

Structural Change with Time to Consume

Technical Appendix (for online publication only)

W.L. Bednar* Nick Prettnar^{23†}

²*Laboratory for Aggregate Economics and Finance, UCSB*

³*Tepper School of Business, CMU*

October 23, 2025

A Personal Care and Sleep Time Data Appendix

In this appendix we highlight the role that personal care and sleep time play in contributing to overall goods time. When personal care and sleep time are excluded from the goods series, we keep only grooming activities, t010201 and t010299. The rationale behind this is that grooming activities generally require utilization of physical goods, while sleep and other unspecified personal care activities may not. Figures A.1a and A.1b are analogous to those in the main text, while Figures A.1c and A.1d show that personal care time has been increasing, though the fraction of personal care time associated with goods time has been relatively stable, with the exception of the COVID-19 pandemic.

*wbednar101@gmail.com

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




Figure A.1: Panels (a) and (b) are analogous to those from Figure 2 in the main text, except personal care plus sleep time have been removed. Panels (c) and (d) show personal care time in hours per day from the ATUS and the fraction of goods time associated with personal care, respectively.

B Model Appendix

B.1 Income and Substitution Effects in Response to Wage Variation

In this appendix we characterize Slutsky equations to analyze income and substitution effects for a model without parametric functional forms for the utility function and consumption technologies. Assume utility is at least four times differentiable, strictly increasing, strictly concave, and homogeneous of degree one. Assume that the consumption technologies have the same properties. Below, we will show that wage variation affects the quantities of market purchases via *both* a total income effect and a substitution effect in models where utility and the consumption technology have the described properties. Similarly, market price variation will affect the allocation of off-market time via two separate channels.

We require the expenditure-minimization dual problem (EMP) associated with the utility-

maximization problem (UMP). To build toward an EMP we also require the Beckerian budget constraint, which can be derived by substituting the time-use constraint into the budget constraint for ℓ :¹

$$\sum_{j \in \{g,s\}} (P_j q_j + w n_j) \leq w \bar{n} + y \quad (\text{B.1})$$

The right-hand side of (B.1) is maximum-possible income net of savings if the consumer devoted all of his/her time to labor-market activities (Becker 1965). Blundell and Macurdy (1999) show that the objective function for the EMP in models featuring at least one off-market time-utilization decision and elastic labor is just the left-hand side of (B.1). This yields the following expenditure function:

$$e(P_g, P_s, w) = P_g q_g + P_s q_s + w(n_g + n_s) \quad (\text{B.2})$$

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

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

Since the prices of n_g and n_s are constrained to be identical, this will impact Shephard's Lemma generated from a home production model. We now describe a version of Shephard's Lemma for the expenditure function in (B.2). Let superscript m index the Marshallian demand functions derived by solving the UMP. Let superscript h index the Hicksian demand functions derived by solving the EMP. Note that Marshallian demands are functions of consumption prices, wages, and capital income net of savings, while Hicksian demands are functions of consumption prices, wages, and utility.

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

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

¹We dispense with i subscripts in this appendix only.

Proof. Note that

$$e(P_g, P_s, w, ; \bar{u}) = P_g q_g^h + P_s q_s^h + w n_g^h + w n_s^h \quad (\text{B.3})$$

where n_j^h are the Hicksian demands for off-market time use in process j . Differentiating (B.3) in w we get

$$\frac{\partial e}{\partial w} = P_g \frac{\partial q_g^h}{\partial w} + P_s \frac{\partial q_s^h}{\partial w} + w \frac{\partial n_g^h}{\partial w} + w \frac{\partial n_s^h}{\partial w} + n_g^h + n_s^h \quad (\text{B.4})$$

Letting λ be the multiplier on the budget constraint, we replace prices with the first-order conditions from the UMP, where we set $w = \frac{\partial u}{\partial n_j^h} \frac{1}{\lambda}$, next to its corresponding Hicksian partial derivative:

$$\frac{\partial e}{\partial w} = \frac{\partial u}{\partial q_g^h} \frac{1}{\lambda} \frac{\partial q_g^h}{\partial w} + \frac{\partial u}{\partial q_s^h} \frac{1}{\lambda} \frac{\partial q_s^h}{\partial w} + \frac{\partial u}{\partial n_g^h} \frac{1}{\lambda} \frac{\partial n_g^h}{\partial w} + \frac{\partial u}{\partial n_s^h} \frac{1}{\lambda} \frac{\partial n_s^h}{\partial w} + n_g^h + n_s^h \quad (\text{B.5})$$

