

Recent Developments in the Economics of Time Use

Mark Aguiar,¹ Erik Hurst,² and Loukas Karabarbounis²

¹Department of Economics, Princeton University, Princeton, New Jersey 08544; email: mark@markaguiar.com

²Booth School of Business, University of Chicago, Chicago, Illinois 60637; email: erik.hurst@chicagobooth.edu, loukas.karabarbounis@chicagobooth.edu

Annu. Rev. Econ. 2012. 4:373–97

First published online as a Review in Advance on April 5, 2012

The *Annual Review of Economics* is online at economics.annualreviews.org

This article's doi: 10.1146/annurev-economics-111809-125129

Copyright © 2012 by Annual Reviews. All rights reserved

JEL codes: D13, E21, E32, J22

1941-1383/12/0904-0373\$20.00

Keywords

consumption, labor supply, trends, life cycle, business cycles

Abstract

The proliferation of new data sets and their harmonization with the older data sets have allowed researchers to make significant progress in our understanding of how individuals allocate their time away from market work. We highlight how these new data can be used to test theories of time use and we review recent developments in long-run trends in time use, life-cycle patterns of expenditures and labor supply, and the allocation of time over the business cycle.

1. INTRODUCTION

How individuals allocate their time away from market work has been a cornerstone of many economic models over the past 50 years. For example, Becker (1965) has put forth a general theory of how one should think about the allocation of time across different activities. Ghez & Becker (1975) show the importance of the substitutability of time between market and nonmarket work in explaining life-cycle profiles of consumption expenditures and work hours. Benhabib et al. (1991) and Greenwood & Hercowitz (1991) develop models where the extent to which market expenditures and market work fall during recessions depends on the willingness of households to substitute between market-produced and home-produced goods. Greenwood et al. (2005) show that innovations in the market of consumer durables allowed women to increase their labor supply in a model where home production is an active margin of substitution. Diamond (1982) and Mortensen & Pissarides (1994) emphasize the time allocated to job search in explaining unemployment dynamics.

Although these models highlight the importance of the time allocated to activities other than market work in explaining a variety of different economic phenomena, the lack of high-quality time use data makes it hard to assess systematically many of the key empirical predictions of these models.¹ During the past decade, however, advances in both the collection of new data and the harmonization of older data with the new data have allowed researchers to make significant progress in our understanding of how individuals allocate their time. The goal of this review is to highlight some recent empirical literature that has used the new (or the newly harmonized) time use data sets.

The article proceeds in five parts. In Section 2, we describe the data that researchers have used to measure time use. In Section 3, we discuss recent developments in analyzing long-run trends in time use. In Section 4, we set up a life-cycle model with home production and highlight how recent empirical work has used the model predictions to explain patterns of the life-cycle behavior of households. In Section 5, we review recent theoretical and empirical work on the allocation of time over the business cycle. In Section 6, we conclude and identify areas where additional research is needed.

Before proceeding, we stress that our goal is to highlight some recent advances in the empirical literature that test existing theories of time allocation. We do not consider this review to be an exhaustive survey of all the relevant work on time use that has taken place in the past few decades. In this sense, our article is complementary to Gronau (1987), who surveys the older (mostly theoretical) literature, and to Juster & Stafford (1991), who discuss in more detail issues related to the measurement of time. Additionally, most of our applications are drawn from US data. There is much interesting work being done with time use data in other countries, and we discuss some of it below.

2. TIME USE DATA

Micro data on the time individuals spend on market work have been consistently measured since at least the 1960s. For example, in the United States, the Current Population Survey (CPS) has collected hours worked in the market for a large representative sample of the

¹This is not to say that these papers did not have empirical content. For example, Ghez & Becker (1975) use data from the Survey of Consumer Expenditures (1960–1961) and the census (1960) to test several of their hypotheses.

US population starting in 1962. The Panel Study of Income Dynamics (PSID), the Health and Retirement Survey, and the National Longitudinal Surveys are panel data sets that provide market work hours for a given individual over time.

Micro data on the extent to which individuals allocate their time to activities other than market work have been much more scarce. Both the PSID and the Health and Retirement Survey provide some nondiary recall data on certain aspects of nonmarket time. However, a number of studies have shown that time diaries provide more accurate estimates for nonmarket time relative to surveys based on recall data.² For instance, Robinson & Godbey (1999) discuss how estimates of time use in studies that focus on recall estimates of time spent in various activities almost always add up to a number that exceeds the total time endowment. We start by reviewing available US time use data sets that are based on time diaries. We then discuss some of the time use data available in other countries.

2.1. US Historical Time Use Data: 1965–1999

Between 1965 and 1999, there were five large nationally representative time use surveys that documented how individuals spend their time away from market. The surveys were the 1965–1966 Americans' Use of Time; the 1975–1976 Time Use in Economic and Social Accounts; the 1985 Americans' Use of Time; the 1992–1994 National Human Activity Pattern Survey; and the 1998–1999 Family Interaction, Social Capital, and Trends in Time Use Study. All surveys used a 24-h recall of the previous day's activities to record time-diary information. Aside from the 1975–1976 survey, the surveys collected diaries for only one individual per household. Below we briefly summarize the other salient features of these surveys.

The 1965–1966 Americans' Use of Time survey was conducted by the Survey Research Center at the University of Michigan. The survey sampled one individual per household in 2,001 households in which at least one adult person between the ages of 19 and 65 was employed in a nonfarm occupation during the previous year. Of the 2,001 individuals, 776 came from Jackson, Michigan. The remainder of the sample was designed to be nationally representative. The time use data were obtained by having respondents keep a complete diary of their activities for a single 24-h period between November 15 and December 15, 1965, or between March 7 and April 29, 1966.

The 1975–1976 Time Use in Economic and Social Accounts survey was also conducted by the Survey Research Center at the University of Michigan. The sample was designed to be nationally representative, excluding individuals living on military bases. Unlike any of the other time use studies, this study sampled multiple adult individuals in a household (as opposed to a single individual per household). The sample included 2,406 adults from 1,519 households. The survey collected up to four diaries for each respondent over the course of a year. However, the attrition rate for the subsequent rounds after the first round was high. The 1975–1981 Time Use Longitudinal Panel Study is a longitudinal data set

²An issue with time diaries (when used to conduct individual-level analysis) is that a large fraction of individuals have zero values for the time spent in many activities. For example, it is reasonable to assume that most parents spend at least some time on child care, but it has been observed that a relatively large fraction of parents report no time on this in their diary day. Some researchers have advocated Tobit methods to deal with the zeros. Stewart (2009) argues that the zeros in time use data arise from a mismatch between the reference period of the data (the diary day) and the period of interest, which is typically much longer. As a result, two-part models and ordinary-least-squares methods may be preferable to Tobit.

that combines additional data collected in 1981 with the earlier 1975–1976 Time Use in Economic and Social Accounts data set. This combined data set consists of 620 respondents (and their spouses if they were married at the time of first contact).

The 1985 Americans' Use of Time survey was conducted by the Survey Research Center at the University of Maryland. The sample of 4,939 individuals was nationally representative with respect to adults over the age of 18 living in homes with at least one telephone. The survey sampled its respondents from January 1985 through December 1985. Part of the survey design was to compare response rates for individuals who were asked to complete the survey via mail relative to individuals who were asked to complete it via telephone or face-to-face interviewing.

The 1992–1994 National Human Activity Pattern Survey was conducted by the Survey Research Center at the University of Maryland and was sponsored by the US Environmental Protection Agency. As a result, this survey also asked detailed questions about the location where different time use activities were taking place. The sample was designed to be nationally representative with respect to households with telephones. The sample included 9,386 individuals, of whom 7,514 were over the age of 18. The survey randomly selected a representative sample for each three-month quarter starting in October 1992 and continuing through September 1994. This survey contained the least detailed demographics of all the time use surveys. Specifically, the survey reports the respondent's age, sex, level of educational attainment, race, labor force status, and parental status. Unfortunately, the survey does not report the respondent's marital status, household income, or the number of children present in the household.

The 1998–1999 Family Interaction, Social Capital, and Trends in Time Use Study, conducted at the University of Maryland's Survey Research Center, covers a small-scale, contiguous-state sample of individuals over the age of 18. Between March 1998 and December 1999, 1,151 individuals were surveyed. One of the goals of the survey was to measure time spent on social interactions and, in particular, the time parents spend with their children.

Even though these time use surveys were conducted during different time periods and had different objectives, their structures are similar.³ As a result, it is easy to harmonize the data across the years. The American Heritage Time Use Study is a data set that harmonizes the above five data sets. In addition, it has started harmonizing the first waves of the American Time Use Survey (ATUS).

2.2. US Recent Time Use Data: 2003–2010

The ATUS was recently developed to provide consistent and periodic data on how US individuals are allocating their time away from market work. The ATUS is conducted by the US Bureau of Labor Statistics (BLS). Covering more than 400 detailed time categories, the ATUS is considered the state of the art in time use surveys. Participants in ATUS, which include individuals over the age of 15, are drawn from the existing sample of the CPS. The individual is sampled approximately 3 months after completion of the final CPS survey. At the time of the ATUS survey, the BLS updates the respondent's employment and demographic information. Respondents are surveyed during every month of the year.

³John Robinson, currently at the University of Maryland, was the principal investigator on all five time use surveys, which explains their similar design.

There were roughly 20,000 respondents in the 2003 wave of the ATUS, whereas the 2004–2010 waves included approximately 13,000 respondents each year (for a more detailed description of the ATUS data set and the activities it records, see Hamermesh et al. 2005).

One difficulty in harmonizing the ATUS data with the other historical time use data within the United States is that the survey structure and subsequent coding of time use activities are dramatically different. Most of the US time use surveys prior to the ATUS categorized time use into roughly 90 different subcategories, whereas the ATUS includes over 400 different subcategories. Individual researchers have taken the raw ATUS data and tried to create classifications that were consistent with the earlier surveys (for an example of harmonizing of the various data sets, see Aguiar & Hurst 2007c).

