Job Polarization, Structural Transformation, and Labor Force Participation*

Alberto Vindas Q.

Arizona State University

avindasq@asu.edu

February 2017

Abstract

This paper investigates the quantitative importance of the labor force participation margin in the labor market outcomes for the United States. During the last 50 years, the labor force has been shifting from producing goods to producing services. In terms of occupations, the routine share decreased, giving way to increases in manual and abstract ones. I argue that these two patterns are related, and that the expansion of the labor force had an important role. I propose and estimate a labor allocation model where goods, market services, and home services use different occupations as inputs. The driving force is productivity growth, which is occupation-specific. Quantitative exercises show that holding this participation channel could significantly slow down polarization and structural transformation, and induce significant displacement within the labor force.

*I am greatly indebted to Berthold Herrendorf, Bart Hobijn, and Gustavo Ventura for useful comments, discussions, and support throughout the development of this project. All remaining errors are my own.

1 Introduction

The labor markets in the United States have changed significantly during the last 50 years. The share of routine occupations, heavy on procedural and repetitive tasks, decreased by 24%. On the other hand, manual occupations (heavy on physical tasks and in-person interactions), and abstract occupations (heavy on problem solving and creative tasks) increased by 26 and 38%. In terms of industries, the share of goods fell by 48%, giving way to an increase in services of 33%.[1] The first pattern is related to the *job polarization* process, the shrinking concentration of employment in routine occupations. The second pattern is related to the *structural transformation* process, the reallocation of economic activity across industries.

This transition has received considerable attention due to its implications on wage inequality and mobility costs. Routine occupations tend to be in the middle of the wage distribution, so job polarization pattern translates into higher wage differentials.[2] Human capital has been shown to have a heavy occupation-specific component, so displacements within the labor force would end up in costly adjustments for the workers.[3]

This paper speaks to the latter point. Some of the job polarization analyses frame this transition in terms of job losses and disappearing routine occupations, which ensue large and costly reallocations.[4] Using data from the Current Population Survey, I argue that this is not the case: the adjustment is mostly through an expansion of the labor force. Between 1968 and 2015 labor non-participation decreased from 33 to 23%, a drop of 30%. This translated to an increase in abstract and manual occupations, which explains the decrease in routine's labor share.

This paper also speaks to the occupation-industry mix in the labor force. In quantitative terms, both the occupation mix within industries, and the industry mix within the economy play an important role in explaining overall polarization. Occupational changes within industries have a stronger effect on the increase in

[1] These are percentages with respect to their initial labor shares, so these don't add up to zero.

[2] Goos & Manning (2007) study this pattern in the United Kingdom, Autor & Dorn (2013) focus on the United States, while Goos, Manning & Salomons (2014) expand the analysis for 16 Western European countries.

[3] Kambourov & Manovskii (2009) document this in the United States using data from the Panel Study of Income Dynamics.

[4] Examples include Autor (2010), Acemoglu & Autor (2011), Jaimovich & Siu (2012), and Mandelman (2016).

abstract occupations, while the shift towards services explains most of the increase in manual occupations.

In this article I ask how important the expansion in the labor force was, quantitatively speaking, for the productive structure of the economy. To answer that, I propose a labor allocation model explaining the occupational and industrial structure of the economy. It bridges some loose ends found in the literature: it incorporates the labor market participation decision, it justifies Baumol's cost disease from an occupational point of view, and gives the polarization process a proper treatment of the forces that take place between and within broad industries.

This model distinguishes between occupations in labor, and industries in consumption. Its building blocks are motivated by four patterns in the data. First, job polarization has played in a smooth, constant fashion during the last 50 years, so a persistent force should be behind these changes. Second, the adjustment implied an expansion of the labor force, which requires including this participation margin. Third, the occupational structure within goods and services differs substantially, so the industrial reallocation channel is of quantitative importance. Fourth, both the goods and services industries have polarized similarly, so the forces behind polarization should not be industry-specific.

On the production side, the model considers home and market production. Home production yields services, and market production yields goods or services. Home production requires one type of occupation, while market production requires manual, routine, and abstract occupations. These are imperfectly substitutable, so these enter as complements in a CES aggregator. As a starting point, labor is perfectly mobile, and has no sector or occupation specificity to it. Productivity growth is at the occupation level, so technical change is neutral across market sectors; industry productivity growth is uneven due to differences in their occupation intensities. Higher growth in routine productivity induces polarization (through complementarity), and higher intensity in routine occupations implies higher productivity growth in the goods industry.

On the preferences side, households supply their labor inelastically to home and market production, and consume goods and a basket of services. The preferences are represented by a CES aggregator, and consistent with the literature, goods and services are price-inelastic. The basket of services is made up by home and market services, which are aggregated through another CES aggregator as good substitutes. Preferences are homothetic and price effects play an important role. Income effects are generated through the interaction of home and market consumption.

To study the quantitative implications of the model, I calibrate it to the United States using data from 1968 to 2015. Productivity growth is the highest in routine occupations, followed by manual, abstract, and home production. The model is successful in reproducing the occupation dynamics within goods and market services, and is able to generate the movement towards market services we see in the data.

Finally, to assess the quantitative importance of the expansion in labor force participation, I perform two counterfactual exercises, exploring factors that could have affected this pattern. The first exercise, inspired by women's insertion into the labor force, freezes labor force participation at its 1968 level. This decreases the production of goods and market services by 2 and 18%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 2%. The second exercise, inspired by the home productivity slowdown reported in Bridgman (2016), has home productivity growing at the rate of market services. This increases the production of goods by 17%, decreases the production of market services by 27%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 7.5%.[5] This illustrates not only the importance of this channel, but that its causes also play a important role.

This article relates to the literature on job polarization, a pattern first noted in the United States in Acemoglu (1999). This occupations approach to the labor market was further studied in Autor, Levy & Murnane (2003), who examined the task content of these occupations and proposed the routinization hypothesis. Goos & Manning (2007) study this same routine-biased technical change for the United Kingdom, and more recently Goos, Manning & Salomons (2014) extend it to more European countries. Autor & Dorn (2013) use a spatial equilibrium model to study this hypothesis in the United States. Jaimovich & Siu (2012) and Cortés et al. (2014) study the United States polarization by focusing on its cyclical patterns. A recent review of the advances in this field is provided in Acemoglu & Autor (2011).

This article is also related to the structural transformation literature, the process of unbalanced growth described very early in Baumol (1967). Kongsamut, Rebelo &

[5] The comparison points are the model's predicted outcomes for 2015.

Xie (2001) study a model where labor reallocation is consistent with the Kaldor facts, and is driven by non-homotheticicities in preferences. Ngai & Pissarides (2007) propose a technology-driven explanation, in line with Baumol's cost disease. More recently, the importance of home production has been gaining attention, given the marketization hypothesis explained by Freeman & Schettkat (2005). Rogerson (2009), for instance, analyzes the market work decision by distinguishing between market and home produced consumption, and Buera & Kaboski (2012) focus on the service sector, distinguishing between home and market-produced services. Ngai & Pissarides (2008) embed this in a structural transformation model, and analyze a long time series to explain the trends in hours of market and home work. A recent review of the advances in this field is provided in Herrendorf, Rogerson & Valentinyi (2014).

The paper is organized as follows. Section 2 introduces the stylized facts behind job polarization. Section 3 presents the model I use: first I explain the environment and market structure, and then I discuss its equilibrium outcomes. Section 4 deals with the quantitative matters: first it explains the estimation procedure, and then it goes over the counterfactual exercises. Finally, section 5 concludes.

2 Data and Stylized Facts

This study focuses on the job polarization process, and the impact that the expansion of the labor force has had on it. As such, I present some facts relating these two, along with the productive structure of the economy. The goal is to establish relevant margins of adjustment in a clear fashion. In brief, the central points I want to illustrate are:

1. The job polarization process has been a smooth, constant reallocation during the last 50 years.
2. This adjustment has implied an expansion of the labor force, not displacements within it.
3. Both the goods and services industries have polarized similarly.
4. The industrial structure in the economy has contributed to overall polarization.

To do this, I use employment data from the Annual Social and Economic (ASEC) supplement to the Current Population Survey. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and labor force status. Appendix A presents a more detailed discussion of this data source and the classifications.

2.1 Job Polarization in the Labor Force

One of the stark, motivating facts is the job polarization that we observe in the labor market: how the labor share of routine occupations has decreased, while that of manual and abstract has increased. In this paper I follow Cortés et al. (2014) to construct the occupational categories, grouped by their task content. Abstract occupations are intense in tasks that require problem solving, judgment, and creativity. Some examples are managers, lawyers, and architects. Routine occupations are heavy on tasks that follow precise, and well understood routinary procedures. Examples include cashiers, machinists and travel agents. Finally, manual occupations rely more on tasks that require flexibility, in-person interactions, and physical adaptability. These include janitors, bartenders, and nursing aides.

Figure 1 shows the shares of these occupations for the 1968-2015 period. Manual occupations have increased slightly by 3 percentage points, routine occupations lost 14 percentage points, and abstract occupations increased 11.

Related to this figure, two observations are in order. First, the job polarization process has been a smooth one over time, and seems to be more related to a longer trend in the economy. Second, the polarization process had begun well before the 1980s. Several of the studies concerning polarization in the labor market focus on more recent decades. Since the CPS data start in 1968, we can go further back in time and state that is has been in place for all this period. This is in line with the findings of Bárány & Siegel (2015), where Census data show the same pattern for an even longer time span.

