Returns to Skills: Does Location Matter? Evidence from India's IT Services Industry

Gagandeep Sachdeva[1]

Abstract

This paper attempts to investigate whether returns to different levels of cognitive skills systematically vary for locations where the Indian Information-Technology (IT) Services Industry is intensely concentrated. Using data from Census of India (2001 and 2011), Indian districts are classified into 'IT-Clusters' and 'non-clusters.' Furthermore, a triple-differences framework is employed on wage data from India's National Sample Survey's Employment and Unemployment Data. I find evidence that the skill premium in wages (i.e, the difference in wages between high-skilled and low-skilled workers) is higher in IT-clusters relative to non-clusters. Furthermore, there is mixed evidence to suggest the possibility of positive spillovers in wages for low-skilled workers associated with being located in IT-Clusters. There is no evidence to suggest that the difference in skill premium between IT-Clusters and non-clusters is increasing between 2004-05 and 2011-12[2]

Keywords: Skill-Biased Technical Progress, Spatial Inequality, Class Inequality, IT-Services Clusters

JEL Codes: J310, J240, O330, O120, R120.

[1]gagandeep@igidr.ac.in

[2]This paper is adapted from my MPhil thesis written at the Indira Gandhi Institute of Development Research between Jan-Dec 2018. I thank my supervisor Dr Sudha Narayanan for her comments and suggestions at various stages of this work. The usual disclaimer applies.

1 Introduction

A central theme in labour economics involves understanding what causes differences in the earnings of individuals. Deriving an explanation from neoclassical production theory is fairly straightforward; differences in wages simply reflect differences in the productivity of individuals. Individuals with varying productivities can be modeled as distinct inputs in production functions, with wage differences being explained by their respective marginal products.

Under this framework, understanding how wage differences evolve over time, therefore, involves inquiring into factors that may selectively influence the productive capacity of one group of workers vis-à-vis other groups. Typically, these factors are modeled as technical change in production functions. The idea of Hicks-neutrality (Hicks, 1963) suggests that the nature of technical change is such that it increases the productivity of inputs in equal proportions, or that technical change is factor neutral. If that were the case, one would correspondingly expect proportionate movements in wages for each input. However, a large body of literature empirically rejects this notion; the experience of the United States in the 20th century has been characterized by improvements in technology, increasing supply of skilled labour, and an increase in wage inequality between skilled and unskilled labour (Acemoglu, 2002).

Labour economists over the last few decades have broadly arrived at a consensus regarding the skill-biased nature of technical progress. Skill-biased technical change (hereafter SBTC), in terms of a production function, implies that improvements in technology increases the productivity of skilled workers disproportionately more than that of unskilled workers. The idea is summarized by Acemoglu, 2002, pg.7:

"The recent consensus is that technical change favors more skilled workers, replaces tasks previously performed by the unskilled, and exacerbates inequality."

Modeling the consequences of SBTC for inequality involves, at its core, modeling wages as returns to skills. A class of models that Acemoglu refers to as 'the canonical model' (Katz and Murphy, 1992, Card and Lemieux, 2001, among others) attempt to model this relationship. The canonical model expresses the skill premium to wages in terms of the degree to which SBTC increases the productivity of skilled workers relative to unskilled workers, and the elasticity of substitution between skilled and unskilled workers in production (assuming them to be strategic substitutes). The family of models that are characterized under the canonical model perform well empirically in accounting for changes in income distribution in the United States (Acemoglu and Autor, 2011).

The importance of skill-biased technical progress in explaining movements in wage inequality across class (as characterized by education and skill) is well-established in literature. However, SBTC may also lead to movements in spatial inequality, if the rate at which technical change occurs differs spatially. Baum-Snow et al. (2018) attempt to investigate the role of spatial differences in SBTC in explaining movements in urban inequality. They posit that urban agglomerations are associated with higher levels of capital-skill complementarity, and there is a greater degree of of skill-bias in production associated with agglomeration economies, and find a positive relationship between city size and the wage gap between skilled and unskilled labour in urban areas in the United States between 1980 and 2007.

The idea of spatial differences in SBTC raises some interesting questions. Specifically, can patterns of SBTC explain movements in both class and spatial inequality? In other words, are skilled labour systematically earning more in certain locations vis-à-vis other locations? This paper attempts to examine some of these questions in the context of the information-technology revolution in India. Global advances in information technology over the last few decades have fueled what has been described as the 'Third Industrial Revolution,' and a few authors have claimed that these advances have also contributed to increases in wage inequality, through the channel of SBTC (Greenwood and Yorukoglu, 1997, Caselli, 1999, among others)

The narrative of India's growth story in the post-liberalization period has been one of services-led growth. The five years preceding the financial crisis of 2008 saw the services sector growing at nearly 10%, and contributing to about half of India's GDP. While these figures characterizing a "Dream Run" (Nagaraj, 2013), have somewhat tapered, the sector has still registered an average growth rate of 8.5% between 2011-12 and 2016-17 (Ministry of Finance, 2018).

Within services, information-technology, communications, and banking and financial services have been major contributors to services output. In particular, the Information-Technology (hereafter IT) and Business-Process-Outsourcing (hereafter BPO) services industries were instrumental in maintaining a services trade surplus, which to some extent offset the goods trade deficit for India. As of 2010, the IT-BPO industry contributed to roughly 6.4% of GDP in India (Chanda, 2012). According to the latest available NASSCOM report summary, the industry has grown to be a USD 154 billion industry, and was projected to grow at over 8% per annum in financial year 2017 (NASSCOM, 2018).[3].

The post-liberalization period in India has also been characterized by increases in both spatial and class inequality. Increasing spatial inequality is well-documented in terms of increasing regional differences in growth of wages (Papola and Kannan, 2017), consumption (Motiram and Sarma, 2014, Subramanian and Jayaraj, 2013), per capita income (Sachs et al., 2002, Krishna, 2004), and wealth (Subramanian and Jayaraj, 2013). In the context of class

[3]The sources of these estimates from NASSCOM are not clear in their reports

inequality, Vakulabharanam (2010) has several interesting insights regarding the nature of urban inequality in India between 1993-94 and 2004-05. Specifically, urban inequality was higher than overall inequality, urban elites' consumption grew more than that for urban unskilled, service-sector professionals recorded maximum gains between 1993-94 and 2004-05, and there has been a polarization between the consumption levels of urban elites and urban unskilled workers.

The motivating idea of this paper is to attempt to explore the possibility of unifying links between spatial and class inequality in urban India, in the context of the IT-revolution in the post-liberalization period. There is anecdotal evidence that the Indian information-technology services industry is spatially concentrated in locations like Bengaluru, Hyderabad and Gurugram. However, there seems to be no literature that attempts a systematic and analytical identification of these locations in India. This is an important step towards attempting to enquire about spatial differences in skill-wage premia in the context of the IT Services industry. Conforming with the framework of Baum-Snow et al. (2018), if it is the case that skill premia are higher in these locations, is it possible to characterize increases in spatial inequality in terms of increasing spatial differences in class inequality?

Subsection 3.1 lays out a framework for identifying locations where the Indian IT-Services industry has concentrated, or 'IT-Clusters.' Once these clusters have been identified, this paper attempts to answer some questions regarding how wage distributions (as characterized by varying returns to skills) differ in these clusters vis-*à*-vis locations that are not characterized as clusters, or 'non-clusters.' Specifically, are high-skilled workers being rewarded more in these clusters, relative to those outside these clusters? How has this 'location premium' to skilled labour evolved over time? Are these clusters only rewarding high skills, or are there spillover effects to unskilled labour as well? Are the observed movements consistent with the trend of increasing spatial and class inequality?

The remainder of this paper is organized as follows. Section 2 presents a review of literature on industrial clustering and on the relationship between skill-biased technical progress and wage movements. Section 3 outlines the empirical methodology involved in attempting to answer the aforementioned research questions. Section 4 discusses and provides the estimation results. Section 5 concludes.

