Aggregate Labor Supply with Isoelastic Hours and Employment

Benjamin Schoefer
(UC Berkeley)

March 2026

This version is an in-progress note.
As an unusual "note," it would be not be suitable for regular full talk, but I'd be happy to give a talk.
Please do not circulate.
Thank you for considering our submission!

Abstract

This paper presents and dissects a parsimonious utility specification for isoelastic aggregate labor supply on both margins: the intensive margin (hours per worker, elasticity ε_h) and the extensive margin (employment, elasticity ε_n). It delivers closed-form labor supply at both margins, with total-hours Frisch elasticity $\varepsilon_H = \varepsilon_h + \varepsilon_n$, and both parameters map directly to observable reduced-form elasticities. The framework nests MaCurdy (1981) (intensive only), Galí (2011b) (extensive only), and indivisible labor as special cases. It lends itself to applications including tax policy analysis, labor wedge decomposition by margin, identification of the margin split from observed hours composition, and demographic heterogeneity extensions.

I thank the students of Econ 202B at UC Berkeley, to whom this material was taught over the years. I also thank Preston Mui for many discussions on these topics during our joint work published as Mui and Schoefer (2025).

1 Introduction

Aggregate labor supply is the product of two distinct margins: how many people work (employment, the extensive margin) and how much they work (hours per worker, the intensive margin). The two margins diverge over the business cycle, in trend, across countries, between demographic groups, and in response to policy.[1] The empirical labor-supply elasticities on both margins, especially the extensive one, remain contested.[2] However, existing applied macro models sit at two extremes: representative-household models typically feature only a single margin, while heterogeneous-agent models that accommodate both feature no tractable elasticity to calibrate.[3] No standard representative-household model appears in use that would provide the applied macroeconomist with a tractable labor-supply specification that separately parameterizes both margins with one easy-to-calibrate elasticity each.

This paper specifies, microfounds, and dissects such a utility function: a simple, additively separable *dual-margin MaCurdy isoelastic* specification, embedded in a frictionless spot labor market. Its two parameters are directly the observable Frisch wage elasticities at each margin: the *average-hours elasticity* ε_h (intensive margin) and the *employment elasticity* ε_n (extensive margin). Their sum gives the total-hours Frisch elasticity: $\varepsilon_H = \varepsilon_h + \varepsilon_n$.

Surprisingly, this paper's specification appears not to have been used in practice. It can be viewed as a coherent merge of the intensive-margin-only MaCurdy (1981) setup with the extensive-margin-only Galí (2011b) setup. Its only close—and nearly identical—cousin is a similar specification in Cho and Cooley (1994)—which has gone largely unused for more than three decades.[4] [The next version will list the surprisingly few instances in which the functional form related to Cho and Cooley (1994) has been used and flesh out this argument better. The point of this note is

[1] For instance, in the United States, hours fluctuations are concentrated at the extensive margin (Hansen, 1985; Rogerson, 1988; Kydland and Prescott, 1991), a composition that varies substantially across countries (Ohanian and Raffo, 2012). Long-run trends in total hours worked also split differently across margins and across countries (Boppart and Krusell, 2020). Young workers drive a disproportionate share of aggregate hours fluctuations (Jaimovich and Siu, 2009; Jaimovich et al., 2013), and women's labor supply is concentrated at the extensive margin (Attanasio, Levell, Low and Sánchez-Marcos, 2018).

[2] Reduced-form evidence points to small intensive-margin elasticities (0.1–0.5; MaCurdy 1981; Altonji 1986; Chetty, Guren, Manoli and Weber 2012), with evidence on the extensive margin similarly pointing to values below 0.5 (Chetty, Guren, Manoli and Weber, 2012; Martínez, Saez and Siegenthaler, 2021; Mui and Schoefer, 2025), although room remains for structural, local, or downward elasticities to be larger (Keane and Rogerson, 2012; Mui and Schoefer, 2025).

[3] The workhorse DSGE models (Christiano, Eichenbaum and Evans, 2005; Smets and Wouters, 2007) typically use a single MaCurdy (1981) Frisch elasticity, even papers explicitly studying the cross-country margin split (Ohanian and Raffo, 2012). Similarly, papers providing reduced-form estimates of margin-specific elasticities (Chetty, Guren, Manoli and Weber, 2012; Martínez, Saez and Siegenthaler, 2021; Mui and Schoefer, 2025) work without a tractable two-margin utility specification. Heterogeneous-agent frameworks (Chang and Kim, 2007; Rogerson and Wallenius, 2009; Chang, Kim, Kwon and Rogerson, 2019; Boppart, Krusell and Olsson, 2023) accommodate both margins but in fully numerical environments. Their aggregate extensive-margin labor supply schedules have no closed-form expression capturing a single elasticity parameter (see Mui and Schoefer, 2025, for an attempt to compare those models in a harmonized way, alongside empirical counterparts)—hence rarely used as labor-supply building blocks. This paper also abstracts from worker flows and participation dynamics with adjustment costs (see, e.g., Kydland and Prescott, 1991; Krusell, Mukoyama, Rogerson and Şahin, 2017).

[4] This is likely because Cho and Cooley's specification obscured the link to elasticities, presented equilibrium rather than Frisch elasticities, and did not dissect the margins or provide a microfoundation for the extensive-margin disutility. See Appendix A for a comparison and discussion.

to acknowledge, flesh out, and slightly improve upon their underappreciated setup.] I hope this crystallized dual-margin MaCurdy version and dedicated advertisement will see wider adoption, permitting macroeconomists to separately study both margins of aggregate labor supply without compromising on tractability.

2 The framework

I present the utility specification, derive Frisch optimality conditions at each margin, and establish the closed-form elasticity results.

