The insurance role of the firm*

Luigi Guiso — EIEF

Luigi Pistaferri — Stanford University

December 14, 2019

Abstract

We review the recent literature on the risk sharing role of the firm. We provide a framework for studying risk sharing between workers and firm owners *vis-à-vis* firms specific shocks of different nature. We show how this framework can be taken to the data to provide estimates of the extent of insurance within the firm. Estimates from a large number of Western countries strongly support the view that in capitalist economies the firm is a large albeit far from complete wage insurance instrument. We quantify the welfare benefits of firm-provided wage insurance, show evidence on how workers react to firms shocks passed through wages, and discuss the future role of the firm as a wage insurance provider.

*This paper was prepared for the invited lecture at the 46th EGRIE annual workshop in Rome.

1 Introduction

A long tradition in economics assigns a central role to the firm as an insurance provider. According to this view risk-tolerant capitalists are more willing than employees to absorb fluctuations in firm performance, and hence end up offering long-term contracts in which employees enjoy a flat compensation scheme in exchange for a higher firm profit rate. In his seminal contribution "Risk, Uncertainty and Profit", Frank Knight (1921) went so far as to argue that: "the system under which the confident and venturesome assume the risk and insure the timid by guaranteeing a specified income... is the enterprise and wage system of industry. Its existence .. is the direct result of the fact of uncertainty". Thus Knight traces the very existence of the firm - an issue that has preoccupied the economics profession since then - to the provision of insurance.

The key insight of Knight's view is that agglomerating production within the boundaries of the firm allows valuable exchanges between the "timid" worker who supplies labor services (the risk-averse party in today's language), and the "venturesome" capitalist (today referred to as the risk-tolerant party), who offers an insured income in exchange for the right to manage the labor services. Subsequent work in the 1970's has formalized this basic intuition in the context of the implicit contract model of Baily (1974) and Azariadis (1975). Implicit in Knight (1921), as well as in this strand of literature, is that markets cannot provide the insurance that risk-averse workers demand due to the presence of frictions in labor and/or financial markets.[1] Clearly, if labor markets were frictionless and workers could move costlessly from job to job, an adverse (specific) shock to their firms could be undone by just moving to another firm. Alternatively, if insurance and credit markets were not subject to informational frictions, workers could rely on them to smooth consumption against unpredictable shocks to labor income. The more recent formal contracting models of the worker-firm owner relation emphasize the differential access of firms and workers to financial markets rather than just differences in preferences: firms (shareholders) can better diversify idiosyncratic risk away and so act as risk-neutral agents in the relationship with workers, who have limited access to financial markets. However, the assumption that firms can fully diversify idiosyncratic risk, and therefore offer full insurance to risk averse workers, is very likely an extreme one. Many firms, particularly the non-listed ones (i.e., the majority in most countries), have highly concentrated ownership. Private equity owners tend then to hold very large stakes in their business, bearing a substantial amount of the firm risk (Fagereng et al 2019; Moskowitz and Vissing-Jorgensen, 2002). Undiversified business ownership weakens the assumption that capitalists are risk-neutral and pulls away from the full insurance benchmark of the implicit insurance model. A similar effect is played by the nature of the shock to the firm. Gamber (1988) shows that in

[1] In implicit contract model, firms smooth worker's consumption across states when insurance markets fail to work due to moral hazard or limited commitment problems. In a similar spirit, the wage contract may reflect opportunities to redistribute factor rewards across time when access to the credit market is limited or too costly. Guiso, Pistaferri and Schivardi (2013) study the role of the firm as an internal credit market.

an implicit insurance model *à la* Azariadis firms may be unwilling to fully insure persistent shocks if they are constrained by bankruptcy risk. More generally, as made clear by many contract theory models of the firm (starting with Holmstrom and Milgrom, 1987) unobservability of workers' effort by the firm owners is an obvious obstacle to insurance provision. When effort is unobserved or not contractible, the supply of insurance by the firm faces a trade off as it conflicts with workers' incentives to put up effort. The extent of wage insurance will then depend crucially on differences in risk tolerance between workers and firms, the marginal cost of providing effort, as well as other attributes of the contracting party such as the variability of firm performance.[2] More generally, it will depend on features of the financial markets and of the societal structure where workers and firms interact and that grant the two parties imperfect access to lending and risk diversification opportunities (Berk and Walden, 2013). The latter include any labor insurance provision from government programs such as unemployment insurance schemes which may substitute for within firm insurance coverage.[3]

In sum, the amount of insurance that a firm can offer to its workers needs to be quantified and, most likely, it is partial and highly dependent on the context examined, and on the characteristics of the workers and those of the firm. A quantitive evaluation is needed in order to assess the welfare benefits of the insurance role of the firm and to evaluate whether and how workers' choices change when full insurance is not available.

In the next section we start providing a framework to characterize the insurance role of the firm, by specifying which shocks can be insured and relating them to the wages that workers receive. In Section 3 we show how this framework can be flexibly taken to the data and, if suitable data on firms performance and workers compensation are available, how to obtain estimates of the degree of wage insurance against fluctuations in firm performance achieved within the firm. Evidence available for several countries suggests that firms offer a substantial amount of insurance with non-negligible welfare gains for the workers. But insurance is far from full and varies across worker and firm types. In Section 4 we discuss how the residual exposure of worker compensation to the firm-specific risk can be used directly and indirectly to shed light on workers' precautionary motives and financial portfolio allocations, documenting these responses empirically. While most of the literature has focused on insurance exchanges between firm owners and workers, the boundaries of the firm and the repeated relation between the worker and the firm - and thus between co-workers - can facilitate insurance and financial exchanges among co-workers. These exchanges are particularly

[2]The pass-through of firm-specific shocks onto wages is the subject of the vast "rent sharing" literature in labor economics. Most of the early literature on rent sharing was based on industry-level data. We comment on recent advances in this literature below.

[3] See Ellul, Pagano and Schivardi (2018) and Pagano (2019) for a thorough survey of the limitations to firm insurance provision. Lamadon (2016) finds that 2/3 of government provided insurance is undone by the crowding out of firm-provided insurance.

valuable against shocks to their pay that do not originate from the firm and that firm owners may be unwilling or unable to insure. We discuss this extended role of the firm as facilitator of insurance exchanges in Section 5. Section 6 concludes by offering some thoughts on the future role of the firm as a wage insurance provider.

2 Framework

In this section we follow Guiso, Pistaferri and Schivardi (2005) and lay down a simple framework that highlights the link between the shocks to the performance of the firm and the variation in the salary of the worker. We highlight the nature and type of the shocks that the firm can insure and define different insurance regimes.

2.1 The Firm

Let the (log) performance of firm j in year t, Π_{jt}, be written as:

$$\Pi_{jt} = X_{jt}^{'}\gamma + h_j + \psi_{jt}$$

where X_{jt} is a vector of controls that define aggregate and industry sources of variation in the performance of the firm. Its purpose is to capture non-idiosyncratic fluctuations in Π_{jt} (e.g., those originating from world- or country-specific business cycles, as well as industry-specific rents). What is left is the firm idiosyncratic component of Π_{jt}. This is modeled as the sum of a fixed firm effect h_j and a firm- and time-varying component ψ_{jt}. The fixed effect h_j captures systematic performance differences across firms reflecting managerial quality and practices (as stressed in Bloom an Van Reenen, 2007), quality of corporate governance, or some exclusive location advantage not available to competitors. What matters in our context is that this component is predictable and (being highly persistent over time) hardly interpretable as a shock. The second component, ψ_{jt}, is an unpredictable shock to the firm performance.

