Assignment Preferences and Performance: Evidence from Pakistan's Civil Service

Adnan Khan Zahra Mansoor Gabriel Tourek

September 9, 2026

Abstract

Aligning tasks to worker preferences is thought to raise productivity, yet little evidence exists beyond the lab, with none from the public sector. We study whether assigning bureaucrats to postings they prefer improves performance using real career outcomes and peer ratings of officers within Pakistan's civil service. We assemble administrative data on bureaucrats between 2008 and 2021 and exploit variation arising from sharp discontinuities in the likelihood of assignment to service groups that differ in how strongly individual officers prefer them. An exogenously stronger match between preferences and assignment increases promotion speed, and officers assigned a group they prefer are rated as more effective by peers, better able to withstand political pressure, and less interested in personal gain. Evidence from early-career performance and surveys with bureaucrats is consistent with intrinsic motivation channels, rather than match to skills, with a limited role for crowd out arising from the prestige attached to generally preferred roles. Our findings shed light on how assigning workers in line with their preferences shapes performance in the public sector, where traditional incentives are constrained.

*Khan (A.Q.Khan@lse.ac.uk) – London School of Economics; Mansoor (zahra.mansoor@bristol.ac.uk) – University of Bristol; Tourek (gtourek@ucdavis.edu) - University of California Davis. We thank Shan Aman-Rana, Alex Eble, Yasir Khan, Diana Moreira, Mazhar Waseem, as well as seminar/conference participants at the University of Pittsburgh and the Behavioral Economics Design Initiative (BEDI) Conference. We are grateful for the collaboration of the Civil Services Academy of Pakistan. For outstanding research assistance, we thank Shabeeh Haider Raza as well as Saniya Abbas, Rylee Ericson, and Abeeha Fatima. For support in data collection, we thank the Centre for Economic Research Pakistan (CERP) Survey Unit. This study has been approved by the School for Policy Studies Research Ethics Committee at Bristol University (IRB approval number SPSREC/22-23/335). It was funded by the International Growth Center (IGC) and the BEDI (University of Pittsburgh).

1 Introduction

An effective bureaucracy is a central component of state capacity. In developing countries, bureaucrats often hold uniquely powerful positions as front-line agents of the state, tasked with administering public services and programs directly to citizens. Emerging evidence suggests public sector workers can explain a large share of variation in government performance (Best et al., 2019; Fenizia, 2022; Bergeron et al., 2021). Identifying determinants of bureaucrats' effectiveness can thus inform strategies for improving the effectiveness of governments.

What we know about these determinants falls along two distinct dimensions. The first examines the role of bureaucrats' *personal attributes* including skill, personality, and dishonesty (Xu et al., 2018; Callen et al., 2018; Hanna and Wang, 2017) as well as public-spiritedness (Ashraf et al., 2020). Improving bureaucrat quality along this margin is judged primarily a question of selection. The second centers on how *organizational features* of bureaucracies, such as rules and incentive structures, shape performance. Recent work, drawn mainly from developing countries, studies monitoring versus autonomy (Duflo et al., 2012; Dal Bó et al., 2020; Rasul and Rogger, 2018; Bandiera et al., 2021) and the effects of monetary and non-monetary incentives on performance (Khan et al., 2016, 2019; Khan, 2023). Less well understood is how organizational features condition the impact of individual traits on performance. In other words, whether the rules of bureaucracy might constrain the ability of public institutions to maximize the gains arising from public workers' productivity-enhancing attributes, or conversely to minimize the losses associated with harmful ones.

This paper explores one dimension in this intersection: the extent to which public institutions match bureaucrats in accordance with their *preferences* for a particular nature of work. Preferences for specific tasks, status, and prestige are considered important determinants of worker productivity (Cassar and Meier, 2018a; Heffetz and Frank, 2008), and the hypothesized between preferences and productivity is an oft-cited motivation for task-matching mechanisms like "rotations" before specialization in fields from management consulting to medicine. Yet there is little direct evidence on the relationship between such preferences and performance beyond the lab (Kamei and Markussen, 2023; Freyer et al., 2025) and none from the public sector, where any gains from matching workers to their preferences could be particularly valuable given constraints on other tools for raising performance. Selection in bureaucracies is often

limited in scope by hiring processes that are weakly tied to expected performance (e.g., nation-wide civil service exams). Rigid wage structures forestall use of high-powered financial incentives. At the same time, selection and assignment processes in civil service often prize "generalists" over specialists, a legacy of British colonial systems that sought to recruit agents "flexible to any task" (Kennedy, 1987). We examine whether the preferences of bureaucrats — for particular postings involving distinct natures of work — shape their performance over the short and long runs.[1]

We study this question in the context of Pakistan's civil service, where an allocation mechanism for postings creates exogenous variation in whether civil servants are assigned to positions they prefer more. We exploit sharp thresholds in the likelihood of assignment to service groups within Pakistan's elite Central Superior Services (CSS) to assess how increasing the strength of the match between bureaucrats' preferred and assigned postings affects their performance.

Ex ante, it is not clear that matching bureaucrats more closely to preferred postings would improve performance. On one hand, preferences over postings — in our setting, specified before exposure to service — may be unrelated to long-run performance if they primarily reflect information constraints, potentially weakened by exposure (Conlon and Patel, 2023), or if preferences are adaptable to environments (Di Falco and Vieider, 2023). On the other, preferences may shape intrinsic motivations for work and, if they persist, could have lasting effects on motivation and productivity (Cassar and Meier, 2018a; Heffetz and Frank, 2008), while, pushing in the opposite direction, the prestige or status attached to preferred positions may crowd out such incentives, a possibility we discuss theoretically in Section 3. Alternatively, preferences can simply reflect suitability for particular roles — e.g., coherence in skills, capabilities, or desired career objectives between preferred and assigned positions — and matching workers to their preferences may raise effectiveness relative to assignments where suitability is misaligned (Burks et al., 2023).

Using an instrumental variables design, we leverage non-linearities in the probability

[1]It is worth noting that how bureaucracies treat candidate preferences when assigning new entrants to roles varies. Pakistan, along with India and Bangladesh, allocate posts in accordance with performance on national civil service exams and preferences, with limited ability to switch service groups post-assignment. Other countries like the United Kingdom and Singapore rely on central matching that does not account for candidate preferences but provide more substantial opportunities for movement across posting types after entry. France transitioned from the South Asian model to the UK/Singapore model in 2022. Japan decentralizes the process, with candidates selecting which ministries they choose to interview with, while lifetime appointments once hired restrict mobility across ministries.

of assignment to postings to estimate the impacts of assignment to postings officers better prefer on measures of career advancement and performance. Drawing on administrative data on bureaucrats' career progression, we first consider promotion and transfers to new roles. These outcomes, which we collect five to ten years after service entry, represent the most consequential events for career advancement and can, at least in part, provide a view into performance. We also consider the negative events of punishment and exit from service. We then use ratings of officers from colleagues in service to assess performance more holistically through peer-rated measures of effectiveness, vulnerability to political pressure, and interest in personal gain.

We find that assignment to a less preferred posting reduces the likelihood of promotion but has no average effects on transfer to new roles. In our preferred specification, which controls for the level of promotion (or transfer) within a service group, assigning a candidate from their fourth to first preferred SG — the average assigned preference-rank being 4.49 — translates into a 7% reduction in the likelihood of promotion, relative to other candidates.

This pattern is consistent with a closer correspondence between a bureaucrat's posting and assignment increasing the *speed* of promotion rather than necessarily the likelihood, at least in the time period our data covers. Promotion in Pakistan's civil service, as in other settings, permits officers to qualify for a move up in rank after a specific number of years in service. Bureaucrats assigned their more preferred postings are more likely to be promoted after five and ten years, though in the intervening years this difference dissipates as their counterparts "catch up" in moves along the hierarchy. Our data captures officers mostly at the beginning of their careers in service:the typical length of service is 25 years and we collect administrative outcomes for up to the first 10 years. However, we find evidence that the strength of the preference-to-posting match can affect the likelihood of reaching the upper rungs of service, to which promotion is not guaranteed given the dearth of positions at the top. Assigning a bureaucrat to a marginally less preferred posting reduces the likelihood he or she is promoted to a Secretary-level position (the highest possible).

We interpret the higher likelihood of advancement that attends assignment to a more preferred posting as capturing higher performance. In other words, speedier promotion corresponds to a "fast tracking" of high performers. Aman-Rana (2025) documents this practice within the CSS's Pakistan Administrative Service, finding that while careers typically follow a predictable trajectory determined by seniority, mandatory train-

ing, and performance evaluations, there also exists a "parallel track of discretionary, fast-track promotions overseen by senior bureaucrats, designed to circumvent the constraints of rigid rules," and that these promotions are largely meritocratic.

To corroborate this interpretation, we collected peer ratings of CSS officers through a survey exercise that asked bureaucrats to evaluate their colleagues. These ratings provide a more well-rounded and assessments of performance, which we evaluate along several dimensions.[2] We find that those marginally assigned to a service group they more strongly prefer are rated as more effective, better able to withstand political pressure, and less interested in personal gain.

Why does assignment to more preferred services raise performance? Surveys with new entrants to service suggest roles for intrinsic motivation, suitability match, and prestige in determining how aspiring officers rank their service group preferences. Job satisfaction ratings among seasoned officers reveal a negative relationship between motivation and assigned preference rank. To help adjudicate among mechanisms, we examine direct performance measures collected during trainings conducted immediately after assignment to service. These short-run measures, captured befire direct exposure to the work within their service, are more likely to capture a combination of ability and motivation-related factors, whereas longer-run measures of performance in service may also reflect the nature of hierarchies within services and personal connections. Assignment to a more preferred posting has positive effects on performance during training. Assigning a candidate from their fourth to first preferred service translates into a 13.17% improvement in a candidate's overall training performance rating. Though short-run in nature, the early presence of these effects suggests scope for persistence throughout careers and aligns with impacts we estimate on longer-run career advancement and peer assessments.

In short, our findings point to a role for preference-match in mediating performance in a representative developing country bureaucracy. The nature of the variation we exploit — thresholds determining marginal match to more preferred postings — limits our ability to evaluate counterfactual assignments that could shed light on the aggregate performance value of prioritizing matching to preferences relative to other assignment schemes. However, the effects we document help to quantify what is (lost or) gained through preference (mis)alignment in the public sector. Acknowledging the existence

[2]The measures we use mirror those used by Bertrand et al. (2020), which source ratings among Indian civil servants, as well as other stakeholders, to assess the effects of career incentives on performance.

of such forces can also inform how complementary tools for motivating performance (e.g., monitoring, rewards) are deployed.

We contribute to two main strands of literature. First, we add to the literature that examines the relationship between preferences over tasks and performance. Previous work documents how workers' self-selection of roles can drive productivity differences across firms (Lazear, 2000). Causal evidence on assignment to preferred roles and performance thus far mostly comes from laboratory settings: e.g., Kamei and Markussen (2023) show via a lab experiment that mismatch between assigned tasks and task preferences undermines performance and produces free riding in teamwork settings, while Freyer et al. (2025) find productivity increases for subjects assigned preferred vs. non-preferred data entry tasks. Our study provides causal evidence from a real-world setting and the first to our knowledge from the public sector, where the impact of assignment along the preference-match dimension may be particularly salient and valuable given constraints on tools for incentivizing performance in bureaucracies. Compared to earlier work, we are also able to study impacts over a long time horizon and examine whether the experienced status of career positions crowds out effort.

We also contribute to the growing literature on bureaucrats in the developing world, which has generated evidence on questions from the selection of quality personnel (Xu et al., 2018; Callen et al., 2018; Hanna and Wang, 2017; Ashraf et al., 2020) to monitoring and incentives (Duflo et al., 2012; Dal Bó et al., 2020; Rasul and Rogger, 2018; Bandiera et al., 2021; Khan et al., 2016, 2019; Khan, 2023). We build on this literature by providing quasi-experimental evidence on the relatively underexamined role of bureaucrats' preferences over the nature of work they do in shaping performance. This despite the prevalence of preference-to-posting matching as a primary mode of assigning public workers in elite civil service. Our work sheds light on the intersection between organizational rules for assigning bureaucrats and the value of personal attributes (their preferences) to raising bureaucrats' effectiveness.

2 Setting

Pakistan's Central Superior Services (CSS) operate the administrative functions of national government, its personnel regarded as the "steel frame" of the nation (Wilder, 2009). Officials within the CSS staff civilian bureaucratic operations from the provincial levels up to the secretariats of the Prime Minister's cabinet, filling roles from sec-

tion officers overseeing the day-to-day functioning of specialized departments to Chief Secretaries informing policy at the highest levels of government.

There are twelve service groups to which newly-recruited bureaucrats can be assigned: Foreign, Police, Administrative, Customs, Inland Revenue, Railways, Military Lands and Cantonments, Commerce and Trade, Postal, Information, Audits and Accounts, and Office Management. Work across these services differs in focus area and subject matter but not as often in the day-to-day: CSS officers primarily fill oversight and management roles, directing the work of lower-ranked civil servants. Though the operations an officer in the Police service may be focused on ensuring security in a district, the nature of his or her role can resemble, in its emphasis on personnel management and translating national-level policies into local realities, that of a Postal officer who oversees mail deliveries and the financial services provided through post office banks.

A position in CSS has traditionally been considered a prestigious one that confers a high salary with perks like government-provided living accommodations, transport, healthcare, and security along with a post-retirement pension. This status has eroded somewhat in recent years with the rise of lucrative private sector alternatives, but a posting with CSS remains one of the most sought-after for college-educated Pakistanis.

2.1 Selection and Assignment Preferences

As a result, selection is competitive. Annual exams are held by the Federal Public Services Commission (FPSC) across the country, with candidates taking the exam in their home ("domicile") province. To qualify for the exam, individuals must be between 21 and 30 years old and possess a college degree, though exceptions are made for those with military experience. The CSS exams are comprehensive and exhaustive and include four parts: written, medical, psychological, and "viva-voce" (verbal interviews) evaluations.[3] Candidates scoring a minimum of 600 points "pass" the written stage and qualify for the psychological and medical evaluations, which assess suitability for service. Finally, a group of current CSS personnel meet candidates for the viva-voce interview, in which they assess communication skills, moral character, and leadership

[3]The written exam involves individual tests on six compulsory subjects — English essay, English comprehension, general science, current affairs, Pakistan affairs, and Islamic studies or comparative study of major religions (for Non-Muslims) — and six optional subjects to be selected from a variety of possible topics, including: computer science, economics, political science, chemistry, applied mathematics, business administration, history, zoology, languages, law, among others.

qualities. Candidates passing all four evaluations are recommended for assignment to service. The exams are competitive: in 2021, 17,240 candidates took the exam, with only 207 (1.2%) being recommended for service.

Newly-recruited officers can be assigned to one of the twelve service groups (SGs). Each year, the Establishment Division of the Federal Government releases the set of entry-level vacancies available in each service group. Allocation to service groups (SGs) then proceeds along candidates' rank in exam score and their preferences over services, which candidates specify during the exam process. Preferences must be unique — i.e., no service groups can be indicated as equally preferred — and can be specified across all 12 SGs. Candidates rank SGs from their most preferred to their least preferred. Figure 1 describes heterogeneity in preferences across preference ranks for the top five preferences. Top preferences are concentrated among the Pakistan Administrative Service, Police Service of Pakistan, and Foreign Service of Pakistan, with approximately 65% of candidates ranking Administrative Service as their most preferred group. Preferences become more diffuse at lower ranks, with Customs and Inland Revenue receiving the fourth and fifth largest share of interest. The least preferred SGs are Railways, Postal, and Information.

2.2 Allocation to Service Groups

In this section, we describe the allocation mechanism through which candidates qualifying for service are assigned to service groups, which provides the variation we exploit to explore the relationship between preferences and performance.

Assignment to a posting within the CSS is determined by (a) a candidates rank as a function of their score on the CSS Exam and (b) the candidate's preferences over SGs. Candidates specify posting preferences before completing the exam, and these preferences correspond to a rank-ordering of SGs from most to least preferred. Lifetime appointments to postings are made on the basis of the entrance exam. Each year, the national government announces a fixed number of new open positions within each SG. Candidates passing the CSS Exam are then assigned sequentially based on their exam scores and in order of their preferences. In other words, the highest-scoring candidate is automatically assigned his or her first-ranked preference, and so on, until the number of open positions in a SG are exhausted.

As a result, this allocation mechanism generates sharp thresholds in exam rank in the likelihood of assignment to a particular SG. To have obtained a more preferred

service, a candidate's exam rank must have been at least one rank higher than the candidate who scored high enough to be assigned to the last available position in a particular SG. The threshold for assignment to a SG is thus the largest value of exam rank (i.e., lowest rank) of the candidate assigned to that SG.

In reality, a quota system means that the allocation mechanism does not proceed purely sequentially down exam rank. Only 7.5% of vacancies each year are filled purely along merit lines (exam rank across all candidates, per the procedure described above). The remaining 92.5% of vacancies are filled according to provincial quotas that correspond to the relative population sizes of provinces. Implemented in 1948 to counter over-representation of Mohajirs (Muslim migrants from Indian post-Partition) in government posts, the portion of posts assigned to each province evolved over time but the current allocation has remained fixed since 1973. Punjab — the most populous province — receives 50% of the vacancies reserved for provincial quotas, and Sindh 19% (7.6% for urban, 11.4% for rural areas), Khyber Pakhtunkhwa 11.5%, Balochistan 6%, Federally Administered Tribal Areas and Gilgit-Baltistan 4%, and Azad Kashmir 2%. To address under-representation of women and minorities (Non-Muslims) in service, a quota reserving 10% of vacancies in each provincial share was included in 2007 and 5% quota in provincial share reserved vacancies for minorities beginning in 2009. However, assignment within each quota-defined set proceeds per the same steps: the candidate with the highest exam rank within the set is assigned to his or her most-preferred SG, and so on.

Though complex, allocation according to these quotas is made straightforward through the Establishment Division's specifying discrete numbers of vacancies across provinces, identity groups (defined by gender and religion), and service groups each year. Figure A1 displays the number of vacancies specified for pure merit, Punjab, and Sindh (rural) for the 2020 CSS exam. For the majority of vacancies, candidates are thus evaluated within the "set" defined by the domicile province, gender, and religion groups to which they belong, with the top ranked candidate in a set being assigned to their most preferred (available) service. The number of available vacancies determines the exam rank threshold for assignment to a particular service group through the application of this process: e.g., the exam rank of the candidate assigned to the last available vacancy in a service defines the cutoff below which joining that SG is no longer possible. Other candidates sharing the same home province and identity must be assigned to another SG. These thresholds provide the quasi-experimental vari-

ation we exploit to study the consequences of assigning a candidate to a SG that is more-preferred relative to a candidate assigned to a less-preferred SG.

Figure A7 describes the effective allocation of officers across service groups and preference ranks for our sample (described in Section 4). The modal preference rank assigned is the first preference, and 20.6% of officers get their most preferred posting, while 63.7% receive a second- through seventh-preferred posting. A roughly equal share — between 8.7 and 13.7% — are allocated to each preference rank in this range. The remaining 16.7% of officers are assigned to a posting between their eighth and twelfth most-preferred service. The largest services are Inland Revenue (17.6% of officers), Office Management (16.5%), and Administrative Service (14.4%). Slightly smaller shares go to Audits and Accounts (9.9%), Police (8.8%), Foreign (8.7%), and Customs (6.4%), while the smallest groups are Information (5.6%), Commerce and Trade (4.7%), Postal (3.4%), Railways (2.1%), and Military Lands and Cantonments (1.7%).

2.3 Advancement

Once assigned at service entry, bureaucrats persist in their SG for the remainder of their career in public service: bureaucrats can climb the hierarchy within service but are unable to switch SGs unless they retake the CSS exam, which is rare. Immediately after assignment to a vacancy, all candidates are trained together at the Civil Services Academy for six months, where they receive training in competencies common across SGs and relevant to all civil service operations. Officers are evaluated on individual subjects over the course of training and ranked again by merit at training conclusion. We use measures of performance during training to explore mechanisms in Section 7. Afterwards, officers then move to service-group-specific training, which can last between one and two years, before being deployed into service.

