2022

Essays on worker heterogeneity and its macro implications

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS & SCIENCES

Dissertation

ESSAYS ON WORKER HETEROGENEITY AND ITS MACRO IMPLICATIONS

by

MARTIN SHU

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2022

Approved by

First Reader

David Lagakos, PhD
Associate Professor of Economics

Second Reader

Pascual Restrepo, PhD
Assistant Professor of Economics

Third Reader

Kevin Lang, PhD
Laurence A. Bloom Professor of Economics

Fourth Reader

Natalia Ramondo, PhD
Associate Professor of Economics

To my parents, and their parents, for everything they mean to me.

And to all friends who ever become a part of my life, for making me who I am.

Acknowledgments

My deepest gratitude goes to my advisors, Pascual Restrepo and David Lagakos, without whom this dissertation would not have been possible. I thank them wholeheartedly for their guidance, support, and encouragement throughout this journey.

I am grateful to Kevin Lang for being a mentor as well as a friend, who provided great advice in shaping the substance and the writing of this dissertation.

I have also benefitted from other faculty members at the Department of Economics, including but not limited to Stephen Terry, Natalia Ramondo, Stefania Garetto, Yuhei Miyauchi, and Siddharth George, in no particular order.

Finally, I thank all my friends within and outside the department, whose support has been invaluable in making this endeavor peaceful, joyful, and memorable.

ESSAYS ON WORKER HETEROGENEITY AND ITS MACRO IMPLICATIONS

MARTIN SHU

Boston University, Graduate School of Arts & Sciences, 2022

Main Advisors: David Lagakos, PhD
Associate Professor of Economics

Pascual Restrepo, PhD
Assistant Professor of Economics

ABSTRACT

This dissertation examines the aggregate and sectoral implications of the self-selection of workers with different skills into sectors and occupations.

The first chapter concerns the low growth rate of labor productivity in services in the United States that accompanied the large expansion of the sector. The slow growth of labor productivity in this sector may be caused by the selection of less productive workers into services as the sector expands. Using panel data for US workers, I document that workers moving into professional services are more productive than an average worker in the sector. On the other hand, workers moving into education, health, and the public sector are less productive than incumbents. I build a generalized quantitative multi-sector Roy model that allows for these rich patterns of selection. Compared to a conventional selection model where new workers arriving at all services sectors are negatively selected by assumption, the estimated model in this paper leads to an effect of selection on labor productivity for professional services that is 14 percentage points higher, and a weaker negative effect for education, health,

and public services in the U.S. between 1989 and 2019. Overall, selection leads to little effect on labor productivity in aggregate services, contrasting sharply with the prediction of a conventional selection model.

The second chapter, joint with Siddharth George, studies the potential gains in aggregate productivity from mitigating the intergenerational persistence in occupations. Workers are much more likely to enter the occupation of their parents around the world. In this chapter, we develop a measure of this dynastic bias for occupations the odds ratio of the probability of choosing an occupation conditional on whether one's father is in that occupation for over 90 countries with data from 275 censuses or national surveys. At the occupation level, we document that the dynastic bias is increasing in the eliteness of occupation but decreasing in the average years of schooling. At the aggregate level, we document that the dynastic bias first rises and then falls with both GDP per capita and the average years of schooling of an economy. To evaluate the aggregate implication of the dynastic bias on productivity, we build a quantitative Roy model of occupational choice with entry barriers into occupations that depend on whether one's parent works in that occupation or not. We calibrate the model and perform counterfactual exercises to infer the potential gains in productivity for all economies in our sample. The entry barriers of occupations are set to match the dynastic bias of each occupation. We then remove all entry barriers and solve for the new equilibrium so that workers pursue their occupations entirely based on their comparative advantage without any income penalty. This frictionless counterfactual leads to labor productivity gains at the aggregate level that average at 8.6% for all censuses in our sample. Focusing on the latest censuses of the 51 non-high-income countries with data after 2000, we find that the average potential gains in labor productivity remain considerable at 8.7%, with 20 countries having potential gains above 10%.

Preface

二十一世纪什么最贵？

人才！

——冯小刚《天下无贼》(2004)，34:07

What's most precious in the twenty-first century?

TALENT!

Feng Xiaogang, *A World without Thieves* (2004), 34:07

It was my first year of middle school when my parents took me to the Shanghai Film Art Center to watch the film *A World without Thieves*, released for the New Year. This famous quote is said by Uncle Li, head of a theft group cast by Ge You, when he notices a theft couple and wishes to test out their skills and recruit them. It soon became a cultural phenomenon that I would bet five dollars that a random person between 30 and 60 years of age now in China knows about this line. The popularity of the quote owes partly to the sensational performance by Ge You, and partly to the contrast that even a theft group recognizes the value of talent.[1]

In the wake of the twenty-first century, it was not so obvious to China how talent should be regarded as the most important factor of production when this country was building up as the assembly line of the world. It was also not so obvious to economists how to think about the implication of talent on productivity or economic growth. The earliest scholarly work I am aware of that explicitly focuses on "talent" and growth is the article by Murphy et al. (1991), in which the authors studied how the flow of talented workers into rent-seeking occupations hurts economic growth. Constrained

[1]The original English subtitle for this line is, "This century's most expensive commodity is talent." I prefer my translation since I consider it inaccurate to take talent as a kind of commodity. In addition, the word "precious" reflects not only the high value of talent but also its scarcity.

by the computational capacity at the time, the authors provided only the results of some exploratory regressions without precise quantitative statements. It was also unclear how the differences in the talent of individuals affected the interpretation of the regression results. Therefore, my reading of the paper is that it is more about the "allocation" of the general labor force instead of "talent".

More recently, the concept of talent has become clearer and more relevant in the growth and macro development literature with the introduction of quantitative heterogeneous workers models. These models trace their intellectual origin back to Roy (1951) and make use of the nice properties of extreme value distributions, borrowing from the international trade literature. The assumptions made with these models capture the key characteristics of what we call talent   it differs by person, and the most talented persons of each profession, occupation, or industry are rare in number. The "allocation of talent" has since then become a much better-defined research object   it is not just about the allocation of workers, but more importantly about which specific workers sort themselves into each category of economic activities.

The fact that I choose to write on this topic reflects the extent to which I agree with the quote. To take one step further, it reflects my approach to the subject of economics. It is deeply engraved in my mind that the centerpiece of economics is human beings, for economic activities are organized by people and ultimately consumed by people. This narrative pushes me to investigate the nature, causes, and implications of the heterogeneity among people along various dimensions. This dissertation serves as a starting point of my exploration along this line, and I do not doubt that more enlightening findings lie ahead.

I hope this monologue creates some resonance with the reader, whether an economist or not. Personally, it is refreshing for me to look back once in a while on how I have come this way and what I continue to struggle for. It is duly appreciated that the reader makes it to this paragraph. Now, I hope the reader finds this treatise useful for whatever purpose. In the event that the reader enjoys the luxury of reading for leisure, I hope that my writing entertains and is worth a cup of tea.

Martin Shu

In the cozy chair of my office,

B30C, 270 Bay State Road, Boston, MA

May 9, 2022

Contents

List of Tables

List of Figures

List of Abbreviations

BEA	......	Bureau of Economic Analysis
BLS	......	Bureau of Labor Statistics
CES	......	Constant Elasticity of Substitution
CPS	......	Current Population Survey
EHP	......	Education, Health, and Public services
GDP	......	Gross Domestic Product
GGDC	......	Groningen Growth and Development Centre
ID	......	Identification
MFG	......	Manufacturing
PPP	......	Purchasing Power Parity
PROF	......	Professional services
TFP	......	Total Factor Productivity
UK	......	United Kingdom, the
US	......	United States, the

Chapter 1

Selection, Structural Transformation, and the Cost Disease of Services

1.1 Introduction

The service sector is becoming increasingly important for the economy, especially in developed countries. According to the World Bank, the employment share of services has risen from 63% to 74% for high-income countries over the past 30 years, while the latest number for the United States is 79%. On the other hand, labor productivity growth of the service sector averages only 1.3% per year between 1989 and 2009, much lower than the annual 3.7% labor productivity growth in manufacturing. These observations lead to concerns that the growth in aggregate labor productivity might continue to drop in the long run alongside the continual reallocation of workers from manufacturing into services a phenomenon known as Baumol's cost disease (Baumol, 1967).

These concerns notwithstanding, the *measured* growth in sectoral labor productivity might not reflect the true extent of technological progress when a large amount of workers reallocate between sectors. An idea going back to Roy (1951) is that the expansion in employment of a sector entails absorption of new workers less skilled than the incumbents due to self-selection, and the skill level of the marginal worker declines as the sector continues to expand. Given that the employment share of the U.S. service sector has increased substantially, the lower skill level of new services workers could have substantially lowered the measured growth of labor productivity in the sector.[1]

This chapter explores the contribution of selection and worker reallocation to sectoral labor productivity trends in the U.S. I show that the selection effect differs across subsectors *within* services. In particular, I document with cross-country data that real output per worker and changes in employment share are negatively correlatetd for manufacturing and the sector of education, health services, and public administration ("EHP" hereinafter), but *positively* correlated for professional services.[2] These patterns lead to questions about whether the employment expansion of a sector necessarily entails reduction in the average skill level of its workers. I then verify the pattern of selection at the micro level by utilizing the panel component of the monthly CPS in the U.S. I find that workers who move from manufacturing into professional services earn about 5% more than incumbent professional services workers, implying that new workers create more value on average. The opposite happens to workers in EHP who reallocate from manufacturing, whose real weekly earnings are about 13% below those of the incumbents, so that expansion of the sector substantially lowers its average skill level. These micro-level observations point to a richer pattern of selection that cannot be explained by the conventional Roy models used in the development and trade literature (e.g. Lagakos and Waugh (2013); Bryan and Morten (2019); Burstein et al. (2019); Hsieh et al. (2019); Galle et al. (2021)). The models used in these studies have restrictive parametric assumtpions, which force newcomers to be less productive than incumbents at expanding sectors.

To incorporate the opposite patterns of selection for professional services and EHP, I build a three-sector Roy model that explicitly describes the relationship between the

[1]This argument also applies to the measured growth of total factor productivity. When employment share rises in a sector, it may have reduced TFP due to the inflow of less productive workers, even if its technology does not regress or if its capital stock does not shrink.
[2]The sector of "professional services" includes finance, real estate, professional, and business services.

absolute advantage of workers in professional services and their comparative advantages in manufacturing and EHP with respect to professional services. Adão (2016) and Alvarez-Cuadrado et al. (2019) illustrate that this relationship is central to pinning down the selection effect on the labor productivity of a sector, instead of the correlation between the level of skills. My model generalizes the two-sector log-linear specification in Adão (2016) into a three-sectors setup and specifies how the absolute advantage varies with the two comparative advantages. I then structurally estimate the model using Simulated Method of Moments, with worker-level panel data from the Outgoing Rotation Group of the monthly Current Population Survey of the U.S. The parameter values imply that the marginal worker's skill in professional services increases with the reallocation of workers from manufacturing, but decreases if the reallocated workers come from EHP. Consequently, the expansion of professional services increases its labor productivity as it draws more productive workers into the sector from manufacturing.

The fully estimated model shows that selection *increases* labor productivity in professional services by 0.9% and decreases labor productivity in EHP by 2.3%. These numbers are substantially different from the predictions of a conventional selection model with indpendent Fréchet distribution of skills. The number for professional services is 14.2% higher than that from the conventional model, reverting a substantially negative effect to a positive number. The 2.3% decrease for EHP is also weaker than the 5.3% decrease predicted by the conventional model. These differences add up to have a 0.4% decrease in labor productivity for aggregate services, which is negligible compared to the 10% decrease the conventional model predicts.

I then apply the model to quantify the contribution of selection to sectoral labor productivity growth in other developed countries. The results show that the magnitudes of the selection effect on labor productivity for professional services and EHP in France, Spain, and the U.K. are similar to those in the U.S. The effect for Japan is stronger in both service sectors. Italy is an exception where selection contributes negatively to labor productivity growth in professional services but positively to EHP, since it is the only country in the sample whose EHP sector shrinks and releases less productive workers to the professional services sector. Nonetheless, the effect on aggregate services remains small.

In a final exercise, I extend the structural model to a general equilibrium environment with the non-homothetic CES preference as in Comin et al. (2021). I examine the outlook of structural transformation in the U.S. for the next 50 years by feeding in the implied sectoral technology growth rates of the U.S. from the past 30 years, comparing what my baseline model predicts with the prediction of a conventional selection model with independent Fréchet distribution of skills. My baseline model achieves more structural transformation in expenditure share with less reallocation of workers compared to the conventional model, since expanding services employment in my model leads to higher gains in human capital than in the conventional model. Consequently, the total welfare gains in my model is more than twice as much as the gains predicted by the conventional model.

My results contrast sharply with Young (2014), who finds a substantially negative elasticity of an industry's "worker efficacy" with respect to its employment share, albeit "estimated imprecisely" with industry-level data. With his estimate, Young (2014) argues that we cannot reject the hypothesis that manufacturing and services in advanced economies experience similar true technology growth, and that the sectoral gaps in measured labor productivity growth are entirely driven by the reallocation of less productive workers. My paper reaches a different conclusion by providing a precise estimate of the selection effect. The use of worker-level panel data alleviates the problem of low availability of data at the industry level and allows me to infer distinct selection effect for different sectors, which was a constraint for Young (2014) that led him to impose one single elasticity for all industries. Moreover, Young (2014) uses annual defense spending-to-GDP ratio as the instrument variable for demand shocks, which does not differentiate the selection effect between the long-run sectoral reallocation of workers and the short-run increase in labor supply from unemployment and non-participation. In contrast, I focus on workers shifting directly between sectors, which alleviates the concern about temporary effects.

My findings expand the current set of implications of structural Roy models used in the literature. Lagakos and Waugh (2013) use a two-sector Roy model with Fréchet marginal skill distributions to explain the cross-country gaps in the ratio of agricultural to non-agricultural labor productivity. Alvarez (2020) extends their framework with labor market frictions and supplies evidence of selection with Brazilian data. Bryan and Morten (2019) and Hsieh et al. (2019) use multivariate Fréchet distribution of skills to evaluate the aggregate productivity gains from reducing labor market frictions, while Pulido and Święcki (2019) and Adamopoulos et al. (2021) use multivariate normal distribution to decompose the effect of selection and distortions on the income gap between agricultural and non-agricultural workers. All of these studies use fully parametric Roy models and find falling labor productivity with the expansion of a sector. This chapter provides empirical evidence of an opposite pattern that new workers increase the labor productivity in their new sectors, and develops a more general parametric framework to quantitatively evaluate the selection effect. Distinctly, I use the income gap between incumbent workers and incoming workers from different sectors to identify the parameters for the skill distribution, since this empirical moment provides information most closely related to the effect of labor reallocation on labor productivity. This method echoes with the recommendation by Heckman and Honoré (1990) to use panel data for identification of a Roy model.

The modeling techniques of this chapter consider the latest development in the literature regarding the skill distribution and the selection effect on labor productivity. Adão (2016) and Alvarez-Cuadrado et al. (2019) argue that the key determinant of the selection effect is the relationship between workers' comparative advantage and absolute advantage in each sector. The specifications of most parametric Roy models in the literature implicitly assume that workers' absolute advantage in a sector increases with their comparative advantage in the same sector, which imposes the restriction that labor productivity of an expanding sector must decrease. Adão (2016) illustrates this point explicitly for independent Fréchet distribution, and the sensitivity test in Lagakos and Waugh (2013) suggests numerically that this is the case for joint Fréchet marginals linked by a copula. While Adão (2016) and Alvarez-Cuadrado et al. (2019) use non-parametric methods and find evidence against the alignment of comparative and absolute advantage in some sectors, my paper provides a parametric method to quantify the contribution of selection to measured labor productivity growth in the presence of the new selection patterns.[3]

Lastly, the selection mechanism explored in this chapter is independent of other factors studied in the literature on labor productivity growth and Baumol's disease. Triplett and Bosworth (2004) contain a review of measurement problems towards services industries in the U.S. that might have led the measured numbers to be lower than the actual ones.[4] Relatedly, the Bureau of Labor Statistics addresses measurement issues on labor input with the labor composition index (Zoghi, 2010), which adjusts labor input by the age, education, and gender of workers. My findings are orthogonal to these adjustments as I show that the selection pattern is not explained

[3]Adamopoulos et al. (2021) note in their appendices how multivariate normal distribution may allow an expanding sector to attract more productive workers, but choose to match a different set of empirical moments than those in this chapter due to the different focus of their research question and found the conventional pattern of selection.

[4]The effect of selection this chapter arrives at, however, does not depend on the measured numbers of labor productivity. It is solely inferred from worker-level income data.

away by observable characteristics of the workers. More recently, Duernecker et al. (2019) and Sen (2021) address the cost disease from the demand side and find that a subset of services industries with high productivity growth are substitutes with other economic activities, so that labor will not flow indefinitely into the most stagnant service industries. Their framework abstracts away from the heterogeneity of workers, hence the selection effect, so that my paper complements their findings from the perspective of labor supply.

1.2 Structural Transformation and Sectoral Labor Productivity Growth

The growth of employment in services concentrates in two broad subsectors in developed countries one being finance, real estate, professional and business services ("professional services"), and the other education, health services, and public administration ("EHP"). The left panel of Figure 1·1 shows the average changes in the employment share of sectors across countries. Professional services and EHP expands by 4.9 and 2.9 percentage points respectively, absorbing the 6.9-percent decline in manufacturing. Given the large volume of labor reallocation in these sectors, the average skill levels of their workers are most likely to be affected by selection. According to the conventional story of selection (Young, 2014), the reallocation of labor would have lowered labor productivity in sectors that expand, and the effect would be stronger for sectors that expand more. Therefore, we are to expect a lower labor productivity growth rate in professional services than in EHP, since professional services expands faster and therefore should have absorbed even more workers less skilled in the sector.

The data, however, do not show the expected patterns. As shown in the right

Figure 1·1: Structural Transformation in Developed Countries, 1989-2009

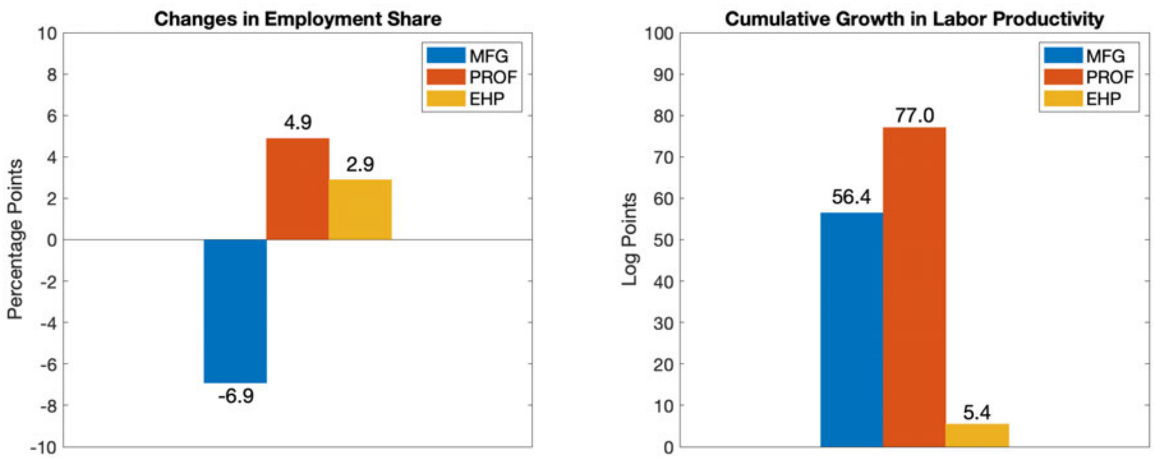

Source: GGDC 10-Sector Database
Numbers are averages of a sector across countries.
MFG: manufacturing; PROF: professional services; EHP: education, health, and public administration.
Sample includes Denmark, France, Italy, Japan, the Netherlands, Spain, Sweden, the U.K., and the U.S.
Labor productivity measured in gross value added in constant national prices per person engaged.

panel of Figure 1·1, labor productivity[5] grows on average by 56 log points for manufacturing, 77 log points for professional services, but only 5 log points for EHP for the same period across countries. These numbers mean that labor productivity has more than doubled for professional services during this period, growing even faster than for manufacturing, when the figure for EHP barely changed. Although the literature is aware that some service activities are the "most progressive" within an economy (Baumol et al., 1985; Duernecker et al., 2019; Sen, 2021), the growth rate for professional services is so large that inevitably raises a second thought about whether labor productivity of this sector is indeed adversely affected by selection.

The correlation between the growth of employment and labor productivity provides a further check for the conventional selection story. If the expansion of pro-

[5]In this chapter, labor productivity is measured in output per worker at constant domestic prices.

fessional services entails absorption of new workers with lower skill levels, then the growth in output per worker should at least be negatively correlated with the expansion of the sector across countries.[6] Figure 1·2 visualizes these correlations by plotting the growth in labor productivity against the growth in employment share for each country-sector pair. Markers with the same shape and color denote the same sector in different countries. The blue squares denote the manufacturing sector, with a clear trend that labor productivity grows more in countries where the manufacturing sector contracts more. For professional services represented by the orange triangles, we observe that labor productivity grows more in countries where the sector expands more. Finally, the correlation is slightly negative for EHP, denoted by the yellow dots. To obtain an exact number, I compute the correlation between the growth of labor productivity and the growth of employment share, weighted by the number of workers in 2009. The correlation is -0.68 for manufacturing and -0.14 for EHP, indicating that the expansion or contraction of these sectors indeed feature reallocation of workers with skill levels below the sectoral average. Meanwhile, the correlation is 0.48 for professional services, which means that greater expansion of the sector is associated with higher growth in measured labor productivity.[7] Opposite to the conventional selection mechanism as in Young (2014), this positive relationship is, at face value, consistent with a story that selection helps improve the labor productivity in professional services.