2 A Review of Literature

2.1 Industrial Concentration

At the outset, it is useful to point out that this paper does not attempt to identify why IT-Services clusters developed in certain locations vis-à-vis others. The location of the clusters is taken as given, and the paper attempts to investigate patterns in wage distributions (as characterized by skills) in IT-Clusters relative to non-Clusters. Having said that, this section briefly reviews some literature concerning industrial concentration, to briefly explore whether there is a theoretical anchor for the IT-Services industry to have evolved in clusters.

To understand why industries should spatially concentrate, one needs to delve into some literature in economic geography. The starting point for this is Krugman (1991). In Krugman's construct, an economy is characterized by two sectors (agriculture and manufacturing), and two types of labour complementary to these sectors. Since agricultural labour is a complement to land, it is not mobile, while manufacturing labour has greater mobility. Krugman posits that the economy will endogenously align itself geographically into a 'manufacturing core' and an 'agricultural periphery,' on account of increasing returns to scale in manufacturing stemming from agglomeration economies. The endogenous realignment occurs on account of the manufacturing core locating itself in regions of higher demand (to minimize transport costs), and regions of higher demand in turn being associated with the location of the manufacturing core. Increasing returns to scale is a crucial assumption behind why manufacturing should concentrate, while iceberg transport costs is a crucial assumption for determining where manufacturing should concentrate.

Another question that emerges from this line of thought is regarding the nature of the industries that form the 'core.' Specifically, which industries have a greater propensity to cluster? From Krugman (1991), one would infer that industries that are in a position to exploit increasing returns to scale from agglomeration would be more prone to clustering. Technological spillovers have been identified as a source of increasing returns in agglomeration economies (Glaeser et al., 1992). There are several papers that support this claim; industries that entail knowledge and technological spillovers have a greater tendency to spatially concentrate themselves (Jaffe et al., 1993, Audretsch and Feldman, 2004, among others). Saxenian (1996) finds that the growth trajectories of Silicon Valley in California and Route 128 in Boston (two major regional centers for innovation and technology firms in the 1970's) diverged because of the former's greater capacity to absorb knowledge spillovers. Saxenian (2000) examines the efficacy of policy initiatives aimed at replicating the Silicon Valley model for the Bangalore IT-Cluster in India.

2.2 Skills and The Labour Market

I begin by briefly describing what Acemoglu (2002) refers to as the canonical model, which draws on the work of Card and Lemieux (2001) and Katz and Murphy (1992), among others. Acemoglu defines a production function with two types of labour inputs; high-skilled labour and low-skilled labour. Each type of labour is associated with a factor-augmenting technological input; high-skill augmenting technology and low-skill augmenting technology. The production function takes the constant elasticity of substitution specification. Assuming competitive labour markets, the wages for high-skilled and low-skilled labour equal the marginal product of each type of labour. There are two key results that the model hypothesizes. Firstly, under the assumption of high-skilled and low-skilled labour being gross-substitutes (i.e, the elasticity of substitution being greater than one), the wages of each type of labour are increasing functions of either kind of factor-augmenting technical input. In other words, ceteris paribus, high(low)-skill augmenting technical progress leads to an increase in wages of both high-skilled and low-skilled labour. Secondly, the skill premium, defined as the ratio of high-skill wages to low-skill wages, is an increasing function of the skill-intensity of technical inputs[4] when high-skilled and low-skilled labour are gross substitutes, and a decreasing function of the relative supply of skills[5].

The central premise for the idea of Skill-Biased Technical Progress is that the skill-intensity of technological inputs is increasing over time. Under this premise, the model implies that if SBTC is not matched with a corresponding increase in the relative supply of skills, the wage premium would increase over time. The aggregate effect is therefore determined by the degree to which an increasing relative supply of skills offsets SBTC. It follows that if the skill-intensity of technological inputs is assumed to increase over time, it must be matched by a proportionate increase in the relative supply of high-skilled labor in order to keep the skill premium at a constant level. This phenomenon is described as "The Race Between Eduction and Technology." by Tinbergen (1974) and Goldin and Katz (2007), among others.

Katz and Murphy (1992) attempt to test the premise of skill-biased technical change, using data on wages and employment for the United States between 1963-87. They express the relative demand for skills (the ratio of the demand for high-skilled labour to the demand for low-skilled labour) in terms of the the skill-intensity of technical inputs (which is a demand-shift factor in their framework), and the skill premium. They reject the hypothesis of a 'steady' demand for skills. Moreover, they find that the magnitude of demand shifts accounted for by changes in industrial structure and trade patterns is quite small in compar-

[4] defined as the ratio of high-skilled augmenting technical input to low-skill augmenting technical input

[5] defined as the ratio of supply of high-skilled labour to supply of low-skilled labour

ison to the observed magnitude of shifts in demand. They conclude that the observed shifts in demand follow a linearly increasing time trend. Their key result is that they establish SBTC as an empirical fact, and characterize it in terms of demand for skills increasing at a constant rate. It follows that any movements in wages and skill premia that occur off this time path are explained by movements in the supply of skills.

2.3 Spatial Differences in Skill Premia

Having characterized skill premia in terms of relative demand for and relative supply of skilled labour, a natural extension in the context of this paper is to explore why there may be spatial differences in skill premia. Baum-Snow et al. (2018) is one such attempt, which has been discussed in Section 1. Davis and Dingel (2019) also model how skill premia are affected by agglomeration driven by idea exchange and local learning. Operating under the premise of large cities being characterized by higher cost of living and richer "idea-exchange environments" relative to small cities, they find that the wages of both high-skill and low-skill workers are higher in large cities relative to small cities (to compensate for higher cost of living). Furthermore, on account of skill/ ability and local learning in idea exchange environments being complements, high-skill workers in larger cities benefit more from richer idea-exchange environments than their counterparts in smaller cities. As a result, skill premia in larger cities is higher than those in smaller cities. They find the positive relationship between skill premia and city population to be true for a sample of 275 cities in the United States, separately for 1990, 2000, and 2007 Census years.

3 Empirical Framework

The empirical framework is designed to test a class of hypotheses regarding how wages of skilled and unskilled labour in urban India differ between locations that can be characterized as IT-clusters, vis-à-vis non-clusters. The first step towards conducting this exercise involves the construction of two crucial variables. The 'Cluster' variable identifies locations as IT-Services clusters or non-clusters, and the 'Skill' variable identifies individuals (occupations) as skilled or unskilled. The following segment lays out a framework for constructing these two variables.

3.1 Identifying IT Clusters and Skilled Jobs

3.1.1 IT Clusters

To the best of my knowledge, there are no studies that attempt to identify the economic activity of the Indian IT industry at a sufficiently disaggregated and localized level. Khomiakova (2007) identifies seven Indian metropolitan cities as the "Big 7 IT-Clusters." These are Bangalore, Hyderabad, Delhi (NCR), Mumbai, Chennai, Pune, and Kolkata. However, this identification seems to be based solely on the presence of a select group of firms in these locations, and is therefore not a data-driven way of identifying locations where the IT-Services industry is concentrated.

The first step towards identifying locations that can be characterized as IT-Clusters, is to choose a spatial unit of analysis. I choose districts to be the ideal spatial unit, as it satisfies the Goldilocks principle; further aggregation from districts abstracts out a significant degree of heterogeneity, and further disaggregation from districts makes identifying localized economic activity quite difficult.

The next steps involve codifying the economic activity that can be characterized as "IT Services" (a term that has been used somewhat informally so far), and identifying a source of data that measures this economic activity at the level of districts. Table B-18 of the Census of India provides district-wise population numbers for individuals employed in industries, codified by National Industrial Classification (hereafter NIC) codes. I choose a set of NIC codes that map to industries that can be classified as IT Services[6]. Using employment numbers for the chosen NIC codes, a location quotient is constructed for each district, which is calculated as follows:

$$LQ_{i,r} = \frac{E_{ir}/E_r}{E_{in}/E_n} = \frac{E_{ir}/E_{in}}{E_r/E_n} \tag{1}$$

E_{ir} represents industry $i's$ level of economic activity (in this case, employment) in district r, E_r represents total level of economic activity in district r, E_{in} represents industry $i's$ aggregate economic activity in the economy, and E_n represents aggregate economic activity across the economy. A district having location quotient greater than one indicates that the share of individuals employed in the industry of interest is higher in this district, as compared to the national average. Therefore, districts with location quotients greater than or equal to one are defined as 'IT-clusters,' and those with location quotients less than one are defined as 'non-clusters.' Table 1 shows how IT clusters have evolved from 2001 to 2011.