Preferences. This is a parsimonious representative-household model with a frictionless spot labor market: a single household with a unit mass of potential workers chooses the employment rate $n \in [0,1]$ (extensive margin) and hours per employed worker $h \geq 0$ (intensive margin), with total hours $H = nh$. The period utility function is

$$U(c,h,n) \;=\; u(c) \;-\; n\,\eta\,\frac{h^{1+1/\varepsilon_h}}{1+1/\varepsilon_h} \;-\; \nu\,\frac{n^{\,1+(1+\varepsilon_h)/\varepsilon_n}}{1+(1+\varepsilon_h)/\varepsilon_n}, \tag{1}$$

an additively separable, isoelastic specification in both margins—the *dual-margin MaCurdy isoelastic* specification—where $u(c)$ is an increasing, concave consumption sub-utility; $\eta, \nu > 0$ are scale parameters; $\varepsilon_h > 0$ is directly the average-hours elasticity; and $\varepsilon_n > 0$ is directly the employment elasticity, as shown below. The hours term $n\eta h^{1+1/\varepsilon_h}/(1+1/\varepsilon_h)$ is the MaCurdy (1981) isoelastic disutility per employed worker; the employment disutility term is its extensive-margin analogue (Galí, 2011b,a; Galí, Smets and Wouters, 2012)—reparameterized so that ε_n appears directly (see below). See Appendix A for the mapping to the Cho and Cooley (1994) parameterization.

All workers face the same wage w and hence supply the same hours if working due to homogeneous intensive-margin disutility. The employment disutility term can be given a microfoundation through heterogeneity in workers' employment fixed costs: if individual employment fixed costs are drawn from a power-law distribution, aggregation yields an isoelastic aggregate employment schedule with Frisch elasticity ε_n (see Galí, 2011b; Mui and Schoefer, 2025). The setup appeals to a utilitarian representative household: the household (planner) dispatches members into employment in ascending order of their employment fixed cost and also chooses the hours h of all employed workers; individual members do not solve separate optimization problems.[5]

[5] A simple one-sector version of the Galí (2011b) setup proceeds as follows. Let there be a unit continuum of workers $j \in [0,1]$ with heterogeneous employment fixed costs $\eta_j = \nu\, j^{(1+\varepsilon_h)/\varepsilon_n}$, indexed in ascending order of cost (a power-law distribution with exponent $(1+\varepsilon_h)/\varepsilon_n$). The household chooses h for all employed workers, yielding the per-worker employment surplus (gross of the employment fixed cost) $S \equiv \lambda w h - \eta\, h^{1+1/\varepsilon_h}/(1+1/\varepsilon_h) = (\lambda w)^{1+\varepsilon_h}/[(1+\varepsilon_h)\,\eta^{\varepsilon_h}]$ at optimal h. Employing workers $j \in [0,n]$ costs $\int_0^n \nu\, j^{(1+\varepsilon_h)/\varepsilon_n}\, dj = \nu\, n^{1+(1+\varepsilon_h)/\varepsilon_n}/[1+(1+\varepsilon_h)/\varepsilon_n]$—precisely the employment disutility term in Equation (1), defined for any $n \in [0,1]$. The FOC with respect to n (Leibniz rule) equates the marginal fixed cost to the surplus, $\nu\, n^{(1+\varepsilon_h)/\varepsilon_n} = S$, which is exactly Equation (7): the household employs worker j iff $\eta_j \leq S$, with the marginal worker satisfying $\eta_n = S$. At $\varepsilon_h = 0$, the exponent reduces to $1/\varepsilon_n$ and the setup collapses to Galí (2011b) with indivisible labor: $u(c) - \nu \int_0^n j^{1/\varepsilon_n}\, dj = u(c) - \nu\, n^{1+1/\varepsilon_n}/(1+1/\varepsilon_n)$, with $h = 1$ following automatically from $h = (\lambda w/\eta)^0 = 1$. See Mui and Schoefer (2025) for a related derivation and empirical application. The isoelastic

Optimality conditions. A household maximizes $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t, h_t, n_t)$ subject to $c_t + a_{t+1} = w_t n_t h_t + (1 + r_t) a_t$. The within-period optimality conditions are (dropping time subscripts):

Consumption:

$$\lambda = u'(c), \tag{2}$$

Intensive margin (average-hours elasticity ε_h):

$$\eta\, h^{1/\varepsilon_h} = \lambda\, w \tag{3}$$

$$\Leftrightarrow \quad h = \left(\frac{\lambda w}{\eta}\right)^{\varepsilon_h} \tag{4}$$

$$\Rightarrow \quad \varepsilon_h = \left.\frac{d \ln h}{d \ln w}\right|_{\lambda}.$$

Extensive margin (employment elasticity ε_n):

$$\nu\, n^{(1+\varepsilon_h)/\varepsilon_n} = \underbrace{\lambda w h - \eta \frac{h^{1+1/\varepsilon_h}}{1 + 1/\varepsilon_h}}_{\equiv S} = \frac{(\lambda w)^{1+\varepsilon_h}}{(1+\varepsilon_h)\, \eta^{\varepsilon_h}} \tag{5}$$

$$\Leftrightarrow \quad n = \left[\frac{1}{(1+\varepsilon_h)\, \eta^{\varepsilon_h}\, \nu}\right]^{\varepsilon_n/(1+\varepsilon_h)} \cdot (\lambda w)^{\varepsilon_n} \tag{6}$$

$$\Rightarrow \quad \varepsilon_n = \left.\frac{d \ln n}{d \ln w}\right|_{\lambda}. \tag{7}$$

where $S \equiv \lambda w h - \eta h^{1+1/\varepsilon_h}/(1 + 1/\varepsilon_h)$ is the employment surplus (net flow value of employment after hours disutility but before the employment fixed cost), equal to $\frac{(\lambda w)^{1+\varepsilon_h}}{(1+\varepsilon_h)\, \eta^{\varepsilon_h}}$ at optimal h. S is the natural allocative object for the extensive margin: the FOC delivers $d \ln n / d \ln S|_{\lambda} = \varepsilon_n/(1 + \varepsilon_h)$.

