LABOR SUPPLY WITHIN THE FIRM*

Michele Battisti   Ryan Michaels   Choonsung Park

November 22, 2017

Abstract

Estimates of labor supply elasticities can be sensitive to the source of identifying variation. This paper's model of production complementarities helps to interpret conflicting evidence. Complementarities attenuate working time adjustments to idiosyncratic, or individual-specific, variation in work incentives. Complementarities do not restrict, however, responses to firm-wide shocks; the latter is mediated by preference parameters. Estimating the model using matched firm-worker data, the paper disentangles production from preference parameters. The Frisch elasticity along the intensive margin is found to be almost 0.5. A quasi-experimental approach, using idiosyncratic variation in work incentives, would instead find an elasticity less than half this.

JEL Codes: J22, J23, J31.

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

*Battisti: Ifo Institute, Leibniz Institute for Economic Research at the University of Munich; Michaels: Federal Reserve Bank of Philadelphia; Park: Korea Institute of Finance. We thank Professor Tattara and the team of researchers at the University of Venice for granting access to the Veneto data. We appreciate feedback from participants at NYU's micro working group; the 2013 meetings of the Canadian Economic Association; the 2013 meetings of the Society for Economic Dynamics (SED); the 2014 North American and European meetings of the Econometric Society; the 2014 European Association of Labour Economists Conference; the 2015 Society of Labor Economists (SOLE) meetings; and seminar participants at Sogang University (Seoul) and the University of Rochester. We are also thankful for comments and encouragement from Yongsung Chang. Battisti gratefully acknowledges financial support by the Leibniz Association (SAW-2012-ifo-3). The views expressed here do not necessarily reflect the views of 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 focused on the extensive margin. But variation along the intensive margin is 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, 2012). At the plant level, data on U.S. manufacturers show that the variance of changes in working time per person is equal to that of employment growth (Cooper et al, 2015).[1]

This evidence on intensive-margin fluctuations appears at odds with implications of the earlier labor supply literature. The data in Cooper et al (2015) imply that a one standard-deviation movement in hours amounts to 96 hours per quarter.[2] Yet Hall (1999) notes that estimates of the Frisch labor supply elasticity (for men) are centered around 0.2 and are often nearly zero. If the latter were right, Hall argues, the deadweight burden of these plant-level hours fluctuations would seem to be implausibly high.

In this paper, we consider a framework that can reconcile this seemingly contradictory evidence on the intensive margin. In this setting, workers are complements in production but have heterogeneous preferences over leisure. Complementarities have important implications for the identification of the intertemporal (Frisch) elasticity of substitution in (intensive-margin) labor supply. For instance, 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 employer-employee matched 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

[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, depending on the detrending procedure (Cacciatore, 2017).

[2] Cooper et al report that the standard deviation of the log change in quarterly hours was 0.18 in their sample period (1972-80). Annual hours per production worker averaged 1,952 during this time, according to the NBER-CES manufacturing database. Thus, starting from a quarterly level of $488 \cong 1952/4$ hours, an 18 log-point increase amounts to a change of 96 hours.

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 supply.[3]

In this paper's model, the firm does have a starring role. The firm and its workers join in long-term employment relationships, bound together by the fact that extensive-margin adjustments are costly. Working time is bargained jointly 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 disperse 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 in Sections 2 and 3 a unique source of panel data. We use a matched worker-firm dataset that tracks the universe of workers and firms in the northern Italian region of Veneto from 1982 to 2001.[4] 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. However, we show that in Italian household survey data, 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. As a result, 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

[3]Of course, there may be some jobs (e.g., taxi driver) that align with what is envisioned by canonical intertemporal labor supply theory, in which the worker has substantial discretion over his schedule (see Farber, 2005).

[4]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.

firms). If the ratio of the former to the latter is small, idiosyncratic variation is 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 (firm-wide) working time of 0.455. This suggests more willingness to substitute effort intertemporally than found in the earlier, seminal life-cycle literature (see MaCurdy, 1981; Browning et al., 1985; Altonji, 1986). It is, however, more in line with recent results 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.

To highlight the implications of our results for empirical analysis, we simulate a simple policy intervention in Section 5. A fraction of a firm's workforce receives the "treatment"—a shift in their own labor supply incentives—but the remainder of the firm's workers do not. We contrast the outcome with the case in which all workers participate in the intervention. Reflecting the role of complementarities, working time declines by 50 to 115 percent more when all employees receive the treatment (depending on how the extensive margin adjusts). Furthermore, if we use the treatment effect in the case where only a fraction of the workforce participates to infer the Frisch elasticity, the implied elasticity is less than half the estimate (0.455) we uncover.

This experiment illustrates that the response of working time to an idiosyncratic event may bear little resemblance to the underlying preference parameter. This is a simple, but important, point, because many influential studies utilized this latter variation. Hall (1999) notes, for instance, that the tepid response of working time in the randomized control trials known as the Negative Income Tax (NIT) experiments greatly informed the consensus on labor supply. Yet this kind of variation—a sample of workers is selected to receive a cash grant—is clearly idiosyncratic to the worker.[5] The same point applies to the seminal life-cycle analyses of MaCurdy (1981) and Altonji (1986), which identify the Frisch elasticity off the response of time worked to an individual's own (predictable) wage changes.

A few recent contributions touch on a number of themes presented here. 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

[5] 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. We discuss the NITs again in Section 5, since the size of the simulated policy intervention is based on the typical NIT.

and a distinct identification strategy, but like these authors, we conclude that idiosyncratic variation in the return to working will typically fail to recover the true willingness of workers to vary their labor input.[6] Chetty (2012) offers another approach to inference, which uses estimated elasticities to bound preference parameters even when the source of the wedge between elasticities and parameters is not explicit. Our approach is complementary: we formalize a specific reason why reduced-form estimates may not identify preference parameters and use this model to recover the parameters. Finally, Rogerson (2011) also argues that coordination may break the link between estimated elasticities and structural parameters, but studies an aggregative model in which workers coordinate their leisure.

The paper proceeds as follows. Section 1 introduces a dynamic labor demand model in which a firm and its worker bargain over working time and wages. 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 empirical work on the intensive margin. Section 6 examines the robustness of our results along several dimensions. We devote special attention to the implications of mismeasuring working time, reflecting the lack of information on daily hours in our Veneto data. Section 7 concludes.

1 Theory

1.1 The setting

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

Preferences. A worker's utility is separable in consumption and leisure. The disutility from time worked h is given by

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

where φ governs the worker's willingness to vary working time and ξ indexes the "distaste" for work. More generally, ξ encompasses any shift in the worker's marginal value of time. For instance, a rise in ξ can capture the case in which a worker is needed at home to temporarily look after a family member.

[6] Our theoretical framework also differs from that in Chetty et al (2011), who assume firms post a single work schedule for all employees. This approach echoes Deardorff and Stafford (1976) and Dickens and Lundberg (1993). Our model leaves room for idiosyncratic factors, so we can accommodate the observed working time changes within the firm.

For tractability, we make several simplifying assumptions concerning ξ. Each worker draws a value of ξ at the start of a period 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 each 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 ξ is unknown to the firm at the time of hire but perfectly observed thereafter. Accordingly, firm and worker can contract (earnings and working time) on ξ. This is plausible if the two anticipate a long-term arrangement that supports the (credible) communication of private information.[8]

In general, shifts in ξ would impinge on consumption. To avoid this complication, we appeal to the now-familiar notion of a "large" family (Merz, 1995; Hall, 2009). Specifically, we assume each individual belongs to one of many large families, each of which deploys members to the (national) labor market where they are randomly matched with firms. A family then pools its workers' earnings, thereby insuring members' consumption against risk that is idiosyncratic to the member (i.e., ξ). As a result, the flow value of working can be written without regard for the degree of risk aversion; it depends only on earnings and the disamenity of supplying labor, $\xi \nu(h)$ (Trigari, 2006).[9]

Production structure. A firm's output is an aggregate over a continuum of jobs, which are (potentially) complements in the production of a final good. Formally, $\gamma(i)$ is output of job i, and

$$\Gamma = \bar{K} Z \left(\int_0^1 \gamma(i)^\rho \, \mathrm{d}i \right)^{\alpha/\rho}, \tag{2}$$

is final output, or revenue, where Z is an index of firm-wide profitability; $\bar{K} \equiv K^{(1-\rho)\alpha/\rho}$ is a normalizing constant that simplifies the algebra to follow; $\alpha \in (0,1)$ is the returns to scale at the firm level; and $\rho \in (-\infty, \alpha)$ determines the elasticity of substitution across jobs, given by $1/(1-\rho)$. Note that under decreasing returns ($\alpha < 1$), the limiting case of perfect substitutes corresponds to $\rho = \alpha$.

Under certain simplifying assumptions, (2) takes a more tractable form. Assuming output $\gamma(i)$ of a job i is proportional to total man-hours on that job and supposing that no worker has a comparative advantage in any one job, one can motivate a simple allocation in which an equal measure $k \equiv 1/K$ of (non-overlapping) jobs is assigned to each type.[10] Accordingly,

[7] We discuss later the possible implications of persistence in ξ.

[8] We will show that earnings are increasing in, and working time decreasing in, ξ, suggesting that low-ξ workers may want to mimic high-ξ workers. Since type is i.i.d., however, reporting a high ξ period after period would suggest that one is not being truthful, enabling the firm to root out bad behavior.

[9] The Bellman equation describing the value of working is reported below—see equation (5).

[10] The Online Theory Appendix compares the solution of the assignment problem to this rule of thumb.

total man-hours on any job i assigned to type ξ are $n_\xi h_\xi / k$, where n_ξ is the measure of workers of type ξ and h_ξ is time input per worker.[11] Equation (2) becomes

$$\Gamma = G(\mathbf{h}, \mathbf{n}, Z; \boldsymbol{\theta}) = Z\left(\sum_{\xi \in \mathcal{X}} (n_\xi h_\xi)^\rho\right)^{\alpha/\rho}, \tag{3}$$

where $\mathbf{n} \equiv \{n_\xi\}$ and $\mathbf{h} \equiv \{h_\xi\}$ are column vectors of, respectively, type-specific employment and working time. Note that for $\rho < \alpha$, time inputs of different types are q-complements, i.e., the marginal product of a type is increasing in the input of any other type.

In our empirical application, we work with a more general version of (3). 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, $\varsigma \equiv (\xi, \theta)$. Again supposing an equal measure $m \equiv 1/M$ of jobs is assigned to each type, equation (3) generalizes to

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

Labor market frictions. Labor market frictions mediate the formation of employment relationships. Following Roys (2016), there is a matching friction that operates at an aggregate level, that is, the pace of job finding (and, job filling) is mediated by aggregate conditions. Since we analyze a firm's problem in the aggregate steady state, we do not elaborate further on matching. There are also employment adjustment costs, which take the form of a per-capita cost of hiring, $\bar{c}$, and firing, $\underline{c}$.

