LABOR SUPPLY WITHIN THE FIRM*

Michele Battisti
Ifo Institute

Ryan Michaels
University of Rochester

Choonsung Park
University of Rochester

December 15, 2014

Abstract

This paper studies individual and firm-wide adjustment in working time and wages. The analysis is framed by a model in which heterogeneous firms and workers bargain working time and earnings. Critically, we allow that workers are complements in a decreasing-returns-to-scale production technology. This implies an incentive to coordinate labor supply within the firm, which compresses working-time adjustments across workers in response to purely idiosyncratic variation in their return from working. This places no restrictions, however, on the response of average working time to firm-wide shocks. The model yields closed-form solutions for working time and earnings bargains, in addition to predictions for plant-wide employment demands. It is assessed using matched firm-worker data from North-East Italy, and then used to revisit earlier findings of a small intertemporal elasticity of substitution.

*Preliminary and incomplete. We appreciate the feedback we have received 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; the 2014 North American and European meetings of the Econometric Society; and seminar participants at Sogang University (Seoul) and Rochester. We are also especially thankful for comments and encouragement from Yongsung Chang. All errors are our own.

Variation in labor input occurs along two margins. The extensive margin involves the commencement and termination of employment spells, whereas the intensive margin consists of adjustments to working time conditional on being employed. Much recent labor market analysis, perhaps most prominently in the search and matching literature, has focused on extensive-margin fluctuations. But variation along the intensive margin is significant. At the aggregate level, fluctuations in working time per person are comparable in magnitude to variation in employment in several European economies (Llosa et al, 2012). In the U.S., it has been long been known that intensive-margin adjustments account for about one-half of the variation in total labor input at a quarterly frequency (Heckman, 1984).[1] Moreover, at individual U.S. manufacturers, the variance of working time per person appears to be equal to that of employment (Cooper et al 2004).

This evidence on intensive-margin fluctuations suggests the importance of identifying the structural model underlying them. The theoretical analysis of observed working time has been dominated to a considerable extent by a framework in which workers are free to supply any desired level of labor to an employer whom they contact.[2] Pencavel (1986) notes, though, that there has long been a sense that this framework is incomplete. In particular, Pencavel says, a worker's labor input is often coordinated by his employer. Relatedly, Hall (1999) contends that, "if an event occurs that is personal to the worker ... it is unlikely that the employer will agree to a reduction in weeks *ad hoc*" (p. 1148). These comments place the *employer* at the center of the theory of intensive-margin labor supply.[3]

In the first section of this paper, we consider a framework for analyzing working time in which the firm does indeed have a starring role. We study a setting in which workers are complements in a firm's production process but have heterogeneous preferences over leisure. 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 then bargained jointly in order to maximize the surplus from participating in the match. The resulting distribution of working time across employees represents a balancing of two interests— productive complementarities and heterogeneity in the disamenity from working. If the former is forceful enough, then employees agree, jointly with their employer, to vary their working time in very similar manner despite having disperse preferences. For instance, an employee's working time may fall by relatively little in this setting in reaction to an increase in her marginal value of time. The reason is that a further reduction in her working time exacts a cost, in terms of degrading her colleagues' productivity, that outweighs the marginal disamenity from working.

[1] At an annual frequency, the contribution of the intensive margin in the U.S. to variation in total hours is smaller, on the order to 20-30 percent.

[2] This is true in life-cycle analyses, such as the influential works of MaCurdy (1981) and Altonji (1986) and, more recently, Pistafferi (2003). It is also true in many studies of experimental (randomized control) data. See, for instance, Burtless and Hausman (1978), Robins and West (1983), and Johnson and Pencavel (1984).

[3] Indeed, Pencavel and Hall both highlight theories in which the firm unilaterally chooses working time. As we shall see,w e observe far too much variation in annual working time fluctuations *within* firms to impose this assumption. Our approach still leaves room for "events personal to the worker" to influence labor supply among the employed.

The theory has several straightforward implications for the dynamics of working time and earnings at the individual and firm level. First, complementarities encourage workers to coordinate changes in their working times, despite (temporary) differences in their marginal values of time. Differences in preferences over leisure are accommodated, instead, in our model by the earnings bargain, which implements the solution to a generalized Nash bargaining game. Intuitively, if a worker's labor input remains high despite an increase in her marginal value of time, she is compensated accordingly. Hence, under complementarities, working-time adjustments *within the firm* are compressed *relative* to earnings growth (again, within the firm).

Second, the reaction of an employee's working time to idiosyncratic events does not predict the reaction of working time to *firm-wide* events, such as changes in demand for a firm's output. Purely idiosyncratic variation in an individual's return from working is unlikely to elicit a substantial response in working time because of complementarities. In contrast, employees' labor supply incentives are affected equally by firm-wide events. As a consequence, firm-wide working time can fluctuate considerably from year to year, depending (only) on the elasticity of intertemporal substitution. In this sense, complementarity "frees up" the model to predict significant changes in firm-wide working time—that is, to infer a nontrivial elasticity of substitution—without implying counter-factually large responses to (idiosyncratic) events personal to the worker.

Relatedly, the model imparts an important lesson concerning the source of variation capable of identifying the intertemporal elasticity of substitution. In particular, it suggests that firm-wide variation serves as a coordinating device, thereby eliciting a response in working time that more accurately reflects a worker's willingness to substitute effort across time. In contrast, the response of working time to a purely idiosyncratic change in the worker's own marginal value of time may bear little resemblance to the underlying structural parameter. This is simple, but important, point, because many influential studies on labor supply utilized precisely this latter variation. Consider, for instance, the series of papers that analyzed the randomized control trials in the U.S. known as the Negative Income Tax (NIT) experiments. In a handful of cities in the late 1960s and early 1970s, a (quasi)-random sample of workers received a cash grant on a sliding scale, with the grant declining in the worker's market earnings. The tepid response of working time to the treatment greatly informed the consensus on labor supply (Hall, 1999). Yet this variation is clearly idiosyncratic to the worker. The same point applies to the seminal life-cycle analyses of MaCurdy (1981) and Altonji (1986).

To assess our interpretation of working-time fluctuations and earnings, we introduce in Section 2 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 over the period, 1982-2001.[4] Critically, the dataset includes information on individual-level working time that is often absent in estab-

[4] In Italy, these data have to be recorded because taxes and social insurance contributions are tied to days worked. Data are reported to the public social security organisation INPS. The dataset has then been accessed, organized and maintained by researchers at the University of Venice. We thank Professor Tattara and the team of researchers at the University of Venice for giving us access to the dataset.

lishment data. Moreover, the reports of working time and annual earnings in the Veneto dataset are arguably more accurate then related measures in household surveys, easing concerns regarding measurement error.

The data report each employee's annual days worked and the number of months she is attached to each of her employers. Unfortunately, we do not observe total working hours. Nevertheless, we find, perhaps surprisingly, that working days is an active margin: among employees who remain with their firms in consecutive years, the standard deviation of working days is about 15. Though this understates variation in total working hours, it may serve as a worthwhile proxy. Supplementing our analysis with household data from the Italian Labor Force Survey, we find that variation in working days can account for the majority of variation in total hours. Digging still deeper, a good deal of this variation in working days appears to take the form of Saturday overtime, according to Giaccone (2009).

Our identification strategy relies on our ability to observe working time and earnings variation *inside* firms. Recall that complementarities compress the distribution of working time adjustments within the firm relative to the distribution of earnings growth (within the firm). This suggests that one especially salient moment of the data is the variance of annual working time changes across workers in the firm relative to the variance of annual earnings growth (again, within the firm). If this is small, 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, we have argued that the intertemporal elasticity of substitution is better informed by firm-wide fluctuations. In particular, the elasticity of substitution is revealed by the size of annual fluctuations in the firm's *average* working time, conditioning on the size of the underlying (firm-level) shocks. The latter shocks are identified, in turn, by the dispersion in the employment growth distribution.

A recent contribution of Chetty et al (2012) touches on a number of themes presented here. They identify evidence of coordination in working time using the "bunching" of taxable income at kinks in the tax-rate schedule. This bunching suggests, as they note, that an individual's working time may respond more to aggregate variation than to idiosyncratic changes in her marginal value of time. We use different data and apply a distinct identification strategy, but the implication is similar, namely, idiosyncratic variation in the return to working may fail to recover the true willingness of workers to substitute effort across time.[5]

The paper proceeds as follows. To frame the discussion, we introduce in section 1 a dynamic labor demand model in which a firm and worker bargain over working time (conditional on the worker participating) and wages. This offers a rich theory of the firm with heterogeneous workers, and characterizes analytically its implications for working time and earnings dynamics. In section 2, we describe our Italian firm-level data and present the key empirical moments that will guide

[5] Rogerson (2011) reaches the same conclusion, but he studies a model in which workers coordinate their leisure. The model is not quite suitable for testing on firm-level data.

our structural estimation. Section 3 carries out this estimation, and investigates the implications of our estimates for the effects of certain policy interventions with respect to labor supply. Section 4 concludes.

1 Theory

This section compares two models of the intensive margin. In the first, workers do not interact in production. We develop in section 1.1 this model's implications for earnings growth and working time fluctuations. Next, section 1.2 characterizes working time and earnings in a model in which workers perform complementary jobs in the production of a final good.

Throughout much of this section, we take as given a firm-worker match. We defer a complete treatment of extensive-margin adjustment under complementarities until the end of this section. For what immediately follows, what matters is that extensive-margin frictions–costs to hire and fire–imply the existence of rents to ongoing firm-worker attachments. Since the firm and worker are in a long-term relationship, we assume they are able to bargain to the efficient allocation of time.[6] Interestingly, Giaccone (2009) summarizes survey evidence for Italy showing that half of respondents report deciding working time by bargaining.[7]

1.1 A benchmark

In this section, the output of each worker is independent of the working time of other employees. That is, there are no complementarities. The firm simply sums up the working times of its employees and combines this total labor input with (formally, multiplies this by) its own technology to yield final output. Since production at the firm level is thus constant returns, we may think of the firm and each of its workers as jointly choosing that individual's working time, h, independently of the schedule of any other worker.

Under constant returns, the marginal product of a worker's time is Z, where Z indexes firm-wide profitability (or, technology). We assume the disutility from labor has the form $\xi\frac{h^{1+\varphi}}{1+\varphi}$, where the "taste" parameter, ξ, indexes the intensity of the disamenity from work, and $1/\varphi$ is the Frisch elasticity of labor supply. For the moment, ξ is the only source of heterogeneity across workers (we discuss innovations to a worker's individual productivity below).

The first-order condition (FOC) for an efficient allocation, h, equates the marginal value of time, given by ξh^{φ}, to the marginal product, Z. The latter should be valued in *utils*. However, we cannot distinguish the marginal value of income from ξ in our data. To see why, let the marginal value of income be ℓ, so the FOC now implies $Z = (\xi/\ell)\, h^{\varphi}$. We identify variation in the *ratio* ξ/ℓ using fluctuations in working time, h, across employees within a firm (that is, for given Z).

[6] The conclusions of section 1.1 hold under alternative protocols for determining working time, as we discuss shortly.

[7] Among this group, half cite "individual" bargaining—negotiations directly with management—whereas the other half cite "collective bargaining".

But in the absence of data on consumption or wealth, we cannot distinguish variation in ξ from fluctuations in ℓ. Thus, in what follows, we normalize $\ell = 1$, but keep in mind that variation in ξ may be interpreted as fluctuations in the marginal value of wealth. Card (1990) has flagged the latter as a promising source of variation in working time.[8]

Now rearranging the FOC, we write the optimal choice of time worked for type-ξ, denote by h_ξ, as

$$h_\xi = \left(\frac{Z}{\xi}\right)^{1/\varphi}. \tag{1}$$

Equation (1) reveals a simple, but important, result: in the absence of coordination, working time reacts symmetrically to firm-wide and worker-specific driving forces. This is a key restriction of the model. To the extent that the data "want" workers to react tepidly to ξ, for instance, the model infers a lower estimate of $1/\varphi$, the intertemporal elasticity of substitution. If we could identify ξ and Z separately in the data, we would test (1) directly. But we do not observe these disturbances. Instead, we show below how to use data on working time and earnings to examine the model (1).

Next, the firm-worker pair bargain over earnings. Defining $W_\xi \equiv w_\xi h_\xi$ as earnings, the Nash wage bargain for a type-ξ worker can be written as[9]

$$\begin{aligned} W_\xi &= \eta\left(Zh_\xi + r\underline{c}\right) + (1-\eta)\left(\mu + \xi\frac{h_\xi^{1+\varphi}}{1+\varphi}\right) \\ &= \eta r\underline{c} + (1-\eta)\mu + \tfrac{1+\eta\varphi}{1+\varphi} Zh_\xi, \end{aligned} \tag{2}$$

where r is the real interest rate, $\eta \in (0,1)$ is an index of worker bargaining power, and μ can be interpreted as the flow value of search (taken to be common across workers). The top line is a convex combination of the worker's contribution to the firm and his outside option. The former consists of the revenue from production, Zh_ξ, plus the separation tax, $\underline{c}$, which the firm avoids as long the match remains intact. The outside option consists of the flow value of search plus the utility, $\xi\frac{h_\xi^{1+\varphi}}{1+\varphi}$, that could be recovered if the worker quit. The second line uses equation (1) to replace h_ξ.

To distinguish (1)-(2) from a model with complementarities, we develop the theories' implications for the joint behavior of working time and earnings. In particular, we illustrate that the two models can leave very different imprints on the variance of working time changes across workers *within* the firm, relative to the variance of their earnings changes within the firm. These differences will provide the leverage we need to identify the degree of complementarity.

To this end, consider a simple thought experiment. Suppose all workers in the firm are initially identical, that is, $\xi = 1$ for all. Then the ξs and Z are drawn. This implies a log change in earnings

[8] Following Merz (1995), the earnings bargains in sections 1.1 and 1.2 are derived assuming the worker is insured by "large" family against job-specific (Z) and idiosyncratic (ξ) earnings risk. Thus, our approach still admits (exogenous) variation in ℓ, but only if it derives from a change in consumption possibilities at some higher aggregate, such as the worker's geographic region.

[9] See Appendix for derivation.

of $\Delta \ln W_\xi = (1 - (\omega/W)) \cdot [\Delta \ln Z + \Delta \ln h_\xi]$,where $\omega \equiv \eta r \underline{c} + (1-\eta)\mu < W$ and W is the initial level of earnings corresponding to $\xi = 1$. Since $\Delta \ln Z$ is common to all employees, the variance of earnings growth *within* the firm is

$$\mathrm{var}(\Delta \ln W_\xi) = (1 - (\omega/W))^2 \,\mathrm{var}(\Delta \ln h_\xi) < \mathrm{var}(\Delta \ln h_\xi). \tag{3}$$

This says that the distribution of earnings growth inside the firm is compressed relative to the dispersion in working time adjustments. As discussed in the Introduction, and shown below, this is precisely the opposite of what the data suggest.

To build intuition for this result, it is instructive to see how one might overturn it. First, take the limiting case of $\eta \to 1$ and $\underline{c} \to 0$. Then the earnings bargain is tied much more closely to working time, $W_\xi = Z h_\xi$. Accordingly, the variances of $\Delta \ln W_\xi$ and $\Delta \ln h_\xi$ across workers within the firm are equal (since Z is common to all). This suggests that, in the more general case of $\eta \in (0,1)$, earnings changes are less sensitive to events at the firm–and, thus, less disperse–because the firm's bargaining power anchors earnings, at least in part, to the outside option μ. We regard $\eta \in (0,1)$ as the more plausible case.

