A THEORY OF WAGE AND PROMOTION DYNAMICS IN INTERNAL LABOR MARKETS

Robert Gibbons
Michael Waldman

Working Paper **6454**

NBER WORKING PAPER SERIES

A THEORY OF WAGE AND PROMOTION DYNAMICS IN INTERNAL LABOR MARKETS

Robert Gibbons
Michael Waldman

Working Paper 6454
http://www.nber.org/papers/w6454

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 1998

We thank George Baker, Bengt Holmstrom, Caroline Minter, and seminar audiences at Chinese University of Hong Kong, Harvard University, Hong Kong University of Science and Technology, Massachusetts Institute of Technology, UC Santa Cruz, University of Chicago, and University of Toronto for helpful conversations. Both authors thank Cornell's Johnson School for financial support. Gibbons also thanks the Center for Advanced Study in the Behavioral Sciences for a fellowship funded in part by NSF grant SBR-9022192. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

A Theory of Wage and Promotion Dynamics
in Internal Labor Markets
Robert Gibbons and Michael Waldman
NBER Working Paper No. 6454
March 1998

ABSTRACT

We attempt to explain employment practices in internal labor markets using models that combine job assignment, on-the-job human-capital acquisition, and learning. We show that a framework that integrates these familiar ideas captures a number of recent empirical findings concerning wage and promotion dynamics in internal labor markets, including the following. First, real wage decreases are a minority of the observations, but are not rare, while demotions are very rare. Second, there is significant serial correlation in wage increases. Third, promotions are associated with particularly large wage increases, but these wage increases are small relative to the difference between average wages across levels of a job ladder. Fourth, on average, workers who receive large wage increases early in their stay at one level of a job ladder are promoted more quickly to the next level. Fifth, individuals promoted from one level of a job ladder to the next come disproportionately, but not exclusively, from the top of the lower job's wage distribution (and arrive disproportionately, but not exclusively, at the bottom of the higher job's wage distribution).

Robert Gibbons
Sloan School of Management
Massachusetts Institute of Technology
50 Memorial Drive
Cambridge, MA 02142
and NBER
rgibbons@mit.edu

Michael Waldman
Graduate School of Business
University of Chicago
1101 East 58th Street
Chicago, IL 60637
mw46@cornell.edu

I. INTRODUCTION

A recurrent theme in the labor economics literature is that the workings of internal labor markets deviate from the predictions of the standard theory of competitive labor markets. Doeringer and Piore (1971), for example, argue that wage rates are often more closely associated with job assignments than with workers' human-capital characteristics. More recently, Medoff and Abraham (1980, 1981) and Baker, Jensen, and Murphy (1988) have also presented evidence on a variety of practices that seem at odds with the standard competitive model. Continuing in this vein, recent papers by Lazear (1992) and Baker, Gibbs, and Holmstrom (1994a,b) provide detailed empirical analyses of particular internal labor markets.

In response to evidence presented by these and other authors (discussed in more detail below), a literature has developed that attempts to provide a theoretical foundation for several common features of internal labor markets. We focus on three principal approaches from this literature: job assignment, on-the-job human-capital acquisition, and learning. Our paper integrates these three theoretical approaches in a natural fashion, in an attempt to explain the main findings of the recent empirical literature. We begin by giving brief descriptions of these three theoretical approaches.

By the job-assignment literature we mean papers that investigate the assignment of workers to jobs when firms consist of a variety of potential job assignments and there is full information about workers and jobs. The models in this literature are typically static and focus on the roles of comparative advantage and the "scale-of-operations effect" (i.e., the idea that higher ability workers should be assigned to jobs where decisions have an impact over a larger scale of operations) in determining the equilibrium assignment. For example, Sattinger (1975) considers a model where comparative advantage determines which workers get the better jobs, while Rosen (1982) and Waldman (1984a) focus on the role of the scale-of-operations effect in assigning workers to different levels of a job ladder. This literature has produced a number of insights concerning the wage distribution both within and across firms; see Sattinger (1993). Because of the static nature of the analyses, however, these models do not provide direct explanations for wage and promotion dynamics.

A second important perspective on the workings of internal labor markets involves on-the-job human-capital acquisition. Since Becker (1964), an extensive literature has developed that investigates the implications of on-the-job human-capital acquisition for age-earnings profiles. Examples include Ben-Porath's (1967) analysis of how the incentive to invest in human capital varies over a worker's career and Hashimoto's (1981) analysis of the financing decision. More recently, a number of authors have begun to combine the human-capital perspective with issues of promotion and (to some extent) job assignment. For example, both Carmichael (1983) and Prendergast (1993) show how promotion to higher paying jobs can be used to provide incentives for efficient human-capital acquisition, while Kahn and Huberman (1988) explore the role of Up-or-Out contracts in providing such incentives.

A third literature has recently developed that investigates the role of learning in (internal and external) labor markets. In this approach, firms are uncertain about a worker's ability when the worker enters the labor force, but gradually learn about the worker's ability during his or her career. Papers that take this approach typically fall into one of two categories. One set of papers assumes that learning is symmetric -- that is, any information generated about a worker's ability during the worker's career is public information. Examples include Harris and Holmstrom's (1982) analysis of insurance, Holmstrom's (1982) model of career concerns, and Farber and Gibbons's (1996) investigation of wage dynamics. The other category is asymmetric learning. These papers assume that a worker's current employer receives better information about the worker's ability than do other potential employers. Examples include Greenwald's (1986) model of adverse selection in the labor market, the analyses of Waldman (1984b, 1990), Ricart i Costa (1988), and Bernhardt and Scoones (1993), in which promotions serve as a signal of ability, and Gibbons and Katz's (1991) application of the adverse-selection and signaling approaches to data on the consequences of layoffs versus plant closings.

In this paper we develop a model that integrates job assignment, on-the-job human-capital acquisition, and learning.[1] We show that a framework that integrates these familiar ideas captures a number of recent empirical findings concerning wage and promotion dynamics in internal labor markets, including but not limited to those of Baker, Gibbs and Holmstrom (hereafter, BGH).[2] In their study, BGH consider the ability of a variety of theoretical approaches to explain their empirical findings. Considering the human-capital-acquisition and learning approaches separately, they conclude that their empirical findings allow them to reject both approaches as explanations for wage and promotion practices in their internal labor market. Our conclusion, in contrast, is that a model that combines a number of existing approaches explains many of BGH's findings. Hence, rather than ruling out existing theoretical approaches to wage and promotion practices in internal labor markets, our interpretation is that the BGH (and other) findings support many of these approaches -- although in a more integrated form than typically appears in the literature.

This paper contributes to a growing literature that attempts to provide a theoretical explanation for a broad pattern of evidence rather than the more standard approach of developing a model that is focused on a single finding. Previous papers in this "broad pattern" spirit include Harris and Holmstrom's (1982) analysis of insurance, MacLeod and Malcomson's (1988) investigation of reputation in a model characterized by both adverse selection and moral hazard, and Demougin and Siow's (1994) model of on-the-job training and screening. The paper closest to ours is Bernhardt (1995), which also considers a model that combines job assignment, on-the-

[1] A number of the papers referred to above combine two or more of these elements. For example, Prendergast (1993) considers both job assignment and human-capital acquisition, while Waldman (1984b) and Bernhardt and Scoones (1993) capture aspects of all three elements. Other papers that combine two or more of the modeling elements we consider include Bernhardt (1995) and Jovanovic and Nyarko (1996). Of these papers, only Bernhardt (1995) is similar to ours in the sense of combining these elements in order to explain a "broad pattern" of empirical findings.

[2]Doeringer and Piore originally utilized the term "internal labor market" to refer to the human-resource practices of a firm (or part of a firm) that employs ports of entry, promotion from within, wages attached to jobs, and job ladders. Their analysis focused mostly on blue-collar settings. Following the more recent literature, we employ the term to refer to the human-resource practices of a firm (or part of a firm) where a significant proportion of the workforce exhibits long-term attachments to the firm. In the current usage, the term can be used to describe either blue- or white-collar settings. Thus, it would be equally accurate to say that we focus on "careers in organizations" as on "internal labor markets."

job human-capital acquisition, and learning. Our theoretical model focuses on symmetric learning while Bernhardt considers the asymmetric-learning case. More importantly, empirical evidence in this area has progressed rapidly and has developed a more detailed account of wage and promotion dynamics. Our paper speaks to many of these new more detailed findings.

The outline for the paper is as follows. Section II reviews the recent empirical evidence on wage and promotion dynamics. Section III combines job assignment and on-the-job human-capital acquisition in a model with a three-level job ladder and full information. Section IV analyzes this model under the assumption of symmetric learning. Section V relates our framework to Medoff and Abraham's findings concerning performance evaluations, and also discusses findings from the recent empirical literature that are inconsistent with the predictions of our symmetric-learning model. Section VI presents concluding remarks.

II. EVIDENCE ON WAGE AND PROMOTION DYNAMICS

The main focus of Sections III and IV is the extent to which our theoretical framework explains the findings of BGH; in Section V.B we discuss the major BGH findings we do not capture. We proceed in this fashion not because BGH is the only empirical study that considers wage and promotion dynamics in internal labor markets. Rather, because various studies have investigated different issues in different environments, it is difficult to assess what combinations of facts are true in any given environment. We therefore focus on a single study, choosing BGH because it is the most comprehensive. The fact that our framework can explain many of the BGH findings is of more interest if these findings are representative of other authors' findings concerning wage and promotion dynamics. Therefore, after describing the BGH findings, we briefly discuss the extent to which their findings are representative of the larger literature. See Gibbons (1997) for a more comprehensive review of this literature.

BGH analyze personnel data from a single firm over a twenty-year time period. An observation in their dataset consists of a managerial employee in a specific year. The information they have for that manager includes the employee's ID number, job title, salary, and performance rating. The first step of their analysis was to construct the firm's internal job ladder, which they do by recording how workers move across job titles, rather than by relying on salary data. By constructing the job ladder in this way, they are able to test for a number of relationships between salary and job level without worrying that the results are an artifact of the way the job ladder was constructed.

