I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16414

Labor Supply within the Firm

Michele Battisti
Ryan Michaels
Choonsung Park

AUGUST 2023

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16414

Labor Supply within the Firm

Michele Battisti
University of Glasgow and IZA

Ryan Michaels
Federal Reserve Bank of Philadelphia

Choonsung Park
Korea Institute of Finance

AUGUST 2023

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Labor Supply within the Firm*

There is substantial variation in working time even within employer-employee matches, and yet estimates of the Frisch elasticity of labor supply can be near zero. This paper proposes a tractable theory of earnings and working time to interpret these observations. Production complementarities attenuate the response of working time to idiosyncratic, or worker-specific, shocks, but firm-wide shocks are mediated by preference parameters. The model can be identified using firm-worker matched data, revealing a Frisch elasticity of around 0.5. A quasi-experimental approach that exploits only idiosyncratic variation would find an elasticity less than half this.

JEL Classification: J22, J31

Keywords: labour supply, production complementarities, Frisch elasticity

Corresponding author:
Michele Battisti
University of Glasgow
Glasgow G12 8QQ
Scotland

E-mail: Michele.Battisti@glasgow.ac.uk

* We are grateful to Lance Lochner for many detailed and incisive comments and to Yongsung Chang and Igor Livshits for helpful discussions and encouragement. We also appreciate feedback from participants at numerous conferences and workshops. We thank Guiseppe Tattara and the team of researchers at the University of Venice for granting access to the Veneto Work History (VWH) files, and Michelangelo Filippi for his help in understanding the Social Security records underlying the VWH. Battisti gratefully acknowledges financial support by the Leibniz Association (SAW-2012-ifo-3). Michaels gratefully acknowledges financial support from the UK Economic and Social Research Council (ESRC), Award reference ES/L009633/1. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. All errors are our own.

Fluctuations in labor input occur along two margins. The extensive margin refers to the formation and termination of employment relationships, whereas the intensive margin describes the choice of working time conditional on being employed. Recent research has largely focused on the extensive margin.[1]

The intensive margin is also active, though. In several European economies, fluctuations in working time per employee are as large as movements in employment (Llosa et al, 2014). In addition, plant-level data from the U.S. shows that working time per person is as variable as employment (Cooper et al, 2015). This variability can seem puzzling in light of earlier findings showing individual working time to be highly inelastic (Hall, 1999; Keane, 2011).

This paper presents a tractable framework for studying working time and earnings of heterogeneous workers within a firm. The firm and its employees join in long-term relationships, bound together by the fact that extensive margin adjustments are costly. In this setting, an employee's working time and earnings reflect her own preferences, the preferences of her colleagues, and the production technology. This framework informs a novel approach to the joint estimation of utility and production function parameters. In addition, our model can shed light on why working time may appear more inelastic in certain circumstances than implied by the utility parameters underlying structural labor supply elasticities.

A key ingredient of the model is that workers are (potentially) complements in production but have heterogeneous preferences over leisure. Under complementarities, variation in a worker's own, *idiosyncratic* labor supply incentives (i.e., leisure preferences) can yield small changes in working time given any utility parameter. Intuitively, the efficient response of working time is attenuated when one's effort is not complemented by higher effort among co-workers. On the other hand, *firm-level* driving forces (i.e., firm productivity or product demand) act, in effect, to coordinate the responses of heterogeneous workers, eliciting the true intertemporal (Frisch) labor supply elasticity. Hence, the model can predict economically significant responses of working time to firm-level variation even if the reaction of individual working time to idiosyncratic events is more tepid.

In Section 1, we characterize the working time, earnings, and employment decisions in the model. An individual's working time is set to maximize the surplus of the firm-worker match. The solution represents a balancing of two forces—production complementarities and preference heterogeneity. The former acts to synchronize working time adjustments, whereas the latter drives them apart. If complementarities are strong enough, less of the dispersion in preferences passes through to working time.

[1] This literature includes search and matching models of unemployment (see Rogerson and Shimer (2011)) and quantitative models of participation (Chang and Kim, 2006; Erosa et al, 2016).

Under complementarities, idiosyncratic variation is instead accommodated by the earnings bargain, which divides the match surplus. If a worker's labor input is hardly adjusted despite an increase in her marginal value of time, she must be compensated for the added disutility. Thus, a telltale sign of strong complementarities is that employees' working time adjustments to preference shocks are compressed *relative to* their earnings growth.

While we have highlighted the role of preferences, the model also allows for a worker-specific component of productivity. Interestingly, complementarities do not drive such a wedge between the elasticities of working time and earnings with respect to idiosyncratic productivity. Thus, the joint dynamics of working time and earnings depend crucially on the mix of idiosyncratic forces. We show how to use the covariance of working time and wages to distinguish between the two shocks, confirming that preference heterogeneity is especially critical to fitting our data.

The model is "closed" by the firm's choice of employment demand. Employment adjustments will tend to crowd out working time, underscoring how an active intensive margin must be supported by frictions on the extensive margin. We consider realistic hiring and firing costs that imply the optimal policy takes an Ss form, which can impart significant inertia to employment (Dixit, 1997).[2] The structure and size of adjustment frictions are based on direct evidence from the labor market we study in our empirical application.

To assess our theory, we turn in Sections 2 and 3 to a rich employee-employer matched dataset. The Veneto Worker History database tracks the universe of workers and firms in the northern Italian region of Veneto from 1982 to 2001. The dataset includes employees' earnings and days worked for each of their employers. Working days is an active margin: in a given year, over 50 percent of workers adjust their days, and among these, the typical change is between 10 and 19 days. Using simple regressions, we can apportion variation in days worked and earnings into firm-wide and idiosyncratic, or work-specific, components. These components have clear counterparts in the model and underlie our structural estimation.

In Section 4, we estimate the model's parameters using the method of simulated moments. Our identification strategy relies on observing the distributions of earnings and working time *inside* firms. A key moment is the variance of working time adjustments within firms relative to the variance of earnings growth (again, inside firms). If this ratio is small, the model infers that preference shocks, manifest in earnings, are being "squeezed out" of working time. Therefore, complementarities must be strong or, more exactly, the elasticity of substitution across jobs is low.

[2] In Dixit (1997), the firm chooses capital and labor. In our case, the factors are distinct types of workers, each of which has a type-specific leisure preference and productivity.

Whereas complementarities can be recovered from within-firm variation, the utility parameter behind the Frisch elasticity is more sharply revealed by *firm-level* fluctuations in working time. Intuitively, an increase in firm productivity elicits higher working time across the board, easing the diminishing returns that a lone worker faces. The remaining deterrent to adjusting working time is the rate of increase in marginal disutility, which can then be inferred from the firm-level response. We find a Frisch elasticity of working time of 0.483. This estimate is near the middle of the range cited by Chetty et al (2011), who review several novel identification strategies that are relatively robust to idiosyncratic (preference) variation. Our relatively more structural approach is complementary to these efforts.

To close Section 4, we simulate policy interventions that draw out the implications of complementarities for labor supply. The policy can be interpreted as a temporary change in the labor income tax rate, though its effect is akin, more simply, to a one-off change in the disutility of labor. The policy is targeted to a fraction of a firm's workforce in one scenario, whereas all workers face it in another. An individual's working time falls twice as much when all her colleagues face the policy, attesting to the importance of complementarities. Conversely, if we used a model without complementarities to interpret results when the policy targets a fraction of workers, we would infer a Frisch elasticity of 0.2, or almost 60 percent smaller than our estimate.

Section 5 evaluates the robustness of our results, with a focus on assessing our measure of working time. Recall that we observe paid days of work rather than total hours. Reassuringly, measures of days and hours worked in household surveys indicate that fluctuations in the former likely account for the vast majority of variation in the latter.

Finally, Section 6 concludes with a few remarks on how our framework could be applied, and extended, to consider other topics where complementarities play a role.

Related literature. Our paper echoes research showing that working time can seem deceptively inelastic if one looks at the "wrong" driving forces. This point has been made in different contexts, including in Imai and Keane's (2004) treatment of on-the-job learning and Rogerson and Wallenius' (2009) model of nonlinear earnings-hours schedules (see also Keane and Rogerson (2012)). In each case, variation in the relevant labor supply incentives (i.e., the present value of work in Imai and Keane) elicits significant responses—the Frisch elasticity is nontrivial—even if working time generally varies (far) less than wages.

Second, our paper is reminiscent of a large literature on hours adjustment frictions. Most recently, Chetty (2012) shows how to bound the Frisch elasticity if the friction can be thought of as a "small" deviation from standard life cycle theory. (Our estimate is at the top end of Chetty's range.) Other analyses are cast in more explicit models of hours constraints in which firms offer fixed hours-wage bundles (Altonji and Paxson, 1988; Dickens and Lundberg, 1993; Chetty et al, 2011).[3] In our model, there are no constraints per se. Rather, coordination of labor supply *emerges* under complementarities.

Finally, there is a smaller literature that specifically considers the interaction of worker heterogeneity and complementarities. In Deardorff and Stafford (1976) and Yurdagul (2017), there is an outside sector (i.e., self employment) to which workers move if hours under complementarities deviate too far from what they would choose on their own. Our model instead highlights the scope for renegotiating earnings (rather than necessarily separating) following idiosyncratic shocks. Indeed, this feature is crucial for interpreting certain moments of the data. In a recent empirical contribution, Labanca and Pozzoli (2022) find that a targeted income tax cut elevates working time only in firms with relatively diffuse hours, e.g., where complementarities appear weak. Our model also predicts that narrowly targeted policies may have little impact if there is a strong incentive to coordinate effort.[4]

1 Theory

1.1 An illustration

It may be helpful to first sketch a simplified version of the optimal working time problem that can still convey the essential message of the paper. We will relax a number of the restrictions later in this section.

In this labor market, firms and workers are heterogeneous. Firms differ with respect to productivity. Workers have heterogeneous preferences over leisure, or, more broadly, different marginal values of time. Any such idiosyncratic variation across workers is a force for diffusion in their labor inputs.

At the firm level, however, suppose that production (potentially) requires the coordination of effort across workers. To formalize this notion, imagine that a firm's output is produced by the execution of a fixed number N of jobs. For simplicity, we treat the firm's workforce as given, and assume that each worker performs one job, e.g., the workforce is also

[3] Whereas hours constraints are generally thought to apply at the firm level, Rogerson (2011) considers a macroeconomic model in which workers at all firms work a uniform level of hours.

[4] Labanca and Pozzoli report that *changes* in hours are also more compressed in firms where the *levels* of hours are less diffuse. We will tend to emphasize changes in working time, because complementarities cannot generally be identified by differences in levels alone, as we discuss in Section 3.1.

N. The firm's output, Γ, is assumed to be given by

$$\Gamma = ZG(\mathbf{h}) = Z\left(\sum_{j=1}^{N} h(j)^{\rho}\right)^{1/\rho}, \tag{1}$$

where Z is firm productivity, $h(j)$ is employee j's working time, and $\mathbf{h} \equiv \{h(j)\}$ is an $N \times 1$ vector. The key structural parameter in (1) is ρ, which determines the elasticity of substitution across jobs. A value of $\rho = 1$ implies jobs are perfect substitutes whereas $\rho = -\infty$ implies perfect complements. Note that (1) exhibits constant returns with respect to $\mathbf{h}$, but diminishing returns with respect to any individual $h(j)$.

Assume a worker's *marginal* disutility of effort has the form, $\xi(j)\, h(j)^{\varphi}$, where $\varphi > 0$ and $\xi(j)$ encompasses any shift in the marginal value of time of the worker who performs job j. For instance, ξ would rise if a worker is needed at home to care for a family member. The utility parameter, φ, is a key object of interest in our paper. To convey the meaning of our findings for the broader literature, we will refer to $1/\varphi$ as *the implied Frisch elasticity* (even though workers are not wage-takers in our model, as we will see).

The firm and its workers choose an allocation of time $\{h(j)\}_{j=1}^{N}$. We suppose the parties bargain to the efficient outcome whereby each worker's marginal value of time (outside the firm) is equated to her marginal product (inside the firm): $\xi(j)\, h(j)^{\varphi} = Z\partial G/\partial h(j)$.[5] Optimal labor input therefore satisfies

$$h(j) = (\Omega(N)\, Z)^{1/\varphi}\, \xi(j)^{-\frac{1}{\varphi+1-\rho}}, \tag{2}$$

where $\Omega(N) \equiv \left(\sum_{i=1}^{N} \xi(i)^{-\frac{\rho}{\varphi+1-\rho}}\right)^{\frac{1-\rho}{\rho}}$.

Equation (2) imparts an important lesson about how to identify the (implied) Frisch elasticity, $1/\varphi$. Consider first the response of working time to a change in $\xi(j)$. In general, this elasticity depends on utility (φ) *and* production (ρ) parameters. The response of $h(j)$ unambiguously reflects φ only in the special case of perfect substitutes, $\rho = 1$. If j is complementary to other jobs, a change in $\xi(j)$ may instead have little impact on $h(j)$. Indeed, as ρ declines toward $-\infty$, the response of $h(j)$ becomes increasingly attenuated *for any* φ. Intuitively, if complementarities are strong, diminishing returns to one's own working time sets in rapidly. Even a slight increase in $h(j)$ following a reduction in $\xi(j)$ can be sufficient to drive j's marginal product into line with the lower marginal cost of effort.

Whereas the influence of the Frisch elasticity is obscured in working time variation driven

[5] In the first order condition, the marginal value of wealth is subsumed under $\xi(j)$. See Section 1.3.

by $\xi(j)$, it clearly shapes the behavior of $h(j)$ in the presence of *firm-level* shocks, Z. When an employee adjusts her working time to a higher Z, her colleagues will match her higher effort, erasing the force of diminishing returns. The remaining deterrent to ramping up working time is the rate of increase in the disutility of effort. Thus, the elasticity of working time depends only on $1/\varphi$. This paper emphasizes that this distinction between idiosyncratic (e.g., ξ) and firm-wide (Z) variation can help shed light on labor supply dynamics.

In what follows, we expand on this set-up along a number of dimensions. First, we demonstrate how to identify φ and ρ in the presence of both worker-specific preference *and* productivity differences. Second, we solve for an earnings bargain and show that, while shifts in $\xi(j)$ are weakly passed through to changes in working time, they are more clearly manifest in the dispersion of earnings changes within the firm. Third, we endogenize employment, N, by introducing decreasing returns to scale at the firm level. We show how working time is shaped by both the returns to scale as well as frictions on employment adjustment.

1.2 The environment

We now describe in detail workers' preferences, firms' production technology, and the structure of the labor market.

Preferences. Utility is separable in consumption and leisure. In line with Section 1.1, the disutility from time worked h is given by

$$\xi\nu(h) \equiv \xi\frac{h^{1+\varphi}}{1+\varphi}, \tag{3}$$

where, to recall, $\varphi > 0$ and ξ represents fluctuations in the worker's marginal value of time.

In general, shifts in ξ can impinge on consumption. To avoid this complication, we assume each individual belongs to one of many large families (Merz, 1995). By pooling members' earnings, a family insures consumption against member-specific risk. The flow value of working will not depend directly on the degree of risk aversion but only on earnings and the cost of supplying labor, $\xi\nu(h)$ (see (5) below and Appendix C).

To preserve tractability, we make several simplifying assumptions concerning ξ. First, ξ is i.i.d. across time and workers: at the start of each period, each worker selects anew a ξ from a K–dimensional set, $\mathcal{X} \subseteq \mathbb{R}^K$. Second, we assume types are drawn after hires have been made, but types are perfectly observed thereafter. Accordingly, firm and worker can contract (earnings and working time) on ξ. Finally, we assume labor is divisible and so (ab)use a law of large numbers (Uhlig, 1996) to eliminate any "noise" in the distribution of types within firm: a deterministic share $\lambda_\xi \in (0,1)$ of a firm's workforce draws type $\xi \in \mathcal{X}$, where $\sum_{\xi\in\mathcal{X}}\lambda_\xi = 1$ and $\frac{1}{K}\sum_{\xi\in\mathcal{X}}\xi$ is normalized to 1.

We revisit these assumptions later in the paper, but a few remarks now are worthwhile. First, the absence of persistence in ξ will have no direct impact on working time or earnings. As we will see, optimal working time is a static condition, and the earnings bargain will take the same form when types are persistent. (We consider implications for the extensive margin in Section 5.) Second, the assumption of perfect information is clearly stylized, but theories of hidden information are not necessarily consistent with our moments. For instance, pooling equilibria under hidden information involve a compression of working time *and* earnings (Levin, 2003), whereas the *relative* variability of earnings is a key feature of the data that emerges naturally within our framework. Finally, the assumption of divisible labor is made to aid analytical tractability.

Production. Whereas (1) assumes each worker performs a unique job, we now suppose that each *type* is assigned a unique *set* of jobs. The organization of production across a discrete number of types allows us to carry over the basic structure of (1) even when labor is divisible. Relative to (1), we also incorporate worker productivity heterogeneity.

A type is now a pair of preference ξ and productivity θ levels. Analogously to preferences, we assume productivity is i.i.d. and drawn from a L–dimensional set $\mathcal{Y} \subseteq \mathbb{R}^L$.[6] A fixed share λ_θ of the initial workforce will have productivity θ, so the fraction with pair (ξ, θ) is $\lambda_{\xi,\theta} = \lambda_\xi\lambda_\theta$. Note that, in total, there are $M \equiv K \times L$ pairs or types.

The efforts of heterogeneous types are combined to produce final output. Total labor input of a type (ξ, θ) is $n_{\xi,\theta}h_{\xi,\theta}$, where $n_{\xi,\theta}$ is the measure of that type's employment and $h_{\xi,\theta}$ is the average supply of time among workers of that type. The type-specific labor inputs are aggregated via a CES production function,

$$\Gamma = ZG(\mathbf{h}, \mathbf{n}) = Z\left(\sum_{\xi\in\mathcal{X}}\sum_{\theta\in\mathcal{Y}}(\theta n_{\xi,\theta}h_{\xi,\theta})^\rho\right)^{\alpha/\rho}, \tag{4}$$

where ρ again reflects the elasticity of substitution across jobs; $\mathbf{n} \equiv \{n_{\xi,\theta}\}$ and $\mathbf{h} \equiv \{h_{\xi,\theta}\}$ are $M \times 1$ vectors; and $\alpha \in (0, 1)$ is the returns to scale.[7] The departure from constant returns

[6] Chang et al (2020) instead consider correlated draws of preference and productivity shocks in order to match their moments on working time and wages. We can fit the comovement of these variables when ξ and θ are i.i.d. (see Section 4.2).

[7] We continue to interpret $(1-\rho)^{-1}$ as the elasticity of substitution across *jobs*, not across types per se. A simple example illustrates how (4) can "inherit" ρ from a more primitive production function. Firm output Γ is an aggregate over jobs $j \in [0, 1]$, $\Gamma = Z\kappa\left[\int_0^1 \gamma(j)^\rho \,\mathrm{d}j\right]^{\alpha/\rho}$, where $\kappa \equiv K^{(1-\rho)\alpha/\rho}$ is a normalizing constant and output of job j, $\gamma(j)$, is proportional to total person-hours. If types are allocated an equal share, $1/K$, of jobs, then $\gamma(j) = Kn_{\xi,\theta}h_{\xi,\theta}$ for any j performed by type (ξ, θ). Substituting for $\gamma(j)$ in Γ yields (4).

$\alpha < 1$ ensures a well-defined notion of firm size, $N \equiv \sum_{\xi,\theta} n_{\xi,\theta}$. Note that since $\alpha < 1$, the limiting case of perfect substitutes refers to $\rho = \alpha$, e.g., ρ must satisfy $\rho \leq \alpha$. Finally, recall that Z is firm productivity, which will follow a first-order Markov process.

Equation (4) is a reduced-form structure, but we believe it gets the "big picture" right. Specifically, it captures the notion that the production of final output requires different jobs to be performed by different workers, each of whom faces her own idiosyncratic circumstances (e.g., ξ). In reality, the fineness of this division of labor surely reflects more primitive forces, such as the costs of coordinating activities across jobs (Becker and Murphy, 1992) and training workers on complex tasks (Costinot, 2009). Nevertheless, these factors are likely to change slowly as new ideas and technologies are gradually developed, whereas we will be interested in year-to-year fluctuations in earnings and working time. Thus, a more explicit microfoundation of (4) would not necessarily alter our characterization of the basic trade-off between coordinating working time and accommodating heterogeneous preferences.

Labor market frictions. Labor market frictions play a subtle but crucial role in the model. The costs of forming, and dissolving, matches imply a surplus to ongoing firm-worker relationships, creating scope for bargaining over earnings and working time. Following Roys (2016), we assume there is a matching friction such that neither firm nor worker can instantaneously replace the other. We assume that, in the firm's problem (see below), the matching friction is channeled entirely through the (broader) cost of hiring a worker, denoted by $\bar{c}$. For instance, if it takes longer to fill a job, the cost of recruiting is higher. Thus, conditional on $\bar{c}$, we do not need to elaborate further on matching. In addition, we assume there is a cost $\underline{c}$ of firing a worker, which can be interpreted in our empirical application as mandated severance.

1.3 Characterization

This section characterizes the choices of working time, earnings, and employment.

1.3.1 Firm and worker objectives

Workers. Consider the surplus from working as type (ξ, θ) at a firm with productivity Z and workforce $\mathbf{n}$. In the current period, the employee receives a return equal to earnings, $W_{\xi,\theta}(\mathbf{n},Z)$, less the disutility of effort, $\xi\nu(h_{\xi,\theta}(\mathbf{n},Z))$ and the value, μ, of non-market time. Next period, productivity, Z', is realized, and the worker draws a (potentially new) type (x, y). If the match surplus is negative given $\mathbf{n}$, separations occur at rate $s_{x,y}$. Putting these pieces together, the surplus from working, $\mathcal{W}_{\xi,\theta}(\mathbf{n},Z)$, is

$$\mathcal{W}_{\xi,\theta}(\mathbf{n},Z) = \begin{array}{c} W_{\xi,\theta}(\mathbf{n},Z) - \xi\nu\left(h_{\xi,\theta}(\mathbf{n},Z)\right) - \mu \\ +\beta\sum_{x,y}\lambda_{x,y}\mathbb{E}\left[\left(1 - s_{x,y}(\mathbf{n}',Z')\right)\mathcal{W}_{x,y}(\mathbf{n}',Z') \mid Z\right]. \end{array} \tag{5}$$

A few remarks on (5) are warranted. First, in the absence of data on non-market incomes and activities, we simply treat μ as a fixed parameter. (See Boerma and Karabarbounis (2021) on identification of heterogeneous non-market values with time use data.) Second, the large-family assumption means that the marginal value of wealth is invariant to worker- and firm-specific shocks and is, therefore, suppressed in (5).[8]

The firm. The firm has an initial workforce N_{-1}. (A subscript $_{-1}$ denotes a one-period lag, and a prime $'$ denotes next-period values.) After productivity, Z, is realized, the firm may choose to hire. We assume a worker's type (ξ,θ) is unknown at the time of hire. After hires (if any) are made, the firm's workforce is denoted by $\mathcal{N}$. Then, all $\mathcal{N}$ workers draw a type, and the firm and some of its workers may choose to separate. Let $n_{\xi,\theta}$ be the measure of type-(ξ,θ) workers retained. It follows that $N = \sum_{\xi,\theta} n_{\xi,\theta}$ is the measure of the workforce used in production and "carried into" next period. Wages and time worked will be determined after separations (if any) are made.

It is helpful to first define the present value of a firm for a given allocation, $\mathbf{n} \equiv \{n_{\xi,\theta}\}$. Let π stand for profit *gross* of firing and hiring costs,

$$\pi(\mathbf{n},Z) \equiv ZG\left(\mathbf{h}(\mathbf{n},\mathbf{Z}),\mathbf{n}\right) - \mathbf{n}^T\mathbf{W}(\mathbf{n},Z),$$

where $\mathbf{n}^T$ is the transpose of $\mathbf{n}$ and $\mathbf{W}$ is the vector of earnings over types, $\mathbf{W} \equiv \{W_{\xi,\theta}\}$. The corresponding present value of the firm is then

$$\tilde{\Pi}(\mathbf{n},Z) \equiv \pi(\mathbf{n},Z) + \beta\int \Pi(N,Z')\,\mathrm{d}F(Z'|Z), \tag{6}$$

where $\beta \in (0,1)$ is the discount factor, F is the distribution function of productivity $Z'|Z$, and Π is the continuation value.

Critically, Π can be written as a function of just two state variables, (N,Z'), despite the heterogeneity across workers. This tractability is purchased by the assumption of i.i.d. types, which implies that we do not have to track individual types of workers over time.

[8] Note that under *in*complete insurance, persistent shocks to Z would shift the marginal value of wealth. Thus, the substitution effect (of a change in Z) would be partially offset by an income effect. In this case, the model would "need" a higher Frisch elasticity to match the variance of firm-wide working time.

The dynamic programming problem may now be written as follows. Consider, first, the problem for a given $\mathcal{N}$. The firm's problem at this stage is to decide separations, and is characterized by the Bellman equation,

$$\Pi^{-}(\mathcal{N}, Z) = \max_{\mathbf{n}} \left\{ \tilde{\Pi}(\mathbf{n}, Z) - \underline{c} \cdot \sum_{\xi,\theta} \left[\lambda_{\xi,\theta}\mathcal{N} - n_{\xi,\theta}\right] \right\}, \tag{7}$$

subject to $\lambda_{\xi,\theta}\mathcal{N} \geq n_{\xi,\theta}$ for each type (ξ, θ). Then, step back one stage and consider the choice of hires, which brings the workforce up to a level, $\mathcal{N}$. Since hires are anonymous, the value of the firm at this stage is

$$\Pi(N_{-1}, Z) = \max_{\mathcal{N}} \left\{ -\bar{c} \cdot [\mathcal{N} - N_{-1}] + \Pi^{-}(\mathcal{N}, Z) \right\}, \tag{8}$$

subject to $\mathcal{N} \geq N_{-1}$.

Note that (7)-(8) imply that a firm may hire and separate workers in the same period. However, for realistic values of $\underline{c}$ and $\bar{c}$, this will not happen: Z must be quite low to warrant separations, in which case no hires are made. Thus, at firms that separate, $\mathcal{N} = N_{-1}$, and at firms that hire, $\mathbf{n} = \boldsymbol{\lambda}\mathcal{N} = \boldsymbol{\lambda}N$.

1.3.2 Working time

The firm and each of its workers jointly choose working time efficiently by equating the employee's marginal disamenity to the marginal value of his time to the firm. The symmetry of workers within a type (ξ, θ) implies that their working time and earnings will be equal. Specifically, solving this first order condition yields the following result.

Proposition 1 *For any individual worker of type (ξ, θ), the efficient choice of working time is given by*

$$h_{\xi,\theta} = (\alpha Z \Omega(\mathbf{n}))^{\frac{1}{\varphi+1-\alpha}} \cdot \left[\theta^{\rho} n_{\xi,\theta}^{\rho-1}/\xi\right]^{\frac{1}{\varphi+1-\rho}},$$
$$\textit{with} \;\; \Omega(\mathbf{n}) \equiv \left(\sum_{x\in\mathcal{X}} \sum_{y\in\mathcal{Y}} \left[y^{\varphi+1} n_{x,y}^{\varphi}/x\right]^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha-\rho}{\rho}}. \tag{9}$$

Equation (9) indicates that the elasticity of $h_{\xi,\theta}$ with respect to Z is $1/(\varphi + 1 - \alpha)$. As we saw in Section 1.1, this is decreasing in φ, or, equivalently, increasing in the Frisch elasticity. What is new is the role of the returns to scale, α. The elasticity of working time is higher at higher α, which signifies that diminishing returns sets in slowly. In what follows, we treat α as known and later parameterize it based on external evidence, as we lack much of the data (revenue, etc.) that would typically be used to estimate it.

Importantly, since the elasticity of $h_{\xi,\theta}$ with respect to Z is independent of type, it applies to all types. In other words, the elasticity of *mean* working time with respect to Z is also $1/(\varphi+1-\alpha)$. This is a simple but crucial point for estimation: although $h_{\xi,\theta}$ is unobserved, variation in average working time within the firm can be measured and informs the identification of φ.

Equation (9) also reveals the role of complementarities in shaping the reaction of working time to idiosyncratic events. To see this most clearly, consider a measure of workers of type (ξ,θ) that is "small" relative to the size of the firm (such that spillovers across types via Ω may be ignored for now).[9] The corollary reports the effects of a perturbation to ξ and θ.

Corollary 1 *(I) The elasticity of working time with respect to* ξ, *given by* $-1/(\varphi+1-\rho)\leq 0$, *tends to zero as* $\rho\to-\infty$ *. (II) The elasticity of working time with respect to* θ, *given by* $\rho/(\varphi+1-\rho)$, *is bounded above by* $\alpha/(\varphi+1-\alpha)>0$ *and approaches* -1 *as* $\rho\to-\infty$.

There are several aspects of Corollary 1 that deserve attention. First, the responses of working time to changes in ξ and Z coincide only if $\rho=\alpha$, which implies that tasks are perfect substitutes. Otherwise, working time adjustments to ξ are attenuated. Indeed, as we saw in Section 1.1, working time is increasingly invariant to ξ as $\rho\to-\infty$, e.g., in the limit where tasks are perfect complements. Crucially, this invariance emerges *regardless* of the value of the Frisch elasticity, $1/\varphi$.

By contrast, the elasticity of working time to θ does not vanish as $\rho\to-\infty$. The reason is that, unlike a shift in ξ, a change in productivity, θ, has a *direct* effect on the worker's output. As a result, in the $\rho\to-\infty$ limit, $\ln h_{\xi,\theta}$ must move virtually one for one to offset shifts in $\ln\theta$ (and, thus, stabilize $\theta h_{\xi,\theta}$), or else the type's marginal product would change precipitously. Still, even though working time is responsive in this limit, the change in $h_{\xi,\theta}$ becomes almost entirely detached from φ.

A final aspect of (9) concerns the relationship between the extensive and intensive margins. One can show that working time is declining in own-type employment but increasing in the employment of other types. These properties follow from concavity of the production function with respect to labor input $(n_{\xi,\theta}h_{\xi,\theta})$ of any one type (ξ,θ) and supermodularity with respect to any two types. In addition, if there is a balanced expansion of employment of all types (e.g., $\Delta\ln n_{\xi,\theta}$ is identical $\forall\ (\xi,\theta)$), $\alpha<1$ implies that $h_{\xi,\theta}$ falls $\forall\ (\xi,\theta)$.

[9] As a result, the comparative static casts a sharper light on the roles of φ and ρ. In general, the elasticities also depend on the size of the group with type (ξ,θ), e.g., the $n_{\xi,\theta}$s in Ω. See Appendix C. In our estimated model, these scale effects are accounted for since the number of types is finite.

Proceeding, the solution to working time (9) enables us to concentrate $\mathbf{h}$ out of the firm's problem. Substituting (9) into the revenue function (4) yields

$$ZG(\mathbf{h},\mathbf{n}) = \hat{G}(\mathbf{n},Z) \equiv \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\frac{\alpha}{\alpha-\rho}\frac{\varphi+1-\rho}{\varphi+1-\alpha}}. \tag{10}$$

Accordingly, period profit can be written as $\hat{\pi}(\mathbf{n},Z) \equiv \hat{G}(\mathbf{n},Z) - \mathbf{n}^T\mathbf{W}(\mathbf{n},Z)$. We now turn our attention to determining earnings $\mathbf{W}(\mathbf{n},Z)$ and employment $\mathbf{n}$.

1.3.3 Earnings

Following Cahuc et al (2008), earnings are determined by splitting the marginal match surplus, awarding a share, $\eta \in (0,1)$, to the worker. One can motivate marginal surplus splitting using the bargaining solution proposed by Stole and Zwiebel (1996) and developed rigorously in Brügemann et al (2019).

The marginal surplus is the sum of the worker's surplus, $\mathcal{W}_{\xi,\theta}(\mathbf{n},Z)$, and the firm's surplus, which has two parts. The first, denoted by $\mathcal{J}_{\xi,\theta}(\mathbf{n},Z)$, is the marginal value of type-(ξ,θ) labor *gross* of adjustment costs. This term is obtained by differentiating the value function (6) (swapping out π for $\hat{\pi}$) with respect to $n_{\xi,\theta}$,

$$\mathcal{J}_{\xi,\theta}(\mathbf{n},Z) \equiv \frac{\partial}{\partial n_{\xi,\theta}}\hat{\pi}(\mathbf{n},Z) + \beta \int \Pi_N(N,Z')\,\mathrm{d}F(Z'|Z),$$

where $N = \sum_{\xi,\theta} n_{\xi,\theta}$. In addition, the firm's surplus accounts for the penalty $\underline{c}$ that the firm incurs if an agreement is not reached, resulting in the worker's separation. Accordingly, the surplus from retaining a worker is $\mathcal{J}_{\xi,\theta}(\mathbf{n},Z) + \underline{c}$, and earnings solve

$$\mathcal{W}_{\xi,\theta}(\mathbf{n},Z) = \eta\left(\mathcal{W}_{\xi,\theta}(\mathbf{n},Z) + \mathcal{J}_{\xi,\theta}(\mathbf{n},Z) + \underline{c}\right). \tag{11}$$

Our solution of (11), reported below, generalizes Cahuc et al (2008) to an environment with endogenous separations and an intensive margin.

Proposition 2 *The earnings bargain for a worker of type (ξ,θ) is given by*

$$W_{\xi,\theta}(\mathbf{n},Z) = \eta\left(\kappa\frac{\partial\hat{G}(\mathbf{n},Z)}{\partial n_{\xi,\theta}} + r\underline{c}\right) + (1-\eta)\left(\kappa\xi\nu\left(h_{\xi,\theta}(\mathbf{n})\right) + \mu\right), \tag{12}$$

where $r \equiv 1-\beta$; $h_{\xi,\theta}(\mathbf{n})$ solves (9); and $\kappa \equiv \frac{\varphi+1-\alpha}{(1-\eta(1-\alpha))(\varphi+1)-\alpha} > 1$ with $\eta\kappa < 1$ and $(1-\eta)\kappa < 1$.

The bargain in (12) is a weighted average of a (i) worker's contribution to the firm and (ii) the value of his non-market time. The former (i) depends on the worker's marginal product, $\partial \hat{G}(\mathbf{n}, Z) / \partial n_{\xi,\theta}$, and the annuitized firing cost, $r\underline{c}$, which he "saves" the firm by continuing the match for another period.[10] The latter (ii) also has two parts: outside of employment, a worker does not incur the disutility $\xi\nu(h_{\xi,\theta}(\mathbf{n}))$, and he avails himself of μ. Notably, solutions of certain *collective* bargaining games also admit a role for individual heterogeneity in the value of non-market time, which is a crucial element of (12) (Taschereau-Dumouchel, 2020). This observation suggests that our theory may serve as a useful framework even when the structure of bargaining is more elaborate, e.g., there are collective and individual elements to it.[11]

1.3.4 Comparing earnings and working time dynamics

Several of the model's key implications for the joint dynamics of earnings and working time can be gleaned from (9) and (12). To this end, it is helpful to write out earnings more explicitly using (9) and (10),

$$W_{\xi,\theta}(\mathbf{n}, Z) = \varkappa \alpha^{\frac{\alpha}{\varphi+1-\alpha}} (Z\Omega(\mathbf{n}))^{\frac{\varphi+1}{\varphi+1-\alpha}} \left(\frac{\theta^{\varphi+1}}{\xi}\right)^{\frac{\rho}{\varphi+1-\rho}} n_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} + \omega, \tag{13}$$

where $\varkappa \equiv \alpha^{\frac{\eta\varphi+(1-\eta)\frac{\varphi+1-\alpha}{\varphi+1}}{(1-\eta(1-\alpha))(\varphi+1)-\alpha}} \in (0,1)$ is increasing in η and $\omega \equiv \eta r\underline{c} + (1-\eta)\mu$. Thus, $\partial\hat{G}/\partial n_{\xi,\theta}$ and $\xi\nu(h_{\xi,\theta})$ in (12) can be collected into a single term summarizing the variable portion of earnings. Note that a higher bargaining power η redistributes weight toward this term (since r in ω is small), amplifying the effect of firm-level and idiosyncratic shocks.

