Federal Reserve Bank of Philadelphia

Working Papers
Research Department

WP 20-27
July 2020
https://doi.org/10.21799/frbp.wp.2020.27

Labor Supply Within the Firm

Michele Battisti
University of Glasgow

Ryan Michaels
Federal Reserve Bank of Philadelphia Research Department

Choonsung Park
Korea Institute of Finance

ISSN: 1962-5361

Disclaimer: This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in these papers are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: https://philadelphiafed.org/research-and-data/publications/working-papers.

LABOR SUPPLY WITHIN THE FIRM*

Michele Battisti   Ryan Michaels   Choonsung Park

July 17, 2020

Abstract

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 mimics the design of earlier studies by exploiting only idiosyncratic variation would find an elasticity less than half this.

JEL Codes: J22, J23, J31.

Keywords: Intensive margin, production complementarities, bargaining, employee-employer data.

*Battisti: University of Glasgow; Michaels: Federal Reserve Bank of Philadelphia; Park: Korea Institute of Finance. We thank Guiseppe Tattara and the team of researchers at the University of Venice for granting access to the Veneto Work History files. We appreciate feedback from participants at numerous conferences and workshops, and are especially thankful for comments and encouragement from Yongsung Chang and Igor Livshits. 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.

Variation in labor input occurs 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 labor market analysis, such as in the search and matching literature, has mostly focused on the extensive margin.

However, variation along the intensive margin is also significant. At the aggregate level, fluctuations in working time per employee are as large as movements in employment in several European economies (Llosa et al, 2014). At the plant level, data from U.S. manufacturers show that working time per person is as variable as employment (Cooper et al, 2015).[1] The size of these fluctuations can appear at odds with the implications of the earlier labor supply literature, whose estimates of the Frisch labor supply elasticity (for men) were centered around 0.2 and often very close to zero (Hall, 1999).

In this paper, we propose a framework that can help interpret this seemingly contradictory evidence on the elasticity of intensive-margin labor supply. In this setting, workers are complements in production but have heterogeneous preferences over leisure. Complementarities have important implications for the identification of the Frisch elasticity of (intensive-margin) labor supply. On the one hand, variation in a worker's own, *idiosyncratic* labor supply incentives yields relatively small changes in working time, since the efficient response is attenuated when one's effort is not complemented by higher effort of co-workers. On the other hand, *firm-wide* variation in the return to working coordinates the responses of heterogeneous workers, revealing the true willingness to substitute effort intertemporally. The model can thus predict more significant changes in firm-wide working time without implying counterfactually large responses to idiosyncratic events. We estimate the model using matched employer-employee data from Northern Italy and show how to recover the structural parameters governing the degree of complementarities and the Frisch labor supply elasticity.

Our approach has been foreshadowed (informally) in several earlier assessments of the labor supply literature. For instance, Pencavel (1986) notes that a worker's labor input is often coordinated by his employer. Relatedly, Hall (1999) contends that, "if an event occurs that is personal to the worker ... it is unlikely that the employer will agree to a reduction in weeks *ad hoc*" (p. 1148). These comments place the *employer* at the center of the theory of intensive-margin labor input.[2] In this paper's model, the firm does have a starring role.

Our theory of earnings and working time is presented in Section 1. The firm and its workers join in long-term employment relationships, bound together by the fact that extensive-

[1] Intensive-margin adjustments account for between one-fifth and one-third of the aggregate variation in U.S. total labor input at a quarterly frequency (Cacciatore et al, 2020).

[2] Of course, there may be some jobs (e.g., taxi driver) that align with what is envisioned by canonical labor supply theory, in which workers have substantial discretion over their schedules (see Farber, 2005).

margin adjustments are costly. Working time is set efficiently so as to maximize the surplus from the match. The resulting distribution of working time across employees represents a balancing of two interests— productive complementarities and heterogeneity in the disamenity from work. If the former is forceful enough, then employees agree, jointly with their employer, to vary their working time in a similar manner despite having different preferences.

Differences in preferences over leisure are accommodated, instead, by the earnings bargain, which is derived from a Nash-like surplus-sharing protocol. If a worker's labor input remains high despite an increase in her marginal value of time, she is compensated accordingly. Hence, under complementarities, the distribution of working-time adjustments across employees *within the firm* is compressed *relative* to the dispersion in earnings growth.

To assess our interpretation of working-time fluctuations and earnings, we introduce a unique source of panel data in Sections 2 and 3. The Veneto Worker History database is a matched employee-employer dataset that tracks the universe of workers and firms in the northern Italian region of Veneto from 1982 to 2001.[3] The dataset includes each employee's annual days worked for each of her 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. Still, the omission of daily hours in our data is arguably concerning. Therefore, we investigate this matter further using Italian household survey data, and find, reassuringly, that fluctuations in days worked account for about 80 percent of variation in total hours, consistent with the prevalence of Saturday overtime in Italy (Giaccone, 2009).

In Section 4, we estimate the model using the method of simulated moments. Our identification strategy relies on observing earnings and working time *inside* firms. Complementarities "squeeze out" the influence of idiosyncratic factors on working time. Instead, these factors are reflected primarily through the (within-firm) dispersion of earnings growth. We can thus infer the strength of complementarities by comparing the variance of working time adjustments across workers within firms to the variance of earnings growth (again, inside firms). If the ratio of the former to the latter is small, idiosyncratic variation is being suppressed in working time. Accordingly, our model infers a high degree of complementarities or, more exactly, a low elasticity of substitution across workers in production.

Whereas we identify complementarities off within-firm variation, preference parameters governing labor supply are more sharply revealed by *firm-wide* fluctuations in working time. Our approach uncovers an estimate of the Frisch elasticity of working time of 0.483. Estimates in the earlier, seminal life-cycle literature were clustered around 0.2 (see MaCurdy, 1981;

[3]In Italy, taxes and social insurance contributions are tied to days worked, which is why data on the latter are reported to the public social security organisation INPS.

Browning et al., 1985; Altonji, 1986). Importantly, our estimate of 0.483 is consistent with the observed covariance of working time and wage changes, which was the fundamental moment of the data on which the life-cycle estimates were based. This finding underscores the importance of isolating the specific driving forces behind working time and earnings. Interestingly, our result is more in line with recent estimates summarized in Chetty, Guren, Manoli, and Weber (2011). In Section 4, we discuss the source of variation used in more recent studies and why we suspect our results align with theirs.

In Section 5, we simulate a simple policy intervention in order to highlight implications of our results for the design of empirical labor supply analyses. A fraction of a firm's workforce is given a "treatment" that represents a shift in their own labor supply incentives. This intervention is inspired by the seminal U.S. Negative Income Tax (NIT) experiments, a series of randomized control trials administered in the 1960s and 1970s.[4] We contrast the outcome from this trial with the case in which all workers participate in the intervention. Reflecting the role of complementarities, working time declines by 40 to 105 percent more when all employees receive the treatment (depending on how the extensive margin adjusts). Furthermore, if we infer the Frisch elasticity using the treatment effect in the case where only a fraction of the workforce participates, the result is less than half the estimate (of 0.483) we uncover.

This experiment illustrates that the response of working time to an idiosyncratic event can be untethered from the underlying preference parameter. This is a simple, but important, point, because many influential studies of labor supply utilized this latter variation. Hall (1999) notes, for instance, that the tepid response of working time in the NIT randomized trials greatly informed the consensus on the Frisch elasticity. Yet this kind of variation—a sample of workers is selected to receive a cash grant—is clearly idiosyncratic to the worker. A similar point applies to the seminal life-cycle analyses of MaCurdy (1981) and Altonji (1986), which identify the Frisch elasticity based on the response of time worked to an individual's own (predictable) wage changes.

Our analysis of intensive-margin labor supply is related to several contributions in the literature. Chetty, Friedman, Olsen, and Pistaferri (2011) identify evidence of coordination in working time using the "bunching" of taxable income at kinks in the tax-rate schedule. We use different data and a distinct identification strategy, but like these authors, we find that idiosyncratic variation in the return to working fails to recover the true willingness of workers to vary their working time.[5] In a different setting, and using a model in which work-

[4] The NIT experiments were run in a handful of U.S. cities in the late 1960s and early 1970s. Participating households received a cash grant that was declining in their earnings. See Section 5 for more on the NITs.

[5] Our theoretical framework also differs from that in Chetty et al (2011), who assume firms post a single

ers coordinate their leisure, Rogerson (2011) likewise finds that variation in work incentives that is idiosyncratic to the worker can fail to elicit the true values of preference parameters. The potential importance of coordination was also a theme of the earlier literature on "hours constraints," which argued that labor supply responses to policy interventions are functions of preference parameters and frictions in adjusting hours (Dickens and Lundberg, 1993). We contribute both a novel way to formalize this idea and new evidence based on matched worker-firm data. Last, Chetty (2012) has offered an approach to inference that uses estimated elasticities to bound preference parameters even when the source of the wedge between elasticities and parameters is not explicit. Our strategy is complementary: we formalize a specific reason why reduced-form estimates may not identify preference parameters and use this model to recover the parameters.

The paper proceeds as follows. Section 1 introduces our model of earnings, working time, and employment demand. In Sections 2 and 3, we describe our data and present the empirical moments used in estimation. Section 4 estimates the model, and Section 5 assesses the implications of our results for earlier empirical work. Section 6 examines the robustness of our results with special attention to the implications of having incomplete measures of working time. Section 7 concludes.

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. The presence of any such idiosyncratic variation across workers serves to drive apart workers' desired 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

work schedule for all employees. The latter approach echoes Deardorff and Stafford (1976). Our model leaves room for idiosyncratic factors in order to accommodate the observed variation in working time.

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. 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.

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 preference 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): $\xi(j)\, h(j)^{\varphi} = Z\partial G/\partial h(j)$.[6] Optimal labor input therefore satisfies

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

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

Equation (2) imparts an important lesson about how to identify the (implied) Frisch elasticity, $1/\varphi$. The elasticity of working time to $\xi(j)$ depends, in general, on preference (φ) and production (ρ) parameters. One can infer φ from variation in $\xi(j)$ only in the special case of perfect substitutes, $\rho = 1$. If j is complementary with other jobs (equivalently, workers), an increase in $\xi(j)$ may have little impact. Indeed, if $\rho = -\infty$, $h(j)$ is *invariant* to $\xi(j)$ *regardless of* φ. Intuitively, under complementarities, it takes only a slight reduction in $h(j)$ to elevate j's marginal product in line with a higher marginal cost of effort.

According to (2), the value of $1/\varphi$ can be more reliably inferred from the response of $h(j)$ to *firm-level* variation, e.g., shifts in Z. Intuitively, firm-level variation serves to coordinate the efforts of heterogeneous workers such that the response of $h(j)$ to Z is mediated only by the preference parameter, $1/\varphi$, which is common across workers. This paper's main message

[6] The reader may note the absence of the marginal value of wealth in this first order condition. We have, implicitly, subsumed preference shifters *and* the marginal value of wealth under $\xi(j)$. We return to this point below.

is that this distinction between idiosyncratic (e.g., ξ) and firm-wide variation can help us understand apparently contradictory evidence on the elasticity of labor supply.

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 firm-wide employment, N, taking account of realistic employment adjustment frictions. Having the opportunity to adjust employment will influence the volatility of working-time.

1.2 The environment

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

Preferences. A worker's 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 shifts in the worker's marginal value of time.

To preserve tractability in a dynamic setting, we make several simplifying assumptions concerning ξ. At the start of each period, each worker draws a ξ from a K–dimensional set, $\mathcal{X} \subseteq \mathbb{R}^K$. These draws are i.i.d. across time and workers. Invoking a law of large numbers, a deterministic share $\lambda_\xi \in (0,1)$ of a firm's workforce will be of "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.[7] 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.

In general, shifts in ξ impinge on consumption. To avoid this complication, we assume each individual belongs to one of many large families (Merz, 1995; Hall, 2009). By pooling members' earnings, a family can insure consumption against member-specific risk (i.e., ξ). Therefore, the flow value of working (see equation (6) below) will not depend directly on the degree of risk aversion but only on earnings and the disamenity of supplying labor, $\xi\nu(h)$ (Trigari, 2006).

[7] Since types are redrawn each period, we do not have to track the distribution of workers across types (see Section 1.3.1). Later, we discuss the possible implications of persistence in ξ.

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, an assumption that facilitates the subsequent analysis. Total labor input of a type ξ is $n_\xi h_\xi$, where n_ξ is the measure of type ξ employment and h_ξ is the average supply of time within type ξ.[8] We assume that final output can then be written as a CES aggregate,

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

where ρ again reflects the elasticity of substitution across jobs; $\mathbf{n} \equiv \{n_\xi\}$ and $\mathbf{h} \equiv \{h_\xi\}$ are $K \times 1$ vectors; and $\alpha \in (0, 1)$ is the returns to scale. The departure from constant returns $\alpha < 1$ ensures a well-defined notion of firm size, $N \equiv \sum_\xi n_\xi$.[9] Note that, since $\alpha < 1$, the limiting case of perfect substitutes now refers to $\rho = \alpha$.

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).[10] A more explicit microfoundation may shed light on the source of these arrangements but would not necessarily enrich our description of the basic trade-off between coordinating working time and accommodating heterogeneous preferences.

In our empirical application, we work with a more general version of (4). In this case, workers also take an i.i.d. draw θ from a L–dimensional set of productivities, $\mathcal{Y} \subseteq \mathbb{R}^L$. A worker's type is now summarized by one of $M \equiv K \times L$ pairs, $\varsigma \equiv (\xi, \theta)$. Equation (4) generalizes to

$$\Gamma = ZG\left(\mathbf{h}, \mathbf{n}; \boldsymbol{\varsigma}\right) = Z\left(\sum_{\xi \in \mathcal{X}} \sum_{\theta \in \mathcal{Y}} \left(\theta n_{\xi,\theta} h_{\xi,\theta}\right)^\rho\right)^{\alpha/\rho}. \tag{5}$$

Labor market frictions. Labor market frictions mediate the formation of employment relationships. Following Roys (2016), there is a matching friction that renders

[8]The symmetry of workers within a type ξ will imply that their working time and earnings are equal.

