Offices or Officers: Sources of Variation in Government Performance*

Prabhat Barnwal† Aaditya Dar ‡ Aprajit Mahajan§ Maitri Punjabi¶

August 2025

Abstract

State performance varies widely across space even within relatively restricted geographies. We document substantial variation in the timely delivery of digitally provided public services within an Indian state and examine the relative contribution of offices and officers to this variation. First, adding officer fixed effects improves model fit by 1.7 percentage points over the 65% explained by office and time fixed effects. Next, using a movers design we show that officers converge almost immediately to their new office's productivity level following a transfer. Finally, decomposing total delivery time into officer time and subordinate staff time, we cannot reject that staff processing time is completely unresponsive to supervising officer changes (arising from transfers), suggesting that office effects work at least in part through (relatively immobile) subordinate staff work routines. The results suggest an important role for place in public service delivery.

*We thank Gulshan Banas, Torsha Chakravorty, Rishita Sankrit, Shivangi Chaturvedi and Kanchana Chiring for excellent research assistance. We are grateful to the Department of Governance Reform, Government of Punjab, India for sharing data. Financial assistance from J-PAL Governance Initiative, International Growth Center and the Digital Identity Research Initiative at the Indian School of Business is gratefully acknowledged.

†Michigan State University, prabhat@msu.edu.

‡Independent researcher, aadityadar@gmail.com.

§Department of Agricultural & Resource Economics, UC Berkeley, aprajit@gmail.com.

¶Michigan State University , mpunjabi@msu.edu

1 Introduction

Government performance varies widely both across countries but also within relatively narrow jurisdictions. Within India, for instance, state performance as measured by health or educational outcomes varies widely (e.g. infant mortality rates vary by a factor of 12 across the best and worst performing states while the analogous gap for functional literacy rates is 1.5).[1] The underlying causes of such sweeping differences are actively debated.[2]

Understanding variation in government performance is crucial for policy. If variation across jurisdictions is driven largely by officers, then policies that directly affect officer incentives and performance could yield large benefits. Conversely, if variation primarily reflects place-based differences—like institutional capacity, political economy factors (including collective action) or resource constraints—such policies by themselves will likely have limited effects, and more innovative approaches may be required. At least two fundamental challenges complicate the study of this question. First, granular data on individual officer performance is typically unavailable. Second, the experimental ideal of randomized officer assignment across locations is practically infeasible.

We address these challenges by combining detailed administrative performance data from Punjab (India) with officer transfer data to separate officer-specific from office-specific components of government performance.[3] Punjab's e-governance platform provides uniquely granular data on bureaucratic processes, digitally tracking citizen service requests through the system. This allows us to observe – for 5.9 million time-stamped citizen applications – which officers and clerks handle citizen service requests and for how long. We combine this data with data on 538 transfers across five departments during a two year period (2022-2023).

We have three main findings. First, officer characteristics contribute minimally to explaining service delivery: adding officer fixed effects improves R^2 by only 1.7 percentage points (pp) beyond the 65% explained by office and period fixed effects.[4] Second, when officers transfer between

[1] For instance, infant mortality (per 1000 live births) in Uttar Pradesh is 50, but it is only 4 in Kerala. Under-five mortality rate shows a comparable gap (Source: National Family Health Survey, 2019–21). Similarly, literacy rates range from 94 percent in Kerala to 62 percent in Bihar, one of the lowest (Source: Census of India, 2011)

[2] Scholars have pointed to variations in political institutions such as state developmental orientation (Kohli, 2004), party system fragmentation (Chhibber and Nooruddin, 2004), bureaucratic reshuffles (Iyer and Mani, 2012), electoral incentives (Ferraz and Finan, 2011; T. Ryan Rogers and Andrew Hall, 2020) and ethnic favoritism (Burgess et al., 2015) as potential contributors. Others have emphasized variation in implementation capacity, documenting that under-resourced and overloaded local bureaucracies undermine effective service delivery Dasgupta and Kapur (2020). See also Ahuja (2019); Vivek (2014) for state-level case studies in India. In part this reflects the difficulty of measuring government functioning cleanly and drawing credible causal links.

[3] In the Indian context, a *transfer* refers to the re-assignment of an officer from one office location to another. The potential reasons for a transfer include the following – periodic rotation norms, administrative and public interest grounds (e.g., filling vacancies or project-specific needs), personnel management (e.g., promotions, training), welfare or hardship grounds (e.g., family or health reasons), and when ordered by legal or regulatory institutions (e.g., higher courts, the Election Commission). See Section 2.1 for more details on transfers.

[4] Office fixed effects thus explain only $1.7/65.0 \approx 2.5\%$ of the variation in service delivery (see Kline, 2024, on the interpretation of such calculations).

offices, they immediately adopt their new office's productivity level with no evidence of gradual adjustment. Third, we separate total delivery time into officer time and subordinate staff time and find that staff processing time—which accounts for the bulk of total processing time—responds minimally to officer changes (following a transfer), suggesting that place effects work at least in part through (relatively immobile) subordinate staff work routines.

We contribute to the growing literature on government performance in four ways. First, we add to the body of work using mover designs to decompose performance variation.[5] Most closely, Fenizia (2022) studies the effect of managers on overall office performance in Italy, while we can examine both office and officer performance separately and do so within a developing country context, providing novel evidence on the office versus officer effects for lower-level bureaucrats. Second, we contribute to the work on public sector management and productivity (see e.g. Rasul and Rogger, 2018) and identify a potential key mechanism for institutional inertia—junior staff workflows remain unresponsive to managerial changes.[6] Third, our results provide an instance of the value of studying the "organizational economics of the state" (Besley et al., 2021), demonstrating that place-based factors play a significant role in routine administrative tasks.[7] In addition, we add to the literature using e-governance data to track government effectiveness over multiple tiers of bureaucracy, place, and time (e.g., Banerjee et al. (2020); Dodge et al. (2025); see Rogger and Schuster (2023) for a broader review).

Our findings have at least two important implications. First, the dominance of place effects suggests increasing the research and policy focus on the institutional environment in which bureaucrats operate. Second, reforms targeting office-level factors—such as technology, management practices, local political economy or resource allocation—may be important complements to officer-level policies in improving service delivery.

The paper proceeds as follows. Section 2 describes the institutional context and data. Section 3 outlines our empirical strategy, Section 4 discusses identification, Section 5 presents the results while Section 6 discusses implications and concludes.

[5] This approach has been applied, most prominently, to understand variation in wages (see Abowd et al., 1999; Card et al., 2013), healthcare utilization (Finkelstein et al., 2016) and teacher effectiveness (Chetty and Hendren, 2018).

[6] This complements the literature on the role of selection, incentives and monitoring in bureaucracies (see e.g. the review in Finan et al., 2015).

[7] A distinct literature examines government performance in episodic "mission mode" activities—time-bound tasks with automatic exit once complete, such as elections, vaccination campaigns, limiting seasonal crop residue burning or census operations (see e.g. Dipoppa and Gulzar, 2024; Kapur, 2020; Khan, 2025)— arguing that the state may perform better on such episodic activities than on routine service delivery. We complement this literature by examining bureaucratic performance in non-mission mode tasks where organizational effectiveness may operate through different mechanisms.

2 Context and Data

2.1 Context

We focus on mid-level bureaucrats in Punjab, a state in Northern India with a population of approximately 28 million (2011 Census). The state has implemented a series of e-governance reforms that digitized a broad set of citizen services and created detailed administrative records enabling precise measurement of service delivery times and responsible officers (Kumar and Singh, 2024). In 2018, the state enacted the Punjab Transparency and Accountability in Delivery of Public Service Act which mandates time-bound delivery of government services with specific stipulated time limits (Government of Punjab, 2018). The act empowers the government to specify the designated officers responsible and time limits for the delivery of specified public services. The Act also mandated electronic service delivery when technically feasible. This legislation creates the legal framework that generates the time-stamped administrative data for our analysis.

Our data covers 952 offices across five major government departments over a two-year period (January 2022–December 2023). These departments—Revenue (e.g. caste and residence certificates), Health (e.g. birth and death registrations), Social Security (e.g. pensions and welfare), Rural Development (e.g. rural area certificates), and Local Government (e.g. special cases of birth and death certificates)—collectively handle a large volume of citizen-government interactions in Punjab through 106 distinct digital services processed via the state's e-governance platform.[8] Table 1a summarizes the number of services and applications processed by officials in each department during our sample period.

The officers in our study occupy mid-level positions with substantial discretion over service delivery. They include Deputy Commissioners (DCs), Sub-Divisional Magistrates (SDMs), Tehsildars, Medical Officers, Block Development Officers, and Executive Officers of Municipal Committees. These positions typically require specialized training and are filled through competitive examinations and/or internal promotions.

Transfers are common and do not appear to be driven primarily by performance considerations.[9] We have provided the list of potential causes for transfers in footnote 3. Iyer and Mani (2012) find that elite Indian civil servants average about 16 months in a posting (well below the recommended tenure guidelines of 2-3 years) with extensive transfers occurring when new political leadership assumes office.[10] In our data, transfers peak after the state elections in February

[8] Best et al. (2023) provide examples of three specific government functions (social security, public procurement, and tax collection) from where existing administrative case data can be repurposed to construct objective measures of bureaucratic performance. In our context, a common e-Governance platform called eSewa has standardized the processing of citizen applications on a large number of services across multiple state government departments.

[9] Recognizing the administrative disruption caused by frequent transfers, the Indian Supreme Court ordered tenure protections for civil servants Supreme Court of India (2013). However, implementation has been limited, with state governments continuing to exercise discretionary transfer powers. Officers rarely exit the public administrative sector prematurely (firing and quitting are rare) and retire around age 60 after 30-35 years of service.

[10] Using data on 2,800 Indian Administrative Service officers over 1980-2004, Iyer and Mani argue that politicians use

2022 (Appendix Figure OA-3) highlighting the importance of political considerations.[11] We also provide some additional evidence below that transfers in our context are unrelated to performance.

2.2 Data

2.2.1 Service (Application) Data

The state government has been rapidly moving public service delivery online.[12] We focus on these online services as they form an increasingly important medium of service delivery and provide detailed time-stamped data for each stage of the process which allows us to measure officer level productivity relatively cleanly. Our primary dataset contains the universe of 5.9 million citizen applications filed between January 2022 and December 2023 across 106 different services in the five departments.

These applications encompass a wide range of citizen services—including residence and health certificates and social security benefits (including caste certificates required for affirmative action purposes) —representing core administrative processes that form a significant portion of routine citizen-government interactions. An analysis of a sample of 10,000 filed citizen grievances reveals that almost half were related to these services and non-delivery by mandated deadlines was the most common complaint.[13]

For each application we observe the filing date, approval date, the approving authority (the officer in charge of the final approval), and processing office. The data tracks processing time for each actor along the chain, allowing us to identify the application's time with the office staff and time with the approving authority. We use the information processing time taken for each citizen request to construct our key performance measure. An alternative measure, the delay rate, uses the information on whether the service was delivered within the stipulated deadline. Appendix C summarizes how we construct our key outcomes.

Service complexity varies significantly both within and across departments. Simple services like income certificate issuance require minimal documentation and single-officer approval, while complex services such as a delayed registration of death involve multiple verification stages, fee collection, and coordination between officials in different departments. Processing deadlines reflect this complexity, ranging from 7 days for straightforward certificate services to 45 days for multi-stage services. The average application in our sample takes 1.60 log days (approximately 5 calendar days) to process, with 18% exceeding mandated deadlines. Substantial variation exists across offices: the 90th percentile office processes applications in half the time of the 10th percentile

transfers as a tool to control officers and find that officers who develop expertise-based reputations experience fewer transfers and lower variability in post importance (relative to those who develop loyalty-based reputations). In our setting we focus primarily on lateral transfers for relatively junior-level officers across equivalent positions.

[11] Removing these transfers from our sample does not affect our results.

[12] Despite this push, some citizen-government interactions require physical presence or complex multi-party coordination and continue to be handled offline.

[13] Based on authors' analysis of data from the Department of Grievance Redressal for 2021.