Initial deployment involves postings to field offices throughout the country at lower-level positions. Newly-assigned bureaucrats begin their career at Basic Pay Scale (BSP) 17, which carries the rank of section officer, whose tasks vary across service but generally involve administrative and organizational work relevant to service operations at the provincial level. The position provides relatively attractive terms: the starting salary at BSP-17 is 30,000 Rs. per month, which translates in PPP terms to an annual salary of approximately USD$8,687. Bureaucrats also receive housing and medical benefits. Advancement in service proceeds according to a promotion formula: bureaucrats become eligible for promotion to BSP-18 after five years of service, where they often

fill roles as Deputy Secretaries at the provincial level. After twelve years, bureaucrats are eligible for BSP-19, becoming in many groups Additional Secretaries with more responsibilities. The pyramid narrows at higher levels, with bureaucrats being eligible for promotion to levels BSP-20 to BSP-22 after additional years of service but subject to available positions and stronger performance assessment. The BSP-22 level is the terminal level in civil service and includes the highest ranking officials in the Central Superior Services like Provincial and Chief Secretaries, among the most powerful positions in the government.

2.4 What Do Preferences Capture?

To identify potential channels through which preferences over SGs may influence performance within our sample, we surveyed a batch of newly-assigned bureaucrats in November 2022, just as they began training at the Civil Services Academy. We collected measures relating to the formation of preferences, candidates' rationale for selecting top preferences, and the relative strength of preferences across preference-ranks. We also collected short-run measures of motivation.

Preferences appear to reflect interest in the tasks in which bureaucrats expect to engage in a SG, opportunities for career growth offered within a SG, and the match of SG tasks to candidate skills. Figure 2 reports the share of candidates rating as "very important" reasons for selecting their top-ranked preference. Being interested in the tasks of the SG is reported most often as an important factor (57% of candidates), with career growth opportunities (46%), and a belief that the SG would best use candidate skills (44%) as close runner-ups. An additional factor in preference selection appears to be prestige: approximately one-third of candidates report "everyone saying an SG is best" or "the SG is prestigious" as being very important reasons for ranking it first. Other inputs to preference determinations — like possessing information about and SG, family preferences, work/life balance, and connections — appear less important.[4] Taken together, we interpret these patterns as suggesting intrinsic motivation (via interest in the work of an SG) and career objectives (via what an SG can offer in terms of advancement) as potential channels that may affect performance. Match to candidate capabilities (via maximizing use of candidate skills) and prestige concerns also appear plausible explanations. We discuss evidence for these mechanisms in Section 7.

[4]Figure A3 summarizes sources of information, showing family information and peers, friends, and classmates seem to inform preferences most strongly.

The relative strength of preferences is also of interest for understanding how assignment to a more preferred service may engage such channels. An agnostic view treats all moves between preference-ranks as equivalent — i.e., that marginal assignment from one's second to first preference-rank is the same as marginal assignment from one's sixth to fifth — yet candidates may have stronger preferences over top-ranked SGs or vice versa have strong aversion to to lower-ranked preferences. In the survey, we asked candidates about the relative strength of their preferences for various combinations of preference-ranks. Figure A2 displays the share of candidates reporting that they strongly prefer a particular preference rank over another. The strongest preference-rank comparisons arise when comparing first to third preferences and second to third preferences, whereas candidates are relatively more indifferent between first and second preference-ranks, as well as third vs. fourth and fourth vs. sixth. We interpret these differences as evidence that treatment effects on performance we identify may draw from variation in marginal assignment to first or second vs. third preferred SGs. We explore heterogeneity by preference rank assigned in Section 6.

Finally, we collected measures of sentiments regarding assignment at training entry. Figure A4 displays the share of candidates reporting particular sentiments. There is a high degree of relief at receiving their assigned SG (>70% of candidates report such feelings), but a non-trivial share (approximately 20%) report disappointment with their assigned SG, specifically that they wish they would have been assigned a more preferred SG or one with a different nature of work.

3 Conceptual Framework

A simple conceptual framework motivates our empirical design. This framework illustrates the ambiguous performance effects of assignment to more preferred postings, identifying the role that prestige may play in crowding out intrinsic motivation.

We repurpose the model of bureaucrats' career incentives in Bertrand et al. (2020) to consider the performance implications of preference-match, drawing on theories of motivated agents from Besley and Ghatak (2005) and Cassar and Meier (2018b). We first introduce intrinsic motivation arising from preference-match then incorporate the role of prestige.

3.1 Preference-Match and Performance

We consider bureaucrats of two types that vary in whether they are matched to a preferred posting: $a = 1$ if an individual preferred assignment to their ultimate post and $a = 0$ if they did not. After assignment, a career spans a junior-level period ($t = 1$), in which effort and performance do not impact advancement, and a senior-level ($t = 2$), promotion to which depends on performance at the junior-level.[5]

We assume bureaucrats can differ in their ability μ_a, receive a fixed wage w_t in each period (with $w_1 \leq w_2$), and exert effort e_{at} to produce output at convex cost $C(e_{at}) = \frac{e_{at}^2}{2}$.[6] We also assume the cost of effort $c = 1$ (Besley and Ghatak, 2005). Bureaucrats receive intrinsic value $m(e_{a1}, a)$ from engaging in the work of their assigned post, which is a function of the effort they expend (e_{a1}) and their preference for assignment to their post (a). We assume that $m(e_{at}, a) = \theta_a e_{at}$ with $\theta_a > 0$ for $a = 0, 1$ and $\theta_1 > \theta_0$.[7] Performance is linearly separable in effort and ability: $p_t(e_{at}, a) = e_{at} + \mu_a$.

Bureaucrats choose effort to maximize their payoff across both periods:

$$\max_{e_{a1}, e_{a2}} V = w_1 + \theta_a e_{a1} - \frac{e_{a1}^2}{2} + \Pi(p_1(e_{a1}, a), a)(w_2 + \theta_a e_{a2} - \frac{e_{a2}^2}{2}) \tag{1}$$

where $\Pi(p_1(e_{a1}, a), a) \in [0, 1]$ is the promotion function that captures the probability that bureaucrat of type a and output level $p_1(e_{a1}, a)$ is promoted to the senior level, assumed to be linear in output conditional on type: $\Pi(p_1(e_{a1}, a), a) = \pi(a)p_1(e_{a1}, a)$, with $\pi(a) = \pi(0) = \pi(1) = \pi$, and thus increasing in performance. We assume that a bureaucrat who is not promoted exits and earns a zero wage.

Proposition 1. With e_{at}^* the optimal choice of effort, initial period performance differences between bureaucrats who prefer their posting and do not will stem from differences in both intrinsic motivation and underlying ability:

[5] In our setting, the senior-level can represent upper-level positions (e.g., Secretary) to which promotion is not guaranteed given limited availability of open roles at the top of the hierarchy but could also stand in for roles (e.g., advancement to higher intermediate levels, directorships, or desirable positions) below the upper-level that take performance into account when selecting among candidates.

[6] These are identical to the **?** setting. While, as in theirs, the goal of our empirical design is to neutralize ability differences, we are instead focused on within period differences in intrinsic payoffs rather than cross-period differences in financial incentives.

[7] Appendix Section A2 details the joint restriction on θ_a, π, w_2, and μ_a that required for the model to be well-defined (Assumption A2).

$$\Delta p_1 \equiv p_1(e^*_{11}, 1) - p_1(e^*_{01}, 0) = (\tilde{\theta}_1 - \tilde{\theta}_0) + (\mu_1 - \mu_0). \quad (2)$$

We define $\tilde{\theta}_a \equiv \theta_a + \pi\frac{\theta_a^2}{2}$, which combines period-specific motivated effort (θ_a) and a dynamic term ($\pi\frac{\theta_a^2}{2}$) reflecting the value attached to intrinsic payoffs in the later period. Appendix Section A2 provides the proof using the model's first order condition and backward induction.[8]

Bureaucrats assigned postings they prefer will exert more effort — and exhibit higher performance — than those assigned to postings they do not. The difference in motivated effort is captured by the $\tilde{\theta}_1 - \tilde{\theta}_0$ term which depends on the gap in intrinsic payoffs arising from the link between preference-to-posting match and motivation. Note that θ_a could also be conceptualized as capturing unique inputs to motivation that Cassar and Meier (2018b) propose are multidimensional, including competence through suitability, interest in the "mission" of an organization, preference for a style of work, or relatedness to skills.

Differences in performance may also arise from differences in ability ($\mu_1 - \mu_0$). In our empirical setting, the exogenous variation in the preference-to-posting match we exploit should make bureaucrats of each type comparable, isolating motivation as the driver of performance differences.[9]

3.2 The Role of Prestige

We next incorporate into our framework the prestige or status attached to a posting, which we show can countervail intrinsic rewards. Let $r \geq 0$ denote a posting's prestige. In our empirical setting, preference-to-posting match and prestige are typically strongly correlated: postings that bureaucrats prefer are in general also more prestigious. We thus focus on the comparison most relevant to our empirical design: *Type I* bureaucrats, assigned to a posting they prefer and which is more prestigious ($a = 1$, $r = r_H$), versus *Type II* bureaucrats, assigned to a less prestigious posting they do not prefer ($a = 0$,

[8] It is worth noting that in a single period, differences in effort are simply the difference in the marginal intrinsic payoff across types ($\theta_1 - \theta_0$). We embed these differences in a dynamic model to illustrate the role of differences in ability.

[9] We also test this assumption by testing for differences in bureaucrat characteristics at entry as well as controlling for these characteristics when estimating performance differences.

$r = r_L$), with $r_H > r_L \geq 0$.[10]

We introduce a simple mechanism that illustrates the conditions under which prestige could lead Type I bureaucrats to exert *less* effort than Type II bureaucrats, despite preferring their post more. Specifically, we allow bureaucrats to obtain value from a posting's prestige independent of the effort they exert (e.g., from simply belonging to a prestigious service), yet prestige reduces the marginal payoff to effort deriving from intrinsic motivation. This mechanism reflects the intuition that status payoffs may displace those from intrinsic motivation: the validation provided by a prestigious posting can erode the returns from exerting effort in a desired role. We incorporate these forces into equation (1) as two distinct terms: a fixed utility rent $\psi(r)$, $\psi'(r) \geq 0$, that a bureaucrat receives each period from holding post within a service of prestige r, regardless of the effort they exert, and a crowd-out term: a bureaucrat of type a in a posting of prestige r has a marginal intrinsic-motivation payoff to effort $\theta_a - \delta r$, where $\delta > 0$ indexes the intensity of crowd-out and is common across types.[11,12]

Proposition 2. If bureaucrats maximize:

$$V = w_1 + \psi(r) + (\theta_a - \delta r)e_{a1} - \frac{e_{a1}^2}{2} + \Pi(p_1(e_{a1}, a), a)\left(w_2 + \psi(r) + (\theta_a - \delta r)e_{a2} - \frac{e_{a2}^2}{2}\right), \tag{3}$$

with $\psi'(r) \geq 0$ and $\delta > 0$ common across types and $r < \min(\theta_0, \theta_1)/\delta$ for both posting types, so that both types retain strictly positive net marginal motivation, then:

(i) *No prestige gap.* If prestige does not vary with preference-match ($r_H = r_L$), Type I bureaucrats exert strictly more effort than Type II bureaucrats, as in the Proposition 1 case: $e_{11}^*(r) > e_{01}^*(r)$.

[10] Examples of assignment to more preferred (according to stated preferences) but less prestigious as well as less preferred but more prestigious postings exist in our data but are rare. These cases may also reflect attempts to "game" the system by ranking less attractive postings as more preferred.

[11] We assume a common δ across types for simplicity as doing so allows us to isolate the prestige crowd-out mechanism with a single free parameter, rather than requiring an additional assumption about whether well-matched or poorly-matched bureaucrats are more susceptible to it. We show in Appendix A2 that the model can be easily extended to allow δ to vary by type.

[12] We are focused on identifying a reasonable case in which bureaucrats posted to more preferred, more prestigious settings may perform worse than those posted to less preferred, less prestigious ones. However, the prestige of a posting could instead magnify the marginal payoffs ($\delta < 0$), in which case officers would exert more effort (perform better) when preference-match is stronger and a post confers higher status. In this case, equation (4) is strictly positive.

(ii) *Prestige gap.* If Type I faces prestige r_H and Type II faces prestige $r_L < r_H$, then

$$e^*_{11}(r_H) - e^*_{01}(r_L) = \Big[(\theta_1 - \theta_0) - \delta(r_H - r_L)\Big]\Big[1 + \frac{\pi}{2}\big((\theta_1 - \delta r_H) + (\theta_0 - \delta r_L)\big)\Big] + \pi\big[\psi(r_H) - \psi(r_L)\big], \tag{4}$$

Type I bureaucrats exert less effort than Type II bureaucrats if and only if

$$\delta(r_H - r_L) > (\theta_1 - \theta_0) + \pi(\psi(r_H) - \psi(r_L))/B, \tag{5}$$

where $B \equiv 1 + \frac{\pi}{2}((\theta_1 - \delta r_H) + (\theta_0 - \delta r_L)) > 1$. Appendix Section A2 provides the proof and discusses the inclusion of performance differences (including μ_a) in addition to effort and allowing for type-specific δ's.

In other words, Type I bureaucrats exert less effort (and perform worse) than Type II bureaucrats when the crowd-out generated by the relative prestige of the Type I posting ($\delta(r_H - r_L)$) exceeds both (a) the relative intrinsic motivation generated by the Type I posting being preferred ($\theta_1 - \theta_0$) and (b) the incentive generated by the persistent status gap between service assignments, ($\pi[\psi(r_H) - \psi(r_L)]/B$), which crowds in effort by making a more prestigious service more valuable to advance within. Because Type I postings are both more preferred and more prestigious, our empirical question is which of these forces dominates: the preference-match channel, which raises effort, or the prestige channel, which could lower it.

To contextualize these tradeoffs, imagine two bureaucrats (A and B) assigned to two distinct service groups. The work they undertake differs little, but both A and B ranked A's assigned service as more preferred ($\theta_1 - \theta_0$ is positive but close to zero). However, A's service is more prestigious ($r_A > r_B$).[13] There are two forces: First, the degree to which the marginal intrinsic-motivation payoff erodes in the more prestigious position is captured by δ. Even though both A and B derive utility from exerting effort in their work, A derives less than B does because A's marginal payoffs are diluted by prestige: the utility gained by, e.g., completing a successful project is less for A because they retain a high status regardless. This modeling choice reflects cases, for

[13] An empirical analog in our setting would be involve comparing bureaucrats assigned to the Pakistan Administrative Service (PAS), the service most often ranked first by candidates and thought to be the most prestigious, versus the Office Management Group (OMG), a service candidates rank on average as their fifth or sixth preference. Bureaucrats' tasks in these services can often be very similar — involving management of personnel and local departments —- despite the prestige gap between the two.

example, where the value attached to status comes to displace intrinsic motivation: after entering a junior role, bureaucrats grow to care more about status than intrinsic value when they become accustomed to receiving status rents from prestige, such as positive recognition by others independent of effort. The second force is a relative boost, absent effort, that A gets from holding a position in a more prestigious service relative to B ($\pi(\psi(r_A) - \psi(r_B))/B$). This expected future payoff incentivizes bureaucrats to exert effort to ensure they retain it by advancing to a senior role within their service.

The tension lies in the relative magnitude of these forces. It is plausible to imagine these forces move in tandem: the greater the gap in status payoffs ($\psi(r_A) - \psi(r_B)$) the greater the crowd-out effect will be (a larger $\psi'(r)$ implies a larger δ). This framing posits the pure status rent and crowd-out effect flow from the same source. The more value a prestigious posting carries on its own, the larger the crowd-out of intrinsic motivation. This co-movement makes condition (5) harder to satisfy.[14] Yet it is not necessary that a high δ implies a high $\psi'(r)$. If the prestige of a posting mainly represents greater material benefits like better amenities, workloads, attractive locations but does not provide a "stamp" of superior identity (e.g., a signal of inherently belonging to "the best"), displacement of intrinsic motivation may be minimal despite a large difference in status payoffs (high $\psi'(r)$, low δ). Conversely, if a prestige posting carries little additional material rewards (low $\psi'(r)$) but marks someone as "elite," this label could undermine effort via self-signaling (high δ). A bureaucrat in a more prestigious service can coast on their title instead of proving to themselves through hard work that they are high-quality. In this case, δ would capture a psychological mechanism by which an elite status undermines via self-signaling the marginal incentive to prove through execution that status to oneself. Whereas $\psi'(r)$ would embody a relative advantage in the tangible rewards.

Applying the framework. How these forces play out in our setting or a comparable one depends on the extent to which these conditions match empirical realities. The goal of our framework is not to argue that the assignment process and nature of differences across service groups in our setting aligns more with one version than another. Instead, we intend to underscore the fundamental ambiguity of how preference-match impacts

[14] In this case, the exogenous portion of the promotion probability π, which does not depend on effort, becomes pivotal. When the probability of advancement to a more senior role is high (e.g., promotion to the next grade in service), the continued status rents a bureaucrat expects loom larger. However, when promotion to a senior post is rarer (e.g., to the highest levels of Secretary, which only 2% of officers in our setting ever attain), then differences in effort will depend more on the size of crowd-out relative to the difference in intrinsic-motivation payoffs.

performance when prestige and preferences correlate.

We use this framework to interpret the effects of preference-match on performance measures we estimate, and in considering mechanisms. Condition (5) predicts that comparisons of postings where prestige differences across assignments are large relative to preference-match differences. For example, a highly prestigious posting that a bureaucrat only somewhat prefers, compared to a low-profile posting they only somewhat dislikes should exhibit larger negative differences in performance when comparing bureaucrats assigned to a more versus less preferred post. We investigate the role of prestige directly in Section 7.2.

4 Data and Sample

We combine administrative records with survey data. We draw on two primary administrative data sources. First, we use CSS exam and allocation records that include individual exam ranks and final service group assignments for thirteen batches (CTPs) of CSS officers between 2008 and 2021, 3,256 officers in total. We collected this information from publicly available records of exam results and service allocations. The Civil Services Academy provided records of officers' preferences over service groups (declared during exam) and demographics like age, gender, religion, domicile, and family characteristics. Together these data allow us to define thresholds — within the quota categories corresponding to officer characteristics — determining assignment to service groups and the alignment between assignment and preferences for all officers in our sample.

Second, we construct a new administrative dataset recording individual career outcomes for officers in our sample that occur between 2013 and 2024. The start date of this period is five years after the first batch in our sample enters service, which corresponds to the time at which officers become eligible for their first promotion. The end date extends three years after the last batch in our sample entered service in 2021. We collected this data from the Gazette of Pakistan, which is a weekly bulletin of notifications pertaining to civil servants. The gazettes record career events for all civil service officers,[15] including promotions, transfers and postings, assignment to additional and acting charges, punishment and awards, completion of trainings, taking

[15]This include CSS officers as well as officers in provincial service that are not members of the CSS as well as officers recruited into national service outside of the CSS exam, who are assigned to lower service rungs.

of leave for personal reasons, and exit from service. Weekly gazettes were circulated only in printed form to government departments throughout Pakistan until 2021, after which they also became available online. We digitized the weekly gazettes for all weeks between 2013 and 2024 then matched career-related notifications to individual CSS officers in our sample. To do so, we first purchased paper copies of gazettes from the Printing Corporation of Pakistan, the institution responsible for printing and circulating paper gazette copies. We then digitally scanned gazette copies and scraped individual notifications, extracting those that included a mention of civil service officers at Basic Pay Scale 17 and above (level of Central Superior Services), contained keywords for career outcomes, and which contained a name that at least partially matched the name of an officer in our sample (i.e., there existed a match between first, last, or other names). We identified a set of notifications potentially relevant to our sample totaling 35,000. Enumerators then manually reviewed each of the selected notifications using the scraped text and referring to the relevant digitized gazette copies. Details for career events were recorded for each notification (see Figure A6 for an example of the information extracted for an individual notification). A randomly selected portion was double-entered to ensure accuracy. Finally, the set of coded notifications was matched to names in our sample of CSS officers. Due to the non-existence of a unique identifier to link notifications to our sample and slight differences in the spelling of names across datasets, a team of research assistants then reviewed each match between an officer in our sample and an individual notification, checking that service group assignment, batch, and other details, including name, were appropriately matched. (Figure A5 provides a visual illustration of this process.) We can ultimately match approximately 14,000 individual notifications to officers in our sample.