The BGH findings we concentrate on are the following. First, real wage decreases are a minority of the observations, but are not rare, while demotions are very rare. Second, there is significant serial correlation in both wage increases and promotion rates (the latter is sometimes referred to as a "fast track"). Third, promotions are associated with particularly large wage increases, but these wage increases are small relative to the difference between average wages across levels of a job ladder. Fourth, on average, workers who receive large wage increases early in their stay at one level of a job ladder are promoted more quickly to the next level. Fifth, individuals promoted from one level of a job ladder to the next come disproportionately, but not exclusively, from the top of the lower job's wage distribution (and arrive disproportionately, but not exclusively, at the bottom of the higher job's wage distribution).

There are a number of other studies that address one or more of these findings. Some of these studies such as Rosenbaum (1984) and Lazear (1992) are similar to BGH in that the analyses are based on the personnel records of one or a small number of firms, while other studies such as Murphy (1985), Topel (1991), and McKue (1996) employ data from a large number of firms.

Two findings with strong support in the literature are that real wage decreases are not rare, and that there is serial correlation in promotion rates. For example, McLaughlin (1994) and Card and Hyslop (1995) both find a significant frequency of real wage decreases, while there is extensive support for serial correlation in promotion rates (e.g., Rosenbaum (1984), Bruderl et al. (1991), and Podolny and Baron (1997)). Further, although we know of no study besides BGH that documents this finding, we take it to be uncontroversial that demotions are quite rare.

The evidence concerning serial correlation in wage increases is less clear cut. Lillard and Weiss (1979) and Hause (1980) find evidence of serially correlated wage changes (or wage residual changes), while Abowd and Card (1989), Topel (1991), and Topel and Ward (1992) do not. One possible explanation relies on differences in the samples studied. BGH, Lillard and Weiss, and Hause all study relatively homogeneous samples (BGH managers in a single firm, Lillard and Weiss American scientists, and Hause young Swedish males), while the other studies investigate large heterogeneous samples such as the Panel Study of Income Dynamics. Perhaps only certain small groups of workers such as managerial and professional workers exhibit such serial correlation. If most groups of workers do not, then the representative cross-sections in Abowd and Card, Topel, and Topel and Ward would not either.

Another finding with strong support in the literature is that promotions are associated with large wage increases. This finding is supported by the studies of Gerhart and Milkovich (1989), Lazear (1992), and McCue (1996), among others. There is also some support for the notion that these wage increases, although large relative to wage increases associated with not being promoted, are small relative to the difference between average wages across levels of a job ladder. Murphy (1985) documents both of these facts. In a study of top executives in 72 large U.S. manufacturing firms, he finds that the average real increase in salary plus bonus for the whole sample was 3.7% but the average increase for a Vice President promoted to President was 20.9%. On the other hand, the average salary plus bonus for Presidents was 60% higher than for Vice Presidents. Main et al. (1993) present similar results.

In contrast to what we described above for the first three BGH findings we focus on, there is little in the literature concerning the remaining two findings. The fourth BGH finding we consider is that, on average, workers who receive large wage increases early in their stay at one level of a job ladder are promoted more quickly to the next level. As far as we know, no other study directly addresses this issue. However, McCue does find something similar -- that is, a high wage today is positively correlated with promotion tomorrow. Although this is not exactly the BGH finding, it suggests that the BGH finding might hold in McCue's dataset.

The fifth BGH finding we consider is that individuals promoted from one level of a job ladder to the next come disproportionately, but not exclusively, from the top of the lower job's wage distribution (and arrive disproportionatly, but not exclusively, at the bottom of the higher job's wage distribution). As for the fourth finding discussed above, we know of no other study that directly addresses this question.

In summary, many of the BGH findings we focus on have broad support in the literature, while all but one of the others have at least weak support. Finally, the remaining BGH finding has simply not been addressed elsewhere.

III. AN ANALYSIS OF JOB LADDERS UNDER FULL INFORMATION

This section develops a model in which firms consist of three-level job ladders. Job-assignment and human-capital considerations determine wage and promotion dynamics. As a benchmark, in this section we analyze this model under full information; the following section provides an analysis under symmetric learning. The technology of production we consider throughout the paper is similar to that investigated in Sattinger (1975), Rosen (1982), and Waldman (1984a): jobs are ranked in terms of the extra value produced by a worker of greater ability. Sattinger, Rosen, and Waldman consider this technology in one-period models characterized by full information and no human-capital acquisition, while our focus is on multiple periods, human-capital acquisition, and (in later sections) learning.

The analysis in this section demonstrates that a dynamic model that combines job assignment and on-the-job human-capital acquisition provides a reasonable starting point for explaining wage and promotion dynamics in internal labor markets. That is, the model captures a number of the findings discussed in Section II. On the other hand, some of the findings are not consistent with the full-information model. We show in Section IV that the model captures these additional findings when learning is added to the framework.

A) The Model

There is free entry into production. All firms are identical and the only input is labor. A worker's career lasts for T periods, $T \geq 5$ (an unusual assumption which we motivate below). In each period, labor supply is fixed at one unit for each worker. Worker i's *innate ability* is denoted θ_i and can be either high or low: $\theta_i \in \{\theta_H, \theta_L\}$. A worker's *effective ability* is a function of the worker's innate ability and the worker's labor-market experience. Let t denote calendar time, and x_{it} the worker's labor-market experience prior to period t (i.e., for a worker in his or her first period in the labor market, prior experience x_{it} equals zero). We assume that worker i's effective ability in period t is given by

$$(1) \qquad \eta_{it} = \theta_i f(x_{it}),$$

where $f'>0$ and $f''\leq 0$.

A firm consists of three different jobs, denoted 1, 2, and 3. If worker i is assigned to job j in period t, then the worker produces

$$(2) \qquad y_{ijt} = d_j + c_j(\eta_{it} + \varepsilon_{ijt}),$$

where d_j and c_j are constants known to all labor-market participants and ε_{ijt} is a noise term drawn from a normal distribution with mean 0 and variance σ^2. Let η' (η'') denote the effective ability level at which a worker is equally productive at jobs 1 and 2 (2 and 3). That is, η' solves $d_1+c_1\eta=d_2+c_2\eta$ and η'' solves $d_2+c_2\eta=d_3+c_3\eta$. We assume that $c_3>c_2>c_1$ and $d_3<d_2<d_1$, and that these parameters are such that $\eta''>\eta'$. Thus, given full information about worker abilities, the efficient assignment rule for period t is to assign worker i to job 1 if $\eta_{it}<\eta'$, to job 2 if $\eta'<\eta_{it}<\eta''$, and to job 3 if $\eta_{it}>\eta''$.

Workers and firms are risk-neutral and have a discount rate of zero. There is no cost to workers from changing firms or to firms from hiring or firing workers. Under these assumptions, there are no benefits to long-term contracts, so we assume that wages are determined by spot-market contracting. Finally, to ease the comparison of the model with the empirical evidence, we restrict attention to wages that are paid in advance of production, as opposed to one-period piece-rate contracts.

At the beginning of each period, all firms simultaneously offer each worker a wage for that period. The worker then works for the firm that offers the highest wage. If there are multiple firms tied at the highest wage, the worker chooses randomly among these firms unless one of these was the worker's employer in the previous period, in which case the worker remains with that firm. This tie-breaking rule is equivalent to assuming a moving cost that is infinitesimally small; as a result, in the full-information case considered in this section and the symmetric-learning analysis of Section IV, there is no turnover in equilibrium.

To reduce the number of cases that need to be considered, we restrict the analysis to parameterizations that satisfy the following conditions. First, $\theta_H f(1)<\eta'$, so that it is efficient for each worker to be assigned to job 1 in the first two periods of the worker's career. (Recall that x_{it} measures *prior* labor market experience, so $x=1$ in the worker's second period in the labor market.) Second, $\eta''-\eta'>\theta_H[f(3)-f(1)]$, so that it is efficient for each worker to be at job 2 for at least two periods before being promoted to job 3 (since $\theta_H>\theta_L$ and $f(3)-f(1)\geq f(x+2)-f(x)$ for all $x\geq 1$). Third, $\theta_L f(T-1)>\eta''$, so that it is efficient for each worker to be on job 3 by the end of the worker's career. For these three conditions to hold simultaneously requires $T\geq 5$.

B) <u>The Full-Information Benchmark</u>

We now analyze the benchmark case of full information. That is, each worker's innate ability (θ_i) is common knowledge at the beginning of the worker's career. If innate ability is common knowledge then in each period a worker is assigned to the job that maximizes the worker's expected output and paid a wage equal to that expected output.[3] Proposition 1 characterizes the

[3]Throughout the analysis we assume that if a firm is indifferent between assigning a worker to either of two jobs then the firm assigns the worker to the higher-level job. This assumption is not crucial for the analysis but simplifies the statements of some results.

assignment rule in this case. We use w_{it} to denote the wage paid to worker i in period t in equilibrium. Proofs are given in the Appendix.

Proposition 1: Suppose each worker's innate ability is common knowledge at the beginning of the worker's career. Then job assignments and wages are given by i)-iii) below.

i) If $\eta_{it}<\eta'$ then worker i is assigned to job 1 in period t and earns the wage $w_{it}=d_1+c_1\eta_{it}$.

ii) If $\eta'\leq\eta_{it}<\eta''$ then worker i is assigned to job 2 in period t and earns the wage $w_{it}=d_2+c_2\eta_{it}$.

iii) If $\eta_{it}\geq\eta''$ then worker i is assigned to job 3 in period t and earns the wage $w_{it}=d_3+c_3\eta_{it}$.

Proposition 1 says that in equilibrium there is a job ladder that workers climb as they gain labor-market experience. That is, the equilibrium job assignments described in Proposition 1 can be restated in terms of the worker's labor-market experience rather than the worker's effective ability: for each θ_i (i=H,L) there exist values x_i' and x_i'', where $x_i''>x_i'$, such that if a worker's innate ability is θ_i then in period t the worker is assigned to job 1 if $x_{it}<x_i'$, to job 2 if $x_i'\leq x_{it}<x_i''$, and to job 3 if $x_{it}\geq x_i''$. There are no demotions in equilibrium because effective ability increases monotonically as a worker gains labor-market experience, so it is never optimal to promote and then demote a worker. For the same reason, a worker's wage rises every period.

