e)] = e\gamma_B^g + -(1 - e)(a) - \frac{c_B e^2}{2} \quad (3)$$

⁶In the interpretation of this model with an altruistic politician, one could consider S as the politician's internalization weight on the service provided to any citizen.

Parameter	Description
PANEL A: MODEL PRIMITIVES	
$c_C^g \sim f_g(\cdot)$	Citizen's cost of complaint.
b	Benefit of service to citizen.
c_B	Marginal cost of bureaucrat's effort
$\gamma_B^g \sim \Gamma_B^g(\cdot)$	Bureaucrat's taste for serving group g citizen
c_P	Marginal cost of politician's oversight
$\gamma_P^g \sim \Gamma_P^g(\cdot)$	Politician's taste for serving group g citizen
S	Politician's baseline reputational benefit from providing service.
PANEL B: ACTIONS	
q	Indicator for citizen complaint.
e	Bureaucrat's effort exerted to serving citizen.
a	Politician's rate of auditing/oversight.
PANEL C: AGGREGATION/MAPPING TO RESEARCH DESIGN	
Δ	Bureaucratic bias toward group x over group y .
η_B	Average bureaucratic taste for group x over group y .
η_P	Average politician taste for group x over group y .
η_Q	Differences in rate of complaint between group x and group y .

Table 1: Description of model parameters.

2.1.1 Sequence

The game proceeds as follows:

1. The bureaucrat chooses an effort level e to provide the service to the citizen.
2. The citizen decides whether or not to lodge a complaint to the politician.
3. The politician chooses the intensity of audits, a . With probability a she overturns the bureaucrat's decision.
4. Payoffs are realized.

I characterize the unique subgame perfect Nash equilibrium (SPNE) in pure strategies. The bureaucrat's allocation strategy sets $e \in [0, 1]$. The citizen's complaint strategy maps the realization of the service provided into a binary decision whether to complain to the politician $q : \{0, 1\} \rightarrow \{0, 1\}$. The politician's audit strategy then maps receipt of a complaint into auditing intensity, $a : \{0, 1\} \times \{0, 1\} \rightarrow [0, 1]$.

2.2 Results

The main results characterize the bureaucrat's equilibrium effort. I solve the model by backward induction, beginning with the politician's audit intensity. The politician's objective is clearly concave in a ; maximizing Equation 1 yields audit intensity:

$$a^* = \frac{S + \gamma_P^g + q}{c_P}. \quad (4)$$

The politician's audit strategy includes two types of oversight. S and γ_P^g represent “police patrols” for underprovision of the service while a complaint, q , represents a “fire alarm” (McCubbins and Schwartz, 1984). The citizen anticipates the politician's response when deciding whether to complain. Specifically, if the bureaucrat has failed to deliver the service, the citizen will complain if:

$$\frac{(S + \gamma_P^g + 1)}{c_P}b - c_C > \frac{S + \gamma_P^g}{c_P}b. \quad (5)$$

This yields an optimal complaint strategy of:

$$q^* = \begin{cases} 1 & \text{if } c_C < \frac{b}{c_P} \\ 0 & \text{if } c_C \geq \frac{b}{c_P}. \end{cases} \quad (6)$$

At higher costs of complaint, c_C , citizens are effectively “frozen out” of contesting the bureaucrat's service provision. Importantly, the politician's audit and citizen's complaint strategies map directly into the bureaucrat's determination of how much effort to exert. Substituting equations (4) and (6) into the bureaucrat's objective and maximizing, the bureaucrat's optimal effort is:

$$e_g^* = \min \left\{ \frac{\gamma_B^g}{c_B} + \frac{S + \gamma_P^g + \mathbb{I} \left[c_C < \frac{b}{c_P} \right]}{c_B c_P}, 1 \right\}, \quad (7)$$

where \mathbb{I} is an indicator function.

Collectively e_g^* , q^* , and a^* characterize the SPNE of the game. In Appendix A1.4, I endogenize

the citizen request for service by assuming that citizens pay a cost proportional to c_C to request the service in the first place. This extension builds upon Kruks-Wisner’s 2018 description of how citizens select into making claims to state services. Including this cost introduces two mechanisms through which service provision changes from the baseline results. Clearly, if costs are sufficiently large relative to the benefits of receiving the service, some citizens opt out, receiving no service. Less obviously, it changes the composition of the portion of a group that requests service. This increase in the conditional probability that a citizen that requests service will complain increases the average bureaucratic effort devoted to citizens that “opt in.”

2.3 Defining and Measuring Bureaucratic Bias

There are two measures of bias implied by the model: bias in effort and inequality in outputs. Bias in effort corresponds to different equilibrium levels of effort across groups. Inequality in outputs corresponds to different levels of ultimate service provision by group (see Appendix A1.2 for a formal definition). I assume that the effort and service afforded to each citizen is independent of the effort and service afforded to other citizens. In the context of service provision, if citizens request services at different times or on different days, this assumption is plausible. Even in environments in which bureaucrats face unmanageable caseloads such that more effort for one citizen implies less effort for another, so long as citizens receive service on a first-come-first-served basis and the order of requests is independent of group membership, aggregate bias can be captured by treating cases independently.

Definition 1. *Bias in effort.* *Bias in effort refers to the difference in expectation of equilibrium effort devoted to a citizen from each group, formally, $\Delta = \mathbb{E}[e_x^*] - \mathbb{E}[e_y^*]$.*

Definition 1 defines bias between groups in the aggregate. I focus on the case in which effort is interior, ($e_g^* < 1$) for all citizens. I characterize bias in terms of aggregate differences by group. Define differences in the expectation of bureaucrat’s tastes as $\eta_B = \mathbb{E}[\gamma_B^x] - \mathbb{E}[\gamma_B^y]$; differences in the expectation of politician’s tastes as $\eta_P = \mathbb{E}[\gamma_P^x] - \mathbb{E}[\gamma_P^y]$; and differences in the probability of complaint as $\eta_Q = F_x(c_C) - F_y(c_C)$. Note that the last term, η_Q , shows how departures from the

assumption of trivial costs of complaint in Prendergast (2003, 2007) are consequential. Complaint-driven bias emerges *because* some citizens are more able to complain than others. When costs of complaint correlate with group membership, bureaucrats exert greater effort to stave off complaints from groups with louder voices.

Proposition 1. *Between-group bias in effort. The aggregate level of bias between groups x and y evaluates to:*

$$\Delta = \underbrace{\frac{\mathbb{E}[\gamma_B^x] - \mathbb{E}[\gamma_B^y]}{c_B}}_{\text{Bureaucrat's Tastes}} + \frac{1}{c_B c_P} \left[\underbrace{\mathbb{E}[\gamma_P^x] - \mathbb{E}[\gamma_P^y]}_{\text{Politician's Tastes}} + \underbrace{F_x\left(\frac{b}{c_P}\right) - F_y\left(\frac{b}{c_P}\right)}_{\text{Complaint-Driven}} \right] = \frac{\eta_B}{c_B} + \frac{(\eta_P + \eta_Q)}{c_B c_P} \quad (8)$$

(Proof in appendix.)

Proposition 1 implies three mechanisms that drive the bias in effort and outcomes. The differences η_P and η_B capture taste-driven biases of the politician and bureaucrat, respectively. The difference η_Q implies the potential for complaint-driven bias, which stems from differences in the *anticipated* rates of complaint across groups. Note that, in contrast to standard models of statistical discrimination in which group membership is observable and correlates with some latent trait (Becker, 1957; Phelps, 1972), I show that this form of discrimination emerges even with complete information. When one group is more able to complain, bureaucrats anticipate the increased probability of oversight by giving better service ex-ante. This is captured through a comparison of the distribution of costs for each group. The stochastic dominance assumption serves as a sufficient condition for complaint-driven bias to emerge on average (in the aggregate).

Of the three sources of bias, complaint-driven and the politician's taste-based biases are driven by oversight of the bureaucrat by the politician. In this sense, both forms of bias are *strategic*. Oversight is biased if $\eta_P + \eta_Q \neq 0$. When oversight is biased, bureaucrats exert differential effort in anticipation of variation in oversight intensity.

2.4 Testable Implications

The model posits three mechanisms underlying bureaucratic bias in effort. Learning which mechanisms are operative helps to understand the distributional consequences of bias in effort for “who (ultimately) gets what” and provides information about what which policy interventions might be successful in combatting bureaucratic biases in service provision. Proposition 2 derives several testable implications that I use to discriminate between types of bureaucratic bias.

Proposition 2. *Tests of the mechanism. Decomposing oversight-driven and non-oversight-driven bias:*

1. *Bias in effort varies in the politician’s cost of auditing if and only if oversight is biased: $\frac{\partial \Delta}{\partial c_P} \neq 0$, if and only if $\eta_P + \eta_Q \neq 0$. For sufficient increases in c_P , bias in effort attenuates toward zero.*

Discriminating politician’s tastes from complaint-driven bias:

1. *The magnitude of bias in effort increases in the between-group differences in ability to complain: $\frac{\partial \Delta}{\partial \eta_Q} > 0$ (< 0) if $\eta_Q > 0$ (< 0).*
2. *The magnitude of bias in effort increases in the between-group differences in the politician’s tastes: $\frac{\partial \Delta}{\partial \eta_P} > 0$ (< 0) if $\eta_P > 0$ (< 0).*

Proposition 2 guides efforts to test the mechanisms described in the model and in Table 2. I proceed in two steps. First, I test for evidence of oversight-driven bias. This distinguishes bias coming from bureaucrats’ tastes (η_B) from the bias that comes from different probabilities of oversight ($\eta_P + \eta_Q$). As in Table 2, oversight-driven bias incorporates both the politician’s tastes and complaint-driven bias. To do this, I examine variation in bias with respect to the politician’s cost of effort, c_P , which should only drive variation in the oversight mechanism.