Labor market frictions play a subtle but crucial role in our analysis. Our empirical strategy is predicated on the idea that complementarities can leave very different imprints on the adjustment of earnings as opposed to working time. Underpinning the role of complementarities are labor market frictions. To see this, suppose workers share the same θ and are perfectly mobile across firms (e.g., there is no matching friction). A law of one wage must then prevail, notwithstanding the presence of complementarities. As a result, changes in earnings will merely mimic movements in working time.[12]

Our analysis of complementarities thus rests on a foundation laid by models of frictional labor markets. Matching frictions pervade labor market analysis (Rogerson and Shimer, 2010), and firing costs, such as mandated severance, are common in European labor markets,

[11] At this point, h_ξ is interpreted as average time worked among workers of type ξ. In the proof of Proposition 1, we establish that every worker of a given type will supply the same time.

[12] We are grateful to the Editor for urging us to address this point.

which is the context of our empirical application (Bentolila and Bertola, 1990).

1.2 Characterization

This section reports the dynamic problems faced by workers and firms and characterizes the choices of working time, earnings, and employment. Note that the timing of events is such that working time and earnings will be bargained between firm and worker *after* employment has been decided. Accordingly, our notation will reflect that working time, h, and earnings, W, are conditioned on employment, $\mathbf{n}$, as well on types $\boldsymbol{\varsigma}$ and productivity Z.

1.2.1 Firm and worker objectives

Workers. Consider the present value of working as type ς at a firm of productivity Z. In the present period, the employee earns a flow return equal to earnings less the disutility of supplying labor, $W_\varsigma(\mathbf{n},Z;\boldsymbol{\varsigma}) - \xi \cdot \nu(h_\varsigma(\mathbf{n},Z;\boldsymbol{\varsigma}))$. In the following period, productivity at the worker's firm, Z', is realized, and the worker draws a type, ς'. A separation will occur if the continuation value of the match is driven down to (at least) $\mathcal{U}'$, the present value of nonemployment. Accordingly, the present value of working is given by

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

Now subtracting $\mathcal{U}$ from $\mathcal{W}_\varsigma$ and rearranging yields an expression for $\mathcal{S}_\varsigma^W \equiv \mathcal{W}_\varsigma - \mathcal{U}$, the surplus from working,

$$\mathcal{S}_\varsigma^W(\mathbf{n},Z;\boldsymbol{\varsigma}) = \begin{array}{c} W_\varsigma(\mathbf{n},Z;\boldsymbol{\varsigma}) - \xi\nu(h_\varsigma(\mathbf{n},Z;\boldsymbol{\varsigma})) - \mu \\ +\beta\sum_{\varsigma'}\lambda_{\varsigma'}\mathbb{E}\left[\max\left\{0, \mathcal{S}_{\varsigma'}^W(\mathbf{n}',Z';\boldsymbol{\varsigma}')\right\}\right], \end{array} \tag{5}$$

where $\mu \equiv \mathcal{U} - \beta\mathcal{U}'$ and $r \equiv 1-\beta$.

A few remarks on (5) are warranted. First, $\mathcal{U}$ reflects, in part, the anticipated value of a *future* job. Since type is i.i.d., the value of a future job is independent of the worker's present type, ς. Hence, $\mathcal{U}$ is not indexed by ς. Since there is no idiosyncratic component to $\mathcal{U}$, it follows that $\mathcal{U} = \mathcal{U}'$ in an aggregate stationary state. Therefore, $\mu = r\mathcal{U}$, which will be treated as a parameter to be estimated.[13] Second, the flow return from working should be written, more generally, as $W_\varsigma - (\xi/\ell)\cdot\nu(h_\varsigma)$, where $1/\ell$ is the inverse of the marginal

[13] Other forms of persistent heterogeneity—for instance, in the value of home producton—would likely render $\mathcal{U} \neq \mathcal{U}'$. We abstract from such considerations and assume throughout that $\mathcal{U}$ is fixed. For further discussion of the present value of nonemployment, see section 4.2.

value of wealth that translates utils, $\xi\nu(h_\varsigma)$, into units of the numeraire. Accordingly, the choices of working time and earnings will hinge on the *ratio*, ξ/ℓ. However, since our data do not measure wealth (or, consumption), we cannot separately identify these two elements. To proceed, we treat ℓ as an i.i.d. draw from a finite dimensional set such that ξ/ℓ satisfies the restrictions on ξ outlined in Section 1.1. We can then suppress ℓ in what follows, though we can still exploit the isomorphism between shifts in ξ and marginal utility to interpret heterogeneity in the data.[14]

The firm. At the start of a period, the firm has a workforce of measure N_{-1}.[15] After productivity, Z, is realized, the firm may choose to hire. We assume hires are anonymous, in that a new worker's type $\varsigma \equiv (\xi, \theta)$ has not been drawn at the point 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 jointly decide to separate. The number of separations of type-ς workers is defined by

$$s_\varsigma = \max\left\{0, \;\; \lambda_\varsigma \mathcal{N} - n_\varsigma\right\}, \tag{6}$$

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 then "carried into" next period). Wages and time worked will be bargained after separations (if any) are made.

It is helpful to proceed by first defining 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,

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

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

$$\tilde{\Pi}(\mathbf{n}, Z) \equiv \pi(\mathbf{h}(\mathbf{n}), \mathbf{n}, Z; \boldsymbol{\varsigma}) + \beta \int \Pi(N, Z')\, \mathrm{d}F(Z'|Z), \tag{7}$$

where $\beta \in (0,1)$ is the discount factor, F is the distribution function of productivity and Π is the continuation value. Note that Π can be written as a function of just two state 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] Card (1990) flagged changes in the marginal value of wealth as a source of variation in working time.

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

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

$$\begin{aligned} \Pi^{-}(\mathcal{N}, Z) &= \max_{\mathbf{n}} \left\{ \tilde{\Pi}(\mathbf{n}, Z) - \underline{c} \textstyle\sum_{\varsigma} s_{\varsigma} \right\} \\ &= \max_{\mathbf{n}} \left\{ \tilde{\Pi}(\mathbf{n}, Z) - \underline{c} \textstyle\sum_{\varsigma} \max\{0, \ \lambda_{\varsigma}\mathcal{N} - n_{\varsigma}\} \right\}, \end{aligned} \tag{8}$$

where we have used (6). Then, step back to the initial stage and consider the choice of hires, which brings the workforce up to a level, $\mathcal{N}$. Since hires are anonymous, the value of the firm at this stage is

$$\Pi(N_{-1}, Z) = \max_{\mathcal{N}} \left\{ \ -\bar{c} \cdot \max\{0, \mathcal{N} - N_{-1}\} + \Pi^{-}(\mathcal{N}, Z) \ \right\}. \tag{9}$$

Note that (8)-(9) allow that a firm may hire and separate workers in the same period. However, for sufficiently (and realistically) high $\underline{c}$ and $\bar{c}$, this will not happen. Intuitively, productivity would have to be very low to warrant any separations, in which case no hires will be made.[16] Thus, at firms that separate, $\mathcal{N} = N_{-1}$.

1.2.2 Working time

We assume that working time is chosen efficiently via bilateral bargaining between a firm and each of its workers. Each worker's marginal disamenity of working time is equated 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*

$$\begin{aligned} h_{\xi,\theta} &= (\alpha Z \Omega(\mathbf{n}; \boldsymbol{\varsigma}))^{\frac{1}{\varphi+1-\alpha}} \cdot \left[\theta^{\rho} n_{\xi,\theta}^{\rho-1} / \xi\right]^{\frac{1}{\varphi+1-\rho}}, \\ \textit{with} \ \ \Omega(\mathbf{n}; \boldsymbol{\varsigma}) &\equiv \left(\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \left[y^{\varphi+1} n_{x,y}^{\varphi} / x \right]^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha-\rho}{\rho}}. \end{aligned} \tag{10}$$

Equation (10) is our first indication of the role complementarities play in shaping the variation in working time. In particular, the response of a type's working time to firm productivity, Z, may be quite different than its response to idiosyncratic forces, ξ and θ.

[16] See the Appendix and the Online Theory Appendix for more.

The elasticity of working time with respect to Z is $1/(\varphi + 1 - \alpha)$. This measures the worker's willingness to vary working time in response to changes in the return to working induced by variation in Z, holding all else equal. Accordingly, this has the interpretation of a Frisch (marginal-value-of-wealth-constant) elasticity. In the limiting case of constant returns, $\alpha = 1$, this elasticity is determined solely by the preference parameter, φ. More generally, it also reflects $\alpha \in (0, 1)$, since agents want to avoid concentrating production in one period relative to another when the technology is subject to diminishing returns.

Consider, next, the reaction of working time to idiosyncratic events, ξ and θ.[17] Specifically, we can imagine reassigning a single worker to another one of the $M - 1$ types, leaving unchanged the preferences and productivities of the remaining workers.[18] Straightforward differentiation then 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 $\rho = -\infty$ (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 if $\rho = -\infty$.*

There are several aspects of the Corollary that deserve attention. First, the reaction 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, the response of working time is almost entirely suppressed if tasks are sufficiently strong complements. Though a fall in ξ reduces the marginal disamenity from working, the marginal product of an individual's effort is negligible holding fixed her colleagues' working time. Since there is almost nothing to be gained by working more, the efficient choice of working time calls for no change to be made. Critically, this invariance obtains *independently* of the preference parameter, φ, that shapes the Frisch elasticity with respect to firm-wide productivity, Z.

Interestingly, the response of working time to θ does not vanish when $\rho = -\infty$. 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 highly complementary, working time is reduced to bring the outputs of this type into line with those of other types. The extent of

[17] 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.

[18] 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 changes. 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.

the change in working time hinges on the extent of complementarities. In fact, if $\rho = -\infty$, the response of working time becomes entirely detached from φ.