Table 1: Evolution of IT-Services Clusters-2001 to 2011

		Census 2011		Total
		Cluster	Non-Cluster	
Census 2001	Cluster	30	28	58
	Non-Cluster	19	505	524
Total		49	533	582

One can infer that there is a significant degree of inertia in the evolution of clusters. Out of 582 districts, 505 have never been clusters, while 30 have always been clusters. Of the remaining 47 districts, 28 were classified as clusters in 2001 and not in 2011, and 19 became clusters in 2011 and weren't in 2001. The class of districts that were clusters in 2001 but didn't remain clusters in 2011 are dropped, for clearer analysis of evolution of wage distributions for IT-clusters vis-$\grave{a}$-vis non-clusters.

[6] these industries are software publishing, telecommunications, computer programming and consultancy, information services and related activities

Figure 1 points to the locations of the identified clusters. The districts shaded yellow have location quotients less than one (non-clusters), and the districts shaded red have location quotients greater than 1 (clusters).

Figure 1: Location of IT Services Clusters: 2001 and 2011

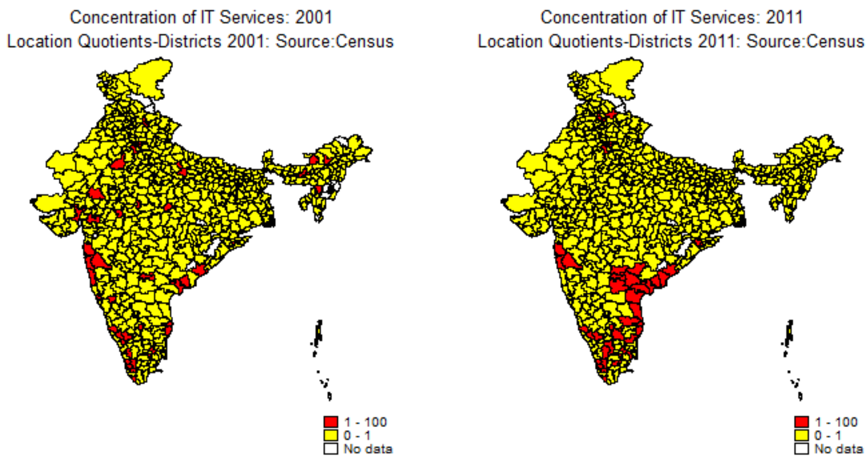

One can observe that between 2001 and 2011, IT Services clusters to some degree converged to southern parts of India, with Andhra Pradesh becoming a prominent destination for these industries between 2001 and 2011. One can also see that a number of the districts that were always IT-clusters are located in Maharashtra, Karnataka, and Kerala.

Based on this methodology, the 'Cluster' variable is defined for the regression equation in four distinct specifications; Cluster1, Cluster2, Cluster3 and Cluster4. This is done to have a deeper understanding of how being located in a district classified as an IT-cluster could influence wages. A district that was a non-cluster in both 2001 and 2011 is referred to as a Never-Cluster, a district that was an IT-cluster in 2001 but not in 2011 is referred to as an Old-Cluster, a district that was a non-cluster in 2001 but an IT-cluster in 2011 is referred to as a New-Cluster, and a district that was an IT-Cluster in both 2001 and 2011 is referred to as an Always-Cluster. These specifications are summarized in Table 2.

Table 2: Specifications of Cluster Variable

	Values	
Variable	1 (Focal)	0 (Comparison)
Cluster1	Always IT-Cluster	Never IT-Cluster
Cluster2	Always IT-Cluster or New IT-Cluster	Never IT-Cluster
Cluster3	New IT-Cluster	Always IT-Cluster
Cluster4	New IT-Cluster	Never IT-Cluster

For the initial analyses, the specification for each of the Cluster variables are binary, for ease of interpretation in a triple-differences framework. Moreover, Old Clusters are not used in any of these specifications, since they may not add much additional value to the analysis of wage distributions between clusters and non-clusters. As a robustness check, a fifth specification is also constructed, referred to as Dynamic Cluster, which takes the following specification:

$$\text{DynamicCluster} = \begin{cases} 0 & \text{if the district is classified as a Never Cluster} \\ 1 & \text{if the district is classified as a Old Cluster} \\ 2 & \text{if the district is classified as a New Cluster} \\ 3 & \text{if the district is classified as a Always Cluster} \end{cases}$$

This specification includes a greater number of cases within a particular estimation, specifically the case of how wages in Old Clusters look relative to Never Clusters, and therefore adds some value to the narrative in terms of the dynamics of how wage distributions have evolved differently in districts that were clusters relative to non-clusters. The results from this specification are discussed in Section 4.2.

3.1.2 Skills

Empirically verifying the results of models concerning SBTC is somewhat challenging, since the characterization of workers as skilled or unskilled isn't straightforward. A typical way of doing so has been to use the level of education as an observable alternative. A class of seminal papers in this field define skilled and unskilled labour as college graduates and high-school graduates respectively (Tinbergen, 1974, Katz and Murphy, 1992, Card et al., 1999, Card and Lemieux, 2001, Goldin and Katz, 2007, among others). In the Indian context, Azam (2009) finds an increase in the tertiary-secondary (college-high school) wage premium for India, driven by demand shifts in favour of workers with tertiary education in the 1980's and 1990's, using data from India's National Sample Surveys.

Apart from levels of educational attainment being used as indicators of skills, there are other alternatives, that attempt to classify skills relative to the context of work. Ramaswamy (2012) uses the classification of workers into manual/production (blue-collar) and non-manual (white collar) jobs as indicators of their level of skill, as reported in India's Annual Survey of Industries (ASI). Using ASI data, he finds a positive relationship between the share of skilled workers in producers' wage bills, and the scale of production, intensity of contract unskilled workers, and an increasing positive relationship between share of skilled workers in producers' wage bills, and capital-output ratio, for Indian manufacturing industries between 1981-2004. Azam et al. (2013) find positive and significant returns to wages for Indian workers endowed with English language skills, using data from the India Human Development Survey (IHDS).

A central issue with classifying workers into skilled and unskilled groups using binary criteria, such as the level of education, and alternative criteria mentioned above, is that doing so abstracts from the heterogeneity of skills within these groups. To account for this, several organizations maintain databases on the skill profiles of the labour force. For instance, the

World Indicators of Skills for Employment is an information system of the OECD that provides data on demand, supply and matching of skills for a set of 214 countries (though not for India). Another useful database is O-NET (United States), which is a survey database, and conducts surveys on the skill requirements of each occupational classification (according to the Standard Occupation Classification (SOC) system). Several papers have used O-NET data, specifically for constructing context-specific skill indices to classify labour as skilled or unskilled, in lieu of using their level of education, and thus augmenting the framework for empirical verification of the canonical model. Autor et al. (2003) construct skill indices from O-NET's Abilities, Work Context, and Work Activities indicators, that reflect the degree of routinization of occupations, and find that technological progress in the form of computerization has led to substitution of workers by computers in occupations that had a greater degree of routinization. Feser (2003) uses data from O-NET's Knowledge Indicators, to construct knowledge-based spatial occupation clusters in the United States. In the Indian context, Balasubramanian (2016) constructs skill indices for each Indian district, using O-NET's Abilities Indicators, to determine the role of the spatial distribution of skills in movements in employment growth between 2001 and 2011. Since O-NET data are collected for each occupation in the United States according to the SOC system, the author maps these occupations to India's National Classification of Occupations (NCO 2004).

