

Why Do Income Effects Exist? Structural Change with Micro Foundations

Bill Bednar¹ and Nick Pretnar ^{*1}

¹*Carnegie Mellon University, Tepper School of Business*

September 25, 2018[†]

Abstract

This paper presents a general model of household consumption where income effects result from changes in the home production process, even though preferences are homothetic. We estimate the model taking market purchases of durable goods, non-durable goods, and services as inputs to an unobserved home production process, yielding empirical results similar to those in Herrendorf, Rogerson, and Valentinyi (2013) and Boppart (2014) but without assuming underlying preferences are non-homothetic. For the United States, estimates suggest that 41% of the decline in real consumption of non-durable goods relative to durable goods and 67% of the decline in real consumption of services relative to durable goods since 1929 can be attributed to relative price changes, while unobserved returns to home production and product quality comprise the difference. Thus, our framework provides a demand-side explanation for the structural change in the composition of U.S. consumption as resulting from unobserved changes in underlying home production processes masked as income effects.

Keywords: aggregate consumption, structural change, durables, services, home production
JEL Classification: D11, D12, D13, E21, R2

*Corresponding Author, Email: npretnar@cmu.edu

[†]Nick acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE1252522.

1 Introduction

It is a well-worn trope that the United States economy has exhibited “structural change” from an industrial economy dominated by the consumption of physical commodities, a.k.a. “goods,” toward a post-industrial one dominated by the consumption of non-physical, but equally tangible “service” commodities. Changes in the shares of expenditure versus shares of consumption since 1929 are an observable symptom of this phenomenon, which can be seen in Figures 1 and 2.¹ Note that relative consumption of durable goods has increased at the expense of services, a consequence of improvements in the real value of durable goods outpacing those of services. The posited reasons for such a transition are myriad, but broadly speaking tend to focus on opposite sides of the market. Our strategy here is to focus on household behavior. We will argue that a strain of the literature positing that income effects in preferences are causing structural transformation in the United States economy fails to account for the role that returns to home production may play. We demonstrate that a household model featuring non-homothetic preferences that accommodate income effects is observationally equivalent to one with homothetic preferences but in which the returns to home production are stochastic and unobserved. Using our formulation, one can estimate the degree to which relative price changes have contributed to structural transformation versus underlying quality improvements due to productivity gains in home production.

From the demand side, many authors have hypothesized that income effects, resulting from consumers having preferences over consumption that yield demands for goods and services which expand at rates that are non-linear in income (i.e., non-homothetic preferences resulting in non-linear Engel curves) have helped drive the broad transition (Echevarria 1997; Laitner 2000; Kongsamut, Rebelo, and Xie 2001; Caselli and Coleman 2001; Herrendorf, Rogerson, and Valentinyi 2013; Boppart 2014). These models, however, all depend on an *assumed* preference structure that captures these non-linearities with fixed (i.e., time independent) preference parameters, the economic meanings of which are often not obvious. As a consequence, such papers fail to address *why* income effects exist in the first place. In this paper, we argue that income effects are a result of unobserved changes to home productivity, and we show that a model accommodating changes to home productivity where preferences in the final home produced commodities are still homothetic provides the same results as the above papers with what we believe to be more robust economic intuition.

On the supply side, the usual explanation for structural change is that total factor productivity (TFP) in the goods sector has increased at faster rates than in the services sector, simultaneously causing goods-sector capital deepening while reducing that sector’s employment share (Baumol,

¹Throughout this paper, our analysis is two-fold, operating on both a consumption set that includes a stock and flow of household durables and one that does not. We are not the first in this literature to consider a model that excludes durable goods (see Herrendorf, Rogerson, and Valentinyi (2013) and Boppart (2014) for example). Often, durables are treated in the same manner as fixed, non-consumable capital assets. By allowing consumers to derive utility from the stock of durables, we are saying that durables are slightly different from fixed capital investment in that a household utilizes durable goods to produce some final consumption commodity in the spirit of Gomme, Kydland, and Rupert (2001).

Blackman, and Wolff 1985; Ngai and Pissarides 2007). Ngai and Pissarides (2007) present a model where differing rates of TFP growth between sectors induce worker migration toward sectors that produce durable capital and/or have the lowest TFP growth rate, a result that is generally consistent with Baumol’s “cost disease” afflicting low-productivity industries (Baumol 1967). Results in Ngai and Pissarides (2007) hold in general equilibrium despite homothetic utility on the demand side and the general independence of underlying preference parameters from sectoral and aggregate TFP. But the analysis of Ngai and Pissarides (2007), along with all of the papers which explain away structural change with non-homothetic preference structures, fails to account for the ever-evolving ways in which households *use* goods and services to achieve productive ends.

For example, Becker (1965) theorizes that households combine market goods with time to produce ultimate, final consumption goods. Here, we use Becker (1965) as a springboard to argue that unobserved returns to home production are due to unobserved quality improvements in market inputs used to produce final consumption goods. Syverson (2017) uses a similar “Beckerian” argument to attribute the observed slowdown in productivity growth to mis-measurement of implicit household value of new products. Our approach is similar to the partial equilibrium approach of Herrendorf, Rogerson, and Valentinyi (2013) who analyze how much income effects play a role in contributing to structural change by comparing predictions from an estimation of a non-homothetic model of household utility against a homothetic counterfactual. We take the analysis of Herrendorf, Rogerson, and Valentinyi (2013) a step further, arguing that any causal inference with regards to structural change being attributed to income effects should also consider changes in home productivity of market goods. Not only do such considerations help answer the question as to “why” structural change may be occurring, but also such considerations provide a micro-founded explanation of why there exist income effects at all.

While we look only at demand-side aggregates, estimating the within-period first-order conditions of a structural model of household preferences over multiple consumption goods, we consider two simultaneously competing mechanisms driving structural change: 1) relative price changes which are correlated with exogenous firm productivity; 2) endogenous technical changes in home productivity resulting from unmeasured quality changes. Like Herrendorf, Rogerson, and Valentinyi (2013) and Boppart (2014), the tractable model we estimate fits the observed change in relative consumption from 1929 to 2017 very well. This should not surprise: econometricians observe neither preferences nor home productivity, so we make assumptions on household behavior to add curvature to utility functions that will mathematically capture the observed data trend. The underlying mechanism, though, must be assumed. We believe that our explanation involving relative changes to home productivity is as plausible as any explanation that posits households prefer services to goods more as they get richer. Even better, our explanation provides a reason for why this may be true. Thus, returns to home production represent a sort of “final” value-added to consumption, value that is added after a product is purchased on the market, and added due to the ability of the household to transform market commodities into some final consumption good they ultimately consume.

This additional value-added represents possibly unobserved changes to market good quality. By unobserved quality changes, we specifically mean changes in the quality of a commodity that price indices fail to account for, thus causing estimates of the real consumption value of that commodity to be biased downwards.² In a partial equilibrium sense, the market value of one commodity relative to some other may indeed be fully captured by the price indices used by both the BEA (personal consumption expenditure indices, PCE) and the BLS (consumer price indices, CPI), yet statisticians fail to observe the often significant ancillary benefits consumption of certain commodities can provide to a household. These benefits are best thought of as intermediate returns achieved after market purchase and within the household, wrought from some intrinsic home production process which utilizes market inputs, existing home capital, and time to produce a final consumption item.

The model we present below is constructed in the spirit of Becker 1965: that is, by using a variant of Gary Becker’s original home production formulation, we can separate the “indirect” price of consumption from the “direct” or market price. In Becker’s original formulation, final consumption is produced using time and market goods, and the indirect price of consumption is thus the opportunity cost of foregone wages earned from market time. More generally we think of the indirect price of consumption as the opportunity cost of foregone rents which acknowledges the fact that consumers use durable capital as inputs to home production in addition to time and other non-durable items. We estimate two versions of a model featuring preferences that have the constant elasticity of substitution property in final, consumed goods, but not necessarily in market goods. To estimate the model, we employ deterministic Hamiltonian Monte Carlo (H-MC) simulation, which is little-used in economics but which is both statistically and computationally efficient when estimating unobserved, possibly unit-root, time trends. We employ this technique to infer how latent changes in home productivity (or product quality) have affected relative consumption in broad aggregates (durables, non-durables, and services) as well as their respective components since 1929.³ Our findings suggest that relative price changes are responsible for approximately 41% of the change in relative non-durable to durable flow consumption and 67% of the change in relative services to durable flow consumption since 1929.

The paper proceeds as follows: Section 2 outlines the theory of the household we build our model around; Section 3 presents model assumptions needed for estimation and the results of that estimation; Section 4 concludes, outlining paths forward for future work.

²See for example Gordon (2016) who discusses shortcomings to the ways commodity bundles are counted and indexed.

³A discussion of the general structure of the H-MC algorithm, as well as the merits of using such a routine, is featured in Appendix C.




Figure 1: Shares of aggregate expenditure and stock of consumption since 1929 in real 2017 dollars. See Appendix B for data details.




Figure 2: Shares of aggregate expenditure and consumption since 1929 in real 2017 dollars — non-durable goods and services only. See Appendix B for data details.