Conditional on finding evidence of oversight-driven bias, I aim to distinguish between politician tastes and citizen propensity to complain, the two components of oversight-driven bias. To

Bias Mechanism	Classification		Case	Implications		
	Bias Type	Oversight-driven		$\frac{\partial \Delta}{\partial c_P}$	$\frac{\partial \Delta}{\partial \eta_Q}$	$\frac{\partial \Delta}{\partial \eta_P}$
<p>Complaint-driven: Citizens from group x are more likely to complain than from group y which draws a higher likelihood of auditing by the politician. The bureaucrat devotes more effort to x in anticipation of higher probability of audit.</p> <p>Politician’s tastes: The politician prefers to audit service to group x more than to group y. The bureaucrat devotes more effort to x in anticipation of higher probability of audit.</p> <p>Bureaucrat’s tastes: Bureaucrat prefers providing service to group x over group y.</p>	Statistical	Yes	$\Delta > 0$	$\neq 0^*$	> 0	0
			$\Delta < 0$	$\neq 0^*$	< 0	0
	Taste-based	Yes	$\Delta > 0$	$\neq 0^*$	0	> 0
			$\Delta < 0$	$\neq 0^*$	0	< 0
	Taste-based	No	$\Delta > 0$	0	0	0
			$\Delta < 0$	0	0	0

Table 2: Summary of the bias mechanisms implied by the theory and the testable implications for distinguishing the mechanisms. *Note that sufficient increases in c_P attenuate oversight-driven bias toward zero.

do this, I examine variation in citizens’ cost of complaint. As the “distance” between petitioner types’ costs of complaint, η_Q , increases, so too should the relative magnitude of complaint-driven bias. This test is able to distinguish between politician tastes and complaint-driven bias when $Cov(\eta_P, \eta_Q)$ is small. I also consider parallel tests with regard to politician incentives that may drive politician tastes, η_P with a parallel logic.

3 Case Context

I measure variation in bureaucratic discretion at national scale in Colombia. Writings on state capacity in Colombia have long focused on a two-century history of civil wars as a cause or consequence of state weakness (Centeno, 2003; González González, 2014). Nevertheless, characterizations of bureaucracy and cross-national indicators (Table A5) typically characterize Colombia’s national bureaucracy as comparatively “Weberian” by regional standards (Evans and Rauch, 1999;

Mayka, 2016). At the national level, bureaucratic capacity in Colombia is believed to vary with the relative shares of “technocrats” and patronage employees (e.g., Schmidt, 1974; Dargent, 2016).

Less is known about municipal bureaucrats, the subjects of this investigation. However, dozens of semi-structured interviews with local and national bureaucrats as well as participant observation in *alcaldías* evidence substantial variance in professionalism, competence, and outputs.⁷

3.1 Municipal Politics in Colombia

Since political decentralization in the late 1980s, Colombia’s 1102 municipalities have assumed responsibility for many services. Decentralization created larger and more professional municipal public administration (Fizbein, 1997). Important for this project, some national programs are now implemented “on the ground” within municipalities by municipal bureaucrats.

Municipalities are governed by an elected mayor and local council (*concejo*) of seven to 45 councilors, according to population. In local elections, parties are weak and the role of ideology is limited. Among mayors that governed from 2016-2019, only 21% could be identified as belonging to “right” or “left” parties per the Fergusson et al. (2021) party classification (see Table A21). In these contexts, the distribution of public and private goods – like the audited services – arguably constitutes the basis of political competition (Rueda and Ruiz, 2020).

3.2 Complaints and Oversight

I posit a fundamental role for citizen complaints as a means of seeking oversight over bureaucrats. Complaints take different forms spanning individual appeals, protests, and even use of intermediaries (i.e., brokers) in different contexts. These modes of complaint may be associated with different distributions of costs in the population. In the present context of social welfare provision in Colombia, individual appeals are the modal form of complaint.⁸ The Colombian Constitution of 1991 (Article 74) and statutory law (e.g., *Ley 1755 de 2015*, Article 14) protect Colombians’ rights to file various forms of complaints and claims and outline guarantees for timely government

⁷See Appendix A2 for a discussion in of the ethical considerations in the different forms of human subjects research described in this paper.

⁸Appendix A7 documents the empirical rarity of protests and the limited reach of brokers in this social policy context.

recourse. Complaints emerge through government entities' complaint systems (usually abbreviated PQRS) or through less formal channels; the distinction is not relevant for the purposes of the theory so long as costs vary across the population. The use of other forms of individual-level complaints is well documented in other Colombian social policy domains, most prominently through the *acción de tutela* in the subsidized healthcare system (Taylor, 2018).

Some complaints are ostensibly handled by other higher-level bureaucrats, while others rise to local mayors. While such mundane complaints are, in principle, hardly newsworthy, mayors do audit the local administration of social programs. There are regular news reports of mayors responding to complaints about the function of social programs, typically by auditing local rolls of beneficiaries.⁹ For such oversight to influence bureaucrats' behavior, threats of reprimand must be perceived. An original survey of street-level bureaucrats in *alcaldías* in Bogotá and Cundinamarca finds that 78% (57/73) of these bureaucrats perceive that a mistaken decision would be punished (with varying severity) and decisions would be reversed (see Appendix A4 for survey information).

While nationwide data on complaints to local entities is not collected, analysis of over 440,000 complaints filed in public entities in Bogotá from January 2017 through June 2018 and compiled by the Veeduría Distrital provides two stylized facts of note. First, virtually all complaints relate to service provision and approximately 125,000 complaints (28 percent) explicitly relate to bureaucrats' actions in service provision. This represents the most common class of complaints. Second, leveraging Bogotá's geographic segregation by class, analysis of complaints at the locality level suggests that there are substantially *more* complaints per capita in richer localities. Figure 2 shows a positive association between the the average class designation (*estrato*) of residential properties in each locality and the per-capita rate of complaint submission to each local *alcaldía*. This pattern emerges despite the fact that service provision is notably better in richer localities – those where complaints are most frequent.¹⁰

⁹Contemporaneous newsworthy investigations include investigation of how a councilor in Mosquera, Cundinamarca made it onto a list of means-tested beneficiaries for social programs (SISBÉN); a scam to stuff the rolls for Adultos Mayor, a subsidy for senior citizens, in Florencia, Caquetá; and a general audit of the SISBÉN rolls in Pitalito, Huila. The first two investigations occurred in response to citizen complaints.

¹⁰The model implies that the if the rich (lower c_C) receive better services such that they do not need to complain, the poor (higher c_C) could complain at higher rates than the rich. To the extent that the rich do complain more frequently

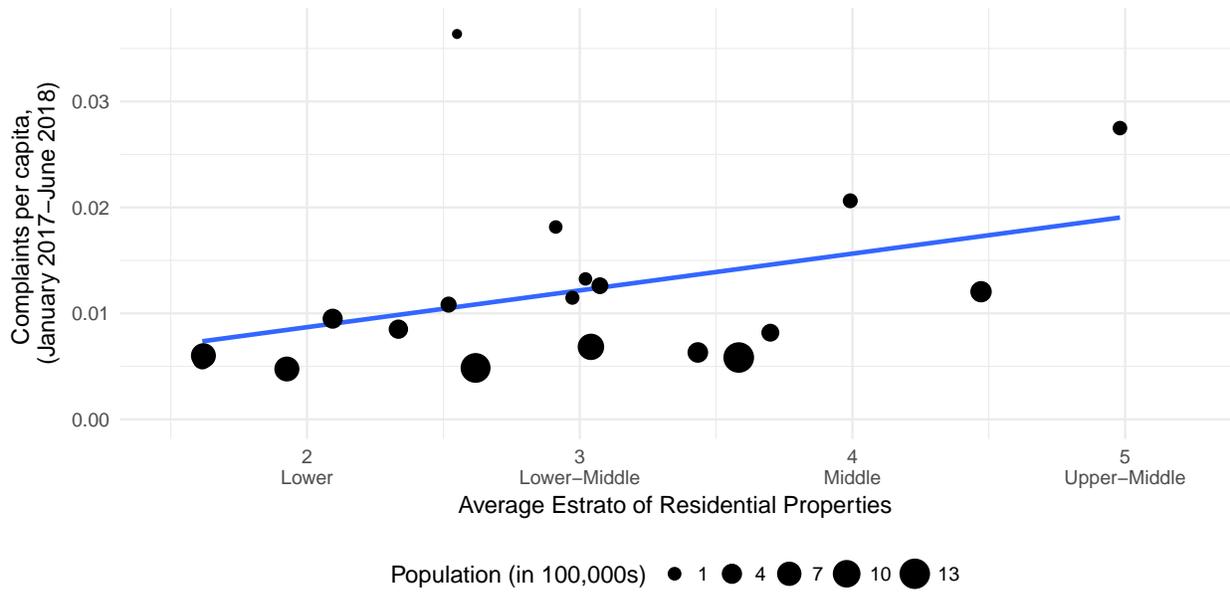


Figure 2: Rate of complaints filed by locality in Bogotá by average wealth of the locality.

While Figure 2 shows rates of individual complaints by aggregate levels of wealth, it does not necessarily establish that rich citizens within these localities complain at higher rates than the poor. Three waves of Latin American Public Opinion Project (LAPOP) data suggest that rates of a related form of complaint-making – protest participation – is similarly increasing in an individual’s socio-economic status (Figure A1). This pattern holds in comparisons between and within surveyed municipalities. As with the analysis of complaints in Bogotá, this correlation is consistent with the idea that richer citizens face lower costs of complaint, particularly given the conventional wisdom that these citizens encounter better treatment by the state in the first place.¹¹

3.3 Audited Social Welfare Services

I audit two nationwide social welfare programs that are administered, in part, by officials embedded in every municipal government. Specifically, the national government agencies that oversee these programs maintain agreements (*convenios*) with each municipality that mandates that local

despite receiving better service, this observation places an upper bound on the possible degree of bureaucrats’ taste-driven bias (see Appendix A1.6).

¹¹The audit experiment affords a test of this conventional wisdom and supports the idea that richer citizens do indeed receive better services.