2.3. International Time Use Data

Many other developed countries have collected cross-sectional time use data during the past 40 years. Much of this data can be downloaded from the Multinational Time Use Study (MTUS). The MTUS is an ex post harmonized cross-time, cross-national comparative time use database constructed from national random-sampled time-diary studies. Currently, the MTUS encompasses over 60 data sets from roughly 20 countries (including the US studies discussed above). Countries with multiple surveys spanning multiple decades include Australia, Canada, Denmark, Finland, France, Germany, Italy, the Netherlands, Norway, and the United Kingdom. Similar to the US data, most of these surveys are based on 24-h recall of the prior day's activities. Many of these studies include all adults in a household and multiple consecutive days.

3. AGGREGATE TRENDS IN TIME USE

The recent progress in the measurement of time use in many developed economies has allowed researchers to analyze how time allocated away from market work has evolved over the past few decades. In this section, we review the trends in nonmarket work (home production), leisure, and child care within the United States and a few other developed economies.

3.1. Trends in US Home Production and Leisure

It has been well documented that in the United States since the late 1960s, time spent on market work for men has been falling, whereas for women it has been increasing steadily (see, e.g., McGrattan & Rogerson 2004). The harmonization of the historical US time use surveys allows researchers to examine what men and women have been doing with their time more broadly during this period. In this subsection, we focus on two broad time use categories: leisure time and nonmarket work time. As seen from the work surveyed below, two robust features of the data have emerged. First, since the 1960s, aggregate nonmarket work time has fallen sharply. Second, during this time period, leisure time has increased. The extent to which leisure time has increased, however, differs across authors depending on how the leisure activities are defined.

Specifically, nonmarket work time includes activities such as cooking, cleaning, laundry, home maintenance, grocery shopping, and obtaining services (e.g., going to the barber or

going to the dry cleaners). Time spent obtaining education or medical services is almost always excluded from the nonmarket work category. Instead, researchers treat such time as investments (either in human capital or health). There is some disagreement about the extent to which time spent on child care should be treated as nonmarket work time. Given this, we choose to treat child care as a separate category and discuss it separately below.

Most researchers define leisure time as the time individuals spend in activities such as watching television, socializing, playing games, talking on the phone, exercising, engaging in sports, using the computer for personal use, and going to the movies. Some researchers lump together personal maintenance time with their leisure measures. These categories include time spent eating, time spent sleeping, and time spent on personal care. The inclusion of these categories in the leisure measure usually makes little difference for the trend in leisure time, given that the trend in the sum of these personal maintenance categories has been constant since the 1960s, as reported, for example, by Aguiar & Hurst (2007c).

Juster (1985), Robinson & Godbey (1999), Aguiar & Hurst (2007c), and Ramey & Francis (2009) use the harmonized time use data sets within the United States to explore trends in leisure and nonmarket work time. Robinson & Godbey (1999) is a book-length treatise comparing the time use data from the 1965–1966, 1975–1976, 1985, and 1992–1994 surveys. Aguiar & Hurst (2007c) harmonize all the earlier time use surveys with the more recent data from the ATUS. A key finding from their paper is that total nonmarket work in the United States fell by nearly 4 h per week between 1965 and 2003 for nonretired individuals between the ages of 18 and 65. For women, the decline in nonmarket work was over 10 h per week. Men, conversely, actually increased their time spent on nonmarket work by 4 h per week during this period. In the same paper, Aguiar & Hurst document that, between 1965 and 2003, average leisure time in the United States increased by roughly 5 h per week for nonretired individuals between the ages of 18 and 65. The increase in leisure was slightly higher for men than it was for women (6.2 h versus 4.9 h per week).

Ramey & Francis (2009) independently harmonize the US time use data and document trends in leisure and home production for the population as a whole and for men and women separately. Similar to Aguiar & Hurst (2007c), Ramey & Francis (2009) also find a large decline in aggregate home production time for prime-age individuals between 1960 and the early 2000s. Ramey & Francis (2009), however, find that there was very little increase in leisure for either prime-age men or women during this time period.⁴

The work of Ramey & Francis (2009) provides an additional innovation above and beyond the work of Aguiar & Hurst (2007c). In particular, Ramey & Francis (2009) incorporate the findings of Ramey (2009) into their analysis, which allows them to compute trends in nonmarket work and leisure prior to 1965. This is an ambitious task given that there are no nationally representative time diaries within the United States prior to 1965. The goal of Ramey (2009) is to use nonrepresentative time use surveys conducted within the United States prior to 1965 to compute the amount of home production done in the United States for an average individual by weighting the nonrepresentative samples appropriately. Using this methodology, Ramey (2009) concludes that between 1900 and 1965, nonmarket work time for women fell by approximately 6 h per week, whereas for men it increased

⁴We refer readers to Ramey (2007) and Aguiar & Hurst (2007a) for a reconciliation of the differences in leisure trends between the two papers. A large part of the debate is whether eating while at market work is considered market work (Aguiar & Hurst) or leisure (Ramey & Francis).

by approximately 7 h per week. Given Ramey's (2009) estimates, Ramey & Francis (2009) state that aggregate leisure increased by an additional 2 h per week for prime-age individuals between 1900 and 1965.

3.2. Trends in US Leisure Inequality

Aguiar & Hurst (2007c, 2009) explore the evolution of other moments of the time use distribution in the United States with the aim of documenting the evolution of leisure inequality. They show that between 1965 and 1985, the average time spent on leisure activities for low-educated households grew at a rate similar to that for high-educated households. Between 1985 and 2003, however, the leisure time of low-educated households grew substantially while that for high-educated households actually contracted. Almost all the differences in leisure across education groups were driven by differences in the time allocated to market work. The growing leisure inequality mirrors the well-documented change in wages and consumption between education groups starting in the early 1980s (see, e.g., Katz & Autor 1999, Attanasio et al. 2004).

3.3. Trends in US Child-Care Time

Parental time spent with children combines elements of parental leisure, parental home production, and parental investments.⁵ For example, when asked to assess the satisfaction they receive from the various activities they perform, individuals consistently rank the time they spend playing and reading to their children as being the most enjoyable (Robinson & Godbey 1999). Such evidence suggests that elements of child-care time are akin to leisure. Alternatively, some child-care services can be purchased directly from the market. As a result, some elements of child-care time are akin to home production.

Kimmel & Connelly (2007) treat child-caregiving time distinctly from time spent on home production, leisure, and paid market work. Their focus is to identify differential responses of these time use categories to demographics and prices. For instance, the authors find that higher-wage mothers devote more time to caregiving; paid work time on weekdays responds positively to higher wages, whereas leisure time and home production weekday time respond negatively.

Guryan et al. (2008) document that the income elasticity of child-care time is strongly positive, with richer and high-educated parents spending much more time with their children than poorer and low-educated parents. In all years of the time use surveys, high-income parents spend less time on home production relative to low-educated parents. In recent time use surveys, high-income parents also allocate less of their time to leisure than low-educated parents. The income elasticity of child-care time stands in sharp contrast (and is of opposite sign) to the income elasticities of both home production time and leisure time.

Sayer et al. (2004) and Aguiar & Hurst (2007c) document that time spent with children has been increasing in the late 1990s and the 2000s relative to the 1960s, 1970s, and 1980s.⁶

⁵For the purposes of this short review, we do not distinguish among time devoted to a child-care activity, time in which a child was present, time in which an adult reported responsibility for the care of a child under 12, and general supervisory time (see Folbre et al. 2005 for a careful definition of the time spent on child care).

⁶These results could be influenced by classification issues. Because of the structure of the ATUS, child care may be over-reported compared with earlier time use surveys.

Ramey & Ramey (2010) document that time spent with children has increased more for high-educated parents relative to low-educated parents. The increasing gap in time spent with children by education has occurred in all categories of child-care time: time spent on basic child care, educational child care, and recreational child care. Collectively, these results show that there are interesting patterns in the trends in child care and the response of child care to changing demographics and prices that should be treated as distinct from the responses in either home production or leisure.

3.4. Time Use Patterns in Other Countries

There is much work that also uses international time use data to measure the trends in time use within other countries or to compare the allocation of time across countries [for some additional recent work using international data, see the 2005 symposium organized by the *European Economic Review* (Hamermesh & Pfann 2005)]. Early work by Juster & Stafford (1991) provides some cross-country analysis of time use. For instance, this work documents that total work hours for men declined substantially in Japan, Norway, and the United States between the 1960s and the 1980s. This decline was a consequence of a small increase in hours spent on nonmarket work, which was more than fully offset by a large decline in hours spent on market work.

Lee et al. (2011) use time diaries from Japan and Korea to analyze the effects of legislated labor demand shocks on time use. The identification of the effects on time use comes from the fact that the exogenous reduction in legislated work hours is more likely to affect workers who are closer to the constraint. Using time diaries from before and after these shocks, they find that the increased time was mostly reallocated to increased leisure and increased personal maintenance, with very little time being absorbed by home production.

Gimenez-Nadal & Sevilla-Sanz (2011) use the MTUS and other country-specific data sets for seven developed countries (Australia, Canada, Finland, France, the Netherlands, Norway, and the United Kingdom) to examine trends in the allocation of time since the mid-1970s. They conclude that leisure time for men increased by roughly 4 h per week over the past few decades in Australia, Finland, and the United Kingdom. Conversely, at the same time, leisure time for men actually declined in France, the Netherlands, and Norway and remained constant in Canada. They find that in most countries, changes in leisure were less pronounced for women relative to men. Additionally, Gimenez-Nadal & Sevilla-Sanz (2011) document that in most countries, decreases in men's market work were offset by increases in time spent on nonmarket work and child care. For women, time devoted to market work increased in almost all countries, and time spent on home production decreased.

