The Anatomy of the Aggregate Labor Supply Elasticity*

Riccardo Fiorito† Giulio Zanella‡

First Draft: September 2008. This Draft: December 2011.

Abstract

We show that the aggregate Frisch elasticity of labor supply can greatly exceed the corresponding individual-level parameter, and we illustrate the "anatomy" of the former in terms of intensive and extensive margins. The methodology consists of using micro data from the PSID to construct a panel of individuals and an aggregate time series obtained by aggregating these individuals each year. These two data sets represent *exactly* the same sample at different levels of aggregation, and we use them to identify the parameters of two distinct MaCurdy-type micro and macro equations. We find a micro elasticity of about 0.1 and a much larger macro elasticity that ranges from 1.1 to 1.7. There is no conflict between the two estimates: the micro one reflects only the intensive margin while the macro one reflects, in addition, the much more volatile extensive margin. Furthermore, aggregation of only continuously employed individuals allows us to provide a reliable estimate of the intensive margin elasticity in the range 0.3–0.4. This implies an extensive margin elasticity in the range 0.8–1.4. These findings suggest that micro evidence is not a benchmark for assessing how large the Frisch elasticity of labor supply should be in a model of the aggregate economy.

JEL Classification Codes: E13, E32, J22
Keywords: labor supply, intensive margin, extensive margin, calibration.

*This paper previously circulated under the title "Labor Supply Elasticities: Can Micro Be Misleading for Macro?" We are grateful to Pierangelo De Pace, Carlo Favero, Francesco Nucci, Giacomo Rondina, Peter Rupert, and in particular Gianluca Violante and two anonymous referees for extremely useful comments as well as to seminar participants at UCSB, the universities of Roma-Tor Vergata and Siena, Italy's Department of the Treasury, and participants at the respective sessions at the SED 2008 Meeting in Boston and the ES 2008 European Meeting in Milan. The usual disclaimer applies.

†University of Siena, Italy. Email: rfiorito@iol.it

‡Corresponding author. University of Bologna, Italy. Email: giulio.zanella@unibo.it

1 Introduction

In this paper we provide an empirical reconciliation of micro and macro labor supply elasticities. Our goal is to show that the aggregate Frisch elasticity of labor supply is much larger than the corresponding individual-level elasticity and yet is fully consistent with it. From analyses of panel data it is well known that individual hours are relatively wage inelastic in the short run. This fact is often regarded as a challenge to the benchmark real business cycle (RBC) model (Prescott, 1986). The RBC methodology rests on Lucas's (1980) suggestion to draw parameters from census information and panel data, but the model requires that aggregate labor be elastic in order to generate fluctuations that are consistent with business cycle facts (Kydland and Prescott, 1982; Prescott, 2006). This is not necessarily a challenge because the micro and macro Frisch elasticities of labor supply are conceptually different objects. The micro parameter is the elasticity of individual hours conditional on being employed whereas the macro parameter is the elasticity of total hours, which reflects variations in hours per worker as well as in the employment level (Heckman, 1993; Browning, Hansen, and Heckman, 1999). In other words, the aggregate elasticity reflects both the intensive and extensive margins.

We contribute to this line of research by estimating mutually consistent micro and macro Frisch elasticities of labor supply for the same population. By "mutually consistent" we mean that the two estimates are based on the same micro data and the same specification; the only difference is the level of aggregation. Specifically, we estimate two distinct, MaCurdy-type (after MaCurdy, 1981) labor supply equations. The first, the micro equation, relates individual hours to the individual wage rate and is estimated using data from the Panel Study of Income Dynamics (PSID). The second, the macro equation, relates aggregate hours to the aggregate wage rate and is estimated in a time series obtained by aggregating the single waves of the PSID each year. We find a micro elasticity of about 0.1, a small value in line with benchmark microeconometric estimates, and a much larger macro elasticity in the range 1.1–1.7.[1] Two additional results are worth emphasizing. First, the use of microdata allows us to decompose the macro elasticity into intensive and extensive margins for the whole sample and for subgroups of interest. When looking at subgroups, we document that individuals who are non–prime age, married or cohabitating with partners who work, and/or low educated are "marginal workers" who make aggregate labor much more elastic than at the individual level. We find mixed evidence about women. Second, by focusing on the special subgroup of continuously employed individuals (for whom the extensive margin is inactive) we

[1] For samples that are more representative of the US population, we estimate the aggregate elasticity to be in the range 0.6–1.1. However—as discussed in detail later in the paper—this range should be regarded as a lower bound for the US economy.

provide an estimate of the intensive margin that presumably removes, via aggregation, most of the bias affecting estimates on disaggregated data. Such an estimate is in the range 0.3–0.4. The fact that this is larger than the corresponding estimate on disaggregated data confirms that the bias is towards zero.

This paper contributes to a large and expanding literature. The first generation of microeconometric estimates of the Frisch elasticity of labor supply ranges from about 0 to about 0.2 for men and from about 0 to about 1 for married women.[2] The macroeconomic evidence is far less numerous and somewhat conflicting. In their seminal paper, Lucas and Rapping (1969) estimate an elasticity of 1.4. At the other extreme, Mankiw, Rotemberg, and Summers (1985) reject the intertemporal substitution hypothesis altogether, but they explicitly focus on the "labor input per member of the adult population" (p. 235)—that is, an intensive margin. The importance of including the extensive margin as well is nicely illustrated by Alogoskoufis (1987), who rejects the intertemporal substitution hypothesis for fluctuations in hours per worker but cannot reject the same hypothesis for fluctuations in aggregate employment. This is not surprising, given that the bulk of the cyclical adjustment of total hours occurs via adjustments in the employment stock (Hansen, 1985; Kydland, 1995; see also our own computations in Section 2).

The necessity of reconciling the large aggregate elasticity assumed in calibration studies with the small elasticity estimated in microeconometric studies led to the development of variants and extensions of the benchmark RBC model in order to better accommodate the data. A precursor is the seminal work of Kydland and Prescott (1982) based on nonseparability of leisure at different points in time. A prominent position in the field is occupied by the indivisible labor model (Rogerson, 1988; Hansen, 1985), where people can either work a fixed number of hours or not work at all. In that model, all labor changes take place at the extensive margin.

A second generation of empirical studies reduces the elasticity gap by arguing that benchmark micro regressions are misspecified. Examples include the omission of: home production (Rupert, Rogerson, and Wright, 2000), actual expectations of wage changes (Pistaferri, 2003), time devoted to accumulating human capital (Imai and Keane, 2004; Wallenius, 2011), nonseparability of consumption and leisure (Ziliak and Kniesner, 2005), and liquidity constraints (Domeij and Flodén, 2006). Although it is very important to obtain a correct estimate of the micro elasticity, our work shows that the elasticity gap, per se, is not an issue. Chetty, Friedman, Manoli, and Weber (2011) review the empirical evidence on the intensive and extensive margins elasticities generated by quasi-experimental variations. The intensive margin elasticity they report (0.3–0.5) is in line with our estimate based

[2] See the surveys of Pencavel (1986), Killingsworth and Heckman (1986), Card (1994), and Blundell and MaCurdy (1999).

on continuously employed workers (0.3–0.4). However, the latter in combination with our estimate of the aggregate elasticity implies an extensive margin elasticity in the range 0.8–1.4, substantially larger than the corresponding estimate reported by Chetty, Friedman, Manoli, and Weber (0.2–0.3).

A few studies have performed a similar exercise starting from a calibrated micro elasticity. Chang and Kim (2006) combine the indivisible labor assumption and the heterogeneity of reservation wages in an incomplete markets model. Assuming an individual elasticity of 0.4, they find an aggregate elasticity of about 1. Rogerson and Wallenius (2009) assume an individual elasticity ranging from 0.05 to 1.25 and find that the corresponding macro elasticity ranges between 2.25 and 3 in a model where the mapping between hours of work and labor services is initially flat. Our paper can be regarded as the empirical counterpart of these calibration studies.

A related work that employs survey data is Gourio and Noual (2009), who use 14 years of monthly observations (National Longitudinal Survey of Youth data from 1979 to 1992) to estimate the aggregate elasticity as the hazard rate of the distribution of reservation wages. Their estimate is 1.3.[3] The basic idea in Gourio and Noual is the same as in Chang and Kim (2006) and in this paper: labor is more elastic at the aggregate than at the individual level because of marginal workers who move into and out of employment in response to wage changes. There are important differences between this work and ours. First, Gourio and Noual shut down the intensive margin, so one cannot tell how their estimate compares with the underlying micro elasticity. Second, estimating two MaCurdy equations allows us to avoid making distributional assumptions and calibrating parameters. Third, the data sources are different. Despite these important distinctions, it is remarkable that the magnitude of the macro estimates that we and they produce are similar.[4]

The remainder of this paper is organized as follows. In Section 2, we illustrate the importance of the intensive and extensive margins and discuss the statistical relations between the micro and macro elasticities we estimate. Section 3 describes the data set, compares key sample statistics with US data from the US Census and the Current Population Survey (CPS), and illustrates our aggregation procedure. Section 4 describes how the parameters of interest are identified, and Section 5 presents our results. Section 6 concludes.

[3] Browning, Hansen, and Heckman (1999) observe that the intertemporal elasticity of substitution is typically larger at higher frequencies than at lower frequencies.

[4] Mulligan (1998) uses *cross-sectional* data from the Current Population Survey to estimate the intertemporal elasticity of substitution for employed men and for all working-age men, employed or not. In the basic specification, he finds elasticities of 0.65 for the former and 1.41 for the latter.

