Estimating government worker skills

Application to worker selection and wage setting in Indonesia

Jonas Gathen*

August 13, 2024

Abstract

This paper provides a new approach to estimate government worker skills that is applicable in settings where government output is unobserved and government wages are uninformative about skill differences. The approach estimates skills from wages in comparable jobs in the private sector, relates these skills to skill-related observables using Machine Learning tools and then predicts government worker skills out-of-sample. I apply the new estimation approach to rich Indonesian household-level panel data from 1988 to 2014, showing two main applications. First, I show that the government worker skill premium has declined continuously over time. While government workers are highly selected, their skills have increased less strongly than skills in the overall population, driven by the most skilled workers ending up in the private sector. Second, I analyze government wage setting: the Indonesian government pays a wage premium of at least 30% conditional on skills, about 1/3 of which is driven by the large gender wage gap in Indonesia's private sector.

*Toulouse School of Economics (TSE), Email: jonas.gathen@tse-fr.eu. I would like to thank my advisors Stéphane Straub, Matteo Bobba, Fabrice Collard and Christian Hellwig for their continued support. For helpful feedback and comments, I also want to thank in no particular order: Olivier de Groote, Thierry Magnac, Mitch Downey, Pascal Lavergne, Guohui Jiang, Johanna Isman, Oscar Fentanes and Jens Wäckerle. At last, I am indebted to the feedback by seminar and conference participants at the IAST Political Economy workshop, the TSE Macro workshop, the EUDN development conference in Passau 2020, Stockholm University's applied micro workshop and the workshop of the World Bank's Bureaucracy Lab. All remaining errors are my own.

1 Introduction

How do relative skills of government workers evolve over the course of development and how does this affect changes in the quality of government services? Does the government manage to select highly skilled workers? Are government wages informative about underlying skills and does the government pay a wage premium? A key ingredient for answering these questions is a good measure of government worker skills. This paper provides a new approach to estimate such skills, showing how to estimate unobserved skills in potential government jobs for any worker in the economy. The paper then applies the estimation approach to rich household-level panel data from Indonesia to study (1) systematic changes in the skills of government workers and their selection and (2) government wage setting over a period of almost 30 years.

Traditional approaches in Economics measure a worker's skill – or marginal productivity of labor – either directly by observing output and making an assumption on the underlying production function (e.g. Attanasio et al. 2020; Chetty, Friedman, and Rockoff 2014) or by drawing on observed wages assuming some form of competitive wage setting process that relates wages to the marginal productivity of a worker (Polachek 2008; Meghir and Pistaferri 2011; Sanders and Taber 2012). For the skill estimation of government workers, both approaches are problematic, because government output is usually unobserved (think of the output of a bureaucracy), government worker wages often follow rigid wage-setting rules that disguise underlying skill differences (e.g. Biasi 2021) and wage competition may not force bureaucracies to change wages.

This paper develops a three-step estimation approach that sidesteps these concerns. The first estimation step extracts a noisy individual-specific estimate of skills – an individual fixed effect – from wages in private sector jobs that are comparable to government jobs. The key assumption here is that there exist comparable private sector jobs for which observed wages may flexibly deviate from competitive wages but are at least a noisy function of underlying worker-specific skills and experience. The second estimation step is about predicting the individual-specific noisy skill estimate using a rich set of skill-related observables. I show how Machine Learning algorithms are particularly useful here to disentangle actual skills from estimation noise and flexible shocks to wages. The key assumption in this step and also the most restrictive assumption of the overall approach is that skills of interest are solely a (flexible) function of skill-related observables. An important requirement of the estimation approach is thus the availability of rich skill-related observables that renders this assumption plausible. In the third and last estimation step, one can then enforce the estimated relationship between individual-specific skills and skill-related observables to obtain an out-of-sample potential skill estimate. The approach allows to estimate skills of government workers in government jobs as well as potential skills in a government-like job for any worker in the economy for whom we observe these skill-related observables. The estimation approach can thus also be used in many other settings in which researchers need to estimate potential job-specific skills for workers who work in other jobs with very different skill requirements, while allowing for flexible selection into jobs based on observables and unobservables, as long as these unobservables are not both skill-related and remunerated.

In the second part of the paper, I apply this estimation approach to the case study of Indonesia between 1988 to 2014 drawing on a large and high-quality representative panel that tracks individuals over time. The data – the Indonesian Family and Life Survey (IFLS) – is particularly suited for the estimation approach as it features (1) a large sample of government and private sector workers across many different occupations, (2) one of the longest wage and employment panels in a developing country context, and (3) an exceptional breadth of skill-related variables such as educational background, national exam scores, test results for self-administered Raven tests, cognition and memory tasks, the big-5 personality traits and elicited risk and time preferences. Drawing on the IFLS, I estimate individual-level skills in government-like jobs using private sector workers, which I reweight based on similar occupations and sectors of work as in the government. To do so, I first purge observed wages from changes in the equilibrium skill price in government-like private sector jobs. The equilibrium skill price captures systematic changes in demand and supply, which may in part be driven by changes in how the government hires workers. The approach identifies equilibrium skill price changes using the flat spot identification approach in Bowlus and Robinson (2012). I then back out individual-level fixed effects from a standard Mincerian regression on wages deflated by the skill price. The next step predicts the noisy individual fixed effects using skill related observables and different Machine Learning algorithms. I find that an off-the-shelf gradient boosted tree algorithm outperforms other algorithms in the Indonesian context and use this estimated model to predict government worker skills for all workers in the economy.

Having estimated government worker skills, I then illustrate two main applications for which skill estimates are key. In the first application, I look at changes in government worker skills, the selection of government workers based on skills and how relative skills of government workers have evolved. I find that Indonesian government workers are strongly positively selected on skills and that their skills have also increased over time as new cohorts with higher skills entered the labor force. However, I find that the skill premium of government workers actually declined over time in comparison to the private sector. I show that a big part of this is driven by an increasing right tail of workers who would be highly skilled in government jobs but who the government does not attract for government jobs. Further, I find evidence that changes in the hiring practices of Indonesian government workers in the process of democratization, decentralization and civil service reforms after the year 2000 have led to small improvements in how the government *de facto* selects government workers. Finally, I show evidence that uneven hiring waves across years lead to differences in cohort-specific government employment shares, which have an adverse effect on government worker selection: in years in which the government hires more, average government worker skills decline, which holds conditional on the jobs that the government opens. This is consistent with the idea that more hiring forces the government to eventually move down the distribution as the top of the distribution thins out.

In the second main application, I show how to use the skill estimates to look at government wage setting and the wage premium of government workers. First, I find evidence that Indonesian government wages are indeed less informative about underlying skills than private sector wages in comparable jobs. I find this by predicting government wages using estimated skills from comparable private sector jobs and by re-estimating skills using government wages instead. However, due to more predictable life-cycle wage progression in the public sector, government wages become more predictive of skills when accounting for experience. Second, I find that the Indonesian government pays a large wage premium of at least 30% compared to similar jobs in the private sector. Does this mean that the Indonesian government is overpaying for workers? While this could be true, I also find strong evidence for more discriminatory wage setting in the private sector. A gender wage gap in the private sector of 35% even after controlling for job, skills and experience accounts for 1/3 of the government wage premium alone.

The paper is structured as follows. Below, I discuss how the new estimation approach and empirical results relate to the existing literature. In the next section, I explain the procedure to estimate government worker skills. This approach is applied to Indonesian data in Section 3, where I also give more details on the local context. Section 4 shows two main applications for Indonesia as examples for the usefulness of the approach. Section 5 discusses extensions and the last section concludes.

Related literature

The contribution of this paper is both technical and conceptual. Technically, the paper contributes to the literature by proposing a novel estimation approach that allows to measure government worker skills for all government workers as well as any other worker in the economy. This differs from two alternative approaches that have previously been proposed in the literature. One common approach – as followed by Dal B'o, Finan, and Rossi (2013), Dal B'o et al. (2017), Besley et al. (2017), and Colonnelli, Prem, and Teso (2020) – is to use (residualized) previous private sector wages of government workers as a measure of government worker skills. Previous private sector employment can be a misleading measure of skills in cases where government workers take their first real job in the public sector or have previously worked in a very different occupation, in which case it might be a better measure of a worker's outside option. More technically, the approach in this paper – similar to Dal B'o et al. (2017) and Besley et al. (2017) – also draws on residuals from observed private sector wages but restricts to a comparable subset of government jobs and deals directly with the econometric concern that residualized wages only give a noisy estimate of individual-level skills, "regularizing" skills using skill-related observables and Machine Learning tools. Whenever rich skill-related observables are available, I thus believe that my approach provides a better way to estimate skills from wages.

The second approach followed in Best, Hjort, and Szakonyi (2023) is complementary to my approach. Best, Hjort, and Szakonyi (2023) consider a context in which individual bureaucrats' output is observable – prices paid by bureaucrats in the government procurement of goods. Similar to using wages, they then also draw on a residualized measure of output and extract an individual-bureaucrat fixed-effect. While Best, Hjort, and Szakonyi (2023) are not directly interested in the individual fixed-effects, they are interested in a variance decomposition for which noisy estimates of fixed-effects would bias their results. They thus propose a "covariance shrinkage" approach that uses bootstrapping to separate variances in the true signal and noise. The two key differences are that, first, the approach in this paper does not require to observe government output but informative wages instead, and second, that to "shrink" the noise in the estimation of individual fixed-effects this paper uses rich covariates that also allow to predict skills for any worker in the economy, while the approach in Best, Hjort, and Szakonyi (2023) uses a shrinkage approach without covariates.

Conceptually, the paper uses estimated government worker skills to contribute to the growing literature on the workings of bureaucracies, the developmental state and how the delivery of government services can be improved (Chong et al. 2014; Finan, Olken, and Pande 2017; Rasul and Rogger 2018; Decarolis et al. 2020). The estimation approach allows for a systematic measurement of government worker skills, a key – but difficult to measure – input in the production of government services. The estimation approach thus complements a number of recent papers that have studied how the government selects government workers (Colonnelli, Prem, and Teso 2020; Dal B'o, Finan, and Rossi 2013; Jia, Kudamatsu, and Seim

2015; Bhavnani and Lee 2019; Estrada 2019), who self-selects into government jobs (Ashraf et al. 2020; Hanna and Wang 2017; Weaver 2021), how the government remunerates workers (see: Finan, Olken, and Pande 2017) and how the government competes with alternative employers for talent (Deserranno, Nansamba, and Qian 2024). In line with Ashraf et al. (2020) and Weaver (2021), I find strong positive selection of government workers based on skills. This also means that large documented government wage premia in developing countries (e.g. see: Finan, Olken, and Pande 2017) are smaller after controlling for selection on skills; for Indonesia, I find a 30% wage premium conditional on skills, which is more than 15 percentage points lower than without controlling for skill selection.

The key novelty with respect to this literature is the broader scope of the estimation approach that allows to study all government workers over a long period of time. The approach is particularly useful in settings where government output is hard to observe, and thus especially relevant for studying higher-level bureaucracies where government output may be hard to define and measure, allowing researchers to move beyond the study of last-mile service delivery (e.g. Chaudhury et al. 2006; Banerjee, Iyer, and Somanathan 2007; Finan, Olken, and Pande 2017). Importantly, the broader scope of the approach allows me to establish at least two novel findings in the literature. First, I show that despite growing absolute skills, relative skills of government workers systematically declined in Indonesia over the past 30 years. I link this finding to the difficulty of the Indonesian government to attract the workers with the highest skills to the government. Second, I show evidence for the detrimental effect of government hiring cycles on the selection of government workers. The evidence is consistent with the idea that in years of outsized hiring, the government needs to move down the skill distribution of the applicant pool to fill all government positions.

A good sign of a new estimation approach is that it raises many interesting questions that can now be studied more rigorously: For example, what are the output or welfare costs of government hiring cycles? Or what drives the relative decline in government skills and does this go in hand with a relative decline in state capacity versus private sector capacity over the course of development? These questions are particularly well-suited for future structural work, for which the estimated government skills in this paper can function as a direct input.

2 Identifying government worker skills

This section outlines the three-step identification approach and its underlying assumptions. The focus is on the baseline identification setup that is also used for the subsequent empirical application, while extensions are left to Section 5.

The first step is about identifying a "noisy" signal of human capital and skills using wages in a suitable private sector job, which Proposition 2.1 formalizes.

Proposition 2.1 (Step 1: Identifying a signal of government-worker skills)**.** *Assuming there exists a subset of "government-like" jobs (potentially in the private sector) for which:*

*(i) (**wage determination process**): observed real hourly wages $W_{i,e,t}$ of individual i with experience e at time t follow:*

$$W_{i,e,t} = P_t * H_{i,e,t} * exp(\epsilon_{i,e,t}) \tag{1}$$

where $H_{i,e,t}$ captures human capital relevant for government jobs, P_t captures the (equilibrium) price of efficiency units of human capital and $\epsilon_{i,e,t}$ is a flexible mean-zero error term that is independently distributed of $(P_t, H_{i,e,t})$, follows a stationary process and has finite variance.