2 Theory of the Household

2.1 Home Production

For illustration, consider a consumer who purchases a dishwasher, a durable commodity whose market value depreciates at the net annual rate δ . Each period, real market value of services derived from using this dishwasher q_t is a fraction $(1 - \delta)$ of its real market value in the previous period:

$$q_t = (1 - \delta)q_{t-1} \tag{1}$$

Yet the dishwasher is utilized to achieve cleanliness, a commodity for which the household derives final utility in the sense of Becker (1965). Denote this final commodity by c_t (c is for cleanliness). c_t is produced using some observed quantity q_t (the dishwasher), some other commodities plus fixed capital assets, all of which we denote by a_t , and time n_t . These inputs are passed through some home production function f_t such that $c_t = f_t(q_t, a_t, n_t)$. This formulation differs from those of Gronau (1977), Graham and Green (1984), Benhabib, Rogerson, and Wright (1991), Rupert, Rogerson, and Wright (1995), Aguiar and Hurst (2007), and Aguiar, Hurst, and Karabarbounis (2013) who all, for the most part, dispense with the idea that market goods *and* time are combined together in some home production function to produce a final commodity. Notable exceptions are McGrattan, Rogerson, and Wright (1993) and Gomme, Kydland, and Rupert (2001) which allow for durable capital to be combined with time toward the production of home consumption commodities. The model structures in these papers are closest in spirit that we have found in the applied literature to the original formulation of Becker (1965). With its provision allowing exogenous technical change to home-based productivity, we utilize the model in Gomme, Kydland, and Rupert (2001) as a springboard toward a richer, but still decidedly theoretical, model of the household production process.

While Gomme, Kydland, and Rupert (2001) allow for capital to complement home labor, Benhabib, Rogerson, and Wright (1991) and Rupert, Rogerson, and Wright (1995) structurally posit that consumers derive final utility from market goods and an additional commodity produced in the home using only non-market and non-leisure time. In their formulation, the home-produced commodity is wrought only by individual effort and without complimentary assistance from the utilization of market goods or fixed capital inputs. The analyses in Aguiar and Hurst (2007) and Aguiar, Hurst, and Karabarbounis (2013) are empirical and do not delve into the structural details of household utility maximization, but the authors' implicit definition of home production is the provision of off-market time for non-leisure-based home activities. Such simplifications seem to violate the spirit of home production as introduced in Becker (1965), that market goods *and* time are utilized simultaneously to produce some final commodity.⁴

⁴Aside from the plaudits hoisted on Gary Becker for his pioneering theories and work on human capital, he gets less credit for his unique interpretations of consumption utility. For illustration, Stigler and Becker (1977) describe

Continuing with our simple illustration of Becker's original theory at work, let us return to the dishwasher example. Suppose we do not observe a_t and n_t . With some assumptions on f_t and the assumption that $c_t, q_t \geq 0$ always, we can safely denote final consumption c_t as equal to the multiplicative combination of q_t and some unobserved residual λ_t

$$c_t = q_t \cdot \lambda_t \quad (2)$$

λ_t captures everything from unobserved quality changes in the dishwasher q_t which improve the home production process to changes in both the household's relative internal and market values of time n_t and other fixed inputs a_t which affect the opportunity cost of home production. This can be thought of as some kind of exogenous productivity residual in the general spirit of Kydland and Prescott (1982) and Gomme, Kydland, and Rupert (2001), though like in Gomme, Kydland, and Rupert (2001) this residual affects only the home production process and even more specifically, a narrow component of the process at that. λ_t also captures the fact that only a fraction of the service flow from using the dishwasher is to achieve cleanliness and that the remaining fraction of the dishwasher's service flow may also be utilized to produce other home commodities. Finally, we must recognize that this is a simple illustration, and cleanliness is not just produced with a dishwasher.⁵

Now extending this to multiple final commodities vectorized by c_t and multiple intermediate and market commodities vectorized by q_t , assume that the market commodities are aggregated into final consumption according to some highly-structured production process f_t and let g_t be a reduced-form representation of f_t . Index the components of c_t by $i \in \{1, \dots, I\}$ and the components of q_t by $j \in \{1, \dots, J\}$. q_t is a vector of market commodities, including durable goods.

Assumption 1. The components of f_t are homogeneous of degree $z \leq 1$ in all inputs.

Assumption 2. The elasticity of home production for final commodity $c_{i,t}$ with respect to the quantity $q_{i,j,t}$ associated directly with the production of $c_{i,t}$ is constant over time for all combinations of i and j , though depends on the combination.

Denote real investment in market commodities by i_t , and let δ be the J -dimensional vector of net depreciation rates for q_t where $\delta_j = 1$ if good j is non-durable.⁶ Let ι denote the unit vector and \mathbb{I}

a model of rational addiction where such final consumption commodities as "music appreciation" and "euphoria" are produced using home production functions and stocks of "euphoria capital" which depreciate or appreciate with increasing consumption, depending on whether the addictive nature of the final commodity is harmful or beneficial. Readings of Stigler and Becker (1977) and Becker and Murphy (1988) lead one to believe that the Nobel laureate thought of utility as something akin to the raw benefit one received from consuming something but not pleasure in and of itself. It seems rather that Becker thought of the experience of "pleasure" as exclusive from the utility derived from its consumption, a rather advanced notion which helps characterize the underlying assumption of modern economics that tastes are neither good nor bad and utility should not be confused with happiness.

⁵One may desire sponges, soap, and perhaps even a maid.

⁶The durability distinction obviously depends on the units of time over which the model is defined. The set of products which would be characterized as durable over the course of a year is a subset of the set of products whose usefulness lasts at least one week after purchase.

the identity matrix. The stock of market commodities held by the household evolves according to

$$\mathbf{q}_t \leq (\mathbf{I} - \boldsymbol{\delta}) \cdot \mathbf{q}_{t-1}^\top \cdot \mathbf{I} \cdot \mathbf{I} + \mathbf{i}_t \quad (3)$$

A household can invest in market commodities at price \mathbf{p}_t . Let y_t denote income net of savings and investment in assets which cannot be directly consumed in period t .⁷ Each period, the household faces the following budget constraint:

$$\mathbf{p}_t^\top \cdot \mathbf{i}_t \leq y_t \quad (4)$$

Assumption 3. All input resources are used up in the production of final consumption commodities. For the case where an input resource is durable, the service flow is entirely utilized in home production.

Let \mathbf{Q}_t be the input matrix with rows indexing the final consumption good (i) and columns indexing the input (j). $\mathbf{q}_{i,t}$ will be a row vector which contains the levels of each $q_{j,t}$ used in the production of $c_{i,t}$.⁸ Note that $q_{j,t}$ may be a durable item, which continues to get used in each period, though its market value declines as discussed above.⁹ What distinguishes a market good from a final consumption good produced in the home is a matter of measurability. Market goods are those which have well-defined formal markets and whose prices and quantities organizations like the BEA and BLS spend resources measuring. Final consumption goods can be thought of as those commodities we consume which are not as readily defined or measured. In the above example, cleanliness takes this role.¹⁰ Lemma 1 states that all market commodities must be used in home production activities and thus cannot be stored. Note that this does not preclude $q_{j,t}$ from representing the flow of services from durable goods. Rather under Assumption 3, Lemma 1 says that the *entire flow* of said services will be utilized in the production of some final commodity, though the stock of the commodity may still yield additional services next period.

Lemma 1. Consider the j^{th} column of \mathbf{Q}_t , $\mathbf{q}_{j,t}$ and let \mathbf{q}_t denote the J -dimensional vector whose components contain the total real amount of the market commodities used in all home production

⁷ y_t includes both returns from labor and external capital investments accrued in period t less the household's reinvestment in productive, non-consumable capital.

⁸A notational aside: bold face font is used to emphasize vectors, capital bold face font for matrices, and non-bold face font for single scalars. When we refer to the quantity $q_{j,t}$ without bold face, we mean the the sum $\sum_{i=1}^I q_{i,j,t}$ where each $q_{i,j,t}$ denotes the amount of $q_{j,t}$ devoted to production of final commodity $c_{i,t}$. When we refer to the quantity $\mathbf{q}_{j,t}$ with bold face, we are referring to the I -dimensional vector whose components are $q_{i,j,t}$.

⁹For our purposes here, the stock of a durable is nonetheless a "market commodity." A dishwasher after all can be sold as a used durable item at some rate less than the purchase price. The decline of the dishwasher's market value is captured by the depreciation rate δ .

¹⁰Of course, there are other examples. Becker famously cites the "seeing of a play" as a final consumption commodity, but attending a play is actually a measurable market commodity (Becker 1965; Stigler and Becker 1977). Perhaps one could think of the "relaxation" resulting from seeing the play as the final commodity, with the individual audience member deriving diminishing marginal utility from this ultimate relaxation which can be produced through a variety of means — attending a play, time spent taking a nap, drinking a glass of wine, etc. The means through which "relaxation" is produced are themselves the market-side inputs to the relaxation production function which is some component of f_t .

activities in period t . Under Assumption 3,

$$\mathbf{q}_t^\top = \boldsymbol{\iota}^\top \cdot \mathbf{Q}_t \quad (5)$$

Proof. See Appendix A. □

As we have mentioned, \mathbf{f}_t and \mathbf{c}_t are unobservable, as are \mathbf{a}_t and \mathbf{n}_t , the vectors of fixed capital and labor inputs to the home production process. However, we can specify \mathbf{g}_t as a quasi-reduced form version of \mathbf{f}_t which we can then compose with some assumed utility representation. Proposition 1 denotes the structure of the components of \mathbf{g}_t which satisfy our assumptions.