To analyze the skill content of occupations, I draw on the work of Balasubramanian (2016), who uses the O*NET database of skill profiles of occupations in the United States. O*NET surveys a sample of workers in a large set of occupations (classified under the SOC system), and the survey includes questions regarding the importance and the extent (level) to which a particular skill is required for each occupation. These skills include cognitive skills (such as mathematical skills, oral and written expression, etc), sensory skills (hand-eye coordination, reaction time etc), physical skills (strength, stamina, etc), among others. For the purposes of this paper, I construct a Cognitive Skills Index by using variables from O*NET's Cognitive

Abilities Descriptors. The index is constructed by taking the geometric means of variable in the Cognitive Abilities descriptors along the Importance scale, for each occupation (classified according to the Standard Occupational Classification (hereafter SOC) system). The approach differs from Balasubramanian (2016) to the extent that only the importance scale is used to construct the indices, unlike the original paper, which uses the product of the importance and the level scales, to introduce a greater degree of variation. This is done keeping in mind the idea that the levels of skills may differ across countries for the same occupation, but the relative importance of skills may be spatially invariant. Further, each occupation in the SOC system is mapped to one in the NCO (2004) system. As a result, one can construct the cognitive skills index for 107 unique occupations, classified at the 3-digit level of NCO (2004). To facilitate a robustness check in the empirical exercise, a Physical Skills Index is also constructed, using the same methodology with O*NET's Physical Abilities descriptors.

Once each NCO code is assigned a value of the skill index, the result is further mapped to the sample of individuals surveyed in the NSS employment and unemployment surveys that satisfy the inclusion criteria specified in Section 3.2. Table 3 provides some summary statistics for each of the eight constructed skill indices, once they are mapped to the surveyed individuals:

Table 3: Summary Statistics: Skill Indices

	mean	sd	min	max
Cognitive Index	2.77	.27	2.04	3.41
Physical Index	1.86	.48	1	2.72
N	49593			

A priori, one expects returns to cognitive skills respond to the concentration of the IT Services industry, since its knowledge-intensive nature might increase demand for labour with a high degree of cognitive skills. To that extent, the analysis will only involve using the cognitive skills index. For classifying individuals as skilled or unskilled, the cognitive skill index is split into quintiles, and the 5th quintile is classified as skilled, while the remaining 4 quintiles are classified as unskilled. Accordingly, the variable takes value 1 if an individual is employed in an occupation for which the value of the skill index belongs to the fifth quintile, and 0 if it belongs to the remaining four quintiles. In extensions of this model, regressions are run separately for subsets of the skill distribution, taking the 5th quintile as skilled, and only one of the remaining four quintiles, each at a time, as unskilled.

3.2 Data and Variables

There are five major sources of data that have been used for constructing the pooled cross-sectional dataset required for this analysis. These are the National Sample Survey Organization's (hereafter NSSO) Employment and Unemployment Surveys, conducted in 2004-05 (61st Round) and 2011-12 (68th Round), Census of India 2001 and 2011, All India Survey of Higher Education (AISHE) survey year 2011, The Reserve Bank of India's (hereafter RBI) quarterly statistics on deposits and credits of scheduled commercial banks (Statement 4A), and O-NET 17.0 Database. Given that districts became the level at which strata were defined in the NSSO surveys following the 61st Round, one can merge data across multiple datasets at the district level.

There are two major reasons why the 61st and the 68th rounds of the NSS were chosen. Firstly, as mentioned above, it was from the 61st round onwards that the NSS samples became representative at the district level, and wage premia within districts reflect to a distribution of wages within each district. Secondly, since Census data are being used to construct location quotients for the IT services industry, these two are the closest rounds to the Census years that satisfy the first two criteria.

Formally, I estimate the following regression:

$$logWages_{i,t} = \alpha + \beta_1 t + \beta_2 Skill_{i,t} + \beta_3 Cluster_{i,t} + \beta_4 Skill_{i,t} \times t + \beta_5 Skill_{i,t} \times Cluster_{i,t} + \beta_6 Cluster_{i,t} \times t + \beta_7 Skill_{i,t} \times Cluster_{i,t} \times t + \delta District_{i,t} \times t + \gamma' X_{i,t} + \epsilon_{i,t} \quad (2)$$

Here, $Wages_{i,t}$ is the reported weekly wages of unit i surveyed in Round t. t is the round in which unit i was surveyed, and takes the value 0 for the 61st Round, and 1 for the 68th Round. $Skill_{i,t}$ is a binary variable indicating whether unit i surveyed in Round t is classified as skilled or unskilled, $Cluster_{i,t}$ is a binary variable indicating whether unit i surveyed in Round t is located in a district classified as a IT-cluster or non-cluster, $District_{i,t} \times t$ is a district-round interaction term, and $X_{i,t}$ is a vector of controls and potentially endogenous variables.

At the outset, it must be stated that this specification does not attempt to establish a causal connection. This is so as one cannot fully address the range of endogenous factors that influence the formation of IT-clusters and the level of wages in IT-clusters relative to non-clusters. For instance, migration of skilled workers who work in IT Services may cause a district to be classified as an IT-cluster, as well as influence the levels of wages in both IT-clusters and non-clusters. The issue of migration is not addressed, and is therefore an omitted variable that prevents one from making any causal claim regarding the impact of IT-clusters on urban wage distributions. Similarly, an IT-cluster could also be correlated to other industries clustering in the same location, and therefore, it would be impossible to isolate the impact of the IT-Industry clustering on wages from the impact of the other industries, which would make the channel of the causal connection unclear. Furthermore, this paper only looks at how wage distributions have differentially evolved for IT-clusters vis-$\grave{a}$-vis non-clusters, and not the demand or supply side factors that may affect those distributions. Factors that influence selection of individuals into the labour market differentially between IT-clusters and non-clusters are not included, and hence one can't claim the resultant relationships to be causal in nature.

However, there are several sources of endogeneity that the specification does address. The triple-differences framework differences out all time-invariant sources of endogeneity. I also include control variables that may explain variation in both the level of skills and earnings of individuals (like education). Furthermore, I account for several district-level factors which may be a cause or a consequence of cluster formation, and may also explain variation in wages. I also account for district-varying trends in inflation and other time-varying factors at the district level.[7] These controls are discussed further on in the section.

Data on wages are used from Block 5.3 of the NSS Employment and Unemployment Surveys. The inclusion criteria for units in the surveys are that they must be classified as residing in urban areas, must be reporting their wages, and be working in an occupation that is reported under a NCO code (NCO 2004 for the 68th Round and NCO 1968 for the 61st Round). In order to analyze the effects of independent variables on wages in terms of percentage changes, the log transformation of reported wages is used. The level of skills of units are inferred from the skill indices constructed for each occupation code in subsection 3.1. The classification of districts into IT-clusters and non-clusters is derived from the identification system developed in subsection 3.1.

The vector $X_{i,t}$ is a vector of control variables that would be significant in explaining variation in our dependent variable, as well as accounting for endogeneity in terms of factors that could influence both wages, and which districts become clusters, or the cognitive skills of individuals. These controls are measured at the level of the individual, household, and district. Firstly, in conformity with the Mincerian earnings function (Mincer, 1974), the

[7] Since the regression equation attempts to test for the convergence of skill premia (as discussed in section 1), one does need attempt to account for baseline differences in levels of wages.

education of unit i surveyed in period t, age of unit i surveyed in period t, and a squared age term are included.[8] These variables come from Block 4 (Household Demographics) of the Employment and Unemployment Surveys. Secondly, to account for the existence of a gender-based wage gap, a gender dummy for the unit is included. At the household level, a variable for the social group to which the household belongs is included, to account for caste-based differences in labour market outcomes (Banerjee and Knight, 1985, among others). Furthermore, a variable for education of the head of household is included, to account for the existence of varying job-market search costs for individuals, which might be reflected in their earnings. In addition, two variables at the district level are added, to account for potential endogeneity of IT clusters. The first is the number of universities in each district, since the presence of IT services firms in a district could result in a higher number of technical universities, which in turn could affect earnings. The number of universities in each district is retrieved from the AISHE Survey for 2011-12. The second is bank credit extended by commercial banks for each district, since the amount of credit flow in a district could be an indicator of the overall financial development of a district, and could therefore influence both the existence of IT firms, and the level of wages.

