The Scandinavian
Journal of Economics

Check for updates

Scand. J. of Economics 126(1), 194–217, 2024
DOI: 10.1111/sjoe.12544

Frisch elasticities in a model of indivisible labor supply with endogenous workweek length

Minchul Yum†

University of Southampton, Southampton, SO17 1BJ, UK
m.yum@soton.ac.uk

Abstract

In this paper, I provide an extension of the classical indivisible labor supply model where a large macro Frisch elasticity is reconciled with a small micro counterpart. Households take as given state-dependent hours per worker – shaped by a nonlinear mapping from hours worked to labor services and employment frictions – and make intertemporal labor supply decisions. In the standard indivisible labor supply model, aggregate fluctuations are independent of the individual preference parameter that governs the intensive-margin elasticity. In my model, however, they are connected through the extensive margin whose elasticity is empirically reasonable and is shaped by the individual preference parameter.

Keywords: Indivisible labor; intensive margin; extensive margin; Frisch elasticity; labor supply
JEL classification: E32; J22

1. Introduction

Models of indivisible labor supply following Rogerson (1988) can reconcile small micro-based individual labor supply elasticities with large aggregate counterparts. This feature is important as it can generate the large volatility of aggregate hours which we observe in the data, while being consistent with smaller estimates of individuals' willingness to substitute labor over time (Keane and Rogerson, 2012). Despite this merit, in pure indivisible labor supply models, aggregate fluctuations (or the macro elasticity) are disconnected from the preference parameter governing the micro elasticity.[1]

In this paper, I present an extension of the classical indivisible labor supply model that circumvents this disconnect by allowing equilibrium hours

†Also affiliated with CEPR.

[1]This is the case not only in stand-in household models (Hansen, 1985) but also in heterogeneous-agent models with incomplete asset markets, pioneered by Chang and Kim (2006), in which aggregate fluctuations depend on the individual distribution but are still independent of the preference parameter governing the underlying individuals' Frisch elasticity.

per worker to be state-dependent. Through the exercises in the spirit of Rogerson and Wallenius (2009) who provide a reconciliation on the micro versus macro steady-state (or long-run) labor supply elasticities, my model reconciles a large macro Frisch (or short-run) labor supply elasticity with a small micro counterpart. Importantly, this reconciliation is achieved in a framework where the aggregate labor supply elasticity depends on the individual preference parameter that shapes the curvature of the utility function along hours worked. This happens through the novel feature of the model that a higher micro Frisch intensive-margin elasticity can raise the volatility of aggregate hours along both intensive and extensive margins.[2]

In the model, the firm solves a dynamic problem in the presence of a nonlinear labor services mapping – which captures set-up costs and fatigue effects (Prescott et al., 2009) – and employment frictions (Hall, 2004). The firm's optimization problem gives rise to the state-dependent workweek length, which is then taken as given by households. I embed this setting into an otherwise standard real business-cycle model. I follow a standard procedure to calibrate the model in steady state except for the degree of the employment adjustment cost in the baseline specification, which is calibrated to match the cyclicality of employment over the business cycle.

I first evaluate the calibrated model using a set of conventional business-cycle statistics. My model outperforms the real business-cycle model with rich household heterogeneity in Chang et al. (2019) in terms of the cyclical volatility of labor markets along different margins.[3] Moreover, the volatility of employment increases with the individuals' willingness to substitute labor supply over time in my model, in contrast to Chang et al. (2019) where this relationship is qualitatively opposite.

I then use the model economy to quantify the relationship between individuals' intensive-margin elasticity and its model-implied aggregate elasticities as in Rogerson and Wallenius (2009). I find that the estimated extensive-margin Frisch elasticity tends to increase with the intensive margin elasticity and is indeed quite sizable, broadly in line with the recent empirical evidence (Fiorito and Zanella, 2012; Peterman, 2016). Therefore, the aggregate labor supply elasticity – the sum of the intensive-margin elasticity and the extensive-margin elasticity by definition

[2]For example, this is in contrast to the model in Chang et al. (2019) where the two margins are essentially substitutes. Specifically, Table 1 reproduces their business-cycle results where a higher individual Frisch elasticity raises the cyclical volatility of hours at the intensive margin while reducing the counterpart at the extensive margin.

[3]Earlier real business-cycle models with both intensive and extensive margins include Kydland and Prescott (1991), Bils and Cho (1994), Cho and Cooley (1994), and Osuna and Ríos-Rull (2003), among others.

(Chetty et al., 2013) – is considerably larger (ranging from 1.4 to 2.8) than the micro elasticity at the individual level (varied from 0.5 to 1.5) in the baseline model.

Rogerson and Wallenius (2009) focus on the micro versus macro steady-state labor supply elasticities, which govern labor responses with respect to permanent tax changes, in a life-cycle model without aggregate uncertainty. My paper focuses on a related yet different object (i.e., the intertemporal elasticity or Frisch elasticity) in a business-cycle environment with aggregate uncertainty. Although both models indicate the importance of the extensive margin, their model's prediction on the magnitude of the extensive-margin elasticity as a function of that of the intensive-margin elasticity is different. Specifically, my model implies that the extensive-margin elasticity generally increases with the individual's intensive-margin elasticity, whereas the extensive-margin elasticity implied by the Rogerson–Wallenius model is essentially unrelated with or slightly decreases with the individual's intensive-margin elasticity.

Erosa et al. (2016) build a rich heterogeneous household environment with both extensive and intensive margins of labor supply, and compute the aggregate labor supply elasticities with respect to transitory and permanent wage changes. While their framework emphasizes life-cycle aspects in a partial equilibrium setting, my paper emphasizes labor supply changes over the business cycle, driven by aggregate productivity shocks in general equilibrium.[4] Nevertheless, their finding that the aggregate labor supply elasticity with respect to temporary wage changes is 1.75, of which 62 percent is due to the extensive margin, is in line with my main results with a moderate size of employment adjustment costs.

The remainder of this paper is organized as follows. In Section 2, I introduce the firm's optimization problem in the presence of the nonlinear mapping in a simple static environment, and then present the main dynamic general equilibrium business-cycle model. In Section 3, I conduct the main quantitative analysis. I first discuss how the model is calibrated and solved, and then present the main quantitative results about its business-cycle performance and the relationship between the preference parameter and the implied aggregate labor supply elasticities. I conclude in Section 4.

[4]Heterogeneous-agent models with both extensive and (divisible) intensive margins are widely used in the quantitative macroeconomics literature. For example, such frameworks have been used to study welfare implications of social security (Kitao, 2014; Peterman and Sommer, 2019), sources of female labor supply changes over time (Jang and Yum, 2022), and welfare state reforms (Guner et al., 2023). Likewise, these papers do not consider stochastic equilibrium with aggregate uncertainty and do not explore how Frisch elasticities are related at the micro and macro levels.

2. Model

I first introduce the nonlinear labor services mapping in a simple setting, and then embed this technological setting into a standard real business-cycle model.