A note on employment, surplus, and wages. To appreciate the employment isoelasticity in the presence of an extensive margin, note that the employment-wage elasticity can be written as:

$$\left.\frac{d \ln n}{d \ln w}\right|_{\lambda} = \left.\frac{d \ln n}{d \ln S}\right|_{\lambda} \times \left.\frac{d \ln S}{d \ln w}\right|_{\lambda} = (1 + \varepsilon_h) \times \frac{\varepsilon_n}{1 + \varepsilon_h} = \varepsilon_n. \tag{8}$$

The first term rightmost term is the surplus-wage elasticity and uses the envelope theorem (so that intensive-margin hours reoptimization has no first-order effect on S):

$$\left.\frac{d \ln S}{d \ln w}\right|_{\lambda} = \frac{\lambda w h}{S} = 1 + \varepsilon_h. \tag{9}$$

This expression reflects a "leverage" effect (i.e., it is ≥ 1) stemming from the convexity of intensive-margin disutility of labor. The prenormalization exponent $(1+\varepsilon_h)/\varepsilon_n$ giving the inverse employment-

schedule for n can in principle exceed 1 for sufficiently high wages; throughout, $n \in (0, 1)$ is assumed, and the analysis pertains to the parameter and wage region satisfying this constraint.

surplus elasticity (following Equation (6)) exactly cancels this effect, to isolate the employment-wage elasticity ε_n in Equation (7):[6]

$$\left.\frac{d\ln n}{d\ln S}\right|_{\lambda} \;=\; \frac{\lambda w h}{S} \;=\; \frac{\varepsilon_n}{1+\varepsilon_h}. \tag{10}$$

Hence, the elasticity parameters reflect the economics inherent to the combination of both active margins and have been engineered to map into clearcut elasticities for wages. Hence, any shifter that shifts S only but not the wage would affect desired employment through $\varepsilon_n/(1+\varepsilon_h)$.

Total-hours elasticity. The total-hours elasticity follows immediately from Equations (3) and (7). Since $H = nh$,

$$\varepsilon_H \;\equiv\; \left.\frac{\partial\ln H}{\partial\ln w}\right|_{\lambda} \;=\; \varepsilon_h \;+\; \varepsilon_n. \tag{11}$$

Frisch elasticities and income effects. The Frisch elasticities hold λ fixed, isolating substitution effects. Income effects operate symmetrically through λ: a fall in λ contracts both margins with the same elasticities $(\varepsilon_h, \varepsilon_n)$. In the special case of quasi-linear utility $u(c) = c$, $\lambda \equiv 1$ and income effects vanish, Frisch and Marshallian elasticities coincide.

Total-hours equivalence. From the perspective of total hours, the utility specification in Equation (1) is equivalent to a one-margin setup:

$$U^{\text{agg}}(c,H) \;=\; u(c) \;-\; \Psi\,\frac{H^{1+1/\varepsilon_H}}{1+1/\varepsilon_H}, \tag{12}$$

which has total-hours elasticity $\varepsilon_H = \varepsilon_h + \varepsilon_n$.

Special cases. The specification nests several existing specifications, including canonical limiting cases. $\varepsilon_n \to 0$ recovers MaCurdy (1981) (pure intensive margin). $\varepsilon_h = 0$ recovers Galí (2011b) (pure extensive margin; the exponent simplifies to $1/\varepsilon_n$; also Galí, 2011a; Galí, Smets and Wouters, 2012). At $\varepsilon_h = 0$ and $\varepsilon_n \to \infty$, the model recovers the Hansen (1985)–Rogerson (1988) indivisible-labor limit: perfectly elastic employment and no intensive-margin choice.[7] The Cho and Cooley (1994)

[6] Intuitively, the convexity of the hours disutility makes S grow faster than w. Think of S as the producer surplus of a worker facing "price" λw per hour and rising marginal cost $\eta h^{1/\varepsilon_h}$: since marginal cost exceeds average cost for a convex cost curve, total disutility $\eta h^{1+1/\varepsilon_h}/(1+1/\varepsilon_h)$ equals only $\varepsilon_h/(1+\varepsilon_h)$ of gross earnings $\lambda w h$, so the surplus S—the area between the price line and the rising marginal cost curve—equals $1/(1+\varepsilon_h)$ of gross earnings: the accounting identity $\lambda w h = (1+\varepsilon_h)S$ holds at the optimum. When w rises, this surplus area expands faster than w itself. At $\varepsilon_h = 0$, the disutility is linear—average equals marginal, no infra-marginal area—and S moves one-to-one with w.

[7] At $\varepsilon_n \to \infty$ with $\varepsilon_h > 0$, the framework generalizes the indivisible-labor setup to allow non-zero intensive-margin elasticity. For instance, Hall (2009) sets $\nu = 0$ (no employment disutility; employment from search-and-matching) and Bils, Klenow and Malin (2018) impose a fixed employment cost ψ mapping to $\nu = \eta\psi > 0$; as $\varepsilon_n \to \infty$ the disutility collapses to the linear term νn and the FOC reduces to $S = \nu$—their employment condition. Hence, the present framework makes clear that an infinite Frisch employment elasticity arises either with $\nu = 0$ or with $\nu > 0$ and $\varepsilon_n \to \infty$. Of course, the framework also crucially relaxes $\varepsilon_n \to \infty$ by introducing the Galí (2011b) employment disutility term, enabling a finite, easy-to-calibrate ε_n.

specification is an equivalent reparameterization of Equation (1); see Appendix A for the precise mapping.