*(ii) (**experience profile**): Human capital $log(H_{i,e,t}) \equiv h_{i,e,t}$ follows a standard Mincerian experience process:*

$$h_{i,e,t} = z_i + \delta_0 * exp_{i,e,t} + \delta_1 exp^2_{i,e,t} \tag{2}$$

with $\delta_0 > 0$ and $\delta_1 < 0$ such that the experience profile is concave, and where $z_i = log(Z_i) = h_{i,0,t}$ denotes individual-specific human capital at labor market entry.

Then:

1. *(**flat-spot identification**): Following Bowlus and Robinson (2012), an unbiased and consistent estimator (for $N \to \infty$) for the path of the equilibrium skill price P_t (up to a level of normalization) is given by within-individual wage changes for workers in their flat-spot (FP) region:*

$$\mathbb{E}_{i \in FP}[w_{i,e,t} - w_{i,e-1,t-1}] = p_t - p_{t-1} \tag{3}$$

2. *(**skill signal identification**): Estimates of the experience parameters (δ_0, δ_1) are consistent and unbiased. However, estimates of individual skills z_i are only unbiased but generically inconsistent in a panel with fixed T due to the incidental parameters problem.*

The proof draws on standard results and is relegated to Appendix A.2. Proposition 2.1 defines the two main measures of skills used throughout this paper: (1) individual-specific human capital $h_{i,e,t}$ that incorporates experience and (2) skills at labor market entry z_i. The former is more relevant for questions related to the stock or distribution of human capital at any point in time, while the latter is more relevant for questions related to the selection of government workers and comparing government workers with different levels of experience. Proposition 2.1 shows how wages in a government-like job can be used to identify signals for both measures of skills. These are only noisy signals, because estimates for $h_{i,e,t}$ and z_i are not consistent and thus still include permanent and temporary components of the error $\epsilon_{i,e,t}$ that drive observed wages but are unrelated to human capital.

In the following, I discuss the two main assumptions. The assumption on the **wage determination process** follows a large human capital literature that links wages to human capital. This nests most models of the labor market, but for example rules out models of labor search and matching where informational asymmetries between workers and employees lead to a permanent disconnect between skills and wages as in Taber and Vejlin (2020).[1] The approach only assumes that wages are partly determined by underlying human capital, allowing the price of human capital to change flexibly over time and for wages to be influenced by many other (unobserved) factors as captured by the error $\epsilon_{i,e,t}$.

Allowing the price of human capital to change flexibly over time is key, as this allows for changes in the supply and demand of government-job-specific human capital in the overall economy. Such changes may be directly driven by changes in government hiring. Similarly, the approach allows for systematic changes in the price of government-like output, which could also be directly influenced by the government, and which will also show up as changes in the skill price. Importantly, the approach requires to focus on a market for government-like jobs where there is a single skill price P_t, but it does not require that the government pays that skill price or that other related markets have the same skill price.

The error $\epsilon_{i,e,t}$ in the wage-determination process allows for different factors that influence wages, such as compensating differentials, deviations from fully competitive labor markets or contractual deviations from spot-market wages. This means that firms can temporarily value individual-specific human capital in a way that differs from the market in order to attract, retain, or discourage specific individuals, or because information is imperfect. Moreover, these decisions can be correlated over time, nesting flexible time dependence of the errors. As will become clear in the second estimation step, the approach can even allow for permanent individual-specific shocks such as permanent employer-specific wage markdowns or permanent non-wage job characteristics (e.g. Lamadon, Mogstad, and Setzler 2022). However, the

[1] It should be noted that permanent characteristics at labor market entry are generally found to be by far the main drivers of earnings inequality (e.g. Huggett, Ventura, and Yaron 2011; Keane and Wolpin 1997; Lamadon, Mogstad, and Setzler 2022; Taber and Vejlin 2020), making it a natural starting point.

more stringent assumption is that errors need to be independent of human capital and the skill price; otherwise, the estimation approach will wrongfully attribute part of the error to changes in the skill price or human capital. This rules out deviations from competitive wages that are systematically correlated with experience as this would lead to inconsistent estimates of the experience profile.

The second main assumption is on the **experience profile**. Here, I take a more stringent parametric restriction, assuming a standard quadratic experience profile. The parametric assumption can easily be relaxed as I show in Section 5. The key economic restriction is that experience profiles need to exhibit a flat-spot at some level of experience in which human capital does not grow further, a condition that finds strong empirical support (e.g. Lagakos et al. 2018) and theoretical support from models of endogenous human capital accumulation that imply decreasing returns to human capital accumulation towards the end of a worker's life cycle (e.g. Magnac, Pistolesi, and Roux 2018). The flat-spot restriction is needed for the skill price identification. Following Bowlus and Robinson (2012), the idea is that the wages of workers for whom human capital is not increasing further over the life cycle reflects only changes in the underlying skill price.

Before proceeding with the second step, two further remarks are in order. First, the approach can only identify relative differences in skills across workers at a single point in time and across workers over time, but the overall level of skills is unidentified. A benefit of focusing on differences is that the setup allows for constant differences in wages, for example when labor or government-like output are priced with a constant markup or markdown. Second, the idea of wages revealing human capital implicitly assumes an underlying production function that defines the marginal product of labor for government-like output: $\frac{\partial Y^G}{\partial h_{i,e,t}}$. A statement such as "a worker is twice as skilled as another worker" thus maps differences in skills to the worker's marginal contribution to government-like output.

Having obtained a noisy estimate of skills, the second estimation step gives a consistent estimate of skills by projecting the skill signals on a rich set of skill-related observables.

Proposition 2.2 (Step 2: Projecting skill signals on skill-related observables)**.** *Denote the skill signals of the previous step by $\hat{z}_i$, which can be written as $\hat{z}_i = z_i + \eta_i$ with η_i being independent of z_i given the assumptions in Proposition 2.1. Further assuming that worker skills in the "government-like" job follow:*

$$z_i = f(X_i) \tag{4}$$

where $f()$ is any Borel measurable function and X_i are observable individual-specific charac-

teristics, then the non-parametric regression of $\hat{z}_i = f(X_i) + \eta_i$ using observable X_i gives a consistent estimate of individual skills z_i in government-like jobs.

Proposition 2.2 draws on the Machine Learning literature, where the idea of flexibly projecting noisy estimates of y on explanatory variables X to obtain a consistent estimate of y is called "regularizing".[2] As in the Machine Learning literature, the goal is to remain as flexible as possible for approximating the relationship between z_i and X_i using flexible non-linear functions for $f()$, while best selecting skill-related variables X that predict skills z_i. In the labor market context, many variables are likely correlated with skills, but ex ante, it is unclear which variables will be the most predictive of skills and what the exact functional relationship looks like, let alone interaction effects between variables, making it an ideal setting for Machine Learning algorithms. The key assumption here is that skills are fully explained by observable skill-related variables, which can be restrictive in many settings. For the Indonesian application below, I use a rich set of variables related to educational background, results from self-administered cognition and intelligence tests, survey responses on risk and time preferences, spoken languages, and various literacy measures. In this setup, the key assumption would, for example, be violated if unobserved variables related to soft skills or mechanical skills are relevant for skills in government jobs, but are only insufficiently correlated with observed variables.

Again, two additional remarks are in order. First, a notable benefit of the approach is that one does not need to disentangle between which variables actually *cause* skills versus which are simply *correlated* with skills as long as the same relationship between $f(X_i)$ and skills holds for other workers in the economy. Second and related, correctly separating skills from permanent and temporary shocks to wages relies on using skill-related observables X_i that are uncorrelated with wage shocks. The following example illustrates practical difficulties with this assumption. Suppose that private sector workers in government-like jobs face a gender wage gap due to pure discrimination. Given the interest in actual skills, we do not want to include gender in X_i because the approach would then wrongly interpret the discriminatory gender wage gap as actual skill differences. However, leaving gender out of X_i will only correctly identify $f(X_i)$ if gender is uncorrelated with $f(X_i)$.

The third and final estimation step, as formalized in Corollary 2.1, exploits the mapping between skills z_i and skill-related observables X_i to obtain consistent skill estimates for any worker in the economy.

[2]The proposition follows from universal approximation results for standard Machine Learning algorithms (e.g. Hornik, Stinchcombe, and White 1989). While large-sample properties of Machine Learning estimators such as consistency are under-studied, consistency has been shown for different estimators (e.g. see Athey and Imbens 2019; Shen et al. 2019).

Corollary 2.1 (Step 3: Predicting skills out-of-sample)**.** *Given a consistent estimate of $f(X_i)$ over the support $S \equiv \text{supp}(X)$, the assumptions in Proposition 2.1 & 2.2, then knowledge of X is sufficient to obtain a consistent estimate of z_j for any worker j in the economy for whom X_j is observed (and for whom X_j has common support with S).*

Corollary 2.1 makes clear that the purpose of the second estimation step is not only to "regularize" noisy fixed effects. In this case, one could have also used approaches such as Ridge Regression that do not rely on covariates and the restrictive assumption that skills are solely a function of skill-related observables (e.g. see: Best, Hjort, and Szakonyi 2023). In combination with the last estimation step, the approach allows to predict skills in government-like jobs for any worker in the economy using skill-related observables. For workers in very different jobs, this allows to predict their "potential skills" in government jobs and for government workers it allows to quantify their actual skills without relying on their wages or requiring to observe government output.

At last, note that the common support assumption in Corollary 2.1 is not needed if one is willing to extrapolate on the functional form for $f(X_i)$. The approach already allows for systematic worker sorting across jobs based on individual and unobserved taste differences and sorting based on X_i. However, only in the case of perfect sorting on X_i some support of X_i will never be observed for the private sector workers in government-like jobs and hence skill predictions for any other workers with X_i outside the support will rely entirely on extrapolations of the estimated functional form $f()$. Common support on X_i can be empirically tested and is not an issue in the empirical application that I study below.

3 Estimating government worker skills in Indonesia

In this section I apply the identification approach to the case of Indonesia, the fourth most populous country in the world, over a period of almost 30 years from 1988 to 2014. I start by giving an overview of the context and underlying data and then go through each estimation step in turn.

3.1 Context & Data

Between 1988 and 2014, Indonesia moved away from a military-ruled, highly centralized authoritarian government under General Suharto (1967-1998) to a relatively consolidated, highly decentralized democracy. Incomes per capita have increased roughly 7-fold, poverty has been dramatically reduced and Indonesia has transitioned to become a middle-income country.

The Suharto regime was characterized by extensive cronyism and patronage under which the bureaucracy was greatly expanded but also seen as a direct political instrument to collect votes and offer political support (Fisman 2001; Hadiz and Robison 2013; Martinez-Bravo, Mukherjee, and Stegmann 2017; Robison and Hadiz 2004). Civil Servants were obliged to become members of the political apparatus, support the party in power and hiring and promotion decisions were made to support regime stability (e.g. Kristiansen and Ramli 2006; McLeod 2008). With the Asian Financial Crisis in 1997/1998 and the subsequent fall of the Suharto regime, Indonesia embarked on the *Era Reformasi*, targeting constitutional, judicial, public financial management and privatization reforms as well as an unprecedented decentralization process. In this new, decentralized system, three-fourths of the Civil Service (including teachers and health workers) are assigned to local governments in contrast to a fraction of this under the Suharto regime. Since the decentralization reforms, the central government can steer civil service hiring by setting overall quotas for the number of civil service jobs, while districts are left with a high degree of discretion as they decide on applicants and applicant requirements. Throughout the 2000s, a number of ministries started bureaucracy reform initiatives that tried to set civil service remuneration on par with the private sector and pushed for more competitive hiring and promotion practices (e.g. Horhoruw et al. 2013). Due to uneven adoption across ministries, in 2010, the government mandated public sector reform for all central and local governments, but only in 2014, a new Indonesian Civil Service law was passed.

I draw on nationally representative data from the Indonesian Family Life Survey – IFLS in short – which is particularly suited for the identification strategy in this paper. Specifically, the IFLS is a large and high-quality household- and individual-level panel dataset that collects exceptionally detailed information on individual's occupations, wages, skills and preferences. It is based on a sample of 7,224 households and 22,347 individuals tracked throughout five waves (1993, 1997-98, 2000, 2007-08, and 2014-2015), representing about 83% of the Indonesian population living in 13 of the nation's 26 provinces in 1993. Due to an intensive focus on respondent tracking, re-contact rates between any two rounds are above 90%, and 87% of the original households were contacted in all five rounds (see: Strauss, Witoelar, and Sikoki 2016).[3] As a comparison, these re-contact rates are as high or higher than most longitudinal surveys in the United States and Europe.