We complement these records with two sources of survey data, collected during trainings and workshops facilitated by the CSA. We first conducted surveys with the 50th CTP batch of officers newly-assigned to service. At the time of survey, these individuals had just completed the 2022 CSS exam. We surveyed 201 bureaucrats via an online tool as they began training at the Civil Services Academy in November 2022, collecting information about the formation of their preferences across SGs, their rationale for preferring particular SGs, and the strength of their preferences. We also collected data on satisfaction and motivation regarding their assignment. We discuss these data in Sectio 2.4 to shed light on mechanisms through which preferences may shape performance.

Second, between March 2025 and September 2026, we conducted detailed surveys with officers from the 37th through 49th batches. These surveys involved collecting data on individual respondent's career trajectories and views of service. Crucially, the survey also collected peer ratings of colleagues. Officers were asked to rate up to 7 colleagues — primarily within the same service group[16] — on effectiveness, ability to withstand political pressure, interest in personal gain, and sensitivity towards the poor. We use these ratings to construct a measure of bureaucrat effectiveness to complement to the administrative data measures. Recruitment of officers for this survey presented a challenge: we organized multiple workshops to deliver training on integrating evidence into policy, after which the survey was administered. All told, we were able to collect peer ratings from 135 officers, covering over 300 peers. While the sample size is limited relative to the career outcomes, we test for representativeness and thus interpret the limited response rate as primarily representing a limitation on statistical power rather than a source of bias.

Additional data sources. We draw on supplementary sources of data that provide additional outcomes. First, we use evaluations conducted by the Civil Services Academy during the introductory training of newly-assigned bureaucrats captures performance on individual subjects — from economics, history, and English, to leadership and public speaking — as well as a merit rank assigned for overall training performance, unique across candidates in a batch. Training data provides these measures for all 3,256 candidates in the exam data. These data provide measures of short-run performance post-assignment that are comparable across service groups. We use this data to explore mechanisms in Section 7.

4.1 Sample

Of the 3,256 candidates assigned to SGs between 2008 and 2021 in the exam and training data, our analysis excludes those who were not assigned according to the allocation mechanism described above. These candidates — around 6.5% — originate from military service. We do not consider them in our analysis as they do not receive an exam rank like those applying through the CSS exam process and do not specify preferences across service groups. Our analysis sample thus comprises all candidates receiving an exam rank and ultimately assigned to service in a given year.

[16] We prioritized ratings of peers within the same service group and batch, but also obtained ratings of those within the same service group and +/- 3 batches.

Summary statistics for our sample are presented in Table A1. Candidates are on average slightly more likely to be male (64.9%), are 26 years old on average, and 24% are married. The average annual income of candidates' fathers is 1.3 million Rs. (USD$4,800 PPP). We use these variables as well as the others listed in Table A1 —- which provide measures of health (O-type blood indicating a lower propensity for illness), skills, and other family characteristics like parents' ages and number of siblings and children — when considering validity of our empirical design in Section 5.1 and control for them in all specifications. The mean preference to which candidates are assigned is 4.49, meaning that the average candidate is assigned to the SG they ranked as 4th or 5th most preferred.

5 Empirical Design

Our empirical design exploits non-linearities in assignment to SGs arising from the allocation mechanism described in Section 2.2. The service vacancies made available each year define 252 possible thresholds, across 12 service groups, 7 provinces, and 3 identity group combinations (male Muslim, female Muslim, and minority).

This set of thresholds contains those which may possibly bind, depending on a candidate's exam score and preferences. In reality, it is likely that one unique threshold will determine marginal assignment to more preferred SG. We draw from the threshold set to construct *preference-rank thresholds*. These thresholds are defined across the ranking assigned by candidates over SGs for ranks $k \in [1, 12]$. To identify the threshold for a preference-rank k, we identify all SGs preferred kth by candidates, the set of which will contain different SGs. We then identify the exam rank threshold from the full set of 252 possibilities that corresponds to the threshold for assignment to each candidate's kth-preferred SG within their "set" (year/batch, province, and identity group). We then center scores within each kth-preferred SG. We repeat this process to identify preference-rank thresholds for ranks 1 through 12.

Our setting resembles that of Rose and Shem-Tov (2021), which consider non-linearities in sentencing outcomes for offenders, in that our running variable — exam rank — is discrete, with multiple thresholds affecting likelihood of assignment to more preferred service groups within distinct comparison sets defined by year, province, and identity. This approach reflects a parameterized regression discontinuity design in the spirit of the one considered in Clark and Del Bono (2016).

We include separate linear slopes in exam rank on each side of preference-rank thresholds and allow for positive shifts ("jumps") in the probability of assignment to a higher preference at the threshold. The thresholds represent boundaries between preference-ranks and thus potential instruments. Figure 3 illustrates changes in the average service preference-rank assigned to candidates around thresholds for preference-ranks first through sixth. Marginally surpassing the threshold in exam rank for preference-ranks first through fifth results in a discontinuous increase in the preference-rank to which a candidate is assigned that is approximately equal to 1 and statistically distinguishable from zero for each of these thresholds. At the sixth preference-rank (and beyond) the average SG preference-rank assigned increases to the right of the threshold, but the change is not statistically significant.

Our strategy uses the five discountinuities where the likelihood of assignment to a more preferred SG is statistically detectable as excluded instruments (the first through fifth preference-rank thresholds). The empirical specification stacks the variation from each discontinuity to estimate a single treatment effect. We represent this approach formally in the two-stage equation system below, mirroring the approach in Rose and Shem-Tov (2021). The first stage (equation (6)) estimates the preference-rank to which a candidate is assigned as a function of exam rank, preference-rank discontinuities, and individual-level covariates. The second stage (equation (7)) estimates the relationship between an outcome, preference-rank assigned, and non-excluded controls. We restrict comparisons to be within candidate "sets," defined by year, province, and identity group.

The first-stage equation is modeled as:

$$\begin{aligned}\text{PreferenceRank}_{ig} = \; & \underbrace{\eta_g^1 + \mathbf{X}_{ig}\mathbf{\Gamma}^1}_{\text{controls}} \\ & + \underbrace{\sum_k 1\{\text{preference}_{ig} = k\}\Big[\sum_l \beta_{lk}^1 1\{r_i \geq l_g\}(r_i - l_g + 0.5) + \gamma_k^1 r_i\Big]}_{\text{linear slopes in rank by preference}} \\ & + \underbrace{\sum_{k,l \in shift} \xi_{kl} 1\{r_i \geq l_g\} 1\{\text{preference}_{ig} = k\}}_{\text{preference discontinuities}} + \varepsilon_{ig}\end{aligned} \tag{6}$$

The second stage equation takes the form:

$$Y_{ig} = \beta_0 \text{PreferenceRank}_{ig} + \underbrace{\eta_g^2 + \mathbf{X}_{ig}\mathbf{\Gamma}^2}_{\text{controls}} + \underbrace{\sum_k 1\{\text{preference}_{ig} = k\}\Big[\sum_l \beta_{lk}^2 1\{r_i \geq l_g\}(r_i - l_g + 0.5) + \gamma_k^2 r_i\Big]}_{\text{linear slopes in rank by preference}} + \varepsilon_{ig} \quad (7)$$

where $\text{PreferenceRank}_{ig}$ represents the preference-rank to which candidate i belonging to set g is assigned. η_g are set (year times province times identity group) fixed effects, and r_i is the exam rank. Thresholds l_g correspond to the preference-rank thresholds within a set g for preference-ranks $k \in [1, 5]$. When estimating changes on either side of thresholds, we recenter by $r_i - 0.5$ to set the threshold halfway between the exam rank values on either side, rather than at either extreme.[17] Covariates X_{ig} include the demographic variables listed in Table A1 (excluding preference assigned). We calculate robust standard errors.

We use this strategy to explore whether marginal assignment to a more preferred SG affects performance.

5.1 Validity Tests

We assess the validity of our instruments by examining whether the discontinuities predict the individual characteristics listed in Table A1.

These tests support the assumption that changes in outcomes across discontinuities reflect the impacts of assignment to more preferred SGs rather than selection along characteristics correlated with performance. Table A2 summarizes balance tests across all the five discontinuities we employ as instruments, displaying p-values on the estimated coefficients.[18] Only 6 out of the 85 (7.1%) comparisons reported in Table A2 are significant at the 10% level, one (1.2%) at the 5% level, and none at the 1%, in line with random chance.

Additionally, Figure A8 displays no visual evidence of discontinuities in density at preference-rank thresholds, and manipulation tests using local polynomial density

[17] We adopt this approach as it is applied by Rose and Shem-Tov (2021) to accommodate a discrete running variable, but, as in their case examining sentencing of offenders, our results are not sensitive to this choice.

[18] We omit gender as it is included as part of the set fixed effect included in the specification estimating these differences.

estimators (Cattaneo et al., 2020) report no statistically significant differences.

6 Results

We estimate impacts using two-stage least squares on career outcomes drawn from the administrative gazette data. We focus on promotions and transfers as these events are the most consequential for career advancement and the events most frequently reported in gazette notifications, but we also consider punishment, leave-taking, and exit from service. We restrict our sample to CSS Batches 37–43 (years 2009–2015), given the reach of our administrative outcomes only extend to 2021, and we observe far fewer promotion events overall for Batches 44 and beyond. Transfers appear more often in these later batches, but to ensure comparability in the promotion and transfer outcomes as applying to the same sample, we restrict the sample in this way for all administrative outcomes. For peer ratings, we consider officers of all batches as there is no time-censoring problem for our peer assessments.

6.1 Promotion

Table 1 reports 2SLS estimates from equation (7) for promotions. The 2SLS estimate using our baseline approach, which includes only demographic variables and set fixed effects as controls (Column 1), corresponds to a 3.75 percentage point (4.83%) reduction in the probability of being promoted when an officer is assigned to a preference-rank one step "lower" (i.e., less preferred, larger preference-rank value) in the sample period we observe. Adding fixed effects for candidates' top three preference rankings only marginally reduces the estimate to -3.65 percentage points (Column 2).[19] In Column 3, we include a control for the average of promotion among previous batches within the same service group. This covariate accounts for differences across service groups in the average likelihood of promotion, arising from underlying differences in their organization. Indeed, there is substantial variation in the likelihood of promotion across

[19]The determination of preferences may itself by correlated with ability and future performance — e.g., if the top-performing candidates are more likely to select the most prestigious services. To account for this possibility we include preference-type fixed effects to restrict comparisons among those with the same preferences over services. The dimensionality of potential SG-preference combinations is large, thus we only define such fixed effects for the top three preference-ranks being identical (i.e., having the same preferences over services for preference-ranks 1 through 3), which are spread across a limited number of services and thus ensure we have sufficient support in each service-preference group.

service groups (Figure A9, Panel A).[20] Thus we view this adjustment as necessary to ensure promotion likelihood is comparable across service groups. We detect a smaller decline in the likelihood of being promoted on average for being assigned to one rank less preferred service group that amounts to to a 2.88 percentage point 3.71% reduction. In our preferred specification (Column 4), we add the leave-one-out mean of the outcome calculated with "set" (year, province, and identity group cells) and service group, which controls for the promotion outcomes of peer officers subject to the same thresholds and ending up in the same service group for each observation in the data. The estimated impact of assignment to a one-rank less-preferred service group is -1.81 percentage points (2.33%) Assigning a candidate from their first to the CSS-wide average fourth preferred SG translates into a 6.99% reduction in the probability of being promoted within the five to ten year window our data captures.

We also examine how promotion likelihood varies with the rank of preference assigned. Columns 5–7 of Table 1 consider assignment to 1st preference relative to lower preference ranks (Column 5), 2nd vs. lower ranks (Column 6), and 3rd vs. lower ranks (Column 7). Assignment to one's most preferred service increases the likelihood of promotion (within our sample window) by 13.9 percentage points (25.6%). Assignment to one's 3rd preference vs. a less-preferred one corresponds to a 10.1 percentage point (18.3%) increase in promotion likelihood ($p = 0.06$), while we observe no statistically significant difference for assignment to one's 2nd preference (vs. lower ranked preference).

Closer correspondence between a bureaucrat's posting and assignment increases the *speed* of promotion rather than necessarily the likelihood, within the time period our data covers. Promotion in Pakistan's civil service, as in other settings, is relatively formulaic, with officers qualifying for a move up in rank after a specific number of years in service. Our data captures officers mostly at the beginning of their careers in service — the typical length of service is 25 years and we collect administrative outcomes for up to the first 10 years. Bureaucrats assigned more preferred postings are more likely to be promoted after five and ten years, though in the intervening years this difference dissipates as their counterparts "catch up" in moves along the hierarchy. Indeed, promotion likelihood does not differ by preference-rank in the first through

[20]Note that the levels of promotion probability in Figure A9 are lower than the reference group mean in Table 1. This is both because (i) the sample in Table 1 is restricted to earlier batches, which are more likely to have reached seniority, and (ii) the reference group corresponds to those assigned more preferred service groups (e.g., preference rank 1 in Column 1), which have a higher likelihood of promotion. Figure A9 pools all observations across batches and preference ranks.

fourth years after service entry then jumps in the fifth year (Figure 4), with the effect persisting in subsequent years.

These effects can be credibly understood as capturing performance differences as "fast tracking" of promotions within Pakistan's CSS is applied to recognize high performers (Aman-Rana, 2025). Additionally, we find evidence that the strength of the preference-to-posting match can affect the promotion *ceiling*, i.e., the likelihood of reaching the upper rungs of service, to which promotion is not guaranteed and for which competition is high given the dearth of positions at the top. Assigning a bureaucrat to a marginally less preferred posting reduces the likelihood they are promoted to a Secretary-level position (the highest possible) (Table A3, Column 8). The effect we detect is small in magnitude but corresponds to a 24.8% reduction in reaching the highest bureaucratic level. Posting to a more preferred service thus also raises the limit on levels of the hierarchy to which officers can aspire.

We assess the robustness of our main effects through several strategies. A primary concern is the possibility of differing practices regarding career paths across services. For example, the Pakistan Administrative Service is anecdotally viewed as "hogging" promotions; indeed, 37% of PAS officers in our sample are promoted compared to a 17% average across all service groups. Likewise, 74% of officers in the Office Management Group are transferred compared to an across-service-group average of 38%. To account for variation in career outcomes arising from differences across service groups, in our preferred specification we control for differences in promotion likelihood across service groups (lagged average for previous batches) and over time (leave-one-out mean within group defined by "set" and service).

This approach, however, may not entirely resolve the issue embedded in our definition of treatment: that assignment to a more preferred service group also entails assignment to a *different* service. If services have different performance evaluation rules or opportunities for advancement, estimates derived using the former strategy may primarily capture the effects of services rather than the strength of preference-to-SG match. This may not be a salient concern for performance outcomes like peer ratings and training evaluations (considered in Section 7) but could be for longer-run outcomes like promotion. To address this issue we implement an alternative strategy that includes versions of service group fixed effects when estimating the impact of assignment to a more preferred SG. This approach requires dropping set fixed effects — as the variation within sets precludes identification of "treated" and "less treated"

(assigned lower preference-rank) pairs in the same SG as there is a unique threshold within set determining assignment to a service. Removing set fixed effects, however, prevents the restriction of comparisons within the set of candidates from whom a particular threshold binds. Thresholds determining assignment to service groups remain the same — corresponding to an individual's set and preferences — but we are unable to compare marginally treated to marginally less treated candidates within the same set. To compensate for this loss of precise variation in treatment comparison, we include independent fixed effects for year, province, and identity group to partial out variation arising from differences across each margin.

Applying this alternative comparison strategy — including service group fixed effects, dropping set fixed effects and including independent year, province, and identity group fixed effects — results in a substantially weaker first stage (with F-statistic of 24.76, compared to 277.6 in our preferred approach) and substantially noisier estimates (Table A4), yet the point estimate remains negative and even slightly larger in magnitude than our main estimate (Column 3).[21] Ultimately, comparisons within service groups rely on comparing officers who end up assigned to the same service group but preferred it relatively more or less but who are drawn from different batches (recruitment years) or provinces and who thus may be difficult to compare for other reasons.

We also estimate a stacked regression discontinuity design using data from the full sample (i.e., no longer restricting to Batches 37-43). We do so to provide a more complete picture of how promotion probability is affected by assigned preference rank. Figure 6 (Panel A) plots the probability of promotion by exam rank distance to thresholds (exam rank of candidate assigned last available vacancy) determining assignment to a more preferred posting. The outcome is the residual from regressing each measure on demographic controls (including dummies for missing values) and the leave-one-out mean of the outcome calculated within batch and service group (and a dummy for the leave-one-out mean being missing). The regression also includes the average of the year-specific outcome in up to nine previous batch years (as well as a dummy for the previous batch average being missing). We stack observations around the preference-rank thresholds for the 1st to 5th preferences, include set fixed effects, and cluster

[21] We also include a version with service group by province fixed effects (to compare officers from the same service and province), which returns a similar estimate (Column 4), and a version with service group X province X batch fixed effects (Column 5). The latter severely restricts comparisons groups to exceedingly few officers with a far weaker first stage ($F = 12.35$). We view the null estimate as uninformative.

standard errors at the candidate level. We estimate that assignment to a marginally more preferred service group results in a 5.21 percentage point increase in promotion likelihood ($p = 0.094$).

6.2 Transfers and Other Career Outcomes

Transfers are common: 63.5% of our full sample (75.9% for Batches 37–43) log a move of some form with the period our data covers. However, we estimate no differences for being assigned a less preferred SG on transfer likelihood across our specifications (Table A5. The estimates are precise in magnitude (ranging between -0.9 to -1.1 percentage points). Over time, it appears that transfer likelihood becomes tied to promotions, which may involve a move up in rank as well as a move of posting within services: transfer likelihood increases in the fifth year after service entry for those assigned their 1st or 3rd preference (relative to lower ranked preferences) (Figure A10).

We also observe no differences of meaningful magnitude nor statistical significance for other career events, which are in general rarer. For assignment of an additional/acting charge, relinquishment of roles, punishment, leave-taking, and exit from service, we detect relatively precisely estimated null effects (Table A7). We do observe a statistically significant decline in the likelihood of being sanctioned or punished (e.g., for corruption or poor performance) for officers assigned their first preference relative to those assigned less-preferred services (Figure A11), consistent with stronger preference-match improving performance.

In comparing the effects across promotions and transfers, it is important to note that transfers can involve moves to new positions within a promotion rank and thus do not correspond one-to-one with promotions. That assignment to a more preferred service translates into a difference in promotion but not transfer likelihood confirms that transfers do not solely derive from promotions (i.e., transfer to a new position upon promotion) and suggests that transfers may be sought in lieu of promotion. We find suggestive evidence for marginal assignment to a lower-ranked preference motivating compensatory movement — i.e., in lieu of promotion — to postings in cities, which are generally thought to be more attractive (Table A6, Column 5).

6.3 Peer Ratings

Table 2 reports 2SLS estimates from equation (7) for measures derived from peer ratings. We apply the strategy from our preferred approach for assessing administrative outcomes: including set and preference fixed effects, the lagged average of the outcome for previous batches in the same service group, and a control for the leave-one-out mean with batch and service group.[22] We first consider an indicator for being rated highly effective, a category to which officers classify 21.3% of their colleagues as belonging. Assignment to a one-rank less-preferred service reduces the likelihood of being rated by peers as highly effects by 7.4 percentage points (34.7%) on average across preference ranks (Column 1). Officers assigned to lower preference-ranks are also judged as less able to withstand political pressure (6 percentage points or 27% more likely, Column 2) and more likely to use their role for personal gain (7.8 percentage points or 70.4% more likely, Column 4). We observe no differences in ratings of sensitivity towards poor citizens (Column 3). Using an index of meausures, rather than indicators, we observe a negative effect of assignment to a less-preferred service on average of approximately 0.1 standard deviations ($p = 0.065$, Column 5).

