(This is a preliminary draft; full version upon request)

Beyond the Promise: Distributional Discordance in China's Pandemic Control

Jingyang Rui*　　Zijin Zhang†

Abstract

Authoritarian governments often distribute resources to secure support, but what do the recipients truly receive beyond the surface of these allocations? In this study, we propose a new distribution theory that accounts for both sincere and insincere distributive outcomes by considering the misaligned incentives between policy designers and implementers. We leverage a unique dataset of 30 million real-time queueing observations at thousands of COVID-19 testing sites in H city, one of the largest metropolises in southeast China. Utilizing a border analysis of queueing time alongside an NLP analysis of government documents, our paper examines the distribution of testing resources across housing communities of different income levels before and after a popular protest which fundamentally shifted the preferences of policy implementers. Our findings reveal that, prior to the protest, high-income groups received covert benefits, low-income groups received theatrical benefits, and middle-income groups received void benefits. After the protest, low-income groups began to receive substantive benefits, while both high- and middle-income groups received covert benefits. Beyond contemporary China, our theory helps explain patterns of performative distribution in a wide range of political regimes.

*PhD Candidate, Department of Political Science, University of Michigan, Ann Arbor. `jyrui@umich.edu`

†PhD Candidate, Department of Technology and Operations, Stephen M. Ross School of Business, University of Michigan, Ann Arbor. `zijinz@umich.edu`

1 Introduction

Authoritarian governments often distribute resources to gain support (Gandhi and Lust-Okar 2009; Magaloni 2006), but what do the recipients receive beyond what it looks like they were given? A grand array of scholarships tackles distributive politics from the designer's angle: how to choose the size and identity of recipients to maximize the support that the distribution can generate (Blaydes 2011; De Mesquita et al. 2005)? What should be distributed (Geddes 2001; Levitsky and Way 2010; Liu 2020)? And when should distribution occur (Greene 2010; Lust 2006)?

When the intentions of the givers are aligned with data on the autocrats' fate following resource distribution, it deepens our understanding of how regimes survive (Albertus, Fenner and Slater 2018; Cansunar 2022; De Mesquita et al. 2005; Hong, Park and Yang 2023; Magaloni 2008; Svolik 2012; Wallace 2014). However, this understanding relies on the assumption that recipients received what the givers claimed to provide. Yet, evidence from the recipients' side has shown that the intended and actual distribution patterns often diverge. Scholars have noted that the provision of public goods in authoritarian regimes frequently becomes mere window-dressing, offering little genuine benefit despite the apparent good intentions of higher-level designers (Ding 2022; Scott 1998). A notable example is the construction of monumental but underused public spaces in Turkmenistan, where grand infrastructure projects such as Ashgabat's white marble buildings serve more as symbols of the regime's power than functional public goods for citizens (Koch 2018). Similarly, in North Korea, high-profile public housing projects in Pyongyang aim to showcase the regime's strength and commitment to its people, yet the housing remains inaccessible to most citizens, with benefits disproportionately directed to elite groups (Buzo 2017). Ding (2022) offers a compelling theoretical framework to explain such performative state actions, identifying public scrutiny and state capacity as key determinants of governance style. Specifically, when public scrutiny is high and state capacity is low, the state tends to engage in performative rather than substantive governance.

Attributing the misalignment between intended and actual distribution patterns solely to state hypocrisy or a lack of capacity, however, may be an oversimplification. Distribution outcomes can falter even when the state has genuine intentions. While central leaders set policy guidelines, the actual distribution of limited resources falls to local bureaucrats, who often face markedly different incentives than the autocrats. For instance, while autocrats prioritize long-term regime survival, bureaucrats are more concerned with short-term job security. In vying for promotion, bureaucrats may cultivate their own clientele, which can lead to the distribution of resources in ways that diverge from the autocrats' goals. Although bureaucrats cannot openly defy their superiors, due to oversight, they may subtly adjust distributions to favor their personal interests. This incentive misalignment within the hierarchical structure creates various opportunities for deviations between the claimed and actual distribution patterns across different groups—window-dressing being only one possibility.

So, how do actual distribution patterns vary across different social groups? A similar phenomenon—the selective enforcement of laws by the state to mobilize voters and signal distributive commitments in the absence of adequate social policies—has been explored in the context of selective forbearance (Holland 2016). However, it is important to recognize that selective forbearance is generally more passive compared to selective distribution. While Holland interprets selective forbearance as a reflection of distributive politics within the framework of institutional weakness, the issue of incentive misalignment persists even in authoritarian regimes with strong institutions and professionalized bureaucracies. These distinctions highlight the need for a theoretical framework that specifically addresses selective and potentially insincere distribution. This issue becomes particularly intriguing when we focus on the distribution of resources to civilian groups in authoritarian regimes—groups that are uniformly denied political power and, consequently, natural access to benefits.

In this paper, we develop a new theory of distribution that examines how actual distribution patterns in authoritarian regimes may deviate from central policy designs. We argue that bureaucrats in these regimes shape their distributional decisions based on the interactive effects of their superiors' preferences and their own need to complete tasks. In our framework, each social group is either preferred by the policy designers (the state) or not, and is either critical to the policy implementers' (the bureaucrats') task completion or not, in a relative sense. Under imperfect supervision—typical in most authoritarian systems—this interaction yields four possible distribution outcomes: substantive distribution, theatrical distribution, covert distribution, and void distribution (see Table 1). When a social group ranks highly in the priorities of both policy designers and implementers, they receive substantive distribution, characterized an advertised distributive intent that matches actual distribution outcomes. If a group ranks highly in the designers' priorities but lowly in the implementers', they experience theatrical distribution, where distribution is designed as highly visible and symbolic without addressing the underlying needs or demands of the population. Conversely, when a group ranks low in the designers' priorities but high in the implementers', they receive covert distribution, meaning the discreet allocation of resources to a social group without the gesture or spectacle of distribution. Lastly, when a group ranks low in both the designers' and implementers' priorities, they encounter void distribution, where no significant resources or attentions are directed toward them.

While this paper does not aim to provide a comprehensive list of which social groups are preferred by central leaders and local bureaucrats, an example of incentive misalignment might be that, in authoritarian regimes, central leaders often prioritize distributing resources to groups that contribute to the regime's long-term survival. In contrast, local bureaucrats tend to favor groups that are crucial for achieving short-term, measurable targets or mitigating immediate threats to their own survival.