Employment & wages: The IFLS data includes detailed employment data for each survey round. In addition to current employment, the survey included questions on previous

[3] Throughout, all results are based on weighting individual observations using provided cross-sectionally representative survey weights that also correct for attrition.

employment. As in Hamory et al. (2021), this allows to create up to a 27-year annual individual employment panel from 1988 to 2014, making it one of the longest employment panel datasets available for developing countries and uniquely positioned to study life cycle wage growth (see: Lagakos and Shu 2023).[4] Employment information captures principal and secondary employment including government jobs. I focus on principal jobs as the job classification throughout this paper. For wages I use real hourly income based on total wages and total hours worked across all jobs, and by deflating nominal values using Indonesia-wide average monthly CPI-based inflation together with the specific survey months to best deal with periods of high inflation.

Skill variables: The IFLS data also contains an exceptional breadth of skill-related variables. I restrict myself to the 28 most important variables based on variable importance metrics in subsequent prediction tasks. These variables are dummies for the level of education, test results for self-administered Raven tests, cognition and memory tasks, the literacy of respondents, relative rankings in standard Indonesian exams, the big-5 personality traits and elicited risk and time preferences. I z-standardize all numeric variables pooling across all individuals.

Definition of government worker: In Indonesia, there are permanent civil service positions called *Pegawai Negeri Sipil (PNS)* as well as temporary civil servant positions (non-PNS). The latter are for example common in the educational sector where 60% of new teacher hires are on temporary contracts (Pierskalla and Sacks 2018). The IFLS includes both permanent and temporary government workers, defining government jobs broadly as jobs with any government office for which one receives remuneration in money or in kind. Appendix A provides further details on the hiring process of government workers in Indonesia.

Job categories, sectors and occupations: The IFLS distinguishes broad job categories (including government workers) from the job's sector at the 1-digit level and an occupational classification at the 2-digit level. To allow for sufficiently large occupational groups for government workers, I recode the occupational classification into nine different 1- to 2-digit occupations that captures most variation in government sector jobs. For example, I leave the 2-digit codes for "teachers" and for "government officials and executives", while I keep the 1-digit code "other service areas".

[4] Employment status and sector of employment are available for each year, but in the fourth and fifth IFLS round, earnings were collected only for the current job. Following Hamory et al. (2021), the earnings measure in this paper is the sum of all wages, profits, and benefits.

Figure 1: Employment shares by broad job category

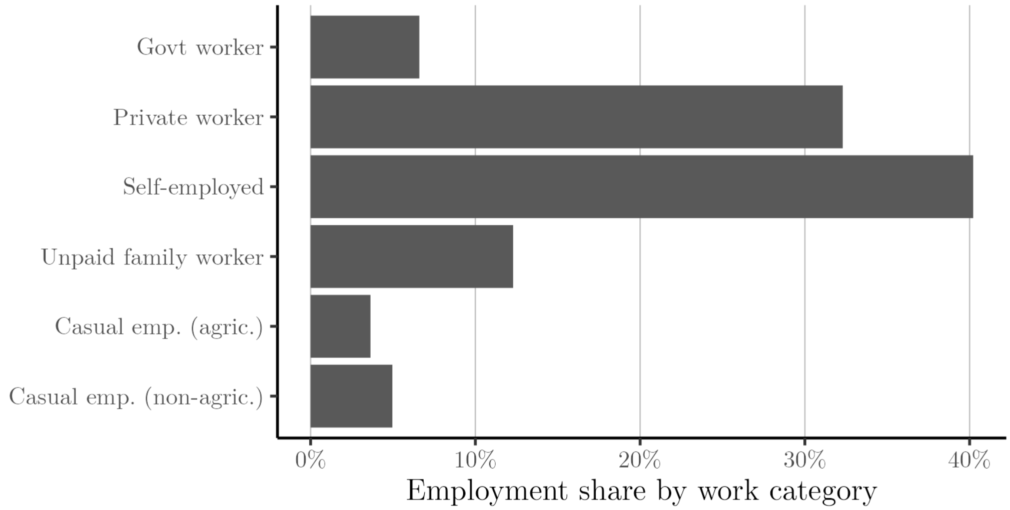

Notes: Barplot of broad job categories for primary occupations over the pooled panel of employment histories 1988-2014. Source: IFLS, Pooled sample: N = 403,626. Unique individuals: n = 36,126

3.2 Descriptives

Figure 1 plots the share of government workers compared to the employment share of the other most common categories of work in developing countries, pooling all waves. On average, about 7% of the workforce are government workers, around 32% of workers have permanent, formal jobs in the private sector, roughly 40% of workers follow some form of self-employment and the remaining 21% of workers are casually employed or unpaid. Compared internationally, government employment in Indonesia is relatively low as it often exceeds 10% of the workforce for middle-income countries and may easily exceed 20% for high-income countries (e.g. Finan, Olken, and Pande 2017). Based on complementary government statistics, up to 50% of the government employees captured in the IFLS data should be non-permanent government workers.[5]

Figure 2 shows the evolution of the government workforce and hiring over time. In line with the stated government policy of reducing government employment, the left panel shows that the share of government workers has been slowly declining since 1988. The decline also shows up in the data through a large drop of hiring for birth cohorts after 1965 as reported in the right panel of Figure 2.

Next, we can ask what types of jobs government workers are doing in Indonesia. Figure

[5]Specifically, permanent government workers only make up roughly 3.5-4% of all employees in Indonesia based on a Civil Service census from the year 2015 reported in Pierskalla et al. (2020) and the total number of employees for the same year taken from Statistics Indonesia.

Figure 2: Evolution of government employment shares

Notes: Evolution of government employment (as primary occupation) as a share of total employment by year from 1988 to 20014 (left) and by 15 equal-sized birth cohort bins (right). Restricted to age range 25 to 58 to account for early government retirement age and tertiary education for government workers. Source: all five waves from the IFLS (1993,1997,2000,2007,2014), Pooled sample: N = 303,368. Unique individuals: n = 29,992. Unique government workers = 2,813.

3 shows the distribution of government jobs across 10 different sectors and compares their shares to the private sector. More than 80% of government sector jobs are in social services, including education and health, with the next biggest sector being agriculture and forestry at around 6%. In comparison, only slightly more than 20% of private sector jobs are in social services, while close to 30% are in manufacturing.

In Figure 4, I provide further details on which occupations government workers hold. More than 30% of government workers are teachers, more than 10% are government officials and executives, and around 15% of government workers work in other service areas. Not surprisingly, the prevalence of the same occupations among private sector workers looks very different: while there are almost no private military/police nor bureaucrats and about twice as many government teachers than private teachers, there are 20x more private laborers and private sales workers and 12x more private agricultural workers.

Another important feature of the data on government employment are transition rates between the private and the public sector, defined as individuals who are observed switching either from a non-government job to a government job (in-mover) or the reverse (out-mover). About 10% of person-year observations are movers of which slightly more move into the government than out (around 53% vs. 47%). Among those who move, 42% are observed to move both in and out of the government.

At last, Table 1 reports how government workers differ from private sector workers and other workers in the economy in terms of observables, including skill variables. The Table also compares government workers to the subsequent control group used in this paper – explained

Figure 3: Employment shares across sectors of work for government vs. private sector workers

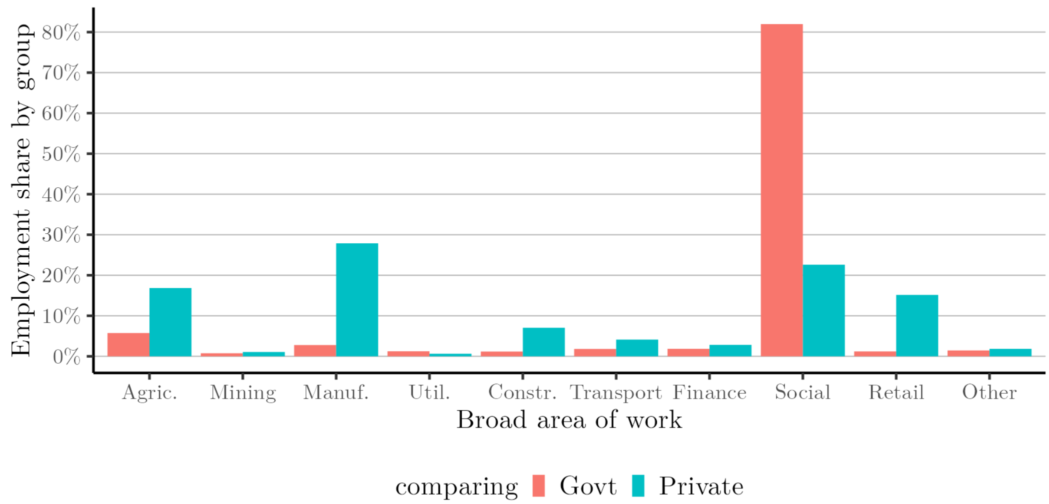

Notes: Pooled data for 1988-2014. Source: IFLS, Total observations = 165,257. Unique individuals: n = 23,377, Govt workers = 3,230. Excludes observations with missing sectoral codes.

in more detail in the next subsection – but which is made up of private sector workers who are reweighted based on the occupations, sector of work, gender and age of government workers. In contrast to private sector workers, government workers are on average slightly more likely to be male, are about 5 years older, earn about 50% higher hourly wages and work slightly less hours. Importantly, even in contrast to the reweighted control group, government workers are on average much more skilled: they are much more likely to have received higher education, perform better on self-administered memory and word ability tasks, and they achieved higher rankings on nationally administered exams. While not reported here, one can also note that these skill measures are mostly increasing for all groups of workers over time.

3.3 Step 1: Estimating "noisy" skills

Following Proposition 2.1, to obtain "noisy" estimates of worker skills in government-like jobs, the first substep is to select a suitable government-like job. As mentioned above, I use private sector workers who are comparable to government sector workers in terms of demographics and broad job-related observables. I do so via a standard propensity score weighting, where weights are derived from a logistic regression predicting the dummy "government worker" using as covariates the occupation, sector of work, age and sex of a worker and restricting to government and private sector workers only. In practice, this means that the highest control weights are assigned to private sector workers who work in social service jobs and have office or service jobs. Note that with larger sample sizes one could also separately apply

Table 1: Main observable differences: government vs. private sector vs. all other workers

Variable	Govt	Control	Private	Other
Share male	0.70	0.71	0.67	0.65
Mean age	39.83	39.60	34.40	42.51
Mean relative wage	1.10	1.09	0.68	1.21
Mean weekly hours	38.94	37.34	39.58	37.71
Share higher educ	0.52	0.24	0.08	0.03
Mean word recall	0.17	-0.08	-0.34	-0.72
Share speak Indonesian	0.38	0.34	0.25	0.15
Mean word ability	0.46	0.00	-0.26	-0.59
Mean math IQ	0.41	0.08	-0.10	-0.18
Mean national ranking (total)	0.04	0.02	-0.05	-0.01
Mean national ranking (language)	0.01	0.02	-0.01	-0.01
Mean national ranking (math)	0.00	0.02	-0.02	-0.01

Details: Based on pooled data and restricting to workers in working age (between 16 to 70 years old), positive work hours and non-missings in age, experience, work hours and income. N = 137,8527. Unique individuals = 24,206. Unique govt workers = 2,342. Unique control workers = 4,812. Unique private sector workers = 13,394. Unique other workers = 8,470. The control group are private sector workers who have worked in social services before. Mean relative wage gives the group-specific mean of the wage divided by the year-specific mean wage across all groups (with means trimmed at the 2.5th and 97.5th percentiles). Apart from dummy variables, all skill-related observables are z-standardized across all individuals.

