Structural Change with Time to Consume

Bednar, W.L. and Pretnar, Nick

11 July 2023

Structural Change with Time to Consume

W.L. Bednar[1] * Nick Pretnar[2†]

[1]*Office of the Comptroller of the Currency*
[2]*University of California, Santa Barbara*

July 11, 2023[‡]

Abstract

Consumers allocate both income toward consumption of goods and services and off-market time toward activities using either goods or services. A model with time to consume embeds rich income effects, which has implications for the causal mechanism driving the rise in the services share of expenditure in the U.S. We estimate that consumers increasingly treat goods as luxuries relative to services, and this may result from the fact that the relative efficiency of using goods versus services, from the perspective of the consumer, has improved. Examining structural change through the model's lens, the rise in the services share of U.S. expenditure over the last 70 years is primarily attributable to the decline in the relative price of goods to services.

JEL Codes: D1, E2, N1, O3
Keywords: structural change, home production, services, income effects, growth, consumption quality

*OCC; 400 7th St. SW; Washington, DC 20219; William.Bednar@occ.treas.gov

†UCSB, LAEF; 2112 North Hall, Santa Barbara, CA 93106-9215; npretnar@ucsb.edu.

‡This research was partially funded by NSF Grant #DGE1252522. Note that the views expressed in this paper do not necessarily reflect the views of the OCC and do not establish supervisory policy, requirements, or expectations. We would like to thank Laurence Ales, Javier Birchenall, Benjamin Bridgman, Lei Fang, Alex Horenstein, Finn Kydland, Diana Mikhail, Alan Montgomery, Hakki Özdenören, Peter Rupert, Ali Shourideh, Stephen Spear, and Kieran Walsh. We would also like to thank seminar participants at the Federal Reserve Banks of Atlanta and Philadelphia, the 2022 UCSB/LAEF Growth, Development, and Structural Transformation Conference, University of Pittsburgh, University of Wuerzburg, University of Miami Herbert Business School, UCSB, and Carnegie Mellon University, as well as participants at the 2020 Society for Non-linear Dynamics and Econometrics Conference, 2019 Society for Economic Dynamics Conference in Saint Louis, Missouri, and the fall 2019 Midwest Macroeconomics Conference at Michigan State University. Additionally, we acknowledge three anonymous referees for helpful comments.

1 Introduction

In this paper, we explore the potential role of consumer time allocation in contributing to the long-run rise in the services share of expenditure in the United States. We begin by documenting several empirical facts regarding long-run trends in consumption allocation and time use in the U.S. We then consider a model that allows consumers to allocate time toward the use of both market goods and services, the purpose of which is to account for channels through which consumption allocation can respond to price and wage variation when allowing for elastic off-market time utilization in multiple different activities. With this model, we explore the forces driving the structural transformation of U.S. expenditure patterns since World War II.

In considering the role that time allocation plays in contributing to structural change, our work is closely linked to those of Bridgman (2016) and Bridgman, Duernecker, and Herrendorf (2018), but with a specific focus on U.S. outcomes. Indeed, similar to Bridgman (2016), we use a home production model to back out and study how in-home productivities associated with using different types of market products have changed over time. While it is common in the structural transformation literature to study how different growth rates in total-factor productivities (TFP) associated with the sectoral production (i.e., firms combining capital and labor to achieve output) of different types of products have driven relative price variation and changing output/expenditure shares, we are the first, to our knowledge, to consider how the *experienced* productivity, from the consumer's perspective, of *using* different kinds of products (i.e., goods versus services) provide insights into the forces driving structural change. Our results will suggest that the efficiency of consumption derived from using goods relative to the efficiency from using serivces increased along the growth path.

We contribute to the structural change literature which explores how non-homothetic preferences over market purchases generate expenditure shares that vary in income, both cross-sectionally and over time. Caselli and Coleman (2001), Kongsamut, Rebelo, and Xie (2001), Matsuyama (2009), Herrendorf, Rogerson, and Valentinyi (2013), Uy, Yi, and Zhang (2013), Boppart (2014), Kehoe, Ruhl, and Steinberg (2018), and Comin, Lashkari, and Mestieri (2021) all consider models where income effects result from non-homothetic curvature in preferences. Alternatively, we demonstrate that under certain conditions on the degree to which market inputs and off-market time are complementary, a Becker

(1965) home production model also features expenditure shares which vary in income.[1] This provides an alternative explanation for why such an income effect may be observed and why income effects, more generally, may be contributing to structural change.

Our work thus follows a suggestion in Kongsamut, Rebelo, and Xie (2001), who first noted that non-homotheticities required in preferences to match data could be attributable to the presence of home production. Additionally, Buera and Kaboski (2009) later suggested that exploring home production models may provide alternative explanations for the causes of structural change. One argument for this is that the rise of the services sector is driven by the movement of consumption toward more skill-intensive output (Buera and Kaboski 2012), because many services are natural substitutes for certain time-intensive home production activities. As our model shows, variation in wages, and thus income, can affect allocations through the interactions between time-use and consumption-expenditure decisions. Our approach thus provides a micro-founded mechanism which can reconcile both increasing cross-sectional services shares of expenditure and an increasing aggregate share along the growth path.

We proceed by outlining several empirical facts regarding the long-run evolution of expenditure and off-market time-use patterns, while also positioning our approach within the literature. Specifically, we show that while the rise in the services share from 1950-1990 can be attributed to increases in both the relative price and relative quantity of services to goods, deceleration in structural change after 1990 is almost entirely attributable to a reversal of the increase in relative quantities. We also show that since the 1970s off-market time that households spend using goods versus services has approximately doubled, which is consistent with cross-sectional evidence that lower-income households spend a smaller share of wallet but a greater share of time using services.

In Section 2, after presenting these facts regarding structural change, we then describe the model used to decompose the drivers of these long-run trends. We show that income effects can be generated by Beckerian home production models because the hourly wage is simultaneously the price of various off-market activities and directly proportional to labor income. We also show that the degree to which time and market purchases are either complementary or substitutable in different kinds of off-market activities dictates how expenditure allocations will respond to wage variation, where even a simple, nested-CES home-production and preference structure can yield rich relationships between allocations, wages, and prices.

[1]Note that Gary Becker's work with Gilbert Ghez, collected in the *The Allocation of Time and Goods Over the Life Cycle*, is often cited as the seminal text in the home production literature. The analyses in Ghez and Becker (1974) are just extensions, however, of the more general, original Becker (1965) model, which is the origin of home production modeling.

We estimate our model in Section 3 using both time-use and expenditure micro data. We use the estimated parameters to quantify the degree to which the rise in the services share can be attributed to changes in relative goods to services prices or rising wages, finding that relative price variation is mostly responsible for structural change since the 1950s. Further, we estimate that while goods were relatively inferior to services in the 1950s and 1960s, they have transitioned to become relative luxuries, helping contribute to the structural change deceleration. We discuss how our estimates suggesting that goods are relative luxuries may be partially attributable to unmeasured quality improvements which may not be embedded in aggregate price and quantity indices. Finally, we show that across the wage distribution the decline in the relative price of goods to services has benefited consumers at lower wages the most, because they consistently spend relatively more on goods than services compared to higher wage earners. In fact, we estimate that the lowest wage earners would need to have received an increase of between 8 and 11% in their 2019 wages to be indifferent to living in a world characterized by the relative goods to services price from 1959.

1.1 Empirical Facts and Background

In this section, we present several well-established empirical facts related to long-run U.S. structural change. We also describe several new facts pertaining to time use and real consumption quantities. While structural change is common across advanced economies, we focus on the U.S. economy. We require a representative panel of granular time-use micro data, which is available for the U.S. Specifically, we want a time-use dataset that contains task classifications which can be readily linked to activities associated with using market products from broad expenditure categories (i.e., goods and services). We build such a dataset by partitioning non-work tasks in the Americans' Use of Time survey (AUT, 1965-1966), the Time Use in Economic and Social Accounts survey (TUESA, 1975-1976), and the American Time Use Survey (ATUS, 2003-2021), into tasks which predominantly involve utilization of market goods versus those involving utilization of market services.[2,3] This allows us to create aggregate and disaggregated measures of how households spend their

[2]For task classifications and descriptions of the data-cleaning processes we deploy, see Technical Appendix B, which is available at the author's website: https://www.npretnar.com/research. All time-use data are available from the ICPSR at the University of Michigan.

[3]Note that we will use the AUT and TUESA datasets for illustrative purposes only. They will only feature into our discussions in this section, and we will restrict our rigorous empirical and quantitative, structural analyses to data taken from ATUS. The reasons for this distinction are two-fold: 1) AUT and TUESA do not have representation weights, but ATUS does; 2) ATUS has a more granular task-classification rubric which is more-easily linkable to expenditure data.

off-market time on different consumption activities.[4]

Granular time-use datasets, such as (but not limited to) AUT, TUESA, and ATUS, are often used to better understand how consumers allocate time toward tasks classified as market work, non-market work, and leisure. However, our endeavor is to understand how consumers allocate time toward tasks associated with different product types. In this vein our work is most closely associated to strands of the time-use literature that analyze the substitutability between expenditure and off-market time, such as Aguiar and Hurst (2005), Aguiar and Hurst (2007), Pretnar (2022), and Fang, Hannusch, and Silos (2022), though we are unique in focusing specifically on the tradeoff consumers face with regards to allocating time toward activities that utilize either market goods or services.

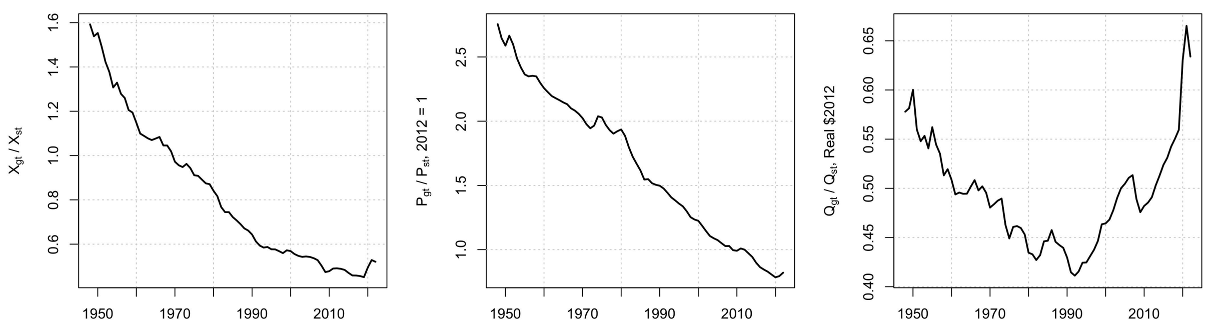

(a) Relative Expenditure (b) Relative Prices (c) Relative Quantities

Figure 1: The ratio of the aggregate nominal value of final goods to services expenditure, X_{gt}/X_{st}, is featured in (a); the relative chain-weighted price of goods to services, P_{gt}/P_{st}, where 2012 = 1 is featured in (b); in (c) we feature relative real quantities in \$2012, Q_{gt}/Q_{st}. Aggregate expenditure, price, and quantity time series depicted here run from 1948-2022.

Figure 1 shows how, on aggregate, relative goods to services expenditure, prices, and real quantities have changed over the last 70-plus years. Aggregate expenditure, price, and quantity series are constructed using NIPA Tables 1.1.3 through 1.1.5. Focusing first on spending and prices, notice that the aggregate nominal consumption value of goods to services, X_{gt}/X_{st}, and aggregate relative goods to services prices, P_{gt}/P_{st}, declined from 1948-2021.[5] These two facts have been well-documented in the literature, with a number of often mutually-exclusive explanations proposed for why relative spending has declined and why relative prices have also fallen. Non-homothetic preferences, such as those in Kongsamut, Rebelo, and Xie (2001), Herrendorf, Rogerson, and Valentinyi (2013), Boppart (2014), and Comin, Lashkari, and Mestieri (2021), can help partially explain the

[4]Our classification rubric closely matches the definitions of market goods and market services from the National Income and Product Accounts (NIPA). Specifically, we match tasks to the expenditure categories in NIPA Table 2.3.5, *Personal Consumption Expenditures by Major Type of Product*.

[5]Aggregate variables will be denoted with capital letters throughout the paper.

shift in demand toward services and away from goods. However, Boppart (2014) finds that over 50% of the decline in relative goods to services expenditure can be explained by substitution effects (i.e., long-run relative price declines) rather than income effects. The findings in Boppart (2014) thus support stories from Baumol (1967), Baumol, Blackman, and Wolff (1985), and Ngai and Pissarides (2007), wherein structural change is driven by differences in sectoral productivity growth rates directly affecting relative market prices and/or the contention in Acemoglu and Guerrieri (2008) wherein structural change is driven by differential sectoral capital deepening.

When looking specifically at the expenditure series, Figure 1a shows that the rate of decline appears to have slowed. The services share of total expenditure increased by 11.5 percentage points from 1950-1970, 10.1 percentage points from 1970-1990, 7 percentage points from 1990-2010, and only 5.4 percentage points from 1999-2019.10 This is one feature of the data we will attempt to explain in the context of our home production model. When accounting for time-to-consume, our findings show that changing household behavior has dampened, not accelerated, structural change because goods, not services, are increasingly treated as luxuries. This presents an alternative mechnaism compared with the attribution of structural change to consumers' increasing desire for services as their incomes rise, as in Kuznets (1966), Kuznets (1973), Kongsamut, Rebelo, and Xie (2001), Herrendorf, Rogerson, and Valentinyi (2013), and Comin, Lashkari, and Mestieri (2021). This result comes from allowing consumers to adjust their off-market time in response to wage and price variation. When accounting for these consumption and time-use complementarities, where such complementarities are dependent upon the product type (i.e., goods versus services), potential income effects driving structural change disappear, and relative prices become the primary driver of changes in relative expenditure.

The idea that the relative price decline seems to be driving the relative expenditure decline makes sense when inspecting the real-quantity ratios in Figure 1c. Notice that the relative \$2012 real value of goods to services fell until 1990 before rising thereafter. The 1990 breakpoint corresponding to the trend reversal is also associated with the deceleration of structural change we previously discussed. For structural change prior to 1990, a crude but intuitive story whereby income effects have driven the decline in X_{gt}/X_{st} is plausible: if the consumption items are normal goods and relative prices and relative quantities are simultaneously falling, this suggests income effects must be greater than substitution effects. However, after 1990 the story would have to change, given relative quantities rise while relative prices fall, suggesting a strong role for substitution effects, absent additional considerations regarding the degree to which possible complementarities between consumers' expenditure and off-market time-use patterns have contributed to aggregate outcomes.

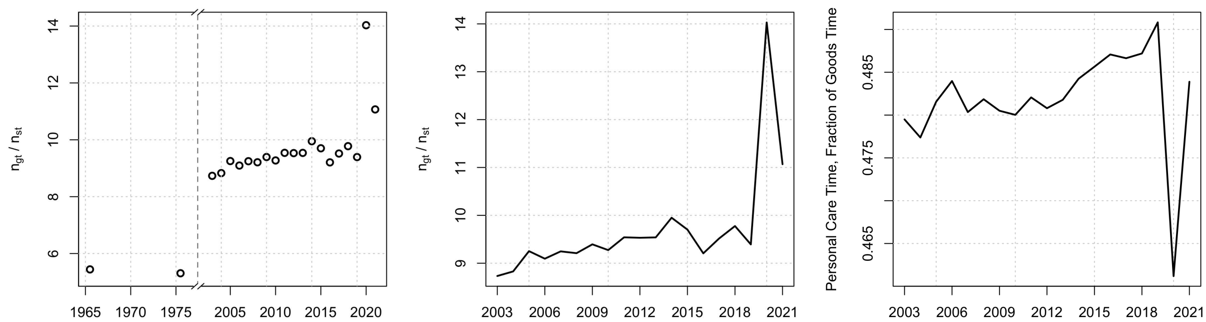

(a) Long-run Ratio (b) Short-run Ratio (c) Personal-care Fraction of Goods Time

Figure 2: In (a) we show the combined AUT-TUESA-ATUS ratio of average off-market time spent using goods versus services, n_{gt}/n_{st}, inclusive of personal-care time and sleep. For AUT and TUESA we average over the years 1965-1966 and 1975-1976 respectively. In (b) we show the time series only from ATUS data for 2003-2021. In (c) we present the personal-care fraction of goods time in ATUS. Note that analogous plots for data where personal-care time and sleep have been removed are featured in Technical Appendix B.