A significant challenge for scholars in understanding actual distribution patterns lies not only in data availability but also in identifying suitable empirical settings for analysis. In this paper, we empirically examine the divergence between claimed and actual distribution patterns across different civilian groups by analyzing the allocation of COVID-19 testing

		Policy Designer's priority ranking	
		Low	High
Policy Implementer's priority ranking	Low	Void Distribution	Theatrical Distribution
	High	Covert Distribution	Substantive Distribution

Table 1. Theoretical framework

resources among various economic classes in H City, a large metropolis and economic powerhouse in southeast China, before and after the outbreak of popular protests against the zero-COVID policy in October 2022. Since the onset of the pandemic, the Chinese central government implemented what was arguably the world's strictest and longest zero-COVID policy, employing measures such as lockdowns, quarantines, and regular mass testing. Beginning in April 2022, a mandatory polymerase chain reaction (PCR) testing policy was enforced, requiring citizens to get tested at designated government-operated testing booths every 24 to 72 hours to maintain access to public spaces. Access to closer testing sites with shorter queues, therefore, became a crucial factor in shaping residents' life satisfaction and, by extension, their attitudes toward the government, as it saved them significant time in the testing process. As a result, the preferential distribution of COVID-19 testing resources can be viewed as a form of "public relief" that eased the lives of certain groups under the regime's extensive surveillance.

Street-level governments were responsible for staffing and funding COVID-19 testing efforts. The quality of resources received by each neighborhood can be measured along two dimensions: testing site density, which refers to the number of testing sites per capita in each housing community, and government responsiveness, which reflects the efficiency with which the government addresses long queues at nearby testing sites by either adjusting staffing levels or increase/change opening times. The time a resident spends on COVID-19 testing consists of two components: traveling time to a testing site and queueing time. Testing site density affects the former, while government responsiveness to busyness affects the latter: A higher density of testing sites does not necessarily reduce queueing time, as many residents may choose to test at the same time (e.g., after work), leading to sporadically long queues throughout the day. queueing time, in turn, can only be decreased through dynamic staffing in response to queueing patterns at the sites. Therefore, it is only when both testing site density and government responsiveness are optimized that a resident's overall time spent on testing can be substantially reduced. Importantly, while testing site density is visible to all users of the testing app, government responsiveness is much less apparent because it requires continuous monitoring of queueing time patterns over days or weeks. As a result, we classify the provision of only testing site density without corresponding government responsiveness as theatrical distribution. The combination of both high testing site density and government responsiveness constitutes substantive distribution. When government responsiveness is provided without a higher density of testing sites, we label this as covert distribution. Lastly, the relative lack of both is categorized as void distribution.

Our analysis of the time an average resident spent on COVID-19 testing reveals significant variations in resource distribution, despite the virus's indiscriminate spread. As shown in Figure 1, when factoring in both walking and queueing times, the average middle-income resident spent nearly twice as much time on each COVID test compared to the average high-income resident. This disparity in time consumption only began to level out after the anti-regime protests in Beijing heightened the risk of systemic unrest for local bureaucrats.

To identify the four types of distribution, we constructed a dataset that records real-time queueing times for every ten minutes at over 7,000 testing sites in H City from September to December 2022, by scraping data from the COVID-19 testing app published by the H City government. Using the geographical locations of these testing sites, we mapped them to 4,654 housing communities, categorized into high-income, middle-income, and low-income neighborhoods based on housing prices. We then compared testing resource density, the visible benefit, and government responsiveness, the invisible benefit, across these neighborhoods, analyzing how local government priorities regarding social groups shifted before and after the outbreak of popular protests in China.

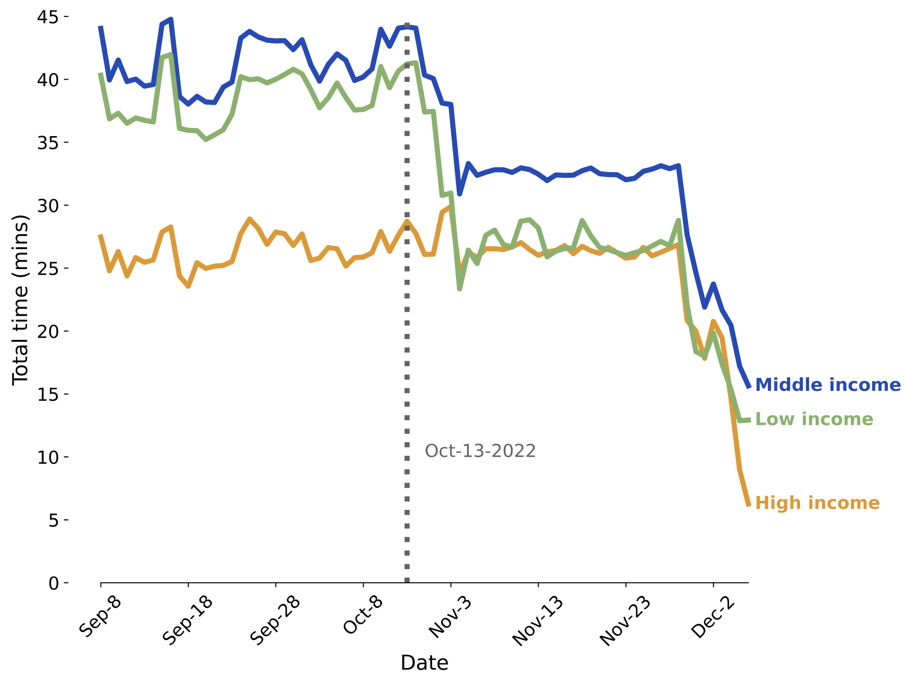

Figure 1. Time consumption on each COVID test by different income groups (Source: Author's calculation, H city, Sep 8 2022-Dec 8 2022)

We chose to focus on economic classes because of their salience in H City, where economic growth is primarily driven by major technology companies and Foreign Direct Investment. According to official statistics, H City is one of China's most populous and economically developed cities and one of the largest immigrant hubs, with some of the highest housing prices in the world and stark wealth inequality. While the average price of a two-bedroom apartment exceeds $900,000, more than 40% of the population lives in crowded urban villages, low-income residences, or public rental housing. This disparity creates incentive misalignments between central leadership and local bureaucrats. The low-income group is the most favored under the current Chinese Communist Party (CCP) propaganda, as the CCP's legitimacy rests heavily on this demographic, evidenced by increases in the central government's poverty alleviation budget and its allocation to COVID-19-affected areas. However, local bureaucrats prioritize groups that help them meet short-term targets, such as economic growth and social stability. In H City, the high-income group (entrepreneurs, investors) is the most critical and irreplaceable for these bureaucrats. Although the middle class (white-collar workers) is also important for economic growth, it is highly replaceable given H City's popularity as a job destination for college graduates. Similarly, the low-income group (blue-collar migrant workers, white-collar workers relatively new to the city) is economically important but also replaceable. A key factor that can shift the prioritization of social groups is their contentious status: when a social group is involved in or threatens to engage in unrest, its priority to the local government may either increase or decrease, depending on the nature of the unrest.