The fact that there are no demotions in equilibrium is roughly consistent with the results of BGH, who found that 0.3% of the year-to-year job movements in their sample were demotions. In contrast, BGH found that approximately 25% of the year-to-year salary changes in their sample were real-wage decreases. Certain years have pronounced effects, however: in each of the high-inflation years of 1979 and 1980, the median real-salary increase was negative at the firm BGH investigated. But even in a typical year, when the average real-salary increase was between 5% and 9%, the fraction of year-to-year salary changes that were real-wage decreases never fell below 5%, and across the nineteen years of their study, the median value for this fraction was 12%.

BGH's evidence on demotions and (especially) real wage decreases does not match the results of our full-information analysis. In contrast, the full-information case does quite well at capturing two of the other findings discussed in Section II. First, the full-information case captures serial correlation in both wage increases and promotion rates. Consider workers with labor-market experience x. The wage increase for a θ_H-worker whose experience increases from x to x+1 will be larger than the corresponding wage increase for a θ_L-worker, and the same ordering will hold for wage increases associated with experience increasing from x+1 to x+2. The logic here is that, because the primary determinant of wage growth is the growth in effective ability, there is serial correlation in wage increases because high-ability workers have faster growth in effective ability at all experience levels.[4]

The model captures serial correlation in promotion rates in that, if η' and η''-η' are both sufficiently large, then high-ability workers are promoted to job 2 more quickly, and also spend less time on job 2 before being promoted to job 3. That is, $x_H'<x_L'$ and $x_H''-x_H'<x_L''-x_L'$.[5] The logic for this result is simple. A worker receives a promotion when effective ability reaches certain absolute levels. Since high-ability workers experience faster growth in effective ability, these workers are promoted to job 2 earlier in their career and are also promoted to job 3 after having spent less time on job 2.

The second finding captured by the full-information case is that wage increases predict promotion. To see this, consider workers with labor-market experience x, where $x<x_H'$. Proposition 1 implies that the wage increase for a θ_H-worker whose experience increases from x-1 to x is $c_1\theta_H[f(x)-f(x-1)]$, while the analogous wage increase for a θ_L-worker is $c_1\theta_L[f(x)-f(x-1)]$. In words, prior to any promotions taking place, high-ability workers receive larger wage increases

---

[4]Serial correlation in wage increases could arise from a concave age-earnings profile, but BGH also find serial correlation in wage-residual changes for white males after controlling for age, education, tenure, and salary. Our model also predicts serial correlation in such wage residual changes.

[5] If x were continuous, then $x_H'<x_L'$ and $x_H''-x_H'<x_L''-x_L'$ would hold for all parameterizations.

than do low-ability workers. In turn, given that high-ability workers are also promoted to job 2 earlier in their careers, we have that wage increases predict promotion.

One of the other findings discussed in Section II concerns large wage increases upon promotion. This finding is captured by the full-information case in a weak form. To see why we say this, consider the average wage increase received by workers who are promoted to job 2. If f'' is close to zero (for all x) then the average wage increase received by workers promoted to job 2 is larger than the average wage increase received by workers who remain in job 1, because increases in effective ability are valued at rate c_1 in job 1 but at rate $c_2 > c_1$ in job 2. The wage increase at promotion is the sum of two parts: c_1 times the worker's increase in effective ability, plus the increased value from assigning a worker with the new effective ability to job 2 rather than job 1. Therefore, the average wage increase at promotion to job 2 will be larger than the average wage increase received by workers who remain on job 1.

The reason we say the full-information case captures the empirical finding of a large wage increase upon promotion in a weak form is that we do not feel that the effect just described is by itself a plausible explanation for the findings of BGH. According to the argument just given, the average wage increase the year after a promotion to job 2 should be larger than the average increase at promotion to job 2 (because in the year after the promotion to job 2 increases in effective ability are valued at rate c_2 rather than at a convex combination of c_1 and c_2). But BGH find that for workers promoted from level 1 to level 2, the average wage increase the year after the promotion is strictly less than the average wage increase at promotion.[6] In Section IV we show that when learning is added to our model this problem disappears: wage increases at promotion can be larger than wage increases both before and after promotion.

The final finding that holds in the full-information case is that, although wage increases upon promotion are large (in a weak sense), they explain only a fraction of the difference between average wages across levels of a job ladder. Consider the difference between average wages at

[6]BGH do not report results that allow us to make this comparison for promotion to any other level.

levels 1 and 2. Because workers at level 2 are on average older than workers at level 1, an important component of this difference is the increased productivity due to human-capital accumulation corresponding to the higher average age at level 2. Since the average wage increase at promotion will capture only one year of this age difference, the difference between average wages at levels 1 and 2 is bigger than the average wage increase at promotion to level 2. In fact, if the difference in average ages across levels is much greater than one, then the average wage increase at promotion can be very small relative to the difference between average wages across levels.

Overall, the analysis of this section yields a number of results that suggest that the full-information case is a reasonable starting point for considering wage and promotion dynamics in internal labor markets. First, consistent with the findings of BGH, there is serial correlation in both wage increases and promotion rates. Second, consistent with BGH, workers who receive large wage increases early in their stay at one level of the job ladder are promoted more quickly to the next level. Third, again consistent with BGH, wage increases upon promotion explain only a fraction of the difference between average wages across levels of a job ladder. But the model does not capture a number of the findings discussed in Section II. For example, the model predicts no demotions and no wage decreases, while BGH find a positive frequency of both demotions and wage decreases, although demotions were quite rare. In addition, the full-information case does not adequately capture the evidence concerning the size of wage increases upon promotion.

In the next section we show that the performance of the model is improved substantially when learning is added. In particular, the addition of learning leaves unchanged most of the predictions of the full-information analysis that match the BGH findings, while changing a number of the other predictions so that they better match the evidence.

IV. AN ANALYSIS OF JOB LADDERS UNDER SYMMETRIC LEARNING

In this section we depart from the assumption of full information. In particular, when a worker enters the labor force, firms are now uncertain about the worker's ability and learn about it only gradually as the worker's career progresses. In this section we consider the case of symmetric learning: at any point in time all firms (and the worker in question) are equally informed about the worker's ability level. In the Conclusion we briefly discuss the case of asymmetric learning, when a worker's current employer has superior information.

Our analysis in this section extends those of Murphy (1986) and Gibbons and Katz (1992), both of which study symmetric learning and assignment to a job ladder, but neither of which allows human-capital accumulation. (See Ross, Taubman, and Wachter (1981) and MacDonald (1982) for related models without clear rankings of jobs.) Because our model also incorporates human-capital acquisition, we capture the idea that workers move up a job ladder as they age.

A) The Model and Preliminary Results

The only difference between this model and the full-information model in Section III is that now there is symmetric learning. At the beginning of a worker's career the worker is known to be of innate ability θ_H with probability p_0 and of innate ability θ_L with probability $(1-p_0)$. Learning takes place at the end of each period when the realization of the worker's output for that period becomes common knowledge. The presence of the noise term ε_{ijt} in the production function (2) implies that learning occurs gradually.

Define $z_{it}=(y_{ijt}-d_j)/c_j=\eta_{it}+\varepsilon_{ijt}$. That is, z_{it} is the signal about the worker's effective ability that the market extracts from observing the worker's output in period t. We refer to z_{it} as the worker's normalized output from period t (i.e., normalized to abstract from job assignment) and to $(z_{it-x},\ldots,z_{it-1})$ as the worker's normalized output history over his x periods of prior labor market experience. Because the signal z_{it} is independent of job assignment, there is no difference in the rate of learning across jobs. Let θ^e_{it} denote the expected innate ability of worker i in period t: $\theta^e_{it}=E(\theta_i \mid z_{it-x},\ldots,z_{it-1})$.

From θ^e_{it} we can compute the expected effective ability of worker i in period t:

(3) $$\eta^e_{it}=\theta^e_{it}f(x_{it}).$$

Given η^e_{it}, wage determination and job assignment proceed as in the full-information case: in each period a worker is assigned to the job that maximizes the worker's expected output and paid a wage equal to that expected output.

Proposition 2: Suppose that at the beginning of a worker's career the worker is known to be of innate ability θ_H with probability p_0 and of innate ability θ_L with probability $(1-p_0)$. Suppose also that learning is symmetric. Then job assignments and wages are given by i)-iii) below.

i) If $\eta^e_{it}<\eta'$ then worker i is assigned to job 1 in period t and $w_{it}=d_1+c_1\eta^e_{it}$.

ii) If $\eta'\leq\eta^e_{it}<\eta''$ then worker i is assigned to job 2 in period t and $w_{it}=d_2+c_2\eta^e_{it}$.

iii) If $\eta^e_{it}\geq\eta''$ then worker i is assigned to job 3 in period t and $w_{it}=d_3+c_3\eta^e_{it}$.

Proposition 2 says that wages and job assignments are now determined by expected effective ability, whereas in the full-information case the actual value of effective ability determined wages and job assignments. (Only the expected effective ability matters, rather than other moments of the distribution, because output is a linear function of effective ability on each job.) In the next sub-section we show that, despite the similarity between Propositions 1 and 2, the symmetric-learning case does a much better job of matching the empirical evidence.

B) Further Analysis

The predictions of the full-information case match some but not all aspects of the BGH evidence. For the analysis of symmetric learning, we begin with the predictions of the full-information model that were shown in Section III to match the evidence well.

In the full-information model, and in the data, there is serial correlation in both wage increases and promotion rates. As shown in Corollary 2.1, for workers in job 1 in periods t and t+1, serial correlation continues to hold in the symmetric-learning case for the wage changes at t+1 and t+2.

Corollary 2.1: Consider a worker in period t who is currently in job 1 earning the wage w_{it}. If the worker is again in job 1 in period t+1, then the conditional expectation of $w_{it+2}-w_{it+1}$ is an increasing function of $w_{it+1}-w_{it}$.

The logic for why there is serial correlation in wage increases in the situation considered in Corollary 2.1 is based on the logic in the full-information case. With symmetric learning, wage increases are determined by the growth in expected effective ability, which in turn is positively related to the change in expected innate ability. At a given experience level, therefore, a large wage increase indicates that an increase in expected innate ability occurred, which means that on average the worker's expected effective ability will grow more quickly in the future (through human-capital acquisition). In turn, there is serial correlation in wage increases because future wage increases are determined by the growth in expected effective ability, so large wage increases now are positively correlated with large wage increases in the future.

