Insurance within the Firm

Luigi Guiso

Universitá di Sassari, Ente Luigi Einaudi, and Centre for Economic Policy Research

Luigi Pistaferri

Stanford University, Stanford Institute for Economic Policy Research, and Centre for Economic Policy Research

Fabiano Schivardi

Bank of Italy and Centre for Economic Policy Research

We evaluate the allocation of risk between firms and their workers using matched employer-employee panel data. Unlike previous contributions, this paper focuses on idiosyncratic shocks to the firm, which are the correct empirical counterpart of the theoretical notion of diversifiable risk. We allow for both temporary and permanent shocks to output and find that firms absorb temporary fluctuations fully but insure workers against permanent shocks only partially. Risk-sharing considerations can account for about 15 percent of overall earnings variability, the remainder originating from idiosyncratic shocks to in-

This paper is a revised version of CEPR Discussion Paper no. 2793. We thank two anonymous referees and the editor for very constructive comments; Piero Casadio, Denis Gromb, Cris Huneeus, Marcel Jansen, Tullio Jappelli, Francis Kramarz, Tom MaCurdy, Claudio Michelacci, John Pencavel, Ed Vytlacil, Sevin Yeltekin, and seminar participants at the Institute for the Study of Labor (Bonn), Northwestern University, the University of Salerno, the University of Turin, the Bank of Italy, the 2001 North American conference of the Econometric Society, the fifth conference of the Society for the Advancement of Economic Theory, and the 2001 conference of the Society for Economic Dynamics for comments on preliminary drafts; and the Istituto Nazionale della Previdenza Sociale for supplying us with the data on workers' compensation. The European Union, the Italian Ministero dell'Universitá e della Ricerca Scientifica, and the Taube Faculty Research Fund at the Stanford Institute for Economic Policy Research provided financial support. The views expressed in this paper are our own and do not involve the Bank of Italy. We are responsible for all errors.

[*Journal of Political Economy,* 2005, vol. 113, no. 5]

dividual workers. Our welfare calculations indicate that firms are an important vehicle of insurance provision.

I. Introduction

The idea that an intrinsic component of the entrepreneur-worker relationship is the allocation of risk has a long tradition in economics, dating back at least to Knight (1921), who ascribes the very existence of the firm to its role as an insurance provider.[1] This view underlies the theory of the firm as an insurance device formalized in the implicit contract model of Baily (1974) and Azariadis (1975): risk-neutral entrepreneurs provide insurance to risk-averse workers and insulate their salaries from adverse shocks to production.

Modern finance theory stresses the differential access of firms and workers to financial markets rather than differences in preferences: firms (shareholders) can 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, could be an extreme one. For one thing, empirical studies of households' portfolios show that investment in equities tends to be concentrated, particularly for private equity owners.[2] For another, the optimal level of insurance may depend on the persistence of the shocks. Gamber (1988) shows that if the provision of insurance in an implicit contract model is constrained by the possibility of bankruptcy, persistent shocks to performance are less likely to be insured than temporary shocks. Finally, modern contract theory has emphasized the role of unobservability of effort as an obstacle to insurance provision by the firm, showing that the optimal level of insurance will depend on specific characteristics of the contracting parties, such as differences in risk aversion, the variability of performance, and other deviations from the simple benchmark (e.g., Holmstrom and Milgrom 1987).

Ultimately, the amount of insurance that firms provide to their workers is an empirical question. Previous empirical work has relied on ag-

[1] As Knight puts it, "the system under which the confident and venturesome assume the risk and insure the doubtful and timid by guaranteeing to the latter a specified income in return for an assignment of the actual results . . . is the enterprise and wage system of industry. Its existence in the world is the direct result of the fact of uncertainty" (1921, 269–70).

[2] Moskowitz and Vissing-Jorgensen (2002) find that U.S. households with private equity ownership invest, on average, almost two-thirds of their private holdings in a single company in which they have an active management interest. In this case, the full-diversification assumption clearly fails, and equity owners might well seek to shift part of the enterprise risk onto their employees.

gregate data to determine this. Given the role of risk diversification, aggregate data are not the most appropriate for a convincing test of the theory. Aggregate shocks are common to all firms and as such are undiversifiable, regardless of technology. Besides, a positive correlation between measures of aggregate profitability and aggregate (or even disaggregated) wages may just reflect the equilibrium response of wages to shifts in the demand for labor and have nothing to do with insurance considerations. In addition, research so far has neglected the distinction between less and more persistent variations in firm productivity. This distinction is, as we shall see, empirically relevant for a correct understanding of the wage-profit relationship

This paper casts the wage insurance test in its correct setting: within the firm. We rely on linked employer-employee longitudinal data with enough information to compute measures of shocks (temporary and permanent) both to the firm and to its employees' wages.[3] To obtain such information, we merge company-level data for a large sample of Italian firms with social security data available for a random sample of their employees. Our data provide a unique opportunity to test risk allocation between firms and workers because they enable us to isolate the *idiosyncratic* shocks faced by firms and workers, which is essential to any convincing test of the insurance hypothesis.

The basic idea behind our test is that, under full insurance, idiosyncratic innovations to wages should be orthogonal to idiosyncratic innovations to firm performance. In other words, with full insurance, wages should be independent of firm performance *after* conditioning on predictable and common components. Under the alternative, our test allows us to measure the amount of wage variability that is attributable to low- and high-frequency shocks to performance and to determine whether the amount of insurance offered varies with the observable characteristics of the firm and of its employees.

We find that firms insure workers fully against transitory shocks but only partially against enduring shocks. Furthermore, the amount of insurance provided varies with worker and firm characteristics in ways consistent with the predictions of standard models of wage insurance. These results are robust to a number of extensions, such as accounting for sample selection and for alternative measures of shocks.

According to our estimates, for the average worker, risk-sharing considerations account for about 15 percent of the overall earnings vari-

[3] The recent development of matched worker-firm data has allowed researchers to control for heterogeneity on both sides of the labor contract (Abowd, Kramarz, and Margolis 1999). Matched data were first used to test the relation between firm performance and wage by Bronars and Famulari (2001). They used a two-year survey of workers and firms to look at the wage-profit relationship in the United States. But their data set is too limited in size and duration to distinguish idiosyncratic risk from market-related risk.

ability, most of which originates from idiosyncratic shocks to the individual. A simple calculation of the social value of wage insurance shows that wages are remarkably well insulated against shocks to firms. In fact, starting from the no-insurance benchmark, we calculate that, for realistic values of risk aversion, workers would be willing to give up about 9 percent of their wages to obtain the insurance the firm actually provides against idiosyncratic shocks to its performance. These results have important macroeconomic implications because they indicate that firms' compensation policy greatly reduces labor income variability, a type of insurance hardly replicable by a market mechanism because of the usual moral hazard problems. This, in turn, has significant effects on households' access to credit and their decisions on such matters as savings, portfolio allocation, and labor supply.

The paper is laid out as follows. In Section II we review the institutional aspects of wage determination in Italy and present our data. Section III characterizes our empirical approach, considering a stochastic specification for firms' performance and workers' earnings. In Section IV we show that in line with the wage insurance hypothesis, there exists a set of orthogonality conditions that can be used to answer a series of empirically relevant questions. In particular, one can determine whether shocks to firms' performance are passed on to wages, and to what extent this depends on whether the shock is transitory or permanent. Section V discusses the main results of the estimates, and Section VI conducts some sensitivity analysis and extensions, insofar as several wage insurance models hold that the degree of insurance depends on specific characteristics of workers and firms. Section VII presents conclusions.

II. The Institutional Background and the Data

A. *Institutional Background*

Like most European countries, Italy is characterized by widespread unionization. Although the wage insurance hypothesis examined by Azariadis (1975) and others is cast in the framework of a competitive labor market with incomplete credit and insurance markets, subsequent work (e.g., Riddell 1981) has shown not only that implicit wage insurance may arise in bargaining models with unions but also that unions may play an important role in making implicit contracts feasible, since they can mitigate the enforcement problems that typically arise (see, e.g., Grossman and Hart 1981; Malcomson 1983). However, if wages were set through a highly centralized bargaining process, there could be little room left for them to react to firm-specific shocks, and this would hamper our empirical strategy for testing the wage insurance model. In fact, Italian industrial relations are based on multitier collective bargaining, with economywide, industrywide, and company-level agreements. The latter, as we illustrate below, provide sufficient room for wages potentially to respond to idiosyncratic shocks.

Collective contracts signed by the three major trade unions (CGIL, CISL, and UIL) have *erga omnes* validity; that is, they apply to all workers covered by the agreement regardless of union membership. Economywide bargaining deals mainly with general employment regulations, such as safety and employment protection. The relevant tiers for wage formation are at the industry and company levels. Agreements at the industry level were signed every three years during the period covered by our data, establishing the contractual minimums for the various job grades. Supplementary components of the compensation package are determined at the company level, where the firm can decide on some components of the compensation unilaterally and can also agree to a company contract with the unions, covering both wage and nonwage matters. Firm-level contracts are not obligatory, and there is no standard rule governing their duration.

The relevance of the firm to wage determination has been evolving with industrial relations.[4] Our data cover the years from 1982 to 1994, when there was a fairly high degree of bargaining decentralization.[5] For this period, the wage bill can be decomposed into the following components:[6]

1. contractual minimums (*minimi tabellari*), established at the industry level;
2. cost-of-living allowance (*indennità di contingenza*), added to the contractual minimum according to the inflation rate;

[4] Until the 1960s, company-level bargaining was not formally recognized. The economic boom encouraged the rise of company-level bargaining, at that time mainly focusing on wages and productivity and essentially autonomous vis-à-vis industrywide agreements. The 1970s saw the greatest development of company-level agreements. Starting toward the end of that decade, the worsening economic crisis and the unions' concern over unemployment led to a gradual reshaping of company-level bargaining. A major restructuring of the industrial relations system came in 1993, following devaluation and the severe recession that ensued. Given that our data cover only up to 1994, the 1993 agreement plays essentially no role, and the industrial relations regime we have to consider is the relatively autonomous one in effect previously.

[5] Iversen (1998) constructs an index of centralization of wage bargaining that combines a measure of union concentration with a measure of the prevalent level of bargaining. The index covers 15 OECD countries from 1973 to 1995. He divides his sample into three groups: centralized (Norway, Sweden, Denmark, Austria, and Finland), intermediately centralized (Netherlands, Germany, Belgium, Japan, and Switzerland), and decentralized (Italy, United Kingdom, France, United States, and Canada), ranked in order of degree of centralization. The index ranges from 0.071 (United States and Canada) to 0.538 (Norway), with Italy having a value of 0.179. For comparison, the United Kingdom has a value of 0.177, France 0.121, and Switzerland 0.25.