Table 4 lists out the variables used, and the data source for each variable.

[8] Age is taken as a proxy for potential experience, since there are no available data regarding work experience available in the NSSO survey

Table 4: Variables and Data Sources

Variable	Data Source	Level
Wages	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 5.3	Individual
Cognitive Skill Index	O-NET 17.0, mapped to NCO Codes, mapped to individuals in NSS Employment and Unemployment Surveys, Rounds 61 and 68, Block 5.1	Individual
Location Quotients for IT Services	Table B18, Census 2001 and 2011	District
General Education	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Technical Education	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Age	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Sex	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Social Group	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Head's Education	NSS Employment and Unemployment Surveys Rounds 61 and 68, Block 4	Individual
Number of Universities	AISHE Survey Year 2011	District
Credit	RBI's Quarterly Statistics on Deposits and Credits of Scheduled Commercial Banks (Statement 4A)	District

The specification expressed in Equation 2 helps in answering the research questions expressed in section 1. They are listed below in the form of the following hypotheses:

Hypothesis 1 For unskilled labour, there is no difference in growth of wages between IT-clusters and non-clusters.

$H_o : \beta_6 = 0$

$H_A : \beta_6 \neq 0$

Under the alternate hypothesis, $\beta_6 > 0$ would imply that wages of unskilled labour are growing at a faster rate in IT-clusters relative to non-clusters, while $\beta_6 < 0$ would imply the opposite.

Hypothesis 2 For skilled labour, there is no difference in growth of wages between IT-clusters and non-clusters.

$H_o : \beta_6 + \beta_7 = 0$

$H_A : \beta_6 + \beta_7 \neq 0$

Under the alternate hypothesis, $\beta_6 + \beta_7 > 0$ would imply that the wages of skilled labour are growing at a faster rate in IT-clusters relative to non-clusters, while $\beta_6 + \beta_7 < 0$ would imply the opposite.

Hypothesis 3 The difference in skill premium between IT-clusters and non-clusters has not changed over time.

$H_o : \beta_7 = 0$

$H_A : \beta_7 \neq 0$

Under the alternate hypothesis, the interpretation of the sign of β_7 is contingent on the sign of β_5. Specifically, if $\beta_5 > 0$, then the skill premium is higher in IT-clusters relative to non-clusters. In this case, $\beta_7 > 0$ would imply that the difference in the skill premium between IT-clusters and non-clusters has increased over time, while $\beta_7 < 0$ would imply that skill premia have spatially converged between IT-clusters and non-clusters. On the other hand, $\beta_5 < 0$ would imply that skill premia are higher in non-clusters relative to IT-clusters. In this case, the opposite inference needs to be made; $\beta_7 > 0$ would imply a convergence in skill premia, while $\beta_7 < 0$ would imply that skill premia have spatially diverged. Finally, if $\beta_5 = 0$, any value of β_7 that is significantly different from zero would imply a spatial divergence of skill premia.

3.3 Summary and Descriptive Statistics

This section provides some descriptions of how wages and some other explanatory variables differ between IT-clusters and non-clusters over time. Table 5 lists out the mean values of these variables, separately for districts that were Never-Clusters and Always-Clusters, and separately for Rounds 61 and 68.

Table 5: Means of Selected Variables

Variable	Round 61		Round 68	
(mean values)	Always IT-Cluster	Never IT-Cluster	Always IT-Cluster	Never IT-Cluster
log (Wages)	6.63	6.65	7.58	7.41
Cognitive Skills	2.73	2.78	2.76	2.77
Universities	225	41.5	204.4	38.3
Credit (Rs million)	0.53	0.07	3.03	0.25
High-School and Above (%)	25.6	27.8	34.0	29.7
Technical Education (%)	8.2	7.7	9.6	7.1

From Table 5, one can make a few key observations. Firstly, there is a small negative difference in mean of log wages between districts that were always IT-clusters and those that were never IT-clusters, for the 61st Round. The same difference becomes positive and larger in magnitude for the 68th Round. Secondly, the level of cognitive skills was on average higher in districts that were never IT-clusters, vis-$\grave{a}$-vis districts that were always IT-clusters, for Round 61. However, the level of skill somewhat declined over time for districts that were never IT-clusters, accompanied by an increase for districts that were always IT-clusters, between Rounds 61 and 68. Thirdly, both average credit and average number of universities are significantly higher for districts that were always IT-clusters, vis-$\grave{a}$-vis districts that were never IT-clusters, for both rounds. Therefore, one must account for the possibility of endogeneity, and for that reason, these two variables are included as additional controls in the regression equation. Finally, the proportion of individuals who have educational attainments including and above the high-school level were more-or less similar between Always-Clusters and Never-Clusters for the 61st Round. However, they recorded a much greater increase for Always-Clusters between 2004-05 and 2011-12, relative to Never-Clusters. Similar trends are observed in the context of technical education related to high levels of skills.

Figure 2 presents boxplots the log transformation of wages separately for the fifth quinitle of the cognitive skills index, and the first four quintiles, for Always-Clusters and Never-Clusters, and for both 61st and 68th Rounds.

Figure 2: Wages against Cognitive Skills Always IT-Clusters and Never IT-Clusters

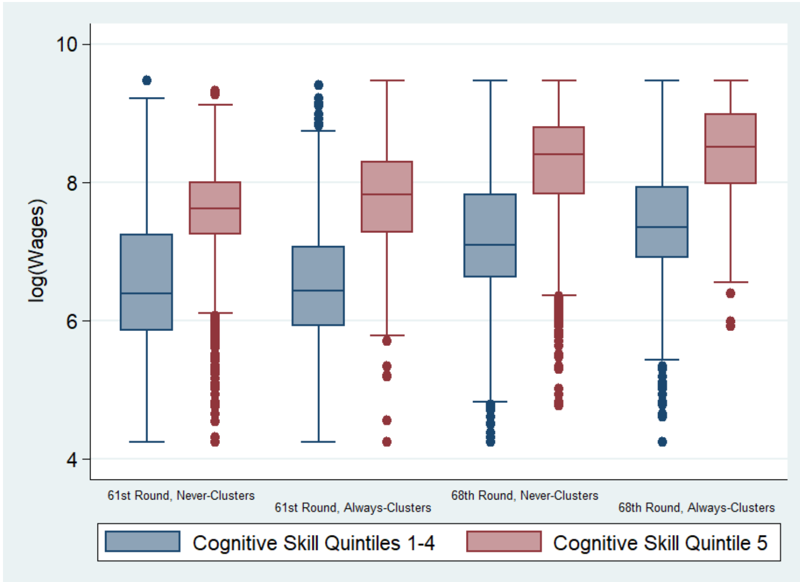

One can clearly infer the existence of a skill premium on average, as observed from the box for the fifth quintile being higher than that for the first four quintiles, consistently across all categories. The boxes for the 68th Round are higher than those for the 61st Round, consistent with an increasing time trend. Moreover, there seems to be a small difference in the set of boxplots between Always-Clusters and Never-Clusters for the 61st Round, which is less pronounced for the 68th Round, indicating the possibility of spatial convergence between Always-Clusters and Never-Clusters between the 61st and 68th Rounds.

Figure 3 plots the absolute change in Average Cognitive Skills between the 61st and 68th Round, separately for Never-Clusters, Old-Clusters, New-Clusters, and Always-Clusters.