Corollary 2.1 does not apply to workers who start on job 1 but whose first wage increase is associated with a promotion, or to workers who start on job 2 or job 3. We can generalize the result if we restrict the analysis to parameterizations for which there is no possibility of demotion.[7]

[7] The reason the possibility of demotion is important is as follows. In the symmetric-learning case, the expected size of next period's wage increase is higher if the worker's current value for expected effective ability is close to either η' or η''. To see why, consider the case where the worker's value for expected effective ability equals exactly η'. In this case, increases in expected effective ability are valued at rate c_2 while decreases (which correspond to demotions) are valued at rate c_1. This asymmetry, which is not a factor in the full-information model (because there workers never experience decreases in expected effective ability), means that the expected size of next period's wage increase is in fact larger than c_2 multiplied by the expected increase in expected effective ability. This feature of the symmetric-learning model is the reason that Corollary 2.1 does not generalize to all cases when there is a positive probability of demotion in either period t+1 or t+2. For example, consider a worker in period t who has x periods of labor market experience and is currently on job 1 earning the wage w_{it}. If the worker is on job 2 in period t+1 and it is possible that the worker will be demoted for period t+2 then, since a higher value for w_{it+1} means η^e_{it+1} is further from η', we have that the conditional expectation of $w_{it+2}-w_{it+1}$ can in fact be a decreasing function of $w_{it+1}-w_{it}$. In contrast, if the worker is on job 2 in period t+1 but there is a zero probability of

That is, if, for any x, $\theta_H f(x)>\eta'$ implies $\theta_L f(x+1)>\eta'$ and $\theta_H f(x)>\eta''$ implies $\theta_L f(x+1)>\eta''$, then for a worker in job j in period t, $w_{it+2}-w_{it+1}$ is an increasing function of $w_{it+1}-w_{it}$. We also impose a no-demotion assumption in deriving another result below. Recall that the no-demotion case is empirically relevant: in the BGH dataset only 0.3% of year-to-year job movements were demotions.

Now consider promotion rates. In the full-information model we found that, if η' and $\eta''-\eta'$ are both sufficiently large, then there is serial correlation in promotion rates. As opposed to the other results that hold in the full-information case, under symmetric learning there is no clear prediction regarding serial correlation in promotion rates. We can show a related but weaker result: given two workers earning the same wage just after promotion to job 2, the worker with less labor-market experience will on average be promoted to job 3 at a younger age. We also conjecture that for many parameterizations of the model there will be serial correlation in promotion rates (i.e., given the same two workers, the one with less experience when promoted to job 2 will on average be promoted to job 3 after less time in job 2). The reason serial correlation is not a general result in our model is that, no matter how informative a worker's history of past outputs is, an extreme value of the next output observation can move the market's belief about the worker's innate ability arbitrarily far. One reason such learning causes problems is that the older (and seemingly less promising) worker of our pair could become a star and earn promotion immediately while the younger (and seemingly more promising) worker could be too young to earn promotion immediately even if he or she achieves the best possible output.

Another prediction of the full-information model that matches the BGH evidence is the prediction that wage increases predict promotion. This prediction continues to hold in the symmetric-learning case in the following two ways. First, workers who receive higher wage increases after their first year of employment are more likely to be promoted at the first experience level at which promotion is possible (i.e., at experience level x_H' as defined in Section

---

demotion for period t+2, then having η^e_{it+1} be close to η' would not be a factor and the conditional expectation of $w_{it+2}-w_{it+1}$ would be an increasing function of $w_{it+1}-w_{it}$.

III.B). Second, BGH find that workers who receive higher average wage increases after their first year of employment are on average promoted to job 2 earlier in their careers. Our symmetric-learning model yields exactly this result for parameterizations such that there are no demotions from job 2 to job 1.

Corollary 2.2: Consider a worker in period t who has zero periods of labor-market experience and therefore is in job 1 and assume a parameterization such that $x_H' < x_L'$.[8] (i) The probability of promotion at experience level x_H' is an increasing function of $w_{it+1}-w_{it}$. (ii) If demotion from job 2 to job 1 is impossible (i.e., for any x, $\theta_H f(x) > \eta'$ implies $\theta_L f(x+1) > \eta'$), then the expected value of the worker's labor market experience when the worker is promoted to job 2 is a decreasing function of $w_{it+1}-w_{it}$.

Corollary 2.2 refers to promotion from job 1 to job 2. BGH also find that workers who receive higher average wage increases after their first year in job 2 spend less time in job 2 before promotion to job 3. The symmetric-learning model also yields a result consistent with this finding for parameterizations such that there are no demotions from job 3 to job 2. Consider a worker in period t who has x periods of labor-market experience, has just been promoted to job 2, and is currently earning the wage w_{it}. If the worker is again in job 2 in period t+1, then the expected value of the worker's labor-market experience when the worker is promoted to job 3 is a decreasing function of $w_{it+1}-w_{it}$.

The logic for why wage increases predict promotion is closely related to the above discussion concerning serial correlation in wage increases. As before, a large wage increase indicates that an increase in expected innate ability occurred, which means that on average the worker's expected effective ability will grow more quickly in the future. Since promotions occur when expected

[8] If $x_H' = x_L'$, then the model yields the degenerate result that all promotions occur at the same value for labor-market experience. Also, under symmetric learning it is possible that the worker will be promoted directly from job 1 to job 3. When this is the case, we define the value of the worker's labor market experience upon promotion to job 2 to be the same as the value upon promotion to job 3.

effective ability reaches certain absolute levels, on average a worker who experiences a large wage increase will need less time to reach the target level of expected effective ability needed for promotion.

We now consider predictions of the full-information case that were shown in Section III to be inconsistent with the BGH evidence. First, the full-information case predicts no demotions and no wage decreases, while the evidence indicates a positive frequency of both, although demotions tend to be quite rare. As opposed to the full-information case, the symmetric-learning case yields predictions concerning demotions and wage decreases that are more consistent with the evidence.

<u>Corollary 2.3</u>: Let x^* be the smallest x such that $\theta_L f(x+1)<\theta_H f(x)$. Consider a worker in period t who has x periods of labor market experience. For $x<x^*$ there are no wage decreases and no demotions (i.e., $w_{it+1}-w_{it}\geq 0$ and the worker is not in a lower level job in t+1 than he or she was in t). For $x\geq x^*$ there exists a positive frequency of wage decreases but fewer (possibly no) demotions. Any demotions that occur are associated with wage decreases.

In the full-information case, a worker's effective ability increases each period, so in equilibrium there are no wage decreases. With symmetric learning, wage decreases occur because a worker's expected effective ability does not necessarily increase every period: because of the learning process, a worker's expected innate ability can fall substantially from one period to the next; if this decrease is sufficiently large it will dominate the increase in expected effective ability due to human-capital accumulation.

Our argument explains not only a positive frequency of wage decreases, but also why wage decreases are a minority of the observations, although not rare, while demotions are very rare. Roughly half the workers will experience a decrease in expected innate ability, so if human-capital acquisition is positive but small then many of these will also experience a decrease in expected effective ability. The result is that wage decreases are a minority of the observations, but may not be rare. To see why demotions might be very rare or nonexistent, consider all workers assigned to job 2 in period t and suppose the parameterization is such that $\eta''-\eta'$ is large. All workers for whom expected effective ability falls will experience a wage decrease, but only those for whom the new value of expected effective ability is below η' will receive a demotion. If $\eta''-\eta'$ is large then few of the workers assigned to job 2 in period t will have values of η^e_{it} that are close to η', which implies that even fewer will have values of η^e_{it+1} that are below η'. Indeed, if $\theta_H f(x)>\eta'$ implies $\theta_L f(x+1)>\eta'$, then there can be a positive frequency of wage decreases but no demotions.[9]

Another important aspect of the symmetric-learning equilibrium is the size of wage increases upon promotion. Recall that we found that in the full-information case, if human-capital acquisition is close to linear then the average wage increase at promotion is larger than the average wage increase prior to promotion, consistent with the evidence. We also found, however, that the average wage increase the period *after* a promotion will be even larger than the average wage increase at promotion, and this is not consistent with the evidence. The symmetric-learning model fixes this problem, as follows.

Consider workers promoted from job 1 to job 2. As in the full-information case, the average wage increase at promotion is larger than the average wage increase prior to promotion, but there are now two reasons for this result. First, as before, increases in expected effective ability are valued at rate c_1 in job 1 and rate $c_2>c_1$ in job 2. Second, in the symmetric-learning case there is a selection effect: among workers who are observationally equivalent at the beginning of a given period, those promoted at the end of that period had a larger increase in expected innate ability than did those not promoted. Because of this selection effect, in the symmetric-learning case it is not necessary for the average wage increase the period after a promotion to be larger than the average wage increase at promotion. It is still the case that, after promotion to job 2, increases in expected effective ability are valued at rate c_2 rather than at a convex combination of c_1 and c_2. This logic also implies that wage increases after promotion should exceed those before

[9] BGH find that, although on average demotions are associated with wage decreases, some demotions are in fact associated with wage increases. Our model does not explain why this would be the case. Two standard explanations for such a finding are coding error (this could be studied by examining job assignments in earlier and later periods) and compensating differentials (not modeled here).

promotion.[10] The main point, however, is that in the symmetric-learning case, the selection effect disappears after promotion to job 2: the expected change in expected innate ability is (essentially) zero in the period after the promotion. Hence, in the symmetric-learning case, the average wage increase the period after a promotion can be either larger or smaller than the average wage increase at promotion.

Although symmetric learning exhibits large wage increases upon promotion in a "stronger" form than does full information, under symmetric learning it is still the case that wage increases upon promotion explain only a fraction of the difference between average wages across levels of a job ladder. The logic is the same as under full information. Since workers at any level j+1 will on average be older than workers at level j, part of the difference between average wages across levels j and j+1 is due to the greater human-capital accumulation of workers at level j+1. In turn, since the average wage increase at promotion to level j+1 only captures one year's worth of human-capital accumulation, the average wage increase at promotion to level j+1 should be significantly smaller than the difference between average wages across the levels.

