

Structural Change with Time to Consume

W.L. Bednar* Nick Pretnar^{23†}

²*Laboratory for Aggregate Economics and Finance, UCSB*

³*Tepper School of Business, CMU*

October 23, 2025[‡]

Abstract

We explore the relationship between household time-use and the long-run rise in the services share of expenditure. Using the American Time Use Survey, we study how households allocate their off-market time between using goods and using services. While relative goods to services market expenditure has declined along the development path, relative off-market time spent using goods versus services has risen. To understand which of rising incomes, declining relative prices, or changing preferences is the primary driver of simultaneous expenditure and time use trends, we estimate a model where consumers have non-separable preferences over market consumption and multiple different off-market activities. We find no role for income effects driving structural change in our preferred specifications. Structural change is almost entirely attributable to both falling relative goods/services prices and the rising relative off-market efficiency of using goods versus services.

JEL Codes: D13, E21, N12, O33

Keywords: structural change, time use, services, income effects, growth, consumption technology

*wbednar101@gmail.com

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

[‡]This research was partially funded by NSF Grant #DGE1252522. We would like to thank Laurence Ales, Javier Birchenall, Benjamin Bridgman, Georg Duernecker, Lei Fang, Alex Horenstein, Finn Kydland, Diana Mikhail, Alan Montgomery, Alessio Moro, Hakki Özdenören, Peter Rupert, Ali Shourideh, Stephen Spear, and Kieran Walsh. We would also like to thank seminar participants at the X Workshop on Structural Transformation and Macroeconomic Dynamics in Cagliari, Italy, 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 multiple anonymous referees for helpful comments.

1 Introduction

What is the primary cause of the rising services share of expenditure? The question is still open, as a number of different theories can reconcile the data. One explanation in the literature attributes the rise to growing incomes, arguing that richer people prefer to buy relatively more services than goods. Yet, along the development path in most advanced economies, we have simultaneously observed *both* incomes and the relative price of services rise. Heretofore, whether the rising services share can be attributed to income effects or substitution effects has been a by-product of model specification.

This paper adds to the debate by also considering a new empirical fact: as the goods to services ratio of relative expenditure has fallen, the relative off-market time spent *using* goods versus services has risen. Leveraging the unique structure of the American Time Use Survey (ATUS) which allows for a convenient mapping of tasks to consumption activities, we show that the patterns of structural change associated with market expenditure (i.e., rising services share) are opposite of those associated with off-market time use (i.e., falling services share). That is, while the average consumer's share of wallet devoted to services has risen, the average consumer's share of off-market time devoted toward actually *using* market services has fallen. Consumers spend relatively less on goods but relatively more of their time using them, while the opposite is true for market services. Our approach thus fits within the mold of an emerging literature seeking to understand how variation in goods prices affects time use (see, e.g., Kopytov, Roussanov, and Taschereau-Dumouchel (2020), Aguiar et al. (2021), Boerma and Karabarbounis (2021), and Fang, Hannusch, and Silos (2025)), though we address the inverse relationship: how does off-market time use impact the allocation of goods and services?

Our empirical investigation inspires a model in which consumers allocate time toward consumption activities in order to derive utility from using different types of market products (i.e., goods and services) associated with such activities. We turn to the general, theoretical formulation of Becker (1965) who imagined each consumption decision as also being associated with a unique time-use decision. In our framework consumers allocate expenditure either to goods or services and allocate their off-market toward *using* either goods or services. Our formulation allows us to investigate how the allocation of off-market time toward different consumption activities is related to a consumer's allocation of expenditure toward different types of market products. The model we write down is rich and flexible: depending on whether goods and off-market time or services and off-market time are substitutes/complements, it can generate all possible patterns of structural change. Thus, with the right time use and expenditure data, the model can falsifiably address whether rising incomes (wages), falling relative goods/services prices, or changes to consumer preferences are most responsible for the simultaneous rise in the services share of expenditure and fall in the services share of off-market time.

In considering the role that time allocation plays in contributing to structural change, our work is closely linked to those of Ngai and Pissarides (2008), Bridgman (2016), and Bridgman,

Duernecker, and Herrendorf (2018). Indeed, similar to Bridgman (2016), we use a model to back out and study how off-market productivities (from the perspective of consumers) 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 add to the literature by considering 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. We estimate that goods have actually become relative more efficient to use off market than services.

We also contribute to the broader 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) 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) that attributing structural change to income effects could be a by-product of model mis-specifications that fail to account for how consumers use the products they buy. Additionally, Buera and Kaboski (2009) later suggested that exploring models with richer preference structures may provide alternative explanations for the causes of structural change.² 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 off-market activities. In our approach variation in wages, and thus income, can affect allocations through the interactions between time-use and consumption-expenditure decisions. We thus provide 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 in Section 2 by outlining several empirical facts regarding the long-run evolution of expenditure and off-market time-use patterns. 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 real quantities (indexes). We also provide a novel way to match time-use data with expenditure by exploiting unique properties of both the American Time Use

¹The original Becker (1965) was later explored in more detail in Ghez and Becker (1974).

²Buera and Kaboski (2009) specifically suggested considering the role that home production played in driving structural change.

Survey (ATUS) and older surveys. We show that since the 1970s off-market time that households spend using goods versus services has approximately doubled, which is consistent with more recent cross-sectional evidence that lower-income households spend a smaller share of wallet but a greater share of time using services.

In Section 3 we describe the model used to decompose the drivers of these trends. We estimate the model in Section 4 using both time-use and expenditure micro data in several separate and independent estimations. All of our estimates and quantitative results are robust to concerns regarding matching time-use activities to consumption expenditure. This is because we estimate the model once using *only* time-use micro data. We then estimate it again using *only* consumption expenditure data, and our parameter estimates and quantitative results are quantitatively similar. Because the estimation using only expenditure data yields similar results to the estimation using only time-use data, this validates *both* our time-use/expenditure matching criteria from Section 2 and our modeling assumptions, splitting the time allocation into two distinct activities each separately associated with a different type of expenditure.

In Section 5 we use the estimated parameters to quantify the degree to which the rise in the services share of spending and fall in the services share of off-market time use can be attributed to rising wages (i.e., income effects), declining relative goods to services prices, or changes to the underlying preference structure. We find that relative-price and preference variation are almost entirely responsible for structural change since the 1950s. In our preferred models we find minimal role for income effects in contributing to the rising services share of expenditure or falling services share of time.

2 Empirical Facts and Background

We present several well-established empirical facts regarding U.S. time-use and expenditure patterns. We focus on the U.S. economy because our exercise requires a representative panel of granular time-use micro data, which happens to be 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 proceed in this section by first discussing how to link activities from time-use data with the utilization of specific types of consumption expenditure. In that same subsection we also present several facts from the data regarding how consumers spend their off-market time utilizing different kinds of market purchases. We conclude this section by reviewing the literature on structural change, as it pertains to the rise in the services share of expenditure.

2.1 Time Use and Expenditure in Micro Data

Can time use patterns suggest to us anything as to what are the drivers of structural change? The idea from Becker (1965) is that a market expenditure decision is *not* the ultimate utility generator. Rather, consumers must also allocate their off-market time toward *using* the different products

they buy in order to derive final utility from them. This is true for *both* tangible goods and intangible services, though the consumption technologies associated with using goods versus services may differ.³ Indeed, we expect market goods and off-market time to be relatively more complementary than market services and off-market time, which we would expect to be relatively more substitutable.⁴

Given this motivation, we naturally consider two granular questions. First, can off-market activities in the time-use data be sufficiently classified so as to even link them to activities associated with primarily using market products from specific sectors (e.g., goods and services)? Second, if the data can be conveniently linked, what do we observe regarding how consumer time-use patterns, classified based on the type of consumption activity with which they are associated, have changed? The answer to the first question is “yes,” which we discuss below. We then provide evidence that the relative amount of off-market time consumers spend using market services has actually *fallen* over the last 20 to 50 years before concluding this section with a review of the structural change literature, as it pertains to variation in sectoral expenditure shares over time.

2.1.1 Classifying Time Diary Activities by Associated Consumption Type

We endeavor to understand how consumers allocate time toward tasks associated with different product types.⁵ In this vein our work is most closely associated to strands of the time-use literature that analyze the substitutability between market expenditure and off-market time, such as Aguiar and Hurst (2005), Aguiar and Hurst (2007), Aguiar et al. (2021), Fang, Hannusch, and Silos (2025), Pretnar (2025), and Albrecht, Phelan, and Pretnar (2025), though we are unique in focusing specifically on the tradeoff consumers face with regards to allocating time toward activities that utilize market goods versus those that utilize services.⁶

To assign off-market time use to activities associated with using either goods or services, we first must establish which expenditures correspond to those using goods or services. We follow the rubric of Boppart (2014) to classify data in the Consumer Expenditure Survey’s Public Use Micro Data (CEX-PUMD) from the Bureau of Labor Statistics (2023) by goods and services. Meanwhile, we get prices from Table 2.4.4 of the National Income and Product Accounts (NIPA), which are divided amongst three broad categories of consumption — non-durable goods, durable goods,

³Becker famously used the example of “attending a play” (an entertainment service). One does not derive utility from merely buying a ticket to a play but also must allocate time toward sitting and attending the play.

⁴That is, there are services that are clear time savers, like hiring a cleaning person or a landscaper. This is *not* to say that the consumption of such services requires zero time (i.e., market services and off-market time are *perfectly* substitutable): a cleaning person or landscaper must (however briefly) be provided information and guidance by the hiring party. For such services the market product and off-market time are considerably substitutable, though not necessarily perfectly so.

⁵Granular time-use datasets are often used to better understand how consumers allocate time toward tasks classified as market work, non-market work, and leisure. See, for example, the activity-classification rubric in Boerma and Karabarbounis (2021), which classifies off-market activities as either those which are laborious or those which are leisurely.

⁶A recent working paper also examines the link between expenditure and time-use allocations as they pertain to the drivers of structural change (Cossu et al. 2023). However, the authors borrow the classification rubric from Aguiar and Hurst (2007) which cannot be cleanly mapped to goods and services expenditure like the rubric we provide here.

and services (Bureau of Economic Analysis 2024). We combine non-durable goods and durable goods prices into a single price index for market goods by following the procedure outlined in On-line Appendix C of Herrendorf, Rogerson, and Valentinyi (2013) and discussed in Whelan (2002) to unwind and then re-compute the chain-weighted indices. With expenditure and price indices categorized, we can proceed to map activities from the time-use datasets to those associated with using goods and those associated with using services.