By the fact that $u(q_g^h, q_s^h, n_g^h, n_s^h) = \bar{u}$ holds for all prices including wages, and given the Hicksian demand functions minimize the Lagrangian for the EMP:

$$\frac{\partial u}{\partial q_g^h} \frac{1}{\lambda} \frac{\partial q_g^h}{\partial w} + \frac{\partial u}{\partial q_s^h} \frac{1}{\lambda} \frac{\partial q_s^h}{\partial w} + \frac{\partial u}{\partial n_g^h} \frac{1}{\lambda} \frac{\partial n_g^h}{\partial w} + \frac{\partial u}{\partial n_s^h} \frac{1}{\lambda} \frac{\partial n_s^h}{\partial w} = 0 \quad (\text{B.6})$$

QED

Lemma B.1 follows directly from the fact that n_g and n_s always have the same price but are separate decisions which will differ from each other due strictly to the structures of $f_g(\cdot, \cdot)$, $f_s(\cdot, \cdot)$, and $u(\cdot, \cdot)$. Note that the fact that the price set is smaller than the choice set will impact Shephard's Lemma in this case. This has implications for the terms of the cross-price responsiveness of market commodities q_j to wages w , which we will characterize under the EMP specified by (B.2).²

We now describe two different ways to characterize the responsiveness of demand to wage variation. Let $\tilde{y}(w) = w\bar{n} + y$. Note that under Assumption B.1 we can write Marshallian demand as a linear function in $\tilde{y}(w)$:

$$q_j^m(P_g, P_s, w, \tilde{y}(w)) = \tilde{y}(w) \Gamma(P_g, P_s, w) \quad (\text{B.7})$$

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

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

$$\frac{dq_j^m}{dw} = \frac{\partial q_j^m}{\partial w} + \frac{\partial q_j^m}{\partial \tilde{y}} \frac{\partial \tilde{y}}{\partial w}$$

²By "cross-price responsiveness" we mean the responsiveness of q_j to wage variation, since the wage is *both* a price and proportional to labor income.

$$\text{where } \frac{\partial \tilde{y}}{\partial w} = \bar{n}$$

so that

$$\frac{dq_j^m}{dw} = \tilde{y}(w) \frac{\partial \Gamma}{\partial w} + \Gamma(P_g, P_s, w) \bar{n} \quad (\text{B.8})$$

Proof. Clearly w enters into both $q_j^m(\dots)$ as an independent argument and in the composite function $\tilde{y}(w)$. Thus, the total derivative of q_j^m with respect to w is obvious. Further, $\tilde{y}(w)$ is linear in w , so $\frac{\partial \tilde{y}}{\partial w} = \bar{n}$. Under Assumption 1 it is clear that

$$\frac{dq_j^m}{dw} = \tilde{y}(w) \frac{\partial \Gamma}{\partial w} \quad \text{and} \quad \frac{\partial q_j^m}{\partial \tilde{y}} \frac{\partial \tilde{y}}{\partial w} = \Gamma(P_g, P_s, w) \bar{n}$$

QED

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

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

Proof. The proof is the standard one, where the version of Shephard's Lemma used is that of Lemma B.1. Note that in equilibrium

$$q_j^m(P_g, P_s, w, e(P_g, P_s, w; \bar{u})) = q_j^h(P_g, P_s, w, \bar{u}) \quad (\text{B.9})$$

Differentiate the above to get the following:

$$\frac{\partial q_j^m}{\partial w} + \frac{\partial q_j^m}{\partial e} \frac{\partial e}{\partial w} = \frac{\partial q_j^h}{\partial w} \quad (\text{B.10})$$

Use Lemma B.1 to replace $\frac{\partial e}{\partial w}$ then rearrange to get the result.

QED

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

affects $w \bar{n}$, i.e. the value of time.