The proliferation of data sets on time use in various countries has also allowed researchers to look at cross-country differences in the allocation of time. Freeman & Schettkat (2005) examine time use data from a number of countries and conclude that there is a very high substitution of time between market and home work across individuals. For instance, they report that in the 1990s, Europeans worked 20% more than Americans in the home sector. Similar patterns are documented by Burda et al. (2008), who compare the allocation of time within the United States to the allocation of time in Germany, Italy, and the Netherlands.

Burda et al. (2007) use time-diary data from 25 countries and demonstrate a strong gender "iso-work" relationship in various advanced economies. Specifically, although men

work more in the market sector than women, and women work more in the home sector than men, these differences tend to balance out in a way such that men and women spend approximately the same time on total (market and nonmarket) work (see also Robinson & Godbey 1999). Burda et al. (2007) argue that the iso-work fact may be explained by a model that incorporates a social norm for leisure. Peer pressure or conformity to common social norms can diminish the incentive to allocate time in response to changing wages and individual tastes. As a result, total work (market plus nonmarket) time and leisure time become more similar across individuals.

In a cross-country study, Alesina & Ichino (2009) use the MTUS data set for Italy (2002), the United States (2003), Spain (2002), and Norway (2000) to estimate how home production affects cross-country comparisons in GDP. Specifically, the authors use two methods to estimate the value of home production. The first one assumes that home production time is paid according to the ongoing market wage, whereas the second one evaluates each hour of home production at the cost at which home services can be bought in the market (measured as the unskilled wage). A robust finding of their analysis is that Italy's position in terms of GDP with respect to comparable countries would improve considerably if official statistics included the imputed value of home production.

4. LIFE-CYCLE CONSUMPTION AND LABOR SUPPLY

The economics literature typically analyzes life-cycle patterns of consumption and work by appealing to models that emphasize only the intertemporal substitution of goods and time. However, various patterns of consumption expenditure and labor supply over the life cycle cannot be explained by intertemporal substitution only. In this section, we start by describing a simple model due to Ghez & Becker (1975) that aims to explain life-cycle expenditures and labor supply by incorporating rich intratemporal substitution patterns between time and goods. We then discuss how the theory can be used to rationalize a number of stylized facts of the life-cycle behavior of households, and we review recent developments in the empirics of consumption and labor supply.

4.1. A Life-Cycle Model with Intratemporal Substitution

A consumer lives for $t = 1, \dots, T$ periods. The consumer self-insures by borrowing and lending at an exogenous and constant interest rate r . Assets a_{t+1} must exceed some lower bound ϕ . The consumer takes all prices in the economy as given.

The consumer derives utility from N commodities C_i , $i = 1, \dots, N$. We denote the period utility function by $U(C_1, \dots, C_N)$. We denote by β the discount factor. Preferences are additively separable across periods. In every period t , the consumer maximizes the expected discounted sum of utilities:

$$\mathbf{E}_t \sum_{s=t}^T \beta^{s-t} U(C_{1s}, \dots, C_{Ns}). \quad (1)$$

Following Becker (1965), we represent the commodities $C_{it} = f^i(H_{it}, X_{it})$ that enter utility as the outputs of production functions that take time H_{it} and market expenditures X_{it}

as inputs.⁷ For instance, a commodity may be watching a television show, which combines services from a durable (the television), a cable subscription, and time. Similarly, another commodity may be a meal, which is produced with groceries and time spent on cooking as inputs. Note that in the former example, time and market goods are complements, whereas in the latter, time and market goods may be substitutes, given the option to purchase food prepared by others. As we see below, the degree of substitutability between time and market inputs in the production function is a key feature that distinguishes various commodities. We denote the price of market good X_{it} by p_{it} .

Let L_t denote the time the consumer spends working in the market. The consumer earns an exogenous (after-tax) wage w_t per unit of time worked in the market. The consumer also earns an exogenous income T_t in each period (e.g., a transfer from the government). We normalize the total time endowment to one in each period.

The sources of uncertainty in the model are the spot wage w_t , the vector of market goods prices $p_t = [p_{1t}, \dots, p_{Nt}]$, and the income transfer T_t . We denote the exogenous state vector by $s_t = (w_t, p_t, T_t)$. The exogenous state vector follows a Markov process whose transition probabilities may vary over the life cycle. Specifically, we denote by $\pi_t(s'|s)$ the probability that in period $t + 1$ the state is $s_{t+1} = s'$ conditional on the state being $s_t = s$ in period t .

This model nests as special cases the standard neoclassical model of consumption (see, e.g., Attanasio & Weber 2010) and the standard neoclassical model of labor supply (see, e.g., Blundell & MaCurdy 1999). Specifically, the standard model is obtained when there are only two commodities, and the first commodity is produced only with market expenditures (consumption) and the second commodity is produced only with time (leisure).

We can represent the consumer's optimization problem in recursive form. In any period $t < T$, the consumer solves

$$V_t(a_t, s_t) = \max_{\{H_{it}, X_{it}, L_t, a_{t+1}\}} U(C_{1t}, \dots, C_{Nt}) + \beta \int V_{t+1}(a_{t+1}, s_{t+1}) \pi_t(s_{t+1} | s_t) ds_{t+1}, \quad (2)$$

subject to the constraints

$$C_{it} = f^i(H_{it}, X_{it}), \quad (3)$$

$$L_t + \sum_{i=1}^N H_{it} \leq 1, \quad (4)$$

$$\sum_{i=1}^N p_{it} X_{it} + a_{t+1} \leq w_t L_t + (1 + r)a_t + T_t, \quad (5)$$

$$a_{t+1} \geq \phi. \quad (6)$$

To solve this problem, we substitute the constraint in Equation 3 into the objective function in Equation 2. Also, we denote by θ_t the multiplier on the constraint in Equation 4, by λ_t the multiplier on the constraint in Equation 5, and by μ_t the multiplier on the constraint in Equation 6. To simplify the exposition, we are assuming that the consumer is in an

⁷For simplicity, we have ruled out joint production; that is, a time or market good used to produce commodity i cannot be simultaneously used to produce commodity j with $i \neq j$.

interior equilibrium with $L_t > 0$, $H_{it} > 0$, and $X_{it} > 0$ for all commodities $i = 1, \dots, N$.⁸ The first-order conditions associated with this problem are

$$\{H_{it}\}: \frac{\partial U}{\partial C_{it}} \frac{\partial C_{it}}{\partial H_{it}} = \theta_t \Rightarrow \frac{\partial U / \partial C_{it}}{\partial U / \partial C_{jt}} = \frac{\partial C_{jt} / \partial H_{jt}}{\partial C_{it} / \partial H_{it}}, \forall i, j = 1, \dots, N, \quad (7)$$

$$\{X_{it}\}: \frac{\partial U}{\partial C_{it}} \frac{\partial C_{it}}{\partial X_{it}} = \lambda_t p_{it} \Rightarrow \frac{\partial U / \partial C_{it}}{\partial U / \partial C_{jt}} = \frac{p_{jt} \partial C_{jt} / \partial X_{jt}}{p_{it} \partial C_{it} / \partial X_{it}}, \forall i, j = 1, \dots, N, \quad (8)$$

$$L_t: \theta_t = \lambda_t w_t \Rightarrow \frac{\partial C_{it} / \partial H_{it}}{\partial C_{it} / \partial X_{it}} = \frac{w_t}{p_{it}}, \forall i = 1, \dots, N, \quad (9)$$

$$a_{t+1}: \lambda_t - \mu_t = \beta(1 + r_t) \int \lambda_{t+1} \pi_t(s_{t+1} | s_t) ds_{t+1}. \quad (10)$$

The first-order condition in Equation 7 states that the consumer allocates time across different commodities in a way that equalizes the marginal rate of substitution between different commodities (the ratio of the marginal utilities) to the ratio of the marginal products of time in the production of these commodities. Similarly, the first-order condition in Equation 8 states that the consumer allocates expenditures across different commodities in a way that equalizes the marginal rate of substitution between different commodities to the ratio of the marginal products of expenditure in the production of these commodities adjusted by the relative price of inputs. The first-order condition in Equation 9 states that the consumer supplies labor in the market up to the point at which the real wage is equalized to the marginal rate of technical substitution between time and expenditure (which, for each good, equals the ratio of the marginal products). Finally, the first-order condition in Equation 10 characterizes the intertemporal allocation of resources. We note that all previous conditions hold regardless of the specific assumption one makes on the structure of asset markets (complete markets, incomplete markets, borrowing constraints, etc.).

For a given marginal value of resources λ_t , wage w_t , and vector of prices $p_t = [p_{1t}, \dots, p_{Nt}]$, Equations 7 and 8 define a system of $2N$ equations in $2N$ unknowns.⁹ Additionally, the time constraint is one more equation that can be used to solve for labor supply L_t , and the production functions provide solutions for the commodities C_{it} . Denote the solution of this system by

$$H_{it}(\lambda_t, w_t, p_t); L_t(\lambda_t, w_t, p_t); X_{it}(\lambda_t, w_t, p_t); C_{it}(\lambda_t, w_t, p_t), \forall i = 1, \dots, N. \quad (11)$$

We examine how time H_{it} and L_t , expenditures X_{it} , and the production of commodities C_{it} respond when the wage w_t or the marginal value of resources λ_t varies over the life cycle. To simplify the exposition and to get a sharper intuition for the results, we assume that the utility function $U(C_1, \dots, C_N)$ is separable across commodities and that the production technology $C_{it} = f^i(H_{it}, X_{it})$ is characterized by constant returns to scale.

⁸Because the consumer optimizes at an interior point, we use the terms “wage” and “opportunity cost of time” interchangeably. If the consumer chooses not to work, however, the opportunity cost of time, now given by the marginal rate of technical substitution between time and expenditures, exceeds the real wage. All our results below can be easily generalized to the case in which the consumer does not work, with the difference that one has to use the marginal rate of technical substitution as the appropriate measure of the opportunity cost of time (instead of the wage).