2 Intensive and extensive margins

The indivisible labor case (Rogerson, 1988; Hansen, 1985) accommodates in an extreme way the evidence that labor supply adjustments at the extensive margin (variations in employment) dominate those at the intensive margin (variations in individual or average hours). As in Hansen (1985), if we denote by n_t the number of individuals employed at time t and by $\overline{h}_t$ the average number of hours these individuals work, then aggregate labor is $H_t \equiv n_t \overline{h}_t$. By taking logs, the variance of the log labor input can be decomposed as follows:

$$var(\ln H_t) = var(\ln n_t) + var(\ln \overline{h}_t) + 2cov(\ln n_t, \ln \overline{h}_t). \tag{1}$$

The share of the total variation that is due to n_t (i.e., the ratio of $var\,(\ln n_t)$ to $var\,(\ln H_t)$ in this equation) provides a measure of the importance of the extensive margin. For Hodrick–Prescott (HP-) detrended quarterly US data ranging from 1955 to 1984, Hansen (1985) finds that employment changes account for 55% of the variation in total hours, whereas hours per worker account for only 20%.[5] We update this empirical exercise using more recent data for five OECD countries. Specifically, we collect annual data for the United States, Canada, Japan, France, and the United Kingdom; we then use these data to compute the standard deviations of detrended log employment, log hours per worker, and log total hours as well as the contribution of the macro extensive and the macro intensive margins to the total variance of detrended log total hours. The results are reported in Tables 1 and 2, for three different filters. Table 1 shows that the extensive margin is systematically more volatile, except for Japan and, when filtering the series through a first difference, France. Table 2 reports the corresponding incidence, computed as in the Hansen decomposition—that is, our equation (1).

Blundell, Bozio, and Laroque (2011) use micro data from France, the United Kingdom, and the United States to construct bounds for the contribution of changes in the participation rate and in hours per worker to changes in hours per person between 1977 and 2007 for different demographic groups. This is a different decomposition than the one discussed previously, yet their analysis also indicates that participation always dominates hours per worker as a source of variation in hours per person in the United States, for all age/gender groups. In the United Kingdom this is true only for women above age 30 and (marginally) for old men; in France, it is true for young and old men as well as for prime-age women.

[5] This pattern is observed in several countries: in HP-filtered, quarterly manufacturing data for 1960–1989, Fiorito and Kollintzas (1994) find that the volatility of employment deviations from the smooth trend always exceeds the corresponding volatility in hours per worker.

Table 1

Standard deviation of the labor input and its components.

	Employment (n)			Hours per worker ($\overline{h}$)			Total hours ($n \times \overline{h}$)		
	Δln	HP$_{100}$	HP$_{6.25}$	Δln	HP$_{100}$	HP$_{6.25}$	Δln	HP$_{100}$	HP$_{6.25}$
USA	1.59	1.43	1.03	0.62	0.56	0.40	1.97	1.78	1.32
Canada	1.67	1.68	1.03	0.59	0.54	0.37	1.96	2.05	1.29
Japan	0.93	0.78	0.47	0.98	0.99	0.62	1.49	1.25	0.88
France	1.09	1.23	0.72	1.18	1.06	0.70	1.60	1.59	0.95
UK	1.62	1.96	1.12	1.22	1.15	0.72	2.50	2.83	1.63

Note: All figures represent standard deviations of detrended log series. The filters are first-differencing (Δln), and Hodrick–Prescott with smoothing parameter 100 (HP$_{100}$) or 6.25 (HP$_{6.25}$). OECD annual data, 1970–2009, total civilian employment

Table 2

Shares of the variance of the labor input (ref. Table 1) accounted for by its components

	Employment (n)			Hours per worker ($\overline{h}$)		
	Δln	HP$_{100}$	HP$_{6.25}$	Δln	HP$_{100}$	HP$_{6.25}$
USA	65.1%	64.5%	60.9%	9.9%	9.9%	9.2%
Canada	72.6%	67.2%	63.8%	9.1%	6.9%	8.2%
Japan	39.0%	38.9%	28.5%	43.3%	62.7%	49.6%
France	46.4%	59.8%	57.4%	54.4%	44.4%	54.3%
UK	42.0%	48.0%	47.2%	23.8%	16.5%	19.5%

Note: The incidence of employment on the total variance of the detrended labor input is computed from equation (1) as $var\left(\ln n_t\right)/var\left(\ln H_t\right)$ and the incidence of hours per worker as $var\left(\ln \overline{h}_t\right)/var\left(\ln H_t\right)$. See note to Table 1 for definition of the filters.

The concurrence of intensive and extensive margins indicates that two elasticities are relevant to understanding aggregate fluctuations of the labor input. By the individual, micro elasticity we refer to the real wage elasticity of individual hours conditional on being employed; by the aggregate, macro elasticity we refer to the real wage elasticity of total hours. It is understood that in both cases the marginal utility of wealth is kept constant, so these are Frisch elasticities. In a regression framework one immediately sees that the wedge between individual and aggregate elasticities reflects the response of employment to changes in wages. Henceforth, we use lowercase for individual variables and uppercase for the corresponding aggregate quantities. Denote by ε and $\mathcal{E}$ the micro and macro elasticities of labor supply, respectively; by w_{it} and W_t individual i's and the aggregate wage rates at time t, respectively; and by h_{it} hours worked by individual i at time t. Consider the following MaCurdy (1981) regressions:

$$\begin{aligned}
\text{(individual)} \qquad \Delta \ln h_{it} &= \kappa + \varepsilon \Delta \ln w_{it} + v_{it}, & (2)\\
\text{(aggregate)} \qquad \Delta \ln H_t &= K + \mathcal{E} \Delta \ln W_t + V_t. & (3)
\end{aligned}$$

Here Δ denotes the first-difference operator, k and K are constants, and H_t denotes total hours worked at time t as before. We postpone until Section 4 the discussion of the structural interpretation of these equations and until Section 3 the definition of the aggregate wage rate. The underlying population elasticities can be written as follows:

$$\varepsilon = \underbrace{\frac{cov\left(\Delta \ln h, \Delta \ln w\right)}{var\left(\Delta \ln w\right)}}_{\text{micro intensive margin}}, \tag{4}$$

$$\begin{aligned}
\mathcal{E} &= \frac{cov\left(\Delta \ln H, \Delta \ln W\right)}{var\left(\Delta \ln W\right)} \\
&= \underbrace{\frac{cov\left(\Delta \ln \overline{h}, \Delta \ln W\right)}{var\left(\Delta \ln W\right)}}_{\text{macro intensive margin}} + \underbrace{\frac{cov\left(\Delta \ln n, \Delta \ln W\right)}{var\left(\Delta \ln W\right)}}_{\text{extensive margin}}. \qquad (5)
\end{aligned}$$

That is, the *micro* elasticity (4) consists of a single term that captures adjustments in the intensive margin (individual hours) by individuals employed in two consecutive periods and for whom $\Delta \ln h$ and $\Delta \ln w$ are observed. We label this the "micro intensive margin". In contrast, the *macro* elasticity (5) is the sum of two terms. The first reflects the covariance between the rates of change in hours per worker and in the aggregate wage rate; thus it represents a "macro intensive margin". The second term reflects the covariance between the rates of change in employment and in the aggregate wage rate, so it represents the extensive margin.

There is an intuitive analogy between the micro and macro intensive margins. The analogy is immediate and precise for the extreme case in which employment is constant and all workers work the same number of hours at the same hourly wage in the cross section. In this case, if we take the aggregate wage rate to be some average of the individual wages, then

$$\frac{cov\left(\Delta\ln h,\Delta\ln w\right)}{var\left(\Delta\ln w\right)}=\frac{cov\left(\Delta\ln\overline{h},\Delta\ln W\right)}{var\left(\Delta\ln W\right)};\tag{6}$$

that is, the micro and macro intensive margins coincide. In this extreme case, employment does not vary and so we would have $cov\left(\Delta\ln n,\Delta\ln W\right)=0$; hence the two elasticities would coincide as well (i.e., $\mathcal{E}=\varepsilon$). In other words, in general there is a wedge between micro and macro elasticities because workers are heterogeneous and because employment varies in response to aggregate wage changes. We will later compare the micro and macro intensive margins by aggregating only individuals who work in every period—a group which we label "continuously employed". This way we can estimate a macro intensive elasticity that is the direct counterpart of the micro elasticity: by considering only continuously employed individuals, one shuts off the extensive margin completely. This also makes clear that the wedge between micro and macro estimates is not due to sample selection—that is, to the fact that we observe the individual wage rate in a random sample only if hours are strictly positive. Correcting the ensuing bias (e.g. by employing Tobit or Heckit estimators) adjusts for the nonrandomness of the sample used in estimation but not for entry and exit decisions. To capture these, one must refer to an aggregate labor supply function (Heckman, 1993). This is what our macro regression does: by taking total hours as the dependent variable, we are able to capture fully the importance of both the extensive and intensive margins. As observed by Browning, Hansen, and Heckman (1999), this is why Lucas and Rapping (1969) found a much larger intertemporal elasticity of substitution of leisure than did the vast majority of empirical studies during the following decades: when considering total hours, Lucas and Rapping included entry and exit. A useful way to illustrate this point is by representing total hours at time t as the sum of four distinct flows generated by four groups, which we label "stayers" (those who work at time t and at time $t-1$), "entrants" (those who did not work at time $t-1$ but do work at time t), "leavers" (those who did work at time $t-1$ but do not work at time t), and "outsiders" (those who work neither at time t nor at time $t-1$):

$$H_t\equiv\underbrace{\sum_{i=1}^{n_t^S}h_{it}}_{\text{stayers}}+\underbrace{\sum_{i=n_t^S+1}^{n_t}h_{it}}_{\text{entrants}}+\underbrace{\sum_{i=n_t+1}^{n_t+n_t^L}h_{it}}_{\text{leavers}}+\underbrace{\sum_{i=n_t+n_t^L+1}^{N_t}h_{it}}_{\text{outsiders}}.\tag{7}$$

Here n_t^S is the number of individuals who are employed at time t and at time $t-1$; n_t is total employment at time t; n_t^L is the number of individuals who were employed at $t-1$ but are not employed at time t; and N_t is population size at time t. The first term on the right-hand side (RHS) is the aggregate labor supply of stayers. The extensive margin is inactive (between times $t-1$ and t) for these individuals, whose behavior is what the micro elasticity captures. The second and third terms represent movements along the extensive margin: people who move into employment (entrants) and out of employment (leavers) in a given period. Finally, the fourth term represents the outsiders.[6] We will later inspect these flows in the data and their correlation with wages to illustrate how specific subgroups contribute to amplifying the aggregate elasticity.