2.1. Optimal workweek length

A firm faces a continuum of households with measure one and maximizes profit by choosing both the employment level n and the schedule of hours for each worker $h(i)$.[5] The production function f has a set of usual properties such as $f(0) = 0$, $f'(\cdot) > 0$, and $f''(\cdot) < 0$. Taking as given the productivity z and hourly wage w, the firm solves

$$\max_{h(i),n\in[0,1]} zf(L) - w\left(\int_0^n h(i)di\right), \tag{1}$$

where L denotes the effective total labor input, $L = \int_0^n g(h(i))di$. The key element is a nonlinear labor services mapping $g(\cdot)$ (Prescott et al., 2009):

$$\begin{aligned} g(h) &= 0 \text{ for } h \in [0, \phi] \\ &= \tilde{g}(\cdot) \text{ for } h \in [\phi, 1], \end{aligned} \tag{2}$$

where $\tilde{g}(\phi) = 0$, $\tilde{g}'(\cdot) > 0$, and $\tilde{g}''(\cdot) < 0$, as depicted in Figure 1. This nonlinear mapping reflects two important features in the relation between actual hours spent and effective labor input: the marginal returns are zero for the first several hours because of set-up costs, and then are decreasing because of fatigue effects (Prescott et al., 2009). In contrast to the linear mapping, the nonlinear mapping leads to the two theoretical properties of the labor demand decision, both of which characterize the optimal length of workweek. The first property is given by Lemma 1.

Lemma 1. *Assume that $g(\cdot)$ is nonlinear and satisfies equation (2). Then the firm, which solves equation (1), optimally chooses the same hours $h \in [\phi, 1]$ for all identical workers.*

First, note that $h(i) \in [0, \phi]$ will never be chosen, as it would give zero marginal services (as well as zero services) while incurring positive marginal costs w. Second, if one compares a constant schedule of hours to any other schedules having the same $\int_0^n h(i)di$, it is always the case that the former (the identical hours) yields higher aggregate labor services (i.e., $\int_0^n g(h(i))di$) than the latter (Jensen's inequality). However, if the g function

[5] The key results in this section do not change when I add capital as another production input. The business-cycle model for the quantitative analysis indeed incorporates capital as well.

Figure 1. Nonlinear labor services mapping

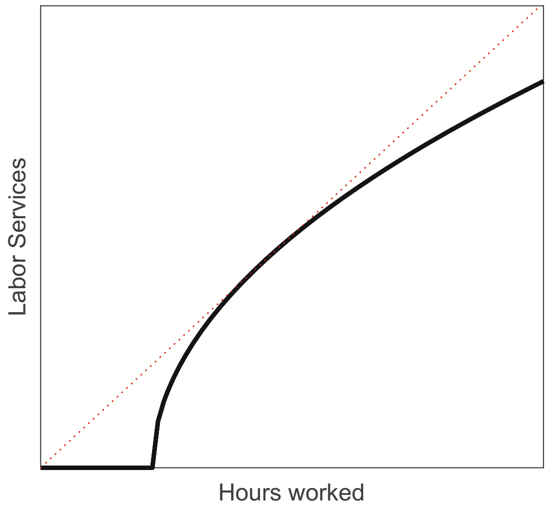

is linear, the firm would be indifferent between any schedules of hours for workers as long as the total labor input L is chosen optimally.

With the identical choice on h, the effective total labor input can be simply expressed as $L = g(h)n$. The firm's profit maximization problem can then be reduced to

$$\max_{h,n\in[0,1]} zf(g(h)n) - whn. \tag{3}$$

The first-order conditions for h and n are

$$zf'(L)g'(h)n = wn,$$
$$zf'(L)g(h) = wh,$$

implying that the optimal $\overline{h}$ is determined by

$$g'(\overline{h}) = \frac{g(\overline{h})}{\overline{h}}, \tag{4}$$

independent of other economic factors such as productivity z and market wage w.[6]

[6]Card (1990) also derives a similar independence result using the same effective total labor input, which incorporates the nonlinear labor services mapping. Prescott et al. (2009) obtain the same independence result in an environment where the same functional form of the nonlinear labor services mapping is embedded in household's labor supply.

Lemma 2. *Assuming that $g(\cdot)$ satisfies equation (2), the schedule of optimal hours reduces to $\overline{h}$, which is determined solely by the function $g(\cdot)$ according to equation (4).*

In sum, the firm would adjust its scale (employment level), while holding workweek hours at this optimal level, independent of other economic conditions such as z. This simple theoretical result echoes the exogenous workweek length of Rogerson (1988).

In a richer dynamic set-up with aggregate uncertainty, this simple theoretical result can be changed. One way pursued in this paper is to introduce quasi-fixity of labor. With this friction, hours per worker are no longer independently determined by the labor services mapping but can vary depending on the aggregate states (e.g., z in the above problem). Next, I describe the environments of the main dynamic stochastic general equilibrium model.

2.2. Equilibrium business-cycle model

I now present the full model economy. This dynamic model will provide the firm with incentives to adjust both hours and employment over the business cycle in the presence of employment frictions. Specifically, I assume that the employment level is predetermined for the next period before a next period productivity shock is realized (Burnside et al., 1993; Shimer, 2010), and that employment adjustment is subject to adjustment costs. The firm discounts future profits using the prices of assets held by households and also perceives the aggregate laws of motion because the firm is a price-taker.

The production function $f(L, k)$ is assumed to be Cobb–Douglas,

$$Y = zf(L, k) = zL^{\alpha}k^{1-\alpha},$$

where $\alpha \in (0, 1)$. The process of the total factor productivity shock z originally follows an AR(1) process in logs,

$$\log z_t = \rho \log z_{t-1} + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma_\epsilon^2),$$

and the expositions henceforth use the corresponding discretized N_z-state Markov chain.

The firm's dynamic problem is given by

$$v(n, z_i, s) = \max_{\substack{h,n' \in [0,1] \\ k \geq 0}} \left\{ z_i f(L, k) - w(z_i, s)hn - r(z_i, s)k - \Phi(n, n') + \sum_{j=1}^{N_z} q_j(z_i, s)v(n', z_j, s') \right\}, \tag{5}$$

subject to

$$L = g(h)n, \tag{6}$$

$$N' = M_1(z_i, s) \text{ and } K' = M_2(z_i, s), \tag{7}$$

where z_i is the today's total factor productivity shock and $s \equiv (N, K)$ denotes endogenous aggregate state variables, which consist of the aggregate employment level and aggregate capital, respectively. Agents take as given the last two perceived laws of motions for the endogenous aggregate state variables, N and K. The variables with a prime denote their values in the next period. $\Phi(n, n')$ is the convex adjustment cost such that $\Phi(n, n) = 0$ for all n.

The nonlinear mapping $g(h)$ is assumed to be

$$\begin{aligned} g(h) &= (h - \phi)^{\eta} \quad \text{if } h \in (\phi, 1] \\ &= 0 \quad \text{if } h \in [0, \phi], \end{aligned} \tag{8}$$

where $\phi \in [0, 1)$ captures the range of unproductive hours at the workplace, and $\eta \in (0, 1]$ captures the extent to which fatigue effects operate for long hours.