Table 1: Mapping Activities to Expenditure

Sector	Expenditure/ Activity	ATUS Codes	CEX Variables
Svcs.	Waiting Assoc. with Eating and Drinking	t110281 — t110289	FDAWAY
Svcs.	Education	t060101 — t069999	EDUCA
Svcs.	Prof., Pers. Care, and Household Svcs.	t080101 — t099999	PERSCA PERINS
	Travel Related to Pers. Care, Household Activities and Helping Others	t180101 — t180499	HEALTH SHELT CASHCO
	Telephone Calls	t160101 — t169989	UTIL
Svcs.	All Other Non-work Travel	t180601 — t189999	OTHVEH PUBTRA
Svcs.	Gov't Svcs. and Civic Obligations	t100101 — t109999	—
Svcs.	Religious and Spiritual Activities	t140101 — t149999	—
Svcs.	Volunteer Activities	t150101 — t159989	—
Goods	Pers. Care Activities Including Sleep	t010101 — t019999	TOTEXP less SVCS.
Goods	Household Activities	t020101 — t020999	TOTEXP less SVCS.
Goods	Caring for and Helping Both Household and Non-household Members	t030101 — t049999	TOTEXP less SVCS.
Goods	Consumer Purchases	t070101 — t079999	TOTEXP less SVCS.
Goods	Eating and Drinking	t110101 — t110199 t119999	TOTEXP less SVCS.
Goods	Socializing, Relaxing, Leisure, Sports, Exercise, and Recreation	t120101 — t139999	TOTEXP less SVCS.

Our time-use data primarily come from the American Time Use Survey (ATUS, 2003-2022)

from the Bureau of Labor Statistics (2024), though we cursorily also use two earlier time use datasets for illustrative purposes only: 1) the Americans’ Use of Time survey (AUT, 1965-1966) from Robinson (1977), and 2) the Time Use in Economic and Social Accounts survey (TUESA, 1975-1976) from Juster et al. (1992). Each of these datasets classifies a respondent’s daily time diary into myriad activities, and none of the time-use datasets are panels, so there is just one observation per respondent. In each dataset we select on respondents who satisfy all of the following criteria: 1) total time sums to 1,440 minutes (1 day); 2) both goods time and services time are strictly positive;⁷ 3) hourly wages are strictly positive; 4) work time is strictly positive.

For ATUS our selection procedure leaves us with 25,951 records over the years 2003 to 2022. ATUS gives us an average of 1,293.85 observations per year after cleaning, with the sample size declining slowly over time: we have 2,639 observations from the first year of the survey (2003), 1,346 by 2010, 895 in 2019, and only 624 in 2022. The AUT data yield 1,241 observations that span 1965 and 1966, while the TUESA data yield 812 observations that span 1975 and 1976.

While each of the ATUS, AUT, and TUESA have different activity classifications, in this section we focus our discussion on how we assign ATUS activities to those involving the utilization of goods versus those involving the utilization of services. This is because, outside the illustrative purposes of this section, AUT and TUESA are not used in any deeper analysis, since neither contain representation weights.⁸ Such an exclusion is problematic for computing time-use summary statistics by income quintile, so we only use AUT/TUESA data to compute economy-wide averages. Thus, we do not use AUT/TUESA data either to understand how agents of different income levels, on average, differ in terms of their time-use behavior nor in our structural estimations. We make our classification scheme for the AUT and TUESA datasets available in this paper’s Supplementary Materials which can be found at the author’s website.⁹

The ATUS classifies consumer time use into hundreds of different activities. In Table 1 we broadly show how we match various activity groups to CEX-PUMD expenditures using the Boppart (2014) rubric. Note that Boppart (2014) selects the services categories and then treats all other expenditure as goods. We do the same. Further, there are some time-use activities for which there are not clear expenditure categories in the CEX, like using governmental services (which are ostensibly free), using religious services or engaging in spiritual activities like meditation, or volunteering. Since these do, indeed, represent utilization of services, we classify the activities as such. Primarily, we do this because in the NIPA tables expenditure by “Non-profits serving households” is classified as services expenditure despite not having a micro-data equivalent entry

⁷We discuss the goods/services assignment further below.

⁸For our analyses requiring micro data (e.g., examining allocations of off-market time for respondents of different incomes and estimation of the structural model in Section 4) we use *only* the ATUS data to recover structural parameters, though we do provide out-of-sample assessments as to the model’s backwards predictability for AUT/TUESA data: it checks out.

⁹Specifically, for the AUT (1965-1966) and TUESA (1975-1976) data we use the merged dataset with common activity codes available from the Inter-university Consortium for Political and Social Research (ICPSR) as ICPSR 7796 (Robinson 1977; Juster et al. 1992). We classify tasks to roughly align with the classification of tasks in ATUS. The full AUT/TUESA dataset along with a table containing the goods/services classification rubric is available in the Supplemental Materials at the author’s website: <https://www.npretnar.com/research>.

in the CEX-PUMD.

The mapping described by Table 1 is actually quite readily apparent when inspecting the detailed descriptions of various ATUS activities. For example, ATUS explicitly records waiting times for many different categories of services (e.g., health care, child care, utilizing vehicular maintenance services, pet care, etc.). But what about labor-saving services whose purpose is to supplant what some call “home production” or “in-home labor”? Consider hiring a gardener: this is a non-trivial time-use entry in the ATUS that falls under “Using lawn and garden services,” which is ATUS category code t090401. Indeed, the designers of ATUS anticipated that consumers may actually *allocate time* toward using these kinds of services, classifying them in the broad category “Household Services (not done by self),” which is classification category t0901xx. The same would hold for monitoring a cleaning person hired to clean one’s home, which is denoted in ATUS as “Using interior cleaning services” (category t090101).

Finally, note that we include most personal care time, as well as sleep, in the goods series. This is because products most associated with these activities are physical goods (lotions, creams, soaps, gels, bedding, etc.). Further, the ATUS allows us to distinguish between time spent using personal care services (e.g., going to the barber or a massage therapist) separately from time spent using personal care goods. 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, in Technical Appendix A we exclude these activities from the goods series construction and provide alternative plots to those in the main text. There is no qualitative difference with respect to the evolution of goods/services time use when excluding these activities.

2.1.2 Time Use Over Time

Let n_{gt} represent average off-market time spent using goods in period t , and n_{st} be the same object for services.¹⁰ Figure 1 presents the ratio n_{gt}/n_{st} over both the long-run (1a) and short-run (1b). Several things are apparent at first glance. First, if one believes that the AUT/TUESA surveys closely match the ATUS, then goods time as a fraction of off-market time has risen since the 1960s and 1970s. Second, goods time as a fraction of off-market time continued to rise from 2003 to 2019, though possibly more slowly than in prior decades. Third, the COVID-19 pandemic was associated with a major spike in goods time relative to services time.

From 1965 to 1976 Americans spent an average of 3.03 hours per day using services and 7.54 hours per day using goods, excluding using goods associated with personal care activities and sleep. They spent an average of 8.78 hours per day on personal care activities associated with using goods at home and sleep. In the 1960s and 1970s Americans worked 4.65 hours per day.

By 2003 total time spent working had fallen from an average of 4.65 hours per day to 3.76 hours per day. Off-market time thus increased from the 1970s to 2003, though all of the increase went toward using more goods. In 2003 consumers spent an average of 2.07 hours per day using

¹⁰Since we have representation weights in the ATUS, we construct weighted averages from that dataset only. The AUT and TUESA observations are unweighted averages.




Figure 1: In (a) we show the combined AUT/TUESA/ATUS ratio of average off-market time spent using goods versus services, n_{gt}/n_{st} . 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 2003-2022.

services, 9.43 hours per day using goods not associated with sleeping and personal care, and 8.68 hours per day using personal care goods and sleeping. By 2019 services time had fallen still more to 1.95 hours per day, while non-personal care goods time had remained relatively flat at 9.31 hours per day, though sleep and personal care time had risen to 8.98 hours per day. To test whether there is a statistically significant trend in the ratio since, 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.

Given we have statistically verified that the ratio n_{gt}/n_{st} from ATUS is subject to a monotonic trend, we can ask if the average annual growth rate in n_{gt}/n_{st} we observe from 2003-2019 or 2003-2022 (including the COVID-19 years) is similar to the average annual growth in n_{gt}/n_{st} from 1976 to 2003. In 1975-1976 consumers spent 5.31 times as much off-market time using goods versus using services. In 2003 this ratio was 8.73, suggesting that, if measurements from AUT/TUESA are similar to those in ATUS then n_{gt}/n_{st} grew at an average annual rate of 1.79% from 1976-2003. Computing average annual growth rates for ATUS from 2003-2019 and 2003-2022, we get values of 0.43% and 1.03% respectively. These calculations suggest that either growth in the goods/services time-use ratio slowed between 1976 and 2003 or the surveys contain measurement differences, though both are possibilities. For these reasons we use the AUT/TUESA data strictly for illustrative purposes, relying on the ATUS for our structural estimations.

The COVID-19 pandemic saw major changes to the ways in which consumers allocate their time.¹¹ In 2020 services time was just 1.43 hours per day on average. This value increased to 1.73

¹¹Note that the survey's structure also changed as a result of the pandemic, which means that 2020 ATUS sampling only took place up to March 17, 2020, and after May 10, 2020, but not in between. This has (minor) implications for the interpretation of 2020 statistical weights in the dataset (Bureau of Labor Statistics 2021).

hours per day by 2022, though still not recovering to its pre-COVID levels. The spike in the ratio n_{gt}/n_{st} observed during 2020 is not just due to a decline in services time but also an increase in goods time. Indeed, while services time fell by half of an hour in 2020, goods time less personal care and sleep time increased by almost 1.5 hours to 10.78 hours per day. Personal care time spent using goods increased from 8.98 hours to 9.22 hours, as work time fell by approximately 1.2 hours from 3.76 hours per day in 2019 to 2.56 hours per day in 2020. By 2022 work time had still not recovered to its 2019 level, and goods time both excluding personal care and personal care time associated with using goods each remained elevated relative to 2019.

As we previously noted, Technical Appendix A contains plots where personal care and sleep time are excluded from the goods series. Exclusion of personal care time appears to amount to mostly a level shift in the time-use ratios. This pattern holds even during the pandemic. Because much of personal care time involves utilization of market goods, and because inclusion/exclusion of personal care time does not seem to qualitatively alter the time-use patterns we document in data, our main structural estimations in Section 4 use data that *includes* personal care and sleep time.¹²

2.2 Rising Services Share of Expenditure

The previous section demonstrates how the composition of off-market time as it pertains to market product utilization has changed over the last 20-50 years: generally speaking, relative goods to services utilization is stronger now than in the past, as consumers spend larger fractions of their off-market time using market goods today than they used to. In this section we now discuss how this fact relates to the well-established patterns of structural change as it pertains to goods versus services market expenditure.

It is well-documented that the services share of expenditure has risen in the U.S. since World War II. Figure 2 shows how, on aggregate, relative goods to services expenditure, prices, and real quantities have changed over the last 75-plus years.¹³ 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 1948-2023.¹⁴ These two facts have been frequently explored 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 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., declining long-run relative goods/services prices) rather than income effects. The findings in Boppart (2014) thus support

¹²In Technical Appendix E we also estimate our main structural model on data without personal care time: the results are the same.

¹³Aggregate expenditure, price, and quantity series are constructed using NIPA Tables 1.1.3 through 1.1.5 (Bureau of Economic Analysis 2024).