The last BGH finding discussed in Section II is that those promoted from one level of the job ladder to the next come disproportionately, but not exclusively, from the top of the lower job's wage distribution (and analogously arrive disproportionately, but not exclusively, at the bottom of the higher job's wage distribution). For example, averaging across all promotions from levels 1 through 4, 66% of workers promoted come from above the median of the lower job's wage distribution, and 72% arrive below the median of the higher job's wage distribution.[11] Corollary

---

[10] BGH present evidence that supports this prediction. In particular, they find that for workers not promoted, the average percentage wage increase is 3.2% for workers at level 1, 4.8% for workers at level 2, 5.4% for workers at level 3, and 6.3% for workers at level 4.

[11] BGH also find a systematic relationship between job level and these percentages. That is, in most cases, the higher is the job from which promotion occurs, the higher is the percentage of promoted workers who come from above the median of the lower job's wage distribution, and the higher is the percentage who arrive below the median of the higher job's wage distribution. Specifically, for promotion from levels 1, 2, 3, and 4, the percentages of workers who come from above the median of the lower job's wage distribution are respectively 60%, 67%, 76% and 92%, while the corresponding percentages for arriving below the median of the higher job's wage distribution are 65%, 73%, 88% and 81%. As discussed in footnote 14, this pattern can also be explained by our model.

2.4 proves a related result: the probability that a worker will be promoted is an increasing function of the worker's percentile in the lower job's wage distribution.

Corollary 2.4: Consider workers in period t who have x periods of labor market experience and are in job j. Let $\pi(w)$ denote the probability that a worker in this group who is currently paid wage w is promoted for period t+1. If $0<\pi(w)<1$, then $\pi(w)$ is increasing in w.

There are two factors that allow our model to capture the empirical findings concerning where workers are promoted from and where they are promoted to: heterogeneity in θ and learning about θ.[12] The first of these factors exists even under full information, as follows. Consider all workers at level j. Since on average workers with low values for expected innate ability accumulate human capital more slowly than workers with high values for expected innate ability, the latter workers will on average have more growth in expected effective ability. Hence, although most promoted workers will come from the top of the lower job's wage distribution, some workers with high values for expected innate ability may come from far below the top of the lower job's wage distribution. Analogously, most promoted workers will arrive near the bottom of the higher job's wage distribution, but some may arrive far above the bottom of the higher job's wage distribution.

Learning about θ produces a similar result. The idea is that, similar to what we argued above for human-capital accumulation, learning makes it possible for a worker's expected effective ability to increase substantially from one period to the next. Thus, there should be observations of workers being promoted from the low end of the lower job's wage distribution, but such

[12] Because Corollary 2.4 considers a set of workers all of whom have the same value for labor market experience, only the second factor plays a role in the proof of the corollary.

observations should be rare. Again, an analogous argument holds for where workers arrive in the higher job's wage distribution.[13] [14]

Overall, the symmetric-learning case does a better job than the full-information case of matching the BGH evidence. On the one hand, the symmetric-learning case leaves unchanged most of the predictions of the full-information case that match the evidence well. For example, in the symmetric-learning case we still have serial correlation in wage increases and wage increases that predict promotion. On the other hand, the introduction of symmetric learning changes a number of the predictions of the full-information case that are a poor match with the evidence. Under symmetric learning we have a positive frequency of both demotions and wage decreases, and can also explain why there would be a significant number of wage decreases but very few demotions. Additionally, symmetric learning does a better job than the full-information case of matching the evidence concerning the size of wage increases at promotion and the positions of promoted individuals in the wage distributions of the sending and receiving jobs.

V. SYMMETRIC LEARNING AND ITS RELATIONSHIP TO ADDITIONAL EMPIRICAL FINDINGS

In Section IV we considered a model involving job assignment, on-the-job human-capital acquisition, and symmetric learning. We demonstrated that a model that integrates these three

---

[13] Another explanation for this finding that is consistent with our approach, but which we do not formally model, is that there is unmeasured heterogeneity concerning job titles in the BGH data set. Suppose that what BGH call a job title actually includes many heterogeneous subtitles, indistinguishable in their data, and that a promoted worker always moves from the very top of one wage distribution to the very bottom of another. If the wage distributions of one pair of subtitles (where there are promotions from the lower subtitle into the higher) are shifted relative to those of other pairs (say, because of compensating differentials) then analysis of the aggregate job titles will yield the BGH pattern.

[14] As discussed in footnote 11, BGH find that in most cases the higher is the job from which promotion occurs, the higher is the percentage of promoted workers who come from above the median of the lower job's wage distribution, and the higher is the percentage who arrive below the median of the higher job's wage distribution. This pattern can be explained by the second factor discussed above. That is, in the type of learning model we consider, the incremental amount that is learned should fall as the worker ages. Hence, since being higher on the job ladder will be positively correlated with age, our second factor predicts both that the percentage of workers promoted from the high end of the lower job's wage distribution should rise as job level rises, and that the percentage promoted to the low end of the higher job's wage distribution should also rise as job level rises.

elements captures a number of recent empirical findings concerning wage and promotion dynamics in internal labor markets. In this section we extend the analysis of our symmetric-learning model in two directions. First, we discuss how our approach relates to evidence concerning performance evaluations. Second, we identify and discuss findings from the BGH study that are inconsistent with the predictions of our symmetric-learning model.

A) Performance Evaluations

In an influential pair of studies in the early 1980s, Medoff and Abraham (hereafter MA) used the personnel records of three firms to explore the human-capital explanation for why wages grow over a career. In their 1980 paper, MA demonstrate that within-job wages are positively related to labor-market experience, but that performance evaluations are either unrelated or slightly negatively related to experience. Because MA interpret performance evaluations as measures of productivity, they conclude that human-capital theory is not consistent with this pair of facts.

To test whether performance evaluations are a good measure of productivity, MA run two further regressions. First, they find that performance evaluations predict future promotions. Second, they (and also BGH) find that performance evaluations predict future wage increases.

In this subsection we build on a discussion in Harris and Holmstrom (1982) to offer an interpretation of performance evaluations that is consistent with both our theoretical model and MA's evidence, and yet preserves human-capital theory as the primary explanation for wage growth over a career. In particular, we propose that supervisors evaluate individuals relative to other individuals with the same labor-market experience. That is, in the terminology of our model, performance evaluations measure expected innate ability (θ^e) rather than current productivity (i.e., expected effective ability, η^e).

Given our interpretation of performance evaluations, all the results described at the beginning of this subsection are easy to understand. The main MA results are that, within a job level, the average wage increases with experience, while the average performance evaluation falls with experience. In our model the wage is an increasing function of expected effective ability, which increases with experience because of human-capital accumulation. Further, performance evaluations within a job level fall with experience because workers with high expected innate ability will be promoted into higher level jobs.[15] The two other MA findings (the second supported also by BGH) are that performance evaluations predict both future promotions and future wage increases. These results also follow easily from our framework because promotions and wage increases are both positively related to expected innate ability.

Two other sources of evidence support our explanation of the MA findings. The first is a meta-analysis of forty studies of age and job performance by Waldman and Avolio (1986). Some of the studies contained direct measures of productivity while others contained supervisory performance evaluations. Waldman and Avolio found that there was a significant positive relationship between direct measures of productivity and age, but a slight negative relationship between supervisory performance evaluations and age. Thus, Waldman and Avolio provide clear evidence that performance evaluations are not unbiased measures of productivity. Furthermore, their results are consistent with MA's empirical findings, but not their conclusion that human-capital acquisition plays a minor role in wage growth. Rather, Waldman and Avolio's findings support our model's explanation for the MA evidence including the important role we give to human-capital acquisition.

The second source of evidence supporting our explanation for the MA findings concerns a comparison between cross-sectional and longitudinal estimates from a single dataset. Our model predicts that in a cross-sectional study the selection effect concerning who earns promotions would be important, so within a job level the average wage should rise with experience while the average performance evaluation should fall. In contrast, consider a longitudinal study that follows a fixed set of workers who remain at the same job level over a specified period of time. Our model still predicts that the average wage should rise with experience. Because the study concerns a fixed set of workers who remain at the same job level, however, most of the selection

[15]MA also present results concerning firm seniority as well as labor-market experience. Because there is no turnover in equilibrium in our model, we do not address these MA findings.

effect concerning who earns promotions will disappear.[16] As a result, the prediction is that there should be a much weaker negative relationship between average performance evaluation and experience.

In their 1981 study Medoff and Abraham perform both cross-sectional and longitudinal regressions on the same dataset and find results consistent with the above discussion (see also Gibbs (1995)). In the cross-sectional estimates they find that within a job level there is a significant positive relationship between the average wage and experience, and a significant negative relationship between the average performance evaluation and experience. In the longitudinal estimates they again find a significant positive relationship between the average wage and experience, but the relationship between the average performance evaluation and experience is either zero or slightly negative.

B) Empirical Findings We Do Not Capture

In Sections IV and V.A we focused on empirical findings in the BGH and MA studies that are consistent with our symmetric-learning approach. In this sub-section we identify and discuss four findings from the BGH study that are not captured by our model. The finding most at odds with our model involves cohort effects. BGH look at cohorts that enter the firm at different dates and find that, even after controlling for observable differences across cohorts, a cohort's average entry-level wage is an important determinant of the cohort's average wage years after entry. In our framework wages are determined by competitive bidding, and the state of the labor market in past years should not be an important determinant of current wages. In contrast, the cohort finding suggests that workers are somewhat insulated from the market, so that competitive bids have only a weak effect on wages.

---

[16]Even in a longitudinal study there should be a small negative relationship between average performance evaluation and experience. The reason is that a fixed set of workers who remain at the same job level over a period of time will on average be workers who experienced negative realizations concerning output. That is, workers who are not promoted will typically be workers for whom learning led to more pessimistic beliefs concerning expected innate ability, and thus we would expect at least a small negative relationship between the average performance evaluation and experience.

A second empirical finding inconsistent with our model involves nominal wage rigidity. Although BGH find a significant percentage of real wage decreases, nominal wage decreases are almost nonexistent in their sample -- fewer than 200 observations in a sample of almost 70,000. Our model provides no rationale for why this should be the case, but an extension of our model does match some related evidence from BGH as follows.

