The Effect of Moral Hazard on Wage Inequality with On-the-Job Search and Employer Competition

Árpád Ábrahám[*] Fernando Álvarez-Parra[†] Susanne Forstner[‡]

PRELIMINARY DRAFT

First version: September 2012

This version: March 2013

Abstract

In this paper we study the impact of moral hazard in labor contracts on the cross-sectional wage distribution, in particular, its effect on the extent of residual wage inequality. The tool of our analysis is a search model with job-to-job mobility and firm competition for workers. In our framework, firms offer long-term contracts to risk-averse workers in the presence of repeated moral hazard. For a quantitative analysis, we calibrate the model to match characteristics of the U.S. labor market derived from micro data from the mid-2000s. We find that, on balance, moral hazard increases residual wage inequality by around six percent. The direct effect of providing incentives through wage variation accounts for a moderate contribution to inequality increase. In addition, moral hazard affects the wage distribution through several indirect effects, as firms adjust the levels of effort implemented and the wage offers made to workers in response to increased effort costs. Through their particularly strong impact on the lower parts of the wage distribution, such effects contribute substantially to the overall rise in inequality. The main reason is that, under moral hazard, low wage workers spend significantly less effort.

Keywords: Wage Dispersion, Job-to-Job Transitions, Repeated Moral Hazard, Incentive Contracts

JEL Codes: J31, J62, D31

[*]European University Institute, Florence. Email: arpad.abraham@eui.eu

[†]CAF Development Bank of Latin America

[‡]European University Institute and IIES, Stockholm University. Email: susanne.forstner@iies.su.se

1 Introduction

To understand the fundamental determinants of earnings inequality is a classic problem in economics. It is well known that a large part of earnings differences cannot be explained by individuals' observable characteristics. Recent contributions to the literature have evaluated carefully how much of residual wage inequality can be attributed to search frictions in the labor market.[1] In this paper, we add to the analysis another reason for why people with the same characteristics receive different wages, namely, incentive pay. If a worker's effort on the job is private information, wage payments need to vary across different levels of match output in order to provide incentives. Combining on-the-job search and employer competition, on the one hand, with a moral hazard problem in the worker-firm relation, on the other, is of interest for two reasons. First, including both mechanisms in a search model allows for a rich structure of wage inequality within and between firms. Second, and more importantly, it can be shown that the two mechanisms are closely linked. Consequently, it is not obvious that the introduction of moral hazard does indeed increase wage dispersion compared to standard models of on-the-job search and employer competition.[2] In addition, the interaction between the two mechanisms may affect the estimation of such models as well as the policy conclusions drawn from analyses based on such frameworks.[3]

As regards the empirical background, the significance of moral hazard for employment relations is indicated by evidence on performance-dependent pay. For example, MacLeod and Malcomson (1998) report for the year 1990 that 24% of young workers in the United States and 34% of workers in the United Kingdom received some form of performance-related compensation. Likewise, job-to-job mobility is documented to be a quantitatively important component of worker turnover. For instance, Fallick and Fleischman (2004) estimate that, for the period from 1994 to 2003, on average each month 2.6% of employed U.S. workers moved to a new employer, and that these job-to-job transitions made up around two-fifths of total monthly separations. Finally, with regard to firm competition, it seems plausible to assume that in case a worker announces his move to another job, his current employer will try to retain him by countering the outside offer.

In terms of modelling, the present paper contributes to the literature by incorporating dynamic moral hazard into an equilibrium analysis of job-to-job mobility with employer competition. In our model, risk-neutral firms offer long-term contracts to risk-averse, ex-ante identical workers in a labor market with search frictions. Output of a worker-firm match is a function of firm-specific productivity and a stochastic productivity component

[1] See Hornstein et al. (2011) for a comprehensive survey of this literature.
[2] See for example Postel-Vinay and Robin (2002).
[3] In particular, the analysis of minimum wage policies as in Lise et al. (2012) may be affected significantly.

that depends on the worker's unobservable effort. The informational friction, internal to the firm, introduces into the worker-firm relationship the well-known trade-off between providing insurance and instilling incentives. The optimal balance between these two objectives implies an efficient level of wage variation across output realizations. In addition, our model features one-sided limited commitment, as firms are assumed to commit to the wage contract, whereas workers may quit to unemployment or leave for another job. With respect to on-the-job search and firm competition for workers, we assume that workers randomly receive alternative job offers, and that it is private information whether or not a worker has received an offer. As soon as an offer is disclosed, the current and the potential future employer enter a Bertrand competition for the worker where offers are made in terms of wage contracts.

In our framework, wage differences between ex-ante identical workers originate from search frictions and the informational friction specific to the moral hazard problem. On the one hand, search frictions lead to differences between workers with different histories of unemployment spells and job offers. For instance, wages differ between employed workers who have received outside job offers and those who have not. Furthermore, variation in productivity levels between firms leads to wage differences between offers. On the other hand, labor contracts prescribe, as an optimal response to the informational friction of unobservable effort, wages that are dispersed across output realizations. This mechanism of providing incentives reduces risk sharing, making a worker's history of productivity shocks relevant for his wage sequence. Accordingly, a direct effect of moral hazard is that of inducing wage differences among workers with different histories of productivity realizations.

The latter argument suggests that the presence of moral hazard in labor contracts unambiguously increases residual wage inequality. This would indeed be the case if the only effect of moral hazard were a mean preserving spread of wages around the levels determined by a model with observable effort. However, in the presence of on-the-job search and employer competition, moral hazard affects the wage distribution also indirectly through a number of channels. Such effects can be traced back to a common source, namely, the fact that incentive provision through utility variation makes worker effort more costly to firms. In other words, providing the same lifetime utility and implementing the same effort level as under effort-observability reduces the profits of a firm. The additional effort costs affect the levels of wage offers to unemployed workers, and to employed workers in the course of firm competition. Moreover, they impact on the dynamics of wages within a job.

The main goal of the present paper is to disentangle and characterize the direct and indirect effects of moral hazard on the wage distribution, and to assess them quantitatively. For the latter purpose, we calibrate our model to key moments of cross-sectional wage dispersion and of individual wage dynamics in the U.S. labor market. In the following paragraphs, we describe the indirect effects of moral hazard and outline mechanisms underlying their impact.

One of the channels through which increased effort costs affect wage dispersion is linked to the maximum levels of lifetime utility firms are willing to offer. The reduction in profits under moral hazard implies that these critical utility levels decrease. At the same time, the lower bound on lifetime utility, which is given by the value of unemployment, remains unchanged. In consequence, the distribution of lifetime utilities in the economy is compressed downwards from the top, leading to a compression of the wage distribution too. One should expect this effect to be quantitatively moderate for the following reason: Using a simplified model, it can be shown that, if effort levels do not change between the two scenarios, average costs to the firm and, hence, average wages have to be equal at the firm's respective break-even points. Therefore, in such a model, the average wage associated with critical utilities remains the same.

The introduction of moral hazard has comparatively much stronger indirect effects on the lower parts than on the rest of the wage distribution. In this connection, a key observation is that in both scenarios firms have to deliver the same level of lifetime utility to workers at the lower bound of the wage distribution, namely, the value of the outside option of unemployment. Under moral hazard, this utility level has to be provided through an output-dependent wage path. Ceteris paribus, average wages of recently unemployed workers therefore need to be higher in order to compensate risk-averse workers for earnings variation. Moreover, due to the mechanism described above, workers expect lower levels of lifetime utility from outside offers. Consequently, their continuation utility decreases, leading to an additional increase in average wages with the first employer after an unemployment spell. The resulting compression effect of moral hazard on the wage distribution emerges clearly in the aforementioned simplified model when effort levels are held constant.

Firms react, however, to the increase in effort costs by lowering the levels of effort prescribed to their workers. If these changes were near proportional at all levels of lifetime utility, they would not impact significantly on wage dispersion because the cost of effort compensation would change uniformly. But the change in effort turns out to be much larger at low levels of lifetime utility, due to the fact that incentives need to be provided through a spread in future lifetime utilities. At low levels of utility (and hence wage), this spread is largely constrained by the requirement that, for any history, workers need to be at least weakly better off working than being unemployed. The large decrease in effort levels in this range of utility implies a parallel decrease in effort compensation to the worker and, ceteris paribus, a significant decrease in the lowest percentile wages. As a result, the indirect effect mediated through changes in effort levels counteracts the above described upward compression of the wage distribution.

Since, through the various channels, moral hazard affects the wage distribution in ways that imply inequality changes in both directions, its overall impact on wage dispersion can be determined only by a quantitative analysis. Using our calibrated model for this assessment, we find that, on balance, the presence of moral hazard in labor contracts increases residual wage inequality by around six percent. Computational experiments suggest that the direct effect of incentive provision leads to a moderate inequality increase due to wage dispersion within groups of workers with the same job offer history. Among the indirect effects of moral hazard, the one most closely related to outside wage offers and firm competition counteracts the direct effect, exerting a modest influence on inequality. In stark contrast to this, non-observability of effort, through other indirect effects, has a particularly strong impact on the lower parts of the wage distribution. The main reason is that, in response to higher effort costs, firms demand significantly lower levels of effort from low wage workers. As a result, within the lower half of the wage distribution, inequality increases more than proportionally, thus contributing substantially to the overall impact of moral hazard.

The present paper closely relates to at least two strands of the literature. First, it contributes to the study of optimal dynamic contracts in the context of labor markets. Thomas and Worrall (1988) provide an early application of dynamic contracts to employment relations under limited commitment, a feature that is present in our model too. Furthermore, our analysis of optimal wage contracts in a setting with repeated moral hazard builds on the seminal work of Rogerson (1985a) and Spear and Srivastava (1987). To the best of our knowledge, there are only a few other papers that study a dynamic moral hazard problem in a labor market context with on-the-job search. Manoli and Sannikov (2005) analyze properties of the optimal contract in a continuous-time environment where bidding-strategies of firms are themselves a device for incentive provision. The authors focus on the analytical characterization of bidding strategies and job changes as well as on ex-post inefficiencies. However, they neither investigate the role of moral hazard in shaping the cross-sectional wage distribution, nor do they provide a quantitative framework for analysis. Tsuyuhara (2011a) and Tsuyuhara (2011b) develop a model of directed on-the-job search with dynamic moral hazard. Within that framework, and similar to our model, both work incentives and job-to-job mobility induce dispersion in labor productivities among ex-ante identical workers. The main focus of the papers is on the longitudinal characteristics of the wage contract and on the response of labor market variables to aggregate shocks. The contribution of moral hazard to observed wage inequality is, however, neither qualitatively nor quantitatively evaluated. Moreover, as opposed to our model, employers are assumed not to react to outside job offers received by their workers.

Second, our paper contributes to the analysis of sources of wage inequality. In a recent article, Hornstein et al. (2011) assess the degrees to which different versions of job search models can account for empirically observed levels of wage dispersion. They find that, in this context, models of on-the-job search with employer competition as in Postel-Vinay and Robin (2002) are among the most promising approaches. In the present paper, we extend this framework by introducing an informational friction leading to moral hazard as an additional source of wage inequality. We find that the overall increase in wage inequality attributable to moral hazard amounts to around six percent. Moreover, inequality increase is comparatively more pronounced within the lower parts of the wage distribution.

The paper is organized as follows: Section 2 presents the model used in our analysis, states the optimal contract problem, and provides two theoretical results. In Section 3 we consider a simplified version of the model and illustrate analytically key parts of the channels through which the presence of moral hazard impacts on the wage distribution. Section 4 returns to the quantitative analysis of the full model. The first part describes in some detail the calibration of our model, while the second part presents the quantitative results together with a discussion of the underlying mechanisms. Section 5 concludes the paper. Some analytical results and a more extensive data description are presented in the appendices.

2 The Model

This section presents the model on which our study is based and states the central contracting problem.[4] It starts with an exposition of the basic environment in which the analysis is carried out. Within this framework, we formulate the optimal contract problem and define as well as characterize a number of key variables pertinent to the problem. Finally, and with a view to the quantitative analysis presented in Section 4, we state two theoretical results and provide a concise definition of equilibrium in the present context.

[4]The framework of analysis builds on Alvarez-Parra (2008).

2.1 The setup

In the present framework, time $t = 0, 1, 2, \ldots$ is discrete with an infinite horizon. The economy is populated by a large number of workers I_w and of firms I_f. Workers are risk-averse and ex-ante identical, whereas firms are risk-neutral and differ with respect to their productivity. Each firm can employ only one worker. Both I_w and I_f are invariant over time as a consequence of the following replacement rule: When a worker dies, he is replaced by a newborn worker who is unemployed, and when a firm goes out of business, it is replaced by a newly established firm. Whenever a new firm is established, it can hire a worker out of unemployment or of employment with another firm. If it fails to do so, it disappears. Finally, workers and firms discount the future at a common discount factor $\beta \in (0, 1)$.

2.1.1 Firms

Each firm is characterized by a productivity level $z \in \mathcal{Z} = \{z_1, z_2, \ldots, z_N\}$ where $z_n < z_{n+1}$. The value is drawn from a distribution with cumulative distribution function $F(z)$ when the firm is established and remains constant over time. Firm output y is a function of both the non-stochastic firm-specific productivity level z and a stochastic worker-specific productivity factor A which depends on the effort ϵ spent by the employed worker over a given period.