⁹This is after substituting $\theta_t = \lambda_t w_t$ from Equation 9 to Equation 7.

First, we consider variations of the opportunity cost of time w_t holding constant λ_t and p_t . We use the following notation. We define the λ_t -constant elasticity of time H_{it} and L_t , expenditures X_{it} , and the production of commodities C_{it} with respect to the opportunity cost of time as

$$\varepsilon_{Hw}^i = \frac{\partial H_{it} w_t}{\partial w_t H_{it}}, \varepsilon_{Lw} = \frac{\partial L_t w_t}{\partial w_t L_t}, \varepsilon_{Xw}^i = \frac{\partial X_{it} w_t}{\partial w_t X_{it}}, \text{ and } \varepsilon_{Cw}^i = \frac{\partial C_{it} w_t}{\partial w_t C_{it}}. \quad (12)$$

We define the elasticity of the output of commodity C_{it} with respect to some input as

$$\varepsilon_{CH}^i = \frac{\partial C_{it} H_{it}}{\partial H_{it} C_{it}} \text{ and } \varepsilon_{CX}^i = \frac{\partial C_{it} X_{it}}{\partial X_{it} C_{it}}. \quad (13)$$

By constant returns to scale, we have $\varepsilon_{CX}^i + \varepsilon_{CH}^i = 1$ for all commodities $i = 1, \dots, N$. Additionally, the elasticities of output with respect to inputs coincide with the implicit income shares of these inputs in the production of output.

We define the intratemporal elasticity of substitution between time H_{it} and expenditures X_{it} in the production of good C_{it} as

$$\sigma^i = d \ln \left(\frac{X_{it}}{H_{it}} \right) / d \ln \left(\frac{\partial C_{it} / \partial H_{it}}{\partial C_{it} / \partial X_{it}} \right) = \varepsilon_{Xw}^i - \varepsilon_{Hw}^i. \quad (14)$$

The last equality of Equation 14 follows from the first-order condition in Equation 9. Finally, we define the elasticity of intertemporal substitution for commodity i as

$$\gamma^i = - \frac{\partial U / \partial C_{it}}{C_{it} (\partial^2 U / \partial C_{it}^2)}. \quad (15)$$

Note that the intertemporal elasticity of substitution γ^i , the intratemporal elasticity of substitution σ^i , and the shares ε_{CH}^i and ε_{CX}^i are allowed to vary by commodity i . To ease the notation, we assume that these parameters are stable over time, but it is straightforward to generalize our analysis to the case under which these parameters vary over time.¹⁰

Differentiating the first-order condition in Equation 8 with respect to w_t (holding constant λ_t) and manipulating the resulting expression, we obtain¹¹

$$\frac{\partial C_{it}}{\partial X_{it}} (\varepsilon_{CH}^i \varepsilon_{Hw}^i + \varepsilon_{CX}^i \varepsilon_{Xw}^i) - \gamma^i H_{it} \frac{\partial^2 C_{it}}{\partial X_{it} \partial H_{it}} (\varepsilon_{Hw}^i - \varepsilon_{Xw}^i) = 0. \quad (16)$$

As $\sigma^i = \varepsilon_{Xw}^i - \varepsilon_{Hw}^i$ and $\varepsilon_{CX}^i - \varepsilon_{CH}^i = 1$, Equation 16 can be solved for the λ -constant elasticity of expenditures with respect to the wage:¹²

$$\varepsilon_{Xw}^i = \varepsilon_{CH}^i (\sigma^i - \gamma^i). \quad (17)$$

Equation 17 states that expenditures for commodity i increase with the wage if the intratemporal elasticity of substitution between time and goods is greater than the intertemporal elasticity of substitution in consumption. The intuition is that as the wage increases, the consumer substitutes away from time and toward market inputs to achieve

¹⁰Note that we also assume that the production of commodity C_{it} is characterized by a stable production function $f^i(H_{it}, X_{it})$. It is conceivable that the efficiency with which time and expenditures are combined into the production of final goods also changes over time.

¹¹For a linear homogeneous function $f(x_1, x_2)$, we have $f_{11} = -x_2 f_{12} / x_1$.

¹²For a linear homogeneous function $f(x_1, x_2)$, we have: $f_{12} = (f_1 f_2) / (\sigma f)$.

a given level of consumption. This is a movement along the production isoquant and is parameterized by σ^i . However, because time is costlier in the current period relative to other periods, the consumer shifts consumption to a period in which the cost of producing consumption (time plus market goods) is lower. This is a parallel movement toward a lower production isoquant and is parameterized by γ^i . The magnitude of the change in expenditures also depends on the share of time in production. In the limiting case in which consumption uses only expenditures as inputs, movements in the opportunity cost of time have no effect on expenditures.

The λ -constant elasticity of time in each commodity with respect to the wage equals

$$\varepsilon_{Hw}^i = \varepsilon_{Xw}^i - \sigma^i = -(\varepsilon_{CX}^i \sigma^i + \varepsilon_{CH}^i \gamma^i). \quad (18)$$

The elasticity is always negative. As the opportunity cost of time increases, the consumer substitutes both intratemporally away from current time toward current expenditures (parameterized by σ^i) and intertemporally away from current time toward time in other periods (parameterized by γ^i).

Using the time constraint in Equation 4, we can obtain the λ -constant elasticity of labor supply with respect to the wage:

$$\varepsilon_{Lw} = -\sum_{i=1}^N \left(\varepsilon_{Hw}^i \frac{H_{it}}{L_t} \right) = \sum_{i=1}^N \left((\varepsilon_{CX}^i \sigma^i + \varepsilon_{CH}^i \gamma^i) \frac{H_{it}}{L_t} \right). \quad (19)$$

Because the time constraint holds along the optimal path, the elasticity of labor supply with respect to the wage is a weighted average of the elasticities of time allocated in the production of the various commodities. As the latter are all negative, the λ -constant elasticity of labor supply is always positive.

Finally, totally differentiating the production functions, we obtain the λ -constant elasticity of consumption of commodity i with respect to the wage:

$$\varepsilon_{CW}^i = \varepsilon_{CH}^i \varepsilon_{Hw}^i + \varepsilon_{CX}^i \varepsilon_{Xw}^i = -\gamma^i \varepsilon_{CH}^i. \quad (20)$$

The elasticity of commodity i with respect to the wage is a weighted average of the elasticity of expenditures and the elasticity of time, both with respect to the wage. Recall that both elasticities are driven by intertemporal and intratemporal substitution. In each case, intertemporal substitution causes a decrease of the inputs when the opportunity cost time increases. Intratemporal substitution, conversely, tends to increase expenditures and to decrease time when the opportunity cost of time increases. Equation 20 shows that these effects cancel out in such a way that only the effects of intertemporal substitution remain. Holding constant the marginal value of resources λ_t , when the opportunity cost of time increases, consumption of commodity i falls as long as the commodity uses time as an input in its production.

Next we consider variations of the marginal value of resources λ_t , holding constant the wage w_t and the price vector p_t . We use the following notation. We define the w_t -constant elasticity of time H_{it} and L_t , expenditures X_{it} , and the production of commodities C_{it} with respect to the marginal value of resources as

$$\varepsilon_{H\lambda}^i = \frac{\partial H_{it}}{\partial \lambda_t} \frac{\lambda_t}{H_{it}}, \quad \varepsilon_{L\lambda} = \frac{\partial L_t}{\partial \lambda_t} \frac{\lambda_t}{L_t}, \quad \varepsilon_{X\lambda}^i = \frac{\partial X_{it}}{\partial \lambda_t} \frac{\lambda_t}{X_{it}}, \quad \text{and} \quad \varepsilon_{C\lambda}^i = \frac{\partial C_{it}}{\partial \lambda_t} \frac{\lambda_t}{C_{it}}. \quad (21)$$

Differentiating the first-order condition in Equation 8 with respect to λ_t and manipulating the resulting expression, we obtain

$$\left(\varepsilon_{CH}^i \varepsilon_{H\lambda}^i + \varepsilon_{CX}^i \varepsilon_{X\lambda}^i\right) - \gamma^i \frac{\partial^2 C_{it}}{\partial X_{it}^2} \frac{X_{it}}{\partial C / \partial X_{it}} \left(\varepsilon_{X\lambda}^i - \varepsilon_{H\lambda}^i\right) = -\gamma^i. \quad (22)$$

Differentiating the first-order condition in Equation 9 with respect to λ_t , we obtain $\varepsilon_{X\lambda}^i - \varepsilon_{H\lambda}^i = 0$. Substituting $\varepsilon_{X\lambda}^i = \varepsilon_{H\lambda}^i$ into Equation 22, we can solve for the w_t -constant elasticities of expenditures, time, and labor supply and consumption of commodities with respect to λ_t :

$$\varepsilon_{C\lambda}^i = \varepsilon_{X\lambda}^i = \varepsilon_{H\lambda}^i = -\gamma^i, \quad (23)$$

$$\varepsilon_{L\lambda} = -\sum_{i=1}^N \left(\varepsilon_{H\lambda}^i \frac{H_{it}}{L_t} \right) = \sum_{i=1}^N \left(\gamma^i \frac{H_{it}}{L_t} \right). \quad (24)$$

Equation 23 shows that when lifetime resources increase (a decrease of λ) for reasons other than changes in wages and prices, consumption of commodity i increases. The consumer increases the production of commodity i by increasing both inputs by the same percent. Intuitively, this follows because the opportunity cost of time w_t/p_{it} is held constant; therefore, the marginal rate of technical substitution in the first-order condition in Equation 9 must also remain constant. With constant returns to scale, this can happen when both inputs increase proportionally as we move to a higher isoquant. Finally, Equation 24 shows that when lifetime resources increase, labor supply decreases, as the consumer increases the time allocated in the production of all other commodities. We now use these theoretical insights to examine a number of issues related to life-cycle expenditures and labor supply.