Proceeding, the solution to working time enables us to concentrate $\mathbf{h}$ out of the firm's problem. Substituting (10) into (4), we can now work with the revenue function,

$$\Gamma = \hat{G}(\mathbf{n},Z;\boldsymbol{\varsigma}) \equiv \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n};\boldsymbol{\varsigma})^{\frac{\alpha}{\alpha-\rho}\frac{\varphi+1-\rho}{\varphi+1-\alpha}}. \tag{11}$$

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

1.2.3 Earnings

Earnings are negotiated each period according to the Stole and Zwiebel (1996) bargain, which was generalized by Cahuc, Marque, and Wasmer (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.[19] The marginal surplus, in turn, is the sum of $\mathcal{S}^W_\varsigma(\mathbf{n},Z) \equiv \mathcal{W}_\varsigma(\mathbf{n},Z) - \mathcal{U}$ 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 concentrating $\mathbf{h}$ out of (7) and differentiating with respect to n_ς,

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

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,

$$\mathcal{W}_\varsigma(\mathbf{n},Z) - \mathcal{U} = \eta\left(\mathcal{W}_\varsigma(\mathbf{n},Z) - \mathcal{U} + \mathcal{J}_\varsigma(\mathbf{n},Z) + \underline{c}\right). \tag{12}$$

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

[19] Brügemann et al (2015) 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.

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

$$W_{\xi,\theta}\left(\mathbf{n}, Z; \boldsymbol{\varsigma}\right) = \eta \left[\kappa \frac{\partial \hat{G}\left(\mathbf{n}, Z; \boldsymbol{\varsigma}\right)}{\partial n_{\xi,\theta}} + r\underline{c} \right] + (1-\eta)\left(\kappa\xi\nu\left(h_{\xi,\theta}\left(\mathbf{n}\right)\right) + \mu\right), \tag{13}$$

where $\kappa \equiv \frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha} \geq 1$, $\mu \equiv r\mathcal{U}$, *and* $h_{\xi,\theta}\left(\mathbf{n}\right)$ *satisfies (10).*

The structure of (13) is intuitive. The bargain is a weighted average of the worker's contribution to the firm and his outside option. The former consists of the worker's productivity plus the annuitized firing cost, $r\underline{c}$, which the worker "saves" the firm by continuing the match.[20] The outside option includes the utility, $\xi\nu\left(h_{\xi,\theta}\left(\mathbf{n}\right)\right)$, that can be recovered by quitting to non-employment and μ, the annuity, or flow, value of non-employment.

Interestingly, (13) shares features with the solutions of collective bargaining games. Taschereau-Dumouchel (2015) shows that the Nash bargaining solution between a firm and its unionized workforce sets a wage for each worker that, like (13), depends on labor productivity and the worker's outside option. Unlike in (13), though, the union-negotiated wage depends only on average, not marginal, product.[21]

1.2.4 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 (10) and (13). To see this, it is helpful to first write out (13) more explicitly using (10) and (11),

$$W_{\xi,\theta}\left(\mathbf{n}, Z; \boldsymbol{\varsigma}\right) = \varkappa\left(\alpha Z \Omega\left(\mathbf{n}; \boldsymbol{\varsigma}\right)\right)^{\frac{\varphi+1}{\varphi+1-\alpha}} \left[\theta^{\varphi+1}/\xi\right]^{\frac{\rho}{\varphi+1-\rho}} n_{\xi,\theta}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} + \omega, \tag{14}$$

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 (14) 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 suppress

[20] 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.

[21] Also unlike in (13), the heterogeneity of outside options in Taschereau-Dumouchel reflects persistent differences in workers' productivities. In our context, differences in worker productivities that persist across employers would render the problem much less tractable.

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 (14). 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, $|\partial \ln W_{\xi,\theta}/\partial \ln \xi|$, *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 but holding Z fixed, we are perturbing earnings and working time *within* a given 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 substitutibility, 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. First, 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 ω. The latter, a weighted average of $r\underline{c}$ and μ, will in fact be dominated by the outside option, μ. 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.[22] We show later that variation in earnings *at the firm level*—that is, changes in average earnings that the model interprets as being due to changes in Z—offers a lever for separately identifying η.

Second, Corollary 2 refers only to a perturbation to ξ. As we saw above (Corollary 1), the reaction of working time to a change in θ is not muted even when $\rho = -\infty$, 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

[22] One may easily verify that a lower η also reduces $\varkappa$, the weight on the first term in (14).

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* $(1-\omega/W_{\xi,\theta})(\varphi+1) > 1$.

The correlation between working time and wage rates is indeed negative in our data (see Section 4). 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.[23]

1.2.5 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.[24] Thus, the observed responses of working time are intertwined with the firm's dynamic employment demand.

To shed light on the optimal labor demand policy, consider the problem of a firm of size $\mathcal{N} = N_{-1}$ (it did not hire in the first stage). We ask if this firm should separate from workers of (arbitrary) 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,

$$\frac{\partial \hat{\pi}\left(\boldsymbol{\lambda} N_{-1}, Z; \boldsymbol{\varsigma}\right)}{\partial n_{\varsigma}} + \beta \int \Pi_1\left(N_{-1}, Z'\right) \mathrm{d}F\left(Z'|Z\right) < -\underline{c}, \tag{15}$$

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 in its arguments, which implies that the marginal value of labor, the left-hand side of (15), 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

[23] Abowd and Card (1987) note that, if labor supply is chosen taking the wage as given, earnings are more volatile than working time if the sole driving force is productivity. But this, and other working time protocols lacking any notion of complementarities, will fail to replicate the relative volatility of earnings growth if there are supply shifts. See Online Theory Appendix.

[24] Consider a firm with workforce N. Recalling (10) 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.

only if) Z falls below $\hat{Z}_\varsigma(N_{-1})$. The type of worker separated first is the type ς for which $\hat{Z}_\varsigma(N_{-1})$ 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 ξ 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 $\left(\hat{Z}_1(N_{-1}),\ \hat{Z}_2(N_{-1}),...\right)$, with the latter listed in decreasing order, such that workers of all types $(\varsigma_1,...,\varsigma_i)$ are separated if and only if $Z < \hat{Z}_i(N_{-1})$.*

In certain cases, we can say more about the mapping from type to threshold. Take the special case in which $\xi \in \mathcal{X} \subseteq \mathbb{R}^K$ is the only source of heterogeneity across workers, and suppose $\lambda_\xi = 1/K$ 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. Thus, its choice is to raise firm-wide employment, and each type's size will be increased in proportion to its share in the population. Thus, starting from a firm size of N_{-1}, the firm assesses whether the marginal value of increasing employment above N_{-1} exceeds the marginal cost of hiring, $\bar{c}$. This obtains if

$$\frac{\partial \hat{\pi}(\boldsymbol{\lambda} N_{-1}, Z; \boldsymbol{\varsigma})}{\partial N_{-1}} + \beta \int \Pi_1(N_{-1}, Z')\, \mathrm{d}F(Z'|Z) > \bar{c},$$

where each cohort size is evaluated at $\mathbf{n} \equiv \boldsymbol{\lambda} N_{-1}$. Again by the supermodularity of the problem, the firm will hire if Z exceeds a certain threshold, denoted by $\hat{Z}_0(N_{-1})$. In principle, the firm could both separate and hire. However, as we argue in the Appendix, one can guarantee that $\hat{Z}_1(N_{-1}) < \hat{Z}_0(N_{-1})$ if $\bar{c} + \underline{c}$ is sufficiently (and realistically) large. For values of Z between $\hat{Z}_1(N_{-1})$ and $\hat{Z}_0(N_{-1})$, the firm does not adjust, e.g., $N = N_{-1}$.

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.

2 Taking model to data

This section begins laying the groundwork for taking our model to microeconomic data on earnings, working time, and employment in Veneto, Italy. Located in the North East, Veneto is one of the largest and richest of Italy's 20 administrative regions.[25] In this section, we introduce the data, 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 Veneto Worker History (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 a unique feature of the VWH and 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 daily wage is around 120 Euros, and on average the number of paid months per worker (per year) is 10.

[25] 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.

Table 2 zeroes in on moments of the distribution of annual changes in paid work days.[26] 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.[27] Moreover, conditional on changing days, the typical size of the change is between 10 and 19, depending on whether some of the largest adjustments are included.

The VWH's measure of paid work days, as valuable as it is, is an incomplete recording of total working time. First, paid days does not equate to days at work; for instance, the former will include paid leaves of absence such as vacation. If the paid time off is taken each year, though, we will measure *changes* in working time correctly. Other forms of leave, such as maternity leave, are only partially compensated, and so will be manifest as changes in working time in our data (see Ray (2008)).[28]

More importantly, the VWH does not capture variation in daily hours. We defer a more extended discussion of this matter until we have presented our main results. At that point, we assess in detail the extent to which our inference may be sensitive to this omission, and conclude that the "bottom line" of our results survives intact.[29]

2.2 Institutional context

Our proposed use of Veneto data requires a brief digression on the institutions of Italian labor markets more generally. 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 so in the relatively high-income region of 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 typically do not 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-

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

[27] We could replicate this degree of inaction ($\Delta h = 0$) by introducing costs of adjusting working time, but this renders the problem less tractable while potentially having little bearing on other structural parameters. Alternatively, one may interpret this inaction as indicative of overhead labor, as we discuss in Section 6.

[28] Later, we will re-run the empirical analysis for the sub-sample of men.

[29] 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 very 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).

half, and the average premia (over industry minima) is about 25 percent (Card et al, 2014).[30] Finally, management can award bonuses to individual workers independent of any union agreement (Dell'Aringa, 1994; Erickson and Ichino, 1995). Among smaller 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 highly heterogeneous, as illustrated by Brusco (1982) in his study of small industrial firms in the North Central region of Emilia-Romagna.

National unions also negotiate weekly and annual hours limits. During the 1980s, working time restrictions—specifically, limits on overtime—were either explicitly eased in union agreements or loosely enforced (Treu et al, 1993; Lodovici, 2000).[31] According to the Bank of Italy's Survey of Household Income and Wealth (SHIW), 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 did recede somewhat during the 1990s, with annual overtime hours (among those working overtime) declining to 180 by 2000.[32] The latter decline likely reflects a mix of union-negotiated limits as well as a reduction in the demand for overtime prompted by the removal of various barriers to part-time and temp work.[33]

3 Estimation Strategy

We will estimate the model of Section 1 by the method of simulated moments (MSM): we specify a list of moments, and select values for the parameters in order to minimize the (weighted) distance between the empirical and model-generated moments. One advantage of MSM in our application is its relatively minimal data requirements. To illustrate, recall that the comovement of Z and average working time at the firm is highly informative of φ. Unfortunately, our dataset does not report firm TFP or revenue per worker, the most obvious proxies for Z. But our data does include other variables, such as employment, whose volatility is informative about the variance of Z. If we can infer the latter, the variance of firm-wide working time then provides substantial identifying information about φ. MSM enables us to harness this information.

[30] 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.

[31] 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).

[32] 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.

[33] 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).