Firms use a production technology according to which output y is given by

$$y = y(z, A) = zA \tag{1}$$

The worker-productivity factor A can take on two different values, depending stochastically on the worker's effort level ϵ:

$$A = \begin{cases} A^+ & \text{with probability} \quad \pi(\epsilon) \\ A^- & \text{with probability} \quad 1 - \pi(\epsilon) \end{cases} \tag{2}$$

where $A^+ > A^-$ and the function $\pi(\cdot)$ is continuous, strictly increasing and strictly concave ($\pi'(\cdot) > 0, \pi''(\cdot) < 0$) with $\pi(\epsilon) \in [0, 1]$ for $\epsilon \in [0, \bar{\epsilon}]$. Firms maximize their expected present-value profits.

2.1.2 Workers

Each worker can be either unemployed or employed by one of the firms active in a given period. For each period, the probability for a worker to die within this period is $(1 - \psi)$. A worker derives utility from consumption and suffers disutility from spending effort while working on his job. The period utility $u(c)$ from consumption c is assumed to be a continuous, strictly increasing and strictly concave function ($u'(c) > 0, u''(c) < 0$) which is bounded from above by zero. By contrast, the period disutility $g(\epsilon)$ from effort is assumed to be a continuous, strictly increasing and convex function ($g'(\epsilon) > 0, g''(\epsilon) \geq 0$) with $g(\epsilon) \geq 0$ for $\epsilon \in [0, \bar{\epsilon}]$. Workers maximize their expected lifetime utility.

While unemployed, a worker enjoys a level of consumption $b > 0$ which is the same for all workers and also invariant over time. Assuming that workers do not save, the unemployed worker's period utility amounts to $u(b)$. In unemployment, the probability of receiving exactly one job offer within a given period is λ_u. The productivity type z of the firm making such an offer is a random draw from distribution $F(z)$. While employed, a worker spends effort ϵ and receives wage w within a time period, implying a period utility of $u(w) - g(\epsilon)$. For the employed worker, the probability of receiving exactly one outside job offer within a given period is λ_e, where, in this case too, the productivity type $\tilde{z}$ of the firm making the offer is drawn from distribution $F(\cdot)$.

2.1.3 Interaction between workers and firms

When a worker and a firm meet, the latter offers the former a long-term contract which is valid as long as the worker does not bring in an outside job offer. The contractual relationship is asymmetric in that the firm commits to the contract, whereas the worker may walk away at any time. Independent of commitment, there is a positive probability δ that within a given period an employment relationship gets destroyed exogenously. The two types of endogenous dissolution of a worker-firm relationship are quitting into unemployment or job-to-job transition. In the event of losing its worker, the firm concerned disappears.

The focus of the present analysis is on one particular feature of the worker-firm interaction, embodied in the assumption that the firm cannot observe the level of effort spent by its worker. Likewise, we assume that an outside offer received on the job cannot be observed by the firm. Finally, we make three additional assumptions about the interaction between firms and workers. First, the long-term contract is a take-it-or-leave-it offer by a firm to a worker. In other words, workers have no bargaining power. Second, if an employed worker reports an outside offer to his employer, the current labor contract is annulled. Subsequently, the current and the potential future employers start competing for the worker by offering new contracts. This competition takes place in the form of a second-price auction between the two firms in terms of expected lifetime utility the respective contracts offer to the worker. And third, a labor contract signed in period t specifies, for all future dates τ and all possible histories of shock realizations until date τ (denoted by $\mathbf{A}^\tau \equiv \{A_j\}_{j=t+1}^{\tau}$), conditional on the worker staying in the contract, a period wage $w_\tau(\mathbf{A}^\tau)$ and a period effort level $\epsilon_\tau(\mathbf{A}^\tau)$.

2.1.4 Timing of events

The timing of events within a model period is as follows: A worker employed by a z-type firm spends effort ϵ and receives period wage w. Then output $y(z, A)$ is produced, revealing the period realization of A. With probability λ_e, the worker is contacted by a firm of type $\tilde{z}$ drawn from distribution $F(\cdot)$. If he reports the outside job offer to his current employer, the two firms start a Bertrand competition by offering new labor contracts which would start from the beginning of the next period. Within the same time, an unemployed worker consumes b and is, with probability λ_u, contacted by a firm for which he may start working at the beginning of the next period. At the end of the current period, existing worker-firm matches are exogenously destroyed with probability δ. A worker dies with probability $(1 - \psi)$ and is replaced by a newborn worker who is unemployed. A firm that has lost its worker disappears and is replaced by a newly established firm.

2.2 The optimal contract

Firms offer to workers long-term contracts that are designed optimally with respect to the non-observability of worker effort and the fact that workers may quit or switch employers. On the one hand, firms offer wages which depend on the history of a worker's output realizations with the firm in order to provide optimal incentives for effort. On the other hand, contracts are designed such that a worker never wants to quit to unemployment and moves only to a competitor firm whose productivity is at least as high as that of the current employer's firm. However, in our framework a firm's future bidding strategy for the case that its worker receives an outside job offer is in general not part of the labor contract.[5] When designing a contract, a firm therefore takes as given both its own and a potential competitor's ex-post optimal bidding strategies together with the optimal strategies of workers for reporting outside offers.

More formally, a labor contract consists of sequences of functions $\{w_\tau(\mathbf{A}^\tau), \epsilon_\tau(\mathbf{A}^\tau)\}_{\tau=1}^{\infty}$ specifying a period wage and effort level for all future dates τ and all possible histories of productivity realizations $\mathbf{A}^\tau$. Thus, the actions prescribed to the worker and the payoffs to both worker and firm at any point in time depend on the whole history of previous productivity realizations, actions, and payoffs. We use the *promised utility approach*,

[5] See Manoli and Sannikov (2005) for a continuous-time framework where firms can commit ex-ante to bidding strategies. In our discrete-time setup, this would unduly complicate the model. Moreover, we believe that ex-ante commitment to bidding strategies is much less plausible than ex-ante commitment to a wage path.

developed, among others, by Spear and Srivastava (1987), to formulate the firm's problem of designing an optimal contract under repeated moral hazard recursively. In this approach, all relevant aspects of history are summarized in a single variable, namely, the expected lifetime utility promised to the worker by the contract.

2.2.1 Preliminaries

In order to be able to describe, in the present setup, an optimal long-term contract between a worker and a firm, we need to expand our notation and introduce a number of additional variables. First, in the analysis different kinds of expected utility for the worker have to be considered. One of them is the expected lifetime utility of a currently unemployed worker, denoted by U^n. Another one, U, is the expected lifetime utility promised to an employed worker under his current labor contract. This utility level, together with the productivity type z of the firm employing the worker, fully characterizes the worker's state. Other expected-utility variables of interest are the continuation values promised to an employed worker, either by his current or by a future employer. Thus, $U^i(U, z)$ denotes the continuation value of the current labor contract at a z-type firm under the conditions that the worker's current realization of specific productivity is A^i, with $i \in \{+, -\}$, and that he did not receive an outside job offer. For notational convenience, we will sometimes use U^i instead of $U^i(U, z)$. By contrast, $U_o(U^i, z, \tilde{z})$ represents the continuation value to the worker when his productivity realization is A^i and he receives an outside offer from a firm of type $\tilde{z}$. The value of $U_o(U^i, z, \tilde{z})$ is equal to the continuation value of the current labor contract if the worker decides not to report the outside offer. If he does, however, report the offer, $U_o(U^i, z, \tilde{z})$ takes on the value of the expected lifetime utility offered by a new contract which is the outcome of Bertrand competition between the two firms.

Switching from the viewpoint of the worker to that of the firm, the value to a z-type firm of a contract delivering lifetime utility U in an optimal way[6] is denoted by $V(U, z)$. The function $V_o(U^i, z, \tilde{z})$ represents the continuation value to the firm when the worker's current productivity realization is A^i and he has received an outside offer from a firm of productivity type $\tilde{z}$. In analogy to utility U_o, the value of V_o is the continuation value to the firm of the current labor contract for the case that the worker does not report the outside offer he has received. If he does, however, report the offer, $V_o(U^i, z, \tilde{z})$ assumes the value to the firm of the outcome of the Bertrand competition with the outside competitor.

With the notation developed up to this point, the value $V(U, z)$ to the z-type firm of

[6] The corresponding optimization problem is defined below.

a contract delivering to the worker lifetime utility U can be expressed as

$$\begin{aligned} V(U,z) &= z\Big[A^{+}\pi(\epsilon)+A^{-}(1-\pi(\epsilon))\Big]-w \\ &+\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[V(U^{+},z)\pi(\epsilon)+V(U^{-},z)(1-\pi(\epsilon))\Big] \\ &+\lambda_e\sum_{\tilde{z}\in\mathcal{Z}}\Big[V_o(U^{+},z,\tilde{z})\pi(\epsilon)+V_o(U^{-},z,\tilde{z})(1-\pi(\epsilon))\Big]f(\tilde{z})\Bigg\} \end{aligned} \tag{3}$$

while the value of a contract to the worker can be written as

$$\begin{aligned} & u(w)-g(\epsilon)+\beta\psi\delta U^{n}+\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[U^{+}\pi(\epsilon)+U^{-}(1-\pi(\epsilon))\Big] \\ &\qquad +\lambda_e\sum_{\tilde{z}\in\mathcal{Z}}\Big[U_o(U^{+},z,\tilde{z})\pi(\epsilon)+U_o(U^{-},z,\tilde{z})(1-\pi(\epsilon))\Big]f(\tilde{z})\Bigg\} \end{aligned} \tag{4}$$

Finally, a recursive contract in the present framework can be defined as follows:

Definition 1. *A recursive contract $\mathcal{C}$ is a collection of functions that, for each pair (U,z) of promised utility U and firm productivity type z, specify a prescribed worker effort $\epsilon(U,z)$, a current period wage $w(U,z)$, and continuation values $\{U^{+}(U,z),U^{-}(U,z)\}$ for the worker who attains productivity realization $A\in\{A^{+},A^{-}\}$ and does not receive an outside job offer.*

2.2.2 Continuation values for the case of an outside offer

If and when an employed worker reports an outside offer to his firm, the current and the potential future employers enter a Bertrand competition for the worker. In this competition, the critical utility level $U^{*}(z)$, defined by the condition of zero expected profits

$$V(U^{*}(z),z)=0 \tag{5}$$

plays a decisive role: A firm of productivity type z is in general willing to offer a level of lifetime utility U up to a maximum equal to the break-even value of $U^{*}(z)$.[7] Therefore, if the lifetime utility promised by the current contract is lower than the break-even utilities of both firms, a worker triggering Bertrand competition will be offered a new contract which promises him the lower of the two critical utility levels.

It should be noted, however, that in the present setup firms may be making losses at some point of the worker-firm match, that is, $U^{*}(z)<U^{i}$. In this case, if the worker

[7] Since $V(U,z)$ is strictly decreasing in U and continuous on the relevant range of the domain, for a given z, the value $U^{*}(z)$ is unique and well-defined by equation (5).

receives an outside offer from a more productive firm which could still make positive profits at utility level U^i, it is in the interest of the current employer that the worker switches jobs. Moreover, transferring the worker to the competitor firm in this way leads to a Pareto-improvement. Once the outside offer has been disclosed, any bidding strategy of the current employer up to $U^*(\tilde{z})$ exclusive can achieve this objective. However, any value lower than U^i would keep the worker from reporting the offer and hence not lead to a quit. We therefore assume that, in such a situation, the current employer bids U^i, which is one of his optimal strategies.[8]

Taking into account these rules, seven different cases of relationships between values of U^i, $U^*(z)$, and $U^*(\tilde{z})$ can be outlined. The resulting classification is presented below. It includes specifications of the worker's decisions on whether to trigger firm competition, on the one hand, and on whether to stay with his current employer or move to the job offered by the competing firm, on the other. It also states the corresponding continuation values for the worker and the current employer.

The first three cases reflect the situation of the competitor firm being more productive than the incumbent firm, so that $U^*(z) < U^*(\tilde{z})$.

Case 1: $U^i \leq U^*(z) < U^*(\tilde{z})$
The worker discloses the offer, moves to the new employer and gets $U_o = U^*(z)$. The firm currently employing the worker disappears ($V_o = 0$).

Case 2: $U^*(z) < U^i \leq U^*(\tilde{z})$
The worker discloses the offer and the incumbent firm offers U^i. The worker moves to the new employer and gets $U_o = U^i$. The incumbent firm disappears ($V_o = 0$).

Case 3: $U^*(z) < U^*(\tilde{z}) < U^i$
The worker does not disclose the offer and stays with his current employer. The worker gets $U_o = U^i$, and the current employer gets $V_o = V(U^i, z)$.

The next two cases cover the situation where the competitor firm and the incumbent firm have the same productivity level.

Case 4: $U^i \leq U^*(z) = U^*(\tilde{z})$
The worker discloses the offer and gets $U_o = U^*(z)$. With probability 1/2 he stays with his current employer who gets $V_o = V(U^*(z), z) = 0$. With probability 1/2 the worker moves to the new firm, in which case the incumbent firm disappears ($V_o = 0$).

[8] Taking into account the worker's optimal reporting strategy, the current employer is in fact indifferent between any bidding strategies in $\left[U^i, U^*(\tilde{z})\right)$. The value chosen simply determines how the gains from the offer are split between the worker and the future employer.