Proposition 1. Suppose an econometrician possesses no information on fixed capital inputs $a_{i,t}$ and time $n_{i,t}$ used in home production. Let $\lambda_{i,j,t}$ be the unobserved residual returns to commodity j in the production of final consumption commodity i . Under Assumptions 1 thru 2 we can implement reduced-form representations $g_{i,t}$ of the component functions $f_{i,t}$ of \mathbf{f}_t as

$$g_{i,t}(\mathbf{q}_{i,t}) = \left[\prod_{j=1}^J (\lambda_{i,j,t} \cdot q_{i,j,t})^{\omega_{i,j}} \right] \cdot \eta_{i,t}(a_{i,t}, n_{i,t}) \quad (6)$$

where $\eta_{i,t}(a_{i,t}, n_{i,t})$ is homogeneous of degree $z \leq 1$. Further, assume $\sum_{j=1}^J \omega_{i,j} \leq 1$ and $0 < \omega_{i,j} < 1, \forall i, j$.

Proof. See Appendix A □

2.2 Preferences & Choices

We have discussed the difficulty in measuring \mathbf{c}_t . To even begin to define what the components of this vector actually are, regardless of the vector's length I , may be even harder. Econometricians also face a choice whether to define \mathbf{q}_t as a vector of broad aggregates like “goods” and “services” for example or to include some dis-aggregated components like “clothing” and “food” in the case of goods or “health care” and “public transit” in the case of services. The degree of granularity with which commodities are defined can have profound effects on economic inferences as we will illustrate here. First, there is the matter of identifying the home production function, which Proposition 1 addresses. Notice that $g_{i,t}(\mathbf{q}_{i,t})$ depends on J , the number of market commodities. Meanwhile, any specification of period t flow utility $u_t(\mathbf{c}_t)$ must map an I -dimensional vector into a one-dimensional space. Generally speaking, household equilibrium conditions will directly depend on I absent usage of clever functional forms for $u_t(\mathbf{c}_t)$.

The general lifetime preference structure we assume here builds on the intertemporal structure of preferences employed in Ogaki and Reinhart (1998). Suppose the demand side of our economy is characterized by a representative household that lives forever and denote the starting period of our analysis as $t = 0$. We assume the household has a time-additively separable utility function with geometric intertemporal time preference at a constant rate β . Further, let γ denote the

intertemporal elasticity of substitution for consumption, and let \mathbb{E}_0 represent the agent's expectations in period zero. Then lifetime utility U is given by

$$U = \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left(\frac{\gamma}{\gamma-1} \right) \left[u_t(c_t)^{1-1/\gamma} - v_t(n_{l,t}) \right] \right\} \quad (7)$$

where $-v_t(n_{l,t})$ is the dis-utility of labor supplied $n_{l,t}$.

Assumption 4. Both $u_t(c_t)$ and $f_t(Q_t)$ are continuously differentiable and strictly concave.

Assumption 5. $u_t(c_t)$ satisfies the Inada condition: $\lim_{c_{i,t} \rightarrow 0} \frac{\partial u_t}{\partial c_{i,t}} = \infty$ for all i and t .

Note that final consumption depends on the household's purchases of market goods through the quasi-reduced form relationship $c_{i,t} = g_{i,t}(q_{i,t})$. We can thus characterize household consumption as resulting from choices of market good purchases. In equilibrium the household choosing market good j solves the standard intratemporal condition equating good j 's marginal utility with that of all other market goods, weighted of course by their respective relative prices. As Becker (1965) illustrates, these prices however contain both implicit and explicit components, since market goods are first transformed into final goods via $f_t(Q_t)$. Consumers thus must choose both how much of good j to purchase on the market and afterwards how much of each good to devote to the production of final commodity i . The implicit prices capture the costs of this allocation process.

To fully characterize the consumer's choice set, note that they can invest in fixed non-consumable capital assets a_t , where we denote the total flow of this investment as $i_{k,t}$. Allow fixed capital to depreciate at rate δ_a and evolve according to

$$a_{t+1} \leq (1 - \delta_a) \cdot a_t + i_{a,t} \quad (8)$$

Capital can of course be used in the home production process or rented out on the market to be used for other means.¹¹ Thus capital has $I + 1$ possible utilizations each of which are associated with a shadow price intrinsic to the household. The same goes for time utilization, which can be allocated to either labor $n_{l,t}$ or various home production pursuits $n_{i,t}$.¹² We normalize total available time to unity and require $1 = n_{l,t} + \sum_{i=1}^I n_{i,t}$ with $n_{l,t} \geq 0$ and $n_{i,t} \geq 0$ for all i and t . Denote \mathbf{a}_t and \mathbf{n}_t as the $I + 1$ -dimensional vectors with the first I components being the quantities of each devoted toward home production and the $I + 1$ component being the quantity of each devoted toward external market production. Also redefine \mathbf{i}_t as a $J + 1$ -dimensional vector whose first J components represent the flow of market purchases with the $J + 1$ component being the

¹¹This is the traditional sense in which capital utilization is modeled as in Solow (1957) — toward the production of market goods.

¹²One may understandably wonder "what about leisure time?" Leisure can be thought of as a final consumption commodity, the production of which takes as inputs market goods like perhaps a novel, running shoes, alcohol, a deck chair and sunscreen, sporting equipment, amongst many other items, as well time and fixed capital. Therefore in our model "leisure" is not just time itself but the outcome of a home production process.

flow of investment in non-consumable, fixed capital assets $i_{k,t}$. Consumers must choose both how to partition their market goods, assets, and time as well as how much to add to their existing stocks of market goods and assets. The full choice set is thus

$$\left\{ \mathbf{Q}_t, \mathbf{a}_t, \mathbf{n}_t, \mathbf{i}_t \right\}_{t=0}^{\infty} \quad (9)$$

In (9) there are $J \times J + 2 \times (I + 1) + J + 1$ choice variables, clearly some of which are redundant. For example, Lemma 1 plus the budget constraint ensures that one column of \mathbf{Q}_t can be identified given the other $J - 1$ columns. Suppose now the consumer attempts to maximize his lifetime utility subject to the budget constraint and laws of motion holding with equality. Propositions 2 and 3 state that the internal household marketplace functions just like any other market place featuring perfect competition. Here, one can think of all the various possible ways in which market goods can be utilized toward household production competing against each other for the best possible, utility-maximizing uses.

Proposition 2. Let i and i' index two distinct final consumption goods which each require strictly positive amounts of market good j as inputs to home production. The internal household relative value of consumption of final good i to i' is equal to the infra-marginal rate of substitution for market good j :

$$\frac{\partial f_{i',t}}{\partial q_{i',j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j,t}} \quad (10)$$

Proof. See Appendix A. □

Proposition 3. Consider some final consumption good $c_{i,t}$ produced by $f_{i,t}(\mathbf{q}_{i,t})$. Suppose $c_{i,t}$ is produced using strictly positive amounts of market goods j and j' . For all internal household production processes, the marginal rate of technical substitution between j and j' must equal their relative market price:

$$\frac{\partial f_{i,t}}{\partial q_{i,j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j',t}} = \frac{p_{j,t}}{p_{j',t}} \quad (11)$$

Proof. See Appendix A. □

2.2.1 Reconciling Observed and Unobserved Choices

Generally speaking an econometrician cannot and does not observe \mathbf{f}_t or the entire matrix \mathbf{Q}_t , let alone the vector \mathbf{a}_t . In some cases such datasets like the American Time Use Survey (ATUS) allow us to observe the vector \mathbf{n}_t , but it is difficult if not impossible to understand the ways in which market goods are split up and divvied out toward household production. In this section we wish to emphasize how various assumptions on household choices commonly made in literature deal-

ing with household economics can have a substantial impact on inference when an econometrician is unable to identify all elements of the household's choice set.

Suppose we, as economists, have data on the J -dimensional vector of consumer purchases q_t . Suppose further that there exists a non-linear shadow process f_t by which the household converts market goods to final consumption. Additionally suppose we, as economists, assume away the necessity of the household to choose how to partition their market purchases amongst these various home production processes and instead consider the reduced-form choice set:

$$\left\{ n_{l,t}, i_t \right\}_{t=0}^{\infty} \quad (12)$$

Assume further we know all initial stocks of consumption q_0 and non-consumable capital assets a_0 along with the depreciation rates of these stocks so that (3) and (8) can be identified. However, absent stronger assumptions on f_t we cannot determine how the quantities of purchased market goods are distributed across the various home production tasks. Let $\alpha_{j,t}$ be an I -dimensional vector describing the distribution of $q_{j,t}$ across the various home production tasks. Given j and t , for each i , $0 \leq \alpha_{i,j,t} \leq 1$ and $\sum_{i=1}^I \alpha_{i,j,t} = 1$. This implies

$$q_{j,t} = q_{j,t} \cdot \alpha_{j,t} \quad (13)$$

As Lemma 2 illustrates, if we assume that households engage in home production but are only privy to market consumption expenditure, absent strong simplifying assumptions on utility and the home production process, the equilibrium conditions can become rather unwieldy. In fact, as an extension of Lemma 2, Corollary 1 directly addresses the fact that the exact quantity of separate home production tasks will generally impact the marginal rate of substitution.