Figure 4: Government employment shares by occupation

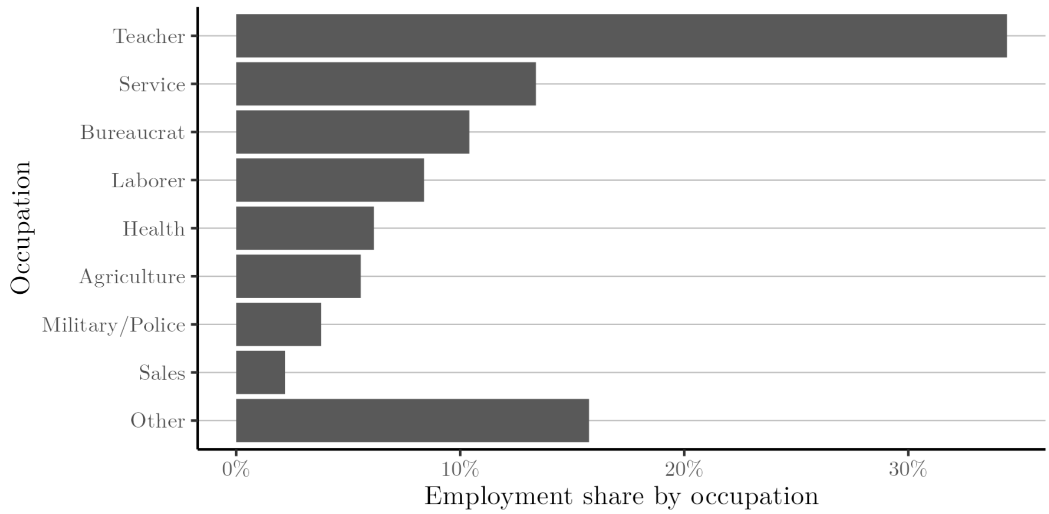

Notes: Barplot of occupations among government workers, 1988-2014. Source: IFLS, Pooled sample: N = 29,855. Unique govt workers = 3,238. Excludes observations with missing occupational codes. Occupation is a combination of 1- and 2-digit occupational codes reported in the IFLS data.

the estimation approach to more disaggregated job categories without pooling across jobs. The level of aggregation that I choose alleviates the difficulty of finding exact private sector counterparts to specialized jobs such as the police. However, the underlying assumption is that even in such jobs, the general skill set needed broadly aligns with the skills needed for social service jobs in the private sector. As explained in more detail in Appendix A, civil servants in Indonesia are also not trained as specialists, but rather go through a common selection and training process that emphasizes more general skills, which is exactly the skills that the estimation approach in this paper seeks to pick up. Further below, I test this and show that the estimation approach indeed picks up on general skills rather than specialized job-specific skills. Another critique may be that the more multi-dimensional nature of government jobs makes performance measures difficult in the public sector (e.g. Finan, Olken, and Pande 2017). However, the same argument can be made for social service jobs in the private sector. And importantly, the skill estimation approach here specifically allows for the idea that measuring performance in a job is difficult and that wages are only noisy signals of underlying human capital and skills. In the results section, I will specifically test the informativeness of wages for underlying skills and compare the private and public sector.

After specifying the estimation sample, the next substep is to purge the observed wage from changes in the equilibrium skill price for government-like jobs in the private sector. Figure 5 plots the estimated skill price drawing on Bowlus and Robinson (2012) and Proposition

Figure 5: Evolution of skill price

Notes: Skill price estimation following Bowlus & Robinson (2012). "Main (Median within)" gives the baseline estimator based on Proposition 2.1, while "Mean within" reports the same estimator using noisier mean changes instead. For the flat-spot region, I use the age range from 52 to 62, consistent with subsequently estimated experience profiles. For all estimates, the level of prices is normalized such that the time-series average is unity.

2.1.[6] Overall, I find little evidence for long-run growth in the skill price despite a 7-fold increase in real GDP per capita over the entire time period. These estimates are in line with a strong increase in the supply of skilled labor that keeps track with increases in demand as the economy grows. The time series can be divided into two periods, separated by the Asian Financial Crisis in 1997. Before the crisis, the estimates indicate a constant to slightly decreasing skill price, explained by an increase in the supply of skilled labor by cohorts affected by Indonesia's large school expansion programs (Duflo 2001, 2004). Upon impact of the Asian Financial Crisis in 1997, the skill price drops strongly as expected by a large drop in the overall demand for government-like jobs in the private sector. In the years after the crisis the skill price recovers strongly, growing by roughly 50% between 2000 and 2014, in line with growing demand as the economy expands and moves further into services.

The last substep following Proposition 2.1 is to obtain a noisy estimate of individual-level human capital by disentangling experience effects from individual-specific skills at labor

[6] Skill price estimates are usually sensitive to the exact flat spot region chosen. I chose an age region from 52 to 62, which is in line with regions usually used in the literature and is consistent with the flat-spot region implied by the subsequently estimated wage experience profile. I use median instead of average within-individual wage changes to reduce the role of outliers, while preserving the identification result in Proposition 2.1.

market entry. Assuming the quadratic experience profile in Equation (2), log wages follow:

$$\tilde{w}_{i,e,t} \equiv w_{i,e,t} - p_t = z_i + \delta_0 * exp_{i,e,t} + \delta_1 exp^2_{i,e,t} + \varepsilon_{i,e,t} \tag{5}$$

Table 6 reports estimates for the parameters (δ_0, δ_1) of the experience profile, which are identified from within-individual changes in wages and experience. Within-individual variation is key to ensure that estimated experience profiles are not biased by systematic composition changes in skills at labor market entry, such as due to less experienced cohorts entering the labor market with better education.[7] The estimated parameters give concave experience profiles with a flat-spot around 32 years of experience, in line with the restriction to 52 to 62 year old workers made for the skill price estimation.

3.4 Steps 2 & 3: Predicting skills

In Step 2 of the estimation approach, following Proposition 2.2, I regress the noisy individual fixed effects flexibly onto a larger set of skill-related observables. I use standard off-the-shelf Machine Learning algorithms for this task. Specifically, I compare three different algorithms with a benchmark of a standard OLS regression with dummies for the educational background (primary, junior secondary, senior secondary and higher education). The three algorithms are LASSO using all variables and all their first-order interaction terms (which includes squared terms), Random Forest and Gradient-Boosted Trees. Due to the high flexibility of Machine Learning algorithms, it is important to avoid overfitting, which leads to noise in small sample estimation. To avoid overfitting, I follow standard practice and use 10-fold cross-validation and train the hyperparameters of each Machine Learning Algorithm via a simple, coarse grid search. Due to overall sample size limitations I avoid additionally separating between validation and test datasets.[8] The final individual-level dataset for "training" the skill estimation algorithm includes 22,597 unique individuals who carry different weights based on their jobs' similarity to government worker jobs.

Table 2 reports the performance of each of the different algorithms in terms of its R^2. I find the GBM algorithm to perform the best with an R^2 of around 17%, while the LASSO

[7]Note that the IFLS only has wage information for two separate years after the year 2000 since the 2007 & 2014 waves do not ask for wages retrospectively. This means that within-worker wage experience profiles are more informed by wage profiles prior to the year 2000. However, the limitation does not preclude to estimate worker fixed effects for workers that enter the labor market after the year 2000 or even workers whose wages are only observed once.

[8]In larger datasets, ensemble methods combining different Machine Learning algorithms could also provide additional gains for prediction tasks and reduces model sensitivity if sufficiently regularized. Model sensitivity usually plays a big issue, but Machine Learning algorithms used on larger datasets can potentially ameliorate the issue.

Table 2: Comparing performance and similarity of different Machine Learning algorithms

Measure	Indiv.	OLS	LASSO	RF	GBM
R^2	22597	0.0555	0.1499	0.1618	0.1731
Corr(OLS,x)	22597	1.0000	0.6857	0.4377	0.6267
Corr(LASSO,x)	22597	0.6857	1.0000	0.7179	0.9307
Corr(RF,x)	22597	0.4377	0.7179	1.0000	0.7687
Corr(GBM,x)	22597	0.6267	0.9307	0.7687	1.0000

algorithm achieves 15% and the Random Forest algorithm about 15%. In comparison to the simple OLS algorithm, the Machine Learning algorithms can explain about three times more variance in the individual fixed effects, highlighting the importance of accounting for non-linearities and additional variables. Interpreting the difference between LASSO & GBM as the importance of non-linearities allows a simple variance decomposition that attributes 80% of the predictive gains compared to OLS to having additional variables and 20% to allowing for additional non-linearities.

So what information are the Machine Learning algorithms picking up and what predicts skills? Focussing on the GBM estimates as the best-performing baseline estimates, Table 6 in Appendix B.1 ranks variables by variable importance measure that weights the relative importance for each covariate in the prediction task taking into account nonlinearity and interactions of the different variables. We can see that the algorithm does identify the various educational background dummy variables as the most important variables for predicting skills, but among the top 15 most predictive variables there are also scores on word cognition tasks, Raven IQ-score, some of the Big-5 measures and the ability to speak the national language. Another way to see the importance of using additional covariates is to look at the correlation of predictions across different algorithms as reported in Table 2. Comparing the predictions of the three different Machine Learning algorithms with the OLS predictions, we can note that the correlation is at most 0.69, indicating that the OLS predictions are missing very important variation that the different Machine Learning algorithms take into account. Secondly, we can compare the different predictions of the Machine Learning algorithms among each other. For example, comparing LASSO with gradient-boosted trees, we find that their correlation is around 0.93, indicating that they both capture similar additional variation.

4 Main empirical results

Taking the individual skill estimates based on the best-performing Machine Learning algorithm, I now show two main empirical applications using estimates of government worker skills. I discuss each in turn.

4.1 Application 1: Changes in skills & selection

In the first application of the skill estimation approach, I zone in on the selection of government workers and changes in skills of government workers and in the overall population. I show how the skill premium declined over time and point to large costs of fluctuations in government hiring for the attraction of talent for the government.

Evolution of absolute and relative skills

Can the government attract the workers who would be most skilled at government jobs? And how does this change over time as the growing private sector competes for talent? Figure 6 plots the evolution of skills in government jobs for government workers, private sector workers, private sector workers in jobs comparable to the government (the control group) and the overall population, revealing three important facts. First, government workers are on average much more skilled: in 1988, they are slightly more skilled than comparable private sector workers and more than 65% more skilled than the overall working population and private sector workers, in line with large skill differences reported in Table 1. Second, as visible on the left plot, average skills increased by roughly 30% across all workers since 1988, again, in line with general increases in education. Third, relative skills of government workers – as reported in the right panel – declined compared to the overall population and particularly in comparison to the average worker in the private sector, but not in comparison to private sector workers who work in similar jobs than government workers.

What drives this decline in the government skill premium with respect to the average private sector worker but not in comparison to private sector workers who are in similar jobs as government workers? I find support for the idea that over the period 1988 to 2014, the government has been able to compete for talent with comparable jobs in the private sector, but increasingly failed to hire workers from new entering cohorts who would have been skilled government workers and lost them to private sector jobs that are different from government jobs.

To show this, Figure 7 reports skills by birth cohort. Overall, more recent cohorts have higher skills, in line with improvements in general education. In particular, 1960 marks a big

Figure 6: Evolution government worker skills

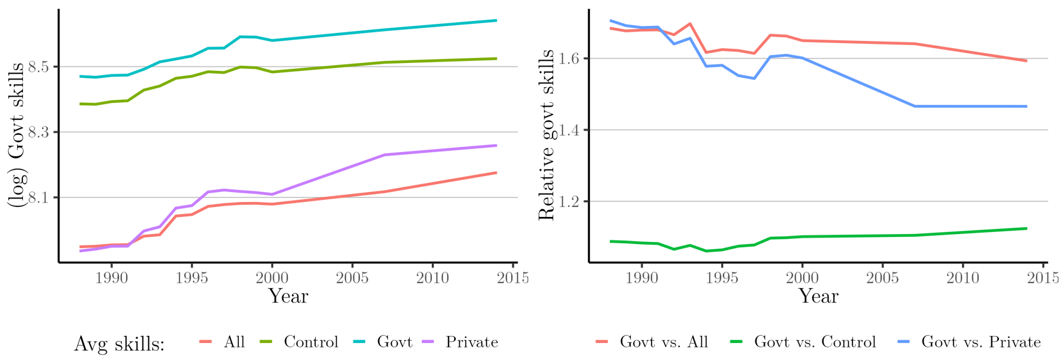

Notes: Results are based on skill estimates using baseline specification with private sector jobs comparable to the government (Control group in Figure) as estimation sample and GBM estimator as ML algorithm. Relative skills are in log differences. Data is pooled across all waves of the IFLS and then plotted by year. Sample restricted to working ages of government workers (25-58).

increase in skills, which coincides with the expansion of schooling as part of the largest school construction program in Indonesian history (see: Duflo 2001). At the same time, the relative skill premium of government workers also strongly declined across birth cohorts compared to all workers and all private sector workers starting around 1960. In Figure 9 in the Appendix, I show that this relative decline after 1960 is not driven by changes in the type of government jobs offered, with Figure 7 looking almost indistinguishable after holding the occupation and sector composition of government jobs fixed over time. At the same time, Figure 2 already showed that government hiring strongly declined in birth cohorts shortly after 1960. In the case that government jobs are in high demand, one would have expected that a decline in supply would have led to an increase in relative skills as long as the government is good at selecting workers on skills. While this effect can indeed be seen for private sector workers in jobs similar to government jobs (the control group), it does not seem to play out for the rest of the private sector. So is the Indonesian government simply becoming worse at selecting other private sector workers? And are there actionable ways to improve the selection of government workers? The next two subsections zone in on these questions in turn.