3 Data and aggregation

Our analysis is based on individual-level data from the PSID, 1968–1997 waves. We adopt the following dating convention: a year in the data set is determined by the period a variable refers to, not the period in which the data were collected (the latter is referred to as "wave"). For instance, the 1968 wave collected labor market data referring to year 1967, the 1969 wave collected labor market data referring to year 1968, and so on. Therefore, in our data set time runs from 1967 to 1996. We stop at the 1997 wave because PSID data were collected annually through 1997 but only every other year afterwards. The biennial frequency after the 1997 wave poses the problem of interpolating missing data. In order to avoid arbitrary choices in interpolation, we use only the annual portion of the panel. Work information is available only for the household "head" (defined in the PSID as the adult male householder, if present) and his "wife" (defined as the female cohabitor). Therefore, our observations fall in one of these two categories. This seemingly minor point is actually important for the interpretation of our estimates in relation to the US economy, as we discuss below. Our final sample consists of 24,461 individuals forming an unbalanced panel that spans 30 years. Depending on the year considered, 48–49% of these individuals are males and 51–52% females.

3.1 Sample selection

The construction of the final data set follows Heathcote, Perri, and Violante (2010; hereafter HPV) who use PSID and other data to study the evolution of economic inequality in the United States. In particular, a record is discarded if the implied hourly wage is below 50% of the federal minimum wage in force during a given year. Wages are set as missing if information on hours is missing, and vice versa. Nominal quantities are deflated using the urban Consumer Price Index and are expressed in constant 2009 dollars. Given the different question at hand, our sample selection criteria depart from HPV in two respects. First, they estimate top-coded

[6]The last two terms on the RHS are, of course, sums of zeros.

wage rates by fitting a Pareto distribution. We prefer not to make any distributional assumption and simply discard observations for which wage information is censored. In the PSID, this is amounts to only a handful of observations (on the order of 0.5–0.7% of employed individuals). Second, we use the so-called core sample of the PSID for our baseline analysis. The core sample is composed of two distinct initial samples of households as well as the new households formed out of these. The two samples are a cross-sectional national representative sample (the SRC sample) and a sample of low-income families located in metropolitan areas (the SEO sample). HPV use the SRC group in order to maximize the representativeness of their sample. Yet for us, maximizing sample size to obtain consistent (relative to the underlying population) means of hours and wages is more important than minimizing the bias from a nonrepresentative sample. In principle, representativeness is irrelevant to our undertaking: we seek to estimate and compare the micro and macro elasticities of a given population—not necessarily the actual US population—using observational data. Such an exercise in itself does not require a representative sample. This is our first level of analysis. However, it is of obvious interest to know whether or not our estimates are reliable for calibrating an aggregate model of the US economy. Toward this end, we also conduct our analysis at a second, more representative level. We do this in two alternative ways. First, we drop the SEO sample and use the SRC sample only, as in HPV. Second, we apply sample weights to the core sample. Two important associated caveats should be kept in mind when viewing the resulting estimates as representative. First, as HPV also notice, the SRC sample was necessarily—by construction—representative at the end of the 1960s, but not necessarily so afterwards. Second, sample weights in the PSID were produced to correct for the initial (1967) unequal sampling probabilities between the SRC and SEO samples. Thereafter, sample weights have been regularly revised—to correct for differential attrition and changing demographics—by comparing them with key proportions of the Current Population Survey. Of course, this ensures representativeness only along the dimensions captured by such key proportions. These caveats are illustrated next.

3.2 Definitions and comparison with Census and CPS data

We report in Figure 1 selected demographic characteristics of the (combined) three PSID samples we employ at the two aforementioned levels of analysis (unweighted core sample, SRC sample, and weighted core sample) along with the corresponding characteristics of the US population as derived from Census data.

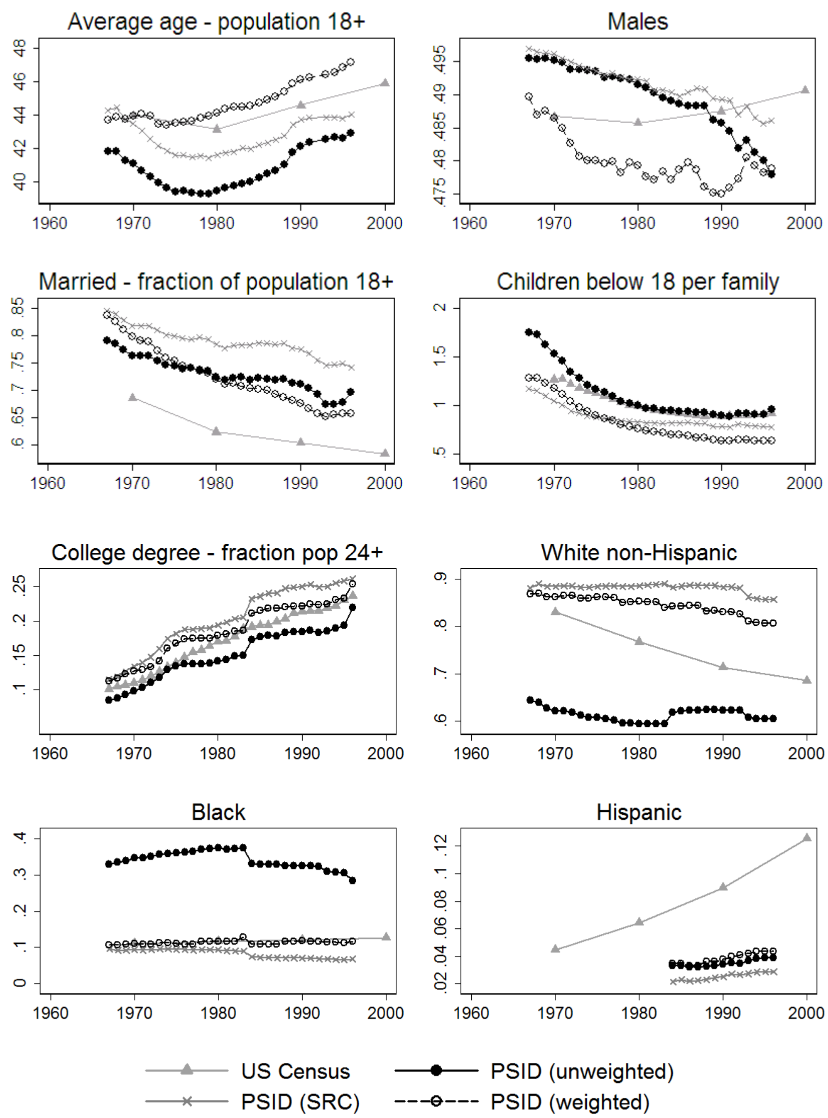

Fig.1. Demographic characteristics in the PSID samples and in the US Census. “Marriage” includes cohabitation in the PSID but refers only to legal marriage in the Census. Hispanics are individuals of any race who report “Hispanic” as ethnic origin. Sources: Bureau of the Census (Current Population Survey) and authors’ calculations from the US Census 1% samples extracted from the IPUMS USA database (Ruggles *et al.*, 2010)

The figure shows that the weighted PSID series, unlike the unweighted one, closely tracks such demographics as age, fraction of college graduates,[7] and (most notably) percentage of Blacks. For other characteristics such as gender and the number of children per family, weighted statistics reproduce national figures well at the beginning of the survey but less so (and in some instances worse than the unweighted sample) during the following years. The most visible discrepancy concerns the fraction of Hispanics and, as a consequence, the fraction of non-Hispanic Whites. The discrepancy for the fraction of adult population that is married is partly explained by the different definitions in the Census and the PSID: contrary to the former, the latter classifies as "married" cohabitating persons of opposite sex, regardless of whether they are legally married or not.[8]

Figure 2 illustrates key labor supply and wage statistics constructed from the CPS and from the PSID samples we use. The variables of interest are annual earnings, annual hours, hourly wages, and the employment rate. Earnings are defined as gross labor income, and hourly wages are constructed as the ratio of annual earnings to annual hours of work for pay.[9] The employment rate is defined in the standard way as the number of workers relative to the population of interest.[10] The figure shows that both the weighted PSID sample and the SRC subsample track the CPS earnings statistics quite accurately in the 1980s except at the two ends of the time interval we consider. In contrast, average annual hours are systematically higher in the CPS than in the PSID, a difference that likely reflects the different wording of the survey questions (hours *usually* worked per week in the CPS versus hours *actually* worked per week in the PSID). As a consequence, there is some discrepancy in the average hourly wage, which is estimated by the ratio of annual earnings to annual hours. The most notable discrepancy between PSID and CPS, however, concerns the adult population employment rate. The reason is that the PSID contains complete labor market information only for household heads and wives—therefore, our analysis has to be based on these individuals—whereas the CPS contains this information for all family members. The higher employment rate in the PSID samples we use implies that household heads and wives are more likely to be employed than other adult family members. In sum, while earnings in the weighted PSID sample are roughly in line with those in the CPS, the former

[7]The survey question about educational attainment is not asked annually to "heads" and "wives". We classify an individual as a college graduate if he or she reports to be a college graduate after age 23 in any wave, and similarly for high school graduates.