2.2 Job Polarization and Labor Force Participation

The job polarization process in the United States has not come about by a decrease in routine occupations. When looking at the entire working-age population, the share

Figure 1: Occupational Job Polarization

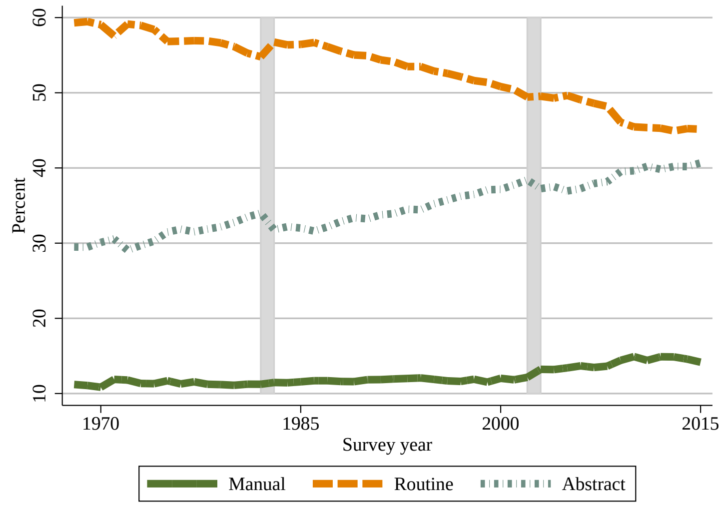

Shaded areas indicate years of major changes in occupational codes. These percentages refer to each occupation's share in employment.
Source: author's calculations using CPS.

of routine occupations has not declined substantially, and does not show a downward trend. This means that job polarization happened through expansions in the labor force, rather than displacements within it.[6]

Figure 2 shows this, by plotting the occupational distribution of the entire working-age population. It includes a fourth category, people out of the labor force, that groups the individuals that do not classify as employed. During this period, labor force participation increased by 10 percentage points, while abstract and manual occupations increased by 11 and 3 percentage points, respectively. On the other hand, routine occupations barely decreased: up until 2007, these had only lost 1.4 percentage points, and were fairly stable throughout the period. It was until the great recession that routine occupations did decrease notably.

It follows that the decrease in routine occupations' labor share depicted in figure 1 is explained by an expansion in the labor force. Moreover, since routine occupations' share in total population is stable over time, displacements in the labor force (or increases of it) can be discarded as the main force behind the fall of routine

[6]Moreover, when looking at life-cycle dynamics this pattern becomes even clearer: new entrants are the ones shaping this transition. These generational transitions are the next step in my research agenda.

Figure 2: Occupational Job Polarization & Labor Non-Participation

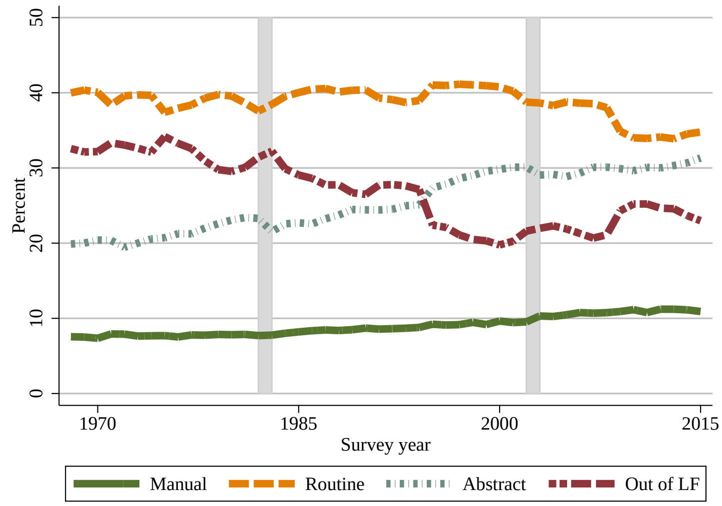

Shaded areas indicate years of major changes in occupational codes. These percentages refer to each category's share in total population.
Source: author's calculations using CPS.

occupations. This is consistent with the findings in Cortés (2016), where panel data show scant evidence of displacement.

2.3 Job Polarization and Structural Transformation

In this last section, I focus on the link between occupations, that work as labor inputs, and the productive structure that demands them. I do this to show that both the occupation mix within industries, and the industry mix within the economy play an important, and complementary role in explaining overall polarization. I also show that the changes within industries follow a similar pace, suggesting that these changes are independent of the industry. Occupational changes within industries account for most of the increase in abstract occupations, while the shift towards services accounts for the increase in manual occupations. For routine occupations, both play a similar role.

I organize the productive structure into two industries: goods and services. Following the standard approach, the goods industry consists of agriculture and manufacturing, which include, among other categories, forestry and construction. The services industry includes categories such as retail, and professional and entertainment services.

Figure 3 shows the shares of each occupation within total employment in each industry. These suggest that the within-industry occupational mix has evolved in a markedly similar manner. There is an almost one-to-one change in the shares of abstract and routine occupations, while manual occupations show little change. In the goods industry, the increase in abstract occupations was of 6 percentage points, matched with a decrease in routine occupations of 6 points as well. Manual occupations only decreased 0.2 percentage points. In the services industry, the change between abstract and routine occupations was of 9 percentage points, while manual occupations decreased 0.4 points.

These numbers imply that within-industry occupational dynamics cannot account for the increase in manual occupations' labor share, but can account for changes in abstract and manual's shares. Between-industry dynamics (the structural transformation process presented in figure 4) should be able to account for the change in manual occupations, and partly to the changes in abstract and routine. Occupational intensities lie behind this reasoning: services is more intensive in manual and abstract occupations, so a shift towards this industry can help explain their labor share increase.

A shift-share analysis can shed light on the relative importance of these two movements. We can express the economy-wide share of each occupation as a weighted average. For period t:

$$p_t(j) = \sum_I s_t(I) p_t(j|I) \tag{1}$$

where $p_t(j)$ is the economy-wide share of occupation j, $s_t(I)$ is the economy-wide labor share of industry I, and $p_t(j|I)$ is the share of occupation j in industry I. The change between period 0 and t can be decomposed into its *between* and *within* components as follows:

$$\Delta p_t(j) = \underbrace{\sum_I \Delta s_t(I) \bar{p}(j|I)}_{\substack{\text{Between industries} \\ \text{effect}}} + \underbrace{\sum_I \Delta p_t(j|I) \bar{s}(I)}_{\substack{\text{Within industries} \\ \text{effect}}} \tag{2}$$

$\bar{p}(j|I)$ is the average between time 0 and t of the conditional occupation share, and

Figure 3: Occupation Shares within Industries

(a) Abstract occupations

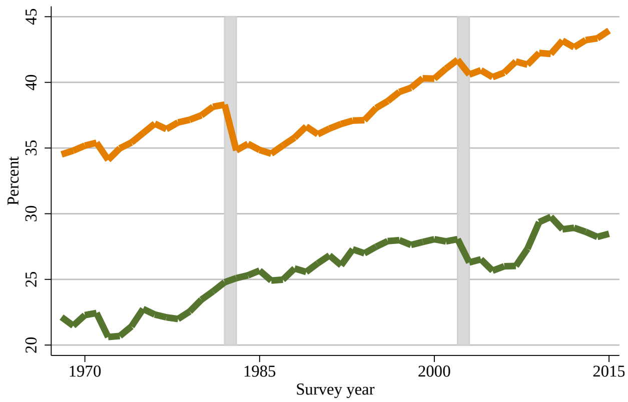

(b) Routine Occupations

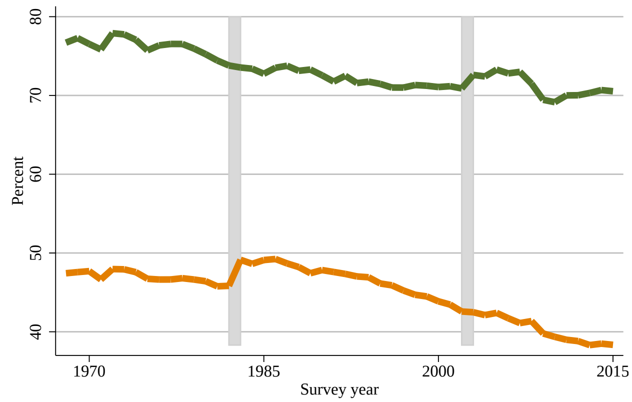

(c) Manual Occupations

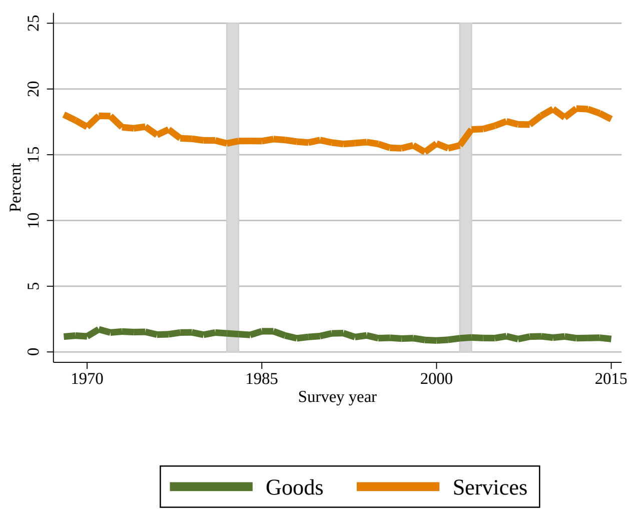

Shaded areas indicate years of major changes in occupational codes. These percentages refer to the share of each occupation in each industry's total labor demand.
Source: author's calculations using CPS.

