

The Rise of Services: the Role of Skills, Scale, and Female Labor Supply

Francisco J. Buera, Joseph P. Kaboski, and Min Qiang Zhao *

August 6, 2017

Abstract

This paper provides a quantitative analysis of the growth in the service share in the U.S. We model households who make decisions on home and market production of services that vary in their skill intensity at any point in time and vary in their optimal scale over time. We also allow for skill- and sector-biased technology progress. The benchmark model fully accounts for the rise in the service share, with the rising scale of services, the rising demand for skill-intensive output, and skill-biased technical change all playing dominant roles. Furthermore, the model with multi-person households confirms that the essential findings of our benchmark model are robust to demographic considerations. It can explain two-thirds of the increase in female labor supply, which also plays a role in services growth.

1 Introduction

Over the past 50 years, the U.S. economy has moved increasingly toward a service-based economy, with the share of services rising roughly from 65 percent to 82 percent

**Affiliations and E-mail addresses:* francisco.buera@chi.frb.org (F. J. Buera, Federal Reserve Bank of Chicago), jkaboski@nd.edu (J. P. Kaboski, University of Notre Dame), and kent_zhao@xmu.edu.cn (M. Q. Zhao, Wang Yanan Institute for Studies in Economics, Xiamen University).

from 1965 to 2010. Several explanations have been proposed to explain this increase, including an increase in the optimal scale of service production and a shift in demand toward more skill-intensive output, which lead to an increase in the proportion of services that are market-produced relative to home-produced. These mechanisms are *qualitatively* consistent with several observations: growth in both the relative price and quantity of services, changes in patterns of home production, and, most importantly, growth in the average scale of service establishments and the shift toward skill-intensive services.¹ This paper develops and calibrates a quantitative model in order to evaluate the explanatory power of these mechanisms.

In the model, specialization plays a key role in the growth of the service economy. Specialized human capital is utilized more efficiently in the market, where workers specialize in production of particular goods or services. The increasing demand for skill-intensive services increases the returns to specialized human capital, so that workers who become skilled earn increasingly higher wages. As the opportunity cost of their time increases, they spend less time in home production and demand increasingly more market services. In addition, specialized intermediate/capital goods give rise to more efficient, larger scale production of services in the market than at home. In this way, a rising efficient scale of services interacts with both labor supply and investment in specialized human capital.

Although these are the forces of most interest, we add the possibility of both sector- and skill-biased technical change as possible explanations. In addition, our benchmark model moves beyond the representative household framework, introducing heterogeneity in the cost of acquiring skills. Finally, we extend the benchmark model by adding demographics that capture the different patterns in the data and incentives of married and single (male and female) households. This serves as a check on whether the quantitative importance of the mechanisms examined in the benchmark model are robust to multi-person households which can exhibit specialization within the household. Specifically, in a married household, one spouse may specialize predominantly in home production, while the other specializes in market production, and these decisions may be linked to decisions about human capital investment as

¹See Buera and Kaboski (2012a,b).

well.

Multi-person household considerations are not only an important theoretical consideration for robustness, but they appear quite empirically relevant. At the beginning of the period in question, women worked disproportionately in home production, while men worked disproportionately in the market. Shifts in female labor supply, due to both changes in the labor supply of married women and changes in marriage rates, are clearly linked to the growth of the service economy (see, for example, Lee and Wolpin, 2006, Ngai and Petrongolo, 2012, and Rendall 2014). As Figures 1 and 2 show, the growth in the service sector quantitatively mirrors the growth in female labor in services (as a percentage of the total labor force), while the decline of the goods sector matches the decline in male labor in goods. All four are roughly linear changes of at least 17 percentage points over the period in question, with the increase in female labor in services increasing 22 percentage points.² These extensions enable us to more closely match important features of the data, but also to assess the importance of female labor supply and demographic changes in explaining the observed patterns of structural change. More importantly, they can help us examine the robustness of the channels associated with the rising scale of services, rising demand for skill-intensive output, and skill-biased technical change when we consider demographics.

We calibrate the benchmark model to the U.S. experience. That is, we choose parameter values to target key facts of the U.S. economy in 1965, as well as the growth between 1965 and 2010 in output, schooling, the relative wage of college-educated workers, and the relative price of services.

Remarkably, despite no free parameters, the calibrated model is able to essentially explain the full growth in the service sector share. Counterfactual analyses allow us to highlight the quantitatively important channels in the model. Skill-biased technical change (SBTC) plays the most important role, accounting alone for over half of the growth in services. Skill-biased technical change increases the service share by

²In comparison, the relative size of the labor force that is female and working in the goods sector decreased by just 4.2 percentage points, while that of males in the service sector dropped by just 1.6 percentage points.

increasing the relative wage and relative quantity of high-skilled workers. The higher relative wage increases the opportunity cost of home production, thereby increasing the demand for market services from high-skilled individuals. The increasing proportion of high-skilled workers magnifies these effects, leading to a further increase in the share of services. The role of skill-biased technical change in the growth of services comes out of the fact that skills are specialized and therefore only productive in the market.

Moreover, the rising skill intensity of demand and the rising scale of services are also quantitatively important, together accounting for about as much growth in the service share as skill-biased technical change. Rising skill-intensity of demand due to non-homothetic preferences has a direct effect on the demand for services, as well as the indirect channels emphasized above for skill-biased technical change. The rising scale of services increases the costs of home production. Alone, these forces account for up to roughly one-third of the growth in services, and together they account for up to two-thirds.

In contrast, sector-biased technical change – the faster productivity growth in manufacturing – leads to a *fall* in the share of market services and a rise in home produced services. This is a unique feature of the model, and it is driven by home production being relatively more intense in manufactured goods. Standard biased productivity explanations for the growth of services assume a low elasticity of substitution between goods and services, so that higher productivity growth in the goods sector increases the growth of the service sector.³ These models predict a rising relative price of services, but a counterfactual decline in relative real quantities. In our model, a unique implication is that biased productivity in manufacturing actually *reduces* the growth of the service sector, since market services economize on intermediate goods/capital relative to home production. In contrast to biased productivity models, which require counterfactually large biases, the benchmark calibration of our model matches the growth in the relative price of services with productivity growth in the service sector that is roughly 0.6 percentage points lower than in the goods sector, relatively comparable to productivity measurements by Jorgensen and Stiroh

³See for example, Ngai and Pissarides (2007) and Baumol (1967).

(2000) over this period.⁴ We do not want to overstate this finding, since the stark assumption in our model may well overstate the true capital-intensity of home production relative to market services. We also likely overstate the traditional channel, however, by eliminating any traditional substitution between goods and services in production, and the channel unique to our model still dominates. We simply stress that this overlooked channel, which does not exist in many other models, appears to be quantitatively important.

We therefore conclude that the channels emphasized in the benchmark model are quantitatively important. The calibrated married household model confirms the essential findings of the benchmark model: fully explaining (actually slightly over-explaining) the rising share of services, with SBTC, the rising scale of services, and the rising demand for skill-intensive output still playing the leading roles. Although accounting for multi-person households somewhat weakens the importance of rising scale and rising demand for skill-intensive output, however, since it allows for partial specialization, the key message is nevertheless robust. Moreover, the model yields important additional insights. The married household model alone can explain 66 percent of the observed growth in female (market) labor supply. Demographics (i.e., the falling share of married couples) play a smaller but still significant role in explaining service share growth, but they explain three-quarters of the growth in female labor supply in the model. Counterfactuals keeping female labor supply fixed show that the endogenous increase in female labor supply, particularly married women, plays a role in the growth in the share of the service sector, but the contribution is not overwhelming.

This paper contributes to several related literatures. First, there has been a recent boom of research in the field of structural change.⁵ A wave of papers focused on simultaneously explaining structural change, including the growth of services, with balanced growth (e.g., Acemoglu and Guerrieri, 2008; Kongsamut, Rebelo, and Xie, 2001; Ngai and Pissarides, 2007). Recent work has begun quantitatively

⁴In Jorgensen and Stiroh (2000), the weighted average of labor productivity growth in the goods sector is 2.07 percent vs. 1.41 percent in the service sector. The analogous TFP growth rates were 0.67 and 0.26 percent.

⁵Herrendorf, Rogerson, and Valentinyi (2013) provide an excellent overview of this literature.

evaluating the standard channels in these theories (e.g., Buera and Kaboski, 2009; Echevarria, 1997; Herrendorf, Rogerson, and Valentinyi, 2013; Duarte and Restuccia, 2010). Our paper contributes a quantitative examination of the role of skill, scale, and female labor supply. Most directly, we build on the work of Buera and Kaboski (2012a, 2012b) by incorporating skill-biased technical change, heterogeneity, and multi-person households and quantify the mechanisms proposed there. Second, there is a micro literature linking skill and female labor supply (e.g., Goldin, 2006; Goldin, Lawrence, and Kuziemko, 2006; Mulligan and Rubinstein, 2008), and among these, Lee and Wolpin (2006) is most closely related since it deals with the service sector directly. We complement their work by focusing on the implications for aggregate output in the context of a general equilibrium model. Finally, we contribute to recent work quantitatively examining the interaction between long run patterns in female labor supply and transformations in home production over development (e.g., Gollin, Parente, and Rogerson, 2002; Greenwood, Seshadri, and Yorukoglu, 2005; Jones, Manuelli, and McGrattan, 2003). Ngai and Pissarides (2008) and Ngai and Petrongolo (2012) are most closely related, since they again address services directly. Again, our emphasis on skill and scale is unique and complements this literature.

The remainder of the paper is organized as follows. The benchmark model is introduced in Section 2 and calibrated and evaluated in Section 3. Section 4 extends the model and quantitative analysis to include multi-person households and discusses the feedback from education to technology. Section 5 concludes.

2 Benchmark Model

This section develops a benchmark model for evaluating growth of the service economy. We extend the model of Buera and Kaboski (2012a, 2012b) to allow for sector-specific technical change, skill-biased technical change, and time-varying efficient scale of services. In order to more easily model demographic changes, we also model heterogeneity across households in the cost of education/acquiring skills.

2.1 Production

There is a continuum of manufacturing goods and services, indexed by their complexity, $z \in [0, \infty)$. Manufacturing goods are produced only in the market, but services can be produced either in the market or at home. Manufacturing goods serve as intermediate input for both home and market production of final services. Technological progress is assumed to be exogenous, sector-specific, and skilled-specific.