Lemma 2. Suppose households take $\alpha_{j,t}$ as given for each j and just choose the total quantity of market goods to purchase each period. For market goods j and j' , the period- t marginal rate of substitution is

$$\left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j,t}} \cdot \alpha_{i,j,t} \right) / \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j',t}} \cdot \alpha_{i,j',t} \right) \quad (14)$$

Proof. See Appendix A. □

Corollary 1. General speaking, the observed rate at which households substitute between market goods depends on the total number of productive tasks for which households utilize each market good.

Proof. See Appendix A. □

3 Toward a Tractable Model of Structural Change

Home production is not usually modeled in such a rich way as we are doing, so there is little precedent for work that makes and subsequently justifies explicit, formal assumptions that simplify the equilibrium conditions generated by the theoretical model. What we mean by “tractability” of course, must be contextualized. Specifically, we seek a model of household consumption that can be used to estimate (with various reasonable assumptions) the degree to which exogenously unmeasured returns to home production have contributed to the structural transformation of the United States economy over the twentieth century from one centered around expenditure on manufactured “goods” (both durable and non-durable) to the present, twenty-first century economy where “services” expenditure dominates. Any demand-side model which accomplishes this will obviously control for changes in the technical productivity of firms. Relative price changes will capture a lot of the change in relative productivities between industries. We seek a model that allows us to separately identify the contributions to structural change of relative market price changes from quality improvements induced by changes in home productivity.

To this point, the be-all-end-all explanation for demand-driven structural change comes from Boppart (2014). The general results in Boppart (2014) complement the supply-side driven results of Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008), and Foellmi and Zweimuller (2008) which all predict, in one way or another, that as $t \rightarrow \infty$ the goods share in the economy goes to 0. In this limit, civilization turns into a nudist colony where everyone eats out at restaurants that do not have stoves to cook with or plates on which to serve you nothing because there is no food.¹³ Of course, neither people nor civilizations can last forever, so a ridiculous asymptotic result may never matter.¹⁴ Nonetheless, these results highlight fundamental flaws in their associated models of structural change: the models cannot accommodate all realistic ways in which an economy could structurally change. Structural change is ultimately, in general equilibrium, driven by changes in relative value, and value is determined separately but simultaneously by both firm and household production processes.

Our formulation differs from those before us as follows. First, we assume households’ underlying, call them “shadow”, preferences for final consumption commodities takes a fairly standard, constant elasticity of substitution (CES) form. Second, we allow deviations from preference homotheticity but via changing returns to home production rather than households’ price sensitivity changing in income. This allows us to say *both how and why* income changes affect the composition of consumption rather than just assuming that they do as others have done.¹⁵ It is not to say

¹³This is because the “goods” sector includes food, clothing, and consumer durables.

¹⁴We beg the reader to note that we understand why these asymptotic results matter for model inference, especially with regards to balanced growth. Our argument here is that the model’s predictions only hold when an economy is behaving in ways similar to the United States economy of the late twentieth and early twenty-first century, and that there is no reason that we should expect the goods sector to ever die entirely. This fact renders the generalizability of such model’s economic inferences somewhat suspect.

¹⁵See for example Echevarria (1997), Herrendorf, Rogerson, and Valentinyi (2013), and Boppart (2014), each of which suppose utility functional forms that are flexible enough to accommodate non-linear Engel curves.

that income effects are themselves not a good explanation for structural change. The underlying question is, why are there income effects? We hope that our interpretations here may be able to help answer this question.

3.1 Simplified Model

Period t flow utility from consumption has the CES form:

$$u_t(c_t) = \left(\sum_{i=1}^I \theta_i c_{i,t}^\rho \right)^{\frac{1}{\rho}} \quad (15)$$

where $\sum_{i=1}^I \theta_i = 1$ and $\sigma = \frac{1}{1-\rho}$ is the intratemporal substitution elasticity for pairs of final consumption.

We do not observe f_t , so we invoke Proposition 1 and specify a quasi reduced-form version $g_t(q_t; a_t, n_t)$. Also, we do not observe a_t and n_t in the home production process. Assumption 6 along with Proposition 1 and Lemma 3 allow us to collapse the home production process into a tractable form to substitute into (15). Specifically, Assumption 6 says that each final good is produced using one and only one market good. In Lemma 3, $\zeta_{i,j,t}$ represents unobserved variation in after-market value-added of $q_{j,t}$ to the household production process.

Assumption 6. For each component i of $g_t(q_t; a_t, n_t)$, all but exactly one $\omega_{i,j} = 0$. If $\omega_{i,j} \neq 0$ then $\omega_{i,j} = 1$. Further, for $i \neq i'$, if $\omega_{i,j} = 1$ then $\omega_{i',j} = 0$.

Lemma 3. Under Proposition 1 and Assumption 6, there exists a unique j such that for each i :

$$g_{i,t}(q_{i,t}; a_{i,t}, n_{i,t}) = \zeta_{j,t} \cdot q_{j,t} \quad (16)$$

where $\zeta_{j,t} = \lambda_{i,j,t} \cdot \eta_{i,t}(a_{i,t}, n_{i,t}) > 0$.

Proof. See Appendix A. □

Now substitute $c_{i,t} = g_{i,t}(\cdot) = \zeta_{j,t} \cdot q_{j,t}$ into (15) and (15) into (7). Proposition 4 provides us with closed forms for the intratemporal first-order necessary conditions.

Proposition 4. Given Lemma 3, the intratemporal first-order necessary conditions characterizing household market consumption can be written:

$$\ln \left(\frac{q_{j,t}}{q_{j',t}} \right) = -\sigma \cdot \ln \left(\frac{\theta_{j'}}{\theta_j} \right) - \sigma \cdot \ln \left(\frac{p_{j,t}}{p_{j',t}} \right) + \ln \left(\frac{\zeta_{j',t}}{\zeta_{j,t}} \right) \quad (17)$$

for all $j \neq j'$.

Proof. See Appendix A. □

Let $\xi_{j,t} = \ln \left(\frac{\zeta_{j',t}}{\zeta_{j,t}} \right)$. Given Proposition 4, for any model as outlined with exactly J market goods then (17) has exactly $\sum_{j=1}^{J-1} j$ representations assuming no two j 's are used together in the same equation more than once. We can safely ignore all but $J - 1$ of these if we treat one good, say $j = 1$ as the base consumption good. Let ξ_t be a $J - 1$ -dimensional vector with the following structure

$$\xi_t = \xi_{t-1} + \epsilon_t \quad (18)$$

ϵ_t is assumed orthogonal, capturing all variation in the home production process due to technical advancement, whether it be in the form of relative improvements to internal household capital or more efficient use of household time, or some unexplained exogenous variation in relative product quality. ξ_t takes a random walk process to account for the observed change in relative expenditure shares as seen in Figures 1 and 2. This specification allows us to quantify directly, in a constant elasticity environment, how much relative price variation has contributed to changes in log-consumption ratios of broad commodity categories verses other factors. In this sense, ζ_t could be thought to represent unobserved quality improvements impacting consumption by affecting marginal returns to use of market goods in home production. ξ_t is then just the relative quality improvement analog of ζ_t .

3.2 Bayesian Sampling Statements

Recall that the goal is to understand how much variation around relative quality/home-productivity verses relative prices has contributed to changes in relative consumption over time. The estimation routine operates on (17) and (18). For our analyses we consider two different specifications where the market goods are broad aggregates. First, let $j \in \{D, ND, S\}$ where D stands for the stock of durable goods, ND market purchases of non-durable goods, and S market purchases of services.¹⁶ Our first analysis operates on a two equation system where the base good is $j' = D$. Second, roll durable goods into the stock of non-consumable assets and let $j \in \{ND, S\}$. The second analysis thus operates on a single equation where we take the base good to be $j' = ND$.

Let \tilde{q}_t denote the $J - 1$ dimensional vector of log consumption ratios on the left hand side of (17) and \tilde{p}_t the corresponding log price ratios. Denote the log share ratios by the vector $\tilde{\theta}$. Consider first the full estimation with agents deriving utility from the service flow of durables. Our full set of observed data is

$$\mathcal{D} = \left\{ \{ \tilde{q}_t, \tilde{p}_t \}_{t=1}^T \right\} \quad (19)$$

Suppose that our orthogonal white noise process is distributed

$$\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{\Omega}) \quad (20)$$

¹⁶Our data series are constructed from the BEA's NIPA data. For a full description of the data see Appendix B.