¹⁴Aggregate variables will be denoted with capital letters throughout the paper.

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. The puzzle as to the causes of structural change still exists.




Figure 2: 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 2017 = 1 is featured in (b); in (c) we feature relative real quantities in \$2017, Q_{gt}/Q_{st} . Aggregate expenditure, price, and quantity time series depicted here run from 1948-2023.

Upon inspecting Figure 2, it becomes apparent that declining relative prices have played a major role in driving down relative expenditure. Some have attributed the long-term rise in the services share to consumers' increasing desire for services as their incomes rise (Kuznets 1966; Kuznets 1973; Kongsamut, Rebelo, and Xie 2001; Herrendorf, Rogerson, and Valentinyi 2013; Comin, Lashkari, and Mestieri 2021). If this were true, we should expect the relative real quantity of goods to services to fall all along the development path. However, when inspecting the real-quantity ratios in Figure 2c, we see that real quantities of *goods*, not services, fall until 1990 before rising again, with only the Great Recession and the COVID-19 pandemic interrupting this more recent trend.

The 1990 break point is also associated with a deceleration of structural change. 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 5 percentage points from 1999-2019. Meanwhile, over the COVID-19 pandemic the services share fell by 2% (2019-2023). For structural change prior to 1990, a crude but intuitive story whereby income effects have helped drive 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 a role for income effects in contributing to structural change. However, after 1990 the story would have to change, given relative quantities rise while relative prices fall. This suggests a strong role for substitution effects in driving structural change. This is because, after 1990 and outside of recessionary periods, we observe real demand for goods growing at a faster rate than real demand for services at the same

time that relative goods/services prices continue to fall. Further, the fact that real demand for goods falls around the most recent recessions also calls into question the story that income effects have primarily driven the rise of the services share: we should expect, instead, that if services consumption is more sensitive to income variation, relative real services (not goods) consumption would fall (not rise) during recessions and rise (not fall) during expansions.

The results, here, compared with the time use results from Section 2.1.2, show that structural change in terms of *real* market consumption looks quite similar to structural change in terms of off-market time utilization, at least since the early 2000s. With respect to activities which use goods, utilization of both market inputs (Q_{gt}) and off-market time (n_{gt}) has been on the rise relative to services. In fact, despite the ratio Q_{gt}/Q_{st} having risen, this is due to Q_{gt} rising *faster* than Q_{st} , not Q_{st} falling. But when looking at rising n_{gt}/n_{st} , services time has been stagnant and even declined in recent years. In a Becker (1965) style model in which both market quantities, Q , and off-market time, n , are used in the production of some kind of final consumption good (or activity), c , at a minimum we should expect the input elasticities of substitution for goods-adjacent activities and services-adjacent activities to be different. Therefore, any model of joint market consumption and off-market time use decisions that aims to reconcile these patterns will need to be sufficiently general so as to allow for different input and output elasticities across different kinds of consumption activities. We now present such a model.

3 Model

The model operates at the household level where, as in Herrendorf, Rogerson, and Valentinyi (2013), we assume households are atomic price-takers and focus on a repeated, static model of household decision making that abstracts from investment decisions.¹⁵ Each household i operates as a collective that buys quantities of goods and services, q_{ij} , $\forall j \in \{g, s\}$, from the market at prices, P_j .¹⁶ They allocate time n_{ij} toward using these purchases in activities, $f_j(q_{ij}, n_{ij})$, to generate final consumption experiences, $c_{ij} = f_j(q_{ij}, n_{ij})$, from which they ultimately derive flow utility, $u(c_{ig}, c_{is})$. Utility is a constant elasticity of substitution (CES) function where both final consumption experiences associated with using goods, c_{ig} , and using services, c_{is} , are given equal weight,¹⁷

$$u(c_{ig}, c_{is}) = (c_{ig}^\rho + c_{is}^\rho)^{\frac{1}{\rho}}. \quad (1)$$

¹⁵This is possible under the same two-stage budgeting assumption required by Herrendorf, Rogerson, and Valentinyi (2013): if intertemporal preferences are assumed to be additively (i.e., strongly) separable then the intratemporal (within-period) allocation of resources toward immediate, utility-generating outcomes is independent of future allocations (Deaton and Muellbauer 1980).

¹⁶The time subscript, t , is dropped to avoid notational clutter. Further, the data we will deploy are cross-sectional in nature: we do not have panels of time-use and expenditure micro data. Aggregate trends we will back out from our quantitative exercises will involve averaging over different cross-sectional panels, year by year. Thus, we maintain the household indices, i , since these will be our units of observation, while encouraging the reader to keep in mind that $i \equiv i(t)$ is time-dependent, as are prices.

¹⁷We give each experience, c_{ij} , equal weight in preferences for identification purposes, which will become more apparent below.

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

As in Becker (1965) and discussed in depth in Aguiar and Hurst (2016), our setup distinguishes between final consumption, c_{ij} , and market expenditure, $x_{ij} := P_j q_{ij}$.¹⁸ Final consumption is the experience that results from combining off-market time and market inputs (i.e., goods and services) together in the consumption technology function, $f_j(q_{ij}, n_{ij})$, in order to achieve a final outcome, c_{ij} , from which a consumer derives ultimate utility.¹⁹ To produce a final consumption experience, c_{ij} , consumers must allocate off-market time, n_{ij} to use the market inputs, q_{ij} , that they have purchased. Note that final consumption experiences, c_{ij} 's, are non-tradable in that markets do not exist for them to be bought and sold, despite the fact that the inputs q_{ij} and n_{ij} into the production of c_{ij} 's are tradable.²⁰

Each market input q_{ij} (i.e., goods or services) is associated with a unique time allocation, n_{ij} , and a unique final consumption experience, c_{ij} . Experiences are generated by combining q_{ij} and n_{ij} in consumption technology functions, which we allow to take the following, flexible CES parameterization:

$$c_{ij} = z_{ij} f_j(q_{ij}, n_{ij}) := z_{ij} \left(\omega_j q_{ij}^{\nu_j} + (1 - \omega_j) n_{ij}^{\nu_j} \right)^{\frac{1}{\nu_j}}, \quad \forall j \in \{g, s\}. \quad (2)$$

Note that goods and services final consumption experiences — c_{ig} and c_{is} — are each associated with a unique technology. Indeed, differences across goods and services with respect to the structures of $f_j(\cdot, \cdot)$ will play a crucial role in determining the degree to which relative expenditure, x_{ig}/x_{is} , responds to either wage growth or changing relative prices over time. Thus, we allow for each technologies' production weight, ω_j , and substitution elasticity, $1/(1 - \nu_j)$ to be unique. Finally, each household is associated with two final consumption productivity parameters, z_{ig} and z_{is} . In Section 4 we will make assumptions regarding the structures of these parameters' statistical distributions in order to allow us to estimate the model and identify the structural parameters ρ , ω_g , ω_s , ν_g , and ν_s from time-use (ATUS) and expenditure (CEX-PUMD) micro data.

The z 's are exogenous productivities that capture the efficiency by which a household can transform market purchases and off-market time into final consumption experiences. Variation in these productivity/efficiency parameters both cross-sectionally and over time may be driven by several different factors. Each z_{ij} captures all variation in the allocation of market expenditure and off-market time that cannot be attributed to price and wage variation. Such variation may be simply unobserved, idiosyncratic cross-sectional heterogeneity, or variation may be systematic,

¹⁸We use x_{ij} to denote price $\times$ market quantity (i.e., expenditure).

¹⁹As some examples the activity "reading a book" requires a book (a good) and time. Meanwhile, the activity "eating out" requires a restaurant (a service) and time (e.g., time spent waiting, eating, etc.). Consuming a meal at home requires that one follow a recipe to combine groceries into a final, prepared meal using in-home durables and time. Or, as in Aguiar et al. (2021), to enjoy a video game, one must spend time actually playing it.

²⁰Boerma and Karabarbounis (2021) suggest that households function like island economies that produce final consumption which cannot be traded, despite that some intermediate inputs, like market goods/services and time, are traded between "islands."

affecting all households in one period more than another. Systematic variation in z 's may be due to a one-time shock, or it could represent a permanent change in the nature of the consumption environment.

To understand how z 's may represent idiosyncratic heterogeneity, consider the example of a professional chef cooking a meal at home versus a lay person. Suppose both the chef and the lay person allocate the exact same amount of time and the exact same amount of market products toward meal preparation. Yet, we would expect that even with the same allocation of inputs, the *quality* of the in-home meal prepared by the professional chef would exceed the quality of the in-home meal prepared by the lay person. Since experiences associated with home-prepared meals are goods experiences (e.g., physical market inputs are required for their production), the quality differential implies that $c_{\text{chef},g} > c_{\text{lay person},g}$. If q_{ig} and n_{ig} are equal for both $i = \text{chef}$ and $i = \text{lay person}$, then the quality differential would imply $z_{\text{chef},g} > z_{\text{lay person},g}$. In this manner the interpretation of each z_{ij} is similar to the interpretation of the productivity parameter in the highly stylized model of Boerma and Karabarbounis (2021), except as applied to activities associated with using goods and services.

Second, systematic changes in the aggregate consumption environment may affect the consumption technology productivities for everyone. For example, regulations requiring restaurants to prohibit smoking make the experience of sitting in a restaurant and eating more enjoyable than in the past. Let c_{is}^{before} be the consumption experience associated with eating in a restaurant before a smoking ban takes place, and let c_{is}^{after} be the experience after the ban. Suppose that restaurants' menus are unchanged by the ban (i.e., prices and market quantities are the same), and that the quality of service they otherwise provide is also unchanged. Further, suppose the time a household spends in the restaurant before and after the ban is the same, so that they wait the same amount of time for their meal in both periods. Then, q_{is} and n_{is} are unchanged between periods, but the quality of the consumption experience has improved because the atmosphere of the restaurant after the smoking ban is less stuffy and generally less toxic. This implies that $c_{is}^{\text{after}} > c_{is}^{\text{before}}$. On average consumers enjoy their restaurant experience more, which amounts to an increase in c_{is} relative to the period before the smoking ban. In this manner a policy change has affected the consumption environment for everyone on aggregate. All else equal, this implies that $z_{is}^{\text{after}} > z_{is}^{\text{before}}$ for all households, i .

It remains only to complete the exposition of the household's problem. Households trade time on the market by supplying hours, ℓ_i , and receiving efficiency wages, w_i , in return. Time use must satisfy:

$$\ell_i + n_{ig} + n_{is} \leq \bar{n}_i \quad (3)$$

where $\bar{n}_i$ is the households' time endowment, which will vary depending on the number of adults in the household.²¹ We assume households operate as a collective, so that a household with one

²¹For example, for a household with one adult, $\bar{n}_i = 24$ hours, whereas for a household with two adults, $\bar{n}_i = 48$ hours.

adult is a single consumption and time-use decision-maker, as is a household with two adults. We make this assumption due to data constraints in the CEX-PUMD and take no stand on how consumption and time allocations are distributed across adults.^{22,23}

The household budget constraint can be written

$$\sum_{j \in \{g,s\}} P_j q_{ij} \leq w_i \ell_i + y_i, \quad (4)$$

where y_i is non-labor income net of savings. As we show below, our estimator is formed entirely around the intratemporal marginal rates of substitution between relative market expenditure and off-market time use. As these conditions are independent of y_i , we only include it in (4) for completeness.