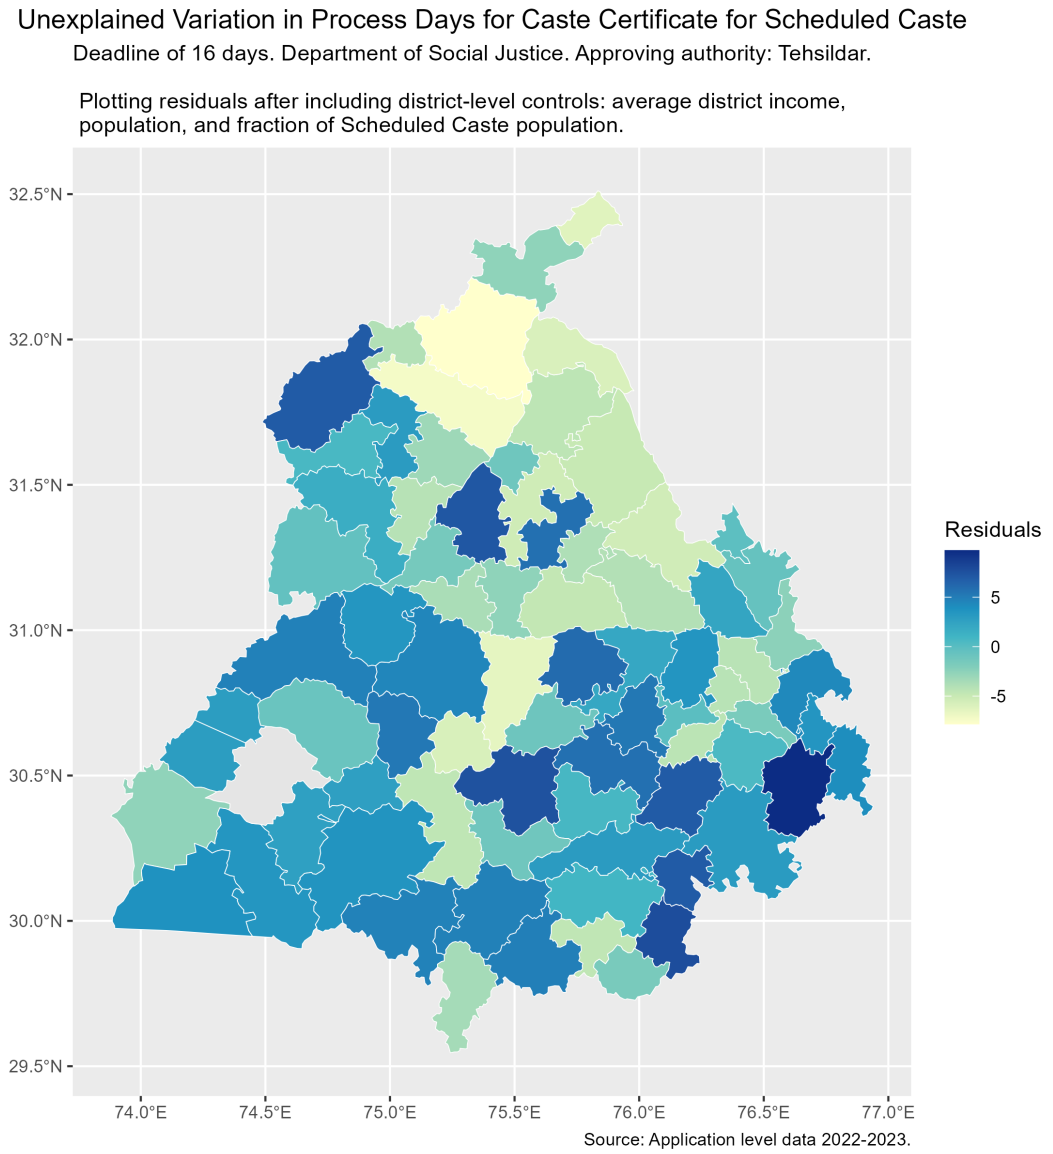
office, highlighting the scope for productivity differences our analysis aims to explain.

To illustrate the service workflow, consider a citizen applying for a Caste Certificate, a document verifying membership in a caste category (Scheduled Caste, Scheduled Tribe or Other Backward Class) required for accessing quotas (affirmative action) in education or employment.[14] The citizen submits an application with supporting documentation by either visiting a service center (called a Sewa Kendra) or online—in both cases the application is digitized and assigned to the relevant SDM (or Tehsildar) office based on the applicant's residence. The office then completes the verification process which may require field investigations by junior officials (patwaris), cross-referencing with government documents or community interviews. Finally, the approving authority—the SDM (or Tehsildar) in this case— approves or rejects the application or in some cases returns it to the applicant with a request for additional documentation.

While the workflow is standardized and uniform across the state (and largely so across the country), there is a large variation in processing time across tehsil offices. Figure 1 shows tehsil-level variation in days taken to process applications for a Scheduled Caste Certificate after controlling for district-level observables. We plot the residuals in processing days after controlling for district population, average income, and fraction of population in the Scheduled Caste category. Even after controlling for these observable district characteristics, substantial variation remains, motivating our investigation into whether this reflects officer quality or place-specific factors.

[14] Punjab's population has the highest fraction of Scheduled Castes among all Indian states (Census 2011).

Figure 1: Geographic variation in public service delivery time



Note: Figure shows a tehsil-level (i.e. sub-district) map of Punjab plotting the residuals from a regression of application processing days for a Caste Certificate (Schedule Caste group) on district-level observables, including average income, population, and fraction of Scheduled Caste population. Map is generated using application-level data for the 2022-2023 period (n=750,538 observations).

2.2.2 Officer and Transfer Data

We merge the services data with comprehensive officer records identifying the 1,455 managers (i.e. approving authorities) in charge of processing applications during this period.[15]

[15]Online services are one component of an officer's duties. Multitasking models (e.g. Holmström and Milgrom, 1991) suggest that differential incentives (between online and offline services) across locations could create place effects through effort substitution between different tasks. Our response is two fold: first, our institutional setting mitigates this

The officer transfer data includes each officer's historical start and end dates in each given posting. We focus on officers that moved across offices during 2022-2023 and impose the following selection criteria: we require officers to serve at least two weeks in both their origin and destination postings to allow sufficient time for productivity assessment. Second, we exclude assignments marked as "additional charge" where officers temporarily oversee multiple positions, to focus on substantive postings. Third, for officers that moved more than once during our sample period, we only consider their first move to simplify our analysis.

This leaves us with a sample of 538 officers who moved across 952 offices during our study period. Table 1b provides a summary of officer characteristics and productivity in our entire sample and the subsample of movers. Transfer data includes exact transfer dates, origin and destination offices, and designation changes. The bulk of transfers (508 out of 538) were within a designation (i.e. lateral moves) and we find similar results when excluding across designation transfers (see Figure OA-9). Average tenure across all postings is approximately 1.9 years, with district-level officers averaging 1.5 years and block or tehsil-level officers averaging 2 years. For movers, average tenure is approximately 1.1 years, with district-level officers averaging 0.9 years and block-level officers averaging 1.2 years. Transfers occur throughout the year but show some seasonal patterns, with larger reshuffles typically occurring mid-year. Our relatively sparse mobility network (sparser than typical firm-worker studies) reflects our institutional context where mid-level officers move much less frequently than private sector workers.

Table 1b compares the 538 movers with the full sample of officers. While movers differ somewhat in observable characteristics—they process slightly fewer applications and have somewhat different delay rates—these differences are relatively modest. We also test whether transfers are correlated with destination productivity and find that officer productivity is weakly negatively correlated with office productivity (see Appendix D for details) providing limited evidence for negative assortative matching. In addition, we show below that there are no "pre-trends" in officer productivity in the weeks leading upto a transfer. The negative correlation between officer and office productivity is not inconsistent with the lack of anticipatory effects. While the negative correlation may reflect (moderate) systematic administrative assignment, the absence of pre-trends suggests the timing of transfers is unrelated to performance shocks. This is consistent with transfer timing reflecting tenure rather than performance.

concern. Punjab's bureaucratic system features uniform wage schedules, standardized performance evaluation criteria, and career advancement processes that do not vary systematically across office locations, making effort substitution an unlikely driver of our observed place effects. Second, to the extent that such substitution exists across locations we interpret it as a *source* of the place-based effect rather than a confound.

Table 1: Summary Statistics

(a) Department Summary

Department	Services	Applications
Revenue	41	2422561
Local Government	29	1425755
Health	51	1412103
Social Security	5	449643
Rural Development	1	95915

(b) Officer Characteristics and Productivity

	Full Sample Mean (SE)	Movers Mean (SE)	Mean difference
Officer tenure			
Posting duration (days)	701.23	418.20	283.04***
	(2.259)	(2.117)	
Productivity measures			
Applications processed per week	76.17	73.40	2.78*
	(0.64)	(0.85)	
Processing days	9.32	9.65	-0.33**
	(0.066)	(0.102)	
Processing days by approving authority	2.60	2.27	0.33***
	(0.020)	(0.029)	
Processing days by office staff	10.83	9.10	1.73***
	(0.068)	(0.097)	
Log(Processing days)	1.60	1.67	-0.07***
	(0.004)	(0.007)	
Delay rate	0.18	0.16	0.013***
	(0.001)	(0.002)	
Log(Delay rate)	-4.27	-4.58	0.32***
	(0.015)	(0.025)	
Number of Officers	1455	538	
Officer-Week Observations	53497	20946	

Note: Table 1a uses application-level data to summarize the number of services and applications processed by each department in our sample. Table 1b summarizes officer posting durations and different productivity measures for the full sample and the subsample of movers using the officer by week level data. Significance levels: * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Our main outcomes are (a) log processing days: the natural logarithm of working days between application filing and approval, and (b) a delay indicator: whether processing time exceeded the government-mandated deadline.[16]

3 Empirical Specification

We use two approaches to decompose total productivity into its office and officer components. We begin with the standard "AKM" model (see e.g. Kline, 2024, for a recent exposition) that decomposes productivity using a set of office and officer fixed effects. Next, we use an event-study (Finkelstein et al., 2016) of changes in productivity around officer transfers. Both sets of results complement each other and we outline each in turn.

3.1 AKM Model

Our first estimating equation is

$$\log(y_{it}) = \tilde{\alpha}_i + \psi_{j(i,t)} + \tilde{\tau}_t + \nu_{it} \tag{1}$$

where $\{\tilde{\alpha_i}\}_{i \in I}$ are officer fixed effects, $\{\tilde{\tau}_t\}_{t \in T}$ are time indicators, $j(i,t) \in J$ is the office location of officer i in period t and $\psi_.$ is a set of office fixed effects.[17]

y_{it} is a measure of productivity associated with officer i in period t. The richness of our data allows us to construct three such measures. First, $y_{it} = y_{j(i,t)}$ is the overall productivity of office $j(i,t)$ when officer i is the supervising officer in the office. This is the total processing time (in days) taken for applications approved in week t. Second, $y_{it} = y^o_{j(i,t)}$ which is processing days taken by officer i in week t and finally $y_{it} = y^s_{j(i,t)}$ is the processing days taken by i's subordinate staff in office $j(i,t)$ in period t.[18] We use the AKM framework to decompose total productivity variation into officer- and office-specific components.

3.2 Event Study

Our second estimating equation uses an event-study of changes in productivity around officer transfers. The key idea is that if officer characteristics drive performance, an officer's productivity should remain unchanged across transfers. Conversely, if office characteristics dominate, officers should adopt their destination office's productivity level upon transfer.

For each moving officer i, we define the "treatment intensity" as the difference in average

[16] For example, if an application for a Scheduled Caste Certificate takes more than the stipulated 16 days to process, it is considered a delayed application.

[17] We will abuse notation and use (I, J, T) to refer both to the sets of officers, offices and time-periods respectively as well as their cardinality.

[18] $y_{j(i,t)} = y^o_{j(i,t)} + y^s_{j(i,t)}$

productivity between destination and origin offices:

$$\delta_i = \log(\bar{y}_{d(i)}) - \log(\bar{y}_{o(i)})$$

where $\bar{y}_{d(i)}$ and $\bar{y}_{o(i)}$ are the mean productivity levels in officer i's destination and origin offices, respectively, calculated using application-level data excluding officer i's own observations. In some specifications we compute mean productivity levels separately for the approving authority and the remaining staff.

Our second estimating equation, estimated only on the sample of transferred officers, is:

$$\log(y_{it}) = \alpha_i + \sum_{r \neq -1} \delta_i \theta_r \mathbb{I}_{\{r = t - t_i\}} + \tau_t + \epsilon_{it} \tag{2}$$

where $\{\alpha_i\}$ are officer fixed effects, $\{\tau_t\}$ are time fixed effects, and θ_r represents the coefficients on the indicators for leads and lags relative to period of the transfer t_i and measures changes in productivity in weeks around the transfer relative to δ_i. The productivity measures are those described above.

In Appendix A we show that Equation (2) can be derived from a stylized utility maximization framework where officers face both effort costs and conformity pressures. The primary coefficient of interest is θ_0, representing the immediate jump in productivity upon transfer. Under purely place-driven productivity variation, we expect $\theta_0 = 1$, indicating officers immediately adopt their new office's productivity level. Conversely, under purely officer-driven variation, we expect $\theta_0 = 0$.

4 Identification

4.1 Identification in the AKM Model

The AKM framework requires a strict exogeneity assumption for consistent estimation via OLS: $\mathbb{E}\left(\nu_{it} | j(i,s) = k\right) = 0 \; \forall i \; \forall (s,t) \in T^2 \; \forall k \in J$. These moment conditions imply that officer transfers are orthogonal to the evolution of productivity.

In addition, estimation also requires understanding the extent of connectedness and officer mobility in the data set.[19] We address connectedness by estimating Equation (1) separately for each of the departments. This approach ensures well-connected mobility networks, as officers typically transfer within rather than across departments. Regarding mobility, our dataset contains 538 officer moves across 952 offices, and the average office pair with moves has one transition during our sample period.

[19] Estimation also usually involves discussions of the transitivity imposed by the linear additive specification of the firm fixed effects (see Kline, 2024, for a discussion).

4.2 Event Study Identification

The strict exogeneity assumption above also provides identification for the event study model, and here we emphasize its implications for transfer timing and the potential for testing for pre-trends.

Transfer Timing Orthogonality: We require that transfer timing is unrelated to destination productivity shocks (recall that we only use the sample of transferred officers). This is plausible in our setting with transfers typically being based on tenure and routine (see also Appendix F) rather than performance. We can allow for some forms of endogenous transfers–for instance movers are allowed to differ from non-movers in their levels of productivity (through officer fixed effects) and we can allow for common trends in productivity around the move (through relative week fixed effects).

Testing the Assumption: The event study framework allows us to test for systematic pre-trends that would indicate endogenous timing (see also Appendix D on matching). Figure 2 plots the estimates of $\{\theta_r\}$ from Equation (2). There is no systematic change in officer productivity in the weeks leading up to the transfer suggesting that transfers are not systematically timed to coincide with productivity shocks. The relatively flat pre-trend pattern is also inconsistent with strategic transfer timing based on performance considerations.

We formally test for pre-trends using a joint F-test of all pre-transfer coefficients (from -12 to -2 in relative weeks from transfer) following the suggestions in Roth (2022). The F-statistic is 1.209 (p=0.275), failing to reject the null of no pre-trends. The mean absolute pre-trend coefficient is 0.0723, which is small relative to the jump at transfer of 0.933.