Equation (13) clarifies the mapping between idiosyncratic events and earnings. Consider, for instance, an increase in the distaste for working, ξ, and suppose jobs are complementary ($\rho < 0$). Since the response of working time is suppressed, workers earn a premium for supplying costly effort, noting the term $\xi^{-\rho/(\varphi+1-\rho)}$ in (13). Indeed, earnings become increasingly responsive as $\rho \to -\infty$. Thus, earnings should be more elastic than working time with respect to ξ, at least if ω, the fixed portion of earnings, is not too large. The following corollary makes this intuition precise.[12]

[10] The worker can use $\underline{c}$ to negotiate a higher wage because the firm is subject to the severance cost as soon as he is hired. This is consistent with the labor contract that was most prevalent in Italy in our sample. See Mortensen and Pissarides (1999) for a discussion of bargaining under severance costs.

[11] To be clear, "collective" here refers to bargaining between a union and a *firm*. See Section 2.1 for more on the structure of bargaining in Italy.

[12] To simplify the presentation, we again suppose the mass of workers with (ξ, θ) is "small". Results similar to Corollaries 2 and 3 obtain in the general case with "large" cohorts. See Appendix C.

Corollary 2 *The (absolute) elasticity of earnings with respect to* ξ, $|\partial \ln W_{\xi,\theta}/\partial \ln \xi|$, *increases relative to the (absolute) elasticity of working time,* $|\partial \ln h_{\xi,\theta}/\partial \ln \xi|$, *as* ρ *falls away from zero and strictly exceeds the latter whenever* $\rho < -(1-\omega/W_{\xi,\theta})^{-1} < -1$.

Corollary 2 holds out the possibility of using data on earnings and working time to infer ρ. If earnings are indeed relatively sensitive to idiosyncratic variation, we should see that the *within-firm* variance of earnings growth exceeds that of working time changes. Corollary 2 indicates that this excess variance of earnings growth is monotonically decreasing in ρ.

The corollary refers only to a perturbation to ξ, however. Unlike ξ, which directly shifts $W_{\xi,\theta}$ via $\xi\nu(h_{\xi,\theta})$, the impact of productivity θ on earnings is channeled more through its effect on working time. In fact, insofar as ρ shapes the elasticity of $h_{\xi,\theta}$ with respect to θ, one can show that ρ affects the response of earnings symmetrically. Thus, the change in earnings *relative to* the change in working time will not depend on ρ.[13]

Clearly, then, the mix of ξ and θ will be critical for interpreting the volatility of earnings and working time. To identify the predominant source of variation (between ξ and θ), we can examine the comovement of the wage rate, $w_{\xi,\theta} \equiv W_{\xi,\theta}/h_{\xi,\theta}$, and working time, $h_{\xi,\theta}$. A change in ξ shifts earnings and working time in *different* directions, which implies that $w_{\xi,\theta}$ and $h_{\xi,\theta}$ move opposite one another.[14] By contrast, changes in θ push earnings and working time in the *same* direction, which means the wage shifts in this direction only if $W_{\xi,\theta}$ moves more than $h_{\xi,\theta}$. This result will obtain as long as there is sufficient curvature in earnings with respect to θ—e.g., φ is large enough (see (13))—and/or the fixed portion of earnings, ω, is small enough. Corollary 3 summarizes this discussion.

Corollary 3 *(I) The responses of working time and the wage to changes in* ξ *are, unambiguously, of the opposite sign. (II) A change in* θ *shifts working time and the wage in the same direction as long as* $\varphi > \left((\omega/W_{\xi,\theta})^{-1} - 1\right)^{-1}$.

The comovement of working time and the wage rate places restrictions on the prevalence of the different shocks. The negative correlation of working time and wages in our data (see Section 3) can be accommodated by variation in ξ *for any* values of ω and φ. By contrast, changes in θ will drive wages and working time in opposing directions only if φ is small, which means the Frisch elasticity is large. Indeed, since $\omega/W_{\xi,\theta} < 1/2$ at our estimated parameter values, the Frisch elasticity would have to exceed one—a claim that our data on (firm-level) working time will not support. Thus, we infer that ξ drives a substantial share

[13]This statement is established in the proof of Corollary 3. See Appendix A.

[14]To be more exact, $W_{\xi,\theta}$ and $h_{\xi,\theta}$ necessarily go in opposing directions if $\rho < 0$. Otherwise, they can move together, but the wage *always* moves opposite working time.

of the movement in earnings and working time, and, in this environment, Corollary 2 implies that the relative variability of earnings conveys critical identifying information about ρ.

1.3.5 Employment demand

Thus far, we have taken employment as given. To complete our analysis, we now describe the solution to the dynamic employment demand problem (see Appendix A for details).

Consider the problem of a firm of initial size N_{-1}. Separation from a given type is optimal if that type's marginal value of labor falls below $-\underline{c}$, where $\underline{c}$ is the cost of termination. Appendix A shows that this happens when firm productivity falls beneath a (type-specific) threshold value of Z. Naturally, the type of worker separated first is the type with the highest threshold, which is denoted by $\hat{Z}_1(N_{-1})$.[15]

If Z falls further, the firm separates from a second type. As it does this, separations from the first type continue. This result reflects the supermodularity of the firm's problem: as the firm reduces labor input of the second type, the marginal value of the first type falls further. The same idea applies when the firm separates from its third type and so forth.

The next piece of the optimal policy is the decision to hire. Given an initial size N_{-1}, the firm hires when Z rises above a threshold $\hat{Z}_0(N_{-1})$ at which point the marginal value of labor exceeds the marginal cost $\bar{c}$. Since the firm hires before types are known, it simply chooses N, and each type's size rises in proportion to its share in the population.

Figure 1 illustrates the labor demand policy with four types, e.g., there are two levels of preferences and productivities. Where $Z > \hat{Z}_0(N_{-1})$, the firm hires, and each type's employment is increased equally. Between $\hat{Z}_0(N_{-1})$ and $\hat{Z}_1(N_{-1})$, N is held at N_{-1}. This space is a *region of inaction* in which shifts in Z do not push the marginal value below $-\underline{c}$ or above $\bar{c}$. As Z falls below $\hat{Z}_1(N_{-1})$, employment of one type is reduced. As Z declines further, another type is separated jointly with all of the other types that were separated prior to it.

A final issue concerns the implications of employment demand for the *intensive* margin. Clearly, in the inaction region where $N = N_{-1}$, working time fully absorbs the effects of changes in Z. Outside of this region, changes in Z have direct and indirect effects on working time. For given employment, the direct effect of, say, a higher Z is stimulative. The indirect effect is channeled through employment adjustments, which dampen the reaction of working time (see Section 1.3.2). However, the direct effect dominates in the calibrated model because the frictions $\bar{c}$ and $\underline{c}$ curtail the *size* of the adjustments to employment. Thus, an increase in Z elevates both working time and employment (of all types).

[15] In certain cases, one can infer this type's *identity*, e.g., its ξ and θ. See Appendix C.

2 Data source

Our data span the universe of private firms in Veneto, Italy during 1982-2001. Located in the North East, Veneto is one of the largest and richest areas of Italy. Among the country's 20 regions, Veneto ranked sixth in income per capita and fifth in population in our sample period, according to data from Istat. We make a case here that Veneto is a reasonable testing ground for our theory and then describe our data, the Veneto Work History (VWH) files.

2.1 Institutional context

As we take model to data, a potential concern for us is the role of national unions in Italy. Unions could, in principle, suppress contracting between firm and worker, which is a pillar of our theory. However, the *de facto* setting of wages and working time is broadly supportive of our modeling approach.

In the North East in particular, decision-making has been reasonably decentralized at the margin. The multi-tiered wage-setting process illustrates this point well. First, union-negotiated sector-wide contracts specify minimum wages, but in high-wage regions like Veneto, these rarely bind (Card et al, 2014). One tier down, workers' representatives at the firm negotiate "add-ons" to sector-wide contracts. These firm-level agreements are common among larger employers—half of firms with at least 20 workers have one—and the average premium (over industry minima) is about 25 percent (Card et al, 2014; Guiso et al, 2005). Notably, firm-level negotiations in Veneto could be mediated by a self-organized committee of employees rather than union representatives. This observation reflects the relatively light touch of unions in the North East (Cattero, 1989). Finally, management awards bonuses to individual workers (Erickson and Ichino, 1995). Among the many employers with less than 20 workers, where firm-level contracts are less common, these individual premia are substantial—as high as 25 percent—and heterogeneous (Brusco, 1982; Cattero, 1989).

Meanwhile, firms generally enjoyed discretion in negotiating working time, at least among full-time employees. Working time rules, including limits on overtime, were often eased in union agreements or loosely enforced, especially during the 1980s (Treu et al, 1993; Lodovici, 2000). Deviations from full-time, open-ended employment contracts were rare, though: part-time work as well as fixed-term contracts, which could be ended after two years at no cost, were uncommon. Consequently, firms could not use a temporary worker to replace the working time of an employee who draws a positive preference shock (a high ξ).

In this setting, where decision-making is generally diffused, we see unions and other national actors as parameterizing the bargaining process rather than deciding firm-level outcomes. For instance, the interplay between national unions and employers' associations is likely to shape worker bargaining power (η), with the latter taken as an input into firm-level negotiations. We return to this point later when we take up a mid 1990s wage-setting accord that arguably enabled more flexible wage bargains, e.g., a higher η (see Section 5.1).

2.2 The Veneto Work History (VWH) files

Our empirical analysis uses the VWH dataset that has been assembled by researchers at the University of Venice. The data are derived from Italian Social Security records, which track earnings and paid days of work for the purpose of calculating social insurance payments. Nearly every private-sector employee in Veneto is covered by the data; public-sector workers and the self-employed are excluded. The full sample contains 22.245 million worker-year observations over the years 1982-2001.

The VWH data has a number of features that recommend it for this analysis. The Veneto data stand out for reporting a measure of working time, namely, a worker's annual paid days with each of his employers. Using paid days and earnings, we can also compute daily wages. Finally, the VWH specifies the calendar months for which a worker received any earnings from an employer, which enables us to track the worker's tenure with a firm.

Table 1 provides a set of summary statistics for the full sample. Average daily earnings were around 120 euros. The average number of paid months per worker (per year) was 10, and, within a paid month, days of work averaged between 23 and 24. There is a good reason why days paid per month is rather high: as Italy transitioned from six- to five-day weeks, its Social Security agency recorded a full week of work as six days in order to treat five- and six-day weeks equally for pension purposes.

Table 2 zeroes in on moments of the distribution of annual *changes* in paid work days. (These estimates pertain to the subsample of "stayers" used in our baseline analysis below.) While many workers do not adjust their days from one year to the next, 33 percent change the number of days worked by more than 10.[16] Moreover, conditional on changing days, the typical size of the change is between 10 and 19, depending on whether some of the largest adjustments are included.

The VWH's measure of paid days does not necessarily equate to days at work, though the link is reasonably tight. For instance, paid days does include leaves of absence paid *by the*

[16] This inaction could reflect a cost of adjusting hours, which is not captured in the model. Such frictions would imply a nontrivial dynamic choice problem for $h_{\xi,\theta}$. Our conjecture is that, since an adjustment cost will eliminate small changes in $h_{\xi,\theta}$, a higher Frisch elasticity may be needed to generate enough variance in working time *conditional* on adjusting.

firm, but if time off is taken each year (i.e., August vacation), we will still correctly measure changes in days at work. Other absences, such as disability, illness, and parental leave, are typically compensated *by the State*, and our data do not record State-remunerated time off as paid days (Filippi et al, 2002). Likewise, spells of temporary layoff in which workers draw State benefits are not recorded as paid days.

More importantly, the VWH does not capture certain sources of variation in paid time worked. The most prominent omission is daily hours. In Section 5, though, we use household survey data to show that variation in paid days is substantial relative to daily hours, and the conclusions from our analysis largely survive intact.[17]

3 Estimation Strategy

We estimate our model by the Method of Simulated Moments (MSM), which selects values for the parameters to minimize the distance between empirical and model-generated moments. Two broad considerations guide our choice of moments.

First, moments derived from first differences, rather than levels, are more likely to robustly identify the structural parameters. For instance, suppose firms operate a Leontief technology but idiosyncratic productivity (θ) is permanent rather than (as assumed in the model) transitory. Since firms will equate efficiency units across jobs, working times will diffuse within a firm and yet be unresponsive to ξ; it is *changes* in working time that are compressed (indeed, equalized across jobs). Thus, first-differenced data correctly convey the degree of complementarities even if the nature of idiosyncratic productivity is mis-specified.

Second, some parameters interact differently with firm-wide, as opposed to idiosyncratic, shocks. For example, recall the significance of Z, but not ξ, to the identification of φ. Therefore, we want to distinguish firm-wide from within-firm components of changes in working time and earnings. In the next subsection, we illustrate how we do this.

3.1 Earnings and working time

We begin by describing moments pertaining to earnings and working time changes, and then relate these to the structural parameters for which they are especially informative.

Empirical framework. Our empirical analysis centers around a simple regression designed to distinguish variation across workers within a firm from firm-wide movements.

[17] As noted above, the data also do not capture the secular trend toward five-day weeks. However, this development would seem to be orthogonal to the economic forces (i.e., intertemporal substitution) in the model and, thus, to the parameters that shape them. See Appendix B for more.

Consider, first, working time. Let $\Delta \ln h_{ijt}$ denote the log change in days worked by employee i in firm j between years $t-1$ and t. We estimate

$$\Delta \ln h_{ijt} = \boldsymbol{\chi}_{it}\mathbf{C}^h + \phi_{jt}^h + \epsilon_{ijt}^h, \tag{14}$$

where $\boldsymbol{\chi}_{it}$ is a row vector of worker characteristics; $\mathbf{C}^h$ is a conformable (column) vector of coefficients; and ϕ_{jt}^h is a firm-year effect. Equation (14) is applied to a subsample of workers employed at the same firm for years $t-1$ and t (see below for more on sample selection). The elements of $\boldsymbol{\chi}_{it}$ consist of a cubic in tenure (measured as of $t-1$) and the change in broad occupation (between $t-1$ and t).[18] These controls help purge the data of observable heterogeneity in work schedules that is not modeled in our theory. The variation captured in ϕ_{jt}^h and ϵ_{ijt}^h is what is used to estimate the structural model.

The parameter ϕ_{jt}^h captures the mean log change in working time across employees in firm j in year t. The variance of ϕ_{jt}^h is thus our measure of the volatility of *firm-wide* working time. From the model's perspective, it is natural to think of shocks to firm productivity as underlying ϕ_{jt}^h, although the latter can reflect other firm-level forcings (see below).

Meanwhile, the residual in (14), ϵ_{ijt}^h, isolates shifts in working time across workers *within a firm*. Our structural model interprets the variation in ϵ_{ijt}^h as being driven by shocks to idiosyncratic preferences and productivity.

We also estimate a regression of the same form for earnings, which relates the log change in annual earnings, $\Delta \ln W_{ijt}$, to observables ($\boldsymbol{\chi}_{it}$) and firm-year effects ϕ_{jt}^W,

$$\Delta \ln W_{ijt} = \boldsymbol{\chi}_{it}^T\mathbf{C}^W + \phi_{jt}^W + \epsilon_{ijt}^W. \tag{15}$$

The meanings of ϕ_{jt}^W and ϵ_{ijt}^W are analogous to their counterparts in (14).

While we have laid out a simple mapping between the structural shocks and estimates in (14)-(15), the connection between the two is likely more subtle in practice. The reason is that the structural shocks behind the residuals, ϵ^h and ϵ^W, do not generally average out among a finite sample of workers in a firm. The effect of a nonzero mean among idiosyncratic draws must be absorbed by the firm-year intercepts. Thus, a simple (quadratic-form) estimator of the variances of ϕ^h and ϕ^W reflects this finite-sample "noise".[19]

Nevertheless, this estimator can be appropriate in our context. To see why, suppose a firm employs a handful of workers, each of whom takes independent draws of ξ. There may

[18] There are four broad occupations: blue-collar workers, who make up 65 percent of the sample; white-collar non-supervisory workers (31 percent); apprentices (3 percent); and managers (1 percent).

[19] Clearly, the first-best strategy is to explicitly incorporate this noise into the structural model. However, this approach deprives us of the tractability afforded by the law of large numbers (see Section 1.2).

be times when workers draw the same ξ. From the perspective of our model, these common draws of ξ do not exaggerate the variance of firm-level events; rather, they *are* a source of *firm-level* variation insofar as they elicit a common, or coordinated, response of workers' labor supplies.[20] Therefore, the diffusion in ϕ^h (as well as ϕ^W) due to such events is properly treated as firm-level variation.

For completeness, Appendix D estimates the variances of firm-year effects based on the finite-sample adjustment in Kline et al (2020), which eliminates any spillovers of idiosyncratic variation into the ϕs. When the structural parameters are re-estimated to fit these revised moments, we recover a Frisch elasticity (of just under 0.4) that is smaller than our baseline estimate (of 0.48), but not dramatically so.

Sample restrictions. Equations (14)-(15) are estimated off a sample of workers who stay at the same firms in consecutive years. We define stayers in year t as workers who were paid for at least one day in all months of the first quarter of year $t-1$ and in all months of the last quarter of year t. Thus, these workers start and end the two-year period with the same employer. After we remove firms in any year t with only one employee—it would be awkward to analyze complementarities with these firms—we are left with 11.8 million worker-year observations. This is our sample of *2-year stayers*.

While the construction of this sample allows for extended nonwork spells, workers' absences from their employers are generally not re-current. For instance, among workers who are not paid for a full month or more in year $t-1$, most are paid for at least one day in every month of the next year. In this sense, these workers appear to have relatively strong attachments to their firms, which underlies our view that changes in their working time can be interpreted as intensive-margin adjustments.

Still, one could consider a tighter definition of stayers, which requires even more consistent participation at the firm. To this end, we also report figures for an alternative sample, which we refer to as the 12/12 stayers. These workers are paid for at least one day in every month over years $t-1$ and t.

Our restriction to stayers may raise concerns about selection bias. Note, however, that we will also select a sample of stayers from our model-generated data to form the relevant moments for MSM estimation. In this sense, we treat the data and model symmetrically, which supports consistent estimation of the structural parameters (Smith, 1993).[21]

[20] Indeed, one can confirm that, if there is a common component to ξ, an increase in it is isomorphic in (9) to a decrease in firm productivity, Z. In this sense, common components of the idiosyncratic draws are not problematic for the theory; they can be subsumed within our lone firm-level forcing variable.

[21] Of course, stayers may differ from the average worker in unmodeled ways. One possibility is that some jobs are more "critical" to production, and firms are more likely to retain workers in these highly

Regression estimates. Table 3 summarizes estimates derived from (14)-(15). The first three rows pertain to within-firm (idiosyncratic) variation. Specifically, the first row reports var $\left(\epsilon^W\right)$; the second shows var $\left(\epsilon^h\right)$; and the third gives the ratio of the two. In the 2-year stayer sample, this ratio is 2.247—idiosyncratic earnings growth is more than twice as variable as idiosyncratic working time changes. The next three rows report the counterparts to these moments at the firm level, namely var $\left(\phi^W\right)$, var $\left(\phi^h\right)$, and the ratio of the two. Note that, given the mean of days worked in Table 1, a value of var $\left(\phi^h\right) = 0.078^2$ represents 1.5-2 days per month. Variation in working time and earnings among the 12/12 stayers is less pronounced, which is unsurprising: the length of their non-working spells in any one year is relatively abbreviated.

A more extended look at the moment, var $\left(\epsilon^W\right)$ /var $\left(\epsilon^h\right)$, is worthwhile. If this ratio is large, idiosyncratic shocks manifest in earnings are not passed through to working time, signifying that complementarities compress variation in working time. Accordingly, this moment will be critical to our strategy for identifying ρ (see below for further discussion).

Table 4 reports values of var $\left(\epsilon^W\right)$ /var $\left(\epsilon^h\right)$ for several sub-samples.[22] Seen through the lens of the model, working time appears to be compressed for broad classes of workers and sectors. These results underscore that our estimate on the *full* sample is reasonably robust, particularly for the 2-year stayers. The ratio var $\left(\epsilon^W\right)$ /var $\left(\epsilon^h\right)$ falls modestly if we drop public service-oriented sectors such as health care but rises modestly if we restrict attention to men. The ratio is also somewhat higher among larger firms.

Appendix G examines estimates of var $\left(\epsilon^W\right)$ /var $\left(\epsilon^h\right)$ in detailed (three-digit NAICS) industry data. Here, clear differences do emerge, but the patterns further underline the value of this ratio as a diagnostic for complementarities. For example, industry estimates of var $\left(\epsilon^W\right)$ /var $\left(\epsilon^h\right)$ are positively correlated with the incidence of "teamwork" in O*NET data. In addition, the ratio is generally higher in industries with larger male-female earnings differentials, a finding reminiscent of Goldin's (2014) observation that the gender gap is pronounced where working time is less tailored to individual circumstances.

Targeting moments. We now review several moments derived from (14)-(15) and sketch how they enable the identification of structural parameters. (Appendix E uses the sensitivity matrix of Andrews et al (2017) to guide a more extended discussion of identification.) The moments are based on the sample of 2-year stayers.

First, as a preliminary matter, note that there is no ex ante heterogeneity in the model

complementary jobs. By this logic, though, firms should compete hard to fill such jobs when they are vacant, which suggests that job *movers* will also work highly complementary jobs. A priori, then, it is unclear that a worker's status as a stayer or mover reveals the complementarity of her job.

[22] In each case, we use all observations to run (14)-(15) but pool ϵ^W_{ijt} and ϵ^h_{ijt} across the relevant sub-sample.

and, thus, no counterpart to the covariates $\boldsymbol{\chi}$ in (14)-(15). In this setting, the unbiased estimate of the firm-year effect is the mean log change of working time (or, earnings) within the firm. Therefore, this simple average is taken as the model analogue to ϕ_{jt}^{h} (or, ϕ_{jt}^{W}). The deviation of a type's outcome from the mean is the analogue to the residuals in (14)-(15).

Proceeding, the first moment is the variance of firm-year working time effects, or var $\left(\phi^{h}\right)$. As a matter of accounting, this variance reflects the elasticity of average working time to (firm-level) shocks as well as the size of the shocks. In the structural model, the sensitivity of average working time to firm productivity, Z, hinges on the Frisch elasticity, $1/\varphi$. Thus, modulo the size of shocks, the moment var $\left(\phi^{h}\right)$ can be highly informative as to φ. We return shortly to discuss how to infer the variance of Z.

Next, we turn to earnings-related moments. According to (13), log changes in earnings reflect (i) fluctuations in the disutility of effort and marginal product and (ii) the pass through of these changes to earnings. The pass through rate in (ii) is shaped, in part, by bargaining power η and applies to *any* change in disutility and marginal product. Therefore, a higher η amplifies fluctuations in both firm-level and idiosyncratic earnings. Furthermore, the volatility of earnings growth in general should rise relative to the variance of working time growth, since a higher η has no direct impact on working time (see (9)).

In addition, the model predicts that complementarities amplify the changes in disutility in (i) stemming from shifts in ξ. By mitigating the decline in working time following a rise in ξ, stronger complementarities (e.g., a "more negative" ρ) induce a larger increase in disutility, and the *idiosyncratic* element of earnings rises more to compensate a worker for supplying effort. This property underlies a monotone mapping between ρ and the ratio of the variances of residuals, var $\left(\epsilon^{W}\right)$ /var $\left(\epsilon^{h}\right)$, providing a clear way of identifying this parameter. This strategy also "frees up" η to be used to target the relative volatility of *firm-level* earnings growth, that is, the ratio var $\left(\phi^{W}\right)$ /var $\left(\phi^{h}\right)$.

To take stock, our strategy to identify parameters φ, η, and ρ is centered around three moments. One reflects firm-level variation in working time, var $\left(\phi^{h}\right)$. The other two moments pertain to the relative variability of earnings at the firm and worker levels.

Given these moments, it remains to consider the volatility of idiosyncratic working time, or var $\left(\epsilon^{h}\right)$. Clearly, this moment bears on the variances of the two idiosyncratic shocks, denoted here by σ_{ξ}^{2} and σ_{θ}^{2}. To pin down *both* parameters, though, we need to supplement var $\left(\epsilon^{h}\right)$ with information beyond (14)-(15), a point we develop in the next section.

3.2 Additional moments

Next, we summarize three additional moments, and discuss their information content for the model's parameters. Table 5 lists all seven moments used in estimation.

To begin, we regress individual changes in log working time on log changes in *daily earnings*, with the latter given by $\Delta \ln w \equiv \Delta \ln W - \Delta \ln h$. The estimated coefficient of -0.169 echoes studies such as Altonji (1986), who uncovered a coefficient of around -0.3. However, standard life-cycle theory implies that one ought to be able to recover the Frisch elasticity from this regression. Earlier results were thought to reflect measurement error (see Borjas (1980)), but even when instrumental variables were used to eliminate the division bias, the regression returned coefficients that were small and often indistinguishable from zero (MaCurdy, 1981; Altonji, 1986). Our estimates using administrative data, which are far less subject to measurement error, reaffirm that this approach can fail to uncover a significantly positive Frisch elasticity.[23]

From our model's perspective, the regression suggests that the covariance of $\Delta \ln h$ and $\Delta \ln w$ heavily reflects idiosyncratic variation. Specifically, the OLS coefficient points toward a crucial role for ξ. Intuitively, ξ acts as a supply shock pushing working time and the wage in opposite directions, whereas θ is more akin to a demand shock, which induces working time and wages to move together (recall Corollary 3). Thus, the projection of $\Delta \ln h$ on $\Delta \ln w$ is informative as to the size of preference *relative* to the size of productivity shocks. This moment, together with the variance of ϵ^h, helps identify both σ_ξ and σ_θ.

The last two moments refer to employment. The first is the standard deviation of employment growth across firms. Note the latter is employment-weighted so it is representative of the volatility faced by a typical worker. The second moment is mean firm size, $\mathrm{E}[N]$. These two moments are conceptually linked to two parameters in particular. The dispersion of employment growth reflects the size of firm-level shocks, and, thus, anchors the choice of σ_Z, the standard deviation of innovations to firm productivity (see below). The size of firms is strongly influenced by workers' outside option, μ. Intuitively, if μ is small, the rents from a match are large, and so more hires are made.

Finally, we have also examined the model's fit with respect to several *non*targeted moments. These include the persistence of average working time and earnings—the model matches these moments rather well—as well as the latter's correlation with employment, which the model somewhat overstates. See Appendix D for a fuller discussion.

[23] One distinction between our analysis and earlier studies is that we observe daily earnings rather than the hourly wage. We return to this point in Section 5.2.

4 Model Estimation

Seven parameters are estimated. They are ρ, which governs the elasticity of substitution across jobs; the utility parameter, φ; worker bargaining power, η; the worker's outside option, μ; and the standard deviations of the shocks, namely, σ_ξ, σ_θ, and σ_Z. We choose the values of other parameters based on outside evidence. In this section, we first report how we set the latter parameters, and then discuss the estimation results.

4.1 Preliminaries

We start with the firm productivity process, which is assumed to follow a geometric AR(1),

$$\ln Z = \zeta \ln Z_{-1} + \varepsilon_Z, \quad \text{with} \ \ \varepsilon_Z \sim N\left(0, \sigma_Z^2\right).$$

To pin down ζ and σ_Z, one could draw from research that studies total factor productivity (in our data, we cannot). At the same time, these parameters will likely have important implications for some of our moments. Our strategy is to "split the difference": we treat σ_Z as a free parameter but fix $\zeta = 0.8$ (Foster et al, 2008).[24] We opt to estimate the former (σ_Z) because it has a more direct impact on the volatility of working time and earnings, which follow static decision rules.

We pin down four more parameters based on external information. First, the returns to scale are set to $\alpha = 0.67$, which will be consistent, given our estimate of η, with Italy's labor share of around three quarters (ILO and OECD, 2015). Second, we choose a discount factor of $\beta = 0.941$, which is consistent with the average annual real interest rate in Italy over our sample. Third, the hiring cost, $\bar{c}$, is equivalent to 2.5 months of average earnings according to a survey of plants in Veneto's neighboring region of Lombardy (Del Boca and Rota, 1998). Finally, the severance cost, $\underline{c}$, represents a little over 7 months of earnings. The latter is a synthesis of multiple separation costs in Italy (see Appendix H).

Idiosyncratic preferences, ξ, and productivities, θ, are independent discrete random variables. We assume that each is drawn from a uniformly weighted three-point distribution. Specifically, $\ln \xi \in \mathcal{X} \equiv \{-\mathsf{X}, 0, \mathsf{X}\}$ and $\ln \theta \in \mathcal{Y} \equiv \{-\mathsf{Y}, 0, \mathsf{Y}\}$, where X and Y are implied immediately by σ_ξ^2 and σ_θ^2, respectively. In total, then, we have 9 pairs of $\varsigma \equiv (\xi, \theta)$, with each cohort equally represented in the population, e.g., $\lambda_{\xi,\theta} = 1/9$ for each (ξ, θ). Appendix

[24] After the completion of this project, we learned of Pozzi and Schivardi (2016), who examine data for three manufacturing sectors in Italy. Their estimates (Tables 2 and 4) yield an autocorrelation of revenue TFP of 0.763, or just slightly less than our choice of ζ. They do estimate a more persistent process for physical TFP, which would imply a higher Frisch elasticity in our model (see Section 5.1).

D argues that our conclusions are likely to hold when $\mathcal{X}$ and $\mathcal{Y}$ are higher dimensional and demonstrates this claim when ξ and θ are drawn from four-point distributions. We also considered alternative shapes for the distributions of ξ and θ, but for reasons discussed in Section 5.3, our moments appear to favor the assumption of uniformity.

Given these choices, and initial guesses for the parameters, we simulate earnings, employment, and working time outcomes within firms. We generate 220 years of data, and compute the moments based on the last 20. Structural parameters are then updated to minimize the equal-weighted quadratic loss between the model-implied and empirical moments.[25]

4.2 Main results

Table 5 summarizes our results. The top panel confirms that the model, which is just identified, perfectly reproduces the targeted moments. The bottom panel reports estimates of the structural parameters. We discuss each of the parameter estimates in turn.

Elasticity of substitution. Our estimate of $\rho = -1.962$ implies an elasticity of substitution across jobs of $(1-\rho)^{-1} \cong 0.338$. To interpret this result, consider the response of working time to a one log point increase in ξ, holding fixed the employment of each type. Given an estimate of φ (see below), and using equation (9), working time declines by approximately $(\varphi + 1 - \rho)^{-1} \cong 0.2$ log points. We further draw out the implications of $\rho < 0$ for making inferences about labor supply behavior in the next subsection.

Frisch elasticity. We find a Frisch elasticity of $1/\varphi = 0.483$. This value is two to three times larger than earlier estimates in the life-cycle literature (see Keane (2011)), but within the range of results in more recent papers (see Chetty et al (2011)) whose research designs are less confounded by idiosyncratic variation. For example, Pistaferri (2003) shows that one can robustly identify the Frisch elasticity in a life-cycle context using the response of hours to *expected* earnings growth, finding $1/\varphi = 0.7$. Pistaferri's strategy can be interpreted within our framework by noting that if ξ is relatively transitory, the expected path of earnings will more clearly reflect firm-level driving forces (Z). There are also a few estimates of the Frisch elasticity based on large-scale policy reforms, although no consensus has emerged. Sigurdsson (2021) finds a Frisch elasticity of just under 0.4 based on a tax holiday in Iceland (see also Bianchi et al (2001)), but Martinez et al (2021) do not find any labor input response to temporary regional tax rate changes in Switzerland.

Worker bargaining power. Our finding of $\eta = 0.407$ implies a flexibility of earnings

[25] When we constructed the moments within the model, we experimented with several different sample sizes. We settled on ten independent panels of 20,000 firms because increases beyond these numbers had almost undetectable impacts on our estimates.

that is within the range of estimates in related research. On the one hand, it is somewhat below the $\eta = 0.52$ in Roys (2016), who estimates a model featuring a similar bargaining problem on French micro data. On the other hand, the elasticity of daily earnings with respect to average product implied by our η is at the top end of estimates in Card et al (2014), who also use the VWH (linked to other company account data). Card et al's results indicate an elasticity between 0.06 and 0.20, whereas our model-implied analogue is 0.22.[26]

Flow outside option. The outside option μ represents 70 percent of mean earnings. To interpret this, it is helpful to see μ as the sum of two parts (although this was unnecessary for estimation): (i) income (i.e., transfers) per period of unemployment; and (ii) the average discounted surplus of *future* employment. Appendix C shows how one can back out (i) using the estimated model, finding it to be almost 50 percent of average earnings. This figure is somewhat higher than unemployment insurance replacement rates in Italy—the latter are likely closer to 40 percent (see Appendix H)—but we take this to be an encouraging result for an "out-of-sample test".

Shocks. Our estimate of σ_Z implies a standard deviation of productivity growth ($\Delta \ln Z$) equal to 0.214. This is similar to estimates based on plant-level TFP in other advanced European economies (Asker et al, 2014). At the same time, the implied standard deviation of $\ln Z$ of 0.338 is somewhat higher than in Foster et al (2008), whose estimates are for the U.S. and centered around 0.23. Finally, the dispersion of $\ln Z$ is slightly less than total idiosyncratic variation as measured by $\sqrt{\sigma_\xi^2 + \sigma_\theta^2} \cong 0.361$.

4.3 The importance of complementarities for labor supply

In this section, the estimated model is used to implement two counterfactual simulations. Labor supply incentives are altered for only a *fraction* of the workforce in one counterfactual but for *all* workers in the other. The difference in outcomes (among participants in both counterfactuals) illustrates the impact of complementarities on the dynamics of labor supply.

Each counterfactual features an unanticipated and temporary change in ξ. While such a experiment may seem a little abstract, we present the counterfactual in this way to maximize simplicity and transparency. In Appendix F, we show that a temporary change in ξ can in fact stand in for a temporary change in a labor income tax rate, e.g., a tax holiday or a one-time tax surcharge.[27] Intuitively, a larger wedge between the worker's marginal product and

[26] Card et al translate their results into estimates of bargaining power, which are substantially smaller than our value of η. However, since their static bargaining framework is unlike our set-up, the more "apples-to-apples" comparison concerns the elasticity of wages with respect to average product.

[27] To be more exact, introducing a labor income tax t is equivalent to scaling ξ and μ by $(1-t)^{-1}$. However,

marginal value of time has the same labor supply implications as a larger ξ. Nevertheless, to make our point, it suffices to directly adjust ξ.

We first perturb ξ for just the median cohort in the firm, $(\xi, \theta) = (1, 1)$. To illustrate, we elevate ξ by 25 log points, but the implied working time *elasticities* are robust to other choices. The affected employees, who make up one-ninth of the workforce, cut their time worked by 5.1 percent. If we viewed this result through the limiting case of $\rho = \alpha = 1$, we would infer a Frisch elasticity of $\Delta \ln h / \Delta \ln \xi = 0.204$ (Corollary 1). In other words, by neglecting complementarities, we would mistakenly infer a Frisch elasticity that is *60 percent less* than our estimate in Section 4.2. Notably, this result is only modestly affected if the cohort size is smaller than one-ninth of the workforce (see Appendix F).

Next, ξ is raised uniformly for *all* workers in the firm. Average time worked falls 9.7 percent, or almost twice as much as when just one cohort is affected. Even if we (wrongly) assumed $\rho = \alpha = 1$, we would recover a Frisch elasticity of 0.39, which is much closer to our estimate. Interestingly, the elasticity of working time in this counterfactual is in line with Sigurdsson's (2021) analysis of Iceland's (nation-wide) tax holiday (again, see Appendix F).

More broadly, these results can help reconcile a nontrivial Frisch elasticity with a relatively muted response to purely idiosyncratic variation in labor supply incentives. A seminal example of the latter is a series of randomized control trials known as the Negative Income Tax (NIT) experiments, which contributed to an earlier consensus on the *in*elasticity of (male) labor supply. The structure of the NITs is too intricate for us to fully capture, but it represented, in short, a temporary shift in tax rates and transfers that were "personal to the worker" (Hall, 1999). Our counterfactuals illustrate how, under complementarities, these outcomes understate the scope for intertemporal substitution. Policy changes that are broadly applied, rather than narrowly targeted, are more likely to elicit working time responses indicative of the underlying preference parameters.

5 Robustness

This section probes the robustness of our estimates. We examine the roles of (a) pre-set parameters and sample periods; (b) measurement error in working time; and (c) the assumed distributions of idiosyncratic types.

the change in μ has little quantitative effect because it operates on working time only indirectly (by modestly altering the choice of employment).