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

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

$$\frac{dq_j^m}{dw} = \underbrace{\frac{\partial q_j^h}{\partial w}}_{\text{Substitution Effect}} - \underbrace{\frac{\partial q_j^m}{\partial \tilde{y}}(n_g + n_s)}_{\text{Income Effect}} + \underbrace{\frac{\partial q_j^m}{\partial \tilde{y}} \bar{n}}_{\text{Time-value Effect}}$$

Proof. Using Lemma B.2 as the baseline, replace the partial derivative of $\frac{\partial q_j^m}{\partial w}$ with the Slutsky equation in Lemma B.3. **QED**

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

$$\frac{dq_j^m}{dw} = \underbrace{\frac{\partial q_j^h}{\partial w}}_{\text{Substitution Effect}} + \underbrace{\frac{\partial q_j^m}{\partial \tilde{y}} \ell}_{\text{Total Income Effect}}$$

Proof. Note that $\ell = \bar{n} - n_g - n_s$. Using Proposition B.1, combine the second and third terms on the right and replace. **QED**

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

B.2 Elasticity of Relative Expenditure with Respect to Wages

In this section we demonstrate that the parameterized version of our nested-CES model has rich and flexible properties which could explain many different kinds of re-allocative processes, from

re-allocative effects which impact households differently across the wage distribution to other re-allocative effects that result from households at different income levels being exposed to product-market inflation in different ways. Further, the re-allocative effects we highlight here are independent of variation, either cross-sectionally or over time, in the productivities, z 's, of the consumption technologies.

We are interested in understanding how heterogeneity in different consumption processes leads to the appearance that re-allocation of expenditure is driven by rising incomes. The allocations of market expenditure and off-market time depend on three relative prices — P_g/P_s , w_i/P_g , and w_i/P_s . This can be seen by inspecting the relative expenditure and time-use conditions from the main text. Typically, in the structural change literature, we only care about how either relative prices affect allocations or income. In this model variation in w_i affects both income *and* relative prices. This necessitates an analysis as to how varying income while holding market prices fixed and vice-versa affects allocations.

We consider three questions. First, how does variation in the structures of the consumption technologies affect the degree to which relative expenditure responds to wages? Further, do changing relative prices have second-order effects on the responsiveness of expenditure allocations to rising wages? What about product-market price inflation relative to wage inflation? To answer these questions, we derive the following elasticity. Take logs and differentiate relative expenditure in $\ln w_i$ to get the elasticity of relative expenditure, x_{ig}/x_{is} , with respect to wages, which we denote $\varepsilon_i^{x,w}$.³

$$\begin{aligned} \varepsilon_i^{x,w} := \frac{\partial \ln(x_{ig}/x_{is})}{\partial \ln w_i} = & \left(\frac{\nu_g - \rho}{\nu_g(\rho - 1)} \right) \frac{(1 - \omega_g)^{\frac{\nu_g}{\nu_g - 1}} \left[\frac{w_i \omega_g}{P_g(1 - \omega_g)} \right]^{\frac{\nu_g}{\nu_g - 1}}}{\omega_g + (1 - \omega_g) \left[\frac{w_i \omega_g}{P_g(1 - \omega_g)} \right]^{\frac{\nu_g}{\nu_g - 1}}} \\ & + \left(\frac{\rho - \nu_s}{\nu_s(\rho - 1)} \right) \frac{(1 - \omega_s)^{\frac{\nu_s}{\nu_s - 1}} \left[\frac{w_i \omega_s}{P_s(1 - \omega_s)} \right]^{\frac{\nu_s}{\nu_s - 1}}}{\omega_s + (1 - \omega_s) \left[\frac{w_i \omega_s}{P_s(1 - \omega_s)} \right]^{\frac{\nu_s}{\nu_s - 1}}}. \end{aligned} \quad (\text{B.11})$$

Notice that the ratio z_{ig}/z_{is} enters the relative expenditure condition log-linearly, so that it does not feature in (B.11). The elasticity $\varepsilon_i^{x,w}$ is different from the elasticity of *absolute* (not relative) expenditure with respect to wages.⁴ Rather, (B.11) describes how wage variation affects the *allocation* of market consumption, holding market prices, P_g and P_s , fixed. Wages will affect relative spending, and thus expenditure shares, because of differences between the marginal rates of technical substitution across the different home production processes. Note that, even if $\nu_g = \nu_s$, as long as the relative intensities of market purchases across production processes are different (i.e., $\omega_g \neq \omega_s$),

³Throughout our exposition ε will be used to denote elasticities.

⁴The elasticity of absolute expenditure on product j with respect to wages is described for a general model in Equation 17 of Aguiar, Hurst, and Karabarbounis (2012). We add to their theoretical contributions by considering how wage variation affects *relative allocations* in a model with consumption time.

then wage variation will non-trivially impact the expenditure allocation.⁵

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




Figure B.1: We present level sets of $\varepsilon_i^{x,w}$, under variation in w_i and P_s , holding P_g fixed. Each panel features a different parameterization, so as to illustrate how the values of the underlying home production elasticities affect the responsiveness of the expenditure allocation to wage variation.