In Figure 2, we present several series describing changes in time-use patterns over both the long-run and short-run. Representation weights are used in the ATUS data, while for AUT and TUESA, we compute simple averages. In all of the plots presented here, we count personal-care time and sleep, which are often left out of home production analyses, as time spent using physical goods. This is because products most associated with these activities are physical goods (lotions, creams, soaps, gels, etc.). To understand the extent to which the inclusion of personal-care time and sleep affect our estimates of the goods-to-services time-use ratio, we exclude these activities from the goods series construction and present those plots in Technical Appendix B. There is no qualitative difference with the evolution of the time series when excluding these activities.

Given the lack of granular time-use data outside of AUT and TUESA datasets, we must infer and extrapolate how household behavior has changed since the mid-twentieth century. In Figure 2a, we extrapolate that between the mid-1970s up to 2003 the ratio of time households spent using goods versus services almost doubled. This is due to the goods share of total time rising while the services share fell. Thus, while relative services expenditure increased over this period, the amount of time consumers spent using services, on average, fell. Looking at ATUS data post-2003, we also observe that the time consumers have spent using goods relative to services continued to slowly rise until 2019, before the COVID spike. It is notable that the COVID pandemic changed time use patterns to a substantially greater degree than expenditure patterns. As this paper is about long-run secular changes, we include the COVID spike data just for completeness, though analyses as to why time-use patterns changed more than expenditure patterns during COVID should be left to another paper.

Focusing on Figure 2b, notice that the ratio of goods-to-services time use appeared to increase leading up to the pandemic. To test whether there is a statistically significant trend in the ratio from 2003-2019 (excluding pandemic years), we run Mann-Kendall tests on the ratio and *reject* the hypothesis that there is no trend present in the goods-to-services time-use ratio (Mann 1945; Kendall 1975). These trends are significant at the 1% level. Thus relative goods-to-services time use increased from 2003-2019. In fact we find that goods time, including personal-care time plus sleep, rose by 10.94 minutes per day for the average consumer while services time fell by 7.58 minutes per day. The average consumer worked 3.36 minutes less per day over this period, as well. Goods time thus increased by 0.76 percentage points while services time fell by 0.53 percentage points, so the changes were small but statistically significant.

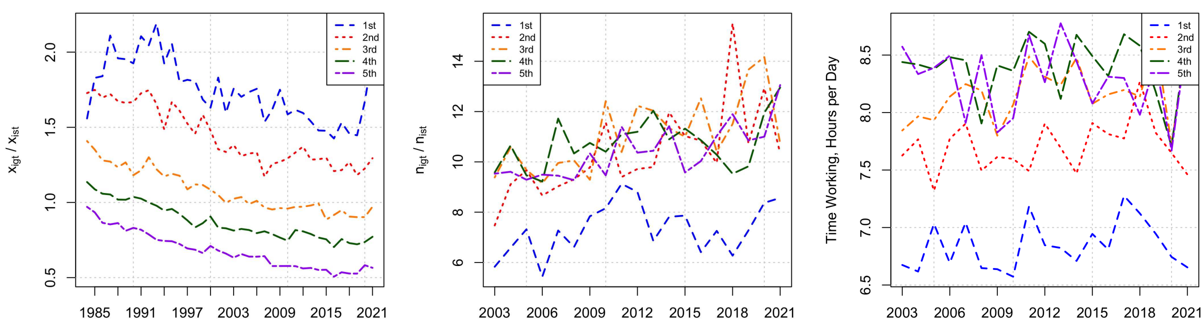

(a) Relative Expenditure (b) Relative Time Use (c) Work Time

Figure 3: Relative expenditure by income quintile from the CEX-LABSTAT database is featured in (a); relative time use by income quintile from ATUS is featured in (b); average work time in hours per day by quintile from ATUS is in (c).

The patterns we observe in aggregate expenditure data and weighted-average time-use data appear to hold when binning consumers by their weighted income quintile. In Figure 3 we show relative goods to services expenditure from the Consumer Expenditure Survey's LABSTAT (CEX-LABSTAT) summary tables in (a), relative goods to services time use, inclusive of personal-care and sleep time, from ATUS in (b), and work time by income quintile also from ATUS in (c).[6,7] Clearly, relative goods to services spending has fallen across income quintiles since 1984, while relative goods to services time use has risen slightly or even remained constant since 2003, depending on the quintile. Further, cross-sectional heterogeneity suggests that higher-income consumers spend relatively more outlay on services than low-income consumers, but low-income consumers spend relatively more time using services than their higher-income peers.

Our empirical results provide an inference regarding household preferences which inform our modeling decisions. As wages increase, households spend more outlay on services but appear to prefer spending more of their time engaged with goods. These facts hold both when examining long-run per-capita averages and cross-sectionally within a single period. Thus, a flexible model which can account for variation in joint spending and time-use allocations, both cross-sectionally and over time, is required to pin down the mechanics driving the shifting composition of demand and time use.

2 Model of the Household

In this section we analyze a Becker (1965) model of household decision making, focusing on the forces which affect the allocation of expenditure across two product types, goods and services. Our approach is partial equilibrium in nature, as households are assumed to be atomic price takers. The goal is to explore how accounting for home production affects household consumption allocations, and thus structural change, in a similar vein as Herrendorf, Rogerson, and Valentinyi (2013), who examine the role of non-homothetic preferences in driving structural change in a partial equilibrium context. Households buy quantities of goods and services, q_j, $j \in \{g, s\}$, from the market at prices, P_j. They use these purchases in activities, f_j, to generate consumption, c_j, from which they ultimately derive flow utility, $u(c_g, c_s)$. The utility function is assumed strictly increasing, strictly concave, and at least twice-continuously differentiable. Further, $u(\cdot,\cdot)$ satisfies Inada conditions in all arguments, so that $u \to -\infty$ as $c_j \to_+ 0, \forall j$.

As in Becker (1965) and discussed in depth in Aguiar and Hurst (2016), we distinguish between consumption and market expenditure. Consumption is the experience that results from combining time and market inputs together to achieve a final outcome from

[6] We partition the CEX expenditure categories following the rubric in Boppart (2014), who uses the CEX Public Use Micro Data (PUMD) to estimate the degree to which relative price variation, as opposed to income effects, has contributed to the rising services share in the U.S. economy. In our quantitative structural analyses later on, we will use granular PUMD data as well, though for illustrative purposes, here, the CEX-LABSTAT suffices.

[7] We do not use the AUT-TUESA dataset in our disaggregated analyses because it does not contain representation weights.

which a consumer derives ultimate utility.[8] To produce final consumption, c_j, consumers must allocate off-market time, n_j, to activities associated with q_j. We assume that the outputs of home production, c_j, are non-tradable to non-household members. To yield c_j market purchases and off-market time are combined in some quasi-concave, at least twice continuously differentiable production activity function:

$$c_j = z_j f_j(q_j, n_j), \qquad \forall j \in \{g, s\} \tag{1}$$

Differences across j in the structures of $f_j(\cdot, \cdot)$ will play a crucial role in determining the degree to which relative expenditure, x_g/x_s, responds to either wage growth or changing relative prices over time.

The parameter z_j is an exogenous in-home technology that captures the efficiency by which a household can transform purchases and time into final consumption. We assume that z_j is realized at the beginning of each period prior to decision making, and is allowed to vary over time along the growth and development path. As a measure of multi-factor productivity, changes in z_j account for changes to c_j not embedded in variation of q_j and n_j.

The changes in this efficiency parameter may be driven by several different factors. First, as a residual, changes to z_j may be driven by unmeasured changes to factors other than market purchases and time that affect final consumption. While it is difficult to determine what such a factor might be in this context, it is not a possibility that can be ruled out. Alternatively, changes in z_j may be driven by either unmeasured quality changes to market purchases or changes in the efficiency of the overall home-production process associated with a particular product type. Constructing such measures of real expenditure directly from nominal data are notoriously difficult, and it is possible for there to be a delay in the recognition of quality improvements or the incorporation of new products into measurements (Gordon 2016; Lebow and Rudd 2021). Additionally, each z_j captures how variation in things like "folk knowledge" (i.e., a households' ability to utilize the products it buys) also affect consumption activities. The residuals are thus functions both of mis-measurement and actual household home-production ability: households of higher ability levels in certain off-market tasks may experience greater returns to engaging in such tasks with the same amount of resources than similar households of lower ability levels. In this manner the interpretation of each z_j is similar to the interpretation of the in-home productivity parameter in the highly stylized model of Boerma and Karabarbou-

[8] With some examples courtesy of Aguiar and Hurst (2016), watching television requires time, a television, and perhaps a cable subscription service, while consuming a meal at home requires that one follow a recipe to combine groceries into a final, prepared meal using utensils and kitchen appliances.

nis (2021), except as applied to activities associated with using goods and services.

Since we are interested in consumption allocations, it will be helpful for us to define the relative efficiency as follows:

Definition 1. *The ratio z_g/z_s is the **relative efficiency of consumption activities** beyond that which is captured by measured market quantities, q_j, and prices, P_j.*

Econometrically, the ratio z_g/z_s is measured as an unobserved residual. This is because we observe expenditure data (i.e., nominal outlay) and from expenditure we can construct price and quantity indices to arrive at indirect measures of both real quantities consumed and relative prices.

Consumers may trade time on the market by supplying hours, ℓ, and receiving efficiency wages, w, in return. Time use must satisfy:

$$\ell + n_g + n_s \leq \overline{n} \tag{2}$$

where $\overline{n}$ is the households' time endowment, which will vary depending on the number of adults in the household.[9] We assume households operate as a collective, so that a household with one adult is a single consumption and time-use decision-maker, as is a household with two adults. We thus take no stand on how consumption and time allocations are distributed across adults.[10] Finally, for completeness, we assume consumers may save and receive dividends from their savings. Let y be capital income net of new investment. The household budget constraint can then be written as

$$\sum_{j \in \{g,s\}} P_j q_j \leq w\,\ell + y \tag{3}$$

Assume now that a household's lifetime utility function is time separable, so that we can invoke two-stage budgeting to focus specifically on the within-period allocation of market expenditure and time, taking intertemporal investment decisions as given.[11] Given y and prices, in each period households choose the vector $(q_g, q_s, \ell, n_g, n_s)$ so as to

[9]For example, for a household with one adult, $\overline{n} = 24$ hours, whereas for a household with two adults, $\overline{n} = 48$ hours.

[10]For an excellent paper examining the intersections of male and female labor force participation, off-market time allocations, and structural transformation, see Ngai and Petrongolo (2017).

[11]See Deaton and Muellbauer (1980) for a deep-dive into the mechanics of two-stage budgeting. In the structural change literature Herrendorf, Rogerson, and Valentinyi (2013) similarly invoke time separability and two-stage budgeting to focus on intratemporal (within-period) allocations.

solve:

$$\begin{aligned} &\max \; u(c_g, c_s) \\ \text{subject to} \quad &(1), (2), \text{ and } (3). \end{aligned} \tag{4}$$

Generally, macroeconomic models such as Kydland and Prescott (1982), Greenwood, Hercowitz, and Huffman (1988), and King, Plosser, and Rebelo (1988) and all of their descendants, impose some degree of separability on market consumption and the decision to allocate time toward market activities versus non-market activities (i.e., what is called "leisure" in these models). These models are useful for analysis of fluctuations in consumption, investment, and labor supply around business cycles and in the context of aggregate patterns of economic growth. Here, however, we assume all off-market time is allocated to a consumption process associated with either goods or services. This omits the possibility of "pure" leisure (i.e. off-market time not associated with a given consumption process). While this assumption may present a potential bias in our analysis in the case that such activities do exist, given that there are few counter examples of time spent outside of work that does not employ some market purchase, we expect that any bias associated with this assumption to be minimal.

Additionally, toward concerns that there exist market services which do not require time to use, note that the designers of ATUS anticipated that consumers may actually *allocate time* toward using services. For example, "Using lawn and garden services," is category code 090401 in the ATUS dataset. Indeed, all "Household Services (not done by self)," are included in classification category 0901xx. Therefore, in the data, there do not seem to be services which consumers could consume without allocating time toward their consumption, regardless of considerations regarding physical constraints. For these reasons in our model, all consumption requires time.

2.1 Household Equilibrium Conditions

Given standard assumptions on the utility function and home production functions, we can write the marginal rate of substitution (MRS) for market products as follows:

$$MRS(q_g, q_s) := \left(\frac{z_g}{z_s}\right)\left(\frac{\partial u}{\partial c_g}\Big/\frac{\partial u}{\partial c_s}\right)\left(\frac{\partial f_g}{\partial q_g}\Big/\frac{\partial f_s}{\partial q_s}\right) = \frac{P_g}{P_s} \tag{5}$$

The condition in (5) shows that the marginal rate of substitution can be decomposed into three separate objects:

i. $\frac{z_g}{z_s}$ is the aforementioned object from Definition 1.

ii. $\frac{\partial u}{\partial c_g} / \frac{\partial u}{\partial c_s}$ is the marginal rate of substitution between different experiences or activities associated with using goods versus services.

iii. $\frac{\partial f_g}{\partial q_g} / \frac{\partial f_s}{\partial q_g}$ is the productivity-neutral ratio of the two separate marginal products of f_j with respect to q_j for the two separate home production processes.

Object (i) is exogenous, while (ii) and (iii) are endogenous objects. The relative consumption efficiency in (i) captures the degree to which variation in the relative experienced efficiency of consumption, beyond that which is embedded in measured quantities and prices, has affected equilibrium outcomes. If goods are indeed benefiting from these exogenous experienced relative efficiency improvements, then we should observe that the ratio $\frac{z_g}{z_s}$ is rising over time. Since $\frac{z_g}{z_s}$ will be a residual term, we thus require a parameterized model in order to estimate (ii) and (iii) and back out (i). We perform such an exercise in Section 3.4.

2.2 Income and Substitution Effects in Response to Wage Variation

In this section we show that in home production models with consumption time, wage variation affects the quantities of market purchases via *both* a total income effect and a substitution effect. Similarly, market price variation will affect the allocation of off-market time via two separate channels. We will thus highlight how income and substitution effects impact consumption allocations and time use in Beckerian models.