5 Results

AKM Estimates: The incremental explanatory power of officer fixed effects (beyond office fixed effects) is consistently small across all three productivity measures. Table 2a displays these results. For overall productivity, adding officer indicators to a model with office fixed effects improves explanatory power by 1.7 percentage points (pp) or 2.6% (the R^2 increases from .650 to .667). The corresponding increase is 1.6pp (1.5%) when examining subordinate staff productivity. Finally, and perhaps most surprisingly the increase is only 3.4pp (4.2%) when examining the officer's own productivity. We see similar results when examining variation explained in the log delay rate (Table OA-1).

Event Study: Officer productivity jumps immediately and substantially at the time of transfer, with the magnitude of the jump closely matching the productivity differential between the origin and destination offices (Figure 2). $\hat{\theta}_0$ is 0.933 for (log) processing days indicating that officer productivity jumps to 93.3% of the destination office productivity level immediately following the transfer (the corresponding figure is similar for delay rates). The confidence interval is $[0.77, 1.10]$ so we cannot reject full convergence immediately after the transfer (a similar statement holds for delay rates).

Dynamics: Both event study panels show remarkably similar dynamics. In the pre-transfer period (weeks -15 to -1), the estimates hover around zero indicating that officer productivity remains stable (relative to their expected post-transfer performance). In the post-transfer period (weeks 1 to 15), the coefficients likewise remain strikingly flat at the new elevated level, showing no evidence of gradual adjustment in either direction.[20]

[20] A potential concern is that officers might inherit backlogs from their predecessors, confounding our interpretation. However, the convergence pattern with—$\theta_0 \approx 1$ and the flat transfer dynamics are inconsistent with a backlog story which would predict more gradual adjustment as inherited applications are processed.

Table 2: Analysis of Variance of Log Process Days

(a) Log(Process Days) of Total Time Taken

	(1)	(2)	(3)	(4)
N	53497	53484	53450	53430
Adj. R sq.	0.030	0.650	0.637	0.667
Adj. R sq. 95% CI		[0.644, 0.656]		[0.661, 0.673]
Time FE	Yes	Yes	Yes	Yes
Office FE	No	Yes	No	Yes
Officer FE	No	No	Yes	Yes
P-value				0.000

(b) Log(Process Days) of Time Taken by Approving Authority

	(1)	(2)	(3)	(4)
N	52731	52718	52683	52663
Adj. R sq.	0.016	0.475	0.480	0.509
Adj. R sq. 95% CI		[0.465, 0.484]		[0.499, 0.518]
Time FE	Yes	Yes	Yes	Yes
Office FE	No	Yes	No	Yes
Officer FE	No	No	Yes	Yes
P-value				0.000

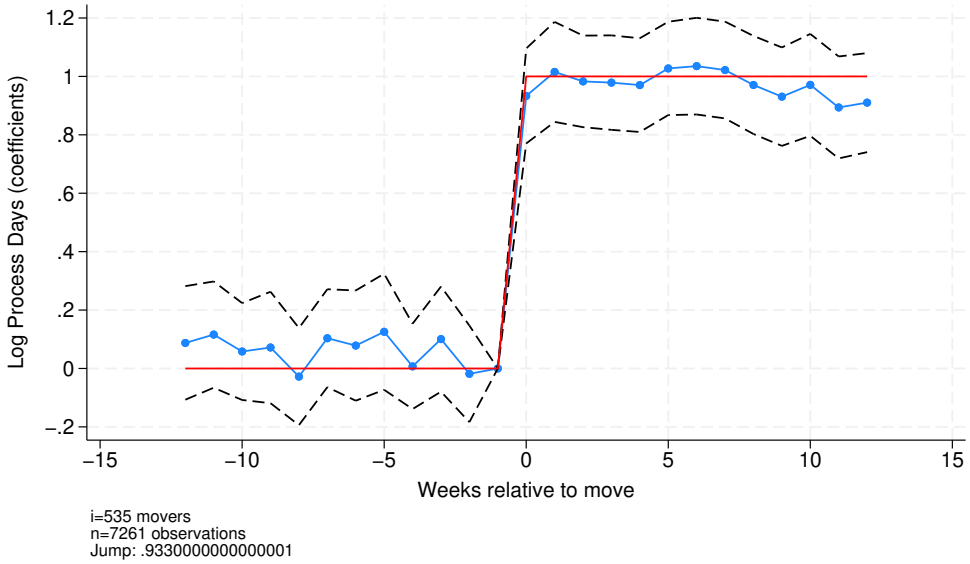
(c) Log(Process Days) of Time Taken by Office Staff

	(1)	(2)	(3)	(4)
N	47821	47804	47767	47746
Adj. R sq.	0.054	0.601	0.587	0.617
Adj. R sq. 95% CI		[0.594, 0.607]		[0.610, 0.624]
Time FE	Yes	Yes	Yes	Yes
Office FE	No	Yes	No	Yes
Officer FE	No	No	Yes	Yes
P-value				0.000

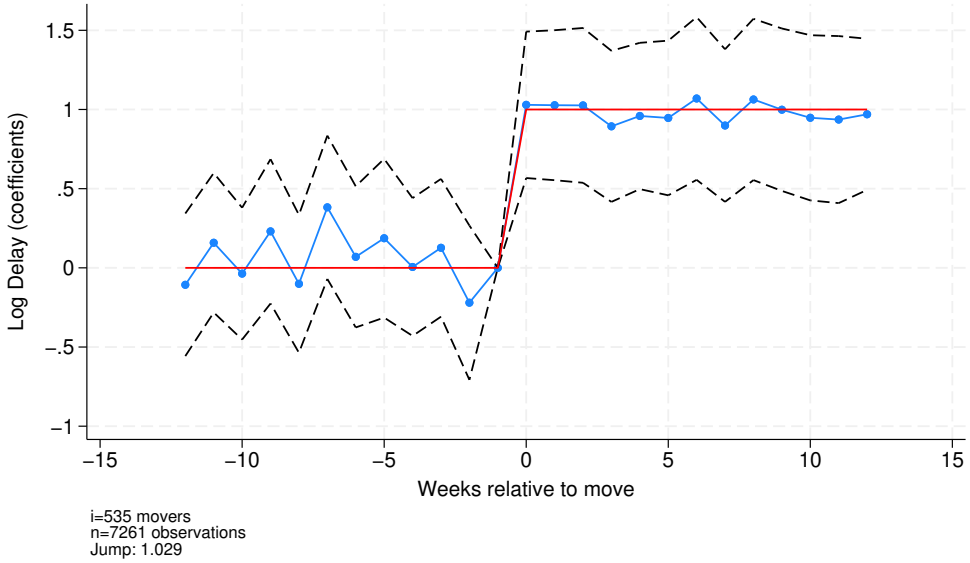
Note: Tables use officer-by-week level data for the entire sample (n=53,497 observations) to analyze the proportion of variation in Log Process Days explained by time, office, and officer fixed effects. We complete this exercise separately for the three outcomes: log of total process days (Table 2a), log of days taken by the approving authority (Table 2b), and log of days taken by office staff (Table 2c). Tables include the p-value of the F-test for whether officer fixed effects are statistically significant in explaining variation in the outcome variable.

Figure 2: Place-based share of overall application processing time

(a) Log Process Days



(b) Log Delay Rate



Note: Figure shows the estimates of the $\theta_{r(i,t)}$ coefficients from our main estimating equation. We normalize the coefficient for the week prior to the transfer to zero. The dashed lines represent the lower and upper bounds of the point-wise 95% confidence intervals. The size of the jump at time zero represents the share of variation attributable to the office. $\hat{\theta}_0$ for the (log) process days is 0.933 while the corresponding coefficient for (log) delay days is 1.029. There are 535 movers in our sample (n=7261 bureaucrat-week observations). Delta is calculated using the difference in (log of) overall time taken to process applications in destination and origin offices, excluding the weeks in which the select approving authority was in each office.

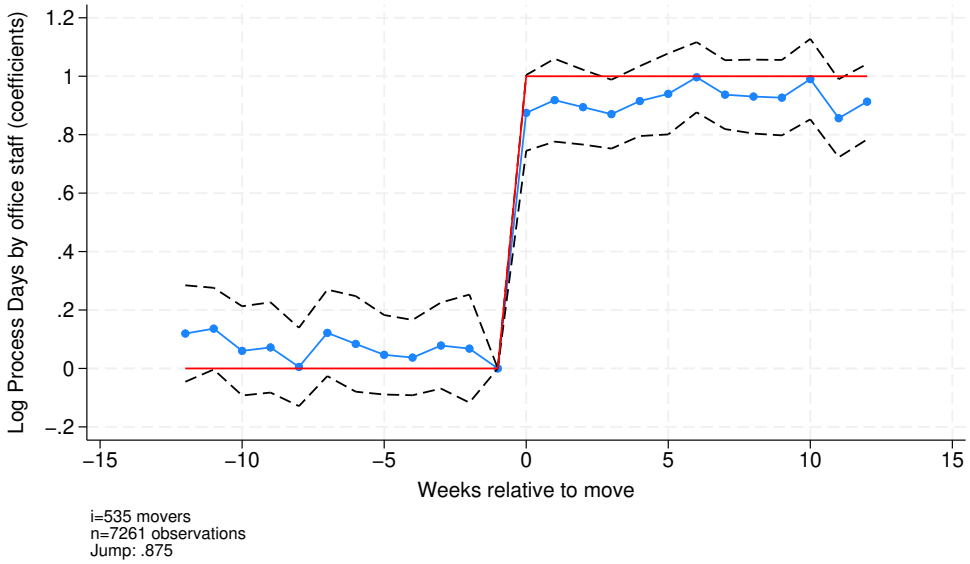
Proximate Mechanisms: We attempt to identify the channels through which place effects operate by disaggregating application processing time into components attributable to different actors in the approval process—time taken by the approving authority (the transferred officer) and time taken by other office staff (the officer's subordinates in the destination office).

We re-estimate Equation (2) using subordinate staff processing time as the left-hand side variable and re-defining δ_i correspondingly using subordinate staff time. Figure 3(a) displays the estimates of $\{\theta_r\}$ from this regression. We estimate $\hat{\theta}_0 = 0.875$ which is statistically indistinguishable from 1 and the remaining coefficients are very similar during the entire post-transfer period. Thus, subordinate staff productivity is unresponsive to changes in the supervising officer (i.e. the transferred officer) both in the short and the long term. Since subordinate staff time comprises the bulk of processing time (Table 1) it is perhaps unsurprising that our estimate of $\hat{\theta}_0$ is similar to that from our main specification. Perhaps most strikingly, when we carry out the same exercise using only the officer's own processing time (Figure 3(b)), $\hat{\theta}_0 = 0.609$ demonstrating that place-based constraints bind even for activities over which officers have presumably the most control (their own time).

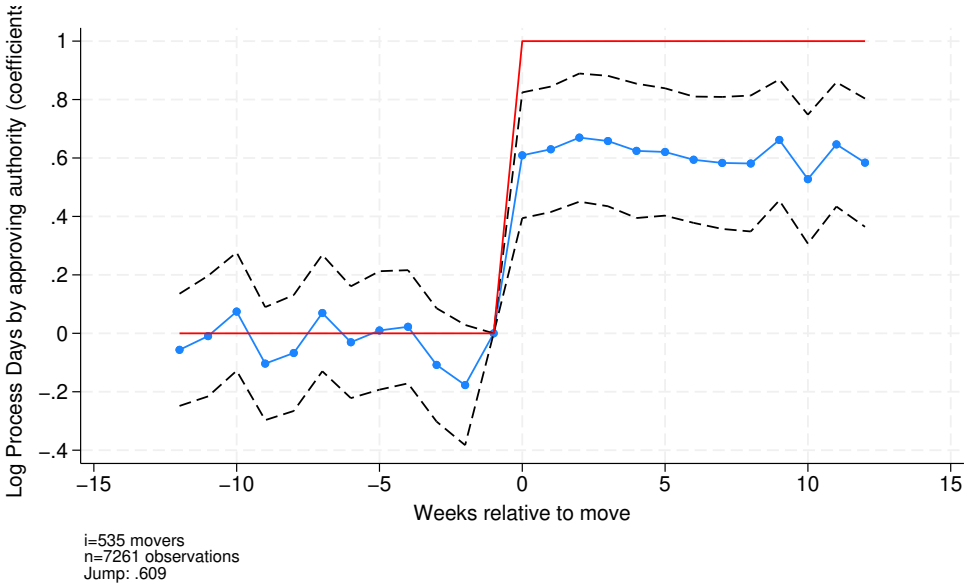
These results suggest that place effects operate in part through subordinate staff maintaining pre-existing productivity levels regardless of the supervising officer. This institutional "stickiness" may reflect several factors: subordinate staff may be more familiar with local procedures and systems, face constant resource constraints or local political pressures that persist regardless of the supervising officer. Subordinate staff are typically less mobile than supervising officers and thus may be more sensitive to local institutional norms as well. However, we lack data to test these specific hypotheses about the underlying mechanisms.

Figure 3: Place-based share of time taken by office staff vs approving authority

(a) Office Staff



(b) Approving Authority



Note: Figure shows the estimates of the $\theta_{r(i,t)}$ coefficients from our main estimating equation, using the log process days taken by office staff as the outcome (and used to calculate delta) (Figure 3(a)) and using the log process days taken by the approving authority as the outcome (and used to calculate delta) (Figure 3(b)). We normalize the coefficient for the week prior to the transfer to zero. The dashed lines represent the lower and upper bounds of the point-wise 95% confidence intervals. The size of the jump at time zero represents the share of variation attributable to the office. $\hat{\theta}_0$ for the (log) process days by office staff is 0.875 while $\hat{\theta}_0$ for the (log) process days by the approving authority is 0.609. There are 535 movers in our sample (n=7261 bureaucrat-week observations). Delta is calculated using the difference in the average of log time taken by office staff (or approving authority) to process applications in destination and origin offices, excluding the weeks in which the select approving authority was in each office.