5.1 Pre-set parameters and sample period

We re-estimated the model given a higher firing cost, $\underline{c}$; a lower persistence of productivity, ζ; and higher returns to scale, α. In another exercise, we fit a more recent subsample of the data. The results are reported in Tables B.1 and B.2 in the Appendix. Taken together, they point to a Frisch elasticity $(1/\varphi)$ between 0.321 and 0.591, and an elasticity of substitution $(1/(1-\rho))$ between 0.280 and 0.454. The midpoints of these ranges are very close to our baseline estimates. The last row of each table presents the counterfactuals: with one exception, working time continues to respond almost twice as much when all workers face the higher ξ.

Larger frictions and less persistent productivity have similar effects. A firing cost of one year's earnings compresses changes in employment, which has two implications. First, since employment crowds out working time (see Section 1.3.2), smaller changes in the former are matched by larger movements in the latter. In addition, the larger firm-level shocks needed to reproduce the observed variance of $\Delta \ln N$ further exaggerate fluctuations in working time as well as earnings. Therefore, a lower Frisch elasticity, $1/\varphi$, and bargaining power, η, are needed to match the data. Like a higher $\underline{c}$, less persistent productivity induces smaller changes in employment: if hiring and firing are costly to reverse, firms attenuate responses to relatively transitory shocks. For $\zeta = 0.6$, which is at the bottom of the range cited in Syverson (2011), the model "needs" a higher σ_Z and lower values of $1/\varphi$ and η.

Many parameters move in the opposite direction when α is raised. We set $\alpha = 0.835$, which is halfway between one (constant returns) and our baseline of $\alpha = 0.67$.[28] A higher α makes labor demand more elastic, and smaller shifts in Z are needed to match the variance of $\Delta \ln N$. It follows that $1/\varphi$ must rise to generate realistic volatility in working time.

We also re-estimated the model over 1994-2001. The results suggest some changes in the contours of wage and working time setting during this period. First, we observe more variance in wage and, thus, earnings growth, perhaps reflecting a 1993 accord among policymakers, employers, and unions that is thought to have enabled more flexible wage bargains (Lodovici, 2000).[29] This change alone implies a higher bargaining power, η but has modest implications for the key preference and production parameters, φ and ρ, that shape the counterfactuals. Table B.2 also shows, though, a coincident decline in working time variability, which is

[28] Our baseline of $\alpha = 0.67$ implicitly treats capital as if it were fixed. Any degree of capital adjustment will imply a (reduced-form) elasticity of output with respect to labor input that exceeds 0.67. Therefore, we consider a higher, rather than lower, α.

[29] To this end, the agreement encouraged firm-level contracts to increase the share of wages tied to firm-level performance. In addition, the accord formally abolished a wage indexation scheme, but the latter had been substantially weakened long before (Manacorda, 2004).

not directly linked to the accord. These moments imply a lower Frisch elasticity, which is needed to reduce $\text{var}(\phi^h)$, and a higher elasticity of substitution, which ensures that $\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$ does not rise too much (when $\text{var}(\epsilon^h)$ falls). Consistent with these changes, there was a separate push by workers around this time to temper firms' overtime use while enhancing individuals' scheduling flexibility (D'Aloia et al, 2006). Finally, a higher elasticity of substitution contributes to a smaller gap between firm-level and idiosyncratic labor supply responses, but even so, the former is 50 percent larger.

5.2 Measurement error in working time

The VWH lacks data on daily hours. As a result, it likely understates the variance of working time and, thus, may overstate the relative variance of earnings growth (since earnings in the VWH do reflect all remunerated time). This subsection examines the implications of mismeasuring these moments for our baseline results.

Our analysis draws on Italy's Labor Force Survey (LFS), which has a uniquely helpful feature: it asks about weekly hours *and* days worked in the survey reference week. We can then match self-reported job stayers across adjacent years and compute the role of days in weekly hours fluctuations. Appendix B shows that the days margin accounts for virtually *all* of weekly hours growth if we include (reference) weeks with no paid days, which reflect, in part, weeks of layoff. In such cases, though, we cannot confirm that the respondent returns to the same job. If we drop weeks with no paid days, the importance of the days margin is reduced by almost half. For the sake of sensitivity analysis, we simply take days to make up roughly 75 percent of hours fluctuations, or the midpoint between these two results.

Appendix B examines the implications of missing one quarter of hours variation. Suppose the latter is distributed across $\text{var}\left(\epsilon^h\right)$ and $\text{var}\left(\phi^h\right)$ in proportion to each moment's share in the total variance of working time in the VWH. The resultant rise in the idiosyncratic variance implies weaker complementarities, which yields a higher elasticity of working time to *idiosyncratic* events. However, the higher implied firm-level variance points to a bigger Frisch elasticity, which amplifies the response to *firm-level* shocks. It turns out that these two changes virtually offset one another: working time still responds twice as much to firm-level as to idiosyncratic events, consistent with results in Section 4.

Relatedly, the lack of hours data means that we cannot separate out the components of daily earnings. The overall elasticity of days to daily earnings can be decomposed into two parts that reflect the comovement of days with (i) *daily hours* and (ii) *hourly earnings.* However, (i) has no counterpart in the model. In Appendix B, we use the LFS to infer the component due to (ii), which has a more natural connection to the comovement of working time and wages in the model. Whereas the overall elasticity is -0.169 (see Table 5), we peg the contribution of (ii) to be -0.130. This result suggests that the overall elasticity largely reflects the comovement of working time and wages.

5.3 Distributions of idiosyncratic types

We now assess the importance of two of our assumptions about idiosyncratic types, namely, ξ and θ are (1) purely transitory and (2) uniformly distributed.

First, the "bottom line" of our results is likely to be robust to the introduction of persistent types. Persistence has a direct impact on neither optimal working time, which is an intra-temporal condition, nor earnings, which takes the same form as in (12) (see Appendix C) when types are persistent. On the extensive margin, persistence would diminish the motive for labor hoarding, leading to more turnover after draws of ξ and θ. Excess turnover would have to be offset in the model by a lower σ_Z, which means that a *higher* Frisch elasticity would be needed to match the volatility of average working time.[30]

Second, Appendix D shows that our moments favor uniform types insofar as modest deviations from uniformity weaken the model's fit. The reason is that, under non-uniform distributions, changes in type often yield changes in own-type employment, $n_{\xi,\theta}$. Since diminishing returns means that $h_{\xi,\theta}$ and $w_{\xi,\theta}$ are *each* declining in $n_{\xi,\theta}$, changes in own-type employment push working time and the daily wage in the same direction. The result is a strongly positive, and highly counterfactual, correlation between $h_{\xi,\theta}$ and $w_{\xi,\theta}$.

6 Conclusion

This paper has argued that production complementarities compress working time adjustments within a firm, "squeezing out" the influence of idiosyncratic, or worker-specific, shocks. As a result, working time elasticities derived from idiosyncratic variation are attenuated relative to the Frisch elasticity. By contrast, firm-level variation acts to coordinate employees' working time decisions and thereby elicits a response more consistent with utility parameters. Indeed, our estimates imply that an identification strategy based on idiosyncratic variation would recover a Frisch elasticity that is biased down by almost 60 percent. More generally, our results suggest that more aggregate-level variation, such as broad-based policy changes,

[30]The effect of persistence on the hiring margin is somewhat harder to predict because it shapes the composition of the pool of potential hires. See Appendix D for a discussion.

is likely to better inform estimation of preference parameters.

Our framework can shed light on other economic questions where complementarities play a role. For instance, it can inform the study of housing wealth effects (on labor supply) when housing prices change unevenly across workers (see Guerrieri et al, 2013). In addition, it can aid in assessing public policies that "target" the intensive margin, such as paid family leave (see Olivetti and Petrongolo (2017) for a review). The cost to the firm of adjusting to a worker's absence depends on the elasticity of substitution across employees. Thus, our framework could inform the cost-benefit analysis of such policies.

At the same time, certain extensions to our model would be worthwhile. A challenging, but valuable, task is to incorporate *imperfect information* over types. Suppose, for instance, that firms have only a noisy signal of the preference shifter, ξ. We conjecture that working time and earnings will respond less to workers' reports of higher ξs, leading to potentially larger *extensive*-margin adjustments. This approach would help bridge the divide between our paper, in which there is substantial scope for renegotiating over ξ, and related models with competitive labor markets and, thus, no space for bargaining (see Yurdagul (2017)).

Another profitable extension addresses the *choice* of complementarities over the long run. Goldin and Katz (2016) have argued that changes in information technology and market structure have supported the adoption of new modes of production with generally weaker complementarities. By integrating the choice of production structure into the firm's problem, our framework could engage long-run trends in working time and earnings as well as the short-run dynamics on which this paper focused.

7 References

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

Altonji, Joseph, and Christina Paxson. 1988. Labor Supply Preferences, Hours Constraints, and Hours-Wage Trade-offs. *Journal of Labor Economics* 6 (2): 254-276.

Andrews, Isaiah, Matthew Gentzkow, and Jesse Shapiro. 2017. Measuring the Sensitivity of Parameter Estimates to Estimation Moments. *Quarterly Journal of Economics* 132 (4): 1553–1592.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. Dynamic Inputs and Resource (Mis)allocation. *Journal of Political Economy* 122 (5): 1013-1063.

Becker, Gary S., and Kevin M. Murphy. 1992. The Division of Labor, Coordination Costs, and Knowledge. *Quarterly Journal of Economics* 107 (4): 1137-1160.

Bianchi, Marco, Björn R. Gudmundsson, and Gylfi Zoega. 2001. Iceland's Natural Experiment in Supply-Side Economics. *American Economic Review* 91 (5): 1564-1579.

Boerma, Job, and Loukas Karabarbounis. 2021. Inferring Inequality with Home Production. *Econometrica* 89 (5): 2517–2556.

Borjas, George. 1980. The Relationship Between Wages and Weekly Hours of Work: The Role of Division Bias. *Journal of Human Resources* 15 (3): 409-423.

Brügemann, Björn, Pieter Gautier, and Guido Menzio. 2019. Intrafirm bargaining and Shapley values. *Review of Economic Studies* 86 (2): 564-592.

Brusco, Sebastiano. 1982. The Emilian Model: Productive Decentralisation and Social Integration. *Cambridge Journal of Economics* 6 (2): 167-184.

Cahuc, Pierre, François Marque, and Etienne Wasmer. 2008. A Theory of Wages and Labor Demand With Intrafirm Bargaining and matching frictions. *International Economic Review* 48 (3): 943-72.

Card, David, Francesco Devicienti, and Agata Maida. 2014. Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data. *Review of Economic Studies* 81 (1): 84-111.

Cattero, Bruno. 1989. Industrial Relations in Small and Medium-sized Enterprises in Italy. In *Industrial Relations in Small and Medium-sized Enterprises: Report to the Commission of the European Communities*, ed. Peter Auer and Helga Fehr-Duda. Luxembourg: Office for Official Publications of the European Communities.

Chang, Yongsung, and Sun-Bin Kim. 2006. From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous-Agent Macroeconomy. *International Economic Review* 47 (1): 1-27.

Chang, Yongsung, Sun-Bin Kim, Kyooho Kwon, and Richard Rogerson. 2020. Cross-Sectional and Aggregate Labor Supply. *European Economic Review* 126.

Chetty, Raj. 2012. Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply. *Econometrica* 80 (3): 969-1018.

Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri. 2011. Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence From Danish Tax Records. *The Quarterly Journal of Economics* 126 (2): 749-804.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review Papers & Proceeding*s 101 (3): 471-75.

Cooper, Russell, John Haltiwanger, and Jonathan Willis. 2015. Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications. *Research in Economics* 69 (1): 37-50.

Costinot, Arnaud. 2009. On the Origins of Comparative Advantage. *Journal of International Economics* 77 (2): 255-264.

D'Aloia, G., G. Olini, and R. Pelusi. 2006. Collective bargaining on working time in Italy, 1990-2005. In *Collective bargaining on working time: recent European experiences*, ed. Maarten Keune and Béla Galgóczi. Brussels: European Trade Union Institute.

Del Boca, Alessandra, and Paola Rota. 1998. How Much Does Hiring and Firing Cost? Survey Evidence from a Sample of Italian Firms. *Labour* 12 (3): 427–449.

Deardorff, Alan, and Frank Stafford. 1976. Compensation of Cooperating Factors. *Econometrica* 44 (4): 671-684.

Dickens, William, and Shelly Lundberg. 1993. Hours Restrictions and Labor Supply. *International Economic Review* 34 (1): 169-192.

Dixit, Avinash. 1997. Investment and employment dynamics in the short run and the long run. *Oxford Economic Papers* 49 (1): 1-20.

Erosa, Andrés, Luisa Fuster, and Gueorgui Kambourov. 2016. Towards a Micro-Founded Theory of Aggregate Labor Supply. *Review of Economic Studies* 83 (3): 1001-1039.

Erickson, Christopher, and Andrea Ichino. 1995. Wage Differentials in Italy: Market Forces, Institutions, and Inflation. In *Differences and Changes in Wage Structures*, ed. Richard B. Freeman and Lawrence F. Katz. Chicago: University of Chicago Press.

Filippi, Michelangelo, Lia Pacelli, and Claudia Villosio. 2002. Appendix A: The data. In *Labour Mobility and Wage Dynamics in Italy*, ed. Bruno Contini. Torino: Rosenberg & Sellier.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability? *American Economic Review* 98 (1): 394-425.

Goldin, Claudia. 2014. A Grand Gender Convergence: Its Last Chapter. *American Economic Review* 104 (4): 1091-1119.

Goldin, Claudia, and Lawrence Katz. 2016. A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation. *Journal of Labor Economics* 34 (3): 705-746.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. Endogenous Gentrification and Housing Price Dynamics. *Journal of Public Economics* 100: 45-60.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. Insurance Within the Firm. *Journal of Political Economy* 113 (5): 1054-1087.

Hall, Robert. 1999. Labor Market Frictions and Employment Fluctuations. In *Handbook of Macroeconomics*, ed. John Taylor and Michael Woodford. Amsterdam: North-Holland.

International Labour Organization and Organization for Economic Cooperation and Development. 2015. The Labor Share in G20 Economies. G20 Employment Working Group, Paris.

Imai, Susumu, and Michael Keane. 2004. Intertemporal Labor Supply and Human Capital Accumulation. *International Economic Review* 45 (2): 601-641.

Keane, Michael. 2011. Labor Supply and Taxes: A Survey. *Journal of Economic Literature* 49 (4): 961-1075.

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

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. Leave Out Estimation of Variance Components. *Econometrica* 88 (5): 1859-1898.

Labanca, Claudio, and Dario Pozzoli. 2022. Constraints on Hours within the Firm. *Journal of Labor Economics* 40 (2): 473-503.

Levin, Jonathan. 2003. Relational Incentive Contracts. *American Economic Review* 93 (3): 835-857.

Llosa, Gonzalo, Lee Ohanian, Andrea Raffo, and Richard Rogerson. 2014. Firing Costs and Labor Market Fluctuations: A Cross-Country Analysis. Unpublished manuscript, Society of Economic Dynamics meetings.

Lodovici, Manuela Samek. 2000. Italy: The Long Times of Consensual Re-regulation. In *Why Deregulate Labour Markets?,* ed. Gøsta Esping-Andersen and Marino Regini. Oxford: Oxford University Press.

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

Manacorda, Marco. 2004. Can the Scala Mobile Explain the Fall and Rise of Earnings Inequality in Italy? A Semiparametric Analysis, 1977–1993. *Journal of Labor Economics* 22 (3): 585-613.

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

Merz, Monika. 1995. Search in the Labor Market and the Real Business Cycle. *Journal of Monetary Economics* 36 (2): 269-300.

Mortensen, Dale, and Christopher Pissarides. 1999. New Developments in Models of Search in the Labor Market. In *Handbook of Labor Economics Vol. 3*, ed. Orley Ashenfelter and David Card. Amsterdam: North Holland.

Olivetti, Claudia, and Barbara Petrongolo. 2017. The Economic Consequences of Family Policies: Lessons from a Century of Legislation in High-Income Countries. *Journal of Economic Perspectives* 31 (1): 205-230.

Pistaferri, Luigi. 2003. Anticipated and Unanticipated Wage Changes, Wage Risk and Intertemporal Labor Supply. *Journal of Labor Economics* 21 (3): 729-754.

Pozzi, Andrea, and Fabiano Schivardi. 2003. Demand or productivity: what determines firm growth? *RAND Journal of Economics* 47 (3): 608-630.

Rogerson, Richard. 2011. Individual and Aggregate Labor Supply with Coordinated Working Times. *Journal of Money, Credit and Banking* 43 (5): S7-S37.

Rogerson, Richard, and Robert Shimer. 2011. Search-Theoretic Models of the Labor Market: A Survey. *Journal of Economic Literature* 43 (4): 959-988.

Rogerson, Richard, and Johanna Wallenius. 2009. Micro and macro elasticities in a life cycle model with taxes. *Journal of Economic Theory* 144: 2277–2292.

Roys, Nicolas. 2016. Persistence of Shocks and the Reallocation of Labor. *Review of Economic Dynamics* 22: 109-130.

Sigurdsson, Jósef. 2021. Labor Supply Responses and Adjustment Frictions: A Tax-Free Year in Iceland. Unpublished manuscript, Department of Economics, Stockholm University.

Smith, Anthony. 1993. Estimating Nonlinear Time-series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics* 8: S63-S84.

Stole, Lars, and Jeffrey Zwiebel. 1996. Intrafirm Bargaining Under Non-binding Contracts. *Review of Economic Studies* 63 (3): 375-410.

Syverson, Chad. 2011. What Determines Productivity? *Journal of Economic Literature* 49 (2): 326-365.

Taschereau-Dumouchel, Mathieu. 2020. The Union Threat. *Review of Economic Studies* 87 (6): 2859–2892.

Treu, Tiziano, G. Geroldi, and M. Maiello. 1993. Italy: Labour Relations. In *Labour Market Contracts and Institutions*, ed. J. Hartog and J. Theeuwes. Bingley: Emerald.

Uhlig, Harald. 1996. A law of large numbers for large economies. *Economic Theory* 8: 41-50.

Yurdagul, Emircan. 2017. Production complementarities and flexibility in a model of entrepreneurship. *Journal of Monetary Economics* 86: 36-51.

Table 1

Summary statistics of Veneto panel

Statistic	Mean	Std. deviation
Paid days per month	23.65	5.25
Job tenure (in months)	53.10	53.71
Daily wage (2003 euros)	121.46	426.76
Total days paid per year	243.88	97.75
Months paid per year	9.96	3.38

Note: A month is “paid” if an employee works at least one day for pay in the month. Moments based on full Veneto panel, 1982-2001. There are 22.245 million worker-year observations.

Table 2

Annual changes in days worked (Δh)

Statistic	Value
Share with $\Delta h = 0$	47.38%
Share with $\lvert\Delta h\rvert > 10$	33.15%
Average $\lvert\Delta h\rvert$ if $\lvert\Delta h\rvert \neq 0$	19.06
Average $\lvert\Delta h\rvert$ if $\lvert\Delta h\rvert \neq 0$, excluding $\lvert\Delta h\rvert > 50$	9.75

Note: This table reports moments of the distribution of annual changes in days worked, denoted by Δh. Statistics are derived from our sample of 2-year stayers, as defined in the main text (see also NOTE to Table 3). There are 11.81 million worker-year observations.

Table 3

Earnings and working time in Veneto

Moment	Interpretation	Stayers: 12/12	2-year
$\sqrt{\text{var}(\epsilon^W)}$	Std. dev. of idiosyncratic component of $\Delta \ln W$	0.162	0.210
$\sqrt{\text{var}(\epsilon^h)}$	Std. dev. of idiosyncratic component of $\Delta \ln h$	0.083	0.140
$\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$	Ratio of idiosyncratic variances	3.798	2.247
$\sqrt{\text{var}(\phi^W)}$	Std. dev. of firm component of $\Delta \ln W$	0.114	0.132
$\sqrt{\text{var}(\phi^h)}$	Std. dev. of firm component of $\Delta \ln h$	0.057	0.078
$\text{var}(\phi^W)/\text{var}(\phi^h)$	Ratio of firm-level variances	3.989	2.885
$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	Projection of $\Delta \ln h$ on $\Delta \ln w$	-0.158	-0.169

Note: W is annual earnings, h is paid days, and w is the daily wage (W/h). The 12/12 stayers are workers paid for at least 1 day in every month in 2 consecutive years. The 2-year stayers are paid for at least 1 day in each of the first 3 months in year $t-1$ and in each of the last 3 months in year t.

Table 4

Estimates of $\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$ for different samples

Sample	12/12 stayers	2-year stayers
Full sample	3.798	2.247
Excluding women	4.282	2.514
Excluding small firms (< 100 workers)	5.080	2.968
Excluding health and education	3.592	2.078

Note: This shows the ratio of the variance of the idiosyncratic component of earnings growth to the variance of the idiosyncratic component of log working time changes for different sub-samples. See Appendix G for an industry-level analysis of the ratio, $\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$.

Table 5

Model fit

Panel A: Moments

Moment	Model	Data
$\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$	2.247	2.247
$\text{var}(\phi^W)/\text{var}(\phi^h)$	2.885	2.885
$\sqrt{\text{var}(\epsilon^h)}$	0.140	0.140
$\sqrt{\text{var}(\phi^h)}$	0.078	0.078
$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	-0.169	-0.169
$\sqrt{\text{var}(\Delta \ln N)}$	0.175	0.175
$\text{E}[N]$	17.130	17.130

Panel B: Parameter values

Parameter	Symbol	Value
Elasticity of substitution across jobs	$1/(1-\rho)$	0.338 [0.0005]
Frisch elasticity of working time	$1/\varphi$	0.483 [0.0006]
Worker bargaining power	η	0.407 [0.0005]
Flow value of non-employment	μ	0.210 [0.0006]
Std. dev. of idiosyncratic preference	σ_ξ	0.294 [0.0004]
Std. dev. of idiosyncratic productivity	σ_θ	0.210 [0.0007]
Std. dev. of shock to firm productivity	σ_Z	0.203 [0.0002]

Note: Data in Panel A refers to 2-year stayers. Estimates are of our baseline model (see Section 4.2). Standard errors are in brackets. Standard errors of $1/(1-\rho)$ and $1/\varphi$ are calculated by the Delta method.

Figure 1: Labor demand policy

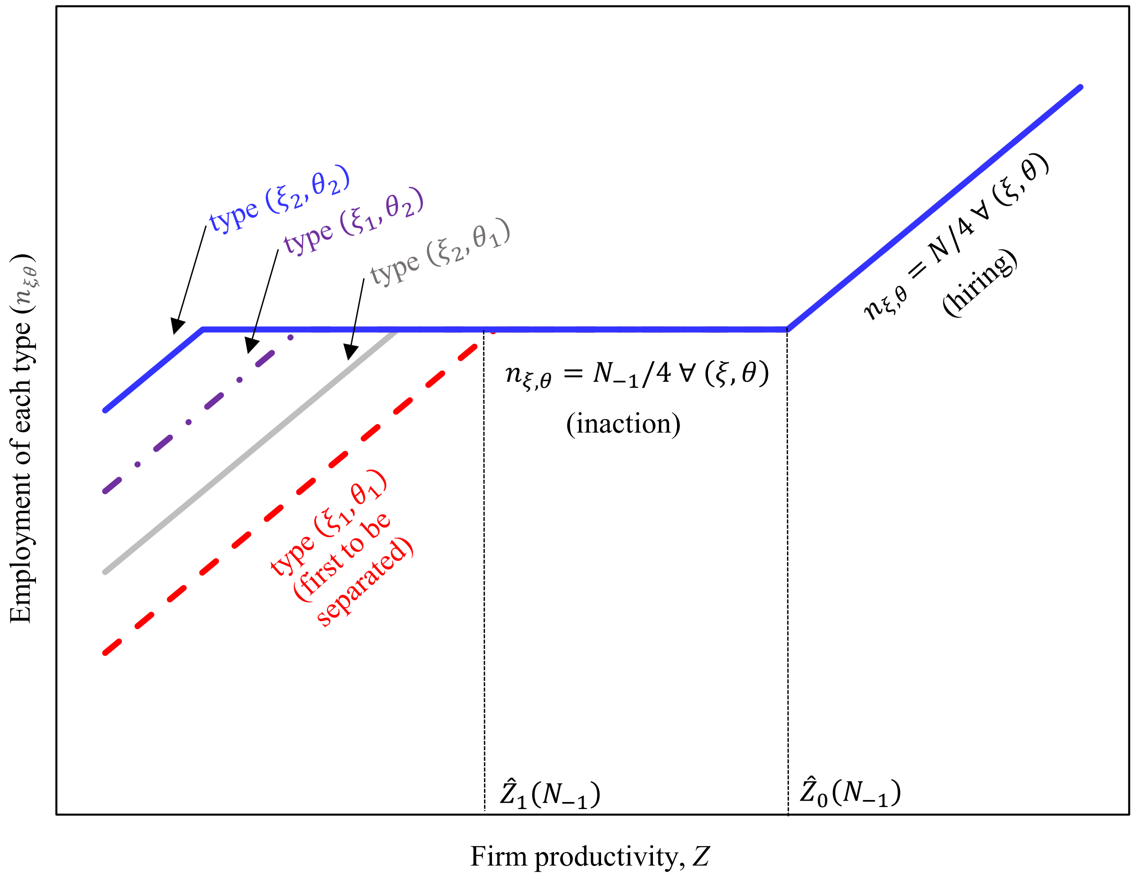

Note: This figure summarizes the optimal employment policy when preferences and productivities are drawn, respectively, from the sets $X \equiv \{\xi_1, \xi_2\}$ and $Y \equiv \{\theta_1, \theta_2\}$. Each type is assumed to be equally likely, e.g., each pair $(\xi, \theta) \in X \times Y$ is drawn with probability 1/4. For $Z > \hat{Z}_0(N_{-1})$, hires are made such that $n_{\xi,\theta} = N/4$ for each (ξ, θ), where N is the (new) level of firm employment. If $Z \in \left(\hat{Z}_1(N_{-1}), \hat{Z}_0(N_{-1})\right)$, the firm does not adjust employment of any type, hence, $n_{\xi,\theta} = N_{-1}/4$ for each (ξ, θ), where N_{-1} is the initial level of firm employment. Separations are carried out at $Z < \hat{Z}_1(N_{-1})$, such that, if the firm separates from type (ξ, θ), it also separates from *every other* type that was separted prior to (ξ, θ).

Online Appendix for "Labor supply within the firm"

Michele Battisti, Ryan Michaels, and Choonsung Park

November 11, 2022

A Main Proofs

A.1 Working time

Proof of Proposition 1. Noting that $n_{\xi,\theta}h_{\xi,\theta}$ represents total working time of type (ξ,θ), it is instructive to first write out the production function (equation (4) in the main text) as

$$\Gamma \equiv Z\left(\sum_{\xi\in\mathcal{X}}\sum_{\theta\in\mathcal{Y}}\left(\theta\int_{\mathtt{I}(\xi,\theta)} h_{\xi,\theta}(i)\,\mathrm{d}i\right)^{\rho}\right)^{\alpha/\rho},$$

where $\mathtt{I}(\xi,\theta) \subset [0,N]$ is the set of workers of type (ξ,θ) and $h_{\xi,\theta}(i)$ is working time of individual $i \in \mathtt{I}(\xi,\theta)$. Equating the marginal value of leisure of individual i, $\xi h_{\xi,\theta}(i)^{\varphi}$, to her marginal product yields,

$$\xi h^{\varphi}_{\xi,\theta}(i) = \alpha Z^{\rho/\alpha}\Gamma^{1-(\rho/\alpha)}\theta^{\rho}\left(n_{\xi,\theta}h_{\xi,\theta}\right)^{\rho-1}.$$

Clearly, each worker $i \in \mathtt{I}(\xi,\theta)$ will choose the same working time. Setting $h_{\xi,\theta}(i) = h_{\xi,\theta}$ simplifies the preceding expression to

$$\xi h^{\varphi+1-\rho}_{\xi,\theta} = \alpha Z^{\rho/\alpha}\Gamma^{1-(\rho/\alpha)}\theta^{\rho}n^{\rho-1}_{\xi,\theta}. \tag{1}$$

Now combining FOCs for types (ξ,θ) and $(x,y) \neq (\xi,\theta)$ implies

$$\frac{\xi}{x}\left(\frac{h_{\xi,\theta}}{h_{x,y}}\right)^{\varphi+1-\rho} = \left(\frac{\theta}{y}\right)^{\rho}\left(\frac{n_{\xi,\theta}}{n_{x,y}}\right)^{\rho-1}.$$

Using this to replace any $h_{x,y} \neq h_{\xi,\theta}$ in Γ, we recover equation (11) in the main text. ■

Proof of Corollary 1. Totally differentiating the solution for optimal working time with respect to $h_{\xi,\theta}$, ξ, and θ yields

$$\mathrm{d}\ln h_{\xi,\theta} = \frac{1}{\varphi+1-\rho}\left(\rho\,\mathrm{d}\ln\theta - \mathrm{d}\ln\xi\right) \;+\; \frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}\frac{\alpha-\rho}{\varphi+1-\rho}\left((\varphi+1)\,\mathrm{d}\ln\theta - \mathrm{d}\ln\xi\right), \tag{2}$$

where

$$\varpi_{\xi,\theta} \equiv \frac{\left[\theta^{\varphi+1} n_{\xi,\theta}^{\varphi}/\xi\right]^{\frac{\rho}{\varphi+1-\rho}}}{\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{Y}}\left[y^{\varphi+1} n_{x,y}^{\varphi}/x\right]^{\frac{\rho}{\varphi+1-\rho}}}.$$

The result stated in the main text refers to the case where the mass of workers of type (ξ,θ) is sufficiently small in the sense that $\varpi_{\xi,\theta} \cong 0$. Accordingly, the response of working time can be approximated by $\frac{1}{\varphi+1-\rho}\left(\rho \mathrm{d}\ln\theta - \mathrm{d}\ln\xi\right)$. The terms, $\rho/\left(\varphi+1-\rho\right)$ and $1/\left(\varphi+1-\rho\right)$, are each increasing in ρ, which implies that each attains its maximum at $\rho=\alpha$ and its minimum at $\rho=-\infty$. ■

Remark 1: Given our baseline parameters (see Table 5), $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi|_{\varpi_{\xi,\theta}=0} = -\left(\varphi+1-\rho\right)^{-1} = -0.199$. The average of $\varpi_{\xi,\theta}$ in model-generated data is 1/9 and is typically no higher than 1/6. Evaluated at $\varpi_{\xi,\theta}=1/9$, we have, more generally, that $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi = -\frac{1}{\varphi+1-\rho} - \frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}\frac{\alpha-\rho}{\varphi+1-\rho} \cong -0.199-0.024 = -0.224$. Now considering a change in θ, $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\theta|_{\varpi_{\xi,\theta}=0} = \rho\left(\varphi+1-\rho\right)^{-1} = -0.389$. More generally, using $\varpi_{\xi,\theta}=1/9$, we have $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\theta = \frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\varphi+1}{\varphi+1-\alpha}\frac{\alpha-\rho}{\varphi+1-\rho} = -0.389+0.074 = -0.315$.

A.2 Employment demand

In what follows, we will rely on two properties of the revenue function, $\hat{G}$, that follow immediately from the fact that $\rho<\alpha$. First, $\hat{G}$ is concave in $\mathbf{n}$, that is, the Hessian, $\nabla^2\hat{G}\left(\mathbf{n},Z\right)$, is negative definite. Second, $\hat{G}$ is supermodular in that $\frac{\partial}{\partial Z}\frac{\partial\hat{G}}{\partial n_{\xi,\theta}}>0$ for any type (ξ,θ) and $\frac{\partial^2}{\partial n_{\xi,\theta} n_{x,y}}\hat{G}\left(\mathbf{n},Z\right)>0$ for any $(\xi,\theta)\neq(x,y)$. We assume these properties of $\hat{G}$ pass to period profit, $\hat{\pi}$, which can be verified after the wage bargain is solved.

Conjecture 1 *The profit function,* $\hat{\pi}\left(\mathbf{n},Z\right)$*, is concave in* $\mathbf{n}$ *and supermodular in* $\left(\mathbf{n},Z\right)$*.*

The next lemma provides a key intermediate result in the characterization of the optimal policy. Since its proof relies on standard techniques, it is deferred until Appendix C.

Lemma 1 *The value function,* Π*, is concave and supermodular, under Conjecture 1.*

We are now prepared to prove Proposition 3. Since this is used to analyze the wage bargain, we present it before the proof of Proposition 2. Note that, for the sake of brevity, we will sometimes use the notation $\varsigma\equiv(\xi,\theta)$ to refer to a type if we do not need to comment further on, specifically, preferences or productivities.

Proof of Proposition 3. The optimal employment level of the first-to-be separated type $\varsigma \equiv (\xi, \theta)$ is dictated by the first-order condition,

$$\frac{\partial \pi \left(n_{\varsigma}, \boldsymbol{\lambda}_{/\varsigma} N_{-1}, Z\right)}{\partial n_{\varsigma}} + \beta \mathbb{E}\left[\Pi_N \left(N, Z'\right)|Z\right] + \underline{c} = 0, \tag{3}$$

where $\boldsymbol{\lambda}_{/\varsigma}$ is a $(M-1) \times 1$ vector of initial employment shares for types other than ς and $N = n_{\varsigma} + \Sigma_{\tau \neq \varsigma} \lambda_{\tau} N_{-1}$. By supermodularity, the left side of (3) is increasing in Z for any n_{ς}. It follows that there is a threshold $\hat{Z}_{\varsigma}(N_{-1})$ such that the firm separates from type ς when Z falls below $\hat{Z}_{\varsigma}(N_{-1})$. At this point, the firm adjusts n_{ς} according to (3), which implies a policy rule $n_{\varsigma} = \mathfrak{n}_{\varsigma}(N_{-1}, Z)$ with $\frac{\partial}{\partial Z}\mathfrak{n}_{\varsigma} > 0$. Note that for $\tau \neq \varsigma$ and $Z > \hat{Z}_{\tau}(N_{-1})$, $n_{\tau} = \lambda_{\tau} N_{-1}$.

At lower values of Z, the firm will separate from a(nother) type $\tilde{\varsigma} \neq \varsigma$, if the marginal value of that cohort falls below $-\underline{c}$ given $n_{\tilde{\varsigma}} = \lambda_{\tilde{\varsigma}} N_{-1}$,

$$\frac{\partial \pi \left(\mathfrak{n}_{\varsigma}(N_{-1}, Z),\ \boldsymbol{\lambda}_{/\varsigma} N_{-1},\ Z\right)}{\partial n_{\tilde{\varsigma}}} + \beta \mathbb{E}\left[\Pi_N \left(N, Z'\right)|Z\right] < -\underline{c}, \tag{4}$$

where $N \equiv \mathfrak{n}_{\varsigma}(N_{-1}, Z) + \Sigma_{\tau \neq \varsigma} \lambda_{\tau} N_{-1}$. Note that since the FOC (3) remains in effect as Z falls below $\hat{Z}_{\varsigma}(N_{-1})$, (4) is evaluated at the optimal size of cohort ς, $\mathfrak{n}_{\varsigma}(N_{-1}, Z)$. Therefore, at lower Z, the left side of (4) declines for two reasons: the direct effect of lower productivity, and the indirect effect of a reduction in a complementary factor, n_{ς}. It follows that, at some lower Z, (4) will take hold, and the firm will separate from type $\tilde{\varsigma}$.

When separations of $\tilde{\varsigma}$-workers begin, the firm continues to separate from type-ς workers. This follows immediately from the supermodularity of the problem: if $n_{\tilde{\varsigma}}$ is reduced, the marginal value of type-ς labor declines, and n_{ς} must be reduced to enforce the FOC (3).

Summarizing, there exist functions $\hat{Z}_{\tilde{\varsigma}}(N_{-1}) < \hat{Z}_{\varsigma}(N_{-1})$ such that the firm separates from *both* type ς and $\tilde{\varsigma}$ workers if $Z < \hat{Z}_{\tilde{\varsigma}}(N_{-1})$. Since type ς is the first type to separate, it is the rank-1 type and denoted by ς_1. Similarly, we refer to $\tilde{\varsigma}$ as the rank-2 type and set $\tilde{\varsigma} \equiv \varsigma_2$. It is straightforward to repeat this analysis for the other types, thereby establishing the ordering of types from rank 1 to rank M. ■

Remark 2: In line with the notation used in Proposition 3, we will henceforth refer to an arbitrary type as type-ς if its rank within the firm is unimportant in the context of the discussion. Otherwise, we will refer to a type as type-j, where j denotes its rank, e.g., rank-1 types are the first to be separated, rank-2 types are separated second, and so on.