For such an analysis we require the expenditure-minimization dual problem (EMP) associated with the utility-maximization problem (UMP) described above. To build toward an EMP we also require the Beckerian budget constraint, which can be derived by substituting the time-use constraint in (2) into (3) for ℓ:

$$\sum_{j \in \{g,s\}} (P_j q_j + w\, n_j) \leq w\, \overline{n} + y \tag{6}$$

The right-hand side of (6) is maximum-possible income net of savings if the consumer devoted all of his/her time to labor-market activities (Becker 1965). Blundell and Macurdy (1999) show that the objective function for the EMP in models featuring at least one off-market time-utilization decision and elastic labor is just the left-hand side of (6). This yields the following expenditure function:

$$e(P_g, P_s, w) = P_g q_g + P_s q_s + w(n_g + n_s) \tag{7}$$

Assumption 1. *The composite function $u\big(z_g f_g(q_g, n_g), z_s f_s(q_s, n_s)\big)$ is strictly increasing, strictly* *concave, continuously differentiable, and homothetic in all of q_g, q_s, n_g, and n_s. As such, Marshallian demand functions are well-behaved and linear in maximum possible income net of savings, $w\,\overline{n} + y$.*

Off-market time utilization is effectively a purchase decision: the consumer gives up a share of possible income that could have been earned working in exchange for more time. In a consumption/leisure model, there are thus two prices — one for each of the two purchasing decisions, consumption and leisure. But with multiple off-market time-use decisions, each weighted in the budget constraint by the same price (i.e., the wage, w), the cardinality of the price vector is less than the cardinality of the quantity vector, which includes time. Indeed, in Beckerian models the price vector is constrained to be one plus the size of the vector of market purchase prices, while the number of time-utilization decisions may grow as much as the modeler sees fit. In our case the consumer faces four effective purchase decisions — q_g, q_s, n_g, and n_s — but only three prices — P_g, P_s, and w.

Since the prices of n_g and n_s are constrained to be identical, this will impact Shephard's Lemma generated from a home production model. We now describe a version of Shephard's Lemma for the expenditure function in (7). Let superscript m index the Marshallian demand functions derived by solving the UMP. Let superscript h index the Hicksian demand functions derived by solving the EMP. Note that Marshallian demands are functions of consumption prices, wages, and capital income net of savings, while Hicksian demands are functions of consumption prices, wages, and utility.

Lemma 1. *Given the expenditure function, $e(P_g, P_s, w)$, Shephard's Lemma for off-market time use and wages is*

$$n_g^h + n_s^h = \frac{\partial e}{\partial w}$$

All proofs are presented in Technical Appendix A.1.[12] Lemma 1 follows directly from the fact that n_g and n_s always have the same price but are separate decisions which will differ from each other due strictly to the structures of $f_g(\cdot,\cdot)$, $f_s(\cdot,\cdot)$, and $u(\cdot,\cdot)$. Note that the fact that the price set is smaller than the choice set will impact Shephard's Lemma in this case. This has implications for the terms of the cross-price responsiveness of market commodities q_j to wages w, which we will characterize under the EMP specified by (7).[13]

We now describe two different ways to characterize the responsiveness of demand to wage variation. Let $\widetilde{y}(w) = w\overline{n} + y$. Note that under Assumption 1 we can write

[12] All technical appendices are available at the author's website: https://www.npretnar.com/research.

[13] By "cross-price responsiveness" we mean the responsiveness of q_j to wage variation, since the wage is *both* a price *and* proportional to labor income.

Marshallian demand as a linear function in $\widetilde{y}(w)$:

$$q_j^m\big(P_g, P_s, w, \widetilde{y}(w)\big) = \widetilde{y}(w)\Gamma(P_g, P_s, w) \tag{8}$$

where $\Gamma(P_g, P_s, w)$ is some continuously differentiable function.

Lemma 2. *The total responsiveness of Marshallian demand for q_j to variation in w is:*

$$\frac{\mathrm{d}q_j^m}{\mathrm{d}w} = \frac{\partial q_j^m}{\partial w} + \frac{\partial q_j^m}{\partial \widetilde{y}}\frac{\partial \widetilde{y}}{\partial w}$$
$$\text{where} \quad \frac{\partial \widetilde{y}}{\partial w} = \overline{n}$$

so that

$$\frac{\mathrm{d}q_j^m}{\mathrm{d}w} = \widetilde{y}(w)\frac{\partial \Gamma}{\partial w} + \Gamma(P_g, P_s, w)\overline{n} \tag{9}$$

Lemma 3. *The Slutsky equations describing the partial responsiveness of demand to wages are*

$$\frac{\partial q_j^m}{\partial w} = \frac{\partial q_j^h}{\partial w} - \frac{\partial q_j^m}{\partial \widetilde{y}}(n_g + n_s), \qquad \forall j \in \{g, s\}$$
$$\text{where} \quad \underbrace{\frac{\partial q_j^h}{\partial w}}_{\text{Substitution Effect}} = \underbrace{\widetilde{y}(w)\frac{\partial \Gamma}{\partial w}}_{\text{Marshallian Responsiveness}} + \underbrace{\Gamma(P_g, P_s, w)(n_g + n_s)}_{-\text{ Income Effect}}$$

Lemma 2 follows directly from the linearity of q_j^m in $\widetilde{y}(w)$. But, note that Lemma 2 characterizes *total* responsiveness. Lemma 3 presents the Slutsky equation describing the responsiveness of market consumption to wage variation. Note that the variation described in Lemma 3 is only partial, in that it fails to account for how a change in w affects both the left-hand side of the Beckerian constraint in (6) but also $\widetilde{y}(w)$, i.e. the right-hand side of the constraint. This is because the Slutsky equation relates the equilibrium outcomes of the EMP with those of the UMP. Since the solution to the EMP describes the variation of demand in prices and utility, holding $\widetilde{y}$ fixed, Shephard's Lemma (i.e., Lemma 1) will naturally fail to account for how variation in w affects $w\,\overline{n}$, i.e. the value of time.

Lemma 2 describes a demand responsiveness relationship that is more sophisticated than the structure in Lemma 3, since it takes into consideration how changes to wages affect the value of a consumer's total time, $\overline{n}$, while Lemma 3 considers only how changing wages affect the value of a consumer's off-market time, $n_g + n_s$. Yet, Lemma 2 says nothing about how the total responsiveness of demand relates to Hicksian substitution and income effects. This brings us to Proposition 1.

Proposition 1. *The total responsiveness of Marshallian demand for q_j to variation in w is the sum of a classic substitution effect, a classic income effect, and a total time-value effect as follows:*

$$\frac{\mathrm{d}q_j^m}{\mathrm{d}w} = \underbrace{\frac{\partial q_j^h}{\partial w}}_{\text{Substitution Effect}} \underbrace{- \frac{\partial q_j^m}{\partial \widetilde{y}}(n_g + n_s)}_{\text{Income Effect}} + \underbrace{\frac{\partial q_j^m}{\partial \widetilde{y}}\overline{n}}_{\text{Time-value Effect}}$$

Corollary 1. *Alternatively, we can write the total responsiveness as a function of a substitution effect and a total income effect:*

$$\frac{\mathrm{d}q_j^m}{\mathrm{d}w} = \underbrace{\frac{\partial q_j^h}{\partial w}}_{\text{Substitution Effect}} + \underbrace{\frac{\partial q_j^m}{\partial \widetilde{y}}\ell}_{\text{Total Income Effect}}$$

Notice that if q_j is observed in the data to increase as wages rise, then Proposition 1 says either the Hicksian substitution effect or the time-value effect must be dominating, since the income effect is negative as long as q_j^m is normal. Corollary 1 provides an alternative characterization of the responsiveness of demand to wage variation that is a function of a Hicksian substitution effect and a total income effect, which accounts for the combined value of the income effect and the time-value effect. By Corollary 1, if $\frac{\mathrm{d}q_j^m}{\mathrm{d}w} < 0$, then the substitution effect must be dominating the total income effect, which is always clearly non-negative, since $\ell \geq 0$ and each q_j is normal. In our quantitative exercises we will decompose the responsiveness of q_j to variation in w using the decomposition from Corollary 1.

2.3 Veblen Goods

Our empirical findings will present summary statistics which suggest both that the nature of consumption has changed over time, and that changes to consumption behavior, specifically with regards to the utility consumers place on experiences associated with using goods, has helped drive structural change. To understand structural change in the context of the changing nature of consumption and off-market time use, we will require the model's version of the own-price Slutsky equation for q_j^m.

The own-price Slutsky equation for q_j^m is the usual one

$$\frac{\partial q_j^m}{\partial P_j} = \frac{\partial q_j^h}{\partial P_j} - \frac{\partial q_j^m}{\partial \widetilde{y}} q_j^m \tag{10}$$

We do not wrap this definition in a lemma nor prove it directly: it is just the standard Slutsky equation. Further, there is no alternative object describing the total responsiveness of q_j to variation in P_j like that featured in Proposition 1.

Definition 2. *A market product is said to be **Veblen** if the responsiveness, $\frac{\partial q_j^m}{\partial P_j}$, is positive, **and** this can be attributed to the substitution effect being positive and dominating the income effect.*[14]

Unlike a Giffen good, which is inferior and also has a positive own-price elasticity, Veblen goods are not inferior and are indeed normal. Consumers simply perceive them as luxuries *because* they have higher prices, which could also, simultaneously, reflect certain luxury goods' trademark and brand valuations that are difficult to quantify.[15]

2.4 Parameterized Model

To engage in more in-depth analyses of our model equilibrium, we deploy flexible parameterizations for u and f_j. For flow utility over activities c_g and c_s we consider a constant elasticity of substitution (CES) function where each activity is given equal weight:

$$u(c_g, c_s) = \left(c_g^\rho + c_s^\rho\right)^{\frac{1}{\rho}}$$

For $\rho \in (0, 1)$ activities associated with goods and services are gross substitutes. When $\rho < 0$ the activities are gross complements. As $\rho \to 0$ the function converges to the Cobb-Douglas utility function, while as $\rho \to -\infty$ the function converges to the Leontief utility function.

In our quantitative exercises we will test the degree to which data suggest that experiences associated with goods or services are complements or substitutes. Our findings will provide evidence that such experiences are indeed complements. Note that there are certainly experiences for which c_g and c_s are substitutes: think, for example, of in-home

[14] Veblen goods are named after Thorstein Veblen's 1899 book, *The Theory of the Leisure Class: An Economic Study of Institutions* (Veblen 1899). Veblen first hypothesized that consumer culture was partially driven by conspicuous consumption and a keeping-up-with-the-Joneses mentality. With respect to Veblen goods, consumers desire them partially because such goods are perceived to be status signals.

[15] For example, Becker and Murphy (2000) argue that Veblen goods are observed because our economic models of consumer behavior fail to account for how social perceptions regarding the quality of certain products change over time.

cleaning activities where one uses soaps and such purchased from a consumer-goods store versus hiring out such cleaning activities to a service provider. By contrast goods and services are complementary in many other aspects: consider communicating with a friend via a smart phone (good) which requires using cellular network data services (service). The experience "Using the smart phone" contains both goods and services as inputs, but this is the case of joint production (simultaneously producing a goods experience and a services experience) which is difficult to quantify and identify in data. Allowing for such joint activities would require us to make an assumption regarding the fraction of goods and services used in those joint activities, which would then require a firm stand as to what constitutes a "joint activity." To avoid having our efforts encumbered by a cascading series of assumptions, we do not take an explicit stand as to what constitutes an activity, maintaining the high-level assumption that there are activities associated with primarily using goods and those associated with primarily using services. In the quantitative section we can then estimate how our model (and its corresponding assumptions) rationalizes the data: if we find $\rho < 0$ then this would suggest many goods and services activities are indeed complementary with each other; if, alternatively, we find $\rho \in (0, 1)$ this would suggest that paying for services is primarily a substitute for using goods.

We consider a flexible CES parameterization for the home production functions:

$$f_j(q_j, n_j) = \left(\omega_j q_j^{\nu_j} + (1 - \omega_j) n_j^{\nu_j}\right)^{\frac{1}{\nu_j}}, \qquad \forall j \in \{g, s\}$$

Production weights, ω_j, and market/off-market substitution elasticities, ν_j, are heterogeneous across the product space. Note that if $\nu_g = \nu_s = 1$ then one can recover a model where "pure" leisure (in the sense in which no market products are utilized) and market consumption are weakly separable (because $\rho < 1$).[16]

2.4.1 Parametric Household Equilibrium Conditions

In this section we derive the parameterized versions of the within-category marginal rates of technical substitution along with the cross-category marginal rates of substitution for market expenditure and off-market time use. We then use the equilibrium first-order conditions to write expressions which characterize relative expenditure, x_g/x_s, and relative

[16] Again, though, this does not make any sense in the context of our analysis, given how tasks are classified in the ATUS data.

time use, n_g/n_s, as functions of market prices, wages, and preference primitives.[17] Assume either that $\nu_j \neq 0, \forall j \in \{g, s\}$, or $\nu_j = 0$. That is, if one process is associated with unit-elastic substitutability between q_j and n_j, then so is the other process.

After solving (4) subject to the time-use and budget constraints, we get parameterized equilibrium expressions for the marginal rates of technical substitution for each process, $j \in \{g, s\}$, and the marginal rate of substitution for market inputs q_g and q_s:

$$\left(\frac{\omega_j}{1-\omega_j}\right)\left(\frac{q_j}{n_j}\right)^{\nu_j - 1} = \frac{P_j}{w}, \qquad \forall j \in \{g, s\} \tag{11}$$

$$\left(\frac{c_g}{c_s}\right)^{\rho-1} \frac{z_g \omega_g q_g^{\nu_g - 1}}{z_s \omega_s q_s^{\nu_s - 1}} \left(\omega_g q_g^{\nu_g} + (1-\omega_g) n_g^{\nu_g}\right)^{\frac{1-\nu_g}{\nu_g}} \times \left(\omega_s q_s^{\nu_s} + (1-\omega_s) n_s^{\nu_s}\right)^{\frac{\nu_s - 1}{\nu_s}} = \frac{P_g}{P_s}, \qquad \nu_j \neq 0, \forall j \tag{12}$$

$$\left(\frac{c_g}{c_s}\right)^{\rho}\left(\frac{q_g}{q_s}\right)^{-1} \frac{\omega_g}{\omega_s} = \frac{P_g}{P_s}, \qquad \nu_j = 0, \forall j \tag{13}$$

Note that (11) is well-defined for all values of $\nu_j < 1$, while (12) is only well-defined if $\nu_j \neq 0$, $\forall j$. (13) describes $MRS(q_g, q_s)$ when home production functions are Cobb-Douglas. We can get similar expressions to (12) and (13) which each describe the marginal rate of substitution for off-market time use, $MRS(n_g, n_s)$:

$$\left(\frac{c_g}{c_s}\right)^{\rho-1} \frac{z_g (1-\omega_g) n_g^{\nu_g - 1}}{z_s (1-\omega_s) n_s^{\nu_s - 1}} \left(\omega_g q_g^{\nu_g} + (1-\omega_g) n_g^{\nu_g}\right)^{\frac{1-\nu_g}{\nu_g}} \times \left(\omega_s q_s^{\nu_s} + (1-\omega_s) n_s^{\nu_s}\right)^{\frac{\nu_s - 1}{\nu_s}} = 1, \qquad \nu_j \neq 0, \forall j \tag{14}$$

$$\left(\frac{c_g}{c_s}\right)^{\rho}\left(\frac{n_g}{n_s}\right)^{-1} \frac{1-\omega_g}{1-\omega_s} = 1, \qquad \nu_j = 0, \forall j \tag{15}$$

Notice that the right-hand sides of (14) and (15) are equal to unity. This follows from the fact that consumers' off-market time-use decisions are all constrained by (2) and that the opportunity cost consumers face with respect to allocating their time away from labor for any arbitrary off-market task is simply w.

[17] We also derive the Marshallian demand functions for each q_j and n_j, which are needed for some of our computational exercises, but we relegate these functions and their derivations to Technical Appendix A.2.

