

Experimentation and Job Choice

Kate Antonovics and Limor Golan* †

July 9, 2008

Abstract

In this paper, we examine optimal job choices when jobs differ in the rate at which they reveal information about workers' skills. We show that informational differences across jobs give rise to experimentation in that workers may be willing to sacrifice current period output in order to take jobs where learning is fast. In addition, we find that while experimentation is the most valuable when workers are young and inexperienced, the optimal level of experimentation is initially small, rises as workers gain experience and then eventually declines. We also find that job transitions and wage growth over the life-cycle are more involved than predicted by existing models. In addition, we use our model to shed light on the importance of job diversity and early career outcomes to future wage growth and on why identical workers (along both observable and unobservable dimensions) may experience distinct career outcomes even in the long run. Finally, we show that our model's prediction are broadly consistent with known patterns of wage dynamics and job transitions.

*Respectively, University of California, San Diego and Carnegie Mellon University.

†We thank Jan Eeckhout, Ronald Goettler, John Kennan, Dirk Krueger, Edward Lazear, Robert Miller, W. Bentley MacLeod, Victor Rios-Rull and Fallaw Sowell for their helpful comments, and Luis Zuluaga and Juan Vera for excellent research assistance. We also benefited from discussions with the participants of the SOLE 2004 meeting, the North American Econometrics Society 2004 meeting, Canadian Economic Theory conference 2005 and the seminar participants at Boston University, the St. Louis Federal Reserve, Stanford University, UCLA, UCSB and the University of Wisconsin . This research is supported by the Berkman Development Fund (Carnegie Mellon).

1 Introduction

When workers enter the labor market, there is uncertainty about their skill set. Over time, some of this uncertainty is resolved by observing workers' on-the-job performance. Since information is valuable in making future job choices, it is natural to think that both workers and firms consider the information value of jobs when making job assignments. Some firms even have job rotation policies explicitly aimed at revealing information about workers' skills. Intel's "Rotation Engineers Program", for example, allows workers to choose 3-4 jobs to rotate in and out of over a one year period. In their literature, Intel states that the program "allows you, as an employee, to accelerate your career growth and find your ideal fit within Intel by exposing you to various business groups and technologies in hardware design, software design, manufacturing and marketing." Thus, jobs are clearly valuable as a source of information about skills. In this setting, workers may be willing to experiment, entering into jobs where their expected current-period output is low, but which reveal a substantial amount of information about their skill.

In order to examine this tradeoff between current-period output and information, we develop a dynamic model of job choice in which workers learn about their skills by observing their on-the-job performance and in which jobs differ in the amount of information they provide about workers' skills. In particular, in our model, for each job, the more output depends upon a given skill, the more information the job reveals about that skill. Thus, for example, a worker will learn the most about whether he or she is a good manager in jobs where performance depends heavily on managerial ability. In this setting, we show that current-period job choices involve making a tradeoff between expected current-period output and information. We refer to this tradeoff as experimentation.

The concept of experimentation has received little attention in labor economics. Thus, one of this paper's primary contributions is to fill this gap. In particular, we analyze how the optimal level of experimentation changes over a worker's career. We then examine how

experimentation affects both job transitions and wage growth. In addition, we use our model to shed light on the importance of job diversity and early career outcomes to future wage growth and on why identical workers (along both observable and unobservable dimensions) may experience distinct career outcomes even in the long run. Finally, we show that our model's empirical predictions are broadly consistent with known patterns of wage dynamics and job transitions.

In our model, workers choose a job in every period to maximize the expected present discounted value of lifetime income. Although the model is formalized as one in which workers choose jobs, we show that under perfect competition (with symmetric information), optimal job assignments and earnings are the same if instead firms choose job assignments. To make the model tractable, we assume that there are two skills and that there is only uncertainty regarding one skill. For example, this might correspond the case in which a worker's ability as a programmer is known, but his or her managerial ability is not. In addition, because we allow workers' output to be correlated across jobs, we cannot make use of solution techniques from the literature on independent multi-armed bandit problems. Thus, we model the worker's problem as a dynamic programming problem and solve it numerically.

Workers value information because it increases the probability that they make correct job choices in the future. Thus, not surprisingly, we find that workers experiment, foregoing expected current period output in order to learn about their skill. Interestingly, however, we also show that the optimal level of experimentation follows a non-monotonic pattern over the life-cycle. At the beginning of a worker's career, when there is considerable uncertainty about skill, the optimal level of experimentation is relatively small, increases as worker's gain experience and then ultimately falls as workers become increasingly certain about their skills. The decline in experimentation at the end of a worker's career is intuitive; as workers gain experience, the degree of uncertainty about their skill level drops and, thus, the value of experimentation falls. The increase in experimentation in the early stages of a worker's career, however, is more puzzling. It is a manifestation

of the of the Radner and Stiglitz (1984) result, that the value of information may not be concave (this result is generalized by Chade and Schlee (2002)). It is driven by the fact that when there is a lot of dispersion in a worker's prior beliefs about his or her skill, marginal increases in information do little to increase the probability that workers are correctly assigned to jobs in the future. Thus, although workers are the most likely to experiment at the beginning of their career, the optimal level of experimentation is initially quite small.

Because the optimal experimentation level is non-monotonic over the life-cycle, we also find that job transitions are more involved than those predicted by standard learning and matching models. In particular, we show that when workers are young, they are likely to choose jobs in which output depends upon multiple dimensions of their skill set. Our model predicts, however, that these jobs are transitory in the sense that no worker remains in them for long. Thus, even though job transition rates in our model generally decline with experience, we show that, for some parameter values, more experienced workers may be more likely to leave these transitory jobs than less experienced workers. Interestingly, this non-monotonic pattern in job transition rates is consistent observed trends in the data. That is, empirical hazard functions computed at monthly frequencies are typically found to increase in the first few months of employment before eventually declining (see, for example, Farber (1999) and Pavan (2007)).

We also examine our model's implications for wage dynamics and show that, unlike standard models of wage growth in which wages increase because of either human capital accumulation or improvements in match quality, in our model, wage growth is also partially driven by the eventual decline in experimentation.¹ In addition, we examine the factors that influence a worker's lifetime earnings and show that job diversity increases both experimentation and lifetime earnings, suggesting that specialization may be subop-

¹One exception is Harris and Holmstrom (1982) who examine an environment in which firms provide risk-averse workers with partial insurance against negative productivity shocks. In their model, wages rises over workers' careers because as uncertainty about worker ability falls, so does the cost of insuring the worker against future wage cuts.

timal when jobs provide information about workers' skills. Further, we show that random productivity shocks can have long-lasting effects on both wages and wage growth. In particular, workers who receive negative productivity shocks may be reassigned to jobs that reveal little about their skill and where wage growth is slow. As a result, luck may lead to different career trajectories even in the long run. Further, since new information has the largest effect on prior beliefs when workers are young, productivity shocks have the largest impact early in worker's career.

In the final part of our paper, we also show that our model's prediction about wage dynamics and job transitions are largely consistent with observed patterns in the data. For example, like many standard learning models, our model predicts that wages grow with experience, wage growth primarily occurs in the early stages of workers' careers, cohort wage dispersion increases over time, job transition rates generally decline with experience and the wage distribution is positively skewed.² In addition, we also draw upon findings in Baker, Gibbs and Holmstrom (1994a and 1994b) to show that our model can account for the fact that 1) job transition rates are initially higher for outside hires than for those promoted from within, 2) wages are serially correlated, 3) workers, particularly those at the bottom of the earnings distribution, often experience real wage decreases in the first several years of employment and 4) labor market conditions at the time workers are hired appear to have lasting effects on wages and wage growth.

Our paper differs from the literature on job matching because these models primarily consider environments in which a worker's productivity is match-specific. Thus, the information a worker receives about his or her productivity on one job reveals nothing about that worker's productivity in other jobs (see, for example, Jovanovic (1979) and Miller (1984)). In this paper, by contrast, we analyze job choices when workers are uncertain about general human capital so that the information workers receive leads them to update their prior beliefs about their productivity in all jobs. In addition, unlike the literature on learning and sorting in which the rate of learning is assumed to be

²These predictions also arise in other models of wage determination.

constant across jobs (see, for example, Gibbons and Waldman (1999) and MacDonald (1982)), we allow the information content of jobs to differ and allow workers to select jobs based on their informational value.