4.2. Retirement Spending

The workhorse model of consumption over the life cycle, the permanent income hypothesis, posits that individuals allocate their resources in order to smooth their marginal utility of consumption across time (see, e.g., Attanasio 1999 for a review). If the marginal utility of consumption depends only on measured consumption, this implies that individuals will save early in their life cycle to maintain a smooth level of consumption at retirement. Hamermesh (1984) was the first to observe that retirees' savings are insufficient to sustain consumption throughout the rest of their lives. If households enter into retirement with low accumulated wealth, their consumption must decline sharply at retirement.

The retirement consumption puzzle refers to the fact that household expenditure falls discontinuously upon retirement. Banks et al. (1998) look at the consumption smoothing of British households around the time of retirement. Controlling for factors that may influence the marginal utility of consumption (such as family composition, age, mortality risk, and labor force participation), they find that consumption falls significantly at retirement. Bernheim et al. (2001) find that total food expenditure declines by 6%–10% between the preretirement and the postretirement period, which leads them to conclude that households do not use savings to smooth consumption with respect to predictable income shocks. Haider & Stephens (2007) use subjective retirement expectations as an instrument to distinguish between expected and unexpected retirements and find a decline in food expenditures ranging from 7% to 11% at retirement.

Aguiar & Hurst (2005) argue that tests of the life-cycle model typically equate consumption with expenditure. However, as stressed by the model above, consumption C_{it} is the output of a home production process that uses as inputs both market expenditures X_{it} and time H_{it} . Equations 17 and 18 show that when the relative price of time falls, individuals will substitute away from expenditures toward time spent on home production. Because retirees have a lower opportunity cost of time than their preretired counterparts, time spent on the production of commodities should increase during retirement. If this is the case, then the drop in expenditure does not necessarily imply a large decrease of actual consumption at retirement.

To test this hypothesis, Aguiar & Hurst (2005) explore how actual food consumption changes during retirement. Using data from the Continuing Survey of Food Intake of Individuals (a data set conducted by the US Department of Agriculture that tracks the dollar value, quantity, and quality of food consumed within US households), they find no actual deterioration of a household's diet as it transitions into retirement. To test the hypothesis that retirees maintain their food consumption relatively constant despite the declining food expenditures, Aguiar & Hurst (2005) use detailed time diaries from the National Human Activity Pattern Survey and from the ATUS and show that retirees dramatically increase their time spent on food production relative to otherwise similar nonretired households. That retirees allocate more time to nonmarket production has been also shown by Hurd & Rohwedder (2006) and Schwerdt (2005).

In light of this evidence, Hurst (2008) concludes that the retirement puzzle "has retired." That is, even though it is a robust fact that certain types of expenditures fall sharply as households enter into retirement, standard life-cycle models with home production are able to explain this sharp fall because retirees spend more time producing goods.¹³ Additionally, as we discuss in the next section, decreases in expenditures are mostly limited to two types of consumption categories: work-related items (such as clothing and transportation expenditures) and food (both at and away from home). When expenditures exclude food and work-related expenses, the measured decreases in spending at retirement are either close to zero or even increasing.

4.3. Life-Cycle Spending

The literature on spending over the life cycle is also large. The typical finding has been that consumption follows a hump-shaped pattern over the life cycle, with consumption being low early in the life cycle, peaking at middle age, and falling sharply at retirement. Some authors have argued that the life-cycle profile represents evidence against the forward-looking consumption-smoothing behavior implied by permanent income models, particularly because the hump in expenditures tracks the hump in labor income (as documented in Carroll & Summers 1991). In this interpretation, decreasing expenditure in the latter half of the life cycle is evidence of poor planning. Other authors argue that the hump-shaped profile of consumption reflects optimal behavior if households face liquidity constraints combined with a need to self-insure against idiosyncratic income risks (see, e.g., Zeldes 1989, Deaton 1991, Carroll 1997, Gourinchas & Parker 2002). Households build up a buffer stock of assets early in the life cycle, generating the increasing expenditure

¹³Hurst (2008) also discusses how health shocks that lead to early retirement can help reconcile the fact that actual consumption falls for a small fraction of households upon retirement.

profile found during the first half of the life cycle. The decline in the latter half of the life cycle is then attributed to impatience once households accumulate a sufficient stock of precautionary savings.

However, Aguiar & Hurst (2008) demonstrate that there is tremendous heterogeneity in the life-cycle patterns of expenditures across different spending categories. In particular, some categories (e.g., food and transportation) display the familiar hump-shaped profile over the life cycle, but other categories display an increasing (e.g., entertainment) or decreasing (e.g., clothing and personal care) profile over the life cycle. This heterogeneity cannot be captured by the standard life-cycle model of consumption that emphasizes only the intertemporal substitution of goods and time.

The heterogeneity in the life-cycle profiles of various expenditures is compatible with Ghez & Becker's (1975) model of life-cycle expenditures. We can use Equation 17 to illustrate this point. Food consumption has a relatively low intertemporal elasticity of substitution and a relatively high intratemporal elasticity of substitution. It should therefore be more likely to covary positively with the opportunity cost of time. Conversely, time and market goods are difficult to substitute in the production of entertainment, and entertainment has a relatively high intertemporal elasticity of substitution. As a result, we should expect expenditures on entertainment to rise as the opportunity cost of time falls. If the opportunity cost of time decreases as workers enter into retirement, then we expect their food consumption to decrease and entertainment expenditures to increase.

We summarize our discussion by stressing that the standard model of consumption focuses only on expenditures responding to movements in lifetime resources, abstracting from movements in the opportunity cost of time. That is, the standard model emphasizes only the elasticity $\epsilon_{X\lambda}^i$ in Equation 23 and assumes that $\epsilon_{Xw}^i = 0$ in Equation 17. This would not be an issue if time were a small share of consumption inputs. However, because time is an empirically prominent input in the production of various commodities, an analysis that abstracts from movements in the opportunity cost of time confounds price and income effects. This is analogous to models of labor supply, for which it has long been recognized that wage changes yield both substitution and income effects. To illustrate this point, consider Equations 20 and 23, which show the response of consumption C_{it} with respect to w_t and λ_t , respectively. If shocks to λ_t (income effects) dominate when consumers are close to retirement, then we expect a decrease in expenditures, time allocated in home production, and consumption. If, however, shocks to w_t (substitution effects) dominate, then the increased time in home production causes an increase in the production of final commodities.

4.4. Labor Supply Elasticities

Most attempts to obtain estimates of the labor supply elasticity are based on models in which any time not spent on the market is allocated to leisure (see, e.g., Blundell & MaCurdy 1999 for a review of the standard neoclassical labor supply model). To map the standard model of labor supply into our framework, consider the simple case with only two commodities. The first commodity, consumption, is produced only with market expenditures, $c(H^1, X^1) = X^1$. The second commodity, leisure, is produced only with time $l(H^2, X^2) = H^2$. Under this case, the λ_t -constant elasticities of consumption, leisure, and labor with respect to the wage are $\epsilon_{cw} = 0$, $\epsilon_{lw} = -\gamma$, and $\epsilon_{Lw} = \gamma(1 - L_t)/L_t$, respectively, where $\gamma = -U_l/U_{ll}$. There is a rich literature that tries to estimate the elasticity of intertemporal substitution in labor supply (related to the parameter γ).

To understand the implications of introducing home production into the standard neoclassical labor supply model for the measurement of the elasticity of labor supply, now consider a simplified version of our model with three (instead of two) goods and $\gamma^i = \gamma$. Along with consumption and leisure, there is also a third good, the home commodity, which is produced according to a function $h(H^3, X^3)$ that uses both time and expenditures as inputs. Under this modification of the standard model, the λ_t -constant elasticity of labor supply with respect to the wage becomes

$$\varepsilon_{Lw} = \gamma \frac{l_t}{L_t} + (\varepsilon_{bX}^3 \sigma^3 + \varepsilon_{bH}^3 \gamma) \frac{h_t}{L_t}. \quad (25)$$

If both the two-good and three-good models have the same equilibrium level of labor supply L_t , then in the model with home production, labor supply is more elastic with respect to wage variations that leave the marginal value of resources constant when $\sigma^3 > \gamma$. To see this, note that

$$\varepsilon_{Lw} = \gamma \frac{l_t}{L_t} + (\varepsilon_{bX}^3 \sigma^3 + \varepsilon_{bH}^3 \gamma) \frac{h_t}{L_t} > \gamma \frac{l_t}{L_t} + (\varepsilon_{bX}^3 \gamma + \varepsilon_{bH}^3 \gamma) \frac{h_t}{L_t} = \gamma \frac{1 - L_t}{L_t}. \quad (26)$$

The main lesson is that augmenting the neoclassical model of labor supply with home production changes the measured responsiveness of labor supply to wage variations. Alesina et al. (2011) adopt preferences such that the elasticity of labor supply is increasing in the amount of time spent on home production.¹⁴ They argue that intrahousehold bargaining that favors men can explain why men take less home duties than women. This in turn explains the well-documented gender gap in labor supply elasticities. In addition, because men now engage in more home production while women engage in less relative to the 1960s, models in which the elasticity of labor supply is increasing in the time spent on home production imply that men's and women's elasticities of labor supplies have converged relative to the 1960s. This indeed seems to be the case (see, e.g., Blau & Kahn 2007).

Empirically, Rupert et al. (2000) construct a synthetic cohort from the cross section of three time use surveys to estimate a structural model of life-cycle consumption and labor supply with home production. They show that estimates of the intertemporal labor supply elasticity based on models that explicitly account for home production are significantly larger than the estimates found in empirical studies that ignore home production. As we discuss in Section 5, the result that labor supply responds more in models that explicitly account for the home sector has important theoretical implications for the comovement of macroeconomic aggregates over the business cycle.