While these correlations and tests provide no definitive evidence for the pattern of selection, they do urge consideration of the possibility that selection might bring in more productive workers into professional services and thereby raise the average skill

[6]This hypothesis assumes that the technological progress in the sector does not differ substantially across countries.

[7]For a formal test, I regress the cumulative change in log real output per worker on the cumulative change in log employment share for professional services across the countries, and reject the null hypothesis that the coefficient is negative at 10% significance level.

Figure 1·2: Growth in Labor Productivity and Employment Positively Correlated in Professional Services

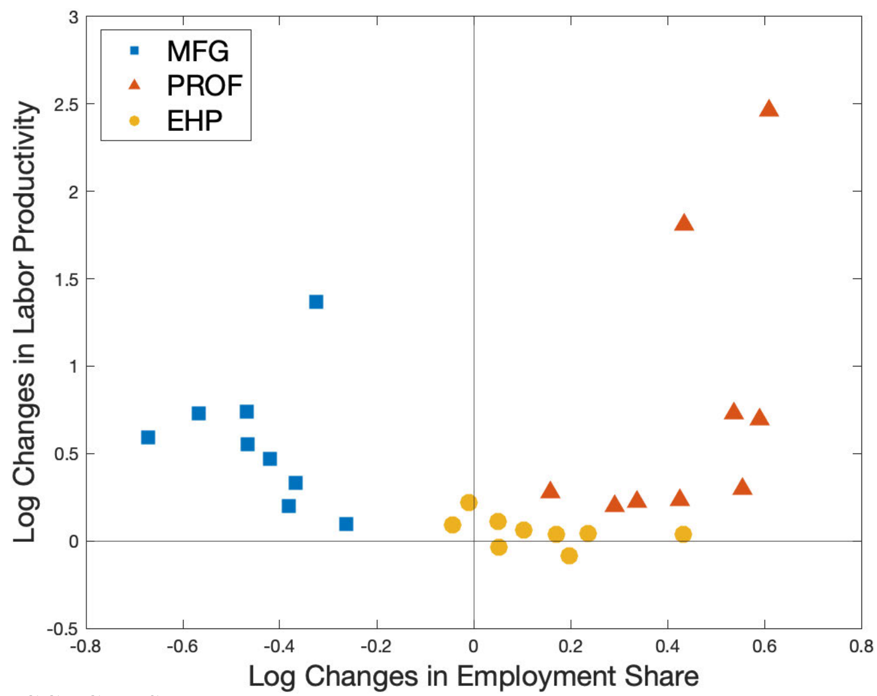

Source: GGDC 10-Sector Database
MFG: manufacturing; PROF: professional services; EHP: education, health, and public services.
Each marker denotes the cumulative change in a country-sector from 1989 to 2009.
Sample includes Denmark, France, Italy, Japan, Netherland, Spain, Sweden, the U.K., and the U.S.

level in the sector. In the following section, I exploit worker-level panel data from the U.S. to infer the exact pattern of selection for the expanding service sectors.

1.3 Micro-Level Patterns of Selection

To examine the patterns of selection at the micro level, I use the Outgoing Rotation Group of the Current Population Survey of the U.S. whose short panel structure allows for tracking the flow of workers switching sectors between consecutive years.[8] The variables I use include year and month of interview, age, gender, race, indus-

[8]Drew et al. (2014) develop a method of creating unique identifiers for all CPS person and household records from 1989 onward, which makes matching observations in the CPS much easier than it used to be.

try, weekly earnings, and the person level ID that identifies the address and person interviewed.[9] The sample spans the period between 1989 and 2020. I drop observations with less than 25 years of age since the choice of sectors is less likely long-run oriented at the very early career stage. The resulting sample contains 1,411,729 successful matches of workers with income and sector information for the current and the last year, averaging 45,540 workers per year.

Using this sample, I first list the employment share of each sector in Table 1.1 to verify that structural transformation affects manufacturing, professional services, and EHP most in the micro data. In three decades' time, the employment share of manufacturing has declined by more than 10 percentage points, while professional services and education, health services, and public administration expand most among the service sectors. The two service sectors sum up to a growth of 9.9 percentage points in their employment share, which roughly offsets the contraction of manufacturing.

Table 1.1: Employment Share in the U.S. (%), 1989-2020

Sector	1989	2020	Changes
Agriculture, Mining & Construction	8.8	9.9	1.1
Manufacturing	22.9	12.4	-10.5
Distribution Services	17.8	16.8	-1.0
Transport Services	6.7	5.8	-0.9
Professional Services	12.6	18.4	5.8
Education, Health & Public Admin.	26.3	30.4	4.1
Personal Services	4.8	6.4	1.6

Source: CPS outgoing rotation group, matched individual panel data.

To paint a rough picture of the selection effect, I take labor income as a proxy for the skill of workers within the same sector and plot the average real weekly earnings[10] by where the workers are in the previous year. Figure 1·3 shows that workers reallocated from manufacturing (MFG) to professional services (PROF) earn more than

[9]Since the CPS tracks addresses instead of people, it is necessary to exclude matches that do not represent the same interviewees across years. I screen the matched panel data by the criteria whether two observations at the same address have the same sex and race, and whether the age difference is between 0 and 2.

[10]Nominal earnings are adjusted to constant 1999 dollars with the IPUMS variable CPI99.

those already working there. Therefore, reallocation of workers from manufacturing increases the average labor productivity of professional services if we take income as reasonably reflecting the productivity of workers within a sector. The story is different for the EHP sector. As shown in Figure 1·4, manufacturing workers switching into EHP earn consistently less than workers already in the sector, so that the reallocation of workers lowers the measured labor productivity growth in EHP.

Figure 1·3: Average Income in Professional Services by Workers' Previous Sector

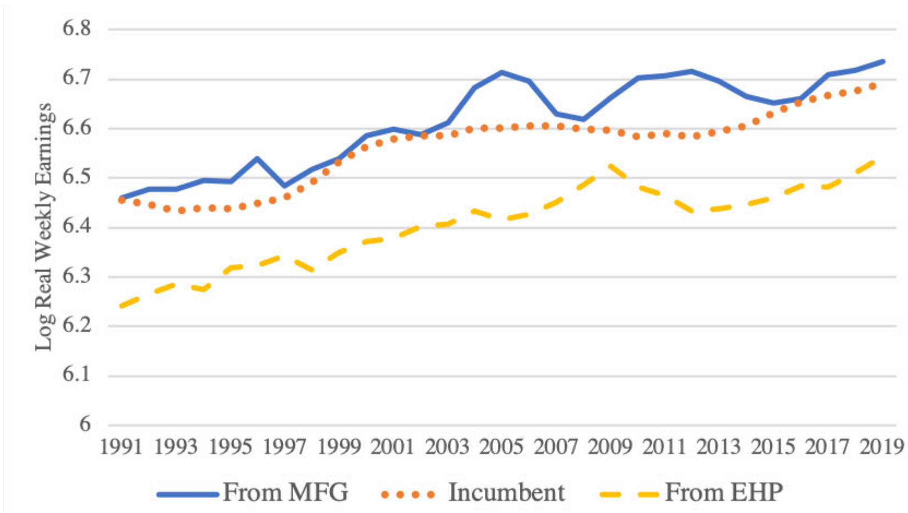

Source: Matched panel of the Outgoing Rotation Group in U.S. Current Population Survey
The levels of earnings are 3-year moving averages.

To formally test the average earnings premium of workers switching sectors over incumbent workers, I regress the log real weekly earnings in professional services and EHP on the worker's previous sector with year fixed effects. The specification is

$$y_{is't} = \sum_{s} \beta_s d_s + \sum_{t} \beta_t d_t + \varepsilon_{is't} \tag{1.1}$$

where $y_{is't}$ is the log real weekly earnings of worker i in sector s' in year t, d_s the dummy variable for the worker's sector in year $t-1$, and d_t the year fixed effect.

Table 1.2 reports the estimated coefficients of the previous-sector dummies with

Figure 1·4: Average Income in Education, Health & Public by Workers' Previous Sector

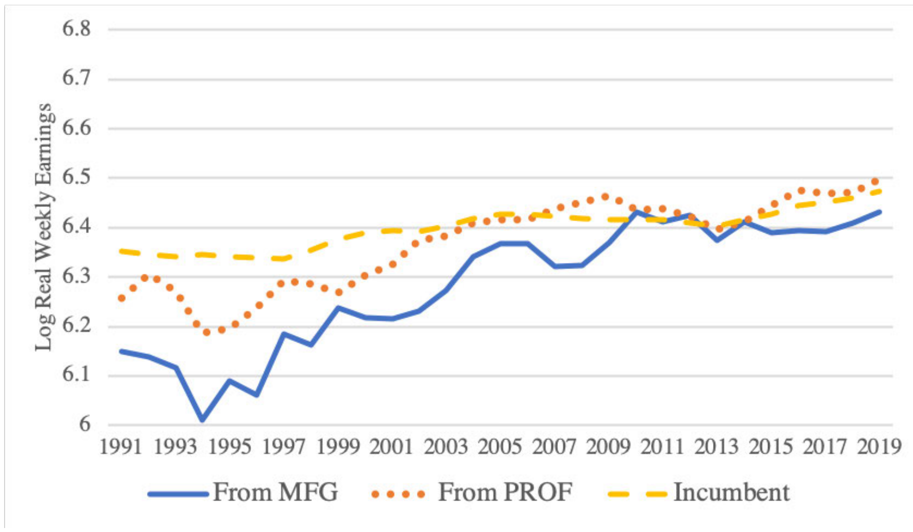

Source: Matched panel of the Outgoing Rotation Group in U.S. Current Population Survey
The levels of earnings are 3-year moving averages.

robust standard errors. The premia of reallocated workers are computed as the differences between the coefficients of the previous sector and the current sector. It can be inferred from the robust standard errors that these premia are all statistically significant. Notably, the premium of manufacturing workers reallocating into professional services is 4.8 log points, implying that these workers earn approximately 4.9% higher than incumbent professional services workers. All other premia are negative.

Based on these empirical observations, it is evident that workers reallocating from manufacturing earn significantly more than incumbent professional services workers. Those reallocating from manufacturing into EHP, however, earn considerably less. These patterns need to be considered for my quantitative exercises, and I present in the next section a generalized Roy model that allows for these rich patterns.

Table 1.2: Log Earnings by Sector of Previous Year

Dependent Variable	Log Earnings in PROF	Log Earnings in EHP
MFG in $t-1$	6.513	6.254
	(0.013)	(0.016)
PROF in $t-1$	6.465	6.340
	(0.010)	(0.010)
EHP in $t-1$	6.307	6.360
	(0.013)	(0.006)
Year Fixed Effect	Yes	Yes
Observations	187,509	399,003
MFG Premium	0.048	-0.106
	(0.009)	(0.015)
PROF Premium		-0.020
		(0.008)
EHP Premium	-0.158	
	(0.008)	

Robust standard errors in parentheses.
For each column, the dependent variable is the log real weekly earnings of a worker in the specified sector in the current year. The independent variables are dummies for whether a worker was working in the sector in the previous year. The premia are computed as the difference of subtracting the coefficient of the dummy for the current sector from the coefficient of the dummy from another sector. For example, the manufacturing premium in professional services (0.048) is the difference between 6.513 and 6.465.

1.4 A Multi-Sector Roy Model of Labor Markets

The economy consists of three sectors - manufacturing, professional services, and a composite sector of education, health services, and public administration, denoted by subscripts m, p, and e respectively. A representative firm in each sector $s \in \{m, p, e\}$ pays wage rate w_{st} in year t to each efficiency unit of labor. There is measure 1 of workers. The total amount of efficiency units employed by a sector is

$$L_{st} = \int_{i \in \Omega_{st}} z_s(i) di, \tag{1.2}$$

with $z_s(i)$ representing the efficiency units supplied by worker i to the sector and Ω_{st} the set of workers employed in the sector. The production function of each sector is

$$Y_{st} = A_{st} L_{st}, \tag{1.3}$$

where A_{st} is the exogenous sector-specific technology. The output of each sector is sold at the price p_{st}, so that each efficiency unit of labor is paid the wage rate $w_{st} = p_{st} A_{st}$ with competitive labor markets.

The supply of labor depends on a selection mechanism where each worker possesses a vector of time-invariant sector-specific skills and chooses to work in the sector that compensates most. Specifically, each worker draws a vector of skills $\mathbf{z} = (z_m, z_p, z_e)$ from a distribution and chooses to work in the sector that gives the highest return by solving the problem

$$\max_{s \in \{m,p,e\}} w_{st} z_s(i) \exp\ (\varepsilon_{st}(i)) \tag{1.4}$$

where $\varepsilon_{st}(i)$ is a transitory idiosyncratic income shock orthogonal across sectors, following a normal distribution $\varepsilon_{st}(i) \sim N(0, \sigma_\varepsilon^2)$. Taking log of the expression, we can write the worker's problem in the form

$$\max_{s \in \{m,p,e\}} \ln w_{st} + \ln z_s(i) + \varepsilon_{st}(i). \tag{1.5}$$

Under this setup, the permanent skills of workers roughly determine the sectoral share of employment in each period and shape the long-term trends of the share in response to changes in wage rates. The transitory income shock allows for workers switching sectors in any direction between two periods.

The comparative advantage in manufacturing with respect to professional services is defined as $s_m(i) \equiv \ln(z_m(i)/z_p(i))$, and the comparative advantage in EHP with respect to professional services is defined as $s_e(i) \equiv \ln(z_e(i)/z_p(i))$. They follow the relationship

$$s_m(i) = \kappa \ln \frac{q_m(i)}{1 - q_m(i)} \tag{1.6}$$

$$s_e(i) = \kappa \ln \frac{q_e(i)}{1 - q_e(i)} \tag{1.7}$$

where $q_m(i), q_e(i) \in (0, 1)$ are independent quantiles of a worker in the distribution of the comparative advantages. When $q_m(i)$ is close to 0, the value of the function $s_m(i)$ corresponds to the comparative advantage of workers least productive in the manufacturing sector with respect to professional services. When $q_m(i)$ is close to 1, the value of the function shows the relative productivity of workers with the highest degree of comparative advantage in the manufacturing sector with respect to professional services. The coefficient κ determines how diverse the working population is in terms of the comparative advantage. A high level of κ means that there is a huge gap between the best and the worst workers in terms of the comparative advantage. When κ is close to 0, workers will have similar levels of comparative advantage.

Conditional on the comparative advantages, the absolute advantage $a(i) \equiv \ln z_p(i)$ follows the Gumbel distribution

$$\{a(i) |\, q_m(i), q_e(i)\} \sim \mathit{Gumbel}\left(\alpha(q_m(i), q_e(i)), \sigma_a\right) \tag{1.8}$$

where

$$\alpha\left(q_m(i), q_e(i)\right) = \alpha_m \ln q_m(i) + \alpha_e \ln q_e(i) \tag{1.9}$$

so that α_m and α_e control the degree of dependence between the absolute advantage and the comparative advantages. When $\alpha_m > 0$, greater comparative advantage in manufacturing implies greater expected skill in professional services. When $\alpha_m < 0$, greater comparative advantage in manufacturing implies weaker expected absolute advantage. Similar interpretation follows for the coefficient α_e. The scale parameter

σ_a governs the dispersion of the conditional distribution. The higher σ_a is, the more dispersed the distribution is. The Gumbel distribution is used here so that the *level* of skills in professional services, $z_p(i)$, conditional on the comparative advantages, follows a Fréchet distribution with shape $1/\sigma_a$ and scale $\alpha(q_m(i), q_e(i))$.

Based on this characterization of sector-specific skills, workers select into sectors with the following criteria in the absence of the idiosyncratic income shocks.[11] When $s_m(i) > \ln w_{pt} - \ln w_{mt}$ and $s_m(i) > s_p(i) + \ln w_{et} - \ln w_{mt}$, worker i chooses to work in the manufacturing sector. When $s_p(i) > \ln w_{pt} - \ln w_{et}$ and $s_e(i) > s_m(i) + \ln w_{mt} - \ln w_{et}$, the worker chooses to work in education and health services. When neither of the two conditions holds, the worker chooses to work in professional services.

In equilibrium, the wage rates pin down the allocation of workers, which in turn determines the percentiles of the comparative advantages of workers indifferent between any two sectors. Consider first a worker indifferent between working in manufacturing and professional services in the absence of idiosyncratic income shocks. This worker would have comparative advantage in manufacturing at the level

$$s_m(i) = \ln w_{pt} - \ln w_{mt}, \tag{1.10}$$

which translates into

$$q_m(i) = \frac{1}{(w_{mt}/w_{pt})^{1/\kappa} + 1} \equiv \bar{q}_{mt}. \tag{1.11}$$

Therefore, the marginal percentile $\bar{q}_{mt}$ increases when the wage rate in manufacturing falls against that in professional services. When α_m is positive, the absolute advantage of reallocated workers from manufacturing to professional services keeps increasing.

We can also define a cutoff percentile $\bar{q}_e$ for the comparative advantage in EHP,

[11]The idiosyncratic income shocks allow the following inequalities to be relaxed by a small amount in any given year to allow for two-way labor flow between any pair of sectors. The shocks do not affect the long-run trends of labor realloation.

below which a worker finds EHP less attractive than professional services:

$$\bar{q}_{et} \equiv \frac{1}{\left(w_{et}/w_{pt}\right)^{1/\kappa} + 1}. \tag{1.12}$$

This cutoff percentile decreases when the wage rate in EHP increases relative to professional services. Put together, these cutoff percentiles determine the comparative advantages of workers who are most likely to switch between sectors. When the long-run trends in relative sectoral wage rates lead to reallocation of workers between sectors, these cutoff percentiles vary accordingly, which in turn determine the expected skill level of the marginal workers who enter or exit professional services.

1.4.1 Implications on the Pattern of Selection

To illustrate the flexibility this model gives to the joint distribution of skills, I use a two-sector version of the model and plot the expected log level of income in Figure 1·5. Consider an economy consisting of only manufacturing and services. Workers' productivity in the service sector is defined as their absolute advantages, and their relative productivity in manufacturing comparative advantages. Workers are ordered with increasing comparative advantage q for this illustration, so that $q = i$. The comparative advantage schedule is defined as

$$s(q) = \kappa \ln \frac{q}{1-q} \tag{1.13}$$

and the corresponding expected absolute advantage schedule as

$$\alpha(q) = \alpha \ln \; q. \tag{1.14}$$

For simplicity, assume that there are no idiosyncratic income shocks for this reduced model.

Figure 1·5 plots a thought experiment of increasing the wage rate in services.

Figure 1·5: Patterns of Selection

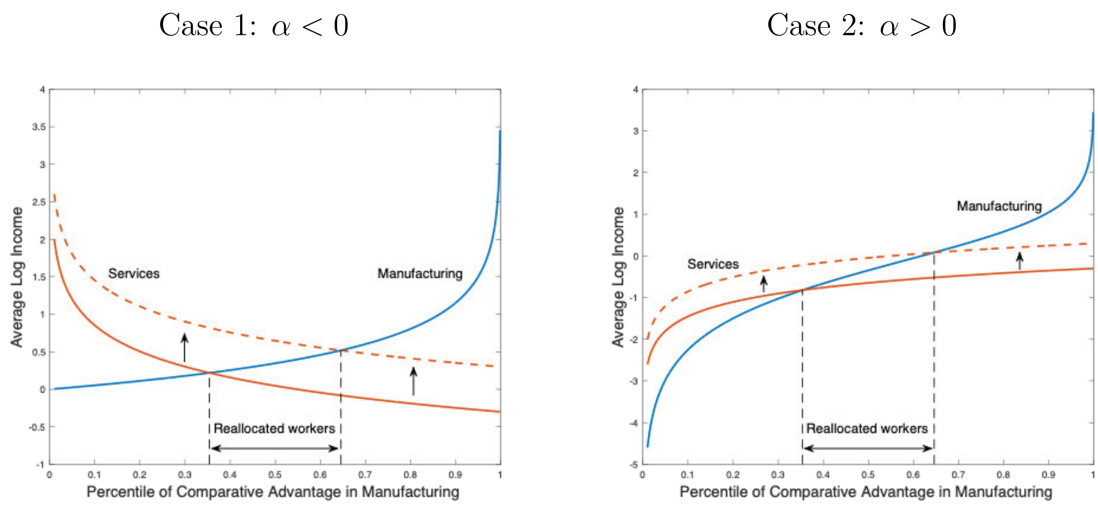

Dashed curves represent potential income from services after a wage increase.