Superficially, the non-monotonic relationship we find between experimentation and uncertainty may seem to contrast with Miller (1984). Miller analyzes a matching model in which the mean and variance of the prior distribution of match quality differ across occupations. Since workers learn about match quality more quickly in high-variance occupations than in low-variance occupations, workers may be willing to enter into occupations in which expected match quality is low as long as the prior variance of match quality is high enough. This tradeoff between expected match-quality and variance is a form of experimentation, and Miller shows that experimentation will primarily take place early in a worker's career. Our model differs, however, from Miller's along two important dimensions. First, in Miller, workers face a binary decision about whether or not to learn (they either enter an occupation or they do not). In our model, on the other hand, workers can choose both whether to learn and how much to learn. As we discuss below, once this distinction is made clear, our results turn out to complement rather than contradict Miller. Second, as mentioned above, unlike Miller, we also focus on uncertainty about general human capital rather than match quality.

Our paper is closely related to several papers in the experimentation literature. A primary difference between our paper and the papers we discuss below is that we examine how the optimal level of experimentation depends on both the mean and variance (precision of beliefs) of the prior distribution of the unknown parameter. Previous models of experimentation have typically assumed that the unknown parameter of interest has a Bernoulli distribution. While this assumption simplifies the analysis (since there is only one parameter to keep track of), it has prevented previous researchers from being able to examine how experimentation separately depends on the mean and variance of the unknown parameter. Since we wish to examine how experimentation depends upon the amount of uncertainty about a worker's skill, we assume that the unknown skill is nor-

mally distributed with a known mean and variance. While this distributional assumption complicates our analysis, it allows us to examine how the optimal level of experimentation changes as workers gain experience and uncertainty about their skill set is resolved.

Moscarini and Smith (2000) develop a model in which the optimal level of experimentation is choice variable. Their model, however, does not have the structure of a bandit model since information costs are explicit. Moreover, the learning process is different in their model than in ours. In particular, in their model, the sample size is chosen each period (the number of i.i.d. signals), while in our model, the i.i.d. noise is fixed across jobs but the degree to which output depends on the unknown skill varies. Pastorino (2006) considers a model of job assignments in which firms choose the rate of learning. In her model, however, learning mainly occurs early in the career because jobs that reveal the least information about a worker's skill are high-level jobs in which the returns to skill are the highest. We make the opposite assumption since we wish to capture the notion that workers learn the most about a particular dimension of their skill in jobs where output depends critically upon that skill. Bolton and Harris (1999) examine the optimal level of experimentation when there are N decision makers who learn about the payoffs of a risky action from the experimentation of both others and themselves. Since agents can experiment simultaneously, Bolton and Harris focus on analyzing players' strategic interactions. Strategic considerations are not an issue in our paper since our model is equivalent to one in which there is a single decision maker. In addition, another set of papers examines the optimal level of experimentation when the unknown parameter of interest changes over time. Keller and Rady (1999) and Rustichini and Wolinsky (1995), for example, examine the problem of a monopolist who faces an unknown, time-varying demand curve. These models differ from ours in that we assume that worker ability is fixed.³

³In all of these models, because the unknown parameter of interest has a Bernoulli distribution, the prior variance of the unknown parameter is not a state variable. Thus, experimentation monotonically increases in the Bellman value, while in our model, the optimal experimentation level is non-monotonic in the Bellman value.

Finally, our paper is also related to a number of papers in the literature on learning that focus on labor markets. Felli and Harris (1996) examine a strategic model of wage determination in which output depends on two different skills, and there is heterogeneity in the productivity of each skill across jobs. Productivity in their model, however, is match-specific, and their focus is on how workers capture some of the returns to information about firm-specific human capital. Ortega (2001) builds a model of job rotation and shows that expected productivity is higher when firms learn about workers through job rotation rather than through fixed job assignments. Ortega does not, however, fully characterize the optimal job rotation policy. Like us, Neal (1999) also examines job choices over the life cycle. In his model, workers search for both good careers and good employers. Since there is no learning in his model, however, the information content of jobs plays no role in workers' decisions. Finally, like us, Jovanovic and Nyarko (1997) consider a model in which uncertainty plays a role in determining how workers move across jobs over their careers. In their model, however, uncertainty surrounds how to perform jobs rather than workers' skill sets.

2 Model

Consider an economy with infinitely-lived, risk-neutral workers and firms with a common discount factor δ . Workers differ in the set of skills that they possess. In principle, this skill set may be multidimensional and include skills such as creativity, diligence, adaptability, etc.. To focus on essentials, however, we examine a simple scenario in which each worker has only two skills: a known skill, k , and an unknown skill, θ , both of which are time-invariant.

Firms in the model offer a variety of jobs, and each job differs in the extent to which output depends upon k and θ . In addition, firms are identical so that a worker's expected output at a given job depends only upon the skill set of the worker and the characteristics of the job. In particular, there are N jobs, each completely characterized by a given value

of α , where α denotes the degree to which output depends on θ , relative to k . Thus, choosing a job in period t is equivalent to choosing a value of α . Given this choice, we assume output in period t is given by

$$y_t = \alpha_t \theta + (1 - \alpha_t)k + \epsilon_t, \quad (1)$$

where ϵ_t is an i.i.d. productivity shock, α_t denotes the value of α chosen by the worker at time t , and $\alpha_t \in \{\alpha^1, \dots, \alpha^N\}$, where $\alpha^1 = 0$, $\alpha^N = 1$ and $\alpha^r > \alpha^s$ for all $r > s$.⁴ Thus, there is one job in which output is only sensitive to θ ($\alpha^N = 1$) and one job in which output is only sensitive to k ($\alpha^1 = 0$). For the rest of the $N - 2$ jobs, the higher is α^j , the more output depends on θ .⁵

Information in the model is symmetric; firms and workers have common priors on θ , k is known to everyone and output is commonly observed. Workers and firms acquire additional information about a worker's unknown skill through successive observations of output. Notice, however, that even though both the known skill, k , and the unknown skill, θ , contribute to output, the contribution of θ cannot be observed separately from the contribution of the i.i.d. shock. Thus, having observed output, workers and firms calculate

$$x_t = \frac{y_t - (1 - \alpha_t)k}{\alpha_t} = \theta + \frac{\epsilon_t}{\alpha_t}, \quad (2)$$

where x_t serves as a signal of the worker's unobserved skill, θ . The noise of x_t is *not* independent of a worker's job choice. In particular, the higher is α_t , the higher is the signal-to-noise ratio and the more information about θ the market is able to extract from

⁴Here we assume that skills are perfect substitutes in production. However, the solution to this problem is equivalent to the solution of one in which there is a Cobb-Douglas production function, where $y_t = \ln \tilde{y}$, $k = \ln \tilde{k}$, $\theta = \ln \tilde{\theta}$, $\epsilon = \ln \tilde{\epsilon}$ and

$$\ln \tilde{y}_t = \ln[\tilde{\epsilon} \tilde{\theta}^{\alpha_t} \tilde{k}^{1-\alpha_t}] = \alpha_t \ln \tilde{\theta} + (1 - \alpha_t) \ln \tilde{k} + \ln \tilde{\epsilon}_t.$$

⁵We assume that set of jobs is exogenously given. Our results below, however, suggest that if firms were costlessly able to choose as many jobs as they want, they would choose a continuum of jobs so that $\alpha \in [0, 1]$.

x_t . Further, this information is valuable since knowledge of θ enables workers to sort into the jobs to which they are best suited.