4.5. Micro Estimates of the Elasticity of Substitution

Our discussion above indicates that the degree to which households are willing to substitute intratemporally between expenditures X_{it} and time H_{it} is important in understanding various patterns of consumption and labor supply. Estimates of this elasticity σ can be obtained by estimating variations of Equation 9. Rupert et al. (1995) use data from the PSID. Most of their estimates suggest that the elasticity exceeds 1. Aguiar & Hurst (2007b)

¹⁴A similar result also holds in this more general model. To be more precise, under the assumption $\sigma^3 > \gamma$ and holding constant the equilibrium level of labor supply L_t , Equation 25 shows that the λ_t -constant elasticity of labor supply with respect to the wage increases in the fraction of time devoted to home production h_t as opposed to leisure l_t .

use data from the ATUS. Assuming that the relevant opportunity cost of time is the marginal rate of technical substitution between time and goods in the shopping technology, they find a value of approximately 1.8. Using PSID data, Gelber & Mitchell (2009) find that, in response to tax shocks, the elasticity of substitution between market and home goods is approximately 1.2 for single men and as high as 2.6 for single women. Finally, using consumer-level data on hours, wages, and consumption expenditure from the PSID and metro-level data on price indices p_t from the BLS, Gonzalez Chapela (2011) estimates a life-cycle model with home production and finds a value of σ in the production of food of approximately 2.

5. BUSINESS CYCLES

One of the most important contributions of the economics of time is in improving our understanding of aggregate fluctuations. The first wave of dynamic general equilibrium models, pioneered by Kydland & Prescott (1982), assumed that total time is allocated into only two activities: market work and leisure. There are good reasons why introducing a third activity, time spent on home production, can make a difference for these models. First, when individuals derive utility from both market-produced and home-produced goods, volatility in goods and labor markets can arise because of relative productivity differences between the two sectors, and not just because of productivity shocks in the market sector. Second, relative price changes cause households to substitute goods and time not only intertemporally between periods, but also intratemporally between the market and the home sector. Intratemporal substitution introduces a powerful amplification channel, which is absent from the standard real-business-cycle model. In fact, in his review of the home production literature, Gronau (1997) writes that “the greatest contribution of the theory of home production in the past decade was in its service to the better understanding of consumption behavior and changes in labor supply over the business cycle.”

The first papers to introduce home production into the stochastic neoclassical growth model were Benhabib et al. (1991) and Greenwood & Hercowitz (1991). We start by describing briefly Benhabib et al.’s (1991) model and some of its main results and extensions (for a more detailed presentation and review of the early literature, see Greenwood et al. 1995). Then we discuss empirical tests of these models. We stress that, as opposed to the empirics on time use trends and time use over the life cycle, the empirical literature on the allocation of time over the business cycle is still in its very early stages of development. This reflects the lack of a long time series data set that can be used to measure time use at business-cycle frequencies.

5.1. Real Business Cycles with Home Production

In Benhabib et al.’s (1991) model, time is discrete and the horizon is infinite, $t = 0, 1, 2, \dots$. The market sector is similar to the standard real-business-cycle model (see, e.g., King & Rebelo 1999). There is a representative household that provides labor services N_t^m and capital services K_{t-1}^m to a competitive, profit-maximizing producer. Final market goods are produced according to the Cobb-Douglas technology:

$$Y_t = \exp(z_t^m)(K_{t-1}^m)^{\alpha_m}(N_t^m)^{1-\alpha_m}, \quad (27)$$

where z_t^m denotes an exogenous technology shock in the market sector and $\alpha_m \in (0, 1)$.

The model departs from the standard real-business-cycle model by introducing a home sector in which goods can be produced. In the home sector, the household good is produced according to a Cobb-Douglas technology that combines time in household activities (N_t^h) with household capital goods (K_{t-1}^h):

$$C_t^h = \exp(z_t^h) \left(K_{t-1}^h \right)^{\alpha_b} \left(N_t^h \right)^{1-\alpha_b}, \quad (28)$$

where z_t^h denotes an exogenous technology shock in the home sector and $\alpha_b \in (0, 1)$.

The representative household has preferences defined over bundles of aggregate consumption C_t and leisure L_t :

$$E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, L_t), \quad (29)$$

where $\beta \in (0, 1)$ is the discount factor. The utility function is specified as

$$U(C_t, L_t) = \frac{(C_t^{1-b} L_t^b)^{1-\gamma} - 1}{1-\gamma}, \quad (30)$$

with $\gamma > 0$ and $b \in (0, 1)$. The difference relative to the standard real-business-cycle model is that here aggregate consumption is a basket of both market-produced and home-produced goods:

$$C_t = \left((1-a)(C_t^m)^\rho + a(C_t^h)^\rho \right)^{\frac{1}{\rho}}, \quad (31)$$

where $a \in (0, 1)$, and $\varepsilon = 1/(1-\rho) > 0$ denotes the elasticity of substitution between market and home consumption goods. To be precise, the elasticity of substitution $\varepsilon = 1/(1-\rho) > 0$ in the utility function in Equation 31 is not the same object as the elasticity of substitution σ between time and expenditures studied in Section 4. At the limit when $\alpha_b \rightarrow 0$, that is, when the home good is produced only with time and not with capital, the two elasticities are equal. Typically, the parameter α_b is calibrated to a low value (e.g., 0.08 in Benhabib et al. 1991), so quantitatively the two elasticities are quite similar.

In the beginning of period t , the household owns a total stock of capital K_{t-1} and invests a total of X_t in new capital goods. Total investment X_t is allocated between the two sectors, $X_t = X_t^m + X_t^h$. We note that capital goods are produced exclusively in the market sector, but they can be used as inputs either in market or in home production. The law of motion for capital stock $j = m, h$, is

$$K_t^j = X_t^j + (1-\delta)K_{t-1}^j. \quad (32)$$

Let w_t be the competitive wage and r_t the competitive rental rate of market capital that the household receives from the firm. The household chooses sequences of consumption, leisure, market and home work, and capital stocks to maximize utility subject to the budget constraint and the time constraint:

$$C_t^m + X_t = w_t N_t^m + r_t K_{t-1}^m, \quad (33)$$

$$L_t + N_t^m + N_t^h = 1. \quad (34)$$

Finally, the resource constraint for the aggregate economy is

$$Y_t = C_t^m + X_t. \quad (35)$$

To close the model, we specify a stochastic process for the technology shocks $Z_t = [z_t^m, z_t^h]'$:

$$Z_t = \mathbf{R}Z_{t-1} + v_t, \quad (36)$$

where $v_t \sim \mathbf{N}(0, \Sigma)$.

The competitive equilibrium of the model is defined as a sequence of quantities and prices such that households maximize their utility subject to the budget constraint, time constraint, and available technology in the home sector; firms maximize their profits subject to the available technology in the market sector; and all markets clear.

Benhabib et al. (1991) show that the real-business-cycle model with home production performs better than the standard real-business-cycle model along a number of dimensions. Specifically, in a calibrated version of their model, one of the main findings is that home production increases the volatility of labor and consumption relative to output. This is because home production introduces an additional margin of substitution toward which market work and market consumption can be directed following exogenous technology shocks. Second, the introduction of technology shocks in the home sector lowers significantly the correlation of productivity with labor hours. This is because technology shocks in the home sector shift the labor supply schedule and tend to create a negative correlation between productivity and hours. This tends to offset the positive correlation induced by technology shocks in the market sector, which shift the labor demand schedule.

However, the model also produces some notable discrepancies relative to the data. As Greenwood & Hercowitz (1991) show, it produces a counterfactual negative correlation between investment in the market sector and that in the home sector. This is because in a two-sector frictionless model, resources tend to flow to the most productive sector. In general, this implies that investment does not increase in both sectors simultaneously following a technology shock in one sector. Greenwood & Hercowitz (1991) show that introducing highly correlated technology shocks between the home and the market sector and increasing the complementarity of time and capital in the production of home goods help address this discrepancy. Chang (2000) shows that adjustment costs in the accumulation of capital help resolve the investment anomaly when time and capital are substitutes in the production of home goods.

The home production model has been used successfully to address a number of additional business-cycle regularities. McGrattan et al. (1997) augment the home production model with fiscal policy and discuss how to estimate the model using maximum likelihood methods. Canova & Ubide (1998) show that home production helps lower the international correlation of consumption. Baxter & Jermann (1999) show how home production can generate “excess sensitivity” of consumption to predictable income changes. Karabarbounis (2011) discusses the determination of real exchange rates in a model with a home sector and shows that home production helps explain why real exchange rates are uncorrelated with the ratio of consumption across countries.

5.2. Time Use over the Business Cycle

Models of aggregate fluctuations with home production typically assume a high degree of substitution of time at business-cycle frequencies to match business-cycle facts. For instance,

in their most preferred specification, Benhabib et al. (1991) set the elasticity of substitution between market and home goods equal to $\varepsilon = 1/(1 - \rho) = 5$. Even though a strong incentive to substitute time intratemporally is central to models with home production, the lack of a long data set covering the time use of individuals for many years has prevented a systematic analysis of time use at business-cycle frequencies.¹⁵ Here we describe some initial attempts to measure the allocation of time over the business cycle.

Burda & Hamermesh (2010) use time diaries from four countries to study the relationship between unemployment and home production. In general, they find that the unemployed spend very little of their additional time on home production activities. However, focusing on US data from the ATUS (2003–2006), they show that individuals residing in metro areas with temporarily high unemployment levels allocate a large fraction (approximately 75%) of a given decrease of market work to home production.