Figure B.1 plots the level sets of $\varepsilon_i^{x,w}$ against variation in w_i and P_s , holding $P_g = 1$ fixed. In each plot we let $\rho < 0$, so that consumption experiences are gross complements.⁶ Figure B.1 is meant to show how variation in relative market prices, P_g/P_s , has second-order effects on the responsiveness of the expenditure allocation to wage variation. In Figure B.1 we consider two parameterizations to demonstrate that the monotonicity of $\varepsilon_i^{x,w}$ depends on differences in the production weights, ω_j 's, across consumption activities: in panel (a) we present a parameterization that yields a monotonic increase in $\varepsilon_i^{x,w}$ in both the ratios w_i/P_g and P_s/P_g ; in panel (b) we present the level sets for an alternative parameterization that yields non-monotonic variation of $\varepsilon_i^{x,w}$. No-

⁵While we have assumed that the consumption technologies are not Cobb-Douglas, this is true even if they are (i.e., $\nu_g = \nu_s = 0$). When the consumption technologies are Cobb-Douglas, the elasticity of relative expenditure with respect to wages is $\varepsilon_i^{x,w} = \frac{\rho(\omega_s - \omega_g)}{\rho - 1}$. This object is non-zero as long as $\omega_g \neq \omega_s$ and $\rho \neq 0$. With Cobb-Douglas consumption technologies, relative spending will change in fixed proportion to wage growth. There are two potential shortcomings of this. First, such a model would not capture cases where consumers across the income distribution may adjust their expenditure shares at different rates in response to proportional wage gains. Second, this would not allow for variation in market prices to affect how the spending allocation responds to wage variation (i.e., second-order, cross-partial effects).

⁶This is the sign of ρ in our preferred quantitative models.

tice, under both parameterizations time use and market purchases are complementary in each process (i.e., $\nu_j < 0, \forall j$), while the final consumption experiences, c_{ij} , are also complementary because $\rho < 0$. Nonetheless, the nature of $\varepsilon_i^{x,w}$ is very different between the two, and in the case of panel (b), monotonicity of $\varepsilon_i^{x,w}$ depends on the value of P_s and, by extension, the relative price of goods and services, P_g/P_s . This can be seen by noting that, in panel (b), middle-income consumers are more elastic when $P_s < 1$, while we observe what appears to be monotonically declining elasticities across the wage distribution when $P_s > 1$.

It is clear that the rate at which consumers re-allocate their market expenditure in response to wage gains depends *both* on the magnitude of the wage gain and the underlying relative price of goods to services. Variation in the relative price clearly also has second-order consequences for the degree to which consumers reallocate expenditure as their wages rise and their time becomes more valuable. Further, consumers across the wage distribution will respond to relative price variation at different rates. Depending on the underlying consumption-technology elasticities, it is even possible to observe non-monotonic responsiveness across the wage distribution of relative spending to relative price variation. This can be seen by inspecting how the level sets vary along the vertical axis in Figure B.1b. Thus, both the parameterization and market prices will determine whether $\varepsilon_i^{x,w}$ is concave or convex in w_i . This has implications for the second-order effects of structural change and whether the transition toward services will accelerate or decelerate as average wages rise.




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

Our results also indicate that if market prices and wages inflate at different rates, then struc-

tural change will either accelerate or decelerate over time depending on the consumption-technology elasticities. Figure B.2 shows this. We fix $w_i = 1$ and simulate $\varepsilon_i^{x,w}$ by proportionately re-scaling P_g and P_s for a set of relative prices. For example, let P^{agg} be some aggregate price level. Let $\theta_i \in \{0.25, 0.5, 1, 1.5, 2\}$ be a scalar, such that $w_i = \theta_i P^{agg}$, and for $\theta_i = 1$, $w_i = P^{agg}$. For a continuum of relative prices, we re-scale the product-level prices each by θ_i and simulate $\varepsilon_i^{x,w}$. Variation between the colored lines in Figure B.2 shows the dependency of $\varepsilon_i^{x,w}$ on inflation in market prices relative to wages. In the left-hand panel as market prices increase relative to wages, the allocation becomes less elastic in wages for all but small values of P_g/P_s . Under the parameterization in panel (a), in an inflationary environment driven by product-market inflation outpacing wage inflation, the rate of structural change will slow. In panel (b) we observe that the allocation elasticity can be non-monotonic in relative inflation, which can be seen by noting that as inflation rises, the allocation elasticity as a function of relative prices becomes more concave. Thus, inflation also has both first- and second-order effects on the allocation's wage responsiveness.