We can re-arrange (11) to get an implicit function, $n_j(q_j)$, that is linear in q_j:

$$n_j(q_j) = q_j \left[\frac{w\omega_j}{P_j(1-\omega_j)} \right]^{\frac{1}{\nu_j - 1}}, \qquad \forall j \in \{g, s\} \tag{16}$$

Using (16) to substitute out instances of n_j in (12) and (13), we can then re-arrange those equations to arrive at expressions for relative expenditure, x_g/x_s, as functions of only prices and parameters:

$$\begin{aligned} \frac{x_g}{x_s} = \left(\frac{P_g}{P_s}\right)^{\frac{\rho}{\rho-1}} \left(\frac{z_g}{z_s}\right)^{\frac{\rho}{1-\rho}} \left(\frac{\omega_g}{\omega_s}\right)^{\frac{1}{1-\rho}} \left(\omega_g + (1-\omega_g)\left[\frac{w\omega_g}{P_g(1-\omega_g)}\right]^{\frac{\nu_g}{\nu_g-1}}\right)^{\frac{\nu_g-\rho}{\nu_g(\rho-1)}} \\ \times \left(\omega_s + (1-\omega_s)\left[\frac{w\omega_s}{P_s(1-\omega_s)}\right]^{\frac{\nu_s}{\nu_s-1}}\right)^{\frac{\rho-\nu_s}{\nu_s(\rho-1)}}, \qquad \nu_j \neq 0, \forall j \end{aligned} \tag{17}$$

$$\begin{aligned} \frac{x_g}{x_s} = \left(\frac{\omega_g}{\omega_s}\right)^{\frac{1}{1-\rho}} \left[\left(\frac{\omega_g}{1-\omega_g}\right)^{1-\omega_g} \left(\frac{\omega_s}{1-\omega_s}\right)^{\omega_s-1}\right]^{\frac{\rho}{\rho-1}} \left(\frac{z_g}{z_s}\right)^{\frac{\rho}{1-\rho}} \\ \times P_g^{\frac{\rho\omega_g}{\rho-1}} P_s^{\frac{\rho\omega_s}{1-\rho}} w^{\frac{\rho(\omega_s-\omega_g)}{\rho-1}}, \qquad \nu_j = 0, \forall j \end{aligned} \tag{18}$$

By inverting (16) we can also substitute out instances of q_j in (14) and (15) to get expressions for relative off-market time use:

$$\begin{aligned} \frac{n_g}{n_s} = \left(\frac{z_g}{z_s}\right)^{\frac{\rho}{1-\rho}} \left(\frac{1-\omega_g}{1-\omega_s}\right)^{\frac{1}{1-\rho}} \left(\omega_g\left[\frac{w\omega_g}{P_g(1-\omega_g)}\right]^{\frac{\nu_g}{1-\nu_g}} + 1 - \omega_g\right)^{\frac{\nu_g-\rho}{\nu_g(\rho-1)}} \\ \times \left(\omega_s\left[\frac{w\omega_s}{P_s(1-\omega_s)}\right]^{\frac{\nu_s}{1-\nu_s}} + 1 - \omega_s\right)^{\frac{\rho-\nu_s}{\nu_s(\rho-1)}}, \qquad \nu_j \neq 0, \forall j \end{aligned} \tag{19}$$

$$\begin{aligned} \frac{n_g}{n_s} = \left(\frac{1-\omega_g}{1-\omega_s}\right)^{\frac{1}{1-\rho}} \left[\left(\frac{\omega_g}{1-\omega_g}\right)^{\omega_g} \left(\frac{\omega_s}{1-\omega_s}\right)^{-\omega_s}\right]^{\frac{\rho}{1-\rho}} \left(\frac{z_g}{z_s}\right)^{\frac{\rho}{1-\rho}} \\ \times P_g^{\frac{\rho\omega_g}{\rho-1}} P_s^{\frac{\rho\omega_s}{1-\rho}} w^{\frac{\rho(\omega_s-\omega_g)}{\rho-1}}, \qquad \nu_j = 0, \forall j \end{aligned} \tag{20}$$

We will use (17) and (19) as our primary estimating equations in Section 3.

By examining all of (17) through (20) it is apparent that differences between ω_j's and ν_j's will determine the degree to which wage variation affects expenditure allocations. Further, such differences will also affect how off-market time allocations respond to relative price variation, P_g/P_s, holding wages fixed. This is because variation across home production activities in both ω_j and ν_j characterizes the heterogeneity consumers face with respect to time-use and market-purchase tradeoffs for products of different types. To summarize, our main points are as follows: 1) when engaging in different consumption activities requires different degrees of time inputs, and labor is elastic, relative spending responds to wage variation; and 2) the time allocation will also depend on the relative price of associated market purchases.

2.4.2 Elasticity of Relative Expenditure with Respect to Wages

We are interested in understanding how heterogeneity in different in-home consumption processes leads to the appearance that re-allocation of expenditure is driven by rising incomes. In this section we analyze how relative expenditure varies in wages and explore several questions. First, how does variation in the structures of home production processes affect the degree to which relative expenditure responds to wages? Further, do changing relative prices have second-order effects on the responsiveness of expenditure allocations to rising wages? What about product-market price inflation relative to wage inflation? To answer these questions, we derive the following elasticity. Take logs and differentiate both (17) and (18) in $\ln w$ to get the elasticity of relative expenditure, x_g/x_s, with respect to wages, which we denote $\varepsilon^{x,w}$:[18]

$$\begin{aligned}\varepsilon^{x,w} := \frac{\partial \ln(x_g/x_s)}{\partial \ln w} = \left(\frac{\nu_g - \rho}{\nu_g(\rho-1)}\right)\frac{(1-\omega_g)\frac{\nu_g}{\nu_g-1}\left[\frac{w\omega_g}{P_g(1-\omega_g)}\right]^{\frac{\nu_g}{\nu_g-1}}}{\omega_g + (1-\omega_g)\left[\frac{w\omega_g}{P_g(1-\omega_g)}\right]^{\frac{\nu_g}{\nu_g-1}}} \\ + \left(\frac{\rho - \nu_s}{\nu_s(\rho-1)}\right)\frac{(1-\omega_s)\frac{\nu_s}{\nu_s-1}\left[\frac{w\omega_s}{P_s(1-\omega_s)}\right]^{\frac{\nu_s}{\nu_s-1}}}{\omega_s + (1-\omega_s)\left[\frac{w\omega_s}{P_s(1-\omega_s)}\right]^{\frac{\nu_s}{\nu_s-1}}}, \qquad \nu_j \neq 0, \forall j\end{aligned} \tag{21}$$

$$\varepsilon^{x,w} := \frac{\partial \ln(x_g/x_s)}{\partial \ln w} = \frac{\rho(\omega_s - \omega_g)}{\rho - 1}, \qquad \nu_j = 0, \forall j \tag{22}$$

$\varepsilon^{x,w}$ is different from the elasticity of *absolute* (not relative) expenditure with respect to wages.[19] Rather, (21) and (22) describe how wage variation affects the *allocation* of market consumption, holding market prices, P_g and P_s, fixed.

Wages will affect relative spending, and thus expenditure shares, because of differences between the marginal rates of technical substitution across the different home production processes. Note that, even if $\nu_g = \nu_s$, as long as the relative intensities of market

[18]Throughout our exposition, ε will be used to denote elasticities.

[19]The elasticity of absolute expenditure on product j with respect to wages is described for a general model in Equation 17 of Aguiar, Hurst, and Karabarbounis (2012). We add to their theoretical contributions by considering how wage variation affects *relative allocations* in a model with consumption time.

purchases across production processes are different (i.e., $\omega_g \neq \omega_s$), then wage variation will non-trivially impact the expenditure allocation. This is true even if the elasticities of substitution within each process are unity, as can be seen by inspecting (22), where the elasticity of relative expenditure with respect to wages, given by $\frac{\rho(\omega_s - \omega_g)}{\rho - 1}$, is non-zero as long as $\omega_g \neq \omega_s$ and $\rho \neq 0$.

One implication of Cobb-Douglas home production functions is that relative spending will change in fixed proportion to wage growth. There are two potential shortcomings of this. First, with Cobb-Douglas, we will not capture cases where consumers across the income distribution may adjust their expenditure shares at different rates in response to proportional wage gains. Second, the Cobb-Douglas case does not allow for variation in market prices to affect how the spending allocation responds to wage variation (i.e., second-order, cross-partial effects).

By contrast the general CES formulation (non-zero ν_j, $\forall j$) is flexible but non-trivial: depending on ν_j, ω_j, $\forall j$, and ρ, as well as the aggregate price level, changing prices and wages will interact in different ways to determine the relative spending elasticity. This means that under certain parameterizations and prices, consumers across the wage distribution will have different first-order consumption responses to wage variation. Further, the degree to which such disparate behavior is observed will vary over time as the price levels of market products and labor rise at different rates. We will illustrate these properties with some examples.

Figure 4 plots the level sets of $\varepsilon^{x,w}$ against variation in w and P_s, holding $P_g = 1$ fixed. In each plot we let $\rho < 0$, so that consumption experiences are gross complements.[20] Figure 4 is meant to show how variation in relative market prices, P_g/P_s, has second-order effects on the responsiveness of the expenditure allocation to wage variation. In Figure 4 we consider two parameterizations to demonstrate that the monotonicity of $\varepsilon^{x,w}$ depends on differences in ω_j's across consumption activities: in panel (a) we present a parameterization that yields a monotonic increase in $\varepsilon^{x,w}$ in both the ratios w/P_g and P_s/P_g; in panel (b) we present the level sets for an alternative parameterization that yields non-monotonic variation of $\varepsilon^{x,w}$. Notice, under both parameterizations time use and market purchases are complementary in each process (i.e., $\nu_j < 0, \forall j$), while the home production outputs, c_j, are also complementary (i.e., $\rho < 0$). Nonetheless, the nature of $\varepsilon^{x,w}$ is very different between the two, and in the case of panel (b), monotonicity of $\varepsilon^{x,w}$ depends on the value of P_s.[21]

[20] This is the sign of ρ in our preferred quantitative models, estimated and analyzed in Section 3.

[21] This can be seen by noting that, in panel (b), middle-income consumers are more elastic when $P_s < 1$, while we observe what appears to be monotonically declining elasticities across the wage distribution when $P_s > 1$.

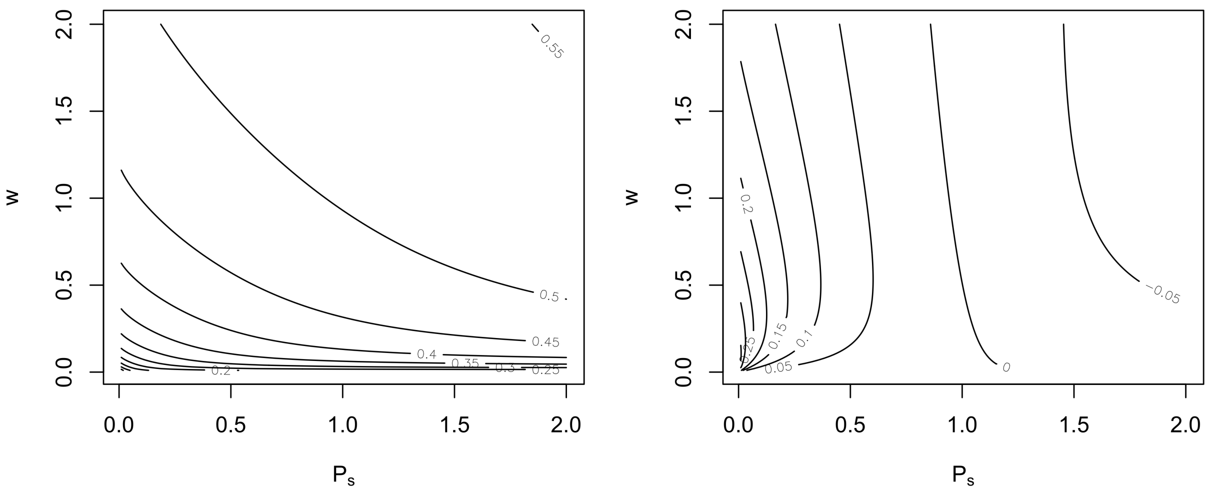

(a) $\omega_g = 0.02$, $\omega_s = 0.003$, $\nu_g = -1$, $\nu_s = -3$, $\rho = -25$ (b) $\omega_g = 0.5$, $\omega_s = 0.25$, $\nu_g = -20$, $\nu_s = -10$, $\rho = -1$

Figure 4: We present level sets of $\varepsilon^{x,w}$, under variation in w and P_s, holding P_g fixed. Each panel features a different parameterization, so as to illustrate how the values of the underlying home production elasticities affect the responsiveness of the expenditure allocation to wage variation.

It is clear that the rate at which consumers re-allocate their market expenditure in response to wage gains depends *both* on the magnitude of the wage gain and the underlying relative price of goods to services. Variation in the relative price clearly also has second-order consequences for the degree to which consumers reallocate expenditure as their wages rise and their time becomes more valuable. Further, consumers across the wage distribution will respond to relative price variation at different rates. Depending on the underlying home production elasticities, it is even possible to observe non-monotonic responsiveness across the wage distribution of relative spending to relative price variation. This can be seen by inspecting how the level sets vary along the vertical axis in Figure 4b. Thus, both the parameterization and market prices will determine whether $\varepsilon^{x,w}$ is concave or convex in w. This has implications for the second-order effects of structural change and whether the transition toward services will accelerate or decelerate as average wages rise.

Our results also indicate that if market prices and wages inflate at different rates, then structural change will either accelerate or decelerate over time depending on the home production elasticities. Figure 5 shows this. We fix $w = 1$ and simulate $\varepsilon^{x,w}$ by proportionately re-scaling P_g and P_s for a set of relative prices. For example, let P^{agg} be some aggregate price level. Let $\theta \in \{0.25, 0.5, 1, 1.5, 2\}$ be a scalar, such that $w = \theta\, P^{agg}$, and for $\theta = 1$, $w = P^{agg}$. For a continuum of relative prices, we re-scale the product-level prices

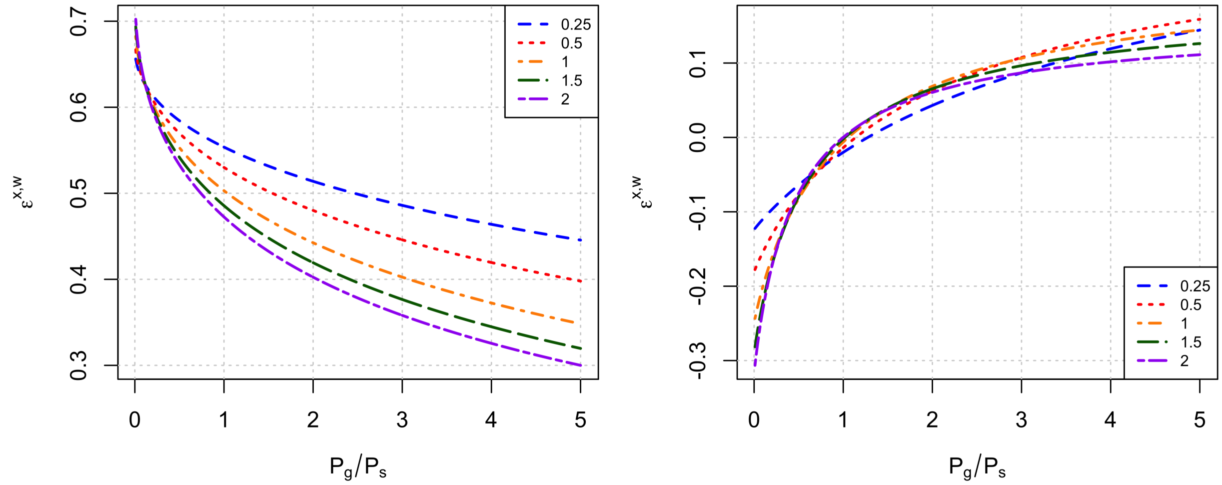

(a) $\omega_g = 0.02$, $\omega_s = 0.003$, $\nu_g = -1$, $\nu_s = -3$, $\rho = -25$ (b) $\omega_g = 0.5$, $\omega_s = 0.25$, $\nu_g = -20$, $\nu_s = -10$, $\rho = -1$

Figure 5: We show how $\varepsilon^{x,w}$ varies in P_g/P_s under different aggregate price levels. For example, the purple line is such that the aggregate price level is $2\times$ that of wages, while the blue line is $0.25\times$ that of wages. In panel (a) as relative prices fall and aggregate inflation increases, the allocation responsiveness falls. In panel (b) a non-monotonic relationship between $\varepsilon^{x,w}$ and market-price inflation is observed.

each by θ and simulate $\varepsilon^{x,w}$. Variation between the colored lines in Figure 5 shows the dependency of $\varepsilon^{x,w}$ on inflation in market prices relative to wages. In the left-hand panel as market prices increase relative to wages, the allocation becomes less elastic in wages for all but small values of P_g/P_s. Under the parameterization in panel (a), in an inflationary environment driven by product-market inflation outpacing wage inflation, the rate of structural change will slow. In panel (b) we observe that the allocation elasticity can be non-monotonic in relative inflation, which can be seen by noting that as inflation rises, the allocation elasticity as a function of relative prices becomes more concave. Thus, inflation also has both first- and second-order effects on the allocation's wage responsiveness.