3.1 Empirical moments

There are two considerations that guide the choice of moments used in estimation. First, we want to distinguish firm-wide (i.e., Z) from the idiosyncratic (i.e., ξ or θ) components of working time and earnings. Using our matched employer-employee data, we can do this using simple least squares regressions. Second, our moments relate to *changes* in working time and earnings, rather than to their levels. To see why, suppose there are fixed differences in productivity across workers. This dispersion can support a non-degenerate distribution in time worked even under perfect complements, since time worked would be set to equate efficiency units across employees. Yet changes in working time would be synchronized, more clearly conveying the extent of complementarities.

3.1.1 Earnings and working time

We begin by developing the moments summarizing earnings and working time changes.

Regression 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 i in firm j in year t, we estimate

$$\Delta \ln h_{ijt} = \boldsymbol{\chi}_{ijt}' \mathbf{C}^h + \phi_{jt}^h + \epsilon_{ijt}^h, \tag{16}$$

where $\boldsymbol{\chi}_{ijt}$ collects the (time-varying) worker characteristics in our data, $\mathbf{C}^h$ is a conformable vector of coefficients, and ϕ_{jt}^h is a firm-year effect. Equation (16) is applied to a sub-sample of workers who stay at a 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).[34] These controls help purge the data of observable persistent heterogeneity in work schedules. The variation then captured in ϕ_{jt}^h and ϵ_{ijt}^h is what is used to estimate the structural model.

The firm-year effect, ϕ_{jt}^h, in (16) measures the log change in firm j's working time relative to the average log change among firms in year t. We interpret ϕ_{jt}^h as reflective of shocks to labor demand at the *firm level*, i.e., changes in Z. Accordingly, the variance of ϕ_{jt}^h is our measure of fluctuations in firm-wide working time, and should be highly informative as to

[34] 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.

the value of φ.

It follows that the residual in (16) isolates variation across workers *within a firm.* We pool the estimated ϵ^hs and calculate var $\left(\epsilon^h_{ijt}\right)$, which we interpret as the variance of idiosyncratic (worker-specific) working time changes.

We can repeat this exercise by replacing $\Delta \ln h$ in (16) with the log change in earnings,

$$\Delta \ln W_{ijt} = \boldsymbol{\chi}'_{ijt}\mathbf{C}^W + \phi^W_{jt} + \epsilon^W_{ijt}. \tag{17}$$

The moment, var $\left(\epsilon^W_{ijt}\right)$ /var $\left(\epsilon^h_{ijt}\right)$, compares the variances of earnings and working time changes within the firm. From our model's perspective, a high degree of complementarities means that idiosyncratic variation in preferences (or productivity) reflected in var $\left(\epsilon^W_{ijt}\right)$ is not passed through to var $\left(\epsilon^h_{ijt}\right)$. Accordingly, the ratio of these two communicates the extent to which working time adjustments are compressed by complementarities, and, thus, provides critical identifying information for ρ.[35]

Sample selection. To estimate (16)-(17), we use a sample of workers attached to their firms for consecutive years. By confining the sample to stayers, we isolate intensive-margin adjustments, i.e., changes in working time within an employment relationship.

To construct our baseline sample, we first identify workers in the year-t cross section who were paid for at least one day in all months of the first (calendar) quarter of year $t-1$ and in all months of the last quarter of year t.[36] We then remove workers employed at firms with only one employee; it would be awkward to discuss complementarities with these firms in the sample. Though such firms make up a substantial share of the population of firms, the number of workers involved is small; we still have well over 11 million observations.

We refer to the workers in our baseline sample as *2-year stayers.* They appear to have relatively strong attachments to their firms insofar their annual absences from their employers are 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.

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

[35] 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.

[36] 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.

The restriction to stayers reduces the sample notably. This seems consistent with data on worker flows in Italy. Contini et al (2009) estimates that in relatively large Italian firms (with at least 20 workers), 36 percent of a firm's workforce exits over two years.[37] Since turnover is lower at larger firms (Idson, 1993), we are not surprised that, if we restrict to 2-year stayers, we drop about half of the sample.

Our selection of stayers can raise concerns insofar as stayers are systematically different than the average worker in ways that we have not modeled. We return to discuss this point at some length in our sensitivity analysis in Section 6.

Estimates. Table 3 summarizes several key moments of the data. The first three rows pertain to within-firm (idiosyncratic) variation. Specifically, the first row reports $\text{var}\left(\epsilon_{ijt}^{W}\right)$; the second shows $\text{var}\left(\epsilon_{ijt}^{h}\right)$; 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 volatile as idiosyncratic working time changes. The next three rows report the counterparts to these moments at the firm level, namely $\text{var}\left(\phi_{jt}^{W}\right)$, $\text{var}\left(\phi_{jt}^{h}\right)$, and the ratio of the two. We note the value of $\text{var}\left(\phi_{jt}^{h}\right)$ in particular (for 2-year stayers). This variation represents 1.5-2 days per month for the typical worker.[38]

Comparing estimates in Table 3 across 2-year and 12/12 stayers reveals clear, but intuitive, differences. For instance, idiosyncratic working-time fluctuations, as captured by $\text{var}\left(\epsilon_{ijt}^{h}\right)$, are larger among the 2-year stayers, which is not surprising: they include employees who can experience longer non-working spells in years $t-1$ or t. Some of this variation in working time fluctuations is also likely reflected in the greater variance of earnings changes. The firm-wide moments are more similar. Finally, we stress that, using either sample, $\text{var}\left(\epsilon_{ijt}^{W}\right)$ substantially exceeds $\text{var}\left(\epsilon_{ijt}^{h}\right)$.

Table 4 reports on sensitivity analysis with respect to the moment, $\text{var}\left(\epsilon_{ijt}^{W}\right)/\text{var}\left(\epsilon_{ijt}^{h}\right)$, which is especially critical to our strategy. This ratio is typically at least 2, and is higher at larger firms that have a freer hand in adjusting wages since union-bargainined minima typically do not bind (Guiso, Pistaferri, and Schivardi 2005). The ratio (for 2-year stayers in particular) rises only slightly if we restrict the sample to men. Also, with the exception of transportation and communication, the ratio also does not vary by much across sector, despite differences in the make-up of industries (i.e., the prevalence of public enterprises in the health and education sectors).[39]

[37] Contini et al estimate that in the 1990s, the separation rate per year was about 20 percent. Therefore, among workers at the start of year $t-1$, $1-(1-0.2)^2 = 36$ percent exit by the end of year t.

[38] The table indicates that a one standard deviation increase in log days is $\text{var}\left(\phi_{jt}^{h}\right)^{1/2} = 0.078$. Since the typical worker puts in about 23 days per month (Table 1), a 7.8 log point increase represents 1.8 days.

[39] To estimate $\text{var}\left(\epsilon_{ijt}^{W}\right)/\text{var}\left(\epsilon_{ijt}^{h}\right)$ for men, we use all firms (and workers) in (16)-(17) but pool ϵ_{ijt}^{W} and

3.1.2 Additional moments

The list of all seven moments that we use in estimation is given in Table 5. The first four refer to results just described. We now summarize the final three, and discuss their information content for the structural parameters.

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.[40] Our finding of a negative association between the two echoes earlier studies including Abowd and Card (1989), who recovered 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.[41]

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. Our regression results using Veneto data thus reaffirm that this approach fails to find any clear evidence of a significantly positive Frisch elasticity.[42]

The final two moments refer to employment. The first is the standard deviation of employment 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, exclusive of single employer firms.

3.2 Identification

Seven parameters are estimated. They are ρ, which governs the elasticity of substitution across jobs, $1/(1-\rho)$; φ, which is a critical input into Frisch elasticities of working time; 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. The moments we aim to reproduce are derived from the sample of 2-year stayers (Table 5).

We now offer some intuition for how the moments identify the parameters. The extent

ϵ_{ijt}^h across only male workers.

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

[41] One distinction between these earlier studies and ours is that we observe daily earnings rather than the hourly wage. We return to this point in Section 6.

[42] 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.

of complementarities influences the dispersion of working time changes within the firm relative to the dispersion in earnings changes (inside the firm). Hence, ρ maps most clearly to $\text{var}\left(\epsilon_{ijt}^{W}\right)/\text{var}\left(\epsilon_{ijt}^{h}\right)$. Second, worker bargaining power, η, helps mediate the reaction of earnings to changes in working time. It thus influences the relative variances of these objects. Since ρ bears most directly on $\text{var}\left(\epsilon_{ijt}^{W}\right)/\text{var}\left(\epsilon_{ijt}^{h}\right)$, bargaining power η is especially informed by the ratio of firm-wide variances, $\text{var}\left(\phi_{jt}^{W}\right)/\text{var}\left(\phi_{jt}^{h}\right)$.

Next, the variances of idiosyncratic preference (ξ) and productivity (θ) are informed in particular by two moments. The size of preference (supply) shocks *relative* to productivity (demand) shocks influences the comovement of working time and the wage, as reflected by the regression of the former on the latter. Negative comovement suggests, for instance, the prominence of "supply-side" idiosyncratic variation (i.e., ξ), which drives working time and wages in opposite directions. In addition, the size of idiosyncratic (worker-specific) movements in working time, as reflected in $\text{var}\left(\epsilon_{ijt}^{h}\right)$, offers further information about the variances of these idiosyncratic shocks.

The final three parameters are φ, σ_Z, and μ. As foreshadowed in Section 1.2, φ influences the amplitude of working time fluctuations at the firm level, conditional on the size of firm-wide shocks, Z. This helps target $\text{var}\left(\phi_{jt}^{h}\right)$. The variance of these latter shocks are, in turn, greatly informed by the dispersion in employment growth across firms, $\Delta \ln N$. Lastly, the outside option, μ, is a critical determinant of the incentive to form new matches: if μ is large, the rents from the match are small, and so fewer hires are made. This indicates that the average size of firms, $\mathbb{E}[N]$, will help pin down μ.

4 Model Estimation

4.1 Preliminaries

To begin, we pre-set values of several parameters. We start with the firm productivity process. Since we lack revenue data, we are inclined to "tie our hands" parameterize this based on results in the firm dynamics literature. However, the parameterization has important implications for moments in our data that we do wish to replicate. Our strategy, then, is to preset some parameters and estimate others. Specifically, we assume firm productivity, Z, follows a geometric AR(1),

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

and fix $\zeta = 0.8$ based on plant-level estimates of total factor productivity (TFP) (see, e.g., Foster et al 2008). But, we treat the standard deviation, σ_Z, as a parameter to be estimated, as discussed below.

Next, we set values for four other parameters for which there is credible external information. Our choice of the severance cost, $\underline{c}$, amounts to seven months of earnings. This is a synthesis of multiple separation costs in Italy (see the Online Data Appendix for calculations). We set the hiring cost, $\bar{c}$, at 5 percent of annual earnings, based on the range of measurements in the literature.[43] Third, we fix $\alpha = 0.667$, which, as we discuss below, is consistent with labor's share in Italy as well as structural estimates off plant-level data (Cooper, Haltiwanger, and Willis, 2015). Lastly, we set the discount factor $\beta = 0.941$, which is consistent with the average annual real rate of interest in Italy over our sample.