Case 5: $\quad U^*(z) = U^*(\tilde{z}) < U^i$

The worker does not disclose the offer and stays with his current employer. The worker gets $U_o = U^i$, and the current employer gets $V_o = V(U^i, z)$.

Finally, the last two cases reflect the situation where the competitor firm is less productive than the incumbent firm, so that $U^*(\tilde{z}) < U^*(z)$.

Case 6: $\quad U^i \leq U^*(\tilde{z}) < U^*(z)$

The worker discloses the offer, stays with his current employer, and gets $U_o = U^*(\tilde{z})$. The current employer gets $V_o = V(U^*(\tilde{z}), z)$.

Case 7: $\quad U^*(\tilde{z}) < U^*(z) \;$ and $\; U^*(\tilde{z}) < U^i$

The worker does not disclose the offer and stays with his current employer. The worker gets $U_o = U^i$, and the current employer gets $V_o = V(U^i, z)$.

Based on these cases, the continuation value U_o to a worker employed by a z-type firm, who would get utility U^i under his current labor contract and has received an outside offer from a $\tilde{z}$-type firm, can be summarized in the following expression:

$$U_o(U^i, z, \tilde{z}) = \max\Big\{U^i(U, z), \min\Big[U^*(z), U^*(\tilde{z})\Big]\Big\} \tag{6}$$

The corresponding continuation value to the current employer is given by

$$V_o(U^i, z, \tilde{z}) = \begin{cases} V(U^i(U, z), z) & \text{if} \quad U^*(\tilde{z}) < U^i(U, z) \\ V(U^*(\tilde{z}), z) & \text{if} \quad U^i(U, z) \leq U^*(\tilde{z}) < U^*(z) \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

2.2.3 The optimization problem

The contract offered by a firm to a worker is designed optimally, given the particular features of the current setup. On the one hand, in response to the informational friction of non-observability of worker effort, a firm prescribes effort levels that the worker would optimally choose under the conditions of the contract. This requirement is introduced into the firm's optimization problem in the form of an incentive-compatibility constraint.

An incentive-compatible effort level ϵ maximizes a worker's expected lifetime utility associated with the contract, that is,

$$\begin{aligned} \epsilon \in \underset{\hat{\epsilon} \in [0, \bar{\epsilon}]}{\operatorname{argmax}} \quad & u(w) - g(\hat{\epsilon}) + \beta\psi\delta U^n \\ & + \beta\psi(1-\delta)\Big\{(1-\lambda_e)\Big[U^+\pi(\hat{\epsilon}) + U^-(1-\pi(\hat{\epsilon}))\Big] \\ & + \lambda_e \sum_{\tilde{z} \in \mathcal{Z}} \Big[U_o(U^+, z, \tilde{z})\pi(\hat{\epsilon}) + U_o(U^-, z, \tilde{z})(1-\pi(\hat{\epsilon}))\Big] f(\tilde{z})\Big\} \end{aligned} \tag{8}$$

The first-order condition corresponding to an interior solution of this maximization problem is

$$
\begin{aligned}
g'(\epsilon) &= \pi'(\epsilon)\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[U^+ - U^-\Big] \\
&\quad +\lambda_e \sum_{\tilde{z}\in\mathcal{Z}} \Big[U_o(U^+, z, \tilde{z}) - U_o(U^-, z, \tilde{z})\Big] f(\tilde{z})\Bigg\}
\end{aligned} \tag{9}
$$

We use equation (9) as the incentive-compatibility constraint (ICC).[9] This relationship, characterizing the set of incentive-compatible effort levels, expresses the standard condition that at the optimum marginal benefit and marginal cost of effort have to be equal. In addition, equation (9) indicates that potential outside offers make the provision of incentives more expensive relative to a situation in which workers cannot search on the job. In other words, the higher the probability of outside job offers (λ_e), the higher is the spread in continuation values from the current contract ($U^+ - U^-$) that is needed to extract a certain level of effort. The reason is that, when the worker receives a relevant outside offer, the firm has only little control over his continuation utility.

On the other hand, in the optimal contract, the problem of lack of commitment of workers has to be taken into account too. In order to prevent workers from quitting to unemployment, firms always promise their workers at least as much utility from the contract as they would get in unemployment.[10] Hence, the participation constraint (PC) requires that the value of a contract to the worker at any point in time weakly exceeds the value of unemployment U^n.

Furthermore, the contract has to deliver at least the expected lifetime utility promised to the worker. This is captured by the following promise-keeping constraint (PKC):

$$
\begin{aligned}
U &\leq u(w) - g(\epsilon) + \beta\psi\delta U^n \\
&\quad +\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[U^+\pi(\epsilon) + U^-(1-\pi(\epsilon))\Big] \\
&\quad +\underbrace{\lambda_e \sum_{\tilde{z}\in\mathcal{Z}} \Big[U_o(U^+, z, \tilde{z})\pi(\epsilon) + U_o(U^-, z, \tilde{z})(1-\pi(\epsilon))\Big] f(\tilde{z})}_{\text{expected continuation value in case of outside offer}}\Bigg\}
\end{aligned} \tag{10}
$$

The possibility of exogenous job destruction, on the one hand, and of outside job offers triggering firm competition, on the other, are the sources of special features of the

[9]The first-order approach is valid under the standard conditions provided by Rogerson (1985b). In particular, given our assumptions on the properties of $\pi(\epsilon)$, the worker's problem of effort choice satisfies the monotone likelihood-ratio and the convex distribution function conditions.

[10]In the present setup, firms whose productivity is too low to make positive profits at U^n will never form a match with a worker. In consequence, in equilibrium all operating firms make positive profits at U^n, and a worker's endogenous quit to unemployment is always inefficient.

promise-keeping constraint in the present setup. First, the value of unemployment U^n enters the constraint as a component which the firm cannot influence at all. Second, and as mentioned above, the expected continuation value for a worker who has received an outside offer is subject to only little control by the firm. Moreover, this component of the promise-keeping constraint varies between firms of different productivity types. The reason is that the productivity level of a worker's current employer determines the upper bound for the increase in lifetime utility the worker can obtain from firm competition. In particular, the higher the current employer's productivity level, the higher is the worker's expected continuation value in case an outside offer arrives. Consequently, a firm with high productivity faces a more relaxed promise-keeping constraint than a low-productivity firm in the sense that, keeping everything else equal, it can satisfy the constraint by paying a lower wage.

Since the value of unemployment U^n affects the promise-keeping constraint and the participation constraint, the term needs to be specified in order to be able to fully describe the contractual problem. Since workers have no bargaining power, a firm offers to an unemployed worker the value of unemployment. Therefore U^n satisfies

$$U^n = u(b) + \beta\psi\Big[(1-\lambda_u)U^n + \lambda_u U^n\Big] \tag{11}$$

and the value of unemployment can be expressed as

$$U^n = \frac{u(b)}{1-\beta\psi} \tag{12}$$

As the utility accruing to a worker from a labor contract has to satisfy the participation constraint, the above expression is also a lower bound for U.

Finally, the contract must guarantee feasibility of delivering the utility promised to the worker. Given that the period utility of an employed worker, $u(c) - g(\epsilon)$, is bounded from above by zero, this value must also be an upper bound on the expected lifetime utility U. In addition, the presence of exogenous separations imposes an even tighter upper bound $\bar{U}$ on promised utility. It can be derived from the expression of a worker's utility from a contract, given by (4), in the following way:

$$\begin{aligned}&\underbrace{u(w)-g(\epsilon)}_{\leq 0} + \beta\psi\delta U^n + \beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[\underbrace{U^+}_{\leq \bar{U}}\pi(\epsilon) + \underbrace{U^-}_{\leq \bar{U}}(1-\pi(\epsilon))\Big] \\ &\qquad + \lambda_e \sum_{\tilde{z}\in\mathcal{Z}} \Big[\underbrace{U_o(U^+, z, \tilde{z})}_{\leq \bar{U}}\pi(\epsilon) + \underbrace{U_o(U^-, z, \tilde{z})}_{\leq \bar{U}}(1-\pi(\epsilon))\Big] f(\tilde{z})\Bigg\}\end{aligned} \tag{13}$$

Given that the whole expression (13) is also bounded from above by $\bar{U}$, this leads to the equation

$$\bar{U} = \beta\psi\delta U^n + \beta\psi(1-\delta)\bar{U} \tag{14}$$

Substituting (12) then yields the expression

$$\bar{U} = \frac{\beta\psi\delta}{1-\beta\psi(1-\delta)}\frac{u(b)}{1-\beta\psi} < 0 \tag{15}$$

The corresponding feasibility constraint (FC) requires that the value of a contract to the worker at any point in time does not exceed the upper bound $\bar{U}$.

Using the components defined and the functions outlined previously, the contractual problem can be stated as the following functional equation problem the optimal contract $\mathcal{C}^*$ has to solve:

$$\begin{aligned} V(U,z) \quad = \quad & \max_{\{w,\epsilon,U^+,U^-\}} z\Big[A^+\pi(\epsilon) + A^-(1-\pi(\epsilon))\Big] - w \\ & +\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[V(U^+,z)\pi(\epsilon) + V(U^-,z)(1-\pi(\epsilon))\Big] \\ & +\lambda_e \sum_{\tilde{z}\in\mathcal{Z}}\Big[V_o(U^+,z,\tilde{z})\pi(\epsilon) + V_o(U^-,z,\tilde{z})(1-\pi(\epsilon))\Big]f(\tilde{z})\Bigg\} \end{aligned} \tag{16}$$

subject to the promise-keeping constraint (PKC)

$$\begin{aligned} U \quad \leq \quad & u(w) - g(\epsilon) + \beta\psi\delta U^n \\ & +\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[U^+\pi(\epsilon) + U^-(1-\pi(\epsilon))\Big] \\ & +\lambda_e \sum_{\tilde{z}\in\mathcal{Z}}\Big[U_o(U^+,z,\tilde{z})\pi(\epsilon) + U_o(U^-,z,\tilde{z})(1-\pi(\epsilon))\Big]f(\tilde{z})\Bigg\} \end{aligned} \tag{17}$$

the incentive-compatibility constraint (ICC)

$$\begin{aligned} g'(\epsilon) \quad = \quad & \pi'(\epsilon)\beta\psi(1-\delta)\Bigg\{(1-\lambda_e)\Big[U^+ - U^-\Big] \\ & +\lambda_e \sum_{\tilde{z}\in\mathcal{Z}}\Big[U_o(U^+,z,\tilde{z}) - U_o(U^-,z,\tilde{z})\Big]f(\tilde{z})\Bigg\} \end{aligned} \tag{18}$$

the participation and feasibility constraints (PC) and (FC)

$$\begin{aligned} U^n \leq U^+ \leq \bar{U} \\ U^n \leq U^- \leq \bar{U} \end{aligned} \tag{19}$$

and

$$w \quad \geq \quad 0 \tag{20}$$

$$\epsilon \quad \in \quad [0,\bar{\epsilon}] \tag{21}$$

2.3 On continuation values in the absence of outside offers

The following two lemmas provide results on the relationship between continuation values U^i, $i \in \{+, -\}$, for the employed worker under his current labor contract, given that he has not received any outside job offer. Lemma 1 establishes that under moral hazard a firm can only extract positive effort from its worker by promising him a strictly higher continuation utility after a high output realization than after a low one. In other words, positive worker effort requires a spread in wages between the two possible output realizations, and therefore a worker's wage depends on the history of productivity shocks.

Lemma 1. *Under the assumption that worker effort is not observable, that is, for the optimization problem (16) with constraints (17) to (21), the following result on continuation values $U^i(U, z)$ holds: If at a given state (U, z) the level of worker effort $\epsilon(U, z)$ is positive, then the worker's continuation value under high productivity realization $U^+(U, z)$ exceeds the corresponding value $U^-(U, z)$ associated with low realization, in short: $\epsilon(U, z) > 0 \Rightarrow U^+(U, z) > U^-(U, z)$.*

Proof. If effort is not observable, the incentive-compatibility constraint (18) must hold. Since $g'(\epsilon) > 0$, the right-hand side of (18) must be positive. As, for any pair $(z, \tilde{z})$, the function $U_o(U^i, z, \tilde{z})$ is non-decreasing in U^i, this requires that $U^+(U, z) > U^-(U, z)$. □

By contrast, Lemma 2 shows that under observable effort, a worker's continuation utility is independent of the output realization, that is, the risk-neutral firm fully insures the risk-averse worker against the uncertainty of effort-dependent productivity.