Proposition 2. *Assume z_g and z_s are fixed. Consider the knife-edge case where $\nu_g = \nu_s = \nu$ and $\omega_g = \omega_s = \omega$. Expenditure shares will respond to wage variation as follows:*

i. If $\nu = 0$, then shares are constant in w.

ii. If $\nu \neq 0$, then shares are constant in w if and only if $P_g = P_s$.

Under all but a knife-edge case for the home production substitution elasticities and input/output intensities, the elasticity of relative spending (and thus expenditure shares) with respect to wages will be non-zero. Proposition 2 describes these knife-edge conditions. It says that neither differences in the within-process elasticities of substitution or input/output intensities are required in order to have relative spending vary in wages and thus income. As long as time and market inputs are not unit-elastic substitutes, then even if each process is associated with the same factor share *and* the same within-process substitution elasticity, the expenditure allocation will vary in wages while holding relative prices fixed.

3 Quantitative Model and Estimation

We recover parameters using U.S. micro data from both ATUS and the CEX's Public Use Micro Data (PUMD), running separate estimations depending on which dataset we target.[22] We do not use the AUT or TUESA data in our estimations, given the datasets lack representation weights.

We have multiple quantitative goals that require precise estimates of the model's underlying structural parameters. We want to assess the degree to which the responsiveness of relative spending to wages varies both cross-sectionally and over time. We also seek to understand how different aspects of household behavior have contributed to long-run structural change, as reflected by the rising services share of expenditure. Next, we engage in several exercises to estimate the degree to which either income or substitution effects appear to dominate for both q_j and n_j in terms of own-price and cross-price responsiveness in an effort to drill down and discover what are the core underlying mechanisms driving structural change. Finally, we consider how consumers across the wage distribution have differentially benefitted from the long-run changes to the relative price of goods to services.

We estimate several versions of the model on both ATUS and PUMD micro data. With the ATUS data we estimate four different models — two including personal-care time and two without personal-care time. With each dataset we separately restrict either $\rho < 0$ or $\rho \in (0, 1)$, and then we assess which model provides the better fit. With the PUMD data we simply estimate two models — one with $\rho < 0$ and one with $\rho \in (0, 1)$. We prefer ATUS models including personal-care time with $\rho < 0$ and the PUMD model with $\rho < 0$ for both theoretical and empirical reasons which will be discussed below. Throughout our quantitative assessments we will focus on results from estimates with $\rho < 0$, relegating the gross-substitutes estimates to Technical Appendix F.

After cleaning, the ATUS dataset contains 25,327 observations. The average household in our sample has a respondent who is 40.51 years old, works an average of 7.62 hours per day, and earns a nominal hourly wage of \$16.49 per hour (\$17.07 in real \$2012). All

[22] That is, we do not synthetically link the two datasets, as is done in Boerma and Karabarbounis (2021).

households feature at least one adult who works and earns positive wages.

Our PUMD dataset is built around the "Interview" component. After cleaning, our PUMD dataset contains expenditure data for 607,467 married households from 1990-2021. We compute hourly wages in PUMD by summing both respondent and spousal work hours, then dividing total before-tax income by this sum, as is done with the PUMD in Boerma and Karabarbounis (2021). Households must work a positive amount of time and spend a non-zero, positive amount on both goods and services in order to be included in the final sample. We further reduce our sample of households by throwing out those who we estimate earn $<$ \$1 per hour and $>$ \$99.99 per hour, where the latter exclusion is made in order to align with ATUS. The average household in our PUMD sample features adults that work a combined 8.09 hours per day and earn a combined average hourly wage of \$22.40 per hour (\$24.89 in real \$2012). In the PUMD the respondent has an average age of 43.70 years with a spouse who is an average 44.85 years of age. In both of our estimations, we deploy the provided ATUS and PUMD sampling weights, which we denote by the letter ζ_{it}, where i indexes households and t years.[23] All weights reside in the unit interval and are such that $\int \zeta_{it}\,\mathrm{d}i = 1$ for all t.

Our estimation routine is of the Bayesian learning variety and focuses on two separate, single-equation likelihood models. Specifically, we use Hamiltonian Monte Carlo (HMC) estimation techniques to recover the posterior distribution of structural parameters given data.[24] We take a Bayesian approach, as opposed to using a frequentist M-estimator, for several reasons. First, the model is highly non-linear in parameters. A Bayesian approach allows us to discipline our estimator with prior distributions that are consistent with theory, so that the data reconciles our estimates under the assumption that the core economic theory underlying our model formulation holds. Second, Bayesian estimators are preferable for various computational reasons. While Bayesian estimators may take longer to run in terms of computer time, they are more computationally robust and require substantially less coding time in order to implement.[25]

We form the likelihood function around logged versions of equations (17) and (19) separately for each of the different estimations (depending on the datasets used and as-

[23] We introduce i and t subscripts for this section only.

[24] We describe HMC integration techniques in Technical Appendix C. For detailed explanations of HMC techniques see Neal (2011), Betancourt and Stein (2011), Gelman et al. (2013b), and Gelman et al. (2013a).

[25] This is especially convenient with highly non-linear models where the parameter space may (and is very likely) multi-modal, so that finding a maximum, as for a maximum likelihood estimator, is difficult. Bayesian models do not require checking for local versus global maxima and thus do not require trying many different starting values or worrying about whether or not an optimization procedure has or has not converged, since the estimated objects of interest are not point estimates (though such estimates can easily be extracted) but distributions of parameters conditional upon modeling assumptions.

sumptions on ρ):

$$\begin{aligned}
&\ln\left(\frac{n_{igt}}{n_{ist}}\right) - \frac{\rho}{1-\rho}\left(\gamma_{z_g/z_s}\cdot(t-\bar{t})\right) - \frac{1}{1-\rho}\ln\left(\frac{1-\omega_g}{1-\omega_s}\right) \\
&\quad - \frac{\nu_g-\rho}{\nu_g(\rho-1)}\ln\left(\omega_g\left[\frac{w_{it}\omega_g}{P_{gt}(1-\omega_g)}\right]^{\frac{\nu_g}{1-\nu_g}} + 1-\omega_g\right) \\
&\quad - \frac{\rho-\nu_s}{\nu_s(\rho-1)}\ln\left(\omega_s\left[\frac{w_{it}\omega_s}{P_{st}(1-\omega_s)}\right]^{\frac{\nu_s}{1-\nu_s}} + 1-\omega_s\right) - \varphi^n_t = \eta^n_{it}
\end{aligned} \tag{23}$$

$$\begin{aligned}
&\ln\left(\frac{x_{igt}}{x_{ist}}\right) - \frac{\rho}{\rho-1}\ln\left(\frac{P_{gt}}{P_{st}}\right) - \frac{\rho}{1-\rho}\left(\gamma_{z_g/z_s}\cdot(t-\bar{t})\right) - \frac{1}{1-\rho}\ln\left(\frac{\omega_g}{\omega_s}\right) \\
&\quad - \frac{\nu_g-\rho}{\nu_g(\rho-1)}\ln\left(\omega_g + (1-\omega_g)\left[\frac{w_{it}\omega_g}{P_{gt}(1-\omega_g)}\right]^{\frac{\nu_g}{\nu_g-1}}\right) \\
&\quad - \frac{\rho-\nu_s}{\nu_s(\rho-1)}\ln\left(\omega_s + (1-\omega_s)\left[\frac{w_{it}\omega_s}{P_{st}(1-\omega_s)}\right]^{\frac{\nu_s}{\nu_s-1}}\right) - \varphi^x_t = \eta^x_{it}
\end{aligned} \tag{24}$$

We assume that the ratio of relative in-home productivities contains a component that is identical and grows at the same rate, γ_{z_g/z_s}, across all households. Households have idiosyncratic errors, η^n_{it} or η^x_{it}, but we allow for time fixed effects, φ^n_t or φ^x_t, since we cannot specify serially-correlated errors due to the non-longitudinal data structure. $\bar{t}$ is the geometric-mean year in the corresponding time series (i.e., the geometric mean year from 2003-2021 for ATUS or 1990-2021 for PUMD). All prices are PCE indices from the BEA.

Table 1: $MRTS_j$ Regression Estimates for Prior Parameters (OLS)

Goods		*Services*	
$\widehat{\phi}_{g0}$	-3.869^{***}	$\widehat{\phi}_{s0}$	-3.904^{***}
	(0.044)		(0.164)
$\widehat{\phi}_{g1}$	-0.897^{***}	$\widehat{\phi}_{s1}$	-1.735^{***}
	(0.016)		(0.060)
$\widehat{\nu}_g$	-0.115	$\widehat{\nu}_s$	0.424
Observations	95	Observations	95
R^2	0.973	R^2	0.901
Residual Std. Error	0.075 (df = 93)	Residual Std. Error	0.292 (df = 93)

Note: *p<0.1; **p<0.05; ***p<0.01

Since our estimator is Bayesian, we must construct prior distributions for the structural parameters. To do so, we turn to intuition and reduced-form estimates using OLS regressions. For the elasticity of substitution between c_{igt} and c_{ist}, we allow for either $\rho < 0$ or $\rho \in (0, 1)$ by running separate estimators and ranking their outputs using an information criterion. For the in-home production substitution elasticities, ν_j, we seek meaningful priors that ensure logged versions of the marginal rate of technical substitution conditions in (11) hold:

$$\ln P_{jt} n_{ijt} - \ln x_{ijt} = \underbrace{\frac{1}{\nu_j - 1} \ln \left(\frac{\omega_j}{1 - \omega_j} \right)}_{\phi_{j0}} + \underbrace{\frac{1}{\nu_j - 1}}_{\phi_{j1}} \ln \left(\frac{w_{it}}{P_{jt}} \right) + \epsilon_{ijt}, \qquad \forall j \in \{g, s\} \quad (25)$$

We only use (25) to *sign* η_j's, as we will estimate these parameters with greater precision in our fully Bayesian model. We use OLS on (25) for each $j \in \{g, s\}$ on a 2003-2021 panel of matched CEX/ATUS data aggregated by income quintile, $i \in \{1, 2, 3, 4, 5\}$, with wages taken from ATUS and consumption prices from PCE series.[26] Both the reduced-form and corresponding structural-parameter prior estimates are presented in Table 1 for all models where personal-care time is included.[27] Notice that $\widehat{\nu}_g < 0$ while $\widehat{\nu}_s > 0$, suggesting that goods and off-market time are in-home complements, while services and off-market time are in-home substitutes. We thus form priors under these assumptions.

The structural priors are:

$$\begin{aligned} &-\rho \in \mathcal{LN}(-0.5, 1), \quad \text{or} \quad \rho \in \mathcal{U}[0, 1], \qquad \nu_s, \omega_g, \omega_s \sim \mathcal{U}[0, 1] \\ &-\nu_g \sim \mathcal{LN}(\mu_{\nu_g}, 1), \quad \text{where} \quad \mu_{\nu_g} = \ln(-\widehat{\nu}_g) - \frac{1}{2} = -2.197 \end{aligned}$$

where $\mathcal{LN}$ denotes a log-normal distribution. Note that, depending on the model, we either assume $\rho < 0$ or $\rho \in (0, 1)$. The fixed-effect coefficients and trend growth rate, γ_{z_g/z_s}, are given flat and improper priors:

$$\varphi_t^k \sim \mathcal{U}(-\infty, \infty), \quad \forall k \in \{n, x\} \quad \text{and} \quad \gamma_{z_g/z_s} \sim \mathcal{U}(-\infty, \infty)$$

All observations are generated from a mean-zero normal distribution:

$$\eta_{it}^k \sim \mathcal{N}(0, \sigma^2), \quad \forall k \in \{n, x\} \qquad \frac{1}{\sigma^2} \sim \text{Gamma}(2, 4)$$

[26] Note, for this regression we match the annual CEX-LABSTAT summary tables by income quintile to weighted averages by year and income quintile from ATUS.

[27] Technical Appendix E contains prior regression estimates for models without personal-care time.

We sample the precision, $1/\sigma^2$, from a gamma distribution with the so-called shape/rate parameterization.[28]

Structural parameter posterior distributions are presented in Table 2 for the $\rho < 0$ models where personal-care time is included. Technical Appendix E contains estimates for models without personal care time (both when $\rho < 0$ and $0 < \rho < 1$), and Technical Appendix F contains estimates for models inclusive of personal-care time where we assume $\rho \in (0, 1)$, but all other parameters are allowed to follow the same prior distributions as described above.

While it is common in the structural change literature, and other quant-macro literature more generally, to match model outputs to aggregate series, we follow Boppart (2014) and estimate elasticities directly on micro data. We do this because our model generates unique income effects, specifically with respect to wages, and the model does not admit a representative consumer, like models with non-homothetic preferences and inelastic time use.[29] The model-fit criteria we thus employ involve assessing how the model fits the 2003-2021 ATUS and 1990-2021 CEX-PUMD micro data.

Figure 6 depicts the error distributions by year for both ATUS and PUMD estimations, where the ATUS estimator includes personal care time. The solid lines represent means across samples by year, while the shaded areas are the 90% confidence regions.[30] Blue values show the error distribution under the assumption that c_g and c_s are substitutes, while red values are for when c_g and c_s are complements.[31] To compute the error distribution we first take the posterior means of parameters as point estimates and then use such point estimates to predict relative time use and relative expenditure for each agent in each sample. The error distributions for all models at the posterior means, including those without personal care time, are in Technical Appendix D.

It is clear that errors are centered around zero because we cannot reject the null hypothesis that the agent-level posterior means of errors are themselves mean-zero across agents in our samples. However, models with $\rho < 0$ have more desirable predictive properties than those with $\rho \in (0, 1)$. First, examining blue lines in both Figure 6a and 6b

[28]This parameterization has probability density function

$$\pi(x) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}, \quad \text{for} \quad x \sim \text{Gamma}(\alpha, \beta)$$

where $\Gamma(\cdot)$ is the gamma function.