2.2 Technologies

Manufactured goods are produced using low- and/or high-skilled labor, L_G and H_G , respectively:

$$G(z, t) = A_G(t) [A_l(z)L_G(z) + \phi(t)A_h(z)H_G(z)]. \quad (1)$$

Here, $A_G(t)$ is a manufacturing good-specific time-varying productivity term, $\phi(t)$ is a time-varying relative productivity between high-skilled and low-skilled workers, and $A_l(z)$ and $A_h(z)$ are time invariant but z -specific productivities of low- and high-skilled labor, respectively. We choose the following functional forms:

$$\begin{aligned} A_G(t) &= e^{\gamma_G t} \\ \phi(t) &= \phi_0 e^{\gamma_h t} \\ A_l(z) &= \frac{1}{z} \\ A_h(z) &= \frac{1}{z^\lambda}, \end{aligned}$$

where γ_G captures the manufacturing specific productivity growth rate and γ_h captures any skill-bias in technological change, respectively. Since z represents complexity, productivities are decreasing in z , but we assume $\lambda \in (0, 1)$, so that high-skilled work has a comparative advantage in more complex output.

Manufactured goods are used as inputs into the production of services. In particular, the production of service z requires one unit of manufactured good z as an intermediate input. Provided that a unit of the intermediate input is used, services of type z are produced with constant labor productivity up to a maximum capacity.

A simple example would be a washing machine that can do a maximum number of loads of laundry per day, with a certain amount of labor required for each load. Denoting the quantity of intermediate goods used as k_s , the (time-varying) maximum capacity as $n(t)$, and the quantity of low- and high-skilled labor by L_S and H_S , the production function is:

$$S(z, t) = \begin{cases} 0 & \text{if } k_S < 1 \\ \min\{n(t), A_S(t)[A_l(z)L_S(z) + \phi(t)A_h(z)H_S(z)]\} & \text{if } k_S \geq 1. \end{cases}$$

The capacity $n(t)$ will reflect the efficient output scale of a productive unit at which market services will be run, which we allow to change over time. In equilibrium, this parameter $n(t)$ will also be strongly related to the number of workers per productive unit. Note that the labor requirements for service z are symmetric to those for manufactured good z , except for the sector-specific term $A_S(t) = e^{\gamma_S t}$, which grows at rate γ_S .⁶

2.3 Firm's Problem

It is assumed that both manufacturing and service firms operate at the minimum average cost curves due to free entry. Making low-skilled labor the numeraire, and denoting the relative price of high-skilled workers as $w(t)$, equilibrium prices of manufactured goods and services are:

⁶The symmetry between the service and manufactured good production function can be strengthened by writing the manufacturing goods technology as:

$$G(z, t) = \begin{cases} 0 & \text{if } k_G < 1 \\ \min\{n_G(t), A_G(t)[A_l(z)L_G(z) + \phi(t)A_h(z)H_G(z)]\} & \text{if } k_G \geq 1. \end{cases}$$

Thus, (1) would arise as the limiting expression for large efficient scale in manufacturing, i.e., as $n_G \rightarrow \infty$.

$$p_G(z, t) = \frac{1}{A_G(t)} \min \left\{ \frac{1}{A_l(z)}, \frac{w(t)}{\phi(t)A_h(z)} \right\} \quad (2)$$

$$p_S(z, t) = \frac{p_G(z, t)}{n(t)} + \frac{1}{A_S(t)} \min \left\{ \frac{1}{A_l(z)}, \frac{w(t)}{\phi(t)A_h(z)} \right\}. \quad (3)$$

The competitive price of services includes two terms, the cost of intermediate goods and the cost of services value-added. The $n(t)$ in the denominator of the first term reflects the fact that intermediate goods are used at their efficient scale in market services.

The minimizations above reflect the choice between low- and high-skilled workers. Given our comparative advantage assumption, they define a threshold, $\hat{z}(t)$:

$$\hat{z}(t) = \left(\frac{w(t)}{\phi(t)} \right)^{\frac{1}{1-\lambda}}.$$

For $z \leq \hat{z}(t)$, firms will hire low-skilled workers. Conversely, when $z > \hat{z}(t)$, firms will hire high-skilled workers instead. The threshold $\hat{z}(t)$ is an increasing function of $w(t)$ and a decreasing function of the skill-biased productivity term $\phi(t)$.

2.4 Households

There is a continuum of infinitesimally-lived households that have preferences over the continuum of services. Households purchase market goods and services, provide labor to market and household production, decide which services to home produce, and whether or not to acquire specialized skills.

2.5 Preferences

Preferences over the continuum of discrete and satiable wants are indexed by the service that satisfies them, z . Define the function $C(z) : \mathbb{R}^+ \rightarrow \{0, 1\}$, which takes the value of 1 if a particular want is satisfied and 0 otherwise. Households can satisfy a particular want by procuring the service directly from the market or by

purchasing the required manufactured goods to home produce the service. Let the function $H(z) : \mathbb{R}^+ \rightarrow \{0, 1\}$ indicate whether want z is home produced. Together the consumption set is defined by the set of indicator functions, $C(z)$ and $H(z)$, mapping \mathbb{R}^+ into $\{0, 1\}$ ². The following utility function represent those preferences over wants and the method of satisfying those wants, i.e., over indicator functions $C(z)$ and $H(z)$:

$$\tilde{u}(C, H) = \int_0^{+\infty} [H(z) + \nu(1 - H(z))] C(z) dz, \quad (4)$$

where $H(z) \leq C(z)$. The parameter $\nu \in (0, 1)$ indicates that a home-produced service yields a greater utility.

These preferences capture growth on the extensive margin (variety of goods consumed) but at the expense of completely abstracting from the intensive margin (the amount of any given variety is fixed). More standard preferences, such as Stone-Geary preferences abstract from the extensive margin, instead modeling all growth on the intensive margin.⁷ Both are highly stylized, but our sequential preferences have the advantage of allowing for more continuous non-homotheticities that lead to a delayed growth in the service sector.⁸

Given that a continuum of wants are satiated sequentially, and production costs, as well as the additional costs of home production are increasing in z , the consumer's problem can be simplified by restricting to the consumption set consisting of step functions:

$$C(z) = \begin{cases} 1 & \text{if } z \leq \bar{z} \\ 0 & \text{if } z > \bar{z} \end{cases}$$

and

$$H(z) = \begin{cases} 1 & \text{if } z \leq \underline{z} \\ 0 & \text{if } z > \underline{z} \end{cases},$$

⁷Buera, Kaboski, and Rogerson (2015) is an example where preferences modeling intensive margin growth are used to quantify structural change. Stokey (1991) is an example of a paper with growth along both the intensive and extensive margin, but her preferences are not appropriate here.

⁸See Buera and Kaboski (2009) for a detailed explanation of the problems with Stone-Geary preferences. Other papers that solve these issues with preferences of intensive margin growth require indirect utility functions, e.g. Boppart (2014), Comin, Lashkari, and Mestieri (2015).

where \bar{z} denotes the most complex want that is satisfied and \underline{z} denotes the most complex want that is home-produced.

The primitive preferences described by (4) can then be represented by preferences over the restricted consumption set expressed by the utility function over two thresholds \underline{z} and \bar{z} :

$$u(\underline{z}, \bar{z}) = \underline{z}(1 - \nu) + \nu\bar{z}, \quad (5)$$

with $0 \leq \underline{z} \leq \bar{z}$. On the margin, there are two ways for agents to increase utility: by increasing \bar{z} to satisfy a new want or by increasing \underline{z} to move a previously market-satisfied want into home production.

2.6 Schooling

The schooling decision involves two choices: $e \in \{l, h\}$. l denotes low-skilled, and h denotes high-skilled. In order to become specialized high-skilled workers, $e = h$, agents must spend a fraction θ of their time endowment acquiring skills. The population is heterogeneous in terms of the time required to acquire specialized skills. More specifically, $\theta \in [0, 1]$ is distributed according to the c.d.f. $F(\theta)$.

2.7 Consumer's Problem

An individual with skill e solves:

$$\begin{aligned} V^e(\theta; t) &= \max_{0 \leq \underline{z}_e \leq \bar{z}_e} (1 - \nu) \underline{z}_e + \nu \bar{z}_e \\ \text{s.t.} \\ &\int_0^{\underline{z}_e} p_G(z, t) dz + \int_{\underline{z}_e}^{\bar{z}_e} p_S(z, t) dz \\ &= w_e \left(1 - \int_0^{\underline{z}_e} \frac{1}{A_S(t)A_l(z)} dz - \theta \mathcal{I}(e) \right), \end{aligned} \quad (6)$$

where $\mathcal{I}(e)$ is an indicator function that equals one if $e = h$ and zero otherwise. The left-hand side of the budget constraint includes expenditures on manufactured goods, which are intermediate inputs into the home production of services, and expenditures on market services. Note that home production of a single unit of service $z \in [0, \underline{z}_e]$ requires paying for an entire manufactured input, $p_G(z, t)$, rather than the $1/n(t)$ units required in market production. The right-hand side is income from market labor, which is the unit time endowment net of home production and schooling time. Note that, because high-skilled workers are specialized, all home production (except for a measure zero) is done with the productivity of low-skilled workers.

At an interior optimum, \underline{z}_e and \bar{z}_e must satisfy the following first order conditions:

$$\mu \left[\left(1 - \frac{1}{n(t)} \right) p_G(\underline{z}_e, t) + \frac{1}{A_S(t)} \left(\frac{w_e}{A_l(\underline{z}_e)} - \min \left\{ \frac{1}{A_l(\underline{z}_e)}, \frac{w}{\phi(t)A_h(\underline{z}_e)} \right\} \right) \right] \geq 1 - \nu \quad (7)$$

and

$$\mu p_S(\bar{z}_e, t) = \nu,$$

where $p_S(z, t)$ has been substituted using (3), and μ denotes the marginal utility of income.

Equation (7) is the marginal condition between home producing or market purchasing a service. The benefit of market services (left-hand side) includes the goods cost savings from the efficient utilization of intermediate goods and the potential labor cost savings from hiring either more productive high-skilled labor or low-wage, low-skilled labor. The cost of market services (right-hand side) is the disutility of market consumption. For any particular z , the goods cost saving will decrease as the price of the manufactured good falls and increase as the efficient scale of services rises. The labor cost savings of market services are higher for high-skilled workers ($w_e = w$). Thus, a shift toward high-skilled workers decreases home production time in favor of market services. Moreover, the labor cost savings are increasing in the relative wage of high-skilled workers w for high-skilled workers, but decreasing for

low-skilled workers ($w_e = 1$), so that increases in the relative wage affect workers differentially.