Lemma 2. *Under the assumption that worker effort is observable, that is, for the optimization problem (16) with constraints (17) and (19) to (21) – and without the incentive-compatibility constraint (18) – the following result on continuation values $U^i(U, z)$ holds: If the firm's value function $V(U, z)$ is concave in U, then the function $U^+(U, z)$, representing the worker's continuation value under high productivity realization, is identical with the corresponding function $U^-(U, z)$ associated with low realization, in short: $V(U, z)$ concave in $U \Rightarrow U^+(U, z) = U^-(U, z) \equiv U'(U, z)$.*

Proof. See Appendix A.1. □

2.4 Stationary equilibrium

In the quantitative analysis to follow, we focus on the properties of stationary equilibria of the present model. Let l denote the labor market status of a worker in a given period, where $l = 1$ if he is employed and $l = 0$ if he is unemployed. Further, let $\mu(l, U, z)$ denote the distribution of workers in a given period over labor market statuses, expected lifetime utilities, and firm productivities. A stationary equilibrium in the present framework can then be defined as follows:

Definition 2. *A stationary equilibrium consists of a value function $V(U, z)$, policy functions $w(U, z)$, $\epsilon(U, z)$, $U^+(U, z)$, and $U^-(U, z)$, functions $U_o(U^i, z, \tilde{z})$ and $V_o(U^i, z, \tilde{z})$ specifying continuation values to employed workers and firms for the case of an outside offer, laws of motion $M : \mu(l, U, z) \to \mu'(l, U, z)$, and a distribution $\mu_s(l, U, z)$ such that:*

(i) Given $U_o(U^i, z, \tilde{z})$ and $V_o(U^i, z, \tilde{z})$, the functions $V(U, z)$, $w(U, z)$, $\epsilon(U, z)$, $U^+(U, z)$, and $U^-(U, z)$ solve the contractual problem (16) subject to constraints (17) to (21).

(ii) The functions $U_o(U^i, z, \tilde{z})$ and $V_o(U^i, z, \tilde{z})$ are consistent with firms' optimal bidding behavior in the case of firm competition, with workers' optimal reporting of outside offers, and with workers' optimal decisions whether to stay with the current employer or move to the new job, and are given by equations (6) and (7).

(iii) The laws of motion M are generated by firms' and workers' optimal decision rules and the specification of exogenous shocks.

(iv) The distribution $\mu_s(l, U, z)$ is consistent with the laws of motion M and is stationary, that is, $\mu_s(l, U, z) = M(\mu_s(l, U, z))$.

3 Analysis of a Simplified Model

In this section, we consider a simplified version of the model in which we can illustrate analytically key parts of the channels through which moral hazard affects the wage distribution. We make the following simplifying assumptions: First, all firms have the same productivity level $z = 1$, hence period output is either $A^+ > 1$ or $A^- < 1$. Second, we assume that effort can only take one out of two values $\{0, \epsilon^*\}$. We set the disutility from effort to $g(0) = 0$ and $g(\epsilon^*) = v > 0$, and the probability of a high worker productivity realization to $\pi(0) = 0$ and $\pi(\epsilon^*) = p > 0$. Third, we assume that firms can only commit to one-period contracts. Fourth, we change the timing within the period so that wages are paid to workers at the end of the period after the output realization has been observed. Finally, to render the analysis pertinent, we assume parameter values to be such that in all cases firms prescribe the high level of effort ϵ^* to their workers.

These changes imply that only two types of job offer histories for workers are relevant for employment contracts: A worker has either received an outside job offer since his last unemployment spell or he has not. Conditional on a worker's job offer history, a one-period contract specifies a wage level for each possible output realization. As a result, there are four equilibrium levels of wages under moral hazard: w_u^+ and w_u^- for workers leaving unemployment, and w_o^+ and w_o^- for workers who have already received an outside offer. When effort is observable, there are only two equilibrium levels of wages, $\hat{w}_u$ and $\hat{w}_o$. As before, U^n denotes the value to a worker of being unemployed. To denote a worker's lifetime utility after having received an outside offer, we use U_o for the case of moral hazard, and $\hat{U}_o$ for that of observable effort. A firm's period expected profits after a worker has received an outside offer is denoted by Π_o under moral hazard, and by $\hat{\Pi}_o$ under observable effort.

We start the exposition by analyzing the problem of a firm whose worker has received an outside offer. Since, in this case, the level of lifetime utility promised to the worker is determined by firm competition, a firm makes zero expected period profits on the contract. Under moral hazard, this condition reads as

$$\Pi_o = p\big[A^+ - w_o^+\big] + (1-p)\big[A^- - w_o^-\big] = 0 \tag{22}$$

and implies that the mean wage is equal to mean output, $pw_o^+ + (1-p)w_o^- = pA^+ + (1-p)A^-$. It should be noted that the same argument implies that the wage under observable effort is equal to mean output, $\hat{w}_o = pA^+ + (1-p)A^-$. In the case of moral hazard, wages moreover need to satisfy the incentive-compatibility constraint

$$\begin{aligned} pu(w_o^+) + (1-p)u(w_o^-) - v + \beta\psi\Big\{\delta U^n + (1-\delta)U_o\Big\} \quad &\geq \\ u(w_o^-) + \beta\psi\Big\{\delta U^n + (1-\delta)U_o\Big\}& \end{aligned} \tag{23}$$

which, in equilibrium, will be satisfied with equality, and can then be reduced to

$$p\big[u(w_o^+) - u(w_o^-)\big] = v \tag{24}$$

These observations lead to the following results:

Lemma 3. *(i) At the top of the wage distribution, moral hazard leads to a mean-preserving spread in wages. (ii) Workers are unambiguously worse off under moral hazard than under observable effort after having received an outside offer.*

Proof. (i) We have shown above that $pw_o^+ + (1-p)w_o^- = \hat{w}_o$. Moreover, (24) implies that $w_o^+ > w_o^-$.

(ii) Given that workers face a mean-preserving spread in period wages and the same prospect in case they lose their job, this result is a trivial consequence of risk aversion. More formally, the value of employment after having received an outside offer, under moral hazard and under observable effort, respectively, are given by

$$U_o = \frac{pu(w_o^+) + (1-p)u(w_o^-) - v + \beta\psi\delta U^n}{1 - \beta\psi(1-\delta)} \quad \text{and} \quad \hat{U}_o = \frac{u(\hat{w}_o) - v + \beta\psi\delta U^n}{1 - \beta\psi(1-\delta)} \tag{25}$$

Part (i) and strict concavity of $u(\cdot)$ imply that $U_o < \hat{U}_o$. □

The intuition behind these results is simple. At the top of the wage distribution, firms need to break even; therefore, the expected wage is equal to expected output in both scenarios. Given that the level of effort is the same, average output and average wage are the same in both economies. Moreover, as long as effort is constant across the two scenarios, these results for workers who have received an outside offer can easily be generalized to a setting with dynamic contracts, different timing of wage payments, and multiple firm productivity levels.

We now turn to the problem of a firm employing a worker who has not yet received an outside offer. Due to the assumption that workers have no bargaining power, the value of unemployment U^n is the same in both scenarios. Under moral hazard, the promise-keeping constraint is given by

$$U^n = pu(w_u^+) + (1-p)u(w_u^-) - v + \beta\psi\Big\{\delta U^n + (1-\delta)\big[\lambda_e U_o + (1-\lambda_e)U^n\big]\Big\} \tag{26}$$

which can be rearranged to yield

$$U^n = \frac{pu(w_u^+) + (1-p)u(w_u^-) - v + \beta\psi(1-\delta)\lambda_e U_o}{1 - \beta\psi\big[\delta + (1-\delta)(1-\lambda_e)\big]} \tag{27}$$

Similar algebra leads to the following expression for U^n under observable effort:

$$U^n = \frac{u(\hat{w}_u) - v + \beta\psi(1-\delta)\lambda_e \hat{U}_o}{1 - \beta\psi\big[\delta + (1-\delta)(1-\lambda_e)\big]} \tag{28}$$

Equations (27) and (28) inform us about the following implication of moral hazard for wages at the bottom of the distribution:

Lemma 4. *The expected wage of workers who have not yet received an outside offer is higher under moral hazard than under observable effort.*

Proof. Since $U_o < \hat{U}_o$, (27) and (28) imply that $pu(w_u^+) + (1-p)u(w_u^-) > u(\hat{w}_u)$. From concavity of $u(\cdot)$ it follows that $pw_u^+ + (1-p)w_u^- > \hat{w}_u$. □

It should be noted that the expected wage of workers leaving unemployment is higher under moral hazard for two reasons. First, as stated in Lemma 3 above, workers face worse future prospects under moral hazard than when effort is observable, that is, $U_o < \hat{U}_o$. In consequence, they require higher contemporaneous utility and therefore a higher expected contemporaneous wage in order to accept a job out of unemployment. However, the average wage at the bottom of the distribution would still be higher under moral hazard, even if it were the case that $U_o = \hat{U}_o$. The reason is that wages need to satisfy an incentive-compatibility constraint corresponding to equation (24), which implies that $w_u^+ > w_u^-$. Due to concavity of utility, $pw_u^+ + (1-p)w_u^- > \hat{w}_u$ must then hold, even if $pu(w_u^+) + (1-p)u(w_u^-) = u(\hat{w}_u)$.

The above results show that introducing moral hazard affects the wage distribution in the simplified model both directly and indirectly. The direct effect consists of the fact that, both for workers who have received an outside offer and for those who have not, wages need to vary with output realizations in order to provide incentives. This wage variation obviously increases wage dispersion relative to the scenario with observable effort. At the top of the distribution, this is the only effect of moral hazard, since the average wage is the same in both scenarios. At the bottom of the distribution, however, additional channels lead to an increase in the average wage of workers who have not yet received an outside offer. This indirect effect reduces the difference in average wages between the two groups of workers. As a result, the overall impact of moral hazard on wage dispersion depends on the relative magnitude of the two effects.

In other words, the indirect effect of moral hazard leads to an upward compression of the wage distribution. This compression ultimately stems from the need to compensate risk-averse workers for wage variation. Interestingly, it is accompanied by a downward compression of lifetime utilities at the top of the distribution. While workers' expected lifetime utility U^n when leaving unemployment is the same in both scenarios, the maximal lifetime utility (after having received an outside offer) is strictly lower in the case of moral hazard. Given the set-up of firm competition, workers at the top of the distribution cannot be compensated for earnings variation by a higher expected wage. Instead, compensation takes place when workers leave unemployment, so that the top-down compression in lifetime utilities translates into an upward compression of wages at the bottom of the distribution.

The potentially most restrictive feature of the present analysis is the assumption of discrete and constant effort levels, both across job offer histories and across scenarios. It is straightforward to see that, if effort were continuous, moral hazard would lead to lower levels of effort, since incentive provision increases the costs of effort to firms. This would decrease the effort compensation component of wages, and thus impact on wage levels. Depending on whether the reduction in effort levels is larger at the top or the bottom of the distribution, this could either reinforce or partially offset the indirect compression effect of moral hazard on the wage distribution.

All the channels of impact of moral hazard described in the present simplified framework appear in the analysis of the full model too. In addition, the assumption of long-term contracts introduces more intricate dynamics of wages within a job even in the absence of outside offers. In the full model we moreover allow for continuous effort choice. This results in differences between the moral hazard and the observable effort scenarios in terms of implemented effort levels, which will in turn lead to differences in wage levels and in the dynamics of wages within a job.

4 Quantitative Analysis

Of the two parts of this section, the first provides a detailed description of the calibration strategy, the arguments supporting the selection of empirical targets, and the calibration results. The second part presents the results of our quantitative analysis, together with an extensive discussion of the mechanisms underlying the various effects of moral hazard on the wage distribution and their implications for residual wage inequality.

4.1 Calibration

In this section we lay out the calibration of our model to U.S. data. Starting from some basic specifications, we move on to a description of the approach to parameter selection and the choice of targets to match, and, finally, present the parameter values obtained by our procedure. A brief closing section outlines the data background to the calibration, including the regression specification used to estimate residual wages which are central to the conceptual framework of our analysis.

4.1.1 Functional forms and distributions

A worker's period utility from consumption is assumed to be given by a constant relative risk aversion (CRRA) utility function, that is,

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma} \tag{29}$$

where $\sigma > 1$ is the coefficient of relative risk aversion. The period disutility from spending effort is given by a power function

$$g(\epsilon) = \epsilon^{\gamma} \tag{30}$$

with exponent $\gamma > 0$, while the probability of a high output realization as a function of effort is given by an exponential function

$$\pi(\epsilon) = 1 - \exp\{-\rho\epsilon\} \tag{31}$$

with parameter $\rho > 0$. We normalize the levels of effort-dependent productivity as follows:

$$\begin{aligned} A^+ &= 1 + \Delta A \\ A^- &= 1 - \Delta A \end{aligned} \tag{32}$$

With regard to firm-specific productivity, we set $\mathcal{Z} = \{z_1, z_2, z_3\}$, with $z_1 < z_2 < z_3$. The distribution over these values from which a newly established firm draws its productivity level is denoted by

$$f(z) = \{f(z_1), f(z_2), 1 - f(z_1) - f(z_2)\} \tag{33}$$

Finally, we normalize firm productivity levels in the following way:

$$\begin{aligned} z_1 &= 1 - \Delta_z \\ z_2 &= 1 \\ z_3 &= 1 + 2\Delta_z \end{aligned} \tag{34}$$

4.1.2 Strategy and results

Given the above assumptions on functional forms and on the distribution of firm-specific productivity levels, values for the following thirteen parameters need to be selected:

1. β, the discount factor,
2. σ, the workers' coefficient of relative risk aversion,
3. ψ, the probability of survival of a worker,
4. δ, the probability of exogenous destruction of a worker-firm match,
5. λ_u, the probability for unemployed workers of receiving a job offer,
6. λ_e, the probability for employed workers of receiving an outside job offer,
7. γ, the power parameter in the function for disutility from spending effort,
8. ρ, the parameter in the expression for the probability of a high realization of worker productivity as a function of effort,

9. ΔA, the parameter determining the difference between high and low realizations of worker productivity,

10. b, the level of consumption while unemployed,

11. Δ_z, the parameter determining differences in firm-specific productivity levels, and

12. $f(z)$, the probability for a newly established firm to draw productivity level z.