Examining how these effects vary with the rank of preference assigned in Figure 5, we find first that the increase in effectiveness is present across assignments to 1st through 3rd preference ranks. Those assigned their 3rd preference (vs. lower ranked preferences) are rated better able to withstand political pressure and as being more sensitive to the poor, while we detect no statistically significant differences for the 1st and 2nd preference assignments. Officers assigned 1st or 2nd preference ranks (vs. lower ranked preferences) are judged less likely to use the role for personal gain. The index measure is only statistically significant for those assigned their 2nd preference.

We assess robustness of the effectiveness measure using the stacked regression discontinuity approach as for promotions (Figure 6, Panel B). We detect a difference of 17 percentage points ($p = 0.051$), which is similar if larger than the estimate from our preferred approach.

[22] This last control differs from the one used in Table 1 Columns 4-7, which controls for the leave-one-out mean within "set" and service (a finer category than batch and service). This is because we lack sufficient observations within sets and services in the peer ratings data and thus adopt the rougher grouping of batch X service.

7 Mechanisms

The evidence from peer ratings aligns with the differences in promotion probability: both are consistent with assignment to a more preferred service generating higher performance. We next consider the mechanisms to which these effects can be attributed, which we classify as falling into two primary categories: intrinsic motivation and suitability for role (e.g., in required skills). We also consider the role prestige may play in mediating both and alternative explanations.

7.1 Intrinsic Motivation vs. Suitability

The patterns of results we document using peer ratings already suggest a role for intrinsic motivation. Officers assigned more preferred service groups are rated as more effective and less interested in personal gain. The survey evidence on inputs to preferences described in Section 2.4 highlights interest in tasks as the primary factor candidates consider when selecting their top-ranked preference, though match to skills and prestige also appear important.

To assess the role of instrinsic motivation versus suitability we rely on two additional sources of information. First, we use survey measures relating to job satisfaction collected from respondents during the peer rating exercise. These measures pertain to the rater themselves, and thus are collected for a smaller sample. The survey asked them to rate their satisfaction with their job along multiple dimensions, including whether the job inspires them, generates enthusiasm, and identification with the organizational mission. Table 3 reports correlations between assigned preference rank and these measures. Focusing on officers assigned their 1st through 5th preference ranks — which comprises the majority and for whom we are able to explore exogenous variation in match to preferences — reveals differences in satisfaction, inspiration, and enthusiasm (Panel A). Marginal assignment to a less-preferred service group is associated with lower values across these measures, while we observe negative but not statistically significant relationships for the other dimensions (sense of accomplishment, personal attachment, identification with mission, willingness to recommend or do extra work). Combining these measures into a index, we find that a one-step move to a lower preference rank is associated with a 0.158 standard deviation decrease in motivation related to one's assigned service. Pooling all respondents (assigned 1st through 12th ranks) attenuates these relationships, through the negative correlation with rank assigned and

inspiration remains statistically significant (Panel B).

Second, we estimate impacts using two-stage least squares on measures of performance during training at the Civil Services Academy (CSA) immediately after assignment to service groups. These measures provide a snapshot of performance that is assessed identically across candidates (they are all trained together, the subject matter is identical), and thus isolates the role of motivation versus, e.g., a match to skills. In other words, if higher performance among those assigned to service groups they prefer more corresponds to a match of candidates to their suitability for the tasks involved in working for that service, we should observe no differences in entry-level training assessments because the subjects that officers are evaluated on do not differ by role during training. Rather, the assessments capture performance across a range of subjects thought to be common across service groups, like leadership, public speaking, as well as general knowledge. If instead we observe differences in performance during training, such effects would be consistent with those assigned more preferred services being more motivated to perform well.

During these trainings, assessments on individual subjects and skills are collected by CSA instructors over the six months of training. At the end of training, each trainee bureaucrat receives a cumulative score reflecting a combination of their performance on individual assessments and a rating of their training performance overall by the CSA Director. As in the CSS exam, these cumulative scores are used to assign a unique CSA merit rank to all trainees.

The CSA scores and merit ranks capture short-run performance that are measured before new bureaucrats have any direct exposure to the work of their service group. These outcomes thus are more likely to capture a combination of ability and motivation-related factors, whereas longer-run measures of performance in service — such as annual performance reviews, task-based performance, and promotions — may also reflect the nature of hierarchies within services and personal connections, in addition to ability and motivation. These measures provide an opportunity to assess whether assignment to more preferred services has an immediate impact on performance.[23]

[23]There is also a more mechanical reason why performance in the CSA training matters for career advancement: the CSA merit ranks — as well as CSS exam ranks —- are in fact used as inputs to placement and promotion decisions throughout a bureaucrat's career. When deciding which role to assign new bureaucrats to upon deployment into service, SG officials higher up consider these markers of ability. These measures continue to influence career advancement in two ways: first, initial placement upon deployment — e.g., into a more powerful role — can advantage an officer to be more speedily promoted; second, the CSS and CSA merit ranks continue to be considered as part of a bureaucrats performance "package" when making promotion decisions in later years. Performance in

Table 4 reports 2SLS estimates from equation (7), and compares them to a corresponding ordinary least squares (OLS) estimate, for CSA training score (Panel A) and merit rank (Panel B), respectively. Both the OLS and 2SLS estimates are negative for score and positive for rank (where higher value rank corresponds to worse relative performance) and using our baseline comparison approach, which includes only demographic variables and set fixed effects as controls (Columns 1 and 2). The baseline OLS estimates are larger than the 2SLS estimates, suggesting that preference-rank assignment (D_{ig}) is positively correlated with unobservable (salutory) determinants of performance.

The baseline specification implies that assigning a bureaucrat to an SG ranked one place lower in their preferences reduces CSA merit rank (score) by 4.123 places (1.467 points). Assigning a candidate from their first to fourth preferred SG — the average assigned preference-rank being 4.49 — thus translates into a 13.17% reduction in merit rank.[24] Applying the first strategy by including fixed effects for the top three most preferred SGs being identical among candidates reduces the magnitude of coefficients only slightly, and they remain statistically significant (Columns 3). The second strategy outlined in Section 6.1 — including service group fixed effects, dropping set fixed effects and including independent year, province, and identity group fixed effects, which restricts comparisons within service groups — results in a substantially weaker first stage (with F-statistic of 57, compared to 433.9 in the baseline approach) but an estimate that is closer to the OLS estimate, though only marginally significant. The consistency of these results across approaches suggest a sizable impact of assignment to more (less) preferred SG on performance in CSA training. Though short-run in nature, the presence of these effects identify scope for persistence throughout careers.

Taken together, we find limited support for a suitability mechanism. In fact, given the similarity of work across service groups — as being focused in personnel management — described in Section 2, it is not surprising that there might be limited scope for the match of an officer's skills to those required by a service to generate differences in performance, given limited heterogeneity in required skills. Instead, evidence from our main findings using peer ratings, correlations from survey evidence, and early-stage training performance are all consistent with intrinsic motivation — an interest and val-

CSA training can thus be informative about the treatment's initial impacts on motivation but also form a component of how treatment echoes through bureaucrats' later careers.

[24]The corresponding change in final CSA score is approximately –1%. Though this change in relative terms is small, the merit rank results demonstrate that even small differences in scores can translate into large differences in candidate rankings.

uation of the content of the work involved in serving in a particular service, or even the status attached to it — driving the observed differences in performance we document in promotion speed and peer-rated effectiveness.

7.2 Prestige

Our framework in Section 3 highlights the role prestige may play in attenuating motivations — arising from a closer match between assigned and preferred service group — to exert effort. It is also true that prestige itself may generate intrinsic motivation if officers are inspired to work hard by the status their role itself possesses.

To assess the role of prestige, we subset officers by the relative prestige of their service groups and compare estimates for impacts on our main outcomes — promotion likelihood and peer-rated effectiveness — across these subsamples. The positive effect of assignment to a more preferred service on promotion and effectiveness is larger among those assigned to the prestigious set (i.e., the five most-preferred service groups: Administrative, Police, Foreign, Customs, and Inland Revenue) (Figure 7). For officers assigned to less prestigious service groups, we estimate negative but small and statistically significant effects.

From the perspective of our framework, these results are consistent with prestige either (*i*) intrinsically motivating effort itself, or (*ii*) the status payoffs from assignment to a more preferred role overshadowing the appeal of "coasting on one's laurels" once assigned to a sought-after service. We are unable to distinguish between these explanations, but at minimum our findings are consistent with prestige not "crowding out" but if anything enhancing intrinsic motivations to perform well.

7.3 Alternative Explanations

Preference strength dissipating over time. The strength of preferences may weaken as officers become accustomed to and comfortable in their positions. The pattern of the promotion results, which persist after the fifth year post-service-entry, argue against this interpretation. More qualitative, respondents to the peer rating survey — drawn from all batches in our sample — were also asked to rate the service they would most prefer to work for and the consistency of their preference over time was overwhelming: 49.3% of officers in other service groups rated the Pakistan Administrative Service as being the best to work for.

Unique status of Administrative service. Dropping officers in the most-preferred, most prestigious service — the Pakistan Administrative Service (PAS) — does not materially alter our findings. Figure A12 shows estimates from samples including and excluding PAS officers. While the effect on promotions is attenuated, the impact on the effectiveness rating is magnified when these officers are excluded. For both outcomes we cannot reject that estimates are same across samples.

8 Conclusion

This paper studies the relationship between assignment to a more preferred posting and performance measures within Pakistan's Civil Service, evaluated using administrative records on career progression and peer ratings. We find that assignment to a service group an officer prefers more results in an acceleration of promotion and higher effectiveness in their role, as evaluated by peers. We interpret these impacts as capturing improvements in performance. We find the most support for an intrinsic motivation channel, as opposed to match to skills, driving these effects, with limited scope for prestige to crowd out such motivation.

Our findings contribute to a growing literature on the determinants of bureaucrats' effectiveness, alongside recognition that their performance translates into the performance of state institutions as a whole. Though our evidence is limited to our setting, the prevalence of preference-matching in assignment mechanisms in not only British-legacy public sectors across the developing world but how roles are assigned to workers more broadly, suggests that the preference-to-task match can be consequential for productivity. We view our findings as proof of the existence of such forces in a representative developing country civil service, with an exploration of the role prestige may play in mediating them. The importance of match to preferences in our and other settings can help inform the deployment of complementary tools, like monitoring or targeted performance incentives. Future work can fruitfully shed light on how alternative assignment mechanisms that emphasize or diminish the priority attached to worker preferences shape the performance of public sector institutions as a whole.

References

Aman-Rana, Shan, "Meritocracy in a Bureaucracy," *Journal of Development Economics*, 2025, *175.*

Ashraf, Nava, Oriana Bandiera, Edward Davenport, and Scott S. Lee, "Losing Prosociality in the Quest for Talent? Sorting, Selection and Productivity in the Delivery of Public Service," *American Economic Review*, 2020, *110* (5).

Bandiera, Oriana, Michael Carlos Best, Adnan Qadir Khan, and Andrea Prat, "The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats," *The Quarterly Journal of Economics*, 2021, *136* (4), 2195–2242.

Bergeron, Augustin, Gabriel Tourek, and Jonathan Weigel, "The State Capacity Ceiling on Tax Rates: Evidence from Randomized Tax Abatements in the DRC," *Working Paper*, 2021.

Bertrand, Marianne, Robin Burgess, Arunish Chawla, and Guo Xu, "Glittering Prizes: Career Incentives and Bureaucrat Performance," *Review of Economic Studies*, 2020, *87* (2), 626–55.

Besley, Timothy and Maitreesh Ghatak, "Competition and Incentives with Motivated Agents," *American Economic Review*, 2005, *95* (3), 616–36.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi, "Individuals and Organizations as Sources of State Effectiveness," *Working Paper*, 2019.

Bó, Ernesto Dal, Frederico Finan, Nicholas Y Li, and Laura Schechter, "Government Decentralization under Changing State Capacity: Experimental Evidence from Paraguay," *Econometrica*, 2020.

Burks, Stephen, Jeffrey Carpenter, Lorenzo Goette, and Aldo Rustichini, "Cognitive skills affect economic preferences, strategic behavior, and job attachment," *PNAS*, 2023, *106* (19), 7745–7750.

Callen, Michael, Saad Gulzar, Ali Hasanain, Yasir Khan, and Arman Rezaee, "Personalities and Public Employee Performance: Experimental Evidence from Pakistan," *Working Paper*, 2018.

Cassar, Lea and Stephan Meier, "Nonmonetary Incentives and the Implications of Work as a Source of Meaning," *Journal of Economic Perspectives*, 2018, *32* (3), 215–38.

_ and _ , "Nonmonetary Incentives and the Implications of Work as a Source of Meaning," *Journal of Economic Perspectives*, 2018, *32* (3), 215–38.

Cattaneo, Matias, Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *15* (531).

Clark, Damon and Emilia Del Bono, "The Long-Run Effects of Attending an Elite School: Evidence from the United Kingdom," *American Economic Journal: Applied Economics*, 2016, *8* (1), 150–76.

Conlon, John and Dev Patel, "What Jobs Come to Mind? Stereotypes about Fields of Study," *Working Paper*, 2023.

Duflo, Esther, Rema Hanna, and Stephen P Ryan, "Incentives Work: Getting Teachers to Come to School," *American Economic Review*, 2012, *102* (4), 1241–78.

Falco, Salvatore Di and Ferdinand M Vieider, "Environmental Adaptation of Risk Preferences," *The Economic Journal*, 2023, *132* (648), 2737–2766.

Fenizia, Alessandra, "Managers and Productivity in the Public Sector," *Econometrica*, 2022.

Freyer, Timo, Jonas Radbruch, Sebastian Schaube, and Louis Strang, "The effect of task (mis)matching and self-selection on intrinsic motivation and performance," *Journal of Economic Behavior and Organization*, 2025, *234*.

Hanna, Rema and Shing-Yi Wang, "Dishonesty and Selection into Public Service: Evidence from India," *American Economic Journal: Economic Policy*, 2017, *9* (3), 262–90.

Heffetz, Ori and Robert Frank, "Preferences for Status: Evidence and Economic Implications," *Handbook of Social Economics*, 2008, *1A* (3), 69–91.

Kamei, Kenju and Thomas Markussen, "Free Riding and Workplace Democracy," *Management Science*, 2023, *69* (7), 3884–3904.

Kennedy, Charles, *Bureaucracy in Pakistan*, Oxford University Press, 1987.

Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken, "Tax farming redux: Experimental evidence on performance pay for tax collectors," *The Quarterly Journal of Economics*, 2016, *131* (1), 219–271.

_ , _ , and _ , "Making moves matter: Experimental evidence on incentivizing bureaucrats through performance-based postings," *American Economic Review*, 2019, *109* (1), 237–70.

Khan, Yasir, "Mission Motivation and Public Sector Performance: Experimental Evidence from Pakistan," *Working Paper*, 2023.

Lazear, Edward, "Performance Pay and Productivity," *American Economic Review*, 2000, *90* (5), 1346–1361.

Rasul, Imran and Daniel Rogger, "Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service," *The Economic Journal*, 2018, *128* (608), 413–446.

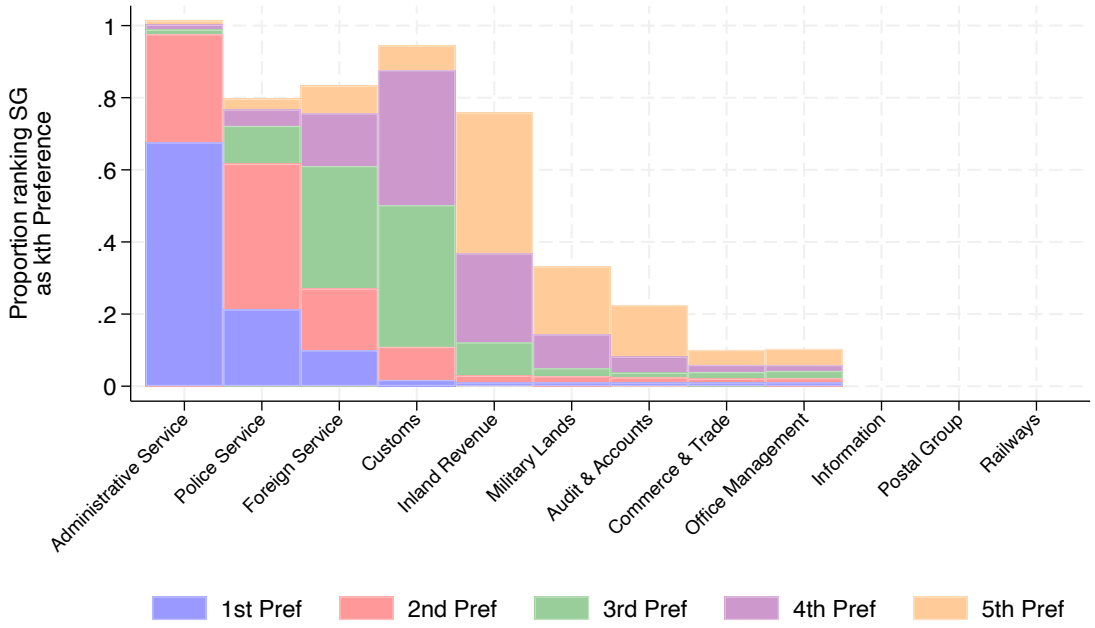
Rose, Evan and Yotam Shem-Tov, "How Does Incarceration Affect Reoffending? Estimating the Dose-Response Function," *Journal of Political Economy*, 2021, *12* (12).

Wilder, Andrew, "The Politics of Civil Service Reform in Pakistan," *Journal of International Affairs*, 2009, *63*, 19–37.

Xu, Guo, Marianne Bertrand, and Robin Burgess, "Hometown Ties and the Quality of Government Monitoring: Evidence from Rotation of Chinese Auditors," *Working Paper*, 2018.