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.²⁴ Given y_i and prices in each period, households choose the vector $\{q_{ig}, q_{is}, \ell_i, n_{ig}, n_{is}\}$ so as to solve:

$$\begin{aligned} & \max u(c_{ig}, c_{is}) \\ & \text{subject to } (2), (3), \text{ and } (4). \end{aligned} \quad (5)$$

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 that off-market time could be allocated to activities that do not involve utilization of any market products. Given that there are few examples of time spent outside of work that is not associated with the utilization of some market purchase, we expect any bias associated with this assumption to be insignificant. Indeed, leisure activities in so far as they are recreational in nature almost always involve utilizing some kind of recreational good or service anyway. Think, for example, about video gaming as in Aguiar et al. (2021), which

²²While we observe how respondents of the ATUS allocate their time toward many different activities, and, to a lesser extent, we also observe how the spouses of *some* respondents allocate their time, in the CEX-PUMD we only observe household-level market expenditure. We thus have no way to split expenditure amongst different household members. This necessitates a model where households are treated as collectives, and time endowments then vary depending on the number of adults in the home.

²³For 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).

²⁴See Deaton and Muellbauer (1980) for a deep-dive into the mechanics of two-stage budgeting. As previously mentioned, 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.

requires an electronic game console, or reading a book, which requires an actual book. Other examples include running in the park, which requires tennis shoes, or attending a football game, which requires the purchase of a ticket to the event from a service provider.

3.1 Allocations of Expenditure and Off-market Time

In this section we derive the conditions which characterize a household's allocations of market expenditure and off-market time. Assume that $v_j \neq 0, \forall j \in \{g, s\}$, so that neither of the consumption technologies are Cobb-Douglas.²⁵ Also, assume that either $\rho < 0$ or $\rho \in (0, 1)$, so that $1/\rho$ is finite, ruling out Cobb-Douglas utility. This assumption on ρ ensures that all variation in expenditure and time-use allocations that cannot be attributed to price and wage variation can instead be attributed to the model's structural residuals, z_{ig} and z_{is} .²⁶

After solving (5) subject to the time-use and budget constraints, we get parameterized expressions for the marginal rates of technical substitution (MRTS) for each consumption process, $j \in \{g, s\}$:

$$\left(\frac{\omega_j}{1 - \omega_j} \right) \left(\frac{q_{ij}}{n_{ij}} \right)^{v_j - 1} = \frac{P_j}{w_i}, \quad \forall j \in \{g, s\}. \quad (6)$$

Expression (6) characterizes the marginal rate of substitution between market products, q_{ij} , and off-market time, n_{ij} , which is unique to either the goods or services consumption experience. Because the consumption technologies are CES, the market-input/off-market-time ratio is log-linear in the ratio of input prices. Because the consumption technologies are neutral a la Hicks (1932), the productivity parameters, z_{ij} , do not enter into the MRTS conditions.²⁷

The cross-category marginal rate of substitution (MRS) for market inputs q_{ig} and q_{is} is

$$\begin{aligned} & \left(\frac{c_{ig}}{c_{is}} \right)^{\rho - 1} \frac{z_{ig} \omega_g q_{ig}^{v_g - 1}}{z_{is} \omega_s q_{is}^{v_s - 1}} \left(\omega_g q_{ig}^{v_g} + (1 - \omega_g) n_{ig}^{v_g} \right)^{\frac{1 - v_g}{v_g}} \\ & \times \left(\omega_s q_{is}^{v_s} + (1 - \omega_s) n_{is}^{v_s} \right)^{\frac{v_s - 1}{v_s}} = \frac{P_g}{P_s}. \end{aligned} \quad (7)$$

To get $MRS(n_{ig}, n_{is})$, which characterizes the goods/services off-market time use tradeoff, we

²⁵When we estimate the model, we find that this assumption holds under the data: that is, we can reject the hypothesis that the data-generating process features Cobb-Douglas consumption-technology functions.

²⁶When $\rho = 0$, relative goods/services market expenditure and relative goods/services off-market time use are independent of the consumption-technology productivities (i.e., the structural residuals). To estimate models with Cobb-Douglas utility one would thus have to introduce non-structural measurement error and could not identify z 's.

²⁷Recall, Hicks-neutrality implies that technological change does not affect the balance of inputs into production.

derive a similar expression to that in (7):

$$\begin{aligned} \left(\frac{c_{ig}}{c_{is}}\right)^{\rho-1} \frac{z_{ig}(1-\omega_g)n_{ig}^{\gamma_g-1}}{z_{is}(1-\omega_s)n_{is}^{\gamma_s-1}} \left(\omega_g q_{ig}^{\gamma_g} + (1-\omega_g)n_{ig}^{\gamma_g}\right)^{\frac{1-\gamma_g}{\gamma_g}} \\ \times \left(\omega_s q_{is}^{\gamma_s} + (1-\omega_s)n_{is}^{\gamma_s}\right)^{\frac{\gamma_s-1}{\gamma_s}} = 1. \end{aligned} \quad (8)$$

Notice that the right-hand side of (8) is equal to unity. This follows from the fact that households' off-market time-use decisions are all constrained by (3) and that the opportunity cost households face with respect to allocating their time away from labor for any arbitrary off-market task is simply w_i .²⁸

We can re-arrange (6) to get an implicit function, $n_{ij}(q_{ij})$, that is linear in q_{ij} :

$$n_{ij}(q_{ij}) = q_{ij} \left[\frac{w_i \omega_j}{P_j(1-\omega_j)} \right]^{\frac{1}{\gamma_j-1}}, \quad \forall j \in \{g, s\}. \quad (9)$$

Using (9), we can exploit the homogeneity of $f_j(\cdot, \cdot)$ to substitute out instances of n_{ij} featured in $c_{ij} = z_{ij} f_j(q_{ij}, n_{ij})$ from the left-hand side of (7). By exploiting this envelope condition, we can then re-arrange (7) and re-write it in terms only of relative expenditure, $x_{ig}/x_{is} = (P_g/P_s) \times (q_{ig}/q_{is})$, market prices, wages, and parameters:

$$\begin{aligned} \frac{x_{ig}}{x_{is}} = \left(\frac{P_g}{P_s}\right)^{\frac{\rho}{\rho-1}} \left(\frac{z_{ig}}{z_{is}}\right)^{\frac{\rho}{1-\rho}} \left(\frac{\omega_g}{\omega_s}\right)^{\frac{1}{1-\rho}} \left(\omega_g + (1-\omega_g) \left[\frac{w_i \omega_g}{P_g(1-\omega_g)} \right]^{\frac{\gamma_g}{\gamma_g-1}}\right)^{\frac{\gamma_g-\rho}{\gamma_g(\rho-1)}} \\ \times \left(\omega_s + (1-\omega_s) \left[\frac{w_i \omega_s}{P_s(1-\omega_s)} \right]^{\frac{\gamma_s}{\gamma_s-1}}\right)^{\frac{\rho-\gamma_s}{\gamma_s(\rho-1)}}. \end{aligned} \quad (10)$$

Notice that the time allocation variables, n_{ig} and n_{is} , no longer feature in (10). Therefore, we can use this equation to estimate the model's structural parameters for only market expenditure data (e.g., from the CEX-PUMD), prices, and wages. Similarly, by inverting (9), we can also exploit the homogeneity of $f_j(\cdot, \cdot)$ to substitute out instances of q_{ij} in (8) to get an expression for relative

²⁸ Another way to think about this is that the implicit price the household pays to spend time on n_{ig} is equal to the price paid on n_{is} . A consequence of this is that they always set

$$z_{ig} \frac{\partial u}{\partial c_{ig}} \frac{\partial f_g}{\partial n_{ig}} = z_{is} \frac{\partial u}{\partial c_{is}} \frac{\partial f_s}{\partial n_{is}}.$$

Thus, allocations of off-market time equate productivity-weighted marginal utilities of final consumption experiences and marginal products of off-market time use between the two consumption experience categories, $j \in \{g, s\}$.

off-market time use that does not feature market expenditure:

$$\frac{n_{ig}}{n_{is}} = \left(\frac{z_{ig}}{z_{is}} \right)^{\frac{\rho}{1-\rho}} \left(\frac{1-\omega_g}{1-\omega_s} \right)^{\frac{1}{1-\rho}} \left(\omega_g \left[\frac{w_i \omega_g}{P_g(1-\omega_g)} \right]^{\frac{\nu_g}{1-\nu_g}} + 1 - \omega_g \right)^{\frac{\nu_g - \rho}{\nu_g(\rho-1)}} \times \left(\omega_s \left[\frac{w_i \omega_s}{P_s(1-\omega_s)} \right]^{\frac{\nu_s}{1-\nu_s}} + 1 - \omega_s \right)^{\frac{\rho - \nu_s}{\nu_s(\rho-1)}}. \quad (11)$$

We will use (10) and (11) as our primary estimating equations in Section 4.

By examining (10) and (11) it is apparent that differences between ρ , ω_j 's, and ν_j 's will determine the degree to which wage variation affects expenditure allocations. 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 final consumption 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. Further, the values of $\nu_g - \rho$ and $\rho - \nu_s$ will determine how much expenditure and time-use allocations actually respond to variation in the wage/price ratios. As can be seen in (10), the sign of ρ will also govern whether relative expenditure is increasing or decreasing in relative market prices.

As long as one of $\nu_g \neq \nu_s \neq \rho$, there will be non-trivial wage effects on the allocations of both expenditure and time. This is because, while the model is indeed nested CES with q_j and n_j for the same type of consumption j in the same inner-nest, q_g and q_s (for example) reside in nests with different elasticities, governed by ν_g and ν_s . This means that the elasticity of substitution between q_g and q_s (n_g and n_s , too) actually *depends* non-trivially on the wage except in extreme cases (e.g., $\nu_g = \nu_s = \rho$). Consider the log of (10), differentiated in the log ratio of P_g/P_s : since both P_g and P_s enter into the large nested objects along with w on the right-side of (10) in non-log-linear fashion, the elasticity of substitution between q_g and q_s will be generally *variable* and non-constant.

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 even if the household's objective function has a homothetic structure, as it does in the full Beckerian case; and 2) the time allocation will also depend on the relative price of associated market purchases. This is because w_i , P_g , and P_s are all prices, even though w_i also scales income. In Technical Appendix B.1 we formally show, for a model with more general functional forms, that wage variation amounts to *both* an income *and* a substitution effect on market demand. Meanwhile, in Technical Appendix B.2 we introduce an elasticity — the elasticity of relative expenditure with respect to wages — to show that the allocation of consumption expenditure varies non-trivially in wages except under knife-edge conditions.