To assess the impact of income on distributional patterns, we utilize a spatial border analysis approach on neighboring housing communities. Due to H City's relatively young and rapidly expanding nature—having been established only 40 years ago—housing communities of different income levels frequently border one another. By matching housing communities from different economic classes that are within 200 meters of each other and fall under the same street-level government, we naturally control for confounding variables, as these neighboring communities share the same resource provider and pedestrian flow. Figure 2 provides an example of a pair comprising a low-income community and a middle-income community, located just 173 meters apart.

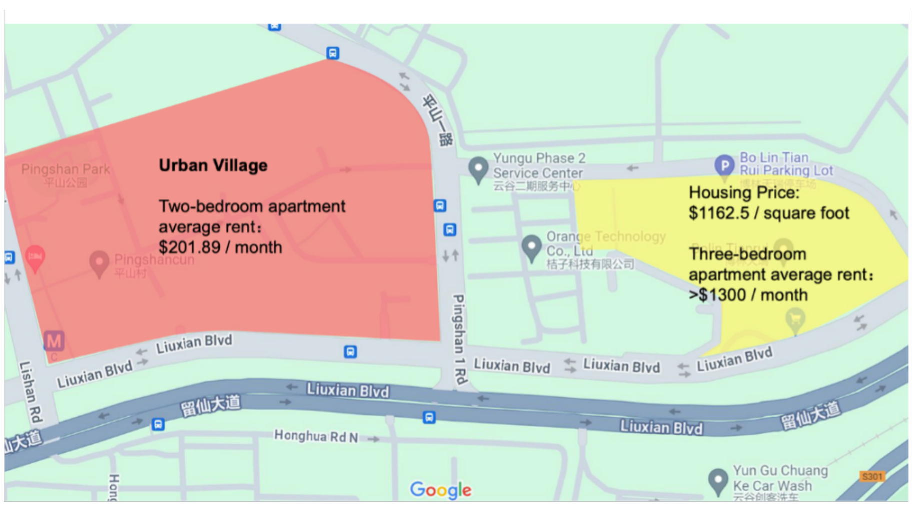

Figure 2. A community pair

While resource distribution toward different social groups often occurs under varying policies, our case presents a rare opportunity to observe the government's distributional intentions across different social groups under a single, uniform policy. This allows us to bypass many confounding factors and issues of incomparability. Furthermore, because COVID-19 testing is fully monopolized by the state, there are no alternative providers of this resource. The quality of COVID testing is also critically important for daily life during the pandemic, and its provision cannot be compensated by other forms of distribution, which strengthens the validity of our identification.

Our empirical results support our theory. We find that, in the absence of protest threats, H City's government provided performative benefits to the low-income group as they were favored by the central government but not local bureaucrats, covert benefits to the high-income group as they were favored by local bureaucrats but not the central government, and void benefits to the middle-income group as they are not the most favored group by either central or local government. However, following the outbreak of protests in Beijing, which heightened the risk of a systemic crisis, the local government increased the priority of all groups. H City's government began distributing substantive benefits to the low-income group and covert benefits to both the low- and middle-income communities, summarized in Table 2 and Table 3.

		Policy Designer's priority ranking	
		Low	High
Policy Implementer's priority ranking	Low	Void Distribution (Middle income)	Theatrical Distribution (Low income)
	High	Covert Distribution (High income)	Substantive Distribution

Table 2. Change of distribution patterns (before protest)

To verify that the empirical patterns reflect the intentions of the bureaucrats, we employ natural language processing (NLP) models to analyze the propaganda on COVID-19 resource distribution in H City. We include all articles related to COVID-19 control published by the 11 city- and district-level government official accounts (gōngzhònghào) on WeChat, China's most popular social media platform, between June 2022 and December 2022. These articles

		Policy Designer's priority ranking	
		Low	High
Policy Implementer's priority ranking	Low	Void Distribution	Theatrical Distribution
	High	Covert Distribution (High income) (Middle income)	Substantive Distribution (Low income)

Table 3. Change of distribution patterns (after protest)

cover narratives of the government's efforts to assist citizens in their fight against COVID-19, including the distribution of testing resources, medical supplies, everyday commodities, and the organization of events aimed at improving living standards during the pandemic. Our analysis reveals that the government was significantly more inclined to mention—and provide detailed accounts of—its support for low-income groups, both before and after the protests. In contrast, high-income groups, despite receiving the highest levels of COVID-19 testing resources, were mentioned far less frequently. This finding supports our theory of theatrical and covert distributional patterns.

2 Data

2.1 Primary Datasets

We have two major sources of data. Our first dataset is collected by scraping data from a COVID-testing app[1] published by H City's government between September 8, 2022 and December 8, 2022. Every 10 minutes from 8am to 10pm daily, we sent individual queries via API to access queueing status of each testing site located throughout the H city.

We construct an original dataset that includes the locations of testing sites (including its district, subdistrict, community, address, latitude, and longitude), the populations with

[1] The app was launched by the local government of H city, which allows users to access the app via WeChat and conveniently check the latest updates regarding the queueing status at nearby testing sites before making a walk-in visit. According to the report from the backend data of this system, it had an average daily visit of 6 million users, with the peak time for resident visits being from 11 AM to 4 PM.

access to these sites, the number of testing counters at each location, and the real-time queueing status. On the app, queueing status is categorized into four levels: free (queueing time $<$ 15 minutes), busy (queueing time 15–30 minutes), packed (queueing time $>$ 30 minutes), and closed. In total, our data consists of 27,216,264 real-time queueing observations across all 6,945 testing sites in H City (excluding "closed"). Our data collection concluded with the end of China's zero-COVID policy. Figure 3 provides a screenshot of the app searching page. This function allows users to select specific districts and subdistricts to locate nearby testing sites at any time. Each testing site's location is marked on the map, and their real-time waiting status is frequently updated by the system.