The economy is populated by a continuum of *ex ante* identical and infinitely lived households on the unit interval. Households have access to a complete set of Arrow securities. The period utility function for each household is given by

$$u(c_t, h_t) = \log c_t - \theta \frac{h_t^{1+(1/\gamma)}}{1 + (1/\gamma)},$$

which implies that individuals' intensive-margin labor supply elasticity is equal to γ. As in Rogerson (1988), households can choose either 0 or $\overline{h}_t$, which is taken as given. As shown in Appendix A, aggregation gives rise to the stand-in household's utility function,

$$U(c_t, n_t) = \log c_t - B(\overline{h}(z_i, s))n_t,$$

where

$$B(\overline{h}(z_i, s)) \equiv \theta \frac{\overline{h}_t^{1+(1/\gamma)}}{1 + (1/\gamma)}. \tag{9}$$

Note that the stand-in household takes as given $\overline{h}_t$, which may be state-dependent.[7] As can be seen above, the stand-in household's utility is linear in n_t, provided that $\overline{h}_t$ is taken as given. When $\overline{h}_t$ is exogenously

[7]In a pure indivisible labor model, this assumption is unnecessary but could have been innocuously imposed, as the workweek is assumed to be fixed in any case.

fixed as in the pure indivisible model (i.e., $\overline{h}_t = \overline{h}$), aggregate fluctuations are independent of the value that we assign to the individual labor supply elasticity. This is because the marginal disutility of employment for the stand-in household, B, is an invariant constant not only in steady state but also over the business cycles.[8]

It is important to note that this disconnect between the individual's parameter and aggregate fluctuations is not due to the indivisibility of labor per se, but the exogenously fixed level of hours. It is easy to see from equation (9) that when $\overline{h}_t$ varies, a different value of γ could matter for the stand-in household through the marginal disutility of employment $B(\overline{h}(z_i, s))$ that changes over the business cycle. For instance, when the stand-in household faces a high $\overline{h}_t$, the marginal disutility of increasing the fraction of workers becomes higher, affecting the optimal labor supply at the extensive margin n_t indirectly.

The stand-in household's dynamic optimization problem can be written as the following functional equation:

$$W(a, k, z_i, s) = \max_{\substack{c \geq 0, k' \in \Gamma(k) \\ n \in [0,1]}} \left\{ \log c - B(\overline{h}(z_i, s))n + \beta \sum_{j=1}^{N_z} \pi_{ij} W(a'_j, k', z_j, s') \right\}, \tag{10}$$

subject to

$$c + \sum_{j=1}^{N_z} q_j(z_i, s)a'_j = a + w(z_i, s)\overline{h}(z_i, s)n + r(z_i, s)k + \Pi(z_i, s) + (1 - \delta)k, \tag{11}$$

$$N' = M_1(z_i, s) \quad \text{and} \quad K' = M_2(z_i, s). \tag{12}$$

Here, π_{ij} denotes the transition probability $\Pr(z' = z_j | z = z_i)$, β is the discount factor, a_j is the Arrow security that pays 1 in state j, q_j is its price, and $\Pi(z_i, s)$ denotes dividends.

In equilibrium, the firm's choice on workweek length should be consistent with the stand-in household's optimality conditions, given prices. The equilibrium definition is standard and is provided in Appendix A.

2.3. Analytical optimality conditions

In this subsection, I derive some analytical results.

[8]Recall that the only thing required to change for different γ is θ, which yields the same B, in order to match the same steady-state total hours in the calibration step.

The stand-in household's first optimality condition is the labor–leisure condition:

$$\frac{1}{c}w(z_i,s)\overline{h}(z_i,s) = B(\overline{h}(z_i,s)). \tag{13}$$

Next, the Euler equation for consumption (or capital)

$$\frac{1}{c} = \beta\sum_{j=1}^{N_z}\pi_{ij}\frac{1+r(z_j,s')-\delta}{c'_j}, \tag{14}$$

and the optimal portfolio choices satisfy

$$q_j(z_i,s) = \beta\pi_{ij}\frac{c}{c'_j}, \tag{15}$$

both of which are standard in RBC models.

As in Khan and Thomas (2003), using equation (15) and denoting by $p(z_i,s)$ the marginal utility of consumption, the firm's functional equation (5) can be rewritten as

$$V(n,z_i,s) = \max_{\substack{h,n'\in[0,1]\\ k\geq 0}}\Bigg\{p(z_i,s)\left[z_i(g(h)n)^{\alpha}k^{1-\alpha} - w(z_i,s)hn - r(z_i,s)k\right.$$
$$\left. -\ \Phi(n,n')\right] + \beta\sum_{j=1}^{N_z}\pi_{ij}V(n',z_j,s')\Bigg\},$$

subject to

$$N' = M_1(z_i,s) \text{ and } K' = M_2(z_i,s),$$

where the firm discounts future profits by β. From this functional equation, if we use the first-order conditions for the static choice variables (i.e., h and k), we can obtain the following optimality conditions with some algebra:

$$k(n,z_i,s) = z_i^{1/[\alpha(1-\eta)]}\left(\frac{\alpha\eta}{w(z_i,s)}\right)^{\eta/(1-\eta)}\left(\frac{(1-\alpha)}{r(z_i,s)}\right)^{(1-\alpha\eta)/\alpha(1-\eta)} n, \tag{16}$$

$$h(z_i,s) = \phi + z_i^{1/[\alpha(1-\eta)]}\left(\frac{\alpha\eta}{w(z_i,s)}\right)^{1/(1-\eta)}\left(\frac{(1-\alpha)}{r(z_i,s)}\right)^{(1-\alpha)/\alpha(1-\eta)}. \tag{17}$$

First, the capital demand is proportional to the firm's current employment level as the firm would need a larger capital stock when more workers are employed, while the demand schedule of hours is independent of the firm's current employment level. With the employment frictions, the schedule of hours that is optimally chosen by the firm would exhibit independence only of its own employment level. Hours per worker now respond to the aggregate state variables, and thus can vary as overall economic conditions change. As the employment level of the firm is predetermined, it might not be at the optimal level after the productivity shock is observed. Thus, the firm now has an incentive to deviate from the optimal workweek length, characterized by equation (4) in Lemma 2.

As the firm's decision on the employment level is dynamic in the presence of the frictions, we have an intertemporal optimality condition for the employment level,

$$\Phi_2(n, n') = \beta \sum_{j=1}^{N_z} q_j(z_i, s) \left[z_j f_1(L', k') g(h) - w'(z_j, s') h' - \Phi_1(n', n'') \right], \tag{18}$$

where the double prime indicates their values two periods later. The left-hand side is the immediate marginal cost due to hiring costs or layoff costs when the firm plans to adjust its employment level next period. This immediate marginal cost must be equal to the expected discounted sum of marginal product of employment net of the two extra terms. The first extra term is the marginal cost of employment that increases the wage bill in the next period. Because the next period employment level becomes a state variable for the next period decision, the marginal reduction in adjustment costs next period should be also accounted for, which is reflected in the last term.

