

The Long-Run Costs of Highly Competitive Exams for Government Jobs

Kunal Mangal*

November 30, 2022

Abstract

Public sector recruitment exams can be highly competitive. Does this competition encourage candidates to develop skills that are useful even outside of the public sector, or do the years spent preparing for the exam impose long-run costs on candidates who fail to get selected? To address this question, I study the impact of a partial public sector hiring freeze in the state of Tamil Nadu, India between 2001 and 2006 on male college graduates. The hiring freeze caused the remaining recruitments to become even more competitive. Candidates responded by spending less time employed, most likely in order to invest more time in exam preparation. A decade after the hiring freeze ended, the affected cohorts have lower employment rates, have delayed forming their own households, and appear to have a lower earning capacity. Elder members of the household compensate by delaying retirement. Together, these results suggest that highly competitive exams result in large numbers of unsuccessful candidates making investments that are ultimately unproductive.

*Contact: kunal.mangal@apu.edu.in; Azim Premji University. This paper was previously circulated as “Chasing Government Jobs: How Aggregate Labor Supply Responds to Public Sector Hiring Policy in India” and “Competitive Exams for Government Jobs and the Labor Supply of College Graduates in India.” I am grateful to my advisors Emily Breza, Asim Khwaja, and Rohini Pande for their support for the project. Rob Townsend provided much appreciated initial encouragement. I also thank Augustin Bergeron, Shweta Bhogale, Michael Boozer, Christina Brown, Deepti Goel, Nikita Kohli, Tauhidur Rahman, Sagar Saxena, Utkarsh Saxena, Niharika Singh, Perdrie Stilwell, Nikhil Sudharsanan, and seminar participants at Harvard and Azim Premji University for thoughtful discussions and comments. This work would not have been possible without the support of K. Nanthakumar, R. Sudhan, and S. Nagarajan of the Tamil Nadu Government, and the staff at the R&D Section of TNPSC. I am also grateful to the many candidates for government jobs who were willing to take the time to share their world with me. Of course, any errors are my own.

1 Introduction

Close to 90% of countries around the world use merit-based exams to select public sector workers (Teorell, Dahlström, and Dahlberg, 2012). Competition for these posts can help create a more professional, better-qualified bureaucracy (Moreira and Pérez, 2022; Colonnelli, Prem, and Teso, 2020). However, in contexts where public sector jobs are substantially more valuable than their private sector counterparts, this competition can reach extreme levels.¹ When the exams are so heavily over-subscribed, candidates can spend years unemployed studying for the exams full time (Mangal, 2021).

Do candidates benefit from participating in these selection exams in the long run? It is possible that exam preparation is an investment in general human capital that has value outside of the public sector. On the other hand, the time spent out of employment, and the disappointment of not getting selected, may have long-term economic and social scarring effects. The net effect of the exam process on candidates who participate in them remains an open empirical question.

To address this question, I study the socio-economic impacts of a policy that increased the competitiveness of public sector recruitment exams.² Between 2001 and 2006, the government of Tamil Nadu, India implemented an unexpected partial hiring freeze. The policy resulted in an 86% drop in vacancies in the sectors of the government that were affected by the policy, but left aggregate demand relatively intact.³ I first study how the shock affected the competitiveness of the remaining posts. I then track how cohorts that were more exposed to the increased competition fared in the long run.

My analysis draws on data from nationally-representative household surveys, gov-

¹For example, in India (Muralidharan, 2015), China (Yu, 2020), Brazil (Simons, 2016), and Southern Europe (Geromichalos and Kospentaris, 2020), selection rates in these exams are often 1% or less.

²Exam preparation may have value outside of the domains I study. For example, it may inculcate a sense of citizenship, improve candidates' knowledge of their rights, or help candidates demand services from their local government. Unfortunately, in this paper I am unable to assess these more subtle potential impacts.

³In Appendix B, I estimate that there were only about 552 vacancies lost per cohort, which is about 0.05% of the cohort size.

ernment reports that I digitized, and administrative data from the Tamil Nadu Public Service Commission (TNPSC), the agency responsible for recruiting the posts that were affected by the hiring freeze. I present three main findings.

First, I show that the affected exams became even more competitive during the hiring freeze. Application levels increased by about 7%, as a result of which these recruitments became 185% more competitive. Competition levels returned close to normal after the freeze was lifted.

How did keeping up with the increased competition during the hiring freeze impact candidates? I estimate causal effects using a difference-in-differences (DiD) design, comparing Tamil Nadu to other states in India and leveraging variation in exposure to the hiring freeze across cohorts. I focus my analysis on male college graduates, both because college graduates are more likely to be affected by the freeze (which improves statistical power), and because low female labor force participation makes it difficult to infer the value of exam preparation from women’s labor market outcomes.⁴

My second main finding is that male college graduates—particularly recent graduates, who were the most exposed to the hiring freeze—responded by decreasing their labor supply. The time they spent employed in the early years of their career fell by 13%. Instead, they switched to unemployment or dropped out of the labor force to enroll in post-graduate education—both of which are symptomatic of spending longer on the “exam track,” either by actively studying for the next recruitment or by waiting for its arrival. Individuals in Tamil Nadu who were not eligible for the exam remained unaffected, and it is unlikely this effect is explained by a shift in aggregate labor demand.

The time spent out of employment could either be an investment with future returns, or the cost of competition. My third finding is that, a decade after the hiring freeze ended, the most affected cohorts show signs of long-term social and economic scarring. These cohorts of men have lower employment rates, have delayed forming their own households, and appear to contribute less to household consumption, suggesting reduced

⁴I will infer whether exams helped develop general human capital by whether earnings outside of public sector employment increased. But when labor force participation is low, as it is for women, then one also has to account for selection into the labor force.

earning capacity. To compensate for the financial shock, elder members of the household delay their retirement.

Throughout the paper, I draw extensively on the seminal ethnographic work of Jeffrey (2010), which provides a detailed account of the lives of candidates studying for government recruitment exams in Uttar Pradesh, India. I rely on Jeffrey’s work to both interpret some of the main results, and to understand their import in richer detail.

This paper contributes to the literature in several ways. First, I provide some of the first empirical evidence linking youth under-employment and the competition for public sector jobs. Economists have long worried that the prospect of a lucrative government job diverts people away from productive activity and towards unproductive exam preparation instead (Krueger, 1974; Muralidharan, 2015; Banerjee and Duflo, 2019). This paper provides causal evidence supporting these concerns.

Second, I contribute to a literature on the social and economic scarring of early career shocks. Past work has shown that entering the labor force during a recession can result in a persistent large, negative effect on individuals’ future career trajectories (Kahn, 2010; Oreopoulos, Von Wachter, and Heisz, 2012). More recent work suggests that these costs extend to household formation, including decreased rates of marriage and increased childlessness (Schwandt and Von Wachter, 2020). This paper shows how pausing recruitment in even a relatively small number of government jobs when competition is very high can generate recession-like effects. This is perhaps because candidates’ willingness to focus exclusively on a specific type of public sector employment meant that the hiring freeze resulted in an individual experience of labor market slack akin to an economy-wide recession.

Finally, I contribute to a literature on optimal public sector recruitment. Previous work has focused largely on how public sector hiring policy influences the productivity of the candidates who are ultimately selected (see Finan, Olken, and Pande (2017) for an overview of this literature). By this logic, an increase in competition is good because it provides the government with more opportunity to select better candidates. Indeed, previous work has highlighted how the quality of the applicant pool increases when public

sector jobs become more valuable (Dal Bó, Finan, and Rossi, 2013). However, the findings in this paper suggest that these potential benefits need to be balanced against the social costs of the selection process itself.

2 The Hiring Freeze

2.1 The Impact on Recruitment

Between November 2001 and July 2006, the Government of Tamil Nadu suspended most recruitment for “non-essential posts” (TN Government Order 212/2001; TN Government Order 91/2006). Doctors, police constabulary, and teachers were explicitly exempted from the hiring freeze (these will be known as the *exempted* posts).

Among the impacted posts, the onus of the freeze fell almost entirely on the unspecialized administrative posts (such as “section officer” or “junior assistant”) that populate the state’s civil service. These posts, which are collectively known as “Group recruitments,” represented 80% of all vacancies and about 93% of all applications among the impacted posts at the time of the freeze.

All the posts impacted by the hiring freeze are recruited by a government agency known as the Tamil Nadu Public Service Commission (TNPSC). I will make extensive use of data from TNPSC reports and administrative files for this analysis.

The impetus for the hiring freeze was a state financial crisis, triggered by a set of pay raises for government employees that were implemented in the late 1990s (The World Bank, 2004). Other states also experienced fiscal crises around the same time, but to the best of my knowledge they did not implement a hiring freeze.⁵ I therefore use the set of states excluding Tamil Nadu as a control group in the empirical analysis. I test the sensitivity of the results to the choice of states included in the control group. To the extent that other states also implemented hiring freezes at the same time, I expect the estimated effects to be attenuated.

The length of the hiring freeze was originally left open-ended.⁶ Ultimately the policy

⁵In India, states set independent recruitment policy and do not coordinate with each other.

⁶Government officials who were working in TNPSC at the time of the freeze told me even they did

was rescinded in 2006, months after the incumbent government was ousted in the 2006 elections. The uncertainty of the length of the hiring freeze has important implications for how candidates would have responded to it, which I discuss in more detail in Section 2.3.

In Figure 1, I confirm that the hiring freeze did have a large effect on the recruitment intensity of impacted posts. The average number of vacancies notified dropped by about 86% during the hiring freeze.⁷ The number of recruitments fell from an average of 37 *per year* to a total of 9 throughout the duration of the hiring freeze. After the hiring freeze was lifted, the government continued to conduct far fewer recruitments for impacted posts, but vacancy levels returned to a level even slightly higher than they were at before the hiring freeze began. Meanwhile, consistent with the letter of the policy, recruitment in the exempted posts remained unaffected (Appendix Figure A.1).