[8]Beginning with the 1983 wave, one can distinguish (to some extent) between marriage and cohabitation in the PSID. However, we ignore this information in order to retain comparability between pre- and post-1983 marital status.

[9]This may lead to "division bias" (Borjas, 1980), an issue to which we return later in the paper.

[10]For comparison with the CPS, we compute the employment rate for adult individuals only because very few household heads and wives are under age 18 in the PSID.

overrepresents people with a higher-than-average attachment to the labor market, which may lead to underestimating the extensive margin relative to the US population. It is important to bear this fact in mind when assessing the relation between our estimates and the underlying parameters for the US economy; although the application of sample weights or the use of the SRC group makes the sample more representative than the unweighted core sample, this is no guarantee that the macro elasticity we estimate using the more representative samples yields a better estimate for the US economy. We present estimates for the three different samples but do not take a stand on this issue since our goal is to show how micro and macro elasticities differ despite being generated by the same data.

Although the statistics presented in this Section refer to the 1967–1996 period, in estimation we do not use the 1992–1996 period. The reason, as detailed in the Appendix, are the important changes occurred in the survey after 1991. Such changes make the 1992–1996 portion of the data not reliable for our purposes.

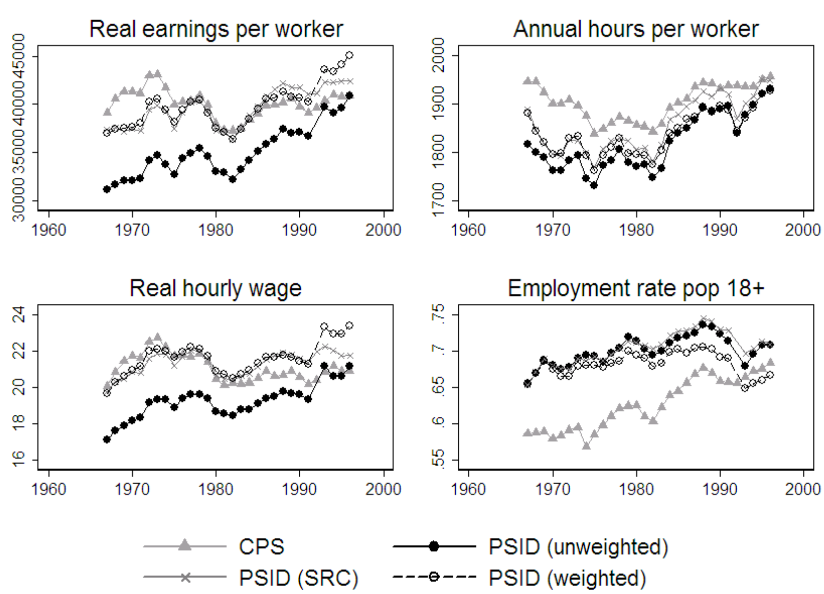

Fig. 2. Earnings, wages, and hours in the PSID and in the CPS. Earnings per worker are gross annual wages and salaries per employed individual. Annual hours are average number of hours worked for pay. The hourly wage is the ratio of earnings per worker to annual hours per worker. Real 2009 US dollars (deflator: CPI-urban).

3.3 Aggregation

The key step in our analysis is aggregation of individual units in the panel to construct the aggregate time series. We take a straightforward approach to aggregation. Since the macro Frisch elasticity refers to the response of aggregate labor, it is natural to aggregate hours by simple summation:

$$H_t = \sum_{i=1}^{n_t} h_{it}, \tag{8}$$

with notation as in Section 2. Since movements in employment, n_t, also reflect variations in sample size, N_t, we will control for such variations as a robustness check. As for the aggregate wage rate, there is no obvious aggregator. The simplest aggregation procedure consists of taking the average, per-worker wage as the relevant aggregate wage rate. That is:

$$W_t = n_t^{-1} \sum_{i=1}^{n_t} w_{it}, \tag{9}$$

where w_{it} is the individual hourly wage at time t. Aggregate W_t is easy to interpret: the wage rate of the representative agent is the average wage. However, heterogeneity suggests weighting individual wages by hours worked. In other words, we might prefer to use the following, alternative aggregate wage rate:

$$W_t' = H_t^{-1} \sum_{i=1}^{n_t} h_{it} w_{it}, \tag{10}$$

which is tantamount to computing the average wage as the ratio of total earnings to total hours (rather than the average ratio of individual earnings to individual hours).

Blundell, Reed, and Stoker (2003) discuss the possible aggregation bias associated with changes in aggregate wage rates expressed as in (9) or (10). There are two sources of such bias that are relevant here: variations in the distribution of hours and variations in the composition of employment. For instance, if some workers work fewer hours at a constant wage rate, then the aggregate wage rate (10) changes, even if the underlying individual wages are unchanged. And if low-wage individuals leave employment while the wage rates of those who remain employed stay the same, then the average wage (9) increases. It is important to make sure that our results are not affected by aggregation bias in wages along these lines. Blundell, Reed, and Stoker use British data from the Family Expenditure Survey to show empirically that predicting wages of nonworkers with the aid of a Heckman-type selection correction procedure—and then taking the average log hourly wage as the measure of the aggregate wage rate—removes virtually all of the aggregation bias affecting aggregates (9) and (10). That is, by employing the log wage rate

$$\ln W_t'' = n_t^{-1} \sum_{i=1}^{n_t} \ln(\widehat{w}_{it}) \tag{11}$$

in the macro regression, where $\widehat{w}_{it}$ is the selectivity-corrected wage rate, Blundell, Reed, and Stoker obtain a time series for aggregate wages that closely tracks the series resulting from correcting W' for aggregation bias. In our analysis we will employ all of these three measures of the aggregate wage—that is, (9), (10), and (11).[11] It turns out that our results are largely insensitive to this choice.

An important advantage of aggregating the individual units is that averaging makes it more likely that measurement errors and other unobservable factors at the micro level will be removed. An example is "division bias", which arises when individual wage rates are estimated as the ratio of earnings to hours (Borjas, 1980). If hours are reported with error, then the estimated individual elasticity is attenuated. Another example is measurement error in reported earnings. Pischke (1995) shows that the dynamics of such error in the PSID cannot be ignored at the individual level. In particular, measurement errors in earnings are serially correlated, which also leads to attenuation bias in panel estimates. Therefore, inferring the intensive margin from aggregated micro data enables us to remove the likely downward bias from the estimated micro elasticity and provide a better estimate of the intensive margin. As we argue later, this is possible if we aggregate only continuously employed workers.

4 Estimation

As anticipated in Section 2, we follow a standard approach in labor economics and estimate the following equations, which we label the micro and macro MaCurdy (1981) equations, respectively:

$$\begin{aligned} \Delta \ln h_{it} &= \kappa + \varepsilon \Delta \ln w_{it} + v_{it}, \qquad (2) \\ \Delta \ln H_t &= K + \mathcal{E} \Delta \ln W_t + V_t, \qquad (3) \end{aligned}$$

with notation as in Section 2. Equation (2) can be derived as a structural equation from a standard life cycle model with time-separable preferences, period utility over consumption c and work h given by $U(c_{it}, h_{it}) = u(c_{it}) - \alpha(1+\varepsilon^{-1})^{-1}(h_{it})^{1+\varepsilon^{-1}}$, and an intertemporal budget constraint. Here ε is the intertemporal Frisch elasticity of labor supply (see Blundell and MaCurdy, 1999). Equation (3), however, has no straightforward structural interpretation in terms of underlying micro pa-

[11] When employing W'' we perform selection correction using Wooldridge's (1995) extension to panel data of the Heckit estimator. See Wooldridge (2010, sec. 19.9.2) for a practical illustration. The variables used to predict employment at the first stage are age, gender, education, race/ethnicity, marital status, and number of children.

rameters. This reflects the well-known fact that, in the presence of heterogeneous wages, generally, individual labor supplies do not aggregate to the labor supply of a representative worker (Muellbauer, 1981). Therefore, we do not attempt here a derivation of $\mathcal{E}$ in terms of ε and other structural parameters. Instead, we limit ourselves to these two observations: (i) $\mathcal{E}$ is a well-defined elasticity of aggregate labor; and (ii) as illustrated in Section 2, there is a statistical relation between $\mathcal{E}$ and ε that has economic meaning in terms of intensive and extensive margins.