The orange curves represent the service sector for which the skill is described by the absolute advantage $\alpha(q)$, and the blue curves depict manufacturing where the skill is the sum of absolute and comparative advantage, $\alpha(q) + s(q)$. The vertical axis describes the level of log income of a worker in each sector, which is equal to the worker's log skill level plus the log wage rate of the sector. The horizontal axis is ordered with increasing q, the percentile of comparative advantage in manufacturing, from left to right. Workers choose to work in the sector that rewards them the highest income so that the intersection of the two curves pins down the share of employment of each sector. Workers to the left of the intersection choose the service sector, while those to the right manufacturing. The dashed curves represent increases in the wage rate in services. The intersections of the curves move to the right to reflect the gains in the employment share by the service sector. Workers between the two intersections reallocate from manufacturing to services.

Case 1 illustrates the income distribution implied by unproductive absorption of workers by services, when $\alpha < 0$. For a worker at a higher quantile q, the decline in absolute advantage is more than compensated for by the increase in comparative advantage. We can see from the figure that reallocated workers earn least in the service sector after the change. As the difference between log income and log skill is just the log wage rate and thus constant across workers in the same sector, these reallocated workers are also the least productive in services. Therefore, their reallocation brings down the average labor productivity of services. Independent Fréchet distribution is a special case of this specification where $\alpha + \kappa = 0$, $\kappa \in (0, 1)$.[12] This means that a model specifying the joint skill distribution as independent Fréchet necessarily implies lower average skill level for an expanding sector.

Given the value of κ, varying the value of α can be thought of as rotating the curves. A higher value of α brings down the left ends and boosts the right ends of the curves. When $\alpha = 0$, it is easy to visualize that the curve for services would be a horizontal line because workers' productivity in the service sector is independent of their comparative advantage in the goods sector. When $\alpha > 0$, workers' skills in both sectors increase with their comparative advantage, as Case 2 depicts. Consequently, workers reallocated into services have better skills in the sector than the incumbent services workers, bringing up the average level.

The discussion above demonstrate the flexibility of modeling selection based on comparative and absolute advantages in terms of generating patterns that lead to different implications for the average skill levels when workers reallocate. The key for achieving this flexibility is the introduction of the parameter α that relaxes the parametric restriction of the typical independent Fréchet distribution.

[12] See Adão (2016) for the derivation in the appendix.

1.5 Quantitative Analysis

1.5.1 Calibration

To obtain the effect of selection on labor productivity implied by the model, I calibrate the parameters to the micro-level observations presented earlier. To align with the setup of the model, I retain only the data of individuals working in manufacturing, professional services, or EHP for both years they are surveyed in the CPS monthly data. This leaves 781,109 observations in the sample. In the reduced sample, 36% of workers work in manufacturing in 1989, 20% in professional services, and 44% in EHP. These numbers become 19%, 29%, and 52% for 2020.

Parameters and Targeted Moments

The parameters in the model are jointly determined. However, each parameter corresponds intuitively to a moment and can be viewed as the major determinant for that moment. This correspondence is described below.

As stated in the last section, α_m and α_e are the most important parameters for the selection mechanism. They determine how workers' productivity in professional services changes with the degree of their comparative advantages in the other sectors. As wages for efficiency units of labor change, the differences between the income of workers switching into professional services and those staying in the sector inform us about whether the new workers are more or less productive than incumbent workers. Therefore, I target the average real weekly income differences between workers switching into professional services and workers staying in the sector.

The parameter κ describes how dispersed comparative advantages are for the working population. The higher κ is, the more different workers are in their productivity within manufacturing or EHP. Consequently, there would be larger earnings gaps between workers reallocating form manufacturing into EHP and incumbent EHP workers. Therefore, I target the average real weekly income differences between workers switching into EHP from manufacturing and workers staying in EHP.

The degree of dispersion of the absolute advantage is governed by σ_a, and the targeted moment is the average variance of income of all workers across the sample period.

Given the key parameters of the joint skill distribution, the standard deviation of the idiosyncratic sector-specific income shocks σ_ε targets the average frequency of sector switches across all years. For each year-survey round, the wage rates for efficiency units of labor target the employment share of the sectors and average log real income.

The calibrated values of the parameters are reported in Table 1.3. The value of α_m is indeed positive, verifying that the skill of new professional services workers from manufacturing is increasing in their comparative advantage in manufacturing. The negative value of α_e means that workers are less productive in professional services if they have stronger comparative advantage in EHP. The values of the targeted moments are listed in Table 1.4.

Table 1.3: Values of Parameters

Parameter	Value	Interpretation
α_m	0.0306	Relationship between comparative advantage in MFG and absolute advantage in PROF
α_e	-0.1665	Relationship between comparative advantage in EHP and absolute advantage in PROF
κ	0.2528	Dispersion of comparative advantages
σ_a	0.4736	Dispersion of absolute advantage
σ_ε	0.0318	Dispersion of idiosyncratic income shocks

In addition to these parameters, the calibrated sectoral wage rates grow by -1.2, 24.4, and 24.9 log points respectively for manufacturing, professional services, and EHP throughout the sample period. The large and similar increases in the two service sectors are in line with the similar percentage gains in their employment share.

Table 1.4: Targeted Moments

Moment	Data	Model
Income premium of MFG workers switching into PROF	0.047	0.047
Income premium of EHP workers switching into PROF	-0.159	-0.159
Income premium of MFG workers switching into EHP	-0.117	-0.117
Frequency of sector switches	0.060	0.060
Variance of log income	0.460	0.460

Recall that the cuttoff percentiles of the comparative advantages of workers follow

$$\bar{q}_{mt} \equiv \frac{1}{(w_{mt}/w_{pt})^{1/\kappa} + 1} \tag{1.15}$$

$$\bar{q}_{et} \equiv \frac{1}{(w_{et}/w_{pt})^{1/\kappa} + 1}, \tag{1.16}$$

the changes in wage rates imply that the marginal worker between manufacturing and professional services in 2019 has much higher comparative advantage in manufacturing compared to the marginal worker in 1989. Meanwhile, the comparative advantage of the marginal worker between professional services and EHP has stayed similar after 30 years.

Implication on Labor Productivity

With the calibrated parameters I examine the implication of the model on sectoral labor productivity changes. The left panel of Figure 1·6 plots the indices of the implied average human capital of each sector since 1989, while the right panel plots the share of employment of each sector. The expansion of professional services contributes positively to its labor productivity growth until 2009, after which the effect becomes negative. The EHP sector has exactly the opposite trend with falling average human capital until 2009 but a rising trend afterwards. Consequently, the selection mechanism enlarges the gap in average human capital between the two service sectors, which reaches its peak in year 2009.

Why does the gap shrink after 2009? We can infer the reason from the changes in

Figure 1·6: Average Human Capital (Simulated) and Employment Share (Matched)

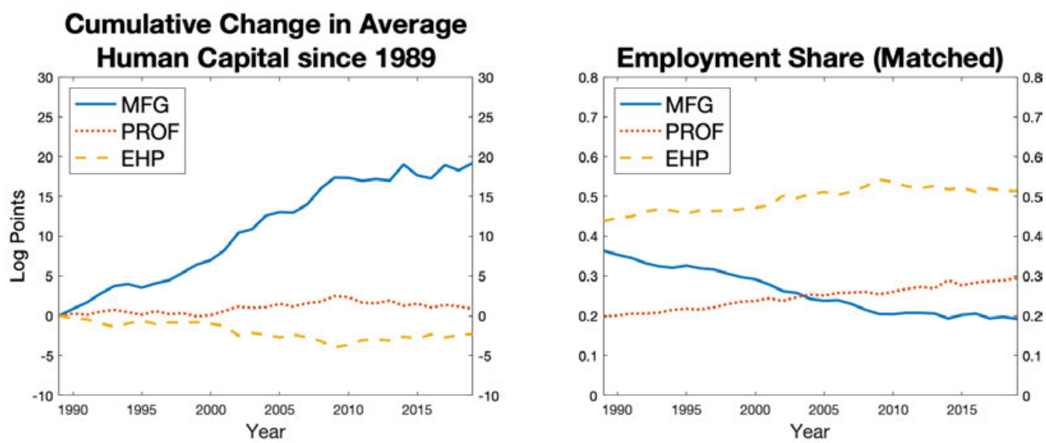

the employment share in the right panel of the figure. Before 2009, manufacturing is the only sector declining in employment share so that the new workers joining professional services are mostly from manufacturing. The pattern of labor reallocation, however, changed after 2009. The employment share of manufacturing declined much more slowly while that of EHP starts to drop. This pattern implies that the comparative advantage in EHP starts to rise for workers moving from EHP into professional services. The negative value of α_e then responds to the increase in q_e and lowers the skill level of the marginal worker between professional services and EHP. Therefore, the professional services sector expands with declining average human capital after 2009.

To compare the implied selection effect on sectoral human capital to observed sectoral labor productivity growth, I turn to data from the Bureau of Economic Analysis for output quantity indices by sector. Figure 1·7 plots the cumulative log changes in the observed sectoral labor productivity, from which we observe that the structural break for manufacturing and professional services in data around the Great Recession is also present in the simulated model.

Figure 1·7: Cumulative Changes in Labor Productivity from Data

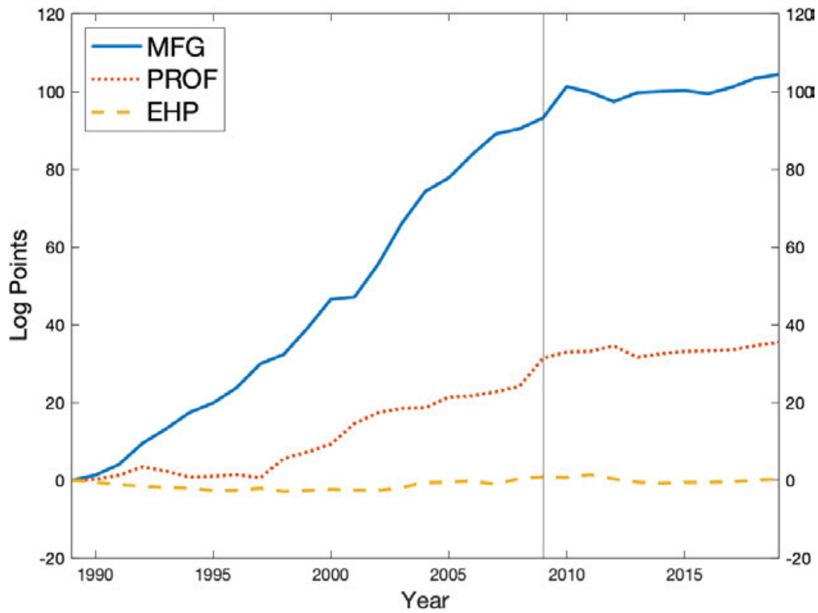

Source: Bureau of Economic Analysis and own calculation.

Table 1.5 summarizes these findings by time period. Panel A shows that for the entire sample period between 1989 and 2019, the calibrated model implies that selection leads to increases of 19.2 log points in labor productivity for manufacturing, 0.9 for professional services, and a decrease of 2.3 log points for EHP. In comparison, the numbers computed from BEA data are 104.4, 35.5, and 0.4 respectively for the same period. Therefore, the underlying true labor productivity growth, or technology growth, is 85.2, 34.7, and 2.7 log points for the three sectors respectively. We can compute with these figures that selection adds 23% and 3% to the technology growth rate for manufacturing and professional services, while downplaying the technology growth in EHP by 85%.

With the growth rates by sector, I can aggregate the selection effect for professional and pubic services into the effect for a single service sector. Weighting the sectoral growth rates by the mean sectoral share of nominal value added, I find that selection only leads to a decrease in labor productivity by 0.4 log point, which is 2% of the

Table 1.5: Effect of Selection on Sectoral Labor Productivity Growth

Cumulative growth in log points	Selection (Model)	Data	Implied Techology Growth
A. 1989-2019			
MFG	19.2	104.4	85.2
PROF	0.9	35.5	34.7
EHP	-2.3	0.4	2.7
Aggregate services	-0.4	21.1	21.5
B. 1989-2009			
MFG	17.4	93.2	75.9
PROF	2.5	31.4	29.0
EHP	-4.0	0.8	4.8
Aggregate services	-0.3	18.3	18.6
C. 2009-2019			
MFG	1.9	11.2	9.3
PROF	-1.6	4.1	5.7
EHP	1.7	-0.5	-2.1
Aggregate services	-0.3	2.2	2.5

21.5-log point true growth rate. Therefore, the gap of labor productivity growth between manufacturing and aggregate services is thus reduced from 83.3 log points to 63.7 log points, roughly corresponding to an 22% reduction in percentage terms, where almost all of the effect comes from workers selecting out of manufacturing.

For the period between 1989 and 2009, as shown in Panel B of Table 1.5, selection has its strongest effect on average human capital, adding 23%, 9%, and -83% to the implied true technology growth. We can see that the effect is much stronger for professional services during this period, while being similar for the other two sectors. Panel C records that the effect of selection reversed for the two services sectors after 2009, due to the change in the direction of labor reallocation between sectors.

The seemingly small effect of selection, however, is a significant result in terms of how it improves out understanding of a selection mechanism from a conventional model. I show this point by comparing my baseline specification to the one in which skills are drawn from independent Fréchet distributions.[13] For this specification, wage

[13]I set the shape parameter of the Fréchet distributions to 3 for this exercise, which is typical within

rates are calibrated to match the employment share of each sector in each year.

Table 1.6 reports the implication of the selection mechanism on cumulative labor productivity growth by the two specifications. Obtaining similar effect for manufacturing, the two specifications however generate qualitatively different patterns for professional services. As shown in panel A of the table, the discrepancy for professional services between the two specifications is 14.2 log points, meaning that the contribution of the selection mechanism to labor productivity growth can be mismeasured by 15 times as much in the opposite direction. Consequently, as shown in panel B, the specification with independent Fréchet skill distribution overpredicts the technology growth in professional services by about one-third of the growth rate under the baseline. Aggregating the two services sectors together under the independent Fréchet specification, selection is shown to decrease measured labor productivity by 10 log points. As a result, the implied growth under the independent Fréchet specification overpredicts the true technology growth by almost 50% compared to the baseline specification.

Table 1.6: Selection Effect on Labor Productivity by Specification

	Growth in log points	
A. Average Human Capital	Baseline	Independent Fréchet
Manufacturing	19.2	21.2
Professional Services	0.9	-13.3
Education and health	-2.3	-5.3
Aggregate Services	-0.4	-10.0
B. Implied True Growth	Baseline	Independent Fréchet
Manufacturing	85.2	83.2
Professional Services	34.7	48.9
Education and health	2.7	5.7
Aggregate Services	21.5	31.1

Figure 1·8 plots the average human capital of the sectors, illustrating the differ-

the range of values found in the literature (e.g. Bryan and Morten (2019) and Hsieh et al. (2019)). The estimate in Young (2014) is equivalent to having a shape parameter of 1.33, which generates even stronger negative selection effect for both service sectors ences in panel A of Table 1.6 more intuitively. The solid lines represent the results obtained from the baseline specification while the dashed lines come from the independent Fréchet specification. The effect of selection on labor productivity growth of manufacturing is very close, but vastly different for professional services and the gap between professional and education and health services. Since the average skill level of a sector is entirely determined by its employment share in an independent Fréchet world, the expansion of professional services only brings down the average quality of workers in the sector, which is qualitatively different from the implication by the general specification of the baseline. The declining average human capital of professional services means that the independent Fréchet specification cannot capture the fact that workers reallocating to professional services from manufacturing earn more than the incumbent professional services workers. Moreover, the Fréchet specification does not capture the trend reversal for professional services around the Great Recession because employment share is a sufficient statistic for average sectoral human capital and thus it does not matter from which sector a worker switches reallocates.

Figure 1·8: Index of Average Human Capital

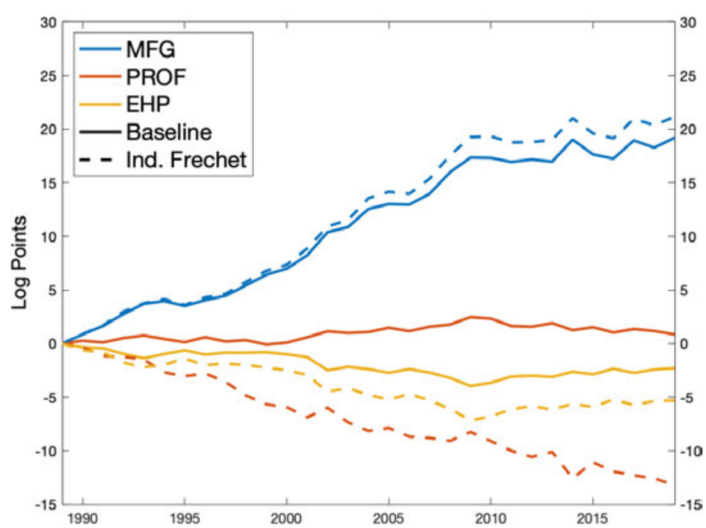

Figure 1·9 plots trends in panel B of Table 1.6 on the implied true technology growth. We can observe that the independent Fréchet specification systematically overpredicts growth in both service sectors, especially for professional services.

Figure 1·9: Index of Implied Sectoral Technology

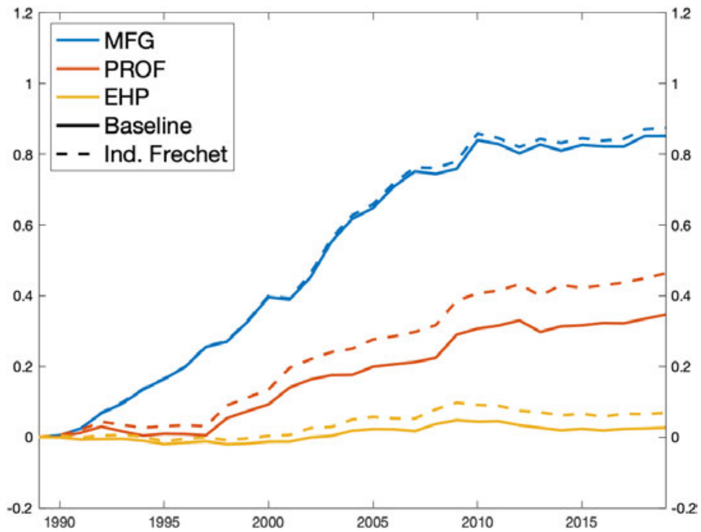

1.5.2 Cross-Country Comparison

The calibrated model can be extended to examine the effect of selection in other developed countries. For this exercise, I use real output and employment data from the GGDC 10-Sector Database for a set of the largest advanced economies, including France, Italy, Japan, Spain, and the United Kingdom, from 1990 to 2010 and feed the data into the model. Specifically, I match the employment share in each sector for each country in the first and the last year in the sample, and examine the effect of selection on labor productivity growth.

Table 1.7 reports the results of this exercise. Overall, the selection effect is largest for the manufacturing sector both in terms of the magnitude and as a percentage of the implied technology growth rate. Specifically, selection adds 13.7 log points to the measured labor productivity growth in France, a similar fraction as in the U.S. The effect is larger for Japan, Spain, and the U.K. For Spain and Italy where technology growth in manufacturing is slow, the effect of selection becomes the dominating force, being larger than the technology growth in Spain and more than 10 times as large in Italy. On average, selection contributes more than 14 log points to measured labor productivity increases across countries.

The effect of selection is mostly positive for professional services and negative for public services, except for Italy. This is due to the fact that Italy is the only country in the sample that experienced a decreasing share in employment for professional services and an increasing share for EHP, so that selection works in the opposite direction for these two sectors in Italy. On the other hand, the two services sectors are mostly affected by the selection mechanism in Japan, where technological progress in professional services is magnified by 21% and that in public services underestimated by 81%. Having adjusted for the selection effect, we see that the technology growth of EHP in Japan becomes the highest among these countries while the raw measure from data is the lowest. In general, the selection effect adjusts up the measured labor productivity growth for education, health services, and public administration sectors commonly thought to have stagnant productivity growth and makes their implied technology growth much closer across countries.

Comparing the selection effect across sectors, we can see that the effect is an order of magnitude larger for manufacturing than for the service sectors, not the mention that the opposite effect in the two service sectors work against each other. These cross-country observations confirm that the selection mechanism closes a fraction of the labor productivity growth gaps between sectors mainly through affecting manufacturing, while the aggregate service sector is largely unaffected.

1.5.3 Quantitative Exercises with General Equilibrium

Having accounted for the changes in labor productivity in the past decades, the selection model is also helpful in shaping our expectations for the future. To answer

Table 1.7: Selection Effect on Labor Productivity across Countries

	Cumulative changes in log points			
A. Manufacturing	Selection Effect	Data	Technology Growth	Selection as % of Tech Growth
France	13.7	72.0	58.3	23%
Italy	8.4	9.0	0.6	1400%
Japan	17.9	60.0	42.1	43%
Spain	13.9	22.0	8.1	172%
United Kingdom	19.4	57.0	37.6	52%
B. Professional services	Selection Effect	Data	Technology Growth	Selection as % of Tech Growth
France	0.6	25.7	25.1	2%
Italy	-2.1	30.6	28.5	-7%
Japan	3.9	22.6	18.7	21%
Spain	0.5	73.1	72.6	1%
United Kingdom	0.9	72.3	71.4	1%
C. Education, health & public	Selection Effect	Data	Technology Growth	Selection as % of Tech Growth
France	-1.2	4.6	5.8	-21%
Italy	2.4	10.1	7.7	31%
Japan	-8.9	2.1	11.0	-81%
Spain	-2.2	5.2	7.4	-30%
United Kingdom	-2.5	3.8	6.3	-40%

Data from GGDC 10-Sector Database.

questions about how selection will affect sectoral labor productivity growth in the coming decades, I construct the demand side of the model and simulate the general equilibrium outcomes with given rates of sector-specific technology growth.