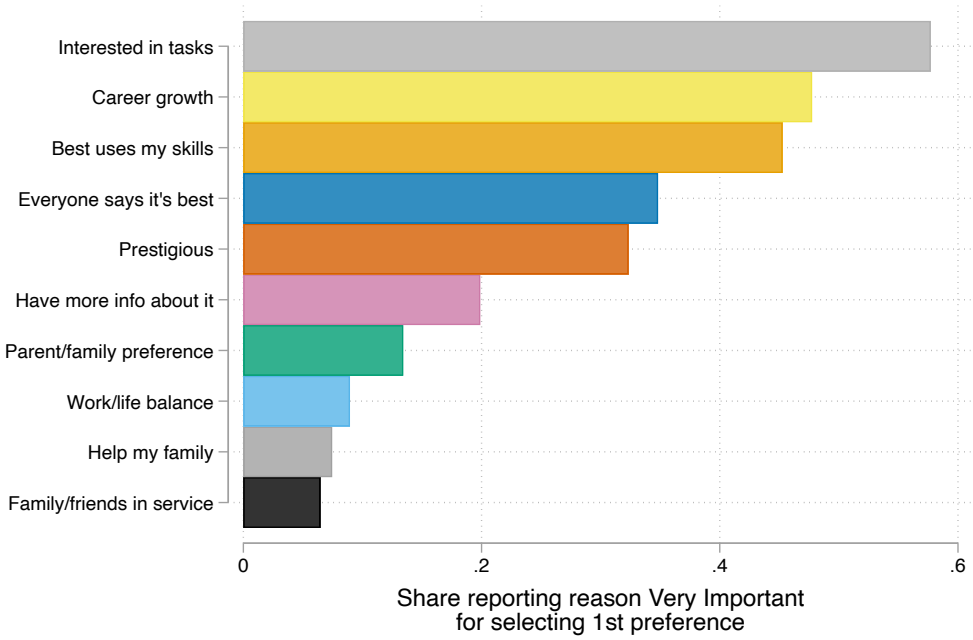
9 Exhibits

Figure 1: Assignment Preferences



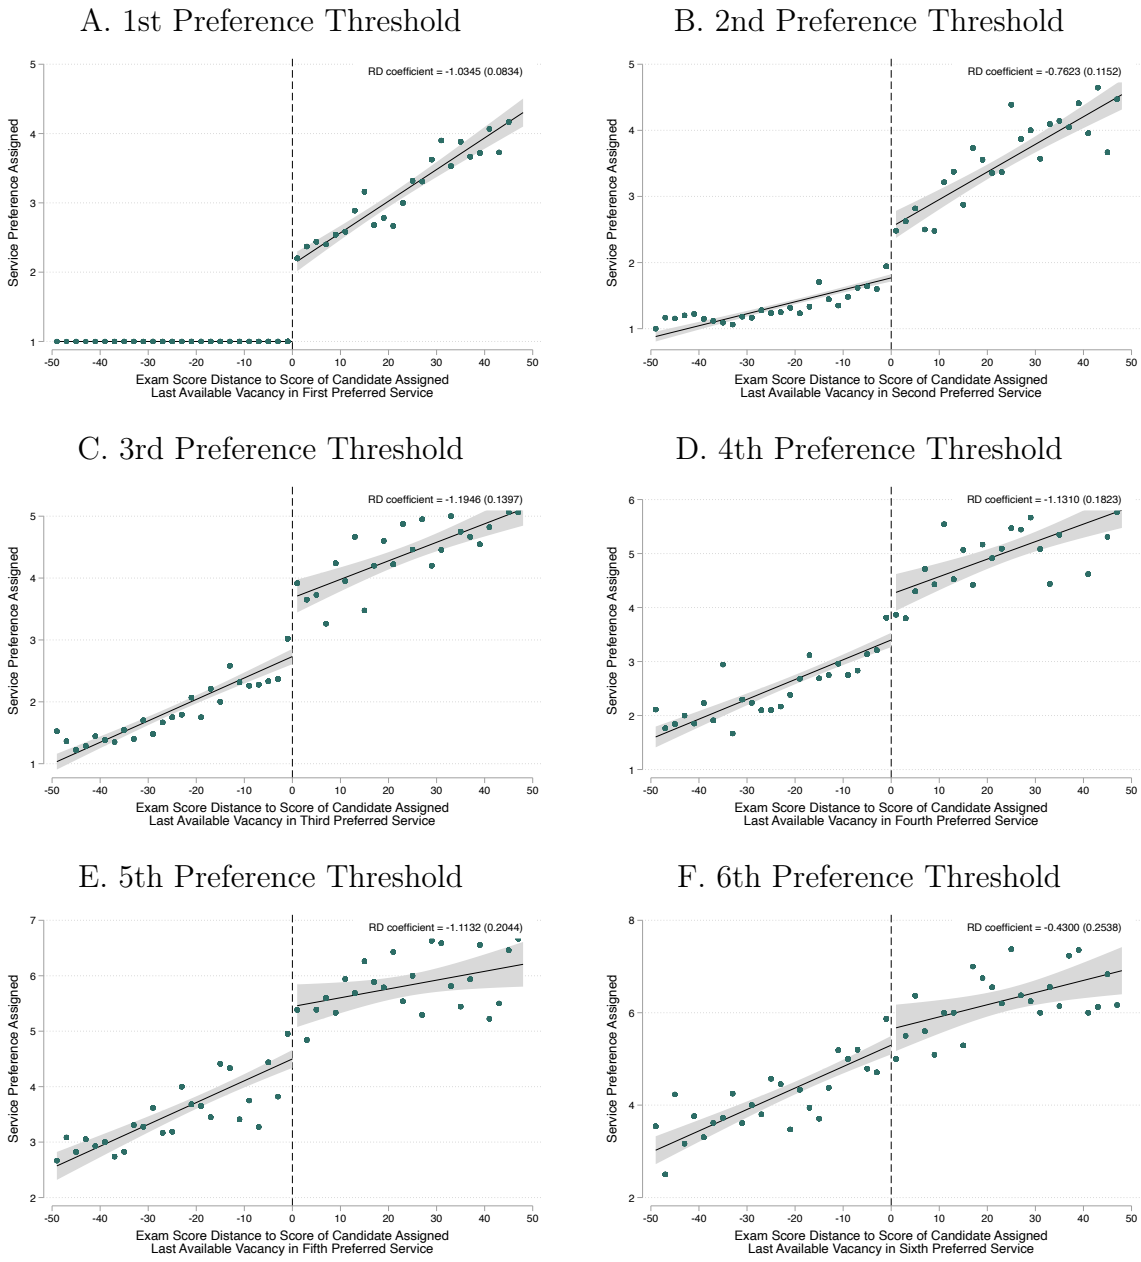
Notes: This figure shows assignment preferences over service groups for the top five preferences listed by candidates during the CSS exam. We discuss this figure in Section 2.

Figure 2: Reasons for Selecting First Preference



Notes: This figure displays the share of candidates reporting a reason as being very important when selecting their first preference. These measures were collected through surveys with candidates, as described in Section 4.

Figure 3: Preference Assigned by Thresholds



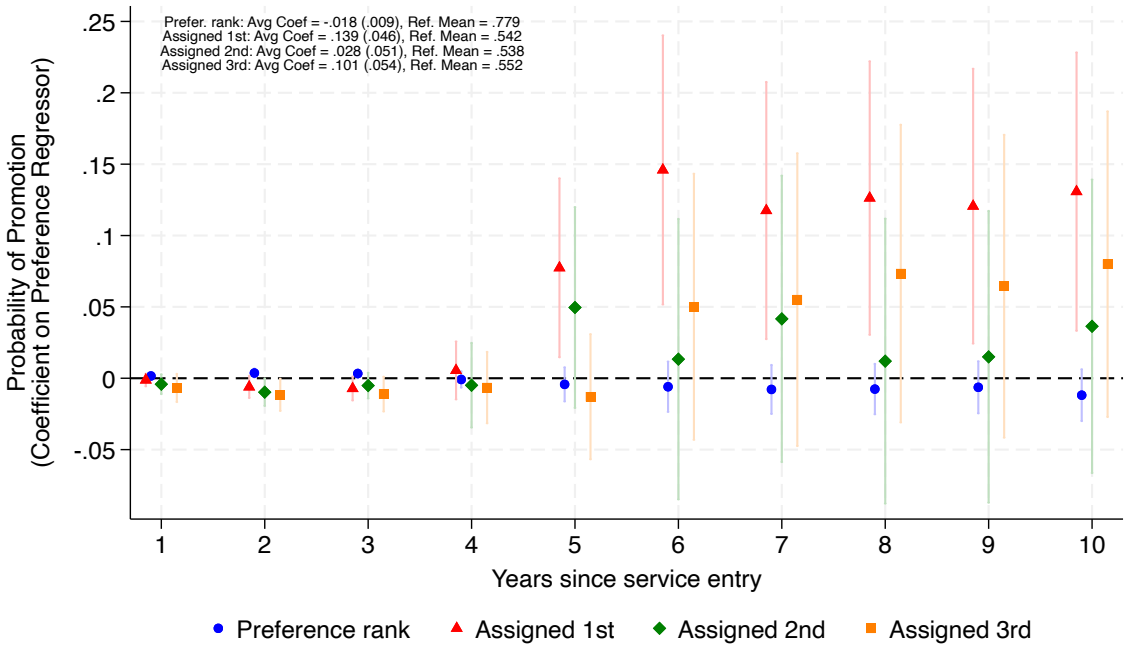
Notes: This figure shows average preference-rank assigned around the preference-rank thresholds described in Section 2. Regression-discontinuity estimates are reported in the top right of each panel. We discuss this figure in Section 5.

Table 1: Treatment Effect on Promotion

	Outcome: Indicator for received a promotion in service grade						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Preference rank assigned	-0.0375***	-0.0365***	-0.0288***	-0.0181**			
	(0.0101)	(0.00994)	(0.00882)	(0.00888)			
Assigned 1st Preference					0.139***		
					(0.0464)		
Assigned 2nd Preference						0.0278	
						(0.0514)	
Assigned 3rd Preference							0.101*
							(0.0537)
Mean among reference group	0.777	0.777	0.777	0.777	0.543	0.537	0.552
Moving pref. 1 to 4 effect in percentages	-14.50	-14.10	-11.11	-6.991			
F-statistic (excluded-instruments)	289.8	292.7	293.5	277.6	16476.0	3478.5	1179.9
N	1259	1259	1259	1259	1259	1021	918
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Preference FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Batch Average Control	No	No	Yes	Yes	Yes	Yes	Yes
Set-Service Leave-One-Out Mean Control	No	No	No	Yes	Yes	Yes	Yes

Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and promotion to a higher service grade (dummy for promotion being reported in administrative gazettes). Column 1 shows estimates from equation (7). Columns 2–7 include fixed effects for candidates' top three preferences over service groups being identical. Columns 3–7 include a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing. Columns 4–7 control for the leave-one-out mean of the outcome calculated within "set" (year by province by identity group) and service group and a dummy for the leave-one-out mean being missing. All columns include set fixed effects and controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage and mean outcome for the reference group: officers' assigned preference rank 1 in Columns 1–4; officers assigned preference rank 2 or lower (Column 5), rank 3 or lower (Column 6), and rank 4 or lower (Column 7). Columns 1–4 report the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.1.

Figure 4: Treatment Effect on Promotion by Years Since Entry



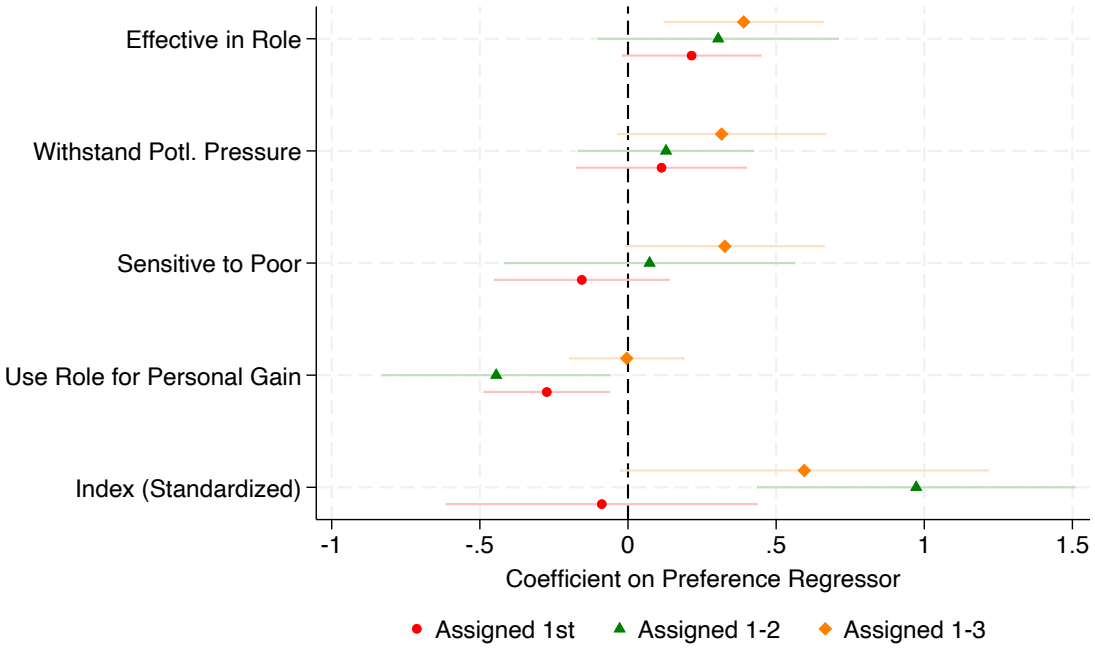
Notes: This figure shows the coefficient estimates by year since service entry from equation (7) for different regressors: preference rank assigned, an indicator for and officer being assigned their first preference, first or second, or first, second, or third preference. The outcome is an indicator for promotion to a higher grade by the year since service entry indicated on the x-axis. Each coefficient is estimated in a separate regression estimating the relationship between the regressor and whether promotion has occurred by a specific year since service entry. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, the average of the year-specific outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, and the leave-one-out mean of the year-specific outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. The text reports the average coefficient standard error, and reference mean (officers assigned 1st preference for the preference rank regressor; officers assigned rank 2 or lower, rank 3 or lower, or rank 4 or lower for assigned 1st, assigned 1-2, and assigned 1-3 regressors, respectively) from a regression pooling all years since service entry with the same controls and fixed effects as in the year-specific regressions. We discuss this figure in Section 6.1

Table 2: Treatments Effects on Peer-Rated Performance

	Outcomes: Indicators for receiving highest rating				
	Effective in Role	Withstand Potl. Pressure	Sensitive to Poor	Use Role for Personal Gain	Index (Standardized)
	(1)	(2)	(3)	(4)	(5)
Preference rank assigned	-0.0740***	-0.0604**	0.00363	0.0781***	-0.0965*
	(0.0227)	(0.0279)	(0.0272)	(0.0223)	(0.0517)
Mean among reference group	0.213	0.224	0.280	0.111	0.0608
Moving pref. 1 to 4 effect in percentages	-104.1	-80.99	3.880	210.9	
F-statistic (excluded instruments)	54.82	29.00	48.70	34.55	56.23
N	368	309	332	286	376

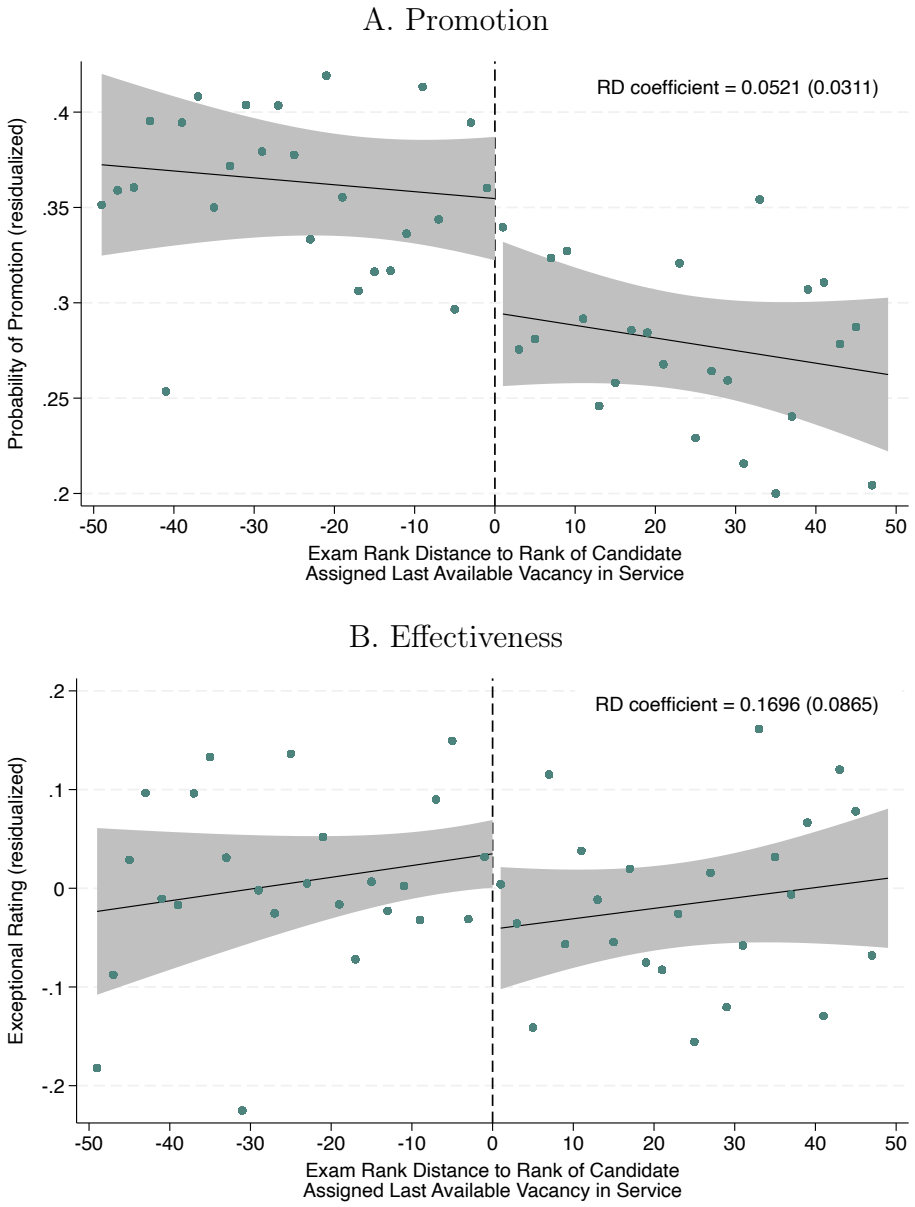
Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and peer-rated measures of officer performance. The outcomes in Columns 1–4 is a dummy for average peer ratings of an officer being higher than the second-lowest category on a 1 to 5 scale, with 5 representing the highest value, for effectiveness in role (Column 1), ability to withstand political pressure (Column 2), sensitivity to the poor (Column 3), and use of role for personal gain (Column 4). The outcome in Column 5 is the standardized index of categories in Columns 1–4, using all values rather than dummies and reversing the personal gain measure such that a higher rating represents lower use of role for personal gain. All regressions include "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, and the leave-one-out mean of the outcome calculated within batch and service group and a dummy for the leave-one-out mean being missing. Robust standard errors are reported. We report the F-statistic for the first-stage, the mean outcome for the reference group (officers' assigned preference rank 1), and in Columns 1–4 the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.3.

Figure 5: Treatment Effects on Peer-Rated Performance by Preference Assigned



Notes: This figure shows the coefficient estimates from equation (7) for different regressors: preference rank assigned, an indicator for and officer being assigned their first preference, first or second, or first, second, or third preference. The outcomes are peer-rated measures of officer performance: indicators for average peer ratings of an officer being higher than the second-lowest category on a 1 to 5 scale, with 5 representing the highest value, for effectiveness in role, ability to withstand political pressure, sensitivity to the poor, and use of role for personal gain. The standardized index of includes the measures of the four rating categories, using all values rather than dummies and reversing the personal gain measure such that a higher rating represents lower use of role for personal gain. Each coefficient is estimated in a separate regression estimating the relationship. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, and the leave-one-out mean of the outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. We discuss this figure in Section 6.3.

Figure 6: Stacked Regression Discontinuity Estimates — Promotion and Peer-Rated Performance



Notes: This figure plots the probability of promotion (Panel A) and a dummy for receiving and exceptional effectiveness rating by peers (Panel B) by exam rank distance to thresholds (exam rank of candidate assigned last available vacancy) determining assignment to a more preferred posting. The outcomes plotted are the residuals from regressing each measure on demographic controls (including dummies for missing values) and the leave-one-out mean of the outcome calculated within batch and service group (and a dummy for the leave-one-out mean being missing). The regression for the promotion residual also includes the average of the year-specific outcome in up to nine previous batch years (as well as a dummy for the previous batch average being missing). Observations around the preference-rank thresholds for the 1st to 5th preferences are stacked. The dark circles represent the mean value bins of 2 exam-ranks. The dark solid lines are linear fits on either side of the threshold, and dark bands are 95 percent confidence intervals. The regression discontinuity coefficient and standard error are estimated from a regression of the residualized outcomes on a dummy for receiving a higher ranked preference-match interacted with the exam-score variables in the stacked data, using a bandwidth of 11 exam-ranks (to accommodate 2-exam-rank bin widths beginning at -1 and 1), including set fixed effects and clustering standard errors at the candidate level. The p-values for the regression coefficients are 0.094 (Panel A) and 0.051 (Panel B). We discuss this figure in Section 6.