The question arises of whether H_t is the appropriate dependent variable in the macro equation. After all, changes in this variable reflect not only changes in sample composition, but also changes in aggregate labor supply. One way of neutralizing the former component is to normalize total hours by sample size. Such a normalization is common in aggregate studies that compare populations whose size changes in space.[12] By analogy, it seems appropriate in our sample because we are comparing waves whose size changes in time because of attrition or other processes like those described in Section 3. Denoting sample size at time t by N_t, we will impose such normalization in a flexible way by including $\Delta \ln N_t$ as an additional RHS variable in equation (3). We will present results from this additional specification as a sensitivity check, since the normalization would apply to the macro equation only and so would violate the micro–macro "isomorphism" that constitutes the key to our empirical exercise.

Equations (2) and (3) are estimated via instrumental variables to account for the endogeneity of wages. Since we implicitly operate in a rational expectations framework, regressor lags are used as instruments. To ensure correspondence between our micro and macro estimates, we use exactly the same instruments in both cases. Thus, the autoregressive terms enter the micro and the macro first-stage equations with exactly the same lags, although aggregation may change the dynamics pertaining to each individual component (Granger and Newbold, 1986). This possibility is another reason to first-difference the data: first differences reduce differences in persistence. We apply the fixed-effects, two-stage least-squares (2SLS) estimator to equation (2). Hence, v_{it} should be interpreted as containing an individual fixed effect, though our notation does not make this explicit. As a consequence, our estimates have the standard interpretation of short-run labor supply responses to productivity shocks. Equation (3) is estimated using the 2SLS estimator. In both cases, the chosen instruments are the second to fifth lags of (respectively) the individual and aggregate wage rates. Instruments are in levels rather than in differences for the efficiency reasons outlined by Arellano (1989).[13]

[12] See, for instance, Rogerson (2006).

[13] With a few exceptions, our results are robust to the chosen number of lags. Inclusions the third, fourth, and fifth lags improves the result of the J-test and so increases the reliability of the instruments.

5 Results

5.1 Full sample

Our main results are reported in Tables 3, 4, and 5. These contain estimates of the individual elasticity (column [1]) and of the aggregate elasticity (columns [2]-[7]) for the unweighted core sample (Table 3), the core sample with application of sample weights (Table 4), and the SRC sample (Table 5). The aggregate elasticity is estimated using three alternative measures of the aggregate wage rate: the log average wage ($\ln W$), the log hours-weighted average wage ($\ln W'$), and the average, selectivity-adjusted log wage ($\ln W''$), as explained in Section 3.3. We also check the sensitivity of the macro estimate to the inclusion of $\Delta \ln (N_t)$ as an additional right-hand-side variable (columns [5]-[7]), as described in Section 4.

Table 3

Micro and macro elasticities: Unweighted PSID core sample

	Individual $\Delta \ln h_{it}$	**Aggregate** $\Delta \ln H_t$					
	[1]	[2]	[3]	[4]	[5]	[6]	[7]
$\Delta \ln (\text{wage})$	0.08*	1.18**	1.20**	1.69**	1.10**	1.01**	1.55**
	(0.049)	(0.49)	(0.47)	(0.62)	(0.44)	(0.39)	(0.54)
$\Delta \ln (N_t)$	—	—	—	—	0.73*	0.82**	0.76**
	—	—	—	—	(0.41)	(0.32)	(0.37)
Constant	– 0.027**	0.02**	0.02**	0.02**	0.01	0.01	0.01
	(0.002)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
J-stat	2.72	1.96	5.18	2.64	5.70	5.89	6.43
p-value	0.44	0.58	0.16	0.45	0.22	0.12	0.17
$\ln (\text{wage})$	$\ln w_i$	$\ln W$	$\ln W'$	$\ln W''$	$\ln W$	$\ln W'$	$\ln W''$
Sample	Core	Core	Core	Core	Core	Core	Core
Weights	No	No	No	No	No	No	No
Obs.	84,985	20	20	20	20	20	20
Individuals	9,161	—	—	—	—	—	—

Note: The sample used is the unweighted core sample (SRC and SEO samples). Column [1]: individual elasticity. Columns [2]–[4]: aggregate elasticity for three alternative measures of the aggregate wage; see equations (9)–(11) in the text for details. Columns [5]–[7]: aggregate elasticity; check of sensitivity to the normalization of aggregate hours by sample size. Standard errors are given in parentheses. * Significant at 10%; ** significant at 5%.

As shown in Tables 3–5, individual-level panel estimates yield a conventionally low individual elasticity of about 0.1.[14] In contrast, the aggregate elasticity from the time series is much larger. Depending on the aggregate wage rate employed, the aggregate elasticity ranges between 1 and 1.7 in the unweighted core sample, between 0.6 an 1.1 in the weighted core sample, and between 0.7 and 1 in the SRC sample.[15] Note that the macro elasticities produced by applying sample weights to the full core sample are virtually identical to those produced by using the SRC, which is consistent with the similarity between the two samples documented in Figures 1 and 2.

Table 4

Micro and macro elasticities: Weighted PSID core sample

	Individual $\Delta \ln h_{it}$	**Aggregate** $\Delta \ln H_t$					
	[1]	[2]	[3]	[4]	[5]	[6]	[7]
$\Delta \ln(\text{wage})$	0.10	0.62**	0.59*	1.10**	0.72**	0.74**	1.12**
	(0.079)	(0.32)	(0.33)	(0.40)	(0.31)	(0.31)	(0.39)
$\Delta \ln(N_t)$	—	—	—	—	0.72	0.72*	0.61
	—	—	—	—	(0.52)	(0.43)	(0.47)
Constant	N/A	0.01**	0.01**	0.01*	0.01	0.01	0.00
	N/A	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)
J-stat	0.28	2.96	5.81	3.81	2.84	6.38	3.63
p-value	0.96	0.40	0.12	0.28	0.42	0.09	0.30
$\ln(\text{wage})$	$\ln w_i$	$\ln W$	$\ln W'$	$\ln W''$	$\ln W$	$\ln W'$	$\ln W''$
Sample	Core	Core	Core	Core	Core	Core	Core
Weights	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	69,429	20	20	20	20	20	20
Individuals	6,955	—	—	—	—	—	—

Note: The sample used is SRC sample. Column [1]: individual elasticity. Columns [2]–[4]: aggregate elasticity for three alternative measures of the aggregate wage; see equations (9)–(11) in the text for details. Columns [5]–[7]: aggregate elasticity; check of sensitivity to the normalization of aggregate hours by sample size. Standard errors are given in parentheses. * Significant at 10%; ** significant at 5%.

[14] Observe that the micro elasticity produced by the weighted regression is similar to the one produced by the unweighted regression but is imprecisely estimated. This confirms that weighting regressions (rather than weighting sample statistics) is not always appropriate and that the unweighted estimate may be preferred (Winship and Radbill, 1994).

[15] Excluded from this range are estimates from specifications in which the J-test rejects the instruments.

Table 5

Micro and macro elasticities: SRC sample

	Individual $\Delta\ln h_{it}$	**Aggregate** $\Delta\ln H_t$					
	[1]	[2]	[3]	[4]	[5]	[6]	[7]
$\Delta\ln(\text{wage})$	0.12**	0.68**	0.49*	1.03**	0.72**	0.54**	1.03**
	(0.056)	(0.32)	(0.29)	(0.38)	(0.31)	(0.26)	(0.36)
$\Delta\ln(N_t)$	—	—	—	—	0.68*	0.67	0.65*
	—	—	—	—	(0.39)	(0.31)	(0.36)
Constant	–0.029**	0.02**	0.02**	0.02**	0.01	0.01**	0.00
	(0.002)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)
J-stat	3.62	4.43	8.59	3.39	4.29	9.20	2.33
p-value	0.31	0.22	0.03	0.34	0.23	0.03	0.51
$\ln(\text{wage})$	$\ln w_i$	$\ln W$	$\ln W'$	$\ln W''$	$\ln W$	$\ln W'$	$\ln W''$
Sample	SRC	SRC	SRC	SRC	SRC	SRC	SRC
Weights	—	—	—	—	—	—	—
Obs.	52,920	20	20	20	20	20	20
Individuals	5,390	—	—	—	—	—	—

Note: The sample used is the core sample (SRC and SEO samples), weighted using sample weights Column [1]: individual elasticity. Columns [2]–[4]: aggregate elasticity for three alternative measures of the aggregate wage; see equations (9)–(11) in the text for details. Columns [5]–[7]: aggregate elasticity; check of sensitivity to the normalization of aggregate hours by sample size. Standard errors are given in parentheses. * Significant at 10%; ** significant at 5%.

Tables 3–5 show that the elasticity of aggregate labor supply greatly exceeds the corresponding individual elasticity in *exactly* the same sample. Hence there is no conflict between a small micro elasticity and a large macro one. In essence, the micro estimate reflects the choice of hours by individuals who are employed in any two adjacent years, whereas the macro estimate includes movements into and out of employment in response to aggregate wage shocks.

The point estimate of the aggregate elasticity in Tables 4 and 5 match the final number resulting from the massive review of evidence on labor supply elasticities summarized by Chetty, Guren, Manoli, and Weber (2011). Their conclusion is that "it would be reasonable to calibrate representative agent macro models to match a Frisch elasticity of aggregate hours of 0.75" (p. 474). Our own analysis indicates that this value is reasonable but conservative: 0.75 is an elasticity towards the bottom of the range we identify using sample weights and the SRC sample.

Furthermore, as discussed in Section 3.2, this range likely bounds below the elasticity of the actual US economy, because the PSID sample of household heads and wives likely underestimates the extensive margin. We next show that the point of departure from Chetty, Guren, Manoli, and Weber is the estimate of the extensive margin elasticity.