I assume that there is a representative household that purchases from different sectors subject to non-homothetic CES preference as in Comin et al. (2021), of which the aggregate consumption index Y_t is implicitly defined by

$$\sum_s \gamma_s^{\frac{1}{\eta}} \left(\frac{Y_{st}}{Y_t^{\xi_s}} \right)^{\frac{\eta-1}{\eta}} = 1 \tag{1.17}$$

where η is the elasticity of substitution between sectors, ξ_s affects the income elasticity of demand, and γ_s is the level shifter for sectoral demand.

Subject to this preference, the representative household minimizes in each period total expenditure

$$E_t = \sum_s p_{st} Y_{st}. \tag{1.18}$$

The demand curves of the sectoral output follow

$$Y_{st} = \gamma_s \left(\frac{p_{st}}{E_t} \right)^{-\eta} Y_t^{(1-\eta)\xi_s}, \tag{1.19}$$

and the market share of a sector is expressed as

$$\omega_{st} \equiv \frac{p_{st} Y_{st}}{E_t} = \gamma_s^{\frac{1}{\eta}} \left(\frac{Y_{st}}{Y_t^{\xi_s}} \right)^{\frac{\eta-1}{\eta}}. \tag{1.20}$$

One property of this non-homothetic CES preference is that the consumption index Y_t can be written as

$$Y_t = \frac{E_t}{p_{mt}} \omega_{mt}^{\frac{1}{1-\eta}} \tag{1.21}$$

as I normalize $\gamma_m = \xi_m = 1$. Therefore, the elasticity of substitution η can be obtained from the equation

$$\ln g_Y = \ln g_E - \ln g_{p_m} + \frac{1}{1-\eta} \ln g_{\omega_m} \tag{1.22}$$

where g_X denotes the net growth of variable X throughout the sample period. In the equation above, the price levels are measured as the quotients of nominal value added divided by real value added. The consumption index Y_t is calculated as the sum of sectoral real value added. The expenditure level E_t is the sum of nominal sectoral value added. The expenditure share ω_{st} is the share of nominal value added of sector s.

The parameters ξ_s governing the income elasticity are inferred from the equation

$$\ln g_{Y_s} = -\eta(\ln g_{p_s} - \ln g_E) + (1-\eta)\xi_s \ln g_Y \tag{1.23}$$

for $s = \{b, p\}$. Lastly, the taste parameters γ_s are set to match the levels of real value added in the last year.

To calibrate the preference parameters, I continue to use data from the BEA for the period 1989 to 2019. I infer price data from the ratio of nominal value added and the quantity index of each sector, and output is adjusted for employment growth of the sum of the three sectors throughout the sample period. Aggregate real output is calculated with the real output by sector of each year and the price level of 2012, which is the base year for the BEA data.

The calibrated parameters are reported in Table 1.8 with their implied income elasticity of demand. These parameter values imply that the demand for professional services is most income elastic, while that for EHP is least income elastic.

Table 1.8: Values of Preference Parameters

η	ξ_m	ξ_p	ξ_e	γ_m	γ_p	γ_e
0.061	1	1.522	0.345	1	0.048	133.8

Lastly, I adjust the sectoral technology A_{st} and the scale of wage rates obtained from the benchmark simulation for 2019 to ensure that the demand for each sector is equal to the supply and that total expenditure is equal to total income.

I use this fully calibrated general equilibrium specification to compare how my model improves our understanding about the prospects of structural transformation in a world with selection. The quantitative experiment is that I extend the growth rates of sectoral technology implied in the previous section for the next 50 years. The implied annual growth rates are 2.84% for manufacturing, 1.16% for professional services, and 0.09% for EHP. This experiment feeds these exogenous growth rates into both my model and a conventional model with independent Fréchet distribution and compare how economic variables differ in 50 years. The results are reported in Table 1.9, with column (1) corresponding to my baseline specification and column (2) the conventional model with independent Fréchet distribution.

Comparing panel A and B, we can observe that my baseline model achieves more structural transformation in expenditure with less reallocation of workers than the conventional model, with the same exogenous technology growth. The reason is that labor reallocation into services is less costly in my model not only that workers reallocating into education, health, and public services are more productive in my model than in the conventional model, the professional services sector in fact *benefits* from labor reallocation. These different effects of selection on human capital are listed in panel C, which ultimately lead to higher growth in output per worker in the service sectors in my baseline specification in panel D. As a result, since labor reallocation in my baseline involves less loss in workers' skills during the structural transformation, the welfare gains in 50 years' time in my baseline more than doubles the gains predicted by a conventional model, as shown in panel E. This set of comparison shows that we are in a world where the effect of selection does not impose as much negative effect on future structural transformation as in the conventional model. Consequently, this can be interpreted as good news since the selection mechanism does not impede future growth as much as we used to think.

1.6 Conclusion

In this chapter, I study the effect of selection on the labor productivity growth of the services sector in the U.S. I conclude that selection is responsible for only an extremely small portion of the low labor productivity growth in services, which starkly contrasts with the hypothesis in the existing literature. I find that selection adds to the true technology growth of labor productivity in professional services by 0.9%, but lowers that in education, health services, and public administration by 2.3%, with the overall effect on aggregated services being only -0.4%. In a way, the quantitative counterfactuals in this chapter show that we need not worry about the adverse effect of selection on the measured labor productivity of the service sector.

Table 1.9: Projection of Structural Transformation, 2019-2069

Cumulative growth	Baseline (1)	Independent Fréchet (2)
A. Change in expenditure share (pct. points)		
MFG	-15.7	-13.6
PROF	6.3	5.6
EHP	9.4	8.0
B. Change in employment share (pct. points)		
MFG	-13.3	-14.2
PROF	4.9	5.4
EHP	8.4	8.8
C. Growth of average human capital (log points)		
MFG	31.8	46.4
PROF	0.2	-5.6
EHP	-3.1	-5.3
Aggregate Services	-1.3	-5.4
D. Growth of output per worker (log points)		
Manufacturing	173.8	188.4
Business Services	59.8	52.2
Public Services	1.3	-0.9
Aggregate Services	21.6	20.4
E. Change in welfare (log points)		
	30.1	14.3

For both columns, the sectoral technology growth rates are exogenous and equal to the annualized growth rates from 1989 to 2019.

On the other hand, Baumol's cost disease of services is still by and large a major theme of structural transformation in advanced economies, especially for education, health, and public services.

I infer the patterns of selection with the help of worker-level panel data. I show that significant gaps exist between the earnings of incumbent workers and newcoming workers in a sector, and argue that these earnings gaps are the relevant empirical moments to match for the quantitative model. Importantly, I document that workers reallocated from manufacturing into professional service earn *more* than incumbent professional services workers, which is a fact that cannot be replicated by conventional selection models in the macro development literature based on independent Fréchet distribution of sklls. To match this fact, I build a generalized Roy model that allows for rich selection patterns. As I determine the parameters of the model with the relevant empirical moments, my model correctly replicates the facts that reallocation of workers from manufacturing *improves* labor productivity in professional services, while lowering that in education, health, and public services.

Focusing on the effect of selection, this chapter abstracts away from other factors of sectoral productivity growth and takes technology as exogenous. Although the drivers of structural transformation are considered independent of the selection mechanism in this chapter, it can be well expected that there be interactive forces between selection and other sources of productivity growth, for example endogenous innovation and spillover of human capital (e.g. Jarosch et al. (2021)), that motivate further studies with richer theoretical frameworks.

Chapter 2

Occupational Dynasties and Cross-Country Productivity Differences

Siddharth George and Martin Shu

2.1 Introduction

The occupational choices of workers are highly correlated with those of their parents. For example, the caste system in India segregates people at birth and imposes a social norm over the types of occupations one may undertake. More generally, most sons of farmers continue to farm in the developing world, where the vast majority of the labor force concentrates in the agricultural sector. On average, the likelihood of a worker entering an occupation is more than 10 times higher around the globe if the worker's father works in that occupation. This gap points towards a potential source of human capital misallocation, where incumbent parents make it easier for their children to enter their occupations, crowding out opportunities for other workers with potentially better skills. As a result, there could be considerable loss of efficiency at the aggregate level.

This paper measures the degree of intergenerational persistence in occupations and its potential impact on aggregate labor productivity across time and space. We document the persistence of occupational outcomes for 90 countries, mostly not of high-income status, from 275 censuses spanning the period from 1960 to 2017. We use what we call the *Dynastic Bias*   the odds ratio of the probability of choosing an occupation conditional on whether one's father is in that occupation to measure the degree of intergenerational persistence in occupations. With the measured bias for each occupation in each census, we calibrate a quantitative Roy model of occupational choice to calculate the potential impacts of these biases on aggregate labor productivity.

At the occupation level, we document opposite correlations of the eliteness of an occupation and the average years of schooling with the dynastic bias. The eliteness of an occupation is approximated by the average living condition of the workers in an occupation, which captures to an extent the advantage an occupation gives its workers in accumulating resources. The eliteness measure is positively correlated with the dynastic bias of the occupation, while the average years of schooling is negatively correlated with the dynastic bias.

At the aggregate level, we document that the dynastic bias varies in a non-monotone way with income and education across countries it first rises and then falls with both GDP per capita and the average years of schooling of an economy. We decompose the dynastic bias and find that the composition of occupational structure accounts for the initial rise, but it is the decline in the unconditional mean of the dynastic biases across occupations that accounts for the lower dynastic bias of richer and better-educated economies.

We exploit the consistent occupational classification of the U.S. historical full count census data over a long period of time to examine the contribution of different broad occupational caterogies to the aggregate-level dynastic bias. We find that professional occupations and craftsmen have particularly high dynastic bias. Sales, services, and clerical occupations, on the other hand, are less dynastic. Surprisingly, farm laborers have the lowest dynastic bias among all broad occupational categories, which is against the common perception that farming is particularly persistent over generations. This is because the probability of workers becoming farm laborers with fathers in other occupations is also high for the agricultural occupations.

To explore the aggregate implication of the dynastic bias on labor productivity, we build a high-dimensional quantitative Roy model of occupational choice with labor market frictions. Workers draw occupation-specific skills and select into occupations with the highest return. Without labor market frictions, the probability of a worker entering an occupation is the same regardless of the occupation of the father. However, labor market frictions cost a fraction of one's income unless one's father is in that occupation. These barriers create incentive for workers to follow the same career path as their fathers, which results in the misallocation of human capital.

We calibrate the model and perform counterfactual exercises to infer the potential gains in TFP for all economies in our sample. The entry barriers of occupations are set to match the dynastic bias of each occupation. We then remove all entry barriers and solve for the new equilibrium so that workers pursue their occupations entirely based on their comparative advantage without any income penalty. This frictionless counterfactual leads to labor productivity gains at the aggregate level that average at 8.6% for all censuses in our sample. Focusing on the latest censuses of the 51 non-high-income countries with data after 2000, we find that the average potential gains in labor productivity remains considerable at 8.7%, with 20 countries having potential gains above 10%.

Our paper contributes to three strands of literature. First, the paper relates to the rising literature of the effect of labor market frictions on aggregate productvity in the presence of self-selection by heterogeneous workers. For example, Bryan and Morten (2019) study the effect of reduced geographic mobility barriers on the aggregate labor productivity of Indonesia. More closely, Hsieh et al. (2019) show that between 20% and 40% of the growth in output per capita in the U.S. since 1960 can be attributed to the improvement in the allocation of talent across occupations. Hurst et al. (2021) take one step further to study occupational sorting based on task content. Specific to the effect of intergenerational persistence is the work by Lo Bello and Morchio (2021), who build a quantitative selection model of occupational choice for the United Kingdom and show that parental networks account for over 70% of the total intergenerational persistence. Their quantitative exercise predicts that shutting down parental networks increases output per worker. Compared with these studies that focus on specific countries, our paper provides a cross-country perspective to illustrate the scope of the potential misallocation problem, especially for developing countries.

One paper of particular relevance to our work is Sinha (2016), which documents intergenerational persistence of occupational outcomes for 30 countries and investigates the effect of the persistence on aggregate productivity. Our work differs from Sinha (2016) in several major aspects. First, the primary data sources in Sinha (2016) are several nationally representative surveys while ours are the censuses from IPUMS-International. We opt for the IPUMS data for much larger coverage over 90 countries across various stages of development, and for much more observations included in each cross-section that considerably reduce measurement errors. Second, we relate our dynastic bias measure to more occupation- and country-level characteristics beyond GDP per capita. Third, Sinha (2016) attributes all the excessive persistence to the financial constraints of education, which limited his quantitative exercises to Tanzania and India where he can obtain measures of financial frictions. In contrast, we stay agnostic about the source of the labor market frictions when making quantitative statements about the potential TFP gains for all countries in our sample.

Our results also shed light on the drivers of social mobility. Intergenerational mobility is a topic of longstanding interest in economics, with Solon (1992) one of the first papers to estimate intergenerational income mobility. Solon (1999) provides an empircal overview on the estimation of intergenerational mobility, while Black and Devereux (2011) survey more recent literature on the causal mechanisms. The rising availability of large-scale administrative data on earnings and ability to link parents and children has spawned a resurgence of interest and research on the determinants of intergenerational mobility. A series of papers have shown that social mobility varies significantly across geography within the U.S. (Chetty et al., 2014). Factors influencing the childhood environment, such as the quality of school and safety, were shown to be key determinants of adult outcomes (Chetty et al., 2016). By contrast, occupational barriers have received less attention in the literature, even though one's ability to enter the middle-class or upper deciles of the income distribution are highly dependent on entering an occupation that pays well. One exception is the work by Long and Ferrie (2013), who compare the intergenerational occupational mobility in the Great Britain to that in the U.S. While these authors find declining mobility in the U.S., our data with much larger sample size suggest that the intergenerational persistence in occupations maintains a slightly downward trend. We contribute to this literature by highlighting the effect of occupational barriers on social mobility.

Third, we contribute to a small literature on dynasties in specific occupations. Several papers have studied political dynasties. This literature, summarised recently by Geys and Smith (2017), documents the private returns from family networks in politics (Smith, 2012; Querubín, 2016). Dal Bó et al. (2009) documents that marginal winners of U.S. House races are more likely to have family members subsequently hold political office, and Cruz et al. (2017) describe the advantages that family networks confer during elections. Several papers study the consequences of dynastic rule. Besley and Reynal-Querol (2017) show that hereditary rulers (like monarchs) can be motivated by career concerns to perform well while in office. George (2020) studies the overall effect of dynastic politics, identifying both positive bequest motives and negative effects of inheriting political capital. A larger literature has investigated family firms. Most research finds that family firms tend to be worse managed (Bloom and Van Reenen, 2007; Bennedsen et al., 2007) and examines theoretical reasons why (Burkart et al., 2003). Relative to these papers, we document the phenomenon of occupational dynasties over a wide variety of occupations and countries. We also highlight the features of dynastic occupations and discuss the consequences for economic development.

2.2 Data and Empirical Patterns

2.2.1 Data and Measurement of Intergenerational Persistence

In this section, we describe the process by which we construct the data set and compute the measure of intergenerational persistence for each occupation and country. We compile data on occupational dynasties based on census data publicly available via IPUMS International. We combine 269 censuses from 89 countries, including rich-countries such as the United States, middle-income countries such as China and Indonesia, and poor countries such as Uganda. We also incorporate 6 waves of the Indian National Sample Survey since no Indian census was available. Overall, our sample includes 275 cross-sections from 90 countries.

For each cross-section (census/survey), we begin by identifying the occupation codes that correspond to undocumented occupations or those that are not linked to any occupation.[1] The next step is to restrict attention to individuals who have a non-missing occupation for themselves and their fathers.[2] At this stage, we lose a significant share of the sample. This is largely because the occupation of a father

[1]We do not count a label as missing as long as it possesses the slightest amount of information. For example, we keep observations with occupation "laborers, not elsewhere classified" but drop those with "no specified occupation".

[2]We do not consider the occupation of mothers due to limited data availability.

would be missing if the son or daughter does not reside with the father. Co-residence rates vary across countries, but it is common for adult children to live separately from their parents. Moreover, co-residence rates may vary across occupations, higher for some (e.g. farmers) but lower for others that may require individuals to migrate to urban areas (e.g. economists). However, we will show that this co-residence problem does not have much effect on our measure of intergenerational persistence at the country level with historical U.S. full count census data.

The third step is to generate a dummy variable for occupational dynasties, which takes value one for all individuals who are in the same occupation as their fathers. Mathematically, let $j(i)$ denote the occupation of individual i and $j'(i)$ the occupation of i's father. The dummy variable is then written as

$$isDynasty(i) = \mathbb{1}(j(i) = j'(i)).$$

With this dummy variable, we define two objects   the *dynastic share* and the *self-made share* for each occupation. Formally,

$$\textit{Dynastic share}_k = \frac{\sum_{i:j(i)=k} isDynasty(i)}{\sum_i \mathbb{1}(j'(i) = k)}$$

and

$$\textit{Self-made share}_k = \frac{\sum_{i:j(i)=k}(1 - isDynasty(i))}{\sum_i \mathbb{1}(j'(i) \neq k)}.$$

Intuitively, the *dynastic share* of occupation k describes what fraction of workers with fathers in this occupation still choose to work in k. The *self-made share* of occupation k describes what fraction of workers without fathers in k end up in this occupation.

Finally, we compute the *dynastic bias* for each occupation. Define

$$\textit{Dynastic Bias}_k = \frac{\textit{Dynastic share}_k}{\textit{Self-made share}_k}$$

as the dynastic bias for occupation k. This quantity is an odds ratio that captures how many times more likely one chooses the occupation of one's father than another individual whose father is not in the occupation. If the children's occupational choices are independent of their fathers' occupations, then the bias should take the value of 1. If occupational choices are biased towards children with fathers already working in the occupations, then this bias would be greater than one.

2.2.2 Dynastic Bias at the Occupation Level

We end up with 27,583 non-zero country-year-occupation level observations as plotted in the histogram in Figure 2·1. This means that each census contains on average 100 distinct occupations. The median observation has its level of dynastic bias at 66. The horizontal axis is organized in log scale as the distribution of the dynastic bias has an extremely long right tail with a skewness of 59. Even as we use the log scale, the distribution still skews to the right with a skewness of 0.36. Notably, the skewness of the distribution does not arise from a small set of countries or censuses. The observations above the 90th percentile belong separately to 178 cross-sections from 68 countries, which account for approximately 70% of the censuses/countries in our sample.

The most striking pattern reflected by the histogram is that almost all occupations have their dynastic biases much greather than one. Occupations such as soccer players in Brazil can have dynastic biases as high as above 8,000. A considerable amount of occupations, mainly those in developing countries, have dynastic biases well above 100. Meanwhile, some occupations commonly perceived as highly persistent over generations turn out to have dynastic biases in the lower half of the distribution. For example, the dynastic bias of lawyers in the U.S. is 66 in 1960 and declines to 20 in 2010. Although this is still a remarkably high number, there remain many more other occupations more dynastic than the legal profession in the U.S. On the other hand,

Figure 2·1: Dynastic Bias of Occupations

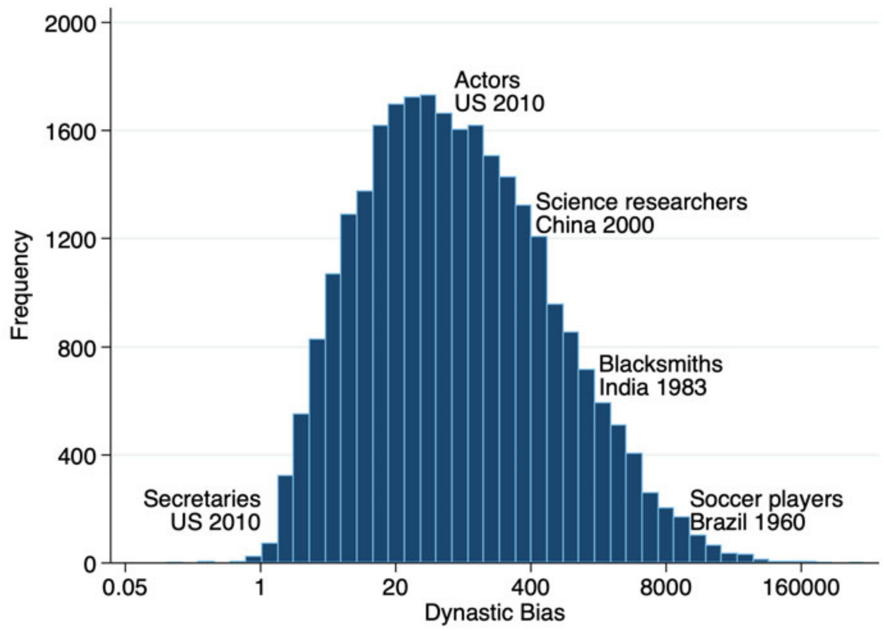

Note: This histogram plots the distribution of country-year-occupation level dynastic bias of 27,583 observations. The level of the dynastic bias is presented in log scale on the horizontal axis.

less biased occupations include secretaries, cashiers, and retail salespersons, whose dynastic biases are only slightly above one.