Undoubtedly, worker wages may be correlated will all these components: they may be correlated with economy or industry-wide fluctuations in firm performance because wages do respond to market forces; this has long been documented in the rent sharing literature.[4] The early evidence as well as more recent contributions based on micro employer-employee data are recently discussed and summarized by Card et al (2018). The rent sharing literature and the one we discuss here differ in interpretation, but not in the

[4]See among others, Blanchflower, Oswald, and Sanfey (1996), Estevao and Tevlin (2003), Christofides and Oswald (1992).

methodology they adopt, and indeed we summarize below estimates from both literatures. In both cases, the existence of "sharing" behavior (risks or rents) is identified by the existence of a relationship between wages and measures of firm profitability or performance (a relationship that is absent in a perfectly competitive setting). The theoretical differences between the two models arise from the nature of market imperfections and the type of contractual arrangements that firms and workers set in place to overcome such imperfections. In the risk sharing case, credit and insurance market imperfections lead to long-term, implicit contract between firms and *individual* workers. In the rent sharing case, the main form of friction is workers' bargaining power, and union contracts are explicitly signed to allow firms and workers to share rents (surplus) arising from the production process.

By definition, industry-level correlations cannot be taken as evidence of firm insurance. Wages tend also to be correlated with the quality h_j of the firm, as shown by another strand of literature started by Abowd, Kramarz, and Margolis (1999). The authors document systematic firm-effects on workers wages and a tendency of high pay workers to match with high-paying firms, consistent with models of assortative matching between highly productive workers and highly productive firms. But again the correlation is best interpreted as reflecting complementarity in production between firm and worker types, rather than insurance.

It is the ψ_{jt} component that can potentially be insured by the firm. The insurability of ψ_{jt} is theory-grounded and fits well with the idea that capitalists can offer insurance to workers because they can potentially fully diversify idiosyncratic innovations to firm performance through, say, the stock or bond markets. We enrich the representation of the firm idiosyncratic innovations by assuming that it is the sum of two shocks:

$$\psi_{jt} = Q_{jt} + f_{jt}^T$$

$$Q_{jt} = Q_{jt-1} + f_{jt}^P$$

a transitory one, f_{jt}^T, and a permanent stochastic component Q_{jt} with innovation f_{jt}^P. The two innovations are assumed to be i.i.d. with mean zero; we denote with σ_ψ^2, $\sigma_{f^T}^2$, $\sigma_{f^P}^2$ the variance of the total idiosyncratic innovations to firm performance, and of the transitory and permanent shocks, respectively. Here the variances are assumed to be time invariant; but later, for some applications, we will let them vary over time. This flexible representation allows firms to offer insurance to their workforce possibly to different degrees depending on the persistence of the shock as predicted by wage insurance models of the firm with financial frictions (e.g. Gamber, 1988). Lamadon (2016) shows that a framework of this type can be generated by a model of the worker-firm interaction with search frictions. The search friction implies that risk averse workers cannot rely on the market to insure their pay through costless job-to-job mobility. Firms instead can offer contracts that shelter workers from productivity shocks by tempering the insurance coverage with the need to give workers incentives to exert effort (which is assumed to be unobservable to the firms).

2.2 Workers

The (log) earnings/wages of a worker i employed at firm j evolves according to the following process:

$$y_{ijt} = Z_{ijt}^{'}\beta + g_i + \omega_{ijt}$$

The first two components are assumed to be deterministic or time-invariant: a vector of variables $Z_{ijt}^{'}$, affecting compensation either because wages respond to market equilibrium forces or due to time varying predictable factors such as age, experience and tenure, and a worker-specific fixed component g_i, reflecting systematic differences in workers' productivity. Unpredictable shocks to workers pay are denoted ω_{ijt}, with variance σ_ω^2 capturing overall volatility in wages. We assume that unpredictable shocks to workers pay are the sum of two components:

$$\omega_{ijt} = v_{it} + F_{ijt} \tag{1}$$

The first component of (1), v_{it}, is a worker-specific term reflecting shocks to workers' pay that are unrelated to the performance of the firm - such as a drop in pay due to (uninsured) illness or a fall/rise in individual worker productivity that is reflected in his/her personal bonus. We assume that this component is itself the sum of a transitory and a permanent shock, respectively, with i.i.d., zero mean, innovations that are uncorrelated with the firm shocks :

$$v_{it} = P_{it} + \eta_{it}$$

$$P_{it} = P_{it-1} + \chi_{it}$$

with $var(v_{it}) = \sigma_v^2$, $var(\eta_{it}) = \sigma_\eta^2$; $var(\chi_{it}) = \sigma_\chi^2$. The idiosyncratic component is not insured by the firm owners, though co-workers can - possibly partially - insure it by leveraging the repeated interactions that arise among them from their common attachment to the firm; we will touch upon this issue in Section 5.

Table 1: Workers' wage shocks

Wage shock component	Definitions	Variances
Total	$\omega_{ijt} = v_{it} + \theta^P Q_{jt} + \theta^T f_{jt}^T$	σ_ω^2
Idiosyncratic	$v_{it} = P_{it} + \eta_{it}$	σ_η^2
	$P_{it} = P_{it-1} + \chi_{it}$	σ_χ^2
Firm risk sharing	$F_{ijt} = \theta^P Q_{jt} + \theta^T f_{jt}^T$	σ_F^2,$\sigma_{f^P}^2$, $\sigma_{f^T}^2$

Notes: The table describes the components of the workers earnings shocks and their variances. All shocks are zero mean and i.i.d.

Table 2: Insurance regimes

Insurance regimes	Values of $\theta^P$and θ^T
Full homogeneous insurance	$\theta^P = 0, \theta^T = 0$
Partial homogenous insurance	$0 < \theta^P = \theta^T < 1$
Partial insurance to permanent, full to transitory	$0 < \theta^P < 1, \theta^T = 0$
Full risk sharing (no insurance)	$\theta^P = 1, \theta^T = 1$

Notes: The table summarizes different wage insurance regimes that capitalists can offer to the employees of their firm.

The second component of (1), F_{ijt}, is related to the firm that employs the worker and we write it as:

$$F_{ijt} = \theta^P Q_{jt} + \theta^T f_{jt}^T$$

F_{ijt} reflects sharing the shocks hitting the firm between the firm capitalists and the workers. It is these shocks that firms can potentially insure. Table 1 summarizes the structure of the wage shocks.

2.3 Insurance regimes

The extent of insurance within the firm for the permanent and transitory shocks to firm performance are respectively summarized by the insurance coefficients θ^P and θ^T. Table 2 describes different insurance regimes that can arise depending on the values of θ^P and θ^T. At one extreme, if θ^P and θ^T are both equal to zero, the capitalists fully insure the workers against unforeseen shocks to the fortunes of their firms, regardless of their persistence. At the other extreme, no insurance is provided and shocks to the firm are equally shared - and thus fully transferred to the workers - which occurs when θ^P and θ^T are both equal to 1. Values of $0 < \theta^P < 1$ and $0 < \theta^T < 1$ characterize regimes of partial insurance and these can be homogenous (if $\theta^P = \theta^T$) or differentiated if firms provide more coverage with respect to one type of shock than another (i.e. $\theta^P \neq \theta^T$).

This framework can be easily generalized to allow for dependence of the degree of insurance on workers (X^w) and firm (X^f) characteristics - such as the worker risk tolerance, his effort observability, or the volatility of firm performance - as well as characteristics of the financial markets where firms and workers interacts or features of the society (X^m) that limit workers and firms diversification of production risk. To allow for this we can write:

$$F_{ijt} = \theta^P(X^w, X^f, X^m)Q_{jt} + \theta^T(X^w, X^f, X^m)f^T_{jt}$$

Below we discuss how to identify empirically the insurance parameters θ^P and θ^T and their dependence on observable worker and firm attributes.