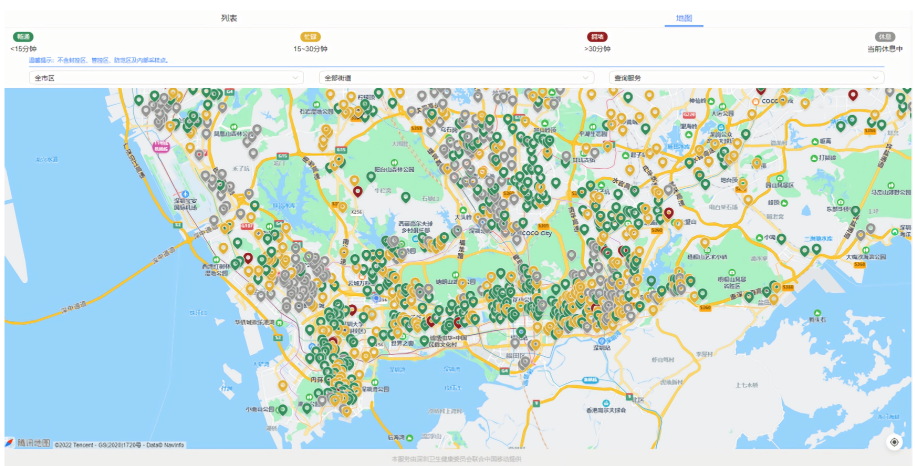

Figure 3. Screenshot of searching page of the mobile system

Our another primary dataset consists of N city's residential real estate and property data along with urban village data. This dataset details H City's housing communities, which includes 3,512 commercial housing communities and 1,142 urban villages. Since this paper focuses on understanding the government's distribution intentions toward civilian groups, we exclude housing communities affiliated with governments, state-owned enterprises, and public institutions from our comparison of distribution across income groups. Distribution to these communities will be used as a comparison in the robustness checks as an exemplification for distribution pattern towards the preferred group (political elites). For the non-village housing communities, we collect data on location, population, and housing prices, while for urban villages, we collect data on location and population. We do not collect housing price data for urban villages, as apartments in these areas are primarily for rental purposes.

2.2 Supplementary Datasets

Additionally, we gather information on the risk level (feng-xian-ji-bie) assigned to each community daily, based on the number of new cases or close contacts. When a community is categorized as high-risk, residents receive at-home testing, rendering the queueing time at their testing sites irrelevant. Therefore, we only include communities during low-risk days in our dataset. We then match housing communities with the testing sites that were exclusively accessible to their residents. Our dataset also combines rich data including daily COVID reports, GPS footprint data, news articles published by the local government, etc., to create a unique and novel dataset.

Our data is an ideal dataset to test authoritarian government's resource allocation strategy towards different economic groups. From official statistics, H city is China's one of the most populous and economically developed cities and one of the largest immigrant cities with world's highest housing prices and high wealth inequality – while an average two-bedroom unit sells for over $900,000, more than 40 percent of the population live in crowded urban villages, low-income residence, and public rental houses. The young and expanding nature of H city means that housing communities of different income levels often neighbour each other, thereby lending us the advantage of conducting a a spatial border analysis on adjunct housing communities.

3 Empirical Methods

We define the low-income group in H city as the residents in the urban villages. Urban villages in H city emerged as rural settlements that were engulfed by Shenzhen's rapid urban expansion. When the government converted rural land for urban use, the villagers retained ownership of their housing, which was often substandard and not built to modern city codes. Urban villages offer low-cost housing, making them attractive to migrant workers, low-income residents, and those who wish to save up for purchasing properties. Rents are significantly cheaper than in surrounding urban areas for the infrastructure is often less developed (see Figure 4). It is worth noting that residents in urban villages do not only consist of the migrant workers—by 2023, over 50 percent were those in what are typically perceived as white-collar and "middle-class" jobs, such as IT workers. Because over 95% of the apartments in these villages are rented, using property price to measure the village residents' income level is irrelevant. Therefore, we assign a dummy variable of 1 = low-income housing communities to the urban villages.

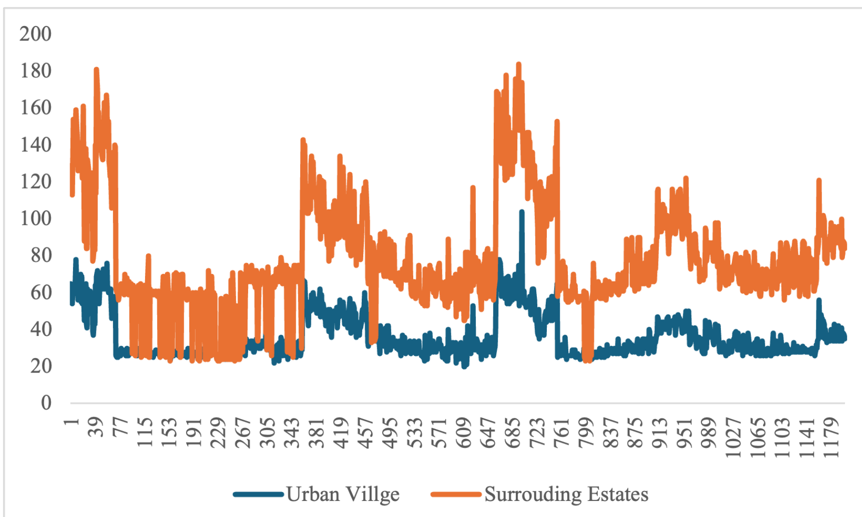

Figure 4. Comparison of rent in urban villages and surrounding estates (Source: Shanhu Data)

We assign income levels to the remaining 60% of the population based on the property prices of housing estates. In our main analysis, for non-village housing estates, we categorize residents in communities with property prices below the 80th percentile (RMB 92,300/USD 13,011.91 per square meter) as middle-income, and those above this threshold as high-income. Thus, the middle-income and high-income groups comprise approximately 48% and 12% of the population, respectively. Variations in the cutoff between middle- and high-income groups are used for robustness checks. In addition to the approximately 40% of residents living in urban villages, another 23.3% of the population rents homes in non-village housing communities. Since the proportion of renters in each non-village housing community is unknown, we use housing prices, rather than rent levels, as a proxy for income levels. We believe this approach does not affect the accuracy of income ranking for two reasons: rents in urban villages are systematically lower than those in all non-village communities, and rent levels are positively correlated with housing prices (insert data here). Apart from the around 40% of residents in urban villages, another 23.3% of the population rent houses in the non-village housing communities. Because the renting proportion in each non-village housing community is unknown, we use the housing price, rather than the rent level of the non-village housing communities to proxy for the income level. We believe that this practice would not interfere with the ranking of the residents' income level for two reasons: rents in urban villages are systematically lower than that in all non-village housing communities and that rent level is positively correlated to housing prices (insert data here).