Table 2.1 reports the regression coefficients of some covariates of the occupation level dynastic bias. We include three independent variables per capita income of the country,[3] average years of schooling of the occupation, and a composite variable reflecting the eliteness of the occupation. The specification for the benchmark regression in column (1) is

$$\ln(Dynastic\ Bias_{ctk}) = \beta_1 Eliteness_{ctk} + \beta_2 YrSchool_{ctk} + \beta_3 \ln(Income_{ct}) + \alpha_c + \varepsilon_{ctk}$$

where the subscripts c, t, and k denote country, year, and occupation respectively. We include country fixed effects α_c and cluster the errors at the country level.

[3]Availability of income data at the occupation level is very limited.

The eliteness measure for each occupation is constructed with the information of workers' living conditions. We extract from each census variables of home ownership, access to electricity, access to piped water, sewage connection, use of toilets, and availability of bath facilities. These variables are recorded at the worker level. We then conduct a principal component analysis with these variables and retain the variations along the first principal component. Within each census, we generate the percentile of the first principal component of each worker and compute the mean of the percentiles by occupation. We name this average percentile at the occupation level the "eliteness" of the occupation, with the idea that more elite occupations allow workers to enjoy better living conditions. In other words, this eliteness measure is an approximation to the socioeconomic status of the workers.

Table 2.1: Occupation Level Covariates of Dynastic Bias

	(1)	(2)	(3)
Eliteness of occupation	0.011		0.0075
	(0.0031)		(0.0025)
Average years of schooling	-0.051	-0.026	
	(0.013)	(0.015)	
Log GDP per capita	-0.41	-0.57	-0.48
	(0.14)	(0.13)	(0.081)
Country FE	Yes	Yes	Yes
R^2	0.22	0.22	0.26
Observations	15,546	16,703	21,747

The dependent variable is the log of dynastic bias.
Observations are at the country-year-occupation level.
Standard errors are in parentheses and clustered by country.

Column (1) of Table 2.1 shows that eliteness and average education level have opposite correlations with the dynastic bias. Specifically, an increase in the eliteness of an occupation by 1 percentile is associated with 0.011 higher log dynastic bias, representing a 1.1-percent higher share of workers with parents in the same occupation than those without. On the other hand, each additional year of schooling of the workers in an occupation is associated with a decline the likelihood of occupational dynasties by 5.1 percent. With country fixed effects, these estimates describe how the dynastic bias correlates with the eliteness and education level of the occupations *within* a country.

The magnitude of the income coefficient is an order larger than that of the occupation-level variables. Since income enters the regression at the country-year level, the estimate should be interpreted in the way that as a country becomes richer, the unconditional average of occupational dynastic bias declines.

Moreover, we observe that the effects of eliteness and education are underestimated by a quarter or half in magnitude when only one of the two variables enters the regression, as shown in Column (2) and (3) of Table 2.1. This means that we need to look beyond explanations based on a single variable and think deeper about how these variables interact to affect the dynastic biases we observe.

2.2.3 Dynastic Bias at the Country-Year Level

We now aggregate the country-year-occupation level dynastic biases. We do so by taking the average of the log biases of an occupation in a census, weighted by the occupational share among the young generation. Mathematically,

$$\ln(\textit{Dynastic Bias}_{ct}) = \sum_k s_{ctk} \ln(\textit{Dynastic Bias}_{ctk}) \tag{2.1}$$

for country c in year t, where s_{ctk} is the share of occupation k in the young generation. This is equivalent to aggregating the level of the dynastic bias into a weighted geometric mean, i.e.

$$\textit{Dynastic Bias}_{ct} = \prod_k \textit{Dynastic Bias}_{ctk}^{s_{ctk}}.$$

This weighting scheme reduces the importance of small occupations, while taking the geometric mean downplays the extremely large biases that may result from imprecise measurement. We plot the histogram of the country-year level biases in Figure 2·2. The median bias levels at 6.8 and the distribution of the log biases becomes much less skewed. As we further aggregate the bias to the world level, we retain the most recent census of each country, weight the log biases with the population of the country from Penn World Table (Feenstra et al., 2015), and arrive at an average level of 10.6. This means that children with their fathers in an occupation are, on average, more than 10 times more likely to enter the occupation compared to those whose fathers are not in that occupation. In leymen's words, it is 10 times easier for one to enter the occupation of one's father.

Figure 2·2: Dynastic Bias at Aggregate Level

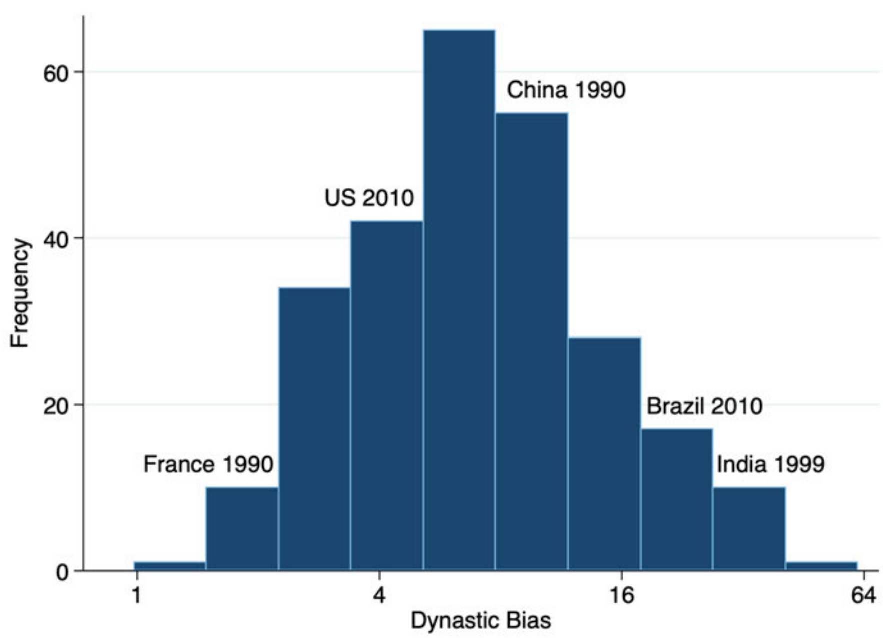

Note: This histogram plots the distribution of country-year level dynastic bias of 275 censuses/surveys from 90 countries. The level of the dynastic bias is presented in log scale on the horizontal axis.

We then take a closer look at the correlation between the dynastic bias and the income level of the countries. Figure 2·3 is a scatter plot of the dynastic bias of each census versus GDP per capita. Each circle in the figure denotes a country-year observation, and the size of the circle represents the population.

Figure 2·3: Dynastic Bias versus Income by Country-Year

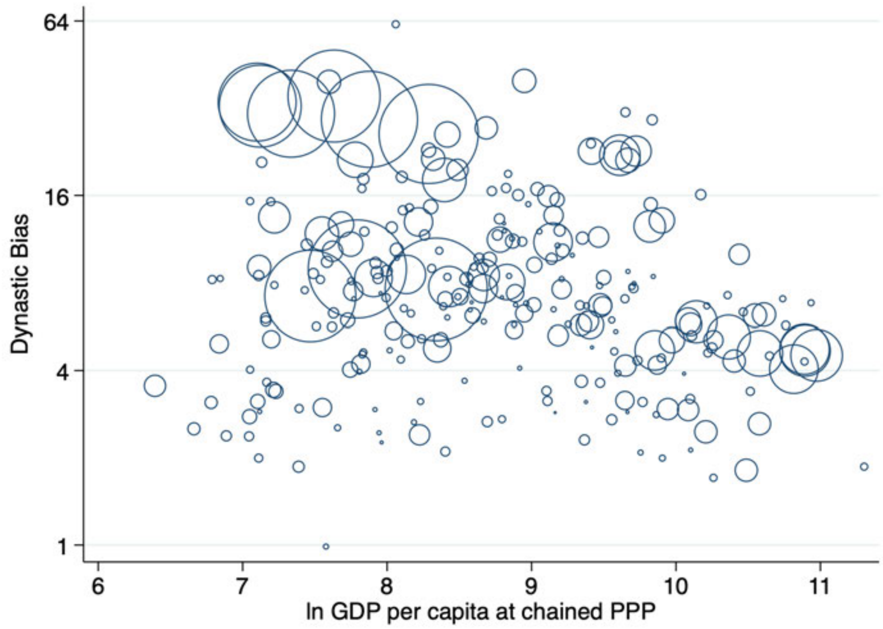

Note: This figure plots the dynastic bias at the country-year level against the log of GDP per capita at constant international prices. The level of the dynastic bias is presented in log scale on the vertical axis. The size of the observations represents populaton of the country at the year of the census/survey.

We can easily recognize countries with the largest populations in Figure 2·3. The several large circles at the top of the graph come from the National Sample Surveys of India, whose dynastic biases are a little below 30. The three large circles below are three censuses from China with the dynastic biases around 8. The several medium-sized circles at the far right of the graph represent the U.S., whose dynastic biases level slightly above 4. Other than one census from Togo that records no significant bias, all other censuses show considerable degree of dynastic bias in the occupational choices of their workers.[4]

[4]The low dynastic bias in Togo comes almost exclusively from a structural difference in classifying agricultural workers across generations.

We can observe from Figure 2·3 that the dynastic bias at the census level is lower for rich countries. Meanwhile, some poor countries also have low levels of dynastic bias. For clearer illustration, we collect the census-level dynastic biases and make a binscatter in Figure 2·4. Interestingly, the income level exhibits a non-monotone association with the dynastic bias as shown by the quadratic fit curve. The relationship is first increasing for the poorer countries and then decreasing for the richer countries.

Figure 2·4: Dynastic Bias by National Income

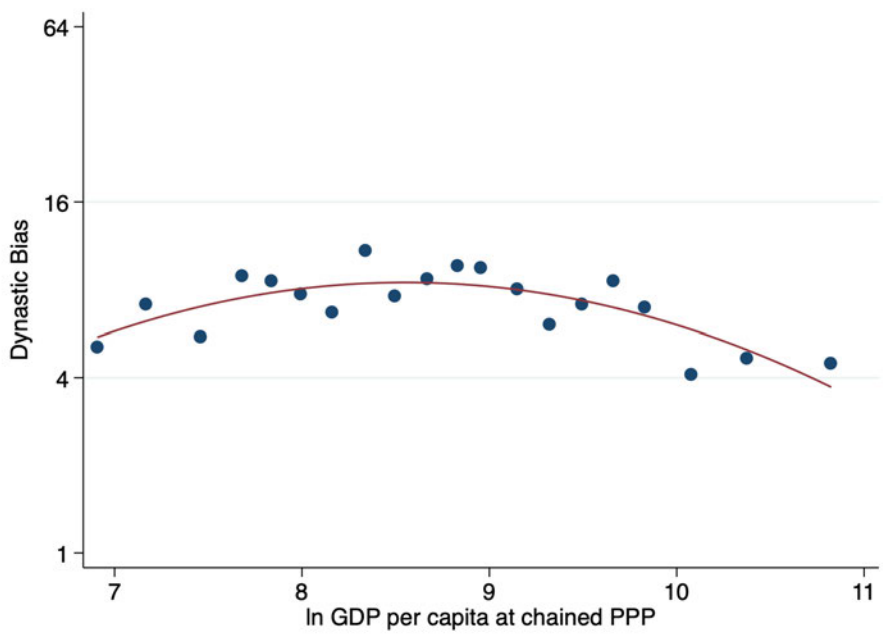

Note: This figure collects the country-year level dynastic bias in 20 bins by log GDP per capita. The fit line is quardratic.

To further examine this non-monotone relationship, we conduct a simple decomposition exercise that isolates the unconditional average of the dynastic biases from any composition effect. Recall that we use Equation (2.1) to construct the country-level dynastic bias, which can be rewritten as

$$\ln(\mathit{Dynastic\ Bias}_{ct}) = \overline{\ln DB}_{ct} + \sum_k s_{ctk} \left(\ln(\mathit{Dynastic\ Bias}_{ctk}) - \overline{\ln DB}_{ct}\right)$$

where $\overline{\ln DB}_{ct}$ denotes the unweighted mean of the log dynastic biases across occupations in country c in year t. This unweighted mean can be called a within-occupation component or a shift component of the dynastic bias, which represents the level of dynastic bias of a random occupation chosen from the country. The second term on the right-hand side can be interpreted as the between-occupation component or the share component, indicating the effect of occupational composition on the country-level dynastic bias.

We plot the two components of the dynastic bias of each census in Figure 2·5. For any observation, multiplying the horizonal and vertical coordinates gives the level of the dynastic bias, so that moving towards the northeast of the graph means an increase in the dynastic bias, represented by the darker color. The vertical axis measures the unconditional mean of the dynastic biases, while the horizontal axis the composition effect. The size of each marker is determined by the GDP per capita of the country at the year of the census.[5]

We observe from Figure 2·5 that richer countries, represented by larger markers, cluster at the bottom-right corner of the graph and have on average lower dynastic biases indicated by the lighter fill of the markers. Poorer countries sit mostly at the upper-left corner. The shift component dominates as it is greater than one for virtually every observation, especially the poor countries. Moreover, the share component is less than one for almost all observations and smaller for poor countries, meaning that large occupations are typically less dynastic and that poor countries would have been even more dynastic if their occupational structure in terms of the dynastic bias is closer to that of the rich countries. In other words, rich countries are less dynastic not because they have more workers in less dynastic occuaptions, but because their occupations are in general less dynastic than those in poor countries.

[5]We calculate GDP per capita with the output-side real GDP at chained PPPs and population retrieved from the Penn World Table, so that the income levels can be compared across countries and over time.

Figure 2·5: Decomposition of Aggregate Dynastic Biases

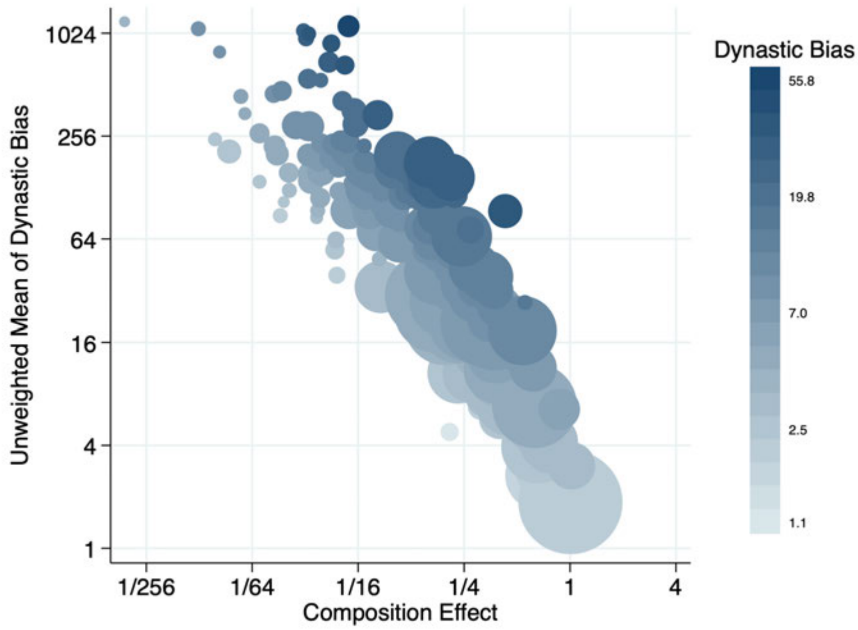

Note: This figure plots the decomposition of dynastic bias at the country-year level. Darker colors represent country-years with higher dynastic biases. The size of the observations represents GDP per capita at constant international prices.

The patterns in Figure 2·5 also shows the driving force for the non-monotone relationship in Figure 2·4, where the observations with the highest dynastic biases are not the poorest countries. We can see from Figure 2·5 that, as we move from poorer to richer countries (i.e. from smaller to larger dots), the dynastic bias first worsens due to a higher share of workers flowing into highly dynastic occupations, i.e. through the composition effect, and then mitigates as the dynastic bias goes down for all occupations.

Education attainment is another factor commonly thought of as an determinant of intergenerational persistence (e.g. Restuccia and Urrutia, 2004). In Figure 2·6 and 2·7, we plot the dynastic bias against the average years of schooling as we did for income, and observe that the relationship between dynastic bias and education is also non-monotone. In particular, the dynastic bias is increasing when the average years of schooling is below 8, and decreases for more years of schooling.

Figure 2·6: Dynastic Bias versus Years of Schooling by Country-Year

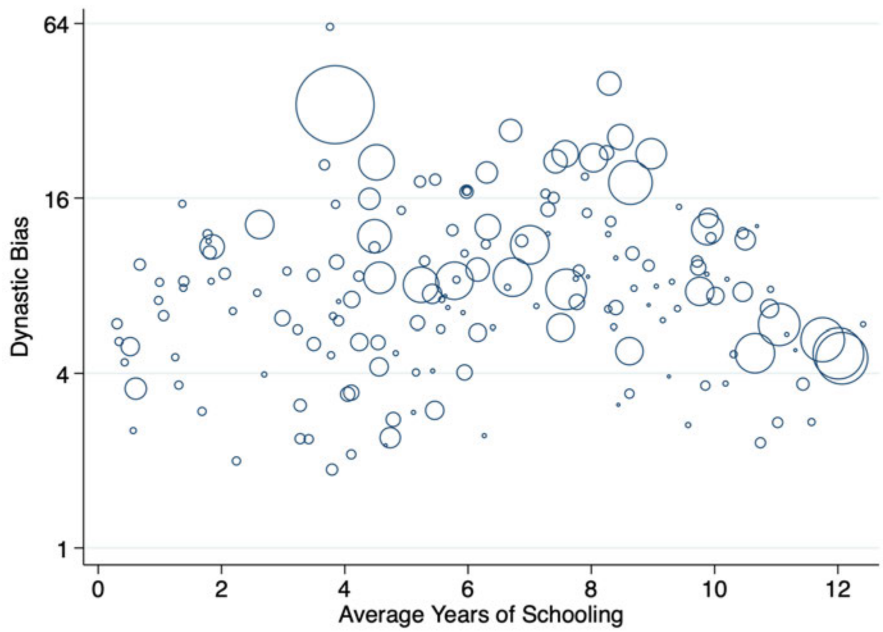

Note: This figure plots the dynastic bias at the country-year level against the average years of schooling. The level of the dynastic bias is presented in log scale on the vertical axis. The size of the observations represents populaton of the country at the year of the census/survey.

Table 2.2 confirms with regressions that these non-monotone relatoinships are statistically significant. Column (1) and (2) of the table correspond to the quardratic fit curve in Figure 2·4 and Figure 2·7 respectively. When estimated together as reported in Column (3), the coefficients for the linear and quardratic terms stay statistically significant, and are not significantly different from the coefficients when income or education is estimated alone in Column (1) or (2).

2.2.4 Dynastic Bias in Historical U.S. Censuses

As we discussed eariler, one limitation of our dynastic measure is that it captures only working children who live with their working fathers. In this section, we address

Figure 2·7: Dynastic Bias by Average Years of Schooling

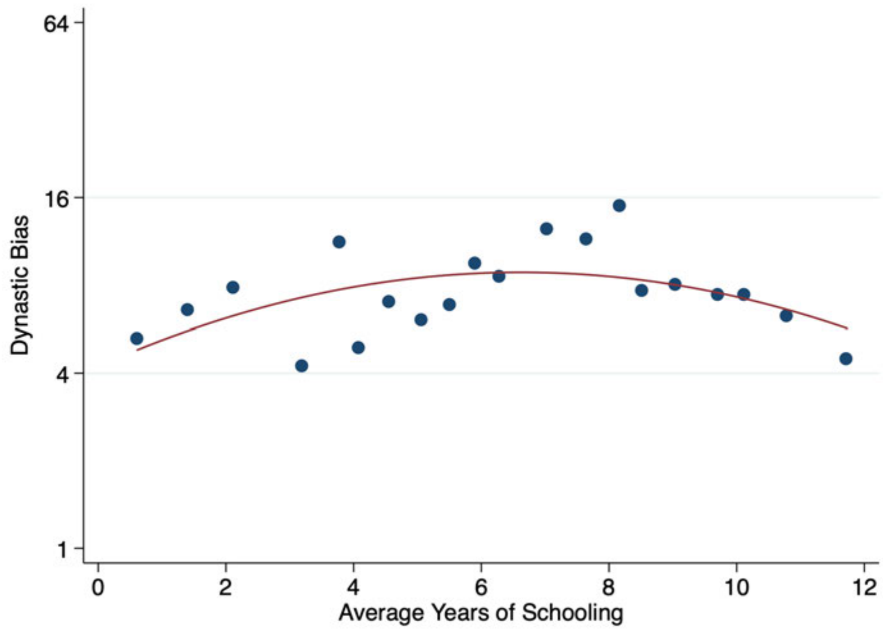

Note: This figure collects the country-year level dynastic bias in 20 bins by the average years of schooling. The fit line is quardratic.

concerns of potential biases in our measure with the help of the full count data of historical U.S. censuses from 1850 to 1940.