[6] See Erickson and Ichino (1995) for further details on wage formation in Italy for the period covered by our data.

3. company-level wage increment (*superminimum*), added to the contractual minimum on a permanent basis (in nominal terms); the increment has a firm-level and a worker-level component;
4. production bonuses (*premi di produzione*), determined at the firm level: these are bonuses and other one-time payments, decided unilaterally by the firm without formal negotiations with the trade unions; they are not permanent;
5. variable compensation (*retribuzione variabile*), determined by a firm-level contract; it can introduce a contingent component in the compensation package.

The importance of firm-specific wage effects depends on the diffusion of company-level bargaining and on the magnitude of the firm-specific wage components. In terms of diffusion of firm-level contracts, the yearly report of CESOS, an association of trade unions, indicates that approximately half of Italian workers were involved in firm-level contract negotiations each year from 1984 to 1994.[7] The likelihood of firm-level contracts increases with firm size (Bellardi and Bordogna 1997). As we shall see, in our sample, large firms are disproportionately represented, implying a greater importance of firm-level contracts than for the average firm.

A breakdown of the wage bill into its various components is not available for the economy as a whole. The data set that is most extensively used to address this issue uses wage formation data in the metal products, machinery, and equipment industry assembled by Federmeccanica, the association of employers for that sector. Though a partial view with respect to our analysis and data, which encompass all the private sectors, these data provide insights that are likely to extend to the economy as a whole because the patterns are induced by quite general institutional features (Rossi and Sestito 2000). Table 1 gives the breakdown of the average wage into the five components discussed above for 1984–94 (approximately the same as our data, which cover the years 1982–94). The industrywide component of the wage declined from 84 percent in 1984 to 77 percent in 1994. This means that between one-sixth and one-quarter of the compensation was firm-specific, with quantitatively important scope for firms to influence the wages of their employees over and above industry-level bargaining.

Taken together, then, the institutional setting and the wage formation process indicate that an important component of the compensation of

[7] The 1994 Bank of Italy survey on manufacturing firms with at least 50 employees found that more than 92 percent of these workers were covered by a firm-level contract in addition to the industrywide agreement and that 40 percent had been involved in a contractual round during the year (information on previous years is not available).

TABLE 1
Wage Bill Decomposition for the Machinery Industry, 1984–94

Year	Contractual Minimum (1)	Indexation (2)	Industry Share (Cols. 1+2) (3)	Super-minimum (4)	Production Bonuses (5)	Variable Component (6)	Firm Share (Cols. 4+5+6) (7)
1984	32.4	51.9	84.3	13.0	2.7	.0	15.7
1985	32.1	51.2	83.3	14.2	2.5	.0	16.7
1986	30.1	51.3	81.4	15.5	3.1	.0	18.6
1987	31.2	50.0	81.2	15.8	3.0	.0	18.8
1988	30.4	47.5	77.9	17.3	3.1	1.7	22.1
1989	29.5	47.8	77.3	16.8	3.4	2.5	22.7
1990	28.1	48.3	76.4	17.9	3.3	2.4	23.6
1991	30.3	48.0	78.3	16.6	3.0	2.1	21.7
1992	31.0	46.8	77.8	17.2	3.0	2.0	22.2
1993	32.7	45.2	77.9	17.3	3.0	1.8	22.1
1994	32.5	44.6	77.1	18.2	2.9	1.8	22.9

Source.—Federmeccanica, the association of employers for the metal products, machinery, and equipment sector.
Note.—Each entry represents the percentage contribution to the wage bill. Col. 1 is the contractual minimum; col. 2 is the indexation component; and col. 3 is the sum of the cols. 1 and 2, which constitutes the industrywide component of the wage. The remaining columns represent the firm-level components: the superminimum (col. 4) is the wage premium above the contractual minimum, the production bonuses (col. 5) are bonuses and other one-time payments, the variable component (col. 6) is determined by firm-level contracts, and col. 7 is the sum of the firm-level components.

employees is determined at the level of the firm. This begs the question of how much the firm insures its employees against firm-specific shocks.

B. *The Data*

We rely on two administrative data sets, one for firms and one for workers. Data for firms are obtained from the Company Accounts Data Service (Centrale dei Bilanci [CB]), and those on workers are supplied by the National Institute for Social Security (Istituto Nazionale della Previdenza Sociale [INPS]). Since for each worker we can identify the firm, we combine the two data sets and use them in a matched employer-employee framework.[8]

The CB data span from 1982 to 1994, which comprises two complete business cycles, and give detailed information on a large number of balance sheet items together with a full description of firm characteristics (location, year of foundation, sector, and ownership structure), plus other variables of economic interest usually not included in balance sheets, such as employment and flow of funds. Company accounts are collected for approximately 30,000 firms per year by the service, which was established in the early 1980s jointly by the Bank of Italy, the Italian Banking Association, and a pool of leading banks to gather and share information on borrowers. Since the banks rely heavily on it in granting and pricing loans, the data are subject to

[8] There is a burgeoning empirical literature on the use of matched employer-employee data sets (see Hamermesh [1999] for an account).

extensive quality controls by a pool of professionals, so measurement error should be negligible. While the CB data are reasonably representative of the entire population in terms of distribution by sector and geographical area (Guiso and Schivardi 1999), the focus on level of borrowing skews the sample toward larger firms: CB reporting firms account for approximately half of total employment and 7 percent of the number of firms in manufacturing.

The INPS provides us with 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. Data are available on a continuous basis from 1974 to 1994. We use the data after 1981 for consistency with the timing of the CB data. The INPS data do not cover self-employment or public employment (public firms are also absent from the CB). The INPS data set derives from employer forms roughly comparable to those collected by the Social Security Administration in the United States.[9] Misreporting is prosecuted.

Both the INPS and the CB data include the employer's tax code, so matching employers with the employees is straightforward. Like other countries' social security data, the Italian INPS data contain detailed information on worker compensation, but information on demographics is scant.

Table 2 reports various descriptive statistics for the firms (panel A) and workers (panel B) using 1991 as an example. We report separate statistics for the whole sample (before matching) and for the sample obtained after matching. From an initial sample of 177,654 firm/year observations, we end up with 122,860, corresponding to 17,272 firms. We exclude firms with intermittent participation (40,225 observations) and those with missing values on the variables used to construct the test (14,569 observations). Since the panel is unbalanced, the firms in this sample may appear from one to 13 years.

The firms range from very small to a workforce of 180,000, the mean staff size is 194, and the median is 60. Most of the firms are in the North (74 percent). Manufacturing accounts for about 75 percent of the final sample and construction for 15 percent, whereas the remaining 10 percent is scattered in the retail and other service sectors. On the average, firms in the matched sample are larger, but their distribution by region and industry is similar to that of the whole sample.

Panel B reports sample characteristics for the workers in the 1982–

[9] While the U.S. administrative data are usually provided on a grouped basis, the INPS has truly individual records. Moreover, U.S. earnings records are censored at the top of the tax bracket, whereas the Italian data set is not subject to top coding.

TABLE 2
Worker and Firm Characteristics

	Mean		Standard Deviation	
	Whole Sample	Matched Sample	Whole Sample	Matched Sample
	A. Firm Characteristics			
Value added	8.27	16.14	123.02	156.56
Number of employees	194	381	2,281	3,526
South	.0917	.1089	.2887	.3115
Center	.1643	.1686	.3706	.3744
North	.7439	.7226	.4365	.4478
Manufacturing	.7722	.8052	.4194	.3961
Construction	.1573	.1231	.3641	.3287
Retail trade	.0257	.0230	.1582	.1499
Services	.0447	.0486	.2068	.2151
	B. Worker Characteristics			
Earnings	13,363	13,929	6,004	6,189
Age	40.83	42.33	9.69	9.01
Male	.7454	.7657	.4357	.4236
Production workers	.6275	.6362	.4835	.4812
Clerical workers	.3556	.3452	.4787	.4755
Managers	.0169	.0186	.1289	.1352
South	.1446	.1113	.3517	.3146
Center	.1837	.1845	.3872	.3879
North	.6717	.7042	.4696	.4565
Manufacturing	.7290	.8197	.4445	.3845
Construction	.1161	.0911	.3203	.2878
Retail trade	.0849	.0408	.2787	.1978
Services	.0700	.0485	.2552	.2147

Note.—Panel A reports summary statistics for the firms in our data set. Panel B shows descriptive statistics for the sample of workers. All statistics refer to 1991. For firms, the "whole sample" is the full CB sample. For workers, it is the largest sample for which information on industry and location can be recovered by merging the worker sample with the firm sample. This sample is larger than the "matched sample" because time-invariant firm characteristics can be extended to years in which we do not directly observe the firm but, through its fiscal identifier, we can infer that the worker was employed by such firm. The matched data set includes only observations for which we have contemporaneous observations on both the worker and the firm. Value added is in millions of euros and earnings in euros. All the other variables, with the exception of age, are indicator variables.

94 INPS sample.[10] We start with an initial sample of 267,539 worker/year observations (including multiple observations in a year when workers change positions within firms or change firms, or when employers change, etc.) and end up with 130,785, corresponding to 23,788 individuals.[11] Since the focus is wage insurance, our sample selection is made with the explicit aim of retaining workers with stable employment and

[10] The INPS data do not report workers' sector of activity and location. This information can be recovered only after matching workers with firms. The "whole sample" of panel B is the largest sample for which a match is obtained and this information recovered. This sample, which we use to estimate the earnings equation, is larger than the "matched sample" because time-invariant firm characteristics (such as sector of activity and location) can be extended to years in which we do not directly observe the firm but, through its fiscal identifier, can infer that the worker was employed by such firm.

[11] Additional observations are lost (for both firms and workers) in the empirical analysis given the dynamic nature of most of our estimators.

tenure patterns. First, we excluded workers younger than 18 or older than 65 (2,652 observations), circumventing the problem of modeling human capital accumulation and retirement decisions. To avoid dealing with wage changes that are due to job termination (quits or dismissals) or unstable employment patterns, we excluded part-time workers, those changing position during the year or with multiple jobs (81,117 observations), and those who worked for less than 12 months (43,750 observations). In this way we isolate the on-the-job aspect of the wage insurance contract, leaving consideration of changes in occupational status to future studies. Moreover, we keep only individuals with nonzero recorded earnings in all years (105 observations lost) and eliminate some outliers (503 observations).[12] Finally, we eliminate those with missing values on the variables used in the empirical analysis (8,627 observations). Since these selections—particularly those that exclude job and firm changes—could affect the results, we check the robustness of our findings when retaining these observations (see Sec. VI.*A*).