Table 3: Correlations Between Preference Rank and Job Satisfaction

Measure:	Satisfied with job	Job inspires	Sense of accomplishment	Personal attachment	Identify with mission	Recommend service group	Enthusiastic about job	Willing to do extra work	Index (standardized)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Assigned Preference Rank 1st to 5th									
Preference rank assigned	-0.0711**	-0.129***	-0.0403	-0.0353	-0.00628	-0.0565	-0.0876**	-0.00812	-0.158**
	(0.0338)	(0.0323)	(0.0337)	(0.0398)	(0.0365)	(0.0394)	(0.0349)	(0.0339)	(0.0648)
Dep. var. mean among pref. 1 receivers	0.913	0.957	0.870	0.739	0.739	0.783	0.870	0.870	0.534
N	75	75	75	74	75	75	75	75	75
B. Assigned Preference Rank 1st to 12th									
Preference rank assigned	-0.0295	-0.0666***	-0.0113	-0.0246	-0.0142	-0.0485*	-0.0297	-0.00294	-0.0706
	(0.0237)	(0.0216)	(0.0205)	(0.0271)	(0.0237)	(0.0251)	(0.0228)	(0.0224)	(0.0507)
Dep. var. mean among pref. 1 receivers	0.913	0.957	0.870	0.739	0.739	0.783	0.870	0.870	0.534
N	93	93	93	92	93	93	93	93	93

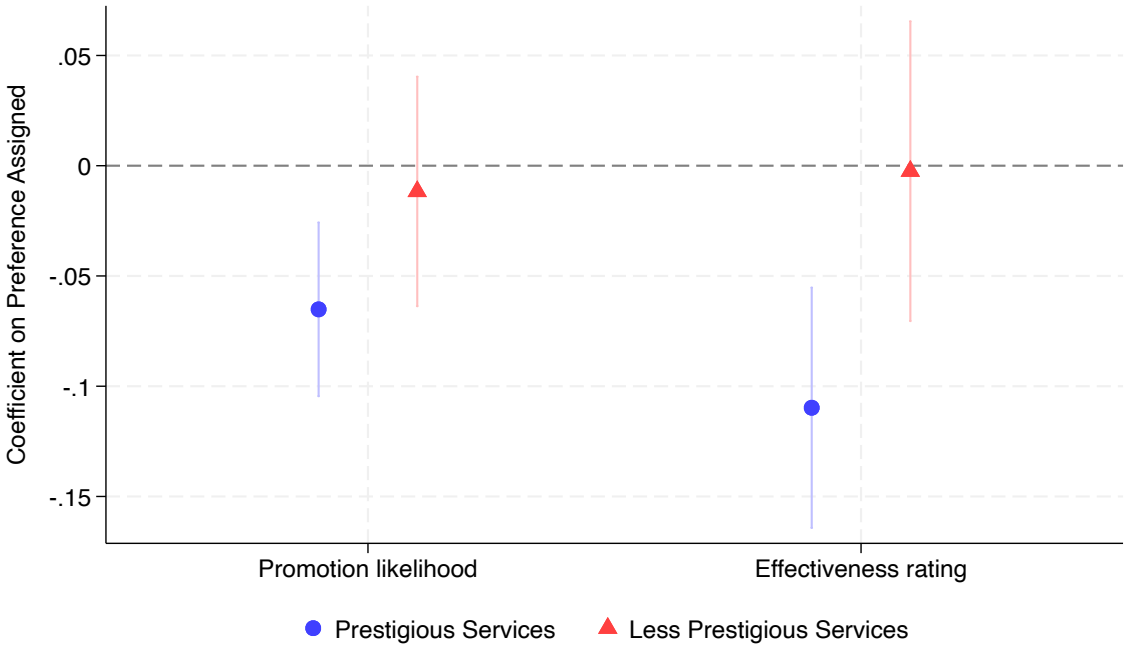
Notes: This table reports correlations between assigned preference rank and measures of job satisfaction collected from respondents to the peer rating survey. Each outcome is an indicator for agreeing or strongly agreeing to a statement asserting sentiments about a respondent's work generating: satisfaction (Column 1), inspiration (Column 2), a sense of accomplishment (Column 3), personal attachment (Column 4), identification with organizational mission (Column 5), willingness to recommend service group (Column 6), enthusiasm (Column 7), and willingness to do extra work beyond what is expected (Column 8). Column 9 considers a standardized index of the measures in Columns 1–8. Panel A considers officers assigned their 1st through 5th preference, and Panel B includes respondents assigned any preference-rank 1st through 12th. All columns include controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the mean outcome for the reference group of officers' assigned preference rank 1. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 7.

Table 4: Treatment Effects on Performance in Training

	(1) OLS	(2) RD	(3) RD	(4) RD
A. Training Exam Score				
Preference rank assigned	-1.976***	-1.467***	-1.408***	-2.078**
	(0.191)	(0.413)	(0.415)	((1.006)
Moving pref. 1 to 4 effect in percentages	-1.674	-1.243	-1.193	-1.761
Dep. var. mean among pref. 1 receivers	354.1	354.1	354.1	59.84
F-statistic (excluded-instruments)		433.2	413.0	56.92
N	2149	1715	1715	1715
B. Training Merit Rank				
Preference rank assigned	5.836***	4.123***	3.912***	5.570*
	(0.599)	(1.328)	(1.344)	(3.166)
Moving pref. 1 to 4 effect in percentages	18.64	13.17	12.50	16.71
Dep. var. mean among pref. 1 receivers	93.92	93.92	93.92	93.92
F-statistic (excluded-instruments)		433.9	413.3	60.1
N	2150	1716	1716	1716
Controls	Yes	Yes	Yes	Yes
Fixed effects	Set	Set	Set,Pref.	Batch, ID, SG

Notes: This table reports the results of estimating equation (7), measuring the relationship between preference assigned — as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference-rank) — and the training exam score (Panel A) and merit rank assigned candidates (Panel B) based on performance during CSA training. The first column shows estimates from and OLS regression of preference-rank assigned on CSA merit rank. Column 2 shows estimates from equation 7. Column 3 includes fixed effects for candidates' top three preferences over service groups being identical. Columns 1–3 include set (year by province by identity group) fixed effects. Column 4 show estimates including service group fixed effects and individual fixed effects for year, province, and identity group. Robust standard errors are reported. We report the F-statistic for the first-stage in Columns 1–3. All columns show the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 7.

Figure 7: Treatment Effect Heterogeneity by Service Prestige



Notes: This figure shows the coefficient estimates from equation (7) where the second-stage regressor is preference rank assigned. The outcomes are promotion likelihood and peer-rated effectiveness (an indicator for average peer ratings of an officer being higher than the second-lowest category on a 1 to 5 scale, with 5 representing the highest value, for effectiveness in role). Each coefficient is estimated in a separate regression estimating the relationship. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, and the leave-one-out mean of the outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. Blue circles are estimates from regressions restricting to the five most prestigious service groups (Administrative, Police, Foreign, Customs, and Inland Revenue). Red triangles are estimates from regressions excluding these service groups. We discuss this figure in Section 7.

Appendix

A1 Additional Exhibits

Figure A1: Civil Service Vacancies, CSS 2020

FEDERAL PUBLIC SERVICE COMMISSION
Aga Khan Road, Sector F-5/1
****

Islamabad, the 5th May, 2021.

PRESS NOTE

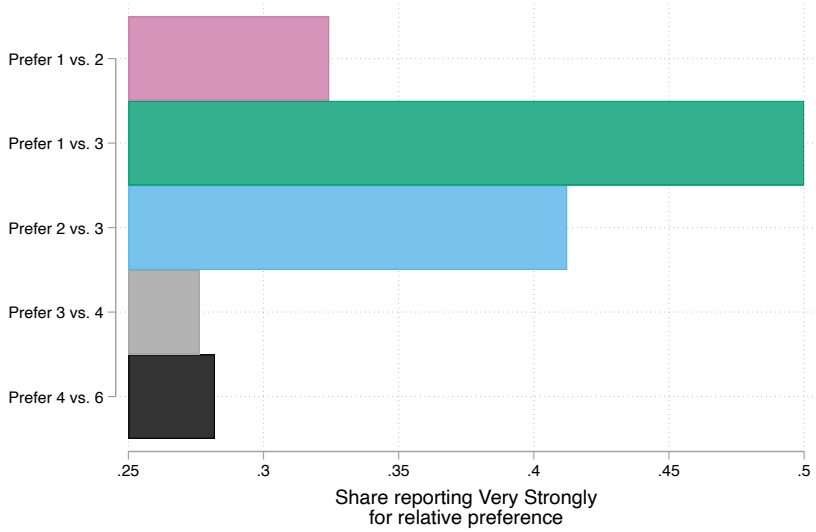
Subject: - **CSS COMPETITIVE EXAMINATION 2020 - NUMBER OF VACANCIES ALLOCATED TO VARIOUS OCCUPATIONAL GROUPS/SERVICES.**

No. F.2/20/2020-CE. In pursuance of Rule 19 of the CSS Competitive Examination Rules 2019 and onwards vide S.R.O No. 1452(I)/2018, number of vacancies, including carried over vacancies, to be filled on the basis of the final result of CSS Competitive Examination 2020 are hereby announced in accordance with Government policy on Merit/ Provincial/ Regional/Women and Minorities quotas. Details are as following: -

QUOTA		FRESH VACANCIES	CARRIED OVER VACANCIES	Total	Grand Total
Merit 7.5%	All Pakistan Merit	PAAS=3, PCS=1, PAS=2, FSP=1, IRS=3, IG=3, OMG=2, PSP=2, POSTG=1	Nil	18+0=18	**18**
Punjab 50%	Open Merit	PAAS=12, PCS=6, PAS=15, FSP=10, IRS=19, IG=13, MLCG=2, OMG=6, PSP=14, POSTG=6, RCTG=1	IG=1	104+1=105	**166**
	Women	PAAS=2, PCS=1, PAS=2, FSP=1, IRS=2, IG=2, OMG=1, PSP=2	Nil	13+0=13	
	Minorities	PAAS=1, FSP=1, IRS=2, IG=1, PSP=1	PAAS=4, CTG=2, PCS=4, IRS=10, IG=2, MLCG=1, OMG=12, PSP=4, POSTG=2, RCTG=1	6+42=48	
Sindh (R) 11.4%	Open Merit	PAAS=4, PAS=3, FSP=1, IRS=5 , IG=4, OMG=1, PSP=4, POSTG=1, RCTG=1	PAAS=1, CTG=3, IG=5, OMG=14, POSTG=2, RCTG=2	24+27=51	**77**
	Women	PAS=1, FSP=1, OMG=1	PAAS=1, CTG=1, IG=1, OMG=5, POSTG=4, RCTG=1	3+13=16	
	Minorities	PCS=1	PAAS=1, CTG=1, IRS=2, IG=1, OMG=2, PSP=1, POSTG=1	1+9=10	

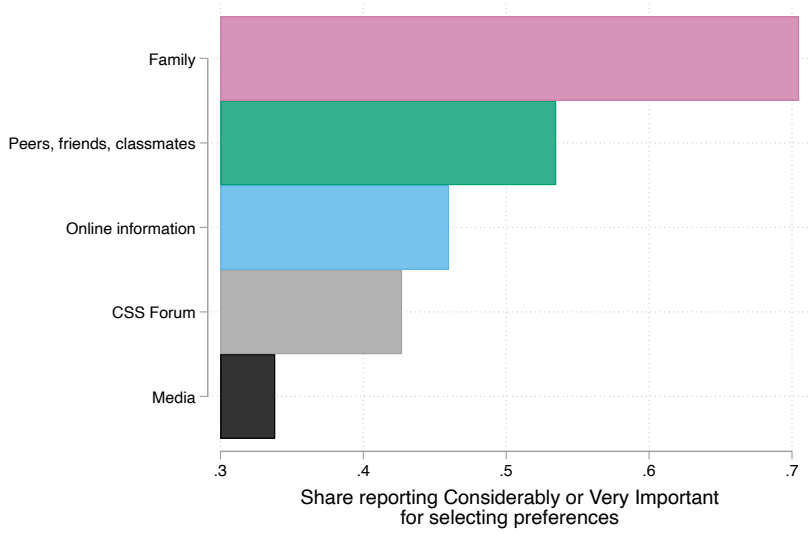
Notes: This figure shows the allocation of vacancies across merit categories, service groups, provinces, and identity groups released by the FPSC for CSS exam-takers in 2020. We discuss this figure in Section 2.

Figure A2: Relative Strength of Preferences



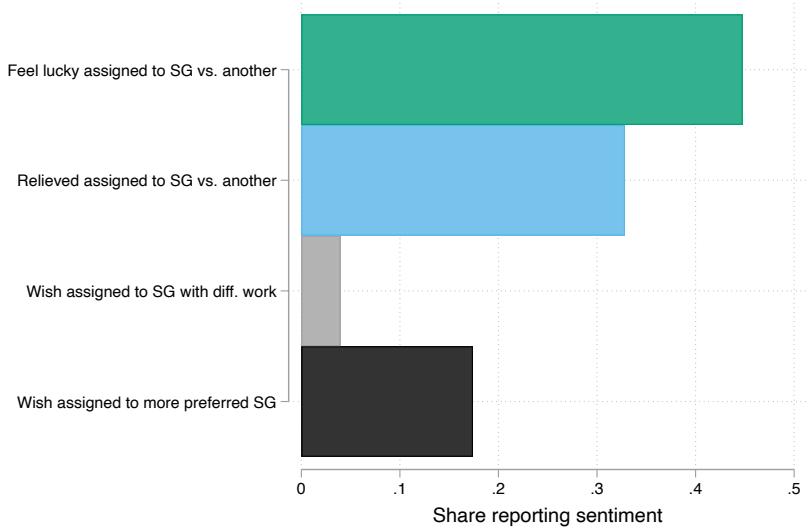
Notes: This figure displays the share of candidates reporting a particular comparison between preference as being very strongly preferred. These measures were collected through surveys with candidates, as described in Section 4.

Figure A3: Information Sources Informing Preferences



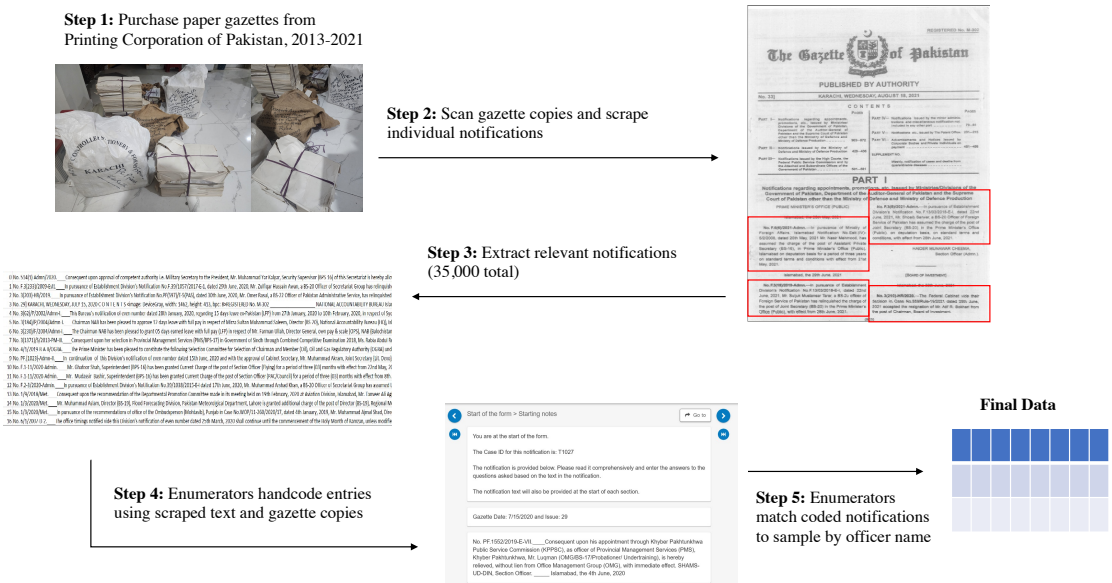
Notes: This figure displays the share of candidates reporting as very important sources for selecting preferences across service groups. These measures were collected through surveys with candidates, as described in Section 4.

Figure A4: Feelings about Assignment



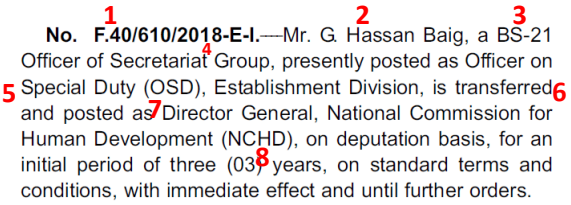
Notes: This figure displays the share of candidates reporting sentiment about their assignment. These measures were collected through surveys with candidates, as described in Section 4.

Figure A5: Data Extraction Process



Notes: This figure describes the process for extracting career outcomes for individual officers from the Gazette of Pakistan. We discuss this figure in Section 4.

Figure A6: Example Entry Coding



1. Notification reference
2. Name of the officer
3. Basic Pay Scale / Grade
4. Current occupation group (OG)
5. Current Designation
6. Key words "transfer", "posted", "appointment"
7. New designation (posting)
8. Transfer/posting duration
9. Date (if applicable)

Basic Information

Serial Number	Notification #	Date (Gazette Date)	Full Name (excl. Title)	BPS/BS Grade	Occupation Group	Designation and/or Department	Type of Notification
12	F.40/610/2018-E-I	15/07/2020	G. Hassan Baig	BS-21	Secretariat Group (SG)	Officer on Special Duty, Establishment Division	Transfers/Posting
							(Select from dropdown menu)

Transfers/Posting

Previous Posting (Service group)	Previous posting (Dept/Org/Ministry/Location)	Previous posting (grade)	New Posting (service group)	New Posting (Dept/Org/Ministry/Location)	New posting (grade)	Length of Posting (if applicable)
Secretariat Group (SG)	Officer on Special Duty, Establishment Division	BS-21	Secretariat Group (SG)	Director General, National Commission for Human Development (NCHD)	BS-21	3 years
(Select from dropdown menu)			(Select from dropdown menu)			

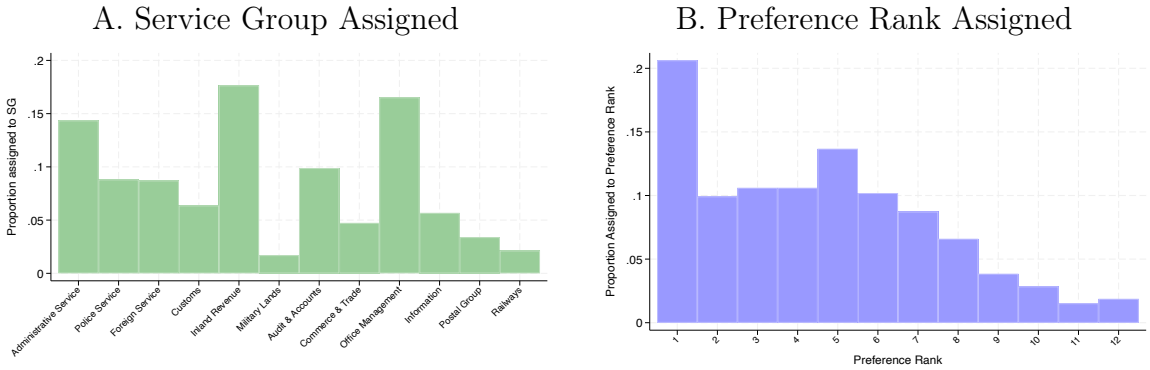
Notes: This figure describes the process for coding of individual entries from the Gazette of Pakistan and the details enumerators recorded. The entry displayed is a transfer/posting. We discuss this figure in Section 4.