Figure 4: Industry Shares

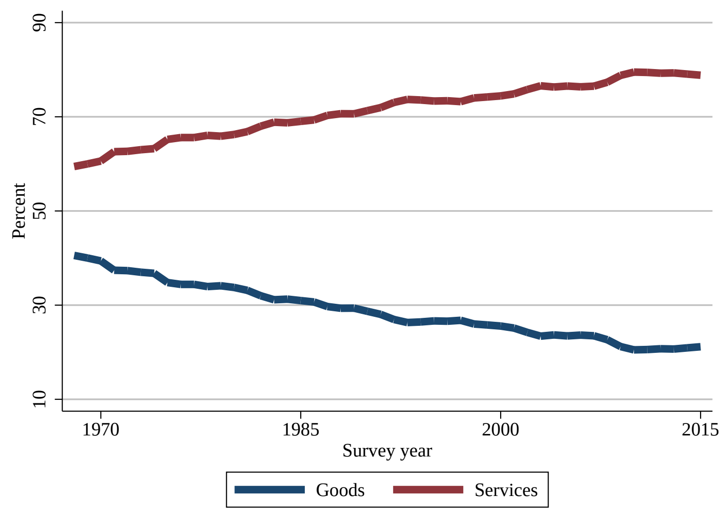

These percentages refer to each industry's share in the labor force.
Source: author's calculations using CPS.

and $\bar{s}(I)$ that of the industry share. The *between* effect refers to the impact of structural transformation, the changes in the productive structure of the economy. The *within* effect refers to the occupational mix inside each industry.

Table 1 shows this decomposition. Between-industry changes are the most important force for manual occupations, while within-industry changes are the stronger force for routine and abstract occupations. Regarding manual occupations, their share in both industries decreases, it is no surprise that the expansion of services accounts for all of its economy-wide increase. Routine occupations show a higher swing in their within-industry shares, which account for almost 60% of their economy-wide decrease. Between-industry changes also have an important role since goods demand routine occupations more intensively: on average their labor share is almost 30 percentage points higher than in services. Finally, in abstract occupations, within-industry changes account for more than 75% of its increase. The between-industry component is weaker because the intensity across industries is lower: on average their labor share in services is 13 percentage points higher than in goods.

The message of this decomposition is that a proper study of the economy-wide polarization must include the industrial structure. These results are in line with those of past studies. Autor & Dorn (2013) argue that the expansion of personal

Table 1: Shift-Share Decomposition of Changes in Occupational Shares

Occupation	Absolute Change (p.p.)			Relative Contribution (%)	
	Total Change	Between Industries	Within Industries	Between Industries	Within Industries
Manual	3.0	3.3	-0.3	110.0	-10.0
Routine	-14.1	-6	-8.2	42.6	58.2
Abstract	11.2	2.7	8.5	24.1	75.9

Source: author's calculations using CPS.

services lies behind the increase in low-skill, manual occupations. Moreno-Galbis & Sopraseuth (2014) go one step further, and attribute this pattern to an aging demographic structure. Tüzemen & Willis (2013) and Bárány & Siegel (2015) also show that within-industry changes are more important for routine and abstract occupations, using both CPS and Census data. All these show similar results when using finer classifications in both industries and occupations, which implies that this coarse classification is not a particular source of concern.

2.4 Discussion

This article is not the first to explore the topics of job polarization, structural transformation, and labor force participation. After all, these patterns have been noticed some years ago, and there have been extensive research efforts trying to explain and understand their consequences. To the best of my knowledge, however, this is the first article that attempts to study the effects of the labor force expansion in the polarization process, which then extends to the overall productive structure. Even though this is one of the major adjustment margins, its effect has gone unnoticed in the polarization literature. One notable exception is Cerina, Moro & Rendall (2017).

The object of study in this article is the allocation of work at the extensive margin.[7] Average hours of work, and therefore leisure, are an interesting phenomenon to explore, but their change does not seem to be the main driving force. Studies using time-use surveys in the United States, such as Aguiar & Hurst (2007), conclude that leisure has indeed been increasing over time. These surveys also show that it has been

[7]The extensive margin refers to the discrete choice of work sector. The intensive margin, on the other hand, refers to the amount of hours worked.

doing so relative to both market and home work, where the latter has traditionally been carried out by women. In this sense, gender differences played a major role in the polarization process, as women have also been behind most of the increase in abstract occupations. Interestingly enough, total hours of work (including market and home work) do not vary significantly across genders, and both have seen nearly identical increases in their leisure time. For an even broader group of countries, Burda, Hamermesh & Weil (2013) document the *iso-work* phenomenon: that total hours of work tend to be the same for both genders. They also document that leisure is positively related to income. This suggests that the intensive margin, at the aggregate level, is more related to the level of income in the country, rather than to its productive structure, and as such is left out of this analysis.

The trade-offs involved between the allocation of time in market and non-market activities go back in time to Becker (1965). In his theory, households are not only consumers, but producer themselves as well. Going further back in time, Kuznets (1941) recognized the relative importance of the activities in the domestic circle, and the tendency of the business system to take over some of these activities. By incorporating the fraction of people out of the labor force, and examining its decrease, I integrate these two ideas into the analysis. More recent vintages of studies, including Freeman & Schettkat (2005), Buera & Kaboski (2012) and Ngai & Petrongolo (2014), follow up on these ideas, which I do as well.

The relationship between polarization and structural transformation has been explored, for instance, descriptively in Tüzemen & Willis (2013), and in the context of a model in Autor & Dorn (2013), and Goos, Manning & Salomons (2014). Their empirical results go in hand with the ones presented here: different industries require occupations differently, and the swings in their labor shares can affect overall polarization. To those results, I add an additional observation: the changes between broad industries are also very similar.

The job polarization process happened through expansions in the labor force, rather than displacements within it. To that extent, it is worth exploring the importance of this channel, quantitatively speaking. One way to do this is to study a labor allocation model, where the agents choose where to work and what to consume. This could help explain why labor force participation has increased, and provide a coherent framework to study counterfactual scenarios. The following sections examine these points.

3 Baseline Model

In this section, I present a static model of labor allocation to study the patterns I presented earlier. I extend the models of Duernecker & Herrendorf (2016) and Ngai & Petrongolo (2014) to account for the ongoing polarization within industries, and the non-participation margin by including a home production sector.

The agents in this model choose between market and non-market work. In the market, firms decide how to allocate their labor into the three market occupations: manual, routine or abstract. In non-market work, agents devote time exclusively to home production. The driving force is *occupation specific* technical progress, and the difference in their growth rates induces the three main results: polarization, structural transformation, and changes in the labor force participation.

Theoretically, this paper's contribution is to provide a general equilibrium model that speaks to some loose ends found in different strands of the literature. From a structural transformation point of view, I link Baumol's cost disease to the ongoing polarization phenomenon, stemming from occupational productivities. From a polarization perspective, I formalize the demand-driven, non-homotheticity channel through which manual occupations have increased, and incorporate the labor force participation margin, which is a key missing link.

3.1 Environment

This is a discrete-time model where time runs forever. On the production side, I follow Ngai & Petrongolo (2014) and study three productive sectors: goods, market services, and home services. To distinguish between the jobs agents are working in, and the industries where these take place, I denote by lowercase j the occupation, and by uppercase I the industry. Then, $j \in \{h, m, r, a\}$, meaning these jobs can be in home production, manual occupations, routine occupations, and abstract occupations. Similarly, $I \in \{G, M, H\}$ denotes the production of goods, of market services, and of home services.

Home services are produced with a linear technology on home labor:

$$Y_{Ht} = A_{ht} N_{Hht} \tag{3}$$

where A_{ht} denotes the efficiency of home production, and N_{Hht} denotes the amount of labor used in home production.

The firms in goods and market services produce with a technology that requires the three types of market occupations: manual, routine, and abstract. These are combined according to a CES aggregator:

$$Y_{It} = \left[\sum_{j \in \{m,r,a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt} N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{4}$$

where $I \in \{G, M\}$ indicates the industry and $j \in \{m, r, a\}$ the occupation. In this setting, N_{Ijt} denotes the input that industry I uses of occupation j, and A_{jt} is the labor efficiency in occupation j. The elasticity of substitution between the labor inputs is $\sigma > 0$, and $\alpha_{Ij} \in (0, 1)$ is the intensity of occupation j in sector I. This productive structure is similar to Duernecker & Herrendorf (2016), where labor efficiency is occupation-specific as opposed to industry specific, which is the standard assumption in the structural transformation literature.

For notational convenience, define the following:

$$N_{It} = N_{Imt} + N_{Irt} + N_{Iat} \qquad I \in \{G, M\} \tag{5}$$

$$N_{jt} = N_{Gjt} + N_{Mjt} \qquad j \in \{m, r, a\} \tag{6}$$

$$\begin{aligned} N_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \end{aligned} \tag{7}$$

where N_{It} is the total amount of labor in industry $I \in \{G, M\}$, N_{jt} is the total amount of labor in occupation $j \in \{m, r, a\}$, and N_t is total market labor, which is clearly equal to the sum of labor over market industries or occupations.