Aguiar et al. (2011) explore ATUS data between 2003 and 2010 to measure how foregone market work is allocated to alternate time uses during both the nonrecessionary period and the recent recession. Given the short time series of the data set, they argue that simply comparing time spent on a given time use category prior to the recession with that during the recession confounds the business-cycle effects on time use. The reason is that during the prerecessionary period (2003–2008), home production and leisure display noticeable trends that are extensions of the trends that started in the 1960s, which we discuss above. Instead, they identify the business-cycle effects on time use from cross-state variation with respect to the severity of the recessions. Looking at the cross-state differences in the changes of market work allows one to control for a common low-frequency trend in time use. With this identification strategy, the main finding is that roughly 25%–30% of the foregone market work hours are allocated to increased home production and approximately 5% to increased child care. Leisure activities absorb a total of approximately 50%–55% of the foregone market work hours, with sleep and television watching accounting for more than 30%. Finally, Aguiar et al. (2011) show that investments in education, civic activities, and health care absorb an important fraction of the decrease in market work hours (more than 10%), whereas job search absorbs a relatively small fraction of the decrease in market work hours. The latter finding is not surprising, given how little time the unemployed spent searching for a job (Krueger & Mueller 2010).

5.3. Macro Estimates of the Elasticity of Substitution

Based on aggregate US data and likelihood methods, McGrattan et al. (1997) estimate the elasticity of substitution between market and home goods, $\varepsilon = 1/(1 - \rho)$, to be slightly less than 2, whereas Chang & Schorfheide (2003) estimate it to be approximately 2.3. Karabarbounis (2011) estimates an elasticity of substitution of approximately 3.4. This parameter is identified from the requirement that the home production model produces a labor wedge with cyclical and long-run moments that match the moments of the labor wedge observed in the data.

¹⁵Ramey (2009) and Ramey & Francis (2009) provide a comprehensive and very long data set with annual observations of aggregate hours spent on home production. Karabarbounis (2011) uses this data set to argue that the home sector is countercyclical to market output. However, this provides only indirect evidence for the cyclical behavior of home production because most of the annual observations for hours in home production in Ramey's (2009) and Ramey & Francis's (2009) data set are imputed based on the aggregate employment rate.

Aguiar et al. (2011) assess quantitatively whether the business-cycle model of Benhabib et al. (1991) generates movements in the allocation of time over the business cycle that are consistent with the evidence they document from the ATUS data (2003–2010). Specifically, they use their estimate of the fraction of foregone market work that is reallocated toward home production in the data to identify the elasticity parameter ε . In a version of the model in which sleep, eating, and personal care are excluded from leisure activities, an elasticity of approximately 2.0–2.5 produces a reallocation of market hours to home hours in the model that matches the actual behavior of households in the data at business-cycle frequencies. In a version of the model in which these activities are included in leisure, the estimated elasticity increases to 3.5–4.0. This is because estimates from the US-state sample show that sleep is one of the least elastic time use categories in the data.

6. CONCLUSIONS AND FUTURE RESEARCH

Above we argue that the proliferation of new data sets and the harmonization of older data have allowed researchers to make significant progress in our understanding of how individuals allocate their time away from market work. We highlight how these new data can be used to test older theories of time allocation and to better inform us about a variety of economic phenomena, including long-run trends in hours in various countries, life-cycle patterns of consumption, labor supply elasticities, and the business-cycle behavior of macroeconomic aggregates.

This strand of research has produced a number of important results. For example, nonmarket work (home production) is falling in the United States over time dramatically. Leisure time is increasing, but the extent to which it increases depends on the classification of activities. Studies from other countries confirm that home production time is decreasing over time for women, whereas the results for leisure are mixed. Moreover, child care should be treated separately from home production or leisure, given its differential response to changes in prices, income, and demographics. There are also rich cross-country differences in the allocation of time other than market work. These differences have been used to improve our understanding of the effects of taxation and labor market institutions on market work time. Models that emphasize the intratemporal substitution of time and goods (as opposed to those that emphasize only intertemporal substitution) predict that spending declines in retirement because retirees, given their low opportunity cost of time, are substituting away from expenditures and toward time in the production of final goods. Indeed, the evidence shows that retired people spend more time in home production and have less expenditures in food consumption relative to similar nonretired people. In addition, the intratemporal substitution of time and goods helps us understand why some spending categories decrease over the life cycle (e.g., clothing), others increase (e.g., entertainment), and others display the familiar hump-shaped profile (e.g., food). The incorporation of home production into models has deepened our understanding of gender differences in labor supply elasticities and can help researchers estimate more reliably the elasticity of labor supply. Dynamic general equilibrium models with home production are successful in amplifying the volatility of hours and expenditure over the business cycle, as long as the elasticity of substitution between market-produced and home-produced goods is sufficiently high. Finally, various micro and macro estimates of the elasticity of substitution between time and expenditures (or between market-produced and home-produced goods) suggest a high value (approximately 2).

Among those topics, the least well-developed area is the study of the allocation of time at business-cycle frequencies. Given that the ATUS presently covers only 8 years of data, standard statistical methods that macroeconomists use to filter time series and make them amenable to business-cycle analysis are not yet applicable. We expect that these methods will become increasingly popular in the next decade among researchers working on time use and business cycles.

Above we review some of the estimates of the key behavioral parameter that governs the willingness of households to substitute intratemporally. Although the identifying assumptions differ across studies, many estimates are in general consistent with each other. However, similar to the returns to market work, it is possible that the returns to nonmarket work change during recessions. Under this scenario, it may not be appropriate to use the elasticities of substitution estimated during nonrecessionary periods to predict the joint movements of market work, nonmarket work, and expenditure during recessions. An important area of future research would be to assess the returns to nonmarket work during recessions.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Dan Hamermesh, Jan Zilinsky, and two anonymous reviewers for useful suggestions and helpful comments.

LITERATURE CITED

- Aguiar M, Hurst E. 2005. Consumption versus expenditure. *J. Polit. Econ.* 113:919–48
- Aguiar M, Hurst E. 2007a. *Comments on Valerie A. Ramey's "How much has leisure inequality really increased since 1965?"*. Work. Pap., Booth Sch. Bus., Univ. Chicago
- Aguiar M, Hurst E. 2007b. Lifecycle prices and production. *Am. Econ. Rev.* 97:1533–59
- Aguiar M, Hurst E. 2007c. Measuring trends in leisure: the allocation of time over five decades. *Q. J. Econ.* 122:969–1006
- Aguiar M, Hurst E. 2008. *Deconstructing lifecycle expenditure*. NBER Work. Pap. 13893
- Aguiar M, Hurst E. 2009. *The Increase of Leisure Inequality: 1965–2005*. Washington, DC: Am. Enterp. Inst.
- Aguiar M, Hurst E, Karabarbounis L. 2011. *Time use during recessions*. NBER Work. Pap. 17259
- Alesina A, Ichino A. 2009. *L'Italia Fatta in Casa*. Milan: Mondadori
- Alesina A, Ichino A, Karabarbounis L. 2011. Gender-based taxation and the division of family chores. *Am. Econ. J. Econ. Policy* 3(2):1–40
- Attanasio O. 1999. Consumption. In *Handbook of Macroeconomics*, Vol. 1, ed. J Taylor, M Woodford, pp. 741–812. Amsterdam: North Holland
- Attanasio O, Battistin E, Ichimura H. 2004. *What really happened to consumption inequality in the US?* NBER Work. Pap. 10338
- Attanasio O, Weber G. 2010. Consumption and saving: models of intertemporal allocation and their implications for public policy. *J. Econ. Lit.* 48:693–751
- Banks J, Blundell R, Tanner S. 1998. Is there a retirement-savings puzzle? *Am. Econ. Rev.* 88:769–88

- Baxter M, Jermann U. 1999. Household production and the excess sensitivity of consumption to current income. *Am. Econ. Rev.* 89:902–20
- Becker GS. 1965. A theory of the allocation of time. *Econ. J.* 75:493–517
- Benhabib J, Rogerson R, Wright R. 1991. Homework in macroeconomics: household production and aggregate fluctuations. *J. Polit. Econ.* 99:1166–87
- Bernheim D, Skinner J, Weinberg S. 2001. What accounts for the variation in retirement wealth among U.S. households? *Am. Econ. Rev.* 91:832–57
- Blau F, Kahn L. 2007. Changes in the labor supply behavior of married women: 1980–2000. *J. Labor Econ.* 25:393–438
- Blundell R, MaCurdy T. 1999. Labor supply: a review of alternative approaches. In *Handbook of Labor Economics*, Vol. 3, ed. O Ashenfelter, D Card, pp. 1559–695. Amsterdam: North Holland
- Burda M, Hamermesh DS. 2010. Unemployment, market work and household production. *Econ. Lett.* 107:131–33
- Burda M, Hamermesh DS, Weil P. 2007. *Total work, gender and social norms*. NBER Work. Pap. 13000
- Burda M, Hamermesh DS, Weil P. 2008. The distribution of total work in the EU and US. In *Working Hours and Job Sharing in the EU and USA: Are Europeans Lazy? Or Americans Crazy?*, ed. T Boeri, pp. 13–91. New York: Oxford Univ. Press
- Canova F, Ubide A. 1998. International business cycles, financial markets and household production. *J. Econ. Dyn. Control* 22:545–72
- Carroll C. 1997. Buffer stock saving and the life cycle/permanent income hypothesis. *Q. J. Econ.* 112:1–56
- Carroll C, Summers L. 1991. Consumption growth parallels income growth: some new evidence. In *National Saving and Economic Performance*, ed. D Bernheim, J Shoven, pp. 305–29. Chicago: Univ. Chicago Press
- Chang Y. 2000. Comovement, excess volatility, and home production. *J. Monet. Econ.* 46:385–96
- Chang Y, Schorfheide F. 2003. Labor-supply shifts and economic fluctuations. *J. Monet. Econ.* 50:1751–68
- Deaton A. 1991. Saving and liquidity constraints. *Econometrica* 59:1221–48
- Diamond P. 1982. Aggregate demand management in search equilibrium. *J. Polit. Econ.* 90:881–94
- Folbre N, Yoon J, Finnoff K, Fuligni A-S. 2005. By what measure? Family time devoted to children in the United States. *Demography* 42:373–90
- Freeman R, Schettkat R. 2005. Marketization of household production and the EU-US gap in work. *Econ. Policy* 20:6–50
- Gelber A, Mitchell J. 2009. *Taxes and time allocation: evidence from single women*. NBER Work. Pap. 15583
- Ghez G, Becker GS. 1975. *The Allocation of Time and Goods over the Life Cycle*. New York: Columbia Univ. Press
- Gimenez-Nadal JI, Sevilla-Sanz A. 2011. Trends in time allocation: a cross-country analysis. *Discuss. Pap.* 547, Univ. Oxford
- Gonzalez Chapela J. 2011. Recreation, home production, and intertemporal substitution of female labor supply: evidence on the intensive margin. *Rev. Econ. Dyn.* 14:532–48
- Gourinchas PO, Parker J. 2002. Consumption over the life cycle. *Econometrica* 70:47–89
- Greenwood J, Hercowitz Z. 1991. The allocation of capital and time over the business cycle. *J. Polit. Econ.* 99:1188–214
- Greenwood J, Rogerson R, Wright R. 1995. Household production in real business cycle theory. In *Frontiers of Business Cycle Research*, ed. F Cooley, pp. 157–74. Princeton, NJ: Princeton Univ. Press
- Greenwood J, Seshadri A, Yorukoglu M. 2005. Engines of liberation. *Rev. Econ. Stud.* 72:109–33
- Gronau R. 1987. Home production: a survey. In *Handbook of Labor Economics*, Vol. 1, ed. O Ashenfelter, R Layard, pp. 273–304. Amsterdam: North Holland
- Gronau R. 1997. The theory of home production: the past ten years. *J. Labor Econ.* 15:197–205