Regarding idiosyncratic preferences, ξ, and productivities, θ, it seems heroic to try to identify the shape of their distributions given the limitations of our data. We instead assume at the outset that each is an independent uniform random variable with a mean of one. The upper and lower bounds of ξ and θ are then implied immediately by the variances, σ_ξ^2 and σ_θ^2, respectively.

To estimate the model, then, we conjecture values of σ_ξ^2 and σ_θ^2 (and other parameters) and discretize the distributions of ξ and θ in order to numerically solve the firm's problem. In light of computational constraints, we set $K = 3$ preferences (ξ) and $L = 3$ productivities (θ). This yields $M = 9$ types of $\varsigma \equiv (\xi, \theta)$, with each cohort equally represented in the population (i.e., $\lambda_{\xi,\theta} = 1/9$ for each (ξ, θ)). Once we have solved the firm's problem, we can simulate earnings, employment, and working time outcomes within each of 10,000 firms and compute the relevant moments. After comparing the model-generated moments to the data, we update our guesses for all parameters, and repeat.

4.2 Main results

Table 5 summarizes our results. The top panel lists the empirical and model-generated moments. The model replicates the moments nearly exactly. This goodness of fit should arguably be demanded from a just-identified model, but it is, still, the first test to be passed, and the model does so. The bottom panel lists MSM estimates of the structural parameters.

Frisch elasticity. Our estimate of φ implies a Frisch elasticity of firm-wide working time, $1/(\varphi + 1 - \alpha)$, of 0.455. This result suggests a greater willingness to vary working time than implied by Macurdy (1981) and the earlier life-cycle literature, more generally.

[43] Our choice is the average of estimates derived from U.S. surveys of employers (see Barron et al (1997) and Hall and Milgrom (2008)) and a survey of French firms described in Abowd and Kramarz (2003).

Recall that in a canonical life-cycle framework, optimal labor supply satisfies $\xi h^{\varphi} = \ell w$, where ℓ is the marginal value of wealth. Thus, this earlier literature sought to recover $1/\varphi$. Estimates centered around 0.15-0.2 and were often statistically indistinguishable from zero.[44] An estimate of 0.15 would imply, in our context, a Frisch elasticity of (firm-wide) working time of just 0.14.

However, our estimate looks a little low relative Pistaferri's (2003) finding of $1/\varphi = 0.7$ (which implies $1/(\varphi + 1 - \alpha) = 0.568$). Pistaferri identifies $1/\varphi$ (in a life-cycle setting) by estimating the response of total hours (in year t) to the survey respondent's *expected* earnings growth (between years $t-1$ and t). Expected earnings growth reflects, in turn, expectations about $\varsigma_t \equiv (\xi_t, \theta_t)$ and Z_t. Our model indicates that if the idiosyncratic element ς_t is transitory, and firm productivity Z_t is persistent, the expected path of earnings will be shaped by the latter. In that case, Pistaferri's estimate may reflect, like ours, the response of working time to *firm-wide* variation.

Elasticity of substitution. Our estimate of $\rho = -1.907$ implies an elasticity of substitution across jobs within the firm of $(1-\rho)^{-1} \cong 0.344$. To convey the meaning of this result in more concrete terms, we can compute 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 $\rho = \alpha$ (see Section 1), working time would adjust by $1/(\varphi + 1 - \alpha) = 0.455$ log points. Our estimate of ρ instead implies a response equal to $(\varphi + 1 - \rho)^{-1} \cong 0.209$. Thus, a worker's reaction to idiosyncratic events is less than half as large as implied by the perfect-substitutes case.

Worker bargaining power. Our estimate of $\eta = 0.452$ is not too different than Roys' (2016) estimate of 0.52, though we bring very different identifying information to bear on η. 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, Pistaferri, and Schivardi (2005). Interestingly, as we discuss below, 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.[45]

Flow outside option. The outside option is estimated to be $\mu = 0.196$. This result can be hard to interpret in the absence of an explicit theory of nonemployment (which was

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

[45] Note that, to the extent firms (in the data) smooth out earnings relative to the underlying shocks, we will mistakenly attribute such compression in earnings growth (relative to changes in working time) to a *lack of* complementarities. Thus, we will over-estimate the elasticity of substitution.

not required for estimation). To assess our finding, suppose the value of transitioning to nonemployment satisfies the Bellman equation,

$$\mathcal{U} = \upsilon + \mathsf{r}\underline{c} + \beta f \sum_{\varsigma'} \lambda_{\varsigma'} \mathbb{E}\left[\mathcal{S}^{W}_{\varsigma'}\left(\mathbf{n}',Z'\right)\right] + \beta \mathcal{U}', \tag{18}$$

where $\mathsf{r} \equiv r + \beta f$, υ is nonemployment income, and f is the probability of matching with a new employer.[46] Assuming that υ is fixed, and recalling that type is i.i.d., it follows that the expected worker's surplus is constant in an aggregate stationary state. Noting that workers receive a share η of the match surplus then yields

$$\mu \equiv r\mathcal{U} = \upsilon + \mathsf{r}\underline{c} + \beta f \frac{\eta}{1-\eta} \sum_{\varsigma'} \lambda_{\varsigma'} \mathbb{E}\left[\mathcal{J}_{\varsigma'}\left(\mathbf{n}',Z'\right)\right]. \tag{19}$$

The expected firm's surplus on the right side can be computed from the estimated model. Setting $f = 0.40$ (Elsby et al, 2013) and using our estimates of μ and η, we can then solve for υ. Dividing this by average earnings (in the model) yields a replacement rate of 49 percent. Reassuringly, though this result was not targeted in estimation, it is not too different from the replacement rate of 58 percent under Italy's unemployment insurance program.[47]

Shocks. Our estimate of σ_Z implies a standard deviation of the log change in firm productivity (i.e., $var\left(\Delta \ln Z\right)$) of 0.198. This is remarkably similar to estimates implied by plant-level TFP in European economies (see "France" and "Spain" in Table 2 of Asker, Collard-Wexler, and De Loecker, 2014). As for idiosyncratic heterogeneity, we find that it is slightly more substantial than firm-level dispersion: the unconditional standard deviation of firm productivity $\left(\text{i.e., } \sqrt{var\left(\ln Z\right)}\right)$ is 0.315, and the standard deviation of the sum of idiosyncratic disturbances is $\sqrt{0.292^2 + 0.219^2} \cong 0.365$.

5 Implications for empirical research

The presence of production complementarities implies that the labor supply response to idiosyncratic variation can yield a downwardly biased estimate of a worker's willingness to substitute effort intertemporally. We illustrate this point quantitatively in this section.

Specifically, we carry out a randomized control trial within the estimated model. A

[46] This can be derived by writing $\mathcal{U}$ as the sum of the lump-sum severance, $\underline{c}$, and a remainder, $\widetilde{\mathcal{U}}$, that can be expressed recursively, $\widetilde{\mathcal{U}} = \upsilon + \beta f \sum_{\varsigma'} \lambda_{\varsigma'} \mathbb{E}\left[\mathcal{W}_{\varsigma'}\left(\mathbf{n}',Z'\right)\right] + \beta\left(1-f\right)\widetilde{\mathcal{U}}'$. Note that the expectation of the future value of working is taken against the distribution of hiring firms.

[47] This is the rate in the first year of an unemployment spell. Benefits are offered beyond the first year to older workers at a reduced rate. See Online Data Appendix for more.

fraction of a firm's workforce is "treated" with a higher distaste, ξ, for work. We then compute the change in working time of the treated group, and compare this to the outcome if the full workforce were treated.

As noted in Section 1, increasing ξ is, in general, equivalent to reducing the marginal value of wealth, ℓ. We can use this isomorphism as a tool for calibration. Specifically, we can use canonical consumer theory to recover the change in ℓ—and thus, the change in ξ—from the behavioral response to the "treatment", which is, in this case, a lump-sum transfer.[48]

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 so that we can illustrate the implications of complementarities within a familiar context; we do not attempt to capture all of the various details of the NITs in our model.[49]

The implied transfer amounts to 37 percent of a participant's initial (pre-NIT) income.[50] To recover the change in ℓ, we assume a marginal propensity to consume out of transitory income of 1/3 (Johnson et al, 2006). Then, recalling (Section 1) that utility is separable in consumption C and effort, we can map from the change in consumption to the change in ℓ by log-differentiating the first-order condition for C to obtain $\Delta \ln C = -(1/\phi)\,\Delta \ln \ell$, where $1/\phi$ is the own-price elasticity. Setting $1/\phi = 1/2$ (Hall, 2009) implies $\Delta \ln \ell \cong -0.25$, which, is equivalent in our setting to increasing ξ by 25 percent.

This treatment is applied to one of the 9 (ξ, θ) cohorts in the firm. The model implies that these employees reduce their time worked by 5.3 percent. If we viewed this reaction through the lens of a model where workers are perfect substitutes (e.g., $\rho = \alpha$), we would mistakenly infer that $(\varphi + 1 - \alpha)^{-1} = 0.053/0.25 = 0.212$. This is *less than half* the size of the Frisch elasticity of firm-wide working time that we estimated in the prior section.

This effect can be contrasted to the change in working time when *all* workers in the firm receive the treatment. To illustrate, first suppose that the designer of the randomized trial

[48] 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 in their firms' productivities can be diversified within the neighborhood, which acts like a "large" family that insures members' consumption (see Section 1). Yet, since the transfer operates neighborhood-wide, it will alter ℓ.

[49] 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.

[50] This is the ratio of the (participation-weighted) average guarantee across NIT trials (see Burtless 1987) to the midpoint of the eligible income range (see footnote 49).

can hold employment fixed. In that case, using (10), we can compute the treatment effect as $\frac{1}{\varphi+1-\alpha} \times \Delta \ln \xi$ =11.4 percent. Thus, the reduction of working time is more than twice as large. More realistically, though, if firms can adjust on both margins of labor demand, this will take some of the burden off adjusting working time. Allowing for employment adjustments, mean working time declines by 7.8 percent. Though smaller, this is 50 percent higher than what we find if only one cohort were treated.

Bearing in mind that our exercise is only a stylized version of the NIT, we can 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, that is, working time conditional on working. Rather, the decline in hours largely reflected longer job search spells (Moffit, 1981; Robins and West, 1983). Thus, the intensive-margin response was appreciably lower and perhaps more in line with our model's predictions.[51]

6 Robustness

This section probes the robustness of our results in several respects. First, in Section 6.1, we investigate the sensitivity of our results to alternative values of the pre-set parameters and sub-samples. Section 6.2 then assesses the implications of shortcomings in our measurement of working time. Section 6.3 focuses in on threats to the identification of complementarities.