Second, suppose we introduce *worker*-specific productivity, in addition to Z. Output of the firm-worker match is now Zz, where z is the idiosyncratic (worker-specific) component of productivity. Assume z is independent of ξ. Then the first order condition for working time implies $\Delta \ln h_\xi = (1/\varphi)[\Delta \ln z - \Delta \ln \xi]$, and one can then show that

$$\mathrm{var}(\Delta \ln W_\xi) = (1 - (\omega/W))^2 \cdot \left[\mathrm{var}(\Delta \ln h_\xi) + \left(1 + 2\varphi^{-1}\right)\mathrm{var}(\Delta \ln z)\right].$$

This says that worker-specific productivity amplifies changes in earnings relative to changes in working time. The reason for this is that the (labor) demand-side shock, z, moves *both* components of earnings–working time and the wage rate–in the *same* direction, leading to larger earnings fluctuations. Thus, the compression in (3) also reflects in part the fact that a (labor) supply-side disturbance, ξ, is the only source of within-firm dispersion. As we show in section 2, the covariance between individuals' log earnings and working time changes is indeed negative. This finding echoes earlier work by, among others, Abowd and Card (1989). Seen through the lens of the model, this result substantially limits the scope of idiosyncratic productivity shocks (z) as driving forces of fluctuations in working time across employees.

We have derived these results under the assumption of efficient bargaining over time inputs. Our conclusions obtain under other, plausible protocols for determining working time. The online Appendix considers two cases: in one, the worker unilaterally chooses her working time, and in the second, the firm unilaterally makes the decision (the so-called "right-to- manage" protocol). In each case, we show that $\mathrm{var}(\Delta \ln W_\xi) < \mathrm{var}(\Delta \ln h_\xi)$, that is, these theories do not compress working time adjustments within the firm.

1.2 A model with complementarities

1.2.1 Working time

This section first describes the production structure under complementarities and the optimal choice of time worked. Our approach is parsimonious, and, as such, involves a number of assumptions that help retain tractability. These assumptions are discussed later in this section.

We envision a production process in which employees in a firm have to "work together" to some degree. A convenient way to formalize this is as follows. Suppose a firm's output is an aggregate over a continuum of *jobs,* which are (potentially) complements in the production of the final good. Formally, $y(i)$ is the output of job i, and final output is,

$$Y = Z\left(\int_0^1 y(i)^\rho \,\mathrm{d}i\right)^{\alpha/\rho}, \tag{4}$$

where Z is a firm-wide profitability shock; $\alpha \in (0,1)$ is the returns to scale at the firm level; and $\rho \in (-\infty, 1)$ is the elasticity of substitution across jobs. If $\rho < 0$, then jobs are gross complements.

Workers differ in their preferences over time worked. In each period, workers draw a type, or taste, for labor. There is a finite number, M, of types, and the set of these types is denoted by $\mathcal{X} \subseteq \mathbb{R}^M$. The share of the workforce which draws $\xi \in \mathcal{X}$ in any period is given by $\lambda_\xi \in (0,1)$, where $\sum_{\xi\in\mathcal{X}} \lambda_\xi = 1$ and the mean, $\frac{1}{M}\sum_{\xi\in\mathcal{X}} \xi$ is normalized to 1. We refer to workers who draw ξ as type-ξ workers.

For now, differences over the preference for labor is the only source of heterogeneity across workers. Thus, all workers of a given type are identical. For this reason, one may analyze the labor supply decision in what follows as if there is a representative worker associated with each type ξ.[10]

As we show below, the firm and worker jointly negotiate the worker's *total* input of time. Conditional on this, we assume the firm is free to allocate the worker's efforts *across* jobs. In principle, this is a nontrivial problem. However, since our data do not speak to this level of detail, we try to streamline the analysis. To that end, assume workers are equally productive at all jobs. In that case, one can imagine that the firm implements a simple "rule of thumb" that allocates an equal measure $m \equiv 1/M$ of (non-overlapping) jobs to each team of workers of type ξ. Equation (4) then simplifies nicely. Let n_ξ denote the mass of the cohort ξ and h_ξ the time worked of each employee of type ξ. If the output $y(i)$ of any job i equals the man-hours supplied, then our restrictions imply that $y(i) = n_\xi h_\xi / m$ for any i in the set of jobs performed by type ξ. Therefore, (4) becomes

$$Y = F(\mathbf{h}, \mathbf{n}, Z) = Z\left(\sum_{\xi\in\mathcal{X}} (n_\xi h_\xi)^\rho\right)^{\alpha/\rho}, \tag{5}$$

[10] More precisely, we conjecture that all workers of a given type will supply the same amount of time worked, and confirm this below.

where $\mathbf{h} \equiv (h_1, \ldots)$ is the vector of time inputs across types and $\mathbf{n} \equiv (n_1, \ldots)$ lists the sizes of each team. (The constant $m^{(1-\rho)\alpha/\rho}$ has been subsumed into Z.) The presence of $\alpha < 1$ here implies decreasing returns to the plant-wide workforce, ensuring a well-posed employment demand problem (see section 1.2.5).

The choice of type ξ working time, h_ξ, is determined by the following static necessary (and sufficient) condition, given a vector of workers $\mathbf{n} \equiv (n_1, \ldots)$. This equates the marginal disutility of time worked to the marginal product of an additional unit of time worked. Given our functional forms, one may show that this collapses to[11]

$$h_\xi = (\alpha Z)^{\frac{1}{\varphi+1-\alpha}} \, \Omega(\mathbf{n}) \cdot \left[n_\xi^{\rho-1}/\xi \right]^{\frac{1}{\varphi+1-\rho}}, \tag{6}$$

where $\Omega(\mathbf{n}) \equiv \left(\sum_{x \in \mathcal{X}} (n_x^\varphi / x)^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha-\rho}{\rho} \frac{1}{\varphi+1-\alpha}}$. Substituting (6) into (5), we obtain

$$Y = \hat{F}(\mathbf{n}, Z) \equiv \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \left(\sum_{\xi \in \mathcal{X}} \left(n_\xi^\varphi / \xi \right)^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha}{\rho} \frac{\varphi+1-\rho}{\varphi+1-\alpha}}, \tag{7}$$

which expresses revenue *after* optimization of time worked.

1.2.2 Earnings

Our bargaining protocol is due to Stole and Zwiebel (1996), which was generalized by Cahuc, Marque, and Wasmer (2008) to the case of heterogeneous workers. When bargaining opens, any worker (of any type) may request a pairwise bargaining session with the firm. Taking as given the participation of the remaining workers, the firm and employee split the surplus between them according to an exogenously given bargaining weight. (The firm may also request a bargaining session with any worker at any time.) If a pairwise session ends in disagreement and the worker exits, the protocol enables the remaining workers to renegotiate their wages. In particular, in the event of breakdown, diminishing marginal product implies that the productivity of remaining workers would rise, and they would request bargaining sessions to re-balance the allocation of the surplus. Bargaining continues in this manner until a "stable" outcome is achieved, that is, until no pairwise sessions are requested.

At the conclusion of bargaining, then, the wage of any worker in type ξ must satisfy the

[11] The problem is to select time worked for each worker of each team ξ. Formally, then, we solve, $Z \left(\sum_{\xi \in \mathcal{X}} \left(\int_0^{n_\xi} h_\xi(i) \, \mathrm{d}i \right)^\rho \right)^{\alpha/\rho} - \sum_{\xi \in \mathcal{X}} \xi \int_0^{n_\xi} \frac{h_\xi(i)^{1+\varphi}}{1+\varphi} \mathrm{d}i$. With respect to worker i on team ξ, the FOC is

$$Z\alpha \left(\sum_{\xi \in \mathcal{X}} \left(\int_0^{n_\xi} h_\xi(i) \, \mathrm{d}i \right)^\rho \right)^{\frac{\alpha-\rho}{\rho}} \left(\int_0^{n_\xi} h_\xi(i) \, \mathrm{d}i \right)^{\rho-1} = \xi h_\xi(i).$$

Since a worker is "small" relative to his team, the left side is identical from the perspective of each member i of any type ξ. Moreover, since team members share the same ξ, it is optimal for each worker of type ξ to work the same time. This enables us to simplify and arrive at (6).

following surplus-splitting condition. The surplus accruing to the worker is given by $\mathcal{W}_\xi - \mathcal{U}$, where $\mathcal{W}_\xi$ denotes the value of working to an employee of type ξ and $\mathcal{U}$ represents the value of not working. The surplus accruing to the firm from an ongoing relationship is $\mathcal{J}_\xi + \underline{c}$, where $\mathcal{J}_\xi$ denotes the value of employing a member of type ξ and $\underline{c}$ is the cost of separating that the firm can avoid paying by continuing to employ the worker. The separation cost could be interpreted as legally mandated severance pay. If wage negotiations break down, and the worker exits, the firm would pay $\underline{c}$.[12] Letting η denote the worker's bargaining weight, the surplus-splitting condition then states,

$$\mathcal{W}_\xi - \mathcal{U} = \eta \left(\mathcal{W}_\xi - \mathcal{U} + \mathcal{J}_\xi + \underline{c} \right). \tag{8}$$

We note that $\mathcal{U}$ is not indexed by ξ. The reason is that ξ is regarded as a temporary shift in a worker's preference over current working time. It has no bearing on a worker's expected future earnings from being matched with a new employer, which is encased in $\mathcal{U}$.[13]

In our setting, this sharing rule yields the following bargain.

Proposition 1 *The solution to the generalized Stole and Zwiebel bargaining game is given by*

$$W_\xi \left(\mathbf{n}, Z \right) = \eta \left[A \frac{\partial \hat{F} \left(\mathbf{n}, Z \right)}{\partial n_\xi} + r \underline{c} \right] + (1 - \eta) \left(A g_\xi \left(\mathbf{n} \right) + \mu \right), \quad \xi \in \mathcal{X} \tag{9}$$

where $A \equiv \frac{\varphi + 1 - \alpha}{(\varphi + 1)(1 - \eta(1 - \alpha)) - \alpha} > 1$ *and* $g_\xi \left(\mathbf{n} \right) \equiv \xi \frac{h_\xi (\mathbf{n})^{1+\varphi}}{1+\varphi}$.

The basic form of the bargain is intuitive. The worker's earnings are a convex combination of a measure of the worker's marginal contribution to the firm's revenue and his outside option. The former consists of both his marginal product and the severance cost, $\underline{c}$, that the worker saves the firm by remaining with this employer. The latter consists of the flow return on job search, denoted by μ, and the utility that could be gained by foregoing labor, $g_\xi \left(\mathbf{n} \right)$.

1.2.3 Testable restrictions

The two models of sections 1.1 and 1.2 have distinct implications for working time and earnings fluctuations inside firms. To see these, there is a special case of (9) that is particularly instructive. Suppose all tastes are equally likely, so $\lambda_\xi \equiv \lambda = 1/M$ for all $\xi \in \mathcal{X}$. Also, assume that the firm, in light of the costs of hiring and firing, chooses to leave firm-wide employment at its initial level, N_{-1}. In that case, $n_\xi = \lambda N_{-1} \equiv n$ for any ξ. Then (6) simplifies to

$$h_\xi = \mathcal{P} \left(N_{-1}, Z \right) \cdot \xi^{-\frac{1}{\varphi + 1 - \rho}}, \tag{10}$$

[12] The value of a worker to the firm, $\mathcal{J}_\xi$, is recovered as part of the characterization of the complete dynamic employment demand problem. We defer a discussion of this problem until later in section 1, in order to retain our focus on the joint behavior of working time and earnings.

[13] It is true that ξ does influence current earnings. As a result, it can affect a worker's unemployment benefit if the latter is tied to her earnings on her last job (as in Italy's *cassa integrazione* program). We assume this source of income risk–essentially, the dispersion in earnings due to variation in ξ–can be fully insured.

where $\mathcal{P}(N_{-1}, Z) \equiv \Xi \left(Z\alpha n^{\alpha-1}\right)^{1/(\varphi+1-\alpha)}$ and $\Xi \equiv \left(\sum_{x\in\mathcal{X}} x^{\frac{-\rho}{\varphi+1-\rho}}\right)^{\frac{\alpha-\rho}{\rho}\frac{1}{\varphi+1-\alpha}}$. Equation (10) consists of firm-wide and idiosyncratic (worker-specific) components. The idiosyncratic term reflects the worker's own (dis)taste for labor. The firm-wide component, $\mathcal{P}(N_{-1}, Z)$, is a combination of two elements. The first, $Z\alpha n^{\alpha-1}$, indicates the marginal product of labor: if this is high at $n \equiv \lambda N_{-1}$, workers supply more time. The other, Ξ, reflects the fact that other teams' tastes influence those teams' hours bargains, and the latter partly drive the choice of this type's choice h_ξ via complementarities.

Our first result is that the effect of a worker's own preferences on working time diminishes as complementarity increases. To see this most clearly, we now take the limit as $\rho \to -\infty$, which implies that workers are perfect complements in production. Equation (10) reveals that

$$\rho \to -\infty \Longrightarrow h_\xi \to \left[Z\alpha n^{\alpha-1}/M\right]^{1/(\varphi+1-\alpha)}. \tag{11}$$

This says that idiosyncratic variation in the marginal value of time, in the form of ξ, has no direct effect on a team's own time worked. Intuitively, under perfect complements, the marginal product of an individual's additional time tends to zero *holding constant* the working time of her co-workers. Accordingly, the efficient bargain calls for no change in her working time in response to variation in her own marginal value of time.

However, an increase in ξ must still be compensated if $\rho \to -\infty$. To see this, we return to (9) and consider the special case where $\lambda_\xi \equiv \lambda = 1/M$ and $n_\xi = \lambda N_{-1} \equiv n$ for all ξ. Then, taking the limit to perfect complements, we can show the following:

$$\rho \to -\infty \Longrightarrow W_\xi(\mathbf{n}, Z) \to \eta\left[\frac{\varphi}{\varphi+1-\alpha}\hat{\xi} AZ\alpha n^{\alpha-1} + r\underline{c}\right] + (1-\eta)\left(A\xi g(n) + \mu\right),$$

where $\hat{\xi} \equiv \xi/M$. Hence, in this limit, it is exclusively the earnings bargain that accommodates idiosyncratic variation in working-time incentives. Specifically, a higher ξ is compensated by higher earnings.

This result contrasts with the case analyzed in section 1.1. In that setting, in which workers are independent in production, there is no force for compressing working-time changes within the firm. This leads to the prediction that dispersion in working-time changes *exceeds* that in earnings growth. In the setting of this section, complementarities in production compress working-time changes within firms, and idiosyncratic variation is accommodated to a greater extent by variation in earnings growth.

Complementarities also have important implications for our understanding of firm-wide working time. The hours bargain (1) in section 1.1 predicts that working time is equally elastic with respect to firm-wide and worker-specific events. As a result, the benchmark model of section 1.1. cannot reconcile competing evidence on labor supply derived from plant-level fluctuations (Cooper et al, 2004) and households' life-cycle hours profiles (MaCurdy, 1981; Altonji, 1986), as discussed in the

Introduction. Complementarities can account for the muted response of working time to employee-specific events, thereby "freeing up" the response of working time to plant-level shocks. This can be see in equation (6). Taking logs of each side, evaluating the result at the mean working time within the firm, and differentiating with respect to Z, we have that

$$\frac{\partial \mathbb{E}\left[\ln h_{\xi}\right]}{\partial \ln Z} = \frac{1}{\varphi + 1 - \alpha}. \tag{12}$$

Hence, the elasticity of firm-wide working time is independent of ρ, and anchored exclusively by φ and α. The first of these, φ, summarizes the worker's willingness to substitute effort intertemporally. The second, α, conveys the firm's elasticity of labor demand.

1.2.4 On modeling working time and earnings

Before turning to the choice of employment, we comment on a number of assumptions we made in connection with modeling the production process, working time and earnings.

Production. We have suppressed any mention of a worker's individual productivity in (7). The theoretical analysis can be extended to include productive heterogeneity across workers, as we show in detail in the online Appendix. In brief, the revenue function (5) takes the same form, but is now an aggregate over the output of each $\varsigma \equiv (\xi, z)$ team, where z denotes the i.i.d. worker-specific productivity draw. That is, a "team" is now characterized by the pair of preference and productivity shocks. The output of a team ς is simply $zh_{\varsigma}n_{\varsigma}$. The quantitative model of section 2.3 does incorporate fluctuations in an individual's productivity, z.