Suppose one were to add to our model a constraint that firms could not offer nominal wage decreases and a firing cost so that firms would not have an incentive to fire a worker whenever the worker's wage exceeded the worker's marginal product. In a model with these two new features, if the information revealed about the worker in period t is very negative, then in period t+1 the firm would be constrained to pay a wage higher than would be predicted by our original model. Further, since the actual wage in t+1 is "too high", the wage increase in t+2 should be very low. BGH find evidence in this spirit: the probability of a zero nominal increase in a given year, conditional on a zero nominal increase in the previous year, is two to three times the unconditional probability.

A third empirical finding that is inconsistent with the predictions of our model is that the wage distributions for adjacent job levels are overlapping. That is, the highest paid workers in level j are paid more than the lowest paid workers in level j+1. This is in contrast to our model where the highest paid workers in level j earn wages just below the wages paid to workers in level j+1. As opposed to the two findings discussed above, we do not feel that overlapping wage distributions pose a problem for our model because there are a number of modifications of our model that would yield overlapping wage distributions, such as the following.

A simple modification to our model that would yield overlapping wage distributions is to introduce unmeasured heterogeneity in job titles (see footnote 13 for a related discussion). The BGH evidence concerns the wage distributions for individual job titles. But what BGH call a single job title could easily include many heterogeneous subtitles, indistinguishable in their data. Suppose, for example, that every promoted worker moves from the top of one subtitle to the bottom of the next, as in our model. If the wage distributions for one pair of subtitles is shifted relative to that of another pair (say, because of compensating differentials) then an analysis of the aggregate job titles will yield overlapping wage distributions.

The last BGH finding that is at least somewhat inconsistent with the predictions of our model is the "green card" effect. Consider the set of workers who are in job j at date t. The green card effect refers to the idea that, for this set of workers, the expected wage increase, $E(w_{t+1}| w_t)-w_t$, is negatively related to the initial wage, w_t. Notice that this result is not as clearly inconsistent with our model as the findings discussed above. That is, because the effect of human capital on effective ability is concave ($f''<0$) and workers higher in the wage distribution at a given job are on average likely to have more labor-market experience, our model predicts at least a small degree of negative correlation between $E(w_{t+1}| w_t)-w_t$ and w_t.

Nevertheless, BGH find a strong green-card effect, and we feel it is unlikely that this result can be completely explained by the concavity of f(x). The correct explanation may be that green card effects are the result of administrative rules and procedures that move wages away from what pay would be in a spot market, but other explanations are possible. Our framework can easily be modified to produce green-card effects capable of matching the BGH findings. For example, in the current specification, on any given job, expected output is a linear function of the worker's effective ability. A simple modification would be to have output on each job be a concave function of effective ability. Our preliminary analysis of this case suggests that this modification would have little effect on the predictions of the model discussed earlier, but would clearly strengthen the negative correlation on any given job between $E(w_{t+1}| w_t)-w_t$ and w_t.[17]

[17] One question of interest is whether the four BGH findings discussed in this subsection are supported by other studies in the literature. The answer is that there is strong support for the finding that nominal wage decreases are rare (see e.g., Kahn (1994) and Card and Hyslop (1995)), and some support for green card effects (see Murphy (1991)). In contrast, as far as we know, cohort effects and overlapping wage distributions have not been addressed elsewhere (although Beaudry and DiNardo (1991) provide evidence that is related to cohort effects).

VI. CONCLUSION

Recently, a number of empirical studies have focused on the workings of internal labor markets, and found results that are inconsistent with theoretical models based on any single factor such as human-capital acquisition or learning. In this paper we developed a model that integrates job assignment, on-the-job human-capital acquisition, and learning. We demonstrated that a model that combines these elements captures many of the findings in this recent literature. Our conclusion is not that our symmetric-learning model provides a perfectly accurate representation of the workings of internal labor markets. Rather, we find the results encouraging, and feel that our model offers a basic framework upon which to build more accurate models of careers in organizations.

One direction in which the analysis might fruitfully be extended would be to incorporate an element of asymmetric learning.[18] In this paper we assumed that all information concerning a worker's ability is public knowledge, so firms are always equally informed about a worker's ability. A number of authors have argued that this assumption is unrealistic: most employment relationships are characterized by an element of asymmetric learning, so that a worker's current employer has better information concerning the worker's ability than do other firms in the market.

Several additional results would arise by incorporating an element of asymmetric learning into our model. For example, asymmetric learning introduces an additional reason why promotions are associated with large wage increases: because the current employer has private information, its promotion decision serves as a signal to the market, so promoted workers receive large wage increases to stop them from being bid away by the market (see Waldman (1984b)). Another result that would arise is a difference in the future outcomes of workers who vary only in terms of their seniority at the firm: asymmetric learning means that a firm should be more certain

[18] Bernhardt (1995) develops a model that incorporates job assignment, on-the-job human-capital acquisition, and asymmetric learning. The difference between our suggestion and Bernhardt's analysis is that we suggest a model characterized by a mix of symmetric and asymmetric learning, while in Bernhardt's analysis learning is purely asymmetric -- there are no public signals. One problem with a specification where learning is purely asymmetric is that wage increases do not predict future promotion probabilities, which is inconsistent with evidence from BGH.

concerning the ability levels of those workers with high seniority and this should translate into predictable differences in the future outcomes of newly hired workers versus those with long-term attachments. In fact, BGH find a result of this sort. For workers entering level 3 between 1970 and 1979, workers who have been at the firm longer have a lower variance concerning the highest level of the firm attained through subsequent promotions.

In summary, we feel the analysis in this paper demonstrates that job assignment, on-the-job human-capital acquisition, and symmetric learning are important building blocks for constructing models of internal labor markets. However, incorporating asymmetric learning and other elements such as the role of incentives will be important for extending our framework so that it more accurately matches the reality of careers in organizations.

APPENDIX

Proof of Proposition 1:

A worker with effective ability η has expected output $Ey_j = d_j + c_j \eta$ in job j. The efficient task assignment is therefore: job 1 if $\eta \leq \eta'$, job 2 if $\eta' < \eta \leq \eta''$, and job 3 if $\eta'' < \eta$. Competition among firms yields both efficient task assignment and wages equal to expected output: $w = d_j + c_j \eta$ for the efficient job j. Q.E.D.

Proof of Proposition 2:

The model has been constructed so that the argument in the symmetric-learning case can parallel the argument in the full-information case: wages again equal expected output each period, now given the observed history from prior periods. Downward-rigid contracts such as those in Harris-Holmstrom are feasible here but have no benefits because workers are risk-neutral. We compute a worker's expected effective ability η^e in (3) and then the worker's expected output in job j as $Ey_j = d_j + c_j \eta^e$. The linearity of the production function (2) is key here: without linearity, expected output would not equal the output of a worker known to have ability η^e. Finally, the fact that (2) reads $d_j + c_j (\eta_{it} + \varepsilon_{ijt})$ rather than $d_j + c_j\eta_{it} + \varepsilon_{ijt}$ means that the signal about ability that can be extracted from output—namely, $z_{it} = (y_{ijt} - d_j)/c_j = \eta_{it} + \varepsilon_{ijt}$—does not vary in its signal-to-noise ratio as a function of j. Thus, there is no way to use task assignment to change the speed of learning about ability, so task assignment is determined by current productive efficiency (i.e., maximizing expected output this period), which in turn is solely a function of the worker's current expected effective ability. Q.E.D.

For expositional simplicity, the proofs of Corollaries 2.1 through 2.4 are presented in the following order: 2.3, 2.4, 2.1, 2.2.

Proof of Corollary 2.3:

We first show that if $\theta_L f(x+1) < \theta_H f(x)$ then $\theta_L f(x+2) < \theta_H f(x+1)$. We do this by showing that $f(x)/f(x+1)$ increases in x, so that $f(x)/f(x+1) < f(x+1)/f(x+2)$, so that $\theta_L/\theta_H < f(x)/f(x+1)$ implies $\theta_L/\theta_H < f(x+1)/f(x+2)$. Taking the derivative of $f(x)/f(x+1)$ shows that this fraction is increasing in x if $f'(x)f(x+1) > f(x)f'(x+1)$, which holds because f(x) is increasing and concave. Therefore, $\theta_L f(x+1) < \theta_H f(x)$ for any $x > x^*$.

For $x < x^*$, we know that $\theta_L f(x+1) \geq \theta_H f(x)$, so any worker's expected effective ability must be higher when that worker has x+1 periods of prior experience than when he or she has x. Consequently, there can be no wage decreases or demotions.

For $x \geq x^*$ we show that there is a positive frequency of wage decreases by showing that beliefs about innate ability can move from arbitrarily close to θ_H to arbitrarily close to θ_L in one period. Let z^x denote the normalized output history $(z_{it-x},\ldots,z_{it-1})$, where we drop the subscript i on z^x for simplicity. Let $p = \text{Prob}\ (\theta = \theta_H | z^x)$ be the probability that the worker has high innate ability given z^x. By Bayes' rule:

$$\text{(A1)} \quad \text{Prob}\ (\theta = \theta | z^x, z_{it}) = \frac{p \cdot h[z_{it} - \theta_H f(x)]}{p \cdot h[z_{it} - \theta_H f(x)] + (1-p) h[z_{it} - \theta_L f(x)]},$$

where $h[\cdot]$ is the density of ε_{ijt} in (2), normal with mean 0 and variance σ^2. But

$$\text{(A2)} \quad \frac{h[z_{it} - \theta_L f(x)]}{h[z_{it} - \theta_H f(x)]} = \exp[-\frac{1}{2\sigma^2}\{[z_{it} - \theta_L f(x)]^2 - [z_{it} - \theta_H f(x)]^2\}]$$

$$= \exp[-\frac{1}{2\sigma^2}\{-2z_{it} f(x)(\theta_L - \theta_H) + f(x)^2(\theta_L^2 - \theta_H^2)\}],$$

which is monotonically decreasing in z_{it}, approaching 0 as $z_{it} \to \infty$ and approaching infinity as $z_{it} \to -\infty$. Therefore, $\text{Prob}\ (\theta = \theta_H | z^x, z_{it})$ approaches one as $z_{it} \to \infty$ and zero as $z_{it} \to -\infty$. That is, a sufficiently strong signal can move the market's belief arbitrarily far.