	SISBÉN	Más Familias en Acción (MFA)
Description	Household index of assets used for means-testing	CCT for compliance with dependant child health, education requirements
Creation	1995	2002
Coverage	70% of population	15-20% of households
Benefit	Ability to enroll in social programs (including MFA) if qualified	Subsidies equivalent to 13-17% of median household consumption (Fiszbein et al., 2009)
<i>Alcaldía</i> tasks	Enrollment and actualization of status (score) through household survey of assets	Program enrollment and monitoring compliance with community participation requirements of the program.
Discretion (Politicization)	High – past manipulation of scores by citizens, local bureaucrats, politicians (Carmacho and Conover, 2011)	Low – comparatively low levels of discretion in implementation, at least by regional standards (De la O, 2015)

Table 3: Attributes of the audited social programs.

alcaldías hire at least one local program officials per program.¹² Local program officials administer enrollment and status actualization services that afford municipal residents access to these social programs. The tasks performed by local officials are available at the request of citizens. The programs, the System for the Identification of Beneficiaries of Social Programs (SISBÉN) and Más Familias en Acción (MFA), provide coverage on a nationwide geographic scale. At the time of the experiment, SISBÉN provides access to at least 24 means-tested programs including MFA. MFA is a CCT program that provides transfers to low-income families conditional on childrens' school attendance and compliance with healthcare visit requirements. Colombia is among world's most unequal countries with a very large low-income population, indicating a large but not universal pool of potential beneficiaries for each program. Existing analyses of the programs suggest very different levels of politicization between the two programs. Table 3 outlines attributes of these programs.

4 Research Design

I examine the sources of bureaucratic biases using a national-scale phone audit experiment that permits measurement of effort and identifies bias in effort, Δ . While I contend that anticipation

¹²The are different program officials for SISBÉN and MFA in each municipality and administration of each program should be seen as distinct processes.

of citizen complaints drive bias in effort, I do not manipulate or measure complaints within the experiment for three reasons. First, given selection into complaint-making, standard experimental estimands are undefined and thus unidentified (Slough, 2021). Second, as in the analysis of complaints in Bogotá, the model does not yield a falsifiable prediction about whether low- or high-cost groups should complain more given differential initial effort by bureaucrats. Finally and most importantly, (written) complaints require markedly more effort by *alcaldía* staff and are ethically unjustifiable given that they would not yield further learning about the theory. I address potential concerns about the “small” measures of effort by contextualizing estimates to actual levels of service provision (enrollment) in Section 7.

4.1 Audit Experiment

The phone audits consist of a call with an informational query, paralleling the assumption of an exogenous request for service in the model. The unit of random assignment is the petition. I utilize a factorial experiment to randomly assign characteristics of petitions. The treatments manipulate the bureaucrat’s and politician’s marginal cost of effort, c_B and c_P , as well as in observable attributes of the petitioner, g . Several aspects of petitions are constant across calls. Given the composition of MFA recipients – mostly mothers – all petitioners were female. While all calls were made from Bogotá, the outgoing phone numbers appeared as standard cell phone numbers. In Colombia, cell phone numbers do not convey geographic location.

4.1.1 Treatments

In order to measure bias, I manipulate identity-based characteristics of the petitioners. These characteristics are rooted in the Colombian social and political context, and serve as the analogue to the groups in the theoretical model. First, I assign the socioeconomic class (*estrato*) of the caller. Since independence (and before), class has represented an organizing feature of Colombian society and political life (Martz, 1997; Sanders, 2004). Given the focus of the social programs audited, I differentiate between low- and lower middle-class callers.¹³ Focus groups with Colombians of different

¹³Within Colombia’s political designation of class, lower class refers to *estratos* 1 and 2. Middle class refers to *estrato* 3.

socioeconomic classes and observation of calls in a government call center suggest several avenues in which the class of a caller can be immediately distinguished by phone. Specifically, callers in the two groups vary in their vocabulary, salutations for figures of authority (bureaucrats), and the framing of questions. While class communicates a variety of features, I assume that on average lower middle-class individuals have relatively greater ability to engage (access) the bureaucracy than lower-class individuals.

The second identity-based treatment focuses on migrant status. This treatment is communicated during the petition through a statement that the individual in need of the service is a recent (internal) migrant. The “resident” condition does not provide this information. Rates of internal migration have long stood among the highest in Latin America and encompass both ordinary and conflict-induced migration (Martine, 1975; Ibáñez and Vélez, 2008).¹⁴ The migrant condition signals two potentially relevant features. First, migrants are apt to have less familiarity with a municipality, suggesting higher costs of access. Second, internal migrants may be less likely to vote, especially given challenges of re-registration in new municipalities (Gaikwad and Nellis, 2017).

Given the nationwide reach of the audits, a third identity-based manipulation aims to reduce the threat of detection. Regional accents in Colombia are quite distinctive. While high levels of migration mean that most Colombians have exposure to those with other regional accents, some accents are rarer than others in different regions. To prevent detection, I randomly assign Bogotano, Paisa, and Costeño regional accents. These accents represent the most widely-spoken accents in Colombia and are collectively spoken by $\approx 60\%$ of the population (see Appendix A9.2).

The final manipulation varies the technical specificity of the petition, capturing the costs of service provision in the model.¹⁵ For both programs, the “easy” version of the question simply asks how a non-enrolled/registered citizen could enter SISBÉN or MFA, respectively. The “technical” version of the questions poses a question about a situation with specific technical program requirement. This manipulation allows me to test sensitivity of bias to costs to the bureaucrat which aids

¹⁴For example, 7.4 million Colombians are internally displaced, representing approximately 15% of the Colombian population (see Appendix A9.3).

¹⁵Under the assumptions of the model, that $c_P > c_B$ this may or may not also increase the politician’s cost of effort c_P commensurately.

Factor	Levels	Mode of Administration	Compliance Rate
EFFORT (COSTS)			
Difficulty of Request	• Easy	Technical specificity of request to petition, as defined by national government partners	99.3%
	• Difficult		99.2%
BIAS			
Regional Accent	• Bogotá	Regional accent of caller employed in interaction with bureaucrats.	99.7%
	• Paisa		98.4%
	• Costeño		98.7%
Socioeconomic Class	• Low	Vocabulary, salutations, and framing of the interaction. [†]	76.7%
	• Lower Middle		79.3%
Stated Migrant Status	• Migrant	One statement in delivery of petition (migrant). No reference to internal migration in resident's call.	97.3%
	• Resident		95.0%

Table 4: Factors and levels employed in the factorial design. Compliance rates are calculated as the proportion of calls correctly classified by double coders out of the number of calls assigned to each level. [†]While the framing of the interaction varied by class, the statement of the petition was stated identically for both classes.

in discriminating between bias mechanisms.

Collectively, all four factors are fully crossed, yielding a $2 \times 3 \times 2 \times 2$ factorial design summarized in Table 4. This yields 24 distinct treatments for each of two programs, though I analyze along the margins (by attribute). Note that the twelve confederates were actresses. All confederates voiced both low- and middle-class petitions. To maximize authenticity, actresses voiced only their own regional accent and calls were divided between four actresses per region of origin. Calls were randomly assigned to each confederate.

All calls were recorded. I hired Colombian coders to listen to all of the recordings to double code call characteristics and responses. Given that the coders were blinded to treatment assignment, this yields one measure of compliance with treatment assignment. I define compliance as a measure of whether coders reported hearing the assigned factors (i.e. if they heard a Costeña petitioner on a call assigned to a Costeño accent). The rates of compliance are reported in the final column of Table 4. A more detailed analysis of compliance is reported in Table A17. While I cannot know what bureaucrats intuited, rates of compliance in the double coding exercise are quite high across all factors and levels, alleviating major concerns.

4.1.2 Sampling, Assignment

The sample of *alcaldías* was selected with two opposing objectives. First, by maximizing the number of petitions made to the same *alcaldía*, I increase statistical efficiency and allow the estimation of within-*alcaldía* treatment effects. Second, I seek to minimize the probability of detection. In order to achieve both objectives, I stratify municipalities into three groups by estimated 2018 population. These entities are municipal *alcaldías* outside of Bogotá and local *alcaldías* in Bogotá. The number of petitions varies by stratum. In the large-population stratum, six petitions were assigned, three each for SISBEN and MFA. In the medium stratum, four petitions were assigned, two per program. In the small stratum, one petition was assigned per program. The distribution of the 1,836 petitions by municipal population category and program appears in Table A14.

Blocking by *alcaldía* ensures maximal within *alcaldía* variation and limits the threat of detection. The blocking procedures are detailed in Table A15. The blocking ensures that each *alcaldía* received equal numbers of low- and middle-class petitioners; equal numbers of easy and difficult questions; and received half the petitions from migrants. To minimize the likelihood of detection, the more specific technical questions were never repeated within an *alcaldía*. This implies that the ratio of easy to technical questions in the large stratum was 2:1. The estimation strategy accounts for these differential probabilities of assignment. Further, no *alcaldía* received more than one call from the same class/accents combination or was asked the same question more than once.

The order of calls was randomly assigned to space out calls to the same *alcaldía* over approximately four weeks. The assignment process for this rollout procedure is documented in Appendix A13.3. In general, first attempts of each call were consistent with the assigned ordering (within morning or afternoon), but repeated attempts complicate this mapping. Finally, the time of day – morning or afternoon – within each *alcaldía*'s hours of service was randomly assigned. Each *alcaldía* received equal numbers of calls at each time. Ultimately, just 6 calls were detected. Analysis of the detected calls yields no systematic patterns (see A15.2).

4.1.3 Outcomes

The audits measure a rich set of behavioral outcomes relating to service provision through the course of the call. Appendix A10 clarifies the sequencing of calls and outcome measurement. To measure service provision, all enumerators filled out an instrument to document the trajectory, outcomes, and information conveyed in each call. Outcome measures are coded from these responses and the double coder records of the calls.