The government selection rule

Motivated by the decline in the relative skill premium of government workers, I now turn to studying the selection of government workers across the entire skill distribution. For example, are changes in average skills of government workers driven by the entry of low skilled workers or the failure to attract the most skilled workers into the government? I show that the Indonesian government is generally successful at selecting more skilled workers, but that this selection fails at the top, with the most skilled workers not ending up in the government.

Figure 7: Government worker skills across cohorts

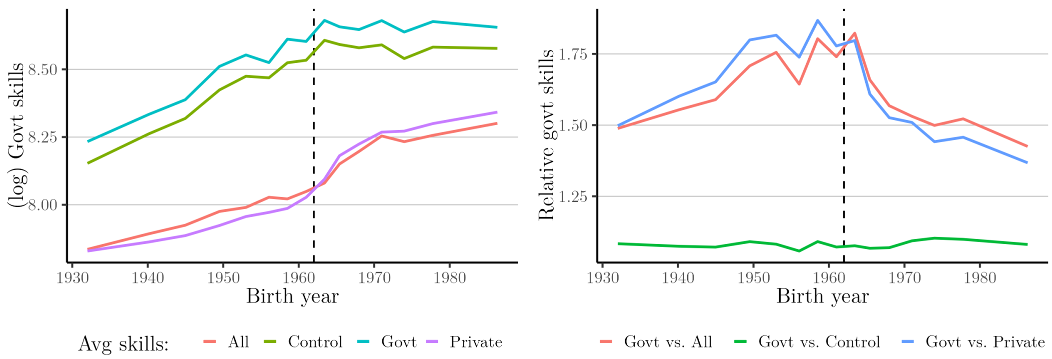

Notes: Results are based on skill estimates using baseline specification with private sector jobs comparable to the government (Control group in Figure) as estimation sample and GBM estimator as ML algorithm. Relative skills are in log differences. Data is pooled across all waves of the IFLS and then plotted by (binned) cohort. Cohort bins are determined by equal-sized bins in pooled data. Sample restricted to working ages of government workers (25-58). Dotted line denotes the first cohort that was treated by the INPRES school construction program studied in Duflo (2001).

I show this by defining a simple reduced-form *selection rule* that denotes the relative probability of being selected into the government conditional on one's skills. In practice, governments – including the Indonesian government – use multi-stage selection processes that include restrictions on who can apply, entry exams and a selection on a combination of observables that are revealed through CVs and information provided during the application process. At the same time, corruption, clientelism and politically-motivated hiring can influence the selection at different stages (Colonnelli, Prem, and Teso 2020; Hanna and Wang 2017; Jia, Kudamatsu, and Seim 2015; Weaver 2021). This holds particularly so for Indonesia, where there is solid evidence that politics has an influence on selection into and promotion patterns within the Civil Service (Pierskalla and Sacks 2018), that there is widespread corruption among bureaucrats (Valsecchi 2016), and where civil service jobs have historically been sold in an auction-like fashion (Kristiansen and Ramli 2006). To study how "good" the government is *de facto* at selecting government workers, I define the reduced-form *selection rule* $s_t(z)$ as follows:

$$s_t(z) \equiv f_t(z|(\text{selected}|\text{applied})) = \frac{f_t(z|\text{applied} \cap \text{selected})}{f_t(z|\text{applied})} \tag{6}$$

where $f_t(z|\text{applied})$ gives the skill distribution of the applicant pool and $f_t(z|\text{applied} \cap \text{selected})$ the skill distribution of newly selected government workers. The selection rule is above unity at any skill level z if the state selects more people with these skills than would be expected under uniform drawing from the conditional skill distribution of applicants. An implicit assumption here is that there are more applicants than government workers to be selected, which is the empirically relevant setting for developing countries and many developed

Figure 8: Estimated government selection rule & changes in the selection rule over time

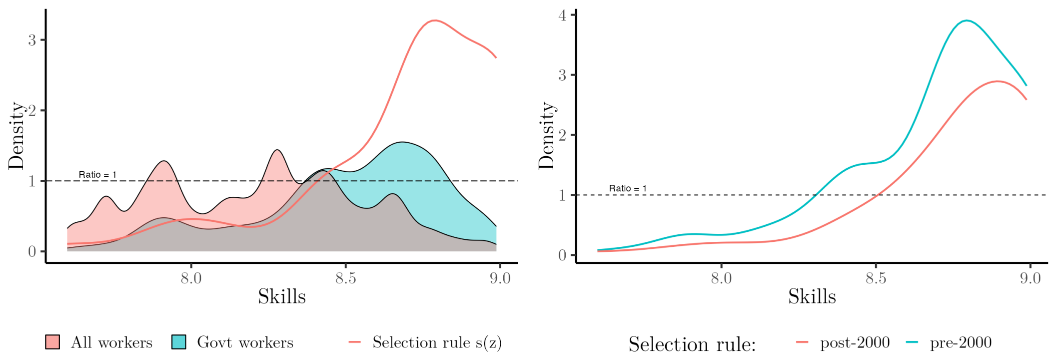

Notes: Both plots estimate selection rules using Equation 9 on the distribution of estimated skills for all workers and for the subset of government workers only, using one observation per worker. Estimates are based on a direct estimator of the density ratio using unconstrained Least-Squares Importance Fitting (uLSIF) taking standard values for the corresponding hyperparameters: lambda (0.2) and sigma (0.1). The left plot pools across all years and workers, while the right plot separately estimates selection rules for workers who entered the labor force in the year 2000 or before, or after 2000.

countries.[9] To estimate the density ratio, I use the sample of selected government workers for the numerator and the sample of all workers for the denominator. In this sense, I study an interesting limit case in which I assume all workers in the economy could potentially take a government job, interpreting the *selection rule* very broadly, serving as a useful benchmark that will be a good approximation in many developing country settings where interest in a safe government position is ubiquitous.[10]

The left plot in Figure 8 shows the estimated selection rule and the conditional skill density functions for government workers and all workers separately.[11] Two main points stand out: First, the government selection rule is mostly increasing. The Indonesian government is much less likely to pick low skilled workers as civil servants compared to their density in the data, which translates into a ratio below 1, and is more likely to pick high skilled workers leading

[9]In 2014, there were around 2.6 million applicants for 100k open civil service positions in Indonesia (see here as of April 22, 2024).

[10]In many contexts, defining the selection rule more narrowly in terms of applicants who actually filled out an application form or have signalled interest in a job would be more appropriate when such information is available. Since the IFLS data unfortunately does not measure "interest in a government job" nor whether a worker applied for a government job in the past, I study the broader measure. However, note that formal application requirements may already exclude highly skilled potential government workers such that a broader definition of the applicant pool can be more interesting in practice. The limit case I study is likely a lower bound for the selection rule because we expect low-skilled individuals to be more likely to take a government-sector job than high-skilled individuals who are more likely to have good outside options.

[11]To estimate the ratio, I use a direct estimator of the density ratio using unconstrained Least-Squares Importance Fitting (uLSIF) proposed by Hido et al. (2011), which is more robust than separately estimating skill distributions for nominator and denominator and then forming the ratio of the two. Similar to density estimation, there are many different direct estimators for density ratios proposed in the literature, which I found to perform very similarly for my application.

to density ratios above 1.[12] Second, the selection rule drops for the most skilled workers. If indeed everyone in the economy would take a government job, the point estimates imply that the government becomes worse at selecting the most qualified workers in the economy. The more likely explanation in this case is that the workers who would be most skilled in a government sector job are actually not interested in taking government jobs.

The right plot in Figure 8 tests whether the selection rule has changed over time and whether this may explain the decline in the skill premium of government workers. Quantitatively, I find that the estimated *selection rule* is remarkably constant over time, despite Indonesia moving from a highly centralized autocratic government with a regime-aligned bureaucracy (Hadiz and Robison 2013; Robison and Hadiz 2004) to a democratic system with a highly decentralized bureaucracy (Blunt, Turner, and Lindroth 2012; Brinkerhoff and Wetterberg 2013) over the entire period. The selection rule shifted downwards over time, in line with a general decline in government hiring (see Figure 2). If anything, the *selection rule* after the year 2000 is only starting to decrease at a higher level of skills, indicating a success of numerous bureaucracy reforms implemented across ministries in the 2000s and a change to a computer-assisted selection that has been shown to have reduced corruption at the selection stage (Kuipers 2023).

What then explains the decrease in the skill premium if it is not changes in the selection rule? The main driving force is a shift in the underlying skill distribution. While the selection rule stays roughly constant, the right tail of the skill distribution grows and more highly skilled workers end up working in the private sector and the rest of the economy.

Skill selection & the costs of government hiring waves

In this final subsection, I highlight the potential costs of fluctuations in government hiring for the skill selection of government workers. Besides its important policy implications, this final application can be seen as a test of the skill selection mechanism. The idea is that there is year-to-year variation in how much the government is hiring, driven in part by political cycles and the discrete nature of legislation. If hiring by the government happens disproportionately for workers at labor market entry, then one would expect differential exposure of birth cohorts to government hiring, as evidenced by observed differences in cohort-specific government employment shares. If skill selection as a mechanism has bite, then conditional on hiring for the same type of job and having similar cohort-to-cohort skill distributions, more treated cohorts should show lower average skills. That is, in years of disproportionately high hiring,

[12] The visible spikes in the skill distributions are due to the importance of the discrete educational background.

the government is more likely to draw from lower parts of the skill distribution, leading to lower average selected skills. The government thus misses out on skilled workers, hurting the quality of the overall government workforce. To test this idea formally, I regress changes in the cohort-level government employment share (denoted by GES) on changes in the skill premium. Throughout, I treat GES as exogenous here and leave more rigorous causal evidence using more plausible exogenous variation in government hiring to future work. I run the following two sets of regressions using data at the aggregated cohort-level and at the individual-cohort level:

$$\Delta \text{skill premium}_c = \alpha + \beta^C \Delta \text{GES}_c + \varepsilon_c \tag{7}$$

$$z_{ic} = \beta_0 \Delta \text{GES}_c + \beta_1 \mathbb{1}_{\{\text{govt?}\}_i} + \beta_2 \Delta \text{GES}_c \cdot \mathbb{1}_{\{\text{govt}\}_i} + \text{controls}_{ic} + \varepsilon_{ic} \tag{8}$$

Table 3 reports regression results and shows that changes in government hiring across cohorts indeed has strong effects on skill selection. Columns 1 & 2 report cohort-level results based on equation (7) using the skill premium in comparison to control workers and in comparison to private sector workers respectively. Despite the small sample size for the aggregate regressions, estimates are in line with increases in hiring leading to lower average selected skills. Columns 3 to 6 then show regression results at the individual-cohort-level. The coefficient β_2 then captures whether average skills are differentially lower among government workers in cohorts that were hired more strongly. Results consistently show that this is the case, which holds within the same job and when restricting to deviations from birth year trends (as in columns 5 & 6). Taking column 6 as the most preferred specification, the estimates imply that increasing the government employment share by its interquartile range (around 4.4pp), would lead to a drop in average government worker skills by around 2.4pp. As discussed in more detail in the next section, this effect accounts for roughly 10% of the overall government worker skill premium. Given more recent fluctuations in aggregate government worker hiring due to annual government hiring freezes in 2015 and 2016, it is possible that these detrimental effects on the selection of government workers have worsened since.

4.2 Application 2: Government wage setting

The second major application looks at the wages of government workers; specifically at the informativeness of government wages and whether government workers are overpriced compared to the private sector.

Table 3: The government skill premium and changes in government hiring intensity

Dependent Variables:	Change Skill Premium		Skill			
Model:	G vs. C	G vs. P	(3)	(4)	(5)	(6)
Variables						
Constant	0.004	0.008				
	(0.013)	(0.029)				
Change govt empl share	-0.608*	-1.304*	0.408***	0.343***	0.399***	0.342***
	(0.319)	(0.732)	(0.101)	(0.072)	(0.069)	(0.050)
Govt worker			0.530***	0.249***	0.523***	0.258***
			(0.031)	(0.013)	(0.031)	(0.014)
Change govt empl share × Govt worker			-0.632**	-0.626***	-0.536**	-0.547***
			(0.263)	(0.190)	(0.223)	(0.154)
Birth year					-0.120	-0.106
					(0.195)	(0.177)
Birth year square					0.000	0.000
					(0.000)	(0.000)
Fixed-effects						
Survey wave			Yes	Yes	Yes	Yes
Occupation				Yes		Yes
Sector				Yes		Yes
Fit statistics						
Observations	59	59	19,971	19,971	19,971	19,971
R^2	0.06	0.05	0.24	0.37	0.29	0.40
Within R^2			0.15	0.03	0.20	0.08

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: All data restricted to workers between the age of 25 and 58, the official age span of government workers. Columns 1 and 2 report results using cohort-level aggregated data with the outcome variable giving the skill premium of government workers compared to control workers (G vs. C) and private sector workers (G vs. P) respectively. Columns 3 to 6 report individual-cohort level regression results. Data in this case aggregates the worker panel by keeping the first observation when an individual is observed. Aggregation across panel waves increases precision and I control for panel waves in the individual-level regressions.