The schooling decision depends on the time cost and the relative wage. Being high-skilled will allow workers to earn a higher wage ($w > 1$), but it will reduce the time endowment to be $1 - \theta$, so the return to becoming high-skilled drops as θ increases. There exists a threshold, $\hat{\theta}(t)$, that equalizes that value of being high- and low-skilled $V_h(\hat{\theta}) = V_l(\hat{\theta})$. For $\theta < \hat{\theta}(t)$, a household will be strictly better off being high-skilled, while for $\theta \geq \hat{\theta}(t)$, a household remains low-skilled.

2.8 Equilibrium

Given $w(t)$, a household decides whether to be high-skilled and chooses the thresholds \underline{z} and \bar{z} . If a household decides to be low-skilled, $\theta \geq \hat{\theta}(t)$, the levels of $\underline{z}_l(t)$ and $\bar{z}_l(t)$ are independent of θ . If a household decides to be high-skilled, $\theta < \hat{\theta}(t)$, the levels of $\underline{z}_h(\theta, t)$ and $\bar{z}_h(\theta, t)$ will increase as θ decreases. Given $w(t)$, each firm sets the prices $p_G(z, t)$ and $p_S(z, t)$ according to (2) and (3), respectively. Summing up, a competitive equilibrium consists of allocations $\hat{\theta}(t)$, $\underline{z}_l(t)$, $\bar{z}_l(t)$, $\underline{z}_h(\theta, t)$, $\bar{z}_h(\theta, t)$, $\hat{z}(t)$ that solve (6) given prices $w(t)$, $p_G(z, t)$, and $p_S(z, t)$ and that are consistent with market clearing for manufacturing goods, services, low and high skilled labor markets.

The model can be solved in two steps recursively. Taking as given the wage $w(t)$, which determines the threshold $\hat{z}(t)$ and the price functions $p_G(z, t)$ and $p_S(z, t)$, the first step is to solve for the schooling threshold ($\hat{\theta}(t)$) and consumption thresholds ($\underline{z}_l(t)$, $\bar{z}_l(t)$, $\underline{z}_h(\theta, t)$, and $\bar{z}_h(\theta, t)$). The price functions are determined by $\hat{z}(t)$ and $w(t)$. The second step is to solve for $w(t)$ from a market-clearing condition given the schooling threshold and consumption thresholds. Then, repeat the first and second steps until convergence.

It can be shown that the disaggregate model can be expressed as a more standard model over aggregate consumption of manufactured goods and services, but the preferences vary with productivity. Moreover, productivity increases that are balanced, in the sense that $A_G(t) = A_S(t)$ and $\phi(t) = 1$, yield growth in the service sector that

is qualitatively consistent with several features of the data (see BK, 2012a, 2012b). First, the growth of services is delayed. At low levels of income, growth leads to new services being consumed in the market but old market services moving to home production as the cost of intermediates falls. This feature is least relevant for the quantitative analysis, since our analysis only covers the period of rising services. Second, and more relevant, the growth of services is driven by the growth of high-skilled services. As incomes continue to rise, demand shifts toward ever more complex output at which specialized high-skilled workers have an ever increasing comparative advantage. Market services increase as these complex services are more difficult to move into home production. In turn, the demand for high-skilled workers increases, and more agents decide to specialize. Given $F(\theta)$, the supply curve for skilled workers is upward sloping. As the relative wage increases, this increases the demand for market services among high-skilled workers, who constitute an ever increasing share of the economy. Third, since manufactured goods are produced in the market for the full range of z consumed, while only high z services are consumed in the market, market services are more intensive in high-skilled labor. *Ceteris paribus*, a rising relative wage w leads to increases in the relative price of services. Finally, the share of services is increasing in their efficient scale of production $n(t)$, which has trended up. This growth in scale in turn decreases labor used in home production in favor of market production, and thereby also increases the incentives for acquiring skill. The following section calibrates the relevant features to quantify the relative importance of these effects.

3 Quantitative Analysis of the Benchmark Model

We calibrate the preferences and technological parameters of the model to match key features of the 1965 U.S. economy and the observed changes between 1965 and 2010 in the market for skilled workers, the importance of home production, the relative price of services relative to manufacturing, and overall GDP. Importantly, in our calibration we do not target the change in the share of services in GDP. We then use the calibrated model to quantify the fraction of the rise of services than

can be accounted for by the exogenous driving forces in our model and the relative contribution of each of these forces.

We need to pin down ten parameters: one preference parameter ν that gives the utility of market services relative to home-produced, one fixed technological parameter λ that captures the comparative advantage of high-skilled workers in more complex output, two parameters describing the distribution of the cost of acquiring skills, a and b , the 1965 and 2010 values for the efficient scale of market services $n(t)$ and the relative productivity of skill workers $\phi(t) = \phi_0 e^{\gamma_n t}$, and the productivity growth of service and manufacturing production γ_S and γ_G . These parameters are chosen to match ten moments from the U.S. data. Four of these moments are for the initial period, 1965: the initial share of service in GDP, the initial share of intermediate manufacturing in service value added, the initial skill premium, and the initial fraction of high-skilled working-age population. Six of these are growth moments between 1965 and 2010: the increase in the fraction of high-skilled working-age population, the increase in the skill premium, the growth in the relative market work hours of high- to low-skilled working-age population, the growth in real per-capita GDP, the growth in the relative price of services to manufacturing, and the growth in the average size of service establishments.^{9,10}

Even though the mapping between some of the parameters of the model and the moments is jointly determined and highly non-linear, it is useful to describe heuristically the calibration by highlighting the moments that are primarily affected by each individual parameter. We start by discussing the parameters primarily determining the demand and supply of skilled workers.

We follow BK (2012a) in viewing college education as the appropriate empirical counterpart to high-skilled workers. The initial relative productivity of high-skilled workers (in low complexity output), ϕ_0 , can vary so that the relative wage in the model matches the college skill premium in 1965 of 1.41. The skill premium data are taken from weekly wage data from the Current Population Survey (CPS), using

⁹Details of data sources and calculations are available from the authors in an unpublished data appendix.

¹⁰In order not to confound our long-run focus with the impacts of the deep 2008 financial crisis and its aftermath, we use simple five-year averages for 2006-2010 as our endpoint targets

male full-time workers between the ages of 21 and 65.

We assume that the distribution of θ , the cost of acquiring skills in the model, follows a Beta distribution, $\beta(a, b)$, which supports θ between 0 and 1. This parametrization assures an interior solution for the fraction of workers acquiring specialized skills. The calibrated distribution can be left-skewed or right-skewed as well as symmetric, depending on the values of a and b . One of these parameters helps us target the fraction of working-age population (aged 21-65 in the CPS) that are college-educated in 1965, 0.23. The other parameter in the beta distribution, together with the parameter λ which captures the comparative advantage of high-skilled workers in more complex output and the skill-biased technical change parameter γ_h , can be varied to match the increase between 1965 and 2010 in the fraction of high-skilled workers (36 percentage points), the increase in the skilled premium (41 percentage points), and the growth in the relative market hours of high- to low-skilled workers (3 percent).

We target several other time trends over the 1965 and 2010 period as well. We choose the rates of technical change in each sector, γ_S and γ_G , to match growth in real GDP per capita (144 percent) and the change in the relative price of services to manufacturing, which increased by 44 percent over the same period. Notice that movement toward more complex and skilled intensive services also generates a rise in the relative price of services to manufacturing. Thus, sector-specific technical change is a complementary force affecting this relative price.

The technology parameter n determines the ratio of intermediate manufacturing inputs to value-added. We choose its initial value n_0 to target this value, which from input-output tables is 0.12 in 1965. Changes in n translate into changes in workers per establishment in services. Given the initial efficient scale value, n_0 , the remaining time series of n is constructed from data on the workers per service establishment. The average service establishment has 1.56 times as many workers in 2010 as in 1965 based on County Business Patterns data.

Finally, given the other parameters, the utility of market services relative to home-produced ν can be varied to pin down the initial share of services, which in 1965 was 0.65.

A summary of parameters and targets is given in Table 1. As shown in the second column of Table 1, the benchmark model is able to hit all the data moments. The calibrated θ distribution is right-skewed (with a larger mass on the smaller values of θ). Although the rising skill premium itself leads to some growth in the relative price of services, targeting the relative prices still requires slightly lower TFP growth in the service sector of 0.0135, which is about two-thirds the TFP growth rate in the manufacturing sector (0.0193).¹¹ This is relatively comparable to productivity measurements by Jorgensen and Stiroh (2000) over this period, which is allowed for by the fact that the rising skill premium accounts for the balance of the growth in the relative price of services. A more standard model of biased productivity growth would need (counterfactually) larger bias sectoral productivities.

The skill-biased productivity growth adds roughly half a percentage point to high-skilled workers productivity annually, which amounts to about 27 percentage points by 2010. The relative wage is 41 percentage points higher in 2010 than in 1965, so the remainder comes from the movement toward more complex goods and the comparative advantage parameter λ .

3.1 Accounting for the Rising Service Share

We now analyze the model's predictions for the change in the service share over time. Note that this is purely an out-of-sample test, since the change in the service share was not targeted by our calibration. We will focus on the predictions for the long-run change between 1965 and 2010. The higher frequency dynamics of this change are not particularly interesting; the model itself is static, we do not account for business cycle factors, and the calibration assumed linear productivity trends. We simply note that the effects occur fairly linearly with increased productivity, so the model matches the relatively stable time trends in the data quite well in this regard, with the exception of the skill premium, which declined in the 1970s before accelerating

¹¹This confirms the results in Buera, Kaboski, and Rogerson (2015), who finds that changes in the skill premium account for a relatively small fraction of sectoral relative price changes.

in the 1980s.¹²

The model does quite well in reproducing this growth in the service share as shown in Table 2. In 1965, the service share in the model matches that in the data (0.650) by construction, i.e., because it is a target in our calibration. In 2010, the model predicts a service share of 0.809, nearly identical to the 0.813 in the data. This is our first important finding: the model is able to fully explain the 16 percentage point increase in the share of services observed in the data. To put this change in perspective, 16 percentage points currently exceeds the total size of the manufacturing sector in 2012.

We now examine which factors are most important in accounting for this increase. We have four exogenous factors that change over time, which we examine in turn. We examine their role by running counterfactuals where either the factor in question is held constant in the model (i.e., the factor is “turned off”) or where the factor in question is the only factor not held constant in the model (i.e., the only factor “turned on”). We turn factors off by keeping the relevant parameters at their calibrated 1965 levels and turn factors on by setting them at their calibrated 2010 levels.