The Census Linking Project (Abramitzky et al., 2020) provides reasonable matches of observations across the historical censuses that can be identified as from the same person. Therefore, we can link workers back in time when they were little and still resided with their fathers to retrieve the occupational information of the fathers. For example, we first identify all child-father pairs in 1850 and retain the occupational information of the fathers. We then use the crosswalk of unique individual IDs from the Census Linking Project to link the children to the 1880 census and obtain the occupations of the working children, who need not live with their fathers in 1880. We then use the children's occupations in 1880 and the fathers' occupations in 1850 to compute the dynastic bias, and compare this bias level to the number computed from the cross-section of the 1880 data with only the working children residing with their

Table 2.2: Country Level Covariates of Dynastic Bias

	(1)	(2)	(3)
Log GDP per capita	2.724		1.945
	(0.513)		(0.769)
Log GDP per capita, squared	-0.159		-0.104
	(0.028)		(0.044)
Years of schooling		0.207	0.157
		(0.046)	(0.055)
Years of schooling, squared		-0.015	-0.014
		(0.004)	(0.005)
R^2	0.09	0.07	0.12
Observations	263	177	170

The dependent variable is the log of dynastic bias.
Observations are at the country-year level.
Robust standard errors are in parentheses.

fathers.

We compare the dynastic biases computed with these two different methods in Figure 2·8 for working children in 1880, 1900, 1910, 1930, and 1940.[6] The horizontal axis marks the log level of the dynastic bias computed from the linked data, while the vertical axis measures that from the cross-sections. The observations scatter around the 45-degree line, which signals reasonable fit between the two methods. For robustness, we repeat this exercise for working children who lived with their fathers 20 years ago. As Figure 2·9 shows, the goodness of fit stays roughly the same.

Next, we aggregate the dynastic biases under different methods to the national level in Figure 2·10 to examine how much the dynastic biases computed from the cross-section deviates from the numbers generated with linked data. The biases computed from the cross-sections are plotted with the solid green curve. Those from the 30-year and 20-year linked data are plotted with the black squares and the blue triangles respectively. We observe that, other than in the late ninteenth century, the levels of the dynastic biases from different methods are very close to each other. The overall

[6]Some years are missing because the full count data of the 1890 census are not available.

Figure 2·8: Dynastic Bias in Different Samples, 30-Year Link

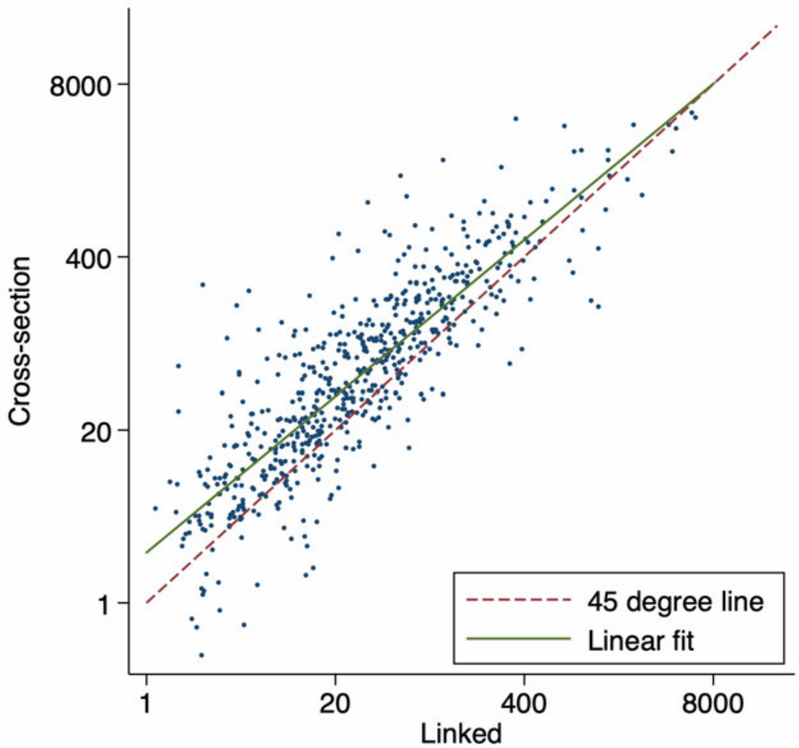

Note: This figure plots the occupation-level dynastic bias derived from different methods. The dynastic bias shown on the vertical axis is derived from the current occupations of coresiding father-child pairs in the cross sections. The dynastic bias shown on the horizontal axis is derived from the current occupation of the child and the occupation of the father who lived with the child 30 years ago. Observations are pooled for the years 1880, 1900, 1910, 1930, and 1940.

trend of the cross-section bias from 1850 to 2015 is slightly decreasing but largely flat, compared to the cross-country differences in the dynastic bias.

As the cross-sectional dynastic biases resembles the estimates from the linked data in the historical U.S. censuses, it is likely that the cross-sectional figures for other countries also reflect the true degree of their intergenerational persistence of occupations reasonably well. We therefore build further analyses on this measure of dynastic bias, taking it as comparable across countries and years.

Our last empirical investigation with the historical U.S. full count census is to

Figure 2·9: Dynastic Bias in Different Samples, 20-Year Link

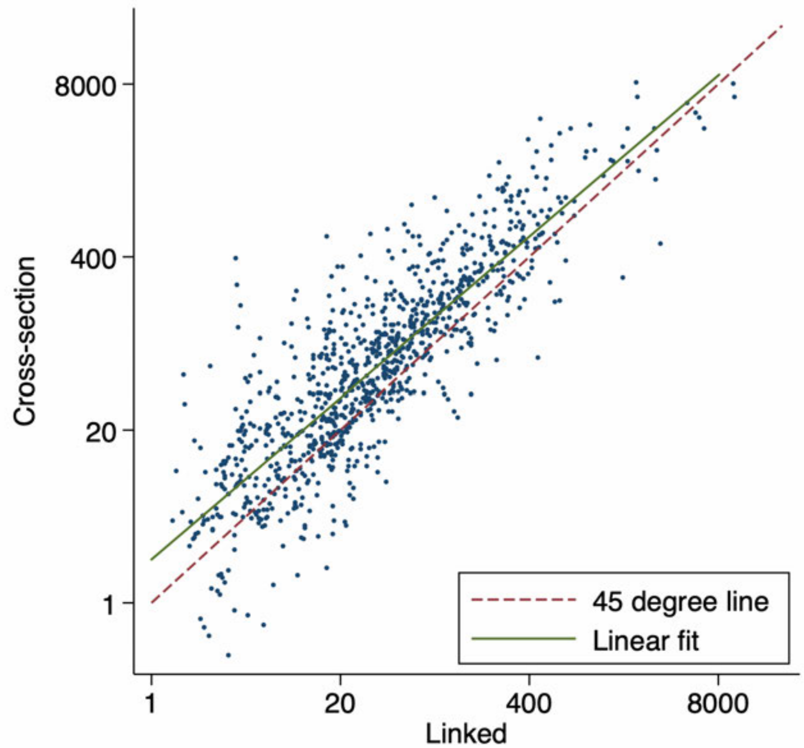

Note: This figure plots the occupation-level dynastic bias derived from different methods. The dynastic bias shown on the vertical axis is derived from the current occupations of coresiding father-child pairs in the cross sections. The dynastic bias shown on the horizontal axis is derived from the current occupation of the child and the occupation of the father who lived with the child 20 years ago. Observations are pooled for the years 1870, 1880, 1900, 1920, 1930, and 1940.

examine the contribution of agricultural occupations to aggregate dynastic bias, taking advantage of the consistent occupational classification in the data throughout the long period of time from 1850 to 1940. The classification groups 269 detailed occupations into 10 groups, two of which are farmers and farm laborers. The former includes farm onwers, tenants, and farm managers, while the latter includes mainly laborers who works in the field.

We conduct two exerecises here, the first to compare the dynastic bias of agricultural occupations to that of the other occupations, and the second to compare

Figure 2·10: Dynastic Bias of the U.S. since 1850

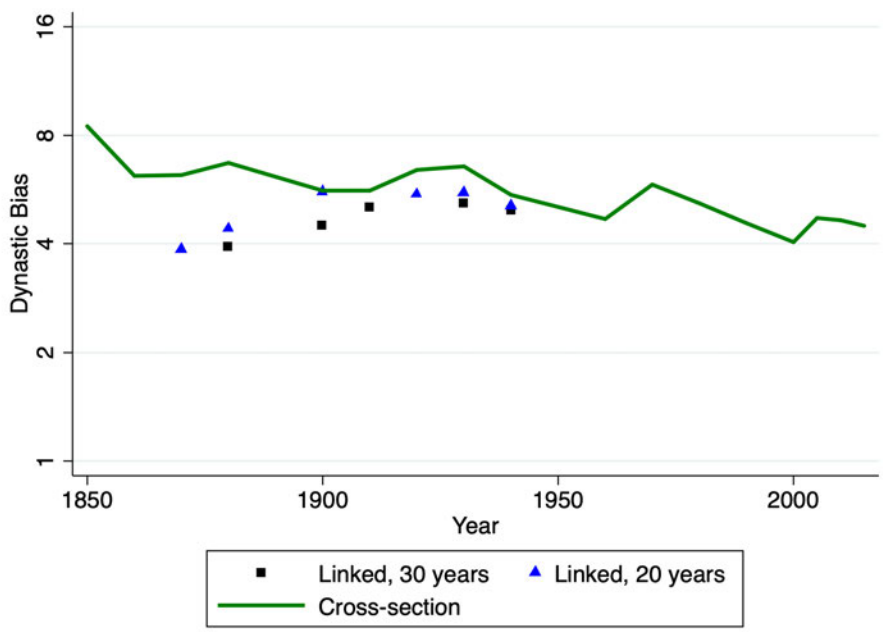

Note: This figure plots the time series of country-level dynastic biases of the U.S from 1850 to 2015. The dynastic bias is presented on the vertical axis in log scale.

the aggregate dynastic bias with and without agricultural occupations. For the first exercise, we run the regression

$$\ln(Dynastic\ Bias_{tk}) = \alpha_{b(k)} + \alpha_t + \varepsilon_{tk}$$

where $\alpha_{b(k)}$ is the fixed effect for the broad category b of occupation k and α_t is the year fixed effect. We weight the regression by the number of workers in each detailed occupation k to downplay the importance of extremely biased but small occupations within each broad occupational category. Surprisingly, the dynastic bias of farm laborers is the lowest among all broad occupational categories. Figure 2·11 shows the dynastic bias of each broad occupational categories estimated from the regression, normalizing that of farm laborers to 1. The standard error bars around the estimates arise from the variations in dynastic biases of the detailed occupations under the same broad category. It is evident that farm laborers only lower the aggregate dynastic bias, while farmers also belong to the half of the broad categories with lower dynastic biases. On the other hand, occupations of professional activities and craftsmenship contributes most to the aggregate dynastic bias.

Now we turn to compare the aggregate dynastic bias with and without the agricultural occupations. The solid green curve in Figure 2·12 denotes the aggregate dynastic biases of the U.S. with farmers and farm laborers. The dashed red curve denotes the dynastic biases when we aggregate without those occupations. It is evident that excluding agricultural occupations only increases the dynastic bias at the aggregate level, while the gap remains unclosed in year 1940. Interestingly, the decline in the dynastic bias without agricultural occupations coincides with the decline in the share of workers in agricultural occupations, implying that the structural transformation out of agriculture may be an important driver that lowers the dynastic bias in the non-agricultural occupations.

Figure 2·11: Dynastic Bias Relative to Farm Laborers

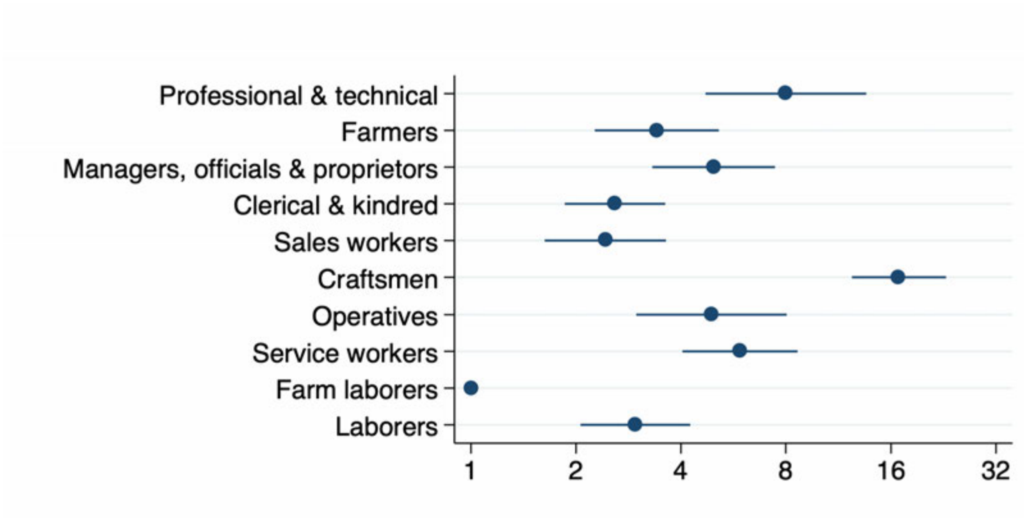

Note: This figure plots the average occupational dynastic bias of each broad occupation category weighted by the number of young workers.The coefficient for farm laborers is normalized to 1. The dynastic bias of other broad occupational groups relative to farm laborers are presented along the horizontal axis in log scale. The bars around the estimates represent robust standard errors. Year fixed effects are applied.

We have shown in this section that the dynastic bias we compute from the censuses is a satisfactory measure of intergenerational persistence of occupations. The dynastic bias correlates with several socioeconomic indicators at the occupation level and the country level, which motivates further investigation of how the bias is associated with less efficient allocation of workers. While clear identification of how the dynastic bias affects allocative efficiency is not feasible for a cross-country setup due to limitation of data and is therefore deferred to further studies, we provide in what follows a preliminary analysis of the aggregate effect of dynastic bias. Specifically, we employ a quantitative model of selection à la Roy (1951) to compute the maximum potential gains from mitigating the dynastic bias.

Figure 2·12: Dynastic Bias with and without Agricultural Occupations

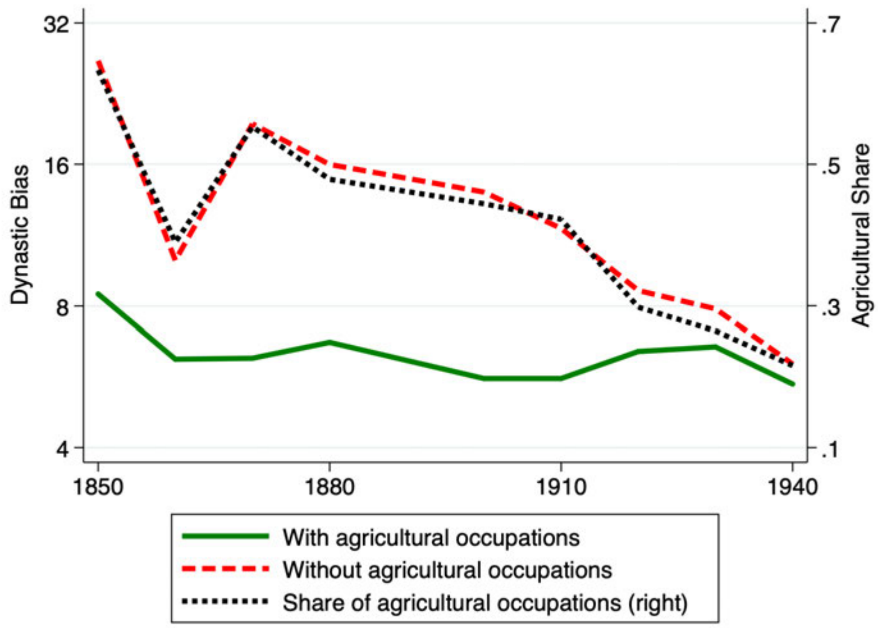

Note: This figure plots the country-level dynastic bias of the U.S. with and without agricultural occupations. The dynastic bias is presented along the left vertical axis. The share of agricultural occupations among the young workers is presented along the right vertical axis.

2.3 Model of Occupational Choice

2.3.1 Production of Final Goods

The economy consists of a continuum of workers and J distinct occupations. To fix notation, we use k for an indeterminant occupation, j for the occupation of workers entering the labor market, and j' for the occupation of the parents. A representative firm produces final goods with a Cobb-Douglas production function in which the human capital of each occupation serves as the inputs. Mathematically,

$$Y = \prod_k L_k^{\alpha_k} \tag{2.2}$$

where L_k denotes the amount of human capital employed in occupation k and α_k, normalized so that $\sum_k \alpha_k = 1$, determines the compensation share of each occupation. The representative firm maximizes profits subject to the prices for human capital, with the problem

$$\max_{\{L_k\}} Y - w_k L_k \tag{2.3}$$

where the price of the final good is taken as the numeraire.

2.3.2 Worker's Occupational Choice

In each period, workers earn their labor income and exit the labor market. Each worker has one child who replaces the worker in the next period. The young generation, upon entering the labor market, draws occupation-specific skills from a joint independent Fréchet distribution

$$\boldsymbol{z} = \{z_1, ..., z_J\} \sim F(\theta)$$

where this distribution has a common shape parameter θ. Specifically, each marginal distribution of the skills follows

$$\Pr(z_j \leq Z) = e^{Z^{-\theta}}$$

As workers enter the labor market, they choose an occupation from which their labor income is highest, subject to labor market frictions $\tau_{jj'}$, where $\tau_{jj} = 1$ and $\tau_{jj'} = \tau_j < 1$ if $j \neq j'$. Specifically, the labor market frictions enter the worker's problem in the form

$$\max_{j} \tau_{jj'} w_j z_j \tag{2.4}$$

where w_j is the wage rate for each unit of skill in j, so that the probability of worker i choosing occupation j with parent working in j' is

$$\pi_{jj'} = \frac{(w_j \tau_{jj'})^{\theta}}{\sum_k (w_k \tau_{kj'})^{\theta}} \tag{2.5}$$

Define the dynastic bias of an occupation as the odds ratio of the probability of workers choosing the occupation with parents also in this occupation to the probability of workers choosing this occupation without parents in the occupation. Explicitly,

$$Bias_j \equiv \frac{\pi_{jj}}{\sum_{j' \neq j} s_{j'} \pi_{jj'} / \sum_{j' \neq j} s_{j'}} = \frac{(1 - s_j)\pi_{jj}}{\sum_{j' \neq j} s_{j'} \pi_{jj'}} \tag{2.6}$$

denotes the dynastic bias of occupation j, where $s_{j't}$ is the share of workers in time t with parents' occupation in j'. The dynastic bias at the economy level is aggregated in the same way as we did in the empirical section. Mathematically,

$$\ln(Bias) = \sum_j s_{j'} \ln(Bias_j). \tag{2.7}$$

2.4 Potential Gains from Lifting Labor Market Frictions

In this section, we compute the potential gains from eliminating the dynastic bias for all economies in our sample. Since we assumed that all dynastic biases are due to labor market frictions, eliminating the biases means removing all the labor market friction. Admittedly, there may exist factors other than labor market frictions that lead to the dynastic bias that do not necessarily distort the allocation of workers. We refrain in this paper from addressing those channels as they are hard to identify with our current data. Therefore, this exercise can be seen as to arrive at an upper bound of the efficiency gains. The result of this exercise will inform us which economies might have benefited/suffered most from the decreases/increases in their dynastic bias, and which economies may have the greatest room for improvement in their allocation of workers along the direction of reducing occupational dynastic biases.

Based on Equation (2.5) and (2.6), our quantitative exercise in this section first calibrates the wage rates w_j and the labor market frictions τ_j to match the occupational share of employment and the dynastic bias of each occupation for each census in our sample, given the share of parents' occupations $s_{j'}$. We then compute the counterfactual scenario where all labor market frictions τ_j are set to 1. The improvement in labor productivity is calculated as the average improvement in each occupation weighted by the expenditure share of each occupation.

Figure 2·13 plots the distribution of the potential labor productivity gains for all 275 censuses in our sample. The histogram shows that the distribution of potential gains is heavily skewed to the right, despite that many economies have potential gains below 5%. While the median gains is 6.8%, 166 of the 269 censuses have potential gains above 5% and 84 of them hit above 10%.

To prevent inaccurate inference from coarse classifications of occupations, we select a subset of the censuses that have at least 40 distinct occupations that employ both the father and the child generation. This sub-sample contains 206 censuses from 72 different countries. Figure 2·14 shows that the distribution of potential gains now shifts slightly to the right. The highest potential gains are 33%, calculated from the 1988 census of Honduras, and the top quartile of the gains are 12%. The median gains are 8.0%, with 149 cross-sections having the potential gains above 5% and 80 above 10%. In a world where the global real growth rates average 3.8% since the 1960s (World Bank 2022), these potential gains point towards a source of considerable growth, given that the dynastic biases in these economies are indeed derived from labor market frictions.

To further investigate the correlation between the potential productivity gains and the observed dynastic biases, we plot the two variables together in Figure 2·15.