6.1 Additional estimation results

We have 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-estimate the model over a certain sub-sample. The results are reported in Tables 6 and 7. Taken together, they point to a Frisch elasticity (of firm-wide) working time between 0.259 and 0.576, and an elasticity of substitution between 0.232 and 0.460.[52]

Higher severance and less persistent productivity push many parameters in the same direction. Severance of one year's earnings compresses changes in employment, and larger firm-wide shocks are required to generate the observed variance of $\Delta \ln N$. Less persistent productivity also induces smaller adjustments in labor demand: if employment changes are

[51] On the other hand, the benefit reduction rate in the NITs likely lowered working time relative to the model (see footnote 49). 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.

[52] We do not report the model-implied moments; these match the data almost exactly.

costly to reverse, firms attenuate responses to transitory shocks. As a result, when we lower ζ to 0.32, which induces a persistence in value-added comparable to Guiso et al (2005), σ_Z must rise to recreate the variance of $\Delta \ln N$.[53] Larger firm-wide shocks, in turn, require a smaller Frisch elasticity and lower bargaining power in order to restrain movements in working time and earnings. The decline in η to 0.231 reduces the elasticity of earnings to average product to 0.37 (from 0.596 in our baseline). Still, earnings remain more responsive than in Guiso et al (2005), who estimate an elasticity closer to 0.1.[54]

Many parameters react in the opposite manner when α is raised. To arrive at our choice of $\alpha = 0.824$, we reinterpret (2) as the reduced form of a monopolistically competitive firm's revenue function where α reflects both returns to scale and the product demand elasticity, ϵ (Cooper et al, 2015). The increase in α from 0.667 (in our baseline) to 0.824 can then be shown to correspond to a doubling of ϵ from 4 (Nakamura and Steinsson, 2008) to 8.[55] This makes labor demand more elastic, which translates into a wider distribution of employment growth. Therefore, σ_Z is lowered to match the variance of $\Delta \ln N$. Working time changes appear larger in light of smaller shocks, which means the Frisch elasticity must be higher: $1/(\varphi + 1 - \alpha)$ is now 0.546.

We have also re-estimated the model over the sub-sample, 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 toward decentralizing wage setting, Table 7 shows that the variance of earnings growth both within the firm and at the firm level is more volatile than in the full sample. Other changes, relative to the full sample, include somewhat smaller fluctuations in working time, and "less negative" comovement of working time and daily earnings. The larger variance of earnings growth at the firm level drives η up to 0.569, and the smaller variance of firm-wide working time changes drives down the Frisch elasticity to 0.315. The increase in η also expands the variance of earnings growth within the firm, because it implies a higher pass through of the idiosyncratic component of marginal product. To offset this effect on $\mathrm{var}\left(\epsilon_{ijt}^{W}\right)/\mathrm{var}\left(\epsilon_{ijt}^{h}\right)$, $1/(1-\rho)$ must rise to 0.46.[56]

[53] Guiso et al's projection of log revenue on its lag yields a coefficient of 0.477. This is what we target.

[54] Guiso et al estimate the response of earnings to a permanent increase in value-added. Roys (2016) proposes a model in which earnings react less to permanent than transitory productivity shocks.

[55] Suppose a fixed measure, χ, of jobs in a firm is done by labor, and the remainder by capital. Let $\mathsf{N} \equiv \left(\int_0^{\chi} y(i)^{\rho}\, \mathrm{d}i\right)^{1/\rho}$ represent the aggregation of jobs done by labor, and assume N and $\mathsf{K} \equiv \left(\int_{\chi}^{1} y(i)^{\rho}\, \mathrm{d}i\right)^{1/\rho}$ are joined to make output, $Y = Z\mathsf{N}^{\tilde{\alpha}}\mathsf{K}^{1-\tilde{\alpha}}$. If product demand is $Y = P^{-\epsilon}$ and if K is chosen optimally, the elasticity of revenue with respect to N is $\alpha \equiv \frac{\epsilon-1}{\epsilon+(1-\tilde{\alpha})/\tilde{\alpha}}$, which is the parameter that appears in (2). Here, $\tilde{\alpha}$ is comparable to labor's share (though somewhat lower, because of bargaining). We recover $\alpha = 2/3$ if we set $\tilde{\alpha} = 2/3$, in line with OECD data, and $\epsilon = 4$.

[56] Interestingly, we estimate a lower degree of complementarity even though $\mathrm{var}\left(\epsilon_{ijt}^{W}\right)/\mathrm{var}\left(\epsilon_{ijt}^{h}\right)$ is higher.

Last, the increase in the comovement of working time and wages requires larger idiosyncratic productivity innovations.

6.2 (Mis)measuring working time

The omission of daily hours in the VWH may have important implications for several moments used in estimation. To see why, suppose workers adjust their number of days and daily hours in the same direction. The VWH data will understate the variance of changes in (firm-wide) total 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 Labor Force Survey (LFS), which asks about hours and days worked, enables us to directly measure the effect of missing hours on the variability of working time. We use LFS data between 1993, when the panel dimension of the micro data becomes available, and 2001. As in our analysis of the VWH, we restrict attention to stayers, who work for the same employer in adjacent years. Among stayers, we calculate the year-over-year change in days worked and weekly hours in the survey reference week.

The LFS indicates days worked to be an active margin of adjustment. In Veneto, 22 percent of workers adjust days per week, whereas 36 percent adjust weekly hours.[57] Thus, the frequency of days adjustment represents up to 60 percent of the incidence of hours adjustment. In fact, the most commonly observed change in weekly hours is 8, which reflects the prevalence of Saturday overtime in Italy (see also Giaccone, 2009). To take this one step further, we decompose log changes in weekly hours across log changes in days and daily hours. Since the adjustment of days involves relatively large changes in weekly hours, the overall contribution of changes in days to hours variation exceeds that implied by the frequency of days changes alone. Indeed, the variance of log changes in days accounts for 78 percent of the variance of weekly hours growth.[58]

What does this result mean for our VWH-based moments? The answer depends on how the "missing" variation in working time is distributed across idiosyncratic and firm-wide

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 $\text{var}\left(\epsilon_{ijt}^{W}\right)/\text{var}\left(\epsilon_{ijt}^{h}\right)$.

[57] The incidence of days adjustment is lower than in our Veneto data (see Table 1) since we observe only the reference week in the LFS, whereas the VWH captures any changes in days through the year.

[58] The European Working Conditions Survey (EWCS) also asks about days and daily hours, but its questions are more qualitative in nature. Our analysis of the EWCS also points to the days margin as a significant source of variation in total hours. Please see the Online Data Appendix for more.

sources. Suppose this variation is attributed in proportion to each source's contribution to the total variance in the VWH. Using the estimates from Table 5, and noting that these components are (by construction) orthogonal, we find that the idiosyncratic piece accounts for three-quarters of the total. Accordingly, we scale the total variance by (1/0.78) and distribute 75 percent of the increase to $var\left(\epsilon^h\right)$. Assuming earnings in the VHW are measured accurately,[59] the ratio of the idiosyncratic variances of earnings growth and working time changes falls to 1.75 from 2.247 in the baseline case.[60] The analogue for the ratio of firm-wide variances is 2.25, down from 2.885.

Another means of assessing the VWH, which does not require observing days worked, is to examine the variance of *hours* changes relative to the variance of earnings changes. To this end, we draw on two surveys administered by the Bank of Italy, each of which contains observations on earnings and hours. The Survey of Household Income and Wealth (SHIW) is a panel of roughly 8,000 households, and the Survey of Industrial and Service (SIS) firms is administered annually to a panel of about 4,150 non-financial companies. The findings from these two surveys can then be compared against the relative variability of *days* worked in the VWH. In the interest of space, this analysis is restricted to the Online Data Appendix. Our conclusion is that the basic message of the LFS holds up, namely, the absence of daily hours data is likely to lead us to overstate the relative variance of earnings growth by 20 to 30 percent.

What do this survey estimates imply for inference of the model's parameters? To consider this, we mark down the ratio of idiosyncratic variances to 1.75, and the ratio of firm-wide variances to 2.25. We assume the VWH measures earnings correctly, so changes to each of the latter ratios imply corresponding changes to the variances of working time adjustments. None of the other VWH moments are altered. We then search for parameter values that best fit the revised moments.

Our results are reported in Table 7. As expected, the model infers a lower degree of complementarities owing to the higher relative variability of working time. In particular, the elasticity of substitution increases to 0.437 (up from 0.344 in the baseline case). An increase in ρ, all else equal, implies less attenuation in the response of working time to idiosyncratic events (i.e., ξ). Consequently, using the latter variation to draw inferences about workers' willingness to substitute labor intertemporally is less misleading. However, the revised moments also include a higher variance of working time changes. The latter implies a higher Frisch elasticity of firm-wide working time, which increases to $(\varphi + 1 - \alpha)^{-1} = 0.53$ (up

[59] Unfortunately, the LFS does not ask about earnings during our sample period.

[60] When we reference the VWH, we again use the full sample (1982-2001).

from 0.455). As a result, the firm-wide elasticity remains *twice as large* as the individual's elasticity with respect to ξ, which is $(\varphi + 1 - \rho)^{-1}$. In this sense, the "bottom line" of our results remains largely intact after correcting for the omission of daily hours.[61]

6.3 On inferring complementarities

This subsection collects a number of concerns regarding model mis-specification. We address them through the lens of their implications for the degree of complementarities.

i.i.d. types. We assumed types were i.i.d. for tractability. To consider the implications of persistent types, first note that the persistence of ς has no direct effect on working time, since the latter is an intra-temporal choice. As for earnings, suppose in particular that ξ is persistent. The earnings bargain has the same form as (13), but μ is now indexed by ξ to reflect that a worker's outside option—the value of searching for new employment—depends on her type (see Online Theory Appendix).[62] Thus, the implications of persistence hinge on the mapping from ξ to μ.[63] If the value of searching were decreasing in ξ (which seems reasonable), earnings would respond less to changes in ξ since earnings are otherwise increasing in ξ (if $\rho < 0$). As a result, our model would "need" a *higher* degree of complementarities to match the relative dispersion of earnings growth in the data.