The results are shown in Table 3. Since the calibration hits the 1965 service share for every simulation, we focus on the overall service increase explained by different simulations and how it differs from either the benchmark simulation, where we turn off factors, or how it differs from zero, where we turn on factors. Turning on factors is giving the 2010 value to the 1965 economy, while turning off factors is effectively giving the 1965 value to the 2010 economy.

The first factor is simply the increase in productivity, which pushes demand for more skill-intensive services because $\lambda < 1$. We call this the income effect. As explained by BK (2012a), this has a direct effect, since for these more complex services, market services are cheaper than home production. It also has an indirect effect by increasing the demand for high-skilled workers. The share of services in consumption is higher for high-skilled workers, since their opportunity cost of home

¹²The literature has typically pointed to the importance of cohort effects, specifically the Baby Boom, in explaining this, while assuming a constant skill bias in technical change (e.g., Katz and Murphy, 1992). These cohort effects are clearly outside the model.

production is higher. Moreover, it is increasing in the skill premium, which captures this opportunity cost. Hence, higher demand for skill leads to both a higher skill premium and more high-skilled workers, both of which contribute to a higher share of services. Quantitatively, Table 3 indicates that this effect accounts for a 5.2 percentage point smaller increase in services when it is the only factor turned off and a 3.8 percent increase when it is the only factor turned on. These effects amount to 33 and 24 percent, respectively, of the total increase in the benchmark model.

The second factor, emphasized by BK (2012b), is the rising scale of services, $n(t)$. Larger scale services lead to a larger cost differential between home and market services because the market economizes on the manufactured inputs, which are a fixed cost. Thus, larger scale services lead to more market services. Quantitatively, this factor is also non-negligible, accounting for a 4.4 percentage point smaller increase in services when it is the only factor turned off and a 6.3 percentage point increase when it is the only factor turned on. These amount to 28 and 40 percent of the total increase, respectively. The difference comes from the fact that cost differences are driven by scale relatively more in 1965, but by skill relatively more in 2010.

These first two factors are unique to this style of model. To see how important these two factors are, we turned them both on and off together. Turning both off together leads to a 10.7 percentage point smaller increase in the service share, while turning both on alone would lead to a 9.7 percentage point increase in the service share. These constitute 67.1 and 61.2 percent of the total increase, respectively. Thus, it appears that these two unique factors play an important role individually, and together they account for the bulk of the increase in the service share.

The third factor, skill-biased technical change, is also very important, however. Skill-biased technical change is certainly part of other models that explain the trends in the skill premium and the supply of skills, but our emphasis on specialized skills being specific to market production implies that skill-biased technical change also leads to the growth in services. The logic is the same as for the indirect channels of the income effect explained above, where the higher demand for skill leads to a higher opportunity cost of home production for high-skilled workers and a higher fraction of high-skilled workers. Skill-biased technical change accounts for a 11.6

percentage point smaller increase in services when it is the only factor turned off and a 9.4 percentage point increase in services when it is the only factor turned on. Both of these are larger than the combined impacts of the first two factors. These amount to 73.2 to 58.9 percent of the total increase in services, respectively.

Thus, if all factors were additive and positive, we would already have over-accounted for the increase in services.

However, the fourth factor, sector-biased technical change, works in the opposite direction. In most biased productivity models (e.g., Ngai and Pissarides, 2007), faster technical change in the manufacturing sector (coupled with an elasticity of substitution between goods and services that is less than one) leads to a rising share of the service sector. In our model, however, one of the benefits of market services is that they save on the cost of manufactured inputs by operating at the maximum scale, $n(t)$. Biased technical change in favor of manufacturing makes these inputs become relatively cheap. As inputs become cheap, the cost savings from using market services disappears. People substitute toward more manufactured goods for home production and fewer market services. Admittedly, the starkness of our model may somewhat overstate this channel, since we assume that home production is $n(t)$ times more capital-intensive than market services for any specific service z , but our extreme choice of no substitutability between goods and services within the (Leontieff) market or home production technologies, also maximizes the magnitude of the more traditional channel. The fact that quantitatively the former dominates the latter, suggests that the channel here, which is an omitted channel in the broader literature, could well be quantitatively important.

In order to quantitatively isolate the effect of biased technical change from overall technical change (i.e. the productivity/income effect of the first factor), we change relative productivity across the sectors without changing the original utility level. Table 3 indeed shows that this factor works to reduce the share of services, but the strength of this factor depends strongly on the presence of others. When it is the only factor turned off, i.e., it is turned off in 2010, it leads to only a 1.8 percentage point greater increase in services. This is because the other factors make the cost savings of market services in 2010 primarily skill-driven rather than goods-driven.

However, when it is the only factor turned on, i.e., if it is turned on in 1965 when the goods cost savings coming from market services are substantially larger, it leads to an 11.0 percentage point decrease in the service share.

The results indicate that the market for skill plays an important role in the rise of services. Table 4 illustrates this more clearly by showing the role of the endogenous increase in the skill premium and the endogenous increase in schooling attainment. We do this by solving and aggregating households' problems at the benchmark equilibrium prices, but keeping either the relative wage fixed, schooling decisions fixed, or both fixed at their 1965 values (Effectively, we model and aggregate a partial equilibrium economy, where goods and labor markets need not clear). We learn two things. First, when both are kept fixed, the increase in the service share is only 4 percentage points, indicating that these labor market adjustments coming from the increased demand for skill are critical. Second, we see that the increase in the skill premium plays an important role in any case, but that the increase in schooling only plays an important role when the skill premium also increases. This is because when the skill premium is high, the share of services in high-skill consumption is much higher than it is in low-skill consumption.

Table 5 examines the impact of these various factors on the other important changes over time that we targeted when calibrating the model: the increase in the fraction skilled, the increase in the skill premium, the growth of measured GDP per capita, and the growth in the relative price of services. The increasing fraction of high-skilled and the skill premium both reflect a rising demand for high-skilled labor, and so the income effect toward skill-intensive output and skill-biased technical change play the leading roles. SBTC is about twice as important, but the income effect still plays a significant role. The relative contribution of SBTC and structural change forces in explaining the rise in the skill premium is in line with the results in Buera, Kaboski and Rogerson (2015), who label the latter forces as skill biased structural change. For the growth in real GDP per capita, naturally the income effect almost exclusively drives things. However, the other factors still play significant roles by moving output between unmeasured home production and measured market services. Finally, the growth in the relative price of services is driven largely by

sector-biased technical change, but scale effects also play a role.

To summarize, the benchmark model has shown that rising scale and the rising demand for skill intensive output stemming from rising incomes are quantitatively important and, overall, the model can explain the observed increase in the service share.

4 Extensions

In the benchmark model, much of the action comes through the rising opportunity cost of home production, i.e., the rising skill premium. In evaluating the robustness of these results, it is critical to investigate whether the implications in the benchmark model hold up in a model with multiple person households, where the opportunity cost of home production may not be the skill premium because one worker can specialize in market production. To address this question, we examine a model with gender-specific differences in home production and married households, which leads to a mechanism for household specialization in Section 4.1. Two additional important channels for services growth came from rising productivity and rising educational attainment. In Section 4.2, we conjecture that the two might be complementary if education increases productivity growth endogenously.

4.1 Married Household Model

It is empirically interesting to disaggregate labor by gender and marital status. According to the Current Population Survey, in 1965 about 12 percent of the population aged 21 to 65 were single women (or widows). By 2010, single women constituted 21 percent of the population. In addition, the market work hours of single females is about 80 percent of the market work hours of their male counterparts during the same period, according to the American Time Use Survey. Moreover, during the same period, the market work hours of married females relative to the market work hours of their male counterparts increased from 0.29 to 0.63, which may be in part explained by the increase in the fraction of married females with high school or

college education. Hence, a greater proportion of single women and the increase in skill intensity among married women could potentially explain a good portion of the increase in the service economy. On the other hand, the existence of married households themselves may weaken the impact of a rising skill premium on the demand for services, since households can specialize.

The production/technology side of the married household model is identical to the benchmark model presented in the previous section, so we only explain the household side of the married household model. There are three types of households in the married household model: single women, single men, and married couples.

As before, each type of household is infinitesimally-lived and they differ in their cost of acquiring skills, $\theta \in [0, 1]$, where $\theta \sim F(\theta)$. We assume that the fractions of each type of household in the overall population are exogenous. We implicitly assume perfect assortative matching among spouses in married couples, which is clearly an abstraction.

Single male and single female households are identical to households in the previous section, except that females have higher productivity in home production, $A_f > A_m$.

A single household solves the identical problem as before, except that home production is gender specific. In equilibrium, for a given ability and education level, single females will spend more time in home production. Because education only pays off in market labor, women may endogenously choose lower education levels even if they have the same inherent ability, θ .

Married couple households, however, require 2 units before satiation, and the consumption, schooling, market labor, and home production decisions are jointly determined between the husband and wife.¹³ In addition, we allow the possibility for the home production of a married couple to economize on intermediate goods

¹³Formally:

$$\tilde{u}(C, H) = \int_0^{+\infty} [H(z) + \nu(1 - H(z))] \cdot C(z) \cdot Q \cdot dz. \quad (8)$$

The additional parameter Q equals 1 if it is a single-person household's preference function and 2 if it is a married couple's preference function.

relative to that of single households. We introduce $n_c > 1$ to parameterize this. The couple's simplified problem becomes:

$$\begin{aligned}
\max_{\tilde{z}_{ee} \leq z_{ee} \leq \bar{z}_{ee}, e_m, e_f \in \{l, h\}} V^{e_m e_f}(\theta) &= 2(1 - \nu)z_{ee} + 2\nu\bar{z}_{ee} \\
& \text{s.t.} \\
\frac{2}{n_c} \int_0^{z_{ee}} p_G(z, t) dz + 2 \int_{\tilde{z}_{ee}}^{\bar{z}_{ee}} p_S(z, t) dz &= w_{e_m}(1 - t_m - \theta I(e_m)) + w_{e_f}(1 - t_f - \theta I(e_f)) \\
t_m &= 2 \int_{\tilde{z}_{ee}}^{z_{ee}} \frac{1}{A_S(t)A_l(z)A_m} dz \geq 0 \\
t_f &= 2 \int_0^{\tilde{z}_{ee}} \frac{1}{A_S(t)A_l(z)A_f} dz \geq 0 \\
1 - t_m - \theta I(e_m) &\geq 0, \quad 1 - t_f - \theta I(e_f) \geq 0, \quad z_{ee} - \tilde{z}_{ee} \geq 0. \quad (9)
\end{aligned}$$

We denote the individual education choices of the husband and wife as e_m and e_f , and their home production times t_m and t_f , respectively.