Even then, though, we continue to abstract from *fixed*, permanent differences in productivity across workers because this source of variation is not used in our empirical application. The reason is that, although permanent heterogeneity supports a non-degenerate distribution of time worked within the firm, the compression–or lack thereof–in this distribution does not help identify complementarities. A very disperse distribution can arise, for instance, even if workers are perfect complements: their time worked in this case would be set in order to equate efficiency units across employees. However, time worked still must move together in response to shocks. Hence, the distribution of *changes* in time worked is more informative with respect to coordination of labor supply, and it is the source of variation used later in our identification strategy.[14]

Next, the derivation of (7) required some treatment of how individual jobs are allocated to workers. In the online Appendix, we characterize the optimal solution to this assignment problem. The rule-of-thumb is optimal only if total man-hours are the same for all types. In that case, the

[14] Even when we allow that workers have different productivities, we treat the worker's productivity as general human capital, that is, one worker is more productive on *all* jobs than another worker. We do *not* model comparative advantage over jobs–for this, see Acemoglu and Autor (2011). We instead focus on the idea that, for a given production process in the short run, workers of all skill levels likely have to coordinate their effort to produce the final output. Over the longer run, the firm will adopt new production technologies and adjust the composition of its workforce. These longer-run adjustments are not the subject of this paper.

shadow value of time is equal across types, and so the firm assigns each type the same measure of jobs. In general, of course, man-hours will differ across types, and so our rule of thumb does suppress one margin of adjustment (the allocation of jobs) that would otherwise be active. We contend that this is a useful simplification because it enables us to focus on the choice of total working time and its distribution within the firm, which are objects that we observe in the data.

Working time. We have chosen a decentralized protocol for determining working time, allowing the firm and its workers to jointly decide the allocation. This contrasts with Deardorff and Stafford (1976), who provide perhaps the first formal treatment of a problem where firms coordinate the labor supplies of heterogeneous workers. Deardorff and Stafford assume the firm enforces a common work schedule across employees. As we shall see, though, there is indeed dispersion of working time changes within the firm, despite evidence of complementarities. This variance within the firm can be easily accommodated within our approach.[15]

Nonetheless, our protocol surely misses some institutional realities of working time determination. For instance, we abstract from unions' involvement in the choice of working time. The unions are indeed often consulted by employers concerning working time changes, especially if the employer proposes reducing working time plant-wide (Giaccone, 2009; Treu, 2007). Moreover, unions have negotiated limits on overtime, though these arrangements still permit 200 or more hours of overtime per year (which amounts to 25 or more eight-hour days). Still, it is hard for us to judge the extent to which unions compress working time changes *within* the firm by removing any role for negotiations between firms and individuals. Accordingly, we have declined here to try modeling a direct role for unions in deciding working time, albeit at some cost in terms of realism.

Finally, there is one clear element of working time in the data that our model is not designed to engage, namely, the persistence in working time at the employee level. As we shall see, there are many workers in our data who do not adjust their working time from year to year. To reproduce this degree of inaction, we could introduce costs of adjusting working time to target this inaction. However, to the extent that these frictions simply convert small, positive working time adjustments into zeros, it is unlikely that they will alter the conclusions of our analysis or affect our identification of the structural parameters. Thus, for the sake of tractability, we abstract from this frictions.

Earnings. We next address a few concerns regarding our application of this bargaining protocol to the Veneto labor market. Italy is often viewed as an economy where collective bargaining is the main mechanism for wage determination. In reality, there are many sources of wage heterogeneity across workers. This reflects, in part, the fact that national regulations are typically silent about compensation levels. Trade union contracts are prevalent, but specify non-binding minimum wages at the industry level. These imply an industry-specific floor for total compensation, but in the region of Veneto, actual compensation is typically higher. This leaves considerable scope for individual bargaining and firm-level agreements, and wage premia are highly heterogeneous across

[15] Deardorff and Stafford also assume the wage is determined competitively. In our setting, costs of extensive-margin adjustments imply rents to ongoing employment relationships. This seems more suitable to a study of firm-level data. Accordingly, earnings in our case are determined as the solution to a bargaining problem, as discussed below.

firms (Erickson and Ichino, 1993; Cingano, 2003). In light of these considerations, we thought Stole and Zwiebel–a decentralized Nash bargain, in effect–seemed a reasonable benchmark.

Still, our approach to wages does neglect evidence that firms in Italy smooth workers' earnings. In an influential paper, Guiso, Pistaferri, and Schivardi (2005) find that workers' earnings change by 0.6 percent per 10 percent change in firm value added. This finding suggests that the bargain (9) may overstate the elasticity of earnings with respect to Z, which indexes firm-wide profitability.[16]

However, this evidence does not necessarily bear directly on the elasticity of earnings to events that idiosyncratically affect an *individual*'s return from market . We could envision a firm insuring a worker against *firm*-level earnings risk even as it agrees to higher earnings for a few employees who supply effort despite having a high idiosyncratic marginal value of time. As we shall see, our identification of complementarities hinges most significantly on the response of earnings to these idiosyncratic events. Moreover, the elasticity of (9) with respect to Z will also not necessarily affect our identification of the Frisch elasticity, $1/\varphi$. The reason is that working time is determined efficiently in the model. Thus, the variance of firm-wide working time, which provides critical identifying information for φ, is not sensitive to the elasticity of the earnings bargain.

1.2.5 Employment demand

Thus far, we have taken total firm employment as given in order to focus attention on the model's implications for the joint dynamics of working time and earnings. But the combination of complementarities and extensive-margin adjustment frictions imply a nontrivial dynamic labor demand problem. Thus, in this section, we shift gears to study the extensive margin of our model.

As we shall see, this dimension of the problem plays two roles in our analysis. First, the solution to the labor demand problem is required to characterize the value of a worker to the firm, denoted above by $\mathcal{J}_\xi$; this was a critical input into the wage bargaining problem. Second, fluctuations along the extensive margin yield another source of variation that sheds light on a key structural parameter of the problem, namely, the size of firm-wide shocks, Z.

Proceeding, we now describe the dynamic labor demand problem. At the start of a period, the firm has a workforce of measure N_{-1}.[17] Firm productivity Z is realized. At this point, the firm may hire at cost $\bar{c}$ per position. The parameter, $\bar{c}$, represents the cost to the firm to recruit and train a worker. We assume hires are anonymous, in that the firm does not observe their type 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. In particular, a share $\lambda_\xi \in (0,1)$ of the workforce draws type $\xi \in \mathcal{X} \subset \mathbb{R}^M$, where $\Sigma_{\xi \in \mathcal{X}} \lambda_\xi = 1$.[18] At that point, the firm and (some of) its workers may jointly decide to separate at cost $\underline{c}$ per separation. Let s_ξ denote the number of separations of type-ξ workers. Then these flows

[16] Our data does not include firm value added, so we can directly discipline (9) using this information.

[17] Througout, the subscript $_{-1}$ denotes the one-period lag, and a prime $'$ is used to denote next-period values.

[18] We remove any uncertainty regarding the realization of the ξs: in each period, a known share of the workforce draws a known type.

out of and into the firm satisfy,

$$\begin{array}{c} s_\xi = \max\{0, \ \ \lambda_\xi \mathcal{N} - n_\xi\} \\ N = \sum_{\xi\in\mathcal{X}} n_\xi \end{array}, \tag{13}$$

where n_ξ is the number of workers of type ξ retained by the firm and N is the size of the workforce taken into next period. After separations are decided, wages and time inputs are bargained.

The timing of events makes it convenient to break the problem into two stages. First, consider a firm with some $\mathcal{N}$. The types, ξ, are revealed, and the firm's remaining choice concerns which workers to retain. Let $\mathbf{n}$ denote a $M \times 1$ vector of the employment levels of each type. We use $\pi(\mathbf{n},Z)$ to stand for flow profit gross of adjustment costs but conditional on the optimization of working time, $\pi(\mathbf{n},Z) \equiv \hat{F}(\mathbf{n},Z) - \sum_\xi W_\xi(\mathbf{n},Z)\, n_\xi$.[19] The problem of this firm can be characterized by the Bellman equation,

$$\begin{array}{c} \Pi^-(\mathcal{N},Z) \equiv \max_{\mathbf{n}} \left\{ \ \tilde{\Pi}^-(\mathbf{n},\mathcal{N},Z) \ \right\} \\ = \max_{\mathbf{n}} \left\{ \pi(\mathbf{n},Z) - \underline{c} \sum_{\xi\in\mathcal{X}} s_\xi + \beta \int \Pi(N,Z')\, \mathrm{d}G(Z'|Z) \right\}, \end{array} \tag{14}$$

where separations, s_ξ, are given by (13) and $G()$ is the distribution function of future productivity conditional on its present value. Second, we take one step back and characterize the firm's choice of hires, which brings its workforce up to the level $\mathcal{N}$. Since hires' types are anonymous to the firm, the value of the firm at this stage is given by

$$\Pi(N_{-1},Z) = \max_{\mathcal{N}} \left\{ -\bar{c} \cdot \max\{0, \mathcal{N} - N_{-1}\} + \Pi^-(\mathcal{N},Z) \right\}. \tag{15}$$

To shed light on the form of the optimal labor demand policy, consider the problem of a firm that has workforce $\mathcal{N} = N_{-1}$ (it does not hire). We ask whether this firm should separate from workers of (arbitrary) type ξ, taking as given the participation of the remaining types.[20] A separation is made if the marginal value of labor, evaluated at N_{-1}, is less than the separation cost. Formally, this implies

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

where $\boldsymbol{\lambda}$ is a $M \times 1$ vector of the shares λ_ξ, and the derivative of π is evaluated at the initial workforce, $\mathbf{n} \equiv \boldsymbol{\lambda} N_{-1}$. The appendix verifies that the marginal value of labor, the left side of (16), is increasing in Z. It follows that there exists a function, $\zeta_\xi(N_{-1})$, such that a type-ξ worker is separated if and only if $Z < \zeta_\xi(N_{-1})$. Therefore, the type of worker to be separated first is the type ξ for which the threshold, $\zeta_\xi(N_{-1})$, is highest.

To gain some insight into the mapping from ξ to the thresholds ζ_ξ, let us consider each of the

[19] The arguments of the earnings function, W, anticipate the result of bargaining.

[20] In the Appendix, we provide conditions under which a firm that is separating will in fact not hire. Accordingly, it will be the case that $\mathcal{N} = N_{-1}$ for a separating firm.

two terms in (16). First, the appendix shows that the flow marginal product, $\frac{\partial\pi(\boldsymbol{\lambda}N_{-1},Z)}{\partial n_\xi}$, has the form,

$$\frac{\partial \pi\left(\boldsymbol{\lambda} N_{-1}, Z\right)}{\partial n_{\xi}}=[\text{terms in } \left(N_{-1}, Z\right)] \times \lambda_{\xi}^{-\frac{(\varphi+1)(1-\rho)}{\varphi+1-\rho}} \xi^{\frac{-\rho}{\varphi+1-\rho}}.$$

If $\rho < 0$, this is increasing in ξ. The reason for this is that individuals with a stronger distaste for work (a higher ξ) supply less working time *conditional on* participation. As a result, these employees' participation is all the more valued if jobs are gross complements ($\rho < 0$). Second, as for the forward value in (16), the assumption of i.i.d. ξ draws implies that the expected future value of any worker is the same regardless of type. It follows that the marginal value of labor, the left side of (16), is increasing in ξ.

It is now straightforward to determine who is separated first. For instance, in the special case where $\lambda_\xi = \lambda = 1/M$ for all ξ, workers with the lowest ξ are separated first. If λ_ξs differ across types, though, one has to take account of the size of the cohort, as reflected by the role of λ_ξ in $\frac{\partial\pi(\boldsymbol{\lambda}N_{-1},Z)}{\partial n_\xi}$. In this case, the first to separate is that group which is both relatively abundant (λ_ξ is large) and whose taste for work is relatively high (ξ is small).

If Z falls further, the firm will seek separations from another type, $x \neq \xi$. As the firm does this, separations from the first type ξ will continue. This follows from the nature of the complementarities ($\rho < 0$): as the firm reduces labor input of a second type x, that further reduces the marginal value of the first type ξ. Thus, the optimal policy prescribes that *both* types are separated in tandem. This intuition underlies the result given below, which is proven in the Appendix. To state the proposition, we use the notation $\xi_1,\ldots,\xi_j,\ldots,\xi_M$ to convey that a type ξ_j is the jth type to be separated.

Proposition 2 *There exists a ranking* $\xi_1, ..., \xi_M$ *and a corresponding set of functions* $\left(\zeta_1\left(N_{-1}\right), \ \zeta_2\left(N_{-1}\right), ...\right)$, *with the latter listed in decreasing order, such that workers of all types* $\left(\xi_1, ..., \xi_i\right)$ *are separated if and only if* $Z < \zeta_i\left(N_{-1}\right)$.

Figure 1 provides two graphical perspectives on the labor demand policy for the case of three types ($M = 3$). In the top panel, initial firm-wide employment, N_{-1}, is fixed, and the profitability index, Z, ranges over the horizontal axis. For a set of Zs, all three types' employment levels are unchanged from their start-of-period values. This reflects the presence of adjustment frictions along the extensive margin, specifically, $\bar{c}$ and $\underline{c}$. The wedge between the marginal costs of upward and downward adjustments implies that there are Zs such that it is optimal to neither hire nor fire–there is an optimal degree of inaction. To the right of this range, all types' employment increases. Since hires are anonymous to the firm, the employer hires a measure of workers that is representative of each type's share in the population. As Z declines to the left of the range of inaction, one type's employment is reduced, while other types' participation remains fixed. In this example, the shares λ_ξ are the same, so the first type to separate is that with the lowest ξ, denoted by ξ_1. As Z falls further, a second type begins to separate jointly with type ξ_1.

In the bottom panel, we show the choice of firm-wide employment, N, as a function of its start-of-period value, N_{-1}. Again, the region of inaction is evident by the portion of the employment demand rule that lies on the 45-degree line where $N = N_{-1}$. To the right of this, $N_{-1} > N$, which means that the firm undertakes separations because N_{-1} is high relative to profitability, Z. In a model with one type of labor, this portion of the labor demand schedule would be flat: conditional on separating, the firm's choice of N would be independent of N_{-1} since the marginal adjustment cost is a constant, $\underline{c}$. In a model with multiple types of workers, however, a firm which is separating from one type will attenuate those separations if there are many complementary types, as indicated by a large N_{-1}. The hiring portion of the employment demand rule is indeed flat, since hires are anonymous to the firm and equally costly. Thus, conditional on hiring, the firm simply expands its size until the marginal value of total employment, N, equals its marginal cost, $\bar{c}$.

This completes our description of the model under complementarities. The next section describes the data we will use to estimate the models of sections 1.1 and 1.2.

2 Veneto Work History Files

In this section, we begin with a brief introduction to our dataset. It includes, uniquely for panel datasets, a direct measure of working time. However, our measure is imperfect. We show, though, that our data likely captures a large majority of annual variation in working time. To conclude our discussion, we summarize a few salient empirical moments that can be used to identify the key structural parameters of the models.

2.1 Data description

Our empirical analysis utilizes the Veneto Worker History (VWH) dataset, which includes virtually all private-sector workers of the northern Italian region of Veneto for years 1982- 2001.[21] The VWH dataset records all employees that have been hired in Veneto for at least one day during the period of observation. The full sample contains around 3.6 million workers and 46 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 number of months worked with each firm. It also gives a worker's annual earnings, from which we can compute the average daily wage.