Our measure of earnings covers regular and overtime pay plus nonwage compensation. We deflate earnings using the consumer price index (1991 prices). For workers with intermittent participation, we treat two strings of successive observations separated in time as though they pertained to two different individuals.

Workers in the whole sample are, on average, 41 years old in 1991; production workers account for 63 percent of the sample, clerical for 36 percent, and managers for about 2 percent. Men make up 75 percent of our sample and residents in the South 15 percent. Finally, average net earnings in 1991 are roughly €13,000. In the matched final sample, the individual characteristics are similar. We end up with 45,446 matched observations, pertaining to 9,203 workers and 4,691 firms. By design, the number of matches per firm increases with firm size. On average, a firm is matched with 1.54 workers in any given year (with a minimum of one and a maximum of 63 matches).

III. Modeling Firms' Performance and Workers' Earnings

Our testing strategy is based on the theory-grounded idea that if firms fully insure their workers against diversifiable risk, idiosyncratic innovations to wages should be orthogonal to idiosyncratic changes in firm performance. To construct the orthogonality test, we need a measure of idiosyncratic shocks to performance and to wages. We also need to separate the dynamics of the idiosyncratic shocks to performance into

[12] An observation is classified as an outlier if (*a*) real earnings are under €500, (*b*) they are below €3,000 *and* the change in log real earnings is less than negative two, or (*c*) they exceed €50,000 *and* the change in log real earnings is greater than two.

transitory and permanent shocks in order to analyze whether insurance depends on the persistence of firm performance. In the sections that follow, we provide a statistical characterization of these processes that allows us to net out all the predictable and aggregate components, thus obtaining an estimate of the idiosyncratic innovations.

Before we move into details, it is useful to contrast our approach with previous empirical work. Papers belonging to one strand in the literature regress individual wages on measures of aggregate profits. Blanchflower, Oswald, and Sanfey (1996) and Estevão and Tevlin (2003) use U.S. industrywide profits drawn from the National Bureau of Economic Research productivity database, and Christofides and Oswald (1992) use Canadian data of similar form. Given the role of risk diversification, however, the use of aggregate profits cannot provide a convincing test of the theory. Aggregate profits reflect economywide shocks that are uninsurable by definition. Besides, a positive correlation between measures of aggregate profitability and individual wages may just reflect the equilibrium response of wages to shifts in the demand for labor, with no bearing on insurance considerations.

Another strand of research relies on firm-specific data. Currie and McConnell (1992) and Abowd and Lemieux (1993) use labor union contracts and regress collectively bargained wages on firm-specific profits. Bargained wages, however, exclude bonuses and other components of pay that make up an important part of wage variability. In the Italian case, the institutional features discussed in the previous section would make collectively bargained wages orthogonal to firm-specific profits even if the firm provides no insurance to workers against firm-specific fluctuations in productivity. A final set of contributions is by Nickell and Wadhwani (1991) and Hildreth and Oswald (1997), who regress firm-specific average wages on firm profitability. While this is a step toward a more disaggregated approach, measures of wages at the firm level do not control fully for individual workers' characteristics. For example, changes in wages might be related to changes in profits because both reflect a change in the composition of the firm's workforce rather than risk sharing.

Finally, it is worth stressing that none of the papers reviewed here looks at the separate effect of permanent and transitory shocks to the firm's performance on the extent of risk sharing between the firm and its workers.[13] The distinction between less and more persistent shocks is, as we shall see, empirically crucial for a correct understanding of the wage–firm performance relationship.

[13] The article by Gamber (1988) is an exception, but his test uses industry-level data and as such is subject to the criticisms given in the text.

A. Firms' Performance

Our measure of firm performance is value added, that is, the volume of contractible output that remains once intermediate inputs have been remunerated (i.e., the sum of pretax profits, wages, and perks). We prefer value added to profits for two reasons. First, value added is the variable that is directly subject to stochastic fluctuations. Second, firms have discretionary power over the reporting of profits in balance sheets, which makes profits a less reliable objective gauge.

Our aim is to isolate changes in firm performance. To do so, we model firm performance according to the following process:

$$(1 - \rho L)y_{jt} = \mathbf{Z}'_{jt}\gamma + f_j + \epsilon_{jt}, \tag{1}$$

where j and t are subscripts for the jth firm at time t; y_{jt} is the logarithm of value added at 1991 prices, deflated by the producer price index; L is the lag operator; $\mathbf{Z}_{jt}$ is a vector of strictly exogenous firm characteristics; f_j is a firm fixed effect; and ϵ_{jt} is the shock against which the amount of insurance provided by the firm will be determined. We include the first lag of value added to capture predictable dynamics (e.g., precommitted sales). The role of $\mathbf{Z}_{jt}$ is to control for nonidiosyncratic (i.e., aggregate, location, and industry-specific) shocks.[14]

Taking the first difference of (1) eliminates the fixed effect and produces

$$(1 - \rho L)\Delta y_{jt} = \Delta\mathbf{Z}'_{jt}\gamma + \Delta\epsilon_{jt}. \tag{2}$$

Since ordinary least squares (OLS) estimates are inconsistent because of the lagged dependent variable on the right-hand side, we use the two-step generalized method of moments (GMM) approach of Arellano and Bond (1991). If ϵ_{jt} were serially uncorrelated, consistent estimates could be obtained using lags of y dated $t - 2$ and earlier as instruments. This hypothesis can be tested using the test of serial correlation of the residuals derived by Arellano and Bond. The test statistic is standard normally distributed under the null hypothesis of no serial correlation. The standard test of overidentifying restrictions (J-test) provides further corroboration of the validity of the specification. Under the null hypothesis that the model is correctly specified, the J-statistic is asymptotically distributed χ^2 with as many degrees of freedom as overidentifying restrictions and is robust to heteroskedasticity of unknown form. Low values of the statistic (high p-values of the test) will signal that the model is correctly specified. In our case, both the test of no serial correlation and the test of overidentifying restrictions reject the hypothesis that

[14] In Sec. VI.*B*, we show that the results are robust to an expanded definition of the vector $\mathbf{Z}$, which includes the firm's capital stock and labor input, while also accounting for their endogeneity.

TABLE 3
Value-Added Regression

Variable	Estimate	Standard Error
Value-added growth at $t-1$	.4769	.0303
Center	.0135	.0033
North	.0172	.0028
Year dummies	624.52	
	[<0.0001]	
Sector dummies	122.15	
	[<0.0001]	
Overidentifying restriction test	59.49	
	[.2830]	
Observations	90,659	

Note.—The table reports the results of a GMM regression for value-added growth at time t. Instruments are constructed using the GMM approach of Arellano and Bond (1991) and include the log of value added dated $t-3$ and earlier. For year and sector dummies, F-statistics are reported; values in brackets are p-values.

ϵ_{jt} is serially uncorrelated (with p-values of 0.1 percent and 1 percent, respectively). We thus instrument Δy_{jt-1} using all the available lags of y dated $t-3$ and earlier. The results of the GMM regression (using all firms in the sample, not just those matched) are reported in table 3. The overidentifying restriction test using these instruments no longer indicates misspecification. Our estimate of ρ is 0.48 with a standard error of 0.03, and region, year, and industry dummies are jointly statistically significant.

We next construct the residual from (2), a consistent estimate of $\Delta\epsilon_{jt}$. The estimated autocovariances $E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-\tau})$, reported in table 4, show that there is no large or statistically significant correlation at lags greater than two, confirming our choice of instruments in table 3.

The autocovariances of $\Delta\epsilon_{jt}$ can now be analyzed to gain knowledge

TABLE 4
Autocovariance Structure of Shocks to Value Added

Order (τ)	All Years	Standard Error
0	.1419	.0052
1	−.0635	.0032
2	.0036	.0010
3	.0004	.0011
4	.0008	.0012
5	.0007	.0011
6	.0001	.0013
7	−.0030	.0015

Note.—The table reports the estimates and the corresponding standard errors of the autocovariances at various orders of the residual of value-added growth in first differences, i.e., estimates of $E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-\tau})$. The data are pooled over all years.

about the dynamic structure of the shocks. This is an important step in the analysis because different dynamic structures will imply different characterizations of the wage determination process. The dynamics of $\Delta\epsilon_{jt}$ depends on (*a*) the number of its variance components and (*b*) the separate dynamics of these components. In terms of the latter, examination of the autocovariances at various lags immediately rules out the presence of an autoregressive component, which would have implied a smooth decline, so that autocovariances would not drop to zero as rapidly as they do. As far as the number of components is concerned, the autocovariance structure of $\Delta\epsilon_{jt}$ is consistent with a single variance component following an MA(2) process. The problem with this representation is that it constrains all shocks to have the same dynamic influence on value added.[15] Most processes, however, contain several shock components of both a stationary and nonstationary nature. This seems particularly important when modeling firms' performance; in fact, firms are likely to experience both transitory shocks (such as temporary changes in demand and machine breakdowns) and permanent shocks (such as persistent demand and technological changes). An alternative representation of the process for ϵ_{jt} that accommodates this fact and is also consistent with the autocovariance structure of $\Delta\epsilon_{jt}$ from table 4 is

$$\epsilon_{jt} = \zeta_{jt} + (1 - \theta L)\tilde{v}_{jt} \tag{3}$$

and

$$\zeta_{jt} = \zeta_{jt-1} + \tilde{u}_{jt}, \tag{4}$$

where ϵ_{jt} is now decomposed into the sum of a random walk and an MA(1) component. We assume covariance stationarity of the disturbances $E(\tilde{u}_{jt}^2) = \sigma_{\tilde{u}}^2$ and $E(\tilde{v}_{jt}^2) = \sigma_{\tilde{v}}^2$ for all t, no serial correlation $E(\tilde{v}_{js}\tilde{v}_{jt}) = E(\tilde{u}_{js}\tilde{u}_{jt}) = 0$ for $s \neq t$, and no cross correlation $E(\tilde{v}_{js}\tilde{u}_{jt}) = 0$ for all s, t. By taking first differences of (3) and using (4), we obtain

$$\Delta\epsilon_{jt} = \tilde{u}_{jt} + (1 - \theta L)\Delta\tilde{v}_{jt}. \tag{5}$$

This process is consistent with the autocovariance structure of $\Delta\epsilon_{jt}$ in table 4 because it implies that $E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-\tau}) \neq 0$ for $|\tau| \leq 2$ and $E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-\tau}) = 0$ for $|\tau| > 2$.