On the consumption side, there are identical households of measure one. These consume goods and a combination of home and market services. The utility level they yield is aggregated according to a nested CES specification:

$$U_t(C_{Gt}, C_{St}) = \left[\omega_G^{\frac{1}{\varepsilon}}(C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}}(C_{St})^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{8}$$

where C_{Gt} and C_{St} denote the consumption of goods and compound services. Their relative weights are ω_G and ω_S, which add up to one, and individually are between zero and one. The elasticity of substitution between goods and compound services is $\varepsilon > 0$.

Compound services are also aggregated through a CES specification:

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}}(C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}}(C_{Ht})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{9}$$

where C_{Mt} and C_{Ht} denote market and home services. Their relative weights are φ_M and φ_H, which also add up to one, and individually are between zero and one. The elasticity of substitution between market and home services is $\eta > 0$.

All households are endowed with one unit of labor in each period. I denote by L_t total labor supply, and the remaining $1 - L_t$ is devoted to home production. Within the market, labor is perfectly mobile across occupations, and has no occupation or sector specificity to it.

The feasibility conditions for the consumption sectors are:

$$Y_{Gt} = C_{Gt} \tag{10}$$

$$Y_{Mt} = C_{Mt} \tag{11}$$

$$Y_{Ht} = C_{Ht} \tag{12}$$

These equations state that what is produced in the goods, market services and home services industries has to be consumed by the households.

Finally, the feasibility condition for the labor market requires that the households' labor supply be equal to the market demand:

$$\begin{aligned} L_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \\ &= N_t \end{aligned} \tag{13}$$

3.2 Decentralized Market Structure

I assume competitive markets for labor, goods and market services, where all agents take prices as given. Firms in the consumption and market services industries face the following profit-maximization problem:

$$\max_{N_{Imt},N_{Irt},N_{Iat}} p_{It}\left[\sum_{j\in\{m,r,a\}} \alpha_{Ij}^{\frac{1}{\sigma}}(A_{jt}N_{Ijt})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} - w_t(N_{Imt}+N_{Irt}+N_{Iat}) \tag{14}$$

where p_{It} is the market price of their output, and w_{jt} the market wages. First order conditions imply:

$$\frac{N_{Imt}}{N_{Irt}} = \frac{\alpha_{Im}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{mt}}\right)^{1-\sigma} \tag{15}$$

$$\frac{N_{Iat}}{N_{Imt}} = \frac{\alpha_{Ia}}{\alpha_{Im}}\left(\frac{A_{mt}}{A_{at}}\right)^{1-\sigma} \tag{16}$$

Equations (15) and (16) describe the relative labor allocations between manual and routine occupations, and abstract and manual occupations, each for industry $I \in \{G, M\}$.

Taking prices as given, the household's utility-maximization problem at time t is:

$$\max_{L_t,C_{Gt},C_{Mt}} \left[\omega_G^{\frac{1}{\varepsilon}}(C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}}(C_{St})^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{17}$$

subject to:

$$p_{Gt}C_{Gt} + p_{Mt}C_{Mt} = w_tL_t$$

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}}(C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}}(C_{Ht})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}$$

$$C_{Ht} = A_{ht}(1-L_t)$$

Households then, maximize their utility subject to their budget constraint in market products, and their technology constraint in household production. This problem can be solved in two steps: the first is to find the optimal allocation between home and market services, and the second one is for the optimal allocation between goods and compound services.

Home services are not traded in the market, meaning there is no market price attached to them. Its opportunity cost, however, is well defined since its alternatives are properly priced in the market. I denote by p_{Ht} this implicit price. The first order conditions to maximize C_{St} imply:

$$\frac{p_{Mt}C_{Mt}}{p_{Ht}C_{Ht}} = \frac{\varphi_M}{\varphi_H}\left(\frac{p_{Mt}}{p_{Ht}}\right)^{1-\eta} \tag{18}$$

Define the following price index:

$$p_{St} = \left[\varphi_H (p_{Ht})^{1-\eta} + \varphi_M (p_{Mt})^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{19}$$

This index can be interpreted as the unit price of the optimal services basket, which is relevant for the decision between the consumption of goods and the composite services basket. For this decision, first order conditions imply:

$$\frac{p_{St}C_{St}}{p_{Gt}C_{Gt}} = \frac{\omega_S}{\omega_G}\left(\frac{p_{St}}{p_{Gt}}\right)^{1-\varepsilon} \tag{20}$$

3.3 Equilibrium Outcomes & Discussion

In this model, the dynamics of labor reallocation are driven by the different rates of occupational productivity growth. How these growth rates end up affecting the final outcomes depend mostly on the elasticities of substitution and the labor intensities for the market industries. In this section I will discuss them and explain the discipline that the facts of section 2 impose on the model. In particular, this requires:

1. Market occupations to be complements in production.

2. Productivity growth to be the highest in routine occupations, followed by manual, abstract and home production.

3. Labor intensity in routine occupations to be higher in goods than in market services.

4. Home and market services to be substitutes in consumption.

5. Goods and services to be complements in consumption.

With the optimality conditions for the firms and households in hand, we can characterize how prices, labor and consumption patterns behave in equilibrium. I will do this by slicing the equilibrium outcomes around five components: polarization, industrial productivity growth, marketization, structural transformation, and labor force participation.

Job Polarization

The first pattern to center on is the job polarization within industries. This results from firms' optimal combination of inputs. Recall that in this model, labor is perfectly mobile and homogeneous, without any type of industry or occupation specificity to it.[8] Recall as well equations (15) and (16), that describe the relative labor demands:

$$\frac{N_{Imt}}{N_{Irt}} = \frac{\alpha_{Im}}{\alpha_{Ir}} \left(\frac{A_{rt}}{A_{mt}}\right)^{1-\sigma} \qquad \frac{N_{Iat}}{N_{Imt}} = \frac{\alpha_{Ia}}{\alpha_{Im}} \left(\frac{A_{mt}}{A_{at}}\right)^{1-\sigma}$$

The expression on the left shows the demand of manual occupations relative to routine occupations, while the expression on the right shows the same for abstract relative to manual occupations. How these relative demands evolve over time depends on their relative productivities and their degree of complementarity. This motivates the following proposition:

Proposition 1 *When labor inputs are complements, the occupation with the highest productivity growth will decrease its employment share, while the occupation with the lowest will increase it.*

Complementarity of the occupations is key for this result.[9] When the tasks that define the occupations are poor substitutes, $\sigma < 1$, and increases in the relative

[8]This assumption is a simplification meant as a first step to understand some of the mechanisms behind the changes in the productive structure. Labor heterogeneity is one of the extensions that I leave for future work.

[9]Recall that in this CES formulation, a unitary elasticity of substitution ($\sigma = 1$) implies a Cobb-Douglas production function. As their substitutability decreases, and $\sigma < 1$, inputs are called gross complements. In the limit case, when $\sigma \to 0$, the function converges to a Leontief technology. When $\sigma > 1$, inputs are gross substitutes, given their higher degree of substitutability. When $\sigma \to \infty$, substitutability is perfect and the production function converges to a linear technology.

productivity of one of the occupations will boost the relative demand of the others. To match its empirical counterpart, this structure requires the growth rate in the productivity of routine occupations to be the highest of the market ones, followed by manual, and finally abstract.

Most of the studies that use different labor inputs in production treat these as good substitutes, rather than complements.[10] The task approach to occupations calls for a different interpretation. Having low substitutability brings the CES technology closer to a Leontief specification, so the production process looks more like one where the amount of manual, routine, and abstract tasks are combined in a determined proportion to get output. Acemoglu & Autor (2011) discuss the case at hand; occupations are considered different factors of production, and increases in the productivity of a certain occupation creates an "excess supply" of it. In this sense, productivity increases can be thought of as occupation replacing. For practical purposes, this structure is very similar to Autor & Dorn (2013). In their model, capital replaces routine occupations, which makes abstract and routine occupations relative complements. In my model, these capital increases would boost routine occupation's productivity growth. Complementarity between these two and manual occupations are introduced through preferences, where they assume services are only produced with manual labor.

Given the complementarity among occupations, this model requires productivity growth to be higher in routine occupations, followed by manual and abstract. This again calls upon the task-based interpretation of occupations, in addition to the "reduced-form" type of productivity of this model. Goos, Manning & Salomons (2014) arrive to this result through an alternative route. Their output technology includes an added layer of depth: the production of tasks. They focus on their cost functions, and estimate negative coefficients on their routine components. As they state, their approach on input prices is observationally equivalent to one where the production function is modeled explicitly, as I do here.

Lastly, an implication of this model is polarization's neutrality across industries. Productivity growth is at the occupational level, and these occupations enter the production function in the same way for both industries (except for the intensities,

[10] Examples of this are Katz & Murphy (1992), that distinguish between educated and non-educated workers, and Caselli (2015), that distinguishes between experienced and inexperienced workers.

that do not vary over time). This neutrality is also present in the data: as discussed in the empirical section, both industries polarize similarly.[11]

Industry Productivity

The second pattern to focus on is the evolution of labor productivity at the industry level. For the goods and market services industries, it follows from equations (15) and (16) that occupation demands can be aggregated to a linear technology in total industry demand. At the optimal occupational demands, the production function (4) can be rewritten as:

$$Y_{It} = \tilde{A}_{It} N_{It} \tag{21}$$

where

$$\tilde{A}_{It} = \left[\alpha_{Im} \left(\frac{1}{A_{mt}} \right)^{1-\sigma} + \alpha_{Ir} \left(\frac{1}{A_{rt}} \right)^{1-\sigma} + \alpha_{Ia} \left(\frac{1}{A_{at}} \right)^{1-\sigma} \right]^{\frac{-1}{1-\sigma}} \tag{22}$$

is the average labor productivity in industry $I \in \{G, M\}$.