3 Applications and extensions

I close this paper by illustrating basic applications and extensions of the two-parameter specification.

Illustrative calibration. Equation (1) is a drop-in utility function for DSGE models: set ε_h and ε_n to empirical targets; calibrate η and ν to steady-state hours and employment. The scale parameters follow in closed form from the two FOCs: $\eta = \bar{\lambda}\,\bar{w}\,\bar{h}^{-1/\varepsilon_h}$ and $\nu = \frac{\bar{\lambda}\,\bar{w}\,\bar{h}}{1+\varepsilon_h}\,\bar{n}^{-(1+\varepsilon_h)/\varepsilon_n}$.

An influential meta-analysis by Chetty, Guren, Manoli and Weber (2012) proposes $\varepsilon_h \approx 0.30$ and $\varepsilon_n \approx 0.35$, giving $\varepsilon_H = 0.65$—well below the values of 2–4 implicitly required by standard RBC models. Drawing on stated preferences rather than equilibrium hours and employment adjustment, Mui and Schoefer (2025) document asymmetry in local labor supply preferences: the local extensive-margin elasticity is roughly 2–3, stays high for downward wage movements, and falls toward 0.5 for upward movements. Besides the validity of the isoelasticity benchmark, an open question is the mapping between reduced-form and structural preference parameters (Keane and Rogerson, 2012; Mui and Schoefer, 2025), and the appropriate choice may depend on whether the analysis is descriptive or welfare-relevant.

The margin split of total-hours fluctuations. Any common impulse $\mu \equiv d\ln(\lambda w)$—whether reflecting a productivity shock, a wealth change, or a policy change—drives $d\ln h = \varepsilon_h\,\mu$ and $d\ln n = \varepsilon_n\,\mu$, hence $d\ln H = \varepsilon_H\,\mu$. The fraction of total-hours adjustment borne by each margin defines the *margin shares*:

$$s_h \equiv \frac{d\ln h}{d\ln H} = \frac{\varepsilon_h}{\varepsilon_H}, \qquad s_n \equiv \frac{d\ln n}{d\ln H} = \frac{\varepsilon_n}{\varepsilon_H}, \tag{13}$$

with $s_h + s_n = 1$ by construction. These shares are identified from observed variation under the assumption that a single common factor $\mu \equiv d\ln(\lambda w)$ drives both margins: in the data, $s_h = \mathrm{SD}(\Delta\ln h)/\mathrm{SD}(\Delta\ln H)$, $s_n = \mathrm{SD}(\Delta\ln n)/\mathrm{SD}(\Delta\ln H)$, and the elasticity ratio $\varepsilon_n/\varepsilon_h = \mathrm{SD}(\Delta\ln n)/\mathrm{SD}(\Delta\ln h)$. The Ohanian and Raffo (2012) cross-country ratios from the pre-1984 period—US: 2.84, Canada: 3.30, UK: 1.74, Euro: 1.06, Japan: 0.66—are direct empirical estimates of $\varepsilon_n/\varepsilon_h$ by country. The additive formula $\varepsilon_H = \varepsilon_h + \varepsilon_n$ implies that any pair summing to the same aggregate elasticity yields identical total-hours responses; the split matters for composition-targeting policies (Saez, 2002; Chodorow-Reich and Karabarbounis, 2016) and welfare evaluation.

A tension arises immediately. Because both margins are driven by the same composite $\mu \equiv \Delta\ln(\lambda w)$, the observed split is a direct read on $\varepsilon_n/\varepsilon_h$—assuming no margin-specific wedge. At the Chetty, Guren, Manoli and Weber (2012) consensus ($\varepsilon_h \approx \varepsilon_n \approx 0.30$), the predicted variance ratio is 1, inconsistent with every country in the Ohanian and Raffo (2012) sample except near the

European average. Hence, either supply elasticities differ substantially across countries and over time, or margin-specific distortions account for the divergence—a question I address next when discussing the labor wedge, a tractable notion of disequilibrium.

Labor wedge decomposition. The two-margin framework delivers a structural decomposition of the household-side labor wedge. Shimer (2009) defines and measures the aggregate wedge but without a margin-specific model. Karabarbounis (2014) decomposes it along the household/firm dimension. Bils, Klenow and Malin (2018) construct empirical counterparts of both margin-specific wedges, though within the extreme case of a Frisch employment elasticity implicitly hardwired to be infinite, rather than a free parameter (see Footnote 7). The present decomposition is on the household side, splitting it further by margin in closed form. Define the *intensive-margin labor wedge* τ_h and *employment-margin labor wedge* τ_n by the proportional gaps in each optimality condition:

$$\begin{aligned} h &= \left(\tfrac{\lambda w(1-\tau_h)}{\eta}\right)^{\varepsilon_h} \\ n &= \left[\tfrac{(\lambda w(1-\tau_n))^{1+\varepsilon_h}}{(1+\varepsilon_h)\,\eta^{\varepsilon_h}\,\nu}\right]^{\varepsilon_n/(1+\varepsilon_h)}. \end{aligned} \tag{14}$$

The two wedges partition into *marginal* and *inframarginal*, both evaluated at fixed λ. The intensive wedge τ_h is the marginal gap: it distorts the optimality condition at the last unit of hours, shifting h with elasticity $\varepsilon_h|_\lambda$. By the envelope theorem, a change in τ_h alone—holding τ_n and λ fixed—leaves the employment surplus S unchanged at the optimum, so employment is unaffected. The employment wedge τ_n is the inframarginal gap: it scales S by $(1-\tau_n)^{1+\varepsilon_h}$ without altering the hours FOC, so h is unchanged. The two wedges thus affect orthogonal margins and are separately identified from (δ_h, δ_n) via Equation (17) below; their fiscal counterparts—the marginal rate τ_m and average rate τ_a—are developed in the tax application below.