We impose the following prior assumptions where normal distributions are parameterized in terms of variances for single dimension and variance/covariance matrices in the multidimensional case:¹⁷

$$- \sigma \sim \mathcal{N}(\sigma_0, 4) \quad (21)$$

$$- \sigma \cdot \tilde{\theta} \sim \mathcal{N}(\theta_0, \sqrt{6} \cdot \mathbb{I}_{(J-1) \times (J-1)}) \quad (22)$$

$$\xi_t \sim \mathcal{N}(\xi_{t-1}, 10 \cdot \mathbb{I}_{(J-1) \times (J-1)}) \quad \text{with } \xi_1 = \mathbf{0} \quad (23)$$

$$\Xi \sim \text{LKJ}(4) \quad (24)$$

$$\chi_j \sim \text{Cauchy}(0, 4) \quad \forall j \in \{1, \dots, J-1\} \quad (25)$$

$$\text{chol}(\Omega) = \text{diag}(\chi) \cdot \Xi \quad (26)$$

In the above $\text{chol}(\cdot)$ denotes the lower triangular Cholesky decomposition of a matrix. We sample the variance/covariance matrices using the LKJ distribution, the details of which are further discussed in Stan Development Team (2016).¹⁸ We fix σ_0 by running separate OLS regressions on each equation in (17) then taking the average over the reduced-form slope of each regression. We fix θ_0 as the share ratio in the first period of observation, which in our case is 1929. With $\xi_1 = \mathbf{0}$, the model is just a time-varying intercept model with a static component $\tilde{\theta}$ and a time-varying component ξ_t for $t > 1$. The full set of unknown structural parameters is

$$\mathcal{P} = \left\{ \sigma, \tilde{\theta}, \{\xi_t\}_{t=2}^T, \Omega \right\} \quad (27)$$

Let \mathcal{H} denote the set of all fixed hyperparameters. Then the posterior distribution of structural parameters given hyperparameters and data is

$$\pi(\mathcal{P} \mid \mathcal{D}, \mathcal{H}) \propto \pi(\mathcal{D} \mid \mathcal{P}, \mathcal{H}) \pi(\mathcal{P} \mid \mathcal{H}) \quad (28)$$

where $\pi(\mathcal{D} \mid \mathcal{P}, \mathcal{H})$ is the probability of observing data \mathcal{D} given the model structure. Note that $\pi(\mathcal{D} \mid \mathcal{P}, \mathcal{H})$ is the product of conditional normal distributions and $\pi(\mathcal{P} \mid \mathcal{H})$ is an analytically intractable prior density. We wish to estimate $\pi(\mathcal{P} \mid \mathcal{D}, \mathcal{H})$. To accomplish this we turn to advanced Monte Carlo techniques, specifically Hamiltonian Monte Carlo as described in Neal (2011). For details see Appendix C.

3.3 Estimation Results

Our estimation results suggest that relative price changes contribute to roughly 41% of the decline in relative non-durable to durable consumption and 67% of the decline in relative services to durable consumption. In a CES preference environment, this result suggests substantial relative

¹⁷For the case where $J = 2$, Ω is just a scalar which we denote ω_1 . Our sampling statements remain the same with the same fixed hyperparameters, except we need not sample Ξ and can simply draw $\omega_1 \sim \text{Cauchy}(0, 4)$ directly.

¹⁸Sampling a Cholesky-factored matrix directly is more efficient in H-MC than exploiting Wishart conjugacy (Stan Development Team 2016). The LKJ sampling procedure is an extension of work in Joe (2006).

returns to home productivity, specifically for durable goods, the denominator in both of the consumption ratios in the three good model (see Table 1). The estimated result also affirms some of the supply-side work, specifically Ngai and Pissarides (2007) suggesting that the between-industry productivity-differential is a major factor driving structural change. In the two goods model we would expect the consumption ratio between services and non-durables to have increased by 123% since 1929 when only accounting for price changes. Instead we observe a 93% increase, suggesting greater returns to home production wrought from the denominator in the ratio — non-durable goods. This, again, is to be expected. As Triplett and Bosworth (2004) and Ngai and Pissarides (2007) point out, the manufacturing sector (i.e. “goods” sector) has experienced relatively higher productivity growth than the services sector over the last half century. Our results suggest that this fact potentially leads to additional gains at the household level that are unaccounted for by available price indices.

Figures 3, 4, and 5 present the posterior fits of our estimation in black and a 95% confidence region in grey where the target dependent variables are respectively the log consumption ratios, $\tilde{q}_{ND/D,t}$, $\tilde{q}_{S/D,t}$, and $\tilde{q}_{S/ND,t}$.¹⁹ The red line shows the counterfactually predicted log consumption ratios under the assumption of no unmeasured quality changes or unmeasured returns to home productivity, so that $\xi_t = 0$ for all t . In this counterfactual we take all other parameter values to reside at their posterior means. Without latent returns to home productivity, we would expect both relative consumption of non-durable goods to durable goods and services to durable goods to fall due to declines in relative prices, though this only explains part of the decline suggesting substantial unmeasured relative returns to home production are contributing to the long-run trend. Table 1 shows that without changes to relative home productivity, we would expect the relative consumption of both ND/D and S/D to have decreased since 1929 by approximately 34% and 45% respectively, less than what is observed.

Figures 6 and 7 show that relative unobserved returns to home production appear to be led by unmeasured quality improvements in goods, assuming a CES preference environment. This makes sense in light of Figure 1 which shows that the share of services expenditure is outpacing the share of total consumption devoted toward services, suggesting home-productivity and quality of services is lagging that for goods forcing consumers to spend relatively more for every unit of real services consumption. This result is confirmed in the two goods model where Figure 8 shows the relative improvement in home productivity of non-durable goods to that of services.

4 Conclusion & Future Work

We have shown, theoretically, why home production matters when analyzing household consumption patterns. Observationally, a model with unobserved returns to home production is equivalent to one where preferences for consumed commodities exhibit non-linear Engel curves.

¹⁹Most figures reside in our estimation appendix, Appendix C. We leave what we consider to be our main results here in the main text.

Table 1: % Change in Relative Consumption Since 1929

| $\tilde{q}_{j/j'}$ | Actual | Under $\xi_t = 0$ | % Due to Prices |
|--------------------|--------|-------------------|-----------------|
| ND/D | -0.83 | -0.34 | 0.41 |
| S/D | -0.67 | -0.45 | 0.67 |
| S/ND | 0.93 | 1.15 | 1.23 |

We have arrived at results of similar magnitude to those in Boppart (2014), while supporting our particular accommodation of income effects with micro-foundations involving home production. We feel that this makes our paper stronger than previous papers that have merely assumed a non-homothetic preference structure without underlying, micro-founded justification.

Still, much work is to be done. Foremost, a more interesting analysis would involve examining the contributions of NIPA “level two” consumption components on structural change, specifically the role that improvements in electronic devices have played in contributing to increases in the relative real consumption of durables over the last 40 years. We are working right now on incorporating a richer analysis like this into a home-production model. Other future work involving home production must wait for the arrival of better data on how households allocate time and resources toward various tasks. To fully understand the long-run transformation of the economy, it is imperative that we think carefully about how and why (i.e., for what ends) households make purchases of market goods, including but not limited to how these goods are utilized toward the home production of an ultimate commodity. We hope that this paper inspires others to think creatively both about ways in which returns to home production may be endogenous to various economic inferences and the impact that omitting such features of household decisions may have on conclusions about important economic aggregates.

A Proofs

Lemma 1. Consider the j^{th} column of Q_t , $q_{j,t}$ and let q_t denote the J -dimensional vector whose components contain the total real amount of the market commodities used in all home production activities in period t . Under Assumption 3,

$$q_t^\top = \iota^\top \cdot Q_t \quad (\text{A.1})$$

Proof. Suppose inequality in (A.1). Note that the components of q_t are $q_{j,t}$, that is total consumption of the j^{th} market commodity. The case where there exists some $q_{j,t} < \sum_{i=1}^I q_{i,j,t}$ clearly violates both the household budget constraint and law of motion for durable commodities. If $q_{j,t} > \sum_{i=1}^I q_{i,j,t}$ then there exists left over service flow, unused in home production violating Assumption 3. \square

Proposition 1. Suppose an econometrician possesses no information on fixed capital inputs $a_{i,t}$ and time $n_{i,t}$ used in home production. Let $\lambda_{i,j,t}$ be the unobserved residual returns to commodity j in the production of final consumption commodity i . Under Assumptions 1 thru 2 we can implement reduced-form representations $g_{i,t}$ of the component functions $f_{i,t}$ of f_t as

$$g_{i,t}(q_{i,t}) = \left[\prod_{j=1}^J (\lambda_{i,j,t} \cdot q_{i,j,t})^{\omega_{i,j}} \right] \cdot \eta_{i,t}(a_{i,t}, n_{i,t}) \quad (\text{A.2})$$

where $\eta_{i,t}(a_{i,t}, n_{i,t})$ is homogeneous of degree $z \leq 1$. Further, assume $\sum_{j=1}^J \omega_{i,j} \leq 1$ and $0 < \omega_{i,j} < 1, \forall i, j$.

Proof. We need to prove that each $g_{i,t}$ satisfies our assumptions. With respect to Assumption 1, under the restrictions $\sum_{j=1}^J \omega_{i,j} \leq 1$ and $0 < \omega_{i,j} < 1, \forall i, j$, note that $g_{i,t}$ is homogenous of degree $z = \sum_{j=1}^J \omega_{i,j}$ which provides us satisfaction of this representation with Assumption 1.