[9]We interpret $(1-\rho)^{-1}$ as the elasticity of substitution across *jobs*, not across types per se. A simple example illustrates how (4) can "inherit" the elasticity of substitution across jobs from a more primitive production function. Firm output Γ is an aggregate over a continuum of jobs $j \in [0, 1]$, $\Gamma = Z\kappa \left[\int_0^1 \gamma(j)^\rho \, \mathrm{d}j\right]^{\alpha/\rho}$, where $\kappa \equiv K^{-\alpha}$ is a normalizing constant and output of job j, $\gamma(j)$, is assumed to be proportional to total "man-hours" on the job. Assume types are allocated an equal share, $1/K$, of jobs. Thus, if type ξ performs job j, the firm sets $\gamma(j) = K n_\xi h_\xi$. Substituting out $\gamma(j)$ using the latter yields (4).

[10]In the model, the i.i.d. nature of ξ means that workers perform different jobs over time. But within a period, (4) embeds a restriction that a type ξ can only execute a subset of jobs.

finite the rates of job finding and job filling. The magnitude of the friction depends only on aggregate conditions; since we analyze a firm's problem in an aggregate steady state, we do not elaborate further on matching. There are also per-worker costs, $\bar{c}$ and $\underline{c}$, of hiring and firing, respectively. Mandated severance, a kind of firing cost, is common in European labor markets, which is the context of our empirical application.

As we shall see, labor market frictions have a subtle but crucial part to play in our analysis. These frictions underpin a non-competitive labor market. As such, they create scope for bilateral bargains over working time and earnings, which, in turn, will reflect the variety of driving forces (ξ, θ, Z) at play. In this setting, it is crucial to trace out how the structural parameters, such as ρ and φ, mediate the effects of these different driving forces on working time and wages.

1.3 Characterization

This section characterizes the choices of working time, earnings, and employment. Note that the timing of events is such that working time and earnings are negotiated between firm and worker *after* employment has been decided. Accordingly, our notation reflects that the working time, h_ς, and earnings, W_ς, of an individual of type ς depend on the distribution of employment across types, $\mathbf{n}$.

1.3.1 Firm and worker objectives

Workers. Consider the surplus from working as type ς at a firm of productivity Z with workforce $\mathbf{n}$, denoted by $\mathcal{W}_\varsigma(\mathbf{n},Z)$.[11] In the current period, the employee earns a flow surplus equal to earnings, $W_\varsigma(\mathbf{n},Z)$, less (i) the disutility of supplying labor, $\xi\nu(h_\varsigma(\mathbf{n},Z))$ and (ii) the flow value, μ, of at-home time. In the next period, productivity at the worker's firm, Z', is realized, and the worker draws a type, $\varsigma' \equiv (\xi', \theta')$. A separation will occur if the present value of the future surplus of the firm-worker match is less than zero. Accordingly, the surplus from working is given by

$$\mathcal{W}_\varsigma(\mathbf{n},Z) = \begin{array}{c} W_\varsigma(\mathbf{n},Z) - \xi\nu(h_\varsigma(\mathbf{n},Z)) - \mu \\ +\beta\sum_{\varsigma'}\lambda_{\varsigma'}\mathbb{E}\left[\max\left\{0, \mathcal{W}_{\varsigma'}(\mathbf{n}',Z')\right\}\right]. \end{array} \tag{6}$$

A few remarks on (6) are warranted. First, μ will be treated as a fixed parameter to be estimated. This can be justified under two restrictions. The first is that the immediate

[11] To limit notation, we do not make explicit the dependence of earnings and working time on the entire vector of values of idiosyncratic types, $\boldsymbol{\varsigma}$.

returns on at-home time, such as unemployment insurance and home production, are independent of type, ς. The second is that type is i.i.d., which means that the anticipated future returns on at-home time, such as the value of job search, also do not depend on the current ς.

Second, since the flow surplus is expressed in units of numeraire, it should be written, more generally, as $(\xi/\ell) \cdot \nu(h_\varsigma)$, where $1/\ell$ is the inverse of the marginal value of wealth. Accordingly, working time and earnings will hinge on the *ratio*, ξ/ℓ. However, our data do not measure wealth (or consumption), so we cannot separately identify these two elements. To proceed, we treat ℓ as an i.i.d. draw such that ξ/ℓ satisfies the restrictions on ξ outlined in Section 1.1. We then suppress ℓ in what follows, though we can still exploit the isomorphism between shifts in ξ and ℓ to interpret heterogeneity in the data.[12]

The firm. The firm has an initial workforce N_{-1}.[13] After productivity, Z, is realized, the firm may choose to hire. We assume a worker's type $\varsigma \equiv (\xi, \theta)$ 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. The number of separations of type-ς workers, s_ς, is defined by

$$s_\varsigma \equiv \lambda_\varsigma \mathcal{N} - n_\varsigma \geq 0, \tag{7}$$

where n_ς is the mass of type-ς workers retained. It follows that $N = \sum_\varsigma n_\varsigma$ is the measure of the workforce used in production (and "carried into" next period). Wages and time worked will be bargained after separations (if any) are made.

Now define the present value of a firm for a given allocation, $\mathbf{n} \equiv \{n_\varsigma\}$. Let π stand for profit gross of firing and hiring costs,

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

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

$$\tilde{\Pi}(\mathbf{n}, Z) \equiv \pi(\mathbf{n}, Z) + \beta \int \Pi(N, Z') \, \mathrm{d}F(Z'|Z), \tag{8}$$

where $\beta \in (0,1)$ is the discount factor, F is the distribution function of productivity Z', and Π is the continuation value. Note that Π can be written as a function of just two state

[12] Card (1990) flagged changes in the marginal value of wealth as a source of variation in working time.

[13] The subscript $_{-1}$ denotes a one-period lag, and a prime $'$ denotes next-period values.

variables, (N, Z'), despite the heterogeneity across workers within a firm. This tractability is purchased by the assumption of i.i.d. types $\varsigma \equiv (\xi, \theta)$, which implies that we do not have to track individual types of workers over time.[14]

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,

$$\Pi^{-}(\mathcal{N}, Z) = \max_{\mathbf{n}} \left\{ \tilde{\Pi}(\mathbf{n}, Z) - \underline{c} \cdot \sum_{\varsigma} [\lambda_{\varsigma} \mathcal{N} - n_{\varsigma}] \right\}, \tag{9}$$

subject to $n_{\varsigma} \leq \lambda_{\varsigma} \mathcal{N}$ for each ς. Then, step back 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

$$\Pi(N_{-1}, Z) = \max_{\mathcal{N}} \left\{ -\bar{c} \cdot [\mathcal{N} - N_{-1}] + \Pi^{-}(\mathcal{N}, Z) \right\}, \tag{10}$$

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

Note that (9)-(10) imply that a firm may hire and separate workers in the same period. In this case, each constraint, $\mathcal{N} \geq N_{-1}$ and $\lambda_{\varsigma} \mathcal{N} \geq n_{\varsigma}$, is slack. However, for realistic values of $\underline{c}$ and $\bar{c}$, this does not happen: Productivity must be quite low to warrant separations, in which case no hires will be made. Thus, at firms that separate, $\mathcal{N} = N_{-1}$.

1.3.2 Working time

We assume that the firm and each of its workers jointly choose working time efficiently by equating the worker's marginal disamenity of working time to the marginal value of his working time to the firm. Solving this first order condition yields the following result.

Proposition 1 *For any individual worker of type $\varsigma \equiv (\xi, \theta)$, the efficient choice of working time is given by*

$$\begin{aligned} h_{\xi,\theta} &= (\alpha Z \Omega(\mathbf{n}))^{\frac{1}{\varphi+1-\alpha}} \cdot \left[\theta^{\rho} n_{\xi,\theta}^{\rho-1} / \xi\right]^{\frac{1}{\varphi+1-\rho}}, \\ \textit{with} \;\; \Omega(\mathbf{n}) &\equiv \left(\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \left[y^{\varphi+1} n_{x,y}^{\varphi} / x\right]^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha-\rho}{\rho}}. \end{aligned} \tag{11}$$

Equation (11) indicates that the elasticity of working time with respect to Z is given by $1/(\varphi + 1 - \alpha)$. The latter is decreasing in the preference parameter φ, which determines the

[14]This tractability is lost in richer models of the production process. For instance, suppose workers specialize in certain jobs. Specialization is a form of persistent heterogeneity, which means that firms will have to monitor the entire distribution of workers across jobs.

rate of change in the marginal disutility of effort. The elasticity is increasing in the returns to scale α, because agents are more willing to "bunch" effort in one period relative to another when diminishing returns sets in more slowly. Clearly, in the limit $\alpha = 1$, the elasticity of firm-wide working time is pinned down exclusively by $1/\varphi$.

Equation (11) also reveals the role of complementarities in shaping the reaction of working time to idiosyncratic events, ξ and θ.[15] Imagine reassigning a single worker to another one of the $M - 1$ types, leaving unchanged the preferences and productivities of the remaining workers.[16] Straightforward differentiation establishes the following.

Corollary 1 *(I) The elasticity of working time with respect to ξ is $-1/(\varphi + 1 - \rho) \leq 0$. In the limiting case of $\rho \to -\infty$ (perfect complements), working time is therefore invariant to ξ. (II) The elasticity of working time with respect to θ is bounded above by $\alpha/(\varphi + 1 - \alpha) > 0$, which obtains if $\rho = \alpha$, and below by -1, which obtains as $\rho \to -\infty$.*

There are several aspects of Corollary 1 that deserve attention. First, the reactions 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, the response of working time is almost entirely suppressed if tasks are sufficiently strong complements—*regardless* of the value of φ.

Interestingly, the response of working time to θ does not vanish as $\rho \to -\infty$. The reason is that, unlike a shift in ξ, a perturbation to productivity, θ, has a *direct* effect on a worker's output. If tasks are gross substitutes, the higher marginal product stimulates an increase in working time. Otherwise, if tasks are complementary, working time is reduced to bring the outputs of this type into line with those of other types. The extent of the change in working time thus hinges on the extent of complementarities. In fact, as $\rho \to -\infty$, the response of working time is, again, entirely detached from φ.

Proceeding, the solution to working time (11) enables us to concentrate $\mathbf{h}$ out of the firm's problem. Substituting (11) into the revenue function (5) yields

$$ZG(\mathbf{h}, \mathbf{n}, Z) = \hat{G}(\mathbf{n}, Z) \equiv \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\frac{\alpha}{\alpha-\rho}\frac{\varphi+1-\rho}{\varphi+1-\alpha}}. \tag{12}$$

[15] Since ξ serves as "shorthand" for ξ/ℓ, we do not refer to $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi$ as a Frisch elasticity. If the source of the variation is ℓ, then this derivative is more akin to the marginal propensity to consume leisure out of changes in (nonwage) income.

[16] In other words, the distribution of employees over types $\mathbf{n} \equiv \{n_{\xi,\theta}\}$ is taken as given, even if the identities of the workers in the types change. Therefore, from the perspective of a single worker, firm-level aggregates, such as $\Omega(\mathbf{n}; \boldsymbol{\varsigma})$, are treated as fixed for this exercise.

Likewise, we use $\hat{\pi}(\mathbf{n},Z) \equiv \hat{G}(\mathbf{n},Z) - \mathbf{n}^T\mathbf{W}(\mathbf{n},Z)$ to denote profits conditional on optimal working time.

1.3.3 Earnings

Earnings are negotiated each period according to the Stole and Zwiebel (1996) bargain, which was generalized by Cahuc et al (2008) to the case of heterogeneous workers. Cahuc et al (2008) abstracted from the intensive margin and assumed a fixed rate of separations (layoffs). Our solution relaxes these restrictions.

Under the Stole and Zwiebel protocol, the wage is set by splitting the *marginal* match surplus, awarding a share, $\eta \in (0,1)$, to the worker.[17] The marginal surplus, in turn, is the sum of the worker's surplus, $\mathcal{W}_\varsigma(\mathbf{n},Z)$, and the firm's surplus, which has two parts. The first, denoted by $\mathcal{J}_\varsigma(\mathbf{n},Z)$, is the marginal value of type-ς labor gross of hiring and firing costs. Since surplus-sharing is carried out *after* $\mathbf{n} \equiv \{n_\varsigma\}$ has been chosen, $\mathcal{J}_\varsigma(\mathbf{n},Z)$ can be calculated simply by differentiating the Bellman equation (8) with respect to n_ς,

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

recalling that $N = \sum_\varsigma n_\varsigma$. As for the second component, note that the firm can be penalized $\underline{c}$ if it and the worker fail to agree, resulting in the worker's separation. Accordingly, the surplus from retaining the worker is $\mathcal{J}_\varsigma(\mathbf{n},Z) + \underline{c}$, and the earnings bargain solves

$$\mathcal{W}_\varsigma(\mathbf{n},Z) = \eta\left(\mathcal{W}_\varsigma(\mathbf{n},Z) + \mathcal{J}_\varsigma(\mathbf{n},Z) + \underline{c}\right). \tag{13}$$

Proposition 2 presents the solution to this surplus-sharing problem.

Proposition 2 *The Stole and Zwiebel bargain for a worker of type $\varsigma \equiv (\xi,\theta)$ is given by*

$$W_{\xi,\theta}(\mathbf{n},Z) = \eta\left[\kappa\frac{\partial\hat{G}(\mathbf{n},Z)}{\partial n_{\xi,\theta}} + r\underline{c}\right] + (1-\eta)\left(\kappa\xi\nu\left(h_{\xi,\theta}(\mathbf{n})\right) + \mu\right), \tag{14}$$

where $\kappa \equiv \frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha} \geq 1$, $r \equiv 1-\beta$ and $h_{\xi,\theta}(\mathbf{n})$ satisfies (11).

The bargain in (14) is a weighted average of the worker's contribution to the firm and the value of his non-market time. The former consists of the worker's productivity and the annuitized firing cost, $r\underline{c}$, which the worker "saves" the firm by continuing the match (for

[17] Brügemann et al (2019) show that splitting the marginal surplus is the outcome of a game in which a firm bargains with each worker in sequence, and the strategic position of workers is symmetric.

another period).[18] The latter also has two parts: outside of employment, the worker does not incur the disamenity $\xi\nu\left(h_{\xi,\theta}\left(\mathbf{n}\right)\right)$, and he avails himself of the flow value, μ. Interestingly, (14) shares features with the solutions of *collective* bargaining games, with a principal difference being that the union wage depends only on average, not marginal, product (Taschereau-Dumouchel, 2019).