3. Quantitative analysis

3.1. Calibration and solution method

I now explain how parameter values are chosen for the following quantitative exercises. The length of a period corresponds to a quarter. The first set of parameter values is chosen using the steady-state equilibrium (Cooley and Prescott, 1995). Specifically, these parameter values are calibrated so that the model in steady state is consistent with the long-run averages of the US data from 1956Q1–2010Q4. To begin, imposing that all of the endogenous variables are constant over time without shocks, equations (13)–(18) characterize the analytical relationship between the variables in steady state. Then, these relations are used to map from the first moments in the data to the parameter values. The quarterly real interest rate of 1 percent gives $\beta = 0.99$. Next, the long-run investment–capital ratio implies the value of δ, and the long-run average capital–output ratio implies α. I choose $\delta = 0.025$ and $\alpha = 1 - 0.36$. These values are commonly used in the equilibrium business-cycle literature, including Kydland and Prescott (1991), Cho and Cooley (1994), and Chang et al. (2019), all of which build a model with both intensive and extensive margins of labor.

I now discuss two parameters in the nonlinear mapping specified in equation (8). Burda et al. (2020) estimate that the average fraction of time at the workplace that employees are not working is 6.9 percent. Accordingly, the baseline value of ϕ is chosen somewhat lower at 5 percent, and the value of ϕ equal to 10 percent of the steady-state hours is also considered as a sensitivity check, as reported in Appendix D. The average fraction of working hours of 39.4/84 in a week pins down $\eta = 0.95$, assuming that the weekly endowment of available hours for work and leisure is 84 hours.

I set $\rho = 0.95$ and $\sigma_\epsilon = 0.007$, commonly used values in the literature. In particular, these values are the same as those used in Chang et al. (2019) whose model outcomes regarding business-cycle properties will be compared to their counterparts from my model.

A variation of γ is necessary for the main exercises to investigate the mapping between individuals and aggregates (Rogerson and Wallenius, 2009). I choose $\gamma = 0.5$, 1.0, and 1.5. Given a value of γ, θ is recalibrated to match the long-run employment–population ratio of 59.6 percent in the US data: θ is equal to 42.0, 13.2, and 8.5 for $\gamma = 0.5$, 1, and 1.5, respectively.

The employment adjustment costs are assumed be quadratic,

$$\Phi(n, n') = \frac{\xi}{2}\left(\frac{n' - n}{n}\right)^2,$$

where $\xi \geq 0$ determines the degree of the employment adjustment cost. A special feature of this parameter is that, in the model, ξ does not affect steady-state prices and quantities, thereby requiring another approach rather than the traditional approach based on steady state.[9] As a higher value of ξ would weaken the link between employment and output over the business cycle, this parameter is calibrated to match the cyclical correlation between employment and output (0.80) in the data. This leads to $\hat{\xi} = 0.040$ for $\gamma = 1.0$, the mid-value of the range considered in this paper.[10] As in Kydland and Prescott (1991), I then consider two alternative values of ξ. The economy with a low ξ (i.e., $\hat{\xi}$ divided by ten) would make the model behave like a pure indivisible labor model, whereas a large ξ (i.e., $\hat{\xi}$ multiplied by ten) would make the model behave like a divisible labor model.

To obtain the equilibrium business-cycle data from the model with aggregate uncertainty, the model is solved numerically. Although an easiest

[9] In general, adjustment costs are known to be hard to estimate. For example, Hall (2004) shows that the estimates of the annual degree of quadratic labor adjustment costs for various industries are quite small with large standard errors.

[10] Alternatively, I also considered recalibrating the adjustment cost parameter for each γ (i.e., $\hat{\xi}_\gamma$). This complicates the interpretation of the following exercises in this paper because a change in γ would also involve a change in ξ, making it difficult to isolate the clean effects of γ.

way might seem to solve the corresponding planning problem, note that households in this economy would have an incentive to affect hours if they were able to do so. Thus, the social planner's problem would yield different allocations than the decentralized equilibrium, except for a special case where the individual supply elasticity is exactly equal to the aggregate elasticity, as shown in Appendix B.

As for computing decentralized equilibria, in a setting where either households or firms have a static problem, we can embed the optimal choices of the static agent into the other agent's dynamic problem by substituting out market prices, and iterate a single value function without considering prices, as in Hansen and Prescott (1995). However, this method cannot be straightforwardly applied here because both the stand-in household and firm face dynamic problems. Therefore, I solve the decentralized equilibrium directly using a nonlinear method for the equilibrium value functions of both agents. In essence, the algorithm iteratively finds the equilibrium laws of motions that are equal to agents' perceived aggregate laws of motion, which are necessary to infer correct prices.

3.2. Business-cycle results

I first present some key business-cycle statistics from the model-generated data. As is standard in the business-cycle literature, statistics in this subsection are based on model-generated data over long (10,000) periods, the first 1,000 periods of which are dropped. The logged variables are detrended using the HP filter with the smoothing parameter equal to 1,600. US data counterparts are computed using the aggregate data from 1956Q1 to 2010Q4 after applying the same procedures.

Table 1 summarizes cyclical volatilities, or percentage standard deviations, of the key macroeconomic variables relative to output. In the baseline specification (Panel A), we see systematic relationships between the individual Frisch elasticity γ and cyclical volatilities in contrast to the pure indivisible labor economy: that is, a higher γ increases the cyclical volatilities of aggregate variables (especially labor market variables). When each individual is more willing to substitute labor intertemporally, the stand-in household that represents those individuals is more likely to accommodate the firm's need to deviate from the optimal workweek hours in the absence of aggregate shocks.

In Panel D, I reproduce the cyclical properties of the model in Chang et al. (2019) who consider the same exercise on the aggregate labor market implications of changing γ, with the same type and size of aggregate shocks. First, it is clear that my baseline model tends to generate the volatility of aggregate labor market variables that is noticeably larger than that from their model along both intensive and extensive margins (and thus in terms of total hours as well). Moreover, in their model, we see that individuals' Frisch

Table 1. Cyclical volatilities relative to output, σ_X/σ_Y

ξ	γ	σ_Y	x					AC/Y
			C	I	h	N	$h \times N$	
US data		1.56	0.60	2.54	0.35	0.64	0.91	
Panel A. Baseline ($\hat{\xi}$)	0.5	1.46	0.36	2.97	0.14	0.51	0.54	1.4e-7
	1.0	1.57	0.35	3.01	0.21	0.54	0.58	1.9e-7
	1.5	1.63	0.34	3.04	0.26	0.55	0.60	2.3e-7
Panel B. Low ($\hat{\xi} \div 10$)	0.5	1.66	0.34	3.05	0.11	0.67	0.66	5.3e-8
	1.0	1.71	0.34	3.07	0.17	0.68	0.67	6.6e-8
	1.5	1.74	0.33	3.08	0.21	0.68	0.67	7.2e-8
Panel C. High ($\hat{\xi} \times 10$)	0.5	1.28	0.34	3.06	0.21	0.25	0.34	2.2e-7
	1.0	1.41	0.33	3.07	0.31	0.29	0.43	3.4e-7
	1.5	1.49	0.33	3.10	0.37	0.30	0.49	4.2e-7
Panel D. CKKR	0.5	1.57	–	–	0.06	0.29	0.35	–
	1.0	1.62	–	–	0.10	0.28	0.37	–
	1.5	1.67	–	–	0.13	0.26	0.38	–

Notes: Numbers are percentage standard deviations of HP-filtered data. The last column (AC/Y) reports the average adjustment costs relative to output. CKKR denotes Chang et al. (2019).

elasticity parameter raises the volatility of hours per worker at the expense of less volatile employment fluctuations. In my model, γ also increases the volatility of employment over the business cycle. This feature is explored further in the next subsection.