Under the assumption that the prior distribution of θ at time t is normal with mean μ_t and variance σ_t^2 and the distribution of ϵ_t is normal with mean zero and variance σ_{ϵ}^2 , the posterior distribution of θ is known to be normal with mean μ_{t+1} and variance σ_{t+1}^2 where

$$\mu_{t+1} = \frac{\mu_t \sigma_{\epsilon,t}^2 + x_t \sigma_t^2}{\sigma_{\epsilon,t}^2 + \sigma_t^2} \quad (3)$$

and

$$\sigma_{t+1}^2 = \frac{\sigma_{\epsilon,t}^2 \sigma_t^2}{\sigma_t^2 + \sigma_{\epsilon,t}^2} \quad (4)$$

and where $\sigma_{\epsilon,t}^2 = \frac{\sigma_{\epsilon}^2}{\alpha_t^2}$. Notice that, while both the mean and the variance of the posterior distribution of θ are influenced by α_t , only the posterior mean is influenced by the signal x_t .

In addition, we know that μ_{t+1} is itself normally distributed with mean, m_{t+1} and variance s_{t+1}^2 given by:

$$m_{t+1} = \mu_t \quad (5)$$

$$s_{t+1}^2 = \frac{\sigma_t^4}{\sigma_t^2 + \sigma_{\epsilon,t}^2}. \quad (6)$$

Thus, the posterior mean of θ follows a martingale, and the more information x_t reveals about θ (the higher is α), the higher is the variance of the posterior mean.

Timing in the model is as follows: at the beginning of each period, workers announce a job choice, firms make take-it-or-leave-it wage offers and each worker accepts an offer. We assume competitive markets and free entry into the labor market.

Thus, given their prior beliefs about θ and given a worker's job choice, α_t , firms pick a wage policy $w_t(\alpha_t, \mu_t, \sigma_t^2)$ to maximize the present discounted value of future profits, and workers choose α_t in each period to maximize the expected present discounted value of lifetime earnings.

3 Equilibrium Wages and Optimal Job Choice

First, we characterize equilibrium wages and then solve for optimal job assignments. We assume spot contracts.

Proposition 1 *In equilibrium, in each period, 1) firms offer wages equal to each worker's expected productivity at the job chosen by the worker in that period, 2) each worker chooses the job that maximizes the worker's expected discounted lifetime productivity, and 3) if firms (instead of workers) made job choices, then firms would also assign workers the job that maximizes the worker's expected discounted lifetime productivity.*

The proof is in Appendix A. The first two points are straightforward. Competition among firms together with spot contracts imply that workers will be paid their expected productivity in every period. In addition, since wages equal workers' expected productivity in every period and since workers want to maximize their expected discounted lifetime earnings, the optimal job choice in each period maximizes workers' expected discounted lifetime productivity. The third point is more subtle. Suppose that firms rather than workers made job choices. Given that all information in the model is symmetric, when firms assign workers to jobs that reveal information about workers' skills, firms provide information to all other firms (and workers) in the market. Under perfect competition, however, when human capital is general, workers earn all the rents from this information. Therefore, if firms offer a worker any job other than the job that maximizes that worker's expected discounted lifetime earnings, then the worker would only accept the job if the firm offered a wage above the worker's expected productivity. This current-period loss cannot be optimal for the firm, however, because competition insures that the firm's future profits will never be greater than zero. Thus, one implication of Proposition 1 is that the solution to our model is the same regardless of whether we allow workers or firms choose jobs.

3.1 The worker's optimal job choice problem

Having established that choosing the job which maximizes expected discounted lifetime productivity can be achieved as an equilibrium outcome under perfect competition, we proceed by characterizing the worker's optimal job choices (taking into account that workers will be paid their expected productivity at each job in every period). The problem boils down to a single-agent dynamic optimization problem.

Since the worker's decision problem in period t is the same as the problem in period 1, except that the worker updates his or her prior beliefs about θ based on the history of productivity signals, $\{x_{t-1}, \dots, x_1\}$, according to Equations (3) and (4), the worker's problem is stationary. Thus, we can write the value function, $V(\cdot, \cdot)$, as the solution to a Bellman equation in which the control variable is α_t and in which the state variables, μ_t and σ_t^2 , describe the prior distribution of θ . That is,

$$V(\mu_t, \sigma_t^2) = \max_{\alpha_t \in \{\alpha^1, \dots, \alpha^N\}} \alpha_t \mu_t + (1 - \alpha_t)k + \delta \int V(\mu_{t+1}, \sigma_{t+1}^2) f(\mu_{t+1} | m_{t+1}, s_{t+1}^2) d\mu_{t+1}, \quad (7)$$

where f denotes the normal probability density function and where the dependence of σ_{t+1}^2 , m_{t+1} and s_{t+1}^2 on α_t is implicit and given in equations 4, 5 and 6.

The first two terms on the right-hand side of (7) represent expected current period output, and the second term represents the continuation value, which incorporates the value of information obtained from observing x_t .

As mentioned above, higher values of α_t are beneficial to workers in the sense that the higher is α_t , the more information on-the-job performance will reveal about a worker's skill. In expectation, this information is valuable because workers can insure themselves against the arrival of negative information about θ by selecting future jobs with $\alpha = 0$, but can take advantage of the arrival of positive information about θ by selecting future jobs with $\alpha = 1$. Thus, the upside benefits of information outweigh the downside risks. In the context of equation 7, the benefit of increasing α_t is reflected in the continuation value and the fact that s_{t+1}^2 , the variance of μ_{t+1} , is increasing in α_t .

Equation 7, however, also makes clear that when $\mu_t < k$, there is a cost associated

with selecting $\alpha_t > 0$. In particular, if $\mu_t < k$, then expected current-period output will be negative when $\alpha_t > 0$, so that obtaining information about θ entails a cost in terms of foregone current-period earnings. Thus, when $\mu_t < k$, workers must weigh the benefit of increasing α_t in terms of expected future output against the cost in terms of expected current-period output. If a worker chooses to forego expected current-period output in order to gain information about θ , then we say that the worker experiments.

Definition 1 *A worker experiments if $\mu_t < k$, but the worker chooses $\alpha_t > 0$.*

3.2 Solution method

A key feature of our model is that the productivity signal, x_t , provides information about a worker's productivity at many jobs. Thus, we cannot appeal to solution techniques developed in the literature on independent multi-armed bandit problems (because the "arms" in our problem are dependent).⁶ Instead, we solve our problem numerically.

Let B denote the space of bounded, continuous, real-valued functions taking their domain on the space, $S = (-\infty, \infty) \times [0, \infty)$. For all $g \in B$, define the mapping $T : B(S) \rightarrow B(S)$ as

$$Tg(\mu, \sigma^2) = \max_{\alpha \in \{\alpha^1, \dots, \alpha^N\}} \alpha\mu + (1 - \alpha)k + \delta \int g(\mu', \sigma'^2) f(\mu' | m, s^2) d\mu', \quad (8)$$

where $\sigma'^2 = \frac{\sigma_\varepsilon^2 \sigma^2}{\sigma^2 + \sigma_\varepsilon^2}$ and where $f(\cdot | m, s^2)$ is the normal probability density function with mean $m = \mu$ and variance $s^2 = \frac{\sigma^4}{\sigma^2 + \sigma_\varepsilon^2}$. As we show in Appendix B, Blackwell's sufficient conditions for T to be a contraction are satisfied. Thus, by the contraction mapping theorem, the value function in (7) is unique and can be obtained by iteration of T .

In our computations, we discretize the prior mean and use linear interpolation and extrapolation to evaluate the value function at arbitrary values of the prior mean. The prior variance is a deterministic function of the job assignment history and is weakly declining over time. Thus, we create a grid by calculating many possible variance points for all possible job assignment histories for large number of periods. As the number of

⁶For further discussion, see Gittins and Jones (1974).

periods becomes arbitrarily large, the variance drops to zero and thus zero is also on the grid. Since the value of the value function is known when the prior variance is zero, we use linear interpolation for variances greater than zero but too small to be on our grid.

4 Productivity Uncorrelated Across Jobs: Benchmark Case

Consider first the benchmark case in which $\alpha_t \in \{0, 1\}$, so that there are only two types of jobs: one in which output only depends on k ($\alpha = 0$) and another in which output only depends on θ ($\alpha = 1$). We refer to this model as the two-job model and call the job in which $\alpha = 1$ the high-level job and the job in which $\alpha = 0$ as the low-level job.