The number of vacancies that were abolished due to the freeze was also small relative to the overall size of the labor force. A back-of-the-envelope calculation suggests that even for the most exposed cohorts, the hiring freeze affected at most 0.6% of the cohort's final occupation attainment. Even accounting for the large wage premium, the drop in average earnings due to the aggregate demand shock is on the order of 0.4% of cohort average earnings. (See Appendix B for the details of these calculations). The direct demand effect of the hiring freeze (i.e. the reduction in labor demand due to less government hiring) is thus negligible.

2.2 Who was impacted by the policy?

In principle, all Indians are eligible to apply for government jobs in Tamil Nadu. However, in practice TNPSC receives very few applications from outside of Tamil Nadu. Even as recently as 2019, when inter-state mobility is likely much higher than in the early 2000s, less than 1% of applicants in the Group recruitments were from states other than Tamil Nadu.

not know how long was going to last.

⁷There were an average of 302 vacancies advertised per fiscal year during the hiring freeze, compared to 2,109 before the freeze.

Government jobs in India have several common eligibility requirements. For entry level posts, candidates must be at least 18 years old, and must have at least a 10th standard education. Higher level posts will require a college degree, or potentially a degree in a specific field. These eligibility requirements imply that individuals without a 10th standard education should not be directly affected by the hiring freeze, at least in the short run. I will use this fact in my empirical strategy when I study the impact of the hiring freeze on candidates.

There is substantial variation in application rates by age and completed education levels. To illustrate, in Figure 2 I plot application rates for men for Group recruitments conducted in FY 2013, the earliest fiscal year for which I have such data.

Note that application rates are highest among recent college graduates. Even if the levels have shifted since the early 2000s, I assume that the relative magnitudes of the application rates are similar at the time of the hiring freeze. It therefore seems reasonable to consider early college graduates to be the cohorts most exposed to the hiring freeze.

2.3 How might the hiring freeze have affected candidates?

The hiring freeze represented a large shock to the value of exam preparation. However, its impact on candidates' exam preparation decisions is ambiguous. Here, I provide intuition for why this may be the case.

On the one hand, the drop in vacancies reduced the current returns to exam preparation. Candidates who were studying full time could therefore choose to abandon or postpone their studies and take up a job in the private sector instead.

On the other hand, the drop in vacancies may not have reduced the surplus of exam preparation by enough to justify a switch to the outside option. Candidates would enjoy surplus from exam preparation if they believed they were infra-marginal to selection at the usual vacancy level. Thus, they could choose to respond to the hiring freeze by studying more instead.

Moreover, the hiring freeze may have *increased* the continuation value of exam preparation. Note that the hiring freeze not only reduced the number of vacancies, it also

reduced the frequency of exams. If candidates learn about their ability from past performance, and they tend to have over-optimistic beliefs about their chance of selection, then a reduction in the speed in which candidates learn about their ability (and hence slow down the speed at which they revise their beliefs downwards) may encourage them to persist longer.

An important consideration is that candidates were making these decisions not knowing how long the hiring freeze would last. If the length of the hiring freeze were pre-announced, then candidates could temporarily suspend their studies and return to them in time to remain competitive in the next exam. However, since the length of the hiring freeze could be announced at any time, candidates who suspended their studies ran the risk of becoming uncompetitiveness when normal hiring levels resumed. Studying continuously can be thought of as a strategy for preserving the option value of remaining competitive when the hiring freeze ended. This value of this option increases when it becomes harder to catch up to other candidates after taking a break in one's studies.

2.4 The Impact on Application Behavior

Even before the hiring freeze, the impacted posts were already quite competitive (Appendix Table A.1), with average selection rates often less than 1%. In this section, I show that during the freeze the competition for the remaining vacancies *increased*.

Data. TNPSC publishes an annual report that lists the notifications that were published during the fiscal year. I digitized this data from the 1992/93 fiscal year to the 2010/11 fiscal year.⁸ For each recruitment, I observe the date of the notification, the post type, the number of vacancies notified, and the number of applications received.

I restrict the sample to recruitments for posts that: i) fell within the purview of the hiring freeze; and ii) were still notified at some point during the hiring freeze. This yields a sample of 57 recruitments: 32 that were notified before the hiring freeze, the 9 exceptions that were notified during the freeze, and 16 that were notified after the hiring

⁸These reports are available online at <https://tnpsc.gov.in/English/AnnualReports.aspx>. The table that I use is located in Annexure IV.

freeze ended.

Regression Model. My main outcomes are : i) number of vacancies offered; and 2) number of applications received. I estimate how these outcomes changed over time while accounting for variation in the types of posts offered using a Poisson regression model:⁹

$$\ln E[y_i | \mathbf{X}_i] = \alpha_{p(i)} + \beta_1 \text{Freeze}_i + \beta_2 \text{After}_i + \Gamma' Z_i \quad (1)$$

where i indexes recruitments, and $p(i)$ indexes the post type. The variable Freeze_i is a dummy for whether the notification date occurred while the hiring freeze was still in effect, and After_i is a dummy for whether the notification date occurred after the freeze was lifted. Z_i includes any specified additional controls. The reference group is the set of posts that were notified before the freeze. To account for possible model mis-specification, I report [White \(1982\)](#) robust standard errors.

Results. The more vacancies that the government offers, the more applications it receives (Appendix Figure [A.2](#)). In this particular sub-sample of posts that were notified both before and during the freeze, the government offered 75% fewer vacancies (i.e. $\exp(-1.36) - 1$) during the freeze compared to before the freeze (Column 1). Given this, we would expect application rates to fall. Instead, application volume during the freeze increased by about 7.5%, which, given the large standard error (s.e. = 52.6 p.p), is a small effect relative to the usual variation in application volume (Column 2).¹⁰

If application volume remained relatively steady while vacancies plummeted, then competition for the remaining vacancies must have increased substantially. Including a control for vacancies in a regression on application volume confirms this intuition (Column 3). Posts that were advertised during the freeze attracted 185% more applications (i.e. $\exp(1.049) - 1$) than the same posts advertising similar number of vacancies did before

⁹A key advantage of the Poisson regression model is that the coefficients can be readily interpreted as a percent effect, i.e. $\beta_j = \left(\partial E[y|x_j] / E[y|x] \right) / \partial x_j$. For a log-linear model estimated via OLS, this is typically true only as long as the errors are homoskedastic ([Silva and Tenreyro, 2006](#)).

¹⁰Since I observe the universe of recruitments conducted during this time period, the standard error does not reflect uncertainty about the effect size itself conditional on this sample.

the hiring freeze. After the freeze ended, the level of competition largely returned to normal.

3 Short-Run Impacts on Candidates

If the government of Tamil Nadu were just like any other employer, we should expect the effect of the hiring freeze on exposed cohorts to be imperceptible, since the drop in hiring is small relative to the overall churn in the market. However, because government jobs are so coveted, the hiring freeze could bend labor supply. In this section, I measure this impact.

3.1 Data

I use data from the National Sample Survey (NSS). The NSS is a nationally-representative household survey conducted by the Government of India. I use all rounds of the NSS conducted between 1993/94 and 2007/2008 (i.e. between the 50th and the 64th rounds) that include a question on employment status.¹¹ By stacking these individual rounds, I obtain a data set of repeated cross-sections. I adjust all estimates using the provided sampling weights.¹²

My main outcome variable is employment status, which takes one of three values: employed, unemployed, or out of labor force. These variables reflect what the NSS terms the household member’s “Usual Principal Status,” i.e. the status in which the individual spent the majority of their time over the past year.¹³ I further divide the out of labor force individuals into those who are enrolled in educational institutions and those who are out of the labor force for other reasons. This is because enrollment in higher education is a common refuge for individuals looking for a supportive environment while they apply

¹¹See Appendix Table A.2 for a summary of the specific rounds used in the analysis.

¹²To ensure that the scale of the weights is consistent across rounds, I normalize them as follows: if w_{ir} are NSS-provided weights for individual i in round r , and there are N_r observations in round r , then the weights I use are: $N_r * w_{ir} / \sum_r w_{ir}$.

¹³In accordance with the NSS’s definitions, individuals are marked as unemployed if they were “available” for work but not working. Note that this definition does not require the individual to be actively searching.

for government jobs (Jeffrey, 2010).

3.2 Empirical Strategy

Sample Construction. My main analysis sample consists of male college graduates who were between the ages of 17 to 35 in 2001, excluding individuals living in Union Territories.¹⁴ The lower age limit is based on the time it usually takes for individuals to complete an undergraduate degree. In India, undergraduate programs typically last at least three years. This implies that individuals would have needed to be at least 17 years in 2001 in order to enter the labor force with a college degree during the hiring freeze. The upper age limit is arbitrary; in Section 3.3 I verify that the results are not sensitive to this choice.

DiD Design. I identify the causal impact of the hiring freeze using a difference-in-differences design that compares Tamil Nadu with the rest of India, and uses variation in exposure to the hiring freeze by cohort.

Variation in exposure to the hiring freeze is generated by a combination of the timing of the freeze and baseline application rates. Everyone who was eligible—including all college graduates—were potentially affected by the hiring freeze. However, the larger the share of the cohort that was intending to apply, the higher the exposure should be. As we saw in Figure 2, application rates are highest among recent college graduates, and decline with age. Individuals who were expected to graduate from college during the hiring freeze should therefore be the most affected. Those who had already graduated college before the the start of the freeze should be less affected, since many of the potential applicants would have already exited exam preparation.

I measure the impact of the freeze by comparing the effected cohorts with older cohorts whose outcomes were measured before the hiring freeze was implemented.

¹⁴I exclude Union Territories from the analysis for two reasons. First, Union Territories are administrative regions that do not have their own state recruitment agencies. Individuals who live in Union Territories typically apply for state-level posts elsewhere. Second, for reasons I outline below, I cluster standard errors at the state-by-cohort level, and I expect the coverage rate of my confidence intervals to deteriorate when the number of observations per clusters varies more widely (MacKinnon and Webb, 2018). Union Territories typically have populations an order of magnitude smaller than state populations.