4 Estimation

We recover parameters using U.S. micro data from both ATUS and the CEX-PUMD, running separate estimations depending on which dataset we target.^{29,30} Because we run two separate and independent estimations — one using ATUS data on the relative time-use equation and another using PUMD data on the relative expenditure equation — our results are robust to concerns pertaining to issues matching time-use allocations with expenditure activities.³¹ Our estimation is also robust to concerns that some time-use allocations cannot be cleanly split up so as to correspond strictly to a goods or services activity. This is because we find that our parameter estimates using only expenditure data (CEX-PUMD) are similar to those estimates extracted from a likelihood function formed *only* around the household’s relative time-use condition (ATUS). If these estimates were substantially different, mis-classification concerns with regards to time use data would be warranted, but they are not.

We estimate the model eight different times — four times with the ATUS and four times with the PUMD. We have two different versions of the ATUS dataset: 1) our main dataset, which includes personal care and sleep time; 2) another dataset where we strip out personal care and sleep time. We have only one version of the PUMD, but estimators are sensitive to prior specifications which differ depending on whether or not personal care time is included in the ATUS. We engage in separate estimations, stripping out personal care time (which includes sleep) in order to understand the degree to which its inclusion in n_{ig} affects our parameter estimates.³² We find that it has no significant effect. With each dataset we also separately restrict either $\rho < 0$ or $\rho \in (0, 1)$, and then we assess which model provides the better fit.

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.³³ 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 computa-

²⁹That is, we do not synthetically link the two datasets, as is done in Boerma and Karabarbounis (2021) and Pretnar (2025).

³⁰We do not use the AUT or TUESA data in our estimations, given the datasets lack representation weights.

³¹After cleaning, the PUMD micro dataset has 613,330 household-level observations from the “Interview” component of the survey.

³²Estimates of the structural model with personal care time excluded are featured in Technical Appendix E.

³³We 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).

tionally robust and require substantially less coding time in order to implement.³⁴

While it is common in the structural change literature, and other quant-macro literature more generally, to match model outputs to aggregate moments and/or time series, we follow the quantitative strategy of 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 which is also true for models with non-homothetic preferences and inelastic time use.³⁵ Estimating on granular micro data, under the assumption that all households are price takers, avoids the aggregation problem which would need to be addressed if calibrating to aggregate series.³⁶ The model-fit criteria we thus employ involve assessing how the model fits the 2003-2022 ATUS and 1990-2022 CEX-PUMD micro data, separately. A likelihood-based estimation method allows us to exploit the entire extent of variation and representation weights in the micro data when recovering structural parameters.

We re-introduce the t subscript in this section. Let $\zeta_{i(t)}^k = \frac{\rho}{1-\rho} \ln(z_{i(t)g}/z_{i(t)s})$ where $k \in \{n, x\}$ encodes whether the observation associated with the residual term, $z_{i(t)g}/z_{i(t)s}$, is from the ATUS or PUMD dataset. The agent-specific, elasticity-weighted relative log-productivities are

$$\zeta_{i(t)}^k = \frac{\rho}{1-\rho} (\gamma_{z_g/z_s}(t - \bar{t})) + \varphi_t^k + \eta_{i(t)}^k, \quad \forall k \in \{n, x\}, \quad (12)$$

where φ_t^k is a year-specific fixed effect term common to all in the sample, while $\eta_{i(t)}^k$ is an idiosyncratic error, assumed to be mean zero. Note that household indices depend on the year of observation, $i(t)$. The implied consumption productivity ratio is $z_{i(t)g}/z_{i(t)s} = \exp\{\zeta_{i(t)}^k(1-\rho)/\rho\}$. The sign of ρ will help determine whether relative productivities grow or decline over time, on average.

We form the empirical likelihood function around logged versions of equations (10) and (11)

³⁴This 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.

³⁵Recall, 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.

³⁶Aggregation problems, however, are likely very minimal, as we discuss in Section 5. This is because normalized wage distributions by dataset are stable over time. Still, we proceed with a likelihood-based model, estimated on granular micro data to avoid any possible aggregation problems while recovering parameters.

separately for each of the different estimations.

$$\begin{aligned} \ln \left(\frac{n_{i(t)g}}{n_{i(t)s}} \right) &= \frac{\rho}{1-\rho} (\gamma_{z_g/z_s}(t-\bar{t})) + \frac{1}{1-\rho} \ln \left(\frac{1-\omega_g}{1-\omega_s} \right) \\ &+ \frac{\nu_g - \rho}{\nu_g(\rho-1)} \ln \left(\omega_g \left[\frac{w_{i(t)}\omega_g}{P_{gt}(1-\omega_g)} \right]^{\frac{\nu_g}{1-\nu_g}} + 1 - \omega_g \right) \end{aligned} \quad (13)$$

$$\begin{aligned} &+ \frac{\rho - \nu_s}{\nu_s(\rho-1)} \ln \left(\omega_s \left[\frac{w_{i(t)}\omega_s}{P_{st}(1-\omega_s)} \right]^{\frac{\nu_s}{1-\nu_s}} + 1 - \omega_s \right) + \varphi_t^n + \eta_{i(t)}^n, \\ \ln \left(\frac{x_{i(t)g}}{x_{i(t)s}} \right) &= \frac{\rho}{\rho-1} \ln \left(\frac{P_{gt}}{P_{st}} \right) + \frac{\rho}{1-\rho} (\gamma_{z_g/z_s}(t-\bar{t})) + \frac{1}{1-\rho} \ln \left(\frac{\omega_g}{\omega_s} \right) \\ &+ \frac{\nu_g - \rho}{\nu_g(\rho-1)} \ln \left(\omega_g + (1-\omega_g) \left[\frac{w_{i(t)}\omega_g}{P_{gt}(1-\omega_g)} \right]^{\frac{\nu_g}{\nu_g-1}} \right) \\ &+ \frac{\rho - \nu_s}{\nu_s(\rho-1)} \ln \left(\omega_s + (1-\omega_s) \left[\frac{w_{i(t)}\omega_s}{P_{st}(1-\omega_s)} \right]^{\frac{\nu_s}{\nu_s-1}} \right) + \varphi_t^x + \eta_{i(t)}^x. \end{aligned} \quad (14)$$

All goods/services prices are PCE indices from the BEA, and wages are taken from their respective micro datasets. Expression (13) is the primary structural estimating equation we use when estimating parameters on time use data and prices, while expression (14) is the equation used when estimating parameters on expenditure data and prices. Each estimator will be run separately and independently.

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 $c_{i(t)g}$ and $c_{i(t)s}$, 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 substitution elasticities, ν_j , we seek meaningful priors that ensure logged versions of the marginal rate of technical substitution conditions in (6) hold:

$$\ln \left(\frac{P_{jt}n_{i(t)j}}{x_{i(t)j}} \right) = \underbrace{\frac{1}{\nu_j - 1} \ln \left(\frac{\omega_j}{1 - \omega_j} \right)}_{\phi_{j0}} + \underbrace{\frac{1}{\nu_j - 1} \ln \left(\frac{w_{i(t)}}{P_{jt}} \right)}_{\phi_{j1}} + \epsilon_{i(t)j}, \quad \forall j \in \{g, s\}. \quad (15)$$

We only use (15) as a preliminary regression to sign ν_j 's, as we will estimate these parameters with greater precision in our fully Bayesian estimator. We use OLS on (15) for each $j \in \{g, s\}$ on a 2003-2022 panel of matched CEX-LABSTAT/ATUS data aggregated by income quintile, $i \in \{1, 2, 3, 4, 5\}$, with wages taken from ATUS and consumption prices from PCE series.³⁷ Both the reduced-form and corresponding structural-parameter prior estimates are presented in Table 2 for

³⁷For this regression (and only this regression) we match the annual CEX-LABSTAT summary tables by income quintile to weighted averages by year and income quintile from ATUS. Note that the CEX-LABSTAT tables feature summary statistics compiled directly by the BLS.

all models where personal care time is included.³⁸ Notice that $\hat{v}_g < 0$ while $\hat{v}_s > 0$, suggesting that goods and off-market time using goods are complements, while services and off-market time using services are substitutes.³⁹ We can reject the notion that $\hat{v}_j = 0$ in all pre-estimates. Note that this amounts to testing whether $\hat{\phi}_{j1} = -1$, which can be rejected at a significance level of 0.01. We thus form priors under the assumption that market goods are complements with off-market time, while market services are substitutes with off-market time.

Table 2: $MRTS_j$ Regression Estimates for Prior Parameters (OLS)

<i>Goods</i>		<i>Services</i>	
$\hat{\phi}_{g0}$	-3.818*** (0.048)	$\hat{\phi}_{s0}$	-3.742*** (0.172)
$\hat{\phi}_{g1}$	-0.898*** (0.016)	$\hat{\phi}_{s1}$	-1.734*** (0.059)
$\hat{v}_g$	-0.114	$\hat{v}_s$	0.423
Observations	100	Observations	100
R ²	0.966	R ²	0.898
Residual Std. Error	0.084 (df = 98)	Residual Std. Error	0.295 (df = 98)

Note:

*p<0.1; **p<0.05; ***p<0.01

Given the results of the preliminary regression shown in Table 2, we can now place prior distributional assumptions on the structural parameters that restrict the signs of the v_g , v_s , and ρ . On each dataset we consider two different models each with separate priors on ρ :

$$-\rho \in \mathcal{LN}(-0.5, 1), \quad (16)$$

$$\text{or } \rho \in \mathcal{U}[0, 1], \quad (17)$$

where $\mathcal{U}$ is a uniform distribution and $\mathcal{LN}$ denotes a log-normal distribution. Prior distributional assumptions on the inner-nest elasticities under the regression results in Table 2 are

$$v_s, \omega_g, \omega_s \sim \mathcal{U}[0, 1] \quad (18)$$

$$-v_g \sim \mathcal{LN}(\mu_{v_g}, 1), \quad \text{where } \mu_{v_g} = \ln(-\hat{v}_g) - \frac{1}{2} = -2.672. \quad (19)$$

In (18) we simply assume that the output elasticities, ω_j 's, and v_s take flat, uniform priors. In (19) we use information about the value of v_g inferred from the regression in (15) to set the mean (location parameter) of the prior distribution, which is assumed to be log-normal. The fixed-effect

³⁸Technical Appendix E contains prior regression estimates for models without personal care time.