A.3 Earnings

Proof of Proposition 2. As stated in the main text, and restated here for convenience, the contribution of a worker of type $\varsigma \equiv (\xi, \theta)$ to the firm, gross of the separation cost $\underline{c}$, is

$$\mathcal{J}_{\varsigma}(\mathbf{n}, Z) \equiv \frac{\partial}{\partial n_{\varsigma}} \hat{\pi}(\mathbf{n}, Z) + \beta \int \Pi_N(N, Z') \, \mathrm{d}F(Z'|Z), \tag{5}$$

where the marginal effect of type-ς labor on period profit is

$$\frac{\partial}{\partial n_{\varsigma}} \hat{\pi}(\mathbf{n}, Z) \equiv \frac{\partial \hat{G}(\mathbf{n}, Z)}{\partial n_{\varsigma}} - \left[W_{\varsigma}(\mathbf{n}, Z) + \frac{\partial W_{\varsigma}(\mathbf{n}, Z)}{\partial n_{\varsigma}} n_{\varsigma} + \sum_{\tau \neq \varsigma} \frac{\partial W_{\tau}(\mathbf{n}, Z)}{\partial n_{\varsigma}} n_{\tau} \right]. \tag{6}$$

The expected marginal value of labor in (5) can be decomposed using Leibniz's rule,[1]

$$\begin{gathered} \int \Pi_N(N, Z') \, \mathrm{d}F \\ = \sum_{j=1}^{M} \int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)} \Pi_N^{j-}(N, Z') \, \mathrm{d}F + \int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)} \Pi_N^0(N, Z') \, \mathrm{d}F + \int_{\hat{Z}_0(N)}^{\infty} \Pi_N^+(N, Z') \, \mathrm{d}F, \end{gathered} \tag{7}$$

where the term Π^{j-}, with $j = 1, ..., M$, denotes the value of the firm in states of the world in which it separates from *all* types indexed by $i \leq j$.[2] The value of the firm in states of the world in which it freezes is given by Π^0. If the firm hires, it is valued at Π^+.

We next describe the marginal value of labor in states of nature in which the firm adjusts. The value of a hiring firm is given by

$$\Pi_N^+(N, Z) = \max_{N'} \left\{ \pi(\boldsymbol{\lambda} N', Z) + \beta \int \Pi(N', Z'') \, \mathrm{d}F(Z''|Z') - \bar{c}[N' - N] \right\}. \tag{8}$$

Accordingly, applying the Envelope theorem yields

$$\Pi_N^+(N, Z') = \bar{c}. \tag{9}$$

To treat the case of separations, consider the state in which the firm separates only from type-1 labor, that is, workers of type ς_1.[3] The composition of the workforce is given by

$$\mathbf{n}^{1-}(N, Z') \equiv \left[\mathfrak{n}_1(N, Z'), \boldsymbol{\lambda}_{/1} N \right],$$

[1] We will often abbreviate $\mathrm{d}F(Z'|Z)$ (and $\mathrm{d}F(Z''|Z')$) by $\mathrm{d}F$.

[2] In (7), we define $\hat{Z}_{M+1}(N) \equiv \min\{Z\}$, the minimum of the support of Z.

[3] Throughout, we assume the firm does not simultaneously hire and fire. As noted in Section 1, firms will not do this in the face of realistic adjustment frictions.

where $\mathfrak{n}_1(N,Z')$ denotes the optimal choice of type-1 labor and $\boldsymbol{\lambda}_{/1} \equiv (\lambda_2,...,\lambda_M)$ are initial employment shares for ranks >1. The value of the firm is then

$$\Pi^{1-}(N,Z') = \hat{\pi}\left(\mathbf{n}^{1-}(N,Z'),Z'\right) - \underline{c}\left[\lambda_1 N - \mathfrak{n}_1(N,Z')\right] + \beta \int \Pi(N',Z'')\,\mathrm{d}F(Z''|Z'),$$

where $N' = \mathfrak{n}_1(N,Z') + \sum_{i=2}\lambda_i N$. By the Envelope theorem,

$$\Pi_N^{1-}(N,Z') = -\lambda_1\underline{c} + \textstyle\sum_{i=2}\lambda_i \mathcal{J}_i\left(\mathbf{n}^{1-}(N,Z'),Z'\right), \tag{10}$$

where

$$\mathcal{J}_i\left(\mathbf{n}^{1-}(N,Z'),Z'\right) \equiv \frac{\partial\hat{\pi}\left(\mathbf{n}^{1-}(N,Z'),Z'\right)}{\partial n_i} + \beta\int \Pi_{N'}(N',Z'')\,\mathrm{d}F.$$

Generalizing from (10), we have that for any state $Z \in \left[\hat{Z}_{j+1}(N),\hat{Z}_j(N)\right]$ with $j \geq 1$,

$$\Pi_N^{j-}(N,Z') = -\Lambda_j\underline{c} + \textstyle\sum_{i=j+1}^{M}\lambda_i\mathcal{J}_i\left(\mathbf{n}^{j-}(N,Z'),Z'\right), \tag{11}$$

where $\Lambda_j \equiv \sum_{i=1}^{j}\lambda_i$, $\mathbf{n}^{j-}(N,Z') \equiv \left[\{\mathfrak{n}_1(N,Z'),..,\mathfrak{n}_j(N,Z')\},\boldsymbol{\lambda}_{/j}N\right]$, and

$$\mathcal{J}_i\left(\mathbf{n}^{j-}(N,Z'),Z'\right) \equiv \frac{\partial\hat{\pi}\left(\mathbf{n}^{j-}(N,Z'),Z'\right)}{\partial n_i} + \beta\int \Pi_{N'}(N',Z'')\,\mathrm{d}F. \tag{12}$$

The marginal value of labor in the "freezing" regime can be obtained as follows. Noting that $\mathbf{n}' = \mathbf{n} = \boldsymbol{\lambda}N$ in this case and differentiating with respect to N yields

$$\Pi_N^0(N,Z') = \sum_{\varsigma\in\mathcal{X}\times\mathcal{Y}}\lambda_\varsigma\frac{\partial\hat{\pi}(\mathbf{n},Z')}{\partial n_\varsigma} + \beta\int\Pi_N(N,Z'')\,\mathrm{d}F. \tag{13}$$

Now recalling (5), evaluating the latter at $\mathbf{n} = \boldsymbol{\lambda}N$, and taking a weighted average of $\mathcal{J}_\varsigma$ across types reveals that Π_N^0 coincides with

$$\Pi_N^0(N,Z') = \sum_{\varsigma\in\mathcal{X}\times\mathcal{Y}}\lambda_\varsigma\mathcal{J}_\varsigma(\boldsymbol{\lambda}N,Z') = \sum_{j=1}^{M}\lambda_j\mathcal{J}_j(\boldsymbol{\lambda}N,Z'). \tag{14}$$

Substituting (9), (11), and (14) into (7) and inserting the result into (5) gives

$$\begin{gathered}\mathcal{J}_\varsigma(\mathbf{n},Z) \equiv \tfrac{\partial}{\partial n_\varsigma}\hat{\pi}(\mathbf{n},Z)\\ -\beta\underline{c}\textstyle\sum_{j=1}^{M}\lambda_j F\left(\hat{Z}_j(N)|Z\right) + \beta\sum_{j=1}^{M}\sum_{i=j+1}^{M}\lambda_i\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\mathcal{J}_i\left(\mathbf{n}^{j-}(N,Z'),Z'\right)\mathrm{d}F\\ +\beta\textstyle\sum_{j=1}^{M}\lambda_j\int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\mathcal{J}_j(\boldsymbol{\lambda}N,Z')\,\mathrm{d}F + \beta\bar{c}\left(1-F\left(\hat{Z}_0(N)\right)|Z\right).\end{gathered} \tag{15}$$

We turn next to the worker's surplus. Using the sharing rule, $\mathcal{W}_{\xi,\theta} = (\eta/(1-\eta))[\mathcal{J}_{\xi,\theta} + \underline{c}]$, we recast equation (6) in the main text in terms of the firm's surplus,

$$\mathcal{W}_{\xi,\theta}(\mathbf{n},Z) = \begin{array}{c} W_{\xi,\theta}(\mathbf{n},Z) - \xi\nu_{\xi,\theta}(\mathbf{n}) - \mu \\ +\beta\frac{\eta}{1-\eta}\mathbb{E}_{Z'}\sum_{i=1}^{M}\lambda_i(1-s_i(\mathbf{n}'(N,Z'),Z'))(\mathcal{J}_i(\mathbf{n}'(N,Z'),Z')+\underline{c}), \end{array} \tag{16}$$

where $\nu_{\xi,\theta}(\mathbf{n}) \equiv \frac{h_{\xi,\theta}(\mathbf{n})^{1+\varphi}}{1+\varphi}$ and s_i is the separation rate of the rank-i type. If the firm does fire type-i labor (e.g., $Z' < \hat{Z}_i(N)$), the type's marginal value, $\mathcal{J}_i$, must be driven to $-\underline{c}$, hence, the surplus is zero. Thus, $s_i(\mathcal{J}_i + \underline{c}) = 0$ must hold in all states. Alternatively, the firm may fire type j but not type $i = j+1$ if $\hat{Z}_i(N) < Z' < \hat{Z}_j(N)$. In the latter case, $\mathcal{J}_i$ is given by (5), with $\mathbf{n}' = \mathbf{n}^{j-}(N,Z')$. If the firm hires (e.g., $Z' > \hat{Z}_0(N)$), equation (8) implies that the average marginal value of labor across types is equated to the marginal cost, $\sum_{i=1}\lambda_i\mathcal{J}_i = \bar{c}$. Otherwise, if the firm freezes all types' employment at $\mathbf{n}' = \boldsymbol{\lambda}N$, then $\mathcal{J}_i$ is given by (5). Collecting these observations, we have

$$\begin{array}{c} \mathbb{E}_{Z'}\sum_{i=1}^{M}\lambda_i\max\{0,\mathcal{J}_i(\mathbf{n}',Z')+\underline{c}\} \\ = \int_{\hat{Z}_0(N)}[\bar{c}+\underline{c}]\,\mathrm{d}F + \int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\left[\sum_{i=1}^{M}\lambda_i\mathcal{J}_i(\boldsymbol{\lambda}N,Z')+\underline{c}\right]\mathrm{d}F \\ +\sum_{j=1}^{M}\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\sum_{i=j+1}^{M}\lambda_i[\mathcal{J}_i(\mathbf{n}^{j-}(N,Z'),Z')+\underline{c}]\,\mathrm{d}F. \end{array} \tag{17}$$

Substituting this into (16) and rearranging yields

$$\begin{array}{c} \mathcal{W}_{\xi,\theta}(\mathbf{n},Z) = W_{\xi,\theta}(\mathbf{n},Z) - \xi\nu_{\xi,\theta}(\mathbf{n}) - \mu \\ +\beta\frac{\eta}{1-\eta}\left\{\begin{array}{c} \underline{c}\sum_{j=1}^{M}\lambda_j\left[1-F\left(\hat{Z}_i(N)|Z\right)\right] + \sum_{j=1}^{M}\sum_{i=j+1}^{M}\lambda_i\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\mathcal{J}_i(\mathbf{n}^{j-}(N,Z'),Z')\,\mathrm{d}F \\ \sum_{j=1}^{M}\lambda_j\int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\mathcal{J}_j(\boldsymbol{\lambda}N,Z')\,\mathrm{d}F + \bar{c}\left[1-F\left(\hat{Z}_0(N)|Z\right)\right] \end{array}\right\}. \end{array} \tag{18}$$

Now inserting (15) and (18) into the sharing rule and using (6) yields

$$W_\varsigma(\mathbf{n},Z) = \eta\left\{\frac{\partial\hat{G}(\mathbf{n},Z)}{\partial n_\varsigma} - \sum_\tau\frac{\partial W_\tau(\mathbf{n},Z)}{\partial n_\varsigma}n_\tau + r\underline{c}\right\} + (1-\eta)(\xi\nu_\varsigma(\mathbf{n})+\mu). \tag{19}$$

The solution to this system of partial differential equations is (Cahuc et al, 2008)

$$W_\varsigma(\mathbf{n},Z) = \eta\left[\kappa\frac{\partial\hat{G}(\mathbf{n},Z;\boldsymbol{\varsigma})}{\partial n_\varsigma} + r\underline{c}\right] + (1-\eta)(\kappa\xi\nu_\varsigma(\mathbf{n})+\mu), \tag{20}$$

where $\kappa \equiv \frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha}$. Using the production function (equation (10) in the main text)) and optimal working time, one can calculate period profit and confirm Conjecture 1. ■

Proof of Corollary 2. Totally differentiating the earnings bargain (equation (12) in the main text)) with respect to $W_{\xi,\theta}$ and ξ yields

$$\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln \xi} = -\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\cdot\left\{\frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha}\right\}, \tag{21}$$

where

$$\varpi_{\xi,\theta} \equiv \frac{\left[\theta^{\varphi+1}n_{\xi,\theta}^{\varphi}/\xi\right]^{\frac{\rho}{\varphi+1-\rho}}}{\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{Y}}\left[y^{\varphi+1}n_{x,y}^{\varphi}/x\right]^{\frac{\rho}{\varphi+1-\rho}}}.$$

Again assuming $\varpi_{\xi,\theta} \cong 0$ and comparing (21) with (2), we have that

$$\left|\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln \xi}\right| \Big/ \left|\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln \xi}\right| = \left(1-\frac{\omega}{W_{\xi,\theta}}\right)|\rho|\,.$$

Clearly, this expression is increasing in $|\rho|$ and exceeds one iff $|\rho| > (1-\omega/W_{\xi,\theta})^{-1}$. Since $\omega/W_{\xi,\theta} < 1$ and $\rho < \alpha < 1$, it follows immediately that ρ must satisfy

$$\rho < -\left(1-\frac{\omega}{W_{\xi,\theta}}\right)^{-1} < -1$$

if earnings are to be more elastic (in absolute terms) than working time. ■

Proof of Corollary 3. Using (2) and (21), the change in the wage rate, $\mathrm{d}\ln w_{\xi,\theta} \equiv \mathrm{d}\ln W_{\xi,\theta} - \mathrm{d}\ln h_{\xi,\theta}$, following a change in ξ (all else equal) is given by

$$\frac{\mathrm{d}\ln w_{\xi,\theta}}{\mathrm{d}\ln \xi} = -\left\{1-\rho\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\right\}\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln \xi}.$$

Since $\rho < \alpha < 1$ and $\omega/W_{\xi,\theta} \in (0,1)$, the expression in brackets must be positive. Thus, the change in $w_{\xi,\theta}$ is of the opposite sign as the change in $h_{\xi,\theta}$. The response of the wage rate to a change in θ is

$$\frac{\mathrm{d}\ln w_{\xi,\theta}}{\mathrm{d}\ln \theta} = \left\{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)(1+\varphi)-1\right\}\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln \theta}.$$

The wage and working time move in the same direction if the leading term is positive. ■

Remark 3: If we use the solution for $\mathrm{d}\ln w_{\xi,\theta}/\mathrm{d}\ln\theta$ and note that $\mathrm{d}\ln W_{\xi,\theta} = \mathrm{d}\ln w_{\xi,\theta} + \mathrm{d}\ln h_{\xi,\theta}$, it immediately follows that $\mathrm{d}\ln W_{\xi,\theta}/\mathrm{d}\ln\theta = (1+\varphi)(1-\omega/W_{\xi,\theta})\,\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\theta$. Thus, the response of earnings relative to working time depends only on φ and $\omega/W_{\xi,\theta}$, as we alluded to in the main text (see Section 1.3.4).

Remark 4: For proofs of Corollaries 2 and 3 with "large" cohorts of types ($\varpi_{\xi,\theta} > 0$), see Appendix C.2.

B Main robustness analysis

This section reports the results of our robustness analysis that were reviewed in Section 5 of the main text. Table B1 reports parameter estimates assuming a higher firing cost, $\underline{c}$; a lower persistence of productivity, ζ; and higher returns to scale, α. Table B2 reports results using moments calculated over the 1994-2001 subsample and, separately, working time moments adjusted for undercounting.

Since the parameter estimates were discussed in the text, we focus our remarks here on fleshing out our treatment of undercounting that underlies the results in Table B2. To recall, our Veneto Work History (VWH) data lack observations on daily hours. We detail below how we use Italy's Labor Force Survey (LFS) to develop a simple adjustment for measuring working time that can be used to quantify the implications of missing daily hours for our parameter estimates. In addition, the VWH does not capture changes in hours due to transitions between five- and six-day weeks (and vice versa), although we will argue that the latter arguably should be excluded from the sample.

Measuring and adjusting for undercounting. The LFS is administered quarterly and asks about weekly hours and days worked in the survey reference week. We use data between 1993, when it becomes possible to link micro data across years, and 2001, the final year of our VWH sample. Each household is scheduled to participate for two consecutive quarters, exits the survey for two quarters, and returns to the sample for two more (consecutive) quarters. A household member in his third (fourth) quarter of participation is referred to as a "stayer" if his employer is the same as it was in the first (second) quarter.

Among stayers, we measure the year-over-year change in weekly hours and days worked per week. We do this in two ways so as to bound the contribution of the days margin. To frame our approach, it is helpful to start with a basic identity: weekly hours is the product of days worked in that week and daily hours. This identity would seem to suggest a simple way to assess the importance of the days margin, namely, compare the variance of the log change in days worked to the variance of the log change in weekly hours. However, this approach can be implemented only if we drop *weeks with zero days* worked.

If we want to incorporate zero-days weeks, we must consider an alternative route. To this end, define the average value of a variable χ across periods $t-1$ and t as $\bar{\chi} \equiv \left(\chi_t + \chi_{t-1}\right)/2$

and define the growth rate of χ as $g_t^\chi \equiv \left(\chi_t - \chi_{t-1}\right)/\bar{\chi}$. This treatment of growth naturally accommodates a realization of $\chi_t = 0$ or $\chi_{t-1} = 0$ (although not both χ_t and $\chi_{t-1} = 0$). Also, let H denote total weekly hours; $\mathfrak{d}$ days worked per week; and $\mathfrak{h}$ daily hours. The change in (the level of) weekly hours can be exactly decomposed as

$$\Delta H_t \equiv H_t - H_{t-1} = \bar{\mathfrak{h}} \times \Delta\mathfrak{d}_t + \bar{\mathfrak{d}} \times \Delta\mathfrak{h}_t.$$

It then follows that

$$\frac{\Delta H_t}{\bar{H}} \equiv g_t^H = \gamma_t^{\mathfrak{d}} + \gamma_t^{\mathfrak{h}} \equiv \frac{\bar{\mathfrak{d}}\bar{\mathfrak{h}}}{\bar{H}} \times \left[g_t^{\mathfrak{d}} + g_t^{\mathfrak{h}}\right] \tag{22}$$

The term $\gamma_t^{\mathfrak{d}} \equiv g_t^{\mathfrak{d}} \times \left(\bar{\mathfrak{d}}\bar{\mathfrak{h}}/\bar{H}\right)$ captures the role of days worked in weekly hours movements. Thus, to assess the role of the days margin, we can now compare the variance of $\gamma_t^{\mathfrak{d}}$ to the variance of g_t^H.

These two estimates will bracket the contribution of days fluctuations to total hours variability. Although we place more weight on the second method, which incorporates zero-days weeks, we recognize that it is important to consider alternatives. For instance, one reason to question the result with zero-days weeks is that, if a worker does not report any days worked in the reference week, we are unable to verify that this individual in fact returned to his firm. This point is of some concern since we seek to measure changes in hours and days only among job *stayers*. For the purpose of sensitivity analysis, we therefore propose to simply use the average of the results with and without zero-days weeks.

To carry out the estimation, we first trim the sample in a few ways. First, we exclude observations in which weekly hours changes reflect holidays, vacations, and off-site job training. This variation in working time generally amplifies the role of days but arguably reflects driving forces that are beyond the scope of the model. In addition, we favor leaving out observations that involve transitions between five- and six-day weeks (and vice versa). Throughout our sample, Italy was in the midst of a longer-run transition toward five-, and away from six-, day weeks. Between 1993 and 2001, the share with five-day weeks rose steadily from 60.5 percent to 64 percent. Concurrently, the share with six-day weeks fell by almost the exact same amount. This *trend* is orthogonal to the economic forces, such as intertemporal substitution, that we model and that are relevant to parameter estimation in our context. Accordingly, we will initially exclude these observations but will report on their implications a little later.

In the end, we have 9,413 observations on annual changes in days and weeks for Veneto workers in the LFS *before* excluding (i) zero-days weeks and (ii) transitions across five- and six-day weeks. Dropping (i) eliminates 7 percent of the sample, reducing it to 8,752 observations. Excluding (ii) then brings the sample down to 7,979 observations.

We can now present our estimates of the role of the days margin in weekly hours movements. We measure the contribution of days growth based on its variance relative to the variance of weekly hours growth.[4] If we exclude zero-days weeks, the variance of days growth represents 54.2 percent of the variance of weekly hours growth. Notably, this estimate is very robust to how we handle potentially spurious reports of working time in the LFS. For this calculation, we drop observations that involve (absolute) changes in daily hours of more than four, a choice that is informed in part by daily overtime regulations in Italy.[5] However, if we simply eliminated the top and bottom 1 percent of weekly hours growth observations, we would obtain almost the exact same result.

Next, if we incorporate zero-days weeks, the variance of days growth slightly exceeds the variance of weekly hours growth (since days and daily hours changes are negatively correlated).[6] For this exercise, we shall simply treat $\text{var}(\gamma^{\mathfrak{d}})$ as interchangeable with $\text{var}\left(g^H\right)$, e.g., $\text{var}(\gamma^{\mathfrak{d}})\,/\text{var}\left(g^H\right) = 1$.

Although we, again, regard the estimate with zero-days weeks as more informative, the sensitivity analysis is based on an (unweighted) average of the estimates with and without zero-days weeks. Thus, we take days to represent $0.5 \times (0.542 + 1) = 77.1$ percent of hours variability.

Importantly, our choice to exclude transitions between five- and six-day weeks (and vice versa) has a reasonably modest effect on our results. Since the VWH does not capture these transitions, one could instead argue that all of the variation in total hours associated with these observations in the LFS should be treated *as if* it reflected daily hours. This operation would serve to shed light on the extent of undercounting by the VWH's measure of days. If we follow this procedure, the variance of log changes in days now represents 45.3 percent (rather than 54.2 percent) of the variance of weekly hours growth in the case where we exclude zero-days weeks. If we include zero-days weeks, the impact of these transitions is negligible. Thus, the midpoint of the range falls only to $0.5 \times (0.453 + 1) = 72.7$ percent.

[4] To exactly decompose the variance of hours growth, one could apportion the covariance of days and daily hours growth across the two components, e.g., the contribution of days would be $\text{var}(\gamma^{\mathfrak{d}}) + \text{cov}\left(\gamma^{\mathfrak{d}}, \gamma^{\mathfrak{h}}\right)$ in the context of (22). However, we simply ask if the hours-equivalent variation of days growth—that is, $\text{var}(\gamma^{\mathfrak{d}})$—is comparable to the variance of hours growth, e.g., if days growth is a good *proxy*. If so, a model that reproduces the variability of days will be consistent with the variability of hours.

[5] These regulations effectively rule out sustained daily overtime in excess of two (let alone four) hours among full-time workers. Meanwhile, transitions between full- and part-time status would have been rare and fairly circumscribed at this time, and so changes of more than four hours—for instance, a transition from 8 to 2 daily hours—strike us as unlikely.

[6] As a result, the sum of $\text{var}(\gamma^{\mathfrak{d}})$ and $\text{cov}\left(\gamma^{\mathfrak{d}}, \gamma^{\mathfrak{h}}\right)$ represents 98 percent of the variance of weekly hours growth.

To trace out the implications of these findings for our baseline estimates, we must consider how the missing variation in total working time is distributed across idiosyncratic and firm-wide sources. We assume it is distributed across $\text{var}\left(\epsilon^h\right)$ and $\text{var}\left(\phi^h\right)$ in proportion to each moment's share in the total variance of working time in the VWH. Using estimates from Table 5 in the main text, and noting that these components are (by construction) orthogonal, we find that the idiosyncratic part accounts for about 75 percent of the total. Accordingly, we scale the total variance in the VWH by (1/0.77) based on our estimates from the LFS and then distribute three quarters of the increase to $\text{var}\left(\epsilon^h\right)$. Assuming VWH earnings are measured accurately, the ratio of the idiosyncratic variances, $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$, falls to 1.742 from 2.247. The analogue for the ratio of firm-wide variances is 2.180, down from 2.885. We then calibrate the model using these two ratios of variances as well as the implied values of $\text{var}\left(\epsilon^h\right)$ and $\text{var}\left(\phi^h\right)$. Other VWH moments, including the corresponding variances of earnings, are unaltered.

Two results of this calibration stand out (see Table B2). The higher variability of idiosyncratic working time implies a higher elasticity of substitution, which increases to $(1-\rho)^{-1} = 0.423$ (up from 0.338 in the baseline case). At the same time, a higher variance of *average* working time implies a higher Frisch elasticity, which is found to be $1/\varphi = 0.590$ (up from 0.483). It follows (see Corollary 1) that the elasticity of an individual's working time with respect to ξ is now (approximately) $(\varphi + 1 - \rho)^{-1} = 0.246$, whereas the elasticity of working time to *firm-level* events equals $(\varphi + 1 - \alpha)^{-1} = 0.494$. Thus, working time still responds twice as much to firm-level as to idiosyncratic events, which is consistent with results in Section 4.

Interpreting daily earnings data. In Section 5.2, we also examined an issue closely related to the omission of hours in the VWH, namely, we can compute daily, but not hourly, earnings. The elasticity of days to daily earnings reflects (i) the comovement of days and *daily hours* and (ii) the comovement of days and hourly *wages*. More precisely, the elasticity of days to daily earnings, which we estimate to be -0.169, can be decomposed according to

$$\begin{aligned} -0.169 &= \mathcal{L} + \mathcal{R} \\ &\equiv \frac{\text{covar}\left(\Delta \ln \text{ daily hours}, \Delta \ln \text{ days}\right)}{\text{var}\left(\Delta \ln \text{ daily earnings}\right)} + \frac{\text{covar}\left(\Delta \ln \text{ hourly wage}, \Delta \ln \text{ days}\right)}{\text{var}\left(\Delta \ln \text{ daily earnings}\right)}. \end{aligned}$$

The term $\mathcal{L}$ on the left reflects the comovement of days and daily hours, whereas the term $\mathcal{R}$ on the right captures the comovement of days and wages. The covariance of days and wages in $\mathcal{R}$ is the correct counterpart to the relationship between working time and wages in the structural model.

We cannot observe $\mathcal{R}$ but we can back out an estimate of it using $\mathcal{L}$. The numerator of $\mathcal{L}$, the covariance of log changes in daily hours and days, is taken from the LFS, whereas the denominator, the variance of daily earnings growth, must be computed from the VWH (since the LFS does not include earnings data during our sample). If we carry out this division, we find that $\mathcal{R} = -0.150$, which is very close to the OLS coefficient, -0.169.[7]

However, such a direct comparison of moments across LFS and VWH is perhaps problematic because of differences in concepts and coverage. A simple way to adjust for these differences is to first express the numerator in $\mathcal{L}$ relative to the variance of log changes in days worked, or $\text{var}(\Delta \ln \text{ days})$, *in the LFS*. We then multiply this ratio by the value of $\text{var}(\Delta \ln \text{ days}) / \text{var}(\Delta \ln \text{ daily earnings})$ *in the VWH*. Note that if $\text{var}(\Delta \ln \text{ days})$ were equal in the two datasets, this procedure would give the same answer as if we divided the covariance term from the LFS by the variance (of daily earnings growth) in the VWH (as we did above). These calculations yield $\mathcal{R} = -0.130$, which is still a relatively modest departure from our estimate of -0.169.

C Additional theoretical results

This section presents five results. The first is the proof of Lemma 1, which is an input into the proofs of Propositions 2 and 3. The second generalizes Corollaries 2 and 3 to allow for "large" type cohorts. The third fleshes out further results on optimal employment demand not covered in Proposition 3. The fourth section derives the worker's surplus (equation (5) in the main text), which is an input into Propositions 1 and 2. We also use this derivation to illustrate how to interpret the flow value of nonwork time, μ, as the sum of two parts: (i) income (transfers) per period of unemployment; and (ii) the expected discounted surplus of future employment. Finally, the fifth section shows that the form of the earnings bargain (Proposition 2) carries over to the case with persistent idiosyncratic preferences and productivities.

C.1 Properties of the value function

Lemma 1 establishes that certain properties of the period profit function pass to the value function. As noted in Appendix A, one can confirm after the derivation of the wage bargain that the profit function does indeed possess these properties. For the purpose of the proof, we shall assume (A1) the domain of Z is a compact subset of the real line and (A2) the

[7] Since this decomposition is naturally cast in terms of logs, we do exclude the zero-days weeks.

conditional distribution function of $Z'|Z$ satisfies the Feller property.

Lemma 1 *Assume that the period profit function* $\pi(\mathbf{n},Z)$ *is concave and supermodular. Then, the value function,* $\Pi(N_{-1},Z)$ *, is concave and supermodular.*

Proof. This proof closely follows Dixit (1997). To establish concavity, combine equations (8)-(10) in the main text,

$$\begin{aligned}\Pi(N_{-1},Z) &= \max_{\mathcal{N},\mathbf{n}} \tilde{\Pi}(\mathcal{N},\mathbf{n},\ N_{-1},Z) \\ &\equiv \max_{\mathcal{N},\mathbf{n}} \left\{-\bar{c}\left[\mathcal{N}-N_{-1}\right]-\underline{c}\sum_{\varsigma}\left[\lambda_{\varsigma}\mathcal{N}-n_{\varsigma}\right]+\pi(\mathbf{n},Z)+\beta\mathbb{E}\left[\Pi(N,Z')\,|Z\right]\right\},\end{aligned} \tag{23}$$

where $\mathcal{N}$ denotes the number of workers after hires (if any) have been made but before separations; $\mathbf{n}\equiv\{n_{\varsigma}\}$ is the vector of employment over types after separations have been decided; and $N\equiv\sum_{\varsigma}n_{\varsigma}$ is the total workforce used in production (and then "carried into" next period). Recall, $\mathcal{N}$ and $\mathbf{n}\in\mathbb{R}^{M}$ must satisfy $\mathcal{N}\geq N_{-1}$ and $\boldsymbol{\lambda}\mathcal{N}\geq\mathbf{n}$. Further, the boundedness of Z (A1) implies a natural upper bound on $\mathcal{N}$ and a natural lower bound on $\mathbf{n}$. Thus, the support of $(\mathbf{n},\mathcal{N})$ is a compact and convex subspace of $\mathbb{R}^{M+1}$.

The induction hypothesis is that $\Pi(N,Z')$ is bounded, continuous, concave, and supermodular. These properties pass to $\mathbb{E}\left[\Pi(N,Z')\,|Z\right]$ by A2 (Stokey and Lucas, 1989) and by the fact that supermodularity is preserved under integration (Topkis, 1998). Further, under our conjecture, period profit is continuous, concave and supermodular (and, by A1, bounded). Trivially, the adjustment cost functions are weakly concave and supermodular. These arguments verify that $\tilde{\Pi}(\mathcal{N},\mathbf{n},\ N_{-1},Z)$ is concave and supermodular. Thus its maximum is supermodular with respect to the state (N_{-1},Z) (Topkis, 1998).

To verify that Π is concave, consider any two states of nature, $q_i\equiv(N_{-1,i},Z_i)$ for $i=1,2$, and let the policies $p_i\equiv(\mathcal{N}_i,\mathbf{n}_i)$ be optimal there. Note that if the state is $\bar{q}\equiv\kappa q_1+(1-\kappa)\,q_2$ for any $\kappa\in(0,1)$, then the policy $\bar{p}\equiv\kappa p_1+(1-\kappa)\,p_2$ is feasible there. Therefore, by the concavity of $\tilde{\Pi}$, it follows that

$$\Pi(\bar{q})\geq\tilde{\Pi}(\bar{p},\ \bar{q})\geq\kappa\tilde{\Pi}(p_1,q_1)+(1-\kappa)\,\tilde{\Pi}(p_2,q_2)=\kappa\Pi(q_1)+(1-\kappa)\,\Pi(q_2),$$

which confirms the concavity of Π.

We have thus shown that the operator defined by (23) maps the space, $\mathcal{C}$, of bounded, continuous, concave, and supermodular functions into itself. It is straightforward to show that this operator is a contraction for any $\beta\in(0,1)$, so the fixed point of (23) is unique and lies in the space, $\mathcal{C}$. ■

C.2 Extensions to Corollaries 2 and 3

Corollary 2. In Corollary 2, we compare the elasticities of earnings and working time with respect to the preference shifter, ξ. For convenience, we restate these comparative statics here:

$$\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\xi} = -\frac{1}{\varphi+1-\rho} - \frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}\frac{\alpha-\rho}{\varphi+1-\rho} \tag{24}$$

and

$$\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi} = -\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\cdot\left\{\frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha}\right\}, \tag{25}$$

where

$$\varpi_{\xi,\theta} \equiv \frac{\left[\theta^{\varphi+1}n_{\xi,\theta}^{\varphi}/\xi\right]^{\frac{\rho}{\varphi+1-\rho}}}{\displaystyle\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{Y}}\left[y^{\varphi+1}n_{x,y}^{\varphi}/x\right]^{\frac{\rho}{\varphi+1-\rho}}}. \tag{26}$$

In the main text, we consider the special, but instructive, case where $\varpi_{\xi,\theta}\cong 0$. We found that the magnitude of the earnings response exceeds the size of the working time response iff (if and only if) $\rho<-\rho_0$, where $\rho_0\equiv\left(1-\omega/W_{\xi,\theta}\right)^{-1}$.

More generally, suppose $\varpi_{\xi,\theta}>0$. Under certain conditions, we can then establish that earnings are more elastic (in absolute terms) than working time iff $\rho<-\rho_1<-\rho_0$, where ρ_1 will be specified below. In other words, when $\varpi_{\xi,\theta}>0$, complementarities must be stronger. Intuitively, if a "large" cohort cuts working time following a rise in its own ξ, other types will follow suit to an extent, thus corroborating cohort ξ's choice and further amplifying the decline in its working time. A more pronounced decrease in ξ's working time attenuates the rise in its marginal disutility of working, restricting the increase in earnings. To ensure that earnings are relatively elastic, one must then impose greater complementarities, which will restrict type ξ's incentive to cut working time in the first place.

To proceed, first suppose that $\rho<0$ is such that $\mathrm{d}\ln W_{\xi,\theta}/\mathrm{d}\ln\xi>0$. Since $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi<0$, we want to establish the conditions under which the former exceeds the absolute value of the latter. Using (24)-(25), we have that

$$\begin{aligned}\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi}-\left|\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\xi}\right| &= \frac{1}{\varphi+1-\rho}\left\{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)|\rho|-1\right\}\\ &\quad -\frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}\frac{\alpha-\rho}{\varphi+1-\rho}\left\{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)(\varphi+1)+1\right\},\end{aligned}$$

which is positive iff

$$|\rho|>\rho_1\equiv\frac{1+\chi\alpha\varpi_{\xi,\theta}}{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)-\chi\varpi_{\xi,\theta}},$$

where

$$\chi \equiv \left(1-\frac{\omega}{W_{\xi,\theta}}\right)\frac{\varphi+1}{\varphi+1-\alpha}+\frac{1}{\varphi+1-\alpha}.$$

It is clear that $\rho_1 > 0$ as long as its denominator is positive, and, given the definition of χ, one can confirm that the latter will hold if

$$\varpi_{\xi,\theta} < \frac{\varphi+1-\alpha}{\varphi+1+(1-\omega/W_{\xi,\theta})^{-1}}. \tag{27}$$

This result says that there is a ρ that satisfies $\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi} > \left|\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\xi}\right|$ as long as the cohort is not too large. This inequality holds in the estimated model. To see why, recall that our results imply $\omega/W_{\xi,\theta} < 1/2 \iff (1-\omega/W_{\xi,\theta})^{-1} < 2$ for any (ξ,θ) (see Section 1.3.4). Therefore, the right side of (27) is bounded below by 1/2, whereas $\varpi_{\xi,\theta}$—roughly, the type's share in output—is much below this. (Indeed, with nine types, $\varpi_{\xi,\theta} \cong 1/9$ in model-generated data.) We can then conclude that

$$\rho_1 \equiv \frac{1+\chi\alpha\varpi_{\xi,\theta}}{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)-\chi\varpi_{\xi,\theta}} > \left(1-\frac{\omega}{W_{\xi,\theta}}\right)^{-1} \equiv \rho_0,$$

which establishes that $\rho < -\rho_1 < -\rho_0 < 0$. As previewed above, complementarities must be stronger when $\varpi_{\xi,\theta} > 0$.