I compare the labor market trajectories of more and less exposed cohorts against the trajectories that we observe among older cohorts whose trajectories are observed before the hiring freeze was implemented. Appendix Figure A.3 summarizes the comparisons across cohorts and time that will be used to identify the impact of the hiring freeze.

To improve the precision of my estimates, I combine cohorts into groups.¹⁵ The *High Exposure* group consists of individuals who were expected to complete their undergraduate degree during the hiring freeze.¹⁶ The *Low Exposure* group includes everyone who would be expected to graduate after the hiring freeze ended, i.e. they were older than 21 in 2001.

The regression that implements these comparisons takes the following form:

$$y_i = \beta_1 [TN_{s(i)} \times \text{HighExposure}_i] + \beta_2 [TN_{s(i)} \times \text{LowExposure}_i] + \delta_1 \text{HighExposure}_i + \delta_2 \text{LowExposure}_i + \zeta TN_{s(i)} + \eta_{t(i)} + \Gamma' \mathbf{X}_i + \epsilon_i \quad (2)$$

In this regression, $\eta_{t(i)}$ captures a fixed-effects for each NSS round, and $TN_{s(i)}$ is an indicator for Tamil Nadu. The vector of controls \mathbf{X}_i include (Tamil Nadu) \times (age) fixed effects, which adjust for any imbalance in the age at which we observe individuals in the comparison group compared to the treated groups.¹⁷

I cluster standard errors at the state-by-cohort level.¹⁸ Although the total number of clusters in the sample is large, sandwich-based estimates of the standard error are still too small because there are very few clusters in the “treated” group for the coefficients of interest (Donald and Lang, 2007; MacKinnon and Webb, 2018), i.e. the

¹⁵Because college completion rates at this time were relatively low, the number of observations within each cohort is small in Tamil Nadu (see Appendix Table A.2).

¹⁶This includes anyone who would have been between the ages of 17 and 21 (inclusive) in 2001. A typical college student in India begins their studies at age 18, and graduates at age 21. Most undergraduate degree programs are three years long.

¹⁷I make sure the age fixed effects include observations from both the comparison group and the treated groups. I therefore drop any observations from before the hiring freeze that were measured at ages older than what I observe for post-freeze individuals.

¹⁸This approach implicitly assumes that treatment (i.e. exposure to the hiring freeze) can be modeled as having been assigned i.i.d. across state-cohort pairs Abadie, Athey, Imbens, and Wooldridge (2017), but allows for possible serial correlation in error terms within clusters over time. This is reasonable if we are willing to believe: 1) the state in which the hiring freeze happened; 2) the year in which the hiring freeze happened; 3) the length of the hiring freeze were all independently and stochastically determined.

variables HighExposure_i and LowExposure_i have non-zero values from only five clusters each. I therefore report confidence intervals using the wild bootstrap procedure outlined in [Cameron, Gelbach, and Miller \(2008\)](#).¹⁹

Identifying assumptions. The key identifying assumption is the “parallel trends” assumption: namely, in the absence of the hiring freeze, treated cohorts would continue to follow the same labor market trajectories as their predecessors.

A standard technique for supporting this assumption is to check for pre-trends. Unfortunately, given the small sample sizes in each individual cohort, the test lacks statistical power.²⁰ Nonetheless, for transparency, Appendix Figure [A.4](#) plots both the pre-trends and separate coefficient estimates for each treated cohort, i.e. I estimate the following specification:

$$y_i = \sum_{c=-5}^4 [TN_{s(i)} \times \text{Freeze}_{t(i)} \times \beta_{c(i)}] + \sum_{c=-14}^{-4} [TN_{s(i)} \times \alpha_{c(i)}] + \sum_{c=-5}^4 [\text{Freeze}_{t(i)} \times \zeta_{c(i)}] + \sum_{c=-14}^{-4} \gamma_{c(i)} + \Gamma' \mathbf{X}_i + \epsilon_i \quad (3)$$

Cohorts c are indexed by their expected year of college graduation relative to the start of the hiring freeze. $\text{Freeze}_{t(i)}$ is an indicator for whether the outcome was measured after the hiring freeze started. The β_c coefficients capture the treatment effects of interest, while the α_c coefficients capture pre-trends. In the absence of pre-trends, we expect α_c coefficients to hover around $E[\alpha_c]$.

The figure plots estimates of $\beta_c - E[\alpha_c]$ and $\alpha_c - E[\alpha_c]$. Since the estimates have high variance, these estimates often wander away from zero. However, there does not appear to be any clear trend. There is a large, anomalous spike at $c = -5$ (i.e. cohorts age 26 in 2001) but the main results are robust to dropping this cohort from the analysis.

The main threat to identification is that there is some other shock to the Tamil Nadu economy that is concurrent with the hiring freeze. To mitigate this concern, I use

¹⁹I confirm through simulations that the confidence intervals generated by the wild bootstrap have the correct coverage rate in this setting (Appendix Table [A.3](#)).

²⁰Appendix Table [A.2](#) provides counts of the number of observations in each cohort in the analysis sample.

individuals with less than a 10th standard education—who are ineligible for exam-based recruitment—as a placebo group. That is, I re-run the regression in Equation (2) on the sub-sample of ineligible men. The employment rates of less and more educated workers tends to be correlated (Appendix Figure A.5), so to the extent there are state-wide shocks, both highly-educated and less-educated workers should be affected. I also formally test whether the coefficients on the ineligible sample are different from the primary college graduate sample. I refer to this test as the triple difference specification.

3.3 Results

Main Results. Table 2 presents the main results. We see that cohorts in the High Exposure group experienced a stark 9 percentage point decline in employment rates (95 % CI: [-0.117, -0.024]). Relative to a base rate of 73%, this effect corresponds to a 13% drop. This effect averages across two margins: 1) a change in the number of people who were ever employed; and 2) a change in the amount of time people spent employed. Without panel data, I am unable to empirically disentangle these two margins.

The decrease in the employment rate is made up for in almost equal measure by increases in employment and reductions in labor force participation (Columns 2-4). The latter category almost exclusively corresponds to individuals reporting their employment status as attending an educational institute—which, for a sample of college graduates, indicates increased enrollment in post-graduate studies.

By contrast, we see no meaningful change in employment status among in the ineligible sample (Panel B). The point estimates are small, and generally insignificant. Correspondingly, the coefficient estimates remain nearly unchanged in the triple difference specification (see Panel C).

The most likely explanation for these results is that candidates spent more time on the “exam track,” either studying full time or waiting for the next recruitment. In theory, time spent unemployed or in school could reflect investments in obtaining a private sector job. However, the fact that applications increased suggests that these investments were not oriented towards the private sector.

Robustness. I probe the robustness of the result in two ways:

Selection of comparison states. I test whether the results are sensitive to the choice of states included in the comparison group. In Appendix Table A.4 I use only the states that neighbor Tamil Nadu in the comparison group (namely Karnataka, Kerala, and undivided Andhra Pradesh). As we would expect, the confidence intervals are much wider when I use few comparison states, but the point estimates are similar.

The lack of sensitivity to the choice of comparison states generalizes: even if I choose a *random* subset of 10 states out of the 24 available in the comparison group, I almost always obtain a same-signed estimate of β_1 (Appendix Figure A.6). This is what we would expect if cohorts experience common shocks and state-specific trends are largely absent in this context.

Selection of comparison cohorts. Dropping cohorts older than age 35 in 2001 is an arbitrary decision. However, it does not impact the coefficient estimates in a meaningful way. I estimate a similar impact on employment rate if I include older cohorts as well (Appendix Figure A.7).

Endogeneity of college completion. The treatment groups are defined conditional on college completion. In theory, college completion may respond endogenously to the hiring freeze. However, I find no evidence this is the case (Appendix Table A.5).

Alternative Interpretations. One might be concerned that the change in employment reflects a demand shock rather than a change in labor supply. The triple difference specification addresses these concerns to some extent. However, if demand shocks had different effects on employment by education level (e.g. because less-educated individuals tend to have less elastic labor supply (Jayachandran, 2006)), then this specification may not fully address the concern.

To aid in distinguishing between demand- and supply-based interpretations of the data, I study the impacts on wage rates.²¹ Consider a simple supply and demand model of the aggregate labor market, in which both curves have finite elasticity. If the decrease

²¹I am grateful to Jaya Wen for this suggestion.

in employment reflects a reduction in aggregate labor supply, then we would expect to observe an *increase* in average wages among the remaining participants in the labor market. Conversely, if the decrease in employment reflects a drop in aggregate labor demand, then we should see a *decrease* in wages.

To assess how wages responded to the hiring freeze, I use earnings data in the NSS, which are available in the rounds in which the Employment module (Schedule 10) was fielded. Household members report the number of days employed in the week prior to the survey, and their earnings in each day. I compute average wages by dividing weekly earnings by the number of days worked in the week. I convert wages and total earnings from nominal to real figures using the Consumer Price Index time series published by the World Bank.

The change in wages will not necessarily show up in the same sets of cohorts or education groups that responded to the hiring freeze. The impact on wages will depend on the elasticity of substitution between different types of workers, and the distribution of reservation wages in the population. I therefore run an omnibus test that remains agnostic about whose wages change. I include all education levels in the sample and estimate a specification of the form:

$$y_i = \beta [TN_{s(i)} \times Freeze_{t(i)}] + \delta Freeze_{t(i)} + \zeta TN_{s(i)} + \eta_{t(i),ed(i)} + \Gamma' \mathbf{X}_i + \epsilon_i \quad (4)$$

In this specification, I combine all post-freeze observations together. I include a separate fixed effect for each NSS round $t(i)$ interacted with the education group $ed(i)$, which is either college graduate, school graduate, or ineligible. The vector of controls \mathbf{X}_i includes age dummies interacted with the Tamil Nadu indicator and education group dummies. I also run a version of this specification separately for each education group.