³⁹Note that v_s encodes the elasticity of substitution between q_s and n_s , not q_s and all off-market time or "time" in general.

coefficients and trend growth rate, γ_{z_g/z_s} , are given flat and improper priors:⁴⁰

$$\varphi_t^k \sim \mathcal{U}(-\infty, \infty), \quad \forall k \in \{n, x\} \quad (20)$$

$$\text{and} \quad \gamma_{z_g/z_s} \sim \mathcal{U}(-\infty, \infty). \quad (21)$$

All observations are assumed to have been generated from a mean-zero normal distribution, which characterizes the marginal (unconditional) distribution of unobserved, individual heterogeneity:

$$\eta_{i(t)}^k \sim \mathcal{N}(0, \sigma^2), \quad \forall k \in \{n, x\} \quad \text{with} \quad \frac{1}{\sigma^2} \sim \text{Gamma}(2, 4). \quad (22)$$

We sample the precision, $1/\sigma^2$, from a gamma distribution with the so-called shape/rate parameterization.⁴¹

Let $\bar{\pi}(\cdot)$ be an arbitrary probability density function (PDF), where (for example) $\bar{\pi}_{\mathcal{N}}$ is the PDF of the normal distribution. Let $\mathcal{P}^k$ for $k \in \{n, x\}$ be the set of parameters, conditional upon the equation — either (13) or (14) — around which we form the likelihood function. For example, when $k = n$, $\mathcal{P}^n = \{\rho, \nu_g, \nu_s, \omega_g, \omega_s, \gamma_{z_g/z_s}, \{\varphi_t^n\}_t, \sigma^2\}$. Let $\mathcal{D}_{i(t)}^k$ be the set of data observations associated with agent $i(t)$ under the targeted estimator, so that when $k = n$ this object is $\mathcal{D}_{i(t)}^n = \{n_{i(t)g}/n_{i(t)s}, w_{i(t)}, P_{gt}, P_{st}, t\}$. Rearrange (13) and (14) by isolating $\eta_{i(t)}^k$'s on the right-hand sides and let the functions $h^k(\cdot)$ be such that

$$\begin{aligned} h^n(\mathcal{D}_{i(t)}^n, \mathcal{P}^n) = & \ln \left(\frac{n_{i(t)g}}{n_{i(t)s}} \right) - \frac{\rho}{1-\rho} (\gamma_{z_g/z_s} (t - \bar{t})) - \frac{1}{1-\rho} \ln \left(\frac{1-\omega_g}{1-\omega_s} \right) \\ & - \frac{\nu_g - \rho}{\nu_g(\rho - 1)} \ln \left(\omega_g \left[\frac{w_{i(t)} \omega_g}{P_{gt}(1-\omega_g)} \right]^{\frac{\nu_g}{1-\nu_g}} + 1 - \omega_g \right) \\ & - \frac{\rho - \nu_s}{\nu_s(\rho - 1)} \ln \left(\omega_s \left[\frac{w_{i(t)} \omega_s}{P_{st}(1-\omega_s)} \right]^{\frac{\nu_s}{1-\nu_s}} + 1 - \omega_s \right) - \varphi_t^n = \eta_{i(t)}^n, \end{aligned} \quad (23)$$

$$\begin{aligned} h^x(\mathcal{D}_{i(t)}^x, \mathcal{P}^x) = & \ln \left(\frac{x_{i(t)g}}{x_{i(t)s}} \right) - \frac{\rho}{\rho - 1} \ln \left(\frac{P_{gt}}{P_{st}} \right) - \frac{\rho}{1-\rho} (\gamma_{z_g/z_s} (t - \bar{t})) + \frac{1}{1-\rho} \ln \left(\frac{\omega_g}{\omega_s} \right) \\ & - \frac{\nu_g - \rho}{\nu_g(\rho - 1)} \ln \left(\omega_g + (1 - \omega_g) \left[\frac{w_{i(t)} \omega_g}{P_{gt}(1-\omega_g)} \right]^{\frac{\nu_g}{\nu_g-1}} \right) \\ & - \frac{\rho - \nu_s}{\nu_s(\rho - 1)} \ln \left(\omega_s + (1 - \omega_s) \left[\frac{w_{i(t)} \omega_s}{P_{st}(1-\omega_s)} \right]^{\frac{\nu_s}{\nu_s-1}} \right) - \varphi_t^x = \eta_{i(t)}^x. \end{aligned} \quad (24)$$

⁴⁰Depending on the integration scheme used to estimate the posterior distribution of parameters conditional upon data, improper priors are fine.

⁴¹This parameterization has probability density function

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

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

Given the normal distribution assumption in (22), the probability density function and likelihood function associated with the observed set of data observations is

$$\bar{\pi}_{\mathcal{N}}(\mathcal{D}_{i(t)}^k; \mathcal{P}^k) = (2\pi\sigma^2)^{-\frac{1}{2}} \exp \left\{ - \left(\frac{h^k(\mathcal{D}_{i(t)}^k, \mathcal{P}^k)}{\sigma \sqrt{2}} \right)^2 \right\}, \quad (25)$$

$$\mathcal{L}(\mathcal{P}^k | \mathcal{D}^k) = \prod_t \prod_{i(t)} \bar{\pi}_{\mathcal{N}}(\mathcal{D}_{i(t)}^k; \mathcal{P}^k), \quad \forall k \in \{n, x\}, \quad (26)$$

where, in a slight abuse of notation, $\mathcal{D}^k$ without the subscript is the full set of data, $\mathcal{D}^k = \{\{\mathcal{D}^k\}_{i(t)}\}_t$. Expression (26) is the likelihood of realizing a particular parameter vector, $\mathcal{P}^k$, given the data $\mathcal{D}^k$ and the data-generating process described by the model.

We can now put together all of the prior distributions with the likelihood function to get an expression for the posterior distribution (represented by a non-conjugate PDF) of parameters given data and model assumptions:

$$\begin{aligned} \bar{\pi}(\mathcal{P}^k | \mathcal{D}^k) &\propto \mathcal{L}(\mathcal{P}^k | \mathcal{D}^k) \bar{\pi}_{\text{Gamma}}(1/\sigma^2) \bar{\pi}_{\mathcal{U}(-\infty, \infty)}(\gamma_{z_g/z_s}) \left(\prod_t \bar{\pi}_{\mathcal{U}(-\infty, \infty)}(\phi_t^k) \right) \\ &\times \bar{\pi}_{\mathcal{LN}}(\gamma_g) \bar{\pi}_{\mathcal{U}[0,1]}(\gamma_s) \bar{\pi}_{\mathcal{U}[0,1]}(\omega_g) \bar{\pi}_{\mathcal{U}[0,1]}(\omega_s) \bar{\pi}_y(\rho), \quad \forall k \in \{n, x\}. \end{aligned} \quad (27)$$

Note that the object $\bar{\pi}_y(\rho)$ in (27) is such that $y \in \{\mathcal{LN}(-0.5, 1), \mathcal{U}[0, 1]\}$, depending on whether we assume that c_g and c_s are complements or substitutes. The distribution characterized by (27) is non-conjugate (i.e., it cannot be analytically/algebraically derived) in the sense that the structure of the model along with the prior assumptions are *not* such that when combined using Bayes rule an analytical expression for a posterior density is yielded. Thus, sophisticated integration techniques are required to simulate (27) and recover an estimate for $\widehat{\mathcal{P}}^k$ — hence, our utilization of Hamiltonian Monte Carlo sampling under the Stan software system (Stan Development Team 2020). Technical Appendix C provides a summary of how the HMC process works to recover the structural parameters.

Estimated posterior means for the structural parameters are presented in Table 3 for estimations where personal care time is included. In Table 3 we show three columns for each estimator, depending on the dataset and assumption on the sign of ρ . The posterior mean is the primary point estimate. “S.E. Mean” is the standard error associated with this point estimate, while “S.D.” is the standard deviation of the full posterior distribution. Dividing the mean by its standard error gives a rough test statistic which can be used to assess the significance of the posterior point estimates: note that all values are significant at the 5% level.

We find that market goods and time are weak complements, while market services and time are moderate substitutes. There are little qualitative differences for $\widehat{\nu}_g$ and $\widehat{\nu}_s$ across estimators or datasets when personal care time is accounted for in Technical Appendix E. Similarly, there are little qualitative differences for $\widehat{\omega}_g$ and $\widehat{\omega}_s$ across estimators or datasets. When we look at $\widehat{\rho}$, notice that PUMD data suggest goods and services experiences are weakly less complementary

Table 3: Posterior Distribution of Parameters

ATUS Data						
	$\rho < 0$			$\rho \in (0, 1)$		
	Mean	S.E. Mean	S.D.	Mean	S.E. Mean	S.D.
ρ	-0.977	0.027	1.286	0.456	0.061	0.264
ν_g	-0.117	0.004	0.161	-0.105	0.007	0.145
ν_s	0.332	0.007	0.215	0.410	0.047	0.243
ω_g	0.514	0.006	0.287	0.470	0.011	0.271
ω_s	0.466	0.010	0.295	0.532	0.052	0.303
γ_{z_g/z_s}	0.015	0.031	0.997	0.234	0.069	0.913
σ	4.095	0.145	2.294	4.238	0.180	2.169
WAIC ^a = 312.476			WAIC ^a = 303.765			
PUMD Data						
	$\rho < 0$			$\rho \in (0, 1)$		
	Mean	S.E. Mean	S.D.	Mean	S.E. Mean	S.D.
ρ	-0.885	0.038	1.060	0.342	0.033	0.215
ν_g	-0.114	0.003	0.143	-0.106	0.006	0.123
ν_s	0.450	0.007	0.288	0.506	0.040	0.294
ω_g	0.492	0.009	0.291	0.494	0.013	0.276
ω_s	0.530	0.007	0.271	0.534	0.011	0.258
γ_{z_g/z_s}	0.042	0.094	1.005	-0.080	0.093	0.902
σ	3.352	0.122	1.701	3.869	0.221	2.120
WAIC ^a = 561.090			WAIC ^a = 655.408			
^a Widely applicable information criterion. See Watanabe (2010).						

and weakly less substitutable, depending on the assumed sign of ρ , than estimates under ATUS data suggest. When $\rho \in (0, 1)$ with PUMD, we find $\rho < \nu_s$, but with ATUS data we find $\rho > \nu_g$. This will have very minor implications for inferences regarding the drivers of structural change in Section 5. The implications are very minor, though, because, as will be seen, we prefer models with $\rho < 0$ for a variety of fitness reasons.

Technical Appendix D also contains plots of posterior distributions relative to their priors, which help assess the degree to which different parameters are identified. Note that in Bayesian models, identification is a “continuous” rather than a “discrete” notion. That is, parameters can

be strongly or weakly identified relative to assumptions about their prior distributions. The consumption technology elasticities, ν_g and ν_s , appear most well-identified relative to their priors, followed by ρ . We find that ρ appears to be slightly better identified when assuming $\rho < 0$ than when assuming $\rho \in (0, 1)$. We show that neither ω_j 's nor γ_{z_g/z_s} appear well-identified in the Bayesian sense, as their posteriors and priors are practically identical. This is not a problem for our estimation, though: it simply means that our point estimates generally inherit the means of our prior distribution, which in the case of ω_g and ω_s is 0.5 (i.e., to produce c_j the weights on q_j and n_j are roughly identical) and in the case of γ_{z_g/z_s} is zero. That $\gamma_{z_g/z_s} \approx 0$ simply says that variation in average productivities over time is almost entirely picked up by the fixed effects, φ_t^k .