Panel B of Table 1 shows that labor adjustment takes more along the extensive margin with low adjustment costs. To understand this, Figure 2 plots the equilibrium decisions by the stand-in household and firm. The right panels show the employment decision for the next period as a function of the current employment level and the productivity z. It clearly shows that a lower ξ enables the firm to rely more heavily on the extensive margin with respect to a higher z. Consequently, the relative volatility of the extensive margin is much larger, resembling the model properties of the pure indivisible labor model. Next, imposing a large degree of adjustment costs, the model should behave like a divisible labor model, as it causes adjustment of labor to occur more along the intensive margin. As shown in Panel C of Table 1, the cyclical volatilities exhibit a well-documented weakness of the divisible labor model. Even a high labor supply elasticity of 1.5 cannot generate substantially higher volatility of total hours, which leads to weaker amplification of the productivity shocks.

Table 2 reports the cyclicality of aggregate variables. When it comes to correlations with output, the most noticeable fact in the data is that the intensive margin of labor (h) is procyclical but less so ($Cor(h, Y) = 0.71$)

Figure 2. Equilibrium decision rules, by the size of employment adjustment costs

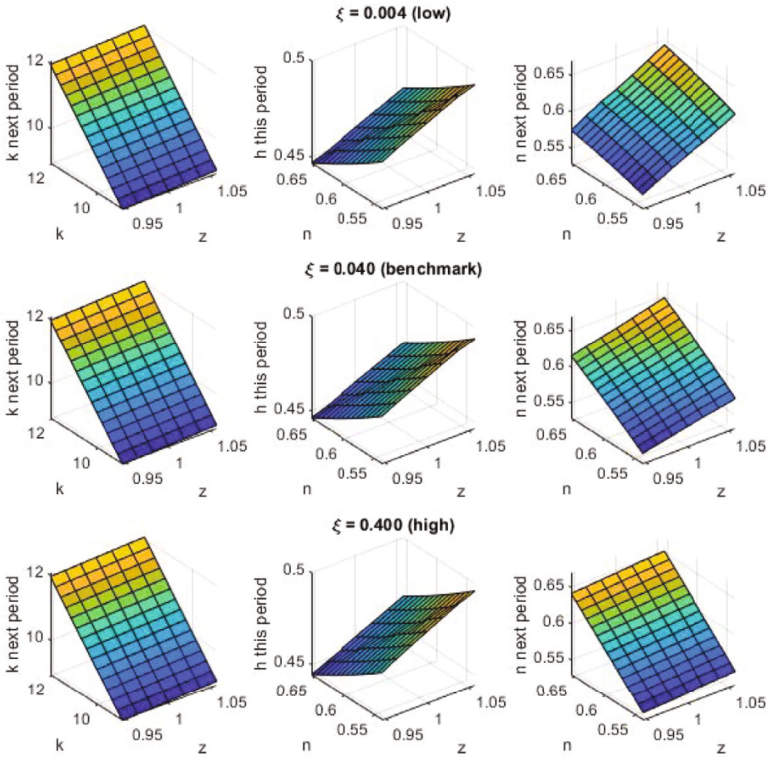

Notes: The left panels show the household's equilibrium decision rule for k' when N is at the steady-state level and $K = k$. The middle and right panels show the firm's equilibrium decision rules for h and N', respectively, when K is at the steady-state level and $n = N$. All figures are from the model with $\gamma = 1.0$.

Table 2. Cyclicality of aggregates: $Cor(x, Y)$

ξ	γ	x				
		C	I	h	N	$h \times N$
US data		0.84	0.92	0.71	0.80	0.84
Panel A. Baseline ($\hat{\xi}$)	0.5	0.93	0.99	0.60	0.79	0.92
	1.0	0.92	0.99	0.56	0.80	0.95
	1.5	0.92	0.99	0.55	0.79	0.97
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.91	0.99	0.32	0.88	0.94
	1.0	0.91	0.99	0.31	0.87	0.95
	1.5	0.91	0.99	0.32	0.85	0.96
Panel C. High ($\hat{\xi} \times 10$)	0.5	0.91	0.99	0.86	0.53	0.94
	1.0	0.91	0.99	0.84	0.57	0.98
	1.5	0.90	0.99	0.83	0.58	0.99

Table 3. Persistence of aggregates: ρ_x

ξ	γ	x					
		Y	C	I	h	N	$h \times N$
US data		0.85	0.85	0.89	0.55	0.91	0.86
Panel A. Baseline ($\hat{\xi}$)	0.5	0.82	0.81	0.83	0.54	0.92	0.93
	1.0	0.82	0.81	0.83	0.48	0.90	0.91
	1.5	0.81	0.81	0.81	0.45	0.90	0.90
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.85	0.82	0.86	0.22	0.83	0.87
	1.0	0.84	0.82	0.84	0.15	0.80	0.87
	1.5	0.83	0.82	0.83	0.11	0.79	0.87
Paenl C. High ($\hat{\xi} \times 10$)	0.5	0.76	0.81	0.75	0.68	0.95	0.89
	1.0	0.76	0.81	0.76	0.67	0.95	0.85
	1.5	0.76	0.82	0.75	0.65	0.94	0.83

compared to the extensive margin ($Cor(N, Y) = 0.80$). Panel A shows that the baseline model replicates this pattern successfully. The other panels also show that this relative magnitude of cyclicality between the two margins is largely shaped by the adjustment costs. Specifically, as ξ increases, which would become similar to a divisible labor model, the intensive margin becomes more procyclical than the extensive margin. Interestingly, the effect of γ on the cyclicality of aggregate variables is found to be quite limited.

Next, as can be seen in Table 3, the model can reproduce the persistence of labor along the two margins remarkably well in the baseline model. Specifically, in both model-generated data and the US data, the intensive margin is quite persistent but is less persistent than the extensive margin, resulting in a very high (but lower than N) persistence of total hours. Figure 3 shows the impulse response functions, which can help us understand these results. In the middle and right panels, I show how equilibrium labor along each margin moves over time from the steady state after the economy is hit by an adverse aggregate productivity shock (-1 percent). The response of the intensive margin is quick and temporary, whereas the equilibrium employment response is sluggish and persistent. Note also that the individual labor supply elasticity γ governs the magnitude of equilibrium labor responses at both intensive and extensive margins, which is consistent with the key cyclical volatility results reported in Table 1.