Are government wages uninformative?

The starting point of the estimation approach was the idea that government worker wages are potentially uninformative about underlying skills, because – as in the case of Indonesia – they follow rigid tenure schedules and allowances with little use of performance incentives or wage dispersion. In fact, I find more nuanced results; government wages are far from uninformative about underlying skills, but I also find that comparable private sector wages are more informative about skills. I show this with two different exercises.

First, I test the relative informativeness of government wages by predicting real hourly wages for government workers and comparable private sector workers using estimated skills. Table 8 in the Appendix reports these results. Comparing Columns 1 & 4 shows that government worker wages are indeed less informative about underlying skills than wages in comparable private sector jobs: the adjusted R^2 is 9.4% vs. 13.1%, a difference of roughly 40%. However, this relative ranking changes once one controls for experience profiles, given a far more deterministic government wage experience schedule. For example, even when enforcing the same experience coefficients as in the comparable private sector (columns 2 & 5), life cycle human capital of government workers predicts 20.5% of the variation versus 15.5% in the private sector.

The second test of the informativeness of government worker wages is to re-estimate government worker skills following all estimation stages but using government wages instead and then looking at the correlation across the two skill measures. As shown in Figure 10 in the Appendix, the correlation between the skill measures is high with an R^2 of around 78%. Assuming that baseline skills are estimated correctly, government wages are thus informative about underlying skills. However, as visible from Figure 10, the variance in skill estimates increases with skills, making skill estimates for the most skilled workers much less informative.

Is there a government wage premium?

The estimation approach allows for a direct measurement of the wage premium of government workers compared to similar jobs in the private sector: conditional on the same level of skills, do government workers earn higher wages? The answer is: yes. Government workers earn a large wage premium of at least 30%. Specifically, Table 4 reports different estimates for the wage premium restricting to government workers and the control group of reweighted private sector workers. Column 1 documents a large unconditional wage premium of 0.6 log points, which translates to roughly 80% higher real hourly wages. The wage premium almost halves when controlling for worker skills (column 4), pointing to strong positive selection on skills. The wage premium also generally reduces when focussing on within-job comparisons by introducing occupation and sector fixed effects (columns 3, 5 & 7). To avoid biasing estimated wage premia due to compositional changes over time – for example, if there are more control group workers in more recent periods – results for columns (2) to (8) include year fixed effects. At last, one may be interested in capturing the wage premium conditional on human capital that incorporates experience. Columns (6) & (7) report wage premia controlling for experience using the life-cycle skill measure $h_{i,e,t}$ that incorporates estimated human capital experience profiles. These estimates lead to the most conservative wage premia of around 0.25 log points, or a wage premium of roughly 30%. The last column considers changes in the wage premium over time and finds that the government wage premium is clearly increasing over time. The increasing wage premium for similar jobs in the private

Table 4: Regression results: Government wage premium

Dependent Variable:	Real hourly wage (log)							
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
Constant	8.562***							
	(0.005)							
Govt job?	0.601***	0.662***	0.537***	0.373***	0.386***	0.236***	0.253***	0.002
	(0.006)	(0.054)	(0.059)	(0.043)	(0.046)	(0.044)	(0.046)	(0.037)
Skill				0.948***	0.705***			
				(0.057)	(0.074)			
Life-cycle skill						1.214***	1.059***	1.057***
						(0.045)	(0.054)	(0.055)
Govt job X year								0.020***
								(0.002)
Fixed-effects								
Year		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation			Yes		Yes		Yes	Yes
Sector			Yes		Yes		Yes	Yes
Fit statistics								
Observations	136,012	136,012	135,106	136,012	135,106	136,012	135,106	135,106
R^2	0.06	0.10	0.18	0.19	0.21	0.25	0.26	0.27
Within R^2		0.07	0.05	0.16	0.09	0.22	0.15	0.15

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Sample restricted to government workers and control group (private sector workers reweighted by jobs in government). Standard errors are clustered at the year level (except for column 1 which assumes IID standard errors). Adding occupation and sector fixed effects reduces the sample size slightly due to missings in the occupation variable.

sector may explain the results of the previous section, namely why the skill premium in these jobs has not deteriorated as for the overall private sector.

What is driving this large observed wage premium? It turns out that an important driver of the government wage premium is the absence of a gender wage gap in the public sector. Women in Indonesia's private sector face a large real hourly wage penalty compared to men in the private sector. Table 5 reports a raw gender wage gap of 0.433 log points, or a roughly 35% (!) wage cut for women compared to men in comparable jobs in the private sector (column 2), which decreases slightly after controlling for the same skills (column 3) or life-cycle skills (column 4). This gender wage is also stable over time (coefficient "Male x year" in column 5). The gender wage gap in the public sector, on the other hand, is much smaller, at about 1/4 of the wage gap in the private sector after controlling for the same job, experience and life-cycle skills (column 4 and using the sum of the coefficients for "Male" & "Govt worker x Male"). That is, in the absence of a differential gender wage gap in the private

Table 5: Regression results: Government wage premium & gender gap

Dependent Variable:	Real hourly wage (log)				
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Constant	8.368***				
	(0.010)				
Govt job?	0.866***	0.752***	0.565***	0.429***	0.171***
	(0.012)	(0.063)	(0.053)	(0.055)	(0.041)
Male	0.273***	0.433***	0.399***	0.320***	0.324***
	(0.012)	(0.023)	(0.020)	(0.018)	(0.032)
Govt job? × Male	-0.377***	-0.320***	-0.266***	-0.251***	-0.241***
	(0.014)	(0.021)	(0.022)	(0.026)	(0.024)
Skill			0.698***		
			(0.073)		
Life-cycle skill				1.035***	1.033***
				(0.057)	(0.057)
Govt job X year					0.020***
					(0.002)
Male X year					-0.001
					(0.002)
Fixed-effects					
Year		Yes	Yes	Yes	Yes
Occupation		Yes	Yes	Yes	Yes
Sector		Yes	Yes	Yes	Yes
Fit statistics					
Observations	136,012	135,106	135,106	135,106	135,106
R^2	0.06	0.19	0.22	0.27	0.27
Within R^2		0.06	0.10	0.15	0.16

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Sample restricted to government workers and control group (private sector workers reweighted to hold comparable jobs than government workers). Standard errors are clustered at the year level (except for column 1 which assumes IID standard errors). Adding occupation and sector fixed effects reduces the sample size slightly due to missings in the occupation variable.

sector, the government wage premium would be roughly 30% lower (the sum of coefficients "Govt job" & "Govt job x Male" in column 4 in comparison to the coefficient for "Govt job" in Table 4 column 7).

5 Extensions

This section discusses main extensions of the skill estimation approach.

5.1 More flexible skill-experience profiles

The estimation approach for government worker skills in this paper allows for far more flexibility in skill experience profiles. A natural extension allows not only for a measure of skills at labor market entry, but also for individual-specific learning capabilities. Such a measure would then also allow to consider selection of workers on learning capabilities. A similar alternative would be to allow for multi-dimensional individual skill estimates. In both cases, the key to estimation would be to construct multiple individual-specific fixed effects and then separately regress these on skill- or learning-related observables. To see how this would look, assume that human capital $h_{i,e,t}$ evolves according to the following factor structure:

$$h_{i,e,t} = z_i * e^{g_i * \delta_e} \tag{9}$$

where z_i are individual time-fixed skills at labor market entry, g_i are individual-specific learning capabilities and δ_e are arbitrary experience-fixed effects. Government worker skills are defined as $z_i = h_{i,0,t}$, assuming that $\delta_0 = 0$, and the correlation between the two different fixed effects is unrestricted. The factor structure allows individual-specific skills and individual-specific wage-experience profiles, but restricts the shape of the latter to a factor structure. This can be seen as a generalization of earlier panel data methods such as the Within estimator which allows only individual-specific levels. The factor structure can match well empirically observed concave wage-experience profiles that differ in slope across individuals. Specifically, it has been well documented that wage-experience profiles are steeper for highly educated than for less educated individuals and that the variance of the slope of wage-experience profiles is increasing with observed skills such as education (Primiceri and Van Rens 2009; Lagakos et al. 2018). Alternatively, one can allow for multiple dimensions of skills by incorporating multiple factors, which are usually restricted to be orthogonal (see Ahn, Lee, and Schmidt 2013; Bai 2009). In Appendix A.5, I provide a more detailed practical guide of how to estimate multiple noisy fixed effects in this setup and discuss Monte Carlo evidence that shows its performance in realistic empirical settings.

5.2 Further heterogeneity in skill prices or jobs

An important limitation of the approach in this paper is that skill estimation requires the assumption that there is a single skill price in the skill estimation sample for a comparable job with the same skill requirements. While subsequent analyses conditioned on narrower within-job comparisons, the skill estimation itself relies on comparing across these jobs to value skills. If one is worried about this assumption, an alternative approach is to specify a model that maps from noisy individual fixed effects to skill-related observables, while controlling for further group fixed effects. These groups could be jobs or also geographic areas to capture different skill prices across locations. The same fixed effects should then be used in any downstream empirical analysis to ensure that results are only based on variation within fixed effects. Another alternative to allow for more price variation is to directly compute group-specific skill prices following a group-specific flat-spot identification. Again, this gives separate normalizations across groups, so that levels in skills cannot be compared across groups anymore, requiring group-specific fixed effects in later analyses.

6 Conclusion

This paper provided a new approach to estimate government worker skills using residualized wages in comparable jobs in the private sector, relating these to skill-related observables using Machine Learning tools and then predicting government worker skills out-of-sample. I then showed two main applications drawing on rich Indonesian household-level panel data. First, I showed evidence for the selection of government workers. Despite growing absolute skills, relative skills of government workers compared to the overall workforce and private sector workers in particular declined consistently over the past 30 years. I linked this finding to the difficulty of the Indonesian government to attract the workers with the highest skills to the government. Furthermore, I showed evidence for the detrimental effect of government hiring cycles on the selection of government workers. The evidence is consistent with the idea that in years of outsized hiring, the government needs to move down the skill distribution of the applicant pool to fill all government positions, leading to lower average skills. In the second main application, I looked at government wage setting and showed that the Indonesian government pays a wage premium of at least 30% conditional on skills, about 1/3 of which is driven by the large gender wage gap in Indonesia's private sector.

A good sign of a new estimation approach is that it raises many interesting questions – both conceptual and theoretically – that can now be studied more rigorously: For example, what are the output or welfare costs of government hiring cycles? Or what drives the relative decline in government skills and does this go in hand with a relative decline in state capacity versus private sector capacity over the course of development? All of these questions are particularly well-suited for future structural work on the functionings of bureaucracies, for which the estimated government skills in this paper can directly be used as inputs. This is just one promising direction where the novel method proposed in this paper could be used to answer outstanding questions in the literature.

References

Ahn, Seung C., Young H. Lee, and Peter Schmidt. 2013. "Panel Data Models with Multiple Time-Varying Individual Effects." *Journal of Econometrics* 174 (1): 1–14.

Anandari, I., and Chaikal Nuryakin. 2019. "The effect of risk preference on choice between public and private sector employment in Indonesia." *International Journal of Business & Society* 20.

Ashraf, Nava, Oriana Bandiera, Edward Davenport, and Scott S Lee. 2020. "Losing Prosociality in the Quest for Talent? Sorting, Selection, and Productivity in the Delivery of Public Services." *American Economic Review* 110 (5): 1355–94.

Athey, Susan, and Guido W Imbens. 2019. "Machine Learning Methods That Economists Should Know About." *Annual Review of Economics* 11 (1): 685–725.

Attanasio, Orazio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina. 2020. "Estimating the Production Function for Human Capital: Results from a Randomized Controlled Trial in Colombia." *American Economic Review* 110 (1): 48–85.

Bai, Jushan. 2009. "Panel Data Models with Interactive Fixed Effects." *Econometrica* 77 (4): 1229–79.

Banerjee, Abhijit, Lakshmi Iyer, and Rohini Somanathan. 2007. "Public Action for Public Goods." *Handbook of Development Economics* 4: 3117–54.

Besley, Timothy, Olle Folke, Torsten Persson, and Johanna Rickne. 2017. "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden." *American Economic Review* 107 (8): 2204–42.