- Guryan J, Hurst E, Kearney M. 2008. Parental education and parental time with children. *J. Econ. Perspect.* 22:23–46
- Haider S, Stephens M. 2007. Is there a retirement consumption puzzle? Evidence using subjective retirement expectations. *Rev. Econ. Stat.* 89:247–64
- Hamermesh DS. 1984. Consumption during retirement: the missing link in the life cycle. *Rev. Econ. Stat.* 66:1–7
- Hamermesh DS, Frazis H, Stewart J. 2005. Data watch: the American Time Use Survey. *J. Econ. Perspect.* 19:221–32
- Hamermesh DS, Pfann GA. 2005. Time-use data in economics. *Eur. Econ. Rev.* 49:1–7
- Hurd M, Rohwedder S. 2006. *Some answers to the retirement-consumption puzzle*. NBER Work. Pap. 12057
- Hurst E. 2008. *The retirement of a consumption puzzle*. NBER Work. Pap. 13789
- Juster T. 1985. Preferences for work and leisure. In *Time, Goods, and Well-Being*, ed. T Juster, F Stafford, pp. 335–51. Ann Arbor: Univ. Michigan Press
- Juster T, Stafford F. 1991. The allocation of time: empirical findings, behavioral models, and problems of measurement. *J. Econ. Lit.* 29:471–522
- Karabarbounis L. 2011. *Labor wedges and open economy puzzles*. Work. Pap., Booth Sch. Bus., Univ. Chicago
- Katz L, Autor D. 1999. Changes in the wage structure and earnings inequality. In *Handbook of Labor Economics*, Vol. 3, ed. O Ashenfelter, D Card, pp. 1463–555. Amsterdam: North Holland
- Kimmel J, Connelly R. 2007. Mothers' time choices: caregiving, leisure, home production, and paid work. *J. Hum. Resour.* 42:643–81
- King R, Rebelo S. 1999. Resuscitating real business cycles. In *Handbook of Macroeconomics*, Vol. 1, ed. J Taylor, M Woodford, pp. 927–1007. Amsterdam: North Holland
- Krueger A, Mueller A. 2010. Job search and unemployment insurance: new evidence from time use data. *J. Public Econ.* 94:298–307
- Kydland F, Prescott E. 1982. Time to build and aggregate fluctuations. *Econometrica* 50:1345–70
- Lee J, Kawaguchi D, Hamermesh D. 2011. *The gift of time*. Work. Pap., Univ. Texas, Austin
- McGrattan E, Rogerson R. 2004. Changes in hours worked: 1950–2000. *Fed. Reserve Bank Minneap. Q. Rev.* 28(1):14–33
- McGrattan E, Rogerson R, Wright R. 1997. An equilibrium model of the business cycle with household production and fiscal policy. *Int. Econ. Rev.* 38:267–90
- Mortensen D, Pissarides C. 1994. Job creation and job destruction in the theory of unemployment. *Rev. Econ. Stud.* 61:397–415
- Ramey G, Ramey V. 2010. The rug rat race. *Brookings Pap. Econ. Act.* 2010:129–76
- Ramey V. 2007. *How much has leisure really increased since 1965?* Work. Pap., Univ. California, San Diego
- Ramey V. 2009. Time spent in home production in the 20th century United States: new estimates from old data. *J. Econ. Hist.* 69:1–47
- Ramey V, Francis N. 2009. A century of work and leisure. *Am. Econ. J. Macroecon.* 1(2):189–224
- Robinson J, Godbey G. 1999. *Time for Life: The Surprising Ways Americans Use Their Time*. University Park: Pennsylvania State Univ. Press. 2nd ed.
- Rupert P, Rogerson R, Wright R. 1995. Estimating substitution elasticities in household production models. *Econ. Theory* 6:179–93
- Rupert P, Rogerson R, Wright R. 2000. Homework in labor economics: household production and intertemporal substitution. *J. Monet. Econ.* 46:557–79
- Sayer L, Bianchi S, Robinson J. 2004. Are parents investing less in children? Trends in mothers' and fathers' time with children. *Am. J. Sociol.* 11:1–43
- Schwerdt G. 2005. Why does consumption fall at retirement? Evidence from Germany. *Econ. Lett.* 89:300–5
- Stewart J. 2009. *Tobit or not Tobit*. Work. Pap. 432, Bur. Labor Stat.
- Zeldes S. 1989. Consumption and liquidity constraints: an empirical investigation. *J. Polit. Econ.* 97:305–46



Contents

Paul Samuelson's Legacy <i>Avinash Dixit</i>	1
Saving Money or Just Saving Lives? Improving the Productivity of US Health Care Spending <i>Katherine Baicker, Amitabh Chandra, and Jonathan S. Skinner</i>	33
International Comparisons in Health Economics: Evidence from Aging Studies <i>James Banks and James P. Smith</i>	57
Rare Macroeconomic Disasters <i>Robert J. Barro and José F. Ursúa</i>	83
Endogenous Extreme Events and the Dual Role of Prices <i>Jon Danielsson, Hyun Song Shin, and Jean-Pierre Zigrand</i>	111
The Distribution of Teacher Quality and Implications for Policy <i>Eric A. Hanushek and Steven G. Rivkin</i>	131
Economic Modeling and Analysis of Educational Vouchers <i>Dennis Epple and Richard Romano</i>	159
Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers <i>Joseph G. Altonji, Erica Blom, and Costas Meghir</i>	185
Credit Constraints in Education <i>Lance Lochmer and Alexander Monge-Naranjo</i>	225
New Perspectives on Statistical Decisions Under Ambiguity <i>Jörg Stoye</i>	257
The Empirics of Firm Heterogeneity and International Trade <i>Andrew B. Bernard, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott</i>	283
Natural Resource Wealth: The Challenge of Managing a Windfall <i>Frederick van der Ploeg and Anthony J. Venables</i>	315

The Economics and Politics of Women's Rights <i>Matthias Doepke, Michèle Tertilt, and Alessandra Voena</i>	339
Recent Developments in the Economics of Time Use <i>Mark Aguiar, Erik Hurst, and Loukas Karabarbounis</i>	373
Life-Cycle Wage Growth and Heterogeneous Human Capital <i>Carl Sanders and Christopher Taber</i>	399
Behavioral Economics and Psychology of Incentives <i>Emir Kamenica</i>	427
The Relationship Between Economic Preferences and Psychological Personality Measures <i>Anke Becker, Thomas Deckers, Thomas Dohmen, Armin Falk, and Fabian Kosse</i>	453
Corruption in Developing Countries <i>Benjamin A. Olken and Rohini Pande</i>	479
A Reduced-Form Approach to Behavioral Public Finance <i>Sendhil Mullainathan, Joshua Schwartzstein, and William J. Congdon</i>	511
Recent Research on the Economics of Patents <i>Bronwyn H. Hall and Dietmar Harhoff</i>	541
Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences <i>Helga Fehr-Duda and Thomas Epper</i>	567
The Theory of Clubs and Competitive Coalitions <i>Myrna Wooders</i>	595
The Promises and Pitfalls of Genoeconomics <i>Daniel J. Benjamin, David Cesarini, Christopher F. Chabris, Edward L. Glaeser, David I. Laibson, Vilmundur Guðnason, Tamara B. Harris, Lenore J. Launer, Shaun Purcell, Albert Vernon Smith, Magnus Johannesson, Patrik K.E. Magnusson, Jonathan P. Beauchamp, Nicholas A. Christakis, Craig S. Atwood, Benjamin Hebert, Jeremy Freese, Robert M. Hauser, Taissa S. Hauser, Alexander Grankvist, Christina M. Hultman, and Paul Lichtenstein</i>	627

Indexes

Cumulative Index of Contributing Authors, Volumes 1–4	663
Cumulative Index of Chapter Titles, Volumes 1–4	665

Errata

An online log of corrections to *Annual Review of Economics* articles may be found at <http://econ.annualreviews.org>