We use the Sitong Bridge protest as the turning point that changed the local government's priority ranking. In Oct 16, a protest in Beijing, the capital of China, called for the end of the zero-COVID policy and even the stepping down of the Chinese leader right before the start of the 20th National Congress of the Chinese Communist Party. Because Chinese social media is subject to extensive censorship delegated to social media companies, many commentors argue that the sudden outburst of dissent caught both the state and local governments at surprise. Despite having occurred in Beijing, the protest, as one of the few high-profile acts of public dissent, heightened concerns about broader unrest in other major Chinese cities. The protest raised fears that similar acts of defiance could erupt in other metropolises, such as Shanghai, Guangzhou, or Shenzhen, where economic disparities, lockdowns, and public discontent with restrictive measures were also present. These cities, being economic powerhouses with large populations and significant migrant communities, were viewed as particularly vulnerable to unrest if frustrations over lockdown policies escalated. This incident eventually became the catalyst for the White Paper Revolution in late November 2022, where wide protests against the zero-COVID policy directly led to the policy's termination by early December 2022. Figure 1 shows that despite the zero-COVID policy being continued to December, the protest served as a powerful cutoff in that resource distribution pattern towards the income groups drastically changed distribution immediately following the protest.

As mentioned in the introduction, we pair all communities whose distance are within 200 meters. Among the 28,183 community pairs, 4,092 pairs include communities from different income groups. Our dependent variables are Site Allocation and Government Responsiveness, each measured separately for the pre- and post-shock periods. Site Allocation is defined as the average number of COVID testing sites allocated to the higher-income community minus the number allocated to the lower-income community in each community pair over the pre- (or post-) protest period. Government Responsiveness is the difference between two dummy variables: the first is set to 1 if a busy testing site in the higher-income community is addressed, and the second is set to 1 if a busy testing site in the poorer community is addressed. A site is classified as busy if the waiting time exceeds 30 minutes, and a busy site is considered addressed if additional testing counters or extended operating hours were provided, reducing the waiting time to below 30 minutes. Our independent variables are as follows: *High_Low* is set to 1 if the higher-income community in a community pair has an average property price above RMB 92,300 (USD 13,011.91) and the lower-income community is classified as an urban-village; *High_Middle* is set to 1 if the higher-income community in the pair has an average property price above RMB 92,300 (USD 13,011.91) and the lower-income community has an average property price below this threshold. Additionally, *Middle_Low* is set to 1 if the higher-income community has an average property price below RMB 92,300 (USD 13,011.91) and the lower-income community is an urban-village. We control for the population size of both the higher-income and lower-income communities in each community pair.

4 Specification

For *Site Allocation*, we calculate the testing sites *exclusively accessible* by residents of each community while controlling for the population. We estimate the following equation using ordinary least square. Because our observations are communities adjacent to each other, fixed effects are not necessary.

$$SiteAllocation_p = \beta_0 + \beta_1 RichPoor_p + \beta_2 MiddlePoor_p + \beta_3 + \gamma X + \epsilon_p$$

where X is the matrix of control variables and p is the subscript for community pair.

For *Government Responsiveness*, we use a Heckman selection model to address for the non-randomness in the busy-status: Only busy sites need responsiveness, but busy sites are not distributed randomly.

In the first stage, we select in the sites that will likely be busy. We use the number of malls near a community, the number of hospitals near a community, the minimal distances between a community and a mall/hospital, the average pedestrian flow near a community derived from the pedestrian flow heat map on Gaode Map in the corresponding days, the distance between a community and the location of district governments to predict the probability that a community will have a busy site. In the second stage, we estimate

$$GovernmentResponsiveness_p = \beta_0 + \beta_1 RichPoor_p + \beta_2 MiddlePoor_p + \beta_3 + \gamma X + \epsilon_p$$

where X is the matrix of control variables and p is the subscript for community pair.

5 Empirical Results

Summary statistics are reported in Table 5. Regression results are reported in Table 4. The coefficients for *High_Middle* in columns (1) and (2) are not significant, indicating no difference in allocation between high- and middle-income communities. However, the coefficient

Dependent Variable	(1) Site Allocation (visible) Pre shock	(2) Post shock	(3) Government Responsiveness (invisible) Pre shock	(4) Post shock
High_Middle	0.013	0.005	0.480***	0.129***
	(0.829)	(0.725)	(0.004)	(0.001)
Middle_Low	-0.724***	-0.684***	0.445	-0.160*
	(0.000)	(0.000)	(0.889)	(0.055)
High_Low	-0.488**	-0.380***	0.464**	0.118
	(0.039)	(0.002)	(0.022)	(0.857)
Population_RicherCommunity	3.47e-05***	3.68e-05***	1.68e-05***	2.54e-05***
	(0.000)	(0.000)	(0.001)	(0.003)
Population_PoorCommunity	-2.65e-05***	-2.43e-05***	-2.96e-05***	-9.38e-06***
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.003	-0.008	0.029	0.039
	(0.742)	(0.250)	(0.107)	(0.102)
Observations	25,574	25,574	2,732	2,732
Adjusted R-squared	0.298	0.271	.	.
Log Likelihood	-21705	-21683	-1654	-2533

Table 4. The allocation of visible and invisible benefits

for *Middle_Low* indicates that on average, a high-income community received 0.724 fewer testing sites than a middle-income community pre-protest and 0.684 fewer post-protest. The coefficient for *High_Low* shows that high-income communities received 0.488 fewer testing sites than low-income communities pre-protest and 0.380 fewer post-protest, which is consistent with our hypothesis. These differences are fairly substantial when compared to average testing site density for housing communities of 1.12 per community before the protest and 1.18 after the protest.

In columns (3) and (4), the coefficient for *High_Middle* shows that busyness in testing sites at high-income communities were 48% more likely to be addressed pre-protest but only 12.9% more likely to be addressed post-protest compared to middle-income communities. However, the coefficient for *Middle_Low* indicates that middle-income and low-income communities were equally (un)likely to have their site busyness addressed pre-protest, but middle-income communities were 16% less likely to have their site busyness than the low-income communites addressed post-protest, a difference that is relatively small. Finally, the coefficient for *High_Low* shows that high-income communities were 46% more likely to have their site busyness addressed pre-protest, but equally likely post-protest. These results are consistent with our hypotheses.