3 Measuring the extent of risk sharing

The framework of the previous section lends itself naturally to be taken to the data with the goal of obtaining estimates of the risk sharing parameters. These estimates can then be used to provide a quantitative evaluation of the economic importance of the firm as a wage insurance provider and estimates of the welfare gains for the workers. Doing this poses some strong data requirements. First, it requires the availability of matched employer-employees data that allow separate measurment of firm performance shocks and wage shocks. Until recently this type of data were not easily accessible and were available to researchers only in a limited number of countries. Today they have become available in a wider set of countries. This is important because it allows testing the external validity of some of the initial findings about the insurance role of the firm and potentially sheds light on features of society that affect wage insurance. In addition to firm performance and worker wage histories, a test of the insurance role of the firm requires an institutional setting where firms can pass through specific shocks to workers, one where at least part of the salary is set at the level of the firm. Clearly, this is not a case in a fully centralized wage setting system where wages are bargained by nation-wide trade unions and apply uniformly to all workers in an industry. Most Western countries, including those with high levels of union coverage, have long departed from this extreme form of centralization. For instance, in most European countries wages are bargained at multiple levels: through nation wide agreements, at the sector level, and at the level of the firm. Below we draw on Guiso, Pistaferri and Schivardi (2005) who were the first to show evidence on insurance within the firm using matched employer-employees data from Italy. This country fulfills the institutional requirement as part of the wages are set (during the sample period) bilaterally between workers and firms at the firm level.

To measure firm performance Guiso, Pistaferri and Schivardi (2005) use firm value added obtained from the Company Accounts Data Service (CADS), a collection of firms balance sheets over the years 1982-1994. Value added is a sensible measure of firm performance as it measures the volume of contractible output between the capitalists and the workers that remains once intermediate inputs have been remunerated. Additionally, value added is the variable that is directly subject to stochastic fluctuations.[5] Wages are measured by employees' earnings histories from the Italian Social Security Administration (INPS), using data for the entire population of workers registered with the social security system whose birthday falls on one of two randomly chosen days of the year. Both the INPS and the CADS data include the employer's tax code, so matching employers with the employees is straightforward.

Using these data it is straightforward to obtain estimates of the firm performance and wage residuals from estimates in first differences (so as to eliminate the firm and worker fixed effects) of equations (1) and (2), that is estimates of $\Delta\psi_{jt}$ and $\Delta\omega_{ijt}$, respectively. It is then immediate to estimate the correlation between the two: 0.0021 (s.e. 0.0003), which implies that part of the shocks to the performance of the firm is indeed transferred to workers - thus allowing to (statistically) reject full insurance. However, this correlation does not reveal anything about the nature of the insurance regime. Because transitory and permanent shocks are not separately observed, one needs an identification strategy to obtain separate estimates of θ^P and θ^T.

3.1 Identification

Guiso, Pistaferri and Schivardi (2005) show that, given the structure of the shocks to the performance of the firm and the workers earnings described in Section 2.1 and 2.2, the insurance parameters θ^P and θ^T are identified by the following moments:

$$\theta^T = \frac{cov(\Delta\omega_{ijt}, \Delta\psi_{jt+1})}{cov(\Delta\psi_{jt}, \Delta\psi_{jt+1})}$$

$$\theta^P = \frac{cov(\Delta\omega_{ijt}, \Delta\psi_{jt-1} + \Delta\psi_{jt} + \Delta\psi_{jt+1})}{cov(\Delta\psi_{jt}, \Delta\psi_{jt-1} + \Delta\psi_{jt} + \Delta\psi_{jt+1})}$$

which are equivalent to running instrumental variable regressions of $\Delta\omega_{ijt}$ on $\Delta\psi_{jt}$ using as instrument $\Delta\psi_{jt+1}$ (to identify insurance with respect to transitory shocks to firm performance) and the moving average $\Delta\psi_{jt-1} + \Delta\psi_{jt} + \Delta\psi_{jt+1}$ (to identify insurance with respect to permanent shocks), respectively.

[5] Using (total or per-worker) sales or TFP measures are valid alternatives; using profits, however, may generate a mechanical form of negative correlation with wages. In recent work, authors have departed from a statistical characterization of what constitutes a shock to firm performance, and searched for "episodes" or "quasi-experiments" in which the firm faces an exogenous, salient innovation. An example is Kline, Petkova, Williams and Zidar (2019), who use as their measure of firm shock the allowances of patent application by US firms, an idea partially based on Van Reenen (1996), who however had only access to industry-level data on innovation for the UK.

Table 3: Estimates of insurance within the firm

Insurance parameter	Estimate	F-test instruments (p-value)
θ^P	0.070***	<0.001
	(0.0128)	
θ^T	0.0049	<0.001
	(0.0045)	

Notes: The table shows estimates of the degree of risk sharing within the firm in response to persistent and transitory shocks to firm performance. Estimates reproduced from Guiso, Pistaferri and Schivardi (2005). (****), (**) and (*) significant a 1%, 5% and 1%, respectively. Standard errors in parenthesis.

3.2 The size of insurance coverage

Table 3 shows estimates of these parameters from the Italian data. The evidence clearly supports a partial insurance regime where capitalists protect workers fully against transitory shocks to firm performance but share permanent risk partially with the workers. The extent of risk sharing implies that a 10% permanent drop in the idiosyncratic firm performance causes a 0.7% drop in workers' pay. This is very far from the full risk sharing regime ($\theta^P = \theta^T = 1$), implying that firms do indeed offer substantial wage insurance.

To shed some light on the role of the firm as insurance provider, Table 4 shows estimates of the variance and corresponding standard deviation of the shocks (and their components) to firm performance and workers earnings estimated on the Italian data. Measured by the standard deviation of the workers' residual earnings, wage volatility is overall 12.7%. Sharing risks with the firm accounts for 13% ($\sqrt{\sigma_F^2/\sigma_\omega^2} = 0.13$) of workers pay volatility - a relatively contained contribution. To appreciate the insurance role of the firm we can compute the standard deviation of earnings in a counterfactual case in which workers were to share completely with capitalists the fortunes of the firm (i.e. if $\theta^P = \theta^T = 1$). As the table shows, in this case the standard deviation of workers earnings would jump to 36.2% - 2.8 times the actual level. Insurance from the capitalists allows workers to reduce quite substantially fluctuations in pay.

3.3 Welfare gain

To get a sense of the welfare gain that firm insurance provides, assume that workers have no other mean to smooth consumption in the face of wage risk (i.e., no access to self-insurance) and (for simplicity) assume no worker idiosyncratic uncertainty (i.e. $\sigma_v^2 = 0$) so that all earnings risk originates from sharing risk with firm owners. Following Lucas's (1987) setup to evaluate the welfare costs of recessions, suppose that workers maximize the expected discounted value of flow utility $U(c) = \frac{c^{1-r}-1}{1-r}$, and receive stochastic earnings Ae^{ω_t}, where r is the degree of workers' relative risk aversion, A is a non-stochastic component, and $\omega_t \sim N(0, \sigma_\omega^2)$, with $\sigma_\omega^2 = (\theta^P)^2\sigma_{f^P}^2 + (\theta^T)^2\sigma_{f^T}^2$. Wage risk depends on the insurance parameters θ^P and θ^T and on the variances of permanent and transitory shocks to the performance of the firm. Following Lucas we assume that the approximate consumption solution is $c_t = (1+g)^t A' e^{-\sigma_\omega^2/2} e^{\omega_t}$ where g is deterministic consumption growth and A' a constant term. The proportional increase in consumption required to leave the consumer indifferent between the full insurance consumption path ($\theta^P = 0$ and $\theta^T = 0$) and the partial insurance path