Finally, at this point, we can verify that $\rho > \rho_1$ ensures that the initial restriction on $\mathrm{d}\ln W_{\xi,\theta}/\mathrm{d}\ln\xi > 0$ holds. Using (25), and recalling that $\omega/W_{\xi,\theta} \in (0,1)$, earnings are increasing in ξ iff

$$\frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha} < 0,$$

which implies that $\rho < 0$ and

$$|\rho| > \alpha\frac{\varpi_{\xi,\theta}\frac{\varphi+1}{\varphi+1-\alpha}}{1-\varpi_{\xi,\theta}\frac{\varphi+1}{\varphi+1-\alpha}} \equiv \underline{\rho}. \tag{28}$$

One can easily confirm that (27) is a sufficient condition for $\rho_1 > \underline{\rho}$, which means (28) holds.

Alternatively, suppose $\mathrm{d}\ln W_{\xi,\theta}/\mathrm{d}\ln\xi < 0$. Then earnings responds relatively more in absolute terms if $\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi} < \frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\xi}$, which implies

$$\rho > \varrho \equiv \frac{1-\alpha\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\frac{\varphi+1}{\varphi+1-\alpha}\varpi_{\xi,\theta}+\alpha\frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}}{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\left(1-\frac{\varphi+1}{\varphi+1-\alpha}\varpi_{\xi,\theta}\right)+\frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}}.$$

The condition (27) will ensure that the right side of this is positive, e.g., it must be that $\rho > 0$. However, ρ also must be less than α (recall, $\rho = \alpha$ corresponds to the limit of perfect substitutes). In fact, one can verify that $\varrho > \alpha$:

$$\frac{1-\alpha\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\frac{\varphi+1}{\varphi+1-\alpha}\varpi_{\xi,\theta}+\alpha\frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}}{\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\left(1-\frac{\varphi+1}{\varphi+1-\alpha}\varpi_{\xi,\theta}\right)+\frac{\varpi_{\xi,\theta}}{\varphi+1-\alpha}} \equiv \varrho > \alpha \Longleftrightarrow 1 > \alpha\left(1-\frac{\omega}{W_{\xi,\theta}}\right).$$

Thus, we obtain a contradiction. We conclude that earnings are more elastic than working time iff ρ is sufficiently small in the sense that $\rho < -\rho_1$.

Corollary 3. Next, we reconsider Corollary 3 in the case where $\varpi_{\xi,\theta} > 0$. The corollary characterizes conditions under which working time and the wage rate, $w_{\xi,\theta} \equiv W_{\xi,\theta}/h_{\xi,\theta}$, move in opposite directions following idiosyncratic shocks. The responses to a change in ξ can be deduced immediately from the preceding analysis. If earnings are increasing in ξ, it follows that the wage rate is increasing in ξ since working time is unambigiously declining in ξ. Alternatively, the wage rate will be declining in ξ—and, thus, of the same sign as working time—if and only if earnings are decreasing in ξ *and* $|\partial \ln W_{\xi,\theta}/\partial \ln \xi| < |\partial \ln h_{\xi,\theta}/\partial \ln \xi|$. However, this scenario requires $\rho > \varrho$, which is infeasible.

Finally, we characterize the response of the wage rate following a shift in θ. The comparative statics in this case are given by,

$$\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\theta} = \frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha}$$

and

$$\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\theta} = \left(1-\frac{\omega}{W_{\xi,\theta}}\right)\cdot\left\{\frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha}\right\}(\varphi+1).$$

We can now evaluate

$$\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\theta}\times\frac{\mathrm{d}\ln w_{\xi,\theta}}{\mathrm{d}\ln\theta} = \left(\left(1-\frac{\omega}{W_{\xi,\theta}}\right)(\varphi+1)-1\right)\cdot\left(\frac{\rho}{\varphi+1-\rho} + \varpi_{\xi,\theta}\frac{\alpha-\rho}{\varphi+1-\rho}\frac{\varphi+1}{\varphi+1-\alpha}\right)^2.$$

If this expression is positive, working time and the wage rate move in the same direction. The relevant condition on φ is $\varphi > (1-\omega/W_{\xi,\theta})^{-1} - 1$, which is the same result presented in the main text for the simpler case where $\varpi_{\xi,\theta} \cong 0$.

C.3 Optimal separation policy

In the main text, we characterize the manner in which firms separate from workers when firm productivity is sufficiently low. Specifically, we established that there is a ranking of the types such that a firm starts shrinking by separating only from one type, leaving employment of the other types in place (for the moment). Here, we shed some light on the *identity* of the type to be separated first.

To begin, it is helpful to write out period profit using the solutions for working time (equation (9)) and earnings (equation (12)) in the main text. We obtain,

$$\hat{\pi}(\mathbf{n}, Z) = (1 - \varkappa)\,\alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\frac{\alpha}{\alpha-\rho}\frac{\varphi+1-\rho}{\varphi+1-\alpha}} - \omega \sum\nolimits_{\xi,\theta} n_{\xi,\theta} \tag{29}$$

where $\varkappa \in (0,1)$ is

$$\varkappa \equiv \alpha \frac{\eta\varphi + (1-\eta)\frac{\varphi+1-\alpha}{\varphi+1}}{(1-\eta(1-\alpha))(\varphi+1) - \alpha}$$

and $\omega \equiv \eta r \underline{c} + (1-\eta)\,\mu$. Since each type is redrawn each period, the future marginal value of a worker of current type (ξ, θ) is identical to that of a worker of any other current type. Thus, when we compare the marginal value of labor among workers of different types, it is sufficient to consider their marginal profitabilities.

To this end, we seek to characterize

$$\frac{\partial \hat{\pi}}{\partial n_{\xi,\theta}} = \frac{\varphi}{\varphi+1-\alpha}(1-\varkappa)\,(\alpha Z \Omega(\mathbf{n}))^{\frac{\varphi+1}{\varphi+1-\alpha}} \left(\frac{\theta^{\varphi+1}}{\xi}\right)^{\frac{\rho}{\varphi+1-\rho}} n_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} - \omega,$$

where we have used the definition of Ω (see (9) in the main text). To clarify matters even further, we can evaluate this at types' initial employment levels, $\mathbf{n} \equiv \boldsymbol{\lambda} N_{-1}$. Accordingly, we have

$$\left.\frac{\partial \hat{\pi}}{\partial n_{\xi,\theta}}\right|_{\mathbf{n}\equiv\boldsymbol{\lambda}N_{-1}} = \frac{\varphi}{\varphi+1-\alpha}(1-\varkappa)\left(\alpha Z \tilde{\Omega} N_{-1}^{\alpha-1}\right)^{\frac{\varphi+1}{\varphi+1-\alpha}} \times \left(\frac{\theta^{\varphi+1}}{\xi}\right)^{\frac{\rho}{\varphi+1-\rho}} \lambda_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} - \omega, \tag{30}$$

where

$$\tilde{\Omega} \equiv \left(\sum_{x,y}\left(\frac{y^{\varphi+1}\lambda_{x,y}^{\varphi}}{x}\right)^{\frac{\rho}{\varphi+1-\rho}}\right)^{\frac{\alpha-\rho}{\rho}}.$$

To determine the first to be separated type, we now inspect how $\partial\hat{\pi}/\partial n_{\xi,\theta}$ varies with $\lambda_{\xi,\theta}$, θ, and ξ.

Straightforward differentiation will confirm a few properties of (30). First, $\partial\hat{\pi}/\partial n_{\xi,\theta}$ is declining in $\lambda_{\xi,\theta}$. Intuitively, a higher $\lambda_{\xi,\theta}$ indicates that the type is relatively abundant. By the concavity of the production function, it is therefore less valued at the margin.

Next, the comparative statics with respect to ξ and θ reflect the sum of two, partially offsetting effects: the direct effect of an increase in θ or ξ on that type's own working time and an indirect effect that reflects that type's influence on the working times of others. The intuition behind the direct effect is clear. A higher-θ type produces more input into final production. If types' inputs are complements, this type will be less valued at the margin. The same point applies to workers of low ξ, since they will supply more time worked (and thus more inputs). The indirect effect is a consequence of supermodularity: when a low ξ supplies more inputs, it calls forth additional effort from other types, which in turn elevates the marginal value of the low-ξ type. The direct effect will dominate when ρ is sufficiently small ("more negative") and/or when the type's share in production is relatively small. Formally, one can verify that $\partial\hat{\pi}/\partial n_{\xi,\theta}$ is declining in θ and is increasing in ξ when (28) holds.

To sum up, the first type (ξ, θ) to be separated is relatively abundant (a high $\lambda_{\xi,\theta}$); relatively productive (a high θ); and values leisure relatively little (a low ξ). Since $\lambda_{\xi,\theta} = \lambda$ for all (ξ, θ) in the estimated model, we can in fact pinpoint the first type to be separated: it is the type that draws the highest θ and lowest ξ.

The same considerations will shape the choice of the next type from which to separate. To see this, consider a perturbation to $n_{x,y}$ of any type (x, y) from which the firm *has not yet separated*. A change to $n_{x,y}$ will imply a change in the employment of the already-separated types. However, by the envelope theorem, the latter's effect on the marginal value of type (x, y) is zero. Therefore, if we compare the marginal values of two not-yet-separated types, it is again sufficient to compare their marginal profitabilities. In that case, our previous discussion indicates that the second-to-be-separated type will have the second-highest value of productivity and the second-lowest value of leisure preferences. In this manner, we can then identify the productivity and preferences of the third-to-be-separated type and so on.

C.4 The worker's surplus

In our model, we imagine that there is a number of "large" families, and each worker is a member of one of them. The families' employed members are distributed (randomly) across a continuum of firms. Thus, the family can pool members' incomes to insure their consumption against worker- and firm-specific risk. In this section, we illustrate how to derive the worker's surplus (equation (5) in the main text) from a family's optimization problem. This exercise applies, and extends, the presentations in Merz (1995) and Trigari (2006).

An employed member of the family has idiosyncratic type $\varsigma \equiv (\xi, \theta)$ and works in a firm with productivity Z and employment across types given by $\mathbf{n}$. We will sometimes refer to $\Phi \equiv (\mathbf{n}, Z)$ as the firm's "type". Let $n_\varsigma(\Phi)$ denote the measure of the family's members with idiosyncratic type ς who work for a firm with type Φ. Let $\mathbf{m}$ denote the distribution of the $n_\varsigma(\Phi)$s across firm and idiosyncratic types.

The family's problem can now be formally described as follows. The value function of the family is given by

$$\mathcal{V}(\mathbf{m}) = u(\mathsf{c}) - \sum\nolimits_\varsigma \int_\Phi n_\varsigma(\Phi)\, \xi\nu(h_\varsigma(\Phi))\, \mathrm{d}\Phi + \beta\mathbb{E}[\mathcal{V}(\mathbf{m}')], \tag{31}$$

where c represents per capita consumption and, to recall, $\nu(h) = h^{1+\varphi}/(1+\varphi)$ describes the disutility over time worked. The budget constraint is

$$\mathsf{c} + \mathsf{b} = \sum\nolimits_\varsigma \int_\Phi n_\varsigma(\Phi)\, W_\varsigma(\Phi)\, \mathrm{d}\Phi + (1-\mathsf{n})\, \underline{W} + (1+\mathfrak{r})\, \mathsf{b}_{-1},$$

where b (b_{-1}) indicates the end (start) of period value of the family's bond portfolio, which pays a risk-free rate $\mathfrak{r}$; n is total family employment, $\mathsf{n} \equiv \sum_\varsigma \int_\Phi n_\varsigma(\Phi)\mathrm{d}\Phi$; and $\underline{W}$ is income that accrues per period of nonemployment.

Finally, consider the law of motion of family employment. The measure of family members of type τ employed at a firm of type ϕ next period is given by

$$n'_\tau(\phi) = \sum\nolimits_\varsigma \int_\Phi (1 - s_\tau(\phi))\, \lambda_\tau n_\varsigma(\Phi)\, g(\phi|\Phi)\, \mathrm{d}\Phi + \mathfrak{p}\mathfrak{g}(\phi)\, \lambda_\tau (1-\mathsf{n}), \tag{32}$$

which says that $n'_\tau(\phi)$ reflects flows to type (τ, ϕ) (i) from among the currently employed as well as (ii) from unemployment. The first term on the right side relates to (i), where $g(\phi|\Phi)$ reports the density of firms with (new) type ϕ given an initial type Φ; and $(1 - s_\tau(\phi))\, \lambda_\tau n_\varsigma(\Phi)$ is the measure of workers of (former) type (ς, Φ) who draw new idiosyncratic type τ and remain with the firm after its transition to type ϕ. The second term relates to (ii), where $\mathfrak{p}$ is the probability of an unemployed family member meeting a hiring firm; and $\mathfrak{g}(\phi)$ is the density of *hiring* firms with type ϕ.[8]

The worker's surplus is the marginal value of an employed worker to the family. Accordingly, suppose an unemployed member with idiosyncratic type ς joins a firm Φ. Differentiating

[8]Strictly speaking, $(1-\mathsf{n})$ in (32) should be premultiplied by the retention rate at hiring firms to account for the possibility that a hiring firm will separate from the worker once it learns his type. In our model, though, we verify that firms never simultaneously hire and fire.

(31) with respect to n_ς gives

$$\frac{\partial \mathcal{V}(\mathbf{m})}{\partial n_\varsigma(\Phi)} = \left(W_\varsigma(\Phi) - \underline{W}\right) u'(\mathsf{c}) - \xi\nu\left(h_\varsigma(\Phi)\right) + \beta\mathbb{E}\left[\frac{\partial n'_\tau(\phi)}{\partial n_\varsigma(\Phi)} \frac{\partial \mathcal{V}(\mathbf{m}')}{\partial n'_\tau(\phi)}\right].$$

This allocation of labor alters flow utility by $\left(W_\varsigma(\Phi) - \underline{W}\right) u'(\mathsf{c}) - \xi\nu\left(h_\varsigma(\Phi)\right)$.[9] The worker earns a market premium of $W_\varsigma - \underline{W}$, which is "priced" in utils according to $u'(\mathsf{c})$. Equivalently, we can translate real goods into utils using the marginal value of wealth (ℓ), noting that the FOC for consumption is $u'(\mathsf{c}) \equiv \ell$. The assignment of the worker to employment also costs a loss of utility equal to $\xi\nu(h_\varsigma)$.

The forward term reflects the marginal future value of assigning the worker to firm type Φ this period. There are two considerations to bear in mind with regard to this term. First, a match formed this period will be maintained in certain states of nature next period, and the family will derive a surplus from the ongoing relationship. However, a worker matched this period is one who *cannot* be matched next period. Thus, the opportunity cost of deploying an unemployed member to firm Φ now is the anticipated surplus afforded by hiring firms next period. Letting F denote the (sub)set of hiring types and using the law of motion (32), both of these elements of the forward term can be seen by writing it out more explicitly,

$$\begin{aligned}\frac{\partial \mathcal{V}(\mathbf{m})}{\partial n_\varsigma(\Phi)} &= \left(W_\varsigma(\Phi) - \underline{W}\right)\ell - \xi\nu\left(h_\varsigma(\Phi)\right) \\ &\quad + \beta\mathbb{E}\left[\sum\nolimits_\tau \int_\phi \left(1 - s_\tau(\phi)\right)\lambda_\tau \frac{\partial \mathcal{V}(\mathbf{m}')}{\partial n'_\tau(\phi)} g(\phi|\Phi)\,\mathrm{d}\phi - \sum\nolimits_\tau \int_{\phi\in\mathsf{F}} \mathfrak{p}\lambda_\tau \frac{\partial \mathcal{V}(\mathbf{m}')}{\partial n'_\tau(\phi)} \mathfrak{g}(\phi)\,\mathrm{d}\phi\right]\end{aligned}$$

The first term inside brackets is the expected future surplus from the match made this period, taking account of the evolution of future firm and idiosyncratic types. The second term inside brackets is the expected surplus from a new match formed next period.

We can now proceed from this expression to what we present in the main text. First, observe that the expected surplus from a new match formed next period is independent of both ς and Φ and define

$$\widetilde{\mathcal{U}}' \equiv \beta \sum\nolimits_\tau \int_{\phi\in\mathsf{F}} \mathfrak{p}\lambda_\tau \frac{\partial \mathcal{V}(\mathbf{m}')}{\partial n'_\tau(\phi)} \mathfrak{g}(\phi)\,\mathrm{d}\phi. \tag{33}$$

This term will only vary with family-level risk. Second, modulo ℓ, the payoffs in the Bellman equation depend directly only on the state of the firm with which the worker is matched—

[9] Note that the family chooses n_ς taking as given the decisions over earnings and working time that will be made at the firm level.

they are independent of $\mathbf{m}$. Accordingly, define $\widetilde{\mathcal{W}}_\varsigma(\mathbf{n}, Z) \equiv \partial \mathcal{V}(\mathbf{m}) / \partial n_\varsigma(\Phi)$ (with the understanding that $\widetilde{\mathcal{W}}$ can vary with ℓ out of steady state). Third, we anticipate that since types are i.i.d., a firm's future employment $\mathbf{n}'$ will depend only on its overall size N (the sum of the elements of $\mathbf{n}$) and Z'. Hence, given N, the density function of $Z'|Z$ is sufficient to track the firm's "type". Putting these pieces together, defining $\mathcal{W}_\varsigma(\mathbf{n}, Z) \equiv \widetilde{\mathcal{W}}_\varsigma(\mathbf{n}, Z) / \ell$, and, correspondingly, $\mathcal{U} \equiv \widetilde{\mathcal{U}}/\ell$, we have

$$\begin{aligned}\mathcal{W}_\varsigma(\mathbf{n}, Z) &= W_\varsigma(\mathbf{n}, Z) - \underline{W} - \frac{\xi}{\ell}\nu(h_\varsigma(\mathbf{n}, Z)) - \mathbb{E}\left[\frac{\ell'}{\ell}\mathcal{U}'\right] \\ &\quad + \beta\mathbb{E}\left[\frac{\ell'}{\ell}\sum\nolimits_\tau \int (1 - s_\tau(\mathbf{n}'(N, Z'), Z'))\lambda_\tau \mathcal{W}_\varsigma(\mathbf{n}'(N, Z'), Z')\,\mathrm{d}F(Z'|Z)\right].\end{aligned}$$

The worker's surplus is now denominated in units of the numeraire. Therefore, the marginal value of time is scaled by $1/\ell$, and the forward terms involve expectations over the relative price of future consumption (e.g., ℓ'/ℓ). However, since the family insures against idiosyncratic and firm risk, the model is solved and estimated under $\ell' = \ell$ and $\mathcal{U}' = \mathcal{U}$ (and so the expectation operator over $\mathcal{U}'$ can be dropped). In this case, we can define the outside option more broadly as $\mu \equiv \underline{W} + \mathcal{U}$, which is the parameter estimated in the main text.

Going one step further, we can back out a value of $\underline{W}$ once we have an estimate of μ and given the structure of (33). Recall that $\mathcal{W}_\tau = (\eta/(1-\eta)) \cdot (\mathcal{J}_\tau + \underline{c})$ by surplus sharing and that the marginal value of labor, $\mathcal{J}_\tau$, at *every* hiring firm equals the hiring cost, $\bar{c}$. It follows from (33) that $\mathcal{U} = \ell^{-1}\widetilde{\mathcal{U}} = \ell^{-1}\beta\mathfrak{p}\mathbb{E}[\ell\mathcal{W}'_\tau] = \beta\mathfrak{p}\eta(\bar{c} + \underline{c})/(1-\eta)$. Therefore, using an estimate of $\mathfrak{p}$ from Elsby et al (2013), we can compute the implied value of $\mathcal{U}$, enabling us to back out the model's prediction for $\underline{W}$. This result is cited in Section 4.2.

Finally, the simple form that (33) takes after accounting for surplus sharing will carry over if idiosyncratic types are persistent, a case that is treated below. The reason is that the worker's surplus at any hiring firm must be proportional to $\bar{c} + \underline{c}$, and, thus, the expected surplus has to be independent of type. Of course, the same surplus is sustained across different types precisely by diffusing working time and earnings across types. Thus, if types are persistent, future working time and earnings will reflect current type, but the value of the match is anchored by $\bar{c} + \underline{c}$.

C.5 Earnings under persistent idiosyncratic types

In this section, we solve for earnings assuming a worker's productivity and leisure preference are persistent and potentially correlated. Recall, workers draw from preferences $\boldsymbol{\xi} \equiv \{\xi_1, ..., \xi_K\}$ and productivities $\boldsymbol{\theta} \equiv \{\theta_1, ..., \theta_L\}$. We let $\varsigma \equiv (\xi, \theta)$ denote a type, of which there are $M \equiv K \times L$ values. The probability of transition from type τ to type ς is denoted by $P_{\tau\varsigma}$. More generally, we let $\mathbb{P} \equiv \{P_{\tau\varsigma}\}$ denote the $M \times M$ transition matrix of types. The rows of $\mathbb{P}$ refer to this period's type, and the columns refer to next period's type.

The first step is to derive the firm's surplus. To this end, we first briefly sketch the firm's labor demand policy rules.

C.5.1 Labor demand

Hiring. At the start of next period, incumbent workers will draw anew from the distribution of types (though they may redraw their original type). Let m_ς denote the measure of type-ς workers after new types have been drawn *but before* the firm has adjusted employment, e.g., $m_\varsigma \equiv \sum_\tau P_{\tau\varsigma} n_\tau$ where n_τ is the measure "carried into" the period.

After the firm observes incumbent workers' types, suppose it wishes to hire. Recall that new hires are anonymous at the time they contact the firm (their types are drawn subsequently). Therefore, a law of large numbers implies that the measure of new workers of type ς equals $\widetilde{\lambda}_\varsigma \Delta N'$, where $\Delta N' \equiv N' - N$ is the growth in the firm's workforce and $\widetilde{\lambda}_\varsigma$ is the share of type ς in the pool of *job applicants*. Note that, when types are persistent, $\widetilde{\lambda}_\varsigma$ will not in general equal the share of type ς in the *population* (which we have denoted by λ_ς elsewhere). The reason is that the pool of applicants will consist of relatively "unattractive" types (e.g., types for whom the marginal match surplus is relatively low). The firm takes $\widetilde{\lambda}_\varsigma$ as given, and one can characterize optimal employment policy for any distribution of $\widetilde{\lambda}_\varsigma$. However, one would have to further characterize the composition of the job applicant pool in order to quantitatively study this version of the model.

To proceed, let $\mathbf{n} \equiv \{n_\tau\}$ denote, more generally, the $1 \times M$ vector of employment across types at the start of next period. The distribution of employment after new types are drawn, but before adjustment, is

$$\mathbf{m} \equiv \mathbf{n}\mathbb{P},$$

also an $1 \times M$ vector. We can then write $\mathbf{n}'$, employment *after* new hires are made, as[10]

$$\mathbf{n}' = \mathbf{m} + \widetilde{\boldsymbol{\lambda}} \Delta N'. \tag{34}$$

It is now instructive to define

$$\tilde{\Pi}(\mathbf{n}', Z') \equiv \pi(\mathbf{n}', Z') + \beta \int \Pi(\mathbf{n}'\mathbb{P}, Z'') \, \mathrm{d}F(Z''|Z'), \tag{35}$$

[10] Again, we assume the firm does not separate in the same period that it hires.

which describes the present value of choosing $\mathbf{n}'$ given Z', gross of hiring costs. As in the main text, $\pi(\mathbf{n}', Z') \equiv \hat{G}(\mathbf{n}', Z') - \sum_\tau n_\tau W_\tau(\mathbf{n}', Z')$ is flow profit after optimizing out working time. The vector of employment at the start of the subsequent period, after types have been drawn (but before adjustment), is $\mathbf{m}' = \mathbf{n}'\mathbb{P}$. Lemma 1 can be extended to verify that Π is supermodular, a property that we will exploit in what follows.

Now defining $\Pi^+(\mathbf{m}, Z')$ as the present value of a hiring firm, we have

$$\begin{aligned} \Pi^+(\mathbf{m}, Z') &= \max_{N'} \left\{ \tilde{\Pi}(\mathbf{n}', Z') - \bar{c}\Delta N' \right\} \\ &\equiv \max_{\Delta N'} \left\{ \pi(\mathbf{n}', Z') + \beta \int \Pi(\mathbf{n}'\mathbb{P}, Z'') \, \mathrm{d}F(Z''|Z') - \bar{c}\Delta N' \right\}. \end{aligned} \tag{36}$$

Using the chain rule and (34), the FOC for $\Delta N'$ is

$$\frac{\partial}{\partial \Delta N'} \tilde{\Pi}(\mathbf{n}', Z') = \sum_\varsigma \widetilde{\lambda}_\varsigma \left[\frac{\partial \pi(\mathbf{n}', Z')}{\partial n'_\varsigma} + \beta \mathcal{E}(\mathbf{n}', Z') \right] = \bar{c}, \tag{37}$$

where

$$\mathcal{E}(\mathbf{n}', Z') \equiv \sum_\tau P_{\varsigma\tau} \frac{\partial}{\partial m'_\tau} \int \Pi(\mathbf{n}'\mathbb{P}, Z'') \, \mathrm{d}F(Z''|Z')$$

is the expected future marginal value of type-ς labor. By supermodularity, the FOC (37) implies a threshold value of productivity $\hat{Z}_0(\mathbf{m})$ defined by

$$\frac{\partial}{\partial \Delta N'} \tilde{\Pi}\left(\mathbf{m}, \hat{Z}_0(\mathbf{m})\right) = \bar{c}$$

such that the firm hires only if $Z' > \hat{Z}_0(\mathbf{m})$. The optimal choice of employment, conditional on $Z' > \hat{Z}_0(\mathbf{m})$, can then be denoted by $\mathbf{n}^+(Z')$.

Firing. The present value of a firing firm is

$$\begin{aligned} \Pi^-(\mathbf{m}, Z') &= \max_{\mathbf{n}'} \left\{ \tilde{\Pi}(\mathbf{n}', Z') - \underline{c} \sum_\tau [m_\tau - n'_\tau] \right\} \\ &\equiv \max_{\mathbf{n}'} \left\{ \pi(\mathbf{n}', Z') + \beta \int \Pi(\mathbf{n}'\mathbb{P}, Z'') \, \mathrm{d}F - \underline{c} \sum_\tau [m_\tau - n'_\tau] \right\}, \end{aligned} \tag{38}$$

where $m_\tau \leq n'_\tau$ (and $\mathrm{d}F$ is used again to abbreviate $\mathrm{d}F(Z''|Z')$). Type ς is the first type to be separated if Z' is such that

$$\left. \frac{\partial \tilde{\Pi}(\mathbf{n}', Z')}{\partial n'_\varsigma} \right|_{\mathbf{n}'=\mathbf{m}} \equiv \left. \frac{\partial \pi(\mathbf{n}', Z')}{\partial n'_\varsigma} + \beta \mathcal{E}(\mathbf{n}', Z') \right|_{\mathbf{n}'=\mathbf{m}} < -\underline{c}.$$

By supermodularity, this implies a threshold $\hat{Z}_1^-(\mathbf{m})$ such that the firm separates from type

ς if Z' falls below $\hat{Z}_1^-(\mathbf{m})$. In this event, the firm chooses n'_ς to satisfy

$$\frac{\partial \tilde{\Pi}\left(\left[n'_\varsigma, \mathbf{m}_{/1}\right],\ Z'\right)}{\partial n'_\varsigma} = -\underline{c},$$

where $\mathbf{m}_{/1}$ denotes the vector of (initial) employment levels for all types *but the first type separated* (ς in this case). The optimal choice of type-ς labor is then denoted by $n'_\varsigma = \mathfrak{n}_\varsigma^-\left(\mathbf{m}_{/1}, Z'\right)$.

Now, suppose Z' is lowered below $\hat{Z}_1^-(\mathbf{m})$. The quantity of n'_ς is adjusted to ensure the marginal value of type-ς labor demand is maintained at $-\underline{c}$. By the supermodularity of $\tilde{\Pi}$, a decline in Z' and in the complementary labor input n'_ς reduces the marginal value of all other types. At the point where

$$\frac{\partial \pi\left(\mathbf{n}', Z'\right)}{\partial n'_\tau} + \beta \mathcal{E}\left(\mathbf{n}', Z'\right) = -\underline{c}, \quad \text{with } \mathbf{n}' \equiv \left[\mathfrak{n}_\varsigma^-\left(\mathbf{m}_{/1}, Z'\right), \mathbf{m}_{/1-}\right]$$

the firm is just indifferent to separation from type $\tau \neq \varsigma$. This implies another threshold $\hat{Z}_2^-\left(\mathbf{m}_{/1}\right)$ such that if $Z' < \hat{Z}_2^-\left(\mathbf{m}_{/1}\right)$, the firm separates from type τ. This decline in type-τ labor demand reduces, in turn, the marginal value of type-ς workers and triggers more separations from the latter type. Thus, the firm separates from *both* types when $Z' < \hat{Z}_2^-\left(\mathbf{m}_{/1}\right)$.

As Z' falls still further, this analysis is repeated for the other types. We thereby derive a list of thresholds $\left\{\hat{Z}_i^-\right\}_{i=1}^M$ such that the firm separates from i types when $Z' \in \left[\hat{Z}_{i+1}^-,\ \hat{Z}_i^-\right]$. In the context of our discussion, type τ is ranked first in the order of types to be separated. Thus, this type is also referred to as the rank-1 type. Likewise, type τ is the rank-2 type.

C.5.2 Earnings bargain

Recall that the surplus sharing rule sets

$$\mathcal{W}_\varsigma(\mathbf{n}, Z) = \frac{\eta}{1-\eta}\left[\mathcal{J}_\varsigma(\mathbf{n}, Z) + \underline{c}\right], \tag{39}$$

where $\mathcal{W}_\varsigma(\mathbf{n}, Z)$ is the surplus of a worker of type ς; and $\mathcal{J}_\varsigma(\mathbf{n}, Z)$ is the firm's marginal surplus vis a vis type-ς workers; and η is the worker's bargaining power.

Firm's surplus. The firm's marginal surplus, $\mathcal{J}_\varsigma(\mathbf{n}, Z)$, is given by

$$\mathcal{J}_\varsigma(\mathbf{n}, Z) = \frac{\partial \tilde{\Pi}(\mathbf{n}, Z)}{\partial n_\varsigma} = \frac{\partial}{\partial n_\varsigma}\left[\pi(\mathbf{n}, Z) + \beta \int \Pi\left(\mathbf{m}, Z'\right) \mathrm{d}F\left(Z'|Z\right)\right], \tag{40}$$

where, to recall, $\mathbf{m} \equiv \mathbf{n}\mathbb{P}$. The forward value in (40) can be decomposed using the labor demand decision rules according to

$$\begin{aligned} \int \Pi\left(\mathbf{m}, Z'\right) \mathrm{d}F = \textstyle\sum_{i=1}^{M} \int_{\hat{Z}_{i+1}^{-}(\mathbf{m})}^{\hat{Z}_{i}^{-}(\mathbf{m})} \Pi^{-}\left(\mathbf{m}, Z'\right) \mathrm{d}F \\ + \int_{\hat{Z}_{1}^{-}(\mathbf{m})}^{\hat{Z}_{0}(\mathbf{m})} \Pi^{0}\left(\mathbf{m}, Z'\right) \mathrm{d}F + \int_{\hat{Z}_{0}(\mathbf{m})} \Pi^{+}\left(\mathbf{m}, Z'\right) \mathrm{d}F, \end{aligned} \tag{41}$$

where $\hat{Z}_{M+1}^{-} \equiv \min\{Z'\}$ and Π^0 denotes the value of the firm in states of the world where it neither fires nor hires. Differentiating (41) with respect to n_ς and using Leibniz's rule,

$$\begin{aligned} \tfrac{\partial}{\partial n_\varsigma} \int \Pi\left(\mathbf{m}, Z'\right) \mathrm{d}F = \int^{\hat{Z}_{1}^{-}(\mathbf{m})} \textstyle\sum_\tau P_{\varsigma\tau} \tfrac{\partial}{\partial m_\tau} \Pi^{-}\left(\mathbf{m}, Z'\right) \mathrm{d}F \\ + \int_{\hat{Z}_{1}^{-}(\mathbf{m})}^{\hat{Z}_{0}(\mathbf{m})} \textstyle\sum_\tau P_{\varsigma\tau} \tfrac{\partial}{\partial m_\tau} \Pi^{0}\left(\mathbf{m}, Z'\right) \mathrm{d}F + \int_{\hat{Z}_{0}(\mathbf{m})} \textstyle\sum_\tau P_{\varsigma\tau} \tfrac{\partial}{\partial m_\tau} \Pi^{+}\left(\mathbf{m}, Z'\right) \mathrm{d}F. \end{aligned} \tag{42}$$

We consider each of the three components of (42) in turn. First, consider the case of inaction. It is easy to see that the present value of the firm in this region, denoted by $\Pi^0\left(\mathbf{m}, Z'\right)$, coincides with $\tilde{\Pi}\left(\mathbf{m}, Z'\right)$. Differentiating with respect to m_τ and recalling (40) yields

$$\frac{\partial}{\partial m_\tau} \Pi^0\left(\mathbf{m}, Z'\right) = \mathcal{J}_\tau\left(\mathbf{m}, Z'\right). \tag{43}$$

Next, consider the separation regime. Suppose types ranked j or lower are separated. If type τ is among the separated, the Envelope condition associated with (38) implies

$$\frac{\partial}{\partial m_\tau} \Pi^{-}\left(\mathbf{m}, Z'\right) = \mathcal{J}_\tau\left(\mathbf{n}^{-}\left(\mathbf{m}_{/j}\right), \mathbf{m}_{/j}, Z'\right) = -\underline{c}, \tag{44}$$

where $\mathbf{n}^{-}\left(\mathbf{m}_{/j}\right) \equiv \left\{\mathfrak{n}_i^{-}\left(\mathbf{m}_{/j}\right)\right\}_{i=1}^{j}$ concatenates the separation policies $\mathfrak{n}_i^{-}$ for ranks $i \leq j$ that are separated. Alternatively, suppose type τ is not separated. By the Envelope theorem, a perturbation to type-τ employment has no first-order effect for ranks $i \leq j$. Therefore, (38) and (40) imply that, in this case,

$$\frac{\partial}{\partial m_\tau} \Pi^{-}\left(\mathbf{m}, Z'\right) = \mathcal{J}_\tau\left(\mathbf{n}^{-}\left(\mathbf{m}_{/j}\right), \mathbf{m}_{/j}, Z'\right) > -\underline{c}. \tag{45}$$

Combining (44) and (45) and defining the separation rate corresponding to type τ by $s_\tau\left(\mathbf{m}, Z'\right) \geq 0$, we can write the marginal value of type-τ labor as

$$\frac{\partial}{\partial m_\tau} \Pi^{-}\left(\mathbf{m}, Z'\right) = -\underline{c} s_\tau\left(\mathbf{m}, Z'\right) + \left(1 - s_\tau\left(\mathbf{m}, Z'\right)\right) \mathcal{J}_\tau\left(\mathbf{n}^{-}\left(\mathbf{m}_{/j}\right), \mathbf{m}_{/j}, Z'\right).$$

Alternatively, noting that optimal separation policy requires $s_\tau\left(\mathcal{J}_\tau + \underline{c}\right) = 0$, we can write the latter equation simply as

$$\frac{\partial}{\partial m_\tau}\Pi^-\left(\mathbf{m},Z'\right)=\mathcal{J}_\tau\left(\mathbf{n}^-\left(\mathbf{m}_{/j}\right),\mathbf{m}_{/j},Z'\right), \tag{46}$$

bearing in mind that this will equal $-\underline{c}$ if type τ is being separated.