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

- i. *If $\nu = 0$ (Cobb-Douglas), then shares are constant in w .*
- ii. *If $\nu \neq 0$, then shares are constant in w if and only if $P_g = P_s$.*

Proof. When consumption technologies are Cobb-Douglas, $\varepsilon_i^{x,w} = \rho(\omega_s - \omega_g)/(\rho - 1)$, which is clearly constant. For condition (ii) suppose first that $P_g = P_s$, then clearly x_{ig}/x_{is} only depends on z_{ig}/z_{is} . Now suppose x_{ig}/x_{is} is constant in w_i , which means that $\varepsilon_i^{x,w} = 0$, implying

$$\frac{\left[\frac{w\omega}{P_g(1-\omega)}\right]^{\frac{\nu}{\nu-1}}}{\omega + (1-\omega)\left[\frac{w\omega}{P_g(1-\omega)}\right]^{\frac{\nu}{\nu-1}}} = \frac{\left[\frac{w\omega}{P_s(1-\omega)}\right]^{\frac{\nu}{\nu-1}}}{\omega + (1-\omega)\left[\frac{w\omega}{P_s(1-\omega)}\right]^{\frac{\nu}{\nu-1}}},$$

which implies $P_g = P_s$. **QED**

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

C Description of Hamiltonian Monte Carlo Sampling

For a thorough and readable, detailed treatment of Hamiltonian Monte Carlo (HMC) sampling, we recommend reading Neal (2011). For an equally thorough, more technical treatment, read Betancourt and Stein (2011). For a brief overview of HMC 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.⁷ 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 $\pi(\mathcal{P} \mid \mathcal{D})$, where $\mathcal{D}$ are data. 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. HMC extends this concept to a sample space with continuous dynamics.⁸ 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 $\mathcal{P}$ (itself, the parameter vector), each of the same dimension. The Hamiltonian equation associated with this system is the sum of “potential” $V(\cdot)$ and “kinetic” $K(\cdot)$ energy:

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

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 $\pi(\mathcal{P} \mid \mathcal{D})$ is defined, the potential energy function for HMC is just the negative log of the right hand side of $\pi(\mathcal{P} \mid \mathcal{D}) = \pi(\mathcal{D} \mid \mathcal{P})\pi(\mathcal{P})$:

$$V(\mathcal{P}) = -\ln [\pi(\mathcal{D} \mid \mathcal{P})\pi(\mathcal{P})] \quad (\text{C.2})$$

Note that we can sample from this using HMC only if $V(\mathcal{P})$ is continuously differentiable over the entire parameter space since HMC operates on Hamilton’s equations which require computation of the gradient vectors for both $V(\cdot)$ 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

⁷The 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 HMC to a discrete choice model (Burda and Daviet 2018).

⁸Recall, “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.

form:⁹

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

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:

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

A dense $\Lambda(\mathcal{P})$ can help account for high local non-linearities in $V(\mathcal{P})$, 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 HMC implementation.¹⁰

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 sufficient warm-up period,¹¹ 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 HMC 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 $(\mathcal{Q}', \mathcal{P}')$.¹² 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

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

HMC is most useful when sampling from $\pi(\mathcal{P} | \mathcal{D})$ 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, es-

⁹This is a more flexible form of the kinetic energy function than that in Neal (2011). See Betancourt and Stein (2011) for more details.

¹⁰Stan is free software for high-performance HMC 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 HMC models, see Gelman et al. (2013a).

¹¹This is akin to the “burn-in” period for an MCMC operating under discrete dynamics.

¹²See Neal (2011) for a detailed description of the iterative “leapfrog” method and how it operates on Hamilton’s equations.

pecially when using Metropolis-Hastings algorithms where high autocorrelation may also require low acceptance probabilities, thus further giving reason for a long sampling chain.

D Main Estimation Appendix

This appendix includes additional details from structural estimations where personal-care time is included in n_{ig} . Figures D.1, D.2, D.3, and D.4 plot prior distributions (pink) against their estimated posteriors (teal). Figure D.5 contains the posterior average idiosyncratic errors over time. Figure D.6 contains the posterior distribution of the fixed effects terms over time.