The posterior average structural residuals, integrated over all agents, are given by $\zeta_{i(t)}^k$ from (12). These are plotted in Figure 3 for both ATUS and PUMD estimations (i.e., posterior distributions by year for agent-level means, $\hat{\zeta}_{i(t)}^k$ for $k \in \{n, x\}$), where the ATUS estimator includes personal care time.⁴² The solid lines represent means across samples by year, while the shaded areas are the 90% confidence regions.⁴³ Blue values show the relative-productivity distribution under the assumption that c_{ig} and c_{is} are substitutes, while red values are for when c_{ig} and c_{is} are complements. To compute the distributions 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, giving us values for $\hat{\eta}_{i(t)}^k$, $\forall k \in \{n, x\}$. Given we have values for $\hat{\rho}$, $\widehat{\gamma_{z_g/z_s}} \cdot (t - \bar{t})$, and $\hat{\varphi}_t^k$ via a combination of parameter estimates and data, with $\hat{\eta}_{i(t)}^k$ in hand, we can then compute $\hat{\zeta}_{i(t)}^k$ for each agent in each sample. In Figure 3 we observe growth in annual average $\hat{\zeta}_{i(t)}^k$ when assuming final experiences are substitutes (blue), though there is little variation over time in $\hat{\zeta}_{i(t)}^k$ when final experiences are complements (red). For any given assumption on ρ , the patterns are qualitatively consistent across both datasets, though variances are larger with estimates from the ATUS dataset.

Analogous plots in Technical Appendix E for data where personal care and sleep time are removed from n_{ig} are starkly different from those shown in Figure 3, despite the posterior distributions of estimated structural elasticities being similar. Indeed, we prefer models where personal care time is included because of how different estimates of $\hat{\zeta}_{i(t)}^k$ are *between* models when it is *not* included. That is, comparing Figures 3a and 3b, notice that red lines ($\rho < 0$) and blue lines ($\rho \in (0, 1)$) have qualitatively similar patterns up to a scalar constant, suggesting that our model is a good approximation of the true data-generating process (DGP). However, in Technical Appendix E we observe that consumption efficiencies have a strong, positive trend estimated on time-use data when $\rho < 0$ but no trend when estimated on expenditure data. Removing personal care time and sleep from n_{ig} thus causes estimates that rely on relative time-use data to diverge from those extracted from relative expenditure data. We focus our attention on models estimated using data *inclusive* of personal care time moving forward.

⁴²We plot posterior average fixed effects, φ_t^n and φ_t^x , as well as posterior average idiosyncratic errors, $\eta_{i(t)}^n$ and $\eta_{i(t)}^x$ separately in Technical Appendix D. We focus strictly on $\zeta_{i(t)}^n$ and $\zeta_{i(t)}^x$ in the main text.

⁴³Since this is a Bayesian model, these are confidence *regions* and not confidence *intervals*, as the posterior distributions need not necessarily be symmetric, though in this case they practically are.




Figure 3: These plots contain the distributions of posterior average elasticity-weighted relative log-productivities, $\zeta_{i(t)}^k$, $\forall k \in \{n, x\}$. The shaded areas are 90% confidence regions, while the solid lines are posterior means in each year of each sample. Note that the ATUS sample goes back to 2003, while the PUMD sample starts in 1990.

To further compare model fit, we also examine different models' WAIC (widely applicable information criterion) statistics, depending both on assumptions regarding ρ and the underlying dataset used to estimate the model. 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.^{44,45} Our comparisons of information criteria yield several model-selection results. Models where $\rho < 0$ (i.e., goods experiences are complementary with services experiences) generally outperform models where $\rho \in (0, 1)$ (i.e., substitutes). This is always true for PUMD models, though for ATUS models WAIC is somewhat of a wash.

We also like models with $\rho < 0$ for reasons beyond their statistical properties. First, we actually believe, ex-ante, that experiences associated with goods and services are likely to be more complementary than substitutable, anyway (i.e., $\rho < 0$). Note that there are certainly experiences for which c_{ig} and c_{is} are substitutes: think, for example, of in-home 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

⁴⁴As 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.

⁴⁵Recall with information criteria, lower is better.

a goods experience and a services experience) which is difficult to quantify and identify in data. Since joint activities are necessarily ones where using goods and services are complementary with each other, then we would expect experiences associated with using goods and services to be generally more complementary with each other in a model where one cannot account for joint activities, and such activities are prevalent. Second, $\rho < 0$ models generate aggregate residuals, z_{jt} 's, that simultaneously reconcile *both* relative aggregate expenditure (from Figure 2a) and relative average off-market time use (from Figure 1a). This will be seen in further detail in Section 5.

5 Quantitative Implications for Structural Change

The degree to which final consumption experiences are complementary or substitutable will determine the degree to which observed relative prices and wages or unobserved relative productivities primarily drive the evolution of economic behavior. Since we do not have micro data going back to mid-century, and while most of the goods to services transition had taken place prior to 1990, we must come up with a way to assess how these different factors have contributed to long-run structural change. To do this we use economy-wide average hourly wages and relative prices, while supposing that the average wage earner in the economy had an expenditure profile the same as that observed in aggregate X_{gt}/X_{st} , as estimated by the Bureau of Economic Analysis (2024) and featured in Figure 2a. We back out the implied relative consumption efficiency, z_{gt}/z_s , that reconciles the long-run evolution of X_{gt}/X_{xt} from the average wage earner.

We focus primarily on the drivers of trend changes to relative expenditure because of the time-use data limitations. Instead, to robustly verify whether model predictions with respect to long-run expenditure variation are consistent with what data observations we do have with respect to time use, we simulate off-market relative time use with in-home productivities backed out from aggregate expenditure data and compare these outcomes to AUT/TUESA and ATUS observations from Figure 1. Keeping with our findings from Section 4, models with $\rho < 0$ yield out-of-sample relative time-use predictions that are consistent with data, while those with $\rho \in (0, 1)$ do not, lending further credence to our contention that *experiences* associated with using goods are *complementary* with those associated with using services.

After assessing the model's aggregate predictions, we engage in several decompositions of long-run relative expenditure and time use by separately turning off the different drivers of structural change (e.g., wages, relative prices, and relative productivities). We then simulate new, counterfactual relative expenditure and relative time use series and compare these to data. This allows us to see what the aggregate goods/services ratios would have looked like had each series not changed since mid-century.

We simulate counterfactuals under four different scenarios: i) fix relative prices, allowing wages and productivities to evolve; ii) fix only wages and also do not allow goods and services prices to inflate, only for their relative prices to change; iii) fix only wages but allow goods and

services prices to inflate according to the aggregate rate of inflation; iv) fix relative productivities but allow wages and prices to evolve according to data. The degree to which the simultaneous decline in relative goods/services expenditure and rise in relative goods/services time use is driven by wage growth, prices, or productivities can thus be understood through the lens of these counterfactual simulations.

The distinction between simulations (ii) and (iii) is subtle but important. In simulation (ii) we fix wages at the value in the initial year of the aggregate time series (1959), and we also deflate P_{gt} and P_{st} by P_t^{agg} , for which we use the real \$2017 aggregate “Personal consumption expenditures (PCE)” price index from line 1 of NIPA Table 2.3.4 (Bureau of Economic Analysis 2024). Simulation (ii) thus isolates the model economy at an income and total-expenditure state that resembles that of the initial year (1959) while allowing inflationary-neutral relative goods/services prices to evolve. The reason we must deflate P_{gt} and P_{st} in order to totally isolate relative-price effects is because the model actually depends on three relative prices — P_{gt}/P_{st} , w_t/P_{gt} , and w_t/P_{st} — which can be seen by inspecting (13) and (14).^{46,47} Fixing w_t alone while allowing P_{gt} and P_{st} to inflate amounts to supposing incomes remained stagnant while goods/services prices grew, which is a statement about what the composition of the economy would look like if the costs of market goods/services rose without corresponding increases to income.

Under our aggregate approach to counterfactuals, we can use the data series from Figure 2b for relative prices, but we still need data for average wages and relative consumption productivities. For average hourly wages we take “Compensation of employees: wages and salaries” from NIPA Table 2.6 and divide by “Hours worked by full-time and part-time employees” from BEA account code A034RC (Bureau of Economic Analysis 2024).⁴⁸ For an estimate of average relative consumption productivities that reconcile relative expenditures, relative prices, and average hourly wages, we simply use the average hourly wage series, as well as relative expenditure as presented in Figure 2a and the prices used to construct Figure 2b, and then evaluate (14) under these series for parameters estimated under various datasets and different assumptions on ρ in order to back out model-implied $\hat{\zeta}_t^x$ for the aggregate average wage earner over the entire 1959 to 2022 time series. Note that we back out $\hat{\zeta}_t^x$ from (14) and then verify that estimated $\{\hat{\zeta}_t^x\}_t$ is close to an analogous estimate of $\{\hat{\zeta}_t^n\}_t$ which would be implied by (13) if we were to have average annual relative time use data going back to 1959.

Figure 4 plots the average relative consumption productivities, normalized to unity in 1959, along the entire 1959-2022 path under the four primary estimations accounting for personal care time.⁴⁹ Note that, under our preferred assumption that $\rho < 0$, the average efficiency of using goods versus services rose until the 1990s before mostly flattening and possibly declining. This

⁴⁶We call wages w_t (with t -subscript) in this section only, since they are aggregate average wages, and there is no heterogeneity.

⁴⁷Of course, the relative time-use condition does not *directly* depend on P_{gt}/P_{st} .

⁴⁸While we show how aggregate expenditure, price, and quantity data have evolved since 1948 in Figure 2, 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.

⁴⁹In Figure 4 we are plotting the actual, estimated relative consumption productivities, $\widehat{z_{gt}/z_{st}} = \exp\{\widehat{\zeta}_t^k(1 - \widehat{\rho})/\widehat{\rho}\}$.




Figure 4: These plots contain the estimates of posterior mean relative consumption productivities as implied by aggregate relative expenditure data. Red values are those backed out from (14) under the assumption $\rho < 0$, while blue values are backed out under the assumption $\rho \in (0, 1)$. We do this with parameters estimated using micro data from both ATUS (a) and PUMD (b).

coincides with the idea that the forces behind pre-1990 structural change are different than those after 1990, given the mechanics underlying the evolution of goods/services demand are different since 1990 compared to mid-century.

Next, for robustness we want to ensure that the productivities we back out from relative expenditure can reasonably reconcile what limited relative time use observations we have. Figure 5 presents predicted relative time use, normalized to unity in 2003 (the first year of the ATUS sample), for parameters estimated both with ATUS (a) and PUMD (b) micro data and the aggregate average productivities presented in Figure 4. Models with $\rho < 0$ match the increase in relative time use observed since mid-century *and* the increase observed since 2003.⁵⁰ Here forward, we present counterfactual results both from models with $\rho < 0$ and $\rho \in (0, 1)$ but in our discussions place more emphasis on results for $\rho < 0$.