Last, suppose the firm hires. An increase in m_τ has, in principle, direct and indirect effects. The direct effect is that a higher m_τ, all else equal, raises n'_τ one for one, as indicated by the law of motion (34). The indirect effect refers to the implication of a higher m_τ for the optimal choice of $\Delta N'$, and how this is channeled through the labor demands of all types. These effects can be seen below, which follows by differentiating (36):

$$\begin{aligned}\tfrac{\partial}{\partial m_\tau}\Pi^+\left(\mathbf{m},Z'\right)=\tfrac{\partial\pi\left(\mathbf{n}^+(Z'),Z'\right)}{\partial n'_\tau}+\beta\mathcal{E}\left(\mathbf{n}^+\left(Z'\right),Z'\right)\\+\left\{\textstyle\sum_\varsigma\widetilde{\lambda}_\varsigma\left(\tfrac{\partial\pi\left(\mathbf{n}^+(Z'),Z'\right)}{\partial m'_\varsigma}+\beta\mathcal{E}\left(\mathbf{n}^+\left(Z'\right),Z'\right)\right)-\bar{c}\right\}\tfrac{\partial\Delta N'}{\partial m_\tau}.\end{aligned}$$

However, the second line must be zero: the expression in brackets {} is the FOC for the choice of $\Delta N'$. Thus, we have

$$\frac{\partial}{\partial m_\tau}\Pi^+\left(\mathbf{m},Z'\right)=\frac{\partial\pi\left(\mathbf{n}^+\left(Z'\right),Z'\right)}{\partial n'_\tau}+\beta\mathcal{E}\left(\mathbf{n}^+\left(Z'\right),Z'\right)\equiv\mathcal{J}_\tau\left(\mathbf{n}^+\left(Z'\right),Z'\right). \tag{47}$$

Substituting (43), (46), and (47) into (42), inserting the resulting expression into (40) taking account of $\mathbf{m}\equiv\mathbf{n}\mathbb{P}$, and using the definition of period profit π, we can write the firm's surplus for type ς as

$$\begin{aligned}\mathcal{J}_\varsigma\left(\mathbf{n},Z\right)=\tfrac{\partial\hat{G}(\mathbf{n},Z)}{\partial n_\varsigma}-W_\varsigma\left(\mathbf{n},Z\right)-\textstyle\sum_\tau n_\tau\tfrac{\partial W_\tau(\mathbf{n},Z)}{\partial n_\varsigma}+\beta\int_{\hat{Z}_1^-(\mathbf{m})}^{\hat{Z}_0(\mathbf{m})}\sum_\tau P_{\varsigma\tau}\mathcal{J}_\tau\left(\mathbf{m},Z'\right)\mathrm{d}F\\+\beta\textstyle\sum_{i=1}^M\int_{\hat{Z}_{i+1}^-(\mathbf{m})}^{\hat{Z}_i^-(\mathbf{m})}\sum_\tau P_{\varsigma\tau}\mathcal{J}_\tau\left(\mathbf{n}^-\left(\mathbf{m}_{/i}\right),\mathbf{m}_{/i},Z'\right)\mathrm{d}F\\+\beta\textstyle\int_{\hat{Z}_0(\mathbf{m})}\sum_\tau P_{\varsigma\tau}\mathcal{J}_\tau\left(\mathbf{n}^+\left(Z'\right),Z'\right)\mathrm{d}F.\end{aligned}$$

Using $\mathbf{n}'\left(\mathbf{m},Z'\right)$ to refer generically to the optimal choice of employment next period, we can condense the forward terms and write this more compactly as

$$\mathcal{J}_\varsigma\left(\mathbf{n},Z\right)=\frac{\partial\hat{G}\left(\mathbf{n},Z\right)}{\partial n_\varsigma}-W_\varsigma\left(\mathbf{n},Z\right)-\textstyle\sum_\tau n_\tau\dfrac{\partial W_\tau\left(\mathbf{n},Z\right)}{\partial n_\varsigma}+\beta\sum_\tau P_{\varsigma\tau}\mathbb{E}_{Z'}\left[\mathcal{J}_\varsigma\left(\mathbf{n}',Z'\right)\mid Z\right].$$

Worker's surplus. The value of a job to a worker of type ς at a firm $(\mathbf{n}, Z)$ is

$$\mathcal{W}_\varsigma(\mathbf{n},Z) = \begin{array}{c} W_\varsigma(\mathbf{n},Z) - \xi\nu_\varsigma(\mathbf{n}) - \mu \\ +\beta \sum_\tau P_{\varsigma\tau}\mathbb{E}_{Z'}\left[(1-s_\tau(\mathbf{n}',Z'))\,\mathcal{W}_\tau(\mathbf{n}',Z') \mid Z\right]. \end{array}$$

where $\xi\nu_\varsigma(\mathbf{n}) \equiv \xi h_\varsigma(\mathbf{n})^{1+\varphi}/(1+\varphi)$ is the disutility from work and μ is flow value of non-work time. Note that the latter is not indexed by type. As shown in Appendix C.4, the expected surplus of future employment, which is encased in μ, is anchored by the hiring and firing costs and independent of type. Now substituting for $\mathcal{W}_\tau(\mathbf{n}',Z')$ using the surplus sharing rule (39) and recalling that the optimal separation policy requires $s_\tau(\mathbf{n}',Z')\,(\mathcal{J}_\varsigma(\mathbf{n}',Z') + \underline{c}) = 0$ yields

$$\mathcal{W}_\varsigma(\mathbf{n},Z) = W_\varsigma(\mathbf{n},Z) - \xi\nu_\varsigma(\mathbf{n}) - \mu + \beta\frac{\eta}{1-\eta}\sum_\tau P_{\varsigma\tau}\mathbb{E}_{Z'}\left[\mathcal{J}_\varsigma(\mathbf{n}',Z') + \underline{c} \mid Z\right].$$

Solution. The earnings bargain can now be obtained by substituting $\mathcal{J}_\varsigma(\mathbf{n}, Z)$ and $\mathcal{W}_\varsigma(\mathbf{n}, Z)$ into (39) and canceling terms. This reveals

$$W_\varsigma(\mathbf{n}, Z) = \eta\left(\frac{\partial \hat{G}(\mathbf{n}, Z)}{\partial n_\varsigma} - \sum_\tau n_\tau \frac{\partial W_\tau(\mathbf{n}, Z)}{\partial n_\varsigma} + r\underline{c}\right) + (1-\eta)\left(\xi\nu_\varsigma(\mathbf{n}) + \mu\right).$$

This is identical to the solution in the main text.

D Further robustness analysis

In this section, we further examine the robustness of our results. Section D.1 concerns various assumptions on the distribution of idiosyncratic shocks ξ and θ. Section D.2 re-examines the estimates of the firm-level components of earnings and working time and their implications for the model's structural parameters. Section D.3 assesses the model's fit against several moments that were not targeted in the structural estimation.

D.1 The distribution of idiosyncratic types

We revisit four features of the distributions of idiosyncratic preferences, ξ, and productivities, θ. First, we consider the implications of our assumption that ξ and θ are purely transitory. Second, we examine a finer discretization of ξ and θ than used in the main text. Third, we deviate from our assumption that both ξ and θ are uniform random variables. Fourth, and finally, we explore a case where ξ and θ are correlated. Note that in each subsection, we examine only a single deviation from the model; outside of this deviation, the assumptions in the paper apply.

D.1.1 Persistence

We have assumed that idiosyncratic preferences, ξ, and productivities, θ, are each uncorrelated over time. The absence of persistence greatly simplifies the model, as it implies that we do not have to track the distribution of workers across types from one period to the next.

Although the case with persistence is computationally intractable, we can still reflect on how its absence is likely to shape model outcomes. We observe, first, that the degree of persistence in ξ and θ has no direct effect on the choice of working time, which solves a static decision problem. Similarly, the earnings bargain takes the same form under persistence, as shown in Appendix C.5.

Finally, there is the question of how persistence in type (ξ and θ) affects employment demand. Consider first the *separation* decision. Higher persistence of types will induce less labor hoarding: since a worker's type is less likely to revert to mean, firms are less inclined to retain a worker if he draws a relatively unattractive, or low-surplus, type. Thus, turnover in the model will rise, and the MSM estimator will infer that firm-level productivity shocks must be smaller to dampen the incentive to separate (see related discussion in Section 5.1 of the main text). The model will then "need" a higher Frisch elasticity to offset the effect of smaller firm-level shocks on the variance of firm-level working time.

Turning to the hiring margin, the effect of persistence can be more subtle, as alluded to in Appendix C.5. Suppose that an individual worker's type is unknown to the firm at the time of hiring, but the firm does know the distribution of types among prospective hires. The firm chooses the *total* number of hires, with the *composition* of hires given by the latter distribution. Under i.i.d. types, the composition of *prospective hires* mirrors the cross section of types in the *population.* Intuitively, if types are i.i.d., past labor market status (i.e., a separation) conveys no information about current type.

This equivalence will no longer necessarily hold when types are persistent. The reason for this is straightforward. The pool of potential hires consists of workers separated from their prior employers (and in search of a job). It follows that, if types are persistent, the pool of prospective hires will be made up disproportionately of unattractive types. We can conjecture about some of the implications of this. For example, firms may end up deferring hiring until productivity (Z) is higher, e.g., setting a higher hiring threshold. However, a fuller analysis requires a careful treatment of the determination of aggregate nonemployment (the pool of job searchers) in market equilibrium. Unfortunately, this extension lies beyond the scope of our current analysis.

D.1.2 Discretization

In the main text, we assume that each of ξ and θ is drawn from a three-point distribution. We now argue that, under a suitable rescaling, our baseline parameter estimates will be largely robust to the choice of three points, and we confirm this claim for the case where each type is drawn from a four-point distribution. Later, in Appendix F, we show that the results of our counterfactual experiment of Section 4.3 are also robust to the addition of more types.

To proceed, it is helpful to slightly rewrite our production function as

$$\Gamma = \tilde{Z} \left(\sum\nolimits_{\xi,\theta} \upsilon_{\xi,\theta} \left(\theta n_{\xi,\theta} h_{\xi,\theta} \right)^{\rho} \right)^{\alpha/\rho}, \tag{48}$$

where the shares, or weights, sum to 1, $\sum_{\xi,\theta} \upsilon_{\xi,\theta} = 1$ and $\tilde{Z}$ is a productivity shifter. When we omit $\upsilon_{\xi,\theta}$ in the main text of the paper, we implicitly assume $\upsilon_{\xi,\theta} = \upsilon = 1/M$ for all (ξ, θ) (where M is the total number of pairs $\varsigma \equiv (\xi, \theta)$). The production function is then equivalent to,

$$\Gamma = Z \left(\sum\nolimits_{\xi,\theta} \left(\theta n_{\xi,\theta} h_{\xi,\theta} \right)^{\rho} \right)^{\alpha/\rho}, \tag{49}$$

where $Z \equiv \tilde{Z} M^{-\alpha/\rho}$. This is how the production function appears in the main text.

Now consider a case with four productivity and preference types, e.g., $M = 4 \times 4 = 16$. Importantly, the *distributions* of θ and ξ (e.g., σ_θ^2 and σ_ξ^2) are held fixed; we simply refine the discretization. Under (48), the weights $\upsilon_{\xi,\theta}$ would naturally be adjusted such that we take a proper average over $\left(\theta n_{\xi,\theta} h_{\xi,\theta} \right)^{\rho}$. Since we use (49), though, we must take care to explicitly scale up the production function by a factor of $(16/9)^{-\alpha/\rho}$ relative to our baseline case where $M = 9$.

When we scale appropriately, and recalibrate the model, we find that the parameter estimates under $M = 16$ are virtually the same as in the $M = 9$ case (see Table D1). In fact, even if we do not rescale, the introduction of more types does not have any meaningful impact on the *dynamics* of earnings, working time, or employment. The only moment that is sensitive to the discretization is the average *level* of firm size. Scaling up by $(16/9)^{-\alpha/\rho}$ very nearly reverses the impact of the change in M on this latter moment; a slight adjustment to μ is all that is needed, as shown in Table D1.

Finally, the assumption that $\upsilon_{\xi,\theta} = \upsilon$ for all (ξ, θ) can in fact be thought of as a normalization. To see this, consider the choices of working time, $h_{\xi,\theta}$, and earnings, $W_{\xi,\theta}$, when $\upsilon_{\xi,\theta}$

is unrestricted. Using (48), we have

$$h_{\xi,\theta} = (\alpha Z \Omega(\mathbf{n}))^{\frac{1}{\varphi+1-\alpha}} \left(\frac{\upsilon_{\xi,\theta}\theta^{\rho}}{\xi}\right)^{\frac{1}{\varphi+1-\rho}} n_{\xi,\theta}^{-\frac{1-\rho}{\varphi+1-\rho}} \tag{50}$$

and

$$W_{\xi,\theta} = \varkappa (\alpha Z \Omega(\mathbf{n}))^{\frac{\varphi+1}{\varphi+1-\alpha}} \left(\frac{(\upsilon_{\xi,\theta}\theta^{\rho})^{\varphi+1}}{\xi^{\rho}}\right)^{\frac{1}{\varphi+1-\rho}} n_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} + \omega, \tag{51}$$

where ω and $\varkappa$ are constants defined in the main text (see Section 1.3.4) and

$$\Omega(\mathbf{n}) \equiv \left(\sum_{\xi,\theta} \left(\frac{(\upsilon_{x,y} y^{\rho})^{\varphi+1} n_{x,y}^{\varphi\rho}}{x^{\rho}}\right)^{\frac{1}{\varphi+1-\rho}}\right)^{\frac{\alpha-\rho}{\rho}}.$$

Critically, in the solutions for working time and earnings, $\upsilon_{\xi,\theta}$ and θ^{ρ} do not enter separately but always as the product, $\upsilon_{\xi,\theta}\theta^{\rho}$. Next, using the solution for working time (50), the production function becomes

$$\Gamma = \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\frac{\alpha}{\alpha-\rho}\frac{\varphi+1-\rho}{\varphi+1-\alpha}}. \tag{52}$$

Again, $\upsilon_{\xi,\theta}$ and θ enter exclusively via $\upsilon_{\xi,\theta}\theta^{\rho}$ (in Ω). It follows from (51) and (52) that only the product $\upsilon_{\xi,\theta}\theta^{\rho}$ is relevant to flow profit, and so the choice of employment will also hinge only on it (rather than on $\upsilon_{\xi,\theta}$ and θ separately). For this reason, $\upsilon_{\xi,\theta}$ and θ cannot be separately identified, and we normalize $\upsilon_{\xi,\theta} = 1/M$ for all (ξ, θ).

D.1.3 Uniformity

We have explored alternatives to our assumption in the main text that ξ and θ are each uniformly distributed. In the end, though, we conclude that certain moments are, in a sense to be made precise, incompatible with meaningful deviations from the uniform distribution.

To proceed, we maintain that ξ is drawn from a three-point distribution but now assume that it is single-peaked such that

$$\Pr(\xi_1) = \Pr(\xi_3) = 0.3 < \Pr(\xi_2) = 0.4,$$

where $\xi_1 < \xi_2 < \xi_3$. This assumption represents a controlled departure from uniformity but will be sufficient to demonstrate our point. Meanwhile, we impose in this exercise that θ is still drawn from a three-point, uniform distribution. In a second exercise, ξ and θ reverse roles:

$$\Pr(\theta_1) = \Pr(\theta_3) = 0.3 < \Pr(\theta_2) = 0.4,$$

with $\theta_1 < \theta_2 < \theta_3$, but we maintain that ξ is uniformly distributed. The assumption of a single-peaked distribution in each case is made only for illustrative purposes. As we will see, our main message does not hinge on this; we have found similar results when we consider bimodal (or, "U" shaped) distributions.

Table D2 reports the values of the structural parameters when we calibrate the model in each of these two cases. What we want to highlight is that the variance of θ collapses even where there are modest deviations from uniformity.

To interpret these results, it is helpful to first re-solve the model under the *baseline* parameters but with the *non-uniform* distributions. We find a strongly positive, and highly counterfactual, correlation between working time, $h_{\xi,\theta}$, and daily earnings, $w_{\xi,\theta}$. To see why, suppose $N = N_{-1}$. Then, under uniform distributions, a worker can change types from $\varsigma \equiv (\xi, \theta)$ to $\varsigma' \equiv (\xi', \theta')$, but the size of his team does not change, e.g., $n_\varsigma = n_{\varsigma'}$. Any deviation from uniformity implies that, in general, $n_\varsigma \neq n_{\varsigma'}$ (even if $N = N_{-1}$). Since h_ς and w_ς are each decreasing in n_ς, changes in the size of one's team push h_ς and w_ς in the same direction, bringing out a positive correlation between the two.

The model has limited means to address this counterfactual implication when the parameters are reestimated. Larger preference (or, "supply") shocks can reduce this correlation, but changes in type size, $n_{\xi,\theta}$, already lead to excess variation in working time; further increases in σ_ξ are problematic. As a result, the model is left to infer much smaller productivity (or, "demand") shocks. Indeed, the values of σ_θ in Table D2 seem implausibly low, and more substantial deviations from the uniform distribution imply infeasible values, e.g., $\sigma_\theta < 0$. In other words, certain moments are incompatible with all but very modest departures from the uniform distribution.

At the same time, other structural parameters of interest are largely unaffected by the alternative distributions that we have considered. To see this, consider again why it is impossible to fit the data by simply raising σ_ξ to reproduce the observed comovement of $h_{\xi,\theta}$ and $w_{\xi,\theta}$ while adjusting other parameters to constrain the volatility of (idiosyncratic) working time. A smaller Frisch elasticity, $1/\varphi$, would have a modest impact on *idiosyncratic* working time fluctuations but hampers the model's capacity for matching changes in *average* working time. A lower ρ could dampen the effect of larger swings in ξ, but it would *amplify* fluctuations in $h_{\xi,\theta}$ due to changes in type size. The reason is that the (absolute) elasticity of $h_{\xi,\theta}$ with respect to $n_{\xi,\theta}$ is decreasing in ρ: if jobs are strong complements, a higher $n_{\xi,\theta}$ implies a sharply lower marginal product (for type (ξ, θ)), leading to a bigger cutback in

$h_{\xi,\theta}$. These observations explain why φ and ρ do not change very much under non-uniform distributions; the model must instead "lean" on σ_θ.

In this setting, an explicit cost of adjusting working time could offer a degree of freedom the model needs. Bigger frictions could dampen idiosyncratic working time fluctuations without facing the same tension that confronts changes in ρ. To the extent that larger frictions also shrink *average* working time, this effect could be counteracted by a higher Frisch elasticity, $1/\varphi$. Thus, a cost of adjustment could potentially support a parameterization where σ_ξ is larger but working time is not excessively volatile. We must leave this extension, which implies a nontrivial dynamic working time problem, for future work.

D.1.4 Independence

Finally, we consider a case where ξ and θ are correlated. We maintain that ξ is uniformly distributed but now introduce a probability mass function, $Q(\theta|\xi)$, that governs the distribution of θ (given ξ). Importantly, the distribution of θ will, in general, no longer be uniform. This implication has significant ramifications for the model, as we saw above.

To proceed to a quantitative analysis, we need to specify $Q(\theta|\xi)$. We propose a highly tractable form that introduces just one new parameter, q, to regulate departures from independence. The form of $Q(\theta|\xi)$ is given by

$$Q(\theta|\xi) = \begin{pmatrix} \Pr(\theta_1|\xi_1) & \Pr(\theta_2|\xi_1) & \Pr(\theta_3|\xi_1) \\ \Pr(\theta_1|\xi_2) & \Pr(\theta_2|\xi_2) & \Pr(\theta_3|\xi_2) \\ \Pr(\theta_1|\xi_3) & \Pr(\theta_2|\xi_3) & \Pr(\theta_3|\xi_3) \end{pmatrix} = \begin{pmatrix} 1/3-q & 1/3+q/2 & 1/3+q/2 \\ 1/3 & 1/3 & 1/3 \\ 1/3+q/2 & 1/3+q/2 & 1/3-q \end{pmatrix}$$

where $\xi_1 < \xi_2 < \xi_3$ and $\theta_1 < \theta_2 < \theta_3$. When $q = 0$, ξ and θ are independent. More generally, $q > (<)\ 0$ implies that the two are negatively (positively) correlated. For instance, if $q > 0$, then $\Pr(\theta_3|\xi_1) > \Pr(\theta_3|\xi_3)$: a higher draw of θ is more likely when ξ is smaller.

To explore the model's behavior under Q, we recalibrated it given a few values of q. We have set $q > 0$ in each case, but the same basic conclusions emerge if $q < 0$.[11] Two results in Table D3 are noteworthy. First, for very slight correlations (i.e., $\text{corr}(\xi,\theta) = -0.06$), the structural parameters are hardly affected. However, when the correlation deviates much further from zero, σ_θ falls dramatically, even as other parameters are reasonably stable. In this sense, Table D3 is reminiscent of Table D2: the model "leans on" very low values of σ_θ

[11] Under complementarities, ξ and θ push working time in the same direction. Therefore, a $q > 0$ (e.g., a higher θ is met by a lower ξ) will tend to dampen variation in working time, whereas a $q < 0$ will have the opposite effect. However, these implications will be second-order, especially given the small values of q that we consider. The first-order effect of a non-zero correlation is to "break" the uniform distribution of types.

to match the moments. Indeed, the table indicates that any correlation smaller than -0.12 would call for $\sigma_\theta < 0$, an *infeasible* choice. The limited scope for $q \neq 0$ suggests to us that our moments are not compatible with meaningful deviations from independent shocks.[12]

D.2 Estimating firm-level variation in earnings and working time

In this section, we present alternative estimates of the variances of firm-level components of earnings and working time. Since firms are of finite size, the (measured) firm-level means of these variables will fluctuate to some degree because of shifts in the (structural) idiosyncratic components. The estimator in Kline et al (2020) implements a correction that, in effect, filters this noise out of the variances of the firm-year effects.

The top panel of Table D4 presents the Kline et al estimates of the standard deviation of firm-year effects. The sample is the two-year stayers, and the baseline estimates refer to the results reported in Table 3 in the paper.

The revised estimates are uniformly smaller. The standard deviation of the firm component of log changes in days declines by a little over 20 percent, whereas it falls by around 12.5 percent in the case of log earnings. As a result, the variance of (the firm component of) earnings growth relative to the variance of (the firm component of) working time growth increases from 2.89 to 3.58.

In the bottom panel of Table D4, we report the structural parameter values that we recover if we substitute the moment estimates in the top panel for our baseline values. (The remainder of the baseline moments are the same.) The first difference to note is that the implied Frisch elasticity declines from 0.48 to 0.38, reflecting smaller estimated average (firm-level) working time fluctuations. A smaller Frisch elasticity reverberates through the model in a few ways. First, since it tempers changes in average working time, it also dampens fluctuations in average earnings. The corresponding decline in the variability of average earnings helps explain why bargaining power falls only modestly despite the use of the Kline et al estimator. Second, a smaller Frisch elasticity also diminishes, to some extent, the variability of the idiosyncratic component of working time. Larger idiosyncratic shocks can counteract this but ramping up σ_ξ and σ_θ will eventually exaggerate fluctuations in (idiosyncratic) earnings growth. This observation helps account for the rise in the elasticity of substitution, $(1-\rho)^{-1}$, which increases the *pass through* of preference shocks to working time relative to earnings.

[12] For the case where $\text{corr}(\xi, \theta) = -0.06$, $q = 0.0396$, whereas $q = 0.0785$ is required to induce a correlation of -0.12.

Finally, it's worth noting that the variance of firm-level shocks hardly changes even though firm-level working time and earnings are less variable. The reason, we suspect, is that the variance of Z is largely anchored by the variance of employment growth, which is unaltered. Accordingly, the revised moments implied by the Kline et al estimator are met with changes to the Frisch elasticity and other parameters rather than a reduction in the variance of firm-level TFP.

D.3 Nontargeted moments

As a final exercise, we examine how well the model fits five moments that were not targeted in estimation. Each of the five pertain to firm-level outcomes. As listed in Table D5, the first three moments are the autocorrelations of annual changes in log paid days, log annual earnings, and log employment. We highlight that these are first differences, rather than levels; the levels are positively correlated, but mean reversion in these series implies that the differences are negatively correlated. The fourth and fifth moments are the elasticities of, respectively, working time and earnings growth with respect to employment growth.

The fit with respect to the first two moments in the table is reasonably good. The model very nearly replicates the autocorrelation of log changes in working time. Given our focus on the intensive margin, this result is quite encouraging. While the model-implied autocorrelation of earnings growth is somewhat too high (equal to -0.27), it is still comparable to its empirical counterpart (equal to -0.17).

With regard to the other three moments, the model-generated values exceed their empirical analogues by more noticeable amounts. For instance, the elasticity of earnings growth with respect to employment growth is 0.28 in the model but almost zero in the data. We have explored how one might obtain a tighter fit with respect to these moments.

One potentially free parameter is the persistence, ζ, of firm-level TFP. In the paper, we set ζ based on estimates in the literature. A more persistent TFP process would have two important implications. First, more persistent changes in TFP elicit larger responses in employment relative to working time and earnings. (Recall that the latter two are effectively static decision rules and only indirectly affected by ζ.) Accordingly, a higher ζ tends to reduce the elasticities in rows (4) and (5) in Table D5. However, a second consequence of higher persistence has counterfactual implications: it *increases* the autocorrelation of employment growth, which is already too high in the model (row (3) in the table).

We have also considered changes in the returns to scale, α, but this has many of the same implications as greater persistence (see Section 5.1 in the paper). There does not appear to be a "free lunch" in either case.

In the end, we would still advocate for our approach whereby ζ and α are set to be consistent with estimates in the literature, and sensitivity analysis is performed for values in the neighborhood of these estimates. We lack the data—namely, revenue and other factors of production—that bears most directly on the values of ζ and α. Therefore, we value the discipline imposed by the external evidence.

E Identification

The mapping between moments and parameter estimates is rarely clear-cut. In the main text, we provided some intuition for why certain moments are likely to be especially informative for certain parameters. In this Appendix, we use the sensitivity matrix proposed by Andrews et al (2017) to guide a discussion of identification. The matrix is reported in Table E1.

Formally, since our model is just-identified, the sensitivity matrix takes a simple form: it is the *inverse* of the Jacobian of the model-based moments with respect to parameters. In practical terms, a column of the matrix measures the change in the parameter estimates given a one percentage point increase in the value of one of the targeted moments. The changes in parameter estimates can be understood as those required to fit the revised moment (that pertains to the column) *while preserving the fit* of the other, unrevised moments.

Several entries in the matrix echo the message of our theoretical analysis. First, a perturbation to $\operatorname{var}\left(\epsilon^W\right) / \operatorname{var}\left(\epsilon^h\right)$ has its most pronounced impact on ρ. A one percentage point increase in the latter moment lowers ρ by 0.051. Thus, as foreshadowed in Section 1 of the main text, a higher value of $\operatorname{var}\left(\epsilon^W\right) / \operatorname{var}\left(\epsilon^h\right)$ signals that idiosyncratic driving forces are being channeled into earnings rather than working time, which is indicative of greater complementarities.[13]

Second, the Frisch elasticity $(1/\varphi)$ is sensitive to changes in firm-wide, but *not within-firm*, working time dynamics. A one percentage point increase in $\sqrt{\operatorname{var}\left(\phi^h\right)}$ lowers φ by 0.011, whereas an increase in its within-firm counterpart, $\sqrt{\operatorname{var}\left(\epsilon^h\right)}$, has a negligible effect on φ. This finding is consistent with the notion that idiosyncratic variation in working time is largely uninformative for the preference parameter, φ.

Third, the responses of σ_ξ and σ_θ to a perturbation in $\sqrt{\operatorname{var}\left(\epsilon^h\right)}$ indicate that greater

[13] The parameter ρ, which is expected to mediate the effects of within-firm variation, is also sensitive to some firm-wide moments, such as $\operatorname{var}\left(\phi^h\right)$. This dissonance is to be expected to an extent. Unless the Jacobian is diagonal, a change in one moment can affect seemingly unrelated parameters, because the latter adjust to offset implications for the other moments.

within-firm variation in $h_{\xi,\theta}$ is accommodated more by preference (ξ) than productivity (θ) heterogeneity. As noted in Section 1, this result reflects that preference dispersion is "needed" more to be consistent with the observed negative covariance of working time and wages.

Several other entries in Table E1 are highly intuitive. For instance, an increase in firm size, $\mathrm{E}[N]$, implies a lower outside option, μ : if μ is small, the rents from a match are large, and so more hires are made. In addition, if there is an increase in the dispersion of employment growth, $\sqrt{\mathrm{var}\left(\Delta \ln N\right)}$, the model infers a higher standard deviation of firm TFP, σ_Z. Finally, since bargaining power η governs the rate at which shocks pass through to earnings, an increase in the variance of either within-firm or firm-level earnings growth implies a higher η.

F Counterfactual exercise

For the counterfactual simulation in Section 4.3 of the paper, we considered an unanticipated and temporary increase in ξ for a specified share of the firm's workforce. In this section, we first report results for variants of this counterfactual. Next, we show that this experiment is closely related to a temporary labor income tax change.

F.1 Sensitivity analysis

We first report a few variants on our main counterfactual exercise.

Targeting fewer workers. When we elevate ξ for just a fraction of the workforce, the smallest share that we can consider is 1/9, or one of the nine (ξ, θ) types. With more types, we could target a smaller share of the workforce, and the latter's labor supply response should be weaker than in our baseline case. Intuitively, if a big cohort cuts its working time, the marginal product of other cohorts falls. As a result, other workers also reduce their labor input, which dampens the marginal product of the targeted group and amplifies the decline in its working time. This feedback is muted when the size of the affected group is small.

However, the *quantitative* effect of an increase in the number of types beyond 9 is pretty limited. We re-run the counterfactual with four values each of ξ and θ, e.g., $M = 16$ pairs of (ξ, θ) (see Appendix D.1.2 for more on this version of the model). Thus, one type represents $1/16 = 6.25$ percent of the workforce. When we raise ξ for the type nearest to $(\xi, \theta) = (1, 1)$, its working time falls 4.8 percent. The latter result marks a small departure from the 5.1 percent decline that we observed in the main text when the experiment was run for one of 9 (ξ, θ) pairs.

The addition of even more types is unlikely to make a significant difference. To see why, consider the special case where the employment distribution across types is fixed. We can then solve analytically for the change in working time of the affected group given any number of types,

$$\mathrm{d}\ln h_{\xi,\theta} = -\frac{1}{\varphi + 1 - \rho}\left\{1 + \frac{\alpha - \rho}{\varphi + 1 - \alpha}\varpi_{\xi,\theta}\right\}\mathrm{d}\ln \xi, \tag{53}$$

where $\varpi_{\xi,\theta} \in (0,1)$ is a weight (defined in (26)) that indicates the type's share in firm output. Now consider again a 25 log-point increase in ξ. If we apply our baseline parameter values (Table 5) given $M = 16$, working time implied by (53) falls 5.3 percent for the pairs of (ξ, θ) nearest $(1, 1)$. The latter outcome is reasonably close to the 4.8 percent decline reported above, even though (53) assumes fixed employment and, strictly speaking, only holds for "small" changes in ξ. As we increase the number of types, $\varpi_{\xi,\theta}$ becomes smaller, and in the $\varpi_{\xi,\theta} \to 0$ limit, $h_{\xi,\theta}$ falls 5 percent according to (53). The difference relative to the $M = 16$ case is only about one-third of a percentage point (5 v. 5.3 percent). This exercise suggests that, even if we extended the counterfactual in the main text to include more than 16 (ξ, θ) pairs, the response of working time would not be much different from what is currently reported.

Temporary v permanent. The extensive margin reacts little to transitory shifts in ξ because of employment adjustment frictions, leaving working time to absorb the shock. It can be instructive to consider how a more persistent change in ξ alters the mix of intensive- and extensive-margin adjustments.

Suppose now that the preference shifter ξ for the median type $(\xi, \theta) = (1, 1)$ increases permanently by 25 log points. When this increase is applied to *all* workers, employment is especially responsive. To see why, recall that hires are made before types are drawn. Thus, hiring depends on the *expected* wage across types. Elevating ξ for just one type has a limited effect on the average wage. However, if ξ is permanently raised across all types, the expected cost of a new hire increases significantly. Since employment crowds out working time, the decline in the former implies that the latter does not fall as much. We now find that the (absolute) response of working time when all workers face a higher ξ is only 40 percent higher than when just one of the 9 (ξ, θ) pairs is affected. By contrast, in the main text (when only a temporary change in ξ was considered), working time fell twice as much when the higher ξ was applied to all workers.

Alternative $\Delta \ln \xi$s. The *elasticity* of working time in the model is highly robust to the size of $\Delta \ln \xi$. To illustrate, we re-ran the counterfactual with (i) $\Delta \ln \xi = 0.1$ and (ii) $\Delta \ln \xi = 0.4$ for the case with 9 pairs of types (ξ, θ). The range of elasticities across these experiments is pretty narrow. When just one pair faces a higher ξ, its elasticity varies from -0.214 (in case (i)) to -0.195 (in case (ii)). The baseline response (given $\Delta \ln \xi = 0.25$) was right in the middle of this range. When all workers face a higher ξ, the elasticity varies from -0.405 (in case (i)) to -0.376 (in case (ii)), and, again, the baseline response was almost exactly the median of the latter two.

F.2 A tax holiday interpretation

Our main counterfactual exercise is closely related to a temporary change in the labor income tax rate. Conveniently, a tax on labor income (e.g., W) can be incorporated into the working time and earnings problems by rescaling the parameters ξ and μ. Reinterpreting μ and ξ accordingly, the firm's employment problem then takes precisely the same form as in the main text. As we shall see, the driving force behind the results in this case remains the shift in ξ; the quantitative impact of a one-off pertubation to μ appears to be slight.

First, we consider how a tax, t, is manifest in the working time decision. Recall that, in the original model, we assumed the firm and worker reach the efficient outcome equating the marginal product (`mpl`) to the marginal value of time (`mvt`). A natural extension to the case with taxes is to consider a solution for working time that satisfies `mvt/mpl`$= 1 - t$, e.g., the tax drives a wedge between the marginal product and the marginal value of time.[14] Since the preference shifter ξ is the leading term of `mvt`, this approach is equivalent to scaling ξ by $1/(1-t)$ in the original working time solution.

Next, we turn to the implications for earnings. To this end, we reconsider the Brügemann et al (2019) alternating offers game underlying the earnings bargain. Now accounting for taxation, the bargain implies that the worker earns a fixed share η of the *pre*-tax joint surplus, $\mathcal{W}/(1-t) = \eta [\mathcal{J} + \mathcal{W}/(1-t)]$, where $\mathcal{W}$ is the after-tax worker surplus and $\mathcal{J}$ is the firm's surplus.[15] Note that $\mathcal{W}$ includes after-tax earnings, whereas non-market time, equal to $\mu + \xi\nu(h)$, is untaxed. Thus, dividing through $\mathcal{W}$ by $1-t$ scales μ and ξ by $1/(1-t)$. Together with the solution for working time, this observation confirms that we can incorporate the tax rate by simply reinterpreting these parameters as $\mu/(1-t)$ and $\xi/(1-t)$.

We are now in a position to consider an experiment patterned after the Icelandic tax holiday analyzed in Sigurdsson (2021). Iceland formerly withheld tax based on the prior

[14] That is, we assume the firm and worker arrive at this outcome, although we do not work out in this paper the bargaining protocol by which this happens.

[15] A corollary of this is that the worker's after-tax surplus is reduced (relative to the $t = 0$ case) by just $\eta t\mathcal{J} < t\mathcal{J}$. Intuitively, the bargain has the firm "cover" some of the worker's losses.

year's taxable income. In late 1986, the Icelandic government announced that it would switch to a withhold-as-you-go tax system in 1988. Income taxes in 1988 would be withheld based on income reported to the tax authority in that year rather than on income earned in 1987. As a result, 1987 income would not be taxed.

Our quantitative analysis of the tax holiday proceeds in two parts. First, given a pre-holiday tax rate of $t = 0.19$, we scale μ and ξ in order to reproduce our baseline moments. Then, we simulate earnings, working time, and employment outcomes when t falls to zero for a single year, after which it reverts to $t = 0.19$. We carry out the simulation under two scenarios: in one, the tax is lowered only for the median cohort (with median values of ξ and θ), whereas in the other, the tax holiday applies to all workers.