Obviously, since $\Delta\epsilon_{jt} \sim$ MA(2) even in the absence of a random walk in the levels, the consistency of (5) with the results in table 4 is a necessary but not sufficient condition to conclude that there is a random

[15] For example, suppose that $\Delta\epsilon_{jt} = \eta_{jt} + \phi_1\eta_{jt-1} + \phi_2\eta_{jt-2}$. Integrating, we get

$$\epsilon_{jt} = \eta_{jt} + (1 + \phi_1)\eta_{jt-1} + (1 + \phi_1 + \phi_2)\sum_{k=2}^{\infty}\eta_{jt-k}.$$

Given that $\lim_{k\to\infty}(\epsilon_{jt+k}/\eta_{jt}) = 1 + \phi_1 + \phi_2$, all shocks have permanent effects unless $\phi_1 + \phi_2 = -1$, in which case they all have only temporary effects.

walk component in the levels. However, one can distinguish between the null $\Delta\epsilon_{jt} = (1-\theta L)\Delta\tilde{v}_{jt}$ and the alternative $\Delta\epsilon_{jt} = (1-\theta L)\Delta\tilde{v}_{jt} + \tilde{u}_{jt}$ by noting that under the null $E[\Delta\epsilon_{jt}(\sum_{\tau=-2}^{2}\Delta\epsilon_{jt+\tau})] = 0$, whereas under the alternative $E[\Delta\epsilon_{jt}(\sum_{\tau=-2}^{2}\Delta\epsilon_{jt+\tau})] = \sigma_{\tilde{u}}^2$ (see Meghir and Pistaferri 2004). If we perform this test using the residuals $\Delta\epsilon_{jt}$, the null hypothesis is rejected quite overwhelmingly (a p-value of less than 0.01 percent), so we conclude in favor of the existence of a non-mean-reverting random walk component in addition to a mean-reverting MA(1) component, as is implied by our characterization of the shocks to firm performance. This characterization has the advantage of allowing for shocks of different persistence, while keeping the dynamic structure of $\Delta\epsilon_{jt}$ as simple as possible.[16]

On the basis of this representation, equation (1) can be inverted and value added rewritten as the sum of a deterministic D_{jt}, a permanent P_{jt}, and a transitory component T_{jt}:

$$y_{jt} = D_{jt} + P_{jt} + T_{jt}, \tag{6}$$

where $D_{jt} = (1-\rho L)^{-1}(\mathbf{Z}'_{jt}\gamma + f_j)$ and $P_{jt} + T_{jt} = (1-\rho L)^{-1}\epsilon_{jt}$. One can show that $P_{jt} = (1-\rho)^{-1}\zeta_{jt}$ and[17]

$$T_{jt} = (1-\rho L)^{-1}[(1-\theta L)\tilde{v}_{jt} - (1-\rho)^{-1}\rho\tilde{u}_{jt}].$$

Taking first differences and premultiplying by $1-\rho L$, we can equivalently write equation (2) as

$$(1-\rho L)\Delta y_{jt} = \Delta\mathbf{Z}'_{jt}\gamma + (1-\rho L)u_{jt} + \Delta v_{jt}, \tag{7}$$

where

$$u_{jt} = \frac{1}{1-\rho}\tilde{u}_{jt} \tag{8}$$

and

$$v_{jt} = (1-\theta L)\tilde{v}_{jt} - \frac{\rho}{1-\rho}\tilde{u}_{jt} \tag{9}$$

are the innovations to the permanent and transitory components of

[16] Implicit in this discussion is the assumption that firms are able to distinguish between temporary and permanent shocks. This is not unreasonable since firms that, say, systematically confuse permanent demand increases for their product for temporary ones will not survive long in a competitive environment and therefore have strong incentives to detect the type of shock they face.

[17] This derivation uses the fact that y_{jt} and ζ_{jt} are cointegrated I(1) and that their cointegration vector is $[1\ -(1-\rho)^{-1}]'$, so that the linear combination $y_{jt} - (1-\rho)^{-1}\zeta_{jt}$ is I(0). The two components P_{jt} and T_{jt} can be found by simple application of the Granger representation theorem. We make no attempt at orthogonalizing P_{jt} and T_{jt} because we identify their separate effect on wages in the unorthogonalized case.

value added, respectively.[18] Note that u_{jt} and v_{jt} are mutually correlated and that v_{jt} is serially correlated.[19]

B. Workers' Earnings

We assume that workers' pay can be described by the following equation:

$$w_{ijt} = \boldsymbol{X}'_{ijt}\delta + \alpha P_{jt} + \beta T_{jt} + h_i + \psi_{ijt}, \tag{10}$$

where the subscript i stands for the ith individual; w_{ijt} is the logarithm of worker compensation;[20] and $\boldsymbol{X}_{ijt}$ denotes a vector of systematic factors that affect individual i's compensation, which can vary across workers, firms, and time (including D_{jt} in [6]). Among other things, $\boldsymbol{X}_{ijt}$ includes region, occupation, industry, and year dummies. These dummies remove any variation in wages due to industrywide and national bargaining (including wage indexation). We let wages depend on the permanent and the transitory components of firm performance through loading factors α and β, respectively. This formulation allows for different levels of insurance against shocks of different persistence, a hypothesis that is testable. We also include workers' fixed effects h_i. Finally, ψ_{ijt} is the stochastic component of earnings unrelated to the firm's fortunes. These idiosyncratic shocks are meant to capture the unobservable component of one's outside wage (individual ability and fluctuations thereof), as well as idiosyncratic changes in labor supply (e.g., prolonged illness, child raising, and family labor supply effects).

Taking first differences to eliminate fixed effects and multiplying by $1 - \rho L$ gives

$$\begin{aligned}(1-\rho L)\Delta w_{ijt} &= (1-\rho L)\Delta \boldsymbol{X}'_{ijt}\delta + \alpha(1-\rho L)u_{jt} + \beta\Delta v_{jt} \\ &\quad + (1-\rho L)\Delta\psi_{ijt} \\ &= (1-\rho L)\Delta \boldsymbol{X}'_{ijt}\delta + \Delta\omega_{ijt}.\end{aligned} \tag{11}$$

Since we have a lagged endogenous variable, (11) must also be estimated

[18] By the definition of permanent and transitory components, the innovation u_{jt} to P_{jt} is such that $\lim_{k\to\infty}(\partial y_{jt+k}/\partial u_{jt}) \neq 0$, whereas the innovation v_{jt} to T_{jt} is such that $\lim_{k\to\infty}(\partial y_{jt+k}/\partial v_{jt}) = 0$.

[19] In particular, $\sigma_{uv} = -\rho\sigma^2_{\tilde{u}}/(1-\rho)^2$ and $\sigma_{vv_{-1}} = -\theta\sigma^2_{\tilde{v}}$. Note also that $\sigma^2_u = \sigma^2_{\tilde{u}}/(1-\rho)^2$ and that $\sigma^2_v = (1+\theta^2)\sigma^2_{\tilde{v}} + [\rho/(1-\rho)]^2\sigma^2_{\tilde{u}}$. Estimation of the primitive parameters ρ, θ, $\sigma^2_{\tilde{u}}$, and $\sigma^2_{\tilde{v}}$ allows us to recover the variances of the innovations σ^2_u and σ^2_v, as well as the covariances σ_{uv} and $\sigma_{vv_{-1}}$.

[20] We let earnings depend on contemporaneous shocks to firm performance; i.e., we assume that wages adjust immediately to changes in performance. In practice, wages might adjust with a lag (think of bonus decisions, which are usually made at the end of the calendar year). Nevertheless, if adjustments are made at a frequency higher than the year (say, quarterly), annual data of the type used here will not detect deviations from the contemporaneous adjustment assumption.

by an instrumental variables procedure.[21] The selection of instruments depends on the extent of serial correlation in $\Delta\omega_{ijt}$. By reasoning analogous to that of the previous subsection, we find that a representation of ψ_{ijt} that is consistent with the data is

$$\psi_{ijt} = \vartheta_{ijt} + (1 - \lambda L)\xi_{ijt} \tag{12}$$

and

$$\vartheta_{ijt} = \vartheta_{ijt-1} + \mu_{ijt}, \tag{13}$$

that is, a permanent plus transitory MA(1) decomposition. Equations (12) and (13) thus imply that $\Delta\omega_{ijt}$ is an MA(3) process. Lagged instruments must be selected accordingly. The first step is to regress Δw_{ijt} on Δw_{ijt-1} and on the exogenous characteristics to estimate ρ and δ.

Our measure of earnings is the log of annual net real earnings. We control for a set of observable individual attributes: a quadratic in age, education (here proxied by a set of occupation dummies), gender, region dummies, sector dummies, and time dummies. Nominal gross earnings are first transformed into nominal earnings net of taxes and social security contributions (using the rules coded in the Italian tax system at each point in time) and then deflated by the consumer price index to 1991 prices. We estimate the earnings growth equation by GMM using all the available lags of w dated $t-4$ and earlier and rely on data for all workers in the whole sample rather than just those in the matched sample.

The results are reported in table 5. The autoregressive coefficient is 0.44 (with a standard error of 0.06), essentially the same as in the firm regression case (a null that is not formally rejected). Earnings growth is higher for men and managers, and increases with age at a declining rate. Year, region, and industry dummies are jointly significant. The overidentifying restrictions test does not signal general misspecification.

We use the residual of this regression to construct a consistent estimate of $\Delta\omega_{ijt}$. We calculate the autocovariances of $\Delta\omega_{ijt}$, pooling over all years, and report the results in table 6. Examination of the estimated autocovariances of the residual component of the rate of growth of earnings reveals that they decline very rapidly after the first order: on average, the autocovariance of order 0 is 0.016, that of order 1, -0.007, and that of order 2, 0.00042.[22] However, they are still statistically different from

[21] Under our hypothesis, the length of the autoregressive process for firm performance carries over to workers' earnings (with the same coefficient if the autoregressive component of wages is exclusively driven by risk-shifting considerations).

[22] These are lower than the estimates for the United States using the Panel Study of Income Dynamics (Meghir and Pistaferri 2004) and those for Italy using the 1995 Survey of Household Income and Wealth (SHIW) (Pistaferri 2001), perhaps reflecting the fact that measurement error is less of a problem in this data set.

TABLE 5
Earnings Equation

Variable	Estimate	Standard Error
Earnings growth at $t-1$	.4360	.0630
Male	.0079	.0008
Age	.0011	.0003
Age^2/100	−.0014	.0003
Production workers	−.0165	.0032
Clerical workers	−.0112	.0026
North	.0040	.0009
Center	.0036	.0010
Year dummies	2,873.04	
	[<.0001]	
Sector dummies	52.25	
	[<.0001]	
Overidentifying restrictions	46.29	
test	[.3780]	
Observations	85,151	

Note.—The table reports the results of a GMM regression for earnings growth at time t. Instruments are constructed using the GMM approach of Arellano and Bond (1991) and include the log of earnings dated $t-4$ and earlier. For year and sector dummies F-statistics are reported; values in brackets are p-values.