Selection bias. In both actual and model-generated data, we use the subsample of stayers to compute many of our moments. If the model is correctly specified, our estimates of the parameters are consistent (Smith, 1993). However, 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 a worker's separation from a firm is indicative of an *absence* of complementarity between his job and others. Then, our sample of stayers will consist of the most complementary jobs; this will confound the inference of ρ. Perhaps an argument in favor of this hypothesis is that a firm competes more aggressively to retain workers in complementary jobs. But, by this logic, a similar firm that seeks to "poach" such a worker to fill a vacancy should also compete aggressively.[64] This latter consideration

[61] More exactly, the ratio of $(\varphi + 1 - \alpha)^{-1}$ to $(\varphi + 1 - \rho)^{-1}$ was 2.17 in the baseline estimation and 2.04 in the case considered here.

[62] If θ is interpreted in this context as match-specific productivity, then it does not persist across labor market states. Hence, ξ would influence $\mathcal{U}$, but not θ.

[63] Though we did not use (19) in estimation, it can offer some guidance here. It shows that μ reflects, under surplus sharing, the hiring firms' anticipated marginal surplus. The effect on the latter of varying ξ is ambiguous: raising ξ makes one's own working time more scarce, which increases the marginal product, but also depresses the working time of others in the firm, which lowers the marginal product.

[64] This assumes the worker will perform a similar job in the new firm, and that the new firm's production

suggests that separations (where a firm poaches a worker) may correspond to jobs with a *high* degree of complementary. A priori, then, it is unclear that separation events systematically reveal the complementarity of the jobs.

Overhead labor. Another factor that might confound our inference of complementarities is the presence of overhead labor. Since the latter does not vary its days (by much), it serves to compress the distribution of days worked movements, from which our model infers that there are complementarities. The concern is that this inference masks a flexible production structure among non-overhead labor. To assess this concern, we drop workers who report 52 weeks of paid work in adjacent years and re-estimate (16)-(17) to recover the idiosyncratic components, ϵ_{ijt}^{W} and ϵ_{ijt}^{h}. This is very generous to the notion of overhead labor, as it drops any worker who participates full time in consecutive years. As anticipated, the amount of compression in the distribution of days worked movements is diminished. And yet, $\text{var}\left(\epsilon_{ijt}^{W}\right) / \text{var}\left(\epsilon_{ijt}^{h}\right)$ is 1.59—well above 1, which is strongly suggestive of a role for production complementarities (Corollary 2).

The wage in data and model. The regression of days worked on daily earnings is critical to our strategy: the negative comovement limits the scope for worker-specific productivity shocks alone to reproduce the moment, $\text{var}\left(\epsilon_{ijt}^{W}\right) / \text{var}\left(\epsilon_{ijt}^{h}\right)$, and so points to a role for complementarities. But, daily earnings conflates movements in daily hours and hourly wages. Therefore, our estimate could reflect the negative comovement of days and daily hours, rather than the comovement of time worked and remuneration per unit time. It would then be inappropriate to map the latter to its counterpart in the model.

We address this concern as follows. The least-squares coefficient from a regression of the log change in days worked on the log change in daily earnings can be decomposed as

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

We use data on days and daily hours from the Italian Labor Force Survey to fill in an estimate for the numerator in the first term in (20).[65] Using the variance of the log change in daily earnings in our Veneto data, we can then calculate the first term, which summarizes the role of daily hours in driving the comovement of daily earnings and days. We estimate this term to lie between -0.045 and -0.10.[66] Taking the midpoint of these and comparing

structure is broadly comparable to the worker's present employer.

[65] We use observations in the LFS for Veneto residents, but results hardly change if we use the full sample.

[66] 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.

to the estimate of -0.169, it seems that shifts in hourly wages do drive the majority of the comovement we are capturing in the Veneto data.

7 Conclusion

This paper has pursued the idea that an individual's labor supply is bound up with the working time choices of her colleagues within the firm. We have developed a tractable theory of earnings, working time, and employment demand that formalizes this idea. In particular, the model expresses the intuition that, if there are sufficiently strong complementaries, working time adjustments across employees inside a firm are compressed, regardless of the true 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 to explore the consequences of failing to control for complementarities in conducting inference about labor supply elasticities. We find that if one estimates 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. Yet, the mapping from the change in house price to the change in labor supply can be tenuous, depending critically on the extent of complementarities. Hence, our framework can be used in this setting to help disentangle the wealth effect from complementarities. Second, the diffusion of matched employee-employer datasets will very likely 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 theory of the production structure.

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.

Abowd, John and Francis Kramarz (2003). "The Costs of Hiring and Separations," *Labour Economics*, 10(5), p. 499-530.

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

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.

Barron, John, Mark Berger, and Dan Black (1997). *On-the-Job Training*. Kalamazoo, MI: W.E. Upjohn Foundation for Employment Research.

Bentolila, Samuel and Giuseppe Bertola (1990). "Firing Costs and Labour Demand: How Bad is Eurosclerosis?", *Review of Economic Studies*, 57(3), p. 381-402.

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, Peter Gautier, and Guido Menzio (2015). "Intrafirm bargaining and Shapley values," NBER Working Paper 21508.

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 (2017). "Hours and Employment Over the Business Cycle: A Bayesian Approach," mimeo, HEC Montreal.

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

Card, David (1994). "Intertemporal Labor Supply: An Assessment," in Christopher Sims, ed., *Advances in Econometrics*. New York: Cambridge University Press.

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.

Cella, Gian Primo and Tiziano Treu (2009). *Relazioni industriali e contrattazione collettiva.* Il Mulino.

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(2), p. 471-75.

Cingano, F. (2003). "Returns to Specific Skills in Industrial Districts," *Labour Economics,* 10(2), p. 149-164.

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.

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

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

Elsby, Michael, Bart Hobijn, and Ayşegül Sahin (2013). "Unemployment dynamics in the OECD," *Review of Economics and Statistics,* 95(2), p. 530-548.

Erickson, C. and A. Ichino (1993). "Wage Differentials in Italy: Market Forces, Institutions, and Inflation," NBER Working Paper 4922.

Farber, Henry (2005). "Is Tomorrow Another Day? The Labor Supply of New York City Cab Drivers," *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," EWCO 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*, 1999, North-Holland, p. 1137-1170.

Hall, Robert and Paul Milgrom (2008). "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98(4), p. 1653-1674.

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 (2012). "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.

Moffitt, Robert (1981). “The Negative Income Tax: Would It Discourage Work?”, *Monthly Labor Review*, April, 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.

Nakamura, Emi and Jón Steinsson (2010). “Monetary Non-Neutrality in a Multi-Sector Menu Cost Model,” *Quarterly Journal of Economics*, 125(3), p. 961-1013.

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.

Pissarides, Christopher (2000). *Equilibrium Unemployment Theory.* MIT Press.

Pistaferri, Luigi (2003). “Anticipated and Unanticipated Wage Changes, Wage Risk and Life Cycle 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 and Robert Shimer (2010). “Search in Macroeconomic Models of the Labor Market,” in Orley C. Ashenfelter and David Card (eds.) *Handbook of Labor Economics*, North-Holland, p. 619–700.

Rogerson, Richard (2011). “Individual and Aggregate Labor Supply with Coordinated Working Times,” *Journal of Money, Credit and Banking* 43(1), p. S7-S37.

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

Smith, Anthony (1993). “Estimating Nonlinear Time Series Models Using Simulated Vector Autoregressions,” *Journal of Applied Econometrics*, vol. 8, S63-S84.

Stole, Lars and Jeffrey Zwiebel (1996). “Intrafirm bargaining under non-binding contracts,” *Review of Economic Studies*, 63(3), p. 375-410.

Taschereau-Dumouchel, Mathie (2015). “The Union Threat,” mimeo, Cornell University.

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

Treu, Tiziano (2007). *Labour Law in Italy.* Kluwer Law International.

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.

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. It is helpful to begin by writing out total time input, $n_\varsigma h_\varsigma$, explicitly, noting that $h_\varsigma = \int_0^{n_\varsigma} h_\varsigma(j)\mathrm{d}j/n_\varsigma$ is average working time per member of type ς. Now differentiating (7) with respect to $h_\varsigma(j)$, individual j's working time, and equating this to his marginal value of time, $\xi h_\varsigma^\varphi(j)$, yields

$$\alpha Z \left(\sum_{x\in\mathcal{X}}\sum_{y\in\mathcal{Y}}(yn_{x,y}h_{x,y})^\rho\right)^{\frac{\alpha-\rho}{\rho}} \theta^\rho (n_\varsigma h_\varsigma)^{\rho-1} = \xi h_\varsigma^\varphi(j) \quad \forall\ j \text{ in } \varsigma \equiv (\xi,\theta). \tag{21}$$

It is evident that all workers of a given type will supply the same amount of time. Accordingly, we can eliminate the index j, setting $h_\varsigma(j) = h_\varsigma$ for all j. Combining FOCs for types (ξ,θ) and $(x,y) \neq (\xi,\theta)$ then yields

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

Using this to substitute for any $h_{x,y} \neq h_{\xi,\theta}$ in (21), and solving for $h_{\xi,\theta} \equiv h_\varsigma$, we recover the solution in the main text. ■

Proof of Corollary 1. Totally differentiating (10) 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 guarantees that $\hat{G}$ is a concave function of $\mathbf{n}$, that is, the Hessian, $\nabla^2 \hat{G}(\mathbf{n}, Z; \boldsymbol{\varsigma})$, 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 that these properties of $\hat{G}$ pass to period profit, $\hat{\pi}$. They can be verified once a solution for the wage bargain is obtained.

Conjecture 1 *The profit function, $\hat{\pi}(\mathbf{n}, Z; \boldsymbol{\varsigma})$, 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 $\xi \neq \varsigma$, if the marginal value of that cohort falls below $-\underline{c}$ given $n_\xi = \lambda_\xi N_{-1}$,

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

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 begins, 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 exists 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 1: In line with our notation from Proposition 3, we will, in what follows, 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 the second to be separated, and so on.

Remark 2: As noted in the main text, it is in principle possible for a firm to hire and subsequently separate. Under certain reasonable conditions, it will not do this, though. In the interest of space, however, the proof of this next claim is omitted here.

Lemma 2 *If $\bar{c} + \underline{c}$ is sufficiently large, then the firm will never hire if it also separates.*

Proof. See Online Theory Appendix. ■

9.3 Earnings

Proof of Proposition 2. As stated in the main text, and restated here for convenience, the marginal contribution of any type-ς worker to the firm, gross of the separation cost $\underline{c}$, is

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

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

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

The expected marginal value of labor in (25) can be decomposed using Leibniz's rule,[67]

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

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 $i \leq j$.[68] 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. If the firm hires, the Envelope theorem implies,[69]

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

To treat the case of separations, return to (8) and consider the state in which the firm separates only from type-1 labor, that is, workers of type ς_1. 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

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

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

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

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

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

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

[68] 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 $Z < \hat{Z}_M (N)$.