5.2 Continuously employed workers: the pure intensive margin

The crucial difference between intensive and extensive margin elasticities can be further appreciated by estimating the macro intensive margin separately. To do this, we consider only individuals who work every year between 1967 and 1991. We call these "continuously employed" workers. By estimating the aggregate elasticity of those individuals in our sample who always work—for whom the extensive margin is therefore inactive—we obtain the pure macro intensive margin, which is the direct counterpart of the micro elasticity from an aggregate perspective. The comparison between the micro and macro elasticity of continuously employed individuals is of special interest because aggregation of this group removes both the effect of employment changes on total hours (i.e., the extensive margin) *and* its compositional effect on the average wage and hours per worker (i.e., part of the aggregation bias). For this exercise we use the unweighted core sample only: the SRC sample would be too small to construct an aggregate for such a special subgroup, and sample weights can be applied reliably only to the full core sample, not parts of it. There are 759 continuously employed individuals in our data. The results are reported in Table 6.

The micro elasticity of these workers turns out to be essentially 0, albeit imprecisely estimated. The macro elasticity obtained from their aggregation—our estimate of the pure macro intensive margin—ranges between 0.3 and 0.4. Why are these two estimates different? One possibility is heterogeneity. Recall from equation (6) in Section 2 that, since the extensive margin is now shut off completely, these two estimates (the micro and the macro intensive margins) would be the same if all these continuously employed individuals worked the same number of hours at the same hourly wage in each cross section. This is not the case, of course. However, the use of $\ln W''$ as a measure of the aggregate wage should remove the residual aggregation bias deriving from such heterogeneity. Another, more relevant possibility is measurement error. We have explained above how inaccuracies in reported earnings and hours lead to "division bias" when estimating the wage rate as the ratio of the two. We also mentioned that these inaccuracies are known to be serially correlated in the PSID. It is well known that such measurement errors lead to identification failure. In the univariate context we consider here, they unambiguously lead to attenuation bias in the micro elasticity—the parameters in the micro regression do have a standard structural interpretation.[16] An important advantage of our aggregation procedure is that these measurement errors (and, possibly, other unobservable factors) at the micro level are removed from the variables used in the aggregate regression.

Table 6
Micro and macro elasticities: continuously employed individuals

	Individual	**Aggregate**		
	$\Delta \ln h_{it}$	$\Delta \ln H_t$		
	[1]	[2]	[3]	[4]
$\Delta \ln(\text{wage})$	–0.02	0.30*	0.42**	0.28*
	(0.071)	(0.17)	(0.20)	(0.17)
Constant	–0.006**	–0.006**	–007**	–0.005*
	(0.003)	(0.003)	(0.003)	(0.003)
J-stat	5.52	3.39	2.87	1.66
p-value	0.14	0.34	0.41	0.65
$\ln(\text{wage})$	$\ln w_i$	$\ln W$	$\ln W'$	$\ln W''$
Obs.	15,180	20	20	20
Individuals	759	—	—	—

Note: The sample used are all individuals in the core sample (SRC and SEO samples) who are continuously employed between 1967 and 1991 (i.e., who report a strictly positive number of hours of work every year). Column [1]: individual elasticity. Columns [2]–[4]: aggregate elasticity for three alternative measures of the aggregate wage; see equations (9)–(11) for details. Standard errors are given in parentheses. * Significant at 10%; ** significant at 5% or better.

A macro intensive margin of continuously employed individuals that is larger than the corresponding micro intensive margin is consistent with the presumption that the latter is subject to a downward bias; this has been shown repeatedly in the literature (see, e.g., Rupert, Rogerson, and Wright; 2000; Pistaferri, 2003; Domeij and Flodén, 2006; Chetty, Friedman, Olsen, and Pistaferri, 2011, and Wallenius, 2011). Our estimate of the intensive margin from aggregated PSID (i.e., 0.3–0.4) is in line with the intensive margin estimate from quasi-experimental data reported by Chetty, Friedman, Manoli, and Weber (2011)—i.e., 0.3–0.5. Yet even if our intensive margin elasticity is larger when estimated at the aggregate level, one still needs a substantial extensive margin to generate macro elasticities of the magnitude

[16] The serial correlation of measurement errors in earnings (and hourly wages) makes our instruments of no help in correcting such bias in the micro regression.

reported in Tables 3–5. Comparing Tables 5 and 7, such extensive margin elasticity ranges from 0.78 to 1.41 in the PSID. The corresponding estimate reported by Chetty, Friedman, Manoli, and Weber ranges from 0.17 to 0.28.

5.3 Subgroups

Because the extensive margin accounts for the discrepancy between micro and macro estimates in our data, it is of interest to know which workers tend to move into and out of employment in response to aggregate wage changes. We can exploit that our time series is derived from micro data and estimate the aggregate elasticity for different demographic and socioeconomic groups of interest. That is, we construct separate time series for these groups and repeat the main empirical exercise.[17] Table 7 shows the results. To conserve space, we report only the estimated coefficient and use a single asterisk to indicate that the estimate is significantly different from zero at the 10% level or better. For this exercise we use the same three measures of the aggregate wage rate used before but with the unweighted core sample only—for the same reasons we used this sample only when analyzing continuously employed individuals.

The table reveals that, in our sample, the micro elasticity of men is virtually zero. In contrast, the micro elasticity of women is positive and significantly different from zero (although still small, 0.34), in line with the findings of Mroz (1987). Prime-age women (i.e., those of age 25–54) have a smaller micro elasticity than women in other age groups. However, we find no consistent evidence that the elasticity of women is also higher at the aggregate level; point estimates are generally lower than those for men and imprecise. The individual elasticities of the remaining groups are very small and always statistically indistinguishable from zero, but the aggregate elasticity is relatively large for individuals not of prime age,[18] married individuals whose partner is employed, and low-educated individuals. In particular, the aggregate elasticity we estimate decreases sharply with education: our point estimate is quite large for individuals who did not complete high school and is negative for college graduates.

[17]When estimating the aggregate elasticity for a subgroup, the relevant aggregate wage rate is the average real hourly wage of individuals belonging to that subgroup.

[18]This is consistent with evidence from other sources. For example, Gourio and Noual (2009) and Wallenius (2011) find that young workers have a higher elasticity than the rest of the population, and French (2005) finds that the labor supply of senior workers is more elastic than average.

Table 7
Micro and macro elasticities: Subgroups

Group	Number of individuals	Individual elasticity
All individuals	9,161	0.08*
Men	5,189	–0.01
Women	3,972	0.34*
Prime-age (25–54) men	4,654	0.01
Prime-age (25–54) women	3,538	0.39*
Non–prime-age men	1,453	–0.07
Non–prime-age women	1,125	0.68
Married, partner works	6,523	0.10
Married, partner does not work	2,721	–0.12
College degree	1,780	0.01
High school degree	4,023	0.05
Less than high school	2,357	0.12

Group	Aggregate elasticity					
All individuals	1.18*	1.20*	1.69*	1.10*	1.01*	1.55*
Men	1.11*	1.40*	1.59*	0.95*	1.13*	1.40*
Women	1.23	4.44	0.36	0.30	0.73	0.42
Prime-age (25–54) men	0.55*	0.69*	0.29	0.65*	0.85*	0.71*
Prime-age (25–54) women	0.05	0.42	0.46	–0.06	0.27	0.41
Non–prime-age men	1.48	1.93*	1.71*	1.16*	1.30*	1.46*
Non–prime-age women	1.43	1.50*	0.83	1.06	1.09	0.93
Married, partner works	2.94	2.96*	2.26*	2.05	0.68	1.22
Married, partner does not work	0.08	0.33	0.35	0.69*	0.74*	0.65
College degree	–1.62*	–1.89*	–1.49	–0.26	–0.67	–0.33
High school degree	0.44	1.52	0.71	1.08	1.90	0.77*
Less than high school	3.04	1.72	–2.13	–0.12	1.95	1.15
Wage	$\ln W$	$\ln W'$	$\ln W''$	$\ln W$	$\ln W'$	$\ln W''$
$\Delta \ln (N_t)$ included	No	No	No	Yes	Yes	Yes

Note: Elasticities for subgroups of the core sample for three alternative measures of the aggregate wage; see equations (9)–(11) for details. The first column in the top panel (number of individuals) contains the number of workers in each subgroup for the 1967-1991 period. The totals of mutually exclusive categories whose composition changes over time may exceed the number of workers. * Significant at 10% or better.

The higher aggregate labor elasticity of these groups can also be seen in the raw data. Figure 3 illustrates the flows that make up the RHS of identity (7) for the following groups: in the upper panel, men versus women and prime-age (PA) versus non–prime-age (Non-PA); in the lower panel, married with working partner (PW) versus married with nonworking partner (PNW) and college graduates (College) versus individuals who never attended college (No college). The series are expressed as shares of the respective sample size every year, so they add up to 1 conditional on a subgroup. The upper panel of Figure 3 shows that, as one would expect, men and prime-age individuals are overwhelmingly employed for any pair of adjacent years (stayers)—much more so than women and non–prime-age individuals. Correspondingly, the flows of men and prime-age individuals into employment (entrants) and out of employment (leavers) are low relative to the other two groups. The figure's lower panel shows that married individuals with a working partner and individuals who did not attend college generate larger flows into and out of employment (although not after the 1970s for the first group) than do married individuals with a partner who does not work and college graduates.