TABLE 6
Autocovariance Structure of Shocks to Earnings

Order (τ)	All Years	Standard Error
0	.01634	.00044
1	−.00712	.00026
2	.00042	.00011
3	.00027	.00011
4	−.00019	.00011
5	−.00011	.00012
6	.00014	.00013
7	−.00015	.00014

Note.—The table reports the estimates and the corresponding standard errors of the autocovariances of the unexplained component of real earnings growth, i.e., estimates of $E(\Delta\omega_{ijt}\Delta\omega_{ijt-\tau})$. Data are pooled over all years.

zero until lag 4, which has a borderline p-value of 9 percent. This is again broadly consistent with our choice of instruments. These autocovariances (and, especially, the cross covariances $E[\Delta\omega_{ij\tau}\Delta\epsilon_{js}]$) are further structured in the next section.

IV. Wage Insurance: Identification

Our strategy is to test for insurance within the firm by looking at the relationship between the residual component of individual wage growth,

$\Delta\omega_{ijt}$, and the two components, transitory and permanent, of the idiosyncratic shocks to firm performance, $\Delta\epsilon_{jt}$. From equation (11), $\Delta\omega_{ijt}$ has the following representation:

$$\Delta\omega_{ijt} = \alpha(1-\rho L)u_{jt} + \beta\Delta v_{jt} + (1-\rho L)\Delta\psi_{ijt}. \tag{14}$$

We assume that $E(\psi_{ijs}u_{jt}) = E(\psi_{ijs}v_{jt}) = 0$ for all s, t. Equation (14) allows for different reactions to temporary and permanent shocks to performance. In this way, the estimates of equation (14) allow for various insurance regimes. If workers share part of the fluctuations in firm performance without distinguishing between transitory and enduring shocks, $0 < \alpha = \beta$; we call this case "homogeneous partial insurance." The insurance regime may result in a different reaction to shocks of a different nature. For instance, workers may bear a substantial portion of the permanent shocks but a small portion of transitory shocks; we call this "heterogeneous partial insurance": $\alpha \neq \beta$, $\alpha > 0$, and $\beta > 0$. Two special cases occur when workers bear only transitory shocks and are insulated from permanent shocks ("permanent full insurance": $\alpha = 0$ and $\beta > 0$) or bear permanent shocks but are insured against transitory ones ("transitory full insurance": $\beta = 0$ and $\alpha > 0$). Finally, workers might be completely sheltered from shocks, temporary and permanent alike ("full insurance": $\alpha = \beta = 0$).[23]

Clearly, full insurance implies that the contemporaneous covariance between shocks to performance and to wage growth, $E(\Delta\epsilon_{jt}\Delta\omega_{ijt})$, is zero. This simple case aside, identification of the parameters is far from trivial. In fact, equation (14) contains the contribution of two unobservable components, u_{jt} and Δv_{jt}. Recall that, from the estimation of (2), we obtain a consistent estimate of $\Delta\epsilon_{jt} = (1-\rho L)u_{jt} + \Delta v_{jt}$, so that only the sum of $(1-\rho L)u_{jt}$ and Δv_{jt} is observed. Without further restrictions, from equation (14) we cannot separately identify α and β.

To see how identification of the relevant parameters is achieved, subtract $\beta\Delta\epsilon_{jt}$ from both sides of (14) to obtain

$$\Delta\omega_{ijt} - \beta\Delta\epsilon_{jt} = (\alpha-\beta)(1-\rho L)u_{jt} + (1-\rho L)\Delta\psi_{ijt}. \tag{15}$$

Multiply both sides by $\Delta\epsilon_{jt+1}$ and take expectations to yield the orthogonality condition

$$E[\Delta\epsilon_{jt+1}(\Delta\omega_{ijt} - \beta\Delta\epsilon_{jt})] = 0, \tag{16}$$

which follows from the fact that, from (8) and (9), $\Delta\epsilon_{jt+1} = \tilde{u}_{jt+1} +$

[23] An interesting extension would be to allow for an asymmetric response of wages to positive and negative transitory and permanent shocks. For instance, one could imagine workers to be insured against downside risk but still take part of the upside movements in value added. Clearly, models that allow for such asymmetric effects are not identifiable. The reason is that we do not separately observe transitory and permanent shocks; it is only (a consistent estimate of) their convolution ($\Delta\epsilon_{jt}$) that is observable.

$(1 - \theta L)\Delta\tilde{v}_{jt+1}$ (with $\tilde{u}_{jt+1}$ and $\tilde{v}_{jt+1}$ mutually and serially uncorrelated); hence, $\Delta\epsilon_{jt+1}$ is orthogonal to the right-hand side of (15). Intuitively, equation (16) tells us that once one filters the unexplained component of earnings growth $\Delta\omega_{ijt}$ by the unexplained component of value-added growth $\Delta\epsilon_{jt}$ (weighted by a factor β, the extent to which wages move in response to changes in the transitory component of value added), what remains is uncorrelated with the future unexplained component of value-added growth. In an OLS regression of $\Delta\omega_{ijt}$ on $\Delta\epsilon_{jt}$, the latter is endogenous because it is correlated with the right-hand side of equation (15) via u_{jt}.[24] However, $\Delta\epsilon_{jt+1}$ (and, more generally, any power $[\Delta\epsilon_{jt+1}]^k$ with $k \geq 1$) is a valid instrument because it is correlated with $\Delta\epsilon_{jt}$ ($E[\Delta\epsilon_{jt}\Delta\epsilon_{jt+1}] = -[1 + \theta]^2\sigma_{\tilde{v}}^2$) and uncorrelated with the error term (the right-hand side of [15]). Equation (16) can be used to identify the first parameter of interest, β.

Identification of α is achieved using a similar strategy and amounts to exploiting the orthogonality condition

$$E\left[\left(\sum_{\tau=-2}^{2} \Delta\epsilon_{jt+\tau}\right)(\Delta\omega_{ijt} - \alpha\Delta\epsilon_{jt})\right] = 0, \tag{17}$$

which identifies α under covariance stationarity.[25] In practice, all instruments of the form $(\sum_{\tau=-2}^{2}\Delta\epsilon_{jt+\tau})^m$ (for $m \geq 1$) are valid under this assumption.[26]

The foregoing is a discussion of the identification of the two insurance parameters α and β. To close the circle on identification, we need to identify the variances of the shock to value-added growth (σ_u^2 and σ_v^2) and the moving average coefficient θ, and the variances of the idiosyncratic component of earnings growth (σ_μ^2 and σ_ξ^2) and the moving average parameter λ. In practice, we start by estimating firm-level parameters using minimum distance (Chamberlain 1984). We do this by choosing the parameters that minimize the distance between the actual moments and the moments predicted by the restrictions (3) and (4).

[24] It is worth noting that OLS estimation provides unbiased and consistent estimation if $\alpha = \beta$. Thus an exogeneity (Wu-Hausman) test for $\Delta\epsilon_{jt}$ can implicitly check whether $\alpha = \beta$.

[25] If covariance stationarity is violated, one can prove that (if $\beta \to 0$, a result consistent with our empirical analysis)

$$p\lim(\hat{\alpha} - \alpha) = \frac{\alpha\rho}{1-\rho}\frac{q_t}{1+q_t},$$

where q_t is the growth of $\mathrm{Var}(\tilde{u}_{jt})$. The hypothesis of covariance stationarity for $\mathrm{Var}(\tilde{u}_{jt})$, which we estimate with the sample analogue of $E[(\sum_{\tau=-2}^{2}\Delta\epsilon_{jt+\tau})\Delta\epsilon_{jt}]$, is not rejected ($p$-value of 71 percent).

[26] The second and third powers of the instruments, for instance, are valid as long as the third and fourth moments of $\tilde{u}_{jt}$ and $\tilde{v}_{jt}$, respectively, are nonzero. See Cragg (1997) for a discussion.

From the autocovariance function of workers' earnings, we can estimate, again by minimum distance, the variance of the transitory and permanent idiosyncratic shocks to wages (σ_ξ^2 and σ_μ^2, respectively) and the moving average parameter λ using the restrictions imposed by (12), (13), and (14).

V. Results

We apply the identification strategy outlined in the previous section to estimate the parameters of interest, α (the sensitivity of earnings to the permanent component of value added) and β (sensitivity to the transitory component). In both cases, our estimating equation consists of an instrumental variables regression of $\Delta\omega_{ijt}$ onto $\Delta\epsilon_{jt}$.

As explained in the previous section, β is identified using $(\Delta\epsilon_{jt+1})^k$ as an instrument and α using $(\sum_{\tau=-2}^{2}\Delta\epsilon_{jt+\tau})^k$. We use the first three powers of the instruments. This gives us two overidentifying restrictions for each equation. These are tested with a standard J-statistic (generalized Sargan test). The power of the instruments in the reduced-form regressions is checked by looking at the p-value of the F-test on the excluded instruments. Finally, an exogeneity (Wu-Hausman) test for $\Delta\epsilon_{jt}$ (Davidson and MacKinnon 1993) is an implicit test for $\alpha = \beta$.

We also comment on the estimates of the variances of transitory and permanent shocks to value added and of idiosyncratic transitory and permanent shocks to wages. Finally, we construct an estimate of the ratio between the standard deviation of wage growth due to the firm-related components and the total standard deviation of wage growth, $\sqrt{E\{[(\Delta\omega_{ijt})^2]|j\}}/\sqrt{E[(\Delta\omega_{ijt})^2]}$, which tells approximately how much wage variability is due to workers sharing the firm's fortunes.[27] This turns out to be a useful way to summarize the evidence.

Table 7 shows the results. From panel A, we see that workers' wages do reflect current shocks to value added: $E(\Delta\omega_{ijt}\Delta\epsilon_{jt}) = 0.0021$ with a standard error of 0.0003. Thus from the outset we reject full insurance. The estimate of $E(\Delta\omega_{ijt}\Delta\epsilon_{jt-1})$ is much smaller and statistically insignificant. For purposes of comparison, the same panel also reports the estimated value of the relevant moments of the shocks to output and wages, which are very close to those estimated for the full sample (see tables 4 and 6).[28]

Panel B reports estimates of α and β along with a variety of diagnostic

[27] The variance of wage growth due to the firm-related components, $E\{[(\Delta\omega_{ijt})^2]|j\}$, is given by

$$E\{[(\Delta\omega_{ijt})^2]|j\} = \alpha[\alpha(1+\rho^2)\sigma_u^2 + \beta(1+\rho)\sigma_{uv}] + \beta[\alpha(1+\rho)\sigma_{uv} + 2\beta(\sigma_v^2 - \sigma_{vv_{-1}})].$$

[28] For example, $E(\Delta\omega_{ijt}\Delta\omega_{ijt})$ is 0.0155 in the matched final sample and 0.0163 in the full sample; $E(\Delta\omega_{ijt}\Delta\omega_{ijt-1})$ is -0.0070 in the matched final sample and -0.0071 in the full sample.