We first look at the prior/posterior plots. The ρ elasticities appear to be less strongly identified in models with $\rho \in (0, 1)$ than when $\rho < 0$, as can be seen by separately comparing Figures D.1 and D.2, followed by Figures D.3 and D.4. In all models the growth rate of relative productivities appears poorly identified. The elasticities of substitution within the consumption technologies are very well identified, which is important

Now, turning to the idiosyncratic errors and fixed effects, notice that for $\rho < 0$, fixed effects are on average zero over both samples. Meanwhile, when final experiences are substitutes, fixed effects decline over time while idiosyncratic errors have a positive average trend. Referring to trends for ζ_i^k from the main text, Figures D.5 and D.6 imply that, when $\rho \in (0, 1)$ the trend in idiosyncratic errors dominates the fixed effects in contributing to a positive average trend for relative consumption efficiencies, z_{ig}/z_{is} . Meanwhile, when $\rho = 0$ it does not appear that either idiosyncratic errors or fixed effects contribute disproportionately to the evolution of ζ_i^k . This is because ζ_i^k appears stationary and approximately mean zero (i.e., we cannot reject the hypothesis that it is mean zero) under estimations from both datasets. Finally, the trend term $\gamma_{z_g/z_s} \cdot (t - \bar{t})$ has little impact on the evolution of ζ_i^k , since, as Figures D.1, D.2, D.3, and D.4 show, it is mean zero in expectation.




Figure D.1: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the ATUS model with $\rho < 0$.




Figure D.2: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the ATUS model with $\rho \in (0, 1)$.




Figure D.3: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the PUMD model with $\rho < 0$.




Figure D.4: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the PUMD model with $\rho \in (0, 1)$.




Figure D.5: These plots contain the posterior average idiosyncratic errors, $\eta_{i(t)}^k$, $\forall k \in \{n, x\}$. The shaded areas are 90% confidence regions, while the solid lines are posterior means in each year of each sample.




Figure D.6: These plots contain the posterior distribution of the fixed effects, plotted over time.

E Models without Personal Care Time

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

<i>Goods</i>		<i>Services</i>	
$\hat{\phi}_{g0}$	−4.711*** (0.057)	$\hat{\phi}_{s0}$	−3.742*** (0.172)
$\hat{\phi}_{g1}$	−0.858*** (0.020)	$\hat{\phi}_{s1}$	−1.734*** (0.059)
$\hat{v}_g$	−0.165	$\hat{v}_s$	0.423
Observations	100	Observations	100
R ²	0.948	R ²	0.898
Residual Std. Error	0.101 (df = 98)	Residual Std. Error	0.295 (df = 98)

Note:

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

Table E.1 contains the prior estimates for the inframarginal regressions on data without personal care time. Note that only the goods regressions change here, since only the goods time-use series is altered; all else remains the same. This is because the inframarginal condition for services inherits the same data as that which we use in the main text; only estimation equations centered around structural conditions x_{ig}/x_{is} and n_{ig}/n_{is} will be affected by the exclusion of personal-care time because they will depend on prior distributions centered around estimates of $\hat{v}_g$, which vary whether or not personal-care time is included in n_{ig} .

Table E.2 contains the posterior point estimates of parameters when personal-care time is excluded from the n_{ig}/n_{is} data. Figure E.1 contains estimates for $\hat{\zeta}_{i(t)}^k$, analogous to the same plots in the main text. Note that the variance of relative productivities for the $\rho < 0$ models when personal-care time is excluded is significantly higher than when it is included. Further, $\hat{v}_g$ is estimated to be significantly less complementary when personal-care time is excluded than in the models from the main text. In the PUMD data this leads the structural residual when $\rho < 0$ to exhibit properties such that it appears to be increasing over time. Further, its profile is significantly different from that estimated by ATUS data, despite the two having similar profiles in the main text. This suggests that our model is likely a bad specification when personal-care time is excluded from n_{ig} , whereas it appears to be a decent specification of the underlying DGP when it is included.

Figure E.2 contains the idiosyncratic distributions, analogous to the same plot in Technical Appendix D when personal-care time is included. Figure E.3 contains the fixed-effect parameter estimates without personal-care time. Notice that the exclusion of personal-care time leads to the fixed-effect distributions being more similar between estimates when final consumption experiences are complementary versus substitutable. Further, the idiosyncratic error distributions

in the $\rho < 0$ case are less precisely centered around zero and possibly more non-stationary than for models estimated with personal-care time included in n_{ig} . Differences in idiosyncratic error distributions across ATUS and PUMD models is what is contributing to the stark differences in estimates of $\hat{\zeta}_{i(t)}^k$ across these models, especially relative to those estimated on data inclusive of personal-care time.