Office Level Heterogeneity: Finally, we examine whether our results differ according to the type of office by examining productivity dynamics for moves between offices at different points in the productivity distribution. We classify offices as high or low productivity using their average performance over the entire sample period.

Figure 4(a) focuses on officers initially working in high-productivity offices. Before their transfers, all officers exhibit similar productivity levels and trends, regardless of their eventual destination. However, their post-transfer trajectories diverge sharply: those transferred to high-productivity offices (blue line) maintain their strong performance, while those moved to low-productivity offices (red line) immediately experience substantial declines in productivity. This immediate convergence to destination office norms occurs despite these bureaucrats high productivity levels at their origin office. Figure 4(b) demonstrates the symmetric pattern for bureaucrats originating in low-productivity offices. Those transferred to high-productivity offices experience immediate performance improvements, while those transferring to low-productivity environments show little change. These patterns are robust across different productivity thresholds for defining high and low performance offices (see Appendix H).

Department Level Heterogeneity: Table 3 demonstrates that the dominance of office over officer effects is consistent across all departments. The incremental explanatory power of officer fixed effects remains modest though with some variation—ranging from 0.7pp to 4.7pp. The event study estimates of immediate convergence ($\hat{\theta}_0$) are consistently close to unity across most departments, ranging from .765 for Revenue to 1.073 in Local Government.[21] The department level results serve two purposes: first, they reinforce our main results that office effects dominate officer effects in determining productivity by showing that these results hold within department as well. Second, they help address potential concerns about connectedness in the AKM model. By demonstrating consistent results across departments with varying numbers of movers and different mobility patterns, we show that our findings are not an artifact of weak identification due to sparse connections between offices and officers.

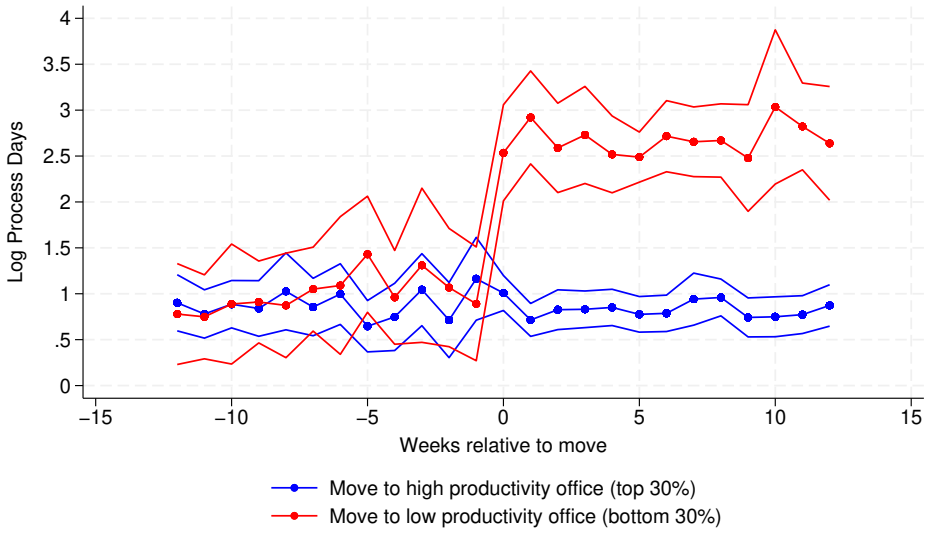
Table 3: Variation explained by offices and officers, by department

Department	Obs	R^2 Time & Office FE	PP ↑ in R^2 from adding Officer FE	Movers	$\hat{\theta}_0$	95% CI
Revenue	15001	74.8%	3.0	241	0.765	[0.538,0.991]
Local Government	7357	48.5%	1.5	90	1.073	[0.627,1.519]
Health	24130	54.6%	0.7	63	0.935	[0.594,1.275]
Social Security	1144	60.9%	4.7	6	0.473	[-0.937,1.882]
Rural Development	5852	33.8%	2.6	135	1.506	[1.068,1.944]
Pooled	53484	65.0%	1.7	535	0.933	[0.771,1.096]

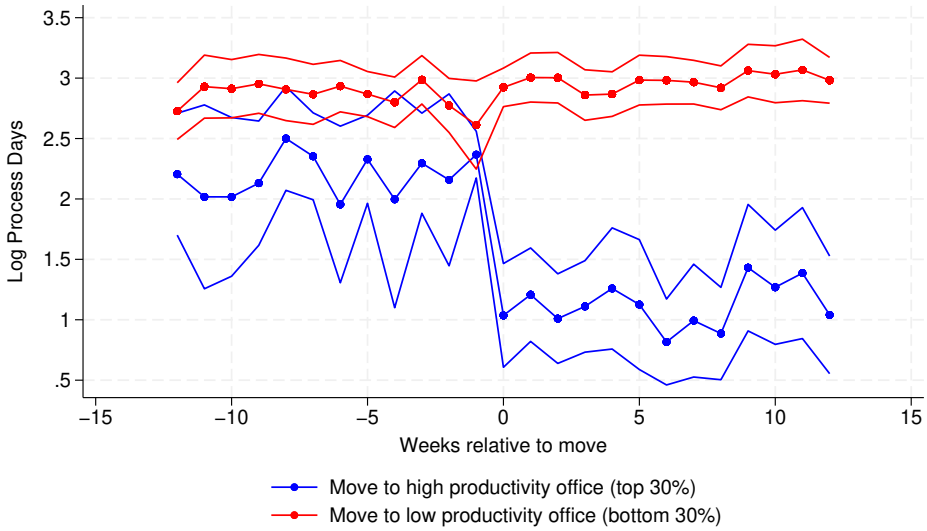
[21] The exception is Social Security, which with only 6 movers exhibits a wide confidence interval including both 0 and 1.

Figure 4: Officer productivity by origin and destination office productivity

(a) Origin: High productivity office



(b) Origin: Low productivity office



Note: These graphs plot weekly average officer productivity over time for officers that move across offices of different productivity levels. We define high productive offices as those that fall in the top 30% in rankings of offices based on $log(\bar{y})$ for the entire sample period and low productive offices as those that fall in the bottom 30%. Figure shows that when an officer moves from a high productive office to another high productive office (4(a) blue), her average productivity remains unchanged. When she moves from a high productive office to a low productive office, however, (4(a) red), her productivity immediately worsens. We show the parallel case for those officers that move from low productivity offices to low/high productive offices. Graphs use manager-by-week level data for the top and bottom 30% of offices (4(a) uses 1,136 observations and 4(a) uses 1,662 observations).

5.1 Discussion

The AKM findings suggest that office characteristics are the primary determinants of productivity—and while officers have slightly more influence over their own processing time than over their subordinates' performance, this effect represents only a modest improvement over the explanatory power of office fixed effects.

It is instructive to contextualize these findings by comparing them to an influential study of the effects of C-suite executives on large firm performance that employed a similar fixed effects design. Bertrand and Schoar (2003) find that adding officer fixed effects increases increases R^2 by over 4 percentage points (pp) on average with particularly large effects for discretionary policies like cost-cutting (37pp). By contrast, our estimates are substantially smaller even for officers' own processing time—an area where officers should have the most control.

The consistency between our two approaches provides complementary evidence for the importance of place effects. The AKM variance decomposition shows that office effects account for approximately 90% of productivity variation, while the event study demonstrates that officers immediately adopt destination office productivity level. These findings reinforce each other: offices dominate in the cross-section (AKM) and we observe immediate convergence in the time series (event study).

This suggests that the bulk of the geographic variation in service delivery stems from office characteristics rather than portable officer characteristics. The rapidity of adjustment is striking—when officers are transferred, they almost immediately adopt the productivity norms of their new institutional environment. This immediacy of adjustment suggests that place-based factors that directly affect productivity such as resources or (informal) administrative procedures or institutional expectations are likely first-order.[22]

6 Conclusion

Understanding the sources of the wide variation in government performance is a first-order policy and research question. Using a comprehensive dataset of digital service requests and officer transfers across Punjab's digital governance system, we employ two complementary approaches to decompose overall government productivity variation into officer and office components.

We find that approximately 90% of the variation in service delivery timing can be attributed to office characteristics, with only 10% attributable to supervising officer effects. When officers transfer between offices, their productivity almost immediately converges to their destination office's average performance level and we cannot reject the null of full convergence. We provide evidence

[22] The immediacy and completeness of convergence provides additional support for our identification assumptions. If transfers were strategically timed to coincide with performance shocks or officer-specific productivity change, we would expect gradual adjustments. Instead, our results suggest that place-based factors are dominant enough to overwhelm officer productivity differences.

that the place effects operate at least in part through subordinate staff workflows, which remain unresponsive to changes in supervising officers. Subordinate staff are less mobile that supervisory staff, and identifying the constraints to their performance is an important next step in uncovering the sources of the strong place-based effects.

Our results also have potentially important implications for public administration reform. The dominance of place effects suggests that policies focused on officers (e.g. incentives) may need to be complemented with policies addressing underlying office-level institutional constraints.[23]

[23] Khan (2025) suggests one approach for government bodies with well defined missions.

References

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High Wage Workers and High Wage Firms," *Econometrica*, 67, 251–333. 2

AHUJA, A. (2019): *Mobilizing the Marginalized: Ethnic Parties without Ethnic Movements*, Modern South Asia, Oxford, New York: Oxford University Press. 1

AKERLOF, G. A. AND R. E. KRANTON (2000): "Economics and Identity," *Quarterly Journal of Economics*, 115, 715–753. OA - 1

BANERJEE, A., E. DUFLO, C. IMBERT, S. MATHEW, AND R. PANDE (2020): "E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in India," *American Economic Journal: Applied Economics*, 12, 39–72. 2

BERNHEIM, B. D. (1994): "A Theory of Conformity," *Journal of Political Economy*, 102, 841–877. OA - 1

BERTRAND, M. AND A. SCHOAR (2003): "Managing with Style: The Effect of Managers on Firm Policies," *The Quarterly Journal of Economics*, 118, 1169–1208. 19

BESLEY, T. J., R. BURGESS, A. KHAN, AND G. XU (2021): "Bureaucracy and Development," Working Paper 29163, National Bureau of Economic Research. 2

BEST, M., A. FENIZIA, AND A. Q. KHAN (2023): "Government analytics using administrative case data," *The Government Analytics Handbook: Leveraging Data to Strengthen Public Administration. World Bank*. 3

BLUME, L. E., W. A. BROCK, S. N. DURLAUF, AND Y. M. IOANNIDES (2011): "Identification of Social Interactions," in *Handbook of Social Economics*, ed. by J. Benhabib, A. Bisin, and M. O. Jackson, North-Holland, vol. 1, 853–964. OA - 1

BRAMOULLÉ, Y., H. DJEBBARI, AND B. FORTIN (2009): "Identification of Peer Effects through Social Networks," *Journal of Econometrics*, 150, 41–55. OA - 1

BURGESS, R., R. JEDWAB, E. MIGUEL, A. MORJARIA, AND G. P. I MIQUEL (2015): "The Value of Democracy: Evidence from Road Building in Kenya," *American Economic Review*, 105, 1817–1851. 1

CARD, D., P. KLINE, AND J. HEINING (2013): "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 128, 967–1015. 2

CHETTY, R. AND N. HENDREN (2018): "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *Quarterly Journal of Economics*, 133, 1107–1162. 2

CHHIBBER, P. AND I. NOORUDDIN (2004): "Do Party Systems Count? The Number of Parties and Government Performance in the Indian States," *Comparative Political Studies*, 37, 152–187. 1

DASGUPTA, A. AND D. KAPUR (2020): "The Political Economy of Bureaucratic Overload: Evidence from Rural Development Officials in India," *American Political Science Review*, 114, 1316–1334. 1

DIPOPPA, G. AND S. GULZAR (2024): "Bureaucrat Incentives Reduce Crop Burning and Child Mortality in South Asia," *Nature*, 634, 1125–1131. 2

DODGE, E., Y. NEGGERS, R. PANDE, AND C. M. T. MOORE (2025): "From Delay to Payday: Easing Bureaucrat Access to Implementation Information Strengthens Social Protection Delivery," Tech. rep., National Bureau of Economic Research. 2

FENIZIA, A. (2022): "MANAGERS AND PRODUCTIVITY IN THE PUBLIC SECTOR," *Econometrica*, 22. 2

FERRAZ, C. AND F. FINAN (2011): "Electoral Accountability and Corruption: Evidence from the Audits of Local Governments," *American Economic Review*, 101, 1274–1311. 1

FINAN, F., B. A. OLKEN, AND R. PANDE (2015): "The Personnel Economics of the State," . 2

FINKELSTEIN, A., M. GENTZKOW, AND H. WILLIAMS (2016): "Sources of Geographic Variation in Health Care: Evidence From PatientMigration," *The Quarterly Journal of Economics*, 131, 1681–1726. 2, 9

GOVERNMENT OF PUNJAB (2018): "The Punjab Transparency and Accountability in Delivery of Public Services Act, 2018," `https://prsindia.org/files/bills_acts/acts_states/punjab/2018/Act16of2018.pdf`, punjab Act No. 16 of 2018, enacted on July 17, 2018. 3

HOLMSTRÖM, B. AND P. MILGROM (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, And Job Design," *Journal of Law, Economics, & Organization*, 7, 24–52. 6

IYER, L. AND A. MANI (2012): "Traveling Agents: Political Change and Bureaucratic Turnover in India," *Review of Economics and Statistics*, 94, 723–739. 1, 3

KAPUR, D. (2020): "Why Does the Indian State Both Fail and Succeed?" *Journal of Economic Perspectives*, 34, 31–54. 2

KHAN, M. Y. (2025): "Mission Motivation and Public Sector Performance: Experimental Evidence from Pakistan," *American Economic Review*, 115, 2343–2375. 2, 20