Table 1 provides a set of summary statistics. On average, workers work nearly 24 days per month (conditional on positive days worked that month). This reflects the prevalence of six-day weeks in this region of Italy in this period. As noted above, the sixth day, in many cases, represents

[21] The region of Veneto, in the North-East of Italy, is one of the largest in Italy (its 2001 GDP ranks third among twenty of Italian regions), and it has a population of around 5 million, or 8 percent of the country's total.

overtime. The average gross daily wage is around 120 Euros, and on average the number of paid months per worker is about 10.

Table 2 zeroes in on moments related to the distribution of annual changes in working days. As noted above, many workers do not adjust days from one year to the next. At the same time, though, 17 percent change working days by more than 10 days. Moreover, conditional on changing days, the typical size of the change is between 10 and 22, depending on whether some of the largest adjustments are included in the sample.

Since our data measure *paid* work days, it is important to be precise about paid leaves of absence. In Italy, workers are typically guaranteed at least 4 weeks of paid vacation. If this is taken each year, we will difference it out in computing annual changes in working time. Other forms of leave will likely show up in our estimates. For instance, a two-parent household is granted 16 months of paternity leave. The parents receive income support from Social Security for the first 11 months of leave.[22] Hence, paternity leave will affect our measure of time worked. For this reason, we repeat our analysis below in a sub-sample with men only, out of concern that the mother's leave in particular might affect our results.

The dataset also provides information on the employment arrangement. For instance, the data identify a worker as full-time or part-time. In addition, the Veneto data reports the type of contract under which a worker was hired. In Italy, a worker may be hired under a permanent contract, which includes restrictions on individual dismissals. Alternatively, beginning in the late 1980s, employers were allowed to hire workers under fixed-term contracts. The latter expired after two years, at which point the employer could dismiss the worker without penalty. However, despite the presence of these more flexible arrangements, it should be noted that part-time and fixed-term contracts were not widespread over our sample. On average, in our data, only 7 percent of workers were part-time and 11 percent of the workforce was employed on a fixed-term contract. In our baseline analysis that follows, then, we typically do not break down the workforce along these lines.[23]

2.2 Measuring working time

Even though the Veneto files stand out for providing *any* information on working time, the absence of total working hours is still worrying. The reason is that our measure of earnings is, implicitly, based on total working hours. This discord between the measurement of working time and earnings affects our analysis in two ways.

[22] The 16 months include 5 months of leave for the mother, to be used for the last 2 months before birth and the first 3 months after birth. Social Security pays 80 percent of the mother's salary over these 5 months. The other 11 months are shared by the parents–each day of leave by either parent counts against this allotment. Social Security pays 30 percent of the parents' salaries for the first 6 of these 11 months. The last 5 months are unpaid (Ray, 2008).

[23] Though formally in use since 1962, fixed-term contracts could be legally applied only in a narrow set of circumstances. This began to change in the late 1980s. Restrictions were further relaxed under a European Union directive, but the latter was not approved by Parliament until 2001. Restrictions on part-time work were also relaxed relatively recently. For instance, in 2000, it became legal to employ a worker at normal weekly hours (40) but for a limited period. See Tealdi (2011) for more.

First, suppose that, in response to *firm-wide* events, workers increase both working days as well as hours per day. In that case, our data understate the variation in the firm-wide component of working time (in the model, the variation due to Z). We refer to this as the *coordinated* response of working time. Understating the coordinated variation would be problematic for our approach, since we rely on this information to identify the intertemporal elasticity of substitution. Specifically, missing variation in hours per day can lead us to understate this elasticity.

The second concern has to do how working time reacts to *idiosyncratic* events, that is, events personal to the worker. This idiosyncratic variation makes up the dispersion we see within firms in working time and earnings fluctuations. Our identification strategy relies on comparing the dispersion in working time adjustments to the dispersion in earnings growth (inside the firm). This comparison is compromised if workers react to idiosyncratic events by varying daily hours rather than working days. The reason is that we do not observe changes in daily hours but these are reflected in annual earnings growth. As a result, our estimates will exaggerate the compression in working time changes, leading to an overestimate of the degree of complementarities.

We do not know of any direct measurements pertaining how coordinated and idiosyncratic working time variation are apportioned between days and daily hours. What we have tried to do is gauge how much total working time variation we are likely missing in our Veneto panel.

To this end, we turn to the Italian Labor Force Survey (LFS). The LFS is administered quarterly, and is a rotating panel: each household is surveyed for two (consecutive) quarters; exits the sample for the next two quarters; and then re-enters for two more quarters. Unlike the U.S. Current Population Survey, the Italian LFS asks households about both their weekly hours and days worked (which is itself arguably suggestive of the role of working days in Italy). Thus, for half of the LFS sample in any (calendar) quarter, we can calculate annual changes in both hours and working days. In total, we use data for 155,820 workers over the period 1993-2001.[24]

Our analysis of the LFS suggests that working days can account for most of the variation in total working hours. We arrive at this by comparing the variances of changes in working days and total weekly hours. Suppose that if a worker adds or reduces her working time by a day, that represents a change of 8 hours of work (the typical length of a workday). Then, the variance of changes in working days accounts for 70 percent of the variance of total hours changes. This holds for the full sample; for full-time workers; and for relatively tenured workers (who have been with their employer for at least a year).

Why does variation in days worked dominant intensive-margin fluctuations? We do not have a definitive answer, but offer one example. Suppose a firm's daily schedule is divided into two 8-hour shifts. Consider how a worker in this arrangement reacts if she experiences a decline in the marginal value of her time. If she is on the second shift, her colleagues leave after 8 hours, and the

[24] Since we can measure annual working time changes in each of a worker's final two quarters in the sample, we have two observations per worker. Also, note that we restrict this sample to include only workers who stay with the same employer across the year. This conforms to our treatment of the Veneto data in the next section. See the discussion there for more.

plant closes. The only way to acquire overtime is to replace an absent worker on the other shift. Unless she is able to work 16 hours in a day, this means an increase in days worked. This example of shiftwork is not merely a curiosity–Giaccone (2009) reports that at least 20 percent of the Italian workforce engages in shiftwork.

Lastly, we mention one final exercise to gauge how well we measure the coordinated component of time worked in particular. The nature of assembly-line production suggests that, in the auto assembly sector, labor supply is likely almost perfectly coordinated. Unfortunately, our industry information is somewhat limited; we do not measure assembly plant activity in Veneto. But, we can consult U.S. data. In the online Appendix, we investigate hours and days worked in the assembly plants of the three large U.S. vehicle manufacturers. We find that employer-wide variation in annual working days is a very good proxy for variation in annual total working hours.[25] This finding, though only suggestive, is encouraging. Recall that our approach is to identify the intertemporal elasticity of substitution using data on firm-wide log changes in working days. It appears the latter might capture reasonably well variation in working time at the employer level.

2.3 Earnings and working time in the Veneto panel

According to the theory of section 1, comparing the variances of idiosyncratic (i.e., within-firm) changes in working time and earnings provides valuable identifying information as to the degree of complementarity. In this section, we discuss how we measure these moments in the data.

Our main empirical analysis centers around a simple regression model. It is 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

$$\Delta \ln h_{ijt} = f\left(\boldsymbol{\chi}_{ijt}\right) + \phi_{jt}^{h} + \epsilon_{ijt}^{h}, \tag{17}$$

where χ_{ijt} is a vector of observables and ϕ_{jt} is a firm-year effect. Note that (17) applies to the subsample of workers who stay at a firm for consecutive years $t-1$ and t. By confining the regression's sample to these stayers, the changes in an individual's working time can be decomposed into changes in a single firm's demand for working time (ϕ_{jt}) and that individual's idiosyncratic labor input (ϵ_{ijt}). If workers switch employers, however, the change in their labor input also reflects cross-sectional differences in firms' labor demands.

The elements of $f\left(\boldsymbol{\chi}_{ijt}\right)$ are measured in year t and consist of a cubic in tenure and dummies for broad occupation (apprentice, clerk, manager, or worker). Tenure and occupation exhaust the set of worker characteristics in our panel.[26] Though somewhat limited, f's inclusion still helps to purge the data of (observable) persistent heterogeneity in work schedules. As discussed above, the model can then be used to infer complementarities based on how workers deviate from their usual

[25] See Bresnahan and Ramey (1994) and Ramey and Vine (2006) for details on the data. The online Appendix discusses how we infer days worked from these data.

[26] More than two-thirds of employees are classified as workers.

work schedules in response to various shocks.

The key term in (17) is the firm-year effect, ϕ^h_{jt}. This measures the change in firm j's average working time relative to the cross-sectional mean of firms in year t. We interpret ϕ^h_{jt} as reflective of idiosyncratic (demand or productivity) shocks to the firm *as a whole*. Accordingly, we take the variance of ϕ^h_{jt} as our measure of fluctuations in firm-wide working time. Recalling (12), this moment will be highly informative as to the value of φ.

It follows that the residual in (17) isolates variation across workers *within firm*. We can then pool the estimated ϵ^hs and calculate the variance of idiosyncratic working time changes using $\text{var}\left(\epsilon^h_{ijt}\right)$. Furthermore, we can repeat this exercise by replacing $\Delta \ln h$ (17) with the log change in earnings, $\Delta \ln W$,

$$\Delta \ln W_{ijt} = f\left(\boldsymbol{\chi}_{ijt}\right) + \phi^W_{jt} + \epsilon^W_{ijt}. \tag{18}$$

Comparing the variances of within-firm working time (using dispersion in ϵ^h_{ijt}) and earnings changes (using ϵ^W_{ijt}), the model can infer the degree of complementarity. Thus, the moment, $\text{var}\left(\epsilon^W_{ijt}\right)/\text{var}\left(\epsilon^h_{ijt}\right)$, will provide critical identifying information for our exercise.

In estimating (17), we measure fluctuations among workers who are attached to a firm for consecutive years. There is no single way of making this notion of attachment precise. Our baseline measure includes workers in the year-t cross section only if they are paid for at least one day in all months of the first quarter of year $t-1$ and in all months of the last quarter of year t. We refer to this latter sample as the 2-year stayers. This measure identifies workers who start and end a two-year period with the firm, while allowing extended absences during the period (due, perhaps, to parental leave or extended downtime at the plant).

There are at least two potential shortcomings of this measure. First, it is questionable whether maternity-related leaves of absence constitute the kind of variation envisioned by the model. Second, the 2-year stayers sample might include relatively peripheral workers, who may remain on the payroll but whose attachment to the firm is far weaker than the median worker. Based on our examination of the data, though, we do not see these as first-order concerns. Below, we report some results for the sub-sample of men, and the moments of interest are not too different. We also do not see much evidence of a large mass of peripheral workers. We would expect that these kinds of workers would be consistently paid for far fewer than 12 months per year. Yet among workers who are not paid for a full month or more in some year $t-1$, most are indeed paid for at least one day in each month of year t. This suggests that, in many cases, the extended absences we observe are indeed temporary episodes for relatively attached workers.

Nevertheless, although we estimate the model off the 2-year stayers, we do include in Table 3 empirical results for an alternative sample. The 12/12 stayers are defined as workers who are paid for at least one day in every month over years $t-1$ and t. This definition of stayers should minimize concerns about peripheral workers noted above.

Table 3 summarizes several key moments of the data. The first moment is the ratio of the variances of idiosyncratic (within-firm) earnings growth to working-time changes, computed according to var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$. The second is the ratio of variances of firm-wide average earnings growth to average working-time fluctuations, calculated according to var $\left(\phi_{jt}^{W}\right)$ /var $\left(\phi_{jt}^{h}\right)$. The third and fourth, respectively, are the standard deviations of the idiosyncratic $\left(\epsilon_{ijt}^{h}\right)$ and firm-wide $\left(\phi_{jt}^{h}\right)$ components of working-time fluctuations. The fifth is the projection of $\Delta \ln h_{ijt}$ on the log change in the daily wage *rate*. This coefficient is negative, echoing results from Abowd and Card (1989), among others. Though this finding in household surveys was sometimes interpreted as reflecting measurement error, we are less concerned about this error in our administrative data.

These five empirical moments are reasonably similar across the 12/12 and 2-year stayers, although working-time fluctuations are larger among the latter. This is not necessarily surprising, as we include in this sample employees who can experience longer non-working spells in years $t-1$ or t. Using either sample, though, the variance of earnings growth within the firm exceeds the variance in working time changes, that is, var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$ > 1 in each case (and indeed, the ratio exceeds 2 in each case). This will drive the model toward inferring a relatively high degree of complementarity.

To conclude our discussion, Table 4 offers a brief sensitivity analysis of the main moment of interest, var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$. For instance, it is natural to worry that many 2-year stayers consist of salaried workers, for whom days paid is a poor measure of working time. We try to assuage this concern by dropping workers who are paid for all working days during the year (full-year workers).[27] Predictably, the extent of compression in working time changes is less striking here, though var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$ remains well over 1.

In addition, we look more specifically at 2-year stayers in large firms (with more than 250 workers). The reason for this is that larger firms reportedly have greater leeway in Italy to conduct firm-level wage bargaining. Union-bargained minimum wages are more likely, for instance, to bind at smaller firms (Guiso, Pistaferri, and Schivardi 2005). Consistent with this claim, there is notably more idiosyncratic variation in earnings growth relative to working-time fluctuations at larger firms.

In Table 4, we also report the moment, var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$, across different industries. One might have expected to see greater compression of working time adjustments in manufacturing, if assembly-line production processes are more common here; the degree of complementarities would seem to be high. But though var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$ is higher in manufacturing than in most industries, it is not especially different from the average. Moreover, this ratio is actually much higher in finance. This might reflect two considerations. First, complementarities do not necessarily depend on the nature of the output (the industry). Rather, it hinges on the gains from specialization in the production process. These gains might well be large at a financial services firm. At a mutual fund, for instance, analysts typically specialize by sector or by region, and each analyst contributes recommendations to a client's portfolios. Presumably, these analysts must work at a similar pace in order to expand the fund's production. Second, industry is correlated with other important drivers

[27] Including paid vacation, full-year workers can record up to $52 \times 6 = 312$ days of paid work per year, amounting to 52 weeks of working six days.

of working time and earnings changes. For instance, finance is the least unionized sector in Italy (Visser, 2013). Thus, if unions' egalitarian objectives lead to a compression of earnings changes, this might reduce var $\left(\epsilon_{ijt}^{W}\right)$ /var $\left(\epsilon_{ijt}^{h}\right)$ in manufacturing relative to in finance.[28]

3 Model Estimation

The dynamic labor demand model consists of optimal working time (6); the employment choice (14)-(15); and the earnings bargain (9). This section seeks to identify six structural parameters of the model. They include: the elasticity of substitution across jobs, ρ; the elasticity of intertemporal substitution in labor supply, $1/\varphi$; and worker bargaining power, η. In addition, there are two parameters that drive dispersion within the firm. One is the variance of idiosyncratic (worker-specific) preference shocks, ξ. Also, as promised above, the quantitative model introduces fluctuations in worker-specific productivity, and so we must identify the variance of this object. Lastly, we will recover the variance of the firm-wide shocks, Z.The remainder of the parameters are fixed outside of the model.

In the first section, we discuss our parameterizations and the identification of the 6 parameters of interest. The next section presents results. In particular, we compare this model with predictions of the baseline theory of section 1.1 in order to highlight the limits of a model that abstracts from complementarities.[29] The final section applies the model to study the effects of certain policy interventions on individual and firm-wide labor input.

3.1 Matching moments

To solve the models, we first fix a few parameters outside of the model. First, we assume the firm-specific revenue shock, Z, follows a geometric AR(1),

$$\ln Z = \vartheta \ln Z_{-1} + \varepsilon^{Z}, \quad \varepsilon^{Z} \sim N\left(0, \sigma^{2}\right).$$

As we are not aware of estimates of ϑ for Italian plants, we set $\vartheta = 0.8$ based on annual U.S. plant-level evidence (Cooper and Haltiwanger, 2006; Foster, Haltiwanger, and Syverson, 2008). However, the variance, σ^2, of the innovation, ε^Z, is a parameter to be estimated, as discussed below.[30] Second, the severance cost, $\underline{c}$, amounts to one year of average earnings. This represents

[28] There were two notable labor laws introduced over our sample, but the moments of interest do not change much across different sub-samples. A 1997 law codified a limit on weekly hours of work of 48. But, most union agreements had put in place such a limit many years before. Moreover, the limit is typically interpreted as an upper bound on average weekly hours over the span of a few months; it is not hard limit on any one week's labor input (Treu, 2007). Another change in labor law was the 1992 dismantlement of the *scala mobile*. This was an indexation scheme that escalated earnings with inflation. However, the *scala mobile* applied uniformly to workers, so it shifted the mean of the distribution of earnings changes without affecting its variance.