For Assumption 2, note that

$$\frac{\partial \ln g_{i,t}}{\partial \ln q_{i,j,t}} = \omega_{i,j} \quad (\text{A.3})$$

□

Proposition 2. Let i and i' index two distinct final consumption goods which each require strictly positive amounts of market good j as inputs to home production. The internal household relative value of consumption of final good i to i' is equal to the infra-marginal rate of substitution for market good j :

$$\frac{\partial f_{i',t}}{\partial q_{i',j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j,t}} \quad (\text{A.4})$$

Proof. The relative value of consumption is the marginal rate of substitution. Let μ_t be the multiplier on the budget constraint. Differentiate $u_t(f_t(Q_t))$ in $q_{i,j,t}$ and $q_{i',j,t}$:

$$\frac{\partial u_t}{\partial f_{i,t}} \cdot \frac{\partial f_{i,t}}{\partial q_{i,j,t}} = p_{j,t} \cdot \mu_t \quad (\text{A.5})$$

$$\frac{\partial u_t}{\partial f_{i',t}} \cdot \frac{\partial f_{i',t}}{\partial q_{i',j,t}} = p_{j,t} \cdot \mu_t \quad (\text{A.6})$$

$$\Rightarrow \frac{\partial u_t}{\partial f_{i,t}} \bigg/ \frac{\partial u_t}{\partial f_{i',t}} = \frac{\partial f_{i',t}}{\partial q_{i',j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j,t}} \quad (\text{A.7})$$

where the left hand side of (A.7) is the marginal rate of substitution of good i to i' . □

Proposition 3. Consider some final consumption good $c_{i,t}$ produced by $f_{i,t}(q_{i,t})$. Suppose $c_{i,t}$ is produced using strictly positive amounts of market goods j and j' . For all internal household

production processes, the marginal rate of technical substitution between j and j' must equal their relative market price:

$$\frac{\partial f_{i,t}}{\partial q_{i,j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j',t}} = \frac{p_{j,t}}{p_{j',t}} \quad (\text{A.8})$$

Proof. Again let μ_t be the budget constraint multiplier. Differentiate $u_t(f_t(Q_t))$ in $q_{i,j,t}$ and $q_{i,j',t}$:

$$\frac{\partial u_t}{\partial f_{i,t}} \cdot \frac{\partial f_{i,t}}{\partial q_{i,j,t}} = p_{j,t} \cdot \mu_t \quad (\text{A.9})$$

$$\frac{\partial u_t}{\partial f_{i,t}} \cdot \frac{\partial f_{i,t}}{\partial q_{i,j',t}} = p_{j',t} \cdot \mu_t \quad (\text{A.10})$$

$$\Rightarrow \frac{\partial f_{i,t}}{\partial q_{i,j,t}} \bigg/ \frac{\partial f_{i,t}}{\partial q_{i,j',t}} = \frac{p_{j,t}}{p_{j',t}} \quad (\text{A.11})$$

□

Lemma 2. Suppose households take $\alpha_{j,t}$ as given for each j and just choose the total quantity of market goods to purchase each period. For market goods j and j' , the period- t marginal rate of substitution is

$$\left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j,t}} \cdot \alpha_{i,j,t} \right) \bigg/ \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j',t}} \cdot \alpha_{i,j',t} \right) \quad (\text{A.12})$$

Proof. The utility function under consideration is $u_t(f_t(q_t; \alpha_t))$. The components of f_t are

$$f_{i,t}(\alpha_{i,t} \cdot q_t^\top \cdot \mathbb{I} \cdot \iota) \quad (\text{A.13})$$

where the argument is the vector

$$(\alpha_{i,1,t} q_{1,t}, \alpha_{i,2,t} q_{2,t}, \dots, \alpha_{i,J,t} q_{J,t}) \quad (\text{A.14})$$

Differentiate $u_t(f_t(q_t; \alpha_t))$ in $q_{j,t}$ and $q_{j',t}$ to get

$$MU(q_{j,t}) = \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j,t}} \cdot \alpha_{i,j,t} \right) \quad (\text{A.15})$$

$$MU(q_{j',t}) = \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j',t}} \cdot \alpha_{i,j',t} \right) \quad (\text{A.16})$$

$$\Rightarrow MRS(j, j') = MU(q_{j,t}) / MU(q_{j',t}) = \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j,t}} \cdot \alpha_{i,j,t} \right) \bigg/ \left(\sum_{i=1}^I \frac{\partial f_{i,t}}{\partial q_{j',t}} \cdot \alpha_{i,j',t} \right) \quad (\text{A.17})$$

□

Corollary 1. General speaking, the observed rate at which households substitute between market

goods depends on the total number of productive tasks for which households utilize each market good.

Proof. This proof is trivial. Note that $MRS(j, j')$ as defined in (A.17) depends on I in the summation notations. \square

Lemma 3. Under Proposition 1 and Assumption 6, there exists a unique j such that for each i :

$$g_{i,t}(q_{i,t}; a_{i,t}, n_{i,t}) = \zeta_{j,t} \cdot q_{j,t} \quad (\text{A.18})$$

where $\zeta_{j,t} = \lambda_{i,j,t} \cdot \eta_{i,t}(a_{i,t}, n_{i,t}) > 0$.

Proof. Proposition 1 allows us to specify the reduced-form production function g_i . Consider the component $g_{i,t}$ from (A.2). Under Assumption 6, there exists a unique j such that

$$g_{i,t}(q_{i,t}; a_{i,t}, n_{i,t}) = \lambda_{i,j,t} \cdot q_{i,j,t} \cdot \eta(a_{i,t}, n_{i,t}) \quad (\text{A.19})$$

Assume the distributions of $\lambda_{i,j,t}$ and $\eta(a_{i,t}, n_{i,t})$ have strictly positive support, so that $\lambda_{i,j,t} \cdot \eta_{i,t}(a_{i,t}, n_{i,t}) > 0$ always. \square

Proposition 4. Given Lemma 3, the intratemporal first-order necessary conditions characterizing household market consumption can be written:

$$\ln \left(\frac{q_{j,t}}{q_{j',t}} \right) = -\sigma \cdot \ln \left(\frac{\theta_{j'}}{\theta_j} \right) - \sigma \cdot \ln \left(\frac{p_{j,t}}{p_{j',t}} \right) + \ln \left(\frac{\zeta_{j',t}}{\zeta_{j,t}} \right) \quad (\text{A.20})$$

for all $j \neq j'$.

Proof. Take the flow utility in (15) as given, and utilize the reduced-form representation $g_{i,t}$ of $f_{i,t}$ from Lemma 3. Note that the internal structure of U allows us to ignore γ , the intertemporal elasticity of substitution (see Ogaki and Reinhart (1998)). Letting μ_t be the budget constraint multiplier and differentiating in $q_{j,t}$ and $q_{j',t}$ yields

$$u_t(c_t)^{1-\rho} \theta_i \cdot (\zeta_{j,t} q_{j,t})^{\rho-1} = p_{j,t} \cdot \mu_t \quad (\text{A.21})$$

$$u_t(c_t)^{1-\rho} \theta_{i'} \cdot (\zeta_{j',t} q_{j',t})^{\rho-1} = p_{j',t} \cdot \mu_t \quad (\text{A.22})$$

$$\Rightarrow \frac{\theta_i}{\theta_{i'}} \left(\frac{\zeta_{j,t} q_{j,t}}{\zeta_{j',t} q_{j',t}} \right)^{\rho-1} = \frac{p_{j,t}}{p_{j',t}} \mu_t \quad (\text{A.23})$$

Without loss of generality, under Lemma 3 we can simply drop the i notation and set $i = j$ and $i' = j'$. Then, taking logs, rearranging, and dividing both sides by $\rho - 1$ yields

$$\ln \left(\frac{q_{j,t}}{q_{j',t}} \right) = -\sigma \cdot \ln \left(\frac{\theta_{j'}}{\theta_j} \right) - \sigma \cdot \ln \left(\frac{p_{j,t}}{p_{j',t}} \right) + \ln \left(\frac{\zeta_{j',t}}{\zeta_{j,t}} \right) \quad (\text{A.24})$$

where $\sigma = \frac{1}{1-\rho}$. \square

B Data

We use the National Income and Product Accounts and Fixed Asset tables from the Bureau of Economic Analysis to construct our data series. The expenditure series are the annual, seasonally unadjusted current dollar expenditure series for non-durable goods and services since 1929. For non-durable goods and services, we arrive at a real consumption series by using the associated non-durable and services chain-price indices, normalizing them to unity for the year 2017. Thus, all real values are in 2017 dollars. For durable goods we use the BEA’s current dollar value estimate for consumer durable goods from the Fixed Asset tables to arrive at the nominal value of durables. Thus, all durable “expenditure” estimates represent the nominal market value of the flow of durable services. Durable “consumption” estimates are arrived at by setting the NIPA annual chain price index for durables to unity for 2017, then dividing the nominal durable consumption series from the Fixed Asset tables by this new, 2017-normalized price index.

C Estimation Details

C.1 Description of Hamiltonian Monte Carlo

For a thorough and readable, detailed treatment of Hamiltonian Monte Carlo (H-MC), we recommend reading Neal (2011). For an equally thorough, more technical treatment, read Betancourt and Stein (2011). For a brief overview of H-MC see Gelman et al. (2013b). Here, we provide the “Reader’s Digest” overview of the sampler’s properties and benefits since there are very few examples of this particular estimation technique being used in the econometrics literature.²⁰ This Appendix is not intended to be a full treatment, only an outline.