Table 4: Variances of firm and worker shocks

Shock definition	Variance estimate	Standard error estimate
Firm shock component		
Total (ψ_{jt})	$\sigma^2_\psi = 0.115$	0.339
Transitory (f^T_{jt})	$\sigma^2_{f^T} = 0.056$	0.237
Permanent (f^P_{jt})	$\sigma^2_{f^P} = 0.059$	0.242
Wage shock component		
Total (ω_{ijt})	$\sigma^2_\omega = 0.0161$	0.127
Total idiosyncratic (v_{it})	$\sigma^2_v = 0.0121$	0.11
Transitory idiosyncratic (η_{it})	$\sigma^2_\eta = 0.009$	0.095
Permanent idiosyncratic (χ_{it})	$\sigma^2_\chi = 0.0014$	0.037
Firm risk sharing ($F_{ijt} = \theta^P Q_{jt} + \theta^T f^T_{jt}$)	$\sigma^2_F = 0.00029$	0.017
Share of wage standard deviation due risk sharing	$\sqrt{\sigma^2_F/\sigma^2_\omega} = 0.13$	
Wage volatility with full risk sharing	$(\sigma^2_\omega - \sigma^2_F) + \sigma^2_\psi = 0.131$	0.362

Notes: The table shows estimates of the variance and standard deviation of different shocks to the worker wage and the firm performance. Estimates rearranged from those obtained by Guiso, Pistaferri and Schivardi (2005).

($0 < \theta^P < 1$ and $0 < \theta^T < 1$) is $r\sigma^2_\omega/2$ (Clark, Leslie, and Symons 1994). This is the risk premium that a worker would be willing to pay to avoid a proportional risk of size σ^2_ω, which provides a first approximation to the cost of risk sharing. Under the (standard) assumption that $r = 3$, and using Guiso, Pistaferri and Schivardi (2005) estimates of the variances of the firm shocks, the proportional increase in consumption that would leave the consumer indifferent between full and no insurance ($\theta^P = \theta^T = 1$) is estimated by GSP to be 0.088. Hence, a worker would be indifferent between the insurance the firm currently provides and *no* insurance if her consumption was raised by about 9 percent - a substantial sum of money. Payment takes the form of a risk premium paid to the owners of capital or, to put it slightly differently, a wage discount that goes to increase average dividends earned by capital owners'. Guiso, Pistaferri and Schivardi (2005) calculation (9% of consumption) provides an estimate of the upper bound of the premium (or the size of the wage discount). The counterpart of the insurance provision is a greater share of national income accruing to corporate shareholders as documented by Hartman-Glaser, Lustig and Zhang (2019).

3.4 Heterogenous effects

Principal-agent contracting models of the firm-worker relation predict that insurance provision varies with both worker and firms characteristics, including workers risk aversion and firms performance volatility. Documenting this type of heterogeneity is important as it help strengthen the interpretation of the correlation in terms of insurance. Table 5 summarizes the findings of Guiso, Pistaferri and Schivardi (2005) along this dimension.[6] It shows that firms provide more insurance to more risk averse workers, and less to their managers (who presumably can stand risk more than white or blue collars). Furthermore, firms with more variable output offer more insurance, while those more exposed to bankruptcy offer less, consistent with theoretical predictions of firms insurance models.

Another, more subtle form of heterogeneity may come from separating wage responses from employment responses. Firms could respond to external shocks by adjusting wages (i.e., limiting insurance to continuing workers) or by adjusting employment (i.e., by laying-off some workers, or adjusting the number of hours worked). This second margin of adjustment is absent in Guiso, Pistaferri and Schivardi (2005), who focus on

[6]See Guiso, Pistaferri and Schivardi (2005) for variable definitions and measurement.

Table 5: Estimates of insurance within the firm

Insurance parameter	Baseline	High risk aversion	Performance volatility	Managers	Bankruptcy
θ^P	0.070***	0.030**	0.083**	0.187**	0.142***
θ^T	Full Insur.	Full Insur.	Full Insur.	Full Insur.	Full Insur.

Notes: The table shows estimates of the degree of wage risk sharing for different groups of workers and firms. Estimates reproduced from Guiso, Pistaferri and Schivardi (2005). (****), (**) and (*) significant a 1%, 5% and 1%, respectively.

continuing workers. This may exaggerate the extent of "insurance" if firms adjust mostly on the employment margin (something that may depend on institutional features of the labor market, such as firing and hiring costs, etc.). A few recent papers have extended the framework by allowing for employment as well as job-to-job mobility (see Friedrich, Laun, Meghir and Pistaferri, 2019).

3.5 Estimates for other countries

Several papers have replicated Guiso, Pistaferri and Schivardi (2005) for other countries, using similar data and methodology. These contributions include: Portugal (Cardoso and Portela, 2009), Germany (Guertzgen, 2014), Hungary (Katay, 2008), Sweden (Lamadon, 2016), Norway (Fagereng et al, 2018), and the US (Lamadon et al, 2019). In Table 6 we report estimates of the wage insurance parameters from these papers. Other papers in the literature do not explicitly distinguish between transitory and permanent fluctuations in performance, and hence estimate the extent of pass-through of firm performance on wages ("rent sharing" parameters). Card et al (2018) provide a survey of this literature and a detailed description of these contributions, with elasticities ranging from 0.01 to 0.16 (with a mean of 0.08).

Table 6 reveals an important empirical regularity: in all countries firms tend to insure fully or (almost completely) transitory shocks to firm specific shocks and to offer partial but substantial insurance with respect to permanent shocks to the firm. Across countries, the range of the elasticity of wages to firm permanent shocks has a median of 0.08, so that around 90% of permanent shocks to the firms are not passed over to workers. For transitory shocks, the median is 0.018.

In sum, at least in western countries the firm appears to be a formidable wage insurance provider. Well diversified capitalists or capitalists with a greater willingness to bear risk offer protection to risk averse workers from large idiosyncratic variations in the performance of the firm. This protection however is neither *gratis* nor perfect. It is not *gratis* as, in exchange for wage insurance, workers willingly give up part of the produced income. It is not perfect because some of the shocks, typically the persistent ones, are (at least in part) passed onto the workers that are retained in the firm, possibly as a consequence of imperfect monitoring of workers effort or because financial market imperfections constrain firms' ability to offer more extensive insurance (Berk and Walden, 2013). This shared risk is very difficult for the workers to diversify away (if it were - given the high wage discount workers seem to pay to capitalists according to the estimates in Section 3.3 - they should share risk with the firm to a higher degree than observed) and thus it acts as a *background risk*. In an attempt to deal with such background risk workers may change some of their actions. In the next section we focus attention to workers' endogenous reactions to background risk and show how,

Table 6: Estimates of insurance within the firm across countries

Country	Authors	θ^P	θ^T
Norway	Fagereng et al. (2018)	0.070^{***}	0.018^{***}
Portugal	Cardoso and Portela (2009)	0.092^{**}	-0.001
Hungary	Kayay (2015)	0.107^{***}	0.055^{***}
Sweden	Lamadon (2016)	0.014^{***}	n.a.
USA	Lamadon et al. (2018)	0.140^{***}	-0.010
Germany	Guertzgen (2009)	0.072	0.022

Notes: The table shows estimates of degree of wage risk sharing for a sample of countries taken from different separate studies that replicate Guiso, Pistaferri and Schivardi (2005) methodology. Guertzgen (2009) interacts the variables of interest with dummies for different industrial relation regimes. We report the average value across the different configurations (no standard errors available). (****), (**) and (*) significant a 1%, 5% and 1%, respectively.

methodologically, the possibility of measuring risk sharing between workers and firms can help identifying several behavioral responses of consumers to risk.