3.3. Aggregate labor supply elasticities

Given the finding that the individual preference parameter γ systematically shapes aggregate labor market fluctuations in the model economy, in this subsection, I quantify this relationship more carefully. Specifically, using

Figure 3. Impulse responses: labor along intensive and extensive margins

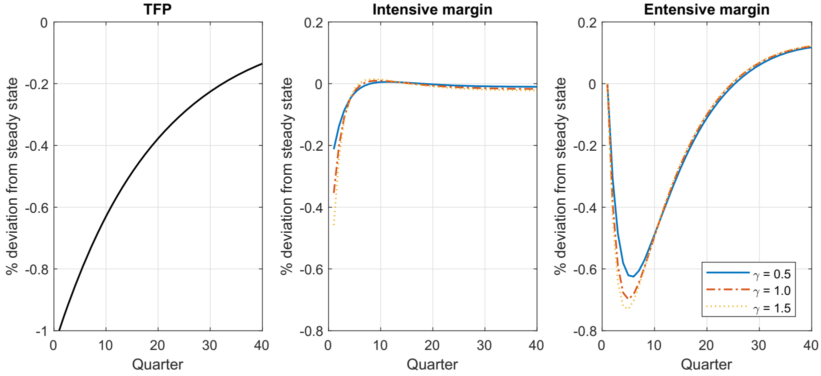

Notes: These figures are obtained from specifications with the benchmark adjustment cost ($\hat{\xi} = 0.04$).

time-series aggregate data generated with different combinations of γ and ξ, as in the previous subsection, I estimate the following equations separately:

$$\log h_t = \alpha_0^h + \alpha_1^h \log w_t + \alpha_2^h C_t + \varepsilon_t^h; \tag{19}$$

$$\log N_t = \alpha_0^N + \alpha_1^N \log w_t + \alpha_2^N C_t + \varepsilon_t^N; \tag{20}$$

$$\log H_t = \alpha_0^H + \alpha_1^H \log w_t + \alpha_2^H C_t + \varepsilon_t^H. \tag{21}$$

These provide the estimates of Frisch elasticity for the intensive margin (α_1^h), the extensive margin (α_1^N), and aggregate hours (α_1^H).[11] Several points are worth noting here. First, I control for consumption because the parameter of interest is Frisch elasticity, which holds marginal utility constant. Second, the above equations can identify labor supply elasticities from my simulated data because the only exogenous shock in my model is the total factor productivity, which shifts labor demands. Finally, it is not actually necessary to estimate the last equation as α_1^H (aggregate) should be equal to the sum of α_1^h and α_1^N (recall $H \equiv h \times N$).

Table 4 shows that the preference parameter γ, which governs the intensive-margin Frisch elasticity of households, is precisely recovered in

[11] According to the macroeconomics literature, γ is typically called the micro labor supply elasticity and α_1^H corresponds to the macro labor supply elasticity (Keane and Rogerson, 2012). In contrast, Chetty et al. (2013) define micro versus macro labor supply elasticities based on the source of data. According to their terminology, α_1^h, α_1^N, and α_1^H are macro elasticities at different margins.

Table 4. Frisch labor supply elasticities

ξ	γ	Intensive margin elasticity, $\hat{\alpha}_1^h$	Extensive margin elasticity, $\hat{\alpha}_1^N$	Aggregate labor supply elasticity, $\hat{\alpha}_1^H$
Panel A. Baseline ($\hat{\xi}$)	0.5	0.50	0.94	1.44
	1.0	1.00	1.13	2.13
	1.5	1.50	1.28	2.78
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.50	0.61	1.11
	1.0	1.00	0.53	1.53
	1.5	1.50	0.50	2.00
Panel C. High ($\hat{\xi} \times 10$)	0.5	0.50	0.53	1.03
	1.0	1.00	0.76	1.76
	1.5	1.50	0.95	2.45

all cases. Although the model does not explicitly allow households to choose desired hours worked due to the indivisibility, the stand-in household's labor supply decision implicitly takes into account the underlying households' desire to substitute labor supply intertemporally, as is evident from equation (9) and its surrounding discussions in Section 2.2.

Table 4 also shows that aggregate labor supply elasticities are substantially larger than the assumed individual intensive-margin elasticities due to the extensive margin (Keane and Rogerson, 2012). Quantitatively, these model-implied extensive-margin elasticities are broadly in line with the recent empirical findings on the extensive-margin Frisch elasticity. For instance, the estimates of Fiorito and Zanella (2012) range between 0.8 and 1.4, and Peterman (2016) finds that contribution of the extensive margin to the aggregate labor supply elasticity is around 0.6–0.7.[12] Notably, the extensive-margin elasticity increases with the individual's intensive-margin elasticity γ, provided that ξ is not counterfactually too low. This result suggests that the individual's preference parameter, γ, governing labor supply elasticity along the intensive margin could also be an important determinant of the extensive margin elasticity.

Finally, we can see that the aggregate labor supply elasticity – the sum of the intensive-margin and extensive-margin elasticities – therefore increases with γ in all cases, showing that the disconnect in pure indivisible models is eliminated. This is due to both the direct effect of the assumed γ and the

[12] These also align with Erosa et al. (2016), who find that the aggregate labor supply elasticity with respect to temporary wage changes is 1.75, of which 62 percent is due to the extensive margin. Using the Survey of Income and Program Participation, Kimmel and Kniesner (1998) find that the extensive-margin elasticity varies from 0.6 (for single men) to 2.4 (for single women).

indirect effect through the extensive-margin elasticity, which is shown to be a function of γ as well.

4. Conclusion

In this paper, I consider an extension of the canonical business-cycle model of indivisible labor supply in which workweek length changes over the business cycle endogenously. In contrast to pure indivisible labor models, this model relates the individual intensive-margin Frisch elasticity to aggregate fluctuations, while maintaining the merit of the pure indivisible labor model that reconciles large aggregate labor supply elasticities with smaller individual labor supply elasticities. This difference is captured by sizable extensive-margin elasticities that also tend to be shaped by the individuals' preference parameter governing the intensive-margin elasticity.

The results make it clear that the reason for the disconnect in pure indivisible labor models is not the indivisibility of labor per se, but the exogenously fixed intensive margin, which makes the variation of the aggregate hours occur only through changes along the extensive margin. The model presented herein also conceptualizes the idea that the extensive-margin elasticity could also reflect the underlying individuals' willingness to substitute labor supply over time at the intensive margin. It should be noted that the model presented in this paper is not meant to serve as a benchmark model that can be used for serious quantitative analyses such as counterfactual and policy analyses. Rather, the model is deliberately kept as simple as possible to deliver the above key messages. Hence, it leaves scope for further research with various elements known to be relevant for the labor market, such as search frictions, rich heterogeneity, and incomplete markets for those interested in quantitative business-cycle studies where those labor supply elasticities are critical.

Acknowledgments

This paper is based on a chapter in my PhD dissertation submitted to the Ohio State University. I thank Aubhik Khan and Julia Thomas for their guidance, and David Blau, Yongsung Chang, Paul Evans, Daeho Kim, Alex Ludwig, Lee Ohanian, Victor Ríos-Rull, and three referees, among others, for helpful comments.