Appendix Table A.6 summarizes these results. For individuals who stayed in the labor market, earnings and wages rose by about 8% in the post-freeze period. This finding suggests that labor supply decreased after the hiring freeze. Moreover, we see consistently positive effects across all education groups (Columns 1 and 2). Finally, we do not see a change in the share of individuals reporting zero earnings (Column 3), which

suggests that these effects reflect changes in wage rates, rather than positive selection into the workforce during the hiring freeze.

4 Long-Run Impacts on Candidates

So far we have seen that candidates responded to the hiring freeze by remaining out of work, most likely to stay competitive in the remaining exams. In this section, I assess the long-run consequences of this investment strategy.

4.1 Data

I use data from the Consumer Pyramids Household Survey (CPHS). The CPHS is a panel survey of Indian households collected by the Centre for Monitoring the Indian Economy (CMIE). The panel includes roughly 160,000 households in each wave. CMIE revisits households every four months. I use all waves of the survey conducted between January 2014 and December 2019. I weight all estimates using the sampling weights provided by CMIE, i.e. the observations are weighted by the probability sampling weight times the non-response factor.²²

I focus on five groups of outcomes: 1) the Attainment of Government Jobs; 2) Occupational Choice in the Private Sector; 3) Income and Expenditure; 4) Household Labor Supply; and 5) Household Formation. In Appendix Table D, I provide details on how I construct each of the variables used in these categories. In most cases, I use each variable for the longest period for which data is available. Note that this means that the samples will differ across measured outcomes and are therefore not necessarily directly comparable.

²²The data are meant to be nationally representative, but recent evidence indicates that the survey may systematically under-sample very poor households (Somanchi, 2021). Moreover, many of the indicators I use were only collected starting in 2017 or later, several years after the panel started. Due to attrition over time, even if the initial sample was fairly representative, there is no guarantee that the sample I work with is still representative. The results should be interpreted with this caveat in mind.

4.2 Empirical Strategy

I use the cohort-based DiD approach from Section 3, adapted to differences between the CPHS and the NSS.²³ The main difference is that, by virtue of when CPHS data was collected, I no longer have access to a comparison group whose outcomes were measured before the hiring freeze. I therefore use older cohorts as my comparison group instead. Specifically, I treat cohorts who were between the ages of 27 to 35 in 2001 (inclusive) as the control group. The evidence from Section 3 suggests that older cohorts should be relatively unaffected by the freeze. The High Exposure and Low Exposure cohorts are defined in the same way as before.

The panel structure of the CPHS data raises the possibility that households attrit endogenously. To mitigate this effect, I include fixed effects at the (first wave) \times (current wave) level, i.e. I only compare observations that entered the sample at the same time, and for whom the same amount of time has elapsed between the current interview and their first interview. As long as the factors that explain sample entry and persistence are common between Tamil Nadu and the rest of India, these fixed effects should ensure that the changing composition does not bias our estimate of the treatment effect for the sample we observe.²⁴

My main regression specification takes the following form:

$$y_{it} = \beta_1 [TN_{s(it)} \times \text{HighExposure}_{c(it)}] + \beta_2 [TN_{s(it)} \times \text{LowExposure}_{c(it)}] \\ + \delta_1 \text{HighExposure}_{c(it)} + \delta_2 \text{LowExposure}_{c(it)} + \zeta TN_{s(it)} + \nu_{g(it)} + \epsilon_{it} \quad (5)$$

where y_{it} is the outcome for individual i measured in month t , and $\nu_{g(it)}$ are the (first wave) \times (current wave) fixed effects.²⁵ As before, I cluster errors at the state \times cohort

²³To construct cohorts, I rely heavily on the age variable. However, in the CPHS, age is measured with substantial error. Throughout the analysis, I use an imputed version of the age variable, rather than the original value. See Appendix C for details.

²⁴If treatment effects are heterogeneous across the population, the changing composition of the sample will still affect the external validity of the estimated impact, depending on how the treatment effect is correlated with presence in the CPHS sample.

²⁵Note that I have also dropped the age controls. This is because the number of years covered by the CPHS is much shorter than that of the stacked NSS data. In the CPHS I do not observe individuals in the comparison cohorts at the same ages as I observe the treated cohorts. This means I cannot separately estimate age and cohort effects.

level and report 95% confidence intervals using the wild bootstrap.

4.3 Results

The results are summarized in Table 3. Consistent with the results in Table 2, the impacts are concentrated on the High Exposure cohorts. I focus my discussion on this group.

The drop in attainment of government jobs in the High Exposure cohorts is about 10 times larger than we would expect from the drop in vacancies alone (Panel A).²⁶ There are several possible explanations for the discrepancy. First, a combination of non-representative sampling and/or non-random attrition in the CPHS may result in the over-representation of individuals affected by the hiring freeze. Second, affected candidates' exam skills may have atrophied during the hiring freeze (despite their intentions), making them less competitive when vacancies returned to normal levels. For now, I am unable to disentangle these explanations.

In the absence of government jobs, what jobs did affected individuals take up instead? Candidates in the High Exposure cohorts shifted primarily out of business, and towards private employment (Panel B). Moreover, even ten years out of the hiring freeze, we still see elevated levels of men who are out of employment, i.e. who report their primary occupation as "Unoccupied."²⁷ These are men in the prime of their earning years (30 - 40), so not only is reduced labor force attachment at this age is particularly concerning for household well-being, it also likely reflects a decline in the supply of scarce skilled labor in the economy.

A key question is whether affected candidates are less productive. The evidence in Panel C provides some suggestive evidence that they are. I do not see any impact on individual income, though the confidence intervals are extremely wide. This is perhaps not surprising since these coefficients include the effect of shifting from business occupations

²⁶Recall, in Appendix B I estimate that due to the drop in vacancies at most 0.5% of the High Exposure cohorts lose a government job on average.

²⁷The CPHS also includes a direct measure of employment status starting in 2016. I do not see any impact using this outcome. The main difference between these outcomes is that nearly everyone who is unoccupied is not employed, but not vice versa. This is because individuals who are transitioning between jobs still report having a primary occupation. Someone calling themselves "unoccupied" can be thought of as a proxy for either long-term unemployment or dropping out of the labor force.

to employment occupations, and business income tends to be more volatile and measured with more error (De Mel, McKenzie, and Woodruff, 2009). In this context, expenditure, which can be consistently measured across business and non-business households, may be a more reliable measure of household members' earning potential. On average, households with affected cohorts consume substantially less per earning member, on the order of -12% (95% CI: [-0.232, 0.027]). This measure does not allow us to directly pinpoint how much each earning member contributes towards household expenditure. However, the fact that *total* household expenditure does not appear to be higher even though the total number of earning members is higher (see Panel D) suggests that at least some of the infra-marginal earning members are earning less. Labor force participation for prime-age men is nearly 100%, which means it very likely that the affected cohorts are the infra-marginal earning members whose earning capacity has fallen.

Household labor supply increases to compensate for the financial shock (Panel D). Some of the effect is due to affected men living in larger household units, suggesting that the affected individuals are more likely to stay in joint families. Troublingly, part of the effect is driven by older household members delaying retirement. In this setting, delaying retirement is a burden, and therefore a costly means of self-insurance. As Chetty and Looney (2006) note, households that successfully insure themselves through costly actions may still suffer first-order welfare losses.

Another symptom of a decline in well-being is a delay in household formation (Panel E). Men in the High Exposure cohorts are substantially less likely to be the head of household, less likely to be married, and more likely to live with their guardians (parents and/or grandparents) instead. In a context with low divorce rates, this strongly suggests that these men were never married in the first place.

Jeffrey (2010) emphasizes just how damaging delayed household formation can be. Concerns about household formation and social status loom large in this population. As he puts it, “the failure to acquire secure salaried work not only jeopardized young men’s social and economic standing but also threatened their ability to marry and thereby fulfill locally valued norms of adult masculinity” (pg. 85). Long-term candidates report feeling

“left-behind,” “failed,” and inferior (pg. 91). The impacts on household formation thus likely hint at very serious damage to affected candidates’ sense of self-worth and their ability to live up to their full potential, in ways that a household survey is unlikely able to capture but are nonetheless critical components of well-being.

5 Conclusion

I find that the Tamil Nadu hiring freeze that was implemented between 2001 and 2006 had far-reaching negative consequences for male college graduates in both the short and long run. These findings suggest several directions for policy and future research.

First, many states in India do not maintain a regular schedule of exams, and struggle to complete the recruitments they do conduct. This effectively limits the number of available vacancies. The evidence from the Tamil Nadu hiring freeze suggests that candidates may still choose not to give up on exam preparation. In this context, a regular and timely testing policy may help reduce under-employment on the margin.

Second, the costs that we observe in this study need to be balanced against an assessment of the potential benefits from high levels of competition. It is possible that the effort that candidates expend during preparation translates in some way into increased performance once selected. This is an important question for future research.

Finally, the results underscore the importance of developing a better understanding of candidates’ application behavior. Why are candidates willing to tolerate such low probabilities of selection? How much of the decision to drop out is driven by learning? A better model of candidate behavior can help us better forecast the impact of future government hiring policies.

References

- Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge (2017), “When should you adjust standard errors for clustering?” Technical report, National Bureau of Economic Research.
- Arellano-Bover, Jaime (2022), “The effect of labor market conditions at entry on workers’ long-term skills.” *Review of Economics and Statistics*, 104, 1028–1045.