Total hours combine Equation (14):

$$H = \frac{C}{\eta^{\varepsilon_h}}\,(\lambda w)^{\varepsilon_H}\,(1-\tau_h)^{\varepsilon_h}\,(1-\tau_n)^{\varepsilon_n}, \tag{15}$$

where $C \equiv [(1+\varepsilon_h)\,\eta^{\varepsilon_h}\,\nu]^{-\varepsilon_n/(1+\varepsilon_h)}$ is a constant. Defining the aggregate wedge τ_H via $H = \frac{C}{\eta^{\varepsilon_h}}(\lambda w)^{\varepsilon_H}(1-\tau_H)^{\varepsilon_H}$ and taking logs yields the *exact* aggregate-wedge decomposition:

$$\ln(1-\tau_H) \;=\; s_h\,\ln(1-\tau_h) \;+\; s_n\,\ln(1-\tau_n), \tag{16}$$

with shares $s_h = \varepsilon_h/\varepsilon_H$ and $s_n = \varepsilon_n/\varepsilon_H$. The aggregate wedge is the log-elasticity-share-weighted average of margin wedges; a country whose hours fluctuations pass mainly through employment (s_n large) has its aggregate wedge dominated by τ_n. Just as the tax formula collapses to $d\ln H = \varepsilon_H\, d\ln(1-\tau)$ under a uniform tax, the aggregate wedge τ_H fully characterizes the distortion when $\tau_h = \tau_n$; the decomposition adds power when $\tau_h \neq \tau_n$.

The two margin wedges are separately identified from hours and employment data. Denote

log-deviations from steady state δ_h, δ_n, $\delta_H = \delta_h + \delta_n$. Forming δ_h/ε_h and δ_n/ε_n and subtracting eliminates the unobserved common shock $\mu \equiv \Delta \ln(\lambda w)$ exactly, giving the *differential wedge*:[8]

$$\ln\left(\frac{1-\tau_h}{1-\tau_n}\right) = \frac{\delta_h}{\varepsilon_h} - \frac{\delta_n}{\varepsilon_n}. \tag{17}$$

Applied to a 1% aggregate shortfall at US composition ($\hat{s}_n = 0.74$) and Chetty elasticities ($\varepsilon_h = 0.30$, $\varepsilon_n = 0.35$), the diagnostic equals +0.012—employment is more distorted than hours; Japan ($\hat{s}_n = 0.40$) gives −0.009, the reverse (all country compositions derived from Ohanian and Raffo 2012). Setting the diagnostic to zero requires $\varepsilon_n/\varepsilon_h = \hat{s}_n/\hat{s}_h$—for the US, $\varepsilon_n \approx 0.85$, more than twice the Chetty consensus. Whether US extensive-margin dominance reflects a high ε_n or a large employment distortion at $\varepsilon_n = 0.35$ is a key open question (Shimer, 2009; Mui and Schoefer, 2025).

Tax policy: marginal vs. average rates. The public-finance distinction between marginal and average tax rates is the fiscal analog of the marginal/inframarginal wedge split in Equation (14): $\tau_m \equiv T'(wh)$ maps to τ_h, and $\tau_a \equiv T(wh)/wh$ maps to τ_n. By the same envelope-theorem argument, τ_m distorts hours with elasticity $\varepsilon_h|_\lambda$ and leaves the after-tax surplus S^{tax} unchanged; τ_a scales S^{tax} by $(1-\tau_a)^{1+\varepsilon_h}$—and by the employment surplus elasticity in Equation (9), this $(1+\varepsilon_h)$ leverage and the employment-surplus elasticity $\varepsilon_n/(1+\varepsilon_h)$ cancel, so ε_n applies to τ_a directly, exactly as ε_n applies to τ_n in Equation (14). With household budget constraint $c_t + a_{t+1} = n_t[w_t h_t - T(w_t h_t)] + (1+r_t)a_t$, the total-hours response is:

$$d\ln H\big|_\lambda = \varepsilon_h \, d\ln(1-\tau_m) + \varepsilon_n \, d\ln(1-\tau_a), \tag{18}$$

the direct tax-policy counterpart of Equation (16), embedding Saez (2002)'s insight in a standard macro isoelastic setting (see, e.g., Eissa and Liebman, 1996; Martínez et al., 2021, for margin-specific estimates). Under a uniform flat tax ($\tau_m = \tau_a = \tau$), Equation (18) collapses to $d\ln H|_\lambda = \varepsilon_H \, d\ln(1-\tau)$—the standard aggregate formula—just as a uniform wedge $\tau_h = \tau_n$ collapses to τ_H in Equation (16).

Nonemployment benefits and employment subsidies. With a nonemployment transfer b per non-working member, the budget constraint becomes $c_t + a_{t+1} = w_t n_t h_t + b_t(1-n_t) + (1+r_t)a_t$, modifying the employment surplus condition to:

$$\nu \, n^{(1+\varepsilon_h)/\varepsilon_n} = S - \lambda \, b, \tag{19}$$

where $S = \frac{(\lambda w)^{1+\varepsilon_h}}{(1+\varepsilon_h)\,\eta^{\varepsilon_h}}$ is the employment surplus at optimal h. The allocative object is now $S - \lambda b$: employment surplus net of the nonemployment value, akin to the opportunity cost of employment in search models (Chodorow-Reich and Karabarbounis, 2016). The intensive-margin condition is

[8] From Equation (14), log-deviations satisfy $\delta_h = \varepsilon_h[\mu + \ln(1-\tau_h)]$ and $\delta_n = \varepsilon_n[\mu + \ln(1-\tau_n)]$, where $\mu = \Delta\ln(\lambda w)$. Dividing: $\delta_h/\varepsilon_h = \mu + \ln(1-\tau_h)$ and $\delta_n/\varepsilon_n = \mu + \ln(1-\tau_n)$. Subtracting eliminates μ exactly, yielding $\delta_h/\varepsilon_h - \delta_n/\varepsilon_n = \ln(1-\tau_h) - \ln(1-\tau_n) = \ln[(1-\tau_h)/(1-\tau_n)]$.

unchanged.