6 Mechanism Testing: NLP Analysis of Government Propaganda

If "higher testing site density allocation with lower government responsiveness" toward low-income groups indicates theatrical benefits, while "lower testing site density allocation with higher government responsiveness" toward high-income groups suggests covert benefits, we would expect to see the following pattern in H City's government propaganda: the government would emphasize its efforts toward the low-income group more both before and after the protest, while mentioning the high-income group far less frequently. To test this, we analyze H City government's daily propaganda on COVID-19 control. We include articles related to COVID-19 published by the city-level and 10 district-level official accounts (gōngzhònghào) on WeChat, China's most popular social media platform, between June 1, 2022, and December 8, 2022. As of December 2022, WeChat had 1.313 billion monthly active users, covering nearly the entire population of China. While the exact number of followers on these WeChat accounts is unknown, many residents rely on the official accounts (gōngzhònghào) of their districts as their primary source of policy information. Attention to these accounts is likely to be especially high due to the frequent policy changes.

We exclude articles focused solely on policy changes or the daily risk-level categorizations of communities as they are not government propaganda, resulting in a final sample of 256 articles. These articles focus on three main themes: the government's efforts to alleviate livelihood difficulties during COVID-19, their organization of stay-at-home events to enrich life under quarantine, and case highlights of specific housing communities or individuals during the pandemic. To understand government propaganda related to different income groups, we first adopted publicly accessible LLMs (large language models) to identify all locations mentioned in the articles. GPT-4 is a cutting-edge language model that excels in generating human-like text and understanding complex language patterns in a generic context. It has been trained on a diverse dataset comprising over 175 billion parameters and outperforms many other open-source LLMs such as LM SOTA and SOTA across multiple human languages, including Chinese. After splitting our documents into about 1000 small sub-passages, we analyze each sub-passage through GPT-4 API along with the tuned interactive message. In total, we identify 1,482 locations (with repeat) after excluding the generic terms such as the province, city, or district names, etc. Using coordinates and names, we then matched these locations to the housing communities and, by extension, the income groups in our samples.

Our preliminary results are as follows. Figure 6 shows the frequency of mentions for housing communities from each income group in government articles. In total, specific housing communities were mentioned 345 times: 221 were urban villages (low-income groups), 107 were middle-income communities, and 17 were high-income communities. This highly disproportionate focus on low-income communities supports our hypothesis that efforts directed toward the low-income group—the group favored under the regime's narrative—are subject to the greatest public display. This result is further verified when we extract all words related to location from the articles. "Villages" (Cun) or its synonyms were mentioned most frequently, while "non-village communities" (Xiaoqu or Huayuan) or its synonyms were mentioned far less, despite the fact that there are nearly three times as many non-village communities as village communities. Additionally, we found that H City's government was more likely to mention specific individuals, residents' names/titles, and precise locations when promoting their efforts toward low-income communities, while they preferred to use more vague language when describing their efforts toward high-income communities. (Visualization is forthcoming!)

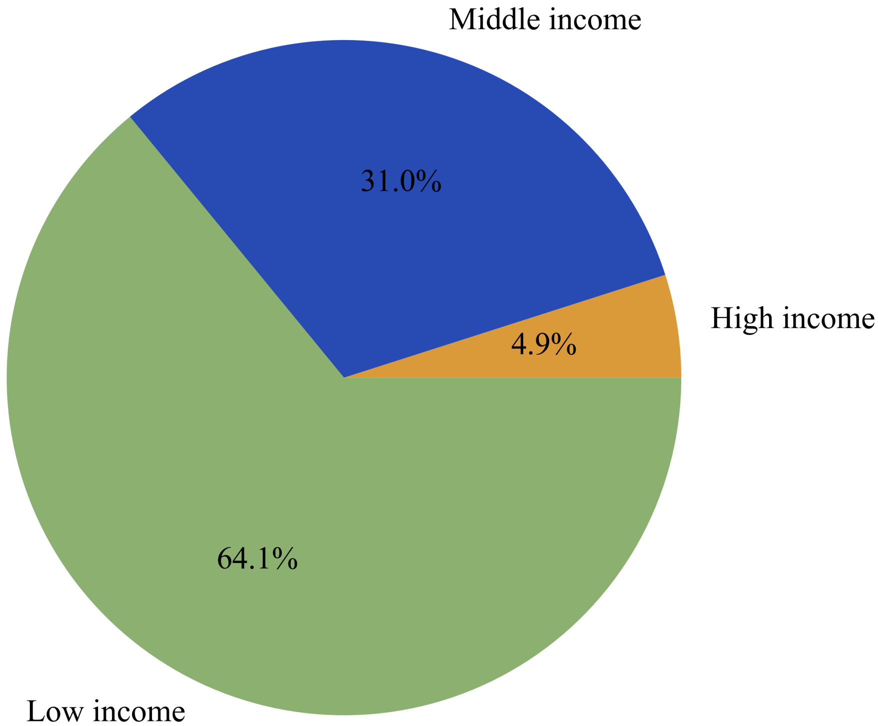

Figure 5. The frequency of mentions for housing communities from each income group in government articles

Figure 6. The frequency of mentions for location-related words

7 Theory

To whom, and when, does an authoritarian government choose to distribute benefits? We highlight "propaganda premium" (i.e., whether a social group is the preferred group under regime propaganda or whether the regime gained power through promising benefits to this group) as a determinant towards whether the government would distribute performative to a group. Our theory proposes that an authoritarian government would prefer to distribute benefits to the groups that has the "propaganda premium". However, this group does not always only receive performative benefits - this depends on the security level enjoyed by the regime. When the regime's survival is fairly secure, propaganda premium raises the distribution priority ranking of a group in narrative but does not raise priority rank in reality. However, when regime survival is under a moderate level of threat, propaganda premium raises the distribution priority ranking of a group in narrative but also in reality, as citing propaganda serves as a quick "band-aid" to regime legitimacy. Hence, when the regime is highly secure, it offers performative benefit to the group with propaganda premium and genuine benefit to the group important for current policy goals; when the regime is less secure, it offers genuine benefit to the group with propaganda premium and genuine benefit to the group important for current policy goals.