KLINE, P. M. (2024): "Firm Wage Effects," . 1, 9, 10

KOHLI, A. (2004): *State-Directed Development: Political Power and Industrialization in the Global Periphery*, New York: Cambridge University Press. 1

KUMAR, S. AND M. SINGH (2024): "E-GOVERNANCE IN PUNJAB: AN ANALYSIS OF STRATEGIES AND ITS IMPACT ON CITIZEN SERVICES," *ShodhKosh: Journal of Visual and Performing Arts*, 5, 2941–2950. 3

RASUL, I. AND D. ROGGER (2018): "Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service," *The Economic Journal*, 128, 413–446. 2

ROGGER, D. AND C. SCHUSTER (2023): *The government analytics handbook: Leveraging data to strengthen public administration*, World Bank Publications. 2

ROTH, J. (2022): "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 4, 305–22. 11

SUPREME COURT OF INDIA (2013): "TSR Subramanian & Ors vs Union of India & Ors," Supreme Court of India, civil Appeal No. 6736 of 2007. 3

T. RYAN ROGERS AND ANDREW HALL (2020): "Political Competition and Public Service Delivery: Evidence from Electricity Provision in India," *American Economic Journal: Applied Economics*, 12, 71–97. 1

VIVEK, S. (2014): *Delivering Public Services Effectively: Tamil Nadu and Beyond*, Oxford, New York: Oxford University Press. 1

Offices or Officers: Sources of Variation in Government Performance

Prabhat Barnwal, Aaditya Dar , Aprajit Mahajan, Maitri Punjabi

Appendix – For Online Publication

Contents

A Deriving the Estimating Equation . . . . . **OA - 1**

B Context: Officer Cadres . . . . . **OA - 2**

C Data Description . . . . . **OA - 2**

D Assortative Matching . . . . . **OA - 3**

E Variation in Service Delivery across India . . . . . **OA - 3**

F Transfer Timing and Variation in Delta . . . . . **OA - 5**

G Additional Results . . . . . **OA - 9**
G.1 Office Productivity by Changes in Officer Productivity OA - 9

H Robustness . . . . . **OA - 13**

A Deriving the Estimating Equation

We present a stylized characterization of the officer's choice of productivity level in period t while in office j by assuming utility is of the form:

$$U_{ijt}(y_{ijt}) \equiv w - \eta_{ijt}(y_{ijt}) - C_j(y_{ijt}) \tag{3}$$

where w is the (fixed) wage.

$\eta_{ijt}(\cdot)$ is the (potentially period- and location-specific) disutility of effort required for productivity y (e.g. number of applications processed in the period) and captures the notion that productivity requires costly effort. We parameterize these utility costs as convex: $\eta_{ijt}(y) = (\tilde{\eta}_i + \tilde{\tau}_t + \tilde{\epsilon}_{ijt})y + \frac{y^2}{2}$

$C_{jt}(y)$ captures the office-level component of the cost of productivity level y and is a reduced form approach to capturing the location effect on productivity. It could reflect conformity pressures in office j—a penalty for deviating from the office norm. This could reflect institutional pressures or social norms within offices (e.g. officers could face social reprecussions for being "too slow" or "too fast" relative to the office norm).[24] It could also reflect location specific constraints to productivity (e.g. resources). Since we cannot distinguish between these components of the location effect, we subsumed them into a single measure. We parameterize the penalty function as $C_{jt}(y) = \frac{\tilde{\phi}}{2}\mathbb{E}_i^e\left(y - \bar{y}_{-ij}\right))^2$ where $\tilde{\phi} > 0$ and $\mathbb{E}_i^e(\cdot)$ is officer i's belief about the average productivity level of all other officers in office j (i.e. excluding officer i's productivity) $\bar{y}_{-ij}$. This quadratic penalty specification (and using the group mean as reference) aligns with canonical models in the social norms (Akerlof and Kranton, 2000; Bernheim, 1994) and social networks (Blume et al., 2011; Bramoullé et al., 2009) literatures.

Appendix A1 in Blume et al. (2011) shows that our estimating equation can be derived as the Bayes Nash Equilibrium of a game in which each officer i in office j maximizes eq. (3) given their beliefs (and that in equilibrium, agent beliefs will agree and be correct). Assuming such a BNE with correct beliefs, we can obtain our estimating equation by characterizing the first-order conditions for maximizing $U_{ijt}(y)$ over y. In particular,

$$\begin{aligned} y_{ijt} &= \frac{1}{1+\tilde{\phi}}\left(\tilde{\phi}\bar{y}_{-ij} - \tilde{\eta}_i - \tilde{\tau}_t - \tilde{\epsilon}_{ijt}\right) \\ &= \eta_i + \phi\bar{y}_{-ij} + \tau_t + \epsilon_{ijt} \end{aligned}$$

where we have relabeled terms to ease exposition.

For a mover i who moves from origin office $o(i)$ to destination office $d(i)$, their productivity (for the relevant period) can be written (where we suppress the exclusion of i from the office means):

$$y_{ijt} = \eta_i + \phi\bar{y}_{o(i)} + \tau_t + \epsilon_{ijt} \quad \text{(for } t \text{ before move)} \tag{4}$$

$$y_{ijt} = \eta_i + \phi\bar{y}_{d(i)} + \tau_t + \epsilon_{ijt} \quad \text{(for } t \text{ after move).} \tag{5}$$

Define δ_i as the difference in average productivity between destination and origin offices :

$$\delta_i \equiv \bar{y}_{d(i)} - \bar{y}_{o(i)} \tag{6}$$

[24] Wages are fixed here so the only relevant trade-off is between the cost of effort $\eta(\cdot)$ and the conformity penalty $C(\cdot)$.

We can re-write the equation in a unified form as:

$$y_{it} = \eta_i + \phi \bar{y}_{o(i)} + \mathbb{I}_{\{r(i,t)>0\}} \phi \delta_i + \tau_t + \epsilon_{ijt} \tag{7}$$

where $\mathbb{I}_{\{r(i,t)>0\}}$ is an indicator for being post-move. Replacing the single indicator function with a set of indicators for leads and lags relative to period of the transfer we obtain our main estimation equation:

$$y_{it} = \alpha_i + \theta_{r(i,t)} \delta_i + \tau_t + \epsilon_{ijt} \tag{8}$$

and where α_i combines η_i and $\phi \bar{y}_{o(i)}$. We can augment this model with covariates x_{ijt} by adding them to $\eta_{ijt}(\cdot)$. Identification for the $\theta_{r(i,t)}$ parameters are discussed in the main text.

B Context: Officer Cadres

The officers in our study represent distinct administrative cadres with varying hierarchical positions and transfer patterns.

Deputy Commissioners (DCs) constitute the senior-most district-level administrators, typically drawn from the Indian Administrative Service (IAS) or promoted from Provincial Civil Service ranks. DCs possess statewide transfer liability and can be posted across any of Punjab's 23 districts.

Sub-Divisional Magistrates (SDMs) serve as the Deputy Commissioner's principal agents at the subdivisional level (within a district), managing administration, development activities, law and order, and public grievances across 93 sub-divisions statewide. SDMs possess broad magisterial powers and coordinate multiple departments within their jurisdiction.

Tehsildars and Naib Tehsildars (T/NT) are revenue officers operating at the tehsil and sub-tehsil levels respectively. They focus primarily on revenue administration including land records management, property mutations, revenue collection, and recovery of government dues. Currently, about 80 Tehsildars operate across 101 tehsil offices while 160 Naib Tehsildars serve across 180 sub-tehsil offices.

Medical Officers and **Block Development Officers** represent specialized functional cadres with more limited geographic mobility. Medical Officers typically transfer within regional health jurisdictions, while Block Development Officers usually serve within specific development blocks or related rural development administrative units.

Executive Officers of Municipal Committees are generally confined to urban local body jurisdictions within their respective districts or municipal regions.

Transfers are intended to reflect a trade-off between preventing local entrenchment and maintaining service continuity. Officers rarely exit the administrative sector voluntarily (dismissals and resignations are uncommon), with most serving until statutory retirement. Average tenure across postings is approximately 1.9 years, with district-level officers averaging 1.5 years and block/tehsil-level officers averaging 2 years.

C Data Description

Officer Transfer Data: This dataset contains an observation for each officer's assignment, including the start and end dates of each assignment and the office ID of the associated office. We calculate each officer's posting

duration by counting the days between the start and end dates for each posting (Table 1b).

Productivity Data: There are two primary datasets that we use to calculate our final productivity measures. We use **application-stage level** data, which includes a timestamped log of each actor's work on an application during its processing history. For example, if an application was seen by a total of four actors, this dataset includes the exact time that each actor received and sent out the application, allowing us to calculate time taken by each of the four actors along the processing chain. We calculate the number of days taken by the final approving authority and the (collective) number of days taken by all other office staff. We also calculate the total number of days taken to process the application using the date the application was filed (and therefore received by the first actor in the chain) and the date it was approved by the final approving authority. We reshape this data to be at the application level and merge it with a separate **service-level** dataset that includes each service's processing deadline. This allows us to measure whether the application took longer to process than the stipulated deadline, giving us a delay indicator for each application.

For our main estimation, we collapse the application-level data to the officer-week level, averaging the productivity measures (processing days and delay rate) before generating the log variables.

D Assortative Matching

We estimate officer fixed effects (using the AKM model with approving authority time taken on the LHS) and office fixed effects (using the AKM model with staff time taken on the LHS). Correlating these fixed effects for the different office-officer pairs in our transfer history data, we find that $\rho = -0.1048$, indicating slightly negative assortative matching of officers to offices. This estimate is statistically significant at the 0.1% level.

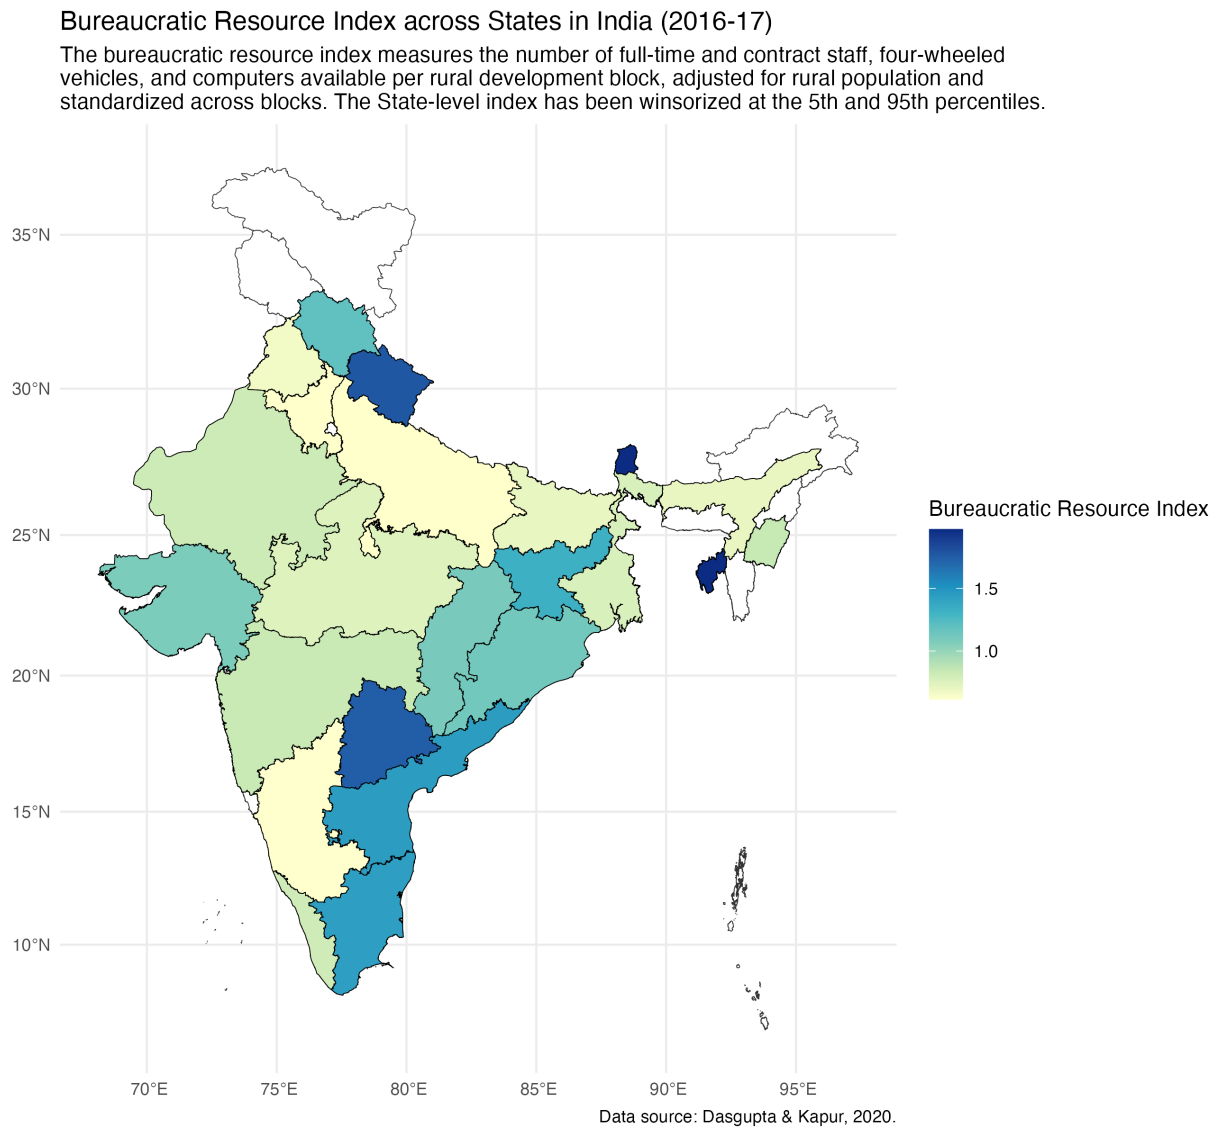
E Variation in Service Delivery across India

Figure OA-1: Service Delivery Deadlines Across Indian States

Services: Issue of...	AP	AS	BR	DL	GJ	KA	KL	MP	MH	OD	PB	RJ	TN	TG	UP	WB
Domicile Certificate	?	?	21	14	?	7	3	7	15	?	15	5	?	6	?	?
BPL Ration Card	?	?	60	45	?	15	0	30	15	?	?	7	?	?	?	?
Caste Certificate	30	15	60	44	?	21	3	30	17	?	7	5	?	30	?	?
Disability Certificate	?	15	21	7	?	7	?	15	7	?	9	21	?	?	?	?
Birth Certificate	6	15	10	7	?	7	?	15	10	?	9	10	?	3	?	?
Death Certificate	6	15	14	7	?	7	?	15	21	?	9	10	?	3	?	?
Income Certificate	7	?	21	14	?	21	6	3	15	?	15	7	15	7	?	?
Driving License	15	?	21	3	?	21	5	15	7	?	?	2	?	?	?	?
Average	**13**	**15**	**29**	**18**	**#DIV/0!**	**13**	**3**	**16**	**13**	**#DIV/0!**	**11**	**8**	**15**	**10**	**#DIV/0!**	**#DIV/0!**

Note: Table summarizes service delivery deadlines for key services across several Indian states.