Figure 3: Growth of Skills Across Different Categories of Clusters

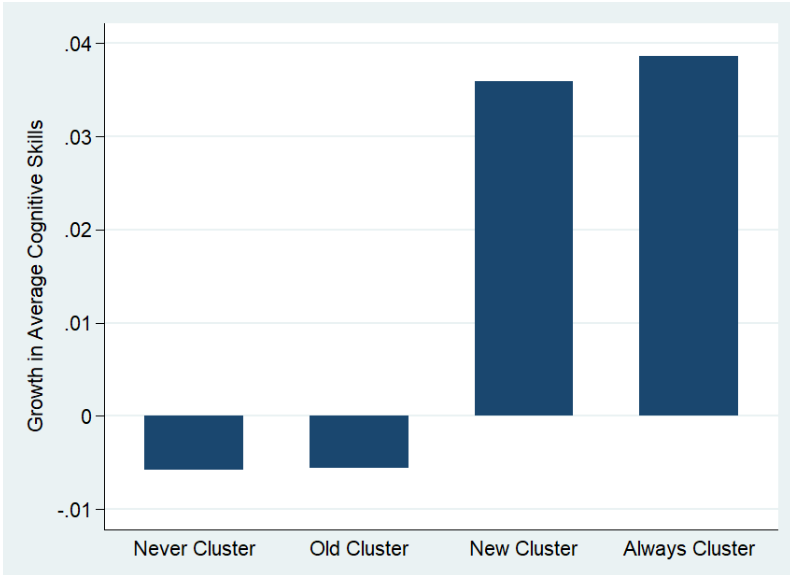

One can observe that the average level of cognitive skills declined for the Never-Clusters and Old-Clusters between the 61st and 68th Round. On the other hand, the average level increased for Always-Clusters and New-Clusters, suggesting an increasing relative supply of skills for these categories.

4 Regression Results

4.1 Preliminary Results

Table 6 presents five different estimation results for variants of Equation (1) specified in Section 1. The $Cluster_{i,t}$ variable takes on the specification represented by Cluster1 in Table 2. The variable COGIndex5 takes the value 1 if a unit belongs to the fifth quintile of the cognitive skill index, and 0 otherwise. In the first column, the equation is estimated for the entire skill distribution, taking the fifth quintile as the treatment group and the remaining four quintiles as the control group. In columns 2, 3, 4 and 5, the equation is estimated for the fifth quintile against the first, second, third, and fourth quintiles respectively. All standard errors are robust, and are clustered at the district level, to account for within-group correlation. Wages are winsorized to account for outliers. Standard person multipliers are used as analytical weights.[9]

Table 6: Estimation Results-Always-Clusters vis-à-vis Never-Clusters

	(1) Skilled: Q5 Unskilled: Q1-Q4	(2) Skilled: Q5 Unskilled: Q1	(3) Skilled: Q5 Unskilled: Q2	(4) Skilled: Q5 Unskilled: Q3	(5) Skilled: Q5 Unskilled: Q4
Round	0.972***	1.075***	1.365***	0.974***	0.734***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
COGIndex5	0.277***	0.834***	0.550***	0.178***	0.192***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Round×COGIndex5	-0.017	-0.072	-0.151**	0.027	0.022
	(0.70)	(0.21)	(0.01)	(0.60)	(0.67)
Cluster1	-0.323***	0.200*	0.084	-0.035	-0.445***
	(0.00)	(0.05)	(0.44)	(0.70)	(0.00)
Round×Cluster1	-0.149***	-0.218*	-0.611***	-0.266**	0.076
	(0.00)	(0.01)	(0.00)	(0.00)	(0.25)
COGIndex5×Cluster1	0.145**	0.090	0.070	0.185***	0.063
	(0.00)	(0.22)	(0.23)	(0.00)	(0.34)
Round×COGIndex5×Cluster1	-0.206*	-0.250*	-0.108	-0.160*	-0.199*
	(0.02)	(0.03)	(0.29)	(0.02)	(0.04)
General Education	yes	yes	yes	yes	yes
Technical Education	yes	yes	yes	yes	yes
District-Round Interaction	yes	yes	yes	yes	yes
Age	0.068***	0.053***	0.062***	0.082***	0.088***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Age^2	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Female	-0.492***	-0.519***	-0.524***	-0.321***	-0.301***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ST	-0.112*	-0.185	-0.132	-0.105	-0.049
	(0.01)	(0.09)	(0.06)	(0.15)	(0.67)
SC	-0.083***	-0.012	-0.114***	-0.082*	-0.050
	(0.00)	(0.71)	(0.00)	(0.01)	(0.19)
OBC	-0.103***	-0.119***	-0.128***	-0.118***	-0.134***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Number of Universities	-0.261	-0.475*	-0.344	-0.478	0.263
	(0.32)	(0.02)	(0.11)	(0.17)	(0.50)
District-Wise Credit (Rs Million)	-0.019	0.717***	0.163	0.490	-0.013
	(0.14)	(0.00)	(0.40)	(0.35)	(0.12)
Education of Head of Household	0.040***	0.039***	0.037***	0.044***	0.060***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
constant	4.791***	4.680***	4.578***	4.299***	4.250***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	43691	15014	18089	17929	14313
R^2	0.644	0.747	0.698	0.633	0.669
Adjusted R^2	0.635	0.728	0.679	0.610	0.642

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

p-values in parantheses

[9]Person multipliers are constructed by multiplying the household multipliers with the number of units in each household that satisfy the inclusion criteria mentioned in Section 3.1

Table 7: Hypotheses 1-3 Corresponding to Results from Table 6

Hypothesis	Skilled: Q5 Unskilled: Q1-Q4	Skilled: Q5 Unskilled: Q1	Skilled: Q5 Unskilled: Q2	Skilled: Q5 Unskilled: Q3	Skilled: Q5 Unskilled: Q4
Hypothesis 1	-0.149***	-0.218*	-0.611***	-0.266**	0.076
	(0.00)	(0.01)	(0.00)	(0.00)	(0.25)
Hypothesis 2	-0.355***	-0.468***	-0.719***	-0.426***	-0.122
	(0.00)	(0.00)	(0.000)	(0.00)	(0.189)
Hypothesis 3	-0.206*	-0.250*	-0.108	-0.160*	-0.199*
	(0.02)	(0.03)	(0.29)	(0.02)	(0.04)

p-values in parantheses

In conjunction with Table 6, Table 7 presents the results for linear restrictions on variables, that reflect the three hypotheses discussed in Section 2. Each column in Table 7 corresponds to the regression results presented in the column of the same number in Table 6. As mentioned earlier, the first four quintiles of cognitive skills have been classified as unskilled. The wages for each of the unskilled groups have grown at a slower rate in Always-Clusters relative to Never-Clusters. However, for the fourth quintile, the difference in wage growth between Always-Clusters and Never-Clusters is insignificant.

The skilled group, i.e, the fifth quintile presents an interesting case. A priori, one would expect the demand for a high degree of cognitive skills to be greater in Always-Clusters vis-à-vis Never-Clusters (and increasing, under the premise of SBTC). However, the wages of the fifth quintile have also grown at a slower rate in Always-Clusters relative to Never-Clusters, possibly due to the relative supply of skills increasing at a faster rate in Always-Clusters vis-à-vis Never-Clusters. between the 61st and 68th Rounds. Figure 3 also seems to suggest this possibility. As in the case of the unskilled group, the difference in wage growth for the skilled group is insignificant, when the comparison group is the fourth quintile. This could possibly be due to imperfections in the mapping of occupations from the SOC system and the NCO system. As a result, one could expect an overlap of occupations with similar skills and reported wages between the fourth and fifth quintiles of cognitive skills. As a robustness check, the estimation is done taking the fourth and fifth quintiles as skilled, and the remaining three quintiles as unskilled. These results are discussed in subsection 4.2.

Finally, the results suggest that the skill premium has decreased more in Always-Clusters relative to Never-Clusters. Table 6 points to the existence of a positive and significant skill premium for the entire sample, which has not significantly changed over time. Therefore, the triple difference coefficient being negative and significant would suggest that the skill premium has declined for Always-Clusters relative to Never-Clusters.

To explore these results in greater detail, the estimation is repeated for different specifications of the $Cluster_{i,t}$ variable as outlined in Table 2. The results for the Cluster2 specification (Always or New-Clusters vs Never-Clusters) is more or less similar to those for the previous specification. For the Cluster3 specification, the results suggest that the wages of both skilled and unskilled groups have increased less in New-Clusters relative to Always-clusters. However, the difference in how the skill premium has changed over time is insignificant between New-Clusters and Always-Clusters. Finally, the Cluster4 specification suggests that there is no difference in the movement of the wages of unskilled or skilled labour between New-Clusters and Never-Clusters, and the skill premium has also not changed differentially between New-Clusters and Never-Clusters.