The b extension is best understood through the employment surplus elasticity in Equation (9). The two policy instruments enter $S - \lambda b$ through different channels: w enters through earnings ($d \ln S / d \ln w = 1 + \varepsilon_h$ by Equation (9), canceling the structural primitive $\varepsilon_n/(1+\varepsilon_h)$ and yielding $d \ln n / d \ln w = \varepsilon_n \cdot S/(S - \lambda b)$), whereas b shifts the threshold directly—$d(S - \lambda b)/db = -\lambda$, no earnings channel—so the structural primitive $\varepsilon_n/(1+\varepsilon_h)$ survives, giving $d \ln n / d \ln b = -\varepsilon_n/(1+\varepsilon_h) \cdot \lambda b/(S - \lambda b)$.[9] The qualitative message is robust: benefits compress the net surplus and raise the stakes of wage changes. The exact amplification factors are functional-form artifacts—adding a common λb to every worker's threshold shifts a power-law distribution by a constant, destroying scale-freeness. I reiterate that the pragmatic goal of the specification is to reverse-engineer a functional form delivering isoelasticity with respect to *wages*. The b extension is a useful qualitative tool but ventures beyond that design principle.

Demographic heterogeneity [under particular revision]. The framework extends to demographic groups by assigning group-specific preference parameters. With population masses N_g sharing the common marginal utility of wealth λ (assuming pooled household income), the household utility is

$$U(c, \{h_g\}, \{n_g\}) \;=\; u(c) \;-\; \sum_g N_g \left[n_g \, \eta_g \, \frac{h_g^{1+1/\varepsilon_h^g}}{1 + 1/\varepsilon_h^g} \;+\; \nu_g \, \frac{n_g^{1+(1+\varepsilon_h^g)/\varepsilon_n^g}}{1 + (1+\varepsilon_h^g)/\varepsilon_n^g} \right], \tag{20}$$

with group-specific elasticities $(\varepsilon_h^g, \varepsilon_n^g)$ and scale parameters (η_g, ν_g). Each group's FOCs deliver $h_g = (\lambda w_g/\eta_g)^{\varepsilon_h^g}$ and, directly parallel to Equation (7), $n_g = [(1+\varepsilon_h^g)\,\eta_g^{\varepsilon_h^g}\,\nu_g]^{-\varepsilon_n^g/(1+\varepsilon_h^g)} (\lambda w_g)^{\varepsilon_n^g}$, with ε_n^g the clean wage elasticity of group-g employment. Defining hours-share weights $\omega_g \equiv N_g n_g h_g / H$, the aggregate margin elasticities are

$$\bar{\varepsilon}_h \;\equiv\; \sum_g \omega_g \, \varepsilon_h^g, \qquad \bar{\varepsilon}_n \;\equiv\; \sum_g \omega_g \, \varepsilon_n^g, \qquad \varepsilon_H \;=\; \bar{\varepsilon}_h + \bar{\varepsilon}_n.$$

Young workers and women concentrate their cyclical adjustment at the extensive margin (Jaimovich and Siu, 2009; Jaimovich, Pruitt and Siu, 2013; Attanasio, Levell, Low and Sánchez-Marcos, 2018), consistent with higher ε_n^g for those groups.

Labor supply and labor demand. This paper characterizes the labor-*supply* side. The equilibrium margin split additionally reflects demand: when n and h are imperfect substitutes in production,

[9] The three elasticities in full:

$$\left.\frac{\partial \ln n}{\partial \ln w}\right|_\lambda = \varepsilon_n \cdot \frac{S}{S - \lambda b}, \qquad \left.\frac{\partial \ln n}{\partial \ln b}\right|_\lambda = -\frac{\varepsilon_n}{1+\varepsilon_h} \cdot \frac{\lambda b}{S - \lambda b}, \qquad \left.\frac{\partial \ln n}{\partial \ln (w-b)}\right|_\lambda = \varepsilon_n \cdot \frac{S}{S - \lambda b} \cdot \frac{w-b}{w}.$$

At $b = 0$ the wage elasticity equals ε_n for any ε_h; at $\varepsilon_h = 0$ the $(w-b)$ elasticity equals ε_n for any b; in general neither recovers ε_n directly.

demand-side forces—hiring costs and adjustment frictions for employment (Llosa, Ohanian, Raffo and Rogerson, 2014), working-time regulations for hours (Schoefer, 2010)—shift the equilibrium composition without altering the supply elasticities ε_h and ε_n themselves. Pairing Equation (1) with a production function that separately features n and h delivers the full general-equilibrium framework; this paper provides the supply-side half.

References

Altonji, Joseph, "Intertemporal Substitution in Labor Supply: Evidence from Micro Data," *Journal of Political Economy*, 1986, *94* (3), S176–S215.

Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos, "Aggregating Elasticities: Intensive and Extensive Margins of Women's Labor Supply," *Econometrica*, 2018, *86* (6), 2049–2082.

Bils, Mark, Peter Klenow, and Benjamin Malin, "Resurrecting the Role of the Product Market Wedge in Recessions," *American Economic Review*, 2018, *108* (4–5), 1118–1146.