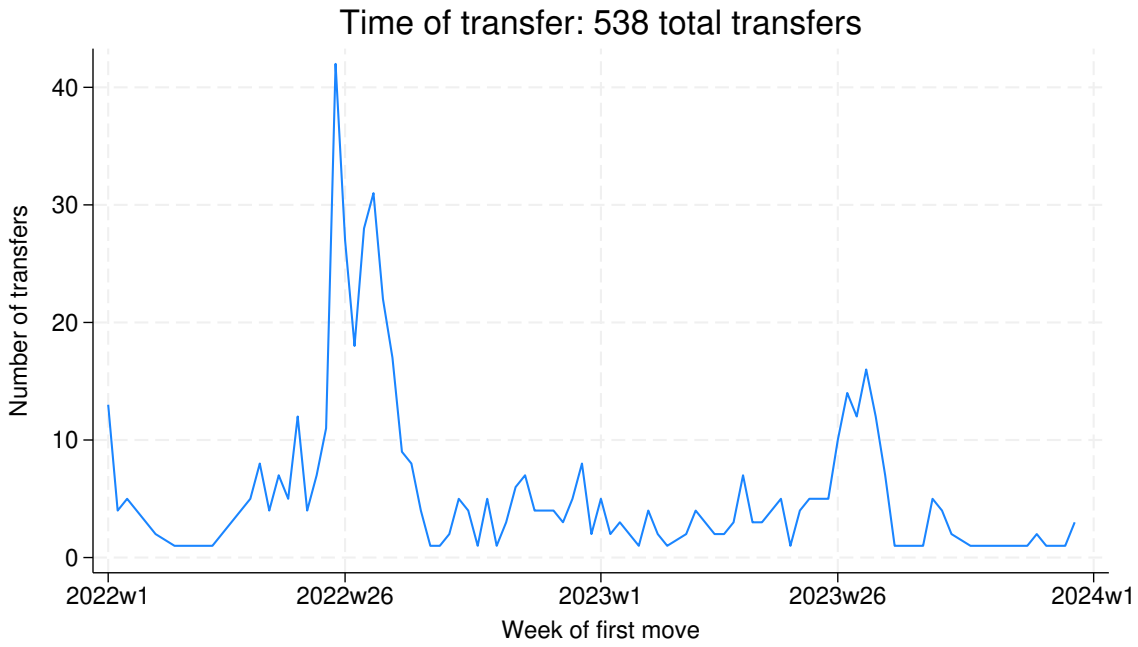
Figure OA-2: Bureaucratic Resource Index Across States



Note: Higher index value represents higher bureaucrat overload. The Bureaucratic Resource Index was constructed by Dasgupta and Kapur (2020) to measure the under-resourcing of civil officers across blocks in India. This index was built using the number of active full-time employees (recruited through contract service procedures), the number of active contract employees, the number of four-wheel vehicles, and computers in each Block development office across 423 blocks across India. They then combined these measures into an index using principal component analysis, which was later normalized (by its in-sample standard deviation across blocks).

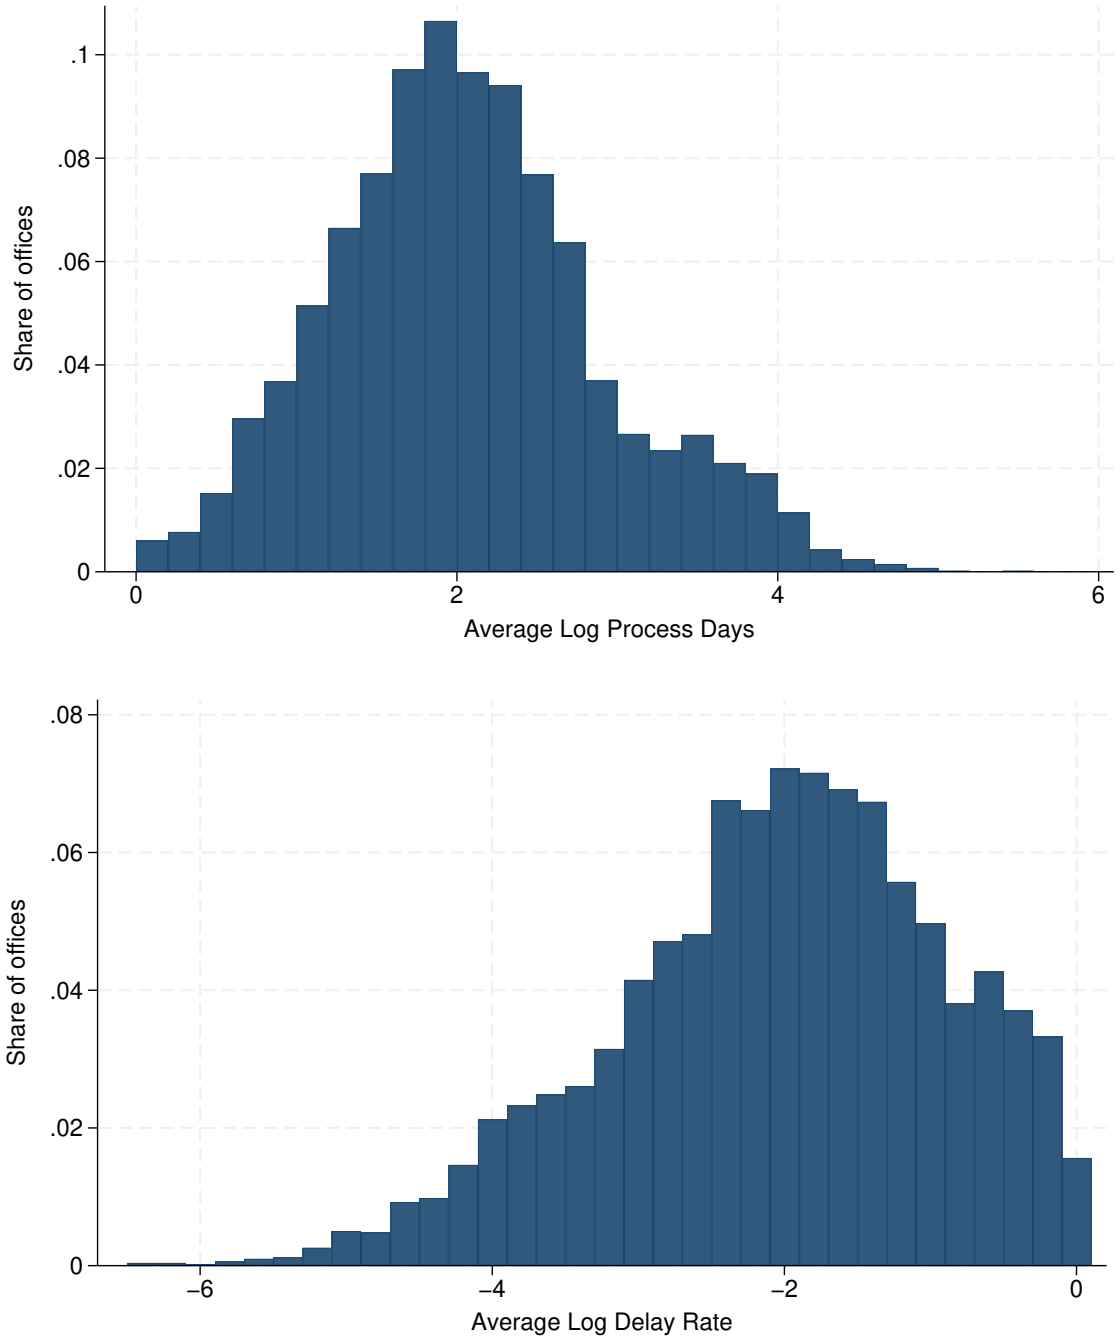
F Transfer Timing and Variation in Delta

Figure OA-3: Transfer Timing



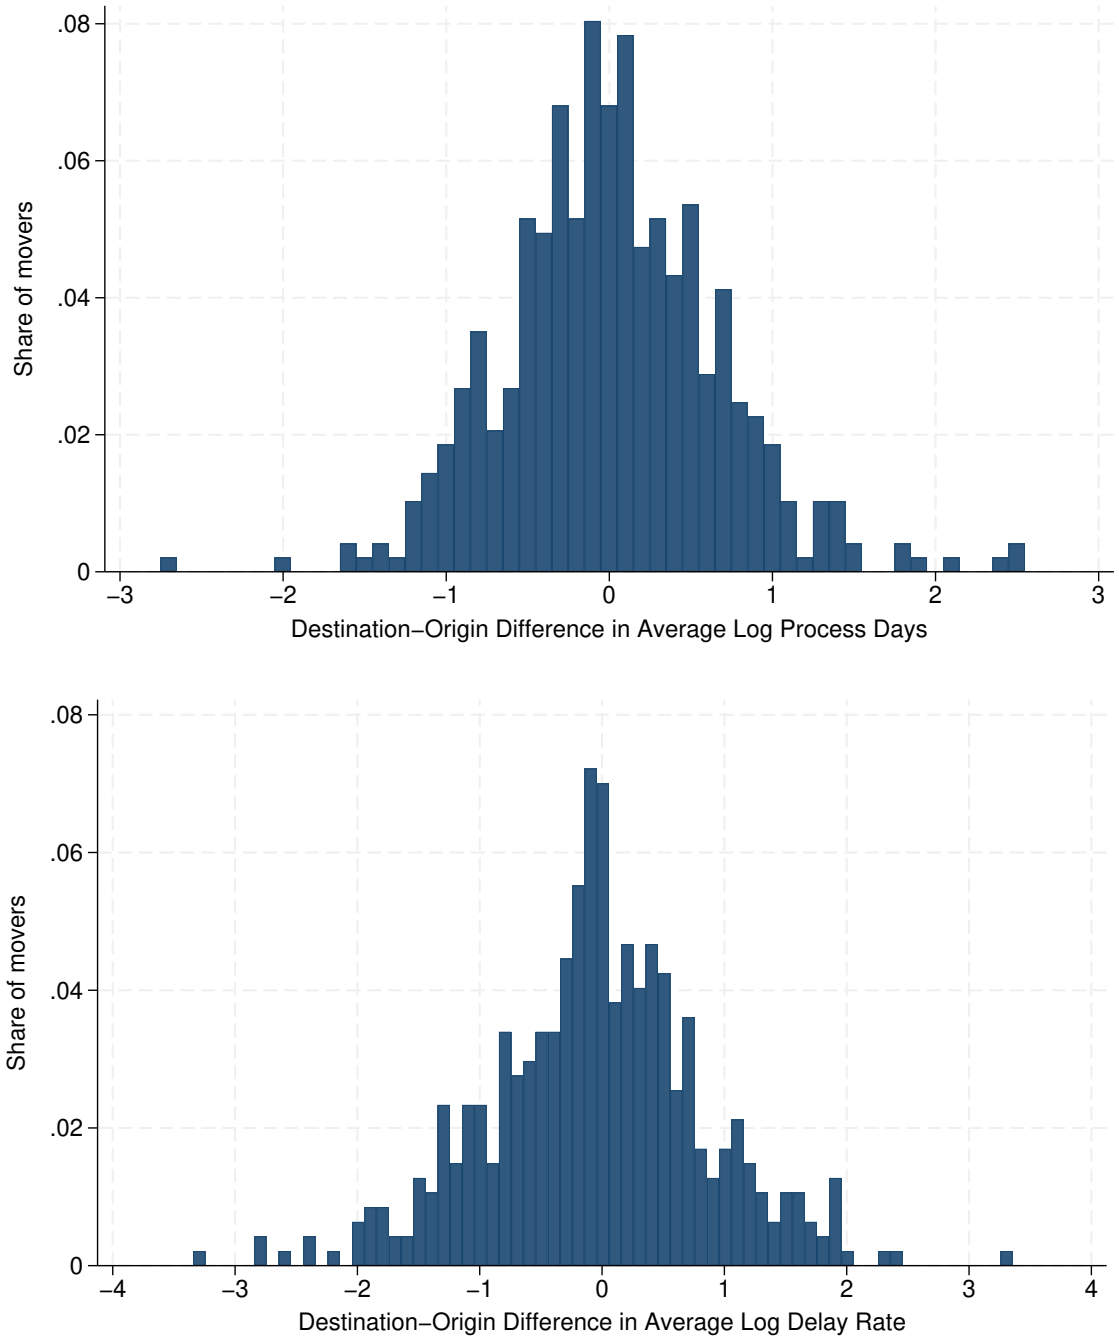
Note: During our sample period, many transfers occurred in the summer of 2022 after a change in the state administration. Excluding this period, there are about 5 transfers per week.