Note that conventional approaches to estimating labor supply parameters fail when earnings are given by (14) and working time by (11). Whereas canonical labor supply theory has a competitive market as its backdrop, labor market frictions in our model create scope for bargaining over working time and earnings. Each bargain reflects, in turn, worker- and firm-specific factors. Accordingly, projecting h_ς on W_ς lacks a structural interpretation. Moreover, it is arguably fruitless to search for an instrumental variable that shifts earnings but not working time: an IV that is orthogonal to ξ, θ, and Z would be irrelevant to h_ς.

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 (11) and (14). To this end, it is helpful to first write out the earnings bargain (14) more explicitly using (11) and (12),

$$W_{\xi,\theta}\left(\mathbf{n},Z\right)=\varkappa\left(\alpha Z\Omega\left(\mathbf{n}\right)\right)^{\frac{\varphi+1}{\varphi+1-\alpha}}\left[\theta^{\varphi+1}/\xi\right]^{\frac{\rho}{\varphi+1-\rho}}n_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}}+\omega, \tag{15}$$

where $\varkappa\equiv\frac{\eta\varphi+(1-\eta)\frac{\varphi+1-\alpha}{\varphi+1}}{(\varphi+1)(1-\eta(1-\alpha))-\alpha}$ and $\omega\equiv\eta r\underline{c}+(1-\eta)\,\mu$. For any $\rho<\alpha$, earnings are increasing in the employment of other types (via $\Omega\left(\mathbf{n}\right)$) and decreasing in own employment.

More importantly, equation (15) sheds light on the mapping between idiosyncratic events, ξ and θ, and earnings. In particular, the earnings bargain can be far more accommodating of idiosyncratic pressures than working time. Consider an increase in the distaste for working, ξ. If tasks are strongly complementary, the efficient choice of working time is to virtually suppress any response to the change in ξ. The workers earn, in return, a premium for continuing to supply effort when doing so is especially costly, as indicated by (15). Thus, a change in ξ passes through to earnings much more so than to working time. The following corollary makes this intuition precise.

Corollary 2 *The absolute size of the log change in earnings with respect to* ξ, $\left|\partial\ln W_{\xi,\theta}/\partial\ln\xi\right|$,

[18] The worker can use $\underline{c}$ to negotiate a higher wage because the firm is subject to the severance cost as soon as the worker 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.

exceeds the absolute size of the log change in working time, $|\partial \ln h_{\xi,\theta}/\partial \ln \xi|$*, if tasks are sufficiently strong complements in the sense that* $\rho < -\left(1-\omega/W_{\xi,\theta}\right)^{-1}$.

Corollary 2 holds out the possibility of using data on earnings and working time to infer ρ. By shifting ξ while holding Z fixed, we are perturbing earnings and working time *within* a firm. Therefore, if we look across workers within a firm and observe more dispersion in earnings changes than in working time adjustments, the model infers that ρ is relatively low. Conversely, under strong substitutability, changes in ξ induce more variation in working time adjustments. These observations suggest that the relative dispersion of earnings and working time changes within the firm can identify the degree of complementarities.

There are, however, a few subtleties in the mapping from ρ to earnings and working time dynamics. Consider again a perturbation to ξ. Corollary 2 shows that the range of ρ over which the response of earnings is amplified (relative to the change in working time) depends on the share of earnings tied down by $\omega \equiv \eta r \underline{c} + (1-\eta)\,\mu$. The latter is, in turn, especially influenced by the outside option, μ (since r is small). Thus, by determining the weight of μ in the earnings equation, worker bargaining power (η) mediates the influence of ρ on earnings and working time.[19]

In addition, Corollary 2 refers only to a perturbation to ξ. As we saw earlier (Corollary 1), the reaction of working time to a change in θ is not muted even as $\rho \to -\infty$, as it is when ξ is perturbed. Indeed, strong complementarities in this case can induce, rather than mitigate, the response of working time to idiosyncratic variation.

Clearly, the mix of these two idiosyncratic forces—ξ and θ—is critical to earnings and working time dynamics. How can we identify the predominant source of variation? A key piece of evidence is the comovement of the wage *rate,* $w_{\xi,\theta} \equiv W_{\xi,\theta}/h_{\xi,\theta}$, and working time, $h_{\xi,\theta}$. A higher ξ (weakly) depresses $h_{\xi,\theta}$ and is compensated by a higher wage, $w_{\xi,\theta}$. In other words, it acts like an inward supply shift. In contrast, a change in productivity θ works like a demand shift, tending to move working time and, as long as ω is again not too large, the wage rate in the same direction. Corollary 3 formalizes this simple idea.

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* $\omega/W_{\xi,\theta}$ *is not too large in the sense that* $\left(1-\omega/W_{\xi,\theta}\right)(\varphi+1) > 1$.

The correlation between working time and wage rates is indeed negative in our data (see

[19] One may easily verify that a lower η also reduces $\varkappa$, the weight on the first term in (15).

Section 4).[20] Corollary 3 indicates that this fact can be accommodated, *for any* values of ω and φ, by variation in ξ. We therefore infer that ξ is likely to comprise a majority of the idiosyncratic variation and, by Corollary 2, working time and earnings changes within the firm convey critical identifying information about ρ.

1.3.5 Employment demand

Thus far, we have taken total firm employment, N, as given. However, if firms can shift along the extensive margin, working time does not have to bear the full burden of adjusting to changes in Z in particular.[21] Thus, the observed responses of working time are intertwined with the firm's dynamic employment demand.

Consider the problem of a firm of size $\mathcal{N} = N_{-1}$ (it did not hire). We ask if this firm should separate from workers of some type $\varsigma \equiv (\xi, \theta)$, taking as given the participation of the remaining types. A separation is made if the marginal value of labor, evaluated at N_{-1}, is less than the separation cost,

$$\frac{\partial \hat{\pi}\left(\boldsymbol{\lambda} N_{-1}, Z\right)}{\partial n_{\varsigma}} + \beta \int \Pi_1\left(N_{-1}, Z'\right) \mathrm{d}F\left(Z'|Z\right) < -\underline{c}, \tag{16}$$

where $\boldsymbol{\lambda}$ is a $M \times 1$ vector of the shares λ_ς, and the derivative of $\hat{\pi}$ is evaluated at the initial workforce, $\mathbf{n} \equiv \boldsymbol{\lambda} N_{-1}$. The appendix verifies that Π is supermodular, which implies that the marginal value of labor, the left-hand side of (16), is increasing in Z. It follows that there exists a threshold, $\hat{Z}_\varsigma\left(N_{-1}\right)$, such that a type-$\varsigma$ worker is separated if (and only if) $Z < \hat{Z}_\varsigma\left(N_{-1}\right)$. The type of worker separated first is the type ς for which $\hat{Z}_\varsigma\left(N_{-1}\right)$ is highest.

If Z falls further, the firm separates from another type, $\tau \neq \varsigma$. As the firm does this, separations from the first type ς continue. This reflects that workers are (q-) complements in production: as the firm reduces labor input of type τ, the marginal value of type ς falls further. Thus, the optimal policy prescribes that *both* types are separated in tandem. This intuition underlies the result given below. To state the proposition, we use the notation $\varsigma_1, \ldots, \varsigma_j, \ldots, \varsigma_M$ to convey that a type ς_j is the jth type to be separated.

Proposition 3 *There exists a ranking* $\varsigma_1, ..., \varsigma_M$ *and a corresponding sequence* $\Big(\hat{Z}_1\left(N_{-1}\right),$

[20] Abowd and Card (1987) note that, if labor supply is chosen taking the wage as given, earnings will be more volatile than working time as long as the driving force behind the wage is productivity. However, in this setting, the wage and working time will also be strongly positively correlated, unlike in our data.

[21] Consider a firm with workforce N. Recalling (11) and supposing $\lambda_{\xi,\theta} = 1/M \ \forall \ (\xi, \theta)$, it follows that $n_{\xi,\theta} = N/M$ and $h_{\xi,\theta} = \left(\alpha Z N^{\alpha-1} \Omega\left(\mathbf{1}; \boldsymbol{\varsigma}\right)/M\right)^{\frac{1}{\varphi+1-\alpha}} \cdot \left[\theta^\rho/\xi\right]^{\frac{1}{\varphi+1-\rho}}$. Thus, the effect of reducing Z on working time can be partially offset by lowering N.

$\hat{Z}_2(N_{-1}), ...\Big)$ *with the latter listed in decreasing order, such that workers of all types* $(\varsigma_1, ..., \varsigma_i)$ *are separated if* $Z \in \left[\hat{Z}_{i+1}(N_{-1}), \hat{Z}_i(N_{-1})\right]$.

In special cases, we can say even more about the separation policy. Suppose $\xi \in \mathcal{X} \subseteq \mathbb{R}^K$ is the only source of heterogeneity across workers, and $\lambda_\xi = 1/K$ for all types. Then, one can show that low-ξ workers will be the first to be separated. Intuitively, high-ξ workers supply less effort conditional on participation, and, as a result, their participation is valued all the more if jobs are complements. If the λ_ξs differ across types, complementarities imply that workers from relatively abundant cohorts (all else equal) will be separated first.

The final piece of the optimal policy is the decision to hire. Recall that the firm hires *before* types are drawn.[22] Thus, the firm simply chooses N, and each type's size will be increased in proportion to its share in the population. The marginal value of increasing employment to N exceeds the marginal cost of hiring, $\bar{c}$, if

$$\begin{aligned} & \frac{\partial \hat{\pi}(\boldsymbol{\lambda} N, Z)}{\partial N} + \beta \int \Pi_1(N, Z') \, \mathrm{d}F(Z'|Z) \Big|_{N=N_{-1}} \\ = \; & \sum_{\varsigma} \lambda_\varsigma \frac{\partial \hat{\pi}(\boldsymbol{\lambda} N_{-1}, Z)}{\partial n_\varsigma} + \beta \int \Pi_1(N_{-1}, Z') \, \mathrm{d}F(Z'|Z) > \bar{c}, \end{aligned} \tag{17}$$

where each cohort size is evaluated at $\mathbf{n} \equiv \boldsymbol{\lambda} N_{-1}$ (and, hence, N is evaluated at N_{-1}). Again by the supermodularity of the problem, the firm will hire only if Z exceeds a certain threshold, denoted by $\hat{Z}_0(N_{-1})$.

Equation (17) presumes the firm does not fire after types are drawn. In principle, though, it may do this: The decision to hire is based on the *expected* marginal value of labor across types, whereas the decision to fire is based on the *lowest* marginal value after types are drawn. In fact, we find that, for plausible values of $\bar{c}$ and $\underline{c}$, the separation condition (16) never holds if Z is so high as to satisfy (17). Hence, $\hat{Z}_1(N_{-1}) < \hat{Z}_0(N_{-1})$. The space between the thresholds defines the inaction region where, $N = N_{-1}$.

Figure 1 illustrates the labor demand policy for a case with four equally likely types ($K = L = 2$ and $\lambda_\varsigma = 1/4 \; \forall \; \varsigma$). There is a middling range of Zs, between $\hat{Z}_1(N_{-1})$ and $\hat{Z}_0(N_{-1})$, over which employment of each type is unchanged. To the right of $\hat{Z}_0(N_{-1})$, the firm hires, and each type's employment is increased equally. As Z declines below $\hat{Z}_1(N_{-1})$, type ς_1 employment is reduced, while other types' participation remains fixed. As Z falls further, a second type is separated jointly with type ς_1, consistent with Proposition 3.

[22] More exactly, the assumption is that firms do not see the type at the moment of hiring. In that case, it is not really critical when the type is drawn (though it must be drawn before bargaining).

2 Data source

Our micro data span the universe of private firms in Veneto, Italy. Located in the North East, Veneto is one of the largest and richest of Italy's 20 administrative regions.[23] In this section, we introduce the Veneto Work History (VWH) files, and make the case that Veneto is a reasonable testing ground for our theory.

2.1 The Veneto Work History files

Our empirical analysis uses the VWH dataset that has been organized and maintained by researchers at the University of Venice. The VWH is a matched worker-firm database that covers the region of Veneto for the years 1982-2001. For virtually every private-sector employee in Veneto, it records each employer for which he worked at least one day. Public-sector employees and the self-employed are excluded. The full sample contains 22.245 million worker-year observations.

The VWH data has a number of features that recommend it for this analysis. Most importantly, the VWH reports for each worker the number of annual days paid and the calendar months worked with each of the individual's employers. The availability of a measure of working time in a matched employee-employer database is critical to our estimation strategy. The VWH files also record each worker's annual earnings, from which we can compute the average daily wage.

Table 1 provides a set of summary statistics for the full sample. On average, workers work between 23 and 24 days per month (conditional on positive days worked that month). This reflects the prevalence of six-day weeks in Veneto in this period; the sixth day, in many cases, represents overtime. The average gross daily wage is around 120 euros, and on average the number of paid months per worker (per year) is 10.

Table 2 zeroes in on moments of the distribution of annual changes in paid work days.[24] 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.[25] 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.

[23] Regional income data from ISTAT, Italy's statistical agency, begin in 1995. Putting these data on a purchasing power parity (PPP) basis using estimates from the Penn World Tables, Veneto's average income per person was $27,433 during 1995-2001, ranking sixth among Italian regions. Its population of 4.45 million (1995-2001) ranked fifth, according to EUROSTAT.

[24] Table 2 pertains to the subsample of workers used in our baseline analysis. See Section 3.1 for details.

[25] One may interpret this inaction as indicative of overhead labor, as we discuss in Section 6.