[29] We do not estimate the theory of section 1.1 but our calibration in this case identifies values of the parameters that replicate key empirical moments.

[30] Our data does not include revenue, which is used by the aforementioned papers to estimate ϑ. One may argue, then, that we should also set σ in line with the evidence in these papers. However, we want to make sure our model

our attempt to synthesize multiple sources of separation costs in Italy.[31] Third, we set the cost of a new hire, $\bar{c}$, at 5 percent of annual earnings, which is in line with the range of estimates in the literature (again, we do not have direct evidence from Italy for this parameter).[32] Lastly, in the absence of much guidance, we set μ (the flow payoff of non-employment) so it is about 50 percent of average earnings. We have considered alternative values–as low as 25 percent and as high as 75 percent–but the effect on the moments of interest is pretty minimal.

One final set of parameters pertains to the number, M, of preference types, ξ, and the distribution of these types. It seems to us heroic to try to identify the shape of the distribution of types, given our data. Thus, we have simply assumed that the ξs are uniformly distributed and set $\lambda_\xi = 1/M$ for all ξ. As for the number of types, we have found that the precise choice of M does not have a material effect on the moments of interest *given* a choice for the variance of types. In other words, as long as we know ξ takes values between, say, 0.75 and 1.25, the choice of M has made little difference to our results. Intuitively, this follows from our standard abuse of the law of large numbers: a deterministic share λ_ξ of the workforce is type ξ, regardless of the size of the employer. In that case, it is possible to replicate a given variance of preference types with any number $M > 1$. And it is this variance which is relevant to our study, as it partly anchors the variances of earnings and working time changes within the firm.[33]

The remaining six parameters are estimated via method of simulated moments (MSM). To be exact, we choose ρ, φ, η, σ, and the variances of idiosyncratic preference and productivity shocks to minimize the distance between the model-implied moments and their analogues listed in Table 3 and derived from our sample of 2-year stayers.

We now discuss our choice of moments and offer some intuition for the mapping from the model's structural parameters to the listed moments. First, as outlined in section 1.2, the choice of the elasticity of substitution across tasks, ρ, bears strongly on the dispersion of working time

replicates the dispersion in employment growth in our Veneto data. Hence, we use σ as a free parameter to target this moment.

[31] In any employer-initiated separation, the firm owes a worker at least 1/13 of a year's earnings per year of tenure. (In our data, mean tenure is roughly 3.5 years.) In addition, in a collective dismissal (where at least 5 workers are separated), the firm must agree to a period of arbitration by the provincial labor office. If no resolution is reached after 2.5 months, the dismissal can take effect. Individual dismissals are often challenged in court. If a judge rules that the dismissal was unjustified, the firm must pay the worker the earings she would have received since the moment of dismissal. In addition, small firms (with less than 15 employees) owe an additional severance of 2-6 months of earnings, and larger firms must reinstate the worker. Based on its experience in employment law disputes, the Belgian law firm Laga reports that these individual dismissal costs add up to roughly 2 years of earnings. Our calibration roughly balances this against the smaller cost associated with collective dismissals.

[32] Our choice is informed by three estimates. First, Barron, et al's (1997) analysis of surveys of U.S. employers imply a hiring cost of almost 1 percent of annual earnings (see also Hagedorn and Manovskii, 2008). Second, estimates from human resource consultants in the U.S. point to a hiring cost of 3.5 percent of annual earnings (Hall and Milgrom, 2008). Lastly, a survey of French employers reveals a hiring cost of 3 percent of annual earnings (Abowd and Kramarz, 2003). The average of these three estimates is 2.6 percent. However, each of these refers to the cost of recruiting, and omits the cost of training. Abowd and Kramarz's (2003) French survey data suggest a training cost of at least 2.4 percent of annual earnings. Adding the recruiting and training costs together yields 5 percent.

[33] The same argument applies to the variance of the idiosyncratic (worker-specific) productivity shock, which is included in the estimated model. We also assume a uniform distribution of productivity.

changes within the firm relative to the dispersion in earings changes. Hence, ρ maps most clearly to $\mathrm{var}\left(\epsilon_{ijt}^{W}\right) \Big/ \mathrm{var}\left(\epsilon_{ijt}^{h}\right)$. Second, the choice of the worker's bargaining power, η, helps mediate the reaction of earnings to changes in working time following firm-wide shocks, which enables us to target $\mathrm{var}\left(\phi_{jt}^{W}\right)/\mathrm{var}\left(\phi_{jt}^{h}\right)$.

Next, the variances of idiosyncratic preference and productivity shocks are informed in particular by two moments. The size of preference (supply) shocks *relative* to the productivity (demand) innovations influences the covariance of working time and the daily wage rate. As we have mentioned, the negative covariance suggests a larger role for supply-side shifts. In addition, the absolute size of our measured idiosyncratic (worker-specific) movements in working time, $\mathrm{var}\left(\epsilon_{ijt}^{h}\right)$, offers further information about the absolute variances of these idiosyncratic shocks.

The final two parameters are φ and σ^2. As foreshadowed by (12), the intertemporal elasticity of working time, $1/\varphi$, influences the magnitude of working time fluctuations at the firm level, conditional on the variance of changes in firm-wide innovations, Z. This helps target $\mathrm{var}\left(\phi_{jt}^{h}\right)$. To identify, the variance, σ^2, of Z, we rely on the extensive margin adjustments. In particular, the size of firm-wide disturbances is greatly informed by the dispersion in employment growth across firms.

3.2 Results

Table 5 summarizes results. The top panel of Table 5 summarizes the empirical and model-generated moments. The fit of the model is excellent–the model is able to replicate the moments nearly exactly.[34] This goodness of fit should arguably be demanded from a just-identified model, but it is, still, the first obvious test to be passed, and the model does so. The bottom panel lists the MSM-based estimates of the structural parameters. There are three sets of remarks we would like to make regarding the parameter estimates.

First, the parameters, φ and η, may be compared with alternative estimates in the literature. The value of the intertemporal elasticity, $1/\varphi$, is somewhat higher than in the classic life-cycle analyses of, among others, MaCurdy (1981) and Altonji (1986), but is actually somewhat lower than the more recent life-cycle results presented in Pistaferri (2003). The latter paper uses another source of Italian data, which enables one to identify the response of working time to variation in survey respondents' expected wage changes. Viewed through the lens of our model, these expected wage movements reflect a combination of idiosyncratic (worker-specific) and firm-wide factors. The former should attenuate his estimate relative to ours *if* our two papers identify the *same* firm-wide variation. It is difficult to know if this latter condition obtains, though, as our models are non-nested.

[34]**NOTE: The reader will note that the model understates the covariance of earnings and working time changes. This is because the present draft omits idiosyncratic productivity shocks. Preliminary work reveals that, as argued above, this demand-side variation enables the model to replicate this moment very precisely.**

As for the bargaining parameter, η, our result suggests that earnings are relatively insensitive to changes in marginal products. As a point of comparison, we consider Roys (2014), who also estimates a model of dynamic labor demand in which earnings are set according to a Stole and Zwiebel bargain. Our two studies differ, though, in terms of what identifying information is brought to bear on η.. We identify η in large part by comparing the variance of firm-wide earnings growth to (firm-wide) working time movements. Roys lacks working time data but observes sales. Estimating his model using French firm-level data, Roys finds an estimate of η of nearly 0.5.

Second, we turn our attention to ρ. The estimation of this parameter enables the model to engage data on within-firm variation in earnings and working time, as summarized by the ratio of $\operatorname{var}\left(\epsilon_{ijt}^{W}\right)$ to $\operatorname{var}\left(\epsilon_{ijt}^{h}\right)$. A model with no complementarities is unable to this. To illustrate this more precisely, we perform a simple exercise. We take the model with *no* complementarities from section 1.1 and choose its five parameters to target the six moments in Table 5. We are particularly interested in the model's implication for $\operatorname{var}\left(\epsilon_{ijt}^{W}\right) \Big/ \operatorname{var}\left(\epsilon_{ijt}^{h}\right)$–is a model with no complementarities nonetheless able to replicate this? We answer, no. This model yields an estimate of this moment equal to 1.24, which is only slightly more than half its empirical counterpart. We view this as supportive of our claim that this moment does indeed help identify production complementarities, and thereby distinguish between competing models.

Third, as we discussed in section 3.1, the estimate of $1/\varphi$ is highly related to the size of firm-specific shocks, Z. In particular, given the observed variation in firm-wide working time, smaller fluctuations in Z imply a higher elasticity, $1/\varphi$. Table 6 illustrates this by reporting estimation results conditional on lower values of σ. Considering a value of $\sigma = 0.15$, the implied elasticity $1/\varphi$ increases to 0.7, which is in line with Pistaferri (2003) and thus near the top end of the range of estimates summarized by Chetty et al (2011).

Interestingly, Table 6 also shows how the change in $1/\varphi$ reverberates through the rest of the model, altering a number of structural parameters. Tracing the effect of changes in σ helps shed light on the workings of the model. If $1/\varphi$ is increased, for instance, the variance of firm-wide working time fluctuations increases relative to the variance of firm-wide earnings growth. Accordingly, from the data's perspective, the model now understates the latter relative to the former. To bring these two moments back into balance, the bargaining parameter, η, must be increased; this increases the pass-through from changes in firm productivity to earnings. But, this higher η also implies greater pass-through from idiosyncratic preference shifts to the earnings bargain. Consequently, the model now understates the extent of within-firm (idiosyncratic) dispersion in working time fluctuations relative to earnings growth. To correct this, the degree of complementarities must be eased, so ρ falls in absolute value.

3.3 Implications for policy analysis

A running theme of our analysis is that working time may react quite differently to idiosyncratic variation, as opposed to firm-wide events. In particular, the labor supply response to idiosyncratic variation can yield a downwardly biased estimate of a worker's willingness to substitute effort intertemporally. In this section, we illustrate this point quantitatively. Using our estimates of ρ and φ, we can undertake a simple experiment designed to mimic a randomized control trial. We randomly "treat" a fraction of a firm's workforce by increasing their "distaste", ξ, for working. As noted in section 1, this is isomorphic in our setting to reducing their marginal utility of income by, for instance, extending a lump-sum transfer. We then compute the change in working time. This result is compared to the outcome of an experiment in which we treat the full workforce of the firm.

The size of the treatment is chosen in the following way. First, we imagine a worker receives a lump-sum transfer. To anchor this in a relevant way, we choose it to be consistent with the size of the typical transfer in the Negative Income Tax experiments in the U.S., expressed as a share of the participant's pre-NIT market income. This approach implies a transfer of 37 percent of income.[35]. Second, we assume a marginal propensity to consume out of transitory income equal to 1/3, based on U.S. evidence in Johnson et al. (2006). This enables us to recover the typical change in consumption. Third, we map from this to the marginal utility of income, ℓ, via the static first-order condition for consumption, under the assumption that preferences are separable with respect to consumption, C, and leisure. In that case, if utility from consumption is isoelastic, $\frac{C^{1-\varsigma}}{1-\varsigma}$, the FOC implies $-\varsigma\Delta\ln C = \Delta\ln\ell$, where ς is the coefficient of relative risk aversion.[36] Following the literature and assuming $\varsigma = 2$, we can then infer the change in $\Delta\ln\ell$, which turns out to be 25 percent. This is equivalent, in our setting, to increasing the distaste, ξ, for working by this amount.

To illustrate the implications of the model with complementarities, we now "treat" 10 percent of the firm's workforce, reducing their marginal utility of income by 25 percent. These workers reduce their working time by 3.6 percent. Interestingly, the model implies a substantial difference between the marginal treatment effects depending on the size of the treated group. For instance, if we treat the full workforce (increasing everyone's marginal utility of income by 25 percent), working time falls by 9.2 percent, or 2.5 times as much as in the earlier experiment. This illustrates the extent to which idiosyncratic variation in the return from working fails to recover the true willingness to

[35] To recall, in a typical NIT trial, each household received a lump-sum benefit, or "guarantee", but a share, r, of this was reduced for each \$1 in market income. Moreover, in each of the four NIT trials, enrollment was restricted to families with income below a threshold, $\hat{y}$. In our calibration, we assume the treated worker has income at the midpoint of the eligible range, $\hat{y}/2$, where $\hat{y}$ is calculated as the mean across the cities weighting by participation (see Table 1 of Burtless, 1987). To compute the transfer in our calibration, we then start with the "guarantee", equal to the weighted average of the midpoints of the ranges in Burtless (1987), and apply a benefit reduction rate, r, of 50 percent (used in most cities) to an individual with market income $\hat{y}/2$. Thus, the transfer is what the individual would receive *conditional* on earning his pre-treatment income, $\hat{y}/2$. Unlike in the NIT, this transfer is *not* adjusted, according to the reduction rate r, based on the individual's labor supply response and subsequent changes in market earnings.

[36] This step has to be finessed slightly. As discussed in footnote 10, our derivation of the earnings bargain assumes that each worker is insured by a "large" family against idiosyncratic earnings changes. But if a transfer is made to a fraction of the workforce, it implies as an idiosyncratic change in ℓ only if families are not too large relative to the size of firms. In that case, one can imagine many families spread over (fewer) firms, so the receipt of a transfer by members of one family will hardly affect the work incentives of these members' firms. This means that our earnings bargain, derived under the assumption of complete markets, should perhaps be thought of as an attempt to approximate the case where markets are incomplete but where incompleteness is in the neighborhood of zero.

intertemporally substitute.[37]

Though a model with complementarities can offer such rich predictions, it is probably the case that it still overstates the reaction of working time to idiosyncratic variation. The published estimates of the effect of the NIT on annual hours werein the range of 6-8 percent, but these appear to have been driven by extended job search spells rather than reduced working time within employment relationships (Moffit, 1981; Robins and West, 1983). The reaction of working time among incumbent workers was smaller. There may be many reasons for this discrepancy between model and data, not the least of which is that we have not tried to fully replicate all dimensions of the NIT (see footnote 25). Staying within our framework, we could further dampen the response of working time to idiosyncratic variation by reducing ρ, and thus amplifying the coordination motive. But this will make it difficult to replicate the extent of within-firm dispersion in working time adjustments. Reconciling this latter dispersion with even smaller elasticities with respect to idiosyncratic variation is a task for future work.

4 Conclusion

TO BE ADDED.

5 References

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002). Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. Center for Economic Studies, U.S. Census Bureau, Technical paper No. TP-2002-06.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). "High Wage Workers and High Wage Firms," *Econometrica* 67 (2), 251 - 333.

Abowd, J.M. and D. Card (1989). "On the Covariance Structure of Earnings and Hours Changes," *Econometrica* 57.

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

Bartolucci, C. and F. Devicienti (2012). "Better Workers Move to Better Firms: A Simple Test to Identify Sorting." Carlo Alberto Notebooks, No. 259.

Bresnahan, Timothy and Valerie Ramey (1994). "Output fluctuations at the plant level," *Quarterly Journal of Economics*, August.

[37] Though our simulation focuses on incumbent workers (who remain with the firm), the NIT also affected participants' employment rates and job search strategies. This latter aspect of the NIT complicates the comparison of earnings across the simulations and NIT data. The reason is that job searchers may have accepted lower-wage employment, because of the benefit reduction rate (Robins and West, 1983). This will mask earnings dynamics among incumbents, for which the model has clearer predictions.