Since the wife has the (comparative and absolute) advantage in home production, she will spend relatively more time at home, while the husband will supply relatively more labor to the market, $t_f > t_m$.¹⁴ Thus, time spent in home production will vary across wives and husbands. Education decisions can differ between the wife and husband. Again, wives may acquire less education, even though they have the same ability θ as their husband because education only pays off for market labor. This force will be stronger among married couples because of the tendency to specialize disproportionately in either home (wives) or market (husbands) production, and it will be stronger the greater is the comparative advantage. We provide a formal statements and characterization of the equilibrium, including proofs of these claims, in the appendix.

¹⁴For simplicity, we define a threshold \tilde{z} with the innocuous assumption that the wife performs all home production below \tilde{z} , and the husband performs all home production between \tilde{z} and \underline{z} . The formulation is equivalent if we define a threshold \tilde{z} such that the husband performs all home production below \tilde{z} , and the wife performs all home production between \tilde{z} and \underline{z} .

The calibration of the model is very similar as well, except that we have added three additional parameters: the relative productivity of men and women in home production, A_m and A_f , respectively; and the number of manufactured goods required per unit of services in married couples, n_c . We use A_m and A_f to match the initial relative market work hours of (1) married women to married men and (2) single women to single men. We choose n_c to match the relative market supply of labor of married women to single women. Details are also relegated to the appendix.

Table 6 shows the models predictions for the service sector growth. The large quantitative magnitude of the channels holds up to introducing the possibility of specialization in two-person households, with the model now predicts an increase of 18.7 percentage points in the service share, which slightly overpredicts the growth in services by 2.4 percentage points. The lower panel examines the decompositions for the cases in which we turn individual factors off. The same essential patterns emerge in this married household model, with the two channels of income effects increasing demand for skill-intensive output and increasing scale combining to account for roughly the same amount of service share increase as SBTC. Overall, the magnitudes of these channels are somewhat smaller in the married household model. The demographic change in household composition contributes 2.9 percentage points to the increase; and when all factors are put together, they overpredict the growth of services. In the married household model, married couples have higher rates of home production and consume smaller shares of market services. The decline in the importance of married households in the population therefore increases the relative size of the service sector through a compositional effect.

Counterfactuals with fixed educational decision or fixed relative wages for men, women, married women and single women (presented in the appendix) highlight the importance of female labor supply for the service share. Keeping the educational decisions or skill premium for men fixed has little impact on the service share (less than 3 percentage points total). The effects for women, however, are more than twice as strong, and both educational choices and the skill premium play some role. Female labor supply decisions are clearly disproportionately important in understanding service growth, and we find a stronger impact for married women's decisions.

In the data, the female market labor supply rose by 37.8 percentage points. Table 7 examines the model’s predictions along this front. The factors we have modeled can account for an 24.8 percentage point increase along this dimension. Recall that we merely targeted the initial relative market work hour ratios in the economy, so the features of the model endogenously driving service economy growth also endogenously explain 66 ($=24.8/37.8$) percent of the observed growth in female market labor supply. The lower panel decomposes the different exogenous factors in the model that drive this endogenous increase. The two forces of higher productivity – demand for skill-intensive services and larger efficient scale of services – are both important and together explain slightly more than 100 percent of the increase. As it did with the service share increase, SBTC again explains roughly the same amount of the female labor supply growth as the two forces. On the other hand, manufacturing-biased technical change now has a large negative effect on female labor supply by lowering the cost of the intermediate goods used intensively in home production, and this offsets more than 100 percent of the increase. Finally, demographic changes lead to an increase in female labor supply through a pure composition effect: Single females have higher labor supply than married females, so the falling proportion of married couples increases overall female labor supply.

As noted, the period we cover is a time of great changes in the family. Our married household model takes these demographic changes as exogenously given. That is, we model neither marriage nor fertility decisions, which are certainly related to specialization and human capital decisions. A recent survey of macro literature on marriage is provided in Greenwood, Guner and Vandenbroucke (forthcoming). Various mechanisms that possibly trigger demographic transitions have been proposed and examined. For example, Ehrlich and Kim (2007) show that an increase in social security tax can shift down the entire dynamic path of family formation and fertility. Galor and Weil (1996, 2000) argue that the main driver of the demographic transition in fertility is the rising demand for human capital. In either case, demographic transitions can lead to increased human capital (Becker et al, 1990). In principle, endogenizing our demographic transition along these lines could generate a feedback effect, contributing to the growth of female labor supply and more importantly to

the overall economic growth. Incorporating such important dynamic feedback effects into a model of structural change is a fertile area for future research.

4.2 Feedback from Education to Technology

In the benchmark model, technological progress increases the demand for high skilled labor, either because skilled labor has a comparative advantage in producing more complex goods or because technological progress is skill biased. It is natural to consider the reverse feedback, in which high skilled labor leads to technological change. This mechanism is at the core of the endogenous growth literature (Lucas, 1988; Romer, 1990). In this Section we briefly discuss an extension of the model where human capital accumulation is the engine of productivity growth.

Following Lucas (1988), we can enrich the model by assuming that there is an external effect from the average skills of the population on the growth rate of productivity. In our simple model with infinitesimally-lived agents, the average skills of the population at any date is a function of the marginal educated type $\hat{\theta}(t)$. In general, the external effects could be sector-specific, rationalizing the differential productivity growth across sectors. These assumptions are captured by the following law of motion for the productivity of the manufacturing good and service sectors

$$\dot{A}_G(t) = A_G(t) \gamma_G(\hat{\theta}(t))$$

and

$$\dot{A}_S(t) = A_S(t) \gamma_S(\hat{\theta}(t)).$$

where $\gamma_G(\hat{\theta})$ and $\gamma_S(\hat{\theta})$ are increasing functions of the marginal educated type $\hat{\theta}$.

Alternatively, we could derive the reduced form relationships between the supply of skills and productivity growth by assuming that the production of new knowledge is intensive in skilled labor, as in Romer (1990).

The analysis of this feedback loop is straightforward in principle. The increase in productivity from higher educational attainment would lead to a further increase in the service sector and the demand for high-skilled labor through the channels we

emphasize. Moreover, if the elasticity of the γ_S function with respect to $\hat{\theta}$ exceeds that of the γ_G function, productivity growth would be service sector biased, which would further increase the demand for high-skilled labor. The opposite would be the case if the goods sector productivity responds more than the service sector productivity.

5 Conclusion

We have shown that a model with home production and market services that vary in their skill intensity is a quantitatively plausible explanation for the observed growth in the share of services in the United States between 1965 and 2010. In particular, the rising scale of services, rising demand for skill-intensive output stemming from income effects, and skill-biased technical change all play quantitatively important roles in the growth of services. These latter two manifest themselves largely through increases in the skill premium and the fraction of the population who are high-skilled.

These results are robust to extending the model to allow for married couples, specialization, and gender-specific labor supply. However, in this model the falling proportion of married couples in the population and, to a lesser extent, the endogenous increase in female labor supply, especially among married women, also play quantitatively important roles. Still, a caveat is that the model only explains about two-thirds of the observed increase in female labor supply, including a small share of the increased labor supply of married women. These results may presumably be driven by forces outside of the model and are promising avenues for future research.

References

- [1] Acemoglu, D. and V. Guerrieri. (2008). “Capital Deepening and Non-balanced Economic Growth,” *Journal of Political Economy*, 116(3), pp. 467-498.

- [2] Baumol, W. J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, 57(3), pp. 415-26.
- [3] Becker, G., K. Murphy and R. Tamura (1990). "Human Capital, Fertility, and Economic Growth," *Journal of Political Economy*, 98 (5), pp. 512-537.
- [4] Boppart, T. (2014). "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82 (November), pp. 2167-2196.
- [5] Buera, F. and J. Kaboski. (2009). "Can Traditional Theories of Structural Change Fit The Data?" *Journal of the European Economic Association*, 7(2-3), pp. 469-477.
- [6] Buera, F. and J. Kaboski. (2012a). "The Rise of the Service Economy," *American Economic Review*, 102 (October 2012), pp. 2540-69.
- [7] Buera, F. and J. Kaboski (2012b). "Scale and the Origins of Structural Change," *Journal of Economic Theory*, 147 (March 2012), pp. 684-712.
- [8] Buera, F., J. Kaboski and R. Rogerson. (2015). "Skill Biased Structural Change," NBER Working Paper No. 21165.
- [9] Comin, D., D. Lashkari, and M. Mestieri. (2015). "Structural Change With Long Run Income and Price Effects," working paper.
- [10] Duarte, M. and D. Restuccia. (2010). "The Role of the Structural Transformation in Aggregate Productivity," *Quarterly Journal of Economics*, 125(1), pp. 129-173.
- [11] Echevarria, C. (1997). "Changes in Sectoral Composition Associated with Economic Growth," *International Economic Review*, 1997, 38(2), pp. 431-452.

- [12] Ehrlich, I. and J. Kim. (2007). “Social Security and Demographic Trends: Theory and Evidence from the International Experience,” *Review of Economic Dynamics* 10, pp. 55-77.
- [13] Galor, O. and D. Weil (1996). “The Gender Gap, Fertility, and Growth,” *American Economic Review*, 86 (3), pp. 374-387.
- [14] Galor, O. and D. Weil (2000). “Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond,” *American Economic Review*, 90 (5), pp. 806-828.
- [15] Goldin, C. (2006). “The Quiet Revolution that Transformed Women’s Employment, Education, and Family,” 2006 Ely Lecture, American Economic Association Meetings, Boston MA.
- [16] Goldin, C., F. Lawrence, and I. Kuziemko. (2006). “The Homecoming of American College Women: the Reversal of the College Gender Gap,” *Journal of Economic Perspectives*, vol. 20, number 4, pp. 133-156.
- [17] Gollin, D., S. Parente, and R. Rogerson. (2002). “The Role of Agriculture in Development,” *American Economic Review*, 92(2), Papers and Proceedings, pp. 160-164.
- [18] Greenwood, J., A. Seshadri, and M. Yorukoglu. (2005). “Engines of Liberation,” *Review of Economic Studies* 72, pp. 109-133.
- [19] Greenwood, J., N. Guner, and G. Vandenbroucke. (forthcoming). “Family Economics Writ Large,” *Journal of Economic Literature*.
- [20] Herrendorf, B., R. Rogerson, and A. Valentinyi. (2013). “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103(7), pp. 2752-2789.
- [21] Jones, L.E., R.E. Manuelli, and E.R. McGrattan. (2003). “Why are Married Women Working So Much?” Staff Report 317, Federal Reserve Bank of Minneapolis.