Best, Michael Carlos, Jonas Hjort, and David Szakonyi. 2023. "Individuals and Organizations as Sources of State Effectiveness." *American Economic Review* 113 (8): 2121–67.

Bhavnani, Rikhil R., and Alexander Lee. 2019. "Does Affirmative Action Worsen Bureaucratic Performance? Evidence from the Indian Administrative Service." *American Journal of Political Science.*

Biasi, Barbara. 2021. "The Labor Market for Teachers Under Different Pay Schemes." *American Economic Journal: Economic Policy* 13 (3): 63–102.

Blunt, Peter, Mark Turner, and Henrik Lindroth. 2012. "Patronage's Progress in Post-Soeharto Indonesia." *Public Administration and Development* 32 (1): 64–81.

Bowlus, Audra J., and Chris Robinson. 2012. "Human Capital Prices, Productivity, and Growth." *American Economic Review* 102 (7): 3483–3515.

Brinkerhoff, Derick W., and Anna Wetterberg. 2013. "Performance-Based Public Management Reforms: Experience and Emerging Lessons from Service Delivery Improvement in Indonesia." *International Review of Administrative Sciences* 79 (3): 433–57.

Chaudhury, Nazmul, Jeffrey Hammer, Michael Kremer, Karthik Muralidharan, and F. Halsey

Rogers. 2006. “Missing in Action: Teacher and Health Worker Absence in Developing Countries.” *Journal of Economic Perspectives* 20 (1): 91–116.

Chetty, Raj, John N Friedman, and Jonah E Rockoff. 2014. “Measuring the Impacts of Teachers i: Evaluating Bias in Teacher Value-Added Estimates.” *American Economic Review* 104 (9): 2593–2632.

Chong, Alberto, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2014. “Letter Grading Government Efficiency.” *Journal of the European Economic Association* 12 (2): 277–98.

Colonnelli, Emanuele, Mounu Prem, and Edoardo Teso. 2020. “Patronage and Selection in Public Sector Organizations.” *American Economic Review* 110 (10): 3071–99.

Dal B'o, Ernesto, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne. 2017. “Who Becomes a Politician?” *The Quarterly Journal of Economics* 132 (4): 1877–1914.

Dal B'o, 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 (3): 1169–1218.

Decarolis, Francesco, Leonardo M Giuffrida, Elisabetta Iossa, Vincenzo Mollisi, and Giancarlo Spagnolo. 2020. “Bureaucratic Competence and Procurement Outcomes.” *The Journal of Law, Economics, and Organization* 36 (3): 537–97.

Deserranno, Erika, Aisha Nansamba, and Nancy Qian. 2024. “The Impact of NGO-Provided AID on Government Capacity: Evidence from Uganda.” *Journal of the European Economic Association*, jvae029.

Duflo, Esther. 2001. “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment.” *American Economic Review* 91 (4): 795–813.

———. 2004. “The Medium Run Effects of Educational Expansion: Evidence from a Large School Construction Program in Indonesia.” *Journal of Development Economics* 74 (1): 163–97.

Estrada, Ricardo. 2019. “Rules Versus Discretion in Public Service: Teacher Hiring in Mexico.” *Journal of Labor Economics* 37 (2): 545–79.

Finan, Frederico, Benjamin A. Olken, and Rohini Pande. 2017. “The Personnel Economics of the Developing State.” In *Handbook of Economic Field Experiments*, edited by Esther Duflo and Abhijit V. Banerjee, 2:467–514. Elsevier.

Fisman, Raymond. 2001. “Estimating the Value of Political Connections.” *American Economic Review* 91 (4): 1095–1102.

Hadiz, Vedi R., and Richard Robison. 2013. “The Political Economy of Oligarchy and the Reorganization of Power in Indonesia.” *Indonesia*, no. 96: 35–57.

Hamory, Joan, Marieke Kleemans, Nicholas Y Li, and Edward Miguel. 2021. "Reevaluating Agricultural Productivity Gaps with Longitudinal Microdata." *Journal of the European Economic Association* 19 (3): 1522–55.

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

Hido, Shohei, Yuta Tsuboi, Hisashi Kashima, Masashi Sugiyama, and Takafumi Kanamori. 2011. "Statistical Outlier Detection Using Direct Density Ratio Estimation." *Knowledge and Information Systems* 26 (2): 309–36.

Horhoruw, Maggy, Tina George Karippacheril, Wahyu Sutiyono, and Theo Thomas. 2013. "Transforming the Public Sector in Indonesia: Delivering Total Reformasi." In *Paper Presented IRSPM Conference,'Contradictions in Public Management Managing in Volatile Times', University of Rome, Tor Vergata.* Vol. 16.

Hornik, Kurt, Maxwell Stinchcombe, and Halbert White. 1989. "Multilayer Feedforward Networks Are Universal Approximators." *Neural Networks* 2 (5): 359–66.

Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *American Economic Review* 101 (7): 2923–54.

Jia, Ruixue, Masayuki Kudamatsu, and David Seim. 2015. "Political Selection in China: The complementary roles of connections and performance." *Journal of the European Economic Association* 13 (4): 631–68. https://doi.org/10.1111/jeea.12124.

Keane, Michael P., and Kenneth I. Wolpin. 1997. "The Career Decisions of Young Men." *Journal of Political Economy* 105 (3): 473–522.

Kristiansen, Stein, and Muhid Ramli. 2006. "Buying an Income: The Market for Civil Service Positions in Indonesia." *Contemporary Southeast Asia*, 207–33.

Kuipers, Nicholas. 2023. "Failing the Test: The Countervailing Attitudinal Effects of Civil Service Examinations." *American Political Science Review* 117 (3): 891–908.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman. 2018. "Life Cycle Wage Growth Across Countries." *Journal of Political Economy* 126 (2): 797–849.

Lagakos, David, and Martin Shu. 2023. "The Role of Micro Data in Understanding Structural Transformation." *Oxford Development Studies* 51 (4): 436–54.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market." *American Economic Review* 112 (1): 169–212.

Magnac, Thierry, Nicolas Pistolesi, and S'ebastien Roux. 2018. "Post-Schooling Human Capital Investments and the Life Cycle of Earnings." *Journal of Political Economy* 126 (3): 1219–49.

Martinez-Bravo, Monica, Priya Mukherjee, and Andreas Stegmann. 2017. "The Non-Democratic Roots of Elite Capture: Evidence From Soeharto Mayors in Indonesia." *Econometrica* 85 (6): 1991–2010.

McLeod, Ross H. 2006. "Private Sector Lessons for Public Sector Reform in Indonesia." *Agenda: A Journal of Policy Analysis and Reform*, 275–88.

———. 2008. "Inadequate Budgets and Salaries as Instruments for Institutionalizing Public Sector Corruption in Indonesia." *South East Asia Research* 16 (2): 199–223.

Meghir, Costas, and Luigi Pistaferri. 2011. "Earnings, Consumption and Life Cycle Choices." In *Handbook of Labor Economics*, 4:773–854. Elsevier.

Pesaran, M. Hashem. 2006. "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure." *Econometrica* 74 (4): 967–1012.

Pierskalla, Jan H., Adam Lauretig, Andrew Rosenberg, and Audrey Sacks. 2020. "Democratization and Representative Bureaucracy - An Analysis of Promotion Patterns in Indonesia's Civil Service, 1980-2015." *American Journal of Political Science* forthcoming.

Pierskalla, Jan H., and Audrey Sacks. 2018. "Personnel Politics: Elections, Clientelistic Competition and Teacher Hiring in Indonesia." *British Journal of Political Science*, 1–23.

Polachek, Solomon W. 2008. *Earnings over the Lifecycle: The Mincer Earnings Function and Its Applications*. Vol. 16. Now Publishers Inc.

Primiceri, Giorgio E., and Thijs Van Rens. 2009. "Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality." *Journal of Monetary Economics* 56 (1): 20–39.

Rasul, Imran, and Daniel Rogger. 2018. "Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service." *The Economic Journal* 128 (608): 413–46. https://doi.org/10.1111/ecoj.12418.

Robison, Richard, and Vedi Hadiz. 2004. *Reorganising Power in Indonesia: The Politics of Oligarchy in an Age of Markets*. Routledge.

Sanders, Carl, and Christopher Taber. 2012. "Life-Cycle Wage Growth and Heterogeneous Human Capital." *Annual Review of Economics*, no. 4.

Shen, Xiaoxi, Chang Jiang, Lyudmila Sakhanenko, and Qing Lu. 2019. "Asymptotic Properties of Neural Network Sieve Estimators." *arXiv Preprint arXiv:1906.00875*.

Strauss, John, Firman Witoelar, and Bondan Sikoki. 2016. *The Fifth Wave of the Indonesia Family Life Survey: Overview and Field Report*. RAND.

Taber, Christopher, and Rune Vejlin. 2020. "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market." *Econometrica* 88 (3): 1031–69.

Valsecchi, Michele. 2016. "Corrupt Bureaucrats: The Response of Non-Elected Officials to Electoral Accountability." *Job Market Paper*.

Weaver, Jeffrey. 2021. "Jobs for Sale: Corruption and Misallocation in Hiring." *American*

Economic Review 111 (10): 3093–3122.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT press.

A Appendix

A.1 Further details on the selection and hiring process of government workers

Entering the Civil Service is remarkably well-defined in Indonesia despite large changes to the Civil Service over the time period of interest as applicants run through a centralized application process. Applicants apply to the specific position or district they are interested in, but still run through a centralized application process and due to frequent rotations and across-country stationing, are likely to end up with a position somewhere else than where they applied if they are admitted. Formal requirements of applying to the Civil Service are that individuals have to be between 18-35 years old, never been imprisoned, not be a member of a political party, be in good physical and mental health and be willing to work in any region in Indonesia. For each job opening there are then additional educational requirements, which are set by the district and which often mean an undergraduate diploma. Since 2012, all applicants have to go through a civil servant enrollment test (*CPNS*), which includes three parts: an administrative selection, a basic competence test (*Seleksi Kompetensi Dasar*) recently administered via a computer-assisted test and a specific field competence selection (*Seleksi Kompetensi Bidang*).[13]

Solely based on aggregate numbers, obtaining a Civil Service job is difficult. In 2014, prior to a 4-year public sector employment moratorium, there were more than 2.6 million applicants for 100,000 available positions, which translates into an acceptance rate of slightly below 4% (see: Anandari and Nuryakin 2019). This is similar to the 1-5% acceptance rates reported in Kristiansen and Ramli (2006) for two Indonesian regions in the early 2000s.[14]

In practice, it is unclear how well the recruitment system in place selects qualified candidates and how this changed over time. Horhoruw et al. (2013) notes that it is unclear whether reform processes since 2001 have actually led to an improvement of hiring practices beyond just a few reform-minded institutions. In Pierskalla and Sacks (2018), the authors draw on teacher censuses to show that changes in the political system after 1998 actually had negative effects on public hiring. They find that increased political competition gave local elites an incentive to use their discretionary control over state hiring to increase patronage efforts as

[13] Under the Suharto regime, the Civil Service system was organized as a military-type organisation where new recruits were not differentiated other than by level of education, which introduced generalist civil servants and abolished further specializations within the bureaucracy (McLeod 2006). Recent reforms have tried to reverse this.

[14] Note, the percentage of acceptances has increased in the author's sample, which would be in line with the model on the evolution of state capacity.

evidenced by election-related increases in the number of contract teachers on local payrolls and increases in civil service teacher certifications. At the same time, Pierskalla et al. (2020) use data on the universe of civil servants to show that civil servants with a postgraduate education are twice as likely to be promoted after 1999 in comparison to before, indicating a combination of composition changes and more performance-related promotion patterns.

Kristiansen and Ramli (2006) draw on in-depth qualitative and quantitative evidence from interviews and focus groups with a non-representative sample of 60 civil servants in two areas of Indonesia to document that personal ties and nepotism are often named as primary reasons for hiring. Moreover, the selling of government jobs is widely practiced. Kristiansen and Ramli (2006) document that all respondents paid for their first Civil Service position and that the average reported price for these jobs is around 2.5 times the official annual initial salary offered. This is slightly higher than the 17 months of salary reported recently in Weaver (2021).[15] There is also some evidence that the average real price for a government position has slowly increased between 1995-2004. Prices are positively correlated with the salary of the job (which in turn is mechanically tied to the education level of the civil servant) and seem to be positively correlated with the ease of rent-seeking possibilities in the specific job offered.[16] This evidence on prices is in line with a competitive auction price for government sector jobs as found in Weaver (2021). In the end, it is unclear how these unlawful hiring practices perform with respect to selecting the most qualified candidates as this depends on the correlation between quality and the ability to pay for a job or the probability of knowing someone important in the bureaucracy. Interestingly, Weaver (2021) finds that for the context they look at, this correlation is highly positive so that the selling of government sector jobs actually leads to a good selection rule in terms of quality of the new hires.