Burtler, Gary and Jerry Hausman (1978). "The Effect of Taxation on Labor Supply: Evaluating the Gary Negative Income Tax Experiment," *Journal of Political Economy*, December.

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*, vol 48 (3), pp. 943-72.

Card, D. (1990). "Labor supply with a minimum hours threshold," *Carnegie Rochester Conference on Public Policy*, 33.

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

Card, D., F. Devicienti, and A. Maida (2010). Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data. NBER Working Papers: 16192.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records," *The Quarterly Journal of Economics* 126 (2), 749-804.

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

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

Cooper, R., J. Haltiwanger, and J. L. Willis (2004). "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications," NBER WP 10297.

Cooper, R., J. Haltiwanger, and J. L. Willis (2007). "Search frictions: Matching aggregate and establishment observations," *Journal of Monetary Economics* 54, S56 - 78.

Cooper, R., and J. Haltiwanger (2006). "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 73(3), p. 611-633.

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

Demekas, D. G. (1994). Labor Market Institutions and Flexibility in Italy - A Critical Evaluation and Some International Comparisons. IMF Working Paper 94/30.

Erickson, C. and A. Ichino (1994). Wage differentials in Italy: market forces, institutions, and ination. NBER Working Paper 4922.

Giaccone, Mario (2009). "Working time in the European Union: Italy," EWCO Comparative Report.

Giannelli, G. C. and C. Braschi (2002). Reducing Hours of Work: Does Overtime Act as a Brake Upon Employment Growth? An Analysis by Gender for the Case of Italy. IZA Discussion

Paper 557.

Guiso, L., L. Pistaferri, and F. Schivardi (2005). "Insurance within the firm," *Journal of Political Economy.*

Hall, Robert E (1999). "Labor market frictions and employment fluctuations," in John Taylor and Michael Woodford (eds.) *Handbook of Macroeconomics*, 1999, North-Holland, pp. 1137-1170.

Hall, Robert E (2013). "Discussion of 'Understanding the Great Recession' by Christiano, Eichenbaum, and Trabant." NBER Conference on Lessons From the Financial Crisis for Monetary Policy, October.

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

Lazear, E. P. and K. L. Shaw (2008). "Wage Structure, Raises, and Mobility: An Introduction to International Comparisons of the Structure of Wages within and across Firms." A National Bureau of Economic Research Conference Report. Chicago and London: University of Chicago Press.

Manski, C. F. (1993). "Identification of endogenous social effects: The reflection problem," *Review of Economic Studies*, 60 (3), p. 531 -542.

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.

Pencavel, J. H. (1986). "Labor Supply of Men: A Survey", pp. 3 - 102. *Handbook of Labor Economics.*

Raffo, A., G. Llosa, L. Ohanian, and R. Rogerson (2012). Firing Costs and Labor Market Fluctuations: A Cross-Country Analysis. mimeo.

Ramey, Valerie and Daniel Vine (2006). "Tracking the source of the decline in GDP volatility: An analysis of the automobile industry," *American Economic Review*, December.

Rogerson, R. (2011). "Individual and Aggregate Labor Supply with Coordinated Working Times," *Journal of Money, Credit and Banking* 43, 7-37.

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.

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

Tattara, G. and M. Valentini (2010). "Turnover and Excess Worker Reallocation: The Veneto Labor Market between 1982-1996," *Labour* 24 (4), 474 - 500.

Tealdi, Cristina (2011). "Typical and atypical employment contracts: the case of Italy," mimeo.

6 Appendix: Derivations

This appendix derives the firm's optimal employment demand policy and earnings bargain under complementarities. The analysis draws on some preliminary results that are summarized below, and proved in the online Appendix.

6.1 The firm's problem

We first establish certain properties of the firm's optimal labor demand policy. To this end, an assumption and a few conjectures (to be verified later) are needed. The revenue function, $\hat{F}$, determines key aspects of the firm's behavior. It, in turn, hinges on the sign of the structural parameter, ρ.

Assumption 1 *The elasticity of substitution across tasks, ρ, is negative: $\rho < 0$. Therefore, tasks are complements in production.*

This has two immediate implications. First, the revenue function, $\hat{F}$, is concave, that is, the Hessian, $\nabla^2 \hat{F}(\mathbf{n}, Z)$, is negative definite. A corollary of this is that $\sum_i \sum_j \hat{F}_{ij}(\mathbf{n}, Z) < 0$, where $\hat{F}_{ij} \equiv \frac{\partial}{\partial n_j} \frac{\partial \hat{F}}{\partial n_i}$ is the effect of team-i labor on the marginal product of team-j.[38] Second, the revenue function, $\hat{F}$, is supermodular, in that $\frac{\partial}{\partial Z} \frac{\partial \hat{F}}{\partial n_j} > 0$ for any $j \in \{1, ..., M\}$ and $\frac{\partial^2}{\partial n_i n_j} \hat{F}(\mathbf{n}, Z) > 0$ for any $i \neq j$.

These restrictions on $\hat{F}$ are not still not quite sufficient to characterize the optimal labor demand policy. We also utilize the conjectures below, which assume that certain properties of $\hat{F}$ pass to period profit, π. They will be verified once a solution for the wage bargain is obtained.

Conjecture 1 *The concavity and supermodularity of $\hat{F}$ extend to period profit, π.*

In addition, we note that Z has the form of "team-neutral" technical change, in that the difference between marginal products, $\frac{\partial \hat{F}(\mathbf{n})}{\partial n_i} - \frac{\partial \hat{F}(\mathbf{n})}{\partial n_j}$, scales with Z. Therefore, a change in Z does not turn one team more productive than another: the *sign* of the difference, $\frac{\partial \hat{F}(\mathbf{n})}{\partial n_i} - \frac{\partial \hat{F}(\mathbf{n})}{\partial n_j}$, is independent of Z. This suggests to us that π will also display this property. We will indeed confirm this below. For now, we leave it as a conjecture.

Conjecture 2 *The firm-wide impulse, Z, is neutral in that, for teams $i \neq j$, the sign of the difference, $\frac{\partial \pi(\mathbf{n})}{\partial n_i} - \frac{\partial \pi(\mathbf{n})}{\partial n_j}$, is independent of Z.*

[38] This is straightforward to see if there are two inputs. In this case, negative definitness requires $\hat{F}_{11}, \hat{F}_{22} < 0$ and $\hat{F}_{11}\hat{F}_{22} - \hat{F}_{12}^2 > 0$. The latter, in turn, implies $\sqrt{\hat{F}_{11}\hat{F}_{22}} > \hat{F}_{12}$. Assume $\hat{F}_{12} > 0$ (or else the corollary follows trivially). We now have that

$$\sum_{i=1}^2 \sum_{j=1}^2 \hat{F}_{ij} = \hat{F}_{11} + \hat{F}_{22} + 2\hat{F}_{12} < \hat{F}_{11} + \hat{F}_{22} + 2\sqrt{\hat{F}_{11}\hat{F}_{22}}$$
$$= -\left\{ \left|\hat{F}_{11}\right| + \left|\hat{F}_{22}\right| - 2\sqrt{\left|\hat{F}_{11}\right|} \cdot \sqrt{\left|\hat{F}_{22}\right|} \right\} = -\left\{ \sqrt{\left|\hat{F}_{11}\right|} - \sqrt{\left|\hat{F}_{22}\right|} \right\}^2 < 0.$$

The next lemma provides a key intermediate result in the characterization of the optimal policy.

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

Proof. See online appendix. ■

We can now characterize the firm's optimal labor demand policy. We begin by assessing whether a firm, with start-of-period workforce N_{-1}, should separate from workers of (arbitrary) type ξ, taking as given the participation of the remaining types. As shown in the main text, a separation is made if the marginal value of labor, evaluated at N_{-1}, is less than the separation cost. Formally, this implies

$$\frac{\partial \pi\left(\boldsymbol{\lambda} N_{-1}, Z\right)}{\partial n_{\xi}}+\beta \int \Pi_{N}\left(N_{-1}, Z^{\prime}\right) \mathrm{d} G\left(Z^{\prime} | Z\right)<-\underline{c}. \tag{19}$$

The supermodularity of the problem implies a threshold, $\zeta_{\xi}\left(N_{-1}\right)$, such that a type-$\xi$ worker is separated for $Z<\zeta_{\xi}\left(N_{-1}\right)$.

The optimal employment level of the first-to-be separated type ξ is then dictated by the first-order condition,

$$\tilde{\Pi}_{n_{\xi}}^{-}\left(\mathbf{n}, N_{-1}, Z\right) \equiv \frac{\partial \pi\left(n_{\xi}, \boldsymbol{\lambda}_{/ \xi} N_{-1}, Z\right)}{\partial n_{\xi}}+\beta \int \Pi_{N}\left(N, Z^{\prime}\right) \mathrm{d} G\left(Z^{\prime} | Z\right)+\underline{c}=0, \tag{20}$$

where $\boldsymbol{\lambda}_{/\xi}$ is a $(M-1) \times 1$ vector of employment shares *exclusive* of the type-ξ share and $N=n_{\xi}+\Sigma_{x \neq \xi} \lambda_{x} N_{-1}$. Again by supermodularity, the left side of (20) is increasing in Z. Hence, the choice of type-ξ employment is, for given N_{-1}, $n_{\xi}=\nu_{\xi}\left(Z ; N_{-1}\right)$, with $\frac{\partial}{\partial Z} \nu_{\xi}>0$. Note that this first-order condition remains in effect as Z falls further below $\zeta_{\xi}\left(N_{-1}\right)$, holding fixed $\mathbf{n}_{/\xi} \equiv \boldsymbol{\lambda}_{/\xi} N_{-1}$.

We next consider how the firm reacts as Z declines further. The result is summarized in Proposition 2, whose proof is given below.

Proof of Proposition 2. As we consider still lower values of Z, the firm will separate from a(nother) type, denoted by $\hat{\xi} \neq \xi$, if the marginal value of that cohort falls below $-\underline{c}$,

$$\frac{\partial \pi\left(\nu_{\xi}\left(Z ; N_{-1}\right),\ \boldsymbol{\lambda}_{/ \xi} N_{-1},\ Z\right)}{\partial n_{\hat{\xi}}}+\beta \mathrm{E}\left[\Pi_{N}\left(N, Z^{\prime}\right) | Z\right]<-\underline{c},$$

where $N \equiv \nu_{\xi}\left(N_{-1}, Z\right)+\Sigma_{x \neq \xi} \lambda_{x} N_{-1}$.[39] (Since the FOC (20) remains in effect, we evaluate this derivative at the optimal size of team ξ, $\nu_{\xi}\left(Z ; N_{-1}\right)$.) To assess this expression, we note from (20) that, at $Z<\zeta_{\xi}\left(N_{-1}\right)$, we can infer the discounted expected marginal value of labor, $\beta \mathrm{E}\left[\Pi_{N}\left(N, Z^{\prime}\right) | Z\right]$, from $-\underline{c}-\frac{\partial \pi\left(n_{\xi}, \boldsymbol{\lambda}_{/ \xi} N_{-1}, Z\right)}{\partial n_{\xi}}$. Substituting this into the preceding inequality, we see that the firm separates from type $\hat{\xi}$ only if

$$\frac{\partial \pi\left(\nu_{\xi}\left(Z ; N_{-1}\right),\ \boldsymbol{\lambda}_{/ \xi} N_{-1},\ Z\right)}{\partial n_{\hat{\xi}}}<\frac{\partial \pi\left(\nu_{\xi}\left(Z ; N_{-1}\right), \boldsymbol{\lambda}_{/ \xi} N_{-1}, Z\right)}{\partial n_{\xi}}. \tag{21}$$

[39] To be clear, the derivative here is taken conditional on $n_{\xi}=\nu_{\xi}\left(N_{-1}, Z\right)$ and is then evaluated at $\mathbf{n}_{/\xi}=\boldsymbol{\lambda}_{/\xi} N_{-1}$.

Since workers of types ξ and $\hat{\xi}$ (who remain at the firm at this instant) contribute equally to the future value of the firm, they are differentiated only by their contribution to current profit. It follows that type $\hat{\xi}$ is separated if its marginal profitability falls below that of the already-separated type, ξ.

It remains to confirm that (21) does indeed take hold as Z falls further below $\zeta_{\xi}(N_{-1})$. Under Conjecture 2, the direct effect of Z on the marginal profitability of any two teams offsets, in that the sign of the difference, $\frac{\partial\pi(\mathbf{n})}{\partial n_{\hat{\xi}}} - \frac{\partial\pi(\mathbf{n})}{\partial n_{\xi}}$ for fixed $\mathbf{n}$, is independent of Z. Hence, what distinguishes the two terms in (21) is how the indirect effect of Z operates. On the one hand, a fall in Z reduces ν_{ξ}, which decreases $\frac{\partial\pi}{\partial n_{\hat{\xi}}}$. On the other hand, since ν_{ξ} decreases, $\frac{\partial\pi}{\partial n_{\xi}}$ is resuscitated. Thus, it must be that, at some Z, $\frac{\partial\pi}{\partial n_{\hat{\xi}}}$ falls to reach $\frac{\partial\pi}{\partial n_{\xi}}$. Equivalently, at Z approaches some sufficiently low threshold $\zeta_{\hat{\xi}}(N_{-1})$, $\frac{\partial\pi}{\partial n_{\hat{\xi}}}\Big/\frac{\partial\pi}{\partial n_{\xi}} \to 1$.

At this juncture, when separations of $\hat{\xi}$-workers begins, the firm continues to separate from type-ξ workers. This follows immediately from the supermodularity of $\tilde{\Pi}$. If $n_{\hat{\xi}}$ is reduced, the marginal value of type-ξ labor declines. It follows that n_{ξ} must be reduced to enforce the FOC (20).

Summarizing, there exists functions $\zeta_{\hat{\xi}}(N_{-1}) < \zeta_{\xi}(N_{-1})$ such that the firm separates from *both* type ξ and $\hat{\xi}$ workers if $Z < \zeta_{\hat{\xi}}(N_{-1})$. Since team ξ is the first team to separate, we refer to it as the rank-1 team and denote its type by ξ_1. Similarly, we refer to $\hat{\xi}$ as the rank-2 team and set $\hat{\xi} \equiv \xi_2$. It is straightforward to repeat this analysis for the other types, thereby establishing the ordering of teams from rank 1 to rank M. ■

In line with our notation from the Proposition, we will, in what follows, refer to an arbitrary team as team-ξ if its rank within the firm is unimportant in the context of the discussion. Otherwise, we will refer to a team as team-j, where j denotes its rank.

We now return to the question of when the firm will hire. The firm hires if the marginal value of a worker, evaluated at $N = N_{-1}$, exceeds $\bar{c}$. Hence, the firm hires if $\frac{\partial\Pi^{-}(\mathcal{N},Z)}{\partial\mathcal{N}}\Big|_{\mathcal{N}=N_{-1}} > \bar{c}$.[40] If the firm chooses not to separate later in the period, it will bargain and produce with $N = \mathcal{N}$ workers. Therefore, in this case, $\frac{\partial\Pi^{-}(\mathcal{N},Z)}{\partial\mathcal{N}}\Big|_{\mathcal{N}=N_{-1}}$ is seen from (14) to be $\frac{\partial\pi(\boldsymbol{\lambda}\mathcal{N},Z)}{\partial\mathcal{N}} + \beta D(\mathcal{N},Z)$, where I have defined $D(\mathcal{N},Z) \equiv \int \Pi_N(\mathcal{N},Z')\,\mathrm{d}G(Z'|Z)$.