Given that our aggregate productivity estimates appear robust, we can now simulate counterfactuals. Figure 6 shows our counterfactual relative-expenditure simulations. Our preferred models with $\rho < 0$ are featured in Figures 6a and 6b. Notice that, when fixing wages (i.e., eliminating long-run income effects) the simulated counterfactuals in these series are hardly different from data. That is, even if wages had remained the same and the economy not grown (i.e., the blue line where both wage inflation and aggregate price inflation are eliminated), structural change still would have occurred. This is because both relative prices and relative consumption productivities are almost entirely driving the relative decline in expenditure, which can be seen by noting that structural change is less pronounced when fixing either of these time series (i.e., the red and

⁵⁰See Technical Appendix F.3 for an analogous, zoomed-in version of Figure 5 only going back to 2003.




Figure 5: This figure plots relative time use, normalized to unity in 2003 (first year of ATUS), as predicted by both models estimated with ATUS (a) and PUMD (b) micro data along with the relative productivities backed out from aggregate relative expenditure data. The black dots in 1965-66 and 1975-76 are data from AUT and TUESA, while all observations after 2003 are from ATUS.

orange lines, respectively).

Figure 7 provides an analog of Figure 6 for relative time use. Recall that one reason we prefer $\rho < 0$ models is because they do a better job of predicting the increase in goods/service time use. We thus focus on Figures 7a and 7b, which counterfactually show that relative price variation (red) has had a strong and significant effect on the relative increase in goods/services off-market time. Thus, since experiences c_{gt} and c_{st} are complementary, as the relative price P_{gt}/P_{st} has fallen so that services have become relatively more expensive, consumers have had to spend more of their resources on market services and, in response, have chosen to substitute their input mix relatively more toward spending time *using* goods. Had prices remained fixed at their 1959 level while income still grew, consumers would have spent relatively *more* time using services. By contrast, fixing income while allowing price inflation (green) does not change the predicted relative time use trajectory. Indeed, product inflation (not wage inflation) plays a more important role in driving structural change to household allocations.

If we believe that $\rho < 0$, then relative to the literature our results are 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 findings suggest that supply-side drivers of relative-price variation have played a significant role in driving down relative sectoral output. As Herrendorf, Rogerson, and Valentinyi (2014) show, when the elasticities of substitution in sectoral production are unity and the factor shares are identical across sectors then relative prices are al-




Figure 6: Here, we plot log relative expenditure from data (black dots) against our four different simulated counterfactuals: fixed relative prices (red); no wage or aggregate price inflation (blue); no wage inflation but allowing for price inflation (green); no relative productivity growth (orange). Panels (a) and (b) contain ATUS and PUMD simulations respectively when $\rho < 0$, and panels (c) and (d) are analogous simulations for $\rho \in (0, 1)$.

ways inversely proportional to relative sectoral TFPs.⁵¹ Thus, if the goods sector is indeed growing faster than services, we would expect the relative price of goods to services to fall, with services eventually coming to dominate nominal output.

By contrast, however, if we believe $\rho \in (0, 1)$ there may be a role for rising incomes in contributing to structural change. Turning now to Figures 6c, 6d, 7c, and 7d, notice that had wages remained fixed and had there been no aggregate price inflation (blue line), relative expenditure would have fallen more slowly and relative time use would have fallen rather than risen. If we

⁵¹Thanks to an anonymous referee for pointing this out.




Figure 7: We plot log relative time use from data (black dots) against the model predictions (grey) and the same four different simulated counterfactuals as in Figure 6. Note that the log series are normalized to zero in 2003 (first year of ATUS). Panels (a) and (b) contain ATUS and PUMD simulations respectively when $\rho < 0$, and panels (c) and (d) are analogous simulations for $\rho \in (0, 1)$.

allow for aggregate price inflation but no wage inflation (i.e., the green line, which essentially represents purchasing-power erosion or an income-effect reversal), then structural change would have gone in the *opposite* direction for both relative spending and time use. Thus, a story where goods/services consumption *experiences* are substitutes yields results that are consistent with research that has attributed structural change to rising incomes (Kongsamut, Rebelo, and Xie 2001; Herrendorf, Rogerson, and Valentinyi 2013; Comin, Lashkari, and Mestieri 2021). However, it is not consistent with the additional data observation we introduce to the structural change puzzle — that as goods/services relative spending has fallen, goods/services relative off-market time

use has risen. To simultaneously reconcile *both* data patterns, we require final consumption experiences associated with goods to be complementary to those associated with services (i.e., $\rho < 0$).

6 Conclusion

We have shown that accounting for differential time-use complementarities in the consumption decision process can impact economic inferences. This is important when considering which mechanisms are most responsible for the structural evolution of an economy from one previously dominated by the consumption of manufactured goods to a services 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. We find little role for income effects in the context of this mechanism under our preferred specifications.

References

- Acemoglu, Daron and Veronica Guerrieri (2008). “Capital Deepening and Nonbalanced Economic Growth”. *Journal of Political Economy* 116.3 (cit. on p. 10).
- 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 p. 12).
- Aguiar, Mark et al. (2021). “Leisure Luxuries and the Labor Supply of Young Men”. *Journal of Political Economy* 129.2, pp. 337–382 (cit. on pp. 1, 4, 12, 14).
- Albrecht, Brian C., Thomas Phelan, and Nick Pretnar (2025). “Time Use and the Efficiency of Heterogeneous Markups”. *Working Paper* (cit. on p. 4).
- Baumol, William (1967). “Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis”. *The American Economic Review* 57.3 (cit. on p. 10).
- 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 p. 10).
- Becker, Gary (1965). “A Theory of the Allocation of Time”. *The Economic Journal* 75.299 (cit. on pp. 1–3, 11, 12).
- Betancourt, Michael and Leo Stein (2011). “The Geometry of Hamiltonian Monte Carlo”. *arXiv:1112.4118* (cit. on p. 18).
- Boerma, Job and Loukas Karabarbounis (2021). “Inferring Inequality with Home Production”. *Econometrica* (forthcoming) (cit. on pp. 1, 4, 12, 13, 18).

- 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. 2, 4, 6, 9, 19, 30).
- Bridgman, Benjamin (2016). “Home productivity”. *Journal of Economic Dynamics and Control* 71, pp. 60–76 (cit. on pp. 1, 2).
- 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).
- Bureau of Economic Analysis (2024). *National Income and Product Accounts (NIPA): Tables 2.4.3, 2.4.4, 2.4.5*. URL: <https://www.bea.gov/national/> (visited on 08/30/2024) (cit. on pp. 5, 9, 27, 28).
- Bureau of Labor Statistics (2021). *American Time Use Survey: Effect of the COVID-19 Pandemic on the 2020 ATUS*. <https://www.bls.gov/tus/notices/2021/covid19tech.htm>. Accessed: 2025-08-18 (cit. on p. 8).
- (2023). *Consumer Expenditure Surveys, Public Use Micro Data (CEX-PUMD)*. Accessed: 2025-08-19. URL: https://www.bls.gov/ce/pumd_data.htm (visited on 08/19/2025) (cit. on p. 4).
- (2024). *American Time Use Survey (ATUS)*. Accessed: 2025-08-18. URL: <https://www.bls.gov/tus/home.htm> (visited on 08/18/2025) (cit. on p. 6).
- 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. 2).
- 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. 2, 9, 10, 32).
- Cossu, Fenicia et al. (2023). “A Theory of Structural Change, Home Production and Leisure”. *Working Paper* (cit. on p. 4).
- Deaton, Angus and John Muellbauer (1980). *Economics and Consumer Behavior*. Cambridge University Press: New York (cit. on pp. 11, 14).
- Fang, Lei, Anne Hannusch, and Pedro Silos (2025). “Luxuries, Necessities, and the Allocation of Time”. *Working Paper* (cit. on pp. 1, 4).
- Gelman, Andrew et al. (2013a). “Computation in R and Stan”. *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. Appendix C (cit. on p. 18).
- (2013b). “Computationally efficient Markov chain simulation”. *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. 12 (cit. on p. 18).
- 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).
- 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. 14).

- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi (2013). “Two Perspectives on Preferences and Structural Transformation”. *American Economic Review* 103.7 (cit. on pp. 2, 5, 9–11, 14, 32).
- (2014). “Growth and structural transformation”. *Handbook of economic growth*. Vol. 2. Elsevier, pp. 855–941 (cit. on p. 30).
- Hicks, John (1932). *The Theory of Wages*. London: Macmillan (cit. on p. 15).
- Juster, F. Thomas et al. (1992). *Time Use in Economic and Social Accounts, 1975-1976*. URL: <https://www.icpsr.umich.edu/web/ICPSR/studies/7580> (visited on 08/31/2024) (cit. on p. 6).
- 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. 2).
- Kendall, M.G. (1975). *Rank Correlation Methods*. 4th ed. Charles Griffin: London (cit. on p. 8).
- 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. 14).
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie (2001). “Beyond Balanced Growth”. *Review of Economic Studies* 68, pp. 869–882 (cit. on pp. 2, 9, 10, 32).
- Kopytov, Alexandr, Nikolai Roussanov, and Mathieu Taschereau-Dumouchel (2020). “Cheap Thrills: the Price of Leisure and the Global Decline in Work Hours”. *Working Paper* (cit. on p. 1).
- Kuznets, Simon (1966). *Economic growth and structure*. Heinemann Educational (cit. on p. 10).
- (1973). “Modern Economic Growth: Findings and Reflections”. *American Economic Review* 63.3, pp. 247–258 (cit. on p. 10).
- Kydland, Finn and Edward Prescott (1982). “Time to Build and Aggregate Fluctuations”. *Econometrica* 50.6 (cit. on p. 14).
- Mann, H.B. (1945). “Nonparametric tests against trend”. *Econometrica* 13, pp. 245–259 (cit. on p. 8).
- 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. 2).
- 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. 18).
- Ngai, Rachel and Barbara Petrongolo (2017). “Gender Gaps and the Rise of the Service Economy”. *American Economic Journal: Macroeconomics* 9.4 (cit. on p. 14).
- Ngai, Rachel and Christopher Pissarides (2007). “Structural Change in a Multisector Model of Growth”. *American Economic Review* 97.1 (cit. on p. 10).
- (2008). “Trends in hours and economic growth”. *Review of Economic Dynamics* 11, pp. 239–256 (cit. on p. 1).
- Pretnar, Nick (2025). “Measuring Inequality with Consumption Time”. *Working Paper* (cit. on pp. 4, 18).
- Robinson, John P. (1977). *How Americans Used Time in 1965*. URL: https://www.ahtusdata.org/ahtus/resources/linked_docs/AHTUS/codebook_6566.pdf (visited on 08/31/2024) (cit. on p. 6).

- Stan Development Team (2020). *Stan Modeling Language: User's Guide and Reference Manual* (cit. on p. 23).
- 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. 2).
- 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. 24, 26).
- Whelan, Karl (2002). "A Guide to U.S. Chain Aggregated NIPA Data". *Review of Income and Wealth* 48.2, pp. 217–233 (cit. on p. 5).