Some of these parameters are selected in accordance with the literature. Others are chosen with the aim of matching empirical statistics on selected features of the U.S. labor market in the mid-2000s by corresponding stationary equilibrium statistics of the model. We obtain empirical statistics for our calibration from the 2004 panel of the Survey of Income and Program Participation (SIPP). The last part of the present section provides some information on the data as well as on measurement and estimation of the variables of interest.

We choose the length of a time period in the model to be one quarter.[11] The discount factor (β) is set at 0.99, a value which corresponds to an annual interest rate of 4%, and the coefficient of workers' relative risk aversion (σ) at the value 2, which is standard in the quantitative macroeconomic literature.

The survival probability of workers (ψ) is set at 0.994, a value which corresponds to an expected length of an individual's working life of forty years. The probability of exogenous destruction of a worker-firm match (δ) is chosen such that the rate of employment-to-unemployment flows (τ_{eu}) in the data is matched. Similarly, the probability for an unemployed worker to receive a job offer (λ_u) is selected with a view to matching the empirical job finding rate (τ_{ue}). In order to obtain these parameter values, we use the following two equations which, conditional on ψ, pin down the flows between employment and unemployment in stationary equilibrium:

$$\tau^{eu} = \delta + (1 - \delta)(1 - \psi) \tag{35}$$

$$\tau^{ue} = \psi\lambda_u \tag{36}$$

Accordingly, δ is set at the value of 0.018, whereas λ_u is set at 0.55.

The remaining parameters on the list are selected with the objective of matching empirical statistics on the frequency of job-to-job transitions, individuals' wage changes both within and between jobs, and the cross-sectional distribution of wages. In view of the empirical rate of job-to-job transitions (τ_{ee}), we assign to the probability for an employed

[11] This is a compromise between using short time periods which correspond more closely to the high frequency of workers' labor market transitions, on the one hand, and computational feasibility of solving the model numerically, on the other.

worker to receive an outside job offer (λ_e) the value of 0.15. The next three parameters, γ, ρ, and ΔA, are most closely connected with costs and benefits of incentive provision when effort is not observable. As was discussed before, incentive provision under moral hazard leads to output-dependent variation of wages within a job. We therefore select values for the above three parameters with a view to matching empirical statistics on within-job wage changes. The targets used in this connection are statistics on differences between two periods in the log wages of workers who stay with the same employer. In particular, we employ as targets the following statistics on wage change within a job: the mean ($\mu(\Delta \ln w^{win})$), the standard deviation ($\sigma(\Delta \ln w^{win})$), and the fraction of negative log wage changes ($P\big[(\Delta \ln w^{win}) < 0\big]$). In this context, a complication arises from the fact that the effects of the power parameters in the function of disutility from effort (γ) and in the function relating the probability of high output realization to the level of effort (ρ) cannot be disentangled. For this reason, we choose to fix γ at a value of 2, that is, to assume quadratic disutility from effort, and to select the other two parameters with the goal of matching statistics on within-job wage changes. As a result, the parameter of the function determining effort-dependence of output realizations (ρ) is set to the value of 3. For the difference between high and low output realizations (ΔA) we select a value of 0.26.

Values for the remaining parameters are chosen in such a way that the model matches empirical statistics of the cross-sectional wage distribution and the distribution of individuals' wage changes upon job-to-job transitions. With a view to matching the average level of wages in the cross-section ($\mu(w)$), we set the level of consumption of unemployed workers (b) to 0.8. The parameter determining differences between firm-specific productivity levels (Δ_z) as well as the productivity distribution ($f(z)$) are closely related to the overall spread of wages in the model economy. Accordingly, when choosing values for these parameters, we aim to match the standard deviation of log wages ($\sigma(\ln w)$) as a measure of wage inequality in the cross-section. Moreover, the parameters in question have strong effects on the dynamics of wages associated with job switching. Since our model framework is capable of producing only a few job-to-job transitions that are associated with wage loss to the worker, we choose to match statistics of the distribution of positive wage changes between jobs. In particular, we use as targets the mean ($\mu(\Delta \ln w_+^{bet})$) and the standard deviation ($\sigma(\Delta \ln w_+^{bet})$) of positive changes in log wages between jobs. As a result, the difference parameter of firm productivity levels (Δ_z) is set to 0.26, while the probability distribution of productivities ($\{f(z_1), f(z_2), f(z_3)\}$) is chosen as $\{0.4, 0.36, 0.24\}$.

Table 4.1 summarizes the parameter values of our calibration. The simulated model statistics together with the empirical targets are reported in Table 4.2. From the latter one can see that the calibrated economy reproduces accurately the empirical flows of workers between labor market states. Both the scale and the dispersion of cross-sectional wages are matched well, as, in particular, the model produces 86% of the empirically observed residual wage inequality. With respect to statistics on within-job wage changes, the model reproduces accurately the fraction of negative log wage changes. Within the framework of our analysis, this means that it captures well the relative probabilities of workers' high and low output realizations. However, the model exhibits too high a mean and too low a standard deviation for within-job wage changes. The reason is that wage gains within a job due to outside offers are large relative to the scale of variation of wages due to the moral hazard problem. Finally, the mean and variation in positive wage changes upon job-to-job transitions are matched reasonably well. As was mentioned in the discussion on the choice of targets, the model produces a very low fraction of negative wage changes associated with job-to-job transitions. To sum up, the calibrated economy reproduces workers' transition rates and moments of the cross-sectional distribution of residual wages very well. Given the focus of our analysis, it does reasonably well in matching statistics of wage changes both within and between jobs.[12]

Table 4.1: Values of model parameters

Parameter	Value	Parameter	Value
β	0.99	ρ	3
σ	2	ΔA	0.35
ψ	0.994	b	0.8
δ	0.018	Δ_z	0.26
λ_u	0.55	$f(z_1)$	0.4
λ_e	0.15	$f(z_2)$	0.36
γ	2	$f(z_3)$	0.24

4.1.3 Data background

We obtain empirical statistics for use in the calibration from the 2004 panel of the Survey of Income and Program Participation (SIPP). The data set contains information at high frequencies of individuals' labor market histories and income sources for a representative

[12]Our numerical analyses suggest that significant improvements in matching statistics on wage changes both within and between jobs can only be achieved by a significant increase in the number of firm productivity levels. A modification of the quantitative model, in which we assume a particular functional form for the distribution of firm productivities and calibrate the distributional parameters to statistics on between-job wage changes and cross-sectional wage dispersion, is work in progress.

Table 4.2: Simulated vs. empirical statistics

Statistic	Model	Data
τ^{eu}	0.0236	0.0238
τ^{ue}	0.5443	0.5484
τ^{ee}	0.0387	0.0391
$\mu(w)$	1.1662	1.1324
$\sigma(\ln w)$	0.4147	0.4824
$\mu(\Delta \ln w^{win})$	0.0148	0.0060
$\sigma(\Delta \ln w^{win})$	0.0851	0.1416
$P\big[(\Delta \ln w^{win}) < 0\big]$	0.3511	0.3550
$\mu(\Delta \ln w_{+}^{bet})$	0.3159	0.2784
$\sigma(\Delta \ln w_{+}^{bet})$	0.2963	0.2580

sample of households in the United States. In particular, the SIPP collects detailed job-specific information for up to two wage and salary jobs per person in a given period (wave). This information allows for a distinction between different jobs that a person has held with different employers over the time span of the panel. In consequence, it is possible to quite reliably set apart job-to-job transitions from other types of labor market transitions. Moreover, the structure of these data allows for a quantitative assessment of the dynamics of wages, both with regard to within-job changes and to changes between jobs upon job-to-job transition.

In the present exercise, we use data covering the period from January 2004 to December 2006 and restrict our attention to information on male workers between 20 and 65 years of age. A further restriction of our sample is that to individuals who have held at least one non-contingent, non-self-employed paid job during the time span of the panel. We classify individuals as employed or unemployed in a given month according to the labor market status they report in the second week of that month. Based on this classification, we compute monthly rates of transition between the states of employment and unemployment. In our measure of job-to-job transitions, we include workers who are classified as employed with two different main employers in two consecutive months, do not return to a previous employer, and have not been without a job and searching or on layoff in between.

For the estimation of residual wages, we further restrict the sample to individuals who are working full-time. Based on this sample, we run a pooled regression of log real hourly wages on the following variables: five educational groups, a quadratic in potential experience, interaction terms between education groups and experience, four region groups, a dummy for being non-white, and year dummies. The residuals from this regression are used to compute statistics of the cross-sectional wage distribution, as well as statistics on wage changes within jobs and between jobs upon job-to-job transitions. Appendix A.2 provides additional information on the data underlying our calibration.

4.2 Computational experiments

The goal of the analysis presented in this section is to provide a quantitative assessment of the overall impact of moral hazard on wage inequality, as well as to identify and characterize the various channels through which this impact is brought to bear on the wage distribution. To this effect, we analyze different versions of a quantitative economy that are based on the model framework outlined in Section 2 and parametrized according to the calibration procedure described in Section 4.1.

4.2.1 The main result

We first provide a quantitative answer to one of the central questions of this paper: What is the overall impact of moral hazard on wage inequality when employed workers engage in job search and firms compete for workers? This answer is derived from a comparison between two scenarios which differ only with respect to the assumptions made about observability of worker effort. The first scenario is identical with that described in Section 2 where, in particular, worker effort is not observable and, in order to extract positive effort from workers, incentives are provided through variation in future utility. We label this the moral hazard (MH) scenario. The contractual problem in this case takes the form of the optimization problem (16) subject to constraints (17) to (21).

For the second scenario, we assume that worker effort is observable and therefore becomes an explicit part of the contractual arrangement between firms and workers. In this observable effort (OE) scenario, since incentive-compatibility is not required, an optimal contract is a solution to the optimization problem (16) subject to constraints (17), (19), (20), and (21). As stated in Lemma 2 of Section 2.3, in this environment a firm fully insures its worker against effort-dependent productivity shocks, that is, $U' \equiv U^- = U^+$.

Table 4.3 presents a comparison of the two stationary equilibria belonging to the above scenarios in terms of various measures of wage inequality. The difference – given in the last column of the table – is stated as the percentage change associated with a transition from the OE to the MH scenario. All statistics shown in Table 4.3 indicate that the overall consequence of introducing a moral hazard problem into the labor contract is an increase in residual wage inequality. The size of inequality increase lies between 3.5% and 6.9% for comprehensive measures such as the standard deviation of log wages, the coefficient of variation, and the Gini coefficient. However, a closer look at changes in wage percentiles reveals that the impact of moral hazard is not even across different parts of the wage distribution. The percentile ratios presented in Table 4.3 show that inequality between the top and bottom five percent of the wage distribution increases by nearly 15%. Moreover, the larger part of this increase comes from a rise in inequality within the lower half of the distribution, as measured by the 50th-to-5th percentile ratio. Finally, the last row in the table reports a change in the mean-min ratio of nearly 60%. Since average wages in the two scenarios are close to equal, this large increase confirms that moral hazard has a particularly strong impact on the lowest parts of the wage distribution.

Table 4.3: Measures of wage inequality: OE vs. MH

	OE	MH	Difference
Std($\ln w$)	0.389	0.415	6.52 %
CV(w)	0.340	0.352	3.50 %
Gini(w)	0.188	0.201	6.87 %
95/05(w)	2.998	3.443	14.84 %
95/50(w)	1.380	1.466	6.28 %
50/05(w)	2.173	2.348	8.06 %
Mean/min(w)	2.329	3.695	58.64 %

As shown in the analysis of the simplified model, moral hazard shapes the wage distribution through a variety of effects. In the next section, we attempt to disentangle and characterize the different channels of impact in the context of the full model, which will shed light on the mechanisms underlying the present results. As a point of departure, in the following paragraphs we discuss differences in firms' optimal policies between the two scenarios and outline their relevance for differences in the cross-sectional wage distributions.

Figure 4.1 shows policy functions of firms (of the z_3-type) for continuation utilities U^i. In the MH scenario, incentive-compatibility requires that workers are – within their current labor contract – rewarded for high output realizations with an increase in utility and penalized for low output realizations with a utility decrease ($U^+_{MH} > U^-_{MH}$). Accordingly, among workers who start working at the same type of firm, lifetime utilities drift apart over time, independently of whether or not they receive outside job offers, merely as a result of differences in their histories of stochastic-output realizations. By contrast, in the OE scenario the continuation utility for a worker within his current labor contract is non-stochastic ($U^{+}_{OE} = U^{-}_{OE} \equiv U'_{OE}$). Moreover, in the case of a z_3-type firm, this utility level is constant over time.[13] The presence of moral hazard, by introducing stochastic dynamics in lifetime utility into labor contracts, thus adds a specific source of wage dispersion to the environment.