- Banerjee, Abhijit V and Esther Duflo (2019), *Good economics for hard times: Better answers to our biggest problems*. Penguin UK.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller (2008), “Bootstrap-based improvements for inference with clustered errors.” *The Review of Economics and Statistics*, 90, 414–427.
- Chetty, Raj and Adam Looney (2006), “Consumption smoothing and the welfare consequences of social insurance in developing economies.” *Journal of Public Economics*, 90, 2351–2356.
- Colonnelli, Emanuele, Mounu Prem, and Edoardo Teso (2020), “Patronage and selection in public sector organizations.” *American Economic Review*, 110, 3071–99.
- Dal Bó, Ernesto, Frederico Finan, and Martín A Rossi (2013), “Strengthening state capabilities: The role of financial incentives in the call to public service.” *The Quarterly Journal of Economics*, 128, 1169–1218.
- De Mel, Suresh, David J McKenzie, and Christopher Woodruff (2009), “Measuring microenterprise profits: Must we ask how the sausage is made?” *Journal of Development Economics*, 88, 19–31.
- Donald, Stephen G and Kevin Lang (2007), “Inference with difference-in-differences and other panel data.” *The Review of Economics and Statistics*, 89, 221–233.
- Finan, Frederico, Benjamin A Olken, and Rohini Pande (2017), “The personnel economics of the developing state.” In *Handbook of Economic Field Experiments*, volume 2, 467–514, Elsevier.
- Geromichalos, Athanasios and Ioannis Kospentaris (2020), “The unintended consequences of meritocratic government hiring.” Unpublished manuscript.
- Jayachandran, Seema (2006), “Selling labor low: Wage responses to productivity shocks in developing countries.” *Journal of Political Economy*, 114, 538–575.
- Jeffrey, Craig (2010), *Timepass: Youth, class, and the politics of waiting in India*. Stanford University Press.
- Kahn, Lisa B (2010), “The long-term labor market consequences of graduating from college in a bad economy.” *Labour economics*, 17, 303–316.
- Krueger, Anne O (1974), “The political economy of the rent-seeking society.” *The American Economic Review*, 64, 291–303.
- Lund, Crick, Kate Orkin, Marc Witte, Thandi Davies, Johannes Haushofer, Judy Bass, Paul Bolton, Sarah Murray, Laura Murray, Wietse Tol, Graham Thornicroft, and Vikram Patel (2019), “The economic effects of mental health interventions in low and middle-income countries.” Working paper.
- MacKinnon, James G and Matthew D Webb (2018), “The wild bootstrap for few (treated) clusters.” *The Econometrics Journal*, 21, 114–135.

- Mangal, Kunal (2021), “How much is a government job in india worth?” Technical Report Working Paper No. 41, Centre for Sustainable Employment.
- Moreira, Diana and Santiago Pérez (2022), “Civil service exams and organizational performance: Evidence from the pendleton act.” Unpublished manuscript.
- Muralidharan, Karthik (2015), “A new approach to public sector hiring in india for improved service delivery.” In *India Policy Forum*, volume 12, Brookings-NCAER.
- Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz (2012), “The short-and long-term career effects of graduating in a recession.” *American Economic Journal: Applied Economics*, 4, 1–29.
- Schwandt, Hannes and Till M Von Wachter (2020), “Socioeconomic decline and death: Midlife impacts of graduating in a recession.” Technical report, National Bureau of Economic Research.
- Silva, JMC Santos and Silvana Tenreyro (2006), “The log of gravity.” *The Review of Economics and Statistics*, 88, 641–658.
- Simons, Lauren M (2016), *Concurso Fever: The Epidemic Sweeping Brazil*. Ph.D. thesis, Norman OK.
- Somanchi, Anmol (2021), “Missing the poor, big time: A critical assessment of the consumer pyramids household survey.” URL <https://doi.org/10.31235/osf.io/qmce9>.
- Teorell, Jan, Carl Dahlström, and Stefan Dahlberg (2012), “The qog expert survey dataset.” *University of Gothenburg: The Quality of Government Institute*.
- The World Bank (2004), “State fiscal reforms in india: Progress and prospects.” Technical Report 28849-IN, URL <https://openknowledge.worldbank.org/handle/10986/7795>.
- White, Halbert (1982), “Maximum likelihood estimation of misspecified models.” *Econometrica*, 1–25.
- Yu, Sun (2020), “Chinese graduates seek shelter in civil service.” *Financial Times*, URL <https://www.ft.com/content/311a3bee-49d0-47fe-9526-9cc746a47c00>.

6 Figures


Figure 1: Recruitment Intensity in Tamil Nadu for Posts Impacted by the Hiring Freeze



Data Source: TNPSC Annual Reports, 1990 to 2010.

Notes: The figure plots measures of recruitment intensity for posts that were not exempted by the hiring freeze, i.e. it includes all posts recruited through merit-based exams in the state government except police/firefighters, medical staff, and teachers. The x -axis is the state government's fiscal year, which runs from April to March of the following calendar year. Red lines mark the beginning and end of the hiring freeze. Fiscal year 2006 is not included in the hiring freeze since most of this fiscal year was not covered by the freeze (the freeze ended in July 2006).

Figure 2: What Fraction of Eligible Men Appear for State-level Competitive Exams?



Data Sources: TNPSC Application Data; Census of India, 2001 and 2011.

Notes: The application rate is calculated as the total number of unique candidates from Tamil Nadu who appeared for “group” recruitment exams in the 2013 fiscal year divided by an estimate of the eligible population in each education \times age group cell from the 2011 Census. Applications across recruitments are considered to be from the same candidate if they have the same name, date of birth, education, gender, and pincode. The group recruitments included in this calculation account for 87% of all vacancies and 91% of all applications in the fiscal year for competitive exams for posts that were not exempted during the hiring freeze. The coarseness of the x -axis corresponds to the age bins provided by the Census. To account for the two year gap between TNPSC application data (the earliest year for which such data are available) and Census data, I multiply the observed count in each education \times age group cell by the average two-year growth rate in that cell from the previous decade. That growth rate is calculated as: $(\text{population}_{2011} / \text{population}_{2001})^{1/5}$.

7 Tables

Table 1: Competition Increased During the Hiring Freeze

	(1) Vacancies	(2) Applications	(3)
Notified During Freeze	-1.360** (0.517)	0.073 (0.526)	1.049** (0.358)
Notified After Freeze	0.813* (0.396)	0.575 (0.379)	0.109 (0.370)
Log Vacancies			0.534*** (0.104)
Dep. Var. Mean Before Freeze	425	61,489	61,489
Post Type FE	X	X	X
N	57	57	57

Data: TNPSC Annual Reports, FY 1992/93 - FY 2010/11.

Notes: The unit of observations is a recruitment. The dependent variable is the number of applications received. The sample is restricted to: i) recruitments that share a post name with a recruitment that was notified during the hiring freeze; and ii) were otherwise impacted by the hiring freeze. Recruitments are classified into before/during/after based on their notification date relative to the timing of the hiring freeze. Columns present coefficient estimates from a Poisson regression (see equation (1) from the main text). [White \(1982\)](#) robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Short-Run Impacts on Labor Supply

	(1)	(2)	(3)	(4)
			Out of labor force	
	Employed	Unemployed	Enrolled in education	Other
<i>Panel A: DiD estimates, college sample</i>				
TN \times High Exposure (β_1)	-.095** [-.177, -.024]	.047 [-.013, .113]	.04 [-.024, .109]	.008 [-.004, .019]
TN \times Low Exposure (β_2)	.031 [-.055, .118]	.013 [-.037, .053]	-.028 [-.083, .02]	-.015 [-.057, .011]
Mean, TN before 2001	.732	.127	.127	.014
Observations	47,998	47,998	47,998	47,998
<i>Panel B: DiD estimates, ineligible sample</i>				
TN \times High Exposure ($\tilde{\beta}_1$)	0 [-.019, .023]	-.004 [-.02, .012]	.002 [-.002, .006]	.003 [-.013, .024]
TN \times Low Exposure ($\tilde{\beta}_2$)	.012 [-.007, .03]	-.009** [-.016, -.001]	-.001 [-.004, .002]	-.002 [-.017, .015]
Mean, TN before 2001	.958	.02	.002	.02
Observations	208,342	208,342	208,342	208,342
<i>Panel C: Triple difference estimates, combined sample</i>				
$\beta_1 - \tilde{\beta}_1$	-.095** [-.176, -.016]	.051* [-.01, .119]	.039 [-.028, .109]	.005 [-.029, .025]
$\beta_2 - \tilde{\beta}_2$.019 [-.064, .109]	.022 [-.031, .063]	-.027 [-.084, .02]	-.014 [-.06, .018]
Observations	256,340	256,340	256,340	256,340

Data: National Sample Survey, 1994 to 2010.

Notes: Panel A presents difference-in-differences estimates of the impact of the hiring freeze on employment status for the main sample of interest. This sample is: 1) men; 2) who are college graduates; 3) who were between the ages of 17 to 35 in 2001. For additional details about the sample construction, see Section 3.2. Coefficients correspond to β_1 and β_2 from Equation (2). Panel B presents a placebo test, estimating Equation (2) on the sample of individuals ineligible for exam-based government jobs, i.e. those with less than a 10th standard education. Panel C differences the coefficients from Panels A and B. 95% confidence intervals in brackets, computed via wild bootstrap with 999 replications, clustered by state \times cohort level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Long-Run Impacts