Recall, in traditional Bayesian MCMC posterior sampling (i.e. Gibbs sampling or Metropolis-Hastings sampling), the state of the sampler at iteration m is the m^{th} draw of the parameter set \mathcal{P}^m from the posterior distribution with density $\pi(\mathcal{P} \mid \mathcal{D}, \mathcal{H})$. Thus a traditional sampler is a Markov Chain operating in a discrete, countable sampling space, with discrete dynamics. For example, the m^{th} draw of a Metropolis-Hastings sampler depends on the $m - 1$ draw, and the econometrician knows this distribution has converged when autocorrelation between draws has been sufficiently minimized, depending on the nature of the estimation problem. H-MC extends this concept to a sample space with continuous dynamics.²¹ In Hamiltonian dynamics commonly employed in physical mechanics, the state of the system depends on both the forward “momentum” of the system \mathcal{Q} and the “position” of the system \mathcal{P} , each of the same dimension. The Hamiltonian

²⁰The only published example in the economics literature we were able to find is Burda (2015). Martin Burda also has a recent working paper providing an application of H-MC to a discrete choice model (Burda and Daviet 2018).

²¹Recall, “dynamics” in this context refer to the dynamics of the parameter sampler, *not* actual time, as in the modeled structural dynamics of the underlying agent’s decision process. One dynamic operates over the sampler dimension m while the other over the data dimension t . This can be confusing when estimating models wrought from a dynamic decision process or which have dynamic time series components.

equation associated with this system is the sum of “potential” $U(\cdot)$ and “kinetic” $K(\cdot)$ energy:

$$H(\mathcal{Q}, \mathcal{P}) = U(\mathcal{P}) + K(\mathcal{Q}) \quad (\text{C.1})$$

The partial derivatives of this equation will determine how the parameter space \mathcal{P} evolves as the sampler proceeds, as well as the rate of this evolution \mathcal{Q} .

Since the “dynamics” of our particular task involve traversing a parameter space over which some posterior density function $\pi(\mathcal{P} \mid \mathcal{D}, \mathcal{H})$ is defined, the potential energy function for H-MC is just the negative log of the right hand side of (28):

$$U(\mathcal{P}) = -\ln \left[\pi(\mathcal{D} \mid \mathcal{P}, \mathcal{H}) \pi(\mathcal{P} \mid \mathcal{H}) \right] \quad (\text{C.2})$$

Note that we can sample from this using H-MC only if $U(\mathcal{P})$ is continuously differentiable over the entire sample space since H-MC operates on Hamilton’s equations which require computation of the gradient vectors for both $U(\cdot)$ and $K(\cdot)$. Let $\#$ denote the cardinality of a countable set. In practice, the kinetic energy function is defined conditional on \mathcal{P} and taken to be a quadratic of the form:²²

$$K(\mathcal{Q} \mid \mathcal{P}) = \frac{1}{2} \sum_{k=1}^{\#Q} \sum_{r=1}^{\#Q} \mathcal{Q}_k \mathcal{Q}_r \Lambda_{k,r}(\mathcal{P}) - \frac{1}{2} \ln \left[\det(\Lambda(\mathcal{P})) \right] \quad (\text{C.3})$$

where $\Lambda(\mathcal{P})$ is what is called the mass matrix and may be constant. It could also be restricted to the identity matrix or simply be diagonal. At most it is a dense symmetric positive definite matrix that represents the variance/covariance of an underlying conditional Gaussian distribution function for the momentum vector:

$$\mathcal{Q} \mid \mathcal{P} \sim \mathcal{N}(\mathbf{0}, \Lambda(\mathcal{P})) \quad (\text{C.4})$$

A dense $\Lambda(\mathcal{P})$ can help account for high local non-linearities in $U(\mathcal{P})$, though can be difficult to compute (Betancourt and Stein 2011). In our estimation, we allow $\Lambda(\mathcal{P})$ to be diagonal and tune $\Lambda(\mathcal{P})$ during a warm-up period using Stan’s H-MC implementation (Stan Development Team 2016).²³

Having defined the objects on which we operate, we can summarize the algorithm described in detail in Neal (2011). Here is where the kinetic energy component of the Hamiltonian helps greatly speed up convergence and reduce autocorrelation in our sampling routine. After a suffi-

²²This is a more flexible form of the kinetic energy function than that in Neal (2011). See Betancourt and Stein (2011) for more details.

²³Stan is free software for high-performance H-MC which utilizes automatic differentiation, LAPACK, and BLAS libraries for computational efficiency and which can be implemented and executed in a number of top-end scientific computing programs like R, Python, Matlab, or executed simply from the shell. See <http://mc-stan.org>. For examples using the Stan language to execute H-MC models, see Gelman et al. (2013a).

cient warm-up period,^{24,25} a sampling step proceeds as follows:

- i. Given the position from the previous iteration \mathcal{P} , draw a new \mathcal{Q} from (C.4). In a sense (C.4) is thus the H-MC analog of what is commonly called a “proposal distribution” for an MCMC Metropolis-Hastings algorithm.
- ii. Given $(\mathcal{Q}, \mathcal{P})$, iterate on Hamilton’s equations using the leapfrog method for L steps to get a pair $(\mathcal{Q}', \mathcal{P}')$.²⁶ This is your proposed new state.
- iii. Similar to the Metropolis-Hastings algorithm, draw a uniform random deviate $u \sim \mathcal{U}[0, 1]$ and accept $(\mathcal{Q}', \mathcal{P}')$ as the next state if

$$u < \exp \{ -H(\mathcal{Q}', \mathcal{P}') + H(\mathcal{Q}, \mathcal{P}) \} \quad (\text{C.5})$$

H-MC is most useful when sampling from $\pi(\mathcal{P} \mid \mathcal{D}, \mathcal{H})$ is computationally burdensome and may require an incredibly long chain of discrete draws in order to achieve convergence. This happens in cases where posterior draws are likely to be highly autocorrelated in a traditional MCMC, especially when using Metropolis-Hastings algorithms where high autocorrelation may also require low acceptance probabilities, thus further giving reason for a long sampling chain.

Testing for convergence of an H-MC run to its posterior distribution is handled by computing what is commonly called the R -hat statistic (\hat{R}),²⁷ coined by Gelman and Rubin (1992) and described in detail in the Stan documentation (Stan Development Team 2016). The statistic describes the ratio of between-sample variance to within-sample variance, thus requiring the modeler to run the H-MC simulation for multiple, parallel and independent chains. As $\hat{R} \rightarrow 1$ the H-MC simulated posterior converges to the underlying model’s true posterior. In general, between-sample variance is greater than within-sample variance, so one usually wants $\hat{R} < 1.2$ for each parameter, including the potential energy $U(\mathcal{P})$, to indicate that the simulation has converged. The H-MC procedure allows users flexibility to choose tuning parameters like step-sizes, adaptation acceptance rates, and total number of warm-up and sampling iterations. Again, this was not meant to be a thorough treatment of H-MC techniques, only an overview. For more information on how to apply and code an H-MC estimation routine, we recommend that curious researchers read the documentation for Stan (Stan Development Team 2016).

C.2 H-MC Detailed Results

We use Stan’s default tuning parameter values except for the adaptation acceptance rate, which we find from simulations is best picked to be $\Delta = 0.85$, as opposed to the default value of $\Delta = 0.8$.

²⁴This is akin to the “burnin” period for an MCMC operating under discrete dynamics.

²⁵This is handled automatically by Stan and discussed thoroughly in the software manual, (Stan Development Team 2016).

²⁶See Neal (2011) for a detailed description of this iterative method and how it operates on Hamilton’s equations.

²⁷This statistic is also referred to synonymously as the “potential scale reduction statistic” (Gelman and Shirley 2011; Stan Development Team 2016).

The H-MC is run in parallel on eight independent chains each with a sampling space of size 1000 for a total posterior sample size of $N = 8000$. Table 2 presents summary statistics for the posterior distribution estimates of main parameters from both the three good model with durables, non-durables, and services and the two good model featuring just non-durable goods and services. We do not present distribution estimates for $\{\xi_t\}_{t=2}^T$. These time series are plotted along with a 95% confidence region in Figures 6, 7, and 8. In Table 2, the last three columns represent distribution percentiles, with a 95% confidence region buttressed on either side by the 2.5-percentile and 97.5-percentile. To say the model has converged, ideally we want $\hat{R} < 1.2$. For the three good model this happens for all main parameters, but not the potential energy $U(\mathcal{P})$. The two good model is plagued by poor convergence over all parameters, suggesting the need to increase the warm-up sequence, increase the adaptation Δ , or sample only over centered values.²⁸ N_{eff} measures the effective sample size of the H-MC. In the asymptotic limit, as the H-MC converges to the true posterior, $N_{eff} \rightarrow N$. Low effective samples means that the model is poorly exploring the parameter space. The three good model, in this regard, is performing significantly better than the two good model, which after 8000 iterations does not appear to have completely converged and thus is plagued by low N_{eff} .