Figure OA-4: Distribution of overall office productivity



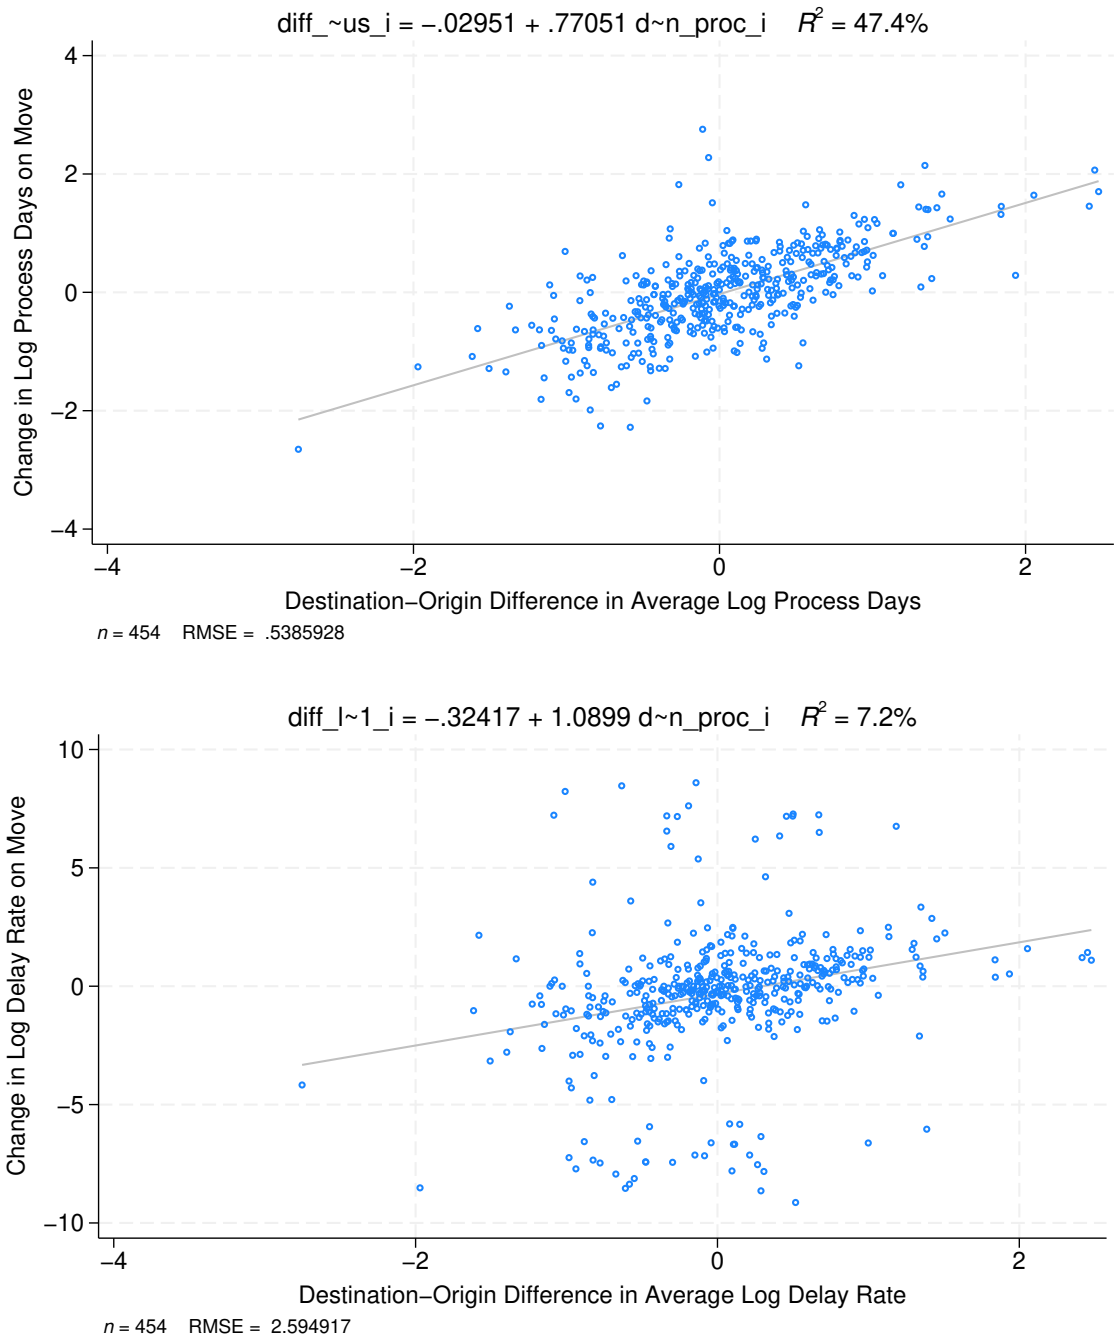
Note: Plots show the distribution of office-level productivity measures used to calculate delta for each mover.

Figure OA-5: Distribution of δ_i



Note: Plots show the distribution of actual delta values that resulted from moves during our sample period.

Figure OA-6: Relationship between δ_i and officer productivity after transfer



Note: Scatter plots show the first stage relationship between δ_i and the change in the officer's average productivity after the move.

G Additional Results

Table OA-1: Analysis of Variance of Log Delay Rate

	(1)	(2)	(3)	(4)
N	53497	53484	53450	53430
Adj. R sq.	0.039	0.424	0.416	0.436
Time FE	Yes	Yes	Yes	Yes
Office FE	No	Yes	No	Yes
Officer FE	No	No	Yes	Yes
P-value				0.000

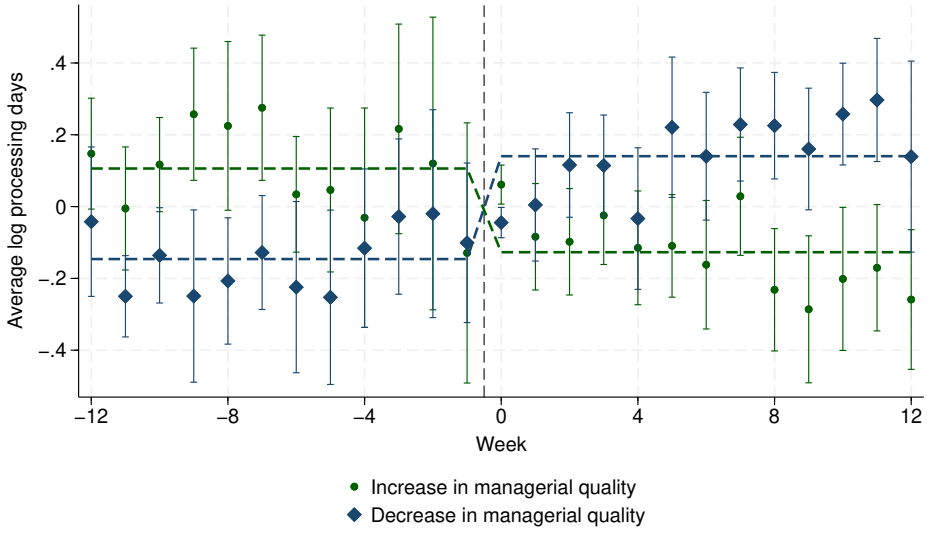
Note: Table uses officer-by-week level data to analyze the proportion of variation in Log Delay Rate explained by time, office, and officer fixed effects. Table includes the p-value of the F-test for whether officer fixed effects are statistically significant in explaining variation in the outcome variable.

G.1 Office Productivity by Changes in Officer Productivity

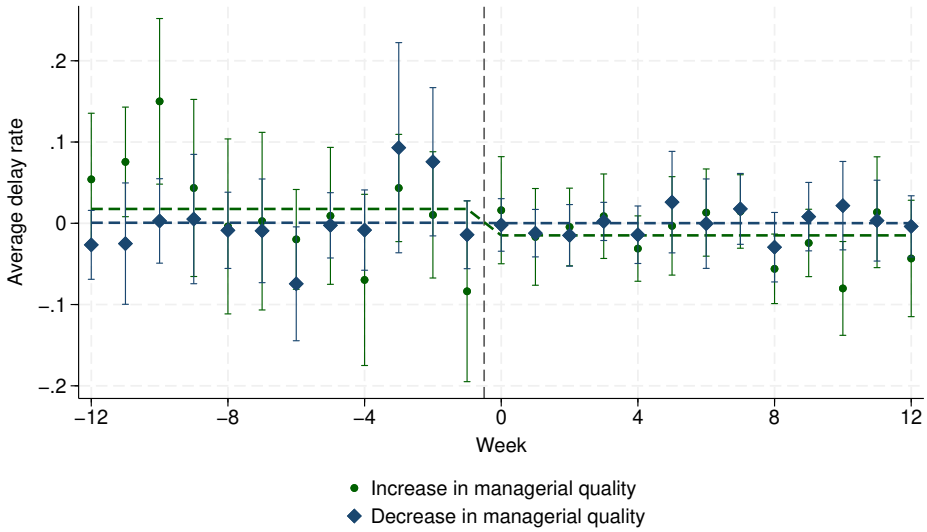
While our primary results suggest that officers do not explain much of the geographic variation in productivity on average, offices may still benefit from being assigned a more productive officer. We investigate how *office-level* productivity is affected by changes in officer quality. We calculate each office's change in officer quality (as measured by officer fixed effects derived from the regression of (log) processing days on officer, week, service and office indicators). To simplify the exposition, we then group these office-level changes in officer quality into quartiles and plot the office-level productivity for the bottom and top quartiles. Figure OA-7 shows changes in overall office-level productivity when the office experiences either an increase (green dots) or decrease (blue diamonds) in officer quality. Offices with an increase in incoming officer quality (i.e. when a higher productivity officer transfers in) show slight improvements in both processing times (a decrease of 0.2 log points) in the weeks following the transfer. Conversely, offices that receive a lower quality incoming officer exhibit marginally worse performance. The results are somewhat noisy, however, and any conclusions are tentative.

Figure OA-7: Office productivity by change in officer quality

(a) Log Process Days



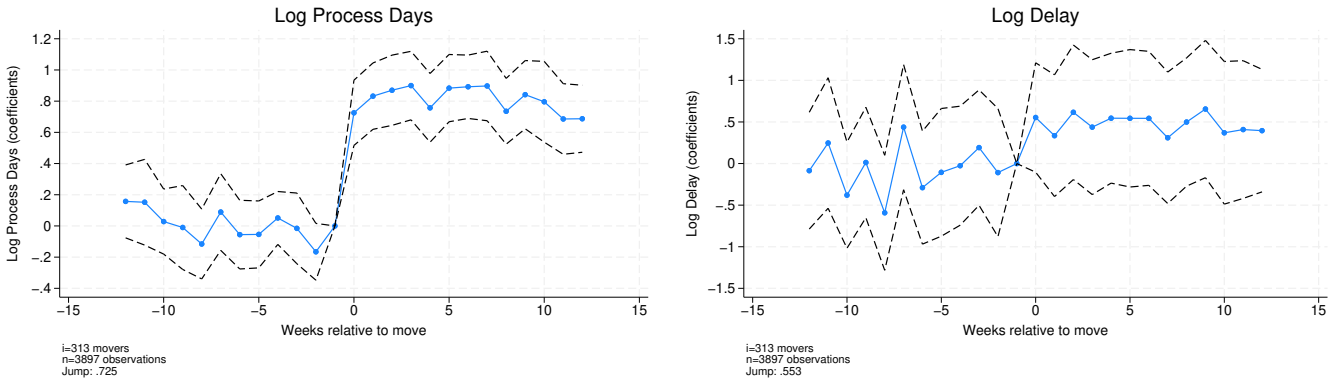
(b) Delay Rate



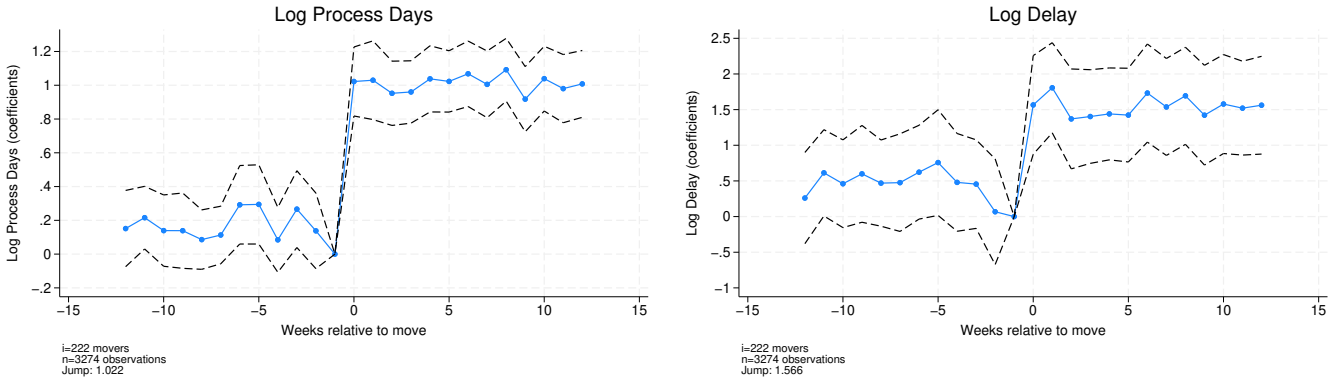
Note: This figure shows what happens to office-level productivity when it experiences either an improvement or a decline in the quality of the officer in charge. We first de-trend productivity by removing time and office fixed effects. We then calculate each office's change in managerial quality using officer fixed effects of the incoming and outgoing officers. We group these office-level changes in managerial quality into quartiles and plot the bottom and top quartiles.