This productivity is a weighted average of each occupation's productivity, when labor inputs are combined optimally. With different productivity growth rates in occupations, growth at the industry level will be non-linear, will vary over industries, and will depend on each industry's occupational intensity. This is summarized in the following proposition:

Proposition 2 *When labor inputs are complements, and with constant productivity growth in occupations, the industry that uses more intensively the occupation with the highest (lowest) productivity growth increases its overall productivity the most (least). Asymptotically, industry productivity growth rates will converge to the rate of the lowest growing occupation.*

The first part of this proposition is fairly intuitive: if an industry relies more heavily in an occupation that has high productivity growth, its productivity gains will be higher. The second part is due to complementarity in production. As the

[11] One way to test this relationship is to take the ratio of (15) and (16) over industries, that should show little variance over time. The coefficients of variation are 0.21 and 0.11, which indicate relative stability.

share of the occupation with the higher productivity growth decreases, so does its contribution to the growth rate of that industry's productivity. Polarization, then, ends up dampening these productivity gains. At some point in the reallocation process, the share of the occupation with the lowest productivity growth will be so high that its effect on industry productivity will be the only discernible one.

Baumol's cost disease lies at the heart of this proposition. In their reappraisal of the unbalanced growth model, Baumol, Blackman & Wolff (1985) discuss that the progressivity or stagnancy of economic activities is caused by the technological advances behind their inputs, which correspond to occupations in this setting. The literature in structural transformation has established, by several measures, that goods-producing industries have had higher productivity growth, compared to services-producing industries.[12] This would require goods to be more intensive in routine occupations, and market services to be more intensive in abstract. The data show that this is clearly the case: on average, goods use 63% more routine inputs than services, and services use 50% more abstract.

For notational completeness, define the industry equivalent of (22) for home services production:

$$Y_{Ht} = \tilde{A}_{Ht} N_{Ht} \qquad \text{where} \qquad \tilde{A}_{Ht} = A_{ht} \tag{23}$$

Marketization

The third pattern to focus on is the marketization of home production, which relates to the reallocation of productive resources from the home sector to the market.

The decision to consume home services is slightly different than that of goods and market services. To do this, households must produce it themselves and give up the market income that they would otherwise earn. Then, the opportunity cost of home production is $p_{Ht} = w_t/\tilde{A}_{Ht}$. In equilibrium, the relative price of home to market services will be inversely related to the sectoral productivities:

$$\frac{p_{Ht}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \tag{24}$$

[12] A review of this literature, and evidence for several countries is presented in Herrendorf, Rogerson & Valentinyi (2014).

Households decide their consumption patterns in services according to (18). This translates into the following labor allocations:

$$\frac{N_{Ht}}{N_{Mt}} = \frac{\varphi_H}{\varphi_M} \left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \right)^{1-\eta} \tag{25}$$

Then, again, the evolution over time of this labor allocation will depend on the degree of complementarity and the relative productivity growth. The conditions to match the labor patterns in the data are summarized in the following proposition:

Proposition 3 *If home and market services are good substitutes, increases in the relative price of home production lead households to substitute its consumption with market services. To do this, they decrease the relative amount of labor dedicated to home production.*

This is the marketization result discussed in Freeman & Schettkat (2005): the United States has seen a shift of traditional household production to the market. In the model this would reflect a decrease in the relative labor allocated to home production. This, again, calls upon questioning how reasonable the assumptions behind this result are.

Firstly, we should analyze the existence of the home sector itself. Kuznets (1941) was well aware of this fact, and pointed out that incomes within the family economy were a prominent missing item in his estimates of national income. His approximations amounted to more than a quarter of national income in 1929. Hill (1985) uses time-use surveys to establish that for married couples in the mid-1970s, time spent on home work was only slightly behind market work. More recently, Aguiar & Hurst (2007) also find that the amount of hours in home production are substantial, compared to market hours, albeit decreasing over time. These observations should be enough to agree with Benhabib, Rogerson & Wright (1991, p. 1185): “models without home production implicitly make the assumption that the willingness or the incentive of individuals to substitute between market and nonmarket activity is small, but this does not seem to be the conclusion one would want to draw from the evidence.”

Secondly, in this model, the time that is not spent in market work is dedicated exclusively to home production. Ngai & Pissarides (2008) posit that in terms of production, only home services remain. They analyze a much longer time period, and use a structure where home work could be devoted to the three sectors in their study: agriculture, manufacturing, and services. By the late 1920s, they conclude that home production in agriculture and manufacturing was practically gone. Thus, I align with their observations, and assume these productions away.

Thirdly, assuming market and home services are good substitutes in consumption ($\eta > 1$) should not come as a controversial issue. Housework, shopping, food preparation, and caring for other people are among the activities that take most of the time in home production, according to time use surveys. These are all activities that can be easily purchased in the modern marketplace, thus their high degree of substitutability.

Lastly, since market and home services are good substitutes, households will tilt their consumption to the sector with lower price, which is the sector with the higher productivity growth rate. The increase in market participation requires a considerable difference between the growth of home and market productivities. Bridgman (2016) presents evidence to suggest that, effectively, productivity growth in the market has outpaced home productivity, in particular during this paper's period of study.

Structural Transformation

The fourth pattern focuses on the reallocation of consumption and productive resources between market industries, in particular from goods to services. Their relative prices are, again, inversely related to their sectoral productivities:

$$\frac{p_{Gt}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Gt}} \tag{26}$$

Despite the homotheticity in the utility function, this model features an income effect due to the home production sector. In equilibrium, the expenditure ratio and labor allocation between market industries are:

$$\begin{aligned}\frac{p_{Mt}C_{Mt}}{p_{Gt}C_{Gt}} =& \frac{N_{Mt}}{N_{Gt}} \\ =& \frac{\omega_S}{\omega_G} \underbrace{\left(\frac{\tilde{A}_{Gt}}{\tilde{A}_{Mt}}\right)^{1-\varepsilon}}_{\text{Price effect}} \underbrace{\left\{\varphi_M^{1-\varepsilon}\left[1+\frac{\varphi_H}{\varphi_M}\left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}}\right)^{1-\eta}\right]^{\eta-\varepsilon}\right\}^{\frac{1}{1-\eta}}}_{\text{Income effect}}\end{aligned} \tag{27}$$

This expression separates the portion driven by price effects, and the portion driven by income effects. The price effect responds to relative productivity between *market* industries, while the income effect responds to relative productivity between *service* industries.

The price effects behave similarly to the canonical structural transformation model of Ngai & Pissarides (2007). In this setting, however, occupation productivity growth is responsible for the growth at the industry level. Complementarity between goods and services (which requires $\varepsilon < 1$) implies that increases in the relative price of services result in increases in its expenditure share.

Income effects, on the other hand, are ultimately induced by marketization. Kongsamut, Rebelo & Xie (2001) introduce income effects with Stone-Geary preferences, and interpret the non-homotheticity term as home production. Here, home production is endogenized, which induces non-homotheticity between *market* goods and services. As home services become comparatively more expensive, households make up for this by switching out of home production, increasing market work and purchasing more services in the market. As market income increases, a bigger share of the expenditure goes to market services, even if the relative price of *market* goods and services were to stay the same.

The presence of price and income effects is consistent with Herrendorf, Rogerson & Valentinyi (2013), and their analysis of structural transformation. By examining both final expenditure and value added data, they conclude that goods and services show up as complements in the preference specification. Moreover, they also find a significant role for income effects, which are endogenized in the model through home production.

Labor Force Participation

The last component to focus on is the net effect of these forces on labor force participation. Structural transformation, with its decrease in the relative price of goods, makes households want to increase their consumption of services. Marketization, on the other hand, makes home production relatively more expensive to market services. Then, participating in the labor force involves a trade-off between home production and market consumption. In equilibrium, the ratio of market employment to non-participation is:

$$\frac{N_{Gt} + N_{Mt}}{N_{Ht}} = \underbrace{\left[1 + \frac{\omega_G}{\omega_S}\left(\frac{\tilde{A}_{St}}{\tilde{A}_{Gt}}\right)^{1-\varepsilon}\right]}_{\text{Structural transformation}} \underbrace{\left[1 + \frac{\varphi_M}{\varphi_H}\left(\frac{\tilde{A}_{Ht}}{\tilde{A}_{Mt}}\right)^{1-\eta}\right]}_{\text{Marketization}} - 1 \tag{28}$$

This expression separates the forces of structural transformation and of marketization. An increase in this ratio implies an increase in labor force participation.

Structural transformation frees up labor that can be used to produce services. This, absent a strong reallocation of consumption within services, would imply a lower participation rate. Home production would end up filling part of this increased demand for services. Marketization has the opposing effect, since it becomes relatively cheaper to consume more services from the market.