Appendix A. Additional materials for Section 2

A.1. Preference aggregation

Following Rogerson (1988), assume that the budget set of individual households is non-convex and that there is perfect employment insurance through lotteries. Thus, although each household can only choose either 0 or $\overline{h}_t$, the stand-in household that chooses the fraction of working population n_t has a convex constraint set. The period expected utility $U : R_+ \times [0,1] \rightarrow R$ for the stand-in household can be written as

$$U(c_t, n_t) = n_t \left(\log c_t - \theta \frac{\overline{h}_t^{1+(1/\gamma)}}{1+(1/\gamma)} \right) + (1-n_t) \left(\log c_t - \theta \frac{0^{1+(1/\gamma)}}{1+(1/\gamma)} \right)$$

$$= \log c_t - \theta \frac{\overline{h}_t^{1+(1/\gamma)}}{1+(1/\gamma)} n_t .$$

A.2. Equilibrium definition

A recursive equilibrium is a set of functions for prices, quantities, and values

$$\left\{ w,\ r,\ (q_j)_{j=1}^{N_z},\ n_h,\ k'_h,\ (a_j)_{j=1}^{N_z},\ h,\ n'_f,\ k_f,\ M_1,\ M_2,\ \Pi,\ W,\ v \right\},$$

such that:

(i) W solves equations (10)–(12), and n_h, k'_h and $(a_j)_{j=1}^{N_z}$ are the associated policy functions for the household;

(ii) v solves equations (5)–(7), and h, n'_f and k_f are the associated policy functions for the firm;

(iii) prices of goods market, labor market, and asset markets are competitively determined; and

(iv) (*consistency*) the individual policy functions are consistent with the perceived aggregate laws of motion, that is, $n'_f(N, z_i, N, K) = M_1(z_i, N, K)$ and $k'_h(K, z_i, N, K) = M_2(z_i, N, K)$ for all z_i, N, K.

Appendix B. Inefficiency of the decentralized equilibrium

Theorem B1. *The decentralized equilibrium yields the planner's allocations only when $\gamma \rightarrow \infty$.*

Proof: Let $\varepsilon = 1/\gamma$. Consider a planner who maximizes the stand-in household with lotteries:

$$V(n, k, z_i) = \max_{\substack{k' \geq 0 \\ n', h \in [0,1]}} \left\{ \log c - \theta \frac{h^{1+\varepsilon}}{1+\varepsilon} n + \beta \sum_{j=1}^{N_z} \pi_{ij} V(n', k', z_j) \right\},$$

subject to

$$c + k' + \Phi(n, n') = z_i F(h, n, k) + (1 - \delta)k.$$

The key is to note that $B(h) \equiv h^{\varepsilon}/(1+\varepsilon)$ is taken as given by households in the decentralized economy.

The three key optimality conditions for the planner are

$$\frac{1}{c} D_2\Phi(n, n') = \beta \sum_{j=1}^{N_z} \pi_{ij} \left\{ \frac{1}{c'_j} \left[z_j D_2 F(h', n', k') - D_1\Phi(n', n'') \right] - \theta \frac{h'^{1+\varepsilon}}{1+\varepsilon} \right\}, \tag{A1}$$

$$\frac{1}{c} = \beta \sum_{j=1}^{N_z} \pi_{ij} \left\{ \frac{1}{c'_j} [z_j D_3 F(h', n', k') + 1 - \delta] \right\}, \tag{A2}$$

$$\frac{1}{c} z_i D_1 F(h, n, k) = \theta h^{\varepsilon} n. \tag{A3}$$

However, recall the optimality condition from the decentralized problem with the labor–leisure conditions (13) gives

$$\begin{aligned} p(z_i, s) D_2\Phi(n, n') &= \beta \sum_{j=1}^{N_z} \pi_{ij} p(z_j, s') \\ &\quad \times \left[z_j D_2 f(h', n', k') - w'(z_j, s')h' - D_1\Phi(n', n'') \right] \\ &= \beta \sum_{j=1}^{N_z} \pi_{ij} \left\{ p(z_j, s') \left[z_j D_2 f(h', n', k') - D_1\Phi(n', n'') \right] - \theta \frac{h'^{1+\varepsilon}}{1+\varepsilon} \right\}, \end{aligned} \tag{A4}$$

which equals equation (A1).

Next, the household's Euler equation (14) can be combined with the first-order condition for k from the firm's problem,

$$\frac{1}{c} = \beta \sum_{j=1}^{N_z} \pi_{ij} \frac{1}{c'_j} \left[z'_j D_3 F(h', n', k') + 1 - \delta \right], \tag{A5}$$

which equals equation (A2).

Finally, the first-order condition for h from the firm's problem can be combined with the labor–leisure condition (13) by eliminating wage:

$$\frac{1}{c} z_i D_1 F(h, n, k) = \frac{1}{1+\varepsilon} \theta h^{\varepsilon} n. \tag{A6}$$

Note that this last equation collapses to equation (A3) if and only if $\varepsilon = 0$. Intuitively, the planner takes account of discrepancy between individual and aggregate labor elasticity. Thus, when the discrepancy collapses to zero, the planner has no margin to improve, and the decentralized equilibrium can produce socially efficient allocations. □

Appendix C. Data

The aggregate labor data I use is based on Cociuba et al. (2018), who obtain data mostly from the US Bureau of Labor Statistics. I define the intensive margin to be hours per worker, namely total hours divided by the number of the employed. The extensive margin is defined as the employment–population ratio. Output is the real GDP (chained 2005 dollars) from the US Bureau of Economic Analysis. All the data series are quarterly and HP-filtered using the smoothing parameter equal to 1,600. The sample periods are from 1956:I to 2010:IV, after eliminating the first and last four quarters of HP-filtered data.

Appendix D. Sensitivity analysis

I also consider a model economy that is calibrated with a different value of ϕ. Specifically, instead of 5 percent of the steady-state hours per worker, I consider 10 percent. Although this is a relatively substantial change in terms of the value of ϕ (an increase of 100 percent), Tables D1–D4 show that the main results are quite robust.

Table D1. Cyclical volatilities relative to output: σ_x / σ_Y

ξ	γ	σ_Y	x					AC/Y
			C	I	h	N	$h \times N$	
US data		1.56	0.60	2.54	0.35	0.64	0.91	
Panel A. Baseline ($\hat{\xi} = 0.081$)	0.5	1.46	0.36	2.98	0.14	0.52	0.54	2.8e-7
	1.0	1.56	0.35	3.01	0.21	0.54	0.58	4.0e-7
	1.5	1.61	0.34	3.04	0.25	0.55	0.60	4.6e-7
Panel B. Low ($\hat{\xi} \div 10$)	0.5	1.66	0.34	3.05	0.10	0.67	0.66	1.1e-7
	1.0	1.71	0.34	3.07	0.16	0.68	0.67	1.4e-7
	1.5	1.74	0.33	3.08	0.20	0.68	0.67	1.5e-7
Panel C. High ($\hat{\xi} \times 10$)	0.5	1.28	0.34	3.06	0.21	0.25	0.33	4.6e-7
	1.0	1.40	0.34	3.07	0.30	0.29	0.43	6.9e-7
	1.5	1.48	0.33	3.10	0.36	0.31	0.48	8.6e-7