The two-job model is a useful starting point for our analysis since the worker's productivity in the high-level job and the low-level job are uncorrelated. Thus, observed output in one job reveals nothing about expected output in the other job. This version of our model is a one-arm bandit problem and can be viewed as a special case of Miller (1984) in which there are only two occupations: one in which expected productivity is known with certainty and another in which expected productivity is unknown. After characterizing the solution to our two-job model, we then discuss optimal job choices when $N > 2$.

Figure 1 plots the value function for the two-job model for various values of μ_t and σ_t when $k = 0$, $\sigma_\epsilon^2 = 1$ and $\delta = 0.90$. As the figure reveals, the value function is increasing in both the mean and the standard deviation of the prior distribution of θ (a formal treatment is in Miller (1984)). The value function simply gives the discounted sum of expected productivity in each period. Therefore, the higher is the expected value of θ , the greater is expected discounted lifetime earnings. In addition, as the prior variance of θ increases, the value function increases because workers are able take advantage of positive information about θ by selecting $\alpha = 1$ while insuring themselves against negative information about θ by selecting $\alpha = 0$. The worker's ability to always select $\alpha = 0$ also implies that the value function is bounded from below.

Next, we describe how μ_t and σ_t^2 affect job choices in the two-job model. Define a marginal worker in the two-job model as a worker for whom the prior mean and prior variance of θ are such that expected lifetime output is the same regardless of whether the worker chooses the low-level job or the high-level job. Let $\tilde{\mu}$ and $\tilde{\sigma}^2$ respectively denote the prior mean and prior variance of a marginal worker.⁷ Note that if a worker optimally chooses the low-level job in one period, then the worker will acquire no additional information about θ and will face the exact same decision-making environment in the next period. Thus, if a worker ever chooses the low-level job, then the worker optimally will choose the low-level job from that point on. As a result, the value of choosing the low-level job is simply $\frac{k}{1-\delta}$, and the prior mean and variance of any marginal worker must be such that $V(\tilde{\mu}, \tilde{\sigma}^2) = \frac{k}{1-\delta}$. Let $\tilde{\mu}(\tilde{\sigma}^2)$ be a function describing how $\tilde{\mu}$ changes with $\tilde{\sigma}^2$.

Proposition 2 *If there is uncertainty, then there is experimentation, and experimentation increases as uncertainty about θ rises. That is, if $\tilde{\sigma}^2 > 0$, then $\tilde{\mu}(\tilde{\sigma}^2) < k$, and $\tilde{\mu}(\tilde{\sigma}^2)$ is decreasing in $\tilde{\sigma}^2$.*

The proof is in Appendix A. Proposition 2 establishes two important facts. First, if $\sigma_t^2 > 0$, then even if $\mu_t < k$, workers may experiment by choosing $\alpha_t = 1$. In a static environment, this choice would be suboptimal since current expected output is higher the low-level job than in the high-level job. However, the high-level job provides information about a worker's unobservable skill, and this information is valuable in making future job choices. Thus, in a dynamic setting, workers may be willing to forego current period wages in order to learn about their skill, θ .

Second, Proposition 2 establishes that as the prior variance of θ increases, the minimum value of μ_t above which workers choose the high-level job falls. To see the intuition behind this, consider a worker with prior mean μ_t , where $\mu_t < k$. Note that the cost to this worker in terms of foregone earnings of choosing $\alpha_t = 1$ is $k - \mu_t$. Thus, the lower is μ_t , the higher is the cost of choosing the high-level job, and if σ_t^2 were zero,

⁷In the two-job case, the “marginal” worker is also the worker for whom the Gittins index on the “uncertain” arm is exactly k .

this worker would clearly choose the low-level job. As σ_t^2 increases, however, so does the option value of the high-level job since as σ_t^2 grows, the worker becomes increasingly likely to learn that $\theta > k$, and the worker can always choose $\alpha = 0$ if he or she believes that $\theta < k$. Thus, as uncertainty about θ increases, so does the value of information and the likelihood experimentation.

Figure 2 illustrates optimal job choices as a function of prior mean and prior variance when $k = 0$, $\sigma_\epsilon^2 = 1$, $\delta = 0.90$, and $\alpha \in \{0, 1\}$. As the figure shows, the boundary line between the region in which the worker chooses the low-level job ($\alpha = 0$) and the region in which the worker chooses the high-level job ($\alpha = 1$) is negatively sloped so that as σ_t rises, workers are willing to take larger and larger current period wage cuts in order learn about themselves. That is, the greater is the uncertainty about the unobservable skill, θ , the more willing are workers to experiment.

5 Productivity Correlated Across Jobs: General Model

In this section, we extend our discussion to include additional jobs. We continue to refer to $\alpha = 0$ as the low-level job and $\alpha = 1$ as the high-level job. In addition, we refer to all jobs with $\alpha \in (0, 1)$ as intermediate-level jobs. The crucial difference between this N -job model and the two-job discussed above is that, in the N -job model, if the worker selects a job with $\alpha > 0$, then observed output in that job will provide information about the worker's expected productivity in *all* jobs with $\alpha > 0$.

We find that our N -job model delivers qualitatively different predictions about optimal job choices than our two-job model. In particular, in the two-job model, for a given value of $\mu_t < k$, the optimal choice of α_t increases as σ_t^2 increases. However, this monotonic relationship between job choice and prior variance no longer holds when we move to the N -job model. In particular, we find that, for certain values of μ_t , the optimal value of α_t is first increasing and then decreasing in the prior variance of θ .

Figure 3 illustrates optimal job choices in the N -job model when $\delta = 0.9$, $k = 0$,

$\sigma_\epsilon^2 = 1$ and $\alpha \in \{0, 0.5, 1\}$. First, as in the two-job benchmark case discussed above, as long as $\sigma_t^2 > 0$, there are workers with $\mu_t < k$ who choose either $\alpha_t = 0.5$ or $\alpha_t = 1$. Thus, as in the two-job model there is a tradeoff between current wages and learning. In addition, similar to the two-job model, holding σ_t^2 constant, the higher is the prior mean of θ , the higher is the optimal choice of α_t .

In contrast to the two-job model, however, the optimal choice of α_t is not always increasing in the prior variance of θ . Put differently, in Figure 3, the frontier along which workers are indifferent between choosing $\alpha_t = 0.5$ and $\alpha_t = 1$ is negatively sloped when the prior variance of θ is relatively low but positively sloped when the prior variance of θ is relatively high.⁸ Thus, in contrast to the two-job model, experimentation is not strictly increasing in the prior variance of θ . Figure 4 demonstrates that this pattern remains when there are five jobs ($\alpha = \{0, 0.25, 0.5, 0.75, 1\}$).

To understand why there is a non-monotonic relationship between α_t and σ_t^2 , recall that when $\mu_t < k$, workers must weigh the benefit of increasing α_t in terms of expected future output against the cost in terms of expected current-period output. Consider, for example, a worker choosing between $\alpha_t = 1$ and $\alpha_t = 0.5$. If $\mu_t < k$, then the worker will only choose $\alpha_t = 1$ if the increase in expected future output outweighs the cost in terms of foregone expected current-period output. Note, however, that foregone expected current-period output (in this case, $0.5(k - \mu_t)$) does not depend on σ_t^2 . Thus, the non-monotonic relationship between σ_t^2 and the optimal choice of α_t must depend on solely on how increases in α_t affect expected future output. In particular, the effect of increasing α_t on expected future productivity must be low both when σ_t^2 is small and when σ_t^2 is large.

It is fairly easy to see why the benefit of increasing α_t is small when there is little uncertainty about θ . First, when σ_t^2 is small, the option value of new information is low

⁸Notice that the frontier between $\alpha_t = 0$ and $\alpha_t = 0.5$ is an “indifference” curve in which the value function is the same along the curve. However, along the frontier, between $\alpha_t = 0.5$ and $\alpha_t = 1$, the value function differs at different points on the frontier. In particular, the higher is the prior mean and prior variance, the higher is the value function.

because new information on θ is unlikely to have a large impact on the posterior mean of θ . Second, the expected loss of output due to incorrect future job assignments is small because the likelihood that θ is much different than μ_t is small.