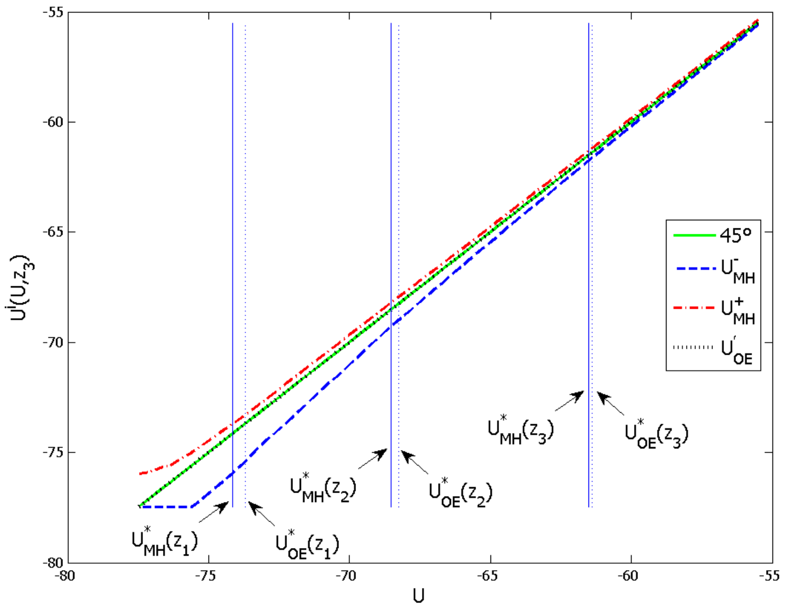

Figure 4.1: Policy functions for continuation utilities, MH vs. OE case, z_3-firms

In parallel to the simplified model, this direct effect of moral hazard is complemented by a number of additional, indirect effects. For all of these, the cost to a firm of extracting the required level of effort plays a central role. A comparison between the MH and the OE scenarios shows that the need to provide incentives makes worker effort more expensive for firms in the former scenario than in the latter. In both scenarios workers have to be compensated for the disutility from spending effort. Under moral hazard, however, risk-averse workers need, in addition, to be compensated for the variation in income arising from incentive provision. These differences in effort costs translate into differences between the wage distributions of the two scenarios through various channels.

One of the component effects of increased effort costs works through changes in critical utility levels. Relatively higher effort costs under moral hazard lead to a relatively lower value to the firm of any worker-firm match. Thus, for all firm types the value function of the MH scenario is below that of the OE scenario, which implies that critical utility levels

[13] For z_1-type and z_2-type firms this is also true for $U < U^*(z)$. Above this value, where the firm is making losses, the potential arrival of outside offers from more productive firms introduces an externality due to which U and U' are not necessarily equal. For reasons of clarity of presentation, we use policy functions of z_3-type firms for whom this externality is not present.

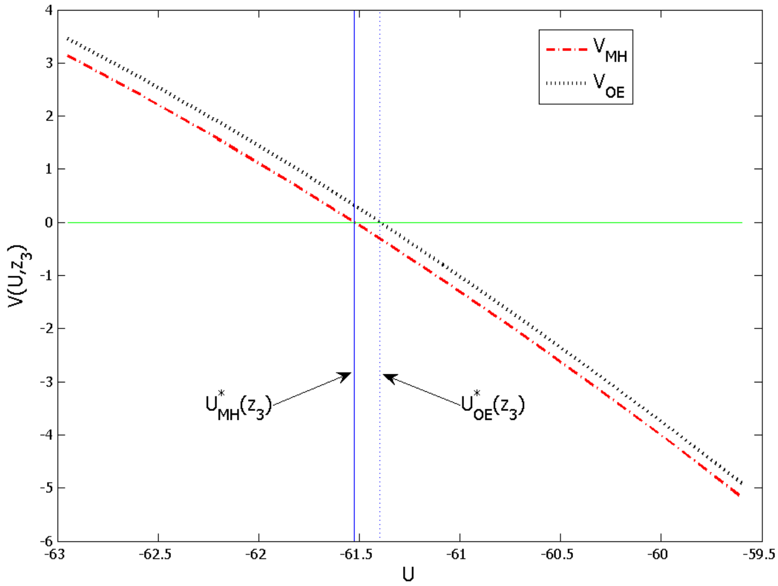

Figure 4.2: Value functions and critical utility levels, MH vs. OE case, z_3-firms

$U^*(z)$ are lower too. Figure 4.2 illustrates the differences between the value functions and critical utilities of z_3-type firms corresponding to the two scenarios. Figure 4.1 depicts the differences in critical utility levels between the MH and the OE scenarios for all firm types, showing that the size of this difference is inversely related to the level of firm productivity. As regards the impact on the wage distribution, the differences in $U^*(z)$-levels are relevant for differences in the levels of wages associated with outside job offers. On the one hand, a decrease in critical utility levels translates into lower wage gains that workers can achieve through outside offers, and therefore into lower maximum wage levels attained through on-the-job search. On the other hand, as demonstrated in the simplified model above, workers need to be compensated for lower continuation values associated with outside offers by higher expected wages when they leave unemployment.

In contrast to the simplified model, in the present framework we allow for continuous choice of worker effort. This implies that, in the full model, another component effect of increased effort costs is associated with the levels of effort that firms seek to extract from their workers. Figure 4.3 shows – again for the case of z_3-type firms – the policy functions for effort in the MH and the OE case. As expected, effort levels are generally lower under moral hazard, where by and large differences between the two scenarios become smaller with increasing levels of lifetime utility. In addition, the figure exhibits a difference in shape between the two policy functions: While in the OE scenario effort levels are monotonically decreasing in lifetime utility throughout, in the MH case effort is increasing over a short interval of low utility values and decreasing over the rest of the domain.

The initial increase in the effort function under moral hazard arises from specific costs

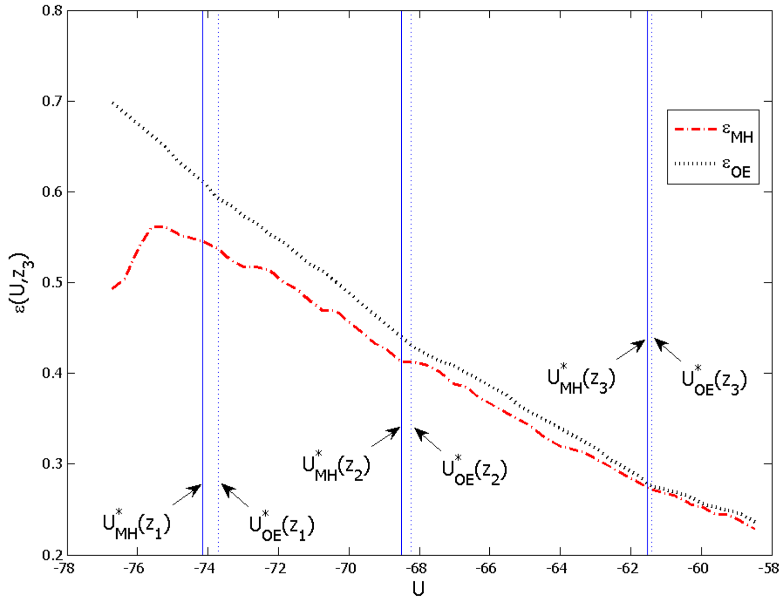

Figure 4.3: Policy functions for effort, MH vs. OE case, z_3-firms

of incentive provision at low levels of lifetime utility. For values close to U^n firms have only little scope to punish workers for low output realizations through utility cuts because in this situation a worker's participation constraint becomes binding easily. Consequently, incentives to extract effort have primarily to take the form of rewarding high output realizations by utility increase. This additional constraint leads to a further rise of the cost of effort to firms. As lifetime utility increases, the cost of the participation constraint declines so that initially effort levels increase in utility. The range of lifetime utility over which the participation constraint is binding can be clearly seen in Figure 4.1. This region, where the function U^-_{MH} is flat, coincides with the interval over which the effort function is increasing.

Leaving everything else unchanged, lower effort levels under moral hazard translate into lower wages at a given level of lifetime utility, since workers have to be compensated less for disutility from effort. Moreover, under moral hazard effort levels have a direct impact on the dynamics of lifetime utility of workers. As the level of effort of a worker determines the respective probabilities of an increase or decrease in utility in the next period, it affects the stochastic process of a worker's wage change within a job over time.

Finally, changes between the two scenarios in firms' policy functions for wages are related to all of the indirect effects of incentive provision. Most immediately, at a given level of lifetime utility, wages increase in order to compensate workers for the additional variation in earnings. At low values of lifetime utility, wages also increase in order to compensate workers for lower continuation values associated with outside offers. They decrease, however, as a consequence of lower effort compensation costs due to lower effort

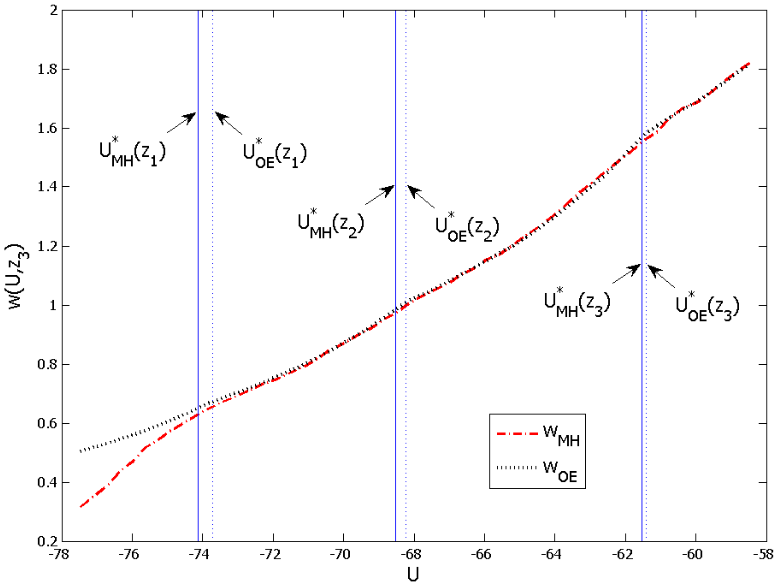

Figure 4.4: Policy functions for wage, MH vs. OE case, z_3-firms

levels prescribed to workers. Finally, wages also adjust to changes in the stochastic process of a worker's lifetime utility arising from lower effort levels. Figure 4.4 presents, again for the case of z_3-type firms, the policy functions for wages in the two scenarios. It clearly shows how the strong decrease in effort levels at low values of lifetime utility translates into significantly lower starting wages under moral hazard than under observable effort. The figure also illustrates that the various forces on firms' wage policies more or less cancel out in the case of z_3-type firms at higher values of lifetime utility.

4.2.2 Decomposition of impact and exploration of mechanisms

This section presents a series of computational experiments that aim at decomposing the overall impact of moral hazard on the wage distribution into different channels. For this purpose, we start from the observable effort scenario and add, one-by-one, the constraints as well as the components of the solution to the firms' optimal contract design problem under moral hazard. Apart from clarifying the mechanisms underlying our main result, these experiments provide an assessment of the direction of impact and the relative quantitative importance of the different effects of moral hazard on wage inequality.

Step 1: Incentive-compatibility

In the first step, we introduce the incentive-compatibility requirement on workers' continuation values. More precisely, we assume that firms prescribe the same levels of effort to workers as in the observable effort scenario, but need to provide incentives through variation in continuation utility. Moreover, firms are not allowed to adjust contemporaneous wage levels, and we assume that critical utility levels are the same as in the observable effort case. Technically, we keep the policy functions for wage and effort as well as the critical utility levels fixed at the OE solution, so that $w_{CF} = w_{OE}(U,z)$, $\epsilon_{CF} = \epsilon_{OE}(U,z)$, and $U^*_{CF} = U^*_{OE}(z)$, and solve (ICC) and (PKC) for the counterfactual policy functions $U^+_{CF}(U,z)$ and $U^-_{CF}(U,z)$ subject to (PC).

Table 4.4 reports changes in statistics on the wage distribution that are associated with a transition from the OE scenario to the present counterfactual scenario (EXP1). Given the setup of the experiment, these changes obviously contain the direct effect of moral hazard which consists in incentive-providing wage variation within a job. The changes in wage percentiles presented in the table show that, both at the top and the bottom of the distribution wage differences within groups of workers with the same job offer history do indeed expand as expected. At the same time, changes in the lowest wage percentiles reveal an upward compression of wages from the lower end towards the middle of the distribution. This compression arises from two forces acting at low levels of lifetime utility. On the one hand, in analogy to the simplified model, risk-averse workers need to be compensated for wage variation by higher expected future wages. On the other hand, incentive provision through variation in future utility is restricted by the worker's participation constraint. Since wage punishments for low output are constrained from below, firms need to reward high output with very large wage raises in order to implement the given levels of effort, thereby increasing a worker's expected future wage. Both forces thus cause wages of workers at low levels of lifetime utility to increase more quickly than in the OE scenario.

In the present experiment the expansion of within-group wage differences, attributable to the direct effect of moral hazard, is dominated by the upward compression arising from indirect effects. As a result, overall wage inequality decreases by nearly ten percent.

Table 4.4: Decomposition step 1: Incentive-compatibility

	OE	EXP1	Difference
Std($\ln w$)	0.389	0.351	-9.94 %
Min(w)	0.504	0.504	0.00 %
P05(w)	0.528	0.600	13.80 %
P10(w)	0.597	0.687	15.03 %
P50(w)	1.147	1.193	4.05 %
P90(w)	1.582	1.663	5.11 %
P95(w)	1.582	1.717	8.56 %
P99(w)	1.582	1.872	18.32 %

Step 2: Critical utility levels

The setting of the second experiment (EXP2) is similar to that of the first one, except that we change the critical utility levels to the values of the moral hazard scenario. This means that we add the reduction in lifetime utility levels associated with outside job offers. Technically, we set $w_{CF} = w_{OE}(U, z)$, $\epsilon_{CF} = \epsilon_{OE}(U, z)$, $U^*_{CF} = U^*_{MH}(z)$, and obtain the counterfactual policy functions $U^+_{CF}(U, z)$ and $U^-_{CF}(U, z)$ by solving (PKC) and (ICC) subject to (PC).