Boppart, Timo and Per Krusell, "Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective," *Journal of Political Economy*, 2020, *128* (1), 118–157.

_ , _ , and Jonna Olsson, "Labor Supply When Productivity Keeps Growing," *Review of Economic Dynamics*, 2023, *50*, 61–87.

Chang, Yongsung and Sun-Bin Kim, "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 2007, *97* (5), 1939–1956.

_ , _ , Kyooho Kwon, and Richard Rogerson, "Individual and Aggregate Labor Supply in Heterogeneous Agent Economies with Intensive and Extensive Margins," *International Economic Review*, 2019, *60* (1), 3–24.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 2012, *27*, 1–56.

Cho, Jang-Ok and Thomas Cooley, "Employment and Hours over the Business Cycle," *Journal of Economic Dynamics and Control*, 1994, *18* (2), 411–432.

Chodorow-Reich, Gabriel and Loukas Karabarbounis, "The Cyclical Behavior of the Price of Nonemployment," *Quarterly Journal of Economics*, 2016, *131* (2), 811–864.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Eissa, Nada and Jeffrey Liebman, "Labor Supply Response to the Earned Income Tax Credit," *Quarterly Journal of Economics*, 1996, *111* (2), 605–637.

Galí, Jordi, "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 2011, *9* (3), 436–461.

_ , *Unemployment Fluctuations and Stabilization Policies: A New Keynesian Perspective*, Cambridge, MA: MIT Press, 2011.

_ , Frank Smets, and Rafael Wouters, "Unemployment in an Estimated New Keynesian Model," *NBER Macroeconomics Annual*, 2012, *26* (1), 329–360.

Hall, Robert, "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 2009, *117* (2), 281–323.

Hansen, Gary, "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 1985, *16* (3), 309–327.

Jaimovich, Nir and Henry Siu, "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility," *American Economic Review*, 2009, *99* (3), 804–826.

_ , Seth Pruitt, and Henry Siu, "The Demand for Youth: Explaining Age Differences in the Volatility of Hours," *American Economic Review*, 2013, *103* (7), 3022–3044.

Karabarbounis, Loukas, "The Labor Wedge: MRS vs. MPN," *Review of Economic Dynamics*, 2014, *17* (2), 206–223.

Keane, Michael and Richard Rogerson, "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom," *Journal of Economic Literature*, 2012, *50* (2), 464–476.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin, "Gross Worker Flows over the Business Cycle," *American Economic Review*, 2017, *107* (11), 3447–3476.

Kydland, Finn and Edward Prescott, "Hours and Employment Variation in Business Cycle Theory," *Economic Theory*, 1991, *1* (1), 63–81.

Llosa, Gonzalo, Lee Ohanian, Andrea Raffo, and Richard Rogerson, "Firing Costs and Labor Market Fluctuations: A Cross-Country Analysis," 2014. 2014 Meeting Papers 533, Society for Economic Dynamics.

MaCurdy, Thomas, "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 1981, *89* (6), 1059–1085.

Martínez, Isabel, Emmanuel Saez, and Michael Siegenthaler, "Intertemporal Labor Supply Substitution? Evidence from the Swiss Income Tax Holidays," *American Economic Review*, 2021, *111* (2), 506–546.

Mui, Preston and Benjamin Schoefer, "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin," *Review of Economic Studies*, 2025, *92* (1), 442–475.

Ohanian, Lee and Andrea Raffo, "Aggregate Hours Worked in OECD Countries: New Measurement and Implications for Business Cycles," *Journal of Monetary Economics*, 2012, *59* (1), 40–56.

Rogerson, Richard, "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 1988, *21* (1), 3–16.

_ **and Johanna Wallenius**, "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory*, 2009, *144* (6), 2277–2292.

Saez, Emmanuel, "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *Quarterly Journal of Economics*, 2002, *117* (3), 1039–1073.

Schoefer, Benjamin, "Regulation and Taxation: A Complementarity," *Journal of Comparative Economics*, 2010, *38* (4), 381–394.

Shimer, Robert, "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics*, 2009, *1* (1), 280–297.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586–606.

Appendix A: Mapping to Cho and Cooley (1994)

[In progress. The next version will list the surprisingly few instances in which the functional form related to Cho and Cooley (1994) has been used and flesh out this argument better. The point of this note is to acknowledge, flesh out, and slightly improve upon their underappreciated setup.]

The Cho and Cooley (1994) specification is an underappreciated early contribution: to the best of this paper's knowledge, it is the only prior work to separately parameterize both labor supply margins in a tractable representative-household utility function. The present paper aims to give

their specification salient acknowledgment. Beyond this connection, the present paper also provides the microfoundation for the extensive margin choice by pointing to a power-law distributed employment fixed cost (Galí, 2011b; Mui and Schoefer, 2025), and identifies the employment surplus S and its wage elasticity (Equation (9))—the mechanism underlying the prenormalization exponent—which Cho and Cooley did not present.

This appendix documents the precise relationship between the utility specification of the present paper and that of Cho and Cooley (1994). It serves three purposes. First, it provides a complete change-of-variables table translating their structural curvature parameters (ν_{CC}, τ) into the Frisch elasticities $(\varepsilon_h, \varepsilon_n)$ used throughout this paper. Second, it verifies algebraically that the Cho and Cooley (1994) equilibrium aggregate elasticity formula reduces to $\varepsilon_h + \varepsilon_n$ in the $\sigma = 1$ special case (quasi-linear utility)—confirming Equation (11) from first principles in their notation. Third, it translates their Table 1 benchmark calibration ($\sigma = 0.8$, $\nu_{CC} = 2$, $\tau = 0.8$) into Frisch elasticities and explains why their reported equilibrium elasticities (0.27 and 1.02) fall below the corresponding Frisch values (0.50 and 1.875): the gap reflects income-effect attenuation along the equilibrium path where λ co-moves with w when $\sigma < 1$.