Because a sufficiently strong signal can move the market's belief arbitrarily far and $\theta_L f(x+1) < \theta_H f(x)$, wage decreases occur with positive probability. If x is such that $\theta_L f(x+1) < \eta' < \theta_H f(x)$ or $\theta_L f(x+1) < \eta'' < \theta_H f(x)$ then demotions also occur with positive probability. Such demotions entail wage decreases: because $\eta > \eta'$ implies $d_2 + c_2\eta > d_1 + c_1\eta$ we have (with some abuse of notation) $w_{it} = d_2 + c_2\, \eta^e_x > d_1 + c_1\, \eta^e_x > d_1 + c_1\, \eta^e_{x+1} = w_{x+1}$, and similarly for $\eta > \eta''$. Of course, the reverse is not true: because there is learning, there may often be wage decreases that occur without demotions. Q.E.D.

<u>Proof of Corollary 2.4</u>:

For simplicity, we consider workers in job 1; analogous arguments apply to job 2. For workers with experience x, the assumption in the Corollary that $0 < \pi(w) < 1$ can be restated as $\theta_L f(x+1) < \eta' < \theta_H f(x+1)$. That is, a worker with x+1 periods of experience remains in job 1 if the belief about her ability is sufficiently pessimistic, but is promoted to job 2 (or 3) if this belief if sufficiently optimistic.

Recall that the prior belief $\text{Prob}(\theta = \theta_H | z^x)$ at the beginning of period t is denoted by p and that the updated belief $\text{Prob}(\theta = \theta_H | z^x, z_{it})$ at the beginning of period t+1 is given in (A1). For simplicity, we will sometimes write the updated belief as q. Because the worker is assigned to job 1 in period t after x periods of experience we know that $[p\theta_H + (1-p)\theta_L]f(x) < \eta'$. To be promoted to job 2 (or 3) for period t+1, the worker's performance in period t, z_{it}, must be sufficiently high that $q \geq p^*_{x+1}$, where $[p^*_{x+1}\theta_H + (1- p^*_{x+1})\theta_L]f(x+1) = \eta'$. That is, given p, z_{it} must satisfy

$$\text{(A3)} \quad \frac{p}{p + (1-p) \cdot \dfrac{h[z_{it} - \theta_L f(x)]}{h[z_{it} - \theta_H f(x)]}} \geq p^*_{x+1}.$$

From (A2), $h[z_{it} - \theta_L f(x)] / h[z_{it} - \theta_H f(x)]$ is monotonically decreasing in z_{it}. Hence, given p, there exists a critical value $z^*_{x+1}(p)$ such that (A3) holds if and only if $z_{it} \geq z^*_{x+1}(p)$. That is, given past performance, there exists a critical value of current performance above which promotion occurs. From (A3), $z^*_{x+1}(p)$ decreases with p. That is, the critical value of current performance above which promotion occurs is lower if past performance has produced a more optimistic belief about innate ability. Furthermore, the conditional probability that z_{it} exceeds an arbitrary cutoff z^* increases in p, because

$$\text{(A4)} \quad \text{Prob}\,(z_{it} \geq z^* | z^x) = \text{Prob}\,(z_{it} \geq z^* | z^x, \theta_H) \cdot p + \text{Prob}\,(z_{it} \geq z^* | z^x, \theta_L) \cdot (1-p)$$

$$= \text{Prob}\,[\varepsilon_{ijt} \geq z^* - \theta_H f(x)] \cdot p + \text{Prob}\,[\varepsilon_{ijt} \geq z^* - \theta_L f(x)] \cdot (1-p)$$

$$= \text{Prob}\,[\varepsilon_{ijt} \geq z^* - \theta_L f(x)]$$

$$+ \; p \cdot \{\text{Prob}[\varepsilon_{ijt} \geq z^* - \theta_H f(x)] - \text{Prob}[\varepsilon_{ijt} \geq z^* - \theta_L f(x)]\},$$

where $\text{Prob}[\varepsilon_{ijt} \geq z^* - \theta_H f(x)] - \text{Prob}[\varepsilon_{ijt} \geq z^* - \theta_L f(x)] > 0$. Because the critical value decreases in p and the probability that z_{it} exceeds an arbitrary value increases in p, the probability of promotion increases in p.

The prior belief p is monotonically related to the current wage: $w = d_1 + c_1[p\theta_H + (1-p)\theta_L]f(x)$. Thus, the critical value of z_{it} such that promotion occurs decreases in w. Furthermore, the probability that z_{it} exceeds an arbitrary cutoff increases in w. For both of these reasons, the probability of promotion increases in w. Q.E.D.

Proof of Corollary 2.1:

The Corollary assumes that the worker is in job 1 for periods t and t+1, but is silent about period t+2. We break the proof into three cases: (1) $\theta_H f(x+2) < \eta'$, so that promotion is impossible after period t+1; (2) $\theta_L f(x+2) > \eta'$, so that promotion is guaranteed after period t+1; and (3) $\theta_L f(x+2) < \eta' < \theta_H f(x+2)$, so that promotion is possible but not guaranteed after period t+1.

In all three cases, we continue to use the notation above: $p = \text{Prob}(\theta = \theta_H \mid z^x)$, so that $w_{it} = d_1 + c_1[p\theta_H + (1-p)\theta_L]f(x)$, and $q = \text{Prob}(\theta = \theta_H \mid z^x, z_{it})$, so that $w_{i,t+1} = d_1 + c_1[q\theta_H + (1-q)\theta_L]f(x+1)$. Thus,

$$\text{(A5)} \quad w_{i,t+1} - w_{it} = c_1[q\theta_H + (1-q)\theta_L]f(x+1) - c_1[p\theta_H + (1-p)\theta_L]f(x)$$

increases in q given p. In each case, therefore, we show that $E(w_{i,t+2} - w_{i,t+1} \mid p, q)$ increases in q given p, and hence that $E(w_{i,t+2} - w_{i,t+1} \mid w_{it}, w_{i,t+1} - w_{it})$ increases in $w_{i,t+1} - w_{it}$ given w_{it}. To express the wage $w_{i,t+2}$, we introduce one last piece of notation: $r = \text{Prob}(\theta = \theta_H \mid z^x, z_{it}, z_{i,t+1})$.

Case (1) is simple, because the worker is certain to be in job 1 for period t+2, so $w_{i,t+2} = d_1 + c_1[r\theta_H + (1-r)\theta_L]f(x+2)$. But beliefs are a martingale, so $E(r \mid q) = q$. Thus, as of the beginning of period t+1, the worker's expected wage increase from t+1 to t+2 is simply the slope c_1 times the expected increase in the worker's effective ability due to human-capital acquisition:

$$\text{(A6)} \quad E(w_{i,t+2} - w_{i,t+1} \mid w_{it}, w_{i,t+1} - w_{it}) = c_1[q\theta_H + (1-q)\theta_L]\,[f(x+2) - f(x+1)]\,,$$

which increases in q.

Case (2) involves one more step than case (1): the worker's expected increase in effective ability from t+1 to t+2 is again $[q\theta_H + (1-q)\theta_L]\,[f(x+2) - f(x+1)]$, but now part of this increase is rewarded at c_1 and the rest at the higher slope c_2. But both of these effects favor workers with larger wage increases: the expected increase in effective ability is larger, and more of this increase is valued at the higher slope c_2. To see this, let $\eta^e_{t+1} = [q\theta_H + (1-q)\theta_L]\, f(x+1)$ denote the worker's expected effective ability at the beginning of period t+1. We know that $\eta^e_{t+1} < \eta'$ because the worker remains in job 1 for period t+1. Thus, in the expected wage increase from t+1 to t+2, the expected increase in effective ability from t+1 to t+2 can be written as all being rewarded at slope c_2 except for the part from η^e_{t+1} up to η', which is rewarded at c_1. That is, (A7) $E(w_{i,t+2} - w_{i,t+1} \mid w_{it}, w_{i,t+1} - w_{it}) =$

$$c_2[q\theta_H + (1-q)\theta_L]\,[f(x+2) - f(x+1)] - (c_2 - c_1)(\eta' - \eta^e_{t+1})\,.$$

The first term increases in q, as in case (1), and $\eta' - \eta^e_{t+1}$ decreases in q.

Case (3) is another step more complicated than case (2), because now it is not certain that the worker will be promoted for period t+2. In this case, it is convenient to rearrange the terms compared to (A7): the expected wage increase can be written as the entire expected increase in effective ability rewarded at slope c_1, plus any increase above η' rewarded at the incremental slope $(c_2 - c_1)$. That is,

(A8) $E(w_{i,t+2} - w_{i,t+1} \mid w_{it}, w_{i,t+1} - w_{it}) =$

$$c_1[q\theta_H + (1-q)\theta_L]\,[f(x+2) - f(x+1)] + (c_2 - c_1)E[g(\eta_{t+2}) \mid w_{it}, w_{i,t+1} - w_{it})\,,$$

where $\eta_{t+2} = [r\theta_H + (1-r)\theta_L]\,f(x+2)$ is the worker's expected effective ability at the beginning of period t+2 and $g(\eta_{t+2}) = \max\{0, \eta_{t+2} - \eta'\}$. As in cases (1) and (2), the first term increases in q. It therefore remains to show that $E(g(\eta_{t+2}) \mid w_{it}, w_{i,t+1} - w_{it}) = E(g(\eta_{t+2}) \mid p, q)$ increases in q given p.

Because $g(\eta_{t+2})$ is an increasing function, it suffices to show first-order stochastic dominance: for any η^* between $\theta_L f(x+2)$ and $\theta_H f(x+2)$, $\text{Prob}(\eta_{t+2} \geq \eta^* \mid p, q)$ increases in q. Because $\eta_{t+2} = [r\theta_H + (1-r)\theta_L]\,f(x+2)$, it suffices to show that for any r* in [0, 1], $\text{Prob}(r \geq r^* \mid p, q)$ increases in q. The argument is analogous to Corollary 2.4 but advanced one period, as follows.

Given q there exists $z^*_{x+2}(q)$ such that $r = \text{Prob}(\theta = \theta_H \mid z^x, z_{it}, z_{i,t+1}) \geq r^*$ if and only if $z_{i,t+1} \geq z^*_{x+2}(q)$. The critical value $z^*_{x+2}(q)$ decreases in q. Furthermore, the conditional probability that $z_{i,t+1}$ exceeds an arbitrary cutoff z^* increases in q. For these two reasons, the probability that $z_{i,t+1}$ exceeds the critical value increases in q. Q.E.D.