Recall that increases in α_t increase the spread of μ_{t+1} , the posterior mean of θ . As is clear from Equation 6, however, the variance of posterior mean of θ is also increasing in σ_t^2 , thus, the effect of an increase in α_t on the spread of μ_{t+1} is small when σ_t^2 is large. When the prior variance is very large (large signal to noise ration), the weight on the signal is high and the beliefs are updated a lot. An increase in experimentation increases the ratio to noise signal, but the marginal increase is small when this ratio is already very large, leading to low value of the additional accuracy in information generated by increase in experimentation intensity. Thus, the increase in experimentation and the resulting information generated has little value as it does not affect the optimal job assignment.

This result is closely related to the nonconcavity result in Radner and Stiglitz (1984). That is, when there is no information (the variance of prior mean is arbitrarily large), the marginal value of information may be very small, as it does not affect the optimal decision for any given signal.

Naturally, the effect of α_t on expected future productivity, also influences the value of experimentation, and as Proposition 3 makes clear, holding expected productivity fixed, the optimal level of experimentation may increase and then decrease as uncertainty about θ falls.

Proposition 3 *For a fixed value of μ_t , when $N > 2$, the optimal choice of α_t may increase and then decrease as σ_t^2 falls.*

Whereas it is possible to show that if α_t is continuous, then the increase in the continuation value from a marginal increase in α_t approaches zero as σ_t^2 approaches infinity, and that α_t approaches zero as σ_t^2 approaches zero we are not able to prove this result analytically in our discrete job model. However, Figures 3 and 4 demonstrate that

it is true.

In our model, one can think of jobs that reveal a lot of information about skill but in which a worker’s expected productivity is low as “risky jobs”, and jobs that reveal very little about a worker’s skill but in which a worker’s expected productivity is high as “safe jobs”. Thus, since uncertainty about θ will fall as workers gain experience, one interpretation of Proposition 3 is that young workers will start out in relatively safe jobs, move toward risky jobs as they gain experience and then return to safe jobs as uncertainty about their skill set dissipates.

As mentioned above, this finding seems to directly contradict Miller (1984) whose results suggest that young (high-variance) workers will gravitate towards risky jobs. However, it turns out that our results complement rather than contradict those of Miller. Recall that Miller’s model differs from ours in that in his model workers can choose whether to learn about their skill, but they cannot choose how much information to gather.⁹ That is, α is restricted to be either zero or one. In contrast, in our model, workers can vary the amount of information they receive about θ by selecting jobs with a range of different values for α . As it turns out, the distinction between a worker’s choice of whether to experiment and a worker’s choice about how much to experiment is important.

To see this, consider again Figure 3. Recall that in this figure $\alpha \in \{0, 0.5, 1\}$. Thus, workers can choose both whether to experiment (whether to pick $\alpha = 0$ or $\alpha > 0$) and how much to experiment (whether to pick $\alpha = 0.5$ or $\alpha = 1$). Notice that the negatively sloped frontier between $\alpha = 0$ and $\alpha = 0.5$ implies that the likelihood of experimentation falls as σ_t^2 falls. Thus, like Miller (1984), we find that workers are less likely to experiment when as uncertainty about θ falls. The frontier between $\alpha = 0.5$ and $\alpha = 1$ in Figure 3 also implies, however, that the optimal *level* of experimentation may actually increase as σ_t^2 falls. Thus, while the likelihood of experimentation declines as uncertainty about

⁹For any given occupation, workers can only choose whether to learn about θ , but not how much to learn.

θ falls, our model also shows that the optimal *level* of experimentation does not vary monotonically with uncertainty about θ . Put differently, although workers are more likely to select jobs that involve risk when they are young, the *level* of risk may initially be small, increase as workers gain experience and then ultimately decline.

Interestingly, these results also highlight an important distinction between traditional human capital investments, which increase skills, and investments in information, which increase knowledge of skills. In particular, since experimentation means sacrificing current period earnings to increase future productivity, it can be thought of as a type of human capital investment, where the optimal level of investment is determined by the prior mean and variance of the skill distribution. However, in standard models of human capital investment in which workers invest in skill rather than information (i.e., Mincer (1958), Ben-Porath (1967)), investment primarily takes place at the beginning of a worker’s career and then declines as the worker gains experience. In contrast, in our model, optimal investment levels are initially small, reach their peak in the middle of a worker’s career and then ultimately decline. Thus, while investments in skill and investments in information are conceptually similar, our paper highlights that they may follow different patterns over the life-cycle.

5.1 Job Assignments Over the Life-Cycle

Having discussed how experimentation changes over the life-cycle, it is also worth directly examining our model’s implications for job choices. To do so, we simulate job choices over time, for workers for whom $\mu_0 = -1$ and $\sigma_0^2 = 9$ in a model where $\delta = 0.9$, $k = 0$, $\sigma_\epsilon^2 = 1$, $\alpha \in \{0, 0.5, 1\}$. As Figure 5 reveals, given this initial mean and variance, all workers start out with $\alpha = 0.5$. Note, however, that this is a transitory job. As workers gain experience and become more certain about whether θ is greater than or less than k , they increasingly sort into specialized jobs in which $\alpha = 1$ or $\alpha = 0$.¹⁰

¹⁰Figure 5 does not reveal the extent of experimentation since some workers who select $\alpha = 1$ do so because they wish to experiment, while others do so because $\alpha = 1$ maximizes their expected current

Proposition 4 *At the beginning of the life-cycle, workers may work in jobs which include a mixture of tasks ($0 \leq \alpha \leq 1$). As they accumulate experience, they sort into more specialized jobs, and in the limit, workers either choose $\alpha = 1$ or $\alpha = 0$.*

The proof is in Appendix A. The intuition for this proposition is clear. The first part follows directly from the optimal solution (see, for example, Figure 3), and the second part of the proposition follows from the fact that as t becomes arbitrarily large, σ_t^2 becomes close to zero for workers who are assigned to $\alpha_t > 0$ (once assigned to $\alpha_t = 0$, workers do not move). Thus, for these workers, the solution approaches the full information solution in which $\alpha_t = 1$ if $\theta > k$, and $\alpha_t = 0$ otherwise.

Since no information is revealed about skill when $\alpha = 0$, however, there will always exist a subset of workers for whom θ is never fully learned, even in the limit.¹¹ As a result, not all job assignments will be efficient. In particular, every period some workers choose $\alpha = 0$ even though $\theta > k$ because of negative prior beliefs about θ . Further, these mistakes are permanent since once a worker chooses $\alpha = 0$, no new information will be revealed about θ and the worker will never learn about his or her true ability. As a result, even good workers can get stuck in bad jobs. As Table 1 reveals, the cumulative probability of finding a good worker in a bad job increases over time since there is always some chance that workers will experience a series of negative productivity shocks that lead them to revise downwards their beliefs about θ .

5.2 Job Transitions

Our model also delivers novel predictions about job transitions. In standard learning and matching models, job transition rates decline with experience. However, as the above discussion suggests, in our model, many workers will begin their careers in transitory jobs. Since no worker will stay in these jobs for very long, for certain parameter values, transition rates out of these jobs may be higher for more experienced than less experienced period output.

¹¹Thus, in the language of Aghion et al. (1991), learning is not adequate.

workers. Proposition 5 clarifies this result.

Proposition 5 *Consider two workers, indexed by n and m , with $\mu_{t,n} = \mu_{t,m}$. Suppose that worker m is more experienced than worker n so that $\sigma_{t,n} > \sigma_{t,m}$, and suppose that both workers are optimally assigned to $0 < \alpha_t < 1$. The transition probability of the more experienced worker may be higher than the transition probability of the less experienced worker.*

Although we cannot prove this result analytically, it is a direct implication of the optimal assignment policy, and Table 2 illustrates this point. Table 2 simulates the transition probabilities for two workers with the same initial prior beliefs about θ ($\mu_t = -0.7$), but one worker is “inexperienced” ($\sigma_t = 5$) while the other is “experienced” ($\sigma_t = 1$). For both of these workers, current period output would be maximized if they selected $\alpha = 0$. As the table shows, however, both workers choose to forgo current period output and instead pick $\alpha = 0.5$ in order to gain information about their unobserved skill. Interestingly, however, in the next period, the more experienced worker is *more* likely to change jobs than the less experienced worker. This result is surprising since in most models of learning and sorting, transition probabilities decline as workers gain experience. In contrast, in this example, the inexperienced worker stays in the job that reveals information about his or her skill, while more the experienced worker stops experimenting and moves into the safe job in which $\alpha = 0$. Thus, in this model, additional information and the concomitant decline in the option value of experimentation, may increase the likelihood that workers change jobs as they gain experience.