The results are reminiscent of the counterfactual presented in the main text. To see this, note that since the (rescaled) ξ declines by $t = 19$ percent, one can express the result in terms of the elasticity of working time with respect to ξ—a statistic that is comparable to what we reported in the first counterfactual. This elasticity is -0.22 when the tax rate is eliminated for only a single cohort but nearly -0.42 in the case of a universal tax holiday. In other words, the working time response nearly doubles when the policy change is universal. This finding is exactly in line with the simulation in the text. Furthermore, the size of the elasticities in each experiment are virtually the same, suggesting that the response to the change in μ is minimal. A lower μ increases the match surplus and will, in principle, stimulate an extensive-margin response. However, given the frictions, a purely transitory change in μ elicits a quantitatively negligible effect.

Interestingly, the model's elasticity under a universal tax holiday is also broadly in line with estimates from Sigurdsson (2021). The author finds that the behavioral response to Iceland's policy implies an (intensive-margin) Frisch elasticity of 0.37, or just slightly below our model's prediction.[16]

G Proxies for complementarities

The moment that we associate most closely with the degree of complementarities is the ratio of the variance of idiosyncratic log earnings changes to the variance of idiosyncratic log working time changes. For brevity, we will refer to this moment as the *ratio of idiosyncratic variances*. In the main text, we pooled data across all firms to calculate this. We now present estimates of the ratio for many detailed industries and compare the results with

[16] Sigurdsson's estimate is based on an analysis of labor earnings, as he judges that the latter's response largely reflects changes in hours rather than wage rates.

other (industry-level) indicators of complementarities.

To begin, Table G1 reports the ratio of idiosyncratic variances at the three-digit NAICS level in the Veneto data. We have data for 67 industries in total. (Below, we discuss how, and why, we map from European industry classifications to the NAICS.) The left panel reports industries whose ratios are in the lower quartile of the distribution (sorted in ascending order), and the right panel reports industries whose ratios are in the top quartile (sorted in descending order).

How might we assess this cross-industry variation? On the one hand, several entries in the table strike us as intuitive. For instance, hospitals (right panel) involve many, highly specialized and interdependent tasks, whereas carriers in the couriers and messengers industry (left panel) would appear to work largely independently of one another.[17] On the other hand, some listings can seem puzzling. For example, mining is in the top quartile of industries, even as oil and gas extraction, an ostensibly related sector, is in the bottom quartile. We hesitate to parse the apparent dissonance reflected in these results. The details of industries' production processes are unobservable to us, which means it is difficult to draw strong conclusions on a priori grounds.

A more fruitful, and systematic, approach is to consider the estimates in Table G1 in relation to alternative indicators of production complementarities. To this end, we draw on surveys that ask workers about the importance of *teamwork* in their jobs. If there is a strong complementarity across employees, the production process is likely to involve relatively teamwork-intensive activity. A "team" implies a degree of cooperation and coordination, which is needed to aggregate efforts across many differentiated and interdependent tasks. Of course, any self-reported measure of teamwork is bound to be an imperfect indicator of complementarities. Still, there is perhaps enough of a connection to complementarity to make this examination worthwhile.

O*NET. Our first source of data on teamwork is the Occupation Information Network, or O*NET, which has been used extensively to characterize the task content of occupations. Drawing in part on Koren and Petö (2020), we examine eight variables that pertain to teamwork-oriented tasks.[18] Responses to questions about the frequency of a task (namely, [A.2] and [A.4] below) range from 1 to 5, with 1 being "never" and 5 being "everyday". The remaining questions concern the importance of the task, and responses range from 1 to 5, with 1 being "not important at all" and 5 being "extremely important". The eight questions

[17] Even in this sector, there is likely to be some coordination between the carriers and office workers who process orders and allocate delivery jobs.

[18] Koren and Petö (2020) use many of these questions to identify occupations that heavily involve personal interactions and which were likely to be most disrupted by social distancing during the COVID-19 pandemic.

that we consider are the following:

[A.1] How important are interactions that require you to work with or contribute to a work group or team to perform your current job?

[A.2] How often does your current job require face-to-face discussions with individuals and within teams?

[A.3] How important are interactions that require you to coordinate or lead others in accomplishing work activities (not as a supervisor or team leader)?

[A.4] How much contact with others (by telephone, face-to-face, or otherwise) is required to perform your current job?

[B.1] How important is developing and building teams to the performance of your current job?

[B.2] How important is guiding, directing, and motivating subordinates to the performance of your current job?

[B.3] How important is coordinating the work and activities of others to the performance of your current job?

[B.4] How important is providing consultation and advice to others to the performance of your current job?

We draw on two samples from O*NET. The first consists of surveys run between 2003 and 2007 (O*NET editions 5.0 through 12.0). The 2003 release was the first O*NET product that surveyed employees rather than exclusively occupational analysts. The survey of employees was extended annually to more occupations, and by June 2007, just 2 percent of responses in the data were still derived from analysts. Our 2003-2007 sample relies only on responses supplied by employees.[19] We also do draw on the earlier analyst database, which was assembled during the middle of the 1990s and released at the end of 1998 (Mariani, 1999). We use the final release three years later, which incorporated the 2000 Standard Occupational Classification (SOC) codes. Note that the first two questions, [A.1]-[A.2], were not included in the analyst survey; they initially appeared in the 2003 release.[20]

The scores for questions [A.1]-[B.4] are aggregated so they can be compared with industry estimates in the Veneto data. We first calculate the average score for each question within each SOC code and use an O*NET-published crosswalk to map from the SOC to Census occupation codes. We then compute the mean score within each Census *industry* using

[19] We do not incorporate surveys after 2007 so as to not extend the sample any further beyond 2001 (the end of our Veneto sample) than necessary.

[20] The wording of [A.1]-[B.4] on the previous page was taken from O*NET questionnaires beginning in 2003. It appears that the wording of [A.1]-[A.4] was slightly different, but substantively similar, in the analyst questionnaires. See Appendix B of Boese et al (2001).

Current Population Survey employment estimates to weight the occupation-level scores. (We draw on CPS data corresponding to the periods covered by the employee and analyst surveys: for the former, we use the years 2003-07, and for the latter, we pool data for 1994-98.) Next, we map Census industry codes to the 2002 NAICS using a crosswalk from the Census Bureau. The three-digit NAICS is the lowest level of aggregation routinely used in the crosswalk and forms the basis for our estimates.[21] Finally, the (NAICS) industry-level scores associated with each question are standardized.

Meanwhile, the industry classification in our Veneto data, NACE Rev 1.1, is cross-walked to the 2002 NAICS using a concordance also published by the Census Bureau. Note that since the Veneto data excludes public administration and private household services, these sectors are dropped from the O*NET-derived dataset. For the sake of symmetry with our O*NET-based measures, the ratios of idiosyncratic variances in the Veneto dataset are also standardized.

It should be noted that our use of U.S. and Italian data implicitly assumes that production processes across the two countries are comparable in a certain sense. Specifically, a clean comparison would require that the intensity of team-oriented tasks in a given industry in the U.S. mirrors that in Veneto. To the extent that this assumption is violated, any association between the two measures will be attenuated.

In Table G2, we now report the correlation coefficient between the ratio of idiosyncratic variances in the Veneto data and various aggregates of the teamwork indicators in O*NET. The left panel reports results based on the 2003-07 O*NET sample, whereas the right panel pertains to the analyst database. In the first row (of each panel), the O*NET measure is the mean of standardized scores across all questions [A.1]-[B.4]. In rows two and three, we use, respectively, the average within group [A] and group [B]. In row four, we use [A.1] specifically, which arguably addresses most directly work in a team or group. We present both unweighted and employment-weighted estimates, where the employment weights are derived from the Veneto data. The standard errors are in brackets.

There are a few noteworthy results in Table G2. First, the weighted correlations are uniformly higher; we return to this point shortly. Second, the two O*NET samples—the 2003-07 surveys of employees and the earlier survey of analysts—yield broadly similar point estimates, though the unweighted correlations between the Veneto and 2003-07 O*NET data are estimated more precisely. Third, there is a significant correlation in the 2003-07 sample between the ratio of idiosyncratic variances and group [A] questions: The unweighted

[21] In other words, if we worked at the four-digit level, we would have to make several imputations for four-digit NAICS industries for which we do not have a Census industry counterpart.

(weighted) correlation with the average score among group [A] questions is around 0.21 (0.45) and statistically significant. The correlation with the teamwork variable, an element of group [A], is of the same order and significant. In other words, if the ratio suggests relatively strong complementarities, the sector is more likely to involve teamwork-intensive activities.

To visualize these results, Figure G1 plots the ratio of idiosyncratic variances and the 2003-07 O*NET score among group [A] questions. One's eye is perhaps drawn to the observations in the northeast corner, where both the ratio and O*NET scores are high. Among these, the largest sector is banking. While the latter stands out, the weighted correlation does not hinge on this single sector. Rather, the figure illustrates that what the employment weighting does is "look past" several, very small sectors whose O*NET scores and standardized ratios are largely uncorrelated.[22] In our view, employment weighting is sensible in this context because it focuses on those sectors that are largely responsible for the value of the moment that we use in our estimation.

In contrast to the relationship shown in Figure G1, the correlation of the Veneto ratio with group [B] questions in both O*NET databases is generally smaller and, among the unweighted estimates, statistically insignificant. However, our reading of group [B] is that it pertains more to management responsibilities, whereas the Veneto-based measure is representative of the task content of the broader industry. For instance, [B.2] asks about "guiding, directing, and motivating subordinates," which is more of an executive function that is not necessarily closely tied to the teamwork-intensity of the employees' tasks.[23]

EWCS. Next, we turn to the Third European Working Conditions Survey (EWCS), which was administered in 2000 to workers in the EU-15 countries.[24] One question in the survey addresses teamwork: "Does your job involve doing all or part of your job in a team?" Unlike in O*NET, which attempts to measure the *intensity* of teamwork, the question here is binary, e.g., responses are "yes" or "no". Our indicator of teamwork in an industry is then the share of workers (in that industry) that responds affirmatively to this question. The industries in this case consist of 23 two-digit NACE sectors, most of which are aggregates of several three-digit NAICS industries. The EWCS does not provide more detailed industry classifications.

[22]In the figure, these sectors have O*NET scores clustered around zero but their standardized ratios range from -1 to 1. The sectors include motion picture and sound recording studios (employment share <0.01%); clothing and clothing accessories stores (share of 0.01%); and amusement, gambling, and recreation industries (share of 0.1%).

[23]Nevertheless, we retain group [B] questions in deference to the literature and so as to discipline the exercise; the questions were not selected to maximize their correlation with our Veneto measures.

[24]The EU-15 countries (in 2000) were Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxemburg, the Netherlands, Portugal, Spain, Sweden, and the UK.

For the sake of comparability with our analysis of Veneto and O*NET data, we isolate the subsample of the EWCS that consists of wage-and-salary employees outside of public administration and private household services. This restriction reduces the size of our sample from just over 21,700 respondents to 16,140. Note that we use all countries in the EWCS for this analysis, as sample sizes are very small if we restrict attention to only Italy. (There are fewer than 15 Italian respondents in a third of the industry cells.)

Figure G2 presents our results. A few points stand out relative to Table G1. First, the coarse industry aggregates in the EWCS can mask considerable variation. For instance, teamwork is prominent in the aggregate transport sector according to the EWCS, but water transport and air transport, taken separately, have different degrees of complementarity by our Veneto-based measure. The same point applies to mining and food/beverage/tobacco manufacturing, both of which consist of industries whose rankings differ noticeably in our Veneto data (i.e., consider food v. beverage/tobacco manufacturing in Table G1).

Second, EWCS's teamwork variable and the ratio of idiosyncratic variances are not that closely aligned. Comparing again with Table G1, the two measures do (roughly) agree on a few industries: health care, publishing/printing, and banking are all ranked relatively highly, whereas hunting/fishing and recreation are ranked near the bottom. However, the two measures differ on others: construction and wood product manufacturing are ranked highly according to EWCS but are near the bottom of the list based on Veneto data, whereas land transport is associated with little teamwork in EWCS but ranks highly by our measure. Not surprisingly, the correlation between the two series is relatively slight. The unweighted correlation is 0.107 and insignificant. The weighted correlation is somewhat higher at 0.176 but still insignificant and only half as large as the correlation with the O*NET teamwork question (see the final row of Table G2).

Although we present EWCS results in the interest of completeness, we would argue that there is likely more "signal" in the O*NET data. For starters, the binary nature of the EWCS question elicits a less precise characterization of the respondent's job, and the industries in the EWCS are far coarser. In addition, the estimates based on O*NET are reasonably robust across two independent survey instruments, namely, the employee and occupational analyst surveys. Finally, the O*NET has been heavily vetted; it is a principal source of data on the task content of jobs and has been used to study employee interaction and on-site work in occupations (see Firpo, Fortin, and Lemieux (2013)).

Gender gaps. Goldin (2014) and Goldin and Katz (2011, 2016) have argued that an important contributor to the gender earnings gap is a difference in the way men and women value working time flexibility. It is often thought that women will seek out positions within an industry in which working time can be more tailored to their personal circumstances (because of, i.e., home-care responsibilities). As a result, if an industry is generally characterized by a high degree of complementarities, there will be many more opportunities for men, disproportionately, to sort into employers where $\rho << 0$. Accordingly, we might expect women's relative earnings to be lower in such sectors.

We can examine this claim with our Veneto data. The *earnings* gap for any industry is defined as the average difference between women's log daily earnings and men's. For consistency with our earlier analyses, we continue to work with 3-digit NAICS industries. In addition, we also explore what we refer to as the gender *participation* gap, which is measured as the female share of the industry's workforce less the male share. The participation gap will indicate whether women sort away from high-complementarities sectors altogether.

Table G3 summarizes the distributions of both gender gaps. On average, female participation and earnings are substantially lower than men's. However, women outnumber men in roughly one quarter of the sectors, including health care; personal care retail; and apparel and leather goods manufacturing. By contrast, the earnings gap is almost uniformly negative across 3-digit NAICS. Even at the 90th percentile, the gap is nearly 15 log points. Finally, the two gaps are negatively correlated: female participation is higher where relative female earnings are lower. For instance, in several health and education sectors, women make up a substantial majority but have relatively low earnings.

Next, we examine the association between the gender gaps and the ratio of idiosyncratic variances, which is our diagnostic for complementarities. Table G4 reports a few results. First, there is not much of a relationship between the ratio of idiosyncratic variances and the participation gap. Indeed, the correlation is almost zero in the unweighted data. At the same time, there is a more robust negative relationship between the ratio of idiosyncratic variances and the earnings gap: the correlation is -0.28, and statistically significant, in the unweighted data. If employment weights are used, the correlation declines to -0.17 and is no longer significant. Thus, there is some (albeit modest) evidence that, in industries where complementarities appear to be stronger, women's relative earnings tend to be lower.

There are a few interesting details lying behind the headline result concerning the earnings gap. The stronger correlation in the unweighted data reflects outcomes in a few small sectors—air transportation and insurance carriers, for instance—where earnings gaps are "more negative" and the ratios of idiosyncratic variances are relatively high. Among the larger sectors, a few stand out as contributing to the weighted correlation. Banking features high ratios and negative earnings differentials, whereas food and wood products manufacturing exhibit just the opposite.

How might we interpret these results through the lens of our model? Of course, the model has no genders or industries. Nevertheless, if we are willing to (informally) consider certain extensions of the model, it is possible to apply some of the lessons of the theory to interpret these results. One important extension is that there must be a distribution of ρs within and across sectors. In addition, we maintain the assumption that women sort into more flexible working time arrangements, e.g., higher-ρ environments. Cubas et al (2020) argue that this result can arise if women value home-produced goods more than men.

Now consider, first, the earnings gap *within* a sector. Under strong complementarities, earnings are especially sensitive to the marginal cost of time, as shown in Figure G3. If preference shocks are symmetric, however, stronger complementarities will imply larger swings in earnings (corresponding to changes in ξ) but not necessarily a different mean. Preference shocks do not have to be symmetric, though. To illustrate, suppose $\xi = 1$ is typical, but there are occasional periods when the worker is "needed" more at home or elsewhere outside of the workplace, e.g., $\xi > 1$. In this case, average earnings could differ across genders within a sector if men sort into the segments of an industry where complementarities are relatively strong, as along the blue schedule in Figure G3.

Next, suppose that the *average* degree of complementarities differs *across* sectors. A lower average ρ (hence, stronger complementarities) can translate into greater inequality (hence, larger differentials) in certain cases. To illustrate, imagine ρ is in practice bounded above by $\rho_m < 0$ (which can be arbitrarily near zero) and assume $-\rho$ follows a Pareto distribution on $[-\rho_m, \infty)$ with shape parameter $\varkappa$. An industry where $\varkappa$ is relatively low has a heavier right tail of $-\rho$s, which also contributes to a lower average ρ. In the context of Figure G3, there are more, very steeply sloped blue schedules. If men are more likely to sort into these industry segments, we could observe larger earnings differentials in sectors with stronger average complementarities.

Although this narrative crucially relies on a certain pattern of sorting within an industry, it does not necessarily predict sharp differences in gender composition across industries. The reason is straightforward: the heterogeneity of ρs within a sector means that if women prefer certain tasks in an industry with a low average ρ (such as health), they can still take up employment in segments where complementarities are relatively weak. This observation could help account for why the ratio of idiosyncratic variances is more clearly related to the earnings gap than to the participation gap.

H Labor market institutions in Italy

This section summarizes salient features of Italian labor market law. First, we review the costs of dismissal and discuss our baseline value of $\underline{c}$, the model's separation cost. In the second subsection, we summarize the benefits available under Italy's unemployment insurance system. The latter information was used to assess our estimate of the outside option, μ, in Section 4 of the main text (see also Appendix C.5).

H.1 Severance costs

H.1.1 A description of dismissal laws[25]

There are three components to Italy's severance payment structure.

TFR. Every separated worker receives a severance, known as *Trattamento di Fine Rapporto* (TFR), equal to 7.4 percent of his total earnings during his tenure with the firm. We do *not* count TFR as a cost of dismissal, because it is enforced regardless of the reason for separation. Rather, TFR is better thought of as a form of deferred compensation. It is, by law, fully funded: in each year, the firm must deposit in a special fund an amount equal to each worker's accrual.[26]

Individual dismissal. There are many aspects to the cost of an *individual* dismissal.

After the firm notifies the worker in writing of its intent to dismiss, a notice period is triggered.[27] At the end of the notice period, the dismissal can legally take effect. In practice, though, firms will pay in leiu of notice, which means they award the worker the earnings he would have accrued during the notice period in exchange for the worker leaving the firm immediately. The notice period is increasing in tenure, but for blue-collar workers, it is fairly short in general; even workers with 20 years of tenure receive just 12 days of notice (OECD, 2004). For white-collar workers, however, the notice period can stretch to several months.

If the worker wishes to contest the dismissal, he must notify the firm in writing within 60 days of the start of notice, and file a case with the local Labor Tribunal within 180 days.

[25]Our description of Italian dismissal regulations draws from Schivardi and Torrini (2008), the European Commission (2006), Baker & McKenzie (2009), Ius Laboris (2009), Deloitte (2012), and Emmenegger (2014). Note that our description pertains to the law as it was written in our sample period (1982-2001).

[26]Many scholars have weighed in against the treatment of TFR as a dismissal cost (Bertola et al 2000; Del Conte et al 2004; Schivardi and Torrini 2008). Recognizing this point, the OECD, which originally included TFR as a dismissal cost in its index of employment protection, removed it in 2004 (OECD 2004).

[27]There could be a slight delay between the delivery of written notification and the start of the notice period. The worker (often in conjunction with his union) can take up to 5 days to respond to the firm's letter in order to, for example, defend the worker's actions if misconduct is alleged.

The failure to meet one of these deadlines means the worker forfeits his right to contest the dismissal.

The worker can challenge his termination on grounds that the firm failed to show either "giusta causa" (just cause) or "giustificato motivo" (valid reason). A just-cause dismissal implies that the worker is guilty of serious malfeasance, such as theft of firm property. Grounds for termination under "giustificato motivo" fall into two categories: (i) "significant non-compliance with contractual obligations" (e.g., failure to follow management's directives, absenteeism, etc.); and (ii) "economic reasons" related to reduced production.

If the court finds against the employer, the costs of the dismissal are substantial. Firms that employ more than 15 workers face the most significant costs.[28]

First, the worker is awarded the earnings he would have received since dismissal. While data in this area are sparse, Macis (2001) reports that, among cases decided in 1995 in Northern Italy, the average time between dismissal and the Labor Tribunal's decision was 15 months.[29] In addition, the firm owes Social Security contributions, which are treated as late payments. Ichino (1996) estimates that Social Security contributions plus late penalties amount to 8 months of earnings.

Second, the firm must offer to reinstate the worker. Alternatively, the employee can waive reinstatement in exchange for a severance of 15 months of earnings. Garibaldi and Violante (2005) report that workers almost always elect the latter option.

Finally, a losing firm is required to reimburse the worker (or, more exactly, his union) for legal fees, which typically amount to 5 months of earnings (Ichino, 1996).[30] Summing up, we find that, if the firm loses at trial, it faces a total cost of 43 months of earnings, which is very close to Garibaldi and Violante's (2005) estimate of 41 months.

Since 1990, individual dismissals at firms with 15 or fewer employees have also been subject to judicial review. If the court finds against such an employer, it can order the firm to pay the worker a severance of between 2.5 and 6 months of earnings (depending on the worker's tenure). Reinstatement is not required at firms with 15 or fewer workers, and the employee is not entitled to foregone earnings.

[28] Strictly speaking, a firm is subject to these relatively steep dismissal costs if it has either 15 or more employees at any one of its worksites or 60 total employees. By far, most firms qualify under the first criteria, so we will take 15 to be the relevant firm-size threshold. In addition, since we do not have access to establishment-level data, we have little choice but to use 15 as the threshold when we later aggregate dismissal costs across size classes.

[29] This is the time to the first hearing of the case. If the worker loses this trial, he could appeal.

[30] Ichino estimated fees of 10-20 million Italian lira. Taking the midpoint of this range and dividing average earnings in 1996 (at the time of Ichino's writing) by 15 million yields the estimate cited in the text. For this calculation, note that OECD's data on average earnings is converted from euros to lira using the exchange rate at the time of the euro's introduction.

Collective dismissal. The notion of a collective dismissal was first codified in Italian law in 1991. The statute was adopted to comply with a European Community directive that required consultation in the event of collective redundancies. We focus our remarks here on the period since 1991. The State's treatment of collective dismissals was less transparent in earlier years; we defer a discussion of the period 1981-1990 until the next section.

A collective dismissal can be invoked only by relatively larger firms that carry out multiple layoffs. Specifically, an eligible firm must employ more than 15 workers and separate from at least 5 of its workers over a span of 120 days.[31]

The protocol of a collective dismissal is as follows. The firm must inform the provincial labor office and the workers' trade union of its intent to dismiss, and note the reason(s) for dismissal. The union may ask to open negotiations with management in order to find an alternative resolution, i.e., it may ask management to look into whether the workers could be assigned to different jobs within the firm. Talks with the union are limited to 45 days, or 1.5 months. If an agreement between firm and union is not reached, the labor office opens a mediation. This process may last up to 1 month. If no other resolution is found, the dismissal takes effect. In total, then, the dismissal can be postponed for up to 2.5 months, during which time the employees in question remain on the payroll.[32]

In the event of dismissal, the firm is required to make contributions to the National Social Security Institute (INPS). The amount of payment depends on the outcome of the negotiations with the union. If the union and management agree to terms of the separation, the firm's contribution per redundant worker is set equal to 2.4 months of the worker's (previous) earnings. If talks between the firm and union failed, the firm is responsible for 7.2 months of earnings.[33]

A few other aspects of collective dismissals should be noted, even if their quantitative importance is hard to pin down. First, labor law requires that firms try to shield workers of high seniority, and workers with many dependents, from collective dismissals. Insofar as this aspect of the law distorts the firm's separation decision, it can be regarded as another kind

[31] To determine eligibility, firm size is measured as average employment during the 6 months preceding the dismissals.

[32] If the number of redundancies is less than 10, the time limit on negotiations with the union and labor office is 38 rather than 75 days. For simplicity, we will treat the limit as 2.5 months, though this will mean that our estimate of the unit cost in a collective dismissal is probably slightly too high.

[33] These payments are derived as follows. If the union and firm reach an agreement, the firm's contribution equals 3 months of the worker's "mobility" allowance, or unemployment benefit (see Appendix H.2 for more on this). Since the allowance is 80% of earnings, the firm's payment is $0.8 \times 3 = 2.4$ months of earnings. If the provincial labor office has to mediate, the contribution is 9 months of the allowance, or $0.8 \times 9 = 7.2$ months of earnings.

of dismissal cost. However, it is much less easily measured. Second, if the firm does not give notice of dismissal or refuses negotiation, the redundant workers can challenge the dismissal in court. If the judge rules in the workers' favor, the firm must offer to reinstate them. To the best of our knowledge, though, legal action in the case of collective dismissals is rare.

H.1.2 Averaging dismissal costs

In our theory, we do not distinguish dismissal costs based on either firm size or the volume of dismissals. To take the model to data, then, we aggregate multiple dismissal costs into a single parameter. In this section, we first develop estimates of the unit cost of individual and collective dismissals for small and large firms and discuss a suitable way to average the two costs *within* each size class. Then, we average costs *across* size classes. Throughout, a "large" ("small") firm is taken to be an employer with more than (no more than) 15 employees.

Large firms—individual dismissals. Although an individual dismissal is very costly *if* a firm loses at trial, it appears that most dismissals are not adjudicated. Based on data from a large financial services firm, for instance, Ichino et al (2003) find that just over 20 percent of individual dismissals are contested.

A firm that seeks to dismiss a worker is more likely to instead negotiate an out-of-court settlement to avoid a lengthy and costly trial. These settlements are presumably shaped by the statutory penalties prescribed by law. Still, the final settlement will be sensitive to a number of factors—the firm's and worker's perception of the probability of victory at trial as well as each party's effective discount rate—that are unobservable to the analyst.

Fortunately, we can draw on estimates of the cost of individual dismissals that are periodically published by the legal services division of Deloitte Ltd.[34] Deloitte's estimates reflect "the average cost which an employer will incur ...[to] reach a final settlement agreement without court interference". We consult the report based on labor laws as of the end of 2011, which excludes the reforms of Italy's dismissal laws passed in June 2012. (We are not aware of substantive changes to the cost of individual dismissal from regular, full-time jobs between the end of our sample and 2011.) In order to standardize dismissal costs across countries, Deloitte reports an estimate for a certain white-collar job with 7 years of tenure at her current employer. In the case of Italy, the cost of dismissal in this scenario was 21 months of compensation.

[34] Deloitte is the one of the largest privately held companies headquartered in the U.S. The estimates are compiled by Laga, Deloitte's legal practice in Belgium. The first report was published in 2009, but it is unavailable to us. It has since published reports in 2012, 2015, and 2018.

We adjust the Deloitte estimate for two reasons. First, it includes the TFR, which amounts to 6 months of compensation. Second, it includes a cost of notice of about 3 months of compensation that is specific to white-collar workers and unrepresentative of the Italian workforce more generally. In our Veneto dataset, 65 percent of workers are classified as blue-collar, and for a blue-collar worker with 7 years of tenure, the cost of notice is roughly 9.5 days, or little less than one-third of a month.[35] Therefore, we compute the payment in leiu of notice as $0.35 \times 3 + 0.65 \times (9.5/30) \cong 1.25$ months of earnings. Accordingly, our estimate of the cost of individual dismissal is $21 - 6 - 3 + 1.25 = 13.25$ months of earnings.

Large firms—collective dismissals. The collective dismissal law of 1991 prescribes that a firm owes a contribution to INPS of at least 2.4 months of earnings per laid-off worker. Any additional costs depend on the outcome of negotiations among the firm, union, and provincial labor office. Although we are unaware of data on these outcomes, we can make a reasoned conjecture as to the firm's typical liability. Recall that negotiations can extend for 2.5 months, during which time the worker remains employed. Thus, there is no reason for a worker to consider a settlement offer worth less than 2.5 months of earnings. The firm, in turn, should not agree to a settlement valued at more than $2.5 + (7.2 - 2.4) = 7.3$ months of earnings, since the latter is the additional cost it would bear if talks carry on for 2.5 months. We suppose that the firm and union bargain to the midpoint between their reservation levels, which implies a settlement worth $0.5 \times (2.5 + 7.3) = 4.9$ months of earnings. Thus, beginning in 1991, we take the unit cost in a collective dismissal to be $2.4 + 4.9 = 7.3$ months of earnings.[36]

Prior to 1991, all dismissals were, strictly speaking, subject to individual dismissal laws, but in practice, many large(r) firms had recourse to State support. A program known as the *Cassa Integrazione Guadagni Straordinaria* (CIGS) was expanded dramatically during the 1980s to accommodate firms' "restructuring and reorganization" (Tronti, 1991). CIGS was originally meant to augment workers' incomes during a relatively short layoff, at the end of which the worker would return to his (initial) employer. In practice, however, a worker's CIGS benefits could be renewed virtually every year, which meant firms could carry out, in effect, collective dismissals without incurring the statutory dismissal costs (D'Harmant Francois and Brunetta, 1987; Reyneri, 1989). Moreover, dismissals under CIGS were not

[35] This is calculated as follows. The OECD reports notice periods of 9 and 12 days for blue-collar workers with, respectively, 4 and 20 years of tenure. We assume that notice increases linearly between these points, which means that a worker with 7 years of tenure is given notice of $9 + (7 - 4)\,\frac{3}{20-4} \cong 9.56$ days.

[36] Recall that if the number of dismissals is less than 10, the time limit on negotiations is 38 days, or about 1.25 months. In this case, the unit cost in a collective dismissal is $2.4 + 0.5 \times (1.25 + 6.05) = 6.05$ months of earnings.

subject to experience-rated taxes (Schioppa, 1988; Del Boca and Rota, 1989). In short, the cost to the firm of placing a worker on CIGS was practically zero.

However, CIGS' coverage was not universal. Eligibility was generally restricted to firms in manufacturing and construction. In Veneto, these sectors accounted for 58% of employment among large firms between 1981-1990. The latter share is sizable, but it remains the case that more layoffs would have been exposed to individual dismissal costs prior to 1991.

In summary, collective dismissals carried essential no cost for some firms, whereas the remainder of employers faced expected dismissal costs of over 13 months of earnings. Note that the *post*-1991 average cost of around 7 months of earnings lies roughly between these two ends of the pre-1991 cost spectrum. What's more, the share of CIGS-eligible employment was just a bit over one-half. Therefore, for the sake of simplicity, we just extend the post-1991 cost of collective dismissal to the earlier period of our sample (1981-1990).

Averaging costs within large firms. Our objective is to average individual and collective dismissal costs into a single per capita cost of separation that can be inserted into the structural model (e.g., $\underline{c}$).[37] It would arguably be ill-advised to weight the two costs by their shares of total dismissals, though. The reason is that Italian firms have a strong motive to avoid individual dismissals. Therefore, even if we never saw an individual dismissal, its cost could be very salient to the firm. In other words, the marginal cost of adjusting that faces a firm *if* it were to make its desired adjustments could often be the individual dismissal cost, even if an individual dismissal is not implemented in the end.

Accordingly, our strategy is to weight individual (collective) dismissals according to how often a shrinking firm would carry out fewer than five (five or more) dismissals *if* dismissal restrictions were removed. This weighting scheme correctly captures the *exposure* of dismissals to each kind of cost (if the desired dismissal were to be implemented). To compute this counterfactual, we use data from the relatively unregulated U.S. labor market. Within a firm size class, we estimate the share of layoffs that involve fewer than 5 workers, which means the layoffs would be treated as individual dismissals in Italy. Our assumption is that, if adjustment costs were eliminated, Italian firms would implement the same individual share of dismissals observed among U.S. firms in the same firm size class. This assumption can be debated, of course. Even so, the U.S. strikes us as the best approximation (for which data is available) of what a dismissal policy would look like in the absence of frictions faced by

[37] Alternatively, Del Boca and Rota (1998) and Rota (2004) advocate for a fixed, rather than per-capita, cost. A fixed cost induces infrequent, lumpy adjustment, consistent with their claim that firms "bundle" separations into a cheaper collective dismissal. We pursue an alternative strategy that emphasizes the statutory form of a dismissal cost—it *is* levied *per* dismissal—and identifies a way to aggregate across multiple costs.

Italian firms.

Our U.S. data are quarterly unemployment insurance (UI) claims records collected for the Short-Time Compensation (STC) Study (Kerachsky et al, 1997). This study includes a sample of 3,415 firms from five states—California, Florida, Kansas, New York, and Washington—and records the number of UI claims charged to each firm in the years 1991-1993. The study designed the sample by first randomly selecting firms among those participating in each state's STC plan, which extends benefits to workers on reduced hours, and then matching a participant to a nonparticipating firm in the same state, industry, and size class.[38] Since our focus is not on STC, we will examine layoff outcomes among *only nonparticipating* firms, although the results do not depend on this restriction.

We proceed in two steps. First, we use the U.S. data to estimate the incidence of individual dismissal within size class. We first identify *layoff events* as quarters of positive UI (initial) claims. We then divide these firm-quarter observations across 9 firm size bins given in Schivardi and Torrini (2004, 2008), the smallest of which (given our focus now on large firms) is 16-20 employees. For each size class, we calculate the individual share of dismissals defined as the share of layoff events that involve fewer than 5 UI claims.[39]

The second step is to aggregate individual dismissal shares across size bins. Specifically, we compute the employment-weighted average of the size-specific individual dismissal shares and find that 54 percent of laid-off workers would be involved in an individual dismissal. Note that the employment weights are from Schivardi and Torrini (2004, 2008) and thus reflect the Italian (rather than the U.S.) firm-size distribution. Thus, our estimate conveys what would happen to Italian workers *now* if dismissal frictions were immediately suspended; we do not consider what dismissal activity would be given a new steady-state size distribution.

We repeat this analysis at a semi-annual frequency, in which case layoff events are defined as 6-month periods in which there are positive UI claims. By this measure, 35 percent of laid-off workers would be involved in an individual dismissal. We combine our quarterly and semiannual estimates to measure the incidence of individual dismissals over a 4-month window (120 days), which is the period length over which collective dismissals are defined in the 1991 law. We then simply linearly interpolate between the two estimates, which yields

[38] Enrollment in STC programs is voluntary. The study's designers defined participation as being enrolled in a STC program in 1992. However, enrolled firms did not necessarily *use* STC subsidies. In our analysis, we define "nonparticipation" as not using STC subsidies *at any point* during 1991-1993. The results are quite similar if we define "nonparticipation" as not being enrolled in a STC plan in 1992.

[39] One shortcoming of using UI claims as a proxy for layoffs is that, in the wake of large-scale layoffs, all of the workers do not necessarily file for UI within the quarter. This observation suggests that quarterly claims data are likely to understate the frequency of collective dismissal.

an individual dismissal share over the 4-month window of 47.7 percent.[40]

Piecing these estimates together, we calculate the unit cost of dismissal for large firms equal to $0.477 \times 13.25 + (1 - .477) \times 7.3 = 10.14$ months of earnings. The first term can be understood as the product of the expected individual dismissal cost and the probability that a layoff would take the form of an individual dismissal. The second term is simply the analogue for collective dismissals.

Small firms. Firms employing no more than 15 workers were not subject to dismissal penalties until 1991. Since 1991, we presume that the prospect of dismissal (typically) leads to a negotiated termination settlement. Deloitte's estimates do not extend to small firms, but they are still potentially informative. Recall that if a dismissed worker from a large firm prevails at trial, he is expected to receive 30 months of earnings (excluding fees for the union, etc.): 15 months of back pay (covering the length of the trial) plus a severance of 15 months of earnings in exchange for waiving his right to reinstatement. Deloitte's estimates imply, however, that a settlement offer of 13.25 months of earnings can head off a trial. In other words, workers from large firms accept settlements equal to $13.25/30 \cong 44.2$ percent of the potential pay-out. We apply this rate to compute the settlement in the case of a small firm. If a dismissed worker from a small firm prevails at trial, he would be awarded between 2.5 and 6 months of earnings. Taking the midpoint of this range and multiplying by 0.442 yields a settlement worth 1.9 months of earnings.

Finally, to obtain an estimate appropriate for the full sample, we "pro-rate" the post-1991 cost. The latter was in effect for 55 percent of our sample period. Accordingly, we set the cost of dismissal at small firms to be $0 \times 0.45 + 1.9 \times 0.55 \cong 1$ month of earnings.