Finally, Figures E.4, E.5, E.6, and E.7 plot the posteriors against the priors for all four models where personal-care time is excluded. Notice that the identification properties that we discussed in the main text and Technical Appendix D are similar even when personal-care time is excluded from n_{ig} .

Table E.2: 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	-0.955	0.025	1.118	0.397	0.015	0.229
ν_g	-0.159	0.005	0.212	-0.174	0.008	0.233
ν_s	0.345	0.006	0.217	0.380	0.011	0.223
ω_g	0.506	0.006	0.289	0.502	0.010	0.290
ω_s	0.468	0.010	0.293	0.488	0.015	0.289
γ_{z_g/z_s}	0.036	0.031	0.945	0.033	0.057	0.927
σ	4.163	0.127	2.153	4.443	0.172	2.480
WAIC = 292.332			WAIC = 312.247			
PUMD Data						
	$\rho < 0$			$\rho \in (0, 1)$		
	Mean	S.E. Mean	S.D.	Mean	S.E. Mean	S.D.
ρ	-0.992	0.146	1.129	0.332	0.036	0.203
ν_g	-0.181	0.006	0.219	-0.162	0.019	0.163
ν_s	0.453	0.012	0.270	0.434	0.022	0.276
ω_g	0.515	0.025	0.286	0.491	0.046	0.302
ω_s	0.544	0.011	0.264	0.544	0.013	0.258
γ_{z_g/z_s}	-0.278	0.266	1.191	0.246	0.110	0.926
σ	3.162	0.118	1.649	3.322	0.137	1.588
WAIC = 594.366			WAIC = 534.331			




Figure E.1: These plots contain the distributions of posterior average 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.




Figure E.2: These plots contain the weighted error distributions over time for estimators on both ATUS and PUMD datasets. The shaded areas are 90% confidence regions, while the solid lines are observation-level means in each year of each sample.




Figure E.3: These plots contain the distributions of fixed-effect parameters, $\phi_t^k, \forall k \in \{n, x\}$. The shaded areas are 90% confidence regions, while the solid lines are posterior means in each year of each sample.




Figure E.4: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the ATUS model with $\rho < 0$.




Figure E.5: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the PUMD model with $\rho < 0$.




Figure E.6: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the ATUS model with $\rho \in (0, 1)$.




Figure E.7: This plot contains prior distributions (pink) plotted against posterior distributions (teal) for the PUMD model with $\rho \in (0, 1)$.

F Quantitative Implications Appendix

F.1 Estimates of Relative Expenditure Elasticity of Wages




Figure F.1: In this figure we plot the relative expenditure elasticity of wages from (B.11), as estimated with micro data using the posterior means of structural parameters as point estimates. All panels are scatterplots with the elasticity on the horizontal axis and the agent's wage on the vertical axis. Panels (a) and (b) contain estimates when final experiences are complements. Panels (c) and (d) contain estimates when final experiences are substitutes.

In this section we demonstrate how relative expenditure appears to respond to wage variation through the lens of the estimated models by exploring the model-implied elasticities in expression (B.11) from parameter estimates and micro data. Recall that the expression for the elasticity in (B.11) is independent of relative productivities, z_{ig}/z_{is} , so that our analyses here do not depend on

observed residuals.

In Figure F.1 we explore the model-implied elasticities for all models where personal-care time is included. In Figures F.1a and F.1b we plot each agent’s wages against his/her estimated elasticity, under the assumption $\rho < 0$. In Figures F.1c and F.1d we show the same plots when $\rho \in (0, 1)$.

Focusing first on cases where $\rho < 0$, we notice a clear pattern, such that the relative expenditure elasticity rises as wages rise, though there is less variation in elasticities between two similarly compensated high-wage earners than between two low-wage earners.¹³ When this elasticity is negative, a 1% increase in wages is associated with an $x\%$ decline in relative x_{igt}/x_{ist} . Low-wage earners thus respond to a wage increase by increasing their relative services consumption, but high-wage earners respond the opposite: they increase their relative goods consumption. This contradicts a story whereby structural change (i.e., the rising services share of spending) is driven by services being perceived as relative luxuries. We include Figures F.2a and F.2b to show that the variation we observe in Figures F.1a and F.1b across our entire sample is not biased by time fixed effects, since the elasticities for the mean, top quintile, and bottom quintile of the wage distribution are fairly flat over both datasets’ sample periods.