	TN \times High Exposure	TN \times Low Exposure	Mean	Individuals	Obs.
<i>Panel A: Attainment of Government Jobs</i>					
Has government job	-0.052** [-0.097, -0.004]	-0.018 [-0.089, 0.065]	0.119	2,601	117,629
<i>Panel B: Occupational Choice in the Private Sector</i>					
Employee	0.080* [-0.011, 0.179]	0.008 [-0.079, 0.101]	0.417	2,601	117,629
Business	-0.074** [-0.126, -0.020]	-0.019 [-0.086, 0.040]	0.258	2,601	117,629
Farmer	0.022 [-0.037, 0.081]	0.014 [-0.051, 0.076]	0.166	2,601	117,629
Daily wage labour	0.003 [-0.030, 0.039]	0.011 [-0.033, 0.056]	0.026	2,601	117,629
Unoccupied	0.022** [0.000, 0.046]	0.004 [-0.008, 0.018]	0.014	2,601	117,629
<i>Panel C: Income and Expenditure</i>					
Log Labor Income	0.056 [-0.083, 0.198]	0.014 [-0.114, 0.148]	9.608	26,834	1,028,793
Log Total HH Expenditure	-0.013 [-0.118, 0.091]	0.000 [-0.098, 0.094]	9.290	27,895	1,066,061
Log Expenditure per earning member	-0.115* [-0.232, 0.027]	-0.075 [-0.227, 0.042]	8.938	27,730	1,060,241
<i>Panel D: Household Labor Supply</i>					
# other employed HH members	0.226*** [0.080, 0.383]	0.087 [-0.075, 0.238]	0.309	4,188	176,824
Fraction other adults employed	0.039 [-0.009, 0.085]	0.006 [-0.037, 0.046]	0.079	4,162	175,974
Has employed HH member aged 55+	0.078*** [0.020, 0.136]	0.012 [-0.015, 0.061]	0.059	4,188	176,824
Fraction HH members 55+ employed	0.070** [0.009, 0.134]	0.022 [-0.013, 0.050]	0.116	1,760	63,488
<i>Panel E: Household Formation</i>					
Head of Household	-0.121*** [-0.185, -0.054]	-0.052 [-0.288, 0.152]	0.677	28,284	281,041
Married	-0.098** [-0.197, -0.012]	-0.016 [-0.092, 0.045]	0.944	471	44,296
Lives with guardian	0.093** [0.018, 0.167]	0.056 [-0.124, 0.249]	0.383	28,284	281,041

Data: CMIE Consumer Pyramids Household Survey, 2014-2019.

Notes: Each row presents results from a separate regression. Columns 2 and 3 correspond to the primary coefficients of interest. See Appendix D for details on variable definitions and construction. 95% confidence intervals in brackets, computed via wild bootstrap with 999 replications, clustered by state \times cohort level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$


Appendix

Table of Contents

A	Additional Figures and Tables	30
B	Estimating the Direct Demand Effect of the Hiring Freeze	42
C	Handling Measurement Error in Age in the CPHS	45
D	Long-Run Outcomes: Variable Construction	47

A Additional Figures and Tables

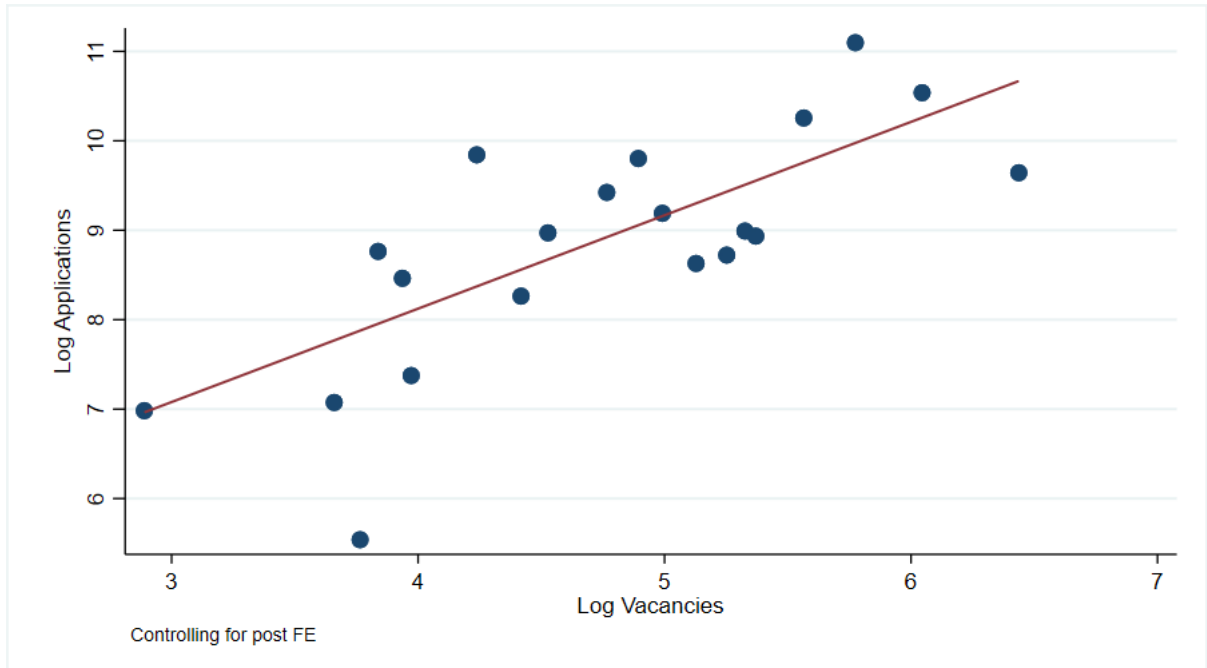
Figure A.1: Recruitment Intensity for Exam-Based State-level Posts Exempted by the Hiring Freeze



Data Source: Panel a) TN Uniformed Service Board; Panel b) TN Public Service Commission; Panel c) TN Teacher Recruitment Board

Notes: This figure plots measures of recruitment intensity for each of the three main categories of posts exempted by the hiring freeze (police and firefighters, medical staff, and teachers, respectively). Note that in the bottom figure, the figure plots the number of *candidates selected*, since data on vacancy notifications are unavailable. The selection year may not correspond with the notification year, since recruitments may take multiple years to complete.

Figure A.2: The Correlation between Applications and Vacancies



Data: TNPSC Annual Reports, 1990 to 2001.

Notes: The figure plots a binned scatter plot and regression line indicating the correlation between the number of vacancies offered in a post and the number of applications received, conditional on the post type. The sample is restricted to: i) recruitments that share a post name with a recruitment that was notified during the hiring freeze; ii) were otherwise impacted by the hiring freeze; and iii) were notified before the hiring freeze.

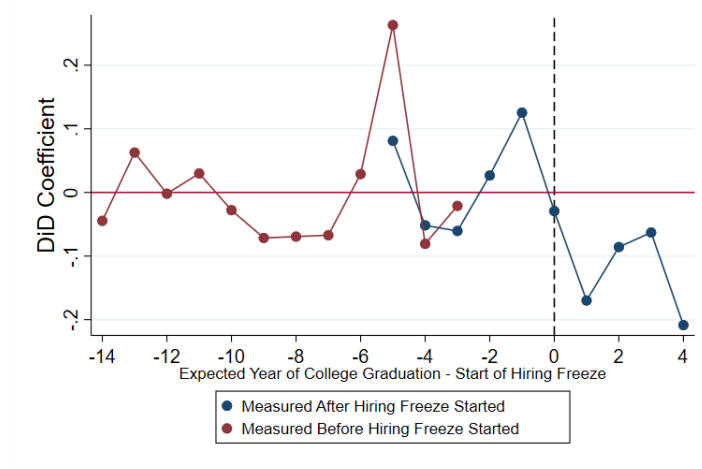
Figure A.3: Empirical Strategy for Estimating Short-Run Impacts

Expected Year of College Graduation - Year of Hiring Freeze	Age in 2001	Outcome measured	
		Post 2001	Before 2001
4	17	Graduated DURING Hiring Freeze	
3	18		
2	19		
1	20		
0	21		
-1	22	Graduated BEFORE Hiring Freeze Started	Comparison Group
-2	23		
-3	24		
-4	25		
-5	26		
-6	27		
-7	28		
-8	29		
-9	30		
-10	31		
-11	32		
-12	33		
-13	34		
-14	35		


Notes: The figure shows how cohorts are grouped for comparison in the empirical strategy used to measure the short-run impacts of the hiring freeze (Section ??). The gray boxes refers to observations that are dropped from the sample.

Figure A.4: Pre-Trends for Short-Run Impacts


(a) Employment



(b) Unemployment



(c) Out of Labor Force

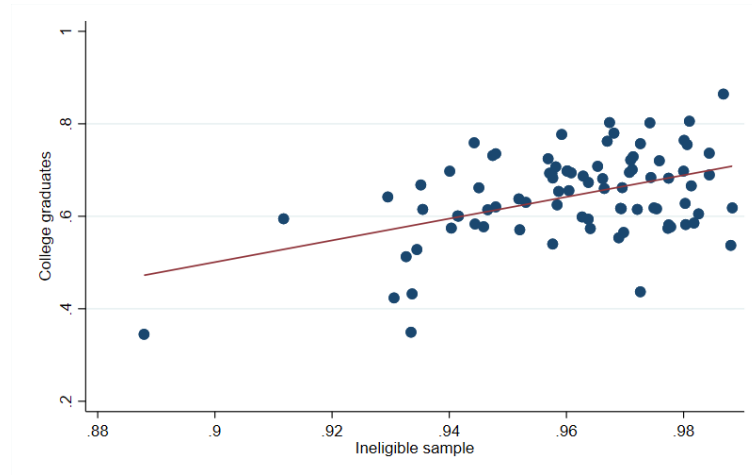


Data Source: National Sample Survey, 1993/94 to 2007/08.


Notes: The figure plots estimates $\beta_c - E[\alpha_c]$ and $\alpha_c - E[\alpha_c]$ from the specification in equation (3). The dashed vertical line separates the High Exposure group from the Low Exposure Group. Confidence intervals are omitted because they are too large to be informative.

Figure A.5: Employment status is correlated between the college-educated and ineligible samples across states and years


Employed



Unemployed




Out of the labor force




Data: National Sample Survey, 50th to 64th rounds (1993/94-2007/08)

Notes: This figure plots the correlation between employment status in the college educated and ineligible samples across states (not including Tamil Nadu) and survey rounds. The x -axis plots the mean of the employment outcome for the ineligible sample, i.e. those with less than a 10th standard education. The y -axis plots the mean for the college-educated sample. The red line plots the regression line. The sample is restricted to large states, defined as those with at least 2500 observations in the sample.