Figure 2·13: Number of Economies by Interval of Potential Gains

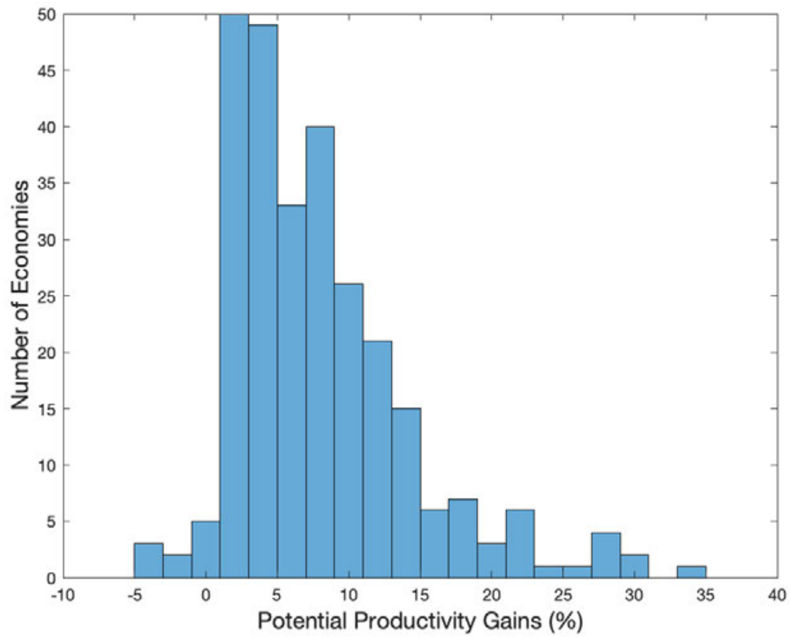

Note: This figure plots the distribution of potential gains in aggregate labor productivity for all censuses in our sample. The gains are calculated with the assumption that all dynastic biases are due to labor market frictions and the gains for each census are the weighted average of labor productivity gains for each occupation from lifting all labor market frictions.

The scatter plot shows a clear positive correlation between the potential gains and the observed biases in log scale. The correlation is 0.74 and the slope of the fitted line is 7.05, meaning that halving the dynastic bias of a country may lead to labor productivity gains as high as 7.05%.

2.4.1 Potential Gains for Developing Countries

In this section, we further reduce the scope of the sub-sample to focus on developing countries that may benefit most from eliminating the dynastic biases. We exclude censuses from countries that the World Bank lists as high-income as of 2020 as well as countries whose last census was conducted before 2000. We then focus on the latest

Figure 2·14: Number of Economies by Interval of Potential Gains with More Than 40 Occupations

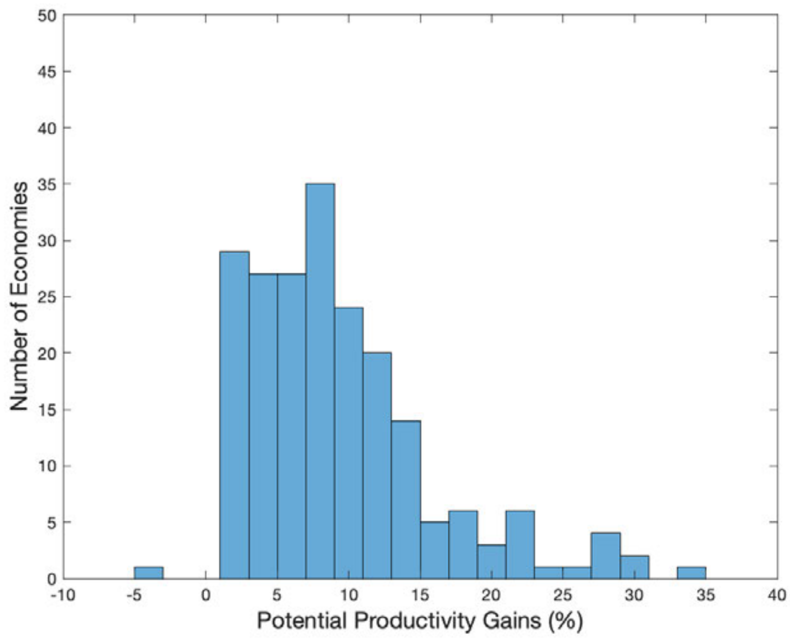

Note: This figure plots the distribution of potential gains in aggregate labor productivity for censuses with more than 40 distinct occupations.

Figure 2·15: Potential Gains versus Dynastic Bias

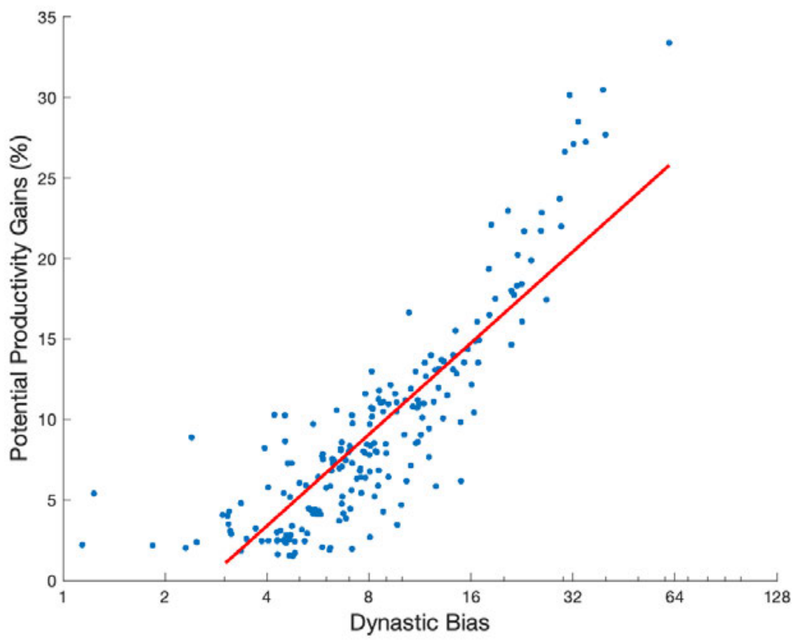

Note: This figure plots the potential gains in aggregate productivity against the aggregate level dynastic biases.

censuses of the remaining 52 developing countries.

We report the results in column (1) of Table 2.3. The median potential gains for the 52 developing countries is 8.1%, where the 2000 census of Thailand leads with the highest number at 27.7%, followed by India, Zambia, and Brazil, whose potential gains are all above 20%. Other populous developing countries with potential gains higher than 10% include Egypt (10.9%), Indonesia (16.5%), Iran (13.7%), and Mexico (11.1%). Twenty-one of these developing countries have potential gains higher than 10%, and 42 of them higher than 5%, including China (7.0%). These statistics show that correcting for the dynastic biases in these developing countries may lead to non-trivial gains in their aggregate labor productivity.

Table 2.3: Summary of Potential Gains from Quantitative Exercises

	(1) Removing all barriers	(2) Matching US bias in 2015	(3) Shift component in (2)
Highest gains	28% Thailand 2000	18% Thailand 2000	16.8% India 2009
Gains at the top quartile	12%	5.7%	6.2%
Median gains	8.1%	3.5%	2.8%
Number of countries with potential gains above 5%	42	15	16
Number of countries with potential gains above 10%	21	9	6
Number of countries in the sample	52	52	52

Nonetheless, it may be unrealistic to implement any reform that completely and permanently eliminates the dynastic biases. As we refer back to Figure 2·3, the least dynastic countries still have their country level bias around 2. Therefore, we take some steps back and explore the potential gains in alternative counterfactuals.

The second counterfactual we explore follows the idea in Hsieh and Klenow (2009). We recognize that there may be some natural level of the dynastic bias, and take the U.S. as a benchmark to explore the gains for developing countries if their intergenerational mobility in occupations becomes more like that of the U.S. Specifically, we multiply all the labor market frictions τ_j in a country by a common factor so that the aggregate dynastic bias of the country matches that of the U.S. in 2015, with the log of the bias at 1.5 and the level of th bias at 4.5.

Column (2) of Table 2.3 reports the results of this exercise. Since we did not change the distribution of labor market frictions within a country, it is not surprising to see that Thailand remains at the top of the potential gains, altough with a smaller amount that is roughly two-thirds of that in the first counterfactual. The potential gains are accounting for a lower fraction of those in the first counterfactual as we move down the distribution. The gains at the top quartile in this counterfactual is less than half of the gains in the first counterfactual, and the ratio becomes even smaller at the median. Yet, these potential gains are still sizable as more than half of the developing countries may still achieve more than 3.5% of labor productivity gains than their current levels.

Following the second counterfactual, we decompose the gains from matching the U.S. dynastic bias into the shift component and the share component as we did earlier in Section 2. For this decomposition, we match the unweighted average of occupation-level log dynastic bias of the developing countries to that of the U.S. in 2015, and compute the implied gains in aggregate productivity. As reported in Column (3), the statistics for the potential gains are not very different from those in Column (2), meaning that the productivity gains are mainly realized through the shift dimension. To illustrate this point more intuitively, we plot the shift component against the total potential gains in Figure 2·16. We can see that the observations scatter closely around the 45-degree line.

Figure 2·16: Decomposition of Potential Gains from Matching with the U.S.

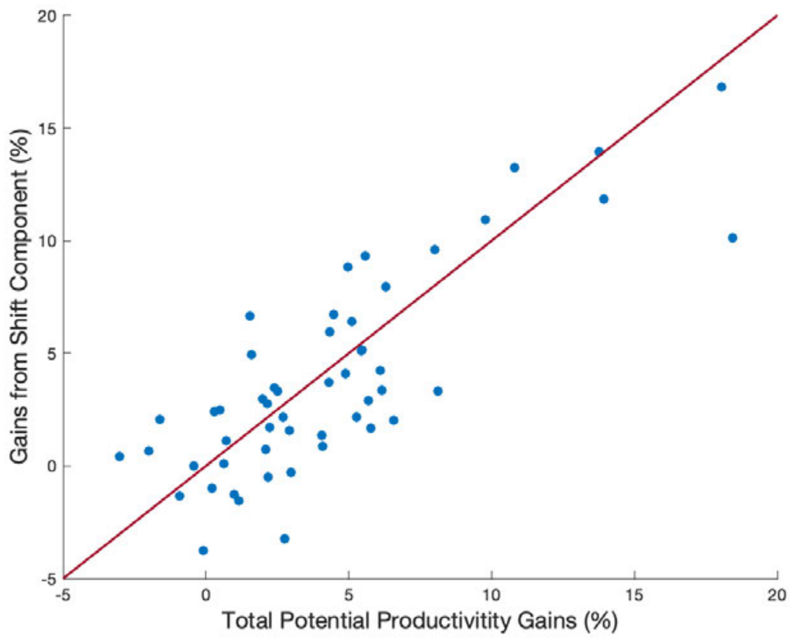

2.5 Conclusion

In this paper, we document empirically the degree of dynastic bias at both the occupation and the country level for 90 countries. We show that the dynastic bias of an occupation is increasing in the eliteness of the occupation but decreasing in the average years of schooling of the occupation within a country. We also show that the degree of dynastic bias across countries first increases and then decreases with both the income level and the average years of schooling.

Quantitatively, we compute the effect of occupational entry barriers that lead to the dyanstic biases on the aggregate level of productivity. We assume that dynastic biases arise due completely to occupational entry barriers, and calibrate the level of the barriers to the occupation-level dynastic bias of each country. Our counterfactual analysis shows that the productivity gains from erasing the dyanstic biases can be as high as 28%, with a median level of 8.1%, for the 52 developing countries in our sample. The gains from matching the dynastic bias with the U.S. level can be as high as 18% with a median of 3.5%.

As we stay agnostic about the source of the occupational entry barriers, our quantitative results may be interpreted as an upper bound for the efficiency loss that comes with the dyanstic biases. We aim to introduce specific mechanisms, such as structural transformation, formal education, intergenerational transmission of human capital, and financial frictions, for the next steps of this research and identify the relative importance of each channel by desciplining the model with the relevant moments from micro level data.

Appendix A

Appendix to Chapter 1

A.1 The Manufacturing Premium in Professional Services beyond Observables

Concerns may arise over the robustness of the new selection pattern, since the earnings premium for workers who reallocate from manufacturing into professional services is relatively small and not significantly positive in a number of years. Moreover, if this earnings premium is easily explained away by observable characteristics of the workers, it would be less obvious how the selection patterns are due to sector-specific skills instead of general skills that can be accounted for by education or experience, such as what the BLS labor composition index adjusts for (Zoghi, 2010). In this section, I further decompose this earnings premium by controlling for observable worker characteristics and show that the earnings gap between new and incumbent workers remain.

Table A.1 presents the composition of workers in professional services by gender and education. The upper panel represents the group of workers switching from manufacturing. We can observe that male workers make up the majority, while half of the workers possess a college degree. Compared to incumbent professional services workers shown in the lower panel, those reallocating form manufacturing involve a similar fraction of workers with college degrees, but substantially more male workers.

To evaluate whether the composition in gender and education fully explains the selection patterns, I run the regression in Section 2 by the subsamples and obtain

Table A.1: Composition of Professional Services Workers by Gender and Education

A. Workers reallocated from manufacturing		
	With college degree	Without college degree
Male	35.0%	29.2%
Female	15.0%	20.8%
Total	50.0%	50.0%
B. Incumbent workers in professional services		
	With college degree	Without college degree
Male	29.5%	17.3%
Female	19.7%	33.5%
Total	49.2%	50.8%

the manufacturing premium for each subgroup of workers, controlling for age and age squared of the workers. Table A.2 reports the estimated manufacturing premium by subsample, where the estimate of 0.045 in column (1) is the effect on all professional services workers. Comparing the estimates in (2) and (3) between workers with and without college degrees, we observe that the 6.8-log point premium for more educated workers is much larger than that for less educated workers, and even greater than the average number. Further examination within the more educated workers shows that the high level of premium exists most prevalently for more educated female workers. While the premium for male workers with college degrees is small in magnitude (1.1 log points) and statistically insignificant, a female professional services worker with college degree who reallocates from manufacturing earns on average 7.8 log points (8%) more than a typical incumbent female professional services worker possessing a college degree. Putting these numbers altogether, my findings echo with the results of Adão (2016) that more educated workers can have above-average earnings when they reallocate.

One could also wonder whether the manufacturing premium can be explained by potential structural differences in the occupational composition between reallocating and incumbent workers. To evaluate this hypothesis, I divide workers by the

Table A.2: MFG Premium in PROF by Subgroup of Workers

	All workers (1)	Without degrees (2)	With degrees (3)	Male with degrees (4)	Female with degrees (5)
	0.045	0.007	0.068	0.011	0.078
	(0.009)	(0.012)	(0.010)	(0.012)	(0.019)
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	187,509	95,954	91,555	53,747	37,808

Robust standard errors in parentheses.
Controlled for age and age squared.

six broad occupation categories implicitly implied in the 1990 Standard Occupation Classification scheme. As reported in Table A.3, managerial, professional, technical, sales, and administrative support occupations account for the overwhelming majority of both reallocating and incumbent workers. Notably, managerial and professional occupations are the only category that has weekly earnings above the sectoral average. Given that the fraction of managerial and professional workers is only 2 percent higher for reallocating workers from manufacturing than incumbent workers, the occupation structure cannot contribute much to the manufacturing premium. In fact, the average log earnings of incoming workers from manufacturing would be 4 log points higher if I adjust the occupation composition of incoming workers to be the same as that of incumbent workers, keeping the average earnings by occupation unchanged. On the contrary, the average earnings of workers reallocated from manufacturing are higher than the earnings of the incumbents *within* the two largest occupations. This observation suggests that the selection pattern persists even after controlling for the occupation structure of reallocating and incumbent workers.

Finally, I regress the log earnings in professional services on workers' previous sectors again, controlling for all factors considered in this section, plus the usual age polynomial. The estimated manufacturing premium turns out to be even larger now with the magnitude of 6.2 log points, compared to the 4.8-log point estimate in Section 2. The premium even becomes statistically significant for workers without

Table A.3: Composition of Professional Services Workers by Gender and Education

A. Occupational Structure (%)	Workers from manufacturing	Incumbent workers
Managerial & professional	48.9	46.9
Technical, sales & admin. support	31.1	43.5
Services	5.1	6.0
Farming	0.3	0.2
Precision production & repairers	5.5	2.0
Operative & laborers	9.0	1.4
B. Log Earnings	**Workers from manufacturing**	**Incumbent workers**
Managerial & professional	6.92	6.86
Technical, sales & admin. support	6.52	6.39
Services	5.77	5.85
Farming	5.56	5.92
Precision production & repairers	6.45	6.40
Operative & laborers	5.93	6.02
Average	6.62	6.57

degrees, while still significantly positive for degree holders, especially female workers. These results show that controlling for observables does not explain away the manufacturing premium in professional services, so that it is relevant to consider the selection mechanism along the sectoral dimension.

Table A.4: MFG Premium, Controlling for Occupation, by Subgroup of Workers

Sample	All workers (1)	Without degrees (2)	With degrees (3)	Male with degrees (4)	Female with degrees (5)
Manufacturing	0.062	0.029	0.063	0.017	0.061
Premium	(0.008)	(0.011)	(0.010)	(0.011)	(0.018)
Occupation FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	187,509	95,954	91,555	53,747	37,808

Robust standard errors in parentheses.
Controlled for age and age squared.

A.2 Correlated Comparative Advantages

In the main quantitative exercise I have assumed independence between the comparative advantages of workers' skills with respective to professional services, which is a handy simplifying assumption. It is possible, however, to fully specify the structure of the skill distribution between each pair of the sectors by consdiering the correlation of the comparative advantages between sectors. In this appendix, I introduce correlation by imposing a Frank copula on the percentiles of workers' comparative advantages. Recall that the percentile of worker i's comparative advantage in manufacturing with respect to professional services is denoted by $q_m(i)$, and that with respect to EHP is denoted by $q_e(i)$. Let these percentiles be correlated by the Frank copula

$$C(q_m, q_e) = -\frac{1}{\rho} \ln \left[1 + \frac{(e^{-\rho q_m} - 1)(e^{-\rho q_e} - 1)}{e^{-\rho} - 1} \right] \tag{A.1}$$

for $\rho \in (-\infty, \infty) \setminus \{0\}$ where ρ controls the degree of dependence. When $\rho = 0$, the comparative advantages in the two sectors are independent. A positive ρ means that if a worker ranks high in the distribution of comparative advantage in the manufacturing sector, then it is likely that the worker also ranks high in EHP.

For the quantitative analysis, the correlation parameter ρ is responsible for the fraction of workers switching between manufacturing and EHP. With a higher value of ρ, workers tend to have more similar skill levels for manufacturing and EHP so that more workers switch between the two sectors in response to idiosyncratic income shocks. When ρ is low, fewer workers reallocate between the two sectors because it becomes harder for idiosyncratic shocks to overcome the income differences between sectors for marginal workers.

The calibrated values of the parameters are reported in Table A.5 and the values of the targeted moments are listed in Table A.6. The value of α_m happens to be a precise 0, meaning that workers' skill in professional services is independent of their comparative advantage in manufacturing. The negative value of α_e means that workers are less productive in professional services if they have stronger comparative advantage in EHP. At first glance, it is not intuitive how these values generate a positive premium for workers moving form manufacturing into professional services. Note that, however, the parameter ρ is negative. This means that when workers move from manufacturing into professional services, the comparative advantage in manufacturing of the marginal worker goes up while that in education and health falls. Therefore, we need to look at the net effect on the absolute advantage. Figure A.1 is a bin-scatter plot of the simulated percentiles of workers' comparative advantages with bin size 50. The fitted line shows us that a 1%-increase in q_m leads to a decrease around 0.4% in q_e on average. Consequently, the decreases in q_e drive up the premium with the negative parameter α_e.

Table A.5: Values of Parameters

Parameter	Value	Interpretation
α_m	0	Correlation between comparative advantage in MFG and absolute advantage in PROF
α_e	-0.1841	Correlation between comparative advantage in EHP and absolute advantage in PROF
κ	0.1723	Dispersion of comparative advantages
ρ	-2.83	Correlation between comparative advantages in MFG and EHP
σ_a	0.4971	Dispersion of absolute advantage
σ_ε	0.0195	Dispersion of idiosyncratic income shocks

Table A.6: Targeted Moments

	Data	Model
Income premium of MFG workers switching into PROF	0.047	0.047
Income premium of EHP workers switching into PROF	-0.159	-0.159
Income premium of MFG workers switching into EHP	-0.117	-0.117
Fraction of workers switching between MFG and EHP	0.0078	0.0078
Frequency of sector switches	0.060	0.060
Variance of log income	0.460	0.460

Figure A.1: Correlation between Percentile of Comparative Advantages

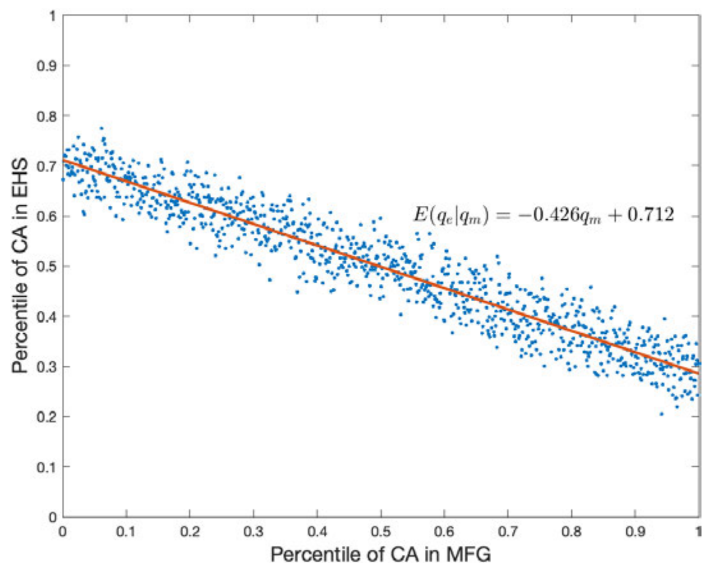

With this parameterization, I plot the implied changes in the average human capital in each sector. As shown in Figure A.2, the contribution of selection to the growth in average sectoral human capital stays similar compared to the case with $\rho = 0$, with the effect greater for professional services. The exact numbers of the selection effect is reported in Table A.7.