Table A1: Candidate Characteristics

Variable	Obs	Mean	SD
Preference assigned	3074	4.49	2.95
Male	2993	0.649	0.478
Age at entry	3194	26.15	2.61
Fathers income (Rs.)	3193	1,304,252	2,924,327
Father living	3137	1.00	0.05
Age father	3126	56.69	15.08
Mother living	3137	0.99	0.09
Age mother	3114	52.33	11.80
Both parents living	3137	0.99	0.09
Eldest child	3194	0.28	0.45
Married	3137	0.24	0.42
Spouse age	740	27.28	5.35
Number of siblings	3137	3.63	2.03
Number of children	3137	0.16	0.52
O-Type blood	3087	0.30	0.46
Computer skills	3056	0.92	0.27
Typing skills	3056	0.77	0.42
Driving skills	3056	0.77	0.42

Notes: This table provides summary statistics for the sample of bureaucrats described in Section **??**. Each measure is drawn from administrative data submitted by applicants during the CSS exam. We discuss this figure in Section 4.

Figure A7: Service Group and Preference Rank Assignments



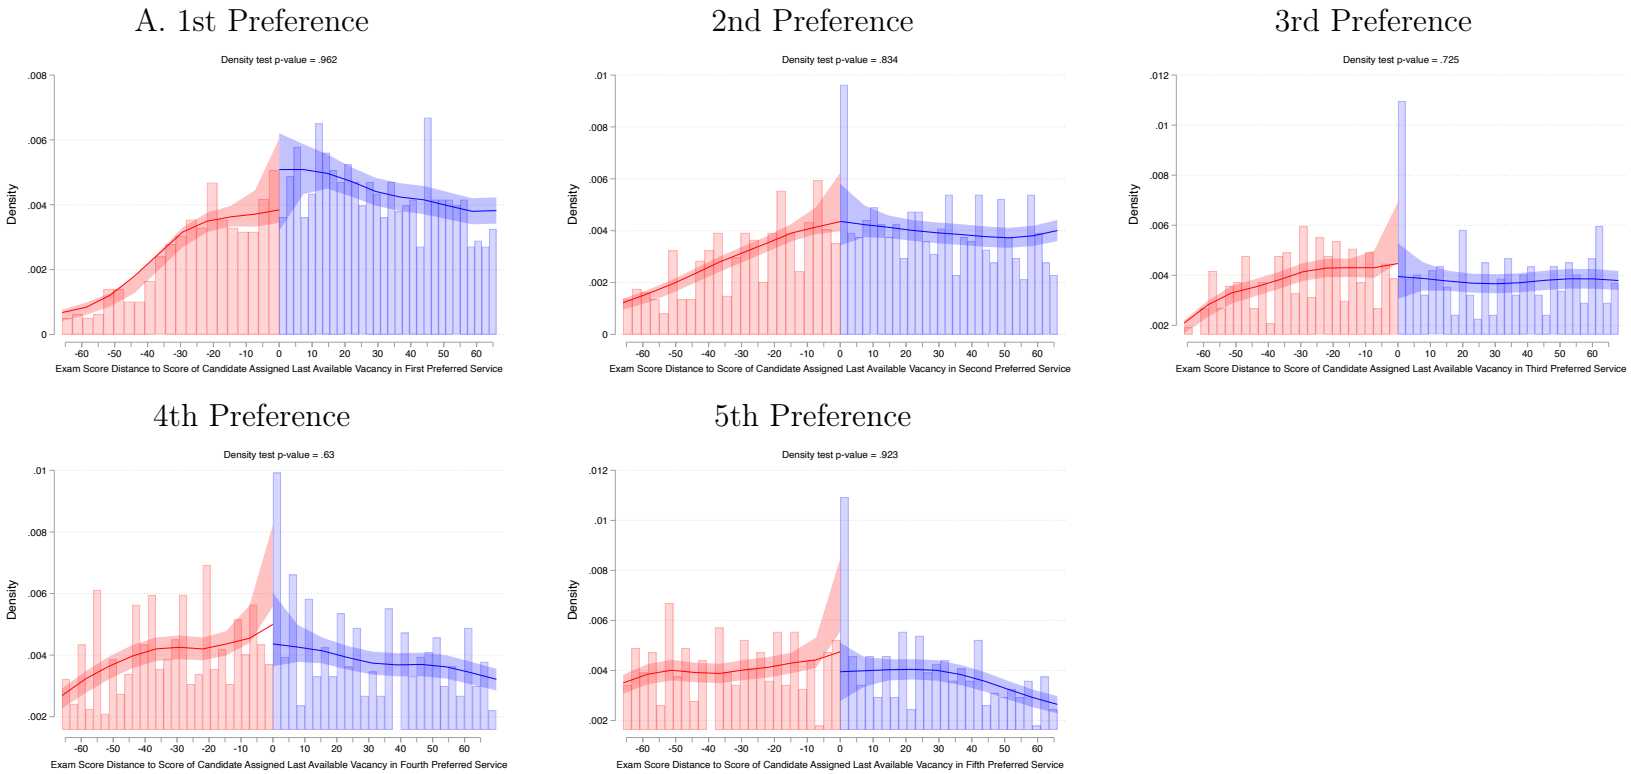
Notes: This figure shows proportions of officers assigned to each service group (Panel A) and preference rank of the assigned service (Panel B). We discuss this figure in Section 2.

Table A2: Balance Tests

Preference Threshold	Age	Father's Income	Log Father's Income	O-Type Blood	Computer Skills	Typing Skills	Driving Skills	Father Living	Age Father	Mother Living	Age Mother	Both Parents Living	Eldest Child	Married	Spouse Age	Number of Siblings	Has Children
First	0.491	0.843	0.884	0.116	0.297	0.486	0.486	0.349	0.294	0.852	0.334	0.586	0.624	0.310	0.806	0.092	0.571
Second	0.389	0.685	0.546	0.055	0.102	0.785	0.785	0.292	0.983	0.235	0.803	0.747	0.535	0.477	0.727	0.060	0.098
Third	0.331	0.665	0.873	0.848	0.577	0.857	0.857	0.350	0.214	0.991	0.460	0.793	0.593	0.765	0.542	0.273	0.909
Fourth	0.507	0.595	0.604	0.281	0.911	0.450	0.450	0.552	0.326	0.120	0.491	0.234	0.900	0.725	0.360	0.876	0.360
Fifth	0.630	0.070	0.025	0.837	0.875	0.460	0.460	0.718	0.765	0.580	0.660	0.646	0.402	0.321	0.691	0.710	0.860

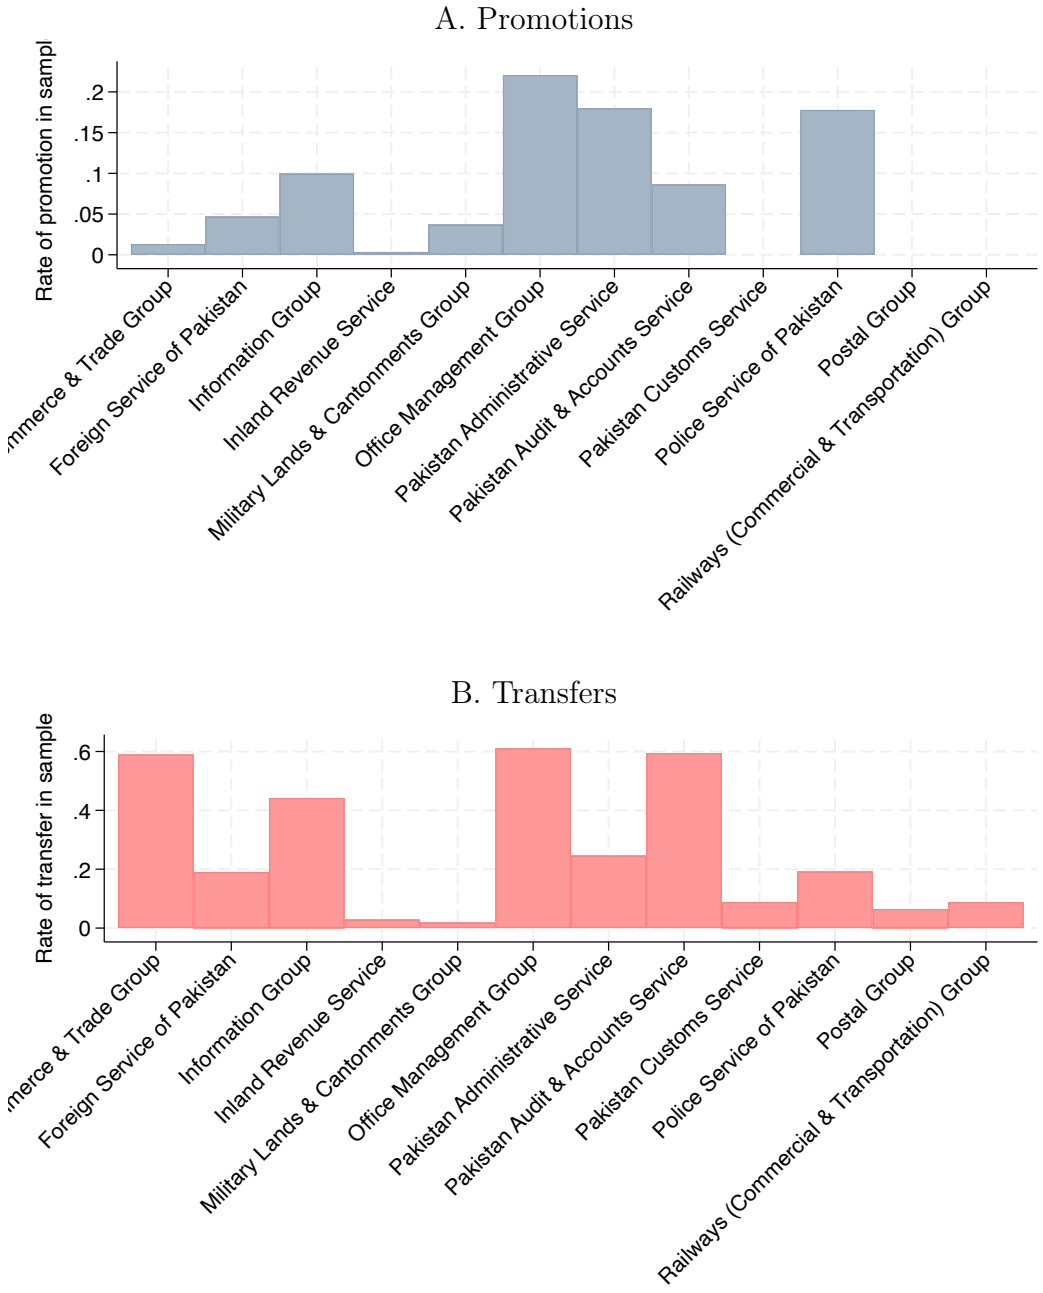
Notes: This table reports the p-values of balance tests for the five preference-rank thresholds used as excluded instruments in equations (6) and (7) for demographic variables summarized in Table A1. We discuss this figure in Section 5.

Figure A8: Density Tests



Notes: This table displays densities on either side of thresholds for the first through third preference-rank thresholds. We report the p-values from manipulation tests using the local polynomial density estimators (Cattaneo et al., 2020) at the top of each panel. We discuss this figure in Section 5.

Figure A9: Promotion and Transfer by Service Group



Notes: This figure reports shows the rate of promotion (Panel A) and transfers (Panel B) by service group for officers in our sample. We discuss this figure in Section 2.

Table A3: Treatment Effect by Promotion Attributes

Outcome:	Years 1 to 5 (1)	Years 6 to 10 (2)	To BS18 (3)	To BS19 (4)	Number Promo. (5)	Time To (Years) (6)	To Director Post (7)	To Secretary Post (8)
Preference rank assigned	-0.00892	-0.0246**	-0.0240**	0.000126	-0.00513	-0.0338	0.00994*	-0.0106**
	(0.00794)	(0.00983)	(0.00939)	(0.00620)	(0.0244)	(0.0428)	(0.00571)	(0.00442)
Mean among reference group	0.217	0.761	0.755	0.0489	1.009	5.957	0.00306	0.0428
Moving pref. 1 to 4 effect in percentages	-12.33	-9.698	-9.516	0.770	-1.525	-1.702	975.5	-73.99
F-statistic (excluded-instruments)	264.2	274.0	292.1	260.2	280.8	212.2	276.6	289.2
N	1259	1259	1259	1259	1259	715	1259	1259
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Preference FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Batch Average Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Set-Service Leave-One-Out Mean Control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and promotion: a dummy for promotion between years 1 to 5 of service (Column 1), between years 6 to 10 of service (Column 2); a dummy for promotion to Basic Pay Scale 18 (Column 3), to Basic Pay Scale 19 (Column 4); number of promotions (Column 5), time in promotion in years (Column 6); a dummy for promotion to a director position (Column 7), and to a secretary position (Column 8). All regressions include include fixed effects for candidates' top three preferences over service groups being identical, a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, a control for the leave-one-out mean of the outcome calculated within "set" (year by province by identity group) and service group and a dummy for the leave-one-out mean being missing, set fixed effects, and controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage, the mean outcome for the reference group: officers' assigned preference rank 1, and the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.1.

Table A4: Robustness for Treatment Effect on Promotion — Service Fixed Effects

	Outcome: Indicator for received a promotion in service grade				
	(1)	(2)	(3)	(4)	(5)
Preference rank assigned	-0.0169***	-0.0181**	-0.0265	-0.0244	0.00631
	(0.00482)	(0.00888)	(0.0276)	(0.0312)	(0.0348)
Mean among reference group	0.777	0.777	0.777	0.777	0.777
Moving pref. 1 to 4 effect in percentages	-6.524	-6.991	-10.23	-9.434	2.437
F-statistic (excluded-instruments)		277.6	24.76	20.64	12.35
N	1547	1259	1259	1259	1259
Individual Controls	Yes	Yes	Yes	Yes	Yes
Preference FE	No	Yes	Yes	Yes	Yes
Alternate FE	Set	Set	SG, Batch, ID, Province	SG X Prov, Batch, ID	SG X Prov X Batch, ID
Lagged Batch Average Control	No	Yes	No	No	No
Set-Service Leave-One-Out Mean Control	No	Yes	No	No	No

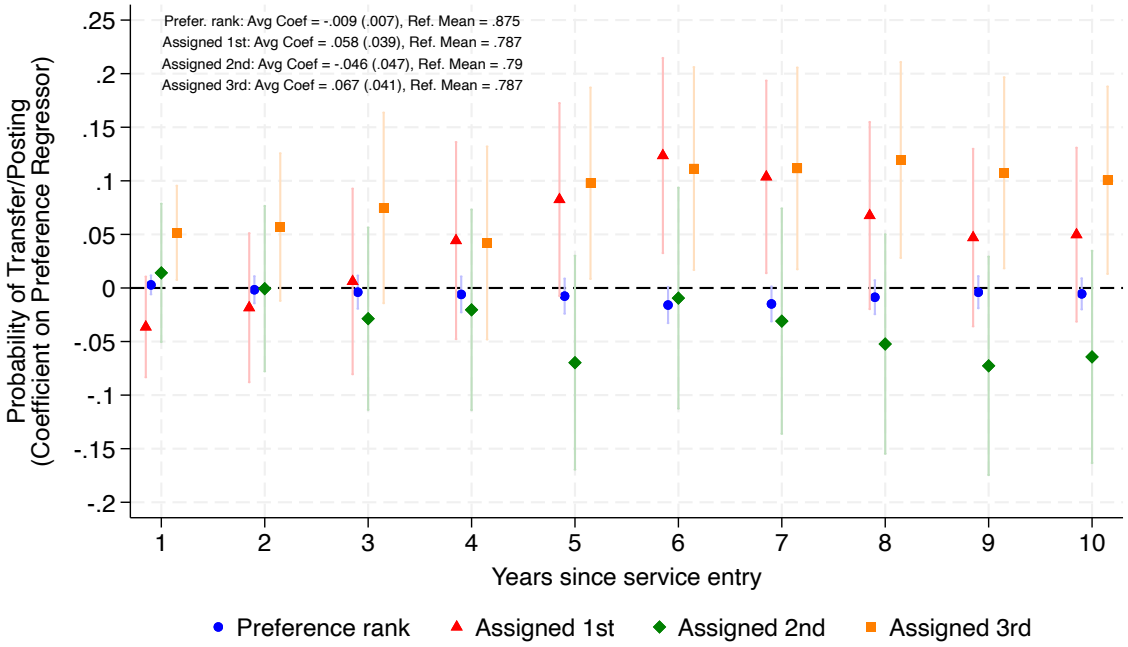
Notes: This table reports the results of estimating equations measuring the relationship between assigned preference rank and assigned preference (Column 1) and as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and promotion to a higher service grade (dummy for promotion being reported in administrative gazettes) (Columns 2–5). Column 1 shows estimates from equation and OLS regression of promotion on preference rank with "set" (year by province by identity group) fixed effects. Columns 2–5 show estimates from equation (**??**). Columns 2–5 include fixed effects for candidates' top three preferences over service groups being identical. Column 2 includes a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, a control for the leave-one-out mean of the outcome calculated within set and service group and a dummy for the leave-one-out mean being missing, and set fixed effects. Column 3 includes fixed effects for service group, batch, identity group, and province. Column 4 includes fixed effects for service group X province, batch, and identity group. Column 5 includes fixed effects for service group X province X batch and identity group. All regressions include controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage and mean outcome for the reference group: officers' assigned preference rank 1. Columns 1–5 of Panel A report the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.1.

Table A5: Treatment Effect on Transfers

	Outcome: Indicator for received a promotion in service grade						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Preference rank assigned	-0.0114	-0.0123	-0.00732	-0.00912			
	(0.00803)	(0.00810)	(0.00719)	(0.00681)			
Assigned 1st Preference					0.0582		
					(0.0391)		
Assigned 2nd Preference						-0.0460	
						(0.0470)	
Assigned 3rd Preference							0.0667
							(0.0414)
Mean among reference group	0.865	0.865	0.865	0.865	0.786	0.789	0.788
Moving pref. 1 to 4 effect in percentages	-3.964	-4.272	-2.538	-3.161			
F-statistic (excluded-instruments)	289.8	292.7	294.9	295.5	17572.4	11713.3	997.0
N	1259	1259	1259	1259	1259	1021	918
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Preference FE	No	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Batch Average Control	No	No	Yes	Yes	Yes	Yes	Yes
Set-Service Leave-One-Out Mean Control	No	No	No	Yes	Yes	Yes	Yes

Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and transfer to another posting within the same service (dummy for transfer being reported in administrative gazettes). Column 1 shows estimates from equation (**??**). Columns 2–7 include fixed effects for candidates' top three preferences over service groups being identical. Columns 3–7 include a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing. Columns 4–7 control for the leave-one-out mean of the outcome calculated within "set" (year by province by identity group) and service group and a dummy for the leave-one-out mean being missing. All columns include set fixed effects and controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage and mean outcome for the reference group: officers' assigned preference rank 1 in Columns 1–4; officers assigned preference rank 2 or lower (Column 5), rank 3 or lower (Column 6), and rank 4 or lower (Column 7). Columns 1–4 report the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.2.

Figure A10: Treatment Effect on Transfers by Years Since Entry



Notes: This figure shows the coefficient estimates by year since service entry from equation (7) for different regressors: preference rank assigned, an indicator for and officer being assigned their first preference, first or second, or first, second, or third preference. The outcome is an indicator for transfer to a different posting within the same service by the year since service entry indicated on the x-axis. Each coefficient is estimated in a separate regression estimating the relationship between the regressor and whether a transfer has occurred by a specific year since service entry. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, the average of the year-specific outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, and the leave-one-out mean of the year-specific outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. The text reports the average coefficient standard error, and reference mean (officers assigned 1st preference for the preference rank regressor; officers assigned rank 2 or lower, rank 3 or lower, or rank 4 or lower for assigned 1st, assigned 1-2, and assigned 1-3 regressors, respectively) from a regression pooling all years since service entry with the same controls and fixed effects as in the year-specific regressions. We discuss this figure in Section 6.2.