Figure A.6: Short-Run Impact: Sensitivity to the Choice of Comparison States



(a) Outcome: Employed




(b) Outcome: Unemployed

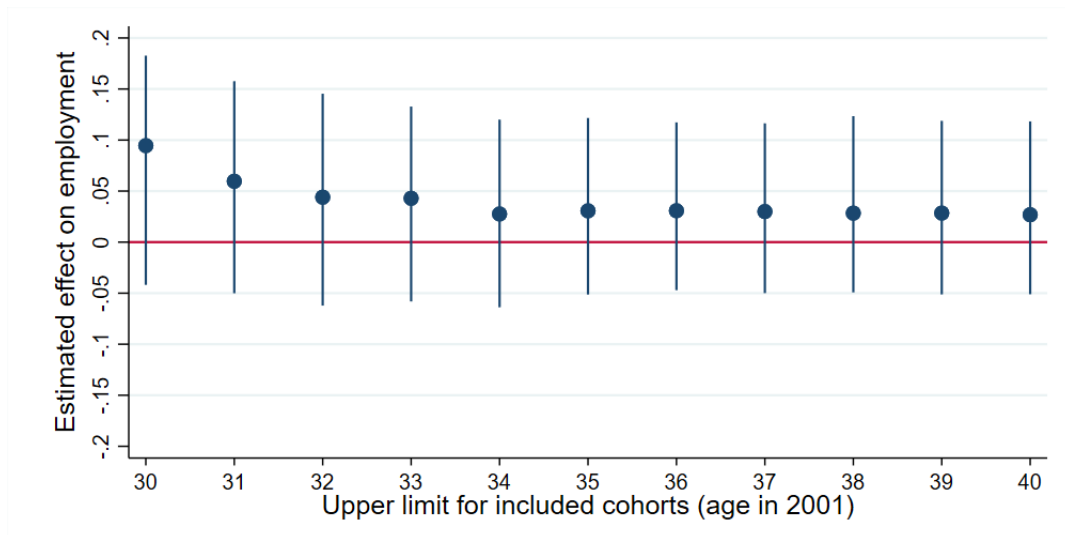
Notes: I randomly sample 10 states from the set of 24 available comparison states. In each of 500 iterations, I re-estimate equation (2) using only the sampled comparison states and Tamil Nadu. The figures plot histograms of the estimates of β_1 ; in the top panel, the outcome variable is employment, and in the bottom panel it is unemployment. A normal distribution is superimposed. The thick red line marks the estimate from Table 2. The dashed black line marks zero.

Figure A.7: Short-Run Impact: Sensitivity to Upper Limit of Cohorts Included in Sample

(a) $TN \times \text{High Exposure } (\beta_1)$



(b) $TN \times \text{Low Exposure } (\beta_2)$



Data: National Sample Survey, 50th to 64th rounds (1993/94-2007/08)

Notes: This figure plots a robustness check for the main results presented in Table 2. I test whether the upper age limit of cohorts included in the sample affects the estimated impact on employment. Each point corresponds to a different estimate for the given upper age limit on the x -axis. Whiskers show 95% confidence intervals, computed via wild bootstrap with 999 replications, clustered by state \times cohort.

Table A.1: Application Intensity in Tamil Nadu

Fiscal Year	Vacancies	Applications	Application Rate	Selection Rate (%)
1992	3,132	430,221	137	0.73
1993	1,076	111,782	104	0.96
1994	647	27,034	42	2.39
1995	770	73,502	95	1.05
1996	1,744	494,048	283	0.35
1997	1,103	173,423	157	0.64
1998	5,492	727,591	132	0.75
1999	3,678	540,611	147	0.68
2000	347	121,035	349	0.29

Data: TNPSC Annual Reports, FY 1992 to FY 2000.

Notes: This table presents statistics on the level of competition for government jobs through the merit-based exam system in Tamil Nadu. The sample is restricted to the sector impacted by the hiring freeze (see Section ?? for details). The application rate is the number of applications divided by the number of vacancies. The selection rate is the reciprocal of the application rate.

Table A.2: Summary of NSS Rounds Included in Analysis Data

Round	Round Start	Round End	Schedule	Sample Size (HH)
50	July 1993	June 1994	10	115,409
51	July 1994	June 1995	1	53,224
52	July 1995	June 1996	1	48,637
53	Jan 1997	Dec 1997	1	51,890
54	Jan 1998	June 1998	1	26,949
55	July 1999	June 2000	10	120,578
56	July 2000	June 2001	1	57,273
57	July 2001	June 2002	1	62,628
60	Jan 2004	June 2004	10	59,159
61	July 2004	June 2005	10	109,601
62	July 2005	June 2006	10	78,879
64	July 2007	June 2008	10	125,578

Data Source: National Sample Survey, 50th to 64th rounds (1993/94 to 2007/08).

Table A.3: Coverage Rate of 95% Confidence Intervals in the Regression Specification Used to Measure Short-Run Impacts

Inference Method	Parameter	
	β_1	β_2
Stata Clustered SE	0.888	0.858
Wild Bootstrap	0.950	0.942

Notes: Table reports the results of simulations that test the coverage rate of different inference methods for the data and main specification used in Section 3. In each of 500 iterations, the outcome variable is changed to a new draw of a Bernoulli random variable that is i.i.d. across observations with a mean of 0.5. The coverage rate measures the fraction of confidence intervals that contain zero.

Table A.4: Short-Run Impacts: Comparison States Restricted to South India

	(1)	(2)	(3)	(4)
	Employed	Unemployed	Out of labor force	
			Enrolled in education	Other
<i>Panel A: DiD estimates, college sample</i>				
TN \times High Exposure (β_1)	-.071 [-.173, .028]	.053 [-.018, .127]	.011 [-.067, .092]	.007 [-.014, .025]
TN \times Low Exposure (β_2)	.063 [-.055, .172]	-.003 [-.067, .058]	-.044 [-.116, .029]	-.016 [-.055, .015]
Mean, TN before 2001	.732	.127	.127	.014
Observations	9,706	9,706	9,706	9,706
<i>Panel B: DiD estimates, ineligible sample</i>				
TN \times High Exposure ($\tilde{\beta}_1$)	.003 [-.023, .032]	-.005 [-.027, .016]	-.004 [-.009, .002]	.006 [-.009, .028]
TN \times Low Exposure ($\tilde{\beta}_2$)	.009 [-.011, .029]	-.008 [-.019, .003]	-.001** [-.003, 0]	.001 [-.016, .017]
Mean, TN before 2001	.958	.02	.002	.02
Observations	42,763	42,763	42,763	42,763
<i>Panel C: Triple difference estimates, combined sample</i>				
$\beta_1 - \tilde{\beta}_1$	-.074 [-.17, .025]	.058* [-.009, .131]	.015 [-.065, .095]	0 [-.029, .026]
$\beta_2 - \tilde{\beta}_2$.054 [-.055, .168]	.006 [-.064, .071]	-.042 [-.114, .03]	-.017 [-.057, .015]
Observations	52,469	52,469	52,469	52,469

Notes: Replication of Table 2, restricting the sample to Tamil Nadu, (undivided) Andhra Pradesh, Karnataka, and Kerala. 95% confidence intervals in brackets, computed via wild bootstrap with 999 replications, clustered by state \times cohort level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.5: Short-Run Impacts: College Completion Rates for Men

	(1)	(2)
TN \times High Exposure (β_1)	.01 [-.006, .031]	-.013 [-.082, .053]
TN \times Low Exposure (β_2)	.011 [-.02, .043]	.054 [-.032, .165]
Sample	Full	High School Graduates
Mean, TN before 2001	.102	.488
Observations	366,273	97,111

Notes: Sample restricted to men between the ages of 17 to 35 in 2001. Table reports estimates using the same specification described in equation (2) from the main text. The outcome is an indicator for college completion. 95% confidence intervals in brackets, computed via wild bootstrap with 999 replications, clustered by state \times cohort level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.6: Short-Run Impacts on Wage Rates

	(1) Log weekly earnings	(2) Log average daily wage	(3) Report zero earnings
<i>Panel A: All education groups</i>			
TN \times Post	.074* [-.001, .154]	.081** [.011, .157]	-.003 [-.039, .033]
Mean, TN before 2001	5.555	3.804	.345
Observations	89,740	89,740	196,456
<i>Panel B: College Graduates</i>			
TN \times Post	.124 [-.177, .485]	.138 [-.181, .509]	-.012 [-.092, .069]
Mean, TN before 2001	6.388	4.469	.38
Observations	9,985	9,985	20,809
<i>Panel C: School Graduates</i>			
TN \times Post	.029 [-.063, .112]	.069* [-.011, .149]	.012 [-.05, .073]
Mean, TN before 2001	5.859	4.004	.392
Observations	21,818	21,818	55,518
<i>Panel D: Ineligible sample</i>			
TN \times Post	.084* [-.007, .181]	.079* [-.006, .165]	-.007 [-.046, .029]
Mean, TN before 2001	5.388	3.684	.325
Observations	57,937	57,937	120,129

Data Source: National Sample Survey, 50th to 64th rounds (1993/94-2007/08).

Notes: Sample restricted to men between the ages of 17 to 35 in 2001. Post is an indicator for whether observations were measured after the hiring freeze started. All income figures are measured in real 2001 INR. 95% confidence intervals in brackets, computed via wild bootstrap with 999 replications, clustered by state \times cohort. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B Estimating the Direct Demand Effect of the Hiring Freeze

Summary. In this section I estimate how large we might expect the impact of the hiring freeze to be on employment rates and earnings in the absence of a supply effect, i.e. the effect arising only from reduced government labor demand. I find that these effects are an order of magnitude smaller than the effects we observe.

Estimation.

1. How many vacancies were lost as a result of the hiring freeze?

I use the data plotted in Figure 1a, i.e. the number of vacancies offered in impacted posts in each fiscal year, focusing attention on observations from either before or during the freeze (FY 1990 - FY 2005).