4 Worker responses

Income shocks that, by their nature, cannot be insured nor diversified away - like the share of the firm volatility that is passed onto workers wages - tend to make workers/consumers more risk averse (under reasonable preference restrictions, see e.g. Kimball, 1990). This induces individuals to modify their behavior by reducing other risks that they can control so as to contain overall risk exposure. Alternatively, prudent workers can respond by piling up precautionary savings to be better able to face unforeseen income drops whenever they occur.

The standard approach to study behavioral responses to uninsurable, background income risk has been to use measures of the individual component of wage variability as an index of background risk. In our set up, this measure is the variance of residual (log) wage $\sigma^2_{\omega t}$ - where we are now allowing this measure to vary over time. However, this approach is likely to be problematic. The key issue is that not all *measured* residual wage variation is risk.[7] As a recent literature makes clear, a possibly non-negligible part of it is just year-to-year fluctuations in wages due to heterogeneity and choice (Cunha, Heckman and Navarro, 2005, Guvenen and Smith, 2013, Low, Meghir and Pistaferri, 2008). As argued by Fagereng, Guiso and Pistaferri (2018), if this is the case estimates of the effect of background income risk on choice are affected by classical measurement error, which would explain why it has typically been difficult to find "large effects".

To illustrate, residual earnings is given by (1):

$$\omega_{ijt} = v_{it} + F_{jt},$$

the sum of the individual idiosyncratic component and the firm component. The key idea is that variation stemming from the v_{it} component can be decomposed into risk (a fraction μ of it) and choice (the remaining $1-\mu$ part), so that:

$$\omega_{ijt} = (1-\mu)v_{it} + \mu v_{it} + F_{jt}$$

[7] Deaton (1991) was probably the first to express skepticism about measured wage volatility from earnings regressions as all reflecting wage uncertainty. He went as far in his simulations as to start from available estimates but then halve them in order to get more plausible values. Even these adjusted values, he argues, still represent "substantial uncertainty in the wage process", suggesting that the overestimation of wage risk by measured volatility is likely substantial.

Only the component $\mu v_{it} + F_{jt}$ strictly qualifies as "unavoidable" risk. The separation of v_{it} in a component that is avoidable and one that is not (with weights μ and $(1-\mu)$, respectively) comes from recognizing that part of what the econometrician identifies as "risk" is variability in earnings that reflects, at least in part, individual choices rather than risk. For instance, time out of the labor market (inducing large swings in earnings across years) could be time invested voluntarily in human capital accumulation. Some volatility can be generated by people choosing to work longer hours, or perhaps to invest in training programs that increase their future productivity, in response to adverse financial market shocks affecting the value of their portfolio. Hence, measured earnings volatility (assuming orthogonality between v_{it} and F_{jt}), will be:

$$\sigma^2_{\omega t} = (1-\mu)^2 \sigma^2_{vt} + \mu^2 \sigma^2_{vt} + \sigma^2_{Ft}$$

The left hand side is measured background risk (allowing for time variation), while the right hand side is the sum of true background risk, $\mu^2 \sigma^2_{vt} + \sigma^2_{Ft}$, plus "measurement" error, $(1-\mu)^2 \sigma^2_{vt}$. Clearly, if one were to use $\sigma^2_{\omega t}$ as a measure of background risk in OLS regressions for consumer responses, this would produce downward biased estimates - with the size of the bias rising in the share of measured variation accounted for by the predictable idiosyncratic component $1-\mu$. Importantly, imperfect observability of background risk makes the assessment of its impact on behavior difficult to assess for two reasons: (1) because of the downward bias in the marginal effect on background risk and (2) because of a correct assessment of the extent of background risk. Indeed, even if one were to know the true marginal effect, the overall economic importance (*marginal effect* $\times$ *background risk*) requires knowledge of both the size of and the marginal response to background risk.

Below we discuss how observability of one component of background risk - the one originating from risk sharing within the firm - can help identify the marginal effect on two important consumer response margins: the portfolio allocation to risky assets and the accumulation of precautionary savings with respect to uninsurable income risk. In the context of portfolio allocation we also discuss how to obtain an estimate of the "true" extent of background risk $\mu^2 \sigma^2_{vt} + \sigma^2_{Ft}$.

4.1 Portfolio choice and background risk

Let the relation between the background risk faced by individual i, B_{it}, and the financial portfolio share in risky assets, S_{it} be described by a simple linear function:

$$S_{it} = X'_{it}\beta + \phi B_{it} + r_i + \varepsilon_{it}$$

where X'_{it} is a vector of controls and r_i and individual fixed effect that captures fixed individual attributes that affect portfolio allocations - such as baseline risk tolerance or financial skill and information. If one could observe B_{it}, OLS estimates would produce consistent estimates of the marginal effect of background risk ϕ and the economic effect would be simply estimated as ϕB_{it}. Since B_{it} is unobserved, the actual regression that can be run is instead:

$$S_{it} = X'_{it}\beta + \phi \sigma^2_{\omega t} + r_i + u_i$$

If the variance term $\sigma^2_{\omega t}$ is an error-ridden measure of B_{it}, the OLS estimate of ϕ would be a downward biased estimate of the true ϕ and the size of the bias can be potentially substantial. Fagereng, Guiso and

Pistaferri (2018) suggest estimating the above regression using an IV strategy where the the variance term $\sigma^2_{\omega t}$ is instrumented with the firm component of the individual background risk, i.e, σ^2_{Ft}. They show that this strategy produces consistent estimates of the marginal effect of background risk. They apply this strategy to Norwegian data, which appear particularly suited for their purposes. Because Norway levies a tax on net worth and because different asset components are taxed at different rates, this makes it possible to use panel data variation in background risk and portfolio composition to identify the effect of background risk on the propensity to invest in risky vs. safe assets. And because intermediaries directly report individual asset holdings to tax authorities, the scope for under- and mis-reporting is greatly reduced. Furthermore, because data are available for the whole population, selective attrition is not an issue. Finally, and critically for the purpose of understanding the economic importance of background risk, matched employers-employees data are available so that one can easily compute the firm component of background risk using the Guiso, Pistaferri and Schivardi (2005) methodology (results for Norway were shown in Table 6).

Table 7 reproduces Fagereng, Guiso and Pistaferri (2018) results. The first column shows the OLS estimates using residual earnings variation ($\sigma^2_{\omega t}$) as a measure of background risk; the second shows reduced form regressions of portfolio allocation on the firm component of wage risk (σ^2_{Ft}); and the third the IV estimates using the latter as instruments. Because in Norway firms share with workers part of the transitory shocks to performance, reduced form and IV regressions include the variances of both permanent and transitory shocks to firm performance as instruments. The OLS estimates show a negative correlation between workers wage volatility and the risky portfolio share - consistent with background risk inducing more prudent behavior. However, the effect is small. The reduced form regressions show negative effects of the firm volatility components on the portfolio share both for the transitory and permanent shocks. Additionally, Fagereng, Guiso and Pistaferri (2018) add a control for a plant closure to capture the effect of unemployment risk, which also affects negatively the risky portfolio share. Interestingly, the IV estimates in the last column show a 20-fold increase in the marginal effect of background risk, with the main coefficient jumping from -0.0202 in the OLS case to almost -0.5 in the IV case. The main inference is that the measurement error effect is substantial, implying that unless it is properly taken into account conclusions about the importance of background risk can be severely misleading.

4.1.1 Importance of background risk

Regarding the assessment of the economic importance of background risk, Fagereng, Guiso and Pistaferri (2018), conduct two additional exercises. First, they document that the effects of background risk are heterogenous across the wealth distribution. Heterogeneity across different levels of wealth arises through two channels: a) a risk-sharing channel, whereby firms tend to insure less wealthy workers by a larger extent, most likely because wealthier workers have a greater ability to self-insure through their own accumulated savings and hence demand less insurance from the firm (and, as a consequence, also obtain higher pay); b) a background risk channel whereby people with more assets are less sensitive at the margin to an increase in background risk. Both channels are sizable: risk sharing coefficients are twice as large for people in the top decile of the wealth distribution compared to those in the bottom decile. Similarly, the marginal effect of background risk is zero among people in the very top percentiles of the wealth distribution and three times as large as that among people in the bottom percentiles.