Continuing our discussion, turn now to cases where $\rho \in (0, 1)$. Notice that Figures F.1c and F.1d show that relative expenditure elasticities have less variance across the wage distribution and are, on average, more elastic than when $\rho < 0$. This can also be observed by inspecting these elasticities over time in Figures F.2c and F.2d. When final experiences are substitutes, even high wage earners have expenditure allocations that are responsive to wage shocks. Further, an increase in wages predicts a decline in relative spending. This is *not* true when final experiences are complements: high wage earners actually *increase* relative goods/services spending in response to a wage increase, once their wages are high enough above a certain threshold.

¹³Readers should not be deceived by these plots. The “thickness” of the lines, which are actually dense clusters of points, result from the fact that relative prices are varying over time, so that a consumer with wage w in period t faces a different market structure compared to a consumer with the same wage in period $t' \neq t$. Still, the variance within a given wage level appears very minimal, suggesting wage differences, not relative prices, are the primary contributor to variation in $\varepsilon_{it}^{x,w}$.





Figure F.2: In this figure we plot the relative expenditure elasticity of wages from (B.11), as estimated with micro data using the posterior means of structural parameters as point estimates. As opposed to Figure F.1, in this figure we show how the elasticities evolve over time at the mean, top quintile, and bottom quintile.

F.2 Micro Wage Distributions Over Time

We include this appendix to ameliorate concerns that the counterfactuals simulated on aggregate, average data fail to capture distributional effects. We address this concern indirectly. We contend that worrying about distributional effects is important only if the wage distribution is somehow changing shape over time. We show, here, that in both ATUS and PUMD, it is not. Further, cross-sectional differences in relative expenditure do not appear to be changing over time, which can be observed by inspecting Figure 2a in the main text. This suggests that distributional implications for structural change, as it results from changes to the income distribution, can be safely ignored.




(a) ATUS




(b) PUMD

Figure F.3: This figure shows the kernel densities of the wage distributions for both the first and terminal years of each sample.



(a) ATUS



(b) PUMD

Figure F.4: This figure shows the average annual growth rate of real wages by wage quantile over the life of the sample — 2003-2022 for ATUS and 1990-2022 for PUMD.

Figure F.3 features density plots of the normalized wage distributions in both the ATUS and PUMD in both the first year (2003 or 1990 in red) and last year (2022 in blue) of each sample. To normalize wages, we divide each wage by the weighted median of the sample in that given year. Notice that these densities are the same in both years in both samples, suggesting the wage distribution is fairly stable. Figure F.4 shows average annual growth rates in wages at different wage quantiles in each sample. The ATUS being of smaller sample size has more variability. Meanwhile, the PUMD features fairly constant wage growth across the distribution except at the

upper tail.

F.3 Predicted Relative Time Use

Figure F.5 shows how well the estimated productivities from aggregate relative expenditure data fit the 2003-2022 ATUS time series of average relative time use. The $\rho < 0$ models perform best. In the main text, looking at the combined AUT/TUESA/ATUS time series, we also find that $\rho < 0$ best reconciles the series.




Figure F.5: This figure shows predicted relative time use under aggregate average wages and aggregate relative productivities estimated using relative expenditure data, normalized to unity in 2003. In the main text we show analogous figures going back to 1959. Here, we show how well the model matches average relative time use from the shortened ATUS time series.

References

- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2012). “Recent Developments in the Economics of Time Use”. *Annual Review of Economics* 4, pp. 373–397 (cit. on p. 7).
- Becker, Gary (1965). “A Theory of the Allocation of Time”. *The Economic Journal* 75.299 (cit. on p. 3).
- Betancourt, Michael and Leo Stein (2011). “The Geometry of Hamiltonian Monte Carlo”. *arXiv:1112.4118* (cit. on pp. 11, 12).
- Blundell, Richard and Thomas Macurdy (1999). “Labor Supply: A Review of Alternative Approaches”. *Handbook of Labor Economics*. Ed. by O. Ashenfelter and D. Card. Vol. 3. Elsevier Science. Chap. 27 (cit. on p. 3).
- Burda, Martin (2015). “Constrained Hamiltonian Monte Carlo in BEKK GARCH with Targeting”. *Journal of Time Series Econometrics* 7.1 (cit. on p. 11).
- 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. 11).
- Gelman, Andrew et al. (2013a). “Computation in R and Stan”. *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. Appendix C (cit. on p. 12).
- (2013b). “Computationally efficient Markov chain simulation”. *Bayesian Data Analysis*. Third. Chapman & Hall. Chap. 12 (cit. on p. 11).
- 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 pp. 11, 12).