- [22] Jorgenson, D. and K. Stiroh. (2000). "Raising the Speed Limit: U.S. Economic Growth in the Information Age," *Brookings Papers on Economic Activity* 1, pp. 125-211.
- [23] Katz, L.F. and Murphy K.M. (1992). "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107(1), pp. 35-78.
- [24] Kongsamut, P., S. Rebelo, and D. Xie. (2001). "Beyond Balanced Growth," *Review of Economic Studies*, 2001, 68, pp. 869-882.
- [25] Lee, D. and L.K. Wolpin. (2006). "Intersectoral Labor Mobility and the Growth of the Service Sector," *Econometrica*, 74(1), pp. 1-46.
- [26] Lucas, R. E., Jr. (1988). "On the Mechanics of Economic Development," *Journal of Monetary Economics* 22 (July), pp. 3-42.
- [27] Mulligan, C. B. and Y. Rubinstein. (2008). "Selection, Investment, and Women's Relative Wages over Time," *Quarterly Journal of Economics*, 2008, 123(3), pp. 1061-110.
- [28] Ngai, L. R. and C.A. Pissarides. (2007). "Structural Change In A Multisector Model of Growth," *American Economic Review*, 97 (1). pp. 429-443.
- [29] Ngai, L. R. and C.A. Pissarides. (2008). "Trends in Hours and Economic Growth," *Review of Economic Dynamics*, 11(2), pp. 239-256.
- [30] Ngai, L. R. and B. Petrongolo. (2012). "Structural Transformation, Marketization and Female Employment," Mimeo London School of Economics.
- [31] Rendall, M. (2014). "The Service Sector and Female Market Work," Mimeo University of Zurich.

- [32] Romer, P. (1990). “Endogenous Technological Change,” *Journal of Political Economy* 98 (October, Part 2), pp. S71-S102.
- [33] Stokey, N. (1991). “Human Capital, Product Quality, and Growth.” *Quarterly Journal of Economics*, 1991, 106(2), pp. 587-616.

Figure 1: Growth of Value-Added Share of Services




Figure 2: Importance of Female Employment in Service Sector Growth

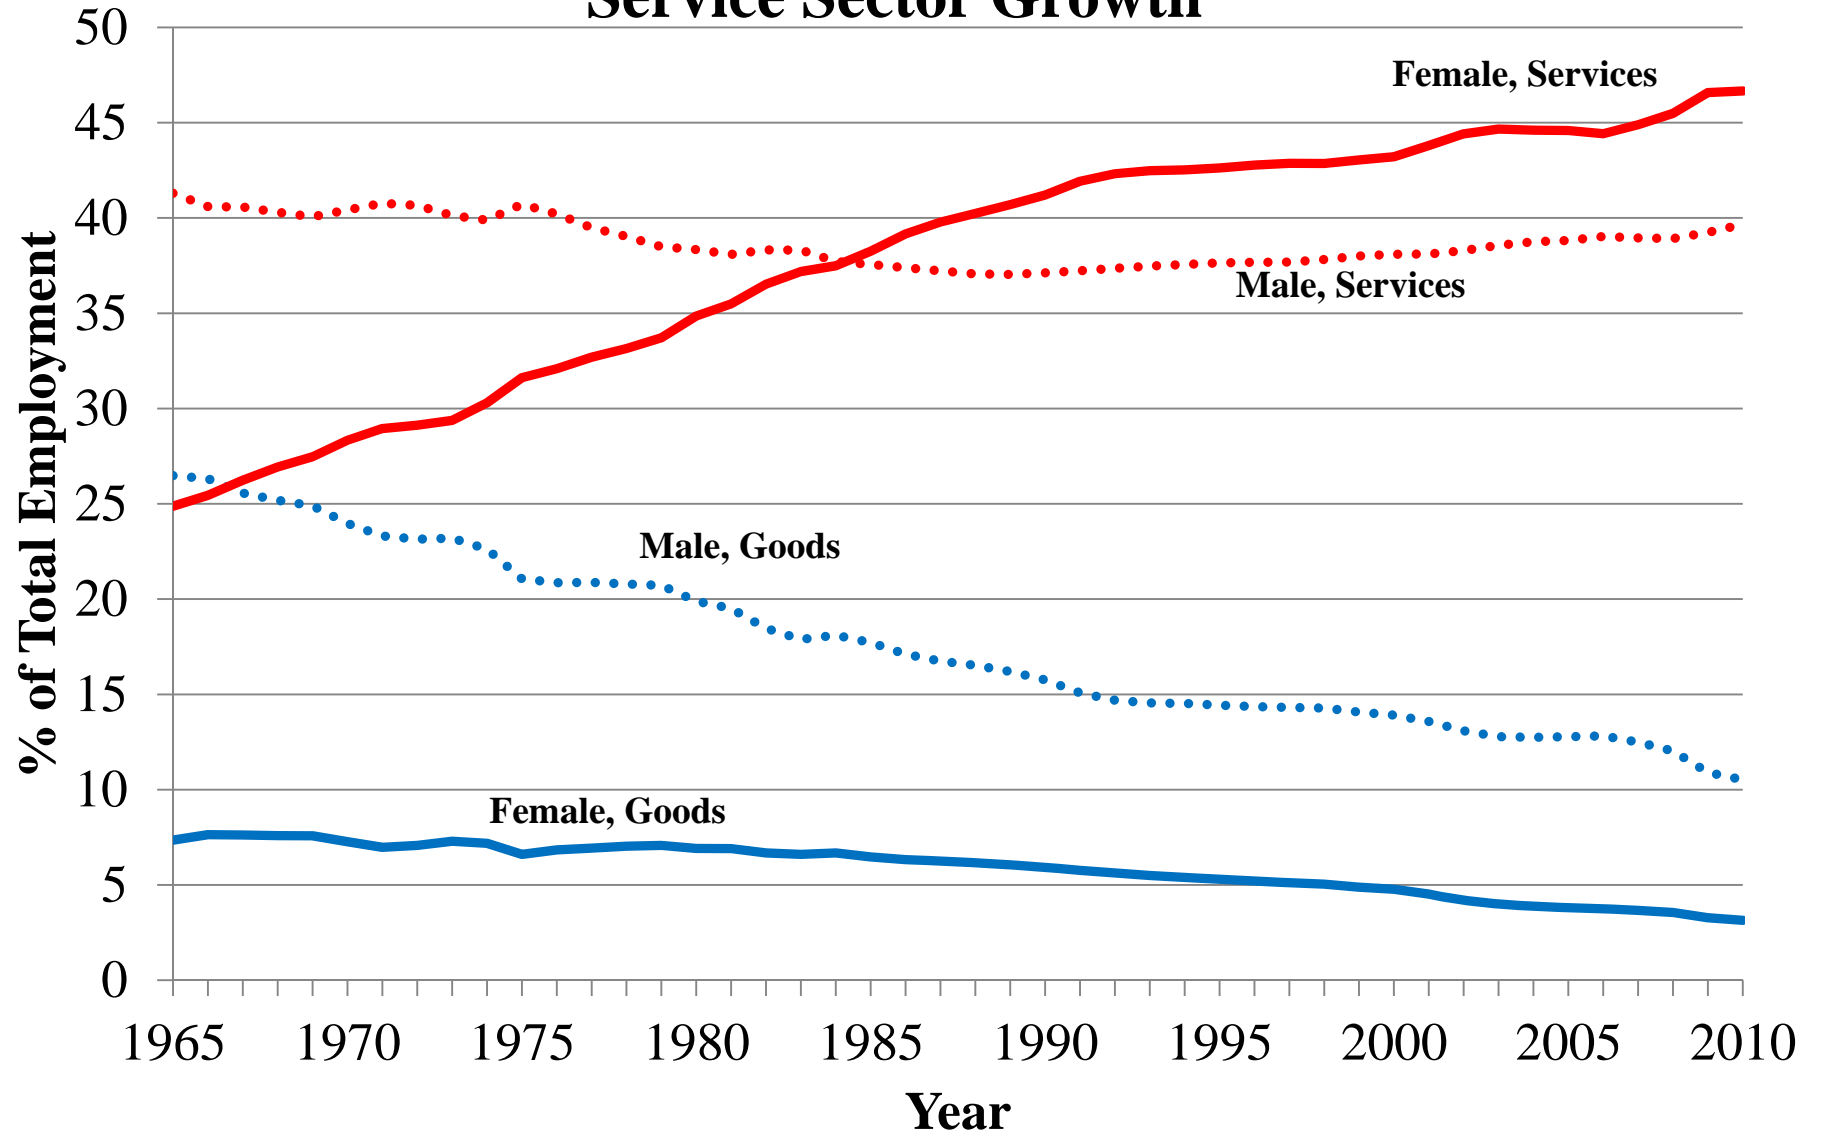


Table 1: Calibration of Benchmark Model

Moment	Data Value	Model Value	Relevant Parameter	Parameter Value
Initial Value (1965) Moments:				
Service Share	0.65	0.65	v	0.67
Intermediate Manufacturing Inputs/Value-Added	0.12	0.12	n_{1965}	8.52
Skill Premium	1.41	1.41	ϕ_0	1.39
High-Skilled Fraction of Population	0.23	0.23	Beta: a	3.61
Growth (1965-2010) Moments:				
Increase in High-Skilled Fraction of Population	0.36	0.36	Beta: b	6.51
Increase in Skill Premium	0.41	0.41	λ	0.70
Growth in Relative Market Work Hours of High to Low Skilled Population	0.03	0.03	γ_h	0.0054
Growth in Real Per Capita GDP	1.44	1.44	γ_S	0.0135
Growth in Relative Price of Services/Manufacturing	0.44	0.44	γ_G	0.0193

Table 2: Service Growth in Benchmark Model

Moment	(Current) Value-Added Service Share		Percentage Point Increase
	1965	2010	
Data	0.650	0.813	16.3
Model	0.650	0.809	15.9

Note: The model matches the data in 1965 by calibration.