Second, they are able to provide an estimate of the share of idiosyncratic wage shocks that qualify as background risk, which they find to be $\mu = 0.2$. Hence, 80% of the measured variation in wages that is not

Table 7: Background risk and portfolio choice

Background risk measure	OLS estimate	Reduced form	IV estimate
$\sigma^2_{\omega it}$	-0.0202^{***}		-0.499^{***}
	(0.0029)		(0.123)
Permanent firm shock: σ^2_{fP}		-0.0033^{**}	
		(0.0012)	
Transitory firm shock: σ^2_{fT}		-0.0028	
		(0.0007)	
Firm closure within 1 year		-0.0112^{**}	-0.0201^{***}
		(0.0050)	(0.0066)
J-test p-value			13%

Notes: The table shows OLS and IV estimates of the effect of background risk on the stock share of financial wealth. Estimates reproduced from Fagereng, Guiso and Pistaferri (2018). Standard errors in brackets; (****), (**) and (*) significant a 1%, 5% and 1%, respectively.

due to risk sharing with the firm is predictable variation due to choice. Total background risk is accordingly estimated at $0.2^2\sigma^2_{vt} + \sigma^2_{Ft}$.

Armed with these estimates they conclude that background risk is an important determinant of portfolio allocations, discouraging stockholding significantly, only for people at the bottom of the wealth distribution. For people at the top (95th percentile of above) it has negligible effects. This has an important implication for assets prices: because most stocks are concentrated in the hands of the wealthy, background risk leaves the aggregate demand for stocks - and thus stock returns - basically unaffected.

4.2 Precautionary savings responses

As has long been recognized, workers can self-protect from uninsurable income shocks by piling up precautionary savings. This is a well studied topic (Kimball, 1990; Carroll, 1992; Dynan, 1992; Guiso, Jappelli and Terlizzese, 1992). Yet, empirical work has faced important issues (see Jappelli and Pistaferri, 2017, for an overview). Identification of a precautionary motive for saving requires measures of risk that satisfy three conditions: exogeneity, observability, and variability (Browning and Lusardi, 1996). The first requires that one needs exogenous variation in the measure of risk to identify the precautionary motive, the second that this measure has to be observable, and the third that it has to vary across households and possibly over time for the same household so that one can study the correlation between consumption growth and (heterogenous) uncertainty. Furthermore, what matters for precautionary savings is all sources of risk that can affect consumption (including for instance the risk of being unable to borrow): this is captured by consumption risk. Income risk, in contrast, is only one source of consumption risk, though obviously a very important one.

Ideally, what one would like to estimate is an Euler equation such as:

$$\Delta log c_{it} = \alpha \sigma^2_{c,it} + X'_{it}\beta + \epsilon_{it}$$

where is c is personal consumption expenditure, $\sigma^2_{c,it}$ the conditional variance of personal consumption growth and X a vector of controls. Under CRRA preferences, α equals 1/2 times the coefficient of relative prudence as defined by Kimball (1990). Thus, if α is estimated consistently, it identifies the strength of the precautionary motive for saving. The problem with this approach is that σ^2_{cit} is not observed. Dynan (1992) solved the problem by replacing the conditional consumption variance with $(\Delta log c_{it})^2$ and instrumenting this

Table 8: Background risk and precautionary savings

	IV estimate	IV estimate
	$(\Delta logc_{it})^2$	$\sigma^2_{\omega it}$
	1.008***	0.650***
	(0.310)	(0.233)
F-stat first stage	7.85	294.9
Hansen J-test p-value	0.396	0.029
Observations	327,518	327,518

Notes: The table shows IV estimates of the precautionary savings effect of background risk. Estimates reproduced from Fagereng, Guiso and Pistaferri (2017). Standard errors in brackets; (****), (**) and (*) significant a 1%, 5% and 1%, respectively.

term with a set of socio-demographic variables that are known at time $t-1$ (such as education, occupation, etc.). Dynan (1992) estimates a very small value of α. The empirical problem with this approach is that the selected instruments either lack power or may be good predictors of changes in the *chosen* pattern of consumption (i.e., they violate the exclusion restriction).

An alternative approach is to replace σ^2_{cit} with earnings volatility $\sigma^2_{\omega it}$ (which can be observed and measured), and estimate

$$\Delta logc_{it} = \alpha\varphi\sigma^2_{\omega it} + X'_{it}\beta + \epsilon_{it}$$

where $1-\varphi$ can be interpreted as a measure of income risk that the individual is able to self-insure, for instance through marriage with a spouse with uncorrelated earnings, which mitigates the effect of income risk on consumption risk. In this case identification of the strength of the precautionary motive and of the self-insurance of earnings risk requires separate identification of the two parameters α and φ. In addition it also requires solving the error-in-variables problem that affects $\sigma^2_{\omega it}$ discussed at the beginning of Section 4. Following the approach in the previous section we can use the firm-related component of earnings risk as an instrument for consumption volatility in the first specification and identify α. We can also instrument residual earnings volatility with the firm components in the second specification and obtain a consistent estimate of $\alpha\varphi$. Next, we can retrieve an estimate of self-insurance using the consistent estimate of α from the first specification of the Euler equation.

Table 8 reports the results of these estimates obtained by Fagereng, Guiso and Pistaferri (2017) using the Norwegian panel data described in the previous section. Instruments work pretty well; not surprisingly, they are more powerful in predicting earnings risk than consumption risk. The strategy also delivers plausible estimates. The IV estimate of α is very close to 1, implying a coefficient of prudence of around 2 and an estimate of $\alpha\varphi = 0.65$, suggesting that individuals can self-insure about 1/3 of their income risk.

4.3 Other potential responses

Workers can adapt to background risk by changing other behaviors over which they have control, helping them to contain overall exposure to risk. Our methodology - based on the exogeneity, observability and heterogeneity of one important source of background risk arising from risk sharing within the firm - can be easily extended to assess the importance of these other channels.

Insurance demand. First, as noticed by Doherty and Schlesinger (1983), Schlesinger (2000), Eeckhoudt and Kimball (1992) and Gollier and Pratt (1996), in the face of background risk individuals can demand more insurance against risks that are formally insurable, such as property or health shocks. Guiso and Jappelli

(1998) provide some evidence of this channel using subjective measures of earnings risk. These measures are subject to some of the concerns that residual earnings measures raise and thus result in downward biased estimates of the effects on insurance demand.

Hedging portfolio. Second, the observation of household portfolios and of the degree of risk sharing within the firm can help address one of the big puzzles in finance: why households do not hold the same risky portfolio, as suggested by the standard CAPM model. One prominent explanation of portfolio heterogeneity, first put forward by Myers (1972) is hedging of human capital risk: individuals hold different portfolios because of the presence of non-tradable assets - notably human capital - that are specific to the individual. If the correlation between the human capital of worker i with stock j differs across workers - then these workers can: a) rely on the stock market to partially hedge the human capital risk; b) the heterogeneity in the patterns of correlations between each worker human capital and each available stock leads workers to invest more heavily - relatively to the market portfolio - in stocks that correlate poorly, possibly negatively, with their human capital, and dilute exposure to stocks that correlate with their human capital. Each worker will have a different portfolio reflecting this heterogeneity. In turn, human capital risks that can reasonably correlate with existing stocks are not those related to worker specific idiosyncratic shocks (these should be orthogonal to any stock in the economy), but those related to workers exposure to the riskiness of their firms. Testing the theory requires properly measuring the latter.