I focus on two classes of outcomes. For the *alcaldías* reached by phone, I provide a mapping of the call through the *alcaldía*. Since dispatchers who answer are not generally program officers, I measure whether a petitioner was provided access to a program officer in order to make the petition. I map the mode of transmission through the bureaucracy to measure the accessibility and navigability of service providers within local bureaucracies. In particular, I measure four outcomes dichotomously: (1) whether the dispatcher identified himself/herself; (2) whether the petitioner was able to make (state) the petition; (3) whether the petitioner was connected to at least a second official; and (4) whether a program officer for SISBÉN or MFA from an ex-ante pre-treatment list was identified.

Most importantly, I measure agents' responses to the petition. I focus on the amount and veracity of information provided relative to the benchmark (correct) answers specified by the national government agencies that oversee each program. Outcomes at this stage also include a measure of red tape: whether an official asked for *extra* requirements not specified by program guidelines and whether petitioners were asked to come "in person" without further guidance. I measure five pre-registered outcomes of interest: (1) whether the correct, complete answer was provided; (2) whether partial information was provided; (3) whether any actionable information was provided (a sum of #1 and #2); (4) whether the petitioner was asked to come to the *alcaldía* in person without further instruction; and (5) whether red tape was solicited. The "come to the *alcaldía*" response merits some clarification. All services require an eventual trip to the *alcaldía*. Arriving without the requisite documents imposes additional costs on the petitioner, regardless of the bureaucrat's

intent.¹⁶

4.1.4 Ethical Considerations

Government audit experiments generally raise three ethical concerns: the use of deception, the protection of subjects, and the waste of time and public resources. I address the concern of deception through a novel model of collaboration with national government agencies. The collaboration included consultation throughout the research design process with the agency overseeing the Colombian bureaucracy at the national level (the Administrative Department of Public Administration) as well as the agencies overseeing SISBÉN (National Department of Planning) and MFA (Department of Social Prosperity). These agencies provided guidance on the programs to be audited, the content of the audits, the correct answers to the audits, and some administrative data. In exchange, I conducted the experiment independently with external funding and produced and presented a policy report to each agency three months after the experiment.¹⁷

Notably, these agencies conduct their own “mystery shopper” audits of employees and contractors periodically, though my collaborators did not recall randomizing shopper attributes. By conducting the audits independently, I provide additional privacy protections to subjects (audited bureaucrats) in a manner that is generally absent in government audits.

In terms of wasting resources, the costs to public entities in Colombia should be weighed against the benefits of this original data and report. The upper bound on the costs to these entities can be quantified quite simply. The answered calls (i.e. those that occupied the time of public employees) total under 200 hours. At the maximum monthly salary for the maximum rank of employee (“*Profesional*”) that would have spoken with a caller, the upper bound on the cost of these calls totals \$2,644 USD.¹⁸ This totals less than 10 months for one employee at the official minimum wage (as of 2018), a common local benchmark.

¹⁶Two plausible interpretations of the “come to the *alcaldía*” response include: (a) political capture is more likely to occur in person than on the phone; or (b) the bureaucrat believes that the petitioner will only understand in person. I remain agnostic between these interpretations but maintain that failure to provide information imposes an additional cost to petitioners.

¹⁷Relevant partners reported an intention to use the findings in a training for local agents, which serves as a debriefing may have increased bureaucrats’ awareness of bias in service provision.

¹⁸Calculated from *Decreto No. 309 de 2018*.

4.2 Estimation

I analyze the experimental data in two steps. First, I measure bureaucratic bias, identifying the parameter Δ . Second, I examine the mechanisms underlying bias using the testable implications in Table 2.

I measure bias by estimating the Average Marginal Component Effect (AMCE) of the randomly-assigned identity treatments. This quantity is the marginal effect of each identity factor, averaged over the joint distribution over other factors. I estimate the AMCEs using Equation 9 using OLS with heteroskedasticity-robust standard errors. Note that these standard errors correspond to the level of treatment assignment: the petition. The set of indicators in the regression model corresponds to the factor levels in the design, here $\mathbf{Z} = \{\text{Afternoon}_i, \text{Technical}_i, \text{Lower Middle Class}_i, \text{Bogotá accent}_i, \text{Costeño accent}_i, \text{and Resident}_i\}$. In Equation 9, ψ_a indicates *alcaldía* fixed effects. κ_p indicates a vector of program (SISBÉN or MFA) fixed effects that are included in all specifications.

$$Y_{ipm} = \sum_{j \in \mathbf{Z}} \beta_j Z_i^j + \kappa_p + \psi_m + \epsilon_{ipm} \quad (9)$$

I account for the differential probabilities of assignment to easy and technical questions across the strata of municipalities using inverse probability weighting (IPW) or *alcaldía* fixed effects. The latter strategy leverages variation only from within the same *alcaldía*. Further, the estimation of the AMCE accounts for the process of selection and the delivery of treatment during the course of the interactions with local government officials. The attributes (factors) in the factorial design were revealed at three distinct points in the call, as depicted in Table 5. This defines three relevant samples: all attempted calls, all answered calls, and all calls in which the petition was delivered. Factors not yet revealed in a given sample are referred to as *placebos*; factors revealed within the sample are referred to as *treatments*; and factors revealed prior to revelation of the a sample are regarded as *pre-treatment covariates*. Point estimates on the treatment variables (in the relevant sample) are causally identified. Taking advantage of the rollout of factors during the course of the call increases statistical efficiency and while avoiding the threat of bias induced by post-treatment

Call Made 1836 Calls 618 Municipalities		→	Call Answered 1194 Calls 466 Municipalities		→	Petition Made 911 Calls 424 Municipalities	
(Time of Day)	✓		(Time of Day)	Not point identified		(Time of Day)	Not point identified
Accent	Not revealed		Accent	✓		Accent	Not point identified
Class	Not revealed		Class	✓		Class	Not point identified
Difficulty	Not revealed		Difficulty	Not revealed		Difficulty	✓
Migrant Status	Not revealed		Migrant Status	Not revealed		Migrant Status	✓

Table 5: Timing of treatment delivery during the process of a call.

sample selection.

With multiple outcomes and high dimensional treatments, the design gives rise to some concerns of limited power, particularly for interaction terms, and of multiple comparisons problems. To alleviate these concerns and make inferences on more theoretically-relevant concepts, I aggregate “up” from the basic AMCE estimates presented here. To test for bias, I make inferences on the basis of F -tests (or the equivalent) on the joint significance of relevant coefficients. To estimate these models, I specify the subset of relevant estimators (β 's in Equation 10) and implement an F -test to test the null hypothesis that all β 's in the subset are equal to zero. I refrain from the use of high-dimensional interactions due to concerns of power. Note, however, that the inclusion of interactions between identity-based characteristics does not improve the predictive power of the models (see Appendix A15.4).

In order to test the mechanisms, I estimate conditional AMCEs with respect to institutional, demographic, or political covariates measuring other parameters of the model. I estimate Equation 10 using OLS with heteroskedasticity-robust standard errors. In this equation, moderators and covariates are represented by the variable X_i (resp. X_a). For a binary moderator, X_i , the conditional AMCEs are estimated by β_j and $\beta_j + \gamma_j$, where j indexes the treatment level.

$$Y_{ipm} = \sum_{j \in \mathbf{Z}} \beta_j Z_i^j + \sum_{j \in \mathbf{Z}} \gamma_j Z_i^j X_i + \kappa_p + \sum_{p \in P} \alpha_p X_i + \epsilon_{ipm} \quad (10)$$

5 Identifying Bias in Bureaucratic Effort

I first estimate the magnitude of bureaucratic bias by socioeconomic class, migrant status, and regional accent. Bias in effort, Δ measures how observed bureaucratic effort differs by randomly assigned petitioner identity characteristics. Given that the regional accent and socioeconomic class (ideally) were revealed as soon as the call was answered, I consider outcomes of process and access as well as the information provision outcomes. I analyze these outcomes on the full sample of answered calls ($n = 1,194$). Logically, the “unrevealed” factors (class, accent, migrant status, and petition difficulty) should be orthogonal to whether or not a call was answered or not. Reassuringly, F -tests of the joint significance of these factors provide no evidence of selection (imbalance) across unrevealed factors on the probability that a call was answered in Table A22.

5.1 Bias in Access to the *Alcaldía*

First, I investigate whether petitioner characteristics influence process of the petition through the *alcaldías* in Table 6. Columns 1-4 report the process of petitioners through the *alcaldía*. I find that dispatchers identified themselves in 85% of calls and do not detect differences as a function of petitioner identity (Column 1). Similarly, in Column 2, 75% of confederates were able to voice the petition to an official (the dispatcher or a program officer); I again do not recover evidence that these rates vary (robustly) with petitioner identity. Collectively, these estimates do not provide evidence that the dispatcher’s handling of calls varied by class or regional accent of the callers.

Within each *alcaldía*, I find no evidence that petitioner identity changes the rate at which a second official (ideally a program officer) is researched. However, Column 4 measures the rate at which the second official identified herself as one of the officials on the pre-treatment administrative lists of MFA and SISBÉN officials collected from national government partners. The results indicate somewhat higher levels of access to these program administrators for the middle-class petitioners relative to lower class petitioners, a difference of approximately 4.9 percentage points. Taken with Column 3, this finding likely emerges from higher levels of identification (by a second official) to middle-class petitioners. The joint test of coefficients on class and accent, however, is

only marginally significant. Collectively, these analyses suggest limited, if any, bias in navigating the *alcaldías* within an initial interaction on account of class or accent. These findings are important for two reasons. First, it suggests the outcomes measuring responses to the petition in Columns 5-10 can be interpreted as measures of bureaucratic effort in responding to questions, rather than differential diversion by the dispatcher. Second, the lack of differences by petitioner identity in Columns 1-3 (outcomes measuring dispatcher behavior) provide no evidence that bureaucrats (at least dispatchers) were “differentially confused” by some petitioner or script characteristics.¹⁹

5.2 Bias in Information Provided

Columns 5-10 examine bias in the responses to the petitions. Note that these responses are not conditional on making a petition; thus failing to receive information comprises both wrong responses and no response. Column 5 provides no evidence of bias in the probability that a petitioner receives a complete, correct response on the basis of the identity attributes. Note, however that baseline levels of correct responses are quite low. To the extent that bias represents the *withholding* of effort or information, there is limited scope to move this outcome. In this context, note that the (small) AMCE estimates on on lower-middle class and resident represent effect sizes of around 20% of this baseline.