8 Contribution

In addition to offering a distribution theory that proposes the distribution of different types of benefits as a strategy, we also contribute to the understanding on selective public goods distribution by proposing an explanation how strong, highly secure authoritarian governments that do not struggle with survive distribute resources. Existing explanations of authoritarian distributioncan be summarized into two camps: while some propose that the regimes offer benefits to supporters of the regime to survive (the winning coalition theory de Mesquita et al. 2003 and the punishment regime theory Magaloni 2006), others propose that they offer the limited resources to those who might be the potential threats of the regime to survive (the co-optation theory Gandhi and Lust-Okar 2009; Svolik 2012) and the squeaky wheel theory Wallace 2014) . We argue that these theories are not sufficient to explain the strategies of strong authoritarian governments that have already established large and stable support base or capability to crush resistance. In addition, we argue that the literature also pays insufficient attention to the upkeep of the provision, which is extremely important for the quality of the provision but is largely invisible to the societal groups other than the recipients. This means that investing in upkeep yields less propaganda values than investing in expanding initial allocation. Thus, understanding the patterns of provision upkeep would help reveal the intentions of the authoritarian regimes in public goods provision

Therefore, by capturing the immediate change in distribution strategy following China's largest popular protest since 1989 and the differential distribution in the "initial provision" and "upkeep" of the public goods, we propose a theory that explains the changes in authoritarian regimes' distribution strategy when its level of security fluctuates. We propose that when regime survival is secure, an aspiring authoritarian regime may desire to realize its policy ambitions. Thus, when the authoritarian regime feels more secure, its main criteria in distribution would be a group's usefulness and irreplaceability in helping to realize the regime's policy goals. It will then try to meet its less-important "survival needs" with the minimum costs by choosing the recipient group that generates the greatest propaganda benefits with performative benefits. However, when it feels less secure, it will distribute genuine benefits to this group to ensure their compliance and support. In our case, although the middle class and the poor are both important for regime survival, the government chooses to offer performative benefits to the poor because they deserves the most care under the Communist narrative when the regime is secure to polish its image with the minimum costs. However, when the protest inflicted a degree of regime insecurity, the government will offer genuine benefits to the the poor. In that respect, our theory and case also highlights the previously overlooked role of ideology in shaping authoritarian governments' distribution decisions.

Beyond the contribution in theory, we empirically contribute to the literature on authoritarian distribution by observing a rare case where the government distribute unevenly to different social groups under a universal provision agenda. Universal provision agendas, like free healthcare, are usually highly uniform across different groups, whereas selective distribution often involve the distribution of different public goods to different groups, a feature that generates incomparability and therefore difficulties in causal identification. Another empirical merit of our case lies in our ability to cleanly identify the change in the attribution of crisis responsibility using the sudden outbreak of the protest under a regime where heavy self-censorship and automatic censorship apparatus makes it difficult the central leader to learn popular opinion before the protest. In addition, our case speaks to the literature on repression-co-optation trade-off. While co-optation and repression are typically presented as a trade-off, we demonstrate that dictators may use public goods distribution as a means to make the implementation of repressive policies more palatable.

This nuanced shift in co-optation strategy highlights the adaptive nature of authoritarian regimes in the face of self-imposed crises and underscores the complex interplay between political survival, legitimacy, and the strategic distribution of resources.

References

Albertus, Michael, Sofia Fenner and Dan Slater. 2018. *Coercive Distribution*. Cambridge University Press.

Blaydes, Lisa. 2011. *Elections and Distributive Politics in Mubarak's Egypt*. Cambridge University Press.

Buzo, Adrian. 2017. *North Korea: The Politics of Regime Survival*. Lynne Rienner Publishers.

Cansunar, Asli. 2022. "Distributional Consequences of Voluntary Provision of Public Goods: Self-Serving Philanthropy in Ottoman Istanbul." *Journal of Politics* 84(2):889–907.

de Mesquita, Bruce Bueno, Alastair Smith, Randolph M. Siverson and James D. Morrow. 2003. *The Logic of Political Survival*. Boston, MA: MIT Press.

De Mesquita, Bruce Bueno, Alastair Smith, Randolph M Siverson and James D Morrow. 2005. *The logic of political survival*. MIT press.

Ding, Iza Yue. 2022. *The Performative State: Public Scrutiny and Environmental Governance in China*. Cornell University Press.

Gandhi, Jennifer and Ellen Lust-Okar. 2009. "Elections Under Authoritarianism." *Annual Review of Political Science* 12(1):403–422.

Geddes, Barbara. 2001. "What Do We Know about Democratization after Twenty Years?" *Opinião Pública* VII(2):224–256.

Greene, Kenneth F. 2010. "The Political Economy of Authoritarian Single-Party Dominance." *Comparative Political Studies* 43(7):807–834.

Holland, Alisha C. 2016. "Forbearance." *The American Political Science Review* 110(2):232–246.

Hong, Ji Yeon, Sunkyoung Park and Hyunjoo Yang. 2023. "In Strongman We Trust: The Political Legacy of the New Village Movement in South Korea." *American Journal of Political Science* 67(4):850–866.

Koch, Natalie. 2018. *The Geopolitics of Spectacle: Space, Synecdoche, and the New Capitals of Asia*. Cornell University Press.

Levitsky, Steven and Lucan A. Way. 2010. *Competitive Authoritarianism: Hybrid Regimes after the Cold War*. Cambridge University Press.

Liu, Dongshu. 2020. The Strategic Balance between the Public and Allies: A Theory of Authoritarian Distribution in China PhD thesis ProQuest Dissertations Publishing.

Lust, Ellen. 2006. "Elections under Authoritarianism: Preliminary Lessons from Jordan." *Democratization* 13(3):456–471.

Magaloni, Beatriz. 2006. *Voting for Autocracy: Hegemonic Party Survival and Its Demise in Mexico.* Cambridge University Press.

Magaloni, Beatriz. 2008. "Credible Power-Sharing and the Longevity of Authoritarian Rule." *Comparative Political Studies* 41(4–5):715–741.

Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed.* Yale University Press.

Svolik, Milan W. 2012. *The Politics of Authoritarian Rule.* Cambridge University Press.

Wallace, Jeremy. 2014. *Cities and Stability: Urbanization, Redistribution, and Regime Survival in China.* Oxford University Press.