In sum, the entrants and leavers groups are both dominated by women, individuals below 25 or above 54, low-educated individuals, and (until the end of the 1970s) married individuals with working partners. For the groups we are considering, Figure 4 shows the net flows (i.e., the difference between entry and exit each year) in terms of percentage-point deviations from the HP trend (e.g., 0.02 on the vertical scale corresponds to 2 percentage points when looking at *net flows*). Net entry/exit flows are important because they correspond to the net effect of the extensive margin on aggregate labor supply. We superimpose on Figure 4 the aggregate wage rate in percentage deviation from its HP trend (e.g., 0.02 on the vertical scale corresponds to 2% when looking at the *wage rate*). Table 8 reports the volatility of the net flows series relative to trend as well as their correlation with the deviations of the wage series. Figure 4 and Table 8 together reveal two facts. First, deviations from trend of net flows and wages are positively correlated. This is what one expects if labor supply is elastic along the extensive margin. Second, the deviations for women, non–prime-age individuals, married individuals with working partners, and individuals without a college degree tend to be both larger and more strongly correlated with wage deviations than those for other groups. As suggested before, these subgroups can be regarded as the marginal workers underlying the large aggregate elasticity.

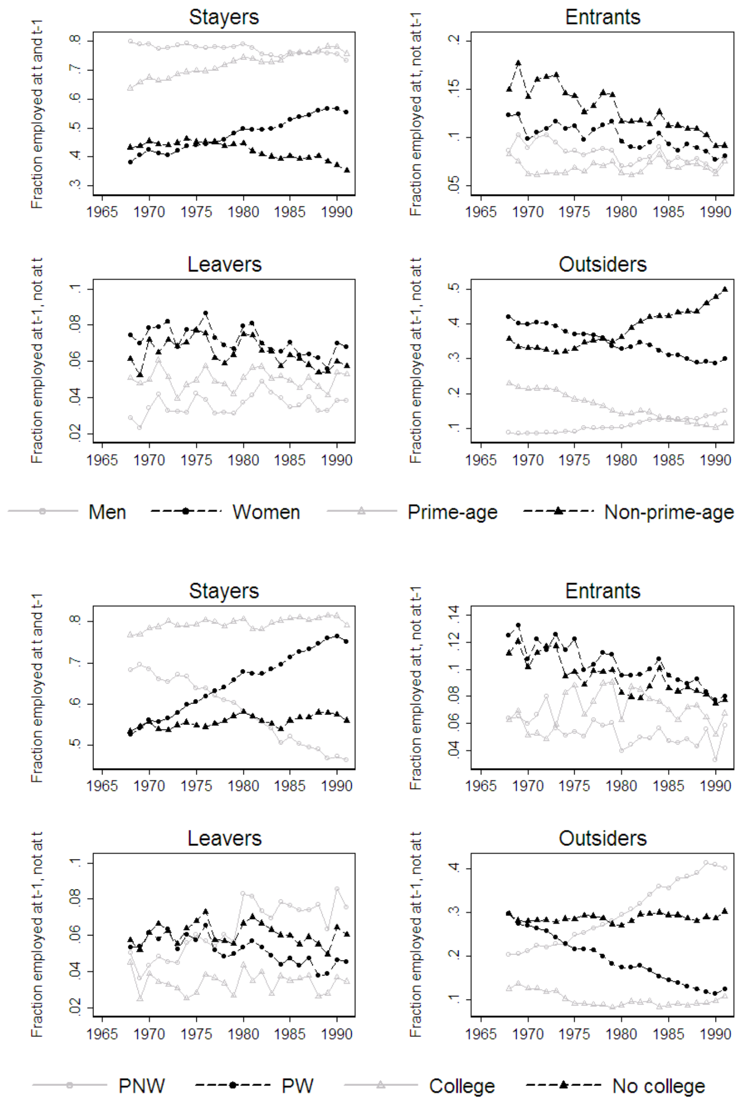

Fig 3. Employment flows. The figure illustrates the incidence of the four groups that constitute the RHS of identity (7) as a percentage of the total for various subgroups. PNW = legally married or cohabitating individuals whose partner is not working; PW = legally married or cohabitating individuals whose partner is working.

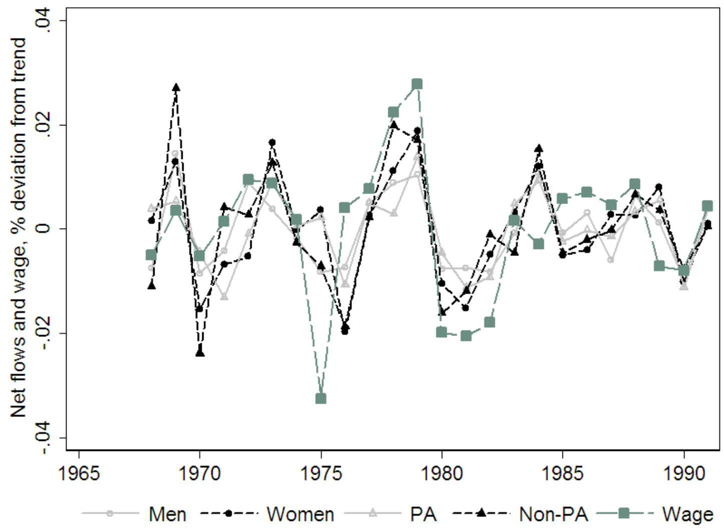

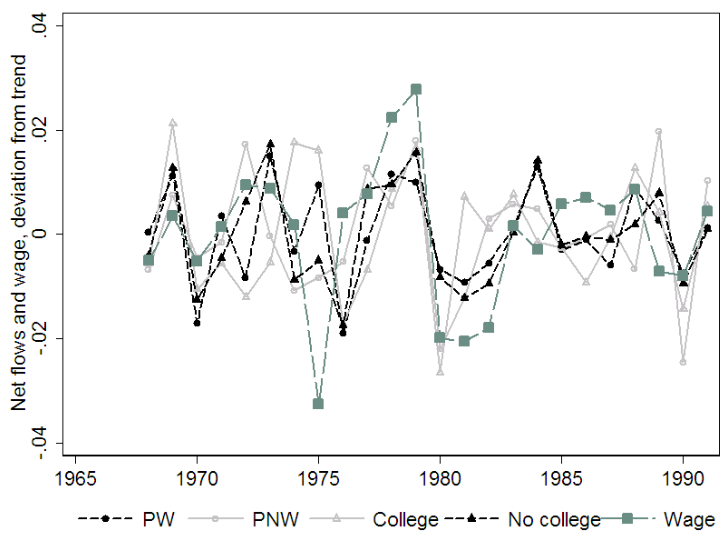

Fig 4. Cyclical behavior of net flows and average wage. The figures illustrate the percentage deviations of net flows (percentage of entrants minus percentage of leavers) of various subgroups from the respective HP trends (smoothing parameter 6.25) as well as the deviation of the average unweighted wage from its own HP trend (smoothing parameter 6.25). PA = prime age, NPA = non–prime age; PNW = partner not working, PW = partner working.

Table 8
Volatility of net flows and correlation with wages

Group	Std. dev.	Correlation
Men	0.007	0.68
Women	0.010	0.44
Prime-age	0.007	0.45
Non–prime-age	0.012	0.55
Married, partner works	0.009	0.28
Married, partner does not work	0.011	0.53
College degree	0.012	0.14
No College	0.010	0.57

Note: The table reports the standard deviation of the percentage deviation of net flows from the HP trend (smoothing parameter 6.25) for various subgroups as well as the correlation between such percentage deviations and the corresponding deviation of the average wage.

6 Conclusions

This paper provides an empirical reconciliation of micro and macro elasticities of labor supply. A conceptual reconciliation has long been offered (see Heckman, 1993) that is based on well-known differences between intensive and extensive margin adjustments over the business cycle. We believe our analysis illustrates this fact empirically in a clean and straightforward way. We estimate MaCurdy (1981) equations for exactly the same samples at two different levels of aggregation and find that aggregation alone leads to a much larger Frisch elasticity than that found in the corresponding panel estimate. The latter ranges between 0.08 and 0.12 whereas the former ranges between 0.6 and 1.7, depending on the sample used and on the way hours and wages are aggregated. These numbers imply a ratio of macro to micro estimates ranging between 6 and 20. There is no conflict between these two estimates, which refer to the same generating data and derive from two "isomorphic" micro–macro specifications. When we aggregate only continuously employed individuals to correct the micro estimates for measurement errors and other individual-level unobservables, the intensive margin elasticity increases by a factor of about 3; the extensive elasticity still explains 60–80% of the micro–macro gap in the baseline estimates.

We did not attempt a structural interpretation of the macro regression in terms of micro parameters. This is an important task for future research, which several authors have already undertaken (e.g., Chang and Kim, 2006; Gourio and Noual, 2009; Prescott, Rogerson and Wallenius, 2009; Rogerson and Wallenius, 2009; Ljungqvist and Sargent, 2011). Instead, we confined ourselves to what we regard as a profound methodological point for macroeconomics and for the RBC approach in particular. Namely, parameter estimates from micro data are not always appropriate for calibrating an aggregate model economy. This caution certainly applies to labor supply: micro and macro labor supply elasticities differ significantly at different levels of aggregation.