There is notable bias in the likelihood of receiving a partial response or any information. Lower middle-class petitioners are substantially more likely to receive a partial response or any information relative to lower class petitioners in Column 6. In Column 7, the point estimate on receipt of any information is 8.1 percentage points and represents a 16 percent increase in the probability of receiving any information relative to the baseline (lower class). There is noisier evidence of a penalty against migrants.

Columns 8 and 9 track two outcomes in which information was not provided. Column 8, “no information” includes any response that did not provide individuals information or invite them to

¹⁹One concern is that because the lower class petitioner scripts were less direct, they may have confused the bureaucrats that answered the phones. There is no evidence that this was the case from these outcomes or qualitatively in the double coding.

	Access/Process				Response to Petition					
	Dispatcher Gave Name	Petition Made	Second Official	Program Officer	Complete	Incomplete	Any Info.	No Info.	Alcaldía Only	Red Tape
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
PANEL A: IPW ESTIMATES										
Lower-Middle Class	0.011 (0.020)	0.014 (0.025)	-0.014 (0.028)	0.043 (0.026)	0.021 (0.018)	0.048* (0.029)	0.069** (0.029)	-0.019 (0.021)	-0.036** (0.017)	0.003 (0.024)
Bogotá Accent	0.029 (0.026)	0.038 (0.031)	0.030 (0.034)	0.027 (0.033)	0.002 (0.022)	-0.005 (0.035)	-0.003 (0.035)	0.026 (0.025)	0.015 (0.019)	-0.056* (0.030)
Costeño Accent	0.034 (0.026)	0.078** (0.031)	-0.020 (0.035)	-0.038 (0.032)	-0.007 (0.022)	0.033 (0.036)	0.026 (0.036)	0.020 (0.026)	0.032 (0.021)	-0.027 (0.030)
Resident	0.015 (0.020)	-0.018 (0.025)	0.015 (0.028)	-0.002 (0.026)	0.027 (0.018)	0.003 (0.029)	0.029 (0.029)	0.007 (0.021)	-0.054** (0.017)	0.036 (0.024)
F-test, <i>p</i> -value	0.472	0.071*	0.445	0.078*	0.462	0.352	0.100*	0.698	0.001***	0.197
Placebo F-test, <i>p</i> -value	0.400	0.761	0.643	0.891						
PANEL B: ESTIMATES WITH ENTITY FIXED EFFECTS										
Lower-Middle Class	-0.015 (0.019)	0.016 (0.025)	-0.021 (0.025)	0.046* (0.024)	0.024 (0.019)	0.057** (0.028)	0.081*** (0.028)	-0.027 (0.021)	-0.039** (0.017)	0.002 (0.026)
Bogotá Accent	0.023 (0.026)	0.002 (0.034)	-0.023 (0.035)	-0.016 (0.034)	-0.006 (0.026)	-0.044 (0.038)	-0.050 (0.038)	0.039 (0.028)	0.014 (0.022)	-0.098*** (0.036)
Costeño Accent	0.021 (0.026)	0.038 (0.034)	-0.063* (0.034)	-0.058* (0.031)	-0.005 (0.026)	-0.016 (0.038)	-0.021 (0.038)	0.025 (0.029)	0.034 (0.023)	-0.071** (0.035)
Resident	0.021 (0.019)	-0.020 (0.025)	0.022 (0.025)	-0.010 (0.024)	0.020 (0.019)	0.021 (0.028)	0.041 (0.028)	0.00005 (0.022)	-0.061*** (0.017)	0.041 (0.026)
F-test, <i>p</i> -value	0.612	0.495	0.193	0.060*	0.597	0.153	0.009***	0.447	0.000***	0.015**
Placebo F-test, <i>p</i> -value	0.173	0.693	0.481	0.840						
PANEL C: ESTIMATES WITH ENTITY + ENUMERATOR FIXED EFFECTS										
Lower-Middle Class	-0.016 (0.019)	0.019 (0.025)	-0.026 (0.025)	0.045* (0.024)	0.028 (0.019)	0.053* (0.028)	0.081*** (0.028)	-0.023 (0.021)	-0.039** (0.017)	0.004 (0.026)
Resident	0.020 (0.019)	-0.017 (0.025)	0.017 (0.025)	-0.014 (0.024)	0.022 (0.019)	0.019 (0.028)	0.041 (0.028)	0.003 (0.021)	-0.061*** (0.017)	0.040 (0.026)
F-test, <i>p</i> -value	0.468	0.263	0.020**	0.061*	0.033**	0.306	0.035**	0.130	0.001***	0.567
Placebo F-test, <i>p</i> -value	0.411	0.246	0.019**	0.166						
Mean DV, Lower Class	0.859	0.758	0.685	0.276	0.104	0.396	0.5	0.153	0.104	0.234
Mean DV, Paisa Accent	0.842	0.723	0.675	0.296	0.119	0.414	0.533	0.124	0.066	0.266
Mean DV, Resident	0.856	0.771	0.668	0.302	0.101	0.428	0.529	0.134	0.108	0.219
Program	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
All Factors	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
DV range	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}	{0, 1}
Observations	1,194	1,194	1,194	1,194	1,194	1,194	1,194	1,194	1,194	1,194

*p<0.1; **p<0.05; ***p<0.01

Table 6: Bias is estimated by the AMCEs of identity-based characteristics on access to bureaucrats and responses to petitions. All specifications are estimated in OLS with heteroskedasticity-robust standard errors. Panel C omits the accent manipulation because regional accents were voiced by natives of each region (not acted); joint tests in this panel refer only to class and migrant manipulations. The F-tests test the joint significance of coefficients on class and accent factors in Columns 1-4 (prior to the petition) and class, accent, and migrant factors in Columns 5-9. The placebo F-test test the joint significance of attributes of the petition, migrant status and technical question, in Columns 1-4.

come to the *alcaldía*.²⁰ These responses included hang-ups, “don’t know”-type responses, and situations in which the bureaucrat stated that they did not want to provide information. It is a relatively rare outcome and disproportionately impacts lower-income callers, though the point estimates and *F*-tests are not significant at conventional thresholds. Column 9 measures whether or not individuals were simply told to “come to the *alcaldía*” without further information. While all services require the person to come to the *alcaldía* with documents, failure to specify these requirements by phone passes the cost onto the citizen. The estimates suggest that lower-middle class individuals are 37.5 percent less likely to receive this response than lower class individuals, while residents are half as likely as migrants to receive the response.

The results in Column 10 indicate disproportionate use of red tape – a request for *extra* requirements – against Paisas relative to both Bogotanas and Costeñas, with sizable point estimates of 0.071 and 0.098, respectively. I do not find evidence that the use of red tape covaries with petitioner class – the primary targeting dimension for these programs.

The observed biases in information provision on the basis of petitioner class merit some additional discussion. It does seem that the class treatment was recognizable; independent coders identified the assigned coding in 77.5% of calls, as reported in Table 4.²¹ While class is necessarily a compound treatment in the Colombian context, analysis of the magnitude of the “complier” AMCE relative to the intent-to-treat AMCE in Table A23 suggests that bias enters through what blinded coders perceive to indicate social class within the calls.²²

Collectively the differences in treatment of lower-middle class versus lower class petitioners track those of residents versus migrants, though the class effects are substantively stronger. However, as indicated in Table 5, migrant status was not revealed until the petition was made. Table A26 reveals that these estimates are conservative and less efficient than estimates of migrant status on the sample of petitions alone. To the extent that these groups are relatively marginalized at

²⁰This outcome was not prespecified.

²¹The blinded coders were given an “I don’t know” option in addition to the two class categories; another 13.5% of calls fell into this category. Only 9% of calls were incorrectly classified.

²²The estimates of the complier AMCE can be seen as an informal test of the excludability assumption as applied to the social class treatment.

least within the experimental comparisons, these comparisons provide some evidence about the dynamics of bias that I explore in the next section.

Beyond the behavioral measures, confederates evaluated their interactions with bureaucrats after each call. These results, reported in Figure A16, suggest that perceptions largely aligned with the behavioral outcomes. Within enumerator and *alcaldía*, enumerators perceived worse service when calling as low-income petitioners, but there is no evidence of differences in respect by petitioner class. The alignment between the behavioral measurements and perceptions of the calls increases confidence in the behavioral measures.

5.3 Does Information Provision Reflect Costly Effort?

I seek to validate that information provision does indeed reflect exertion of costly effort. I consider the total amount of time spent on the call (mean: 4.83, standard deviation: 6.32 minutes). Because the scripts for the petitions varied in length of delivery across the identity characteristics, I lack the ability to identify differences across petitioner identities on this outcome. This represents an excludability violation: observation of a longer call could mean the bureaucrat spent more time answering the question or that the petition took longer to present.

Instead, I show that the length of calls is increasing in the amount of information provided (correct, partial, or no information). I first aim to purge differences in the length of calls due to variation in the experimental scripts. To do so, I fit a regression of logged call length (in minutes) on the experimental factors, a program indicator, and enumerator fixed effects with IPW. I then compare the distribution of residuals from this regression across the three types of outcomes. Figure A16 depicts the distribution of residualized call length by the amount of information provided as empirical CDFs (ECDFs). The graph indicates that the cumulative length of contact for petitions providing no information was substantially shorter than the length of those providing some information. On average, petitions receiving no information were 1.17 minutes shorter ($p < 0.01$) than calls providing partial information and 1.21 minutes ($p < 0.01$) shorter than calls providing complete answers. These differences represent effects of approximately 25 percent of the mean for calls with no information, suggesting that information provision does indeed reflect bureaucratic

effort.