TABLE 7
SENSITIVITY OF EARNINGS TO VALUE-ADDED SHOCKS
A. AUTOCOVARIANCES

Workers' Autocovariances		Firms' Autocovariances		Worker-Firm Cross Covariances	
$E(\Delta\omega_{ijt}\Delta\omega_{ijt})$	.0155 (.0006)	$E(\Delta\epsilon_{jt}\Delta\epsilon_{jt})$	.0978 (.0083)	$E(\Delta\omega_{ijt}\Delta\epsilon_{jt})$	.0021 (.0003)
$E(\Delta\omega_{ijt}\Delta\omega_{ijt-1})$	−.0070 (.0004)	$E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-1})$	−.0453 (.0055)	$E(\Delta\omega_{ijt}\Delta\epsilon_{jt-1})$	−.0003 (.0002)

B. INSTRUMENTAL VARIABLES ESTIMATES

	Permanent Shock	Transitory Shock
Sensitivity	.0686 (.0250)	.0049 (.0045)
Observations	24,956	40,337
J-test (*p*-value)	.5794	.4616
F-test (*p*-value)	<.0001	<.0001
Exogeneity test [*p*-value]	24.16 [.0001]	

C. EWMD ESTIMATES

Firm-Level Parameters		Worker-Level Parameters	
$\sigma^2_{\tilde{u}}$	.0160 (.0030)	σ^2_{ξ}	.0014 (.0005)
$\sigma^2_{\tilde{v}}$	.0426 (.0045)	σ^2_{μ}	.0091 (.0009)
θ	−.0705 (.0418)	λ	−.0998 (.0857)
Ratio	.1476		

NOTE.—In panel A, $E(\Delta\omega_{ijt}\Delta\omega_{ijt-\tau})$ is an estimate of the autocovariance of wage growth shocks of order τ, $E(\Delta\epsilon_{jt}\Delta\epsilon_{jt-\tau})$ an estimate of the autocovariance of value-added growth shocks of order τ, and $E(\Delta\omega_{ijt}\Delta\epsilon_{jt-\tau})$ an estimate of the cross covariance of wage and value-added growth shocks of order τ. Panel B reports the instrumental variables estimates of the sensitivity of wages to value-added shocks (α for permanent shocks and β for transitory shocks). See the text for more details on the instruments used in each regression. The *J*-test is the test of overidentifying restrictions. The *F*-test is the test of joint insignificance of excluded instruments. The exogeneity test tests the null that in an OLS regression of $\Delta\omega_{ijt}$ on $\Delta\epsilon_{jt}$, the latter is exogenous. Panel C reports EWMD estimates of the variance components $\sigma^2_{\tilde{u}}$, $\sigma^2_{\tilde{v}}$, σ^2_{μ}, and σ^2_{ξ} (the variances of value-added permanent shocks, value-added transitory shocks, wage permanent shocks, and wage transitory shocks, respectively) and of the moving average parameter θ (for value added) and λ (for earnings). The ratio is calculated as $\sqrt{E\{[(\Delta\omega_{ijt})^2]\,|\,j\}}/\sqrt{E[(\Delta\omega_{ijt})^2]}$ and measures the amount of wage variability attributable to value-added shocks. Standard errors of the various estimates are reported in parentheses.

tests. There is a substantial difference in impact between permanent and transitory shocks: wages do respond to permanent shocks, but the hypothesis of full insurance against transitory shocks cannot be rejected. The estimated value of β is economically small (point estimate 0.005) and not statistically different from zero (a standard error of similar magnitude). The estimated value of α is 0.069, about 15 times larger, with a small standard error of 0.025.[29] The exogeneity test on $\Delta\epsilon_{jt}$ rejects

[29] Given that we regress wage shocks against firm shocks, which are common to all individuals working in the same firm, we correct standard errors assuming that errors are not independent within a firm (see Moulton 1986).

the null hypothesis that $\alpha = \beta$. The test statistic displays a p-value below 0.1 percent. Thus, while we cannot reject the hypothesis of "full transitory" insurance, "full permanent" insurance can be ruled out. The J-test of overidentifying restrictions has a high p-value in both cases, which indicates that the models are not misspecified. This result also implies that our instrumental variables estimates are not subject to measurement error bias.[30] The power of the instruments is not a concern, as is shown by the low p-value of the F-test in the reduced-form regressions (which checks that the excluded instruments are jointly insignificant).

To allow for an evaluation of the amount of insurance involved, we use equally weighted minimum distance methods (EWMD) to estimate the variances of idiosyncratic shocks to value added and (conditioning on these and the estimated insurance parameters α and β) the variances of idiosyncratic shocks to earnings.[31] Results are reported in panel C. The estimate of $\sigma^2_{\tilde{u}}$ is 0.016 (with a standard error of 0.003), and the estimate of $\sigma^2_{\tilde{v}}$ is 0.0426 (with a standard error of 0.0045). It follows that the variance of the permanent shock σ^2_u is 0.0585 and the variance of the transitory shock σ^2_v is 0.0561. These are both sizable and imply standard deviations of about 24 percent each. The estimate of the moving average parameter is -0.07, which is just barely significant.

Next, we estimate the parameters of the idiosyncratic part of earnings, that is, after filtering the variability that is due to the amount of insurance provided by the firm. The EWMD-estimated variances of idiosyncratic shocks to wages are much smaller than the firm counterpart: σ^2_μ, the variance of permanent shocks, is 0.0091 (standard error 0.0009), and σ^2_ξ is 0.0014 (standard error 0.0005). The moving average parameter is -0.10 (with a standard error of 0.09, which makes it insignificant).

In summary, our findings imply that a 10 percent permanent change in firm performance induces about a 0.7 percent permanent variation in earnings for those employed at the same firm on a continuing basis. To get a sense of the economic significance of this effect, for a worker receiving the average salary (€13,363 per year), a permanent 10 percent decrease in value added of the firm would permanently lower take-home pay by €90.

[30] We can prove that our identification strategy for the sensitivity of wages to the permanent component of firm value added is robust to the presence of (a) serial autocorrelation in $\tilde{u}_{jt}$, (b) correlation between $\tilde{u}_{jt}$ and $\tilde{v}_{jt}$, and (c) classical measurement error in value added. However, in all these cases the sensitivity of wages to the transitory component would be estimated with bias because of the presence of invalid instruments. Proofs of these claims are available on request from the authors. The test of the overidentifying restrictions indicates that these are not serious concerns; otherwise instrument validity would have been rejected.

[31] An alternative would be to use a generalized least-squares procedure (optimal minimum distance). Our choice is dictated by the evidence presented in Altonji and Segal (1996) that EWMD dominates optimal minimum distance even for moderately large sample sizes.

The variability in compensation induced by risk shifting depends on $\sqrt{\alpha^2(1+\rho)^2\sigma_u^2}$ (we ignore the reaction to transitory shocks, which is virtually zero, both economically and statistically). Since the overall standard deviation of the shocks to wage growth is 0.1245, one can infer that less than 15 percent of the total earnings variability can be explained by firm-specific risk (see the last row of panel C of table 7); the rest is related to worker-specific shocks.[32]

To get a rough gauge of the social value of the firm as an insurance provider, suppose that no other devices were available to smooth consumption in the face of wage risk and (for simplicity) assume no worker idiosyncratic uncertainty. Adapting the simple characterization suggested by Lucas (1987) to evaluate the welfare costs of recessions, suppose that workers maximize $U(c_t) = c_t^{1-r}/(1-r)$ and receive stochastic earnings Ae^{ω_t}, where r is the degree of workers' relative risk aversion, A is a nonstochastic component, and $\omega_t \sim N(0, \sigma_\omega^2)$, with $\sigma_\omega^2 = (1-\rho)^{-2}(\alpha-\beta\rho)^2\sigma_{\tilde{u}}^2 + \beta^2\sigma_{\tilde{v}}^2$ (this is the variance of the innovation in ω_t, i.e., $\text{Var}\,[\omega_t - E_{t-1}\omega_t]$). Wage risk depends on α and β, the insurance parameters; ρ, the autoregressive coefficient; and the variances $\sigma_{\tilde{u}}^2$ and $\sigma_{\tilde{v}}^2$.

Lucas assumes that the approximate consumption solution is

$$c_t = (1+\kappa)(1+g)^t e^{-\sigma_\omega^2/2} e^{\omega_t},$$

where g is the growth rate of consumption per period and κ a scaling parameter. The proportional increase in consumption required to leave the consumer indifferent between the full-insurance consumption path ($\alpha = \beta = 0$) and the partial insurance path ($\alpha > 0$ and $\beta > 0$) is $r\sigma_\omega^2/2$ (Clark, Leslie, and Symons 1994). This provides a reasonable first approximation to the cost of risk sharing. Under the assumption that $r = 3$, $\rho = 0.4769$, $\sigma_{\tilde{u}}^2 = 0.0160$, and $\sigma_{\tilde{v}}^2 = 0.0426$ (values consistent with our empirical findings), the proportional increase in consumption that would leave the consumer indifferent between full and no insurance ($\alpha = \beta = 1$) is $r(\sigma_{\tilde{u}}^2 + \sigma_{\tilde{v}}^2)/2 = 0.0879$. According to our estimates, the consumer is fully insured against transitory shocks ($\beta = 0$) and partially insured against permanent shocks ($\alpha = 0.069$). We calculate that workers would be indifferent between the insurance the firm currently provides and no insurance if their consumption were raised by about 9 percent. Starting with an average consumption of €13,363, say, and no wage insurance, a consumer would be willing to sacrifice as much as €1,169 to get the insurance that firms provide, according to our estimates, and just €6 more for full insurance. So the firm proves to be a very substantial provider of insurance for individuals.

[32] Using longitudinal matched data for France, Abowd et al. (1999) conclude that most of the unexplained variation in earnings *levels* is due to individual, not firm, effects. We complement their evidence, showing that most of the unexplained variation in earnings *growth* is also due to worker-specific effects.

VI. Sample Selection, Robustness, and Insurance Heterogeneity

The estimates presented in table 7 are obtained under a number of assumptions concerning sample selection and the measurement of workers' earnings and of shocks to firms. Here we show that they are robust to specific treatment of possible selection induced by the matching of workers with firms as well as to alternative measures of worker compensation and firm shocks. Finally, we consider the issue of insurance heterogeneity.