Appendix

Variable	Obs	Mean	Std. dev.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
(1) Allocation Pre shock	25,574	-0.004	0.675	1.000								
(2) Allocation Post shock	25,574	0.008	0.662	0.831	1.000							
(3) Response Pre Shock	2,732	0.012	0.437	0.649	0.631	1.000						
(4) Response Post Shock	2,732	0.010	0.298	0.306	0.318	0.354	1.000					
(5) High__Middle	25,574	0.090	0.286	0.026	0.049	0.218	0.006	1.000				
(6) Middle__Low	25,574	0.023	0.151	-0.303	-0.288	0.151	-0.084	-0.048	1.000			
(7) High__Low	25,574	0.001	0.035	-0.090	-0.104	0.395	-0.058	-0.011	-0.005	1.000		
(8) Population__HigherIncomeCommunity	25,574	1750.655	1700.302	0.116	0.122	0.102	0.053	0.085	-0.013	0.000	1.000	
(9) Population__LowerIncomeCommunity	25,574	1781.577	4237.694	-0.311	-0.300	-0.155	-0.111	-0.030	0.629	0.249	0.035	1.000

Table 5. Descriptive statistics

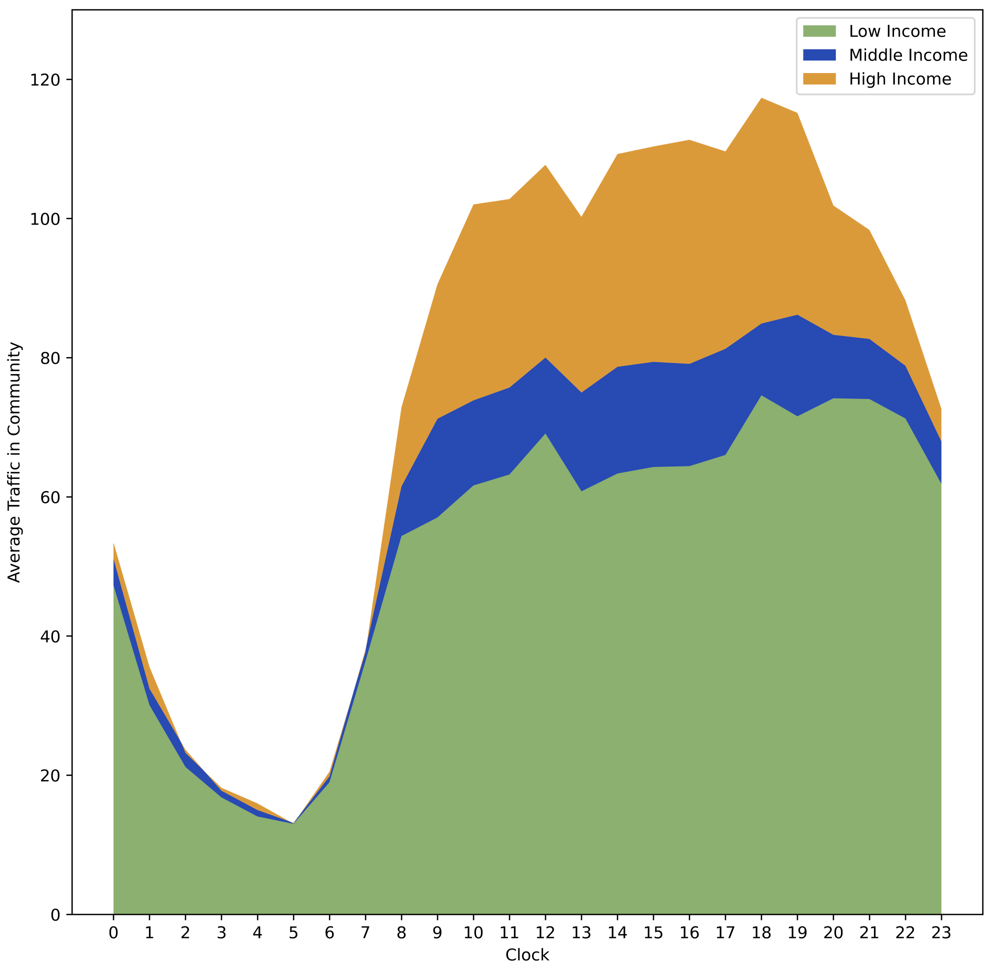

Figure 7. Average footprint traffic appeared in communities of each income group by clock

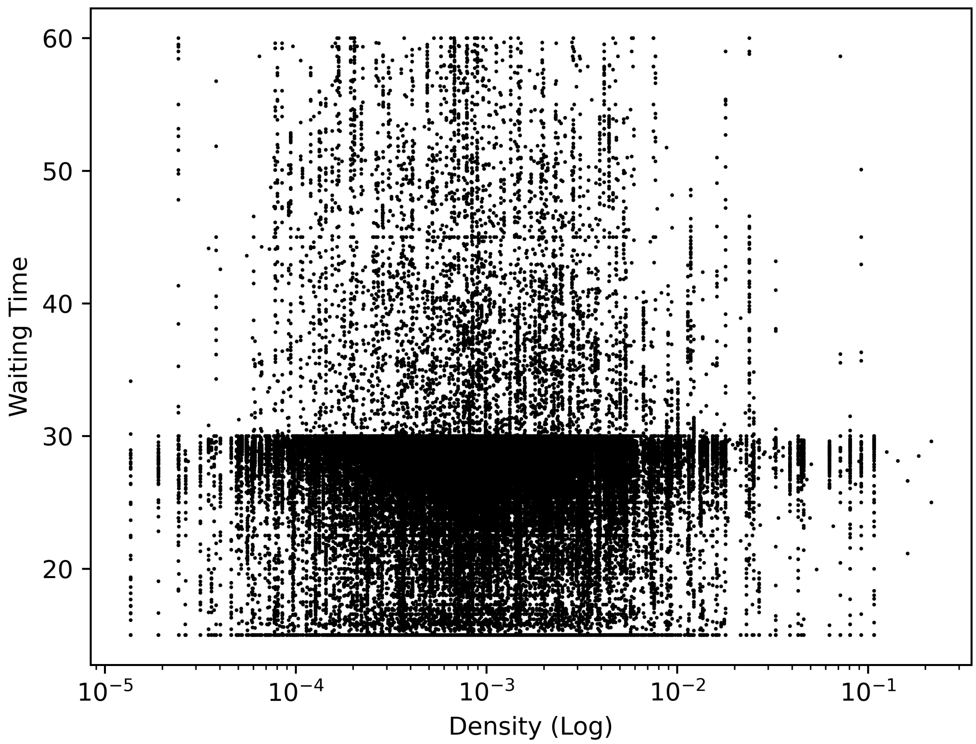

Figure 8. Correlation between testing site density and the average waiting time (each dot represents a housing community-day pair)