6 Examining the Mechanism

The evidence of bias in information provision against lower class petitioners motivates analysis of the three mechanisms suggested by the model: bureaucrats' tastes, politicians' tastes, and complaint-driven bias. Disentangling the mechanisms is critical to developing appropriate policy responses to remedy bureaucratic biases and reduce resultant inequalities in access to state services. Appendix A1.7 lays out policy interventions that could reduce bias from each mechanism, showing that each mechanism generates distinct policy implications. The experiment measures bureaucrats' effort, meaning I do not measure complaints or oversight directly. Indeed, the question is not whether biased oversight leads to biased service outcomes, but how the likelihood of bureaucratic anticipation of oversight yields bias in effort. I use the model to identify the conditions under which oversight-driven bias should be magnified.

To conduct this analysis, I proceed in two steps, following Proposition 2. First, I endeavor to separate bureaucrats' taste-based bias from oversight-based bias (composed of politicians' taste-based and complaint-driven bias). Then, I tease apart complaint-driven bias and the politician's taste-based bias. These reduced-form tests follow directly from the comparative statics presented in Table 2.

6.1 Political Oversight Drives Bias

In order to disentangle bureaucratic taste-based bias from oversight-driven, bias, I consider how measured bias covaries with the politician's marginal cost of effort, c_P . As effort becomes more costly to the bureaucrat and audits become more costly for the politician, the bureaucrat's effort should decline. Per Proposition 2, increases in these costs should also attenuate bias.

I consider two possible measures of c_P within the observed data. I consider the experimental manipulation of an easy vs. technical petition as a shock to c_B and c_P . An easy petition requires less effort to answer, corresponding to lower costs. This is borne out in the lower rates of accurate response and in anecdotal accounts of bureaucrats searching for answers in program documentatio.

The compound shock to c_B and c_P merits caution in interpretation. To this end, I also examine variation in bias across the programs as an observational test that leverages SISBÉN’s politicization (relative to MFA). I operationalize politicization through a smaller c_P for SISBÉN than for MFA. It is less costly for a politician to intervene in SISBÉN for two reasons. First, MFA is a more technical program to administer with more stringent requirements for entrance and specific conditionalities to maintain access. Second, given the rates of enrollment in Table 3, MFA is a smaller program (by construction) and enrollment is therefore a less routine task of *alcaldías*, with fewer potential opportunities for politicians to learn program procedures and requirements.²³ Proposition 2 suggests that if observed biases are oversight-driven, we should expect to observe more bias on easy and SISBÉN petitions, respectively.

Consistent with this logic, Figure 3 suggests that bias in information provision is detectable *only* on easy petitions and for SISBÉN petitions. In the left panel, bias is most strongly apparent in easy petitions for the provision of *any* information ($p < 0.002$). Indeed, the difference in these conditional AMCEs on class in the provision of any information is substantively large at 10.1 percentage points and statistically significant at the $\alpha = .1$ level in a two-tailed test ($p = 0.079$). In the right panel, there is clear class-based bias in the the provision of any information for SISBÉN ($p = 0.013$) but no evidence of bias in the administration of MFA ($p = 0.421$). The difference-in-difference estimate on the interaction between class and program is sizable at 6.7 percentage points, but is not statistically significant at conventional thresholds ($p = 0.24$).

Caution in interpretation of these findings is warranted. If a technical petition increases both c_B and c_P , attenuation of bias could be driven by changes in either parameter (as opposed to both). Note that per Proposition 2, a simultaneous increase in both costs should increase the relative contribution of the bureaucrat’s taste-based bias to the estimated bias in effort. Stated another way, if the oversight-driven biases were entirely absent, the reduction of bias observed in the data corresponds to a very large shock to c_B . Importantly, service provision is not commensurately

²³It may also be the case that politicians value the provision of SISBÉN to any constituents over the provision of MFA, corresponding to a higher S for SISBÉN than for MFA. As long as bureaucratic effort remains interior ($e^* < 1$), Proposition 1 shows that such differences do not affect bias in effort. The rates of information provision in Table 6 suggest that the assumption of interior effort is reasonable in this setting.

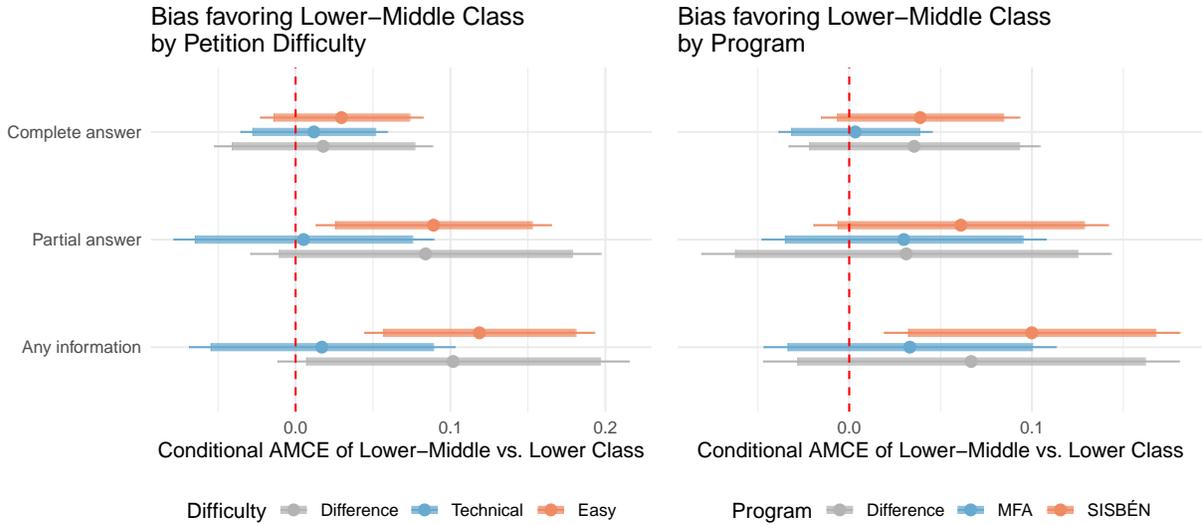


Figure 3: Sensitivity of bias in information provision to the cost of effort (difficulty of the petition) (left) and program (SISBÉN or MFA) (right). 90% (thick) and 95% (thin) confidence intervals constructed on heteroskedasticity-robust standard errors.

driven to zero (Table 6). As such, there is scope to observe taste-based bias, but I do not detect any. This test provides no evidence against the oversight-driven interpretation of findings.

The finding is strengthened by the observational finding on differences levels of bias across programs. Bias attenuates substantially for MFA, where politicization is lower, relative to SISBÉN, where politicization is higher. These tests provide evidence that oversight, not simply bureaucrats’ tastes, drives observed biases. Formally, it suggests that $\Delta_O \neq 0$.

6.2 Bias Occurs where Differences in Costs of Complaints is Greatest

Given the evidence that measured biases in bureaucratic effort are oversight-driven, I now seek to disentangle whether these biases are complaint-driven or generated by politicians’ tastes, following Proposition 2. Recall that complaint-driven bias is produced by the differences in the costs of complaint between groups of citizens, as measured by η_Q . I first motivate an operationalization of η_Q using municipal poverty rates and then consider several operationalizations of politicians’ tastes, η_P , to examine how measured bias in bureaucratic effort (Δ) varies in these parameters.

The lower- and lower-middle class petitioner profiles were fixed across all municipalities in the experiment. However, the *relative* standing of these profiles – particularly the lower-middle

class profile – varies substantially across Colombian municipalities. Using survey data and census microdata, I show that the relative position (rank) of a lower-middle class individual within a municipality is increasing in the municipal poverty rate (Figures A2 and A3). In other words, in a high poverty municipality, a lower-middle income petitioner is ranked higher than they would be in a low-poverty municipality. If costs of complaint are a function of relative position, the differences in costs of complaint for low- and lower-middle income petitioner profiles should be greater in high-poverty municipalities. In Figure A5, I provide support for the assumption that costs of complaint-making – again measured using reported protest participation as a proxy – are decreasing in the rank of an individual within their municipality. Collectively, these two assumptions support the operationalization of η_Q with municipal poverty rates.

Proposition 2 suggests that the magnitude of bureaucratic bias should increase in municipal poverty rates (η_Q) if complaint-driven bias is operative. Figure 4 examines bias in information provision as a function of the portion of residents in poverty as per the multidimensional index of poverty, calculated from the 2005 census.²⁴ The figure shows that anti-poor bias emerges against poor petitioners only in poorer places. The bias is restricted to the enrollment questions (left column) and reception of partial information or the *alcaldía* only response, as described above.

To subject these graphical intuitions to a more rigorous test, I run a series of regression analyses in Table A30. I bin the poverty index into terciles to reduce functional form assumptions on the moderator. Because poverty and population are strongly negatively correlated ($\rho = -.61$ in the sample), I include an interactive binned population control with deciles of the estimated 2018 population in a second estimator. Both the moderator, municipal poverty, and the (demeaned) population decile bin controls are interacted across the whole design (all factors and the program indicator).

This analysis suggests that bias against lower-class individuals (the baseline) is worse in poorer places. There is no evidence of bias in the lowest tercile (the municipalities with the lowest poverty

²⁴This index is compiled by the Departamento Administrativo Nacional de Estadísticas (DANE) at the level of rural and urban populations within each municipality. I take the weighted average where weights correspond to the share of urban and rural residents in the population.