Averaging costs across size classes. As a final step, we calculate the employment-weighted average of costs across large and small firms. Schivardi and Torrini (2004, 2008) estimate that, in the period 1986-1998, 68 percent of the Italian workforce was, on average, employed in firms with more than 15 workers. Thus, we conclude that the unit cost of dismissal in Italy was $.68 \times 10.14 + (1 - .68) \times 1.03 = 7.225$ months of earnings.

H.2 Unemployment benefits

This section reviews our calculation of the replacement rate implied by unemployment insurance in Italy.[41] Our results are used in Section 4.2 to assess the estimated value of μ.

[40] The *realized* share of individual dismissals in Italy is likely to be less than half of this, as we note in Appendix H.2.

[41] Our discussion draws heavily from Rosolia and Sestito (2012), Salvatore (2010), and Bertola and Garibaldi (2003).

Our calculation of the replacement rate reflects an average across two unemployment benefit programs. A third program, the *Cassa Integrazione Guadagni Ordinaria*, (CIGO), is not considered here because it pays benefits exclusively to individuals on temporary layoff (for up to 13 weeeks). As a result, CIGO benefits do not map neatly to μ. Whereas workers under CIGO have not formally severed ties with their employers, μ in the model refers to the flow value of nonwork time available to a worker in the event of permanent separation.

Ordinary benefit. The Ordinary Unemployment Benefit (OUB) is paid to any unemployed worker who meets two requirements: (1) he or she has recorded at least 52 weeks of work in the last two years; and (2) he or she is now available for work. The ordinary benefit was set as a fixed nominal payment when it was introduced in 1955. As it was updated only infrequently, inflation tended to erode its value. Indeed, the implied replacement rate was virtually zero by the start of our sample period (1981). In 1987, the OUB was recast as an explicit replacement rate and was raised in stages. By 1995, it equaled 30 percent of average earnings, a level at which it remained until it was increased to 40 percent in 2001 (the final year of our sample). Throughout virtually all of our sample period, the ordinary benefit was available for no more than 6 months.[42]

Mobility benefit. The Mobility Indemnity (MI), or mobility benefit, is available to unemployed workers who have been separated as part of a collective dismissal. In the first year of unemployment, the mobility allowance replaces 80 percent of a worker's former annual earnings up to a statutorily set maximum benefit. In practice, many workers draw the maximum, which implies a replacement rate more like 67 percent in the first year of an unemployment spell. Workers under age 40 are only eligible for one year of mobility benefits. Workers over age 40 can draw mobility benefits for a second year at a (maximum) replacement rate of 60 percent, and workers over age 50 are eligible for a third year at a (maximum) replacement rate of 40 percent.[43]

The MI was introduced in 1991 as part of a reorganization of Italy's unemployment insurance system. We noted in Appendix H.1 that, prior to 1991, benefits under *Cassa Integrazione Guadagni Straordinaria* (CIGS) were available almost indefinitely (Brunello and Miniaci, 1997; Brugiavini, 2009). Since MI's introduction in 1991, the CIGS program has been time-limited and re-oriented toward workers who will likely be recalled to their employers after a long layoff. Unemployed workers (again, from large firms) who have cut ties with their employers are now served by MI.

[42] The duration was extended to 9 months for workers over age 50 in 2001.

[43] This mobility benefit schedule applies to regions in northern Italy, such as Veneto. The mobility benefits are more generous in the south, where unemployed workers can collect 80 percent of average earnings for the first *two* years, and workers over age 50 can collect 60 (40) percent of earnings in the third (fourth) year.

We calculate the annual replacement rate of unemployment benefits as a weighted average of the ordinary and mobility benefit rates available in the first year of an unemployment spell. The mean OUB-implied annual replacement rate over our sample was 19 percent. Since the OUB lasts only half of a year, though, we pro-rate this to be 9.5 percent. The effective replacement rate of the Mobility Indemnity is 67 percent. We treat the CIGS as the predecessor of MI prior, and assume again a replacement rate of 67 percent. Since a worker may not draw both OUB and MI and since the latter is more generous, we assume that MI-eligible workers receive only it.

Thus, we take an average of 0.19/2 and 0.67, where the latter is to be weighted by the share of employment that would receive CIGS/MI in the event of layoff. From the perspective of our theory, this share is the correct concept for the weight. Recall that μ is the (expected) payoff available to a typical employed worker if he exits employment. Accordingly, the weight applied to CIGS/MI reflects the (ex ante) probability of its receipt if laid off rather than the (ex post) share of the unemployed who receive it.

The CIGS/MI share can be calculated as the product of two terms. The first is the probability that a (randomly selected) worker is employed by a large firm and, thus, potentially subject to collective dismissal. The second is the probability that such a worker, if laid off, is involved in a collective dismissal. The first part is simply the share of employment in large firms, which is 68 percent. The second part is the collective dismissal share of layoffs. Note that for our purposes here, we want to measure the *observed* collective dismissal share rather than the counterfactual used in Appendix H.1. The reason is that the replacement rate that a worker anticipates reflects the actual propensity for collective dismissals.

Unfortunately, there are no official estimates of dismissals by type over sample period (to our knowledge), so we try to pull together a reasonable guess from other sources. Based on Boeri and Jimeno's (2005) estimates, the total layoff rate in 1993-95 appears to be about 1 percent (per year).[44] To remove individual dismissals from this total, we draw on Macis' (2001) data, which reports the number of legal challenges against individual dismissals in this period. We then inflate the latter by a factor of 5 in light of Ichino et al's (2003) estimate that just one-fifth of such dismissals are ever disputed in court. These calculations suggest that around 20 percent of layoffs were individual dismissals.[45] It follows that the CIGS/MI

[44] This estimate refers to indefinite/permanent (rather than temporary) workers, consistent with the model's treatment of employment contracts.

[45] Macis reports the number of disputes only through 1993 but the number of judgments through 1995. We assume that the judgment share of disputes in 1994-95 equals its 1989-93 average, so we can estimate disputes over 1993-95, the same period covered by Boeri and Jimeno (2005). It is important to align the dates between these two estimates because the Italian labor market was relatively weak in 1993-95, and dismissals were somewhat elevated.

share is $0.68 \times (1 - 0.20) = 0.544$, and, thus, the replacement rate is

$$(1 - 0.544) \times 0.095 + 0.544 \times 0.67 = 0.408.$$

We conclude that the average replacement rate is likely to be around 40 percent.[46]

References

Andrews, Isaiah, Matthew Gentzkow, and Jesse Shapiro. 2017. Measuring the Sensitivity of Parameter Estimates to Estimation Moments. *Quarterly Journal of Economics* 132 (4): 1553–1592.

Baker & McKenzie. 2009. *Worldwide Guide to Termination, Employment Discrimination, and Workplace Harassment Laws*. Chicago: Baker & McKenzie.

Bertola, Giuseppe, Tito Boeri, and Sandrine Cazes. 2000. Employment Protection in Industrialized Countries: The Case for New Indicators. *International Labour Review* 139 (1): 57-72.

Boeri, Tito, and Juan F. Jimeno. 2005. The Effects of Employment Protection: Learning from Variable Enforcement. *European Economic Review* 49: 2057–2077.

Boese, Ron, Phil Lewis, Pam Frugoli, and Karen Litwin. 2001. Summary of O*NET 4.0 Content Model and Database. National Center for O*Net Development, Raleigh, NC.

Brugiavini, Agar. 2009. Welfare Reform and Labour Supply in Italy. Working Paper 2009:29, Institute for Labour Market Policy Evaluation, Uppsala, Sweden.

Brunello, Giorgio, and Raffaele Miniaci. 1997. Benefit Transfers in Italy: An Empirical Study of Mobility Lists in the Milan Area. *Oxford Bulletin of Economics and Statistics* 59 (3): 329-347.

Cubas, German, Chinhui Juhn, and Pedro Silos. 2022. Coordinated Work Schedules and the Gender Wage Gap. Unpublished manuscript, Temple University.

D'Harmant Francois, Antonio, and Renato Brunetta. 1987. The Cassa Integrazione Guadagni. *Labour* 1 (1): 15-56.

Del Boca, Alessandra, and Paula Rota. 1989. Wage Compensation and Employment: Some Empirical Evidence from Microdata. *Giornale degli Economisti e Annali di Economia* 48 (3/4): 147-157.

[46] Martin's (1996) tabulations suggest a much lower replacement rate, for two reasons. First, Martin neglects CIGS pre 1991. Second, he weights the MI replacement rate based on the share of *unemployed* workers who draw the benefit. We instead weight by the share of *employed* workers "exposed" to MI/CIGS (in the event they become unemployed).

—. 1998. How Much Does Hiring and Firing Cost? Survey Evidence From a Sample of Italian Firms. *Labour* 12 (3): 427-449.

Del Conte, Maurizio, Carlo Devillanova, Stefano Liebman, and Silvia Morella. 2004. Measurability of Employment Protection Systems. Legal Studies Research Paper 06-03, Institute of Comparative Law, Bocconi University.

Deloitte. 2012. *A comparative look at dismissal costs and issues across Europe*. Zaventem: Deloitte Global Services Limited.

Dixit, Avinash. 1997. Investment and employment dynamics in the short run and the long run. *Oxford Economic Papers* 49 (1): 1-20.

Elsby, Michael, Bart Hobijn, and Ayşegül Şahin. 2013. Unemployment Dynamics in the OECD. *Review of Economics and Statistics* 95: 530-548.

Emmenegger, Patrick. 2014. *The Power to Dismiss: Trade Unions and the Regulation of Job Security in Western Europe*. Oxford: Oxford University Press.

European Commission. 2006. *Termination of Employment Relationships: Legal Situation in the Member States of the European Union*. Brussels: The Publications Office of the European Union.

Firpo, Sergio, Nicole Fortin, and Thomas Lemieux. 2013. Occupational Tasks and Changes in the Wage Structure. Unpublished manuscript, Vancouver School of Economics.

Garibaldi, Pietro, and Giovanni Violante. 2005. The Employment Effects of Severance Payments With Wage Rigidities. *Economic Journal,* 115: 799-832.

Goldin, Claudia. 2014. A Grand Gender Convergence: Its Last Chapter. *American Economic Review* 104 (4): 1091-1119.

Goldin, Claudia, and Lawrence Katz. 2011. The Cost of Workplace Flexibility for High-Powered Professionals. *The ANNALS of the American Academy of Political and Social Science* 638: 45-67.

—. 2016. A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation. *Journal of Labor Economics* 34 (3): 705-746.

Ichino, Andrea, Michele Polo, and Enrico Rettore. 2003. Are Judges Biased by Labor Market Conditions? *European Economic Review* 47: 913-944.

Ichino, Pietro. 1996. *Il Lavoro E Il Mercato: Idee per un diritto del lavoro maggiorenne*. Milan: Mondadori.

Ius Laboris. 2009. *Collective Redundancies Guide*. Brussels: Ius Laboris.

Kerachsky, Stuart, Rebecca London, Deana McCanne, Karen Needels, Walter Nicholson, and Stephen Walsh. 1997. *Evaluation of Short-Time Compensation Programs.* Final Report, Mathematica Policy Research, Inc., Princeton, NJ.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. Leave Out Estimation of Variance Components. *Econometrica* 88 (5): 1859-1898.

Koren, Miklós, and Rita Pető. 2020. Business Disruptions from Social Distancing. *Covid Economics* 2: 13-31.

Mariani, Matthew. 1999. Replace with a database: O*NET replaces the Dictionary of Occupational Titles. *Occupational Outlook Quarterly* 43 (1): 3-9.

Martin, John. 1996. Measures of Replacement Rates For the Purpose of International Comparisons: A Note. *OECD Economic Studies* 26: 99-115.

Macis, Mario. 2001. Il Mercato del Lavoro e la Giurisprudenza in Materia di Licenziamenti (Italia, 1989-1998). *Diritto delle Relazioni Industriali*, 2: 269-284.

Merz, Monika. 1995. Search in the Labor Market and the Real Business Cycle. *Journal of Monetary Economics* 36 (2): 269-300.

Organisation for Economic Co-operation and Development. 2004. *OECD Employment Outlook*. Paris: OECD.

Reyneri, Emilio. 1989. The Italian Labor Market: Between State Control and Social Regulation. In *State, Market, and Social Regulation: New Perspectives on Italy*, ed. Peter Lange and Mariano Regini. Cambridge: Cambridge University Press.

Rota, Paula. 2004. Estimating Labor Demand With Fixed Costs. *International Economic Review* 45 (1): 25-48.

Rosolia, Alfonso, and Paolo Sestito. 2012. The Effects of Unemployment Benefits in Italy: Evidence From an Institutional Change. Working Paper No. 860, Banca D'Italia, Rome, Italy.

Salvatore, Livdia. 2010. Italy: Wage Guarantee Fund, Solidarity Agreements, Bilateral Bodies and Wage Guarantee Fund in Derogation. Report reference number EF10636, Eurofound, Dublin, Ireland.

Schivardi, Fabiano, and Roberto Torrini. 2004. Threshold Effects and Firm Size: The Case of Firing Costs. Discussion Paper No 633, Centre for Economic Performance, London, UK.

—. 2008. Identifying the Effects of Firing Restrictions Through Size-contingent Differences in Regulation. *Labour Economics* 15: 482-511.

Schioppa, Fiorella Padoa. 1988. Underemployment Benefit Effects on Employment and Income Distribution: What We Should Learn from the System of the Cassa Integrazione Guadagni. *Labour* 2 (2): 101-124.

Sigurdsson, Jósef. 2021. Labor Supply Responses and Adjustment Frictions: A Tax-Free Year in Iceland. Unpublished manuscript, Department of Economics, Stockholm University.

Stokey, Nancy, Robert Lucas, and Edward Prescott. 1989. *Recursive Methods in Economic Dynamics.* Cambridge: Harvard University Press.

Trigari, Antonella. 2006. The Role of Search Frictions and Bargaining for Inflation Dynamics. Working Paper No. 304, Innocenzo Gasparini Institute for Economic Research, Bocconi University.

Tronti, Leonello. 1991. Employment Protection and Labour Market Segmentation: the Fiftieth Anniversary of the Cassa Integrazione Guadagni. *Labour* 5 (1): 121-146.

Topkis, Donald. 1998. *Supermodularity and Complementarity.* Princeton: Princeton University Press.

Table B1

Robustness analysis, I

Parameter	Baseline	Larger separation cost	Less persistent revenue	Higher returns to scale
$1/(1-\rho)$	0.338 [0.0006]	0.328 [0.0005]	0.280 [0.0005]	0.368 [0.0005]
$1/\varphi$	0.483 [0.0006]	0.397 [0.0005]	0.336 [0.0005]	0.591 [0.0008]
η	0.407 [0.0005]	0.348 [0.0004]	0.281 [0.0004]	0.474 [0.0005]
μ	0.210 [0.0006]	0.243 [0.0006]	0.286 [0.0006]	0.149 [0.0004]
σ_ξ	0.294 [0.0004]	0.334 [0.0005]	0.378 [0.0006]	0.260 [0.0003]
σ_θ	0.210 [0.0007]	0.218 [0.0007]	0.211 [0.0007]	0.208 [0.0006]
σ_Z	0.203 [0.0002]	0.234 [0.0002]	0.265 [0.0003]	0.147 [0.0002]

Addendum: The following are the percent changes in (i) working time of a single type (ξ, θ) when only its ξ is raised; and (ii) average working time within a firm when all ξs are raised:

	(i) -5.10 (ii) -9.74	(i) -4.52 (ii) -8.14	(i) -3.84 (ii) -6.99	(i) -6.00 (ii) -12.73

NOTE: This shows results of the sensitivity analysis of Section 5.1. The larger separation cost is one year of earnings. To induce less persistent revenue, the AR(1) parameter in the firm productivity process is reduced to 0.6 from a baseline of 0.8. The higher returns to scale refers to an α of 0.835. Standard errors are in brackets. For a description of the counterfactual, whose results are summarized in the Addendum, see Section 4.3.

Table B2

Robustness analysis, II

Panel A: Empirical moments

Moment	Baseline	1994-2001 subsample	Adjusted working time estimates
$\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$	2.247	3.269	1.742
$\text{var}(\phi^W)/\text{var}(\phi^h)$	2.885	5.946	2.180
$\sqrt{\text{var}(\epsilon^h)}$	0.140	0.125	0.159
$\sqrt{\text{var}(\phi^h)}$	0.078	0.061	0.089
$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	-0.169	-0.059	-0.169
$\sqrt{\text{var}(\Delta \ln N)}$	0.175	0.184	0.175
$\text{E}[N]$	17.130	16.760	17.130

Panel B: Parameter values

Parameter	Baseline	1994-2001 subsample	Adjusted working time estimates
$1/(1-\rho)$	0.338 [0.0006]	0.454 [0.0012]	0.423 [n.a.]
$1/\varphi$	0.483 [0.0006]	0.321 [0.0009]	0.590 [n.a.]
η	0.407 [0.0005]	0.521 [0.0013]	0.354 [n.a.]
μ	0.210 [0.0006]	0.153 [0.0011]	0.116 [n.a.]
σ_ξ	0.294 [0.0004]	0.332 [0.0010]	0.248 [n.a.]
σ_θ	0.210 [0.0007]	0.287 [0.0019]	0.295 [n.a.]
σ_Z	0.203 [0.0002]	0.225 [0.0004]	0.198 [n.a.]
Addendum: Results of counterfactual (%)	(i) -5.10 (ii) -9.74	(i) -4.55 (ii) -6.70	(i) -6.19 (ii) -11.42

NOTE: Adjusted working time is corrected for undercounting total hours (see Section 5.2). Since the adjustment is based on auxiliary data, standard errors are not computed.

Table D1

The discretization of idiosyncratic types

Parameter	$3\ (\xi) \times 3\ (\theta)$ Model	$4\ (\xi) \times 4\ (\theta)$ Model
$1/(1-\rho)$	0.338	0.338
$1/\varphi$	0.483	0.483
η	0.407	0.407
μ	0.210	0.201
σ_ξ	0.294	0.294
σ_θ	0.210	0.210
σ_Z	0.203	0.203

NOTE: This presents parameter estimates for the baseline model where ξ and θ can each take one of three values (the "3×3" model) and for an extended version of the model where each type can take one of four values (the "4×4" model).

Table D2

Non-uniformly distributed types

Parameter	Baseline (ξ and θ are uniform)	ξ is single peaked, θ is uniform	ξ is uniform, θ is single-peaked
$1/(1-\rho)$	0.338	0.361	0.322
$1/\varphi$	0.483	0.484	0.484
η	0.407	0.406	0.406
μ	0.210	0.191	0.227
σ_ξ	0.294	0.286	0.293
σ_θ	0.210	0.056	0.014
σ_Z	0.203	0.203	0.203

NOTE: This table presents parameter estimates for cases in which ξ or θ is not a uniform random variable. When ξ is single-peaked, its distribution is given by $\Pr(\xi_1) = \Pr(\xi_3) = 0.3$ and $\Pr(\xi_2) = 0.4$, where $\xi_1 < \xi_2 < \xi_3$. The case in which θ is single-peaked is treated symmetrically.

Table D3

Correlated idiosyncratic types

Parameter	Baseline ($\xi \perp \theta$)	$\mathrm{Corr}(\xi, \theta) = -0.06$	$\mathrm{Corr}(\xi, \theta) = -0.12$
$1/(1-\rho)$	0.338	0.336	0.367
$1/\varphi$	0.483	0.481	0.480
η	0.407	0.408	0.408
μ	0.210	0.213	0.187
σ_ξ	0.294	0.301	0.298
σ_θ	0.210	0.183	0.039
σ_Z	0.203	0.203	0.203

NOTE: This table presents parameter estimates when idiosyncratic types ξ and θ are correlated. Please see the text in Appendix D for a description of the probability mass function that governs the distribution of θ given ξ.

Table D4

Alternative estimates of standard deviations of firm-year effects

Panel A: Moments

Moment	Baseline	Kline et al
Log change in paid days	0.078	0.061
Log change in earnings	0.132	0.116

Panel B: Parameter values

Parameter	Baseline	Kline et al
$1/(1-\rho)$	0.338	0.400
$1/\varphi$	0.483	0.377
η	0.407	0.391
μ	0.210	0.175
σ_ξ	0.294	0.338
σ_θ	0.210	0.263
σ_Z	0.203	0.200

NOTE: The top panel presents estimates of the standard deviations of firm-year effects for paid days and annual earnings using the bias correction in Kline et al (2020). The bottom panel presents structural parameter values when the Kline et al estimates of the variances of firm-year effects are used in estimation.

Table D5

Non-targeted moments

Moment	Data	Model
$\mathrm{Corr}(\Delta \ln h,\ \Delta \ln h_{-1})$	-0.149	-0.166
$\mathrm{Corr}(\Delta \ln W,\ \Delta \ln W_{-1})$	-0.275	-0.169
$\mathrm{Corr}(\Delta \ln N,\ \Delta \ln N_{-1})$	-0.048	0.215
$\mathrm{Cov}(\Delta \ln h, \Delta \ln N) / \mathrm{Var}(\Delta \ln N)$	0.008	0.155
$\mathrm{Cov}(\Delta \ln W, \Delta \ln N) / \mathrm{Var}(\Delta \ln N)$	0.020	0.277

NOTE: This table assesses the model's fit against five moments that were not targeted in estimation. Each moment refers to a firm-level outcome, e.g., h is average working time at the firm.

Table E1

Sensitivity Matrix

Moment / Parameter	$\frac{\text{var}(\epsilon^W)}{\text{var}(\epsilon^h)}$	$\frac{\text{var}(\phi^W)}{\text{var}(\phi^h)}$	$\sqrt{\text{var}(\epsilon^h)}$	$\sqrt{\text{var}(\phi^h)}$	$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	$\sqrt{\text{var}(\Delta \ln N)}$	$\text{E}[N]$
ρ	-5.110	3.473	-0.136	1.033	-0.740	-0.249	-0.081
φ	-2.151	1.914	-0.050	-1.120	-0.295	0.451	0.330
η	0.420	0.079	0.031	0.251	0.062	-0.091	-0.057
μ	0.468	-0.250	0.007	-0.140	0.069	0.070	-0.068
σ_ξ	0.003	0.136	0.298	-0.110	0.057	0.040	0.031
σ_θ	-0.388	0.226	0.195	0.013	-0.084	0.007	0.004
σ_Z	-0.135	0.118	0.008	0.078	-0.020	0.069	0.022

NOTE: The matrix is scaled such that each cell expresses the semi-elasticity of the parameter estimate (along rows) with respect to the moment (along columns).

Table G1

Ratio of idiosyncratic variances across industries

Lowest-ratio industries		Highest-ratio industries	
Fishing, Hunting and Trapping	1.240	Hospitals	4.311
Water Transportation	1.291	Air Transportation	4.136
Food Services and Drinking Places	1.311	Credit Intermediation	3.879
Animal Production	1.525	Clothing and Accessories Stores	2.942
Petroleum and Coal Products Mfg.	1.526	Motion Picture and Sound Recording	2.810
Wood Product Mfg.	1.602	Non-store Retailers	2.794
Wholesale trade, Nondurable Goods	1.625	Beverage and Tobacco Product Mfg.	2.689
Miscellaneous Store Retailers	1.625	Primary Metal Mfg.	2.463
Oil and Gas Extraction	1.631	Publishing Industries (except Internet)	2.341
Leather and Allied Product Mfg.	1.650	Mining (except Oil and Gas)	2.274
Apparel Mfg.	1.655	Chemical Mfg.	2.178
Couriers and Messengers	1.655	Paper Mfg.	2.168
Food Mfg.	1.665	Waste Mgmt. and Remediation Srvcs.	2.157
Plastics and Rubber Products Mfg.	1.672	Transit, Ground Passenger Transportation	2.103
Performing Arts and Spectator Sports	1.674	Computer and Electronic Product Mfg.	2.094
Amusement, Gambling, and Recreation	1.677	Utilities	2.089
Professional, Scientific, and Technical Srvcs.	1.678	Religious, Civic, and Professional Org.	2.076

NOTE: This presents the ratio of idiosyncratic earnings growth to idiosyncratic working time changes, as estimated in the Veneto Work History (VWH) files.

Table G2

Correlation between O*NET and Veneto Work History files

O*NET measure	2003-07 surveys of employees		Analyst database	
	Unweighted	Weighted	Unweighted	Weighted
Avg. of [A] & [B]	0.170 [0.169]	0.355 [0.003]	0.131 [0.300]	0.327 [0.008]
Avg. of [A]	0.209 [0.089]	0.455 [0.000]	0.187 [0.137]	0.299 [0.016]
Avg. of [B]	0.101 [0.416]	0.223 [0.069]	0.098 [0.437]	0.334 [0.007]
[A.1] (Teamwork)	0.228 [0.063]	0.425 [0.000]	N.A.	N.A.

NOTE: Standard errors are in brackets. Where applicable, regression weights are industry employment in the VWH. The analyst database does not include [A.1]-[A.2]. Therefore, the average score within group [A] is the mean of [A.3]-[A.4], and the average score of groups [A] and [B] is the mean of [A.3]-[B.4].

Table G3

Gender gaps in Veneto

Moment	Participation gap	Log earnings gap
10^{th} percentile	-79.3	-40.3
25^{th} percentile	-56.3	-32.2
50^{th} percentile	-38.8	-27.0
75^{th} percentile	0.6	-22.3
90^{th} percentile	38.8	-14.1
Mean	-27.1	-27.4
Correlation	-0.35	

NOTE: This table presents the distribution of gender gaps across industries. An industry's participation gap is the female share of the industry's workforce less the male share (and expressed in percentage points). The log earnings gap is women's log daily earnings less men's. An industry is a 3-digit NAICS sector.

Table G4

Correlation between gender gaps and ratio of idiosyncratic variances

Gender gap	Unweighted	Weighted
Participation	0.016 [0.901]	-0.157 [0.209]
Earnings	-0.284 [0.021]	-0.167 [0.180]

NOTE: This presents the across-industry correlation of gender gaps and the ratio of idiosyncratic earnings growth to idiosyncratic working time changes, as estimated in the Veneto Work History (VWH) files. Weights are industry employment in the VWH.

Figure G1

Ratio of idiosyncratic variances and O*NET teamwork scores across industries

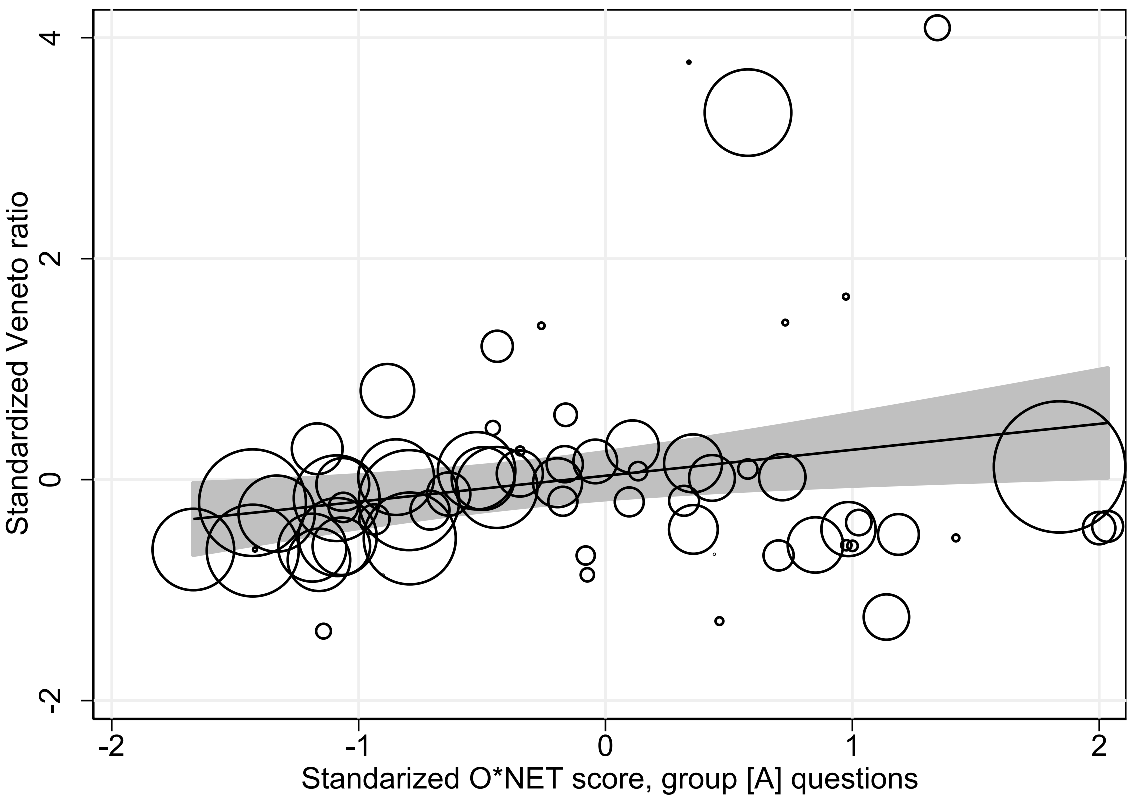

NOTE: This figure uses scores from the group [A] questions in the 2003-07 O*NET surveys and estimates of the ratio of idiosyncratic variances from the VWH. Please see the text (Appendix G) for a description of the group [A] questions.

Figure G2

Teamwork shares in the EWCS industries

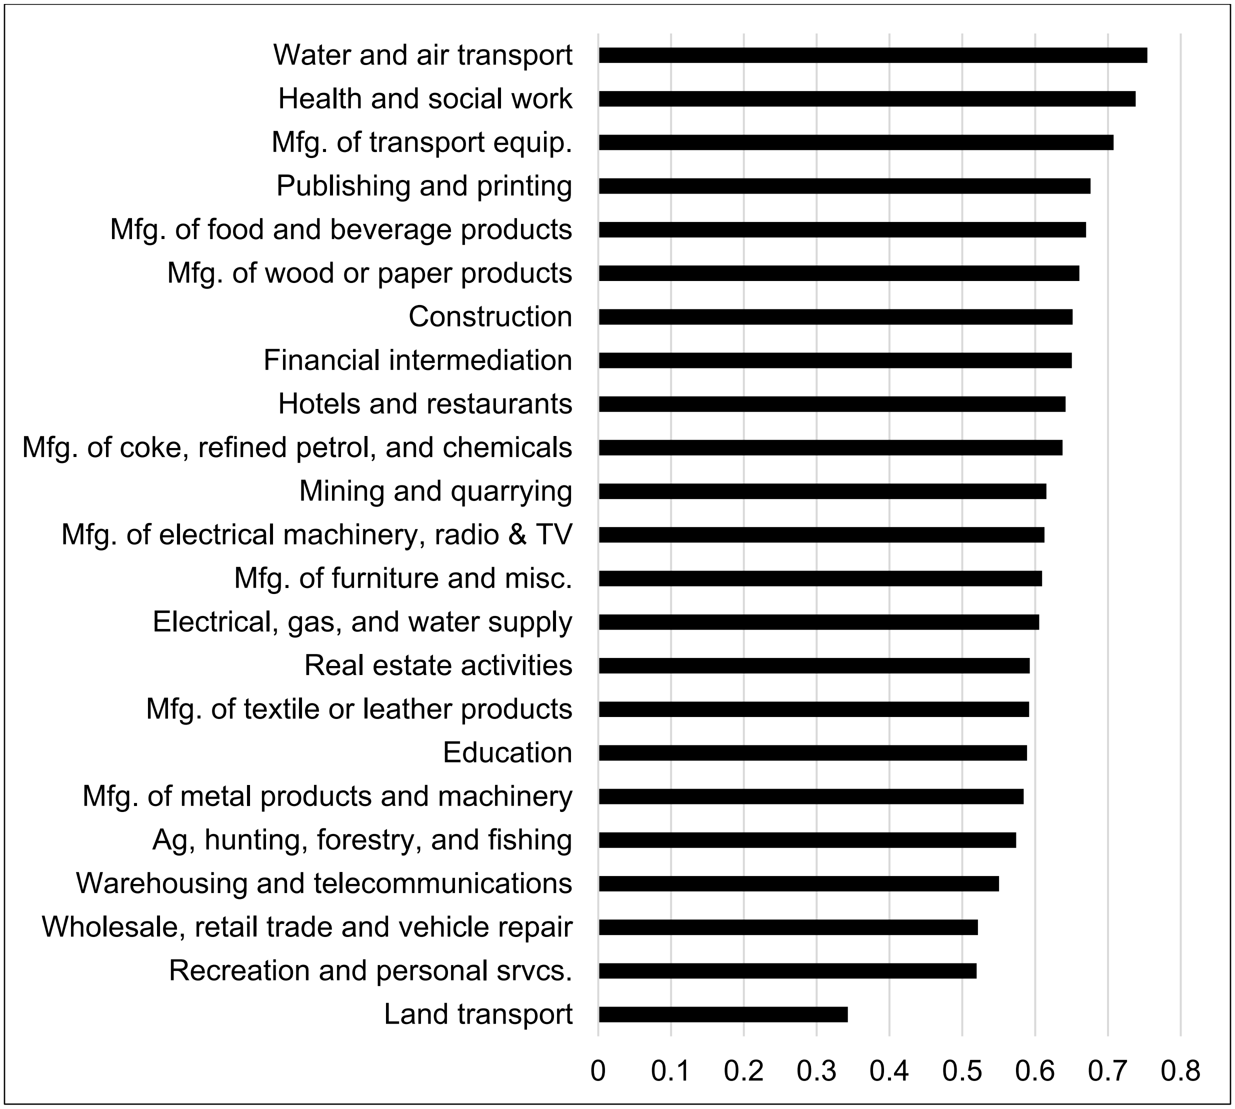

NOTE: This figure reports the share of European Working Conditions Survey (EWCS) respondents by industry who respond affirmatively to the teamwork question.

Figure G3

The shape of earnings under different degrees of complementarities

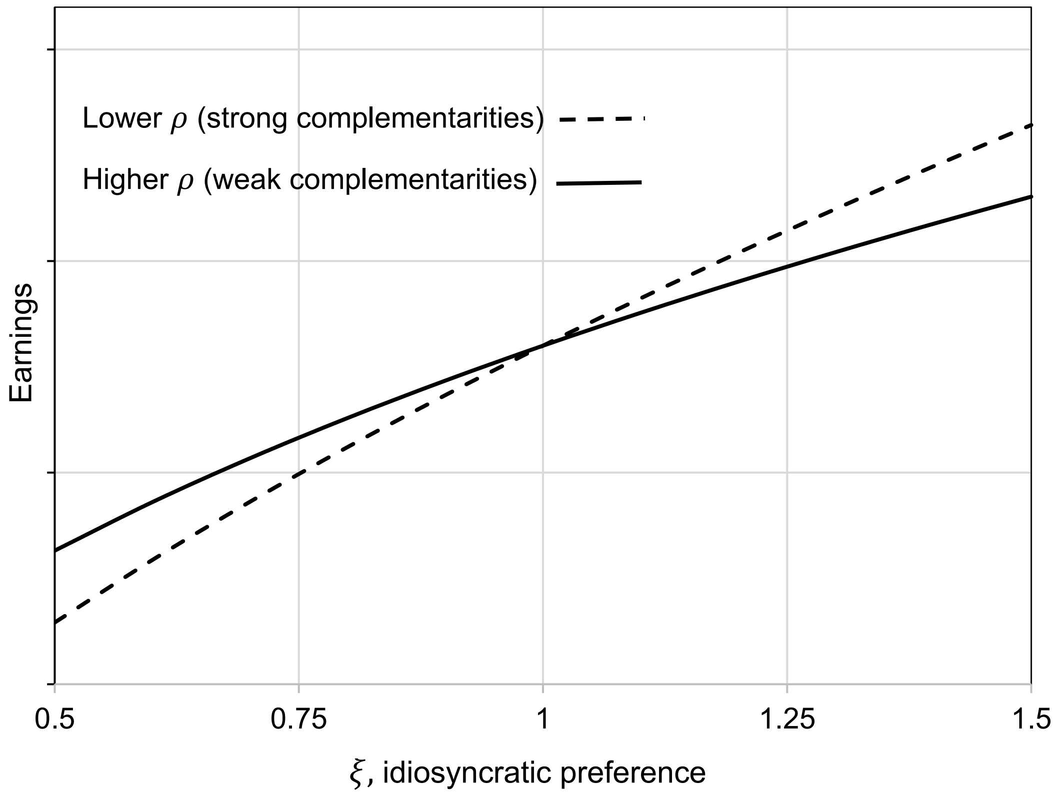

NOTE: This figure illustrates the shape of the earnings bargain over ξ, given different values for complementarities.