Table 2: H-MC Posterior Results

| 3 Good Model | \hat{R} | N_{eff} | Mean | S.D. | 2.5% | 50% | 97.5% |
|-------------------------|-----------|-----------|-------|-------|---------|-------|--------|
| $\tilde{\theta}_{ND/D}$ | 1.04 | 224 | 0.80 | 1.33 | -1.70 | 0.78 | 3.42 |
| $\tilde{\theta}_{S/D}$ | 1.03 | 256 | 1.03 | 1.71 | -2.15 | 0.99 | 4.49 |
| $-\sigma$ | 1.02 | 276 | -0.34 | 0.99 | -2.24 | -0.37 | 1.67 |
| $\Omega_{1,1}$ | 1.09 | 68 | 0.73 | 0.84 | 0.02 | 0.43 | 3.13 |
| $\Omega_{2,1}$ | 1.01 | 1152 | 0.00 | 0.23 | -0.48 | 0.00 | 0.48 |
| $\Omega_{2,2}$ | 1.11 | 86 | 0.71 | 0.90 | 0.03 | 0.38 | 3.30 |
| $U(\mathcal{P})$ | 1.30 | 25 | 1.53 | 76.86 | -130.03 | -3.08 | 178.78 |
| 2 Good Model | \hat{R} | N_{eff} | Mean | S.D. | 2.5% | 50% | 97.5% |
| $\tilde{\theta}_{S/ND}$ | 1.30 | 11 | 1.20 | 1.14 | -1.02 | 1.13 | 3.56 |
| $-\sigma$ | 1.31 | 11 | 1.44 | 1.94 | -2.07 | 1.38 | 5.67 |
| $\sqrt{\omega_1}$ | 1.39 | 11 | 0.54 | 0.36 | 0.05 | 0.47 | 1.42 |
| $U(\mathcal{P})$ | 1.64 | 7 | 34.91 | 74.59 | -76.01 | 22.04 | 224.32 |

²⁸For the current version, our sampling distributions are un-centered. See Stan Development Team (2016) for convergence improvements with normalized sampling statements.




Figure 3: The actual data points are the dots, the fitted H-MC posterior mean is the black line running through the dots, and the grey area is a 95% confidence region for the posterior predicted values. The red line represents the counterfactually predicted time series when $\xi_{ND/D,t} = 0, \forall t$.




Figure 4: The details are the same as above except the relative series we are looking at is that of services to durables S/D .




Figure 5: Here, the relative series is from the two good model, S/ND .




Figure 6: Relative home-productivity/quality residual, $\xi_{ND/D,t} = \ln(\zeta_{D,t}/\zeta_{ND,t})$ with 95% confidence region.




Figure 7: Relative home-productivity/quality residual, $\xi_{S/D,t} = \ln(\zeta_{D,t}/\zeta_{S,t})$ with 95% confidence region. The decline in relative productivity is less stark than in the *ND* case.




Figure 8: Relative home-productivity/quality residual in two good model, $\xi_{S/ND,t} = \ln(\zeta_{ND,t}/\zeta_{S,t})$ with 95% confidence region. Non-durable goods have become relatively more productive than services.

References

- Acemoglu, Daron, and Veronica Guerrieri. 2008. "Capital Deepening and Nonbalanced Economic Growth". *Journal of Political Economy* 116 (3). (Cit. on p. 12).
- Aguiar, Mark, and Erik Hurst. 2007. "Life-Cycle Prices and Production". *The American Economic Review* 97 (5). (Cit. on p. 5).
- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use During the Great Recession". *The American Economic Review* 103 (5). (Cit. on p. 5).
- Baumol, William. 1967. "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis". *The American Economic Review* 57 (3). (Cit. on p. 2).
- Baumol, William, Sue Anne Batey Blackman, and Edward Wolff. 1985. "Unbalance growth revisited: asymptotic stagnancy and new evidence". *The American Economic Review* 75 (4): 806–817. (Cit. on p. 1).
- Becker, Gary. 1965. "A Theory of the Allocation of Time". *The Economic Journal* 75 (299). (Cit. on pp. 2, 3, 5, 7, 9).
- Becker, Gary, and Kevin Murphy. 1988. "A Theory of Rational Addiction". *Journal of Political Economy* 96 (4). (Cit. on p. 6).
- Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations". *Journal of Political Economy* 99 (6). (Cit. on p. 5).
- Betancourt, Michael, and Leo Stein. 2011. "The Geometry of Hamiltonian Monte Carlo". *arXiv:1112.4118*. (Cit. on pp. 21, 22).
- Boppart, Timo. 2014. "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences". *Econometrica* 82 (6). (Cit. on pp. 1, 2, 12, 17).
- Burda, Martin. 2015. "Constrained Hamiltonian Monte Carlo in BEKK GARCH with Targeting". *Journal of Time Series Econometrics* 7 (1). (Cit. on p. 21).
- Burda, Martin, and Remi Daviet. 2018. "Hamiltonian Sequential Monte Carlo with Application to Consumer Choice Behavior". *University of Toronto, Department of Economics Working Paper #618*. (Cit. on p. 21).
- Caselli, Francesco, and Wilbur Coleman. 2001. "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation". *Journal of Political Economy* 109 (3). (Cit. on p. 1).
- Echevarria, Cristina. 1997. "Changes in Sectoral Composition Associated with Economic Growth". *International Economic Review* 38 (2). (Cit. on pp. 1, 12).
- Foellmi, Reto, and Josef Zweimuller. 2008. "Structural change, Engel's consumption cycles and Kaldor's facts of economic growth". *Journal of Monetary Economics* 55. (Cit. on p. 12).

- Gelman, Andrew, and Donald Rubin. 1992. "Inference from Iterative Simulation Using Multiple Sequences". *Statistical Science* 7 (4). (Cit. on p. 23).
- Gelman, Andrew, and Kenneth Shirley. 2011. *Inference from Simulations and Monitoring Convergence*. Ed. by Steve Brooks et al. (Cit. on p. 23).
- Gelman, Andrew, et al. 2013a. *Computation in R and Stan*. (Cit. on p. 22).
- . 2013b. *Computationally efficient Markov chain simulation*. (Cit. on p. 21).
- Gomme, Paul, Finn Kydland, and Peter Rupert. 2001. "Home Production Meets Time to Build". *Journal of Political Economy* 109 (5). (Cit. on pp. 1, 5, 6).
- Gordon, Robert. 2016. *The Rise and Fall of American Growth*. Princeton University Press. (Cit. on p. 3).
- Graham, John, and Carole Green. 1984. "Estimating the Parameters of a Household Production Function with Joint Products". *The Review of Economics and Statistics* 66 (2). (Cit. on p. 5).
- Gronau, Reuben. 1977. "Leisure, Home Production, and Work – the Theory of the Allocation of Time Revisited". *Journal of Political Economy* 85 (6). (Cit. on p. 5).
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi. 2013. "Two Perspectives on Preferences and Structural Transformation". *American Economic Review* 103 (7). (Cit. on pp. 1, 2, 12).
- Joe, Harry. 2006. "Generating random correlation matrices based on partial correlations". *Journal of Multivariate Analysis* 97. (Cit. on p. 15).
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie. 2001. "Beyond Balanced Growth". *Review of Economic Studies* 68:869–882. (Cit. on p. 1).
- Kydland, Finn, and Edward Prescott. 1982. "Time to Build and Aggregate Fluctuations". *Econometrica* 50 (6). (Cit. on p. 6).
- Laitner, John. 2000. "Structural Change and Economic Growth". *Review of Economic Studies* 67. (Cit. on p. 1).
- McGrattan, Ellen, Richard Rogerson, and Randall Wright. 1993. "Household Production and Taxation in the Stochastic Growth Model". *Federal Reserve Bank of Minneapolis: Staff Report #166*. (Cit. on p. 5).
- Neal, Radford. 2011. *MCMC using Hamiltonian dynamics*. Ed. by Steve Brooks et al. (Cit. on pp. 15, 21–23).
- Ngai, Rachel, and Christopher Pissarides. 2007. "Structural Change in a Multisector Model of Growth". *American Economic Review* 97 (1). (Cit. on pp. 2, 12, 16).
- Ogaki, Masao, and Carmen Reinhart. 1998. "Measuring Intertemporal Substitution: The Role of Durable Goods". *Journal of Political Economy* 106 (5). (Cit. on pp. 8, 20).

- Rupert, Peter, Richard Rogerson, and Randall Wright. 1995. "Estimating Substitution Elasticities in Household Production Models". *Economic Theory* 6 (1). (Cit. on p. 5).
- Solow, Robert. 1957. "Technical Change and the Aggregate Production Function". *The Review of Economics and Statistics* 39 (3). (Cit. on p. 9).
- Stan Development Team. 2016. *Stan Modeling Language: User's Guide and Reference Manual*. (Cit. on pp. 15, 22–24).
- Stigler, George, and Gary Becker. 1977. "De Gustibus Non Est Disputandum". *The American Economic Review* 67 (2). (Cit. on pp. 5–7).
- Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the US Productivity Slowdown". *Journal of Economic Perspectives* 31 (2). (Cit. on p. 2).
- Triplett, Jack, and Barry Bosworth. 2004. *Productivity in the U.S. Services Sector: New Sources of Economic Growth*. Brookings Institution Press. (Cit. on p. 16).