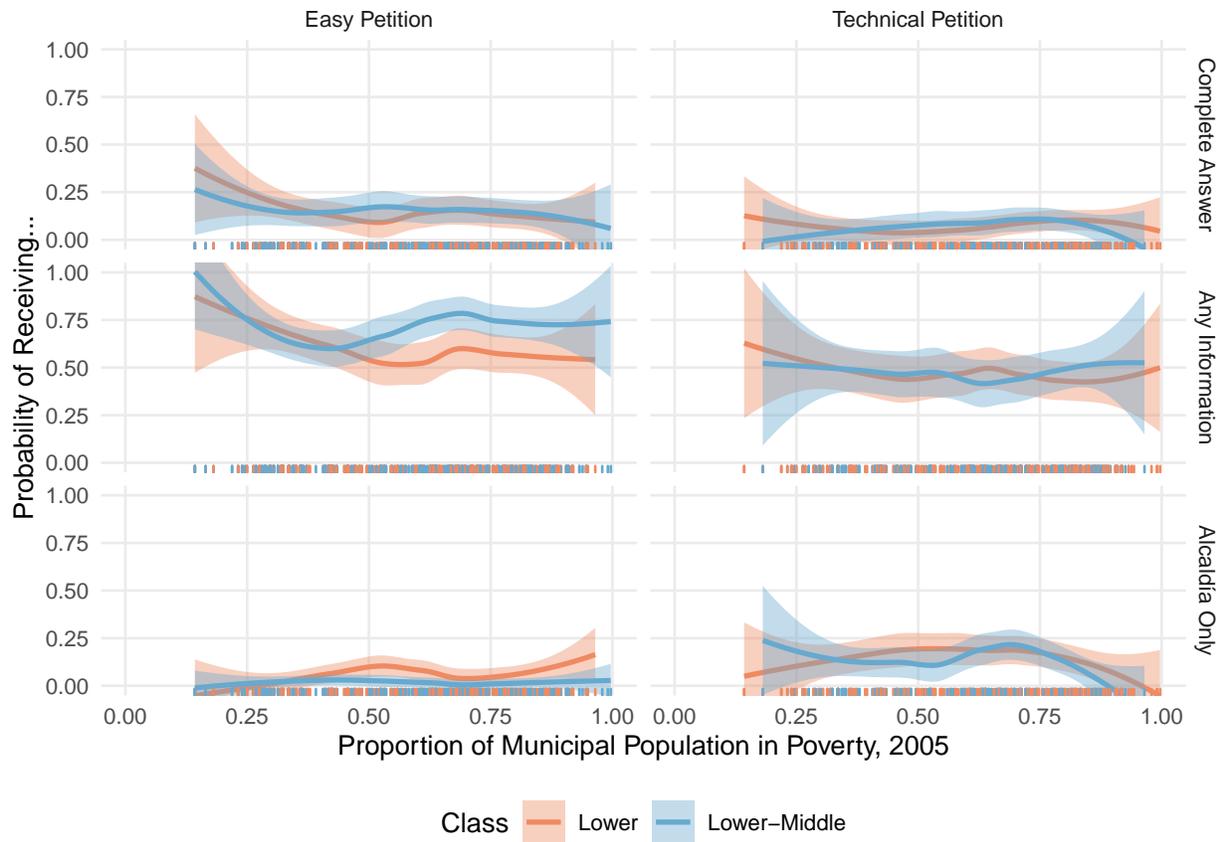


Figure 4: Heterogeneity in level of class-based bias by the level of municipal poverty. Column 1 examines average marginal effects on “easy” (enrollment) questions while Column 2 examines average marginal effects on technical questions. Lines are estimated by Loess regression. The shaded regions are 95% confidence intervals.

rates) for any outcome. Instead, bias against the poor is increasing in the middle-poverty and high-poverty terciles. I find clear, statistically significant evidence of bias in the high-poverty tercile for the receipt of partial information. Further there is suggestive evidence that differential application of the “*alcaldía* only” outcome against poor individuals is driven by poorer municipalities. These findings are robust to other operationalizations of poverty including rates of secondary education (2005).

Could variation in municipal poverty rates instead capture some aspect of politician tastes instead of differences in cost of complaint? For this to be the case, politicians in high-poverty municipalities would need to prefer to serve lower-middle over lower-class citizens and politicians in low-poverty municipalities would need to be neutral (unbiased). While I cannot eliminate this possibility, several observations are useful to consider. First, there is no evidence that service favors the median voter in each municipality. If service were to favor the median voter, the poor should do the *best* in the highest poverty places; these are the places that they do the *worst*. From an elected politician’s perspective, therefore, overseeing worse service provision for poorer individuals in high-poverty municipalities works against the median voter.

Other explanations of bias in terms of politician tastes do not account for the observed incidence of bureaucratic bias in effort only in high-poverty municipalities. While there may be a disproportionate incentive to politicize social programs to claim credit or buy votes in poorer places with more potential recipients, it is not clear why such opportunities to claim credit would yield *unequal* service provision, as opposed to simply less information provision by bureaucrats. If politicians aimed to usurp social service provision responsibilities from bureaucrats as in Weitz-Shapiro (2012), we would expect these practices to lead to with lower *levels* of service provision by bureaucrats – the measured outcome. Importantly, as is evident in Figure 4, there is no evidence that bureaucrats provide less information to lower-middle class petitioners in high-poverty municipalities, they simply withhold information from poor petitioners. Further, as I document in Table A32, while clientelism practices are highly regional in Colombia, the emergence of bureaucratic bias against the poor in high-poverty municipalities similarly obtains in *within* region

and department analyses. Finally, I leverage the municipal classification of electoral risk including clientelism, corruption, and electoral violence by Colombia’s Mission of Electoral Observers (MOE) to show that these patterns persist when interactively controlling for these features (Misión de Observación Electoral, 2018). Given these analyses, unless politician tastes vary systematically in unmeasured ways with the degree of poverty in a municipality, there is little evidence supportive of politicians’ taste-based bias driving the emergence of bias only in high-poverty municipalities.

In examining other plausible measures of politicians tastes (η_P), I find no evidence that bureaucratic bias varies in any of several measures of political competition in Figures A21 and A22. Competition could drive politician incentives to provide public service to all, implying higher S (Camacho and Conover, 2011) or lead to greater effort to serve some fraction of the electorate, γ_P^g . Per Equation 1, the former should not drive bureaucratic bias in effort, but the latter might (if faction is correlated with class). Further, in Figure A22, there is no evidence that bias against lower class petitioners varies in mayoral ideology, another potential measure of politicians’ tastes for serving different groups of citizens.

I therefore argue that the most plausible interpretation of the finding is that where differentials in relative ability to complain between the lower- and lower-middle income treatment conditions are theoretically the strongest, levels of bias against the “quieter” group group is strongest. Lower-middle class individuals are relatively more empowered to complain in places where a majority of the population is poor. Structurally, this analysis suggests that $\eta_Q > 0$.

7 Discussion: Bias in Effort and Inequality in Outputs

To what extent does bias in information provision map onto inequality in public service outputs? Bureaucratic bias in effort is important because of its link to inequality in citizen access to public services. While the experiment allows for measurement of bureaucratic effort in information provision, for practical and ethical reasons, confederates did not try to obtain the service. To connect effort outcomes to actual service provision, I use pretreatment data on SISBÉN registration from across Colombian municipalities as reported in Figure 1. I examine the correspondence between

rates of enrollment (outputs) and the experimental measures of bias.²⁵ I investigate whether we observe higher rates of bureaucratic bias in municipalities where SISBÉN is under-provided than in municipalities in which it is plausibly administered according to program guidelines.

Table 7 suggests that bias presents in precisely in the municipalities in which under enrollment of plausible beneficiaries is the strongest concern. There is strong evidence of bias in information provision in the base category (under-enrolled) municipalities. This bias is substantively, and for some outcomes, significantly attenuated in municipalities with ostensibly “intended” enrollment. These results are robust to redefinition of the “plausible enrollment” category (see Appendix A20). While it is evident that under-enrollment occurs in poorer places, in Panel B the results are robust to controlling interactively for municipal poverty, population, budgets, and internal fiscal capacity (following Ch et al.’s (2018) measure). Additional robustness tests in Figure A23 suggest that this association is robust to a wide array of covariate specifications. This finding that the bias in effort measured in the experiment correlates with public service outputs is consistent with the model’s prediction that bias in bureaucratic effort yields inequality in service provision. These results are also consistent with the theoretical extension of endogenous requests for service. In places with where prospects for service are lowest, Colombians that anticipate poor service may opt out of pursuing SISBÉN registration altogether. These findings complement existing explanations for unevenness in the provision of social programs in Latin America (De la O, 2015; Niedzwiecki, 2018) by providing an explanation for within-municipality inequality in the provision of these services.

8 Conclusion

Observers of Latin American social policy regularly identify unevenness in the implementation of programs intended to reduce inequality (Holland, 2015; Niedzwiecki, 2018). I argue that disparities in administration can emerge in the course of everyday processes of service provision even without such political directives. By characterizing service provision as a strategic relationship

²⁵MFA data by municipality is not publicly available. However, aside from IDPs and indigenous Colombians, SISBÉN is used to qualify for MFA. As such, under-enrollment of SISBÉN should predict under-enrollment of MFA.

between a politician, a bureaucrat, and a citizen, I contribute a new mechanism through which political inequality in ability to draw oversight from a politician leads to inequality in access to social programs. Empirically, I show that bias against lower-class petitioners in the provision of information is substantial, but occurs only where oversight is most likely and in municipalities where inequalities in voice are apt to be strongest.

How generalizable is complaint-driven bias? Under the general scope condition that citizens have an opportunity to complain about service provision by bureaucrats, we should expect complaint-driven bias to be most prominent in settings where inequalities in ability to complain are largest and where bureaucratic insulation from political oversight is most limited (see Table A1). Highly unequal societies are therefore apt to exhibit particularly high levels of complaint-driven biases. Further, complaint-driven bias emerges when bureaucrats are sensitive to political oversight, which is particularly likely when bureaucratic insulation is limited (i.e., outside of civil service systems). The evidence presented in this paper suggests that these biases emerge even in the harder case of much more subtle comparisons: between two relatively disadvantaged – but ranked – groups (lower and lower-middle classes) and across two social programs with different levels of local political interference. As such, estimated levels of bureaucratic bias may understate the scope of complaint-driven bias in Colombia and other highly unequal societies.

The complaint-driven bias presented here presents new policy implications for reducing inequality in access to services. Citizen-focused interventions to minimize disparities in the costs of engaging the bureaucracy while reducing these costs for all hold promise for improving service provision and reducing complaint-driven bias. Recall that the prospect of complaint increases bureaucrats' effort, thereby improving service provision. Eliminating complaint-driven oversight would eliminate complaint-driven bias, but at a cost to overall levels of service provision. Instead, sufficient reductions to barriers to complaint can improve equity by reducing between-group disparities in rates of complaint while promoting more efficient service provision.