The changes in wage percentiles reported in Table 4.5 show that the reduction in critical utility levels affects both the bottom and the top of the wage distribution. First, similar to the mechanism demonstrated in the simplified model, workers at low levels of lifetime utility need to be compensated for the relatively worse prospects associated with outside job offers by higher expected continuation values. This again causes low wage workers to move up faster and leads to an upward compression of wages from the bottom of the distribution. Second, workers who have received outside offers attain lower levels of lifetime utility and thus lower wages. Consequently, wages are compressed downwards from the top of the distribution. Both indirect effects at work in the present experiment thus counteract the direct effect of moral hazard. However, the overall reduction in wage inequality associated with the decrease in critical utility levels is modest.

Table 4.5: Decomposition step 2: Critical utility levels

	EXP1	EXP2	Difference
Std($\ln w$)	0.351	0.349	-0.51 %
Min(w)	0.504	0.504	0.00 %
P05(w)	0.600	0.601	0.17 %
P10(w)	0.687	0.688	0.09 %
P50(w)	1.193	1.188	-0.44 %
P90(w)	1.663	1.655	-0.45 %
P95(w)	1.717	1.712	-0.32 %
P99(w)	1.872	1.869	-0.18 %

Step 3: Effort levels

In the third experiment (EXP3), the change in effort levels between the OE and the MH scenarios is added. Since we still keep firms' policy functions for wages fixed at the observable effort solution, this counterfactual scenario captures mainly the effect of lower effort levels on the wage distribution that operates through lower probabilities of wage increases within a job. Technically, we impose $w_{CF} = w_{OE}(U, z)$, $\epsilon_{CF} = \epsilon_{MH}(U, z)$, $U^*_{CF} = U^*_{MH}(z)$, and obtain the counterfactual policy functions $U^+_{CF}(U, z)$ and $U^-_{CF}(U, z)$ by solving (PKC) and (ICC) subject to (PC).

As can be seen from the changes in wage percentiles reported in Table 4.6, the indirect effect working through changes in lifetime utility dynamics caused by lower effort levels also affects both the top and the bottom of the wage distribution. On the one hand, lower probabilities of wage raises at high values of lifetime utility lead to a slight downward compression of wages from the top. On the other hand, as illustrated in Figure 4.3, the reduction in effort levels in response to moral hazard is much larger at low values of lifetime utility. Therefore, workers who have not yet received an outside job offer, on average, stay much longer at low wage levels than in the previous scenario. These changes in wage dynamics at low values of lifetime utility partly reverse the upward compression of wages observed in the first two experiments. Moreover, the strength of impact on the lower parts of the wage distribution is clearly dominant in the present experiment. Overall, the changes in workers' dynamics of lifetime utility associated with lower effort levels lead to an increase in wage inequality by nearly seven percent.

Table 4.6: Decomposition step 3: Effort levels

	EXP2	EXP3	Difference
Std($\ln w$)	0.349	0.372	6.77 %
Min(w)	0.504	0.504	0.00 %
P05(w)	0.601	0.591	-1.68 %
P10(w)	0.688	0.604	-12.17 %
P50(w)	1.188	1.165	-1.92 %
P90(w)	1.655	1.647	-0.50 %
P95(w)	1.712	1.692	-1.14 %
P99(w)	1.869	1.790	-4.21 %

Step 4: Wage levels

The last step completes the transition from the OE to the MH scenario by adjusting firms' wage policies to the solution under moral hazard. As discussed in Section 4.2.1, changes in wage policies between the OE and the MH scenarios are associated with a variety of indirect effects of moral hazard. The changes in wage percentiles reported in Table 4.7 indicate that, among these effects, the downward adjustment of wages in response to lower effort compensation costs is dominant. At high levels of lifetime utility, this again leads to a downward compression of wages. However, the large decrease in wage levels at low values of lifetime utility leads to a strong expansion of wages at the bottom of the distribution. The adjustment of wage levels thus amplifies the inequality-increasing effect of reduced effort levels in the lower parts of the wage distribution. Analogous to the previous experiment, the impact at low levels of lifetime utility is dominant, and overall wage inequality increases by around 11%.

Table 4.7: Decomposition step 4: Wage levels

	EXP3	MH	Difference
Std($\ln w$)	0.372	0.415	11.36 %
Min(w)	0.504	0.316	-37.38 %
P05(w)	0.591	0.489	-17.27 %
P10(w)	0.604	0.591	-2.12 %
P50(w)	1.165	1.149	-1.39 %
P90(w)	1.647	1.647	0.00 %
P95(w)	1.692	1.685	-0.46 %
P99(w)	1.790	1.786	-0.24 %

4.2.3 A brief summary of results

The results of our quantitative analysis can be summarized in a few points. First, on balance, the presence of moral hazard in labor contracts increases residual wage inequality by around six percent. Second, computational experiments intended to disentangle partial effects and to explore the underlying mechanisms suggest that the direct effect of moral hazard, attributable to incentive-providing utility variation, produces within-group wage dispersion which accounts for a moderate contribution to inequality increase.[14] Third, some of the indirect effects working through channels outlined above are found to counteract the direct effect of moral hazard. Fourth, the indirect effect that is most closely linked to outside wage offers and firm competition opposes the direct effect, exerting a modest influence on inequality. Fifth, the lack of observability of effort, through other indirect effects, has a particularly strong impact on the lower parts of the wage distribution. It

[14] That the direct effect of moral hazard is quantitatively moderate is first of all suggested by the fact that in the first experiment (EXP1) it is dominated by counteracting indirect effects. Furthermore, in an additional counterfactual experiment not presented here we start from the moral hazard scenario and remove the incentive-compatibility requirement on continuation values while keeping wage and effort policy functions and critical utility levels fixed at the moral hazard solution. The difference in inequality between this counterfactual and the MH scenario is around three percent.

originates mainly from the fact that, under moral hazard, firms demand – in response to higher effort costs – significantly lower levels of effort from low wage workers. This results in a more than proportional increase of inequality within the lower half of the wage distribution, which contributes substantially to the overall impact of moral hazard on residual wage inequality.

5 Concluding Remarks

In the present paper we study moral hazard as a source of wage dispersion, using a search model that features job-to-job mobility with firm competition for workers. More specifically, we focus on the particular informational friction – internal to the firm – of worker effort being unobservable to the employer. This friction creates a moral hazard problem within worker-firm relations which eventually affects the wage distribution. Against this background, the specific goal of our analysis is two-fold: First, we quantify the overall impact of moral hazard on residual wage inequality. Second, we identify and quantify different components of the total effect and explore the underlying mechanisms, taking into account potential interactions between search frictions and the particular informational friction of unobservable effort.

An intuitive argument for why non-observability of worker effort should increase wage dispersion starts from the reduction in risk-sharing that moral hazard is known to bring about. In the context of labor contracts, this reduction takes the form of incentive-creating wage variation as the optimal response to the moral hazard problem. The implied output-dependence of wages should in turn lead to higher wage dispersion. In an environment in which wage dynamics arise from both the optimal design of long-term contracts and from job-to-job mobility under employer competition, the direct effect of incentive provision is, however, complemented by a number of indirect effects. These indirect channels of impact arise from an increase in the cost of worker effort to the firm when incentives need to be provided through spreads in future wages. Since some of the channels affect the wage distribution in opposite directions, the overall impact of moral hazard on wage dispersion needs to be determined by a quantitative analysis.

Based on a calibration of our model to statistics on wage dispersion and individual wage dynamics in the U.S. labor market, we find that, on balance, the presence of moral hazard increases residual wage inequality by around six percent. An exploration of the underlying mechanisms suggests that the direct effect of incentive provision leads to a moderate inequality increase due to wage dispersion within groups of workers with the same job offer history. Among the indirect effects, the one most closely related to outside wage offers and firm competition counteracts the direct effect, exerting a modest influence on inequality. By contrast, other channels of impact produce particularly strong indirect effects on the lower parts of the wage distribution. More specifically, due to a large decrease of effort levels in response to a rise in effort costs, inequality increases more than proportionally within the lower half of the distribution.

The nature of our findings suggests a specific extension of the analysis and a general remark on implications for related policy studies. On the one hand, we analyze in our study only differences across identical workers. Given that the strength of impact of moral hazard depends on the level of worker productivity, which may differ between occupations or sectors, it would be a natural extension to incorporate in the model worker heterogeneity in these dimensions. This modification would allow for a joint analysis of the impact of moral hazard on both within-group and between-group wage inequality and may provide interesting implications for the differences in wage dispersion along these lines. On the other hand, the strong interaction between incentive provision, effort levels, and wages in our framework indicates that incorporating moral hazard in the estimation of equilibrium models of on-the-job search and employer competition could change the results considerably. Related to this, an important next step would be to study how these interactions change the policy implications one can draw from analyses based on such models. In particular, since the effect of moral hazard is particularly strong in the lower parts of the wage distribution, the assessment of the impact of minimum wage and unemployment insurance policies might be affected significantly.

References

Alvarez-Parra, F. A. (2008). *Essays in Macroeconomics: Environments with Informational and Financial Frictions*. PhD thesis, University of Rochester.

Eckstein, Z. and Nagypal, E. (2004). The Evolution of U.S. Earnings Inequality: 1961-2002. *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(2):10–29.

Fallick, B. and Fleischman, C. A. (2004). Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows. Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System (U.S.).

Hornstein, A., Krusell, P., and Violante, G. L. (2011). Frictional Wage Dispersion in Search Models: A Quantitative Assessment. *American Economic Review*, 101(7):2873–98.

Katz, L. F. and Autor, D. H. (1999). Changes in the Wage Structure and Earnings Inequality. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3A. North Holland, Amsterdam.

Lise, J., Meghir, C., and Robin, J.-M. (2012). Matching, Sorting, and Wages. University College London, mimeo.

MacLeod, W. B. and Malcomson, J. M. (1998). Motivation and Markets. *The American Economic Review*, 88(3):388–411.

Manoli, D. and Sannikov, Y. (2005). Competitive Contracting and Employment Dynamics. University of California at Berkeley, mimeo.

Menzio, G., Telyukova, I. A., and Visschers, L. (2012). Directed Search over the Life Cycle. University of Pennsylvania, mimeo.

Nagypal, E. (2008). Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions. Northwestern University, mimeo.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica*, 70(6):2295–2350.

Rogerson, W. P. (1985a). Repeated Moral Hazard. *Econometrica*, 53(1):69–76.

Rogerson, W. P. (1985b). The First-Order Approach to Principal-Agent Problems. *Econometrica*, 53(6):1357–1367.

Spear, S. E. and Srivastava, S. (1987). On Repeated Moral Hazard with Discounting. *The Review of Economic Studies*, 54(4):599–617.

Thomas, J. and Worrall, T. (1988). Self-Enforcing Wage Contracts. *The Review of Economic Studies*, 55(4):541–553.

Tsuyuhara, K. (2011a). Repeated Moral Hazard with Worker Mobility via Directed On-the-Job Search. University of Calgary, mimeo.

Tsuyuhara, K. (2011b). Work Incentives and Labour Mobility: Quantitative Assessment of Macroeconomic Implications. University of Calgary, mimeo.

A Appendix

A.1 Proof of Lemma 2

Suppose there is a state (U, z) for which the optimal solution (w, ϵ, U^+, U^-) satisfies $U^+ \neq U^-$. Consider the case in which $U^+ > U^-$ (the opposite case follows the same arguments). One of the following mutually exclusive cases must apply:

1. $U^+ \notin \{U^*(z_n)\}_{n=1}^N$ and $U^- \notin \{U^*(z_n)\}_{n=1}^N$
2. $U^+ \in \{U^*(z_n)\}_{n=1}^N$ and $U^- \notin \{U^*(z_n)\}_{n=1}^N$
3. $U^+ \notin \{U^*(z_n)\}_{n=1}^N$ and $U^- \in \{U^*(z_n)\}_{n=1}^N$
4. $U^+ \in \{U^*(z_n)\}_{n=1}^N$ and $U^- \in \{U^*(z_n)\}_{n=1}^N$

Case 1. Consider the alternative candidate $(w, \epsilon, \tilde{U}^+, \tilde{U}^-)$ where w and ϵ are the same as in the optimal solution, but $\tilde{U}^+ = U^+ - \eta^+$ and $\tilde{U}^- = U^- + \eta^-$. The scalars (η^+, η^-) are positive and will be defined such that the alternative candidate gives the worker the same level of lifetime utility U as the optimal solution. Define z_i as the highest z_n satisfying $U^*(z_n) \leq U^-$, and z_j as the highest z_n satisfying $U^*(z_n) \leq U^+$. Then the following relationships hold:

- $z_i \leq z_j$
- $U^*(z_i) < U^- < U^*(z_{i+1})$ and $U^*(z_j) < U^+ < U^*(z_{j+1})$

For sufficiently low (η^+, η^-), it is the case that

$$U^*(z_i) < \tilde{U}^- < U^*(z_{i+1}) \quad \text{and} \quad U^*(z_j) < \tilde{U}^+ < U^*(z_{j+1}) \tag{37}$$

The requirement that the alternative candidate gives the same utility to the worker implies that

$$\eta^+ = \left(\frac{1 - \pi(\epsilon)}{\pi(\epsilon)} \frac{[1 - \lambda_e + \lambda_e F(z_i)]}{[1 - \lambda_e + \lambda_e F(z_j)]} \right) \eta^- \tag{38}$$

The values (η^+, η^-) satisfy conditions (37) and (38).