Cho and Cooley (1994) write period utility as

$$U_{CC}(c, n, e) \;=\; \frac{c^{\sigma}}{\sigma} \;-\; \frac{a}{\nu_{CC}+1}\, n^{\nu_{CC}+1} \cdot e \;-\; \frac{b}{\tau+1}\, e^{\tau+1}, \tag{A.1}$$

where (translating their notation to ours): their n = hours per employed worker → our h; their e = employment rate → our n; $\sigma \in (0, 1)$ is a consumption curvature parameter; (ν_{CC}, τ) are structural curvature parameters; and (a, b) are scale parameters (note: b here is Cho and Cooley's scale parameter, distinct from the nonemployment transfer of Section 3). The mapping to Equation (1) of the present paper is:

$$\eta = a, \qquad \frac{1}{\varepsilon_h} = \nu_{CC}, \qquad \nu = b, \qquad \frac{1+\varepsilon_h}{\varepsilon_n} = \tau, \tag{A.2}$$

or equivalently,

$$\varepsilon_h = \frac{1}{\nu_{CC}}, \qquad \varepsilon_n = \frac{1+\varepsilon_h}{\tau} = \frac{1+1/\nu_{CC}}{\tau}. \tag{A.3}$$

The mapping $\tau = (1+\varepsilon_h)/\varepsilon_n$ reveals the structural meaning of Cho and Cooley's curvature parameter: τ is precisely the prenormalization exponent that cancels the employment surplus elasticity in Equation (9). Since wages transmit to the employment surplus S at rate $d \ln S / d \ln w|_{\lambda} = 1 + \varepsilon_h$, the exponent $(1 + \varepsilon_h)/\varepsilon_n$ in the employment disutility delivers $d \ln n / d \ln S = \varepsilon_n/(1 + \varepsilon_h)$, so that these two factors cancel and ε_n emerges as the clean Frisch wage elasticity of employment. Cho and Cooley derived their model in terms of (ν_{CC}, τ) without framing it through the employment surplus S, its wage elasticity, or the prenormalization insight; these interpretations are contributions of the present paper.

Aggregate elasticity formula. Cho and Cooley (1994) derive the aggregate elasticity formula

$\varepsilon_{H,CC} = \sigma(1+\nu_{CC}+\tau)/(\tau R)$ with $R = \nu_{CC}+(\tau+\nu_{CC}+1)(1-\sigma)/\tau$. Setting $\sigma = 1$ gives $R = \nu_{CC} = 1/\varepsilon_h$, so:

$$\varepsilon_{H,CC}\big|_{\sigma=1} \;=\; \frac{1+\frac{1}{\varepsilon_h}+\frac{1+\varepsilon_h}{\varepsilon_n}}{\frac{1+\varepsilon_h}{\varepsilon_n}\cdot\frac{1}{\varepsilon_h}} \;=\; \frac{(1+\varepsilon_h)\left(\frac{1}{\varepsilon_h}+\frac{1}{\varepsilon_n}\right)}{\frac{1+\varepsilon_h}{\varepsilon_n\varepsilon_h}} \;=\; \varepsilon_n\varepsilon_h\left(\frac{1}{\varepsilon_h}+\frac{1}{\varepsilon_n}\right) \;=\; \varepsilon_h+\varepsilon_n, \tag{A.4}$$

confirming Equation (11). With $u(c) = c^\sigma/\sigma$, marginal utility is $\lambda = u'(c) = c^{\sigma-1}$. At $\sigma = 1$, utility is quasi-linear, $u(c) = c$ (see Section 2), so $\lambda = 1$ is constant regardless of the consumption level; holding λ fixed (the Frisch experiment) therefore coincides exactly with the full equilibrium response. For $\sigma < 1$, utility is strictly concave (coefficient of relative risk aversion $= 1-\sigma > 0$); as wages rise along the equilibrium path, consumption rises, $\lambda = c^{\sigma-1}$ falls (since $\sigma - 1 < 0$), and income effects attenuate the equilibrium elasticities below their Frisch values. Cho and Cooley (1994) consider $\sigma \in (0,1)$ throughout; the $\sigma \to 1$ limiting case and the Frisch-elasticity parameterization are not part of their analysis.

Benchmark calibration and Frisch elasticities. Cho and Cooley (1994) Table 1 uses $\sigma = 0.8$, $\nu_{CC} = 2$, $\tau = 0.8$. Translating to Frisch elasticities:

$$\varepsilon_h = \frac{1}{\nu_{CC}} = \frac{1}{2} = 0.50, \qquad \varepsilon_n = \frac{1+\varepsilon_h}{\tau} = \frac{1.50}{0.80} = 1.875. \tag{A.5}$$

These Frisch elasticities—$\varepsilon_h = 0.50$ and $\varepsilon_n = 1.875$—are substantially above the micro-consensus ranges of $\varepsilon_h \approx 0.30$ and $\varepsilon_n \approx 0.35$ (Chetty, Guren, Manoli and Weber, 2012; Mui and Schoefer, 2025). The equilibrium elasticities reported by Cho and Cooley (1994) (0.27 and 1.02, with aggregate 1.29) are lower because $\sigma = 0.8 < 1$ attenuates the Frisch response along the equilibrium path where λ co-moves with w, introducing an income-effect attenuation that grows as σ falls below 1. Even the equilibrium employment elasticity of 1.02 exceeds the modern consensus by a factor of roughly three, illustrating that the Cho and Cooley parameterization embeds implicitly high extensive-margin flexibility—a fact that was latent in their specification but was never extracted or connected to the empirical literature on employment elasticities.