[29]Recall, wages are both a price *and* scale income, so wage variation is associated with a non-trivial substitution effect as well as an income effect.

[30]Note that since this is a Bayesian model, these are confidence *regions* and not confidence *intervals*, as the posterior error distribution need not necessarily be symmetric.

[31]Recall, Technical Appendix F includes the posterior parameter distributions when $\rho \in (0, 1)$.

Table 2: Posterior Distribution of Parameters, $\rho < 0$

ATUS Data							
Parameter	Mean	S.D.	2.5%	25%	50%	75%	97.5%
ρ	-0.9607	1.1730	-4.0573	-1.1334	-0.5923	-0.3031	-0.0823
ν_g	-0.1143	0.1693	-0.4687	-0.1356	-0.0687	-0.0354	-0.0102
ν_s	0.3502	0.2199	0.0164	0.1623	0.3335	0.5143	0.7939
ω_g	0.5087	0.2922	0.0314	0.2500	0.5114	0.7656	0.9801
ω_s	0.4643	0.2985	0.0187	0.1953	0.4480	0.7178	0.9690
γ_{z_g/z_s}	0.0582	0.9873	-1.8316	-0.6288	0.0509	0.7333	2.0065
σ	3.8811	1.8365	1.7846	2.6416	3.4249	4.5766	8.9103
Information Criteria:	$WAIC^a = 264.796$						
PUMD Data							
Parameter	Mean	S.D.	2.5%	25%	50%	75%	97.5%
ρ	-0.9579	1.2049	-3.9047	-1.2031	-0.6115	-0.3060	-0.0863
ν_g	-0.1174	0.1453	-0.5002	-0.1431	-0.0745	-0.0363	-0.0099
ν_s	0.4377	0.2844	0.0079	0.1930	0.4170	0.6833	0.9620
ω_g	0.5044	0.2956	0.0263	0.2457	0.4984	0.7522	0.9795
ω_s	0.5338	0.2770	0.0263	0.3057	0.5505	0.7661	0.9729
γ_{z_g/z_s}	0.0445	0.9150	-1.6901	-0.5890	0.0544	0.6679	1.8284
σ	3.0894	1.5564	1.4554	2.1376	2.6680	3.5554	7.2841
Information Criteria:	$WAIC^a = 534.168$						

[a] Widely applicable information criterion. See Watanabe (2010).

note that errors appear to be non-stationary over time, despite the fact we can still reject the mean-zero hypothesis. Meanwhile, when $\rho < 0$ the errors are stationary, and in the case of the PUMD data well-centered around zero. Second, this exercise suggests that the best-fitting model is that which extracts parameters from expenditure data and equilibrium conditions using the PUMD under the assumption that $\rho < 0$. This model features mean-zero errors without a consistent bias, while such a bias appears to exist with the ATUS data. In summary, for these reasons here, we prefer models with $\rho < 0$ and will focus on results generated by such models as we move forward.

To also compare performance across models featuring either different datasets or dif-

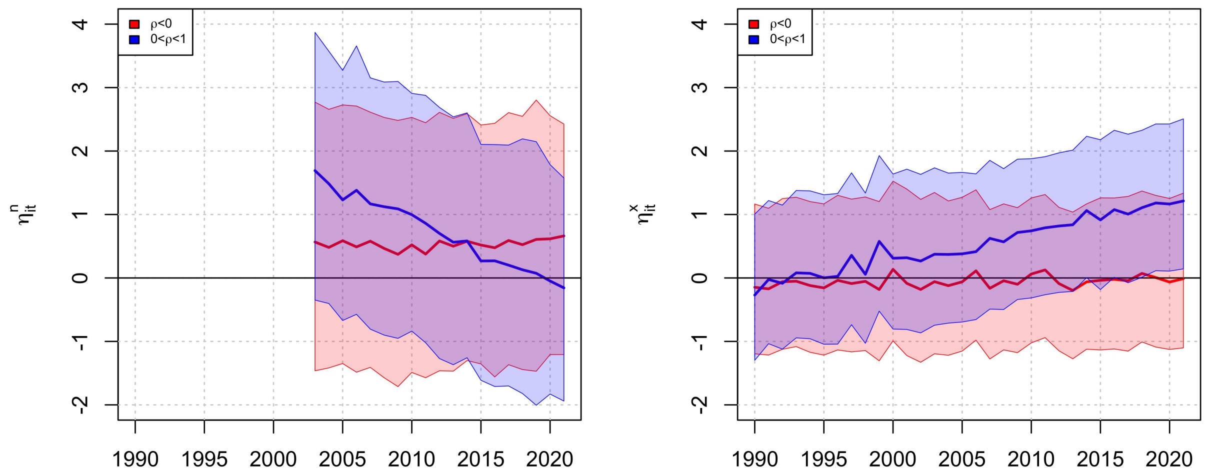

(a) Error Distribution by Year (ATUS)

(b) Error Distribution by Year (PUMD)

Figure 6: These plots contain the weighted error distributions over time for estimators on both ATUS and PUMD datasets. The shaded areas are 90% confidence regions, while the solid lines are observation-level means in each year of each sample. Note that the ATUS sample goes back to 2003, while the PUMD sample starts in 1990.

ferent assumptions on ρ, we examine each model's WAIC (widely applicable information criterion) statistics. Watanabe (2010) shows that assessing model performance with WAIC is asymptotically equivalent to comparing Bayesian models via an information-loss estimation extracted from a cross-validation procedure.[32,33]

Our comparisons of information criteria yield several model-selection results. First, models where $\rho < 0$ (i.e., goods experiences are complementary with services experiences) outperform all models where $\rho \in (0, 1)$ (i.e., substitutes). Second, ATUS models estimated on data with personal-care time outperform models estimated without personal-care time. Third, all ATUS models are more informative than their PUMD counterparts: when $\rho < 0$, the ATUS model with personal-care time is approximately two-times more informative than the corresponding PUMD model, despite being biased. We will examine inferences from parameters estimated on both datasets moving forward.

The posterior distributions of structural parameters feature substantial overlap across models, including the distributions of γ_{z_g/z_s}. We also find that fixed-effect distributions

[32]As an information criterion, WAIC is preferable to AIC (Akaike information criterion) and BIC (Bayesian information criterion) as it is robust to statistical models that feature underlying hidden Markov processes and possible autocorrelation. Further, selection based on WAIC is asymptotically equivalent to cross validation (Watanabe 2010). But, just like every other information criterion, WAIC still provides a measure of the theoretical distance between so-called "true" models of the data-generating process and that of the specification.

[33]Recall with information criteria, lower is better.

in both models are centered around zero (i.e., we cannot reject the null that they are not mean-zero) with relatively high dispersion.[34] Some additional discussion around estimates of γ_{z_g/z_s} is warranted. The posterior distributions of γ_{z_g/z_s} suggest that we cannot reject the hypothesis that there is *no* trend growth/decline in relative in-home productivities, at least via a micro-data lens. Given z_g/z_s is an unobserved residual, in many of our calculations below, especially those pertaining to analysis of aggregate data, we will simply set $\gamma_{z_g/z_s} = 0$ and focus on the model's particular implications for relative-price and wage variation impacting expenditure allocations.

3.1 Relative Expenditure Elasticity of Wages

From here forward all of our analyses operate on the $\rho < 0$ models where personal-care time has been included in estimates using ATUS data. In this section we demonstrate how relative expenditure appears to respond to wage variation through the lens of the estimated models by exploring the model-implied elasticities in (21) from parameter estimates and micro data. Recall that the expression for the elasticity in (21) is independent of relative productivities, z_g/z_s, so that our analyses here do not depend on observed residuals.

In Figure 7 we explore the model-implied elasticities in several formats. In Figures 7a and 7b we plot each agent's wages against his/her estimated elasticity. We notice a clear pattern, such that the relative expenditure elasticity rises as wages rise, though there is less variation in elasticities between two similarly compensated high-wage earners than between two low-wage earners.[35] When this elasticity is negative, a 1% increase in wages is associated with a 1% decline in relative x_{igt}/x_{ist}. Low-wage earners thus respond to a wage increase by increasing their relative services consumption, but high-wage earners respond the opposite: they increase their relative goods consumption. This contradicts a story whereby structural change (i.e., the rising services share of spending) is driven by services being perceived as relative luxuries. We include Figures 7c and 7d to show that the variation we observe in Figures 7a and 7b across our entire sample is not biased by time fixed effects, since the elasticities for the mean, top quintile, and bottom quintile of

[34] For ATUS the average over all of the posterior average-annual fixed-effect coefficients is 1.244, and the average annual standard deviation of the fixed-effect coefficients is 6.234. For PUMD these values are -0.005 and 4.957 respectively. Posterior distributions of fixed-effect parameters are in Technical Appendix D.

[35] Readers should not be deceived by these plots. The "thickness" of the lines, which are actually dense clusters of points, result from the fact that relative prices are varying over time, so that a consumer with wage w in period t faces a different market structure compared to a consumer with the same wage in period $t' \neq t$. Still, the variance within a given wage level appears very minimal, suggesting wage differences, not relative prices, are the primary contributor to variation in $\varepsilon_{it}^{x,w}$.

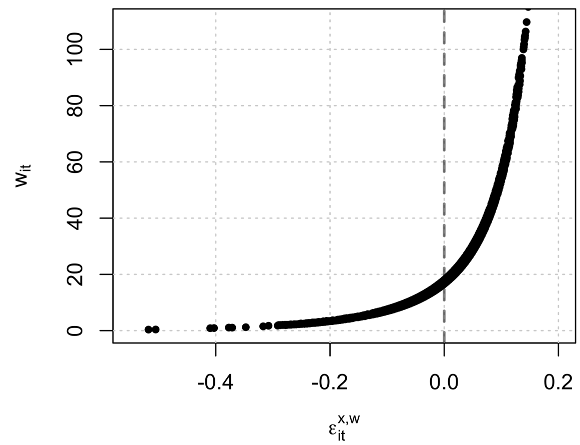

(a) Elasticities and Wages (ATUS)

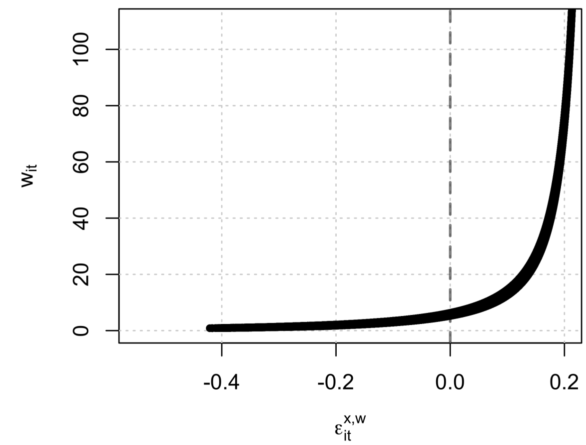

(b) Elasticities and Wages (PUMD)

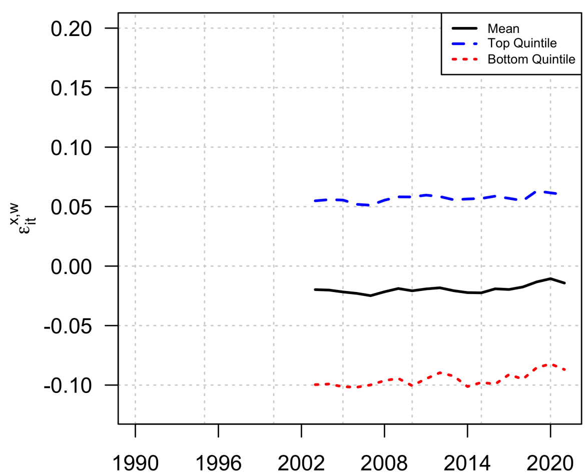

(c) Elasticities by Quintile (ATUS)

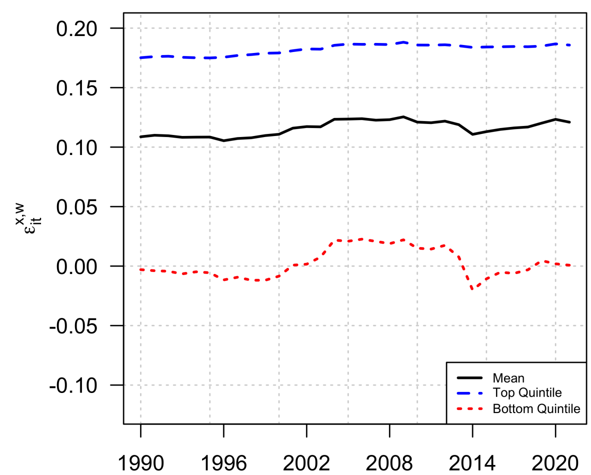

(d) Elasticities by Quintile (PUMD)

Figure 7: In this figure we plot the relative expenditure elasticity of wages from (21), as estimated with micro data using the posterior means of structural parameters as point estimates. Panels (a) and (b) are scatterplots with the elasticity on the horizontal axis and the agent's wage on the vertical axis. Panels (c) and (d) show the evolution of the average elasticity over time and over all agents in our samples (black line), those in the top wage quintile (blue line), and those in the bottom wage quintile (red line). Panels (a) and (c) use estimates from the model featuring ATUS data, while (b) and (d) use estimates from the model featuring PUMD.

the wage distribution are fairly flat over both datasets' sample periods.

3.2 Structural Change with Aggregate Data

In this section we use the posterior distributions of parameters to understand mechanically how changing relative prices and growth in average real wages have contributed to the rise in the services share of expenditure since the mid-twentieth century.[36] For our decompositions we fix each variable of interest one at a time and simulate the time series. We plot these simulations in Figure 8. Note that when we are fixing relative prices, the counterfactual time series reflects the impact of rising wages on relative spending. When wages are fixed, we see how relative price changes affect outcomes. Assume throughout these exercises that relative in-home productivities (z_g/z_s) are simply fixed and not changing over time, (i.e., $\gamma_{z_g/z_s} = 0$). We will eventually estimate these productivities in Section 3.4. The exercises here are thus net of possible productivity improvements and designed to understand the leverage the model's estimated structural parameters (specifically, ρ, ν_j, and ω_j) have on inferences pertaining to the development path (i.e., the goods to services transition) as it relates specifically to relative-price and wage variation.

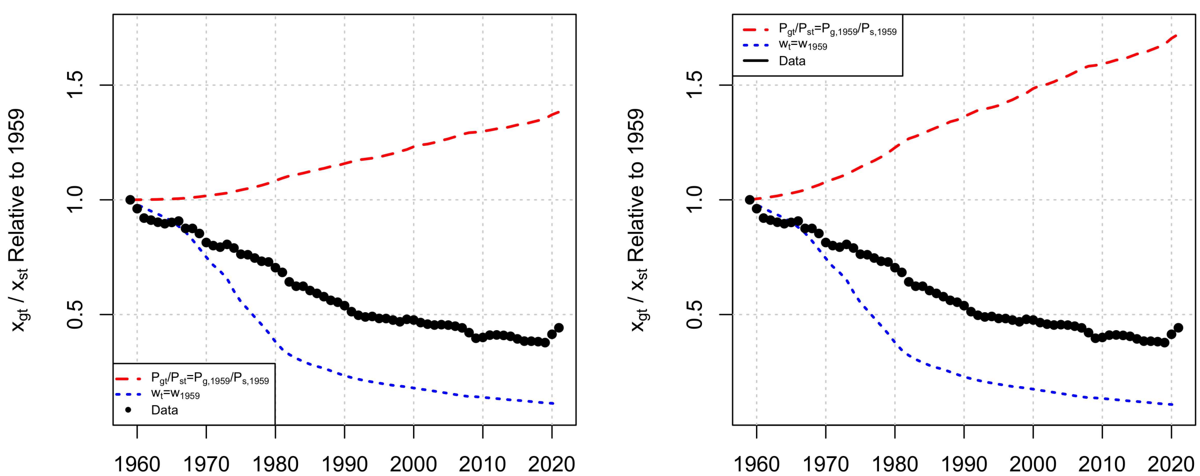

(a) ATUS Parameters

(b) PUMD Parameters

Figure 8: This figure presents posterior-predicted relative expenditure (black lines), along with counterfactual predictions, where fixed 1959 relative prices are in red and fixed 1959 real wages are in blue. The legend is in panel (a). All series are normalized to 1 in 1959.