Suppose one is able to match workers with the firms they work for, and can observe workers' complete stock portfolios, as is the case in the Norwegian data. Then hedging theory can be tested by estimating the following model:

$$\omega_{ijst} - \omega_{st}^{M} = \eta(cov(\theta f_{jt}, r_{st})) = \eta^{'}(cov(f_{jt}, r_{st}))$$

where ω_{ijst} denotes the share invested by worker i, employed by firm j in the traded stock s and ω_{st}^{M} the share of stock s in the market portfolio; f_{jt} is the shock to firm j performance and r_{st} the return on stock s. If there are N traded stocks, there will be N covariances for each firm j and $\eta' = \eta\theta$ is the *measured* hedging response by the workers in firm j which reflects both the extent of the hedging motive, η, as well as the degree of risk sharing within the firm, θ. Using the Fagereng, Guiso and Pistaferri (2018) approach one can obtain an estimate of θ and estimate η' from the equation above, and then retrieve the hedging motive η.

Interestingly, this strategy matches very well with the one of Berk and Walden (2013). Their point is that if firms provide already a lot of insurance, the demand for hedging through the stock market diminishes greatly. But not because the hedging motive is moot but because firms offer already substantial insurance.

A third way through which workers can react to the background risk that arises from partial insurance of firms performance shocks is to share these risks dynastically - for examples through transfers from parents of the two spouses. The Norwegian data allow to link individuals intergenerationally and coupled with observation of uninsured shocks to worker's firms can be used to assess the relevance of this channel as well.

Ongoing research by Fagereng, Guiso and Pistaferri will provide evidence for all these additional reactions. Knowledge on the relevance of all the various margins individuals respond to uninsurable risks will shed light on people's ability to cope with background risk.

5 Risk sharing among co-workers: the extended insurance role of the firm

So far we have identified the insurance role of the firm as facilitating an insurance exchange between capitalists and workers, whereby the former insulate the latter from firm performance volatility in exchange for a lower wage. The exchange is informal but is sustained by the repeated nature of the labor contract that binds workers to the firm. As we have noted, besides being exposed to firm shocks, workers' earnings are also subject to idiosyncratic shocks, the component v_{it} in the residual earnings representation

$$\omega_{ijt} = v_{it} + \theta^P Q_{jt} + \theta^T f_{jt}^T$$

Variation in earnings due to v_{it} arises from shocks to the productivity of the worker, such as those due to health or demographic shocks. These represent an independent source of consumption volatility which can hamper consumer welfare if left uninsured. The firm can potentially help smooth the effect on consumption of these and of other consumption-relevant shocks. However, this may occur not through implicit contracts between capitalists and workers, but rather through exchanges occurring among co-workers. The repeated interaction of individual workers with their firm produces, at the same time, also repeated interactions among workers employed at the same firm, which come to know each other because of intense daily interplay. The enduring relationship between workers and firms, possibly strengthened by the build up of a corporate identity (Akerlof and Kranton, 2005), can greatly facilitate co-workers' insurance exchanges with respect to workers' specific earnings and consumption shocks. A sound study of this role of the firm is still, as far as we know, missing and its empirical relevance has yet to be established.

6 Discussion and conclusion: the future insurance role of the firm

Technological and organizational progress and competitive pressures are changing the insurance role of the firm. More jobs are becoming stand-alone ones and important segments of the production process are outsourced or parcelized through vertical or horizontal dis-integration. For instance, in the 2004-2010 period, micro-businesses (1 to 4 employees) in the US created, on net, about 5.5 million jobs while large businesses (those with greater than 500 employees) lost 1.8 million jobs. Consistent with this, Katz and Krueger (2019) show that contingent workers, such as independent contractors and freelancers, make up an increasing share of the workforce. This is partly a reflection of a tendency of large public companies with well diversified shareholders to outsource to smaller businesses with concentrate ownership (see for instance Segal and Sullivan, 1997, and Weil, 2014). But because the first have a much greater capacity to offer wage insurance than the second, outsourcing results in a diminished role of the firm in the economy as a supplier of wage insurance.[8]

As documented by Benabou and Tirole (2016), competitive pressure is enhancing the reliance on high-powered incentives within the firm, requiring workers to share risks to a larger extent than in the past. Lemieux, MacLeod, and Parent (2009) document that between the late 1970s and the 1990s, the fraction of jobs with wages paid on the basis of performance increased from 38% to 45% and for salaried workers from 45% to 60%. Similar trends are observed for the UK, as documented by Bloom and Van Reenen (2010). Importantly, as Tirole and Benabou (2015) argue, the same competitive pressure that spread the "bonus culture" weakens intrinsic incentives, undermining the firm worker relation and thus the sustainability of

[8] Goldschmidt and Schmieder (2017) document a similar trend in domestic outsourcing in Germany and they argue that this explains a rise in earnings inequality, a feature that is consistent with a greater exposure of workers to firm shocks. .

implicit insurance contracts. All these forces suggest that the Knight-Azariadis-Bailey view of the firm as a generous insurance provider will probably fade out as work becomes more contingent.

On the other hand, technological developments are creating new job types that provide direct buffers against firm-specific shocks. The so called Gig economy jobs (i.e., driving a Uber or Lyft) have a unique feature: they can be accessed flexibly and used as a source of income by salaried workers to compensate for negative shocks to earnings in their main occupation. Fos et al (2019) document that the introduction of the Gig economy changes the way workers respond to job loss. Following a lay off, fewer apply for unemployment insurance, rely less on debt and experience fewer delinquencies. All this suggests that there is less of need for the firm to provide insurance as the labor market itself does.

In sum, the technological forces that are compressing the role of the firm as an insurance provider are also expanding the effectiveness of the market in allowing workers to absorb adverse shocks to their firms and their earnings.

References

[1] Abowd, John M., Francis Kramarz, and David N. Margolis (1999), "High Wage Workers and High Wage Firms." Econometrica 67(3), 251–333.

[2] Akerlof, George A. and Rachel E. Kranton (2005), "Identity and the Economics of Organizations", Journal of Economic Perspectives, 19 (1) : 9–32

[3] Azariadis, Costas (1975), "Implicit Contracts and Underemployment Equilibria." Journal of Political Economy 83 (December): 1183–1202.

[4] Baily, Martin N. (1974), "Wages and Employment under Uncertain Demand." Review of Economic Studies 41 (January): 37–50.

[5] Benabou, Roland and Jean Tirole (2016), "Bonus Culture: Competitive Pay, Screening, and Multitasking", Journal of Political Economy, 124 (2): 305-370

[6] Berk, Jonathan B., and Johan Walden (2013), "Limited Capital Market Participation and Human Capital Risk," Review of Asset Pricing Studies, 3 (1), 1-37.

[7] Blanchflower, David G., Andrew J. Oswald, and Peter Sanfey (1996), "Wages, Profits, and Rent-Sharing." Quarterly Journal of Economics, 111(1), 227–51.

[8] Bloom, Nick and John Van Reenen (2007), "Measuring and Explaining management Practices Across Firms and Countries", The Quarterly Journal of Economics, CXXII (4): 1351-1408.

[9] Bloom, Nick, and John Van Reenen (2010), "Human Resource Management and Productivity." Discussion Paper no 982 May, Centre for Economic Performance, London School Economics.

[10] Browning, Martin and Lusardi, Annamaria (1996) Household Saving: Micro Theories and Micro Facts. Journal of Economic Literature, 34, 1797-1855.

[11] Cardoso, Ana R. and M. Portela (2009), "Micro Foundations for Wage Flexibility:Wage Insurance at the Firm Level", Scandinavian Journal of Economics, 111(1): 29–50.