[15] The author does not share the country of study to provide additional security for their survey respondents.

[16] Among the usual rent-seeking possibilities are various forms of contract kickbacks, payment from staff in exchange for positions and hiring on projects, loan accounts structured to earn interest by the agency, provision of ghost services, inflated invoicing in collusion with contractors, procedures for tax avoidance, irregular payments for health and education services, bribes to police officers and judges, and speed money to obtain formal papers and permits (World Bank 2003; Vian 2005; Chapman 2005; Azfar 2005).

A.2 Proof of Proposition 2.1

I start with the **flat-spot identification**. Unbiasedness is given by:

$$\mathbb{E}_{i\in\text{FP}}[w_{i,e,t} - w_{i,e-1,t-1}] = p_t - p_{t-1} + \underbrace{\mathbb{E}_{i\in\text{FP}}[(h_{i,e,t} - h_{i,e-1,t-1})]}_{\text{=0 due to flat spot assumption}} + \underbrace{\mathbb{E}_{i\in\text{FP}}[\epsilon_{i,e,t} - \epsilon_{i,e-1,t-1}]}_{=0} \quad (10)$$

Consistency of the estimator is given by:

$$\frac{1}{N_{FP}} \sum_{i\in\text{FP}} [w_{i,e,t} - w_{i,e-1,t-1}] \xrightarrow{N\to\infty} p_t - p_{t-1} + \underbrace{\mathbb{E}_{i\in\text{FP}}[(h_{i,e,t} - h_{i,e-1,t-1})]}_{\text{=0 due to flat spot assumption}} + \underbrace{\mathbb{E}_{i\in\text{FP}}[\epsilon_{i,e,t} - \epsilon_{i,e-1,t-1}]}_{=0} \quad (11)$$

The estimator combines estimated year-to-year changes in prices, giving an unbiased and consistent estimate of the entire price path (up to a level of normalization).

Given unbiased and consistent estimates of the skill price process, one can construct an unbiased estimate of individual skills using a standard-within estimator for:

$$\widetilde{w_{i,e,t}} - \overline{\widetilde{w_{i,\cdot}}} = \delta_0 * \left(exp_{i,e,t} - \overline{exp_{i,\cdot}}\right) + \delta_1\left(exp^2_{i,e,t} - \overline{exp^2_{i,\cdot}}\right) + \left(\epsilon_{i,e,t} - \overline{\epsilon_{i,\cdot}}\right) \quad (12)$$

However, as is known as the incidental parameters problem, z_i cannot be consistently estimated from the above as long as $T \not\to \infty$ (e.g. see: Wooldridge 2010, Chp. 10).

Table 6: Regression results: Experience profile

Dependent Variable:	log wage (deflated by skill price)	
Model:	(1)	(2)
Variables		
Constant	8.193***	
	(0.015)	
experience	0.029***	0.031***
	(0.002)	(0.004)
experience square	-0.001***	-0.001***
	(0.000)	(0.000)
Fixed-effects		
Individual		Yes
Fit statistics		
Observations	20,538	20,538
R^2	0.01	0.65
Within R^2		0.01

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Sample are comparison workers (private sector workers in social services).

A.3 Additional results for skill estimation

Table 7: Relative variable importance

Variable	Importance
HigherEduc	100.000000
PrimaryEduc	87.501313
writeLetterIndo	45.192794
SeniorSecondEduc	37.335084
speakIndonesian	35.646910
CountWordRecall	15.582123
RelativeTotalScore	15.163991
RavenIQ	12.486136
RiskPref1	9.377688
RelativeIndonScore	8.609193
JuniorSecondEduc	8.514743
CountBackwards	8.267882
Big5Neu	6.720625
RelativeMathScore	5.824077
MathIQ	3.660663
Big5Ext	3.504978
Big5Open	3.114057
RiskPref2	2.753223
WordAbil	2.403554
RiskPref3	1.260320

Details:

Based on best-performing GBM algorithm and restricting to top 20 variables. Variable importance is based on traversing the tree and recording how much the metric (R2 here) changes every time a given variable is used for splitting. One then takes the average reductions across all base-learners for each variable and normalizes the most important variable to 100

Figure 9: Government worker skills across cohorts (keeping government jobs fixed)

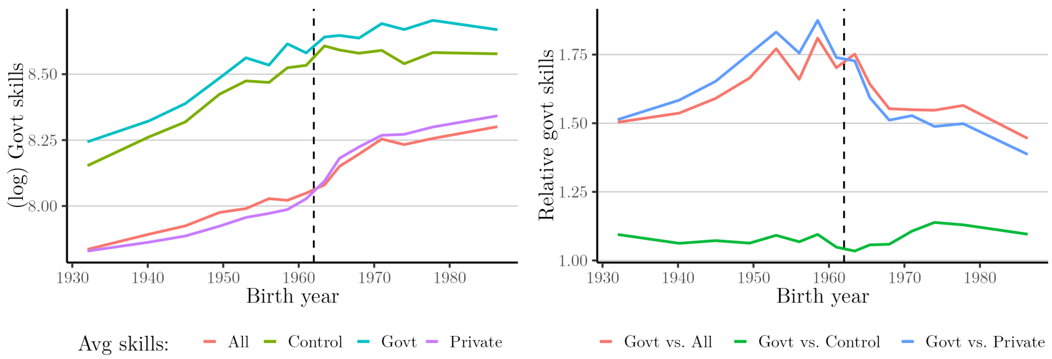

Notes: Results are based on skill estimates using baseline specification with private sector jobs comparable to the government (Control group in Figure) as estimation sample and GBM estimator as ML algorithm. Relative skills are in log differences. Data is pooled across all waves of the IFLS and then plotted by (binned) cohort. Cohort bins are determined by equal-sized bins in pooled data. Sample restricted to working ages of government workers (25-58). For government skill estimate across cohorts, condition on job sector and occupation fixed effects, holding both constant for estimates of average skill changes over cohorts. Dotted line denotes the first cohort that was treated by the INPRES school construction program studied in Duflo (2001).

A.4 Additional empirical results

This part of the Appendix provides additional results for Section 4 that are in part referenced in the main text.

Figure 9 shows absolute and relative government worker skills across (binned) cohorts, now additionally holding the composition of government jobs fixed. The difference between Figure 9 & Figure 7 thus gives the importance of compositional changes in the skill intensity of government jobs over time in explaining changes in government worker skills. Differences are very small in magnitude, indicating that changes in the composition of government jobs in terms of sector and occupation cannot explain large changes in the relative skills of government workers.

Table 8 tests the informativeness of government wages in comparison to wages in similar private sector jobs.

Table 8: Regression results: Informativeness of government wages

Dependent Variable:	Real hourly wage (log)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Constant	0.993***	-2.680***	-1.609***	-0.491***	-2.007***	-2.154***
	(0.222)	(0.204)	(0.215)	(0.066)	(0.070)	(0.072)
Skill	0.955***		1.157***	1.095***		1.243***
	(0.026)		(0.025)	(0.008)		(0.008)
Skill life cycle		1.311***			1.213***	
		(0.023)			(0.008)	
Experience			0.078***			0.044***
			(0.003)			(0.001)
Experience square			-0.001***			-0.001***
			(0.000)			(0.000)
Fit statistics						
Observations	13,130	13,130	13,130	122,882	122,882	122,882
R^2	0.09	0.20	0.22	0.13	0.16	0.16
Adjusted R^2	0.09	0.20	0.22	0.13	0.16	0.16

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Columns 1-3 focus on public worker wages only, while Columns 4-6 focus on comparison workers (private sector workers reweighted to hold similar jobs to government workers). Columns 1 and 4 predict wages using skills only. Columns 2, 3, 5 and 6 instead use measures of life cycle skills. Columns 2 and 5 enforce previously estimated experience coefficients, while Columns 3 and 6 reestimate experience profiles separately for the private sector and government.

A.5 Estimation details: factor structure for experience profile

The factor structure introduced in Section 5.1 can be estimated using the two-stage estimator of Pesaran (2006).[17] Pesaran (2006) proposes to use cross-sectional averages to estimate experience factors δ_e. For this, one can first demean the series to get:

$$\tilde{\tilde{w}}_{i,e} \equiv \tilde{w}_{i,e} - \tilde{w}_{i,\bullet} = g(z_i)(\delta_e - \delta_\bullet) + (\epsilon_{i,e} - \epsilon_{i,\bullet}) \tag{13}$$

Using the economic structure of the problem gives $\delta_0 = 0$, so that identification of $(\delta_e - \delta_\bullet)$ also separately identifies the two terms. Using the fact that in the assumed data-generating

[17] Using Monte Carlo simulations, I checked that the Pesaran (2006) estimator performs much better in estimating the factor structure than alternative estimators such as a Within estimator and two different Concentrated Maximum Likelihood estimators. This performance even holds with unbalanced and potentially non-stationary panel data as observed in real-world applications. I am happy to share these results upon request. Based on the existing econometric literature, this is a novel result, because I am not aware of any studies that have looked at the performance of factor model estimators for the individual-level estimates itself (in contrast to treatment effect parameters that are estimated in the presence of individual-level effects).

Figure 10: Correlation of baseline skill estimates and skill estimates based on government wages

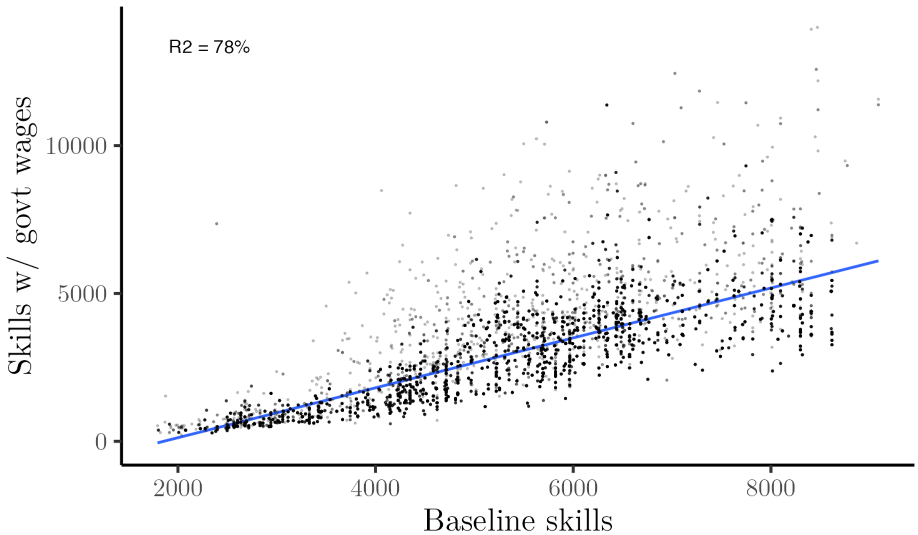

Notes: Baseline skill estimates are based on (reweighted) private sector workers as estimation sample and GBM estimator as ML algorithm (as explained in text). Skill estimates based on government wages go through the same estimation steps using all government workers for estimation sample instead (that is, the same skill price estimator and the GBM estimator).

process: $lim_{n\to\infty}\tilde{\tilde{w}}_{\bullet,e} = \overline{g(z)}(\delta_e - \delta_\bullet)$, we can write:

$$\frac{\tilde{\tilde{w}}_{\bullet,e}}{\tilde{\tilde{w}}_{\bullet,0}} \to 1 - \frac{\delta_e}{\delta_\bullet} \tag{14}$$

$\tilde{\tilde{w}}_{\bullet,e}$ gives as many equations as there are experience levels. However, with the previous restriction of $\delta_0 = 0$, we lose one restriction which requires to directly use: $\tilde{\tilde{w}}_{\bullet,e} \approx \overline{g(z)}(\delta_e - \delta_\bullet)$. Given that we are not directly interested in the estimates of δ nor $g(z_i)$, we can instead choose any non-zero normalization to obtain the same estimates of individual skills $log(z_i)$. For any normalization of $\overline{g(z)}$, we obtain an estimate of δ. We can then include these in the following experience-series regression for each individual to estimate $log(z_i)$ and $g(z_i)$:

$$\tilde{w}_{i,e} = log(z_i) + \widehat{\delta}_e g(z_i) + \epsilon_{i,e} \tag{15}$$

The individual-level skill estimates are the constant of the regression, giving the estimate $\widehat{log(z_i)}$. The derived estimate of private sector skills is generically inconsistent in a panel with fixed T due to the incidental parameters problem. This is generally true of any estimator that gives individual-specific estimates. For example, a standard Within-estimator also gives inconsistent estimates of the individual levels as long as the time dimension is not tending to infinity. Intuitively, the setup only allows us to extract a noisy signal from the even noisier wage data.