[36]While we show how aggregate expenditure, price, and quantity data have evolved since 1948 in Figure 1, here, since we require robust average annual wage data, which we compute using total labor income and total labor hours, we start at 1959 when consistent measurements of hours for both hourly and salaried workers become available. Specifically, we use "Hours Worked by Full-time and Part-time Employees" and let labor income be "Compensation of Employees, Received: Wage and Salary Disbursements," both from the BEA. To get real wages we divide nominal wages by the aggregate price level, which we take to be the NIPA Personal Consumption Expenditure price from NIPA Table 2.4.4.

Focusing first on the red lines (fixed relative prices, $P_{gt}/P_{st} = P_{g,1959}/P_{s,1959}$, $\forall t$), note that for all models structural change would have actually gone in the opposite direction (red lines above the black lines implying a transition away from services and towards goods) had relative prices not fallen to the extent we observe in data.[37] Relative to the literature this result is most consistent with findings in Boppart (2014), who uses a flexible, non-homothetic preference structure to identify whether substitution versus income effects are responsible for the rising services share, finding substitution effects appear to dominate. Our finding thus suggests that either differences in sectoral TFP growth (Baumol 1967; Baumol, Blackman, and Wolff 1985; Ngai and Pissarides 2007) and/or differences in sectoral capital deepening (Acemoglu and Guerrieri 2008), which each contribute to general equilibrium variation in relative prices, have thus played an important role in driving down relative sectoral output. As Herrendorf, Rogerson, and Valentinyi (2014) show, when the elasticities of substitution in sectoral production are unity, relative prices are always inversely proportional to relative sectoral TFPs. Thus, if the goods sector is indeed growing faster than services, we would expect the relative price of goods to services to fall.

By contrast, however, had real wages never grown, the structural transition from goods to services would have actually been *more* extreme (blue lines below black lines) over the entire observation period. This runs in sharp contrast to research that has attributed structural change to rising incomes (Kongsamut, Rebelo, and Xie 2001; Herrendorf, Rogerson, and Valentinyi 2013; Comin, Lashkari, and Mestieri 2021). In fact we observe that wages contribute mostly to second-order dampening effects, whereby rising wages have acted as a buffer *against* goods to services structural change, while the role of relative prices in driving structural change appears first order.

3.3 Model-implied Income and Substitution Effects

In this section we will show that in the mid-twentieth century, services consumption was indeed increasing in wages and goods consumption was decreasing in wages, consistent with the theory of development proposed by Kuznets (1966) and Kuznets (1973). But as the U.S. economy grew, the sign of the goods elasticity of consumption with respect to wages flipped, and by 2019 almost entirely caught up with the corresponding services elasticity. This trend coincided with goods consumption becoming increasingly more of an apparent luxury in a process we are calling "Veblenization." We define this term in Definition 3.

[37] For an illustration of how relative prices have declined in data, see Figure 1b.

Definition 3. ***Veblenization*** *refers to the process by which the own-price elasticity of a market product transforms from negative to positive due to the increasing dominance of a positive substitution effect.*

Meanwhile, services actually became more sensitive to own-price variation, so that consumers came to treat services as relatively *less* luxurious compared to goods.

In this section and the following one we will argue that the relative quality consumers experience from using goods has grown faster than that of using services. Further, these relative changes to the consumption experience are likely not captured in measurements of aggregate quantity and price indices, contributing to the model-implied Veblenization of q_g and relatively increasing (in absolute terms) own-price sensitivity of q_s over time. We will thus seek to understand structural change by focusing on how second-order effects have driven these changing price sensitivities. Together, these indirectly competing forces helped drive the goods to services transition, and subsequently slowed it down as the relative price of goods to services continued to fall.

Tables 3 and 4 present the aggregate average (for the average wage-earner in any given year) Marshallian elasticity, ε^m, Hicksian substitution elasticity, ε^h, and total income effect in quasi-elasticity form, $-\varepsilon^{inc}\vartheta$, over 20-year increments for 1959, 1979, 1999, and 2019 as estimated with parameters from the given models (i.e., ATUS and PUMD respectively) while using aggregate average real wage, capital income net of savings, and price data.[38,39] Again, for now, we maintain the assumption that $\gamma_{z_g/z_s} = 0$ since we cannot reject this hypothesis so that inferences here do not account for possible productivity changes. In the next section we will back out and study aggregate productivities using aggregate data. Note that the total income effects corresponding to market-product price elasticities (e.g., responsiveness to variation in P_g and P_s) are identical to the Slutsky income effect. However, for responsiveness to wage variation, the total income effect is that described in Corollary 1. The tables are divided into four parts, where the variable of interest whose elasticity is being estimated is featured at the top of each section. On the left-hand side of each section the prices P_g, P_s, and w are listed. From these tables, we observe several facts: 1) goods own-price elasticities have become more inelastic over

[38]The data series we use are the aggregate average ones from Section 3.2, except we must add data for capital income net of savings per person per day. See Technical Appendix B.3 for how we arrive at such a data series. All elasticities are model-implied along the equilibrium path, taking wages, prices, and capital income net of savings as given. We pass wages, prices, and capital income net of savings through the estimated Marshallian demand functions to get model-predicted values for q_j and n_j, as well as model-predicted income and substitution effects. The parameterized Marshallian demand functions and income and substitution effects are featured in Technical Appendix A.2 and A.3, respectively.

[39]Note that $\varepsilon^m = \varepsilon^h - \varepsilon^{inc}\vartheta$, where ϑ is the total budget share (i.e., share of $w\overline{n} + y$) of either q_j or n_j — whichever's elasticity we are estimating. This object, along with the underlying elasticities, are non-constant and will vary in prices and wages. Thus, it also varies over time.

time and actually turned positive, so that q_g appears to be a Veblen luxury; 2) in the 1950s and 1970s, q_g was relatively inferior, since q_g would fall as w rose, but by 2019 q_g came to be increasing in w; 3) the nature of the responsiveness of n_g to P_g has also changed over time, so that n_g is now complementary with q_g from a demand (but not a home production) perspective;[40] 4) demand for services quantities is normal and has become more sensitive to own-price variation since 1959; 5) consumption of both goods and services responds increasingly positively to rising wages; 6) consumers generally substitute away from off-market time as wages rise, but over time this substitution effect has weakened for both n_g and n_s.

In 1959, goods consumption would have declined as P_g rose, due to the total income effect outweighing the substitution effect. Further, at this time, q_g was decreasing in wages, while the same is true for n_g. Both q_g and n_g have seen the signed elasticity levels of their responsivenesses to P_g and w increase over time. With respect to n_g, a wage increase in 2019 had a smaller impact on a household's time allocation than it would have in the past, while consumption of market goods would rise in response to a wage increase. With respect to q_s and n_s, however, the opposite trend over time is observed: q_s has become *more* own-price elastic due to a strengthening substitution effect, while n_s has become *less* own-price elastic (i.e., with respect to w) due to both weakening substitution and income effects.

The strength of home production complementarities in c_g and c_s can be seen by inspecting the cross-price responsivenesses of both q_g and n_g with respect to P_s and q_s and n_s with respect to P_g. Elasticities for q_g and n_g move in lock-step with each other, declining over time, as the inputs q_g and n_g appear to go from being substitutes to q_s in mid-century to complements in 2019. As P_s rises both q_g and q_s will fall, holding all else fixed, but as P_g rises, q_g will rise and q_s will fall. Then, over time as the relative price continues to fall, structural change will slow as a result of inflation affecting both P_g and P_s. This is because q_g will continue to rise as P_g rises but q_s will fall as P_s rises. Increases in w will only partially offset the negative inflationary responsiveness. Veblenization can help explain the second-order effects of slowing structural change — at least some products that comprise the goods series are luxuries, causing the decline in relative goods to services spending to start to reverse itself despite relative goods to services prices still falling.

[40] It is important to distinguish between the two different types of complementarities — home production complementarities versus classical demand complementarities. Home production complementarities describe how q_g and n_g relate to each as inputs in the production of c_g, while classical demand complementarities simply describe the cross-price responsiveness of n_g to q_g, where such responsiveness depends on home production complementarities, income, and other prices.

Table 3: Income and Substitution Effects, ATUS Model

Goods Consumption, q_g

	1959			1979			1999			2019		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.551	-0.236	-0.315	-0.442	-0.122	-0.320	-0.165	0.100	-0.265	0.116	0.299	-0.183
P_s	0.072	0.284	-0.212	-0.049	0.131	-0.180	-0.342	-0.193	-0.149	-0.639	-0.498	-0.141
w	-0.180	-0.628	0.447	-0.052	-0.463	0.411	0.151	-0.251	0.403	0.359	-0.076	0.435

Services Consumption, q_s

	1959			1979			1999			2019		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.141	0.173	-0.315	-0.143	0.177	-0.320	-0.140	0.125	-0.265	-0.134	0.049	-0.183
P_s	-0.649	-0.436	-0.212	-0.740	-0.560	-0.180	-0.844	-0.695	-0.149	-0.892	-0.752	-0.141
w	0.251	-0.197	0.447	0.359	-0.052	0.411	0.460	0.057	0.403	0.498	0.063	0.435

Goods Time, n_g

	1959			1979			1999			2019		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	0.032	0.347	-0.315	0.136	0.456	-0.320	0.467	0.732	-0.265	0.830	1.013	-0.183
P_s	0.072	0.284	-0.212	-0.049	0.131	-0.180	-0.342	-0.193	-0.149	-0.639	-0.498	-0.141
w	-1.078	-1.525	0.447	-0.949	-1.360	0.411	-0.746	-1.149	0.403	-0.539	-0.974	0.435

Services Time, n_s

	1959			1979			1999			2019		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.141	0.173	-0.315	-0.143	0.177	-0.320	-0.140	0.125	-0.265	-0.134	0.049	-0.183
P_s	0.678	0.890	-0.212	0.618	0.799	-0.180	0.546	0.695	-0.149	0.506	0.647	-0.141
w	-1.288	-1.736	0.447	-1.180	-1.591	0.411	-1.079	-1.482	0.403	-1.041	-1.476	0.435

Table 4: Income and Substitution Effects, PUMD Model

Goods Consumption, q_g

	1959			*1979*			*1999*			*2019*		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.625	-0.339	-0.286	-0.504	-0.190	-0.313	-0.227	0.051	-0.278	0.061	0.260	-0.199
P_s	0.241	0.508	-0.267	0.063	0.288	-0.224	-0.285	-0.109	-0.176	-0.635	-0.474	-0.160
w	-0.377	-0.905	0.528	-0.171	-0.623	0.452	0.110	-0.295	0.405	0.392	-0.035	0.427

Services Consumption, q_s

	1959			*1979*			*1999*			*2019*		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.141	0.145	-0.286	-0.141	0.173	-0.313	-0.137	0.141	-0.278	-0.130	0.069	-0.199
P_s	-0.587	-0.319	-0.267	-0.750	-0.525	-0.224	-0.938	-0.762	-0.176	-1.019	-0.859	-0.160
w	0.220	-0.308	0.528	0.380	-0.073	0.452	0.557	0.152	0.405	0.630	0.203	0.427

Goods Time, n_g

	1959			*1979*			*1999*			*2019*		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.016	0.270	-0.286	0.078	0.391	-0.313	0.390	0.668	-0.278	0.757	0.956	-0.199
P_s	0.241	0.508	-0.267	0.063	0.288	-0.224	-0.285	-0.109	-0.176	-0.635	-0.474	-0.160
w	-1.272	-1.800	0.528	-1.066	-1.518	0.452	-0.785	-1.190	0.405	-0.503	-0.930	0.427

Services Time, n_s

	1959			*1979*			*1999*			*2019*		
Price	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$	ε^m	ε^h	$-\varepsilon^{inc}\vartheta$
P_g	-0.141	0.145	-0.286	-0.141	0.173	-0.313	-0.137	0.141	-0.278	-0.130	0.069	-0.199
P_s	0.924	1.192	-0.267	0.804	1.029	-0.224	0.665	0.840	-0.176	0.599	0.759	-0.160
w	-1.558	-2.086	0.528	-1.399	-1.851	0.452	-1.222	-1.627	0.405	-1.148	-1.575	0.427

There are two channels, not necessarily mutually exclusive, that could explain the *observed* Veblenization of goods. First, measurement error that fails to account for quality improvements to some durable goods, consistent with that found in the empirical analyses of Bils and Klenow (2001) and Bils (2009), might bias our elasticity estimates. Second, those same quality improvements that may or may not be accounted for in the data could be perceived by consumers as status symbols and thus luxuries, leading to actual Veblenization in response to technological improvements. Both explanations, however, are not mutually exclusive because they each suggest that technological improvements to the quality of goods have driven changes to consumer behavior.

Consider the different kinds of products embedded in the BEA's separate goods and services expenditure series. While goods expenditure is dominated by non-durables, the durables share of total goods expenditure steadily rose from 26% in 1959 to 34% in 2019. Moreover, the durables sub-component of goods contains many categories of consumption associated with substantial technological growth and innovation, such as motor vehicles, household appliances, video and audio processing equipment, televisions, computers, and phones. Given technological advancement in these consumption categories has been swift, quantifying price and quantity levels for these categories is also notoriously difficult, so that quality gains associated with these categories often go years without being fully integrated into the published quantity and price series.[41] It is plausible that both measurement error and actual Veblenization are simultaneously driving the observed trend in the goods own-price elasticity and also second-order dampening of structural change.

3.4 The Relative Efficiency of Using Goods Versus Services

The results from the last section suggest that possible underlying and unmeasured quality improvements to goods may be contributing to our results. But these results do not actually account for unobserved productivity changes. What, exactly, do aggregate data via the model's lens suggest about how z_g/z_s has changed over time?

We now examine the model-implied parameterized decomposition of $MRS(q_g, q_s)$ from (5). Again, we will pass PCE price indices along with average annual real wages through $MRS(q_g, q_s)$ in order to back out aggregate time series for $\frac{z_g}{z_s}$ from both the ATUS and PUMD models. The procedure we deploy is as follows.

First, we take aggregate real quantities ($2012) as data, specifically real goods and services outlay per-capita per-day. Note that we do not have a full time series of off-market

[41] This is what Robert Gordon refers to as the long "gestation period" after an invention before the returns to such an invention are accounted for in productivity and thus price data (Gordon 2016).

time use, so we pass our per-capita daily values of q_g and q_s, along with average hourly wages and prices, through the implicit function $n_j(q_j)$ featured in (11). We also assume that the model-implied equilibrium path corresponds exactly to the data-generating process, so that in equilibrium it must be that $MRS(q_g, q_s) = \frac{P_g}{P_s}$. Under this assumption, each period's relative in-home productivity can be readily backed out of a parameterized version of (5). The productivities we thus show here are model-implied productivities, assuming that real quantities per-capita per-day are exactly generated by a data-generating process described by the model, where aggregate goods and services prices and average hourly wages are taken as given.