Figure OA-8: Place-based share of overall application processing time, by type of transfer

(a) Do not change districts



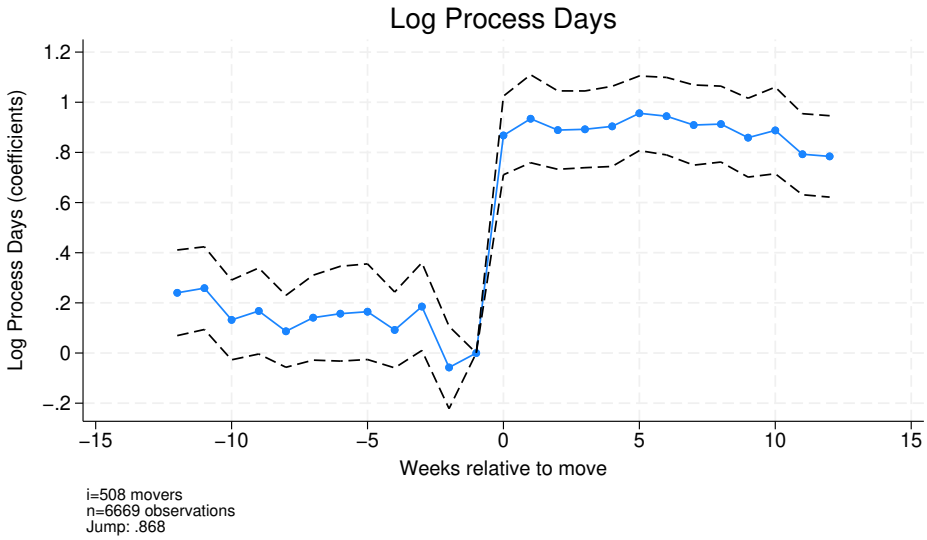
(b) Change districts



Note: Graphs plot the event study results for separate estimation on movers that (a) do not change districts and (b) change districts.

Figure OA-9: Overall place-based share, those that do not change designations

(a) Log Process Days



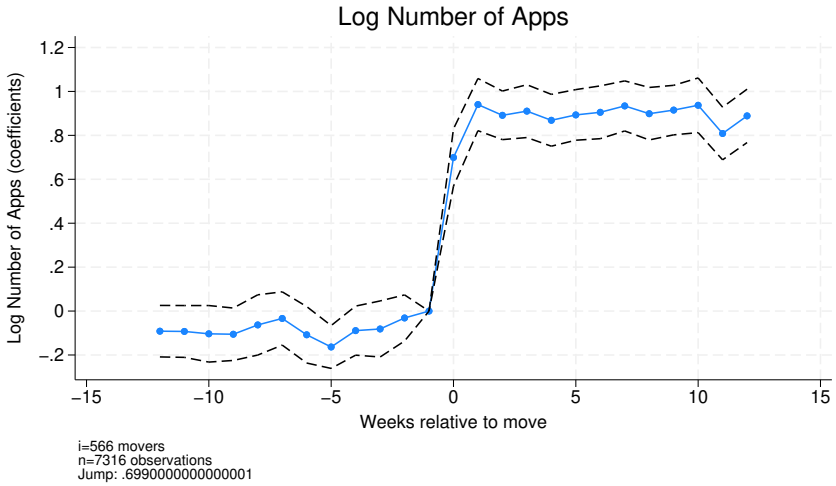
(b) Log Delay Rate



Note: Graphs plot the event study results for the subsample of movers that do not change designations during a transfer.

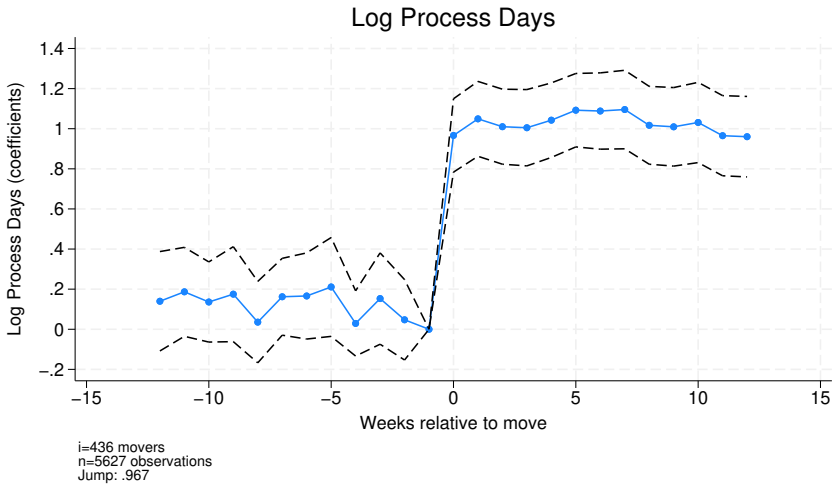
H Robustness

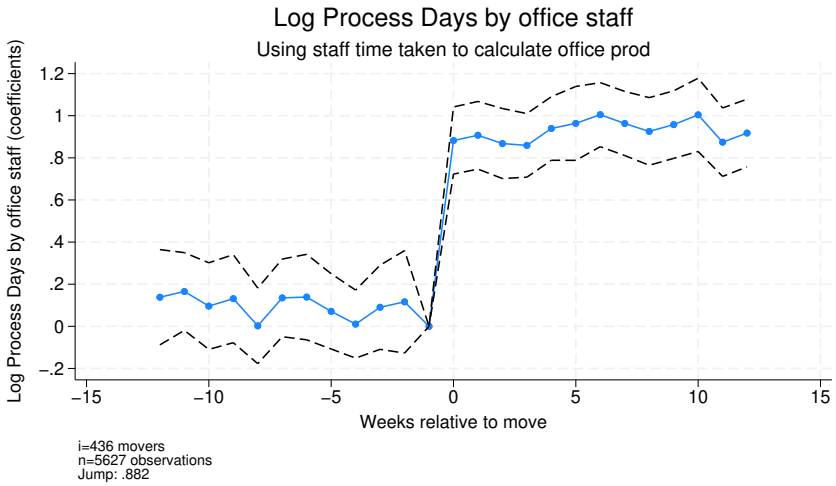
Figure OA-10: Testing officer effect on the number of applications

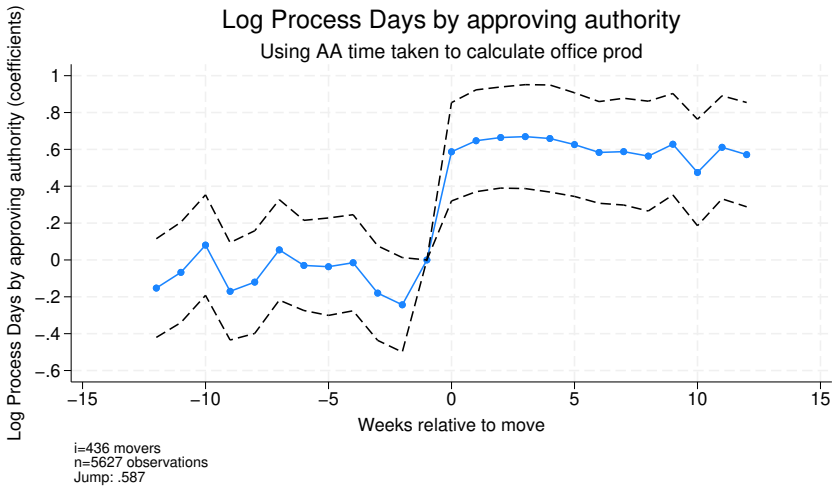


Note: Graph shows the event study results using Log Number of Applications handled by that office as the outcome.

Figure OA-11: Dropping spike of transfers in summer 2022

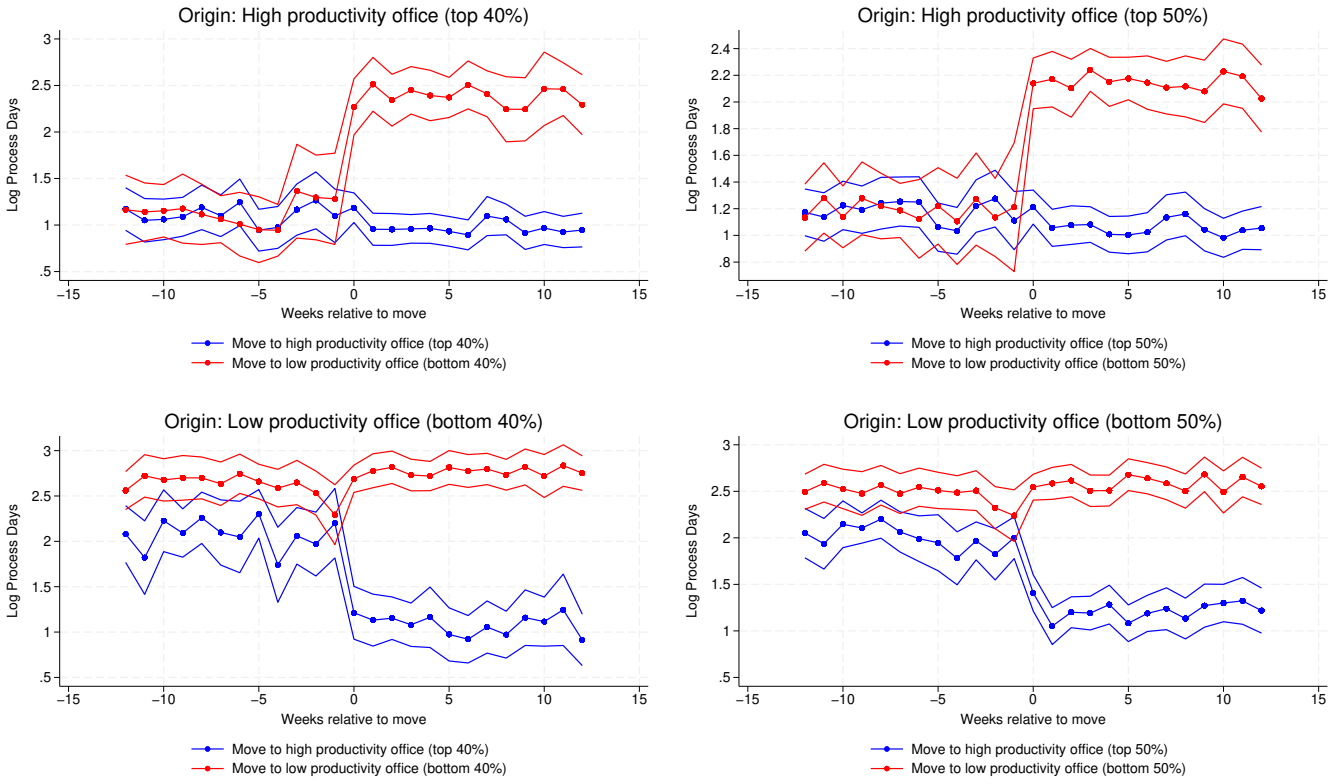






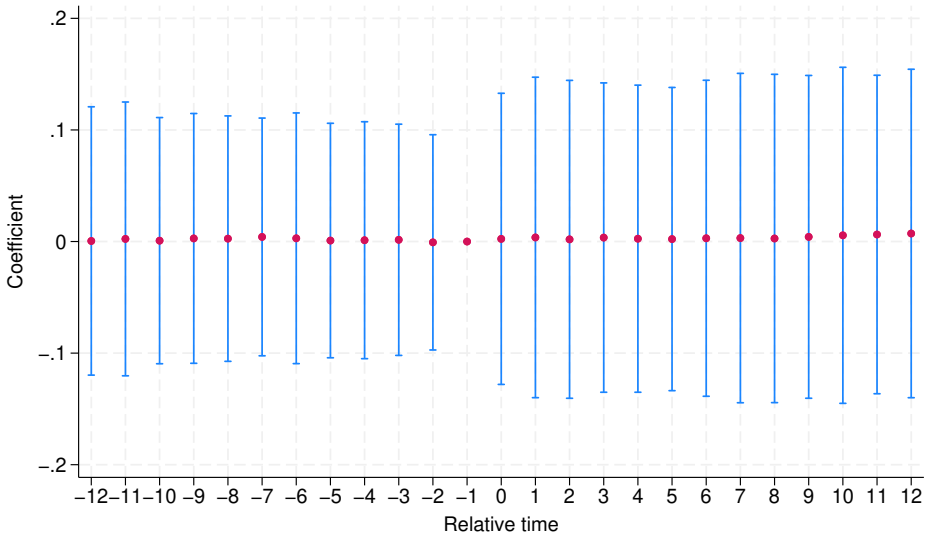
Note: Graphs show the event study results on the sample excluding transfers that occurred during the spikes seen in summer 2022.

Figure OA-12: Changes in officer productivity by origin and destination types



Note: These graphs plot weekly average officer productivity over time for officers that move across offices of different productivity levels. We test two additional definitions of high and low productive offices (top/bottom 40% and top/bottom 50%) using $log(\bar{y})$ for the entire sample period. Figure shows that when an officer moves from a high productive office to another high productive office (blue lines in top graphs), her average productivity remains unchanged. When she moves from a high productive office to a low productive office, however, (red lines in top graphs), her productivity immediately worsens. We show the parallel case for those officers that move from low productivity offices to low/high productive offices in the bottom graphs.

Figure OA-13: Placebo Test



Note: As a placebo test, we examine the event study results from randomly assigning delta values to our sample of movers. We first calculate all possible delta values for within-department moves. Then we randomly select deltas (selected without replacement) from this distribution and assign it to the 538 movers in our existing sample. Running the event study estimation on this placebo delta group, we obtain and store the coefficients. We conduct this exercise 1000 times and calculate average coefficients and 95% confidence intervals. The plot shows the resulting event study plot with coefficients of zero in both pre and post transfer periods.