The VWH's measure of paid days, as valuable as it is, is still an incomplete record of total time worked. First, paid days do not equate to days at work; the former includes paid leaves of absence such as vacation. If the paid time off is taken each year (e.g., August vacation), we would still correctly measure *changes* in working time. Forms of leave, such as maternity leave, that are only partially compensated will imply changes in working time in our data (see Ray, 2008). More importantly, the VWH does not capture variation in daily hours. We defer a more extended discussion of this matter until Section 6, where we assess in detail how our inference may be sensitive to this omission. We conclude that the "bottom line" of our results survives intact.[26]

2.2 Institutional context

Our use of data from Veneto requires a brief digression on the institutions of Italian labor markets. Though these institutions do influence earnings and working time, our reading of the evidence is that decision-making is, at the margin, reasonably decentralized, particularly in relatively high-income regions like Veneto. This supports our modeling approach.

There are three layers of wage bargaining in Italy. At the top, national unions negotiate minimum wages for broad industries, but in the relatively high-wage region of Veneto, these rarely bind (Card et al, 2014). One layer down, union representatives at the firm negotiate "add-ons" to national contracts, which specify firm-performance-related premia. In 1995, 37.5 percent of workers in North Eastern Italy were covered under a firm-level agreement (ISTAT, 2000). Among firms with at least 20 workers, this share is nearly one-half, and the average premia (over industry minima) is about 25 percent (Card et al, 2014).[27] Finally, management can award bonuses to individual workers independent of any union agreement (Dell'Aringa and Lucifora, 1994; Erickson and Ichino, 1995). Among firms with less than 20 workers, where firm-level contracts are rare, these individual premia are significant—as high as 25 percent (Cattero, 1989)—and heterogeneous, as illustrated by Brusco (1982) in his study of the North Central region of Emilia-Romagna.

National unions also negotiate weekly and annual hours limits. During the 1980s, however, working time restrictions—specifically, limits on overtime—were either explicitly eased

[26] The data also distinguish between part- and full-time workers and fixed-term and permanent contracts. We do not break down the workforce along these lines because the average shares of part-time and fixed-term workers were rather modest in our sample period. Restrictions on fixed-term contracts, which could be terminated after two years without penalty, were relaxed by Parliament in 2001 (Tealdi, 2011).

[27] Consistent with these observations, Card et al find that, in medium-sized and large Veneto firms, wages are responsive to fluctuations in firm value-added.

in union agreements or loosely enforced (Treu et al, 1993; Lodovici, 2000).[28] The Bank of Italy's Survey of Household Income and Wealth (SHIW) indicates that nearly 30 percent of workers recorded positive overtime in 1989, and, among these workers, average annual overtime hours were 220—equivalent to about 27 8-hour days. Overtime hours receded somewhat in the 1990s, with annual overtime hours (among those working overtime) falling to 180 by 2000.[29] The latter decline may have reflected union-negotiated limits as well as a reduction in the demand for overtime as barriers to part-time and temp work were lowered.[30]

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. The first is robustness. There are, unavoidably, sources of variation in the data that are absent from the model. We seek moments that are (relatively) independent of such variation and, thus, have a clearer link to their model analogues. This reasoning underlies why our moments relate to *changes* in working time and earnings rather than levels. In the data, levels are likely shaped by permanent heterogeneity that has no counterpart in the model. The second consideration is that firm-wide (i.e., Z) and idiosyncratic (i.e., ξ or θ) driving forces in the model are mediated by different parameters. Accordingly, we want to distinguish the firm-wide from the within-firm components of observed 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 developing the moments summarizing earnings and working time changes.

Empirical framework. Our empirical analysis centers around a simple regression model designed to distinguish variation across workers within a firm from firm-wide movements in working time. Letting $\Delta \ln h_{ijt}$ denote the log change in days worked for employee

[28] Overtime equals the number of weekly hours in excess of "normal" weekly hours. Since at least the early 1970s, union agreements have typically set normal weekly hours to be 40 (Treu et al, 1993).

[29] The unemployment rate in Italy was around 9.5 percent in both 1989 and 2000, suggesting that the fall in overtime was not due to a decline in labor demand. See Online Data Appendix for more on the SHIW.

[30] Part-time labor, though still rare, was used more often after payroll taxes on part-time wages were cut in 1994 (Watanabe, 2014). A 1997 law legalized temporary work agencies (Destefanis and Fonseca, 2007).

i in firm j between years $t-1$ and t, we estimate

$$\Delta \ln h_{ijt} = \boldsymbol{\chi}'_{ijt}\mathbf{C}^h + \phi^h_{jt} + \epsilon^h_{ijt}, \tag{18}$$

where $\boldsymbol{\chi}_{ijt}$ collects the (time-varying) worker characteristics in our data, $\mathbf{C}^h$ is a conformable vector of coefficients, and ϕ^h_{jt} is a firm-year effect. Equation (18) is applied to a subsample of workers who are employed at the same firm for consecutive years $t-1$ and t (see below for more on sample selection). The elements of $\boldsymbol{\chi}_{ijt}$ consist of a cubic in the worker's tenure (measured as of $t-1$) and the change in broad occupation (between $t-1$ and t).[31] These controls help purge the data of observable persistent heterogeneity in work schedules. The variation then captured in ϕ^h_{jt} and ϵ^h_{ijt} is what is used to estimate the structural model.

The firm-year effect, ϕ^h_{jt}, in (18) measures the log change in firm j's working time relative to the average change among all firms in year t. We interpret ϕ^h_{jt} as reflecting shocks to labor demand at the *firm level*, i.e., changes in Z. Accordingly, the variance of ϕ^h_{jt} is our measure of changes in firm-wide working time, and should be particularly informative as to φ.

It follows that the residual in (18), ϵ^h_{ijt}, isolates variation across workers *within a firm*. We interpret the variance of ϵ^h_{ijt} as reflecting idiosyncratic shifts in worker-specific preferences, ξ, and productivity, θ.

We also estimate the counterpart to (18) for earnings, which relates the log change in earnings, $\Delta \ln W_{ijt}$, to observables ($\boldsymbol{\chi}_{ijt}$) and firm-year effects,

$$\Delta \ln W_{ijt} = \boldsymbol{\chi}'_{ijt}\mathbf{C}^W + \phi^W_{jt} + \epsilon^W_{ijt}. \tag{19}$$

The moment, $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$, thus compares the variances of earnings and working time changes within the firm. From the perspective of the model in Section 1, a high degree of complementarities means that idiosyncratic variation in preferences (or productivity) reflected in $\text{var}\left(\epsilon^W\right)$ is not passed through to $\text{var}\left(\epsilon^h\right)$. Accordingly, the ratio of these two, $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$, should provide critical identifying information for ρ.[32]

Sample selection. Equations (18)-(19) are estimated off a sample of workers who stay at the same firm in consecutive years. We refer to our baseline sample as *2-year stayers*. In any year t, the sample consists of workers who were paid for at least one day in all

[31] Initial tenure helps control for the possibility that more tenured workers have less variable work schedules. As for occupation, we measure four broad categories. Blue-collar workers make up 65 percent of the sample; "clerks", or white-collar non-managerial workers, make up 31 percent; managers comprise about 1 percent; and apprentices, or interns, make up 3 percent.

[32] It is tempting to try inferring ρ from how an individual's working time responds to firm-wide working time. However, the Online Theory Appendix shows that this moment is surprisingly uninformative about ρ.

months of the first quarter of year $t-1$ and in all months of the last quarter of year t.[33] We remove firms in any year t with only one employee, as it would be awkward to analyze complementarities with these firms in the sample. We are left with 11.8 million worker-year observations.[34]

Though 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 more consistent participation at the firm. To this end, we also analyze 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.[35]

Regression estimates. Table 3 summarizes estimates derived from (18)-(19). The first three rows pertain to within-firm (idiosyncratic) variation. Specifically, the first row reports $\text{var}\left(\epsilon^W\right)$; the second shows $\text{var}\left(\epsilon^h\right)$; and the third gives the ratio of the two. In the sample of 2-year stayers, 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 $\text{var}\left(\phi^W\right)$, $\text{var}\left(\phi^h\right)$, and the ratio of the two. Note that the value of $\text{var}\left(\phi^h\right) = 0.078^2$ for 2-year stayers represents 1.5-2 days per month.[36] 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 abbreviated. Still, for both groups of stayers, $\text{var}\left(\epsilon^W\right)$ substantially exceeds $\text{var}\left(\epsilon^h\right)$.

Table 4 reports on sensitivity analysis with respect to the moment, $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$, which is especially critical to our strategy. This ratio is typically at least 2. It is higher at larger firms, which tend to pay higher wages and are, therefore, relatively unbound by

[33] One could instead select stayers based on the number of months of employment in adjacent years, regardless of where in a year those months lie. Among workers who draw pay for (any) nine months in years $t-1$ and t, the values of the moments are similar to those reported in Table 3.

[34] The restriction to stayers reduces the sample by about half, which seems consistent with Contini et al's (2009) analysis of large Italian firms. Their estimate of a 20% annual separation rate implies that, among workers at the start of year $t-1$, $1-(1-0.2)^2 = 36\%$ exit by the end of year t. The latter is a lower bound on *average* turnover, since turnover is lower at larger firms (Idson, 1993).

[35] Our selection of stayers may raise concerns about external validity if stayers are systematically different than the average worker in ways that we have not modeled. We return to this point in Section 6.

[36] Since the typical worker puts in about 23 days per month (Table 1), a one-standard deviation increase in log days of 7.8 log points represents 1.8 days.

union-bargained minima (Guiso et al, 2005). The ratio (for 2-year stayers in particular) rises only slightly if we restrict the sample to men. Finally, with the exception of transportation and communication, the ratio also does not vary by much across sectors, despite differences in the make-up of industries (i.e., the prevalence of public enterprises in the health and education sectors).[37]

Targeted moments. From this set of results, we select four moments to use in the estimation of the model of Section 1. They are: the standard deviations of idiosyncratic and firm-level movements in working time, $\sqrt{\text{var}\left(\epsilon^h\right)}$ and $\sqrt{\text{var}\left(\phi^h\right)}$, respectively; and the ratio of dispersion in earnings growth to working time changes both within the firm, as measured using $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$, and at the firm-level according to $\text{var}\left(\phi^W\right)/\text{var}\left(\phi^h\right)$. The moments are based on the sample of 2-year stayers.

3.2 Additional moments

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

First, we project $\Delta \ln h_{ijt}$ on the log change in daily earnings, with the latter given by $\Delta \ln w_{ijt} \equiv \Delta \ln W_{ijt} - \Delta \ln h_{ijt}$. The estimated coefficient on $\Delta \ln w_{ijt}$ is -0.169.[38] Our finding of a negative association between the two echoes earlier studies including Abowd and Card (1989), whose estimates imply a coefficient of -0.3. Though these earlier results have sometimes been attributed to division bias (Borjas, 1980; Hercowitz, 2009), we are less concerned about measurement error in our administrative data.

Interestingly, in the perfect-foresight life cycle framework of MaCurdy (1981), the loading on daily earnings in this regression is in fact the Frisch elasticity of working time. MaCurdy's estimates are small and often insignificantly different from zero.[39] Our regression results using Veneto data thus reaffirm that this approach fails to find any clear evidence of a significantly positive Frisch elasticity.[40]

The final two moments refer to employment. The first is the standard deviation of em-

[37] To estimate $\text{var}\left(\epsilon^W_{ijt}\right)/\text{var}\left(\epsilon^h_{ijt}\right)$ for men, we use all firms (and workers) in (18)-(19) but pool ϵ^W_{ijt} and ϵ^h_{ijt} across only male workers.

[38] This result reflects variation within the firm: we uncover virtually the same estimate using only the idiosyncratic portion of working time (ϵ^h_{ijt}) and daily earnings (the latter is, analogously, the residual in a regression of daily earnings on firm-year effects).

[39] MaCurdy's sample does not necessarily consist of stayers. But, his selection of workers—prime-age white males in stable marriages—are more likely to be in long-lived employer-employee relationships.

[40] One distinction between our analysis and MaCurdy (and related studies) is that we observe daily earnings rather than the hourly wage. Strictly speaking, then, our results are comparable only if daily hours are fixed. We return to this point in Section 6.

ployment growth across firms. This is calculated from the employment-weighted distribution of employment growth, so that it is representative of the employment volatility faced by a typical worker. The final moment is mean firm size among firms with at least two employees.

4 Model Estimation

Seven parameters are estimated. They are ρ, which governs the elasticity of substitution across jobs, $1/(1-\rho)$; the preference parameter, φ; worker bargaining power, η; the worker's outside option, μ; and the variances of idiosyncratic preferences and productivities, σ_ξ^2 and σ_θ^2 respectively, as well as the variance, σ_Z^2, of innovations to firm-wide productivity, 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. Firm productivity, Z, 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 analyses that can examine total factor productivity (TFP) data (we cannot). At the same time, though, there are moments of our data that should be informative with respect to these parameters. Our strategy has been to "split the difference": We fix $\zeta = 0.8$ based on plant-level estimates of TFP (Foster et al, 2008), but treat the standard deviation, σ_Z, as a free parameter.

Next, we set values for four parameters for which there is credible external information. First, the discount factor is set to $\beta = 0.941$, which is consistent with the average annual real rate of interest in Italy over our sample. Second, we fix $\alpha = 0.67$, which is consistent with OECD data on labor's share in Italy.[41] Third, we set the hiring cost, $\bar{c}$, at about 2.5 months of earnings based on survey evidence in Del Boca and Rota (1998).[42] Finally, our choice of the severance cost, $\underline{c}$, amounts to a little over 7 months of earnings. This is a synthesis of multiple separation costs in Italy (see the Online Data Appendix for calculations).

[41] We interpret the omission of capital in (5) to mean that capital is fixed. In that case, α will be closely related to labor share, though the two are not identical because of worker bargaining power ($\eta > 0$).

[42] Del Boca and Rota survey 61 manufacturing firms in the region of Lombardy, which is just west of Veneto. The typical firm had between 16 and 49 workers. Clearly, the cost of hiring may differ by firm size and sector. However, the survey offers the only direct estimate of hiring costs in Italy of which we are aware.

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 our view, it is unlikely that the model could discriminate between uniform and non-uniform weighting; we select the former because it is particularly simple. In total, then, we have 9 types of $\varsigma \equiv (\xi, \theta)$, with each cohort equally represented in the population, e.g., $\lambda_{\xi,\theta} = 1/9$ for each (ξ, θ).

Given these choices, and guesses for the structural parameters to be estimated, we simulate 10 datasets, each consisting of the earnings, employment, and working time outcomes of the nine types within $20,000$ firms.[43] Outcomes are simulated over 220 years, but we compute moments from only the final 20 years. The moments are computed within each simulated dataset, and the average across datasets is treated as the model's analogue to the empirical estimates. The structural parameters are chosen to minimize the equal-weighted quadratic loss between model-implied and empirical moments.[44]

4.2 Main results

Table 5 summarizes our results. The top panel lists the empirical and model-generated moments. The model replicates the moments exactly (as might be expected from a just-identified model). The bottom panel lists MSM estimates of the structural parameters.

Frisch elasticity. Our estimate of φ implies a Frisch elasticity of $1/\varphi = 0.483$. This result suggests a greater willingness to vary working time than implied by MaCurdy (1981) and the earlier life-cycle literature. Estimates of $1/\varphi$ centered around 0.15–0.20 and were often statistically indistinguishable from zero.[45] In fact, as we noted above, we would infer a negative value of $1/\varphi$ if we employed the estimation strategy of MaCurdy in our data.

One exception in the life-cycle literature is Pistaferri (2003), who finds $1/\varphi = 0.7$. Interestingly, Pistaferri's finding is not necessarily inconsistent with the message of this paper. Pistaferri identifies $1/\varphi$ by estimating the response of total hours (in year t) to the survey respondent's *expected* earnings growth (between years $t-1$ and t). In our setting, expected earnings growth reflects the laws of motion of $\varsigma_t \equiv (\xi_t, \theta_t)$ and Z_t. Our model indicates that

[43] Clearly, our dataset consists of more firms. However, given that we simulate and average among 10 panels, the moments were virtually invariant to the addition of more firms. Furthermore, given a panel size of 20,000 firms, the simulation of more panels also had a negligible effect on the moments.

[44] The model is just-identified, and will reproduce the observed moments exactly (see below). Therefore, the use of the identity weighting matrix is without loss of generality.

[45] See Table 1 in MaCurdy (1981) and Tables 1-2 and 4 of Altonji (1986). We draw from specifications that control for year effects.

if ς_t is transitory and Z_t is persistent, the expected path of earnings will be shaped by the latter. In that case, Pistaferri's estimate may in fact reflect the response of working time to *firm-wide* variation.

Elasticity of substitution. Our estimate of $\rho = -1.962$ implies an elasticity of substitution across jobs within the firm of $(1-\rho)^{-1} \cong 0.338$. To convey this result in more concrete terms, consider the reaction of working time to a (one log point) change in an individual's preference ξ, holding fixed employment of each type. If workers were perfect substitutes such that $\rho = \alpha$, working time would adjust by $(\varphi + 1 - \alpha)^{-1} = 0.417$ log points. Our estimate of ρ instead implies a response less than half as large, $(\varphi + 1 - \rho)^{-1} \cong 0.199$.

Worker bargaining power. Our estimate of $\eta = 0.407$ is in the neighborhood of Roys' (2016) estimate of 0.52 though we use very different data. Roys' French firm-level panel lacks data on working time but includes revenue, enabling him to identify η using the comovement of wages and sales per worker. On the other hand, our estimate of η implies that earnings are more responsive to average product than found in Guiso et al (2005). Interestingly, our estimate of η declines if we re-parameterize the process of Z to induce a persistence in revenue that is comparable to that measured by Guiso et al (see Section 6).

Flow outside option. The outside option is estimated to be $\mu = 0.210$, which represents 70 percent of average earnings in the model. To put the latter in context, note that μ encompasses the entire amortized return to non-working time, including both the instantaneous gains such as unemployment insurance (UI) income as well as the anticipated future gains from searching for new work. By itself, UI replaces nearly 50 percent of earnings in the first year of unemployment in Italy.[46] Accordingly, we view a μ equal to 70 percent of earnings as being quite plausible.

Shocks. Our estimate of σ_Z implies a standard deviation of the log change in firm productivity (i.e., $\mathrm{var}(\Delta \ln Z)$) of 0.214. This is similar to estimates implied by plant-level TFP in European economies (see "France" and "Spain" in Table 2 of Asker et al, 2014). As for idiosyncratic heterogeneity, we find that it is slightly more substantial than firm-level dispersion: the standard deviation of firm productivity $\left(\text{e.g., } \sqrt{\mathrm{var}(\ln Z)}\right)$ is 0.338, and the standard deviation of the sum of idiosyncratic disturbances is $\sqrt{0.294^2 + 0.210^2} \cong 0.361$.

4.3 Identification

The map between moments and parameter estimates is rarely clear-cut in structural models. We can use the sensitivity matrix proposed by Andrews et al (2017), however, to guide a

[46] Benefits are offered beyond the first year to older workers at a reduced rate. See the Online Data Appendix for the details of our calculations.

discussion of the identification of the model's parameters. In the sensitivity matrix, which is reported in Table 6, each column measures the response of the parameter estimates to a one percentage point perturbation of a moment. Formally, since our model is just-identified, the sensitivity matrix takes a particularly simple form: it is the *inverse* of the Jacobian of the model-based moments with respect to parameters.[47]

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 .051. Thus, as foreshadowed in Section 1, a higher value of $\operatorname{var}\left(\epsilon^W\right)/\operatorname{var}\left(\epsilon^h\right)$ signals that within-firm driving forces ξ and θ are being channeled into earnings rather than working time, which is indicative of greater complementarities.[48]

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 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 induce the observed negative covariance of working time and wages. Indeed, the table confirms that an increase in this covariance implies a higher σ_ξ and a lower σ_θ.

Several other entries in Table 6 are highly intuitive. For instance, an increase in firm size, $\mathrm{E}[N]$, implies a lower outside option, μ : if μ is large, the rents from a match are small, and so fewer hires are made. In addition, if there is an increase in the dispersion of employment growth, $\sqrt{\operatorname{var}\left(\Delta \ln N\right)}$, 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 variation implies a higher η.

[47] Heuristically, consider a first order expansion of the moments $\mathbf{m}$ around parameter estimate $\mathbf{p} = \mathbf{p}_0$. If we invert the latter, we have $\mathbf{p} - \mathbf{p}_0 \equiv \mathrm{d}\mathbf{p} = \mathbb{J}\left(\mathbf{p}_0\right)^{-1}\mathrm{d}\mathbf{m}$, where $\mathbb{J}\left(\mathbf{p}_0\right)$ is the Jacobian of $\mathbf{m}$ w.r.t. $\mathbf{p}$ and $\mathrm{d}\mathbf{m} \equiv \mathbf{m} - \mathbf{m}\left(\mathbf{p}_0\right)$ represents a perturbation to the moments. The sensitivity matrix is $\mathbb{J}\left(\mathbf{p}_0\right)^{-1}$.

[48] 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. A column of the sensitivity matrix catalogues changes in parameter values that are required to fit the revised moment *and* preserve the fit of other moments. Thus, 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.

5 Implications for empirical research

Under production complementarities, the observed labor supply response to idiosyncratic variation represents a downwardly biased estimate of a worker's willingness to substitute labor intertemporally. We illustrate this point quantitatively in this section.

Specifically, consider a randomized control trial in which a fraction of a firm's workforce receives a "treatment," which is taken to be a lump-sum transfer. The transfer reduces the marginal value of income, ℓ. Recall that a reduction in ℓ has the same implications in our model as an increase in ξ (Section 1.2). Hence, we can simulate the randomized trial by simply increasing ξ for a share of the firm's workforce.[49] We then compute the change in working time among employees whose ξ is increased, and compare this outcome to a case in which the entire workforce is treated. The difference between the two outcomes illustrates the role of complementarities, since the latter will lead to a more dampened response of labor input to idiosyncratic treatments.

To proceed, suppose a transfer is distributed to a (small) number of workers at a firm. The size of the transfer is based on a typical grant in the U.S. Negative Income Tax (NIT) experiments, a set of (quasi-) randomized trials carried out in the late 1960s and early 1970s to study the labor supply response to transfer programs (see Burtless 1987 for an overview). We calibrate to the much-studied NITs in order to illustrate the implications of complementarities within a familiar context, even if we cannot capture all of the details of the NITs in our model.[50]

The implied transfer amounts to 37 percent of a participant's initial (pre-treatment) income.[51] We assume a marginal propensity to consume out of transitory income of 1/3 (Johnson et al, 2006). Since utility is separable in consumption C and effort (Section 1), we can map from the change in C to the change in ℓ using the first order condition for consumption. Letting ϕ denote the coefficient of relative risk aversion (evaluated at pre-treatment consumption), we have that, to a first order, $\Delta \ln \ell = -\phi \Delta \ln C$. Setting $\phi = 2$ following Hall (2009) implies $\Delta \ln \ell \cong -0.25$, which is equivalent in our setting to increasing ξ by 25 percent.

[49] It can be helpful to imagine that the transfer is distributed to workers who live in a small region (e.g., a neighborhood) but work within a larger (local) labor market. Earnings risk due to changes in workers' types and firms' productivities can be diversified within the region, which acts like a large family that insures members' consumption (see Section 1). Yet, since the transfer operates region-wide, it alters ℓ.

[50] Enrollment in NIT trials was restricted to families with relatively low (pre-trial) income. A participating household then received a maximum allotment, or guarantee, which was reduced by 50¢ per $1 of earnings. In the model, the transfer is neither conditioned on initial earnings nor subject to a reduction rate.

[51] This is the ratio of the (participation-weighted) average guarantee across NIT trials to the midpoint of the eligible income range (see Burtless, 1987).

This treatment is applied to one of the nine (ξ, θ) cohorts in the firm. The affected employees cut their time worked by 5.1 percent. To understand the implications of this result, suppose we viewed it through the lens of the limiting case of our model where $\rho = \alpha = 1$. Drawing on Corollary 1, we would infer a Frisch elasticity of $\Delta \ln h / \Delta \ln \xi = 0.051/0.25 = 0.204$. In other words, by neglecting complementarities, we would mistakenly infer a Frisch elasticity *less than half* as large as what we estimated in Section 4.

Now consider the case in which *all* workers in the firm receive the treatment. It is instructive to imagine at first that the designer of the randomized trial can hold employment fixed. Using (11) and substituting in our parameter estimates (including $\varphi = 1/0.483$), we would predict a decline in average working time of $\Delta \ln \xi / (\varphi + 1 - \alpha) = 10.4$ percent, or more than twice as much as in the case where one cohort is treated. Allowing for employment adjustments will take some of the burden off adjusting working time. Average working time then declines by 7.3 percent, which is still over 40 percent higher than what we find if only one cohort is treated.

Although our exercise is a stylized version of the NIT, it can be instructive to compare our results to the labor supply responses in the actual trials. Annual hours worked among men fell 7 percent (Burtless, 1987), but this almost surely overstates the change in intensive-margin labor supply. Rather, the decline in hours largely reflected longer job search spells (Moffitt, 1981; Robins and West, 1983). Thus, the intensive-margin response was lower and arguably more in line with our model's predictions.[52]

6 Robustness

This section probes the robustness of our estimates in several respects. In Section 6.1, we consider alternative choices for the pre-set parameters and subsamples. Section 6.2 assesses shortcomings in our measurement of working time. Lastly, Section 6.3 considers various other threats to the identification of production complementarities.

6.1 Pre-set parameters and sample periods

We re-estimated the model given a higher severance, $\underline{c}$; a lower persistence of productivity, ζ; and higher returns to scale, α. In another exercise, we re-estimated the model over a more recent subsample. Results are reported in Tables 7 and 8. Taken together, they point to a

[52] On the other hand, the reduction rate in the NITs (see footnote 50) likely lowered working time relative to the model. Note that, within the model, a reduction rate would not necessarily diminish the role of complemenarities, since it would operate regardless of the number of workers treated.

Frisch elasticity ($1/\varphi$) between 0.243 and 0.590, and an elasticity of substitution ($1/(1-\rho)$) between 0.223 and 0.454.[53] The final row of each table presents the results of the NIT-like experiment. With one exception, working time responds by 30 to 45 percent more to a firm-wide, as opposed to individualized, treatment.[54]

Higher severance and less persistent productivity have similar implications. Severance of one year's earnings compresses changes in employment. Larger firm-wide shocks can recreate the observed variance of $\Delta \ln N$, but a smaller Frisch elasticity, $1/\varphi$, and a lower bargaining power, η, are then needed to restrain movements in working time and earnings. Less persistent productivity also induces smaller adjustments in labor demand: if employment changes are costly to reverse, firms attenuate responses to transitory shocks. We set $\zeta = 0.329$, which induces a persistence in value-added comparable to Guiso et al (2005). Again, the model infers a higher σ_Z, and lower values of $1/\varphi$ and η. Notably, the smaller η pushes the elasticity of earnings to average product in the direction of Guiso et al (2005).[55]

Many parameters move in the opposite direction when α is raised. An α of 0.835—halfway between 1 (constant returns) and our baseline of $\alpha = 0.67$—makes employment more variable.[56] Therefore, a lower σ_Z is needed to match the variance of employment growth, which implies that $1/\varphi$ must be increased to induce the observed variance of working time.

We also re-estimated the model using data only for the subsample 1994-2001. This covers a period since the Italian government signed the Tripartite Agreement with employer and worker organizations. Consistent with the agreement's push to decentralize wage setting, Table 8 shows that the variance of earnings growth both *within* the firm and at the *firm level* is more volatile than in the full sample. Therefore, a higher η is needed, but one parameter cannot replicate two moments. By matching firm-level earnings dynamics, the higher η implies a counterfactually high ratio $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$. As a result, $1/(1-\rho)$ must rise to 0.454.[57] Note that the combination of a higher elasticity of substitution and a lower Frisch elasticity in this case implies that, in the NIT experiment, the difference between the partial and firm-wide treatments is relatively small.

[53] The model-implied moments (not shown) match the data exactly.

[54] These results take account of the response of employment to the treatment.

[55] The model-implied elasticity falls from 0.55 to 0.326; the elasticity is closer to 0.1 in Guiso et al.

[56] Our baseline of $\alpha = 0.67$ implicitly treats capital as if it were fixed. If the two factors are complements, any degree of capital adjustment will imply a (reduced-form) elasticity of output with respect to labor input that exceeds 0.67. Accordingly, we consider a higher, rather than lower, α.

[57] Interestingly, we estimate weaker complementarities even though $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$ is higher. Despite the relationship between the latter moment and ρ, one must still take into account the implications of changes in other parameters, such as η, for $\text{var}\left(\epsilon^W\right)/\text{var}\left(\epsilon^h\right)$.

6.2 Measurement error in working time

The omission of daily hours in the VWH may have important implications for several moments used in estimation. To illustrate, suppose workers adjust their number of days and daily hours in the same direction. The VWH data will understate the variance of changes in working time, leading us to underestimate the Frisch elasticity. Crucially, this understatement is not mirrored in earnings, which reflect changes in days and daily hours. As a result, estimates of the relative variance of earnings growth may be overstated.

We assess the quantitative importance of measurement error in working time using several data sources. The Italian Labour Force Survey (LFS) asks both about hours and days worked, and therefore enables us to directly measure the effect of missing hours on the variability of working time.[58] Among stayers, who work for the same employer in adjacent years, we calculate the year-over-year log change in days worked and weekly hours in the survey reference week. Strikingly, we find that the variance of log changes in days accounts for 78 percent of the variance of weekly hours growth. Thus, we might expect the VWH files to understate aggregate working time variation by 20-30 percent.

To be more precise about the implications for our VWH-based moments, we must consider how the "missing" variation in total working time is distributed across idiosyncratic and firm-wide sources. Suppose it is attributed in proportion to each source's contribution to the total variance in the VWH. Using estimates from Table 5, and noting that these components are (by construction) orthogonal, we find that the idiosyncratic part accounts for 75% of the total. Accordingly, we scale the total variance by (1/0.78) and distribute 75% of the increase to $var\left(\epsilon^h\right)$. Assuming VWH earnings are measured accurately, the ratio of the idiosyncratic variances of earnings growth to working time changes falls to 1.75 from 2.247 in the baseline case. The analogue for the ratio of firm-wide variances is 2.25, down from 2.885.

Another way of assessing the VWH, which does not require observing days worked, is to examine survey data on the variance of hours changes relative to the variance of earnings changes.[59] We draw on two surveys administered by the Bank of Italy, the Survey of Household Income and Wealth (SHIW) and the Survey of Industrial and Service (SIS) firms. The findings from these two surveys can then be compared against the *relative* variability of working time in the VWH. Note that the SIS enables us to examine, more specifically, the relative variance of firm-level working time, which has a counterpart in our VWH estimates.

[58] We pool LFS data between 1993, when the panel dimension of the micro data becomes available, and 2001, the final year of our VWH sample. For more on the LFS, and other survey data used in this subsection, see the Online Data Appendix.

[59] The LFS did not collect earnings data during our sample period, so we cannot compute both earnings and hours moments in the LFS.

We find that the absence of daily hours data leads us to understate the relative variance of working time changes by 20 to 30 percent, echoing results from the LFS. (For details, please see the Online Data Appendix.)

To examine the implications of these survey estimates for our structural parameters, we re-estimate the model assuming a ratio of idiosyncratic variances of 1.75, and a ratio of firm-wide variances equal to 2.25 (as in the LFS). We assume the VWH measures earnings correctly, so changes to the latter ratios imply corresponding changes to the variances of working time adjustments. No other VWH moments are altered.[60] Table 8 reports results. The higher relative variability of working time implies an increase in the elasticity of substitution to 0.437, up from 0.338 in the baseline case. This means that the elasticity of an individual's working time with respect to ξ is now $(\varphi + 1 - \rho)^{-1} = 0.249$, up from 0.199. At the same time, a higher variance of working time implies a higher elasticity of working time to firm-level events, which means that $(\varphi + 1 - \alpha)^{-1}$ *increases* to 0.485 (up from 0.416). Thus, working time still responds nearly twice as much to firm-level as to idiosyncratic events, which is consistent with results in Section 4.

Lacking data on hours raises one final, related challenge: we cannot compute hourly earnings. Rather, we measure daily earnings, which can conflate movements in daily hours and hourly wages. Therefore, the elasticity of days to daily earnings in the VWH could reflect the comovement of days and daily hours rather than the relationship between working time and remuneration per unit time. To see this, note that the elasticity of days to daily earnings, which we estimate to be -0.169, can be decomposed according to

$$-0.169 = \frac{Covar\left(\Delta \ln \text{ daily hours}, \Delta \ln \text{ days}\right)}{Var\left(\Delta \ln \text{ daily earnings}\right)} + \frac{Covar\left(\Delta \ln \text{ hourly wage}, \Delta \ln \text{ days}\right)}{Var\left(\Delta \ln \text{ daily earnings}\right)}. \tag{20}$$

The first term in (20) summarizes the role of daily hours in shaping the comovement of daily earnings and days. The numerator can be estimated using data on days and daily hours in Veneto from the Italian LFS, whereas the denominator can be taken from the VWH. The result lies between -0.045 and -0.10.[61] Comparing the midpoint of these estimates (-0.073) to -0.169 suggests, reassuringly, that fluctuations in hourly wages likely drive at least the majority of the comovement between days and daily earnings in the Veneto data.

[60] In particular, we retain the covariance of days and daily earnings in the VHW.

[61] The estimates differ depending on whether we use, respectively, usual daily hours or average daily hours in the reference week. One can argue for usual hours if "usual" is interpreted as average hours that year. This is in fact the concept that maps to the annual Veneto data.

6.3 Model mis-specification

This subsection collects a number of concerns regarding model mis-specification and their implications for the estimated degree of complementarities.

i.i.d. types. We have assumed types $\varsigma \equiv (\theta, \xi)$ are i.i.d.. Introducing persistence would have no direct effect on working time, since the latter is an intra-temporal choice. As for earnings, suppose that ξ is persistent, specifically. Reassuringly, the earnings bargain retains the form in (14) (see Online Theory Appendix). Any effects of persistence in ξ would be channeled through the outside option, μ; since μ encompasses the expected value of job search, it will depend on the (persistent) disamenity of work.[62] A fuller treatment of job search is beyond the scope of this paper, but we can note that, if μ were decreasing in ξ, earnings would respond less to idiosyncratic factors (since earnings are otherwise increasing in ξ when $\rho < 0$). Thus, the model would imply a smaller relative variance of earnings growth, $\text{var}\left(\epsilon^W\right) / \text{var}\left(\epsilon^h\right)$, which would have to be countered by *lowering* ρ.

Selection bias. We use the subsample of stayers to compute many of our moments. If stayers are different from the average worker in ways that are not modeled, then our inference can be distorted. To illustrate, suppose there is heterogeneity in complementarities across jobs within the firm—a feature we do not model. In particular, imagine that workers who separate have jobs that are *less* complementary with others. Our sample of stayers will then imply a downwardly biased estimate of ρ. An argument in favor of this hypothesis is that a firm competes more aggressively to retain workers in complementary jobs. However, by this logic, a similar firm that seeks to "poach" such a worker to fill a vacancy should also compete aggressively. This latter consideration suggests that separations may correspond to jobs with a *high* degree of complementarity.[63] A priori, then, it is unclear to us that separation events systematically reveal the complementarity of the jobs. This important matter must await more evidence on the relation between job types and worker flows.

Overhead labor. Insofar as overhead labor does not vary its days (by much), it will compress the distribution of changes in days worked, and lead us to infer a high degree of complementarities. A concern is that this inference masks a flexible production structure among non-overhead labor. To assess this, we drop workers who report 52 weeks of paid work in adjacent years and re-estimate (18)-(19) to recover the idiosyncratic components, ϵ^W and ϵ^h. This exercise is generous to the notion of overhead labor, as it drops any worker

[62] If θ is interpreted as match-specific productivity, then it does not persist across labor market states. Hence, ξ would influence the value of nonemployment, but not θ.

[63] This assumes the worker will perform a similar job in the new firm, and that the new firm's production structure is broadly comparable to the worker's present employer.

who participates full time in consecutive years. As anticipated, the distribution of days worked movements is compressed. Still, $\mathrm{var}\left(\epsilon^{W}\right)/\,\mathrm{var}\left(\epsilon^{h}\right)$ is 1.59—well above 1, which is still suggestive of a role for production complementarities (Corollary 2).

7 Conclusion

This paper has pursued the idea that, in the presence of production complementarities, an individual's labor supply is bound up with the working times of her colleagues. We developed a tractable theory of earnings, working time, and employment demand that formalizes this idea. The model expresses the intuition that, if there are strong complementarities, working time adjustments across employees inside a firm are compressed, regardless of the Frisch elasticity of labor supply. The Frisch elasticity is better informed in this setting by variation at the firm level; intuitively, firm-wide productivity movements serve to coordinate employees' working time and elicit the true elasticity.

We then showed how to estimate the model's structural parameters using moments from a matched employer-employee dataset from Veneto, Italy. Using the model's estimates, we carried out a simple counterfactual exercise to explore the consequences of failing to control for complementarities in conducting inference about labor supply elasticities. We find that if one infers the Frisch elasticity using only variation in labor supply incentives idiosyncratic to a worker, the estimate will be biased down by more than 50 percent.

We see a number of ways to further advance this line of research. First, complementarities are likely to mediate labor supply responses in many settings; our analysis of the NITs (and similar interventions) is just the "tip of the iceberg." For instance, suppose house prices increase unevenly across neighborhoods within a local labor market (see Guerrieri et al, 2013). As a result, the change in the marginal value of wealth can differ substantially among workers within a given firm. This means the mapping from housing price changes to working time changes depends critically on the extent of complementarities. Hence, our framework can be used to help disentangle the wealth effect from complementarities. Second, the increased use of matched employee-employer datasets can provide clues as to how our framework can, and should be, extended. For instance, the German LIAB Longitudinal Model, which reports workers' occupations as well as days worked, can be used to inform a richer production structure that takes explicit account of the types of jobs in firms.

8 References

Abowd, John and David Card (1989). "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), p. 411-445.

Abowd, John and David Card (1987). "Intertemporal Labor Supply and Long-Term Employment Contracts," *American Economic Review*, 77(1), p. 50-68.

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

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

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

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

Browning, Martin, Angus Deaton, and Margareth Irish (1985). "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle," *Econometrica,* 53(3), p. 503–43.

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

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

Burtless, Gary (1987). "The Work Response to a Guaranteed Income: A Survey of Experimental Evidence," in Alicia Munnell, ed., *The Income Maintenance Experiments: Lessons for Welfare Reform*, Federal Reserve Bank of Boston.

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), p. 943-72.

Cacciatore, Matteo, Giuseppe Fiori, and Nora Traum (2020). "Hours and Employment Over the Business Cycle: A Structural Analysis," *Review of Economic Dynamics.*

Card, David (1990). "Labor Supply With a Minimum Hours Threshold," *Carnegie Rochester Conference on Public Policy*, 33 (Autumn), p. 137-168.

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

Cattero, Bruno (1989). "Industrial Relations in Small and Medium-sized Enterprises in Italy," in Peter Auer and Helga Fehr-Duda (eds.), *Industrial Relations in Small and Medium-sized Enterprises: Report to the Commission of the European Communities*, Luxembourg: Office for Official Publications of the European Communities, p. 175-210.

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

Chetty, Raj, J. N. Friedman, T. 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), p. 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 & Proceedings,* 101(3), p. 471-75.

Contini, Bruno, Andrea Gavosto, Riccardo Revelli, and Paolo Sestito (2009). "Job Creation and Destruction in Italy," in Ronald Schettkat (ed.), *The Flow Analysis of Labour Markets*, Routledge, p. 195-214.

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), p. 37-50.

Deardorff, Alan and Frank Stafford (1976). "Compensation of Cooperating Factors," *Econometrica,* 44(4), p. 671-84.

Dell'Aringa, Carlo and Claudio Lucifora (1994). "Collective Bargaining and Relative Earnings in Italy," *European Journal of Political Economy*, 10(4), p. 727-747.

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), p. 427–449.

Destefanis, Sergio and Raquel Fonseca (2007). "Matching Efficiency and Labour Market Reform in Italy: A Macroeconometric Assessment," *Labour*, 21(1), p. 57–84.

Dickens, William and Shelly Lundberg (1993). "Hours Restrictions and Labor Supply," *International Economic Review*, 34(1), p. 169-92.

Erickson, Christopher and Andrea Ichino (1993). "Wage Differentials in Italy: Market Forces, Institutions, and Inflation," in Richard B. Freeman and Lawrence F. Katz (eds.) *Differences and Changes in Wage Structures*, University of Chicago Press, p. 265-306.

Farber, Henry (2005). "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy,* 113(1), p. 46-82.

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

Giaccone, Mario (2009). “Working Time in the European Union: Italy.” Eurofound European Working Conditions Observatory Comparative Report.

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

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

Hall, Robert (1999). “Labor Market Frictions and Employment Fluctuations,” in John Taylor and Michael Woodford (eds.) *Handbook of Macroeconomics*, North-Holland, p. 1137-1170.

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

Hercowitz, Zvi (2009). “Estimating Micro-Data Measurement Errors in Hours Worked and Hourly Wages,” mimeo, Tel-Aviv University.

Idson, Todd (1993). “Employer Size and Labor Turnover,” Columbia University Discussion Paper No. 673.

Istituto Nazionale Di Statistica (ISTAT) (2000). *La Rilevazione Sulla Flessibilità Del Mercato Del Lavoro Nel Periodo 1995-96.* Rome.

Johnson, David, Jonathan Parker and Nicholas Souleles (2006). “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96(5), p. 1589-1610.

Llosa, Gonzalo, Lee Ohanian, Andrea Raffo, and Richard Rogerson (2014). “Firing Costs and Labor Market Fluctuations: A Cross-Country Analysis,” mimeo.

Lodovici, Manuela Samek (2000). “Italy: The Long Times of Consensual Re-regulation,” in Gøsta Esping-Andersen and Marino Regini (eds.) *Why Deregulate Labour Markets?* Oxford University Press.

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

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

Moffitt, Robert (1981). “The Negative Income Tax: Would It Discourage Work?,” *Monthly Labor Review*, 104(4), p. 23-27.

Mortensen, Dale and Christopher Pissarides (1999). “New Developments in Models of Search in the Labor Market,” in Orley Ashenfelter and David Card (eds.) *Handbook of Labor Economics* vol 3: p. 2567-2627. Amsterdam: Elsevier Science.

Pencavel, J. H. (1986). “Labor Supply of Men: A Survey”, in Orley C. Ashenfelter and Richard Layard (eds.) *Handbook of Labor Economics*, North-Holland, p. 3-102.

Pistaferri, Luigi (2003). “Anticipated and Unanticipated Wage Changes, Wage Risk and Intertemporal Labor Supply,” *Journal of Labor Economics*, 21(3), p. 729-54.

Ray, Rebecca (2008). “A Detailed Look at Parental Leave Policies in 21 OECD Countries,” mimeo, Center for Economic Policy Research.

Robins, Philip and Richard West (1983). “Labor Supply Response.” In *Final Report of the Seattle-Denver Income Maintenance Experiment, Vol. 1, Design and Results*. Washington, D.C.: U.S. Government Printing Office.

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

Roys, Nicolas (2016). “Persistence of Shocks and the Reallocation of Labor,” *Review of Economic Dynamics*, Vol. 22, p. 109-130.

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

Taschereau-Dumouchel, Mathieu (2019). “The Union Threat,” mimeo, Cornell University.

Tealdi, Cristina (2011). “Typical and Atypical Employment Contracts: The Case of Italy,” MPRA Paper No. 39456.

Treu, Tiziano, G. Geroldi, and M. Maiello (1993). “Italy: Labour Relations,” in J. Hartog and J. Theeuwes (eds.) *Labour Market Contracts and Institutions*, Elsevier Publishers.

Trigari, Antonella (2006). “The Role of Search Frictions and Bargaining for Inflation Dynamics,” IGIER Working Paper No. 304.

Watanabe, Hiroaki (2014). *Labour Market Deregulation in Japan and Italy: Worker Protection Under Neoliberal Globalisation.* London and New York: Routledge.

9 Appendix

9.1 Working time

Proof of Proposition 1. Noting that $n_{\xi,\theta}h_{\xi,\theta}$ represents total working time of type $\varsigma \equiv (\xi, \theta)$, it is instructive to first write out (5) 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, $\xi h_{\varsigma}(i)^{\varphi}$, to her marginal product yields,

$$\xi h_{\varsigma}^{\varphi}(i) = \alpha Z^{\rho/\alpha} \Gamma^{1-(\rho/\alpha)} \theta^{\rho} (n_{\varsigma} h_{\varsigma})^{\rho-1}.$$

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

$$\xi h_{\varsigma}^{\varphi+1-\rho} = \alpha Z^{\rho/\alpha} \Gamma^{1-(\rho/\alpha)} \theta^{\rho} n_{\varsigma}^{\rho-1}, \ \varsigma \equiv (\xi, \theta). \tag{21}$$

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 the latter to substitute for any $h_{x,y} \neq h_{\xi,\theta}$ in (21), we recover (11) in the main text. ∎

Proof of Corollary 1. Totally differentiating (11) with respect to $h_{\xi,\theta}$, ξ, and θ yields

$$\mathrm{d}\ln h_{\xi,\theta} = \frac{\rho}{\varphi + 1 - \rho} \mathrm{d}\ln\theta - \frac{1}{\varphi + 1 - \rho} \mathrm{d}\ln\xi. \tag{22}$$

The elasticities with respect to θ and ξ are each increasing in ρ. Accordingly, each attains its maximum at $\rho = \alpha$ and its minimum at $\rho = -\infty$. ∎

9.2 Employment demand

In what follows, we need a weak restriction on the revenue function, $\hat{G}$.

Assumption 1 *The parameter, ρ, satisfies $\rho < \alpha$.*

This has two implications. First, it ensures that $\hat{G}$ is concave in $\mathbf{n}$, that is, the Hessian, $\nabla^2\hat{G}(\mathbf{n},Z)$, is negative definite. Second, it implies that $\hat{G}$ is supermodular, in that $\frac{\partial}{\partial Z}\frac{\partial\hat{G}}{\partial n_\varsigma} > 0$ for any type ς and $\frac{\partial^2}{\partial n_\varsigma n_\tau}\hat{G}(\mathbf{n},Z) > 0$ for any $\varsigma \neq \tau$. 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, $\hat{\pi}(\mathbf{n},Z)$, is concave in $\mathbf{n}$ and supermodular in $(\mathbf{n},Z)$.*

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

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

Proof. See Online Theory Appendix. ■

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.

Proof of Proposition 3. The optimal employment level of the first-to-be separated type ς is dictated by the first-order condition,

$$\frac{\partial\pi\left(n_\varsigma, \boldsymbol{\lambda}_{/\varsigma}N_{-1},Z\right)}{\partial n_\varsigma} + \beta\mathbb{E}\left[\Pi_N\left(N,Z'\right)|Z\right] + \underline{c} = 0, \tag{23}$$

where $\boldsymbol{\lambda}_{/\varsigma}$ is a $(M-1)\times 1$ vector of employment shares excluding the type-ς share and $N = n_\varsigma + \Sigma_{\tau\neq\varsigma}\lambda_\tau N_{-1}$. By supermodularity, the left side of (23) 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 (23). This yields a labor demand policy rule $n_\varsigma = \mathfrak{n}_\varsigma(N_{-1}, Z)$, where $\frac{\partial}{\partial Z}\mathfrak{n}_\varsigma > 0$.

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

$$\frac{\partial\pi\left(\mathfrak{n}_\varsigma(N_{-1},Z),\ \boldsymbol{\lambda}_{/\varsigma}N_{-1},\ Z\right)}{\partial n_{\tilde{\varsigma}}} + \beta\mathbb{E}\left[\Pi_N\left(N,Z'\right)|Z\right] < -\underline{c}, \tag{24}$$

where $N \equiv \mathfrak{n}_\varsigma(N_{-1},Z) + \Sigma_{\tau\neq\varsigma}\lambda_\tau N_{-1}$. Note that since the FOC (23) remains in effect as Z falls below $\hat{Z}_\varsigma(N_{-1})$, (24) is evaluated at the optimal size of cohort ς, $\mathfrak{n}_\varsigma(N_{-1},Z)$. Therefore, at lower Z, the left side 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, (24) 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 (23).

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: 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.

9.3 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

$$\mathcal{J}_\varsigma(\mathbf{n}, Z) \equiv \frac{\partial}{\partial n_\varsigma}\hat{\pi}(\mathbf{n}, Z) + \beta \int \Pi_N(N, Z')\,\mathrm{d}F(Z'|Z), \tag{25}$$

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

$$\frac{\partial}{\partial n_\varsigma}\hat{\pi}(\mathbf{n}, Z) \equiv \frac{\partial \hat{G}(\mathbf{n}, Z)}{\partial n_\varsigma} - \left[W_\varsigma(\mathbf{n}, Z) + \frac{\partial W_\varsigma(\mathbf{n}, Z)}{\partial n_\varsigma} n_\varsigma + \sum_{\tau \neq \varsigma} \frac{\partial W_\tau(\mathbf{n}, Z)}{\partial n_\varsigma} n_\tau\right]. \tag{26}$$

The expected marginal value of labor in (25) can be decomposed using Leibniz's rule,[64]

$$\begin{aligned} &\int \Pi_N(N, Z')\,\mathrm{d}F \\ &= \textstyle\sum_{j=1}^{M} \int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)} \Pi_N^{j-}(N, Z')\,\mathrm{d}F + \int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)} \Pi_N^0(N, Z')\,\mathrm{d}F + \int_{\hat{Z}_0(N)}^{\infty} \Pi_N^+(N, Z')\,\mathrm{d}F, \end{aligned} \tag{27}$$

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 $i \leq j$.[65] 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.

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

[65] We define $\hat{Z}_{M+1}(N) \equiv \min\{Z\}$, the minimum of the support of Z. The firm then separates from all types if $Z < \hat{Z}_M(N)$.

If the firm hires, the Envelope theorem, as applied to (10), implies that

$$\Pi_N^+ (N, Z') = \bar{c}. \tag{28}$$

To treat the case of separations, return to (9) and consider the state in which the firm separates only from type-1 labor, that is, workers of type ς_1.[66] The composition of the workforce is given by

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

where $\mathfrak{n}_1 (N, Z')$ denotes the optimal choice of type-1 labor conditional on adjusting and $\boldsymbol{\lambda}_{/1} \equiv (\lambda_2, ..., \lambda_M)$ is the vector of employment shares exclusive of type-1 labor. 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,

$$\Pi_N^{1-} (N, Z') = -\lambda_1 \underline{c} + \textstyle\sum_{i=2} \lambda_i \mathcal{J}_i \left(\mathbf{n}^{1-} (N, Z'), Z'\right), \tag{29}$$

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

$$\Pi_N^{j-} (N, Z') = -\Lambda_j \underline{c} + \textstyle\sum_{i=j+1}^{M} \lambda_i \mathcal{J}_i \left(\mathbf{n}^{j-} (N, Z'), Z'\right), \tag{30}$$

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

$$\mathcal{J}_i \left(\mathbf{n}^{j-} (N, Z'), Z'\right) \equiv \frac{\partial \hat{\pi} \left(\mathbf{n}^{j-} (N, Z'), Z'\right)}{\partial n_i} + \beta \int \Pi_{N'} (N', Z'') \, \mathrm{d}F. \tag{31}$$

The marginal value of labor in the "freezing" regime, $\Pi_N^0 (N, Z')$, can be obtained as follows. Forwarding (9)-(10) one period, setting $s'_\varsigma = 0 \; \forall \; \varsigma$ and $\mathcal{N} = N_{-1}$, noting that $\mathbf{n}' = \mathbf{n} = \boldsymbol{\lambda} N$ in this case, and differentiating with respect to N yields

$$\Pi_N^0 (N, Z') = \sum_{\varsigma \in \mathcal{X} \times \mathcal{Y}} \lambda_\varsigma \frac{\partial \hat{\pi} (\mathbf{n}, Z')}{\partial n_\varsigma} + \beta \int \Pi_N (N, Z'') \, \mathrm{d}F. \tag{32}$$

Now recalling (25), evaluating the latter at $\mathbf{n} = \boldsymbol{\lambda} N$, and taking a weighted average of $\mathcal{J}_\varsigma$

[66] 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.

across types reveals that Π_N^0 coincides with

$$\Pi_N^0(N, Z') = \sum_{\varsigma\in\mathcal{X}\times\mathcal{Y}} \lambda_\varsigma \mathcal{J}_\varsigma(\boldsymbol{\lambda}N, Z') = \sum_{j=1}^{M} \lambda_j \mathcal{J}_j(\boldsymbol{\lambda}N, Z'). \tag{33}$$

Substituting (28), (30), and (33) into (27) and inserting the result into (25) gives

$$\begin{gathered}\mathcal{J}_\varsigma(\mathbf{n}, Z) \equiv \tfrac{\partial}{\partial n_\varsigma}\hat{\pi}(\mathbf{n}, Z) \\ -\beta\underline{c}\sum_{j=1}^{M}\lambda_j F\left(\hat{Z}_j(N)|Z\right) + \beta\sum_{j=1}^{M}\sum_{i=j+1}^{M}\lambda_i\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\mathcal{J}_i\left(\mathbf{n}^{j-}(N, Z'), Z'\right)\mathrm{d}F \\ +\beta\int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\lambda_j\sum_{j=1}^{M}\mathcal{J}_j(\boldsymbol{\lambda}N, Z')\,\mathrm{d}F + \beta\bar{c}\left(1 - F\left(\hat{Z}_0(N)\right)|Z\right),\end{gathered} \tag{34}$$

where we have used

$$\sum_{j=1}^{M}\Lambda_j\left[F\left(\hat{Z}_j(N)|Z\right) - F\left(\hat{Z}_{j+1}(N)|Z\right)\right] = \sum_{j=1}^{M}\lambda_j F\left(\hat{Z}_j(N)|Z\right).$$

We next characterize the employee's surplus. Using the surplus-sharing condition, $\mathcal{W}_{\xi,\theta} = (\eta/(1-\eta))[\mathcal{J}_{\xi,\theta} + \underline{c}]$, we can recast (6) in terms of the firm's surplus,

$$\mathcal{W}_{\xi,\theta}(\mathbf{n}, Z) = \begin{gathered} W_{\xi,\theta}(\mathbf{n}, Z) - \xi\nu_{\xi,\theta}(\mathbf{n}) - \mu \\ +\beta\tfrac{\eta}{1-\eta}\mathbb{E}_{Z'}\sum_{i=1}^{M}\lambda_i\max\left\{0, \mathcal{J}_i\left(\mathbf{n}'(N, Z'), Z'\right) + \underline{c}\right\},\end{gathered} \tag{35}$$

where $\nu_{\xi,\theta}(\mathbf{n}) \equiv \frac{h_{\xi,\theta}(\mathbf{n})^{1+\varphi}}{1+\varphi}$. If the firm fires type-i labor (e.g., $Z' < \hat{Z}_i(N)$), equation (9) implies that the type's marginal value, $\mathcal{J}_i$, must be driven to $-\underline{c}$, hence, the surplus is zero. Note, though, that 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 (25), with $\mathbf{n}' = \mathbf{n}^{j-}(N, Z')$. If the firm hires (e.g., $Z' > \hat{Z}_0(N)$), equation (10) 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 (25). Collecting these observations, we have

$$\begin{gathered}\mathbb{E}_{Z'}\sum_{i=1}^{M}\lambda_i\max\left\{0, \mathcal{J}_i(\mathbf{n}', Z') + \underline{c}\right\} \\ = \int_{\hat{Z}_0(N)}[\bar{c} + \underline{c}]\,\mathrm{d}F + \int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\left[\sum_{i=1}^{M}\lambda_i\mathcal{J}_i(\boldsymbol{\lambda}N, Z') + \underline{c}\right]\mathrm{d}F \\ + \sum_{j=1}^{M}\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\sum_{i=j+1}^{M}\lambda_i\left[\mathcal{J}_i\left(\mathbf{n}^{j-}(N, Z'), Z'\right) + \underline{c}\right]\mathrm{d}F.\end{gathered} \tag{36}$$

Substituting this into (35) and rearranging yields

$$\begin{gathered}\mathcal{W}_{\xi,\theta}(\mathbf{n}, Z) = W_{\xi,\theta}(\mathbf{n}, Z) - \xi\nu_{\xi,\theta}(\mathbf{n}) - \mu \\ +\beta\tfrac{\eta}{1-\eta}\left\{\begin{gathered}\underline{c}\sum_{j=1}^{M}\lambda_j\left[1 - F\left(\hat{Z}_i(N)|Z\right)\right] + \sum_{j=1}^{M}\sum_{i=j+1}^{M}\lambda_i\int_{\hat{Z}_{j+1}(N)}^{\hat{Z}_j(N)}\mathcal{J}_i\left(\mathbf{n}^{j-}(N, Z'), Z'\right) \\ \sum_{j=1}^{M}\lambda_j\int_{\hat{Z}_1(N)}^{\hat{Z}_0(N)}\mathcal{J}_j(\boldsymbol{\lambda}N, Z')\,\mathrm{d}F + \bar{c}\left[1 - F\left(\hat{Z}_0(N)|Z\right)\right]\end{gathered}\right\}.\end{gathered} \tag{37}$$

Now inserting (34) and (37) into (13) and using (26), we have that

$$W_{\varsigma}\left(\mathbf{n},Z\right)=\eta\left\{\frac{\partial\hat{G}\left(\mathbf{n},Z\right)}{\partial n_{\varsigma}}-\sum_{\tau}\frac{\partial W_{\tau}\left(\mathbf{n},Z\right)}{\partial n_{\varsigma}}n_{\tau}+r\underline{c}\right\}+\left(1-\eta\right)\left(\xi\nu_{\varsigma}\left(\mathbf{n}\right)+r\mathcal{U}\right). \tag{38}$$

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

$$W_{\varsigma}\left(\mathbf{n},Z\right)=\eta\left[\kappa\frac{\partial\hat{G}\left(\mathbf{n},Z;\boldsymbol{\varsigma}\right)}{\partial n_{\varsigma}}+r\underline{c}\right]+\left(1-\eta\right)\left(\kappa\xi\nu_{\varsigma}\left(\mathbf{n}\right)+\mu\right), \tag{39}$$

where $\kappa\equiv\frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha}$. Using (12) and the solution for working time, one can calculate period profit and confirm Conjecture 1. ■

Proof of Corollary 2. Totally differentiating the earnings bargain (15) with respect to $W_{\xi,\theta}$ and ξ yields

$$\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi}=-\left(1-\frac{\omega}{W_{\xi,\theta}}\right)\frac{\rho}{\varphi+1-\rho}. \tag{40}$$

Then, recalling from (22) that the partial effect of ξ on $h_{\xi,\theta}$ is given by $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi=-\left(\varphi+1-\rho\right)^{-1}$, it follows that

$$\left|\frac{\mathrm{d}\ln W_{\xi,\theta}}{\mathrm{d}\ln\xi}\right|>\left|\frac{\mathrm{d}\ln h_{\xi,\theta}}{\mathrm{d}\ln\xi}\right|\Leftrightarrow\left|-\rho\right|>\left(1-\frac{\omega}{W_{\xi,\theta}}\right)^{-1}.$$

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 (22) and (40), 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 $\rho<\alpha$ and $\omega/W_{\xi,\theta}\in(0,1)$, the leading term in this expression 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

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

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

Table 1: Summary statistics of Veneto panel

Statistic	Mean	Std. Dev.
Average days per month per year	23.65	5.25
Job tenure (in months)	53.10	53.71
Average daily wage (2003 euros)	121.46	426.76
Total days worked per year	243.88	97.75
Average no. of months paid per year	9.96	3.38

NOTE: This summarizes aspects of the 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 $\|\Delta h\| > 10$	33.15%
Average $\|\Delta h\|$ if $\|\Delta h\| \neq 0$	19.06
Average $\|\Delta h\|$ if $\|\Delta h\| \neq 0$, excluding $\|\Delta h\| > 50$	9.75

NOTE: This table reaports 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 panel

Moment	**Interpretation**	**Data: 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-wide variances	3.989	2.885
$\dfrac{\text{cov}(\Delta\ln h, \Delta\ln w)}{\text{var}(\Delta\ln w)}$	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 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
Including only the following sectors:		
Wholesale and retail trade	3.921	2.005
Construction	2.286	1.714
Manufacturing	3.490	1.968
Transportation and communication	5.057	3.052

NOTE: This shows the ratio of the variance of the idiosyncratic component of earnings growth to that of log working time changes for different sub-samples.

Table 5: Model fit

Panel A

Moment	Model	Data (2-year stayers)
$\mathrm{var}(\epsilon^W)/\mathrm{var}(\epsilon^h)$	2.247	2.247
$\mathrm{var}(\phi^W)/\mathrm{var}(\phi^h)$	2.885	2.885
$\sqrt{\mathrm{var}(\epsilon^h)}$	0.140	0.140
$\sqrt{\mathrm{var}(\phi^h)}$	0.078	0.078
$\mathrm{cov}(\Delta \ln h, \Delta \ln w)/\mathrm{var}(\Delta \ln w)$	-0.169	-0.169
$\sqrt{\mathrm{var}(\Delta \ln N)}$	0.175	0.175
$\mathrm{E}[N]$	17.130	17.130

Panel B

Parameter	Symbol	Value
Elasticity of substitution across tasks	$1/(1-\rho)$	0.338 [0.0005]
Frisch elasticity of working time	$1/\varphi$	0.483 [0.0008]
Worker bargaining power	η	0.407 [0.0006]
Flow return on non-employment	μ	0.210 [0.0007]
Std. dev. of idiosyncratic preference	σ_ξ	0.294 [0.0008]
Std. dev. of idiosyncratic productivity	σ_θ	0.210 [0.0009]
Std. dev. of shock to firm productivity	σ_Z	0.203 [0.0002]

NOTE: This presents estimates of our baseline model. Standard errors are in brackets. Standard errors of $1/(1-\rho)$ and $1/\varphi$ are calculated via the Delta method..

Table 6: Sensitivity matrix

	$\frac{\text{var}(\epsilon^W)}{\text{var}(\epsilon^h)}$	$\frac{\text{var}(\phi^W)}{\text{var}(\phi^h)}$	$\sqrt{\text{var}(\epsilon^h)}$	$\sqrt{\text{var}(\phi^h)}$	$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	$\sqrt{\text{var}(\Delta \ln N)}$	$\text{E}[N]$
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 7: Robustness analysis, I

Parameter	I Baseline results	II Larger separation cost	III Less persistent revenue	IV Higher returns to scale
$1/(1-\rho)$	0.338 [0.0005]	0.328 [0.0005]	0.223 [0.0004]	0.367 [0.0004]
$1/\varphi$	0.483 [0.0008]	0.397 [0.0006]	0.243 [0.0005]	0.590 [0.0008]
η	0.407 [0.0006]	0.347 [0.0005]	0.198 [0.0004]	0.474 [0.0005]
μ	0.210 [0.0007]	0.243 [0.0007]	0.365 [0.0009]	0.149 [0.0004]
σ_ξ	0.294 [0.0008]	0.335 [0.0009]	0.483 [0.0012]	0.260 [0.0007]
σ_θ	0.210 [0.0009]	0.218 [0.0009]	0.212 [0.0008]	0.207 [0.0008]
σ_Z	0.203 [0.0002]	0.234 [0.0003]	0.330 [0.0005]	0.148 [0.0002]
Addendum : In an NIT trial, the change in (i) working time of (ξ,θ) pair when only that pair is treated; and (ii) average working time within firm when all pairs are treated:				
	(i) -5.11% (ii) -7.25%	(i) -4.55% (ii) -6.37%	(i) -3.04% (ii) -4.35%	(i) -5.74% (ii) -8.29%

NOTE: This shows the results of the sensitivity analysis of Section 6. In column II, the separation cost equals 1 year of earnings. In column III, the persistence of firm productivity is lowered to target the estimated persistence of value-added in Guiso et al (2005). In column IV, the returns to scale are raised to α=0.835. Standard errors are in brackets.

Table 8: Robustness Analysis, II

Panel A

Moment	**Baseline**	**1994-2001 subsample**		**Adjusted working time est.**	
	Model	Model	Data	Model	Data
$\text{var}(\epsilon^W)/\text{var}(\epsilon^h)$	2.247	3.269	3.269	1.754	1.750
$\text{var}(\phi^W)/\text{var}(\phi^h)$	2.885	5.946	5.946	2.249	2.250
$\sqrt{\text{var}(\epsilon^h)}$	0.140	0.125	0.125	0.159	0.159
$\sqrt{\text{var}(\phi^h)}$	0.078	0.061	0.061	0.088	0.088
$\frac{\text{cov}(\Delta \ln h, \Delta \ln w)}{\text{var}(\Delta \ln w)}$	-0.169	-0.059	-0.059	-0.169	-0.169
$\sqrt{\text{var}(\Delta \ln N)}$	0.175	0.184	0.184	0.177	0.175
$\text{E}[N]$	17.130	16.760	16.760	17.130	17.130

Panel B

Parameter (Symbol)	**Baseline**	**1994-2001 subsample**	**Adjusted working time est.**
$1/(1-\rho)$	0.338 [0.0005]	0.454 [0.0012]	0.437 [na]
$1/\varphi$	0.483 [0.0008]	0.321 [0.0010]	0.577 [na]
η	0.407 [0.0006]	0.521 [0.0014]	0.360 [na]
μ	0.210 [0.0007]	0.153 [0.0011]	0.108 [na]
σ_ξ	0.294 [0.0008]	0.332 [0.0011]	0.252 [na]
σ_θ	0.210 [0.0009]	0.287 [0.0020]	0.306 [na]
σ_Z	0.203 [0.0002]	0.225 [0.0005]	0.199 [na]
Addendum : Effects of NIT trial on working time (see Table 7)	(i) -5.34% (ii) -7.79%	(i) -4.49% (ii) -5.43%	(i) -6.17% (ii) -8.16%

NOTE: The adjusted working time estimates are derived from the full sample (1982-2001), but include a correction for under-counting total hours variation. See Section 6.2 for details. Since the adjustments are based on out-of-sample data, standard errors are not computed.

Figure 1: Labor demand policy

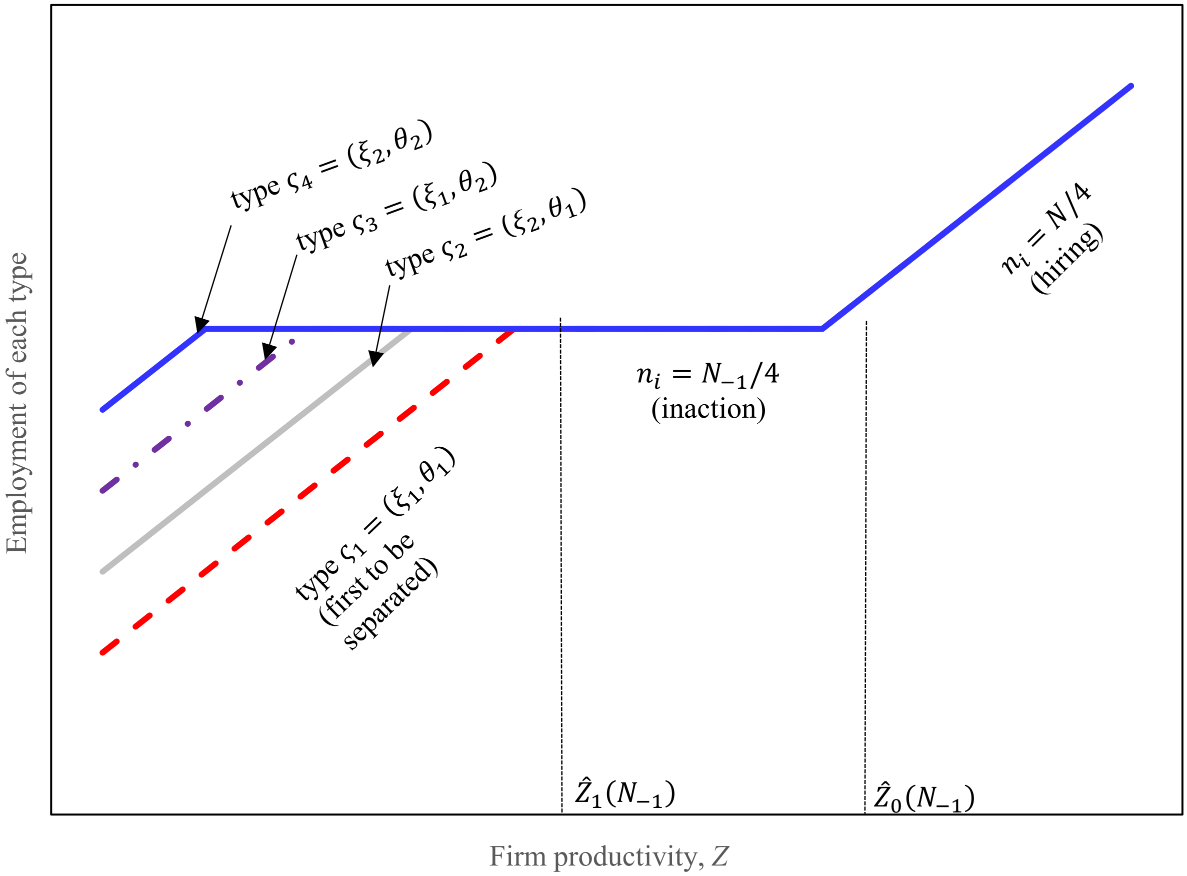

NOTE: This summarizes the optimal employment policy for the case of four equally likely types, e.g., the share λ_i of any type i equals 1/4. For high Z, employment of all four types is increased and, since λ_i=1/4 for each i, employment of each type, n_i, equals 1/4 of firm-wide employment, N. For a middling range of Zs, the firm does not adjust employment of any type, hence, $n_i = N_{-1}/4$. Separations are carried out at low Z such that, if the firm separates from type ς_i, it will continue to separate from this type if it also separates from type ς_j, $j>i$.