4.2 Robustness Checks

For checking the robustness of the results, firstly, the $Cluster_{i,t}$ variable is recoded with 1.5 being the cutoff value for the location quotient. The results are mostly similar to the ones described above. Secondly, as a placebo test, the $Skill_{i,t}$ variable is recoded to take the fifth quintile of the physical skills index as 'skilled,' and the remaining four quintiles as 'unskilled.' On doing so, the location premium that was observed for different levels of cognitive skills between IT-Clusters and non-clusters vanishes. Thirdly, the $Cluster_{i,t}$ variable is recoded, with districts that are classified as "Big 7 IT-Clusters" according to Khomiakova (2007) code as 1, and the remaining districts coded as 0. The results are similar to those expressed in Table 6.

Fourthly, the Dynamic Cluster variable (as constructed in Section 3.1.1) is used instead of the original specifications outlined in Table 2. The results largely conform with those presented initially. One interesting observation that comes up is that one unskilled group (1st quintile of the cognitive skill index) earns more in Old Clusters relative to Never Clusters. However, the skilled group (fifth quintile of the cognitive skill index) has no significant difference in earnings between Old Clusters and Never Clusters. This probably suggests that there were positive spillovers for unskilled groups in old clusters, and they once behaved in the same manner that we observe the new and always clusters to be behaving, but skilled groups may have possibly migrated out of old clusters towards the new clusters. This hypothesis is further supported by the fact that the average level of cognitive skills declined for old clusters between 2004-05 and 2011-12, which is also implied from Figure 3.

Furthermore, to check for whether education and skills are proxies for each other, I estimate the model with an alternative skill variable, which codes individuals having a diploma/certificate course, and a graduate degree or above as skilled, and the remainder as unskilled. The results suggest that while there exists a "higher education premium" to wages on average, it is not significantly different between clusters and non-clusters (in any specification), as suggested by an insignificant second difference term. This seems consistent with the idea that while education may generally be rewarded in the labour market, using levels of education as proxies for specific cognitive skills abstracts out heterogeneity within education-groups, which is discussed in Section 3.1.2.

Finally, as described in Section 5, the $Skill_{i,t}$ variable is recoded to include the fourth and fifth quintiles of the cognitive skills index as skilled, and the remaining three quintiles as unskilled. the results for Always-Clusters vs Never-Clusters are presented in Table 8:

Table 8: Estimation Results-Always-Clusters vis-à-vis Never-Clusters
Fourth and Fifth Quintiles: Skilled

	(1) Skilled: Q4-Q5 Unskilled: Q1-Q3	(2) Skilled: Q4-Q5 Unskilled: Q1	(3) Skilled: Q4-Q5 Unskilled: Q2	(4) Skilled: Q4-Q5 Unskilled: Q3
Round	1.005***	0.853***	1.094***	0.953***
	(0.00)	(0.00)	(0.00)	(0.00)
COGIndex45	0.165***	0.373***	0.273***	-0.002
	(0.00)	(0.00)	(0.00)	(0.96)
Round × COGIndex45	-0.063	-0.088	-0.159***	-0.013
	(0.10)	(0.07)	(0.00)	(0.74)
Cluster1	-0.339***	-0.306*	-0.316***	-0.239**
	(0.00)	(0.01)	(0.00)	(0.00)
Round × Cluster1	-0.203***	0.035	-0.339***	-0.226***
	(0.00)	(0.67)	(0.00)	(0.00)
COGIndex45 × Cluster1	0.093*	0.030	-0.002	0.147***
	(0.01)	(0.64)	(0.96)	(0.00)
Round × COGIndex45 × Cluster1	-0.039	-0.087	0.032	-0.027
	(0.48)	(0.33)	(0.65)	(0.62)
constant	4.755***	4.630***	4.551***	4.215***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	43691	22109	25184	25024
R^2	0.641	0.706	0.673	0.629
Adjusted R^2	0.632	0.691	0.658	0.612

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

all controls used in Table 6 are included. p-values in parentheses

Table 9: Hypotheses 1-3 Corresponding to Results from Table 8

Hypothesis	Skilled: Q4-Q5 Unskilled: Q1-Q3	Skilled: Q4-Q5 Unskilled: Q1	Skilled: Q4-Q5 Unskilled: Q2	Skilled: Q4-Q5 Unskilled: Q3
Hypothesis 1	-0.203***	0.035	-0.339***	-0.226***
	(0.00)	(0.67)	(0.00)	(0.00)
Hypothesis 2	-0.242***	-0.052	-0.308***	-0.253***
	(0.00)	(0.42)	(0.00)	(0.00)
Hypothesis 3	-0.039	-0.087	0.032	-0.027
	(0.48)	(0.33)	(0.65)	(0.62)

p-values in parentheses

The results suggest that wages have grown less for both skilled and unskilled groups in Always-Clusters relative to Never-Clusters between the 61st and 68th Rounds. The key difference between the results of this estimation and the one expressed in Table 6 is that the skill premium has not changed differentially for Always-Clusters relative to Never-Clusters.

Repeating the estimations for different classifications of the $Cluster_{i,t}$ variable (see Table 2), the results suggest a similar story for the Cluster2 and Cluster3 specifications, with the first skill quintile being an exception; thee is no significant difference in the evolution of the first quintile's wages between IT-Clusters and non-clusters. Finally, there are no significant differences in the evolution of wages of unskilled labour, skilled labour, and the skill premium between New-Clusters and Never-Clusters.

4.3 Discussion

In general, the results seem to suggest that the time path of wages has not systematically differed between IT-clusters and non-clusters. A priori, one might expect the IT-clusters to be associated with a higher demand for high levels of cognitive skills, inferring from Baum-Snow et al. (2018). However, given that the average level of skills has increased in Always-Clusters and New-Clusters (and decreased in Never-Clusters and Old-Clusters), one reason for this observation could be an increasing supply of high skills in clusters. This could be because of increasing education levels in IT-clusters relative to non-clusters (as evidenced by the increasing proportion of individuals with higher education attainments, and a higher number of universities in clusters), or by the migration of both skilled and unskilled labour from non-clusters to IT-clusters. This narrative helps in explaining some of the obtained results. Another explanation for the results could be that the Financial Crisis of 2008 disproportionately affected wages in IT-clusters. Sahoo et al. (2013) find that the income elasticity of demand for India's services exports was high between 1980-2011, which implies that the crisis could have disproportionately affected the earnings of high-skilled workers in IT-clusters relative to non-clusters, and trickled down to the earnings of low-skilled workers as well. This might help in explaining some of the results wherein the wages of both high-skilled and low-skilled workers have grown at a slower rate in Always-Clusters relative to Never-Clusters. Moreover, the skill premium is never found to be lower in IT-clusters relative to non-clusters. Consequently, the triple-difference term being negative or insignificant implies that there is no evidence of divergence of skill premia between Always-Clusters and Never-Clusters.

One limitation of the study is that it fails to take into account movements in the demand side for skills, and instead rests on the premise of higher (or at least equal) demand for cognitive skills in IT-clusters relative to non-clusters. However, operating under this premise gives an explanation for the results that is consistent with the framework of Katz and Murphy (1992). It is difficult to comment on the welfare implications of growth driven by services like IT, since the IT Services industry is still in a dynamic phase of growth, and it is too soon to say whether the locations where it is concentrated have stabilized. However, one can say with a reasonable degree of confidence that the growth of this industry in clusters has not exacerbated spatial inequality between clusters and non-clusters, as inferred from the results. An extension of this work would involve understanding patterns in wage distributions once the industry achieves some form of a spatial equilibrium.

5 Conclusion

The post-liberalization period in India has been characterized by episodes of high growth driven by services, at the helm of which has been the Information-Technology revolution. Past and present governments have identified the role of this sector as a driver of growth, and its tendency to grow in clusters has been promoted through initiatives like Software Technology Parks of India. However, growth of this form has also coincided with increases in spatial inequality, and greater polarization along the lines of region and class (Motiram and Sarma, 2014, Vakulabharanam, 2010, among others). This paper endeavours to consolidate these three facts under a unified framework. Specifically, it attempts to see if increases in spatial inequality can be explained in terms of increasing differences in class inequality between spaces associated with the concentration of the IT Services industry.