We emphasize, however, that this increasing likelihood of a job change only holds for workers who are currently in intermediate-level jobs and for a narrow range of parameter values. Eventually, job transition rates in our model will decline because, as workers gain experience, new information has an increasingly small effect on beliefs and, as a result, job transition rates will fall. This is important since empirical evidence suggests that job transition rates tend to decline with experience. Nonetheless, our model can account for

the fact that the empirical hazard functions for job change are not strictly monotonic. Empirical hazard functions computed at monthly frequencies, for example, are typically found to increase in the first several months of employment before eventually declining (see, for example, Farber (1999) and Pavan (2007)).

5.3 Experience-Earnings Profiles

The non-monotonic between experience and experimentation means that wage growth in our model is driven by two factors. First, wages grow as workers gain information about the true value of θ and sort into the job at which their expected productivity is the highest. Second, since experimentation involves a loss in expected current period wages, wage growth later in a worker's career is also driven by the eventual decline in experimentation.

To show our model's predictions for wage growth, we simulate the wage distribution across time assuming $\delta = 0.9$, $k = 1$, $\sigma_\epsilon^2 = 1$, $\alpha \in \{0, 0.5, 1\}$, $\mu_0 = -1$, $\sigma_0^2 = 9$. Figure 6 shows the wage at different percentiles of the wage distribution for a given cohort of workers 10 periods into the future. Similar to most models of learning, Figure 6 shows that cohort wage dispersion increases over time. In contrast to standard models, however, our model also predicts that the bottom tail of the earnings distribution (any wage less than k) disappears as workers stop experimenting.

5.4 The Importance of Jobs

Our model suggests that economy-wide output is higher the larger is the number of jobs. To see this, Figure 7 shows the difference in the expected discounted value of lifetime earnings between the three-job model and the two-job model. As the figure reveals, this difference is always positive. Thus, workers can expect to earn more (and produce more) when there are more jobs to choose from.

As the number of jobs increases, expected earnings increases for two related reasons. First, workers are more likely to experiment. Second, among those who experiment,

the cost of experimentation is lower. To see this, Figure 8 compares job choices in the specialized economy in which $\alpha_t \in \{0, 1\}$ with job choices in the three-job economy in which $\alpha_t \in \{0, 0.5, 1\}$. The black area shows the region in which workers choose $\alpha_t = 0$ in the two-job model but $\alpha_t = 0.5$ in the three-job model. Thus, in an economy with more jobs, workers are more likely to choose jobs which reveal information about θ . In addition, the gray area shows the region in which workers choose $\alpha_t = 1$ in the two-job model but $\alpha_t = 0.5$ in the three-job model. Thus, compared to the two-job model, in the three-job model, workers

sacrifice less in terms of foregone current earnings in order to learn. The fact that expected lifetime earnings increase as the number of jobs in the economy increases implies that productivity may be lower in specialized economies where on-the-job performance depends only on one skill than in diverse economies where jobs draw upon multiple dimensions of a worker's skill set.

5.5 The Importance of Early Labor Market Outcomes

Finally, our model suggests that i.i.d. productivity shocks, especially those early in a worker's career, have a persistent effect on earnings. It is a common feature of all learning models that past output realizations affect current beliefs about workers' skills, and on average, workers who receive positive productivity shocks ($\epsilon_t > 0$) will have higher wages than those who receive negative productivity shocks ($\epsilon_t < 0$), at least for any finite time horizon. In contrast with the previous literature, however, in our model, negative productivity shocks have longer-lasting effects on wages than do positive productivity shocks because workers who receive negative productivity shocks are more likely to choose jobs in which α_t is relatively low and in which little new information is revealed about the workers' skills. Since information helps workers optimally sort into jobs, workers who receive negative productivity shocks will not only have lower wages but also slower wage growth than will those who receive positive productivity shocks. Thus, experimentation serves as a propagation mechanism. Further, this effect is especially

pronounced early in a worker's career since new information has the largest effect on beliefs when there is considerable uncertainty about θ .

The second column of Table 1, which shows the probability of being incorrectly assigned to $\alpha = 0$ when $\theta > k$ in each period, starkly illustrates of this point. As the table reveals, the probability of these incorrect assignments is the greatest early in a worker's career. Further, these mistakes are permanent because once a worker is assigned to the low-level job, no additional information about that worker's skill is revealed and wages cease to grow. Thus, a worker's labor market experiences, especially early on, can have permanent impact on his or her career trajectory. This suggests that the period of time shortly following an individual's entry into the labor market may be critical to his or her future success. Similarly, our model offers an explanation for why workers who have identical observable and unobservable characteristics (i.e. same k and same θ) may have different career paths.

6 Empirical Support for Experimentation

In order to isolate the effect of experimentation, our model abstracts from human capital accumulation, asymmetric information and a number of other factors known to influence wages and career paths. As a result, we do not expect our model to present a complete explanation for career paths and wage dynamics.¹² Nonetheless, it is important to ask whether our model is broadly consistent with the stylized facts. Indeed, most of our model's predictions match known empirical patterns: wages grow with experience, wage growth primarily occurs in the early stages of workers' careers, cohort wage dispersion increases over time, job transition rates decline with experience and the wage distribution is positively skewed. Of course, other models also are able to capture these trends. Our model, however, also is able to account for a number of patterns in the data that existing models have difficulty explaining.

¹²Papers by Demougin and Siow (1996), Gibbons and Waldman (1999), Gibbons and Waldman (2006) more directly take on the task of accounting for broad patterns in the data.

To demonstrate this, we draw heavily upon results from Baker, Gibbs and Holmstrom (1994a and 1994b), who analyze wage dynamics and job transitions over a twenty year period (1968-1988) at a single firm in the service industry. These data are particularly well-suited for examining our model's predictions. First, the data focus on managerial employees for whom job-relevant skills are likely to be hard-to-observe and for whom on-the-job performance is likely to be important in determining career paths. Second, because Baker, Gibbs and Holmstrom (hereafter, BGH) are able to identify the hierarchical structure of jobs within the firm, job transitions are relatively easy to identify. This contrasts with many standard data sets in which job transitions are hard to classify, especially since, as our model highlights, promotions may not always entail pay increases and demotions may not always entail pay cuts. Finally, the data are longitudinal and include salary information. As a result, it is possible to examine whether both career paths and wage dynamics are consistent with experimentation.

BGH divide the hierarchical structure of jobs within the firm they examine into 5 broad levels, where employees typically start at Level 1 and work their way up. BGH compare promotion patterns for workers who enter Level 2 by being promoted from Level 1 to those who are hired into Level 2 from outside the firm.¹³ Interestingly, those hired into Level 2 from outside the firm are typically promoted more quickly into Levels 3 and 4 than are incumbent workers. The rapid promotion rate of the outside hires, however, does not last. After 5 years with the company, those hired into Level 2 from outside the firm are, on average, in lower-level positions than incumbent workers. In addition, outside hires are more likely to be demoted than incumbent workers. BGH suggest, therefore, that the firm may be using the hierarchical structure of the firm to learn about the ability level of outside hires.

In the context of our model, these patterns are consistent with experimentation. Although we cannot observe the variance of the firm's prior beliefs for new hires versus incumbents, BGH report that, on average, outside hires have more labor market expe-

¹³See Table III on page 899 in BGH 1994a.

rience than incumbents, suggesting that there should be more information about the general human capital of outside hires than incumbents. Given the non-monotonic relationship between experimentation and experience, our model thus predicts that outside hires may be placed in higher-level jobs that reveal more information about their skills than similarly able incumbent workers. This may partly explain why those hired into Level 2 from outside the firm were not instead hired into Level 1, despite the fact that their long-term career outcomes within the firm imply that they have lower average ability than incumbent workers. In addition, our model also predicts that the transition rate out of intermediate-level jobs may be higher for more experienced than less experienced workers and, thus our model can also account for the fact that job transition rates (which include both promotions and demotions) are initially higher for outside hires than incumbents.