Let V denote the value of the contract to the firm under the optimal solution, and $\tilde{V}$ the value associated with the alternative candidate. Then

$$\tilde{V} - V = \beta\psi(1 - \delta) \Big\{ \big[aV(\tilde{U}^+) + bV(\tilde{U}^-) \big] - \big[aV(U^+) + bV(U^-) \big] \Big\} \tag{39}$$

where

$$a = \pi(\epsilon)\left[1-\lambda_e+\lambda_e F(z_j)\right] > 0 \tag{40}$$

$$b = (1-\pi(\epsilon))\left[1-\lambda_e+\lambda_e F(z_i)\right] > 0 \tag{41}$$

Let $\alpha = \frac{a}{a+b}$, then $\tilde{V} > V$ reduces to

$$\alpha V(\tilde{U}^+) + (1-\alpha)V(\tilde{U}^-) > \alpha V(U^+) + (1-\alpha)V(U^-) \tag{42}$$

By construction, due to requirement (38),

$$\alpha\tilde{U}^+ + (1-\alpha)\tilde{U}^- = \alpha U^+ + (1-\alpha)U^- \tag{43}$$

This can be seen from

$$\alpha\tilde{U}^+ + (1-\alpha)\tilde{U}^- = \alpha U^+ + (1-\alpha)U^- - \alpha\eta^+ - (1-\alpha)\eta^- \tag{44}$$

and

$$-\alpha\eta^+ - (1-\alpha)\eta^- = -\frac{a}{a+b}\eta^+ + \frac{b}{a+b}\eta^- = \frac{a}{a+b}\underbrace{\left[-\eta^+ + \frac{b}{a}\eta^-\right]}_{=0 \text{ by } (38)} = 0 \tag{45}$$

Therefore, if $V(U,z)$ is concave in U, since $U^+ - U^- > \tilde{U}^+ - \tilde{U}^-$ condition (42) holds. Thus, $\tilde{V} > V$ and $U^+ \neq U^-$ cannot be optimal.

Case 2. Define $(z_i, z_j, \tilde{U}^+, \tilde{U}^-)$ as in Case 1. The following relationships hold:

- $z_i \leq z_j$
- $U^*(z_i) < U^- < U^*(z_{i+1})$ and $U^*(z_j) = U^+ < U^*(z_{j+1})$

For sufficiently low (η^+, η^-), it is the case that

$$U^*(z_i) < \tilde{U}^- < U^*(z_{i+1}) \quad \text{and} \quad U^*(z_{j-1}) < \tilde{U}^+ < U^*(z_j) < U^*(z_{j+1}) \tag{46}$$

The requirement that the alternative candidate gives the same utility to the worker implies that

$$\eta^+ = \left(\frac{1-\pi(\epsilon)}{\pi(\epsilon)}\frac{\left[1-\lambda_e+\lambda_e F(z_i)\right]}{\left[1-\lambda_e+\lambda_e F(z_{j-1})\right]}\right)\eta^- \tag{47}$$

The values (η^+, η^-) satisfy conditions (46) and (47).

Let V denote the value of the contract to the firm under the optimal solution, and $\tilde{V}$ the value associated with the alternative candidate. Recall that the incumbent keeps the worker if the outside competitor is not willing to offer a higher lifetime utility than he has promised to deliver. Then

$$\tilde{V} - V = \beta\psi(1-\delta)\Big\{\big[a'V(\tilde{U}^{+}) + bV(\tilde{U}^{-})\big] + \underbrace{\big[V(U^{*}(z_j)I(z > z_j) - V(U^{+})\big]}_{c \geq 0} - \big[a'V(U^{+}) + bV(U^{-})\big]\Big\} \tag{48}$$

where $a' = \pi(\epsilon)\left[1 - \lambda_e + \lambda_e F(z_{j-1})\right] > 0$ and b is defined as in Case 1. Let $\alpha' = \frac{a'}{a'+b}$. Recall that in the present case $U^{*}(z_j) = U^{+}$, hence the term c is non-negative. Then $\tilde{V} > V$ reduces to

$$\alpha' V(\tilde{U}^{+}) + (1-\alpha')V(\tilde{U}^{-}) > \alpha' V(U^{+}) + (1-\alpha')V(U^{-}) \tag{49}$$

By construction

$$\alpha' \tilde{U}^{+} + (1-\alpha')\tilde{U}^{-} = \alpha' U^{+} + (1-\alpha')U^{-} \tag{50}$$

Hence, given that $U^{+} - U^{-} > \tilde{U}^{+} - \tilde{U}^{-}$, condition (49) holds and $U^{+} \neq U^{-}$ cannot be optimal.

Case 3. Here $(\eta^{+}, \eta^{-}, V, \tilde{V})$ are defined as in Case 1. Then $\tilde{V} > V$ and $U^{+} \neq U^{-}$ cannot be optimal.

Case 4. Here $(\eta^{+}, \eta^{-}, V, \tilde{V})$ are defined as in Case 2. Then $\tilde{V} > V$ and $U^{+} \neq U^{-}$ cannot be optimal.

□

A.2 Data

The Survey of Income and Program Participation (SIPP) is a longitudinal survey of representative households in the United States, administered by the U.S. Census Bureau. The survey focuses on collecting data at high frequencies on individuals' income sources and amounts, their labor market status as well as eligibility for and participation in government programs. The SIPP consists of a set of partially overlapping panels, each between two and four years in duration, starting from 1984. Public use data sets from the SIPP published by the U.S. Census Bureau are provided on the homepage of the NBER.[15]

Over the length of one panel, households are interviewed every four month. At each interview, a detailed monthly labor market history (employers, hours, earnings, job characteristics, employment turnover) for each member of the household over the preceding four months is collected, with some variables being recorded even at a weekly frequency. In particular, detailed information for up to two jobs the individual has held over those four months (referred to as the *wave*) are recorded.

The SIPP 2004 panel covers the period from October 2003 to December 2007. Data are released only in core wave files, and longitudinal sampling weights, constructed at the end of data collection, are provided in separate files. We restrict our attention to observations from January 2004 to December 2006 – January 2007 observations are included for transition rates and wage changes between months – for four reasons. First, this avoids the problem of the sample size for the first and the last three months of any SIPP panel being much smaller due to the rotating design of data collection. Second, the sample period stops early enough before the onset of the current financial and economic crisis. Third, there were no changes in the federal minimum wage during the sample period, as such changes occurred in September 1997 and in July 2007. Fourth, there are suitable longitudinal weights available for the waves corresponding to the sample period. We drop all observations for individuals whose entries on person characteristics (gender, age, and race) are inconsistent over time, or who were in the Armed Forces at some point during the panel span. Furthermore, we restrict the sample to male workers between the age of 20 and 65 years who were employed at least in one month during the panel span in a job that is neither self-employment nor family work without pay. Our basic sample contains information on 6,444 individuals for whom we can use the longitudinal weight variable *lgtpnwt3*. The corresponding weighted sample size in each month of our sample period is around 66 million people, and the average age is 40 years.

We use the weekly record of a person's labor market status in the second week of a month, *rwkesr2*, in order to categorize individuals as employed, unemployed or not in

[15] http://www.nber.org/data/sipp.html

the labor force for a given month. Our assignment to labor market states is similar to Nagypal (2008). In particular, a person is taken to be employed if the status is *"with a job or business - working"* or *"with a job or business - not on layoff, absent without pay"*. He is recorded as unemployed if the status is *"with a job or business - on layoff, absent without pay"* or *"without a job or business - looking for work or on layoff"*, and as not in the labor force if it is *"without a job or business - not looking or on layoff"*. Based on this classification, the average participation rate in our sample is 93%, while the average unemployment rate is 3.6%.[16] The average monthly rate of transitions from employment to unemployment is 23.3%, and the rate of transitions from employment to unemployment is 0.8%.

The SIPP survey collects information on up to two wage and salary jobs per person for a given wave. Job-specific information is recorded in the following way: A person's most important job is recorded as job 1 throughout the wave. If an individual held more than one job within the wave, the second most important is recorded as job 2, even if the two jobs were never held simultaneously. Whether a person held a particular job during a given month or week within the wave can be inferred from the recorded starting and ending dates of each job. Taking into account only jobs that are neither unpaid family work nor in the Armed Forces, we determine an individual's main job in a given month by the following sequence of priorities: (i) job 1 if it was held in week 2, (ii) job 2 if it was held in week 2, (iii) job 1 if it was held at some point during the given month, and (iv) job 2 if it was held at some point during the month.

Based on our definitions of a person's labor market status and main job, we construct a measure of the number of job-to-job transitions between two consecutive months which comprises all workers who were employed in the second week of both months, were not unemployed in any of the weeks between, held main jobs with different employer identification numbers in the two months, held each of the main jobs in the second week of the respective month, and did not return to a job previously recorded as their main job. The average fraction of employed workers in our sample who make a job-to-job transition between the current and the following month is 1.32%.[17]

Regarding the estimation of residual wages, we further restrict our sample to workers employed in full-time jobs, that is, jobs for which they report to be usually working 35 hours or more per week. We first impute a worker's real hourly wage at his main job in

[16] These figures compare to an average participation rate of 86% and unemployment rate of 4.6% among all men aged 20 to 64 years reported by the Bureau of Labor Statistics (from CPS data) for the same time period. See Nagypal (2008) for a discussion on differences between the CPS and the SIPP regarding the categorization of a person's labor market status which lead to lower records of unemployment in the SIPP.

[17] Our definition of job-to-job transitions closely corresponds to the one used in Menzio et al. (2012), but their definitions of a person's main jobs and labor market status are different.

a given month, since the SIPP records some crucial variables only once per wave. For around one-half of the observations in our sample a regular hourly wage rate pertaining to the whole wave is reported for the main job. For the remaining observations, we calculate an average hourly wage over the wave from total earnings in this job over four months, the number of hours typically worked in the job, and the total number of weeks employed in the job throughout the wave. For both types of hourly wages, we use the annual CPI from the BLS to express real hourly wages in constant 2004 dollars. Moreover, in accordance with related empirical studies, we exclude wage observations that fall below one-half of the nominal minimum wage rate as well as observations above $211, a value which corresponds to the 99.9th-percentile of real wages in our sample.[18]

Furthermore, we construct variables reflecting a number of worker characteristics. Following Eckstein and Nagypal (2004), we assign individuals to one of five education groups, based on the highest grade or degree obtained. The five categories correspond to high school dropouts, high school graduates, workers with some college education, college graduates, and post-graduate degree holders. For those individuals who report implausible changes in education levels over the panel span (a decrease in education, or a sharp increase that is not associated with temporary school enrolment or with gaps in observations), we impute the number of years of education completed, based on the person's most frequently reported level out of ten finer education categories. We use these finer categories also to calculate a person's potential labor market experience as age minus years of education minus five. Finally, we construct dummy variables for being non-white as well as for all of the four large regions of the United States as classified by the U.S. Census Bureau.

We obtain our estimates of residual wages by running a pooled regression of log real hourly wages on the five broad education groups, a quadratic in potential experience, interaction terms between education groups and the quadratic in experience, the non-white dummy, the region dummies as well as year dummies. The estimated standard deviation of log residuals is 0.48, while the mean and standard deviations of the exponentiated residual wages are 1.13 and 0.72, respectively.

Using the residuals from the above regression, we calculate statistics of workers' log wage changes both within and between jobs. Our observations for wage changes within a job include all workers from the wage sample who held the same main job across two waves of the panel during the sample period. However, since the distribution of log residual wage changes within a job has very long, flat tails on both sides, we exclude estimates below and above three standard deviations from the mean. These cut-offs correspond to decreases by more than 50% and increases by more than 90% in a worker's wage between two

[18]See for example Katz and Autor (1999) for similar sample restrictions.

waves, and imply dropping around 2% of the sample. The mean and standard deviation of the remaining observations are 0.006 and 0.14, respectively, and the fraction of negative changes is 35.5%.[19]

The observations for wage changes between jobs includes all workers from the wage sample who experienced a job-to-job transition as defined above between two consecutive months during our sample period. Again, we eliminate extreme values by trimming the sample at three standard deviations below and above the mean. This step excludes observations with decreases by more than 75% and increases by more than 290%, and reduces the sample by around 2%. The mean of the remaining observations is 0.028, and the standard deviation is 0.36. Moreover, 39.7% of the log wage changes associated with job-to-job transitions are negative.[20] For the subsample of positive wage changes used in the calibration, the mean and standard deviation are 0.28 and 0.26, respectively.

[19] For comparison, before dropping outliers the mean of log wage changes within a job is 0.007, the standard deviation is 0.22, and the fraction of negative changes is 35.8%.

[20] Before dropping outliers the mean of log wage changes between jobs is 0.023, the standard deviation is 0.45, and the fraction of negative changes is 40%.