It needs to be reiterated that the results should not be interpreted as causal connections. Having said that, the results suggest that the observed increases in spatial inequality cannot be explained in terms of widening spatial differences between class inequality within IT-clusters relative to outside IT-clusters. There are a few caveats that one needs to keep in mind while making this claim; the empirical construct only attempts to analyze differences in urban wages, and thereby can not explain movements in spatial inequality with their roots in rural areas. Moreover, the NSS under-represents the top quantiles of wage-skill distributions; it may be the case that the earnings of the top quantiles have grown at a faster pace within IT-clusters relative to non-clusters, and the inclusion of these individuals in the sample may lead to different results.

This paper raises several questions in the context of spatial differences between IT-clusters and non-clusters. One obvious extension of this work would be to see if one observes patterns in migration that are consistent with the observed movements in wages. Specifically, if (expected) wage differentials are the sole reason behind migration, as predicted by the Harris-Todaro model (Harris and Todaro, 1970), migration patterns consistent with the obtained results should point to an outflow of labour from IT-clusters, since wages of both skilled and unskilled labour have grown at a slower pace in IT-clusters relative to non-clusters. If the opposite trend is observed, what are the reasons behind migration into clusters?

This raises a second question that can be an extension of this work. What kind of employment is being generated in IT-clusters, and how does it differ from non-clusters? It may be the case that being located in an IT-clusters may not reward workers in terms of higher wages, but the formal and informal economies in IT-Clusters may have greater avenues for absorbing labour migrating into clusters.

References

Acemoglu, D. (2002), 'Technical Change, Inequality, and the Labor Market', *Journal of Economic Literature* **40**(1), 7–72.

Acemoglu, D. and Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, *in* 'Handbook of Labor Economics', Vol. 4, Elsevier, pp. 1043–1171.

Audretsch, D. B. and Feldman, M. P. (2004), Knowledge Spillovers and The Geography of Innovation, *in* 'Handbook of Regional and Urban Economics', Vol. 4, Elsevier, pp. 2713–2739.

Autor, D. H., Levy, F. and Murnane, R. J. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *The Quarterly Journal of Economics* **118**(4), 1279–1333.

Azam, M. (2009), India's Increasing Skill Premium: Role of Demand and Supply, Technical report, IZA Discussion Papers 2968.

Azam, M., Chin, A. and Prakash, N. (2013), 'The Returns to English-Language Skills in India', *Economic Development and Cultural Change* **61**(2), 335–367.

Balasubramanian, I. (2016), 'Local Skill Concentrations and District Employment Growth: A Simultaneous Equation Approach for India', *The Annals of Regional Science* **56**(2), 491–511.

Banerjee, B. and Knight, J. B. (1985), 'Caste Discrimination in The Indian Urban Labour Market', *Journal of Development Economics* **17**(3), 277–307.

Baum-Snow, N., Freedman, M. and Pavan, R. (2018), 'Why Has Urban Inequality Increased?', *American Economic Journal: Applied Economics* **10**(4), 1–42.

Card, D., Kramarz, F. and Lemieux, T. (1999), ‘Changes in the Relative Structure of Wages and Employment: A Comparison of the United States, Canada, and France’, *Canadian Journal of Economics* pp. 843–877.

Card, D. and Lemieux, T. (2001), ‘Can Falling Supply Explain the Rising Returns to College for Younger Men? A Cohort-Based Analysis’, *The Quarterly Journal of Economics* **116**(2), 705–746.

Caselli, F. (1999), ‘Technological Revolutions’, *American Economic Review* **89**(1), 78–102.

Chanda, R. (2012), ‘Services-Led Growth’, *in* ‘The New Oxford Companion to Economics in India’, Vol. 2, OUP, pp. 624–633.

Davis, D. R. and Dingel, J. I. (2019), ‘A Spatial Knowledge Economy’, *American Economic Review* **109**(1), 153–70.

Feser, E. J. (2003), ‘What Regions Do Rather Than Make: A Proposed Set of Knowledge-Based Occupation Clusters’, *Urban Studies* **40**(10), 1937–1958.

Glaeser, E. L., Kallal, H. D., Scheinkman, J. A. and Shleifer, A. (1992), ‘Growth in Cities’, *Journal of Political Economy* **100**(6), 1126–1152.

Goldin, C. and Katz, L. F. (2007), ‘The Race Between Education and Technology: The Evolution of US Educational Wage Differentials, 1890 to 2005’, Technical report, National Bureau of Economic Research: Working Paper 12984.

Greenwood, J. and Yorukoglu, M. (1997), ‘1974’, *in* ‘Carnegie-Rochester Conference Series on Public Policy’, Vol. 46, Elsevier, pp. 49–95.

Harris, J. R. and Todaro, M. P. (1970), ‘Migration, Unemployment and Development: A Two-Sector Analysis’, *The American Economic Review* **60**(1), 126–142.

Hicks, J. (1963), *The Theory of Wages*, Springer.

Jaffe, A. B., Trajtenberg, M. and Henderson, R. (1993), 'Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations', *The Quarterly Journal of Economics* **108**(3), 577–598.

Katz, L. F. and Murphy, K. M. (1992), 'Changes in Relative Wages, 1963–1987: Supply and Demand Factors', *The Quarterly Journal of Economics* **107**(1), 35–78.

Khomiakova, T. (2007), 'Information Technology Clusters in India', *Transition Studies Review* **14**(2), 355–378.

Krishna, K. (2004), 'Patterns and Determinants of Economic Growth in Indian States', Technical report, Indian Council for Research on International Economic Relations: Working Paper 144.

Krugman, P. (1991), 'Increasing Returns and Economic Geography', *Journal of Political Economy* **99**(3), 483–499.

Mincer, J. (1974), 'Schooling, Experience, and Earnings', *Human Behavior & Social Institutions* **2**.

Ministry of Finance (2018), *Statistical Appendix, Economic Survey, 2018*, Oxford University Press, Government of India.

Motiram, S. and Sarma, N. (2014), 'Polarization, Inequality, and Growth: The Indian Experience', *Oxford Development Studies* **42**(3), 297–318.

Nagaraj, R. (2013), 'India's Dream Run, 2003-08', *Economic and Political Weekly* **48**(20), 39–51.

NASSCOM (2018), *Strategic Review, 2018*, Mumbai.

Papola, T. and Kannan, K. (2017), 'Towards an India Wage Report', *International Labour Organization Asia-Pacific Working Paper Series ISSN 2227-4391* .

Ramaswamy, K. (2012), 'Wage Inequality in Indian Manufacturing: Is it Trade, Technology or Labour Regulations?', Technical report, Indira Gandhi Institute of Development Research: WP-2008-021.

Sachs, J. D., Bajpai, N. and Ramiah, A. (2002), 'Understanding Regional Economic Growth in India', *Asian Economic Papers* **1**(3), 32–62.

Sahoo, P., Dash, R. K. and Mishra, P. P. (2013), 'Determinants of India's Services Exports', *Institute of Economic Growth Working Paper 333* .

Saxenian, A. (1996), *Regional Advantage*, Harvard University Press.

Saxenian, A. (2000), Bangalore: The Silicon Valley of Asia?, *in* 'A Conference on Indian Economic Prospects: Advancing Policy Reform', Center for Research on Economic Development and Policy Reform, Stanford, California.

Subramanian, S. and Jayaraj, D. (2013), 'The Evolution of Consumption and Wealth Inequality in India: A Quantitative Assessment', *Journal of Globalization and Development* **4**(2), 253–281.

Tinbergen, J. (1974), 'Substitution of Graduate by Other Labour', *Kyklos* **27**(2), 217–226.

Vakulabharanam, V. (2010), 'Does Class Matter? Class Structure and Worsening Inequality in India', *Economic and Political Weekly* pp. 67–76.