Table A6: Treatment Effect by Transfer Attributes

Outcome:	Years 1 to 5 (1)	Years 6 to 10 (2)	Number Promo. (3)	Time to (Years) (4)	To City (5)
Preference rank assigned	-0.0114	-0.00337	-0.130*	-0.00997	0.0203**
	(0.00928)	(0.00773)	(0.0691)	(0.0654)	(0.00916)
Mean among reference group	0.771	0.847	2.505	4.386	0.208
Moving pref. 1 to 4 effect in percentages	-4.420	-1.193	-15.60	-0.682	29.32
F-statistic (excluded-instruments)	264.3	282.1	246.3	190.3	207.1
N	1259	1259	1259	733	1259
Individual Controls	Yes	Yes	Yes	Yes	Yes
Preference FE	Yes	Yes	Yes	Yes	Yes
Lagged Batch Average Control	Yes	Yes	Yes	Yes	Yes
Set-Service Leave-One-Out Mean Control	Yes	Yes	Yes	Yes	Yes

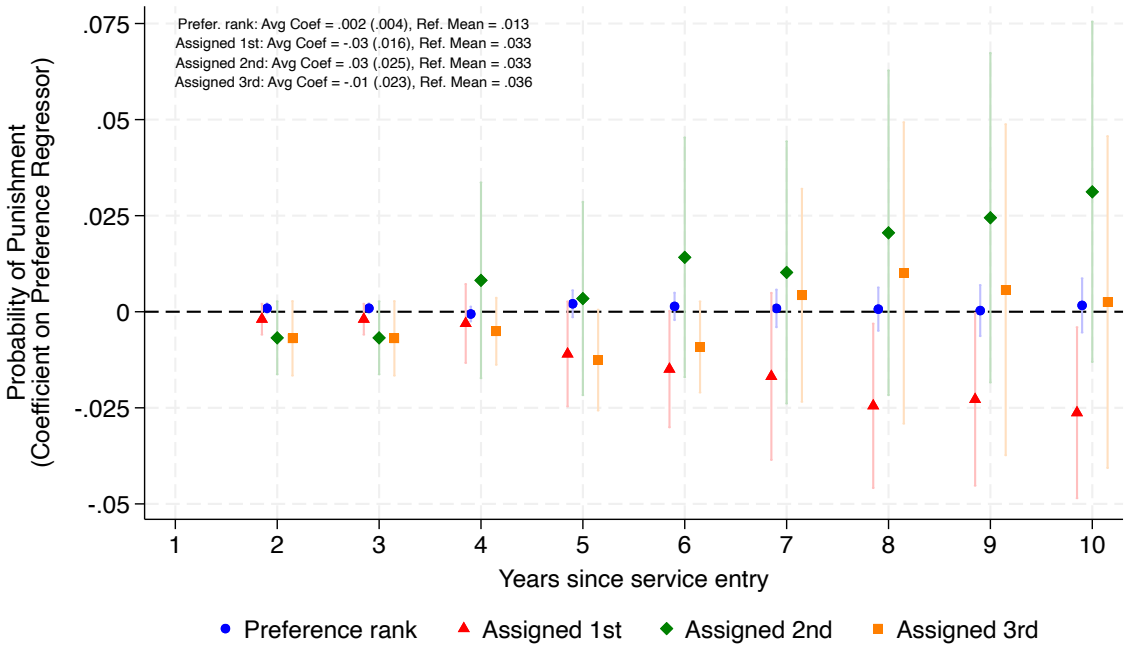
Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and transfers: a dummy for transfer between years 1 to 5 of service (Column 1), between years 6 to 10 of service (Column 2); number of transfers (Column 3), time in transfer in years (Column 4); and a dummy for promotion to a posting in a major city (Column 5). All regressions include include fixed effects for candidates' top three preferences over service groups being identical, a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, a control for the leave-one-out mean of the outcome calculated within "set" (year by province by identity group) and service group and a dummy for the leave-one-out mean being missing, set fixed effects, and controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage, the mean outcome for the reference group: officers' assigned preference rank 1, and the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.2.

Table A7: Treatment Effects on other Administrative Outcomes

Outcome:	Punished	Took Leave	Exited Service	Addtl. Charge	Relinq. Charge	Training Cert.
	(1)	(2)	(3)	(4)	(5)	(6)
Preference rank assigned	0.00146	-0.00995	0.00306	0.00485	0.00584	-0.00217
	(0.00419)	(0.00832)	(0.00232)	(0.00532)	(0.00786)	(0.00352)
Mean among reference group	0.0122	0.135	0.0122	0.0245	0.232	0
Moving pref. 1 to 4 effect in percentages	35.80	-22.18	75.05	59.51	7.540	
F-statistic (excluded-instruments)	293.4	273.5	291.3	276.2	229.2	286.1
N	1259	1259	1259	1259	1259	1259
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes
Preference FE	Yes	Yes	Yes	Yes	Yes	Yes
Lagged Batch Average Control	Yes	Yes	Yes	Yes	Yes	Yes
Set-Service Leave-One-Out Mean Control	Yes	Yes	Yes	Yes	Yes	Yes

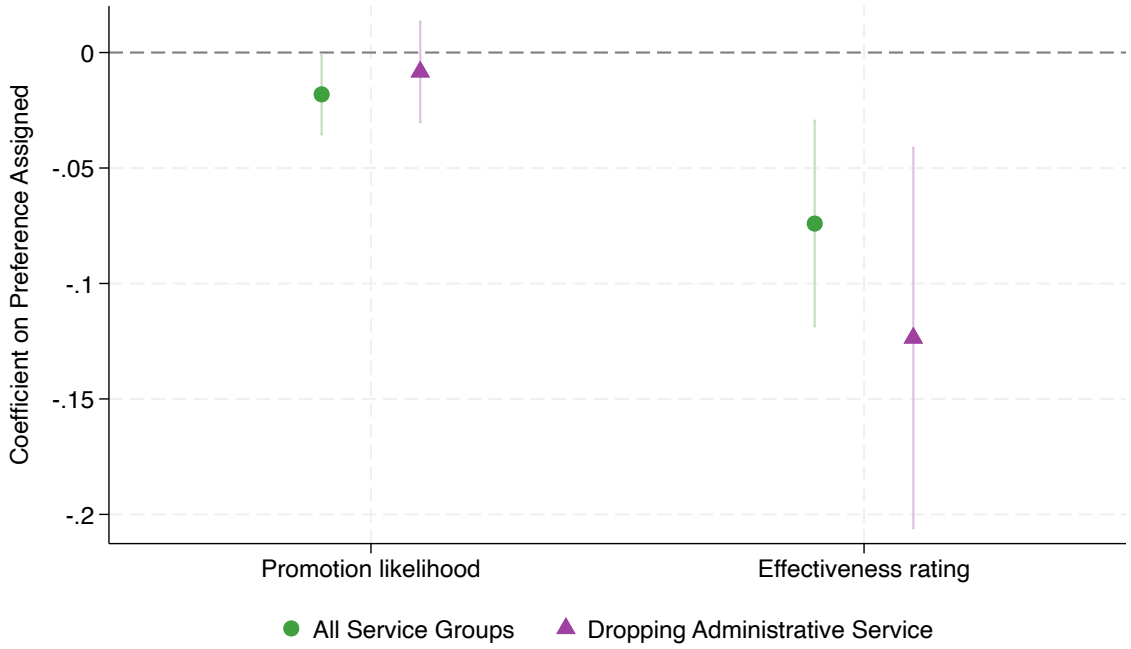
Notes: This table reports the results of estimating equation (7), measuring the relationship between assigned preference rank as a function of the excluded instruments (discontinuities in likelihood of assignment to a preference rank) and other outcomes from the administrative data (gazettes): a dummy for receiving a punishment of any type (Column 1), taking leave from service (Column 2), exiting service (Column 3), receiving and additional charge (Column 4), relinquishing a charge (Column 5), and receiving a course certificate for completing training (Column 6). All regressions include include fixed effects for candidates' top three preferences over service groups being identical, a control for the average of the outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, a control for the leave-one-out mean of the outcome calculated within "set" (year by province by identity group) and service group and a dummy for the leave-one-out mean being missing, set fixed effects, and controls for individual characteristics measured at service entry. Robust standard errors are reported. We report the F-statistic for the first-stage, the mean outcome for the reference group: officers' assigned preference rank 1, and the effect of moving a candidate from their first to fourth preference, which involves multiplying the coefficient by three. ***, **, and * indicate significance at the 1, 5, and 10% levels. We discuss this table in Section 6.2.

Figure A11: Treatment Effect on Punishment by Years Since Entry



Notes: This figure shows the coefficient estimates by year since service entry from equation (7) for different regressors: preference rank assigned, an indicator for and officer being assigned their first preference, first or second, or first, second, or third preference. The outcome is an indicator for punishment of any type within the same service by the year since service entry indicated on the x-axis. Each coefficient is estimated in a separate regression estimating the relationship between the regressor and whether a transfer has occurred by a specific year since service entry. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, the average of the year-specific outcome in up to nine previous batch years as well as a dummy for the previous batch average being missing, and the leave-one-out mean of the year-specific outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. The text reports the average coefficient standard error, and reference mean (officers assigned 1st preference for the preference rank regressor; officers assigned rank 2 or lower, rank 3 or lower, or rank 4 or lower for assigned 1st, assigned 1-2, and assigned 1-3 regressors, respectively) from a regression pooling all years since service entry with the same controls and fixed effects as in the year-specific regressions. We discuss this figure in Section 6.2.

Figure A12: Robustness to Dropping Administrative Service



Notes: This figure shows the coefficient estimates from equation (7) where the second-stage regressor is preference rank assigned. The outcomes are promotion likelihood and peer-rated effectiveness (an indicator for average peer ratings of an officer being higher than the second-lowest category on a 1 to 5 scale, with 5 representing the highest value, for effectiveness in role). Each coefficient is estimated in a separate regression estimating the relationship. Each regression includes "set" (year by province by identity group) fixed effects, fixed effects for candidates' top three preferences over service groups being identical, controls for individual characteristics measured at service entry, and the leave-one-out mean of the outcome calculated within set and service group and a dummy for the leave-one-out mean being missing. Green circles are estimates from regressions including all service groups. Purple triangles are estimates from regressions excluding officers assigned to the Pakistan Administrative Service. We discuss this figure in Section 7.

A2 Conceptual Framework: Appendix

In this section we describe the assumptions of the conceptual framework discussed in Section 3 and provide proofs of the framework's propositions.

Assumption 1. For $\Pi(\cdot)$ to be a valid probability at the optimal effort levels characterized below, parameters must satisfy the following.

Parameters $(\pi, w_2, \theta_0, \theta_1, \mu_0, \mu_1)$ are such that

$$\Pi(p_1(e^*_{a1}, a), a) = \pi \left[\theta_a + \pi \left(w_2 + \frac{\theta_a^2}{2} \right) + \mu_a \right] \in [0, 1] \quad \text{for } a = 0, 1,$$

where e^*_{a1} is the optimal choice of period-1 effort characterized in Proposition 1 below.

Given $\pi, \theta_a, w_2 \geq 0$ and μ_a not too negative, the lower bound $\Pi \geq 0$ will hold. The binding restriction is the upper bound, $\pi \left[\theta_a + \pi(w_2 + \theta_a^2/2) + \mu_a \right] \leq 1$, and since the

bracketed term is increasing in θ_a, it is sufficient to verify this holds at $a = 1$. This forms a joint restriction on the promotion-probability slope π, the senior-level wage w_2, and the levels of θ_a and μ_a; however, this restriction does not imply a bound of the form $\theta_a < 1$. For example, $\theta_a < 1$ for both types does not bound Π if π or w_2 is large, and conversely $\Pi \in [0, 1]$ can hold even when $\theta_a > 1$ if π is small.

Proof of Proposition 1. We begin with backward induction. At the senior level ($t = 2$), the bureaucrat will maximize:

$$\max_{e_{a2}} w_2 + \theta_a e_{a2} - \frac{1}{2} e_{a2}^2 \tag{8}$$

subject to $e_{a2} \geq 0$. FOC: $\theta_a - e_{a2} = 0$, and $e_{a2}^* = \theta_a$, or purely a function of intrinsic payoffs. In period 1, the bureaucrat will choose e_{a1} to maximize V in equation (1). With $\frac{\partial \Pi}{\partial e_{a1}} = \pi$, the FOC for is the following:

$$e_{a1}^* - \theta_a = \pi(w_2 + \theta_a^2 - \frac{\theta_a^2}{2})$$

$$e_{a1}^* = \theta_a + \pi(w_2 + \theta_a^2 - \frac{\theta_a^2}{2}) = \theta_a + \pi(w_2 + \frac{\theta_a^2}{2})$$

Then:

$$\Delta p_1 \equiv p_1(e_{11}^*, 1) - p_1(e_{01}^*, 0) = \theta_1 - \theta_0 + \pi \frac{\theta_1^2 - \theta_0^2}{2} + (\mu_1 - \mu_0)$$

We define $\tilde{\theta}_a \equiv \theta_a + \pi \frac{\theta_a^2}{2}$.

To interpret $\tilde{\theta}_a$, note that $e_{a1}^* = \tilde{\theta}_a + \pi w_2$: since w_2 is common to both types, with $\tilde{\theta}_a$ the type-specific component of period-1 effort, net of the baseline career-concerns effort that even a bureaucrat with $\theta_a = 0$ would exert simply to try to capture the wage w_2. $\tilde{\theta}_a$ has two parts. The first, θ_a, is the static (or myopic) level of motivated effort. This is identical to e_{a2}^*, i.e., what a bureaucrat with no career concerns at all ($\pi = 0$) would choose through just intrinsic motivation. The second, $\pi\theta_a^2/2$, is a dynamic amplification: $\theta_a^2/2$ is the net surplus, above the wage w_2, that a bureaucrat of type a obtains at the senior level by optimally choosing $e_{a2}^* = \theta_a$. $\pi\theta_a^2/2$ is the senior-level surplus weighted by the marginal effect of first period effort combined with the probability of reaching the second period. Bureaucrats with higher θ_a work harder in period 1 not only because they enjoy effort now, but because promotion is worth more to them once they get there. This reinforcement between intrinsic motivation and career concerns makes $\tilde{\theta}_a$ convex in θ_a ($\partial^2 \tilde{\theta}_a / \partial \theta_a^2 = \pi > 0$) rather than tracking θ_a one-for-one, a property we return to in Proposition 2.[25] Then:

[25] Formally, $\tilde{\theta}_a$ is increasing in the intrinsic payoff, $\partial \tilde{\theta}_a / \partial \theta_a = 1 + \pi\theta_a > 0$, and in the strength of career concerns, $\partial \tilde{\theta}_a / \partial \pi = \theta_a^2/2 > 0$. It is also decreasing in the marginal cost of effort, though this is not visible in equation (9) because we have normalized the cost coefficient to $c = 1$. Reintroducing a general convex cost $C(e_{at}) = ce_{at}^2/2$ gives $e_{a2}^* = \theta_a/c$ and $\tilde{\theta}_a(c) = \theta_a/c + \pi\theta_a^2/(2c^2)$, for which $\partial \tilde{\theta}_a / \partial c = -\theta_a/c^2 - \pi\theta_a^2/c^3 < 0$, confirming the claim once c is allowed to vary.

$$\Delta p_1 \equiv p_1(e_{11}^*, 1) - p_1(e_{01}^*, 0) = (\tilde{\theta}_1 - \tilde{\theta}_0) + (\mu_1 - \mu_0). \tag{9}$$

Proof of Proposition 2. At the senior level, $\psi(r)$ enters additively, so it does not affect the second period first-order condition: $e_{a2}^*(r) = \theta_a - \delta r$, as in Proposition 1 with θ_a replaced by $\theta_a - \delta r$. The value of this choice is $w_2 + \psi(r) + (\theta_a - \delta r)^2/2$. In period 1, with $\partial\Pi/\partial e_{a1} = \pi$, the first-order condition is:

$$e_{a1}^*(r) - (\theta_a - \delta r) = \pi \left(w_2 + \psi(r) + \frac{(\theta_a - \delta r)^2}{2} \right),$$

so that

$$e_{a1}^*(r) = (\theta_a - \delta r) + \pi \left(w_2 + \psi(r) + \frac{(\theta_a - \delta r)^2}{2} \right).$$

Case (i) is straightforward: with a common r for both types, $\psi(r)$ and πw_2 are identical across types and cancel in the comparison, leaving the same sign as in Proposition 1.

For case (ii), writing $x_1 \equiv \theta_1 - \delta r_H$ and $x_0 \equiv \theta_0 - \delta r_L$ (both strictly positive under the restriction on r above):

$$\begin{aligned} e_{11}^*(r_H) - e_{01}^*(r_L) &= (x_1 - x_0) + \frac{\pi}{2}(x_1^2 - x_0^2) + \pi\big[\psi(r_H) - \psi(r_L)\big] \\ &= (x_1 - x_0)\left[1 + \frac{\pi}{2}(x_1 + x_0)\right] + \pi\big[\psi(r_H) - \psi(r_L)\big], \end{aligned}$$

which gives equation (4). This is strictly negative if and only if $(x_1 - x_0)\big[1 + \frac{\pi}{2}(x_1 + x_0)\big] < -\pi[\psi(r_H) - \psi(r_L)]$; dividing through by the strictly positive bracketed term and rearranging gives condition (5), proving (ii).

Incorporating ability into Proposition 2. It is worth noting how Proposition 2 would read if it were stated directly in terms of performance $p_1 = e_{a1} + \mu_a$, as in Proposition 1, rather than in terms of effort alone. Adding ability back in gives:

$$\begin{aligned} \Delta p_1 \quad &\equiv p_1(e_{11}^*(r_H), 1) - p_1(e_{01}^*(r_L), 0) \\ &= [(\theta_1 - \theta_0) - \delta(r_H - r_L)]\left[1 + \tfrac{\pi}{2}\big((\theta_1 - \delta r_H) + (\theta_0 - \delta r_L)\big)\right] \\ &\quad + \pi\big[\psi(r_H) - \psi(r_L)\big] + (\mu_1 - \mu_0), \end{aligned}$$

so the performance-based analog of condition (5) (the condition under which Type I under-performs Type II, not only under-exerts effort) is

$$\delta(r_H - r_L) > (\theta_1 - \theta_0) + \frac{\pi\big[\psi(r_H) - \psi(r_L)\big] + (\mu_1 - \mu_0)}{1 + \frac{\pi}{2}\big((\theta_1 - \delta r_H) + (\theta_0 - \delta r_L)\big)}. \tag{10}$$

If Type I bureaucrats also have an ability advantage ($\mu_1 > \mu_0$), condition (10) is harder to satisfy than condition (5). Prestige-driven crowd-out would need to overcome the preference-match gap, the persistent-rent gap, and Type I's ability advantage before performance, not just effort, would be higher for Type II. If instead $\mu_1 < \mu_0$, condition

(10) is easier to satisfy, and Type I could under-perform Type II even with a smaller prestige gap than effort alone would require. Under our identifying assumption — i.e., that preference-to-posting match is orthogonal to ability at entry, so $\mu_1 = \mu_0$ in expectation — condition (10) collapses back to condition (5), and the effort- and performance-based comparisons are identical. We state Proposition 2 itself in terms of effort because it isolates the crowding-out and career-concerns mechanisms; however, condition (10) is the empirically testable object once ability differences are allowed for.

Type-specific δ's. We can also allow δ to vary by type, δ_a: in this case, the left-hand side of condition (5) becomes $\delta_1 r_H - \delta_0 r_L$ in place of $\delta(r_H - r_L)$. Type I under-performs Type II whenever the crowd-out actually experienced at the prestigious posting, $\delta_1 r_H$, exceeds the crowd-out avoided at the less prestigious one, $\delta_0 r_L$, by more than the preference-match gap $\theta_1 - \theta_0$ and the gap in expected second period status rents $\pi(\psi(r_H) - \psi(r_L))/B$. If $\delta_1 > \delta_0$ — i.e., the most intrinsically motivated bureaucrats are more susceptible to displacement by extrinsic recognition — the crossover condition is easier to satisfy than under the common-δ case, since Type I's motivation is crowded out disproportionately. If instead $\delta_1 < \delta_0$, Type I's preference-match advantage is comparatively insulated from crowd-out, and the condition is harder to satisfy than the common-δ benchmark. The common-δ case is thus an intermediate case rather than a claim that susceptibility to crowd-out is necessarily type-independent.