APPENDIX

The reason why we only use waves 1968–1992 only is the substantial number of changes that occurred in the survey design starting with the 1993 wave. In the 1993 wave (and in this wave only) hourly wages are available only for workers paid by the hour. For those not paid by the hour, we need to estimate the wage rate directly using the ratio of annual labor earnings to annual hours; in all the other waves, this estimate is produced by the PSID staff, who also corrects errors and imputes missing values. Although we applied the same selection criteria described in Section 3, the result is an artificial upward jump in the wage rate for year 1992 that is evident in Figure 2. Furthermore, from the 1994 wave onward, sample size underwent substantial variations. Figure 5 illustrates the number of individuals in our sample (sample adult population) and the employment level; the two vertical lines mark years 1968 and 1993, respectively. With the exception of 1967 (which is not used for estimation), the dynamics of adult population in the PSID looks like a normal demographic process. This happens endogenously because when the offspring of a PSID household forms his or her own family, this new unit is included in the survey. However, the dynamics takes a strange appearance after the 1993 wave. First, the 1994 wave included nearly a thousand "new" households that had been lost to attrition during the previous 25 years but had been searched for and brought back into the survey that year. For this reason (the presence of relatively many people with low attachment to the survey) we observe that the attrition rate during the following two years is higher than the rate at which new households are formed from existing PSID households. Finally, in 1997 the core sample was substantially reduced owing to a general restructuring of the survey. These changes induce variations in aggregate labor supply that have nothing to do with responses to wage changes. In light of these issues, we decided to discard the five waves from 1993 through 1997—i.e. years 1992-1996.

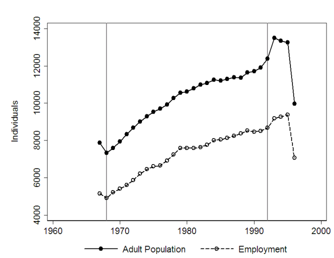

Fig. 5. Sample size and number of workers in the PSID. The solid line is the number of adult individuals (household heads and wives) present in the survey; the dashed line is the number of such adult individuals who are employed.

REFERENCES

Alogoskoufis, G., 1987. On intertemporal substitution and aggregate labor supply. *Journal of Political Economy* 95 (5), 938-960.

Altonji, J., 1986. The intertemporal substitution in labor supply: evidence from micro data. *Journal of Political Economy* 94 (3), S176-S215.

Arellano, M., 1989. A note on the Anderson-Hsiao estimator for panel data. *Economics Letters* 31 (4), 337-341.

Blundell, R., Bozio, A., Laroque, G., 2011. Extensive and intensive margins of labour supply: working hours in the US, UK and France. *IFS Working Paper* 01/11.

Blundell, R., MaCurdy, T., 1999. Labor supply: a review of alternative approaches. In Ashenfelter, O., Card, D. (editors), *Handbook of Labor Economics*, vol. 3, Elsevier, 1559-1695.

Blundell, R., Reed, H., Stoker, T., 2003. Interpreting aggregate wage growth: the role of labor market participation. *American Economic Review* 93 (4), 1114-1131.

Borjas, G., 1980. The relationship between wages and weekly hours of work: the role of division bias. *Journal of Human Resources* 15 (3), 409-423.

Browning, M., Hansen L., Heckman, J., 1999. Micro data and general equilibrium models. In Taylor, J., Woodford, M. (editors), *Handbook of Macroeconomics*, vol. 1, Elsevier, 103-204.

Card, D., 1994. Intertemporal labor supply: an assessment. In Sims, C. (editor), *Advances in Econometrics, Sixth World Congress*. Cambridge University Press.

Chang, Y., Kim, S., 2006. From individual to aggregate labor supply: a quantitative analysis on a heterogeneous agent macroeconomy. *International Economic Review* 47 (1), 1-27.

Chetty, R., Friedman, J., Olsen, T., Pistaferri, L., 2011. Adjustment costs, firm responses, and labor supply elasticities: evidence from Danish tax records. *Quarterly Journal of Economics* 126 (2), 749-804.

Chetty, R., Guren, A., Manoli, D., Weber, A., 2011. Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and

Extensive Margins. *American Economic Review: Papers & Proceedings* 101 (3), 471-475.

Domeij, D., Flodén, M., 2006. The labor-supply elasticity and borrowing constraints: why estimates are biased. *Review of Economic Dynamics* 9 (2), 242-262.

Fiorito, R., Kollintzas, T., 1994. Stylized facts of business cycles in the G7 from a real business cycles perspective. *European Economic Review* 38 (2), 235-269.

French, E., 2005. The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies*, 72 (2), 395-427.

Gourio F., Noual, P., 2009. The marginal worker and the aggregate elasticity of labor supply. *Working paper*, Boston University.

Granger, C., Newbold, P., 1986. *Forecasting Economic Time Series*. Second edition, Academic Press.

Hall, R., 1980. Labor supply and aggregate fluctuations. *Carnegie-Rochester Conference Series on Public Policy* 12 (1), 7-33.

Hansen, G., 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16 (3), 309-327.

Heathcote, J., Perri, F., Violante, G. L., 2010. Unequal we stand: an empirical analysis of economic inequality in the United States: 1967-2006. *Review of Economic Dynamics* 13 (1), 15-51.

Heckman, J., 1993. What has been learned about labor supply in the past twenty years?. *American Economic Review* 83 (2), 116-121.

Imai, S., Keane, M., 2004. Intertemporal labor supply and human capital accumulation. *International Economic Review* 45 (2), 601-641.

Killingsworth, M., Heckman, J., 1986. Female labor supply: a survey. In Ashenfelter, O., Layard, R. (editors), *Handbook of Labor Economics*, vol. 1, Elsevier, 103-204.

Kydland, F., 1995. Business cycles and aggregate labor market fluctuations. In Cooley, T. (editor), *Frontiers of Business Cycle Research*, Princeton University Press, 126-156.

Kydland, F., Prescott, E., 1982. Time to build and aggregate fluctuations. *Econometrica* 50 (6), 1345-1370.

Ljungqvist, L., Sargent T., 2011. A labor supply elasticity accord? *American Economic Review: Papers & Proceedings* 101 (3), 487-491.

Lucas, R., 1980. Methods and Problems in Business Cycle Theory. *Journal of Money, Credit and Banking* 12 (4), 696-715.

Lucas, R., Rapping, L., 1969. Real wages, employment and inflation. *Journal of Political Economy* 77 (5), 721-754.

MaCurdy, T., 1981. An empirical model of labor supply in a life-cycle setting. *Journal of Political Economy* 89 (6), 1059-1085.

Mankiw, N. G., 1989. Real business cycles: a new Keynesian perspective. *Journal of Economic Perspectives*, 3 (3) 79-90.

Mankiw, N. G., Rotemberg, J., Summers, L., 1985. Intertemporal substitution in macroeconomics. *Quarterly Journal of Economics* 100 (1), 225-251.

Mroz, T., 1987. The sensitivity of an empirical model of married women's hours of work to economic and statistical assumptions. *Econometrica* 55 (4), 765-799.

Muellbauer, J., 1981. Linear aggregation in neoclassical labour supply. *The Review of Economic Studies* 48 (1), 21-36.

Mulligan, C., 1998. Substitution over time: another look at life-cycle labor-supply. *NBER Macroeconomics Annual*, 75-134.

Pencavel , J., 1986. Labor supply of men: a survey. In Ashenfelter, O., Layard, R. (editors), *Handbook of Labor Economics*, vol. 1, Elsevier, 3-102.

Pischke, J., 1995. Measurement error and earnings dynamics: some estimates from the PSID validation study. *Journal of Business &Economic Statistics* 13 (3), 305-314.

Pistaferri, L., 2003. Anticipated and unanticipated wage changes, wage risk, and intertemporal labor supply. *Journal of Labor Economics* 21 (3), 729-754.

Prescott, E., 1986. Theory ahead of business cycle measurement. *Federal Reserve Bank of Minneapolist Quarterly Review* 104, 9-22.

Prescott, E., 2006. Nobel Lecture: The transformation of macroeconomic policy and research. *Journal of Political Economy* 114 (2), 203-235.

Prescott, E., Rogerson, R., and Wallenius, J., 2009. Lifetime aggregate labor supply with endogenous workweek length. *Review of Economic Dynamics* 12 (1), 23-36.

Rogerson, R., 1988. Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21 (1), 3-16.

Rogerson, R., 2006. Understanding differences in hours worked. *Review of Economic Dynamics*, 9, 365-409.

Rogerson, R., Wallenius, J., 2009. Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory* 144 (6), 2277-2292.

Ruggles, S., Alexander, J. T., Genadek, K., Goeken, R., Schroeder, M. B., Sobek, M., 2010. *Integrated Public Use Microdata Series*: Version 5.0 [Machine-readable database]. Minneapolis: University of Minnesota.

Rupert, P., Rogerson, R., Wright, R., 2000. Homework in Labor Economics: Household Production and Intertemporal Substitution. *Journal of Monetary Economics* 46 (3), 557-579.

Summers, L., 1986. Some skeptical observation on real business cycle theory. *Federal Reserve Bank of Minneapolis Quarterly Review* 10, 23-27.

Wallenius, J., 2011. Human capital accumulation and the intertemporal elasticity of substitution of labor: How large is the bias? *Review of Economic Dynamics* 14 (4) 577-591.

Winship, C., Radbill, L., 1994. Sampling weights and regression analysis. *Sociological Methods & Research* 23 (2), 230-257.

Wooldridge J., 1995. Selection corrections for panel data models under conditional mean independence assumptions. *Journal of Econometrics* 68 (1), 115–132.

Wooldridge J., 2010. *Econometric analysis of cross section and panel data*, 2nd edition. MIT Press.

Ziliak J. Kniesner, T., 2005. The effect of income taxation on consumption and labor supply. *Journal of Labor Economics* 23(4), 769-796.