A straightforward estimate of the average number of vacancies lost per fiscal year of the hiring freeze is given by:

$$\begin{aligned} E[vacancies_t | freeze_t = 1] - E[vacancies_t | freeze_t = 0] &= 302 - 2109 \\ &= -1808 \end{aligned}$$

where $freeze_t = 1$ for $t \geq 2001$ and 0 otherwise.

Because the time series is short, and because $vacancies_t$ is highly skewed, one might be worried about the influence of outliers. I therefore also estimate the change in the *median* number of vacancies offered, assuming the following data generating process:

$$\log(vacancy_t) = \alpha + \beta freeze_t + \epsilon_t \quad \epsilon_t \sim N(0, \sigma^2) \quad (B.1)$$

This model implies that the median loss in vacancies per fiscal year can be estimated as $\exp(\hat{\alpha} + \hat{\beta}) - \exp(\hat{\alpha}) = -1283$.


Thus, over the course of five fiscal years, the hiring freeze caused an estimated loss of around 6415 to 10,545 vacancies.

2. How many vacancies were lost by each cohort?

The total loss in vacancies is distributed across cohorts. What fraction of available vacancies accrue to any one cohort? To answer this question, I use data from all group exams conducted by TNPSC between 2013 and 2017 (a period of 5 fiscal years, like the hiring freeze) to estimate the share of vacancies captured by each cohort. This is the earliest period for which I have such data.

The results of this calculation are summarized in Figure B.1 below. We see that most vacancies are captured by recent college graduates, consistent with their high application rates and intense preparation.

Figure B.1: Fraction of vacancies accruing to each cohort, FY 2013 - FY 2017



Over five fiscal years, we see that no cohort captures more than 8.75% of the available vacancies. Among male college graduates—the focus of the main analysis—no more than 5% are captured.

Allowing for the possibility that male college graduates captured a larger share of vacancies in the early 2000s, I estimate the loss in vacancies to individual cohorts of male college graduates to be:

$$10,545 \times 0.0875 = 923$$

3. How does the loss in vacancies compare to the size of the labor force in each cohort?

The 2011 Census indicates that there were 484,027 male college graduates between the ages of 30-34. This is the age category that is closest to the High Exposure group, for whom we see the largest effects. This means there were about $484,027 / 5 = 96,805$ male college graduates in each individual cohort in the High Exposure cohorts.

A loss in 923 vacancies means that about

$$923/96,805 = 0.0095$$

or about 1% of the most affected cohort was delayed in obtaining or did not obtain a government job through the competitive exam system.

This is the effect of losing 5 years of vacancies. Since most college graduates did not plan on applying before they graduated (as Figure 2 suggests), the actual impact of the freeze would only be felt for the number of years in the freeze since college graduation. In the High Exposure group, cohorts lost between 1 to 5 years worth of vacancies in about equal proportion. Thus, the average affect across this group

would be on the order of

$$\frac{1}{5} \sum_{p=1}^5 [.0095/p] = 0.0057$$

4. What is the average loss in income for each vacancy?

Finan et al. (2017) estimate the Indian public sector wage premium to be 71.2 log points (Table 1, Column 3). Even with this large premium, the demand effect on earnings would only be about

$$71.2 \times 0.0057 = 0.406 \text{ log points}$$

which is about 0.406%.


C Handling Measurement Error in Age in the CPHS

Summary. Age is a critical variable in the analysis in Section 4, since it defines which individuals belong to which cohort. In this appendix, I first present evidence that there is substantial measurement error in this variable. I then discuss the imputation procedure that I use to adjust for this error.

Evidence of measurement error. In each wave of the survey, CMIE captures the age of each household member. This allows me to track how the age of each individual in the sample evolves over the course of the panel. Since birthdays are roughly uniformly distributed, and since CMIE conducts three survey waves per year, roughly one-third of the sample should complete a birthday between each wave.


To check whether this is the case, I compute, for the sample collected in each month, the average difference in age for each respondent from the previous wave. These results are presented in Figure C.1 below. The red line marks $1/3$, which is where the average should lie if measurement error is *on average* close to zero. It appears this is the case until September 2016. In October and November of 2016, age increments too slowly; thereafter, the age increments too fast.

Figure C.1: Average change in age between waves



Before September 2016, the increase in age is not just correct on average across the sample, it is also correct individual-by-individual. If age is correctly measured, then we should see that about $2/3$ of the sample has the same age across waves, and $1/3$ of the sample increments by 1. In Figure C.2 below, I confirm that this is the case. Only about 1% of observations due not fit into this expected pattern.

Figure C.2: Average change in age between waves



The CPHS documentation explains the reason for this pattern. Before September 2016, enumerators were forced to mark the age consistent what was reported in the previous survey. After September 2016, this constraint was removed, in an attempt to impede the propagation of prior measurement error.

How measurement error affects cohort classification. There are 28,284 individuals in the analysis sample. Of these, 26% cannot be cleanly classified into one of the three cohort groups (High Exposure, Low Exposure or Comparison) using the age provided in the CPHS. These ambiguities cannot be resolved without making a few assumptions.

Imputation Procedure. For each individual, I find a sequence of ages that increment correctly according to calendar time and fit the data best.

My imputation procedure makes two main assumptions. First, I assume I have a set of independent measurements of age, $\{a_{it}\}$, where i indexes individuals and t indexes time in months. Before September 2016, the observed ages are not independent, since enumerators were required to increment the age based on the previous value. Accordingly, in the period before September 2016 I use only the first observed age for the imputation procedure.²⁸

Second, I assume that the measurement error ϵ_{it} is zero on average, i.e. I assume a data generating process of the following form:

$$a_{it} = \gamma_i + t/12 + \epsilon_{it} \quad E[\epsilon_{it}] = 0 \quad (\text{C.1})$$

Given these assumptions, a natural way to impute age is to estimate γ_i and compute the fitted values of equation (C.1), i.e. $\hat{a}_{it} = \hat{\gamma}_i + t/12$. A straightforward estimator of γ_i is given by

$$\hat{\gamma}_i = \frac{1}{T_i} \sum_t (a_{it} - t/12) \quad (\text{C.2})$$

where T_i is the number of data points used in the imputation procedure for individual i .

²⁸In other words, if I have N_{i1} observations before September 2016, and N_{i2} observations after on or after September 2016, then I use $\min\{1, N_{i1}\} + N_{i2}$ observations for the imputation procedure.

D Long-Run Outcomes: Variable Construction

This section provides details on how each of the variables used in Section 4 were constructed.

Attainment of Government Jobs

Each of the variables in this panel are indicator variables for a specific type of government job.

Variable Label	Description	First month	Frequency
Has government job	Has a job whose place of work is a government office and is on a permanent contract	May 2017	Wave

Occupational Choice in the Private Sector

The CPHS provides a variable called “nature of occupation” that categorizes occupations into 22 groups. I coarsen the CPHS’s occupational categories into six groups as in the following table, and exclude anyone that was marked as a government employee according to the criteria in Panel A.

Variable Label	Description	First month	Frequency
Employee	Is in a white collar occupation, a Home-based Worker, Support Staff, an Industrial Worker, a Manager, or a Non-Industrial Technical Worker	May 2017	Wave
Business	Is a Businessman, Self-employed entrepreneur or professional, Small Trader/Hawker, or Social Worker/Activist	May 2017	Wave
Farmer	Is a Small Farmer or Organized Farmer	May 2017	Wave
Daily wage labor	Is an Agricultural Laborer or Wage Laborer	May 2017	Wave
Unoccupied	Is a Student, Unoccupied, Retired, or a Home Maker	May 2017	Wave

Income and Expenditure

Individual labor income is notoriously difficult to measure in India since so many households are involved in collective enterprises. I construct an income measure that accounts for household enterprises as follows:

$$\text{Individual Labor Income} = \text{Wage Income} + \frac{(\text{Business Profits} + \text{Imputed Consumption})}{N_{\text{collective}}} \quad (\text{D.1})$$

where $N_{collective}$ is the number of household members with occupations in either farming or business in that wave.

The number of earning members in the household is determined by the number who are not unoccupied, i.e. whose occupation is not “Student,” “Unoccupied,” “Retired/Aged,” or “Home Maker.”

All income and expenditure series are adjusted to 2014 real values.

Variable Label	Description	First month	Frequency
Log Labor Income	Log of individual labor income, computed according to equation (D.1)	Jan 2014	Month
Log Total Household Expenditure	CPHS calculates total household consumption across a wide range of individual items, including individual food items, schooling, clothing, durable, and more. They then provide an estimate a monthly total consumption that adjusts for potential recall bias using the panel structure of the data	Jan 2014	Month
Log Total Household Expenditure per earning member	This is the household expenditure used in the previous row divided by the total number of earning members	Jan 2014	Month

Household Labor Supply

In this section, a household member is counted as employed based on their employment status. Employment status was first collected in January 2016. CPHS measures employment with a one day recall.

Variable Label	Description	First month	Frequency
# other employed HH members	Total number of other household members who were employed	Jan 2016	Wave
Fraction other adults employed	Total number of other households members who were employed divided by the total number of other household members for whom the employment status question was asked	Jan 2016	Wave
Has employed HH member aged 55+	Is 1 if there is any household member age 55+ who is also employed	Jan 2016	Wave
Fraction HH members 55+ employed	Fraction of household members who are 55+ who are employed	Jan 2016	Wave

Household Formation

Variable Label	Description	First month	Frequency
Head of Household	The determination is made by the surveyor. According to the documentation, the surveyor is encouraged to nominate the person who “has the largest say in major decisions of the household” and “holds veto power.”	Jan 2014	Wave
Married	Is 1 if married, 0 otherwise	Jan 2019	Wave
Lives with guardian	This variable is 1 if: 1) the individual is the son of the head of household; or 2) the individual is the grandchild of the head of household, and the son or daughter of the head-of-household is present. This variable equals 0 otherwise.	Jan 2014	Wave