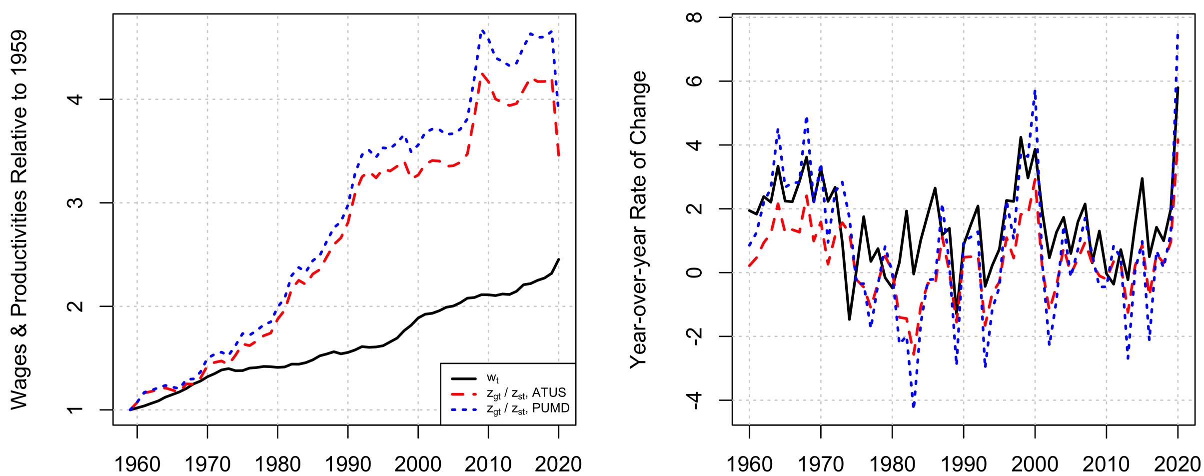

(a) Levels, $1959 = 1$ (b) Growth Rates

Figure 9: This figure presents posterior-predicted relative in-home productivities, $\frac{z_g}{z_s}$ and wages. In panel (a) we plot the levels of each time series, normalized to unity in 1959. In panel (b) we show the year-over-year growth rates of each time series. The legend is in panel (a).

Figure 9 show how growth in relative efficiency, normalized to unity in 1959, has outpaced wage growth in panel (a), while panel (b) shows the year-on-year growth rates. Recall, a rising time series for $\frac{z_g}{z_s}$ indicates that the unmeasured, consumption efficiency of goods has improved faster than that of services. We find that the relative consumption efficiency associated with using goods versus services increased by between 321.07% (ATUS) and 365.60% (PUMD) from 1959 to 2019 for the average wage earner, depending on the parameters. Second, increases in relative efficiency are mildly, negatively correlated with wage gains (Pearson's correlation coefficient of -0.27 with ATUS parameters and -0.22 with PUMD parameters). Indeed, wage gains were strongest from 1959-1970, but weakened in the 1970s and 1980s, despite relative productivity increasing by between

138% (ATUS) and 148% (PUMD) from 1970-1990. The correlation is weak though: wage gains were also sluggish in the 2010-2019 period, and so was relative productivity growth.

Estimates of z_g/z_s can help explain why the Veblenization phenomenon described in Section 3.3 is observed: over time, activities associated with using goods have become of increasingly higher quality, as measured by the home production efficiency of *using* such products. Consumers now can afford to engage in activities associated with more products of ever-higher quality than they could in the past, and market prices/quantities apparently fail to reflect the quality-adjusted value from increases to consumption efficiency implied by our model.

Note that the elasticity of x_{gt}/x_{st} with respect to z_{gt}/z_{st} is $\rho/(1-\rho)$, which is negative when $\rho < 0$ (i.e., our preferred models). That is, rising z_{gt}/z_{st} will force relative expenditure to decline, as consumers allocate relatively more expenditure toward the less-productive consumption activity. This is a direct result of the fact that c_g and c_s are complements. Thus, we estimate that z_{gt}/z_{st} is rising over time, and this is theoretically consistent with the observed decline in relative expenditure under our modeling assumptions.

3.5 The Distributional Implications of Structural Change

We conclude our quantitative exercises by asking, how much would we have had to pay consumers in 2019 to live in a world where the relative price of goods to services was that of the year 1959?[42] The goal is to understand how consumers across the wage distribution have either differentially benefitted or been harmed by changing input costs to home production (i.e., changing market prices). To answer this question we first compute the utility value of model-predicted final consumption for all consumers in the final year of both the ATUS and PUMD datasets. We then compute the utility value of consumption under a counterfactual relative-price regime and solve for the wage level required to make the consumer indifferent between the counterfactual regime and the model-predicted regime. We find that all 2019 consumers from the ATUS sample would have required at least a 10.86% increase to their hourly wages to live in a world characterized by 1959's goods to services relative price, while PUMD consumers would have required at least a 7.84% increase. In Figure 10 we can see that the percentage compensation required for consumers across the wage distribution appears to be slightly non-monotonic and that lower-income consumers have benefitted the most from structural change relative to their middle- and

[42] Analyses here are truncated at 2019 to focus on how structural change has affected welfare over the long-run. We thus chose to compare the last year before the 2020-2021 recession to the economic environment in 1959.

high-income peers.

(a) ATUS Parameters

(b) PUMD Parameters

Figure 10: Panel (a) plots the compensation percentages required to make consumers indifferent between a 2019 and 1959 relative price regime for ATUS parameters and consumers against observed log wages. Panel (b) presents the same plot for PUMD parameters and consumers with 2018 baseline prices.

This exercise demonstrates the degree to which consumers across the wage distribution have benefitted from structural change that has affected the relative market value of goods and services. The fact that services have become relatively more expensive as consumers have become richer is a *net positive* to welfare, as we estimate that consumer preferences are such that they appear to enjoy relatively cheaper goods more than they are negatively impacted by rising relative services prices. Why might this be? Consider, again, the kinds of goods that are associated with the strongest technological advancement: appliances, electronics, televisions, computers, and cell phones. Low-income consumers benefit as much as high-income consumers from having access to these technologies. The results suggest that the advanced technologies of the past (i.e., air conditioning, microwave ovens, washing machines, and dishwashers in 1959) may have been more out of reach for low-income consumers relative to the advanced technologies of today (i.e., computers, smartphones, and flat screen televisions).

We thus estimate that structural change has led to changes in the consumption environment that has enhanced the experienced utility of consumers of all incomes. This result, therefore, suggests that measuring inequality by looking only at income and wealth may not be capturing the whole story as to how consumers across the income distribution benefit from the advancement of consumer products. A model with off-market time utilization and consumption, like ours, which can quantify the *experienced* value of consumption suggests that the structural transformation of the U.S. economy over the last 60 years, while associated with increased income inequality, may actually have led to more equitable outcomes of *consumption experience.* This runs in contrast to results in Aguiar and Bils (2015) but helps confirm the finding from Krueger and Perri (2006), and more recently Pretnar (2022), that inequality of consumption is not as stark as income inequality.

4 Conclusion

We have shown that accounting for differential time-use complementarities in the consumption decision process can impact economic inference. This is especially true when considering which mechanisms are most responsible for the structural evolution of the U.S. economy from one previously dominated by the consumption of manufactured goods to today's service economy. The results suggest that the increase in the services share of expenditure appears to be a consequence of both efficiency gains in goods production that have driven down relative prices, as well as relative improvements to the efficiency of the consumption of goods versus services. Further, welfare differences between different income quintiles may not be as stark as income- or wealth-based measures would suggest, since it appears that both poorer and richer consumers, relative to middle-income consumers, have benefitted from the structural evolution of the U.S. economy.

References

Acemoglu, Daron and Veronica Guerrieri (2008). "Capital Deepening and Nonbalanced Economic Growth". *Journal of Political Economy* 116.3 (cit. on pp. 5, 34).

Aguiar, Mark and Mark Bils (2015). "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105.9, pp. 2725–2756 (cit. on p. 43).

Aguiar, Mark and Erik Hurst (2005). "Consumption versus Expenditure". *Journal of Political Economy* 113.5 (cit. on p. 4).

— (2007). "Life-Cycle Prices and Production". *The American Economic Review* 97.5 (cit. on p. 4).

— (2016). "The macroeconomics of time allocation". *Handbook of Macroeconomics*. Vol. 2. Elsevier, pp. 203–253 (cit. on pp. 8, 9).

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2012). "Recent Developments in the Economics of Time Use". *Annual Review of Economics* 4, pp. 373–397 (cit. on p. 20).

Baumol, William (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis". *The American Economic Review* 57.3 (cit. on pp. 5, 34).

Baumol, William, Sue Anne Batey Blackman, and Edward Wolff (1985). "Unbalanced growth revisited: asymptotic stagnancy and new evidence". *The American Economic Review* 75.4, pp. 806–817 (cit. on pp. 5, 34).

Becker, Gary (1965). "A Theory of the Allocation of Time". *The Economic Journal* 75.299 (cit. on pp. 1, 2, 8, 12).

Becker, Gary S. and Kevin M. Murphy (2000). *Social Economics: Market Behavior in a Social Environment*. Belknap Press (cit. on p. 16).

Betancourt, Michael and Leo Stein (2011). "The Geometry of Hamiltonian Monte Carlo". *arXiv:1112.4118* (cit. on p. 25).

Bils, Mark (2009). "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *The Quarterly Journal of Economics* 124.2, pp. 637–675 (cit. on p. 39).

Bils, Mark and Peter J. Klenow (2001). "Quantifying Quality Growth". *American Economic Review* 91.4, pp. 1006–1030 (cit. on p. 39).

Blundell, Richard and Thomas Macurdy (1999). "Labor Supply: A Review of Alternative Approaches". *Handbook of Labor Economics*. Ed. by O. Ashenfelter and D. Card. Vol. 3. Elsevier Science. Chap. 27 (cit. on p. 12).

Boerma, Job and Loukas Karabarbounis (2021). "Inferring Inequality with Home Production". *Econometrica (forthcoming)* (cit. on pp. 9, 24, 25).

Boppart, Timo (2014). "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences". *Econometrica* 82.6 (cit. on pp. 1, 4, 5, 8, 28, 34).

Bridgman, Benjamin (2016). "Home productivity". *Journal of Economic Dynamics and Control* 71, pp. 60–76 (cit. on p. 1).

Bridgman, Benjamin, Georg Duernecker, and Berthold Herrendorf (2018). "Structural transformation, marketization, and household production around the world". *Journal of Development Economics* 133, pp. 102–126 (cit. on p. 1).

Buera, Francisco and Joseph Kaboski (2009). "Can Traditional Theories of Structural Change Fit the Data?" *Journal of the European Economic Association* 7.2-3, pp. 469–477 (cit. on p. 2).

— (2012). "The Rise of the Service Economy". *The American Economic Review* 102.6 (cit. on p. 2).

Caselli, Francesco and Wilbur Coleman (2001). "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation". *Journal of Political Economy* 109.3 (cit. on p. 1).

Comin, Diego, Danial Lashkari, and Martí Mestieri (2021). "Structural Change with Long-run Income and Price Effects". *Econometrica* 89.1, pp. 311–374 (cit. on pp. 1, 4, 5, 34).

Deaton, Angus and John Muellbauer (1980). *Economics and Consumer Behavior*. Cambridge University Press: New York (cit. on p. 10).

Fang, Lei, Anne Hannusch, and Pedro Silos (2022). "Bundling Time and Goods: Implications for Hours Dispersion". *Working Paper* (cit. on p. 4).

Gelman, Andrew et al. (2013a). "Computation in R and Stan". *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. Appendix C (cit. on p. 25).

— (2013b). "Computationally efficient Markov chain simulation". *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. 12 (cit. on p. 25).

Ghez, Gilbert R. and Gary S. Becker (1974). *The Allocation of Time and Goods Over the Life Cycle*. Columbia University Press: New York (cit. on p. 2).

Gordon, Robert J. (2016). *The Rise and Fall of American Growth*. Princeton University Press (cit. on pp. 9, 39).

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman (1988). "Investment, capacity utilization, and the real business cycle". *American Economic Review* 78.3, pp. 402–417 (cit. on p. 11).

Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi (2013). "Two Perspectives on Preferences and Structural Transformation". *American Economic Review* 103.7 (cit. on pp. 1, 4, 5, 8, 10, 34).

— (2014). "Growth and structural transformation". *Handbook of economic growth*. Vol. 2. Elsevier, pp. 855–941 (cit. on p. 34).

Kehoe, Timothy, Kim Ruhl, and Joseph Steinberg (2018). "Global Imbalances and Structural Change in the United States". *Journal of Political Economy* 126.1, pp. 761–796 (cit. on p. 1).

Kendall, M.G. (1975). *Rank Correlation Methods*. 4th ed. Charles Griffin: London (cit. on p. 7).

King, Robert, Charles Plosser, and Sergio Rebelo (1988). "Production, Growth and Business Cycles I. The Basic Neoclassical Model". *Journal of Monetary Economics* 21, pp. 195–232 (cit. on p. 11).

Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie (2001). "Beyond Balanced Growth". *Review of Economic Studies* 68, pp. 869–882 (cit. on pp. 1, 2, 4, 5, 34).

Krueger, Dirk and Fabrizio Perri (2006). "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory". *Review of Economic Studies* 73.1, pp. 163–193 (cit. on p. 43).

Kuznets, Simon (1966). *Economic growth and structure*. Heinemann Educational (cit. on pp. 5, 34).

— (1973). "Modern Economic Growth: Findings and Reflections". *American Economic Review* 63.3, pp. 247–258 (cit. on pp. 5, 34).

Kydland, Finn and Edward Prescott (1982). "Time to Build and Aggregate Fluctuations". *Econometrica* 50.6 (cit. on p. 11).

Lebow, David E. and Jeremy B. Rudd (2021). "Measurement Error in the Consumer Price Index: Where Do We Stand?" *FEDS Working Paper No. 2001-61* (cit. on p. 9).

Mann, H.B. (1945). "Nonparametric tests against trend". *Econometrica* 13, pp. 245–259 (cit. on p. 7).

Matsuyama, Kiminori (2009). "Structural Change in an Interdependent World: A Global View of Manufacturing Decline". *Journal of the European Economic Association* 7.2-3, pp. 478–486 (cit. on p. 1).

Neal, Radford (2011). "MCMC using Hamiltonian dynamics". *Handbook of Markov Chain Monte Carlo*. Ed. by Steve Brooks et al. Chapman & Hall. Chap. 5 (cit. on p. 25).

Ngai, Rachel and Barbara Petrongolo (2017). "Gender Gaps and the Rise of the Service Economy". *American Economic Journal: Macroeconomics* 9.4 (cit. on p. 10).

Ngai, Rachel and Christopher Pissarides (2007). "Structural Change in a Multisector Model of Growth". *American Economic Review* 97.1 (cit. on pp. 5, 34).

Pretnar, Nick (2022). "Measuring Inequality with Consumption Time". *Working Paper* (cit. on pp. 4, 43).

Uy, Timothy, Kei-Mu Yi, and Jing Zhang (2013). "Structural change in an open economy". *Journal of Monetary Economics* 60, pp. 667–682 (cit. on p. 1).

Veblen, Thorstein B. (1899). *The Theory of the Leisure Class: An Economic Study of Institutions*. London: Macmillan Publishers (cit. on p. 16).

Watanabe, Sumio (2010). "Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory". *Journal of Machine Learning Research* 11, pp. 3571–3594 (cit. on pp. 29, 30).