[12] Carroll, Christopher D. (1992), "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence", Brookings Papers on Economic Activity, 1992(2), 61-156.

[13] Dynan, Karen (1992), "How Prudent are Consumers?", Journal of Political Economy, 101, 6, 1104-1113.

[14] Christofides, Louis N., and Andrew J. Oswald (1992), "Real Wage Determinationand Rent-Sharing in Collective Bargaining Agreements." Quarterly Journal of Economics 107(8), 985–1002.

[15] Clark, Kenneth, Derek Leslie, and Elizabeth Symons (1994), “The Costs of Recessions.” Economic Journal 104 (1), 20–36.

[16] Cunha, Flavio, James J. Heckman, and Salvador Navarro (2005), “Separating Uncertainty from Heterogeneity in Life Cycle Earnings”, Oxford Economic Papers, 57(2), 191-261.

[17] Deaton, Angus (1991), “Saving and Liquidity Constraints”, Econometrica, 59 (5): 1221-1248.

[18] Doherty, Neil A and Schlesinger, Harris, (1983), “Optimal Insurance in Incomplete Markets”, Journal of Political Economy, 91(6), 1045-54.

[19] Eekhoudt Louis, and Miles Kimball (1992), "Background Risk, Prudence and the Demand for Insurance", in Contributions to Insurance Economics, G. Dionne ed. London: Kluwer Academic Press.

[20] Ellul, Andrew, Marco Pagano and Fabiano Schivardi (2018), “Employment and Wage Insurance within Firms: Worldwide Evidence”, The Review of Financial Studies, Volume 31, Issue 4 : 1298–1340,

[21] Estevao, Marcello, and Stacey Tevlin (2003), “Do Firms Share Their Success withWorkers? The Response of Wages to Product Market Conditions.” Economica 70 (11), 597–617.

[22] Fagereng, Andreas, Guiso, Luigi, Davide Malacrino and Luigi Pistaferri (2019), “Heterogeneity and Persistence in Returns to Wealth”, Econometrica (*forthcoming*).

[23] Fagereng, Andreas, Guiso, Luigi, and Luigi Pistaferri (2017), “Risk, Self-Insurance, and Precautionary Saving Response”, American Economic Review P&P, 107 (5): 393-397.

[24] Fagereng, Andreas, Guiso, Luigi, and Luigi Pistaferri (2018), “Portfolio choices, firm shocks and uninsurable wage risk”, Review of Economic Studies, 85 (1): 437–474

[25] Fos, Vyacheslav, Naser Hamdi, Ankit Kalda and Jordan Nickerson (2019), “Gig-Labor: Trading Safety Nets for Steering Wheels”, Boston College working paper .

[26] Friedrich, Benjamin, Lisa Laun, Costas Meghir, and Luigi Pistaferri (2019), “Earnings Dynamics and Firm-Level Shocks”, NBER Working Paper 25786.

[27] Gamber, Edward N (1988), "Long-term Risk-Sharing Wage Contracts in an Economy Subject to Permanent and Temporary Shocks," Journal of Labor Economics, University of Chicago Press, 6(1), 83-99.

[28] Goldschmidt, Deborah and Johannes F. Schmieder (2017), “The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure,” The Quarterly Journal of Economics, 132 (3): 1165–1217.

[29] Gollier, C. and J.W. Pratt (1996). “Risk Vulnerability and the Tempering Effect of Background Risk,” Econometrica 5, 1109–1123.

[30] Guertzgen, Nicole (2014), “Wage Insurance within German Firms: do Institutions Matter?”, Journal of the Royal Statistical Society, 177(2): 345-369.

[31] Guiso, Luigi and Tullio Jappelli (1998), "Background Uncertainty and the Demand for Insurance Against Insurable Risks," The Geneva Risk and Insurance Review, 23(1), 7-27.

[32] Guiso, Luigi, Tullio Jappelli and Daniele Terlizzese (1992), ”Earnings uncertainty and precautionary saving”, Journal of Monetary Economics, 30, 2, 307-337.

[33] Guiso, Luigi, Luigi Pistaferri and Fabiano Schivardi (2005), “Insurance within the Firm”, Journal of Political Economy, 113(5), 1054-1087.

[34] Guiso, Luigi, Luigi Pistaferri and Fabiano Schivardi (2013), “Credit within the Firm”, Review of Economic Studies, 80, 211-247.

[35] Guvenen, Fatih and Anthony Smith (2014), "Inferring Labor Income Risk and Partial Insurance from Economic Choice", Econometrica, 82(6), 2085–2129.

[36] Hartman-Glaser Barney, Hanno Lustig and Mindy X. Zhang (2019), “Capital Share Dynamics When Firms Insure Workers”, Journal of Finance *(forthcoming)*.

[37] Holmstrom, Bengt and Paul Milgrom (1987), “Aggregation and Linearity in the Provision of Intertemporal Incentives”, Econometrica, 55 (2), 303-28.

[38] Jappelli, Tullio and Luigi Pistaferri (2017), *The Economics of Consumption*, Oxford University Press: New York.

[39] Kátay, Gábor (2015), “Do Firms Provide Wage Insurance Against Shocks?”, Scandinavian Journal of Economics, 118(1): 105-128.

[40] Katz, Lawrence F. and Alan B. Krueger (2019), “The Rise and Nature of Alternative Work Arrangements in the United States”, 1995-2015,” Industrial and Labor Relations Review, 72 (2):382-416.

[41] Kimball, Miles (1990), “Precautionary Saving in the Small and in the Large,” Econometrica, January 1990, 58 (1), 53–73.

[42] Kline, Patrick, Neviana Petkova, Heidi Williams, and Own Zidar (2019), “Who Profits from Patents? Rent-Sharing and Innovative Firms”, Quarterly Journal of Economics, 134, 1343–1404.

[43] Knight, Frank H. (1921), *Uncertainty and Profit*. University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship.

[44] Lamadon, Thibaut (2016), "Productivity Shocks, Long-Term Contracts and Earnings Dynamics", University of Chicago working paper.

[45] Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2019), "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market", NBER Working Paper 25954.

[46] Lemieux, Thomas, W. Bentley MacLeod and Daniel Parent (2009), "Performance Pay and Wage Inequality", Quarterly Journal of Economics, 124(1): 1-49.

[47] Low, Hamish, Costas Meghir and Luigi Pistaferri (2010), "Wage Risk and Employment Risk over the Life Cycle", American Economic Review, 100(4), 1432–1467.

[48] Lucas, Robert E., Jr. (1987), *Models of Business Cycles*. Oxford: Blackwell.

[49] Myers, David (1972), "Non-marketable assets and capital market equilibrium under uncertainty". In "Studies in the Theory of Capital Markets", ed. M. C. Jensen. Praeger, 223–48.

[50] Moskowitz, Tobias J. and Annette Vissing-Jorgensen (2002), "The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle"? American Economic Review 92 (4): 745–778.

[51] Pagano, Marco (2019), "Sharing Risk within the Firm: A Primer", Foundations & Trends in Finance, *forthcoming*.

[52] Schlesinger, Harris (2000), "The Theory of Insurance Demand". In: Dionne G. (eds) Handbook of Insurance. Huebner International Series on Risk, Insurance, and Economic Security, vol 22. Springer, Dordrecht.

[53] Segal, Lewis M. and Daniel G. Sullivan (1997), "The Growth of Temporary Services Work," Journal of Economic Perspectives, 11 (2): 117–136.

[54] Van Reenan, John (1996), "The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies," Quarterly Journal of Economics, 111,195–226.

[55] Weil, David (2014), "The Fissured Workplace", Cambridge, MA: Harvard University Press.