Now, to confirm that the firm will in fact not separate from any workers after the draws of ξ, Proposition 1 implies that it is sufficient to verify that it does not wish to separate from the first team in the order, whose taste for work is represented by ξ_1. The latter condition can be inferred from (16). Therefore, if a firm hires, then it will *not* separate if (and only if)

$$\bar{c} < \sum_j \lambda_j \frac{\partial\pi(\mathbf{n},Z)}{\partial n_j} + \beta D(\mathcal{N},Z) \quad \Rightarrow -\underline{c} < \frac{\partial\pi(\mathbf{n},Z)}{\partial n_1}\bigg|_{\mathbf{n}=\boldsymbol{\lambda}\mathcal{N}} + \beta D(\mathcal{N},Z), \tag{22}$$

[40] And if it hires, it will do so until the marginal value of labor is driven down to $\bar{c}$, which gives the first-order condition.

where I have used the chain rule, $\frac{\partial \pi(\boldsymbol{\lambda}\mathcal{N},Z)}{\partial \mathcal{N}} = \sum_\xi \lambda_\xi \frac{\partial \pi(\mathbf{n},Z')}{\partial n_\xi}$. Clearly, since the forward value is held in common, this collapses to the requirement that

$$\bar{c} < \sum_j \lambda_j \frac{\partial \pi(\mathbf{n},Z)}{\partial n_j} \Rightarrow -\underline{c} < \left. \frac{\partial \pi(\mathbf{n},Z)}{\partial n_1} \right|_{\mathbf{n}=\boldsymbol{\lambda}\mathcal{N}}.$$

This will hold if $\bar{c}$ is sufficiently large relative to $-\underline{c}$ and if the dispersion in ξ is not too great. The latter will guarantee that the mean $\sum_\xi \lambda_\xi \frac{\partial \pi(\mathbf{n},Z')}{\partial n_\xi}$ is never too much lower than $\frac{\partial \pi(\boldsymbol{\lambda}\mathcal{N},Z)}{\partial n_1}$.

Assuming (22), the labor demand policy of the firm is to separate according to Proposition 1 and to hire only if $Z > \zeta_0(N_{-1})$, where ζ_0 is the critical point at which the marginal value of labor, assessed at N_{-1}, is just equal to the cost of hiring. Formally, then, $\zeta_0(N_{-1})$ solves the indifference relation, $\left. \frac{\partial \Pi^-(\mathcal{N},\zeta_0(N_{-1}))}{\partial \mathcal{N}} \right|_{\mathcal{N}=N_{-1}} = \bar{c}$.

6.2 Wage bargaining

To recall, Stole and Zwiebel protocol yields a surplus sharing rule,

$$\mathcal{W}_\xi - \mathcal{U} = \eta\left(\mathcal{W}_\xi - \mathcal{U} + \mathcal{J}_\xi + \underline{c}\right). \tag{23}$$

Proposition 1 asserts that this sharing rule yields the wage bargain (9) in the main text. The proof of this proposition is the following.

Proof of Proposition 1. We first determine the value of a worker to the firm. Consider a firm with an array of workers n, and assume that the cost of hiring any new workers among the n has been sunk. The value of this firm can be assembled by combining (14) and (15) to obtain

$$\begin{aligned} \Pi(N_{-1},Z) = \mathcal{B}\left[\tilde{\Pi}(\mathcal{N},\mathbf{s},\ N_{-1},Z)\right] \equiv \max_{\mathcal{N},\mathbf{s}} \tilde{\Pi}(\mathcal{N},\mathbf{s},\ N_{-1},Z) \\ \equiv \max_{\mathcal{N},\mathbf{s}} \left\{-\bar{c}\max\{0,\mathcal{N}-N_{-1}\} + \pi(\boldsymbol{\lambda}\mathcal{N}-\mathbf{s},Z) - \underline{c}\sum_{\xi\in\mathcal{X}} s_\xi + \beta \mathrm{E}\left[\Pi(N,Z')\,|Z\right]\right\}, \end{aligned} \tag{24}$$

where $\mathcal{B}$ is the Bellman operator and $N \equiv \sum_{\xi\in\mathcal{X}} n_\xi = \sum_{\xi\in\mathcal{X}} \lambda_\xi \mathcal{N} - s_\xi$. Using (24), we can then derive the marginal contribution of a type-ξ worker by differentiating with respect to n_ξ. Note that, in accordance with (23), this is defined *gross* of the separation cost $\underline{c}$ that the worker saves the firm by remaining within the match. We have that,

$$\mathcal{J}_\xi(\mathbf{n},Z) \equiv \pi_\xi(\mathbf{n},Z) + \beta \int \Pi_N\left(N,Z'\right) \mathrm{d}G\left(Z'|Z\right), \tag{25}$$

where $\pi_\xi(\mathbf{n},Z)$ is the marginal effect of type-ξ labor on period profit:

$$\pi_\xi(\mathbf{n},Z) \equiv \frac{\partial \hat{F}(\mathbf{n},Z)}{\partial n_\xi} - \left[W_\xi(\mathbf{n},Z) + \frac{\partial W_\xi(\mathbf{n},Z)}{\partial n_\xi} n_\xi + \sum_{x\neq\xi} \frac{\partial W_x(\mathbf{n},Z)}{\partial n_\xi} n_x\right]. \tag{26}$$

To characterize (25), we begin by observing that the expected value of the firm can be decomposed according to the expression,[41]

$$\begin{array}{rc} & \int \Pi\left(N, Z'\right) \mathrm{d}G \\ = & \sum_{j=1}^{M} \int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \Pi^j\left(N, Z'\right) \mathrm{d}G \\ & +\int_{\zeta_1(N)}^{\zeta_0(N)} \Pi^0\left(N, Z'\right) \mathrm{d}G+\int_{\zeta_0(N)}^{\infty} \Pi^{+}\left(N, Z'\right) \mathrm{d}G. \end{array}$$

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 $i \leq j$. The ordering of the types from 1 to M follows their ranking described in Proposition 1, and the sequence, $\left\{\zeta_j\right\}_{j=1}^{M}$, represents the corresponding thresholds governing separation.[42] 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 now differentiate the left side of the preceding expression with respect to N. The optimal labor demand policy implies that, at any threshold ζ_j with $j \in \{1, M\}$, the firm is indifferent to separating from the marginal type ξ_j. It follows that $\Pi^{j-1}\left(N, \zeta_j(N)\right) = \Pi^j\left(N, \zeta_j(N)\right)$. Likewise, at ζ_0, the firm is indifferent between freezing and hiring, so $\Pi^0\left(N, \zeta_0(N)\right) = \Pi^+\left(N, \zeta_0(N)\right)$. Therefore, by Leibniz's rule, we have

$$\begin{array}{rc} & \int \Pi_N\left(N, Z'\right) \mathrm{d}G \\ = & \sum_{j=1}^{M} \int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \Pi_N^j\left(N, Z'\right) \mathrm{d}G \\ & +\int_{\zeta_1(N)}^{\zeta_0(N)} \Pi_N^0\left(N, Z'\right) \mathrm{d}G+\int_{\zeta_0(N)}^{\infty} \Pi_N^{+}\left(N, Z'\right) \mathrm{d}G. \end{array}$$

Next, the Envelope theorem reveals the marginal value of labor in states where separations or hires are undertaken. In the case of hires, we apply the Envelope theorem to (15), which yields

$$\Pi_N^{+}\left(N, Z'\right) = \bar{c}. \tag{27}$$

To treat the case of separations, return to (14) and consider the state in which the firm separates only from type-1 labor (that is, workers with taste ξ_1). The value of the firm in this state is

$$\Pi^1\left(N, Z'\right) = \pi\left(\nu_1\left(N, Z'\right), \boldsymbol{\lambda}_{/1} N,\ Z'\right) - \underline{c}\left[\lambda_1 N - \nu_1\left(N, Z'\right)\right] + \beta \int \Pi\left(N', Z''\right) \mathrm{d}G,$$

where $\nu_1\left(N, Z'\right)$ denotes the optimal choice of type-1 labor conditional on adjusting; $\boldsymbol{\lambda}_{/1} \equiv \left(\lambda_2, ..., \lambda_M\right)$ is the vector of labor shares *exclusive* of type-1 labor; and $N' = \nu_1\left(N, Z'\right) + \sum_{i=2} \lambda_i N$. By the Envelope theorem, we now have that

$$\Pi_N^1\left(N, Z'\right) = -\lambda_1 \underline{c} + \sum_{i=2} \lambda_i \widehat{\mathcal{J}}_{i,1}\left(N, Z'\right), \tag{28}$$

[41] For the sake of brevity, we will often abbreviate $\mathrm{d}G\left(Z'|Z\right)$ by $\mathrm{d}G$.

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

where

$$\widehat{\mathcal{J}}_i^1\left(N, Z'\right) \equiv \frac{\partial \pi\left(\nu_1\left(N, Z'\right), \boldsymbol{\lambda}_{/1} N,\, Z'\right)}{\partial n_i} + \beta \int \Pi_{N'}\left(N', Z''\right) \mathrm{d}G\left(Z'' | Z'\right).$$

Equation (28) says that the marginal value of a worker is a weighted average of the separation cost, which applies if the worker is revealed to be type 1, and the marginal values of labor types $i > 1$, who are not separated. This result stems from the fact that the marginal value of labor, $\Pi_N^1\left(N, Z'\right)$, is assessed at the start of the period *before types are revealed.* Hence, an additional worker exacts a cost on the firm, $-\underline{c}$, with probability λ_1 but otherwise contributes a marginal increase in firm value $\widehat{\mathcal{J}}_i^1\left(N, Z'\right)$ with probability λ_i. Generalizing from (28), we have that for any state $Z \in \left[\zeta_{j+1}\left(N\right), \zeta_j\left(N\right)\right]$ with $j \geq 1$,

$$\Pi_N^j\left(N, Z'\right) = -\Lambda_j \underline{c} + \textstyle\sum_{i=j+1}^M \lambda_i \widehat{\mathcal{J}}_i^j\left(N, Z'\right), \tag{29}$$

where $\Lambda_j \equiv \sum_{i=1}^j \lambda_i$ and

$$\widehat{\mathcal{J}}_i^j\left(N, Z'\right) \equiv \frac{\partial \pi\left(\boldsymbol{\nu}_j\left(N, Z'\right), \boldsymbol{\lambda}_{/j} N,\, Z'\right)}{\partial n_i} + \beta \int \Pi_{N'}\left(N', Z''\right) \mathrm{d}G. \tag{30}$$

We now place (27) and (29) into the expression for $\int \Pi_N\left(N, Z'\right) \mathrm{d}G$. The resulting expression can be simplified slightly by first noting that

$$\sum_{j=1}^M \Lambda_j \left[G\left(\zeta_j\left(N\right)|Z\right) - G\left(\zeta_{j+1}\left(N\right)|Z\right)\right] = \sum_{j=1}^M \lambda_j G\left(\zeta_j\left(N\right)|Z\right).$$

Accordingly, we obtain

$$\begin{array}{rc} & \int \Pi_N\left(N, Z'\right) \mathrm{d}G\left(Z'|Z\right) \\ = & -\underline{c} \sum_{j=1}^M \lambda_j G\left(\zeta_j\left(N\right)|Z\right) + \sum_{j=1}^M \sum_{i=j+1}^M \lambda_i \int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \widehat{\mathcal{J}}_{i,j}\left(N, Z'\right) \mathrm{d}G \\ & + \int_{\zeta_1(N)}^{\zeta_0(N)} \Pi_N^0\left(N, Z'\right) \mathrm{d}G\left(Z'|Z\right) + \bar{c}\left(1 - G\left(\zeta_0\left(N\right)\right)|Z\right). \end{array} \tag{31}$$

The last step is to assess the marginal value of labor in the "freezing" regime, $\Pi_N^0\left(N, Z'\right)$. This is obtained by forwarding (24) one period, setting $s'_\xi = 0\ \forall\ \xi$ and $\mathcal{N} = N_{-1}$, noting that $\mathbf{n}' = \mathbf{n} = \boldsymbol{\lambda} N$ in this case, and differentiating with respect to N. We obtain

$$\Pi_N^0\left(N, Z'\right) = \pi_N\left(\boldsymbol{\lambda} N, Z'\right) + \beta \int \Pi_N\left(N, Z''\right) \mathrm{d}G\left(Z'' | Z'\right), \tag{32}$$

where

$$\pi_N\left(\boldsymbol{\lambda} N, Z'\right) \equiv \frac{\partial \hat{F}\left(\boldsymbol{\lambda} N, Z'\right)}{\partial N} - \sum_j \frac{\partial}{\partial N}\left[W_j\left(\boldsymbol{\lambda} N, Z'\right) \lambda_j N\right].$$

is the marginal effect of current employment on next-period profit. It is helpful to write out this latter term. Setting $\boldsymbol{\lambda}N = \mathbf{n}$, using the chain rule, and reorganizing, we have

$$\begin{aligned}
\left.\frac{\partial \hat{F}(\boldsymbol{\lambda}N,Z')}{\partial N}\right|_{\boldsymbol{\lambda}N=\mathbf{n}} &= \sum_j \lambda_j \frac{\partial \hat{F}(\mathbf{n},Z')}{\partial n_j} \\
\sum_j \frac{\partial}{\partial N}\left[W_j\left(\boldsymbol{\lambda}N,Z'\right)\lambda_j N\right]\Big|_{\boldsymbol{\lambda}N=\mathbf{n}} &= \sum_j \lambda_j \left[W_j\left(\mathbf{n},Z'\right) + \frac{\partial W_j(\mathbf{n},Z')}{\partial n_j} n_j + \mathcal{D}_j\left(\mathbf{n},Z'\right)\right],
\end{aligned}$$

where $\mathcal{D}_j\left(\mathbf{n},Z'\right) \equiv \sum_{i\neq j} \frac{\partial W_i(\mathbf{n},Z')}{\partial n_j} n_i$. Now returning to (25), evaluating at $\mathbf{n} = \boldsymbol{\lambda}N$, taking a weighted average of J_ξ across teams, and comparing the result to (32), one can confirm that

$$\sum_j \lambda_j \mathcal{J}_j\left(\boldsymbol{\lambda}N,Z'\right) = \Pi_N^0\left(N,Z'\right).$$

Making this substitution in (31) and then inserting the result into (25),

$$\begin{aligned}
\mathcal{J}_\xi\left(\mathbf{n},Z\right) &\equiv \pi_\xi\left(\mathbf{n},Z\right) \\
-\beta\underline{c}\sum_{j=1}^M \lambda_j G\left(\zeta_j\left(N\right)|Z\right) &+ \beta\sum_{j=1}^M\sum_{i=j+1}^M \lambda_i \int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \widehat{\mathcal{J}}_i^j\left(N,Z'\right)\mathrm{d}G \\
+\beta\int_{\zeta_1(N)}^{\zeta_0(N)} \sum_{j=1}^M \lambda_j \mathcal{J}_j\left(\boldsymbol{\lambda}N,Z'\right)\mathrm{d}G &+ \beta\bar{c}\left(1 - G\left(\zeta_0\left(N\right)\right)|Z\right).
\end{aligned} \tag{33}$$

We next characterize the surplus to the worker from participating in the firm. The instantaneous return on working is given by earnings less the cost of exerting effort, $W_\xi\left(\mathbf{n},Z\right) - g_\xi\left(\mathbf{n}\right)$, where $g_\xi\left(\mathbf{n}\right) \equiv \xi\frac{h_\xi(\mathbf{n})^{1+\varphi}}{1+\varphi}$ is the disutility from labor to a worker of type ξ. In the next period, a worker may taken on any one of the M types. Conditional on drawing some $\xi_j \in X$, the worker may be separated if $Z' < \zeta_j\left(N\right)$. Accordingly, we have that the present value of working at a firm $\left(\mathbf{n},Z\right)$ is

$$\mathcal{W}_\xi\left(\mathbf{n},Z\right) = \begin{array}{c} W_\xi\left(\mathbf{n},Z\right) - g_\xi\left(\mathbf{n}\right) \\ +\beta\sum_{i=1}^M \lambda_i \left[\ \mathbb{E}_{Z'}\left[\sigma_i\left(N,Z'\right)\cdot\mathcal{U} + \left(1-\sigma_i\left(N,Z'\right)\right)\cdot\mathcal{W}_i\left(\mathbf{n}',Z'\right)\right]\ \right], \end{array}$$

where σ_j is the probability that an individual worker on the team is separated. Rearranging terms, we may write the preceding expression in terms of the surplus from market work, $\mathcal{S}_\xi^W\left(\mathbf{n},Z\right) \equiv \mathcal{W}_\xi\left(\mathbf{n},Z\right) - \mathcal{U}$,

$$\mathcal{S}_\xi^W\left(\mathbf{n},Z\right) = W_\xi\left(\mathbf{n},Z\right) - g_\xi\left(\mathbf{n}\right) - r\mathcal{U} + \beta\sum_{i=1}^M \lambda_i \mathbb{E}_{Z'}\left[\left(1-\sigma_i\left(N,Z'\right)\right)\mathcal{S}_i^W\left(\mathbf{n}',Z'\right)\ \right], \tag{34}$$

where $r \equiv 1 - \beta$.