In the data, there is a sizable increase in labor force participation, most of which is used to fill the increased demand for abstract occupations. This means that the forces of marketization are considerably stronger than those of structural transformation.

Summary of the Model

A brief summary of this model starts with productivity growth rates at the occupational level, since their differences are the source of the reallocation patterns. The growth rate is higher in routine occupations, followed by manual, abstract, and finally home production. The interaction between these productivities, technology, and preferences yields five outcomes:

Polarization: firms demand more workers in the abstract and manual occupations, relative to routine, because labor inputs are complements in production.

Industry Productivity Growth: productivity growth is higher in goods because it is more intensive in routine occupations than market services.

Marketization: households work more in the market and substitute home with market services because of increasing opportunity costs of home production.

Structural Transformation: households demand more services because the relative price of goods decreases.

Labor Force Participation: participation increases because the effect of marketization dominates the effect of structural transformation on services.

4 Quantitative Results

This section explores the quantitative side of the model. First, I describe the estimation procedure, and analyze its results. With the estimation in hand, I conduct two counterfactual exercises to assess the importance of the increase in labor force participation.

4.1 Estimation

In this section, I explain briefly how to calibrate the model, and the moments I use. To begin with, I choose the two elasticities in the utility function based on previous studies. Based on Duernecker & Herrendorf (2016), I set $\varepsilon = 0.05$ (between goods and compound services). Based on Rogerson (2009) and Ngai & Petrongolo (2014), I set $\eta = 2.3$ (between market and home services).[13]

With these restrictions, there are eleven time-invariant parameters and four terminal conditions to determine: six occupation intensities (three for each market industry), the labor elasticity of substitution in market production, the four final occupation productivities, and the four preference weights in the consumption (for goods and services, and for home and market services). Assuming constant growth

[13]This elasticity is in the high end of the estimates available. It was purposely chosen as such, because these come from studies looking at the substitution between home and *total* market goods. This selection attempts to make up for goods being included.

rates, there is only need to look at the initial and final years, which are denoted here by $t = 0$ and $t = T$.[14] I back out their estimates from US data following these steps:

1. Impose the normalization $A_{m0} = A_{r0} = A_{a0} = A_{h0} = 1$.

2. Use the initial market occupation shares N_{Ij0} to solve for α_{Ij}.

3. Use the initial home and market services shares N_{H0}, N_{M0} to solve for φ_H and φ_M.

4. Use the initial market goods and services shares N_{G0}, N_{M0} to solve for ω_G and ω_S.

5. Use the final employment shares in the market services industry N_{MjT} to solve for final relative productivities $(A_{aT}/A_{rT})^{1-\sigma}$ and $(A_{mT}/A_{rT})^{1-\sigma}$.

6. Use the final relative employment share N_{MT}/N_{ST} to solve for the labor elasticity of substitution σ.

7. Use the growth factor of real per capita GDP to solve for A_{rT}.

8. Use the final home and market services shares to solve for A_{hT}

Further details of this procedure are discussed in Appendix C. Table 2 shows the time-invariant parameters of the model, and table 3 the occupation and industry productivity estimates.

These results are in line with those explained in the discussion section, so the qualitative predictions remain. Now we can comment on their quantitative side. As expected, growth in all occupation-specific productivities is positive. Productivity growth is such that by 2015, a worker in routine occupations is 36% more productive than a worker in manual occupations, and twice as much than a worker in abstract occupations. This goes in line with the routinization hypothesis, but established as a force working since (at least) the beginning of the sample. The elasticity of substitution between occupations is considerably lower than other estimates. Again, this stems from considering occupations as different factors of production, which is induced by the task-oriented grouping.

[14]Notice this same procedure could be used to infer a more detailed productivity path on a yearly basis, but the smooth labor share paths suggest this is a reasonable assumption.

Table 2: Model Parameters

	Parameter	Value	Source
ε:	Elasticity of substitution in utility b/w goods and combined services	0.05	Herrendorf, Rogerson & Valentinyi (2013)
η :	Elasticity of substitution in utility b/w home and market services	2.30	Ngai & Petrongolo (2014)
ω_G:	Preference weight for goods in utility	0.27	Initial goods and services relative price
ω_S:	Preference weight for combined services in utility	0.73	Initial goods and services relative price
φ_H:	Preference weight for home services in utility	0.45	Initial home and market services labor shares
φ_M:	Preference weight for market services in utility	0.55	Initial home and market services labor shares
σ :	Labor elasticity of substitution in production b/w market occupations	0.37	Final employment shares in market services
α_{Ga}:	Intensity of abstract occupations in market goods production	0.22	Initial industry-specific occupation shares
α_{Gr}:	Intensity of routine occupations in market goods production	0.77	Initial industry-specific occupation shares
α_{Gm}:	Intensity of manual occupations in market goods production	0.01	Initial industry-specific occupation shares
α_{Ma}:	Intensity of abstract occupations in market services production	0.35	Initial industry-specific occupation shares
α_{Mr}:	Intensity of routine occupations in market services production	0.47	Initial industry-specific occupation shares
α_{Mm}:	Intensity of manual occupations in market services production	0.18	Initial industry-specific occupation shares

All parameters, except the first two elasticities, are estimated from CPS data. See section 4.1 and appendix C for more details.

Table 3: Productivity Estimates

		Year		Average
		1968	2015	Growth Rate
A_h:	Home production	1	1.04	0.08%
A_a:	Abstract occupations	1	1.27	0.51%
A_r:	Routine occupations	1	2.63	2.08%
A_m:	Manual occupations	1	1.93	1.41%
$\tilde{A}_G$:	Market goods industry	1	2.16	1.65%
$\tilde{A}_M$:	Market services industry	1	1.87	1.34%

Table 4 shows the model's predictions with the estimated parameters. By design, it is able to match all of the labor shares in 1968, and the relative labor shares within services (for the three occupations in market services, plus the ratio between market and home services). The model is fairly successful at reproducing the relative occupation shares in the goods industry, and a little less so in reproducing the decrease in the goods share. The normalized root-mean-squared-error (RMSE) shows that for a model with constant growth rates, it is still satisfactory in capturing the main forces driving the changes in labor markets.

4.2 Counterfactual Exercises

One of the main goals of this paper is to assess the importance of the increase in labor force participation. The model at hand allows us to do this, and study its effects on structural transformation, on overall polarization, and on sectoral output. To do that, I perform two experiments. In the first one, I restrict the model so that the labor force participation stays at its 1968 level. In the second one, I increase the productivity growth in home production to shut down the marketization channel. The comparison point of these exercises are the output levels and labor allocations predicted by the baseline model in 2015. The first year of my sample coincides with the years when labor force participation started to increase. Juhn & Potter (2006) report that between 1948 and 1968 the participation rate remained relatively stable, and after that it increased. This provides a convenient turning point to perform counterfactual exercises.

Table 4: Actual and Fitted Labor Shares

	1968		2015		Normalized
	Data	Model	Data	Model	RMSE
Share of total population in:					
Goods production	27.3	27.3	16.3	21.2	14.5%
Market services production	40.1	40.1	60.7	57.2	8.4%
Home services production	32.6	32.6	23.0	21.6	8.7%
Share of goods labor demand in:					
Manual occupations	1.2	1.2	1.0	1.3	18.5%
Routine occupations	76.7	76.7	70.5	67.9	1.9%
Abstract occupations	22.1	22.1	28.5	30.8	5.2%
Share of market services labor demand in:					
Manual occupations	18.1	18.1	17.7	17.7	8.9%
Routine occupations	47.4	47.4	38.3	38.3	6.5%
Abstract occupations	34.5	34.5	43.9	43.9	3.8%

Normalized root-mean-squared-error is calculated using the average of the time series.

Freezing Labor Force Participation

The expansion of the labor force has drawn a considerable amount of attention recently, and has been associated with the insertion of women into the workplace. Several interpretations have been presented to explain this. One of them focuses on the social attitudes towards women's work: Fernández, Fogli & Olivetti (2004) discuss the gradual transformation of the family model (where working mothers set an example for future generations), while Goldin (2006) discusses several other social changes fueling the "quiet revolution" that made women think in terms of lifelong careers. In the spirit of these articles, I freeze labor force participation at its level in 1968. One interpretation of this counterfactual would be to consider how different the productive structure would have been, had societal attitudes towards women's work not changed as much. The results of this are in table 5.

Holding labor force participation fixed decreases output in both market sectors, disproportionately so for market services. This is not surprising, since this experiment deliberately holds down the inputs for market production, and restricts households to consume more home services than they would otherwise want. Within the workforce, this slows down structural transformation: the labor share of goods ends

Table 5: Freezing Labor Force Participation, Counterfactual Results

	Baseline model	Counterfactual prediction	Change (%)
Labor force participation	0.78	0.67	-13.9
Goods output	0.46	0.45	-1.8
Market services output	1.07	0.87	-18.4
Share in labor force of:			
Goods industry	0.27	0.31	14.2
Services industry	0.73	0.69	-5.2
Share in total population of:			
Manual occupations	0.10	0.09	-18.0
Routine occupations	0.36	0.32	-11.8
Abstract occupations	0.32	0.27	-15.0

This are the model's predictions for 2015. The counterfactual exercise freezes the labor force participation rate at its 1968 level, without assuming different rates of productivity growth. Percent changes are reported with respect to the baseline model's predictions.

up being 14% higher, despite market productivities remaining the same. This division of labor between goods and services makes the productive structure look like the model's prediction for 1999. Focusing on job polarization, the forces acting *within* both industries would still take place, so the labor force would polarize to a similar degree. The main difference is that in this case, the entirety of the adjustment takes place *within* the labor force, instead of through an expansion of it. Then, worker displacement would be a significant source of adjustment.