Figure A.2: Average Human Capital (Simulated) and Employment Share (Matched)

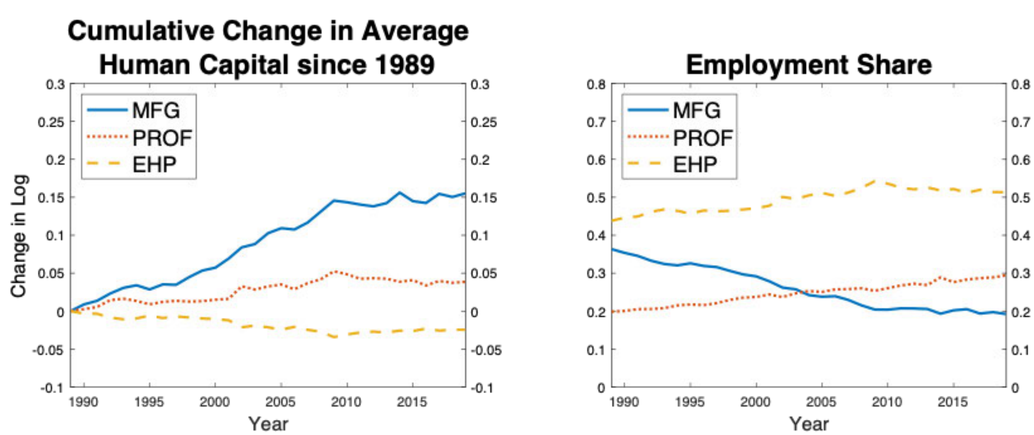

Table A.7: Effect of Selection on Sectoral Labor Productivity Growth

Cumulative growth in log points	Selection	Data	Impied True Growth
Manufacturing	15.5	104.4	88.9
Professional services	3.9	35.5	31.6
Education and health	-2.5	0.4	2.8
Aggregate services	1.3	21.1	19.8

A.3 Earnings of Workers Leaving Manufacturing

One robustness check to the model is to see how the earnings of workers leaving manfuacturing compare to the earnings of staying manufacturing workers. It turns out that, while a worker leaving manufacturing for EHP earns 11.8% less than the staying manufacturing workers, those leaving for professional services earn 14.5% more. This is inconsistent with the baseline model prediction that those who leave manufacturing for professional services should have lower skills in manufacturing. In this section, I introduce an extension of the model to capture this additional moment from data.

Recall that the comparative advantages of workers are determined by their percentiles in the distribution, but this time with sector-specific dispersion κ_s for $s \in \{m, e\}$:

$$s_m(i) = \kappa_m \ln \frac{q_m(i)}{1 - q_m(i)}, \tag{A.2}$$

$$s_e(i) = \kappa_e \ln \frac{q_e(i)}{1 - q_e(i)}. \tag{A.3}$$

The absolute advantage of worker i in professional services now has conditional mean in a slightly different form

$$a(q_m(i), q_e(i)) \sim Gumbel(\kappa_m \ln(1-q_m(i)) + \alpha_m q_m(i)(1-q_m(i)) + \alpha_e \ln q_e(i), \sigma_a) \tag{A.4}$$

where the relative size of α_m to κ_m determines how likely workers switching between manufacturing and professional services would have earnings higher than incumbent workers.

Figure A.3 plots the level of the expected income of workers in manufacturing and professional services against workers' percentile in their comparative advantage in manufacturing, when $\kappa_m = 5$ and $\alpha_m = 1$. This case is similar to the conventional selection pattern where reallocating workers have lower income in both sectors.

Figure A.3: Expected Income Lower for Marginal Workers, $\kappa_m = 5$, $\alpha_m = 1$

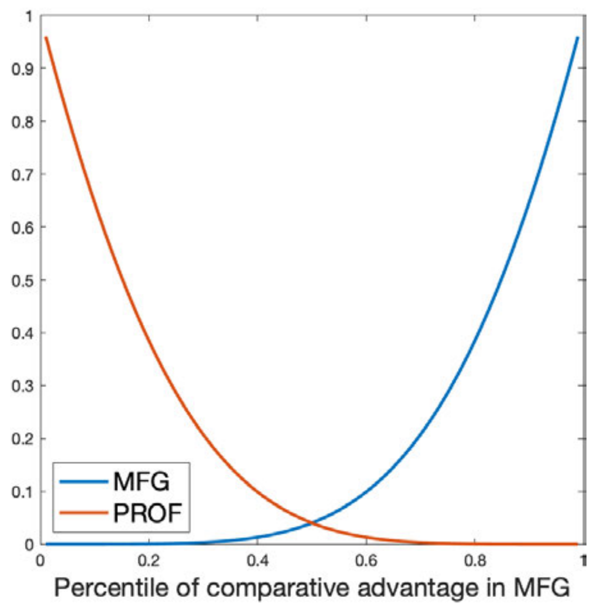

When α_m is sufficiently larger than κ_m, the case arises in which workers switching sectors have above-average income. The left panel of Figure A.4 plots the case when $\kappa_m = 1$ and $\alpha_m = 5$, and we can see that workers at the middle of the distribution could earn more than incumbent workers in either sector and will reallocate in response to wage shocks.

This specification accommodates the higher earnings of the marginal workers between manufacturing and professional services while retaining the property of the

Figure A.4: Expected Income Higher for Marginal Workers, $\kappa_m = 1$, $\alpha_m = 5$

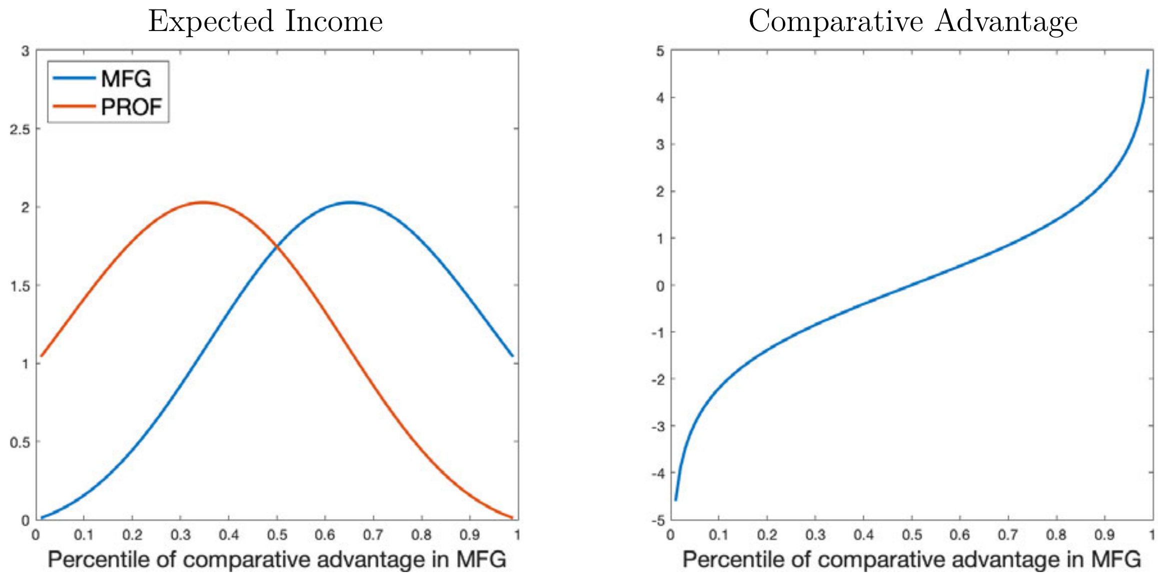

dispersion of comparative advantages in the baseline specification, as illustrated in the right panel of Figure A.4. Note that the pattern of reallocating workers earning more in both sectors is not sustainable for large enough shocks in wage rates. As shown in Figure A.5, when the wage rate in professional services becomes sufficiently larger than the wage rate in manufacturing, those reallocating from manufacturing into professional services will have above-average earnings in manufacturing but below-average earnings in professional services. This means that the pattern of selection collapses into one described by the baseline specification that the reallocating workers are the most productive in the contracting sector but less productive in the expanding sector.

Another way to reconcile the model with data is to look at the earnings gap after controlling for observable worker characteristics. Regressing the log real earnings of manufacturing workers on their sector in the next year, controlling for gender, college degree, and fixed effects of year and current occupation, the coefficient for

Figure A.5: Selection Pattern Collapsing into Baseline Case

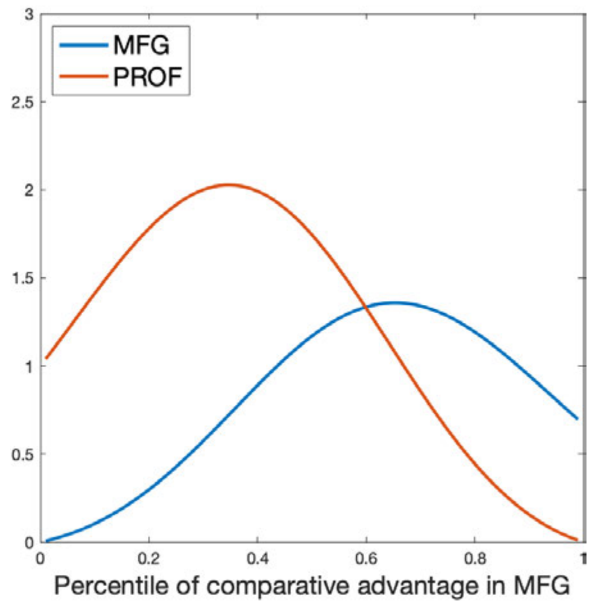

staying manufacturing workers is 4 log points[1] higher than that for workers leaving for professional services. Therefore, the patterns of selection reflected in the calibrated baseline model is consistent with the patterns of selection in data within gender, education, and broad occupation categories.

[1]With robust standard error of 0.8 log point.

Appendix B

Appendix to Chapter 2

Table B.1: Potential Gains of Each Country

Country	Year	Dynastic Bias	(1) Removing all barriers	(2) Matching US bias in 2015	(3) Shift component in (2)
Argentina	2001	1.1	2.2 %	-0.4 %	0.0 %
Armenia	2011	10.6	11.9 %	6.6 %	2.0 %
Belarus	2009	4.3	1.6 %	-0.1 %	-3.8 %
Benin	2013	11.8	12.7 %	5.6 %	9.3 %
Bolivia	2012	11.6	11.0 %	5.7 %	2.9 %
Botswana	2011	8.0	6.8 %	2.9 %	1.6 %
Brazil	2010	22.0	20.2 %	13.8 %	14.0 %
Cambodia	2013	12.2	14.0 %	6.3 %	8.0 %
Cameroon	2005	8.5	11.3 %	4.3 %	6.0 %
China	2000	7.5	7.0 %	2.4 %	3.5 %
Costa Rica	2000	5.8	4.1 %	1.0 %	-1.3 %
Cuba	2012	12.6	5.9 %	4.1 %	1.4 %
Ecuador	2010	12.0	9.4 %	5.8 %	1.7 %
Egypt	2006	11.2	10.9 %	6.1 %	4.2 %
El Salvador	2007	6.6	8.1 %	2.2 %	-0.5 %
Fiji	2014	4.7	5.2 %	0.2 %	-1.0 %
Ghana	2000	5.0	6.0 %	0.6 %	0.1 %
Guatemala	2002	9.7	11.1 %	5.5 %	5.1 %
Guinea	2014	6.2	6.8 %	1.6 %	4.9 %
Haiti	2003	5.4	4.4 %	0.5 %	2.5 %
Honduras	2001	14.2	14.0 %	8.0 %	9.6 %
India	2009	25.9	22.9 %	18.1 %	16.8 %
Indonesia	2005	18.1	16.5 %	10.8 %	13.3 %
Iran	2011	13.1	13.7 %	8.1 %	3.3 %
Jamaica	2001	6.0	5.8 %	1.1 %	-1.5 %
Jordan	2004	9.0	7.9 %	4.1 %	0.9 %

Country	Year	Dynastic Bias	(1) Removing all barriers	(2) Matching US bias in 2015	(3) Shift component in (2)
Malawi	2008	1.8	2.2 %	-3.0 %	0.4 %
Malaysia	2000	8.4	8.0 %	3.0 %	-0.3 %
Mali	2009	9.3	12.1 %	5.0 %	8.8 %
Mauritius	2011	8.8	4.3 %	2.1 %	0.7 %
Mexico	2015	12.4	11.1 %	6.2 %	3.4 %
Mongolia	2000	8.5	5.9 %	2.7 %	2.2 %
Morocco	2014	5.5	4.1 %	0.7 %	1.1 %
Mozambique	2007	3.1	3.5 %	-1.6 %	2.1 %
Nicaragua	2005	10.2	11.2 %	5.5 %	5.1 %
Panama	2010	7.8	7.9 %	2.7 %	-3.2 %
Papua New Guinea	2000	5.5	9.7 %	1.5 %	6.6 %
Paraguay	2002	18.9	17.5 %	9.8 %	11.0 %
Peru	2007	9.1	6.4 %	2.5 %	3.3 %
Philippines	2010	7.6	6.4 %	2.0 %	3.0 %
Puerto Rico	2000	3.1	3.1 %	-0.9 %	-1.3 %
Romania	2011	14.9	6.2 %	4.3 %	3.7 %
Rwanda	2012	8.6	11.8 %	5.1 %	6.4 %
Senegal	2013	8.7	11.0 %	4.5 %	6.7 %
South Africa	2007	11.4	9.0 %	5.3 %	2.2 %
Thailand	2000	39.9	27.7 %	18.4 %	10.1 %
Turkey	2000	6.7	7.1 %	2.1 %	2.8 %
Uganda	2014	3.0	4.1 %	-2.0 %	0.6 %
Venezuela	2001	6.7	8.6 %	2.2 %	1.7 %
Vietnam	2009	4.7	7.3 %	0.3 %	2.4 %
Zambia	2010	18.4	22.1 %	13.9 %	11.9 %
Zimbabwe	2012	9.1	11.0 %	4.9 %	4.1 %

Bibliography

Abramitzky, R., Boustan, L., Eriksson, K., Pérez, S., and Rashid, M. (2020). Census Linking Project: Version 2.0.

Adamopoulos, T., Brandt, L., Leight, J., and Restuccia, D. (2021). Misallocation, Selection and Productivity: A Quantitative Analysis With Panel Data From China. *Econometrica*.

Adão, R. (2016). Worker Heterogeneity, Wage Inequality, and International Trade: Theory and Evidence from Brazil.

Alvarez, J. A. (2020). The Agricultural Wage Gap: Evidence from Brazilian Microdata. *American Economic Journal: Macroeconomics*, 12(1):153 173.

Alvarez-Cuadrado, F., Amodio, F., and Poschke, M. (2019). Selection and Absolute Advantage in Farming and Entrepreneurship.

Baumol, W. J. (1967). Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis. *American Economic Review*, 57(3):415 426.

Baumol, W. J., Blackman, S. A. B., and Wolff, E. N. (1985). Unbalanced Growth Revisited : Asymptotic Stagnancy and New Evidence. *American Economic Review*, 75(4):806 817.

Bennedsen, M., Nielsen, K. M., Pérez-González, F., and Wolfenzon, D. (2007). Inside the Family Firm: The Role of Families in Succession Decisions and Performance. *The Quarterly Journal of Economics*, 122(2):647 691.

Besley, T. and Reynal-Querol, M. (2017). The Logic of Hereditary Rule: Theory and Evidence. *Journal of Economic Growth*, 22(2):123 144.

Black, S. E. and Devereux, P. J. (2011). Recent Developments in Intergenerational Mobility. In Ashenfelter, O. C. and Card, D., editors, *Handbook of Labor Economics*, volume 4B, pages 1487 1541. Elsevier B.V.

Bloom, N. and Van Reenen, J. (2007). Measuring and Explaining Management Practices across Firms and Countries. *The Quarterly Journal of Economics*, 122(4):1351 1408.

Bryan, G. and Morten, M. (2019). The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia. *Journal of Political Economy*, 127(5):2229 2268.

Burkart, M., Panunzi, F., and Shleifer, A. (2003). Family Firms. *Journal of Finance*, 58(5):2167 2202.

Burstein, A., Morales, E., and Vogel, J. (2019). Changes in Between-Group Inequality: Computers, Occupations, and International Trade. *American Economic Journal: Macroeconomics*, 11(2):348 400.

Chetty, R., Hendren, N., and Katz, L. F. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4):855 902.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *The Quarterly Journal of Economics*, 129(4):1553 1623.

Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural Change With Long-Run Income and Price Effects. *Econometrica*, 89(1):311 374.

Cruz, C., Labonne, J., and Querubín, P. (2017). Politician Family Networks and Electoral Outcomes: Evidence from the Philippines. *American Economic Review*, 107(10):3006 3037.

Dal Bó, E., Dal Bó, P., and Snyder, J. (2009). Political Dynasties. *Review of Economic Studies*, 76(1):115 142.

Drew, J. A. R., Flood, S., and Warren, J. R. (2014). Making Full Use of the Longitudinal Design of the Current Population Survey: Methods for Linking Records across 16 Months. *Journal of Economic and Social Measurement*, 39(3):121 144.

Duernecker, G., Herrendorf, B., and Valentinyi, Á. (2019). Structural Change within the Service Sector and the Future of Baumol's Disease.

Feenstra, R., Inklaar, R., and Timmer, M. (2015). The Next Generation of the Penn Word Table. *American Economic Review*, 105(10):3150 3182.

Galle, S., Rodríguez-Clare, A., and Yi, M. (2021). Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade. *Review of Economic Studies*.

George, S. E. (2020). Like Father, Like Son? The Effect of Political Dynasties on Economic Development.

Geys, B. and Smith, D. M. (2017). Political Dynasties in Democracies: Causes, Consequences and Remaining Puzzles. *Economic Journal*, 127(605):F446 F454.

Heckman, J. J. and Honoré, B. E. (1990). The Empirical Content of the Roy Model. *Econometrica*, 58(5):1121.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica*, 87(5):1439 1474.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and Manufacturing TFP of China and India. *Quarterly Journal of Economics*, 124(4):1403 1448.

Hurst, E., Rubinstein, Y., and Shimizu, K. (2021). Task-Based Discrimination.

Jarosch, G., Oberfield, E., and Rossi-Hansberg, E. (2021). Learning from Coworkers. *Econometrica*, 89(2):647 676.

Lagakos, D. and Waugh, M. E. (2013). Selection, Agriculture, and Cross-Country Productivity Differences. *American Economic Review*, 103(2):948 980.

Lo Bello, S. and Morchio, I. (2021). Like Father, Like Son: Occupational Choice, Intergenerational Persistence and Misallocation.

Long, J. and Ferrie, J. (2013). Intergenerational Occupational Mobility in Great Britain and the United States Since 1850. *American Economic Review*, 103(4):1109 1137.

Murphy, K. M., Shleifer, A., and Vishny, R. W. (1991). The Allocation of Talent: Implications for Growth. *Quarterly Journal of Economics*, 106(2):503 530.

Pulido, J. and Święcki, T. (2019). Barriers to Mobility or Sorting? Sources and Aggregate Implications of Income Gaps across Sectors and Locations in Indonesia.

Querubín, P. (2016). Family and Politics: Dynastic Persistence in the Philippines. *Quarterly Journal of Political Science*, 11(2):151 181.

Restuccia, D. and Urrutia, C. (2004). Intergenerational Persistence of Earnings : The Role of Early and College Education. *American Economic Review*, 94(5):1354 1378.

Roy, A. D. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3(2):135 146.

Sen, A. (2021). Structural Change within the Services Sector, Baumol's Cost Disease, and Cross-Country Productivity Differences.

Sinha, R. (2016). Intergenerational Occupational Mobility and Labor Productivity.

Smith, D. M. (2012). *Succeeding in Politics: Dynasties in Democracies*. PhD thesis, University of California, San Diego.

Solon, G. (1992). Intergenerational Income Mobility in the United States. *American Economic Review*, 82(3):393 408.

Solon, G. (1999). Intergenerational Mobility in the Labor Market. In Ashenfelter, O. C. and Card, D., editors, *Handbook of Labor Economics*, volume 3A, pages 1761 1800. Elsevier B.V.

Triplett, J. E. and Bosworth, B. P. (2004). *Productivity in the U.S. Services Sector*. Brookings Institution Press, Washington, D.C.

Young, A. (2014). Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services. *American Economic Review*, 104(11):3635 3667.

Zoghi, C. (2010). Measuring Labor Composition: A Comparison of Alternate Methodologies. In Abraham, K. G., Spletzer, J. R., and Harper, M., editors, *Labor in the New Economy*, pages 457 485. University of Chicago Press.

VITA

[redacted]