Experimentation is also consistent with BGH's finding that wage increases are serially correlated. According to our model, workers who receive positive productivity shocks are not only more likely to experience wage increases but they are also more likely to enter into jobs that quickly reveal information about their skills. As a result these workers are likely to experience relatively rapid future wage growth, thus leading to serial correlation.

Although the serial correlation in wages is also consistent with the hypothesis that workers are heterogeneous and that able workers accumulate human capital more quickly than less able workers, a pure human capital accumulation model has a hard time explaining the consistent pattern of real salary declines found by BGH.¹⁴ In particular, BGH report that for at least the bottom 25 percent of the salary distribution, real salaries decline in the first few years of employment and often do not recover for many years. Among the 1975 cohort of entrants into Level 1, for example, 15 percent were still below their real starting salary after 10 years of employment. Further, these real salary declines are the most pronounced for those at the bottom end of the salary distribution. Both of these regularities are consistent with our model, which predicts that workers,

¹⁴See Figure III on page 941 in BGH 1994b.

especially those with low expected ability, may be assigned to jobs that do not maximize their current-period expected output. Rather, they may be assigned to jobs that quickly reveal information about skill. The result is that wages for these workers may actually decline. As firms accumulate information, however, this experimentation will ease, and wages will again increase. The result is that experimentation may lead to a temporary dip in real wages. Figure 6, for example, which plots predicted wages over time in our model for workers at various percentiles of the wage distribution, reveals a pattern of real wage declines for those at the bottom 10 percent of the wage distribution that bears a remarkable similarity to the patterns found in BGH.

Finally, as mentioned above, our model implies that early career outcomes not only have long-lasting effects on wages but also on wage growth. As a result, experimentation can serve as a propagation mechanism. This feature of our model may partially explain the lasting cohort salary differentials found in BGH. In particular, BGH note that differences across cohorts in initial salary do not disappear over time. Rather, cohort wages appear to grow in parallel.¹⁵ Moreover, BGH note that cohort-level differences in starting salaries are not easily explained by changes in the composition of workers over time.¹⁶ Rather, starting salaries seem to follow industry-level trends and seem likely to reflect broader market conditions at the time workers enter the firm.¹⁷

In the context of our model, one possible explanation for the persistence of these cohort wage differentials is that labor market conditions at the time of entry effect initial job assignments and, as a result, the rate at which future wages grow. This view is bolstered by the fact that differences in starting salary are related to differences in initial job assignments. In particular, when starting salaries are low, most workers enter the firm at Level 1, and during periods when the starting salary is relatively high, the fraction of workers who enter the firm at Level 2 and higher goes up. If low-level jobs

¹⁵See Figure II on page 934 in BGH 1994b.

¹⁶Though it is hard to rule out the possibility that cohorts are changing along unobservable dimensions.

¹⁷Beaudry and DiNardo (1991), Devereux (2002), Kahn (2006) and Oyer (2006) also find evidence of cohort effects.

reveal less information about workers' skills than do high-level jobs, then the salary differences associated with initial assignments will persist. A handful of other papers also propose explanations for cohort effects. Of these, our model's explanation is the most similar to Devereux (2002), who argues that the rate of human capital accumulation may differ across jobs and to Gibbons and Waldman (2006), who argue that workers may accumulate job-specific human capital. Our explanation differs from these other explanations primarily in that we assume that there are differences across job in the rate at which workers learn about their skills rather than differences in the actual accumulation of those skills.

7 Conclusion

There exist very few papers that examine experimentation in the labor market. In order to help fill this gap, in this paper, we develop a dynamic model of job choice in which some jobs provide more information about workers' skills than others. Surprisingly, our model predicts that young, experienced workers will not always choose the job that provides the most information about their skill. Instead, we find that the optimal level of experimentation is initially low, increases as workers gain experience and then eventually falls as uncertainty about workers' skills is resolved. To our knowledge, no other model of job choice predicts this non-monotonicity in experimentation.

In addition, our model delivers novel predictions for both job transitions and wage growth. In particular, our results suggest that some jobs are transitory and that, for some parameter values, transition rates out of these jobs may actually be higher for more experienced workers than for less experienced workers. In addition, we show that wage growth over the life-cycle is partly driven by the eventual decline in experimentation. Further, our model predicts that a diverse range of jobs increases expected lifetime earnings, suggesting that highly specialized economies in which jobs make use of only one skill may be suboptimal in an environment in which on-the-job performance reveals in-

formation about workers' skills. We also find that random productivity shocks can have long-lasting effects on both wages and wage growth, especially early in a worker's career. This suggests that a worker's initial experiences in the labor market can have long-lasting and even permanent effects on his or her career trajectory. It also suggests that luck can lead otherwise identical workers to experience sharply different career paths. Finally, we show that our model's predictions are broadly consistent with well-known trends in the data, even those that are not easily accounted for by existing models.

Appendix A

Proof of Proposition 1

1. Competition and spot contracts imply that workers will never accept a wage w_t such that $w_t < E(y_t)$ and that firms will never offer a wage w_t such that $w_t > E(y_t)$. Thus, $w_t = E(y_t)$ for all t .
2. This follows directly from the fact that workers choose α_t to maximize the present discounted value of their lifetime earnings and $w_t = y_t$ for all t .
3. Suppose that instead of workers announcing job choices and firms making wage offers, firms simultaneously offer (α_t, w_t) and workers either accept or reject. Let α^* denote the solution to

$$\max_{\alpha_t \in \{0, \dots, 1\}} \alpha_t \mu + (1 - \alpha_t)k + \delta E_t[V(\mu_{t+1}, \sigma_{t+1})|\alpha_t]$$

Suppose a firm offers a worker (α', w') where $\alpha' \neq \alpha^*$. A worker will accept α' if and only if

$$w' + \delta E_t[V(\mu_{t+1}, \sigma_{t+1})|\alpha'] > \alpha^* \mu + (1 - \alpha^*)k + \delta E_t[V(\mu_{t+1}, \sigma_{t+1})|\alpha^*].$$

Optimality implies that

$$\alpha' \mu + (1 - \alpha')k + \delta E_t[V(\mu_{t+1}, \sigma_{t+1})|\alpha'] < \alpha^* \mu + (1 - \alpha^*)k + \delta E_t[V(\mu_{t+1}, \sigma_{t+1})|\alpha^*].$$

Thus, (α', w') will be accepted if and only if $w' > \alpha' \mu + (1 - \alpha')k$. However, if firms make this wage offer, then their expected profits will be negative since workers will never accept any future wage offers less than their expected productivity.

Proof of Proposition 2

When $\alpha_t \in \{0, 1\}$, our model is a special case of the model in Miller (1984). Further, Miller's results establish that $V(\mu_t, \sigma_t^2)$ is strictly increasing in both arguments. Thus, since $V(\tilde{\mu}, \tilde{\sigma}^2) = \frac{k}{1-\delta}$ by assumption, $\tilde{\mu}$ must be decreasing in $\tilde{\sigma}^2$. Further, if $\tilde{\sigma}^2 = 0$, then workers will choose $\alpha = 1$ if $\theta > k$ and $\alpha = 0$ otherwise. Thus, $\tilde{\mu} = k$.

Proof of Proposition 3

This proposition is demonstrated in Figures 3 and 4.

Proof of Proposition 4

1. This result is demonstrated in Figures 3 and 4.
2. Suppose that in time t it is optimal for a worker to choose α' where $0 < \alpha' < 1$ and stay in that job forever. We know that at time $t + n$

$$\mu_{t+n} = \frac{\mu_t \sigma_{\bar{\epsilon},t}^2 + \sigma_t^2 \sum_{i=0}^{n-1} x_{t+i}}{\sigma_{\bar{\epsilon},t}^2 + n\sigma_t^2}.$$

Further, by the Law of Large Numbers, we know that

$$\lim_{n \rightarrow \infty} \mu_{t+n} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} x_{t+i} = \lim_{t \rightarrow \infty} \theta + \frac{1}{n} \sum_{i=0}^{n-1} \frac{\epsilon_{t+i}}{\alpha'} = \theta.$$

Thus, the worker's problem approaches the benchmark case in which θ is known. However, when θ is known, workers will optimally choose $\alpha = 1$ if $\theta > k$ and $\alpha = 0$ otherwise. Thus, in the limit, it cannot be optimal for the worker to stay in a job in which $0 < \alpha' < 1$.