Table D2. Cyclicality of aggregates: $Cor(x, Y)$

ξ	γ	x				
		C	I	h	N	$h \times N$
US data		0.84	0.92	0.71	0.80	0.84
Panel A. Baseline ($\hat{\xi} = 0.081$)	0.5	0.93	0.99	0.59	0.79	0.92
	1.0	0.92	0.99	0.56	0.80	0.95
	1.5	0.92	0.99	0.55	0.80	0.96
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.91	0.99	0.32	0.88	0.94
	1.0	0.91	0.99	0.31	0.87	0.95
	1.5	0.91	0.99	0.32	0.86	0.96
Panel C. High ($\hat{\xi} \times 10$)	0.5	0.91	0.99	0.86	0.53	0.93
	1.0	0.91	0.99	0.84	0.57	0.97
	1.5	0.90	0.99	0.83	0.59	0.98

Table D3. Persistence of aggregates: ρ_x

ξ	γ	x					
		Y	C	I	h	N	$h \times N$
US data		0.85	0.85	0.89	0.55	0.91	0.86
Panel A. Baseline ($\hat{\xi} = 0.081$)	0.5	0.82	0.81	0.83	0.54	0.92	0.93
	1.0	0.82	0.81	0.83	0.48	0.90	0.92
	1.5	0.81	0.81	0.82	0.45	0.90	0.90
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.85	0.82	0.86	0.22	0.83	0.87
	1.0	0.84	0.82	0.85	0.15	0.80	0.87
	1.5	0.83	0.82	0.84	0.11	0.79	0.87
Panel C. High ($\hat{\xi} \times 10$)	0.5	0.76	0.81	0.75	0.68	0.95	0.89
	1.0	0.76	0.81	0.76	0.67	0.95	0.85
	1.5	0.76	0.82	0.76	0.66	0.95	0.83

Table D4. Frisch labor supply elasticities

ξ	γ	Intensive margin elasticity, $\hat{\alpha}_1^h$	Extensive margin elasticity, $\hat{\alpha}_1^N$	Aggregate labor supply elasticity, $\hat{\alpha}_1^H$
Panel A. Baseline ($\hat{\xi}$)	0.5	0.50	0.96	1.46
	1.0	1.00	1.17	2.17
	1.5	1.50	1.35	2.85
Panel B. Low ($\hat{\xi} \div 10$)	0.5	0.50	0.62	1.12
	1.0	1.00	0.56	1.56
	1.5	1.50	0.54	2.04
Panel C. High ($\hat{\xi} \times 10$)	0.5	0.50	0.55	1.05
	1.0	1.00	0.79	1.79
	1.5	1.50	1.00	2.50

Supporting information

Additional supporting information can be found online in the supporting information section at the end of the article.

Replication codes

References

Bils, M. and Cho, J-O. (1994), Cyclical factor utilization, *Journal of Monetary Economics 33*, 319–354.

Burda, M. C., Genadek, K. R., and Hamermesh, D. S. (2020), Unemployment and effort at work, *Economica 87* (347), 662–681.

Burnside, C., Eichenbaum, M., and Rebelo, S. (1993), Labor hoarding and the business cycle, *Journal of Political Economy 101*, 245–273.

Card, D. (1990), Labor supply with a minimum hours threshold, *Carnegie-Rochester Conference Series on Public Policy 33*, 137–168.

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

Chang, Y., Kim, S-B., Kwon, K., and Rogerson, R. (2019), Individual and aggregate labor supply in heterogeneous agent economies with intensive and extensive margins, *International Economic Review 60*, 3–24.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2013), Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities, in M. Eichenbaum, E. Hurst, and V. Ramey (eds), *NBER Macroeconomics Annual 2012*, Vol. 27, University of Chicago Press, Chicago, IL, 1–56.

Cho, J-O. and Cooley, T. F. (1994), Employment and hours over the business cycle, *Journal of Economic Dynamics and Control 18*, 411–432.

Cociuba, S. E., Prescott, E. C., and Ueberfeldt, A. (2018), US hours at work, *Economics Letters 169*, 87–90.

Cooley, T. F. and Prescott, E. C. (1995), Economic growth and business cycles, in T. F. Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton University Press, Princeton, NJ, 1–38.

Erosa, A., Fuster, L., and Kambourov, G. (2016), Towards a micro-founded theory of aggregate labour supply, *Review of Economic Studies 83*, 1001–1039.

Fiorito, R. and Zanella, G. (2012), The anatomy of the aggregate labor supply elasticity, *Review of Economic Dynamics 15*, 171–187.

Guner, N., Kaygusuz, R., and Ventura, G. (2023), Rethinking the welfare state, CEPR Discussion Paper 16275.

Hall, R. E. (2004), Measuring factor adjustment costs, *Quarterly Journal of Economics 119*, 899–927.

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

Hansen, G. D. and Prescott, E. C. (1995), Recursive methods for computing equilibria of business cycle models, in T. F. Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton University Press, Princeton, NJ, 39–64.

Jang, Y. and Yum, M. (2022), Nonlinear occupations and female labor supply over time, *Review of Economic Dynamics 46*, 51–73.

Keane, M. P. (2011), Labor supply and taxes: a survey, *Journal of Economic Literature 49*, 961–1075.

Keane, M. P. and Rogerson, R. (2012), Micro and macro labor supply elasticities: a reassessment of conventional wisdom, *Journal of Economic Literature 50*, 464–476.
Khan, A. and Thomas, J. K. (2003), Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?, *Journal of Monetary Economics 50*, 331–360.
Kimmel, J. and Kniesner, T. J. (1998), New evidence on labor supply: employment versus hours elasticities by sex and marital status, *Journal of Monetary Economics 42*, 289–301.
Kitao, S. (2014), Sustainable social security: four options, *Review of Economic Dynamics 17*, 756–779.
Kydland, F. E. and Prescott, E. C. (1991), Hours and employment variation in business cycle theory, *Economic Theory 1*, 63–81.
Osuna, V. and Ríos-Rull, J.-V. (2003), Implementing the 35 hour workweek by means of overtime taxation, *Review of Economic Dynamics 6*, 179–206.
Peterman, W. B. (2016), Reconciling micro and macro estimates of the Frisch labor supply elasticity, *Economic Inquiry 54*, 100–120.
Peterman, W. B. and Sommer, K. (2019), How well did social security mitigate the effects of the Great Recession?, *International Economic Review 60*, 1433–1466.
Prescott, E. C., Rogerson, R., and Wallenius, J. (2009), Lifetime aggregate labor supply with endogenous workweek length, *Review of Economic Dynamics 12*, 23–36.
Rogerson, R. (1988), Indivisible labor, lotteries and equilibrium, *Journal of Monetary Economics 21*, 3–16.
Rogerson, R. and Wallenius, J. (2009), Micro and macro elasticities in a life cycle model with taxes, *Journal of Economic Theory 144*, 2277–2292.
Shimer, R. (2010), *Labor Markets and Business Cycles*, Princeton University Press, Princeton, NJ.

First version submitted June 2022;
final version received June 2023.