Table 3: Decomposing Service Share Increase: Counterfactuals

Simulation	Percentage Point Difference	Percent of Benchmark
All Factors (Benchmark)	15.9	-
No Income Effect	5.2	32.8%
No Scale Effect	4.4	27.9%
No Income or Scale Effects	10.7	67.1%
No Skill-Biased Technical Change	11.6	73.2%
No Sector-Biased Technical Change	-1.8	-11.2%
Only Income Effect	3.8	23.9%
Only Scale Effect	6.3	39.4%
Only Income and Scale Effects	9.7	61.2%
Only Skill-Biased Technical Change	9.4	58.9%
Only Sector-Biased Technical Change	-11.0	-69.4%

Note: The "Percentage Point Difference" is the total percentage point increase in the current-value, value-added service share for "All Factors" and the "Only" simulations in the lower panel. For the "No" simulations, it is the difference between the "All Factors" service share increase and the service share increase under the specified "No" simulation.

Table 4: Effect of Skill Premium and Educational Choices

Counterfactual Percentage Point Increase in Service Share	Skill Premium	
	Fixed at 1965 Value	Benchmark Skill Premium
Fixed at 1965 Choices	3.2	6.8
Benchmark Educational Choices	10.2	15.9

Note: The simulations show the increase in the current-value value-added service share between 1965 and 2010 under counterfactual simulations when fixing the skill premium and educational choices at either the benchmark 2010 values or the initial 1965 values. Budget constraints are imposed, but market clearing conditions are not. The simulations are thus aggregations of partial equilibrium household decisions.

Table 5: Counterfactual Impacts on Other Changes Over Time

Moment	ΔFraction of High Skilled	ΔSkill Premium	Growth in Real GDP per Capita	Growth in Relative Price of Services
All Factors/Data	0.36	0.41	1.44	0.44
Income Effect	0.12	0.14	1.38	0.06
Scale Effect	0.02	0.00	0.06	0.13
Income and Scale Effects	0.14	0.14	1.42	0.18
Skill-Biased Technical Change	0.26	0.27	0.31	-0.02
Sector-Biased Technical Change	-0.04	0.02	-0.13	0.32

Note: These effects are calculated as the difference between the increase or growth of the specified moment in the model with all factors and the data moment in the simulation where the specified effect is turned off in 2010. Thus, these are comparable to the "No" effects in Table 3.

Table 6: Service Share Growth and Decomposition in Extended Model

Simulation	Percentage Point Increase in Service Share	
	Percentage Point Difference	Percent of Benchmark
Full Simulation	18.7	-
Income Effect	3.2	17.3%
Scale Effect	4.0	21.4%
Income and Scale Effects	7.7	41.2%
Skill-Biased Technical Change	7.0	37.7%
Sector-Biased Technical Change	-0.2	-1.0%
Demographic Change	2.9	15.7%

Note: These effects are calculated as the difference between the current-value, value-added service share increase in the extended model with all factors and the service share increase in the simulation where the specified effect is turned off in 2010. Thus, these are comparable to the "No" effects in Table 3.

Table 7: Decomposing Female Market Labor Increase in Extended Model: Counterfactuals

Simulation	Percentage Point Increase in Female Market Labor Supply	
	Percentage Point Difference	Percent of Benchmark
All Factors (Benchmark)	24.8	-
Income Effect	12.0	48.5%
Scale Effect	11.3	45.7%
Income and Scale Effects	26.0	105.2%
Skill-Biased Technical Change	24.0	96.8%
Sector-Biased Technical Change	-31.3	-126.6%
Demographic Change	19.3	78.0%

Note: The comparable moment in the data is 37.8 percentage points. The "Percentage Point Difference" is the total percentage point increase in the Female Market Labor Supply for "All Factors", while each specified "Effect" is the difference between the "All Factors" increase and the increase when the specified effect is turned off. That is, these are the female market labor supply analogs to the "No" effects in Table 3.

A Characterization of a Married Couple's Problem

In the model with married and single households, a competitive equilibrium consists of $w(t)$, $\widehat{\theta}_m$, $\widehat{\theta}_f$, $\widehat{\theta}_1$, and $\widehat{\theta}_2$, $\widehat{z}(t)$, the price functions $p_G(z, t)$ and $p_S(z, t)$, and the consumption thresholds $(\underline{z}_{l,m}(t), \underline{z}_{l,f}(t), \bar{z}_{l,m}(t), \bar{z}_{l,f}(t), \underline{z}_{h,m}(\theta, t), \underline{z}_{h,m}(\theta, t), \bar{z}_{h,m}(\theta, t), \bar{z}_{h,f}(\theta, t), \underline{z}_{ll}(t), \bar{z}_{ll}(t), \underline{z}_{hl}(\theta, t), \bar{z}_{hl}(\theta, t), \underline{z}_{hh}(\theta, t), \bar{z}_{hh}(\theta, t))$. The model can be solved in two steps recursively in a fashion very similar to in the benchmark model.

We state two formal propositions:

Proposition 1 *Given $A_m < A_f$, the schooling choice of lh (low-skilled husband, high-skilled wife) will always be dominated by hl (high-skilled husband, low-skilled wife).*

Proof. [Proof of Proposition 1] The budget constraint for a schooling choice, lh , is denoted as follows:

$$\frac{2}{n_c} \int_0^{\underline{z}_{lh}} p_G(z, t) dz + 2 \int_{\underline{z}_{lh}}^{\bar{z}_{lh}} p_S(z, t) dz = (1 - 2 \int_{\underline{z}_{lh}}^{\bar{z}_{lh}} \frac{e^{-\gamma st} z}{A_m} dz) + w(1 - 2 \int_0^{\bar{z}_{lh}} \frac{e^{-\gamma st} z}{A_f} dz - \theta). \quad (10)$$

The budget constraint for a schooling choice, hl , is denoted as follows:

$$\frac{2}{n_c} \int_0^{\underline{z}_{hl}} p_G(z, t) dz + 2 \int_{\underline{z}_{hl}}^{\bar{z}_{hl}} p_S(z, t) dz = w(1 - 2 \int_{\underline{z}_{hl}}^{\bar{z}_{hl}} \frac{e^{-\gamma st} z}{A_m} dz - \theta) + (1 - 2 \int_0^{\bar{z}_{hl}} \frac{e^{-\gamma st} z}{A_f} dz). \quad (11)$$

Assume \underline{z}^* , \bar{z}^* and \widetilde{z}^* are the optimal allocations for the schooling choice, lh .

Case One: $\frac{1}{A_m} \leq \frac{w}{A_f}$

Under lh , the wife faces a higher opportunity cost in home production, which

implies the following inequality:

$$\begin{aligned} 2 \int_{\tilde{z}^*}^{\underline{z}^*} \frac{e^{-\gamma s t z}}{A_m} &> 2 \int_0^{\tilde{z}^*} \frac{e^{-\gamma s t z}}{A_f} dz \\ \frac{(\underline{z}^*)^2 - (\tilde{z}^*)^2}{A_m} &> \frac{(\tilde{z}^*)^2}{A_f}. \end{aligned} \quad (12)$$

Next, we define \tilde{z}' as follows:

$$2 \int_{\tilde{z}'}^{\underline{z}^*} \frac{e^{-\gamma s t z}}{A_m} = 2 \int_0^{\tilde{z}^*} \frac{e^{-\gamma s t z}}{A_f} dz. \quad (13)$$

Equation (13) can be simplified as follows:

$$\frac{(\tilde{z}')^2}{A_m} = \frac{(\underline{z}^*)^2}{A_m} - \frac{(\tilde{z}^*)^2}{A_f}. \quad (14)$$

It is easy to show that $\tilde{z}' > 0$ and $1 - 2 \int_0^{\tilde{z}'} \frac{e^{-\gamma s t z}}{A_f} dz > 0$ given $\underline{z}^* \geq \tilde{z}^*$, $A_f > A_m$, $1 - 2 \int_{\tilde{z}^*}^{\underline{z}^*} \frac{e^{-\gamma s t z}}{A_m} dz \geq 0$ and the inequality (12). Next, we subtract the RHS of (10) from the RHS of (11), with $(\underline{z}^*, \tilde{z}^*)$ for the schooling choice lh and $(\underline{z}^*, \tilde{z}')$ for the schooling

choice hl :

$$\begin{aligned}
& (1 - 2 \int_0^{\tilde{z}'} \frac{e^{-\gamma st} z}{A_f} dz) - (1 - 2 \int_{\tilde{z}^*}^{\tilde{z}^*} \frac{e^{-\gamma st} z}{A_m} dz) \\
&= e^{-\gamma st} \frac{(\tilde{z}^*)^2 - (\tilde{z}^*)^2}{A_m} - \frac{e^{-\gamma st} (\tilde{z}')^2}{A_f} \\
&= e^{-\gamma st} \frac{(\tilde{z}^*)^2 - (\tilde{z}^*)^2}{A_m} - e^{-\gamma st} \frac{A_m}{A_f} \left(\frac{(\tilde{z}^*)^2}{A_m} - \frac{(\tilde{z}^*)^2}{A_f} \right) \\
&= e^{-\gamma st} ((\tilde{z}^*)^2 - (\tilde{z}^*)^2) \left(\frac{1}{A_m} - \frac{1}{A_f} \right) - e^{-\gamma st} (\tilde{z}^*)^2 \left(\frac{1}{A_f} - \frac{A_m}{A_f} \frac{1}{A_f} \right) \\
&> e^{-\gamma st} (\tilde{z}^*)^2 \frac{A_m}{A_f} \left(\frac{1}{A_m} - \frac{1}{A_f} \right) - e^{-\gamma st} (\tilde{z}^*)^2 \left(\frac{1}{A_f} - \frac{A_m}{A_f} \frac{1}{A_f} \right) \\
&= 0.
\end{aligned} \tag{15}$$

Line 3 of (15) follows from (14), and Line 5 of (15) follows from the inequality (12). The inequality (15) shows that a married couple can consume $(\underline{z}^*, \bar{z}^* + q; q > 0)$ if their schooling choice is (hl) instead of (lh) .

Case Two: $\frac{1}{A_m} > \frac{w}{A_f}$

The wife will always have a comparative advantage in home production regardless of schooling choices. It is easy to show that the schooling choice (hl) dominates the schooling choice (lh) if $1 - 2 \int_0^{\tilde{z}^*} \frac{e^{-\gamma st} z}{A_f} dz \geq 0$, or equivalently $\underline{z}^* \leq \sqrt{e^{\gamma st} A_f}$.