Increasing Home Productivity

Another explanation for the change in labor force participation focuses on the productivity in the household sector. In this vein, the second experiment keeps the labor force from expanding, but because of very different reasons. Bridgman (2016) suggests that productivity growth in the home sector had a considerable slowdown, coinciding with the initial years of this CPS sample. He also reports that previous to these years, home productivity had been growing at a similar pace to market activities. This is consistent with participation being stable before 1968, and increasing after this break. This begs the question: if the slowdown in home productivity was

Table 6: Increasing Home Productivity, Counterfactual Results

	Baseline model	Counterfactual prediction	Change (%)
Labor force participation	0.78	0.66	-15.5
Goods output	0.46	0.53	16.7
Market services output	1.07	0.78	-27.3
Share in labor force of:			
Goods industry	0.27	0.37	38.0
Services industry	0.73	0.63	-14.1
Share in total population of:			
Manual occupations	0.10	0.08	-26.2
Routine occupations	0.36	0.33	-9.9
Abstract occupations	0.32	0.26	-18.3

This are the model's predictions for 2015. The counterfactual exercise shuts down the marketization channel by increasing home production's productivity at the same pace as market services. Percent changes are reported with respect to the baseline model's predictions.

the sole cause of marketization, what would have happened without it? That is, how different would the productive structure look like if productivity in the home production sector had grown at the same rate as market services? The outcome of this experiment is in table 6.

The results in terms of output are more dramatic in this case: production in goods increases, while in market services it decreases even more. Households are much more efficient in producing their own services, so they do not increase their reliance on the market. This explains the greater hit that market services take. Labor force participation decreases; since I am purposely shutting down the marketization channel, only structural transformation has an effect on the consumption of total services (equation (28) shows precisely this). These effects add up to goods having an even higher share in market employment, which increases by almost 40%. The resulting division of labor makes the productive structure look like the model's prediction for 1977. In terms of polarization, the intra-industry reallocation would still take place. In this case, however, lower growth in market services leads to less polarization than in the first counterfactual. As in the first counterfactual, all the reallocation would happen within the labor force.

These two exercises illustrate how important the expansion in the labor force is for both structural transformation and polarization. In the two cases, both structural transformation and polarization are "slowed down" by the induced dynamics of home production. This raises the question of, for instance, to what extent the differences in the labor market structure of the United States and some European countries are driven by the channels highlighted in this model. Prescott (2004) explores the issue of differences in hours worked by looking at taxes. Home production could play a significant role in this setting. If cultural preferences favor home production, despite a lower productivity in it, it could look more like the first counterfactual exercise. If, on the other hand, European households did not experience a significant decrease in their home productivity growth, it could look more like the second counterfactual. These two cases show very different implications for growth and the causes for the differences in their productive structures. Polarization slows down, but the road it takes is quite different. Unfortunately, this model is unable to speak of mobility costs by design. Displacement costs would play an important role in terms of welfare, and could possibly slow down the adjustment process. This interesting avenue of research is left for the near future.

5 Conclusions

In this article I study three recent trends in the United States labor market: job polarization, structural transformation, and increases in labor force participation. With the goal of quantifying the importance of the labor participation margin, I propose a labor allocation model to explain this occupational and industrial structure.

Quantitatively, the model implies higher growth in the occupation of routine occupations, followed by manual, abstract, and finally home production. It is able to reproduce the occupational structure within industries, and the shift towards market services.

Counterfactual exercises suggest that this expansion is very important. Holding constant the attitudes that allowed the expansion in the labor force decreases the production of goods and market services by 2 and 18%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 2%. The second exercise, inspired by the home productivity slowdown reported in Bridgman

(2016), has home productivity growing at the rate of market services. This increases the production of goods by 17%, decreases the production of market services by 27%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 7.5%.

Future work includes extensions to this model to incorporate the age structure of the economy through an overlapping generations setting, considering the gender component to explain the wide differences in their labor market outcomes, and inducing labor heterogeneity to include an analysis of wage polarization.

References

Acemoglu, Daron. 1999. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." *American Economic Review*, 89(5): 1259–1278.

Acemoglu, Daron, & David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." , ed. O. Ashenfelter & D. Card Vol. 4 of *Handbook of Labor Economics*, Chapter 12, 1043–1171. Elsevier.

Aguiar, Mark, & Erik Hurst. 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *The Quarterly Journal of Economics*, 122(3): 969–1006.

Autor, David. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings." The Hamilton Project and Center for American Progress.

Autor, David, & David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–97.

Autor, David, Frank Levy, & Richard Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics*, 118(4): 1279–1333.

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

Baumol, William, Sue Anne Batey Blackman, & Edward N. Wolff. 1985. "Unbalanced Growth Revisited: Asymptotic Stagnancy and New Evidence." *The American Economic Review*, 75(4): 806–817.

Becker, Gary. 1965. "A Theory of the Allocation of Time." *The Economic Journal*, 75(299): 493–517.

Benhabib, Jess, Richard Rogerson, & Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy*, 99(6): 1166–1187.

Bridgman, Benjamin. 2016. "Home Productivity." *Journal of Economic Dynamics and Control*, –.

Buera, Francisco, & Joseph Kaboski. 2012. "The Rise of the Service Economy." *The American Economic Review*, 102(6): 2540–2569.

Burda, Michael, Daniel Hamermesh, & Philippe Weil. 2013. "Total Work and Gender: Facts and Possible Explanations." *Journal of Population Economics*, 26(1): 239–261.

Bárány, Zsófia, & Christian Siegel. 2015. "Job Polarization and Structural Change." Sciences Po Sciences Po publications 2015-07.

Caselli, Francesco. 2015. "Experience-Biased Technical Change." Centre for Macroeconomics (CFM) Discussion Papers 1517.

Cerina, Fabio, Alessio Moro, & Michelle Petersen Rendall. 2017. "The Role of Gender in Employment Polarization." Centre for Macroeconomics (CFM) Discussion Papers 1704.

Cortés, Guido Matias, Nir Jaimovich, Christopher Nekarda, & Henry Siu. 2014. "The Micro and Macro of Disappearing Routine Jobs: A Flows Approach." National Bureau of Economic Research, Inc NBER Working Papers 20307.

Cortés, Guido Matías. 2016. "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data." *Journal of Labor Economics*, 34(1): 63 – 105.

Duernecker, Georg, & Berthold Herrendorf. 2016. "Structural Transformation of Occupation Employment." University of Mannheim & Arizona State University Working Paper.

Fernández, Raquel, Alessandra Fogli, & Claudia Olivetti. 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *The Quarterly Journal of Economics*, 119(4): 1249–1299.

Freeman, Richard, & Ronald Schettkat. 2005. "Marketization of Household Production and the EU-US Gap in Work." *Economic Policy*, 20(41): 7–50.

Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review*, 96(2): 1–21.

Goos, Maarten, & Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *The Review of Economics and Statistics*, 89(1): 118–133.

Goos, Maarten, Alan Manning, & Anna Salomons. 2014. “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring.” *American Economic Review*, 104(8): 2509–26.

Herrendorf, Berthold, Richard Rogerson, & Ákos Valentinyi. 2013. “Two Perspectives on Preferences and Structural Transformation.” *American Economic Review*, 103(7): 2752–89.

Herrendorf, Berthold, Richard Rogerson, & Ákos Valentinyi. 2014. “Growth and Structural Transformation.” In *Handbook of Economic Growth*. Vol. 2 of *Handbook of Economic Growth*, Chapter 6, 855–941. Elsevier.

Hill, Martha. 1985. “Patterns of Time Use.” In *Time, Goods, and Well-Being.* , ed. Thomas Juster & Frank Stafford. University of Michigan.

Jaimovich, Nir, & Henry Siu. 2012. “The Trend is the Cycle: Job Polarization and Jobless Recoveries.” National Bureau of Economic Research Working Paper 18334.

Juhn, Chinhui, & Simon Potter. 2006. “Changes in Labor Force Participation in the United States.” *Journal of Economic Perspectives*, 20(3): 27–46.

Kambourov, Gueorgui, & Iourii Manovskii. 2009. “Occupational Specificity Of Human Capital.” *International Economic Review*, 50(1): 63–115.

Katz, Lawrence, & Kevin Murphy. 1992. “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *The Quarterly Journal of Economics*, 107(1): 35–78.

Kongsamut, Piyabha, Sergio Rebelo, & Danyang Xie. 2001. “Beyond Balanced Growth.” *The Review of Economic Studies*, 68(4): 869–882.

Kuznets, Simon. 1941. *National Income and Its Composition, 1919-1938.* National Bureau of Economic Research, Inc. with Lillian Epstein & Elizabeth Jenks.

Mandelman, Federico. 2016. “Labor Market Polarization and International Macroeconomic Dynamics.” *Journal of Monetary Economics*, 79: 1 – 16.

Moreno-Galbis, Eva, & Thepthida Sopraseuth. 2014. “Job Polarization in Aging Economies.” *Labour Economics*, 27(C): 44–55.