A. *Sample Selection*

Recall from Section II.*B* that our sample selection eliminates those who change position within the firm and those who have multiple employers over the period of observation (including those with multiple contemporaneous jobs in a given year). This was justified by our attempt to focus on workers with stable employment and tenure patterns. Moreover, we use an unbalanced panel of firms and workers. Finally, by its very nature our matched data set will tend to overrepresent workers in large firms and those with characteristics that make them more likely to be employed in large firms. These criteria could bias our estimates, though the direction of bias is not obvious a priori. For example, compared to movers, stayers might be more willing to trade wage insurance for job insurance, in which case their wage would be more sensitive to firm-specific shocks, so that excluding movers from the sample would induce an upward bias in the estimated parameters. On the other hand, movers might be less risk-averse than stayers and thus more willing to accept wage variability; in this case, their exclusion would bias downward the estimated sensitivity of wage to performance. For firms, those that enter and exit might be more constrained financially and hence less able to provide wage insurance. Finally, the extent of insurance could be systematically different by firm size.

To assess the importance of these potential sources of bias, we reestimate the model (*a*) including workers who change position within a firm because of promotion, demotion, or related events; (*b*) including those with multiple employers over the sample period (this accounts for people who quit or are laid off and move to a different employer);[33] (*c*) working with a balanced panel of firms and workers (continuously observed from 1982 to 1994; this checks whether firms that are in the sample for fewer years offer different insurance than those that are

[33] Since our objective is to study annual earnings growth within a firm (not across firms), we treat employment spells in different firms as different individual-firm pairs. We correct the standard errors of our estimates by assuming that the error of the wage equation is correlated across observations pertaining to the same individual.

observed for the entire sample period); and (*d*) using probability weights to account for the oversampling of workers in large firms.[34]

The results, reported in columns 2–5 of table 8, are very similar to those reported in column 1 (the baseline specification) for all the samples. The sensitivity to permanent shocks is slightly lower in all cases, except the weighted regression, but it is statistically indistinguishable from that for the baseline sample. In the balanced sample the point estimate is less precise because of the much lower number of observations. The sensitivity to transitory shocks is unchanged and is never statistically different from zero. We conclude that our results are robust to sample selection issues.

B. *Robustness*

One concern is that we have relied on net earnings as a measure of workers' compensation; though theory offers no guidance as to whether firms insure gross or net earnings, it could be argued that firms and workers contract on gross earnings. Column 6 of table 8 shows results when gross earnings are used to measure compensation. This change has little effect on the estimates: the sensitivity to permanent shocks is slightly larger than in the baseline, as one would expect given the smoothing effect of progressive taxation, whereas sensitivity to transitory shocks is small and statistically insignificant.

In our regressions we have related individual workers' wage shocks to shocks to their firm's value added. It might be argued that wage shocks should be related to the firm's productivity shock, and our controls may not be enough to capture variation in value added induced by changes in labor (or capital) inputs. We address this important issue

[34] We construct the weights following the suggestion in Wooldridge (2002). First, we use the random sample of INPS workers and estimate a logit model for the probability that an observation is in the matched data set as a function of observable characteristics (an age polynomial, gender, region dummies, sector dummies, and a polynomial in firm size). In principle, firm size is observed only for workers who can be matched with firms in the CB data set. In practice, the INPS also collects some limited information on the firms employing the workers in its database (in particular, number of employees, industry, and location). We use this to complement the information on firm size available for nonmatched workers and take the average firm size over the period the individual is observed in the INPS data set as our measure of firm size in the logit regressions. A few individuals (corresponding to 3.6 percent of the sample) have no observations on this variable and are therefore excluded from the procedure. Call $\hat{p}_{it}$ the estimated probability of inclusion in the matched sample (the predicted value of the logit regression, or propensity score). We obtain the estimates in col. 5 of table 8 by weighting the observations by the inverse of the estimated probability of inclusion, or $\hat{p}_{it}^{-1}$. Our procedure delivers consistent estimates under the conditions provided in theorem 4.1 of Wooldridge (2002). In our particular application the crucial requirement is that the firm shock does not explain selection into the matched sample conditioning on the variables that we use in our logit regressions (among which, crucially, we have included firm size). Note that we ignore issues of worker and firm panel attrition, so the analysis is illustrative rather than full-blown.

TABLE 8

Sensitivity of Earnings to Value-Added Shocks: Accounting for Sample Selection and Shock Definition

	Baseline (1)	Including Changing Position (2)	Including Movers (3)	Balanced Sample (4)	Using Weights (5)	Gross Earnings (6)	Value Added per Worker (7)	Controlling for Labor and Capital (8)
	A. Sensitivity to Permanent Shocks							
Sensitivity	.0686	.0568	.0532	.0506	.0875	.0775	.0777	.0711
	(.0250)	(.0241)	(.0218)	(.0379)	(.0350)	(.0270)	(.0405)	(.0256)
Observations	24,956	28,380	31,975	7,297	24,956	24,956	20,663	24,935
J-test (*p*-value)	.5794	.8441	.8162	.6767	.3789	.5679	.6853	.8750
F-test (*p*-value)	<.0001	< .0001	<.0001	<.0001	<.0001	<.0001	.0006	<.0001
	B. Sensitivity to Transitory Shocks							
Sensitivity	.0049	.0079	.0080	.0033	.0064	.0056	.0090	.0019
	(.0045)	(.0054)	(.0050)	(.0116)	(.0064)	(.0060)	(.0058)	(.0055)
Observations	40,337	45,903	52,697	10,435	40,337	40,337	35,142	40,302
J-test (*p*-value)	.4616	.4636	.4278	.5304	.3207	.4356	.1422	.5158
F-test (*p*-value)	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001	<.0001

Note.—The table reports the instrumental variables estimates of the sensitivity of wages to value-added shocks (panel A pertains to permanent shocks and panel B to transitory shocks). See the text for more details on the instruments used in each regression. The *J*-test is the test of overidentifying restrictions. The *F*-test is the test of joint insignificance of excluded instruments. Standard errors of the various estimates are reported in parentheses.

in the last two columns of table 8. Column 7 shows the estimates when we use value added *per worker* instead of total value added and filter it as usual with our controls (time, location, and sector dummies) to obtain a measure of the firm-specific shocks. We use employment at the beginning of the period to avoid endogeneity issues and so lose some observations in the process. Estimated parameters confirm the baseline findings.

However, even this measure may be inappropriate to capture productivity shocks. In a deeper sense, what one is ultimately interested in is the effect of shocks to the production function once variation in output due to changes in capital and labor is accounted for. We do this in column 8 of table 8, where we control for both the log of labor and the log of capital in the value-added regression. Our measure of capital is the producer price index–deflated firm's book value of stock of capital in structure, machinery, and equipment. The value of equipment is missing for a few firms, which explains the reduction in sample sizes. To account for the endogeneity of production factors, we instrument labor and capital with their lagged values (dated $t-3$) and proceed as in Section III.*A* to obtain a measure of the idiosyncratic shocks to firm performance. The results, shown in column 8 of table 8, confirm the baseline estimates: the response of wages to permanent shocks is 0.071 (standard error 0.026), whereas the sensitivity to transitory shocks is even smaller than in the baseline case and not statistically different from zero, implying that the results are not the reflection of mismeasured shocks to the production function. This suggests that idiosyncratic variations in output are largely due to idiosyncratic shocks to productivity rather than changes in the amount of utilized primary inputs.

C. *Insurance Heterogeneity*

Once the full-insurance hypothesis is rejected, one can ask whether the level of insurance varies systematically with worker and firm characteristics. The implicit contract model implies that the parameter linking wages to performance decreases with workers' risk aversion and increases with the firm's. The model proposed by Gamber (1988), which allows for bankruptcy, indicates that wage insurance decreases with the probability of default because a closer correlation of wages to performance reduces the probability of being at the corner of a binding bankruptcy constraint. The principal-agent model (an extension of the implicit contract model that allows for unobservable effort and accordingly turns on incentive as well as insurance problems, as in Holmstrom and Milgrom [1987]) posits that the compensation scheme will be more dependent on performance the more sensitive firm performance is to worker effort. We should therefore expect that the firm offers less insurance to those employees whose effort is more relevant to performance, that is, managers. In addition, principal-agent models predict that the link between wages and performance will be stronger, the more precise the signal the principal obtains on the agent's effort. When the underlying performance of the firm is noisy and the signal less precise, we should expect more insurance.

We thus extend the framework described by equations (16) and (17) to allow the sensitivity of wages to firm performance to vary with workers' risk aversion, occupation dummies, bankruptcy risk, and historical performance variability (as an inverse indicator of the precision of the signal).

To classify individuals by risk aversion, we use outside information from the Bank of Italy's 1995 SHIW. The household head is offered a hypothetical security and asked to report the highest price he would be willing to pay to purchase it, from which a measure of the Arrow-Pratt index of risk aversion can be obtained.[35] We use a matching procedure based on characteristics observable in both samples and impute a measure of relative risk aversion to all the workers in the INPS data set. To reduce misclassification error due to the imputation procedure, we construct an indicator for high risk aversion (a dummy equal to one for workers with an imputed degree of risk aversion above the cross-sectional median).[36] We proxy bankruptcy risk with the frequency of defaults at the provincial level[37] and take historical performance variability as an inverse indicator of signal precision. Variability is measured by the standard deviation of log real value added over the period observed.

Our instrumental variables estimation strategy is modified, allowing the sensitivity coefficients α and β to depend on observable worker and firm characteristics. This amounts to including interactions of these variables with value-added growth shocks, $\Delta\epsilon_{j,t}$, and augmenting the set of instruments by adding the interactions of the original ones with the relevant worker and firm characteristics.

The results are shown in table 9. The indicator for high risk aversion is associated with a statistically significant lower sensitivity of wages to permanent shocks to performance (i.e., more insurance and a lower

[35] See Guiso and Paiella (2001) for details on the wording of the question and the construction of the measure.

[36] Direct use of the imputed risk aversion variable in levels gives qualitatively similar results, although less precisely measured.

[37] We have two measures of the frequency of default for each province (the Italian territory is divided into 95 provinces) and for each year: the number of firms that defaulted on a bank loan divided by the total number of borrowers at the beginning of the period and the value of defaulted bank loans divided by the initial value of total outstanding loans. We summarize the information contained in these variables by using factor analysis and extracting the first principal component for each province-year pair (see Greene 1997, 424–27). Qualitatively similar results are obtained using either variable separately or a simple average of two.