<u>Proof of Corollary 2.2</u>:

The worker has zero periods of experience, so $w_{it} = d_1 + c_1[p_0\theta_H + (1-p_0)\theta_L]f(0)$. We assumed in Section III.A that $\theta_H f(1) < \eta'$, so every worker spends at least two periods in job 1. Thus, the next wage will be $w_{i,t+1} = d_1 + c_1[q_0\theta_H + (1-q_0)\theta_L]f(1)$, where q_0 denotes $\text{Prob}(\theta = \theta_H \mid z_{it})$. The wage change $w_{i,t+1} - w_{it}$ therefore increases in q_0 given p_0.

Because of human-capital acquisition, the worker will eventually be promoted to job 2. The lowest experience at which promotion can occur is x_H' (defined in Section III.B as the lowest value of x such that $\theta_H f(x) > \eta'$). Recall from Corollary 2.4 that a worker with experience x_H' will not be in job 1 if the belief $p = \text{Prob}(\theta = \theta_H | z^x)$ at $x = x_H$' exceeds the critical value p* that solves $[p^*\theta_H + (1-p^*)\theta_L]f(x_H') = \eta'$. Part (i) of the Corollary claims that the probability that a worker with experience x_H' will not be in job 1 is an increasing function of the first wage change, $w_{i,t+1} - w_{it}$. So it suffices to show that the probability that p exceeds p* is increasing in q_0. But this is a special case of Corollary 2.4, extended to allow an arbitrary number of periods between the realization of q_0 after one period of experience and the assessment of whether $p \geq p^*$ after x_H' periods of experience.

Part (ii) invokes a new assumption: that demotions are impossible. Under this assumption, the workers who are not in job 1 at any given date are precisely those who have been promoted by that date. That is, no workers are promoted earlier but then reassigned to job 1 on the date in question. Recall that a worker with experience x is not in job 1 if $p = \text{Prob}(\theta = \theta_H | z^x) \geq p_x^*$, where $[p_x^*\theta_H + (1-p_x^*)\theta_L]f(x) = \eta'$. Thus, the probability that a worker with experience x and first wage change $w_{i,t+1} - w_{it}$ is not in job 1 is $\text{Prob}(p \geq p_x^* \mid q_0)$. This probability increases in q_0, for the same reasons as given in part (i). Thus, considering all the experience levels x from x_H' to x_L', the cumulative probability of promotion $\text{Prob}(p \geq p_x^* \mid q_0)$ is higher at every x for a higher value of q_0, in the sense of first-order stochastic dominance. Hence, the expected date of promotion is lower for a higher value of q_0. Q.E.D.

REFERENCES

Abowd, J. and D. Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57:411-45.

Baker, G., M. Gibbs, and B. Holmstrom. 1994a. "The Wage Policy of a Firm." *Quarterly Journal of Economics* 109:881-919.

__________, __________, and __________. 1994b. "The Internal Economics of the Firm: Evidence from Personnel Data." *Quarterly Journal of Economics* 109:921-55.

__________, M. Jensen, and K.J. Murphy. 1988. "Compensation and Incentives: Practice vs. Theory." *Journal of Finance* 43:593-616.

Beaudry, P. and J. DiNardo. 1991. "The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Microdata." *Journal of Political Economy* 99:665-88.

Becker, G. 1964. Human Capital. New York: NBER.

Ben-Porath, Y. 1967. "The Production of Human Capital and the Life-Cycle of Earnings." *Journal of Political Economy* 75:352-65.

Bernhardt, D. 1995. "Strategic Promotion and Compensation." *Review of Economic Studies* 62:315-39.

__________ and D. Scoones. 1993. "Promotion, Turnover, and Preemptive Wage Offers." *American Economic Review* 83:771-91.

Brüderl, J., A. Diekmann, and P. Preisendörfer. 1991. "Patterns of Intraorganizational Mobility: Tournament Models, Path Dependency, and Early Promotion Effects." *Social Science Research* 20:197-216.

Card, D. and D. Hyslop. 1995. "Does Inflation 'Grease the Wheels of the Labor Market'?" Princeton University Industrial Relations Section Working Paper #356, December.

Carmichael, L. 1983. "Firm-Specific Human Capital and Promotion Ladders." *Bell Journal of Economics* 14:251-58.

Demougin, D. and A. Siow. 1994. "Careers in Ongoing Hierarchies." *American Economic Review* 84:1261-77.

Doeringer, P. and M. Piore. 1971. *Internal Labor Markets and Manpower Analysis*. Lexington, MA: Heath Lexington Books.

Farber, H. and R. Gibbons. 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111:1007-47.

Gerhart, B. and G. Milkovich. 1989. "Salaries, Salary Growth, and Promotions of Men and Women in a Large, Private Firm." in *Pay Equity: Empirical Inquiries* R. Michael, H. Hartmann, and B. O'Farrell (eds.). Washington, DC: National Academy Press.

Gibbons, R. 1997. "Incentives and Careers in Organizations." Chapter 1 in Volume II of *Advances in Economics and Econometrics: Theory and Applications.* D. Kreps and K. Wallis (eds.). Cambridge University Press.

Gibbons, R. and L. Katz. 1991. "Layoffs and Lemons." *Journal of Labor Economics* 9:351-80.

__________ and __________. 1992. "Does Unmeasured Ability Explain Inter-Industry Wage Differences?" *Review of Economic Studies* 59:515-35.

Gibbs, M. 1995. "Incentive Compensation in a Corporate Hierarchy." *Journal of Accounting and Economics* 19:247-77.

Greenwald, B. 1986. "Adverse Selection in the Labor Market." *Review of Economic Studies* 53:325-47.

Harris, M. and B. Holmstrom. 1982. "A Theory of Wage Dynamics." *Review of Economic Studies* 49:315-33.

Hashimoto, M. 1981. "Firm-Specific Human Capital as a Shared Investment." *American Economic Review* 71:475-81.

Hause, J. 1980. "The Fine Structure of Earnings and the On-the-Job Training Hypothesis." *Econometrica* 48:1013-30.

Holmstrom, B. 1982. "Managerial Incentive Problems -- A Dynamic Perspective." in *Essays in Economics and Management in Honor of Lars Wahlbeck*. Stockholm: Swedish School of Economics.

Jovanovic, B. and Y. Nyarko. 1996. "Stepping Stone Mobility." National Bureau of Economic Research Working Paper #5651, July.

Kahn, C. and G. Huberman. 1988. "Two-Sided Uncertainty and Up-or-Out Contracts." *Journal of Labor Economics* 6:423-44.

Kahn, S. 1994. "Evidence of Nominal Wage Stickiness from Microdata." Mimeo, Boston University.

Lazear, E. 1992. "The Job as a Concept." In *Performance Measurement, Evaluations, and Incentives*. W. Bruns (ed.). Boston: Harvard Business School Press.

Lillard, L. and Y. Weiss. 1979. "Components of Variation in Panel Data: American Scientists 1960-1970." *Econometrica* 47:437-54.

MacDonald, G. 1982. "A Market Equilibrium Theory of Job Assignment and Sequential Accumulation of Information." *American Economic Review* 72:1038-55.

MacLeod, B. and J. Malcomson. 1988. "Reputation and Hierarchy in Dynamic Models of Employment." *Journal of Political Economy* 96:832-54.

Main, B., C. O'Reilly, and J. Wade. 1993. "Top Executive Pay: Tournaments or Teamwork?" *Journal of Labor Economics* 11:606-28.

McCue, K. 1996. "Promotions and Wage Growth." *Journal of Labor Economics* 14:175-209.

McLaughlin, K. 1994. "Rigid Wages?" *Journal of Monetary Economics* 34:383-414.

Medoff, J. and K. Abraham. 1980. "Experience, Performance, and Earnings." *Quarterly Journal of Economics* 95:703-36.

__________ and __________. 1981. "Are Those Paid More Really More Productive?" *Journal of Human Resources* 16:186-216.

Murphy, K.J. 1985. "Corporate Performance and Managerial Remuneration: An Empirical Analysis." *Journal of Accounting and Economics* 7:11-42.

Murphy, K.J. 1986. "Incentives, Learning, and Compensation: A Theoretical and Empirical Investigation of Managerial Labor Contracts." *Rand Journal of Economics* 17:59-76.

__________. 1991. "Merck & Co., Inc. (A)." *Harvard Business School Case* #9-491-005.

Podolny, J. and J. Baron. 1997. "Make New Friends and Keep the Old?: Social Networks, Mobility, and Satisfaction in the Workplace." Forthcoming *American Sociological Review*.

Prendergast, C. 1993. "The Role of Promotion in Inducing Specific Human Capital Acquisition." *Quarterly Journal of Economics* 108:523-34.

Ricart i Costa, J. 1988. "Managerial Task Assignment and Promotions." *Econometrica* 56:449-66.

Rosen, S. 1982. "Authority, Control, and the Distribution of Earnings." *Bell Journal of Economics* 13:311-23.

Rosenbaum, J. 1984. *Career mobility in a corporate hierarchy*. New York: Academic Press.

Ross, S., P. Taubman and M. Wachter. 1981. "Learning by Observing and the Distribution of Wages." in *Studies in Labor Markets* S. Rosen (ed.). Chicago: The University of Chicago Press.

Sattinger, M. 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica* 43:455-68.

__________. 1993. "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature* 31:831-80.

Topel, R. 1991. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy* 99:145-76.

____________ and R. Ward. 1992. "Job Mobility and the Careers of Young Men." *Quarterly Journal of Economics* 107:439-80.

Waldman, D. and B. Avolio. 1986. "A Meta-Analysis of Age Differences in Job Performance." *Journal of Applied Psychology* 71:33-38.

Waldman, M. 1984a. "Worker Allocation, Hierarchies, and the Wage Distribution." *Review of Economic Studies* 51:95-109.

__________. 1984b. "Job Assignments, Signalling, and Efficiency." *Rand Journal of Economics* 15:255-67.

__________. 1990. "Up-or-Out Contracts: A Signaling Perspective." *Journal of Labor Economics* 8:230-50.