To assess the forward term in (34), consider the expected continuation value in the event that the worker is a member of the team with rank 1 (that is, her "taste" is ξ_1). By surplus sharing, we can write this as

$$\mathbb{E}_{Z'}\left[\left(1-\sigma_1\left(N,Z'\right)\right)\mathcal{S}_1^W\left(\mathbf{n}',Z'\right)\ \right] = \frac{\eta}{1-\eta}\mathbb{E}_{Z'}\left[\left(1-\sigma_1\left(N,Z'\right)\right)\cdot\left[\text{firm's surplus}\left(N,Z'\right) + \underline{c}\right]\ \right].$$

The firm's surplus in this expression can be derived as follows. By Proposition 2, the worker can be separated in any state $Z' < \zeta_1(N)$. Moreover, as shown in (20), the marginal value of type-1 labor in any such state must be $-\underline{c}$. If, on the other hand, $Z' \in [\zeta_1(N), \zeta_0(N)]$, the employer "freezes" and earns $\mathcal{J}_1(\boldsymbol{\lambda}N, Z')$, given by (25). Lastly, if $Z' > \zeta_0(N)$, the firm hires, and the marginal value of labor must be $\bar{c}$. Putting these pieces together and noting that $\sigma_1(N, Z') = 0$ if $Z' \geq \zeta_1(N)$, we have

$$\mathbb{E}_{Z'}\left[\left(1-\sigma_1\left(N,Z'\right)\right)\mathcal{S}_1^W\left(\mathbf{n}',Z'\right)\ \right] = \frac{\eta}{1-\eta}\left\{\int_{\zeta_1(N)}^{\zeta_0(N)} \mathcal{J}_1\left(\boldsymbol{\lambda}N,Z'\right)\mathrm{d}G + \int_{\zeta_0(N)} \bar{c}\mathrm{d}G + \underline{c}\int_{\zeta_1(N)} \mathrm{d}G\right\}. \tag{35}$$

Next, consider the expected continuation value of a rank-2 team member. The only difference with respect to (35) is that the worker is not subject to separation in states $Z' \in (\zeta_2(N), \zeta_1(N))$. Therefore, as the firm "freezes" rank-2 labor in this state, it earns a surplus (gross of $\underline{c}$) equal to $\widehat{\mathcal{J}}_2^1(N, Z')$, as given by (30). Applying this line of reasoning to higher-ranked teams, we see that, for any team i,

$$\begin{gathered}\mathbb{E}_{Z'}\left[\left(1-\sigma_i\left(N,Z'\right)\right)\mathcal{S}_i^W\left(\mathbf{n}',Z'\right)\ \right] \\ = \tfrac{\eta}{1-\eta}\left\{\textstyle\sum_{j=1}^{i-1}\int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \widehat{\mathcal{J}}_i^j\left(N,Z'\right)\mathrm{d}G + \int_{\zeta_1(N)}^{\zeta_0(N)} \mathcal{J}_i\left(\boldsymbol{\lambda}N,Z'\right)\mathrm{d}G + \int_{\zeta_0(N)} \bar{c}\mathrm{d}G + \underline{c}\int_{\zeta_i(N)} \mathrm{d}G\right\}.\end{gathered}$$

Substituting these results into (34) and collecting terms,

$$\begin{gathered}\mathcal{S}_\xi^W\left(\mathbf{n},Z\right) = W_\xi\left(\mathbf{n},Z\right) - g_\xi\left(\mathbf{n}\right) - r\mathcal{U} \\ +\beta\tfrac{\eta}{1-\eta}\left\{\begin{array}{c}\underline{c}\sum_{i=1}^M \lambda_i\left[1-G\left(\zeta_i\left(N\right)|Z\right)\right] + \sum_{i=1}^M \lambda_i \sum_{j=1}^{i-1}\int_{\zeta_{j+1}(N)}^{\zeta_j(N)} \widehat{\mathcal{J}}_i^j\left(N,Z'\right)\mathrm{d}G \\ +\sum_{i=1}^M \int_{\zeta_1(N)}^{\zeta_0(N)} \lambda_i\mathcal{J}_i\left(\boldsymbol{\lambda}N,Z'\right)\mathrm{d}G + \bar{c}\left[1-G\left(\zeta_0\left(N\right)|Z\right)\right]\end{array}\right\}\end{gathered} \tag{36}$$

We are now prepared to present the wage bargain. Recalling the surplus-splitting rule, $\mathcal{S}_j^W(\mathbf{n},Z) = \frac{\eta}{1-\eta}[\mathcal{J}_j(\mathbf{n},Z) + \underline{c}]$, substituting from (33) and (36),and canceling terms, we have that, for the type-j team,

$$W_j\left(\mathbf{n},Z\right) = \frac{\eta}{1-\eta}\left[\pi_j\left(\mathbf{n},Z\right) + r\underline{c}\right] + g_j\left(\mathbf{n}\right) + \mu,$$

where $\mu \equiv r\mathcal{U}$. Substituting for $\pi_j(\mathbf{n},Z)$ using (26) and rearranging, this becomes

$$W_j\left(\mathbf{n},Z\right) = \eta\left\{\frac{\partial\hat{F}\left(\mathbf{n},Z\right)}{\partial n_j} - \sum_{i=1}^M \frac{\partial W_i\left(\mathbf{n},Z\right)}{\partial n_j}n_i + r\underline{c}\right\} + (1-\eta)\left(g_j\left(\mathbf{n}\right)+\mu\right). \tag{37}$$

Cahuc, Marque, and Weismer (2008) demonstrate how to solve a system of partial differential equations like that in (37). The solution is

$$W_j\left(\mathbf{n},Z\right) = \eta r\underline{c} + (1-\eta)\mu + \int_0^1 \vartheta^{\frac{1-\eta}{\eta}}\left\{\frac{\partial\hat{F}\left(\vartheta\mathbf{n},Z\right)}{\partial n_j} + \frac{1-\eta}{\eta}g_j\left(\vartheta\mathbf{n}\right)\right\}\mathrm{d}\vartheta.$$

The marginal product, $\frac{\partial \hat{F}(\vartheta \mathbf{n},Z)}{\partial n_j}$, can be computed using (7), which yields $\frac{\partial \hat{F}(\vartheta \mathbf{n},Z)}{\partial n_j} = \vartheta^{-\frac{(1-\alpha)(1+\varphi)}{\varphi+1-\alpha}} \frac{\partial \hat{F}(\mathbf{n},Z)}{\partial n_j}$. Next, using the hours bargain (6), we have $h_j(\vartheta \mathbf{n}) = \vartheta^{-\frac{1-\alpha}{\varphi+1-\alpha}} h_j(\mathbf{n})$, and $g_j(\vartheta \mathbf{n}) = \vartheta^{-\frac{(1-\alpha)(1+\varphi)}{\varphi+1-\alpha}} g_j(\mathbf{n})$. Substituting in and integrating, we have

$$W_j(\mathbf{n},Z) = \eta \left[A \frac{\partial \hat{F}(\mathbf{n}, Z)}{\partial n_j} + r\underline{c} \right] + (1-\eta)(A g_j(\mathbf{n}) + \mu), \tag{38}$$

where $A \equiv \frac{\varphi+1-\alpha}{(\varphi+1)(1-\eta(1-\alpha))-\alpha}$ and, using (7) with $\Omega(\mathbf{n}) \equiv \left(\sum_{x \in \mathcal{X}} (n_x^{\varphi}/x)^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha-\rho}{\rho} \frac{1}{\varphi+1-\alpha}}$,

$$\frac{\partial \hat{F}(\mathbf{n},Z)}{\partial n_j} = \frac{\varphi}{\varphi+1-\alpha} (\alpha Z)^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\varphi+1} \xi_j^{\frac{-\rho}{\varphi+1-\rho}} n_j^{-(1+\varphi)\frac{1-\rho}{\varphi+1-\rho}}.$$

Lastly, rewriting the earnings bargain as

$$W_\xi(\mathbf{n},Z) = \frac{\varphi}{\varphi+1} \frac{(\varphi+1-\alpha)+\eta\alpha}{(1-\eta)(\varphi+1-\alpha)+\eta\alpha\varphi} (\alpha Z)^{\frac{\varphi+1}{\varphi+1-\alpha}} \Omega(\mathbf{n})^{\varphi+1} \xi^{\frac{-\rho}{\varphi+1-\rho}} n_\xi^{-\frac{(1+\varphi)(1-\rho)}{\varphi+1-\rho}} + \eta r\underline{c} + (1-\eta)\mu,$$

substituting into period profit and simplifying, we have that

$$\begin{aligned} &\pi(\mathbf{n},Z) \equiv \hat{F}(\mathbf{n},Z) - \textstyle\sum_{\xi \in \mathcal{X}} W_\xi(\mathbf{n},Z) n_\xi \\ = &\tfrac{\varphi+1-\alpha}{\varphi+1} \tfrac{(1-\eta)(\varphi+1-\alpha)}{(1-\eta)(\varphi+1-\alpha)+\eta\varphi\alpha} \alpha^{\frac{\alpha}{\varphi+1-\alpha}} Z^{\frac{\varphi+1}{\varphi+1-\alpha}} \left(\textstyle\sum_{\xi \in \mathcal{X}} \left(n_\xi^{\varphi}/\xi \right)^{\frac{\rho}{\varphi+1-\rho}} \right)^{\frac{\alpha}{\rho} \frac{1+\varphi-\rho}{\varphi+1-\alpha}} - (\eta r\underline{c} + (1-\eta)\mu) N, \end{aligned}$$

where $N \equiv \sum_{\xi \in \mathcal{X}} n_\xi$. One can now confirm easily that Conjectures 1 and 2 are confirmed. ■

Table 1: Summary statistics

Variable	Mean	Std. Dev.
Average days per month per year	23.740	7.050
Job tenure (in months)	41.959	42.363
Average daily wage (2003 Euros)	120.720	422.320
Total days worked per year	244.170	98.340
Average number of months paid	9.970	3.390

NOTE: Number of observations = 22.689 million workers.

Table 2: Distribution of annual changes in days worked (h)

Share with $\Delta h = 0$	57.37%
Share with $\Delta h > 10$	17.56%
Avg $\lvert\Delta h\rvert$ if $\Delta h \neq 0$	22.63
Avg $\lvert\Delta h\rvert$ if $\Delta h \neq 0$, excluding $\lvert\Delta h\rvert > 50$	10.37

NOTE: Statistics above refer to our sample of two year stayers, as defined in the main text.

Table 3: Earnings and working time in Veneto panel

Moment	Data	
	12/12 stayers	2-year stayers
$\frac{var(\epsilon^W)}{var(\epsilon^h)}$	**2.724**	***2.123***
$\frac{var(\phi^W)}{var(\phi^h)}$	2.496	*2.403*
$\sqrt{var(\epsilon^W)}$	0.128	*0.192*
$\sqrt{var(\phi^W)}$	0.095	*0.126*
$\sqrt{var(\epsilon^h)}$	0.078	*0.132*
$\sqrt{var(\phi^h)}$	0.060	*0.081*
$\frac{cov(\Delta \ln h, \Delta \ln w)}{var(\Delta \ln w)}$	-0.243	*-0.189*

NOTE: The 12/12 and 2-year stayers are two samples of workers who remain with the same firm in 2 consecutive years. Workers in the 12/12 stayers are those paid for at least 1 day in every month in the adjacent 2 years. The 2-year stayers are workers who 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. For reasons discussed in the text, the model is fit to the 2-year stayers (moments in italics).

Table 4: Alternative estimates of $var(\epsilon^W)/var(\epsilon^h)$

Sample	12/12 stayer	2-year stayer
Baseline sample (full sample)	2.724	2.123
Excluding women	3.110	2.433
Excluding public sector	2.394	1.917
Excluding full-year workers	1.647	1.562
Including only stayers in large firms	4.038	2.904
Including only specific sectors		
Manufacturing	2.261	1.807
Construction	1.719	1.593
Commerce	2.175	1.750
Finance and Banking	7.673	5.610

Table 5: Model fit

PANEL A

Moment	Model	Data
$\frac{var(\epsilon^W)}{var(\epsilon^h)}$	2.130	2.123
$\frac{var(\phi^W)}{var(\phi^h)}$	2.390	2.403
$\sqrt{var(\epsilon^h)}$	0.132	0.132
$\sqrt{var(\phi^h)}$	0.082	0.081
$\frac{cov(\Delta \ln h, \Delta \ln w)}{var(\Delta \ln w)}$	-0.390	-0.189
$\sqrt{var(\Delta \ln N)}$	0.355	0.355

PANEL B

Parameter	Value
Std dev of idiosyncratic preference shock	0.56
Std dev of idiosyncratic productivity shock	na*
Worker bargaining power, η	0.123
Intertemporal elasticity of substitution, $1/\psi$	0.480
Elasticity of substitution across tasks, ρ	-4.26
Std dev of Z, σ	0.25

* This is omitted in the present draft. See text for a discussion of what this dimension of heterogeneity will imply in terms of the moments of interest.

Table 6: Robustness of certain parameters to alternative values of σ

Parameter	$\sigma = 0.15$	$\sigma = 0.20$	$\sigma = 0.25$
$1/\psi$	0.68	0.54	0.48
η	0.425	0.25	0.123
ρ	-2.77	-3.5	-4.26

Figure 1: Labor demand policy

PANEL A

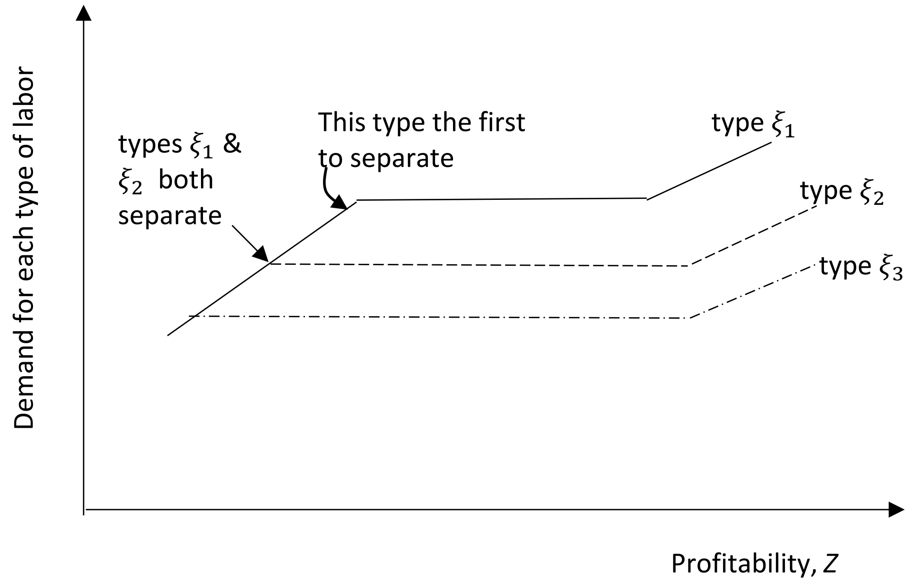

PANEL B