Proof of Proposition 5

This proposition is demonstrated in Table 2.

Appendix B

The mapping T satisfies Blackwell's sufficient conditions.

- i. Monotonicity: It is clear by examination of (8) that if $g'(\mu, \sigma^2) > g(\mu, \sigma^2)$, then $Tg'(\mu, \sigma^2) > Tg(\mu, \sigma^2)$.
- ii. Discounting: It can be seen from (8) that $[T(g + c)](\mu, \sigma^2) = Tg(\mu, \sigma^2) + \delta c$. Since $\delta \in (0, 1)$, discounting is satisfied.

References

- [1] Aghion, Philippe; Bolton, Patrick; Harris, Christopher and Jullien, Bruno. “Optimal Learning by Experimentation,” *Review of Economic Studies*, July 1991, 58(4), pp. 621-654.
- [2] Baker, George; Gibbs, Michael and Holmstrom, Bengt. “The Internal Economics of the Firm: Evidence from Personnel Data,” *Quarterly Journal of Economics*, November 1994, 109(4), pp. 881-919.
- [3] Baker, George; Gibbs, Michael and Holmstrom, Bengt. “The Wage Policy of a Firm,” *Quarterly Journal of Economics*, November 1994, 109(4), pp. 921-55
- [4] Beaudry, Paul and DiNardo, John. “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data,” *Journal of Political Economy*, August 1991, 99(4), pp. 665-88.
- [5] Ben-Porath, Yoram. “The Production of Human Capital and the Life Cycle of Earnings,” *The Journal of Political Economy*, August 1967, 75(4), pp. 352-365.
- [6] Bolton, Patrick and Harris, Christopher. “Strategic Experimentation,” *Econometrica*, March 1999, 67(2), pp. 349-374.
- [7] Chade, Hector and Schlee, Edward. “Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information,” *Journal of Economic Theory*, December 2002, 107(2), pp. 421-452.
- [8] Demougin, Dominique and Siow, Aloysius. “Careers in Ongoing Hierarchies,” *American-Economic-Review*, December 1994, 84(5), pp. 1261-77.
- [9] Devereux, Paul. “The Importance of Obtaining a High-Paying Job,” *mimeo*, University College Dublin, 2002.

- [10] Farber, Henry. "Mobility and Stability: The Dynamics of Job Change in Labor Markets," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*. Elsevier, 1999, pp. 2439-2483.
- [11] Felli, Leonardo and Harris, Christopher. "Learning, Wage Dynamics, and Firm-Specific Human Capital," *Journal of Political Economy*, August 1996, *104*(4) pp. 838-868.
- [12] Gibbons, Robert and Waldman, Michael. "A Theory of Wage and Promotion Dynamics Inside Firms," *Quarterly Journal of Economics*, November 1999, *114*(4), pp. 1321-58.
- [13] Gibbons, Robert and Waldman, Michael. "Enriching a Theory of Wage and Promotion Dynamics Inside Firms," *Journal of Labor Economics*, January 2006, *24*(1), pp. 59-107.
- [14] Gittins, J. and Jones, D. "A Dynamic Allocation Index for Sequential Design of Experiments," in Joseph Gani, Karoly Sarkadi and Istvan Vincze, eds., *Progress in Statistics*. Amsterdam: North-Holland, 1974, pp. 161-173.
- [15] Harris, Milton and Holmstrom, Bengt. "A Theory of Wage Dynamics," *Review of Economic Studies*, July 1982, *49*(3), pp. 315-33.
- [16] Jovanovic, Boyan. "Job Matching and the Theory of Turnover," *Journal of Political Economy*, October 1979, *87*(5), pp. 972-990.
- [17] Jovanovic, Boyan and Nyarko, Yaw. "Stepping Stone Mobility," *Carnegie Rochester Conference Series on Public-Policy*, June 1997, *46*(0), pp. 289-325.
- [18] Kahn, Lisa, "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy" *mimeo*, Harvard University, 2006.
- [19] Keller, Godfrey and Rady, Sven. "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, July 1999, *66*(3), pp. 475-507.

- [20] MacDonald, Glenn. "A Market Equilibrium Theory of Job Assignment and Sequential Accumulation of Information," *American Economic Review*, December 1982, 72(5), pp. 1038-55.
- [21] Miller, Robert. "Job Matching and Occupational Choice," *Journal of Political Economy*, December 1984, 92(6), pp. 1086-1120.
- [22] Mincer, Jacob. "Investment in Human Capital and Personal Income Distribution," *The Journal of Political Economy*, August 1958, 66(4), pp. 281-302.
- [23] Moscarini, Giuseppe and Smith, Lones. "The Optimal Level of Experimentation," *Econometrica*, November 2001, 69(6), pp. 1629-1644.
- [24] Neal, Derek. "The Complexity of Job Mobility Among Young Men," *Journal of Labor Economics*, April 1999, 17(2), pp. 237-261.
- [25] Ortega, Jaime. "Job Rotation as a Learning Mechanism," *Management Science*, October 2001, 47(10), pp. 1361-1370.
- [26] Oyer, Paul. "Initial Labor Market Conditions and Long-Term Outcomes for Economists," *Journal of Economic Perspectives*, Summer 2006, 20(3), pp. 143-60.
- [27] Pastorino, Elena. "Career Mobility Under Uncertainty and Learning," *mimeo*, University of Iowa, August 2006.
- [28] Pavan, Ronni. "The Role of Career Choice in Understanding Job Mobility," *mimeo*, University of Rochester, July 2007.
- [29] Radner, Roy and Stiglitz, Joseph. "A Nonconcavity in the Value of Information," in M. Boyer and R. Kihlstrom, eds., *Bayesian Models of Economic Theory*, Elsevier, 1984, pp. 33-52.
- [30] Rustichini, Aldo and Wolinsky, Asher. "Learning About Variable Demand in the Long Run," *Journal of Economic Dynamics and Control*, July-September 1995, 19(5-7), pp. 1283-92.

Figure 1: Value Function

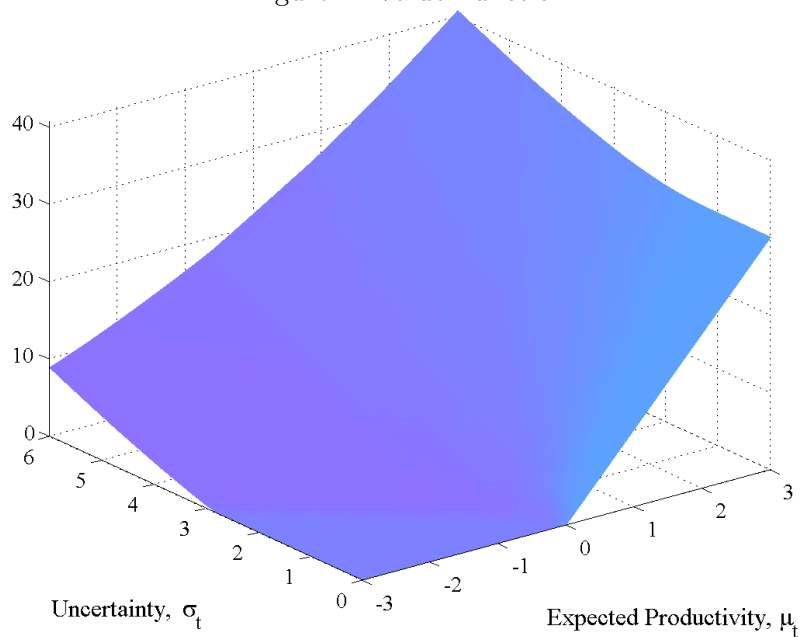


Figure 2: Optimal Job Choice: Two-Job Model