TABLE 9
SENSITIVITY OF EARNINGS TO VALUE-ADDED SHOCK: ACCOUNTING FOR INSURANCE HETEROGENEITY

	Sensitivity to Permanent Shocks (1)	Sensitivity to Transitory Shocks (2)
$\Delta\epsilon_{j,t}$	.1096	.0151
	(.0324)	(.0144)
	[.0213]	[.1947]
$\Delta\epsilon_{j,t}$ × high risk aversion	−.0832	−.0120
	(.0366)	(.0154)
	[.0157]	[.2468]
$\Delta\epsilon_{j,t}$ × manager	.0778	.0132
	(.1197)	(.0166)
	[.0237]	[.2572]
$\Delta\epsilon_{j,t}$ × s.d.[ln (VA_{jt})]	−.0268	−.0040
	(.0129)	(.0038)
	[.0604]	[.3575]
$\Delta\epsilon_{jt}$ × bankruptcy index	.0327	−.0027
	(.0388)	(.0100)
	[.0118]	[.2474]
Observations	24,956	40,337
J-test (*p*-value)	.3257	.2863

NOTE.—Standard errors corrected for province clustering are reported in parentheses; the partial R^2 for the reduced-form regression is reported in brackets (see Shea 1997). See the text for more details on the instruments used in each regression. The *J*-test is the test of overidentifying restrictions.

value of α). In the same direction, managers receive less insurance than other employees, but standard errors are high and prevent reliable inference, arguably because of the small number of observations (scarcely 1 percent of the sample).

In terms of firms, consistent with the basic agency model, those with higher variability in performance do provide more insurance: the coefficient is negative and statistically significant. We interpret this as evidence that incentive schemes are less effective, the noisier the relation between effort and performance, supporting one of the fundamental implications of agency theory. Finally, we find that a higher probability of default is associated with a smaller amount of permanent insurance, in line with the theoretical prediction, but the standard error is high.

To get a sense of the results in table 9, consider a highly risk-averse production worker employed in a firm with a historical performance variability of 17 percent (the median variability of value added in the matched sample) located in a province with median failure rates. For this worker, the sensitivity to firm permanent shocks is 0.027. When these characteristics are retained—but risk aversion is allowed to vary—the coefficient rises as high as 0.11. For employees of a firm with the same characteristics but a larger standard deviation—say, 50 percent—the coefficient declines to 0.018. For employees of a firm with similar characteristics but in a province at the seventy-fifth percentile of failure rates, the coefficient rises to 0.038. As extended versions of the implicit contract model predict, therefore, changes in worker and firm characteristics may impart a wide range of variability in α.

Note finally that the *p*-value of the *J*-test does not indicate misspecification of the model (33 percent) and that in all cases the power of the instruments (as measured by the partial R^2 of the reduced-form regressions) is great enough to identify the relevant parameters and dismiss the possibility of finite sample bias.

We repeat this exercise for the sensitivity of earnings to transitory shocks. In accordance with the results reported in table 7, neither worker nor firm characteristics matter. This implies that insurance of transitory shocks to value added is pervasive, even after conditioning on workers' and firms' characteristics.

VII. Conclusions

We offer empirical evidence on the extent of wage insurance within the firm, based on a matched employer-employee data set for Italy spanning from the early 1980s to the mid-1990s. We find that full insurance against temporary idiosyncratic shocks is provided, whereas enduring disturbances to output are only partially—though substantially—insured. Our estimates also show that the sensitivity of workers' wages to permanent shocks is negatively correlated with the workers' risk aversion and the overall variability of firms' performance, as predicted by standard models of wage insurance. The effects of other characteristics, such as the probability of bankruptcy and occupation dummies, have the expected sign but are not well measured. Our results appear robust to a number of sensitivity checks we conduct on sample selection and the definition of variables. We believe that this is strong evidence for wage insurance mechanisms within the firm, with different types of firms providing insurance packages tailored to different types of workers.

Although the full-insurance paradigm is rejected, the extent of insurance coverage is substantial. For the average worker the standard deviation of wage growth shocks is about 12 percent, and risk sharing contributes only 1.8 percentage points to it. If temporary shocks were transferred to workers in the same proportion as permanent shocks, the variability of earnings would increase by about 15 percent. If no wage insurance were provided—implying that wages vary in the same proportion as the firm's value added—the standard deviation of wages would be as high as 40 percent. These figures, along with a simple calculation of the social value of wage insurance, showing that workers would be willing to pay up to 9 percent of their income to get the current amount of insurance, imply that the firm is a very effective insurance provider.

For future work, it would be useful to extend our analysis to countries with different institutional settings and different degrees of financial market development, which might influence the role of the firm as insurance provider. Moreover, the large amount of wage insurance we estimate begs the question of how firms provide workers with appropriate incentives. It would be of interest to assess the relative importance of different devices for eliciting effort, such as promotions and threat of termination. Finally, our estimates may be only a partial characterization of the insurance role of the firm. Besides offering insurance against idiosyncratic shocks to its performance, the firm may facilitate the emergence of insurance mechanisms against worker-specific shocks. These shocks, unlikely to be formally insured in the market for the usual moral hazard reasons, may be at least partially diversified within the firm, since workers' repeated interaction at the workplace favors risk pooling.

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (March): 251–333.

Abowd, John M., and Thomas Lemieux. 1993. "The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada." *Q.J.E.* 108 (November): 983–1014.

Altonji, Joseph G., and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *J. Bus. and Econ. Statis.* 14 (July): 353–66.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Rev. Econ. Studies* 58 (April): 277–97.

Azariadis, Costas. 1975. "Implicit Contracts and Underemployment Equilibria." *J.P.E.* 83 (December): 1183–1202.

Baily, Martin N. 1974. "Wages and Employment under Uncertain Demand." *Rev. Econ. Studies* 41 (January): 37–50.

Bellardi, Lauralba, and Lorenzo Bordogna. 1997. *Relazioni industriali e contrattazione aziendale: Continuità e riforma nell'esperienza italiana recente.* Milan: Angeli.

Blanchflower, David G., Andrew J. Oswald, and Peter Sanfey. 1996. "Wages, Profits, and Rent-Sharing." *Q.J.E.* 111 (February): 227–51.

Bronars, Stephen G., and Melissa Famulari. 2001. "Shareholder Wealth and Wages: Evidence for White-Collar Workers." *J.P.E.* 109 (April): 328–54.

Chamberlain, Gary. 1984. "Panel Data." In *Handbook of Econometrics*, vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North-Holland.

Christofides, Louis N., and Andrew J. Oswald. 1992. "Real Wage Determination and Rent-Sharing in Collective Bargaining Agreements." *Q.J.E.* 107 (August): 985–1002.

Clark, Kenneth, Derek Leslie, and Elizabeth Symons. 1994. "The Costs of Recessions." *Econ. J.* 104 (January): 20–36.

Cragg, John G. 1997. "Using Higher Moments to Estimate the Simple Errors-in-Variables Model." *Rand J. Econ.* 28 (special issue): S71–S91.

Currie, Janet, and Sheena McConnell. 1992. "Firm-Specific Determinants of the Real Wage." *Rev. Econ. and Statis.* 74 (May): 297–304.

Davidson, Russell, and James G. MacKinnon. 1993. *Estimation and Inference in Econometrics.* New York: Oxford Univ. Press.

Erickson, Christopher L., and Andrea C. Ichino. 1995. "Wage Differentials in Italy: Market Forces, Institutions, and Inflation." In *Differences and Changes in Wage Structures*, edited by Richard B. Freeman and Lawrence F. Katz. Chicago: Univ. Chicago Press (for NBER).

Estevão, Marcello, and Stacey Tevlin. 2003. "Do Firms Share Their Success with Workers? The Response of Wages to Product Market Conditions." *Economica* 70 (November): 597–617.

Gamber, Edward N. 1988. "Long-Term Risk-Sharing Wage Contracts in an Economy Subject to Permanent and Temporary Shocks." *J. Labor Econ.* 6 (January): 83–99.

Greene, William H. 1997. *Econometric Analysis.* 3rd ed. Englewood Cliffs, NJ: Prentice Hall.

Grossman, Sanford J., and Oliver D. Hart. 1981. "Implicit Contracts, Moral Hazard, and Unemployment." *A.E.R. Papers and Proc.* 71 (May): 301–7.

Guiso, Luigi, and Monica Paiella. 2001. "Risk Aversion, Wealth and Background Risk." Discussion Paper no. 2728 (March), Centre Econ. Policy Res., London.

Guiso, Luigi, and Fabiano Schivardi. 1999. "Information Spillovers and Factor Adjustment." Discussion Paper no. 2289 (November), Centre Econ. Policy Res., London.

Hamermesh, Daniel S. 1999. "LEEping into the Future of Labor Economics: The Research Potential of Linking Employer and Employee Data." *Labour Econ.* 6 (March): 25–41.

Hildreth, Andrew K. G., and Andrew J. Oswald. 1997. "Rent-Sharing and Wages: Evidence from Company and Establishment Panels." *J. Labor Econ.* 15 (April): 318–37.

Holmstrom, Bengt, and Paul Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55 (March): 303–28.

Iversen, Torben. 1998. "Wage Bargaining, Central Bank Independence, and the Real Effects of Money." *Internat. Organization* 53 (Summer): 469–504.

Knight, Frank H. 1921. *Risk, Uncertainty and Profit.* New York: Houghton Mifflin.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles.* Oxford: Blackwell.

Malcomson, James M. 1983. "Trade Unions and Economic Efficiency." *Econ. J.* 93 (suppl.; March): 51–65.

Meghir, Costas, and Luigi Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity." *Econometrica* 72 (January): 1–32.

Moskowitz, Tobias J., and Annette Vissing-Jorgensen. 2002. "The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle?" *A.E.R.* 92 (September): 745–78.

Moulton, Brent R. 1986. "Random Group Effects and the Precision of Regression Estimates." *J. Econometrics* 32 (August): 385–97.

Nickell, Stephen, and Sushil Wadhwani. 1991. "Employment Determination in British Industry: Investigations Using Micro-Data." *Rev. Econ. Studies* 58 (October): 955–69.

Pistaferri, Luigi. 2001. "Superior Information, Income Shocks, and the Permanent Income Hypothesis." *Rev. Econ. and Statis.* 83 (August): 465–76.

Riddell, W. Craig. 1981. "Bargaining under Uncertainty." *A.E.R.* 71 (September): 579–90.

Rossi, Fulvio, and Paolo Sestito. 2000. "Contrattazione aziendale, struttura negoziale e determinazione decentrata del salario." *Rivista di Politica Economica* 90 (October/November): 129–83.

Shea, John. 1997. “Instrument Relevance in Multivariate Linear Models: A Simple Measure.” *Rev. Econ. and Statis.* 79 (May): 348–52.

Wooldridge, Jeffrey M. 2002. “Inverse Probability Weighted M-Estimators for Sample Selection, Attrition, and Stratification.” *Portuguese Econ. J.* 1 (August): 117–39.