The concept of complaint-driven bias suggests that inequalities in political voice reduce the efficacy of the state programs to combat inequality. I posit that mundane processes of service

provision may contribute to inequality traps. This mechanism complements literature linking economic inequality to political inequality through more explicit conflict between the interests of elites and non-elites (Acemoglu et al., 2008; Acemoglu and Robinson, 2008). Future research should develop these arguments by integrating processes of production or implementation – not simply budgets – into the study of distributive politics. To this end, broader consideration of relationships between three actors – politicians, bureaucrats, and citizens – may advance our understanding of accountability, inequality, and redistribution.

	Complete (1)	Incomplete (2)	Any Information (3)
PANEL A: CONDITIONAL AMCE ON CLASS BIAS BY MUNICIPAL SISBÉN ENROLLMENT TYPE			
Lower-Middle class	0.083* (0.045)	0.105* (0.062)	0.188*** (0.062)
Lower-Middle class × Plausible enrollment	-0.073 (0.049)	-0.086 (0.072)	-0.159** (0.071)
Conditional effect, Plausible Enrollment Munis.	0.01 (0.022)	0.018 (0.037)	0.029 (0.035)
PANEL B: WITH MUNICIPAL COVARIATES			
Lower-Middle class	0.088* (0.053)	0.063 (0.080)	0.152* (0.078)
Lower-Middle class × Plausible enrollment	-0.071 (0.062)	-0.036 (0.094)	-0.106 (0.091)
Conditional effect, Plausible Enrollment Munis.	0.018 (0.024)	0.027 (0.04)	0.045 (0.037)
Interactive poverty control	✓	✓	✓
Interactive population control	✓	✓	✓
Interactive municipal budget control	✓	✓	✓
Interactive municipal tax capacity control	✓	✓	✓
Mean, Lower Class and Plausible Enrollment	0.108	0.45	0.558
Mean, Lower Class and Under Enrollment	0.084	0.379	0.463
Program Indicator	✓	✓	✓
Enumerator Indicator	✓	✓	✓
All Factors	✓	✓	✓
Observations	901	901	901

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 7: Relationship between bias in information provision and underprovision of SISBÉN. OLS estimates of the conditional AMCE of class by municipal SISBÉN enrollment type. The sample includes places that are under enrolled or plausibly enrolled as intended. Standard errors are clustered by municipality ($n = 366$) because the enrollment is measured at the municipal level.

References

- Acemoglu, Daron, and James A. Robinson. 2008. "Persistence of Power, Elites, and Institutions." *American Economic Review* 98 (1): 267–293.
- Acemoglu, Daron, Maria Angelica Bautista, Pablo Querubín, and James A. Robinson. 2008. "Economic and Political Inequality in Development: The Case of Cundinamarca, Colombia". In *Institutions and Economic Performance*, ed. Elhanan Helpman. Cambridge: Harvard University Press pp. 181–245.
- Banks, Jeffrey S. 1989. "Agency Budgets, Cost Information, and Auditing." *American Journal of Political Science* 33 (3): 670–699.
- Becker, Gary S. 1957. *The Economics of Discrimination*. 2nd edition ed. Chicago, IL: University of Chicago Press.
- Bussell, Jennifer. 2019. *Clients and Constituents*. New York: Oxford University Press.
- Bussell, Jennifer. 2020. "Shadowing as a Tool for Studying Political Elites." *Political Analysis* 28: 469–486.
- Camacho, Adriana, and Emily Conover. 2011. "Manipulation of Social Program Eligibility." *American Economic Journal: Economic Policy* 3: 41–65.
- Centeno, Miguel Angel. 2003. *Blood and Debt: War and the Nation-State in Latin America*. University Park, PA: Penn State University Press.
- Ch, Rafael, Jacob Shapiro, Abbey Steele, and Juan F. Vargas. 2018. "Endogenous Taxation in Ongoing Internal Conflict: The Case of Colombia." *American Political Science Review* 112 (4): 996–1015.
- Dargent, Eduardo. 2016. *Tchnocracy and Democracy in Latin America*. New York: Cambridge University Press.
- De la O, Ana. 2015. *Crafting Policies to End Poverty in Latin America*. New York: Cambridge University Press.
- Distelhorst, Greg, and Yu Hou. 2017. "Constituency Service under Nondemocratic Rule: Evidence from China." *Journal of Politics* 79 (3): 1024–1040.
- Evans, Peter B., and James E. Rauch. 1999. "Bureaucracy and Growth: A Cross-National Analysis of the Effects of "Weberian" State Structures on Economic Growth." *American Sociological Review* 64 (5): 748–765.
- Fergusson, Leopoldo, Pablo Querubín, Nelson A. Ruiz, and Juan F. Vargas. 2021. "The Real Winner's Curse." *American Journal of Political Science* 65 (1): 52–68.
- Fiszbein, Ariel, Norbert Schady, Francisco H.G. Ferrerira, Margaret Grosh, Niall Keleher, Pedro Olinto, and Emmanuel Skoufias. 2009. *Conditional Cash Transfers: Reducing Present and Future Poverty*. Washington DC: The International Bank for Reconstruction and Development/The World Bank.

- Fizbein, Ariel. 1997. "The Emergence of Local Capacity: Lessons From Colombia." *World Development* 25 (7): 1029–1043.
- Gaikwad, Nikhar, and Gareth Nellis. 2017. "Do Politicians Discriminate Against Internal Migrants? Evidence from Nationwide Field Experiments in India." Working paper.
- Gailmard, Sean, and John W. Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51 (4): 873–889.
- Garay, Candelaria. 2016. *Social Policy Expansion in Latin America*. New York: Cambridge University Press.
- Golden, Miriam, and Brian Min. 2013. "Distributive Politics Around the World." *Annual Review of Political Science* 16: 73–99.
- González González, Fernán E. 2014. *Poder y Violencia en Colombia*. Bogotá, D.C.: ODECOFICINEP.
- Hemker, Johannes, and Anselm Rink. 2017. "Multiple Dimensions of Bureaucratic Discrimination: Evidence from German Welfare Offices." *American Journal of Political Science* Forthcoming.
- Holland, Alisha C. 2015. "The Distributive Politics of Enforcement." *American Journal of Political Science* 59 (2): 357–371.
- Holland, Alisha C. 2017. *Forbearance as Redistribution: The Politics of Informal Welfare in Latin America*. New York, NY: Cambridge University Press.
- Ibáñez, Ana María, and Carlos Eduardo Vélez. 2008. "Civil Conflict and Forced Migration: The Micro Determinants and Welfare Losses of Displacement in Colombia." *World Development* 36 (4): 659–676.
- Kruks-Wisner, Gabrielle. 2018. *Claiming the State: Active Citizenship and Social Welfare*. New York, NY: Cambridge University Press.
- Lipsky, Michael. 2010. *Street-Level Democracy: Dilemmas of the Individual in Public Service*. 2nd edition ed. New York, NY: Russell Sage Foundation.
- Martine, George. 1975. "Volume, Characteristics and Consequences of Internal Migration in Colombia." *Demography* 12 (2): 193–208.
- Martz, John D. 1997. *The Politics of Clientelism in Colombia*. New Brunswick, NJ: Transaction Publishers.
- Mayka, Lindsay R. 2016. "Colombia's Surprising Resilience." *Journal of Democracy* 27 (3): 139–147.
- McClendon, Gwyneth. 2016. "Race and Responsiveness: An Experiment with South African Politicians." *Journal of Experimental Political Science* 3 (1): 60–74.

- McCubbins, Matthew D., and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* 28 (1): 165–179.
- Misión de Observación Electoral. 2018. Mapas y factores de riesgo electoral: Elecciones nacionales Colombia 2018. Technical report Misión de Observación Electoral Bogotá, D.C.: .
- Niedzwiecki, Sara. 2018. *Uneven Social Policies: The Politics of Subnational Variation in Latin America*. New York: Cambridge University Press.
- Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* 62 (4): 659–661.
- Prendergast, Canice. 2003. "The Limits of Bureaucratic Efficiency." *Journal of Political Economy* 111 (5): 929–958.
- Prendergast, Canice. 2007. "The Motivation and Bias of Bureaucrats." *American Economic Review* 96 (1): 180–196.
- Rizzo, Tesalia. 2018. "When Clients Exit: Breaking the Clientelistic Feedback Loop." Working paper.
- Rueda, Miguel R., and Nelson A. Ruiz. 2020. "Political Agency, Election Quality, and Corruption." *Journal of Politics* 82 (4): 1256–1270.
- Sanders, James E. 2004. *Contentious Republicans: Popular Politics, Race, and Class in Nineteenth-Century Colombia*. Durham, NC: Duke University Press.
- Schmidt, Steffen W. 1974. "Bureaucrats as Modernizing Brokers? Clientelism in Colombia." *Comparative Politics* 6 (3): 425–450.
- Slough, Tara. 2021. "On Theory and Identification: When and Why We Need Theory for Causal Identification." Working paper available at http://taraslough.com/assets/pdf/theory_id.pdf.
- Taylor, Whitney K. 2018. "Ambivalent Legal Mobilization: Perceptions of Justice and the Use of the Tutela in Colombia." *Law & Society Review* 52 (2): 337–367.
- Ting, Michael M. 2017. "Politics and Administration." *American Journal of Political Science* 61 (2): 305–319.
- Weber, Max. 2009. *From Max Weber: Essays in Sociology*. 1st edition ed. London: Routledge.
- Weitz-Shapiro, Rebecca. 2012. "What Wins Votes: Why Some Politicians Opt Out of Clientelism." *American Journal of Political Science* 56 (3): 568–583.
- White, Ariel R., Noah L. Nathan, and Julie K. Faller. 2015. "What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials." *American Political Science Review* 109 (1): 129–142.