We still need to show that (hl) dominates (lh) when $\underline{z}^* > \sqrt{e^{\gamma st} A_f}$ in the following steps. First, under (lh) , the time constraint of the high-skilled wife will be binding, which implies $1 - \theta = \frac{e^{-\gamma st}}{A_f} (\tilde{z}^*)^2$. Then, the RHS of (10) can be simplified as follows:

$$RHS_{lh}(\underline{z}^*, \tilde{z}^*) = 1 - 2 \int_{\tilde{z}^*}^{\tilde{z}^*} \frac{e^{-\gamma st} z}{A_m} dz = 1 - \frac{e^{-\gamma st}}{A_m} (\tilde{z}^*)^2 + \frac{A_f}{A_m} (1 - \theta). \tag{16}$$

Next, we define \tilde{z}' as $\sqrt{e^{\gamma st} A_f}$. With $\underline{z}_{hl} = \underline{z}^*$ and $\tilde{z}_{hl} = \tilde{z}'$, the RHS of (11) can be simplified as follows:

$$RHS_{hl}(\underline{z}^*, \tilde{z}) = w(1 - 2 \int_{\tilde{z}}^{\underline{z}^*} \frac{e^{-\gamma st} z}{A_m} dz - \theta) = w \left(1 - \frac{e^{-\gamma st}}{A_m} (\underline{z}^*)^2 + \frac{A_f}{A_m} - \theta \right). \quad (17)$$

By subtracting $RHS_{lh}(\underline{z}^*, \tilde{z}^*)$ from $RHS_{hl}(\underline{z}^*, \tilde{z})$, we obtain:

$$RHS_{hl}(\underline{z}^*, \tilde{z}) - RHS_{lh}(\underline{z}^*, \tilde{z}^*) = (w-1) \left(1 - \frac{e^{-\gamma st}}{A_m} (\underline{z}^*)^2 + \frac{A_f}{A_m} \right) + \theta \left(\frac{A_f}{A_m} - w \right) > 0. \quad (18)$$

The inequality (18) shows that a married couple can consume $(\underline{z}^*, \bar{z}^* + q; q > 0)$ if their schooling choice is (hl) instead of (lh) . Therefore, the schooling choice (hl) always dominates the schooling choice (lh) . ■

Proposition 2 *Given $A_f > A_m$ and $w > 1$, at least one spouse will fully specialize. If the husband works both at home and in the market, his wife will fully specialize in home production. If the wife works both at home and in the market, her husband will fully specialize in market production.*

Proof. [Proof of Proposition 2] Given a schooling choice, the Kuhn-Tucker conditions that characterize the optimum, \underline{z}^* , \bar{z}^* and \tilde{z}^* are

$$\mu \left(\frac{e^{-\gamma st} \underline{z}^*}{A_m} w_{e_m} + \frac{1}{n_c} p_G(\underline{z}^*, t) - p_S(\underline{z}^*, t) \right) + \eta_1 \frac{e^{-\gamma st} \underline{z}^*}{A_m} - \eta_3 = 1 - \nu, \quad (19)$$

$$\nu = \mu p_S(\bar{z}^*, t), \quad (20)$$

$$\mu w_{e_m} \frac{e^{-\gamma st} \tilde{z}^*}{A_m} - \mu w_{e_f} \frac{e^{-\gamma st} \tilde{z}^*}{A_f} = \eta_3 + \eta_2 \frac{e^{-\gamma st} \tilde{z}^*}{A_f} - \eta_1 \frac{e^{-\gamma st} \tilde{z}^*}{A_m}, \quad (21)$$

$$1 - t_m - \theta I(e_m) \geq 0, \eta_1 \geq 0, (1 - t_m - \theta I(e_m)) \eta_1 = 0, \quad (22)$$

$$1 - t_f - \theta I(e_f) \geq 0, \eta_2 \geq 0, (1 - t_f - \theta I(e_f)) \eta_2 = 0, \quad (23)$$

$$2(\underline{z}^* - \tilde{z}^*) \geq 0, \eta_3 \geq 0, 2(\underline{z}^* - \tilde{z}^*) \eta_3 = 0, \quad (24)$$

where η_1 , η_2 and η_3 are the Kuhn-Tucker multipliers associated with the inequality constraints (22), (23) and (24), respectively, and μ is the marginal utility of income. Condition (21) characterizes the division of home production \tilde{z}^* . It is driven by the comparative advantage in home production, depending on the values of A_f and A_m . If $A_f > A_m$, the wife will have a comparative advantage in home production, and the schooling choice of (lh) will never be chosen according to Proposition 1. Then, given $A_f > A_m$ and $w > 1$, the LHS of (21) should be positive unless \tilde{z}^* is equal to zero.

However, it is never optimal to choose $\tilde{z}^* = 0$, which can be proved by contradiction. Suppose that $\tilde{z}^* = 0$. If $\underline{z}^* > 0$, a married couple can easily improve the outcome by setting $t_m = 0$ and $t_f = 2 \int_0^{\underline{z}^*} \frac{e^{-\gamma s t} z}{A_f} dz$, so the optimal value of \underline{z} under $\tilde{z}^* = 0$ should be zero, which implies $\eta_1 = 0$. Then, in order to satisfy Condition (21), η_3 has to be zero as well. Given $\eta_1 = 0$ and $\eta_3 = 0$, Condition (19) can be simplified as follows:

$$\mu \left(\frac{e^{-\gamma s t} \underline{z}^*}{A_m} w_{e_m} + \frac{1}{n_c} p_G(\underline{z}^*, t) - p_s(\underline{z}^*, t) \right) = 1 - \nu. \quad (25)$$

The LHS of (25) is zero, but the RHS of (25) is positive, which leads to a contradiction. Therefore, \tilde{z}^* , \underline{z}^* , and the LHS of (21) should be positive.

Next, we prove that $t_f > t_m$ again by contradiction. Suppose that $t_f \leq t_m$. Since $\tilde{z}^* > 0$, it is easy to show that $t_m \geq t_f > 0$, which implies $\eta_3 = 0$. Then, Condition (21) will require η_2 to be positive. In other words, $1 - t_f - \theta I(e_f) = 0$. Given that the schooling choice of (lh) is never chosen and $t_f \leq t_m$, it will always violate either the budget constraint (LHS of the budget constraint $>$ RHS of the budget constraint $= 0$) or the husband's time constraint ($1 - t_m - \theta I(e_m) < 0$), which leads to a contradiction. Therefore, $t_f > t_m$.

If $t_m > 0$, then $\eta_3 = 0$. In order to keep the RHS of (21) to be positive, η_2 has to be positive, which implies that the wife will not work in the market. If $1 - t_f - \theta I(e_f) > 0$, then $\eta_2 = 0$. In order to keep the RHS of (21) to be positive, η_3 has to be positive, which implies that the husband will not work at home. ■

B Additional Results of Married Household Model

The calibration involves three additional parameters: the relative productivity of men and women in home production, A_m and A_f , respectively; and the number of manufactured goods required per unit of services in married couples, n_c . We calibrate the model using the same approach as in the benchmark model, but adding the following three target moments in 1965. We use A_m and A_f to match the initial relative market work hours of (1) married women to married men and (2) single women to single men. We choose n_c to match the relative market supply of labor of married women to single women. These targets require that men are roughly one-fourth as productive as women in home production, but then women are still only two-third as productive at home relative to their productivity in the market. Finally, we require only small returns to scale coming from home production. The calibrated $n_c = 1.173$ implies that married couples purchase 1.7 manufactured goods to produce two units of services. They therefore do not economize as well on manufactured goods as the market does.

In addition, we change the composition of household types (married, single) to match the changes in their composition in the data. The proportion of people in married couples fell from 80 percent in 1965 to just 59 percent in 2010. Correspondingly, the proportion of single households doubled from 20 to 41 percent.

Table A.1 summarizes the calibration for the extended model. As with the benchmark case, we are able to hit all the data moments. Indeed, comparing Table A.1 with Table 1, the same patterns hold, with a rightward-skewed θ distribution and similar productivity parameters.

Table A.2 examines the importance of the endogenous education and skill premium responses in the married household model, results we refer to in Section 4.1. This analysis is analogous to the exercise in Table 4, where we aggregate households' problems at the benchmark equilibrium prices, but keep either the relative wage fixed, schooling decisions fixed, or both fixed at their 1965 values. In Table A.2, we look at the impacts of doing this separately for various subpopulations. The impact

of the skill premium for men has a relatively small effect on the service share, while the impact of men's educational choices are negligible. In total, the two together account for less than 3 percentage points. The effects for women, however, are more than twice as strong, and both educational choices and the skill premium play some role. Female labor supply decisions are clearly disproportionately important in understanding service growth. Examining this more closely, when we look at single women and married women separately in the lower panels of Table A.2, we see that the impact on married women is somewhat larger.

Table A.1: Calibration of Extended Model

Moment	Data Value	Model Value	Relevant Parameter	Parameter Value
Initial Value (1965) Moments:				
Service Share	0.65	0.65	v	0.581
Intermediate Manufacturing Inputs/Value-Added	0.12	0.12	n_{1965}	8.518
Skill Premium	1.41	1.41	Φ_0	1.430
High-Skilled Fraction of Population	0.23	0.23	Beta: a	3.596
Relative Market Work Hours (Married Female/Married Male)	0.29	0.29	A_f	0.678
Relative Market Work Hours (Single Female/Single Male)	0.80	0.80	A_m	0.165
Relative Market Work Hours (Married Women/Single Women)	0.38	0.38	n_c	1.173
Growth (1965-2010) Moments:				
Increase in High-Skilled Fraction of Population	0.36	0.36	Beta: b	6.889
Increase in Skill Premium	0.41	0.41	λ	0.647
Growth in Relative Market Work Hours of High to Low Skilled Population	0.03	0.03	γ_h	0.0059
Growth in Real Per Capita GDP	1.44	1.44	γ_S	0.0112
Growth in Relative Price of Services/Manufacturing	0.44	0.44	γ_G	0.0171

Table A.2: Effect of Skill Premium and Educational Choices on Service Share

	Counterfactual Percentage Point Increase in Service Share	Skill Premium	
		Fixed at 1965 Value	Benchmark Skill Premium
All Men	Fixed at 1965 Choices	15.9	18.5
	Benchmark Educational Choices	15.9	18.7
All Women	Fixed at 1965 Choices	12.8	14.5
	Benchmark Educational Choices	15.9	18.7
Single Women	Fixed at 1965 Choices	14.9	17.4
	Benchmark Educational Choices	15.9	18.7
Married Women	Fixed at 1965 Choices	13.8	16.0
	Benchmark Educational Choices	15.9	18.7

Note: The simulations show the increase in the current-value value-added service share between 1965 and 2010 under counterfactual simulations when fixing the skill premium and educational choices for the stated subpopulation at either the benchmark 2010 values or the initial 1965 values. Budget constraints are satisfied, but market clearing conditions are not imposed. The simulations are thus aggregations of partial equilibrium household decisions.