Ngai, Rachel, & Barbara Petrongolo. 2014. “Gender Gaps and the Rise of the Service Economy.” Centre for Macroeconomics (CFM) Discussion Papers 1404.

Ngai, Rachel, & Christopher Pissarides. 2007. "Structural Change in a Multi-sector Model of Growth." *American Economic Review*, 97(1): 429–443.

Ngai, Rachel, & Christopher Pissarides. 2008. "Trends in Hours and Economic Growth." *Review of Economic Dynamics*, 11(2): 239–256.

Prescott, Edward. 2004. "Why do Americans Work so Much More than Europeans?" *Quarterly Review*, 28(1): 2–13.

Rogerson, Richard. 2009. "Market Work, Home Work, and Taxes: A Cross-Country Analysis." *Review of International Economics*, 17(3): 588–601.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, & Matthew Sobek. 2015. "Integrated Public Use Microdata Series: Version 6.0." Minneapolis: University of Minnesota [Machine-readable database].

Tüzemen, Didem, & Jonathan Willis. 2013. "The Vanishing Middle: Job Polarization and Workers' Response to the Decline in Middle-Skill Jobs." *Economic Review*, 98(1): 5–32.

Appendix A Data Sources

The data I use covers the 1968-2015 period, and comes from the employment data in the Annual Social and Economic (ASEC) supplement to the Current Population Survey. I accessed these databases from the IPUMS-CPS project, an integrated set of data from the Current Population Surveys that goes through a convenient harmonization process. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and labor force status.

I construct the industry categories by grouping into goods and services. The goods industry includes both the manufacturing and agriculture sectors, which encompass manufacturing, construction, mining, agriculture, forestry, and fishery industrial classifications. The industrial classifications of the services industry are transportation, communications, public utilities, wholesale trade, retail trade, finance, insurance, real estate, business and repair services, personal services, entertainment and recreation services, professional and related services, and public administration.

To construct the occupation categories, I follow Cortés et al. (2014). They classify occupations based on two criteria: whether the tasks they involve are primarily manual or cognitive, and whether these are of a routine nature or not. Jaimovich & Siu (2012) note that there is a ranking in terms of wages: non-routine cognitive earn the highest while non-routine manual the lowest. In terms of tasks, non-routine cognitive tend to be high-skilled, while non-routine manual tend to be low-skilled. Routine manual and routine cognitive tend to be middle-skilled, so I group them together, in a similar fashion to Cortés (2016). I end up with three occupation groups then: non-routine manual, routine, and non-routine cognitive. Due to their association with the skills required, in the rest of the article I refer to these as manual, routine, and abstract occupations.

With every decennial Census, the occupation classifications are revised. These imply discrete jumps in their structure; even with the coarse grouping I use the changes are visible. The biggest changes were made with the 1983 and 2003 Censuses, which explain some of the shifts in the figures presented later. Both the harmonization processes from the IPUMS project and the analyses in Cortés et al. (2014) are careful enough to try and minimize these effects. In terms of the longer time trends, these reclassifications do not alter overall patters, and do not represent a significant concern.

Appendix B Detailed Tables

Table 7: Occupational Job Polarization

	1968	2015
Manual	11.2	14.2
Routine	59.3	45.2
Abstract	29.5	40.7

These percentages refer to each occupation's share in employment.
Source: author's calculations using CPS.

Table 8: Job Polarization and Labor Force Participation

	1968	2015
Manual	7.6	10.9
Routine	40.0	34.8
Abstract	19.9	31.3
Out of LF	32.6	23.0

These percentages refer to each category's share in the total population.
Source: author's calculations using CPS.

Table 9: Occupation Shares within Industries

	Goods		Services	
	1968	2015	1968	2015
Manual	1.2	1.0	18.1	17.7
Routine	76.7	70.5	47.4	38.3
Abstract	22.1	28.5	34.5	43.9

These percentages refer to the share of each occupation the industry's labor demand.
Source: author's calculations using CPS.

Table 10: Industry Shares

	1968	2015
Goods	40.6	21.2
Services	59.4	78.8

These percentages refer to the share of each industry in the labor force.
Source: author's calculations using CPS.

Appendix C Estimation Procedure

In this section, I explain in further detail how to match the data to the model's parameters. Recall the assumption of constant growth rates in productivity; because of that I only need to look at the initial and final years. These are denoted by $t = 0$ and $t = T$.

To get the market occupation intensities in production (α_{Ij}), I use equations (15)

(16). For the initial year, these imply:

$$\alpha_{Ir} = \alpha_{Im}\frac{N_{Ir0}}{N_{Im0}} \qquad\qquad \alpha_{Ia} = \alpha_{Im}\frac{N_{Ia0}}{N_{Im0}} \tag{29}$$

These two yield the intensities, since all three add up to one.

To get the relative weights in the consumption of market and home services (φ_H and φ_M), and in the consumption of goods and services (ω_G and ω_S), I use equations (25) and (27) in a similar fashion:

$$\varphi_H = \varphi_M\frac{N_{H0}}{N_{M0}} \qquad\qquad \omega_G = \omega_S\varphi_M\frac{N_{G0}}{N_{M0}} \tag{30}$$

To get the elasticity of substitution in the production function, I first rewrite the average labor productivity (22):

$$\begin{aligned}\tilde{A}_{It} &= A_{rt}\left\{\alpha_{Ir}\left[1+\frac{\alpha_{Im}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{mt}}\right)^{1-\sigma}+\frac{\alpha_{Ia}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{at}}\right)^{1-\sigma}\right]\right\}^{\frac{-1}{1-\sigma}}\\ &= A_{rt}\left\{\alpha_{Ir}\left[1+\frac{N_{Imt}}{N_{Irt}}+\frac{N_{Iat}}{N_{Irt}}\right]\right\}^{\frac{-1}{1-\sigma}}\end{aligned} \tag{31}$$

Then, I use equation (27) in the final period, substituting in equations (25) and (31):

$$\frac{N_{MT}}{N_{GT}} = \frac{\omega_S}{\omega_G}\left[\frac{\alpha_{Mr}}{\alpha_{Gr}}\frac{\left(1+\frac{N_{MmT}}{N_{MrT}}+\frac{N_{MaT}}{N_{MrT}}\right)}{\left(1+\frac{N_{GmT}}{N_{GrT}}+\frac{N_{GaT}}{N_{GrT}}\right)}\right]^{\frac{1-\varepsilon}{1-\sigma}}\left(\varphi_M\frac{N_{ST}}{N_{MT}}\right)^{\frac{1-\varepsilon}{1-\eta}} \tag{32}$$

Applying logarithms and rearranging leads to my estimate of σ.

To get relative occupational productivities, I use the occupation shares in period T, and equations (15) and (16) for the market services industry. Rearranging, these give:

$$\frac{A_{mT}}{A_{rT}} = \left(\frac{\alpha_{Mm}}{\alpha_{Mr}}\frac{N_{MrT}}{N_{MmT}}\right)^{\frac{1}{1-\sigma}} \qquad\qquad \frac{A_{aT}}{A_{rT}} = \left(\frac{\alpha_{Ma}}{\alpha_{Mr}}\frac{N_{MrT}}{N_{MaT}}\right)^{\frac{1}{1-\sigma}} \tag{33}$$

To get the productivity levels in market occupations, I use data from the Bureau of Economic Analysis. In particular, I take the real gross domestic product per capita time series (A939RX0Q048SBEA) to establish that this has grown by a factor of 2.23

between 1968 and 2015. To reproduce this growth pattern, I match this factor with market production evaluated at year 0's prices. Substituting (31) into the production function (4):

$$
\begin{aligned}
2.23 &= \frac{p_{M0}Y_{MT} + p_{G0}Y_{GT}}{p_{M0}Y_{M0} + p_{G0}Y_{G0}} \\
&= \frac{A_{rT}}{A_{r0}} \frac{N_{MT}\left[\alpha_{Mr}\left(1+\frac{N_{MmT}}{N_{MrT}}+\frac{N_{MaT}}{N_{MrT}}\right)\right]^{\frac{-1}{1-\sigma}} + N_{GT}\left[\alpha_{Gr}\left(1+\frac{N_{GmT}}{N_{GrT}}+\frac{N_{GaT}}{N_{GrT}}\right)\right]^{\frac{-1}{1-\sigma}}}{N_{M0}\left[\alpha_{Mr}\left(1+\frac{N_{Mm0}}{N_{Mr0}}+\frac{N_{Ma0}}{N_{Mr0}}\right)\right]^{\frac{-1}{1-\sigma}} + N_{G0}\left[\alpha_{Gr}\left(1+\frac{N_{Gm0}}{N_{Gr0}}+\frac{N_{Ga0}}{N_{Gr0}}\right)\right]^{\frac{-1}{1-\sigma}}}
\end{aligned} \tag{34}
$$

This yields the growth factor of productivity in routine occupations. With this I reconstruct the other productivity levels for market occupations.

Finally, from equation (25) in the final year, I get the productivity level in home production:

$$
\frac{A_{rT}}{A_{hT}} = \left[\frac{\varphi_M}{\varphi_H}\frac{N_{HT}}{N_{MT}}\right]^{\frac{1}{1-\eta}} \left[\alpha_{Mr}\left(1+\frac{N_{MmT}}{N_{MrT}}+\frac{N_{MaT}}{N_{MrT}}\right)\right]^{\frac{1}{1-\sigma}} \tag{35}
$$