[69] This presumes the firm does not also fire (after types ς are drawn). We make this assumption throughout. As noted in Section 1, firms will not hire and, simultaneously, fire in the face of realistic adjustment frictions.

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

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

The marginal value of labor in the "freezing" regime, $\Pi_N^0(N, Z')$, can be obtained as follows. Forwarding (8)-(9) 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

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

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

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

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

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

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, $\mathcal{S}_{\xi,\theta}^W = (\eta/(1-\eta))\left[\mathcal{J}_{\xi,\theta} + \underline{c}\right]$, we can recast (5) in terms of the firm's surplus,

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

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., $Z' < \hat{Z}_i(N)$), 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., $Z' > \hat{Z}_0(N)$), the average marginal value of labor across types is equated to the marginal cost, $\sum_{i=1} \lambda_i \mathcal{J}_i = \bar{c}$.[70] 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

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

Substituting this into (35) and rearranging yields

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

Now inserting (34) and (37) into (12) and using (26), we have that, for a worker of type $\varsigma \equiv (\xi, \theta)$,

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

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

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

where $\kappa \equiv \frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha}$. Using (11) and the solution for working time, one can calculate period profit and confirm Conjecture 1. ■

Proof of Corollary 2. Totally differentiating the earnings bargain (14) with respect to $W_{\xi,\theta}$ and ξ yields

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

Recalling (see (22)) the response of working time to change in ξ, $\mathrm{d}\ln h_{\xi,\theta}/\mathrm{d}\ln\xi = -(\varphi + 1 - \rho)^{-1}$, one can see that

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

[70] $\mathcal{J}_i = -\underline{c}$ is the first order condition corresponding to the firing firm's problem (8). In the case the firm hires, $\sum_{i=1} \lambda_i \mathcal{J}_i = \bar{c}$ is the first order condition corresponding to the problem (9).

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 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

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

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

Table 1: Summary statistics of Veneto panel

Variable	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 number of months paid	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 (h)

Share with $\Delta h = 0$	47.38%
Share with $\|\Delta h\| > 10$	33.15%
Avg $\|\Delta h\|$ if $\Delta h \neq 0$	19.06
Avg $\|\Delta h\|$ if $\Delta h \neq 0$, excluding $\|\Delta h\|>50$	9.75

NOTE: Statistics refer to our sample of 2-year stayers, as defined in the main text (see also Note to Table 3). There are 11.810 million worker-year observations.

Table 3: Earnings and working time in Veneto panel

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

NOTE: W is annual earnings, h is working time, and w is the daily wage (W/h). The 12/12 stayers at a firm 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: Alternative estimates of $\mathrm{var}(\epsilon^W)/\mathrm{var}(\epsilon^h)$

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 & 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)
$\mathrm{var}(\epsilon^W)/\mathrm{var}(\epsilon^h)$	2.244	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
$\mathrm{cov}(\Delta \ln h, \Delta \ln w)/\mathrm{var}(\Delta \ln w)$	-0.170	-0.169
$\sqrt{\mathrm{var}(\Delta \ln N)}$	0.175	0.175
$\mathrm{E}[N]$	17.131	17.130

PANEL B

Parameter	Symbol	Value
Elasticity of substitution across tasks	$1/(1-\rho)$	0.344 [0.0006]
Frisch elasticity of working time	$1/(\psi+1-\alpha)$	0.455 [0.0008]
Worker bargaining power	η	0.452 [0.0006]
Flow return on non-employment	μ	0.196 [0.0006]
Std. dev. of idiosyncratic preference	σ_ξ	0.291 [0.0004]
Std. dev. of idiosyncratic productivity	σ_θ	0.219 [0.0008]
Std. dev. of shock to firm productivity	σ_Z	0.189 [0.0002]

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

Table 6: Sensitivity analysis, I

Parameter	I] Baseline results	II] Larger separation cost	III] Less persistent revenue	IV] Higher returns to scale
Elasticity of sub. across tasks	0.344 [0.0006]	0.330 [0.0006]	0.232 [0.0004]	0.373 [0.0005]
Frisch elasticity of working time	0.455 [0.0008]	0.369 [0.0005]	0.259 [0.0003]	0.576 [0.0009]
Worker bargaining power	0.452 [0.0006]	0.369 [0.0004]	0.231 [0.0003]	0.517 [0.0006]
Flow return on non-employment	0.196 [0.0006]	0.237 [0.0005]	0.342 [0.0007]	0.138 [0.0003]
Std. dev. of idiosyncratic preference	0.291 [0.0004]	0.341 [0.0005]	0.456 [0.0008]	0.262 [0.0003]
Std. dev. of idiosyncratic productivity	0.219 [0.0008]	0.228 [0.0008]	0.218 [0.0007]	0.218 [0.0007]
Std. dev. of shock to firm productivity	0.189 [0.0002]	0.224 [0.0002]	0.291 [0.0003]	0.140 [0.0002]

NOTE: This shows results of the sensitivity analysis of section 6. In column 2, the separation cost is set to equal 1 year of earnings. In column 3, the persistence of firm productivity is lowered to target the estimated persistence of value-added in Guiso et al (2005). In column 4, the returns to scale is raised to α =0.824. Standard errors are in brackets.

Table 7: Sensitivity Analysis, II

PANEL A

Moment	Baseline	1994-2001 subsample		Adjusted working time est.	
	Model	Model	Data	Model	Data
$\mathrm{var}(\epsilon^W)/\mathrm{var}(\epsilon^h)$	2.244	3.267	3.269	1.754	1.750
$\mathrm{var}(\phi^W)/\mathrm{var}(\phi^h)$	2.885	5.945	5.946	2.249	2.250
$\sqrt{\mathrm{var}(\epsilon^h)}$	0.140	0.125	0.125	0.159	0.159
$\sqrt{\mathrm{var}(\phi^h)}$	0.078	0.061	0.061	0.088	0.088
$\frac{\mathrm{cov}(\Delta \ln h, \Delta \ln w)}{\mathrm{var}(\Delta \ln w)}$	-0.170	-0.059	-0.059	-0.169	-0.169
$\sqrt{\mathrm{var}(\Delta \ln N)}$	0.175	0.184	0.184	0.177	0.175
$\mathrm{E}[N]$	17.131	16.760	16.760	17.130	17.130

PANEL B

Parameter (Symbol)	Baseline	1994-2001 subsample	Adjusted working time est.
$1/(1-\rho)$	0.344 [0.0006]	0.460 [0.0012]	0.437 [na]
$1/(\psi+1-\alpha)$	0.455 [0.0008]	0.315 [0.0010]	0.531 [na]
η	0.452 [0.0006]	0.569 [0.0015]	0.404 [na]
μ	0.196 [0.0006]	0.144 [0.0011]	0.102 [na]
σ_ξ	0.291 [0.0004]	0.329 [0.0010]	0.250 [na]
σ_θ	0.219 [0.0008]	0.297 [0.0021]	0.312 [na]
σ_Z	0.189 [0.0002]	0.211 [0.0004]	0.185 [na]

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

Figure 1: Labor demand policy

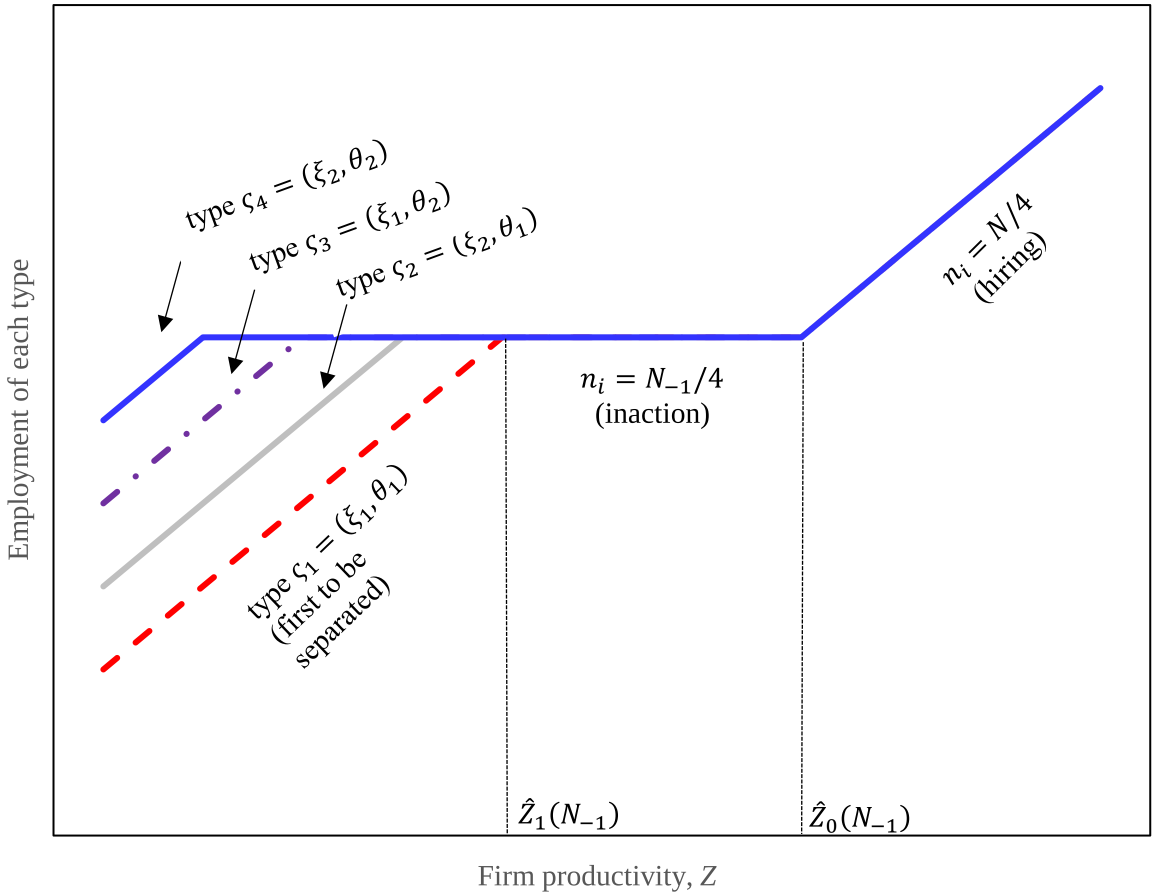

NOTE: This summarizes the optimal employment demand policy for the case of four equally likely types. e.g., the share λ_i of any type i equals 1/4. For high Z, employment of all four types is increased and, since λ_i=1/4 for each i, employment of each type, n_i, equals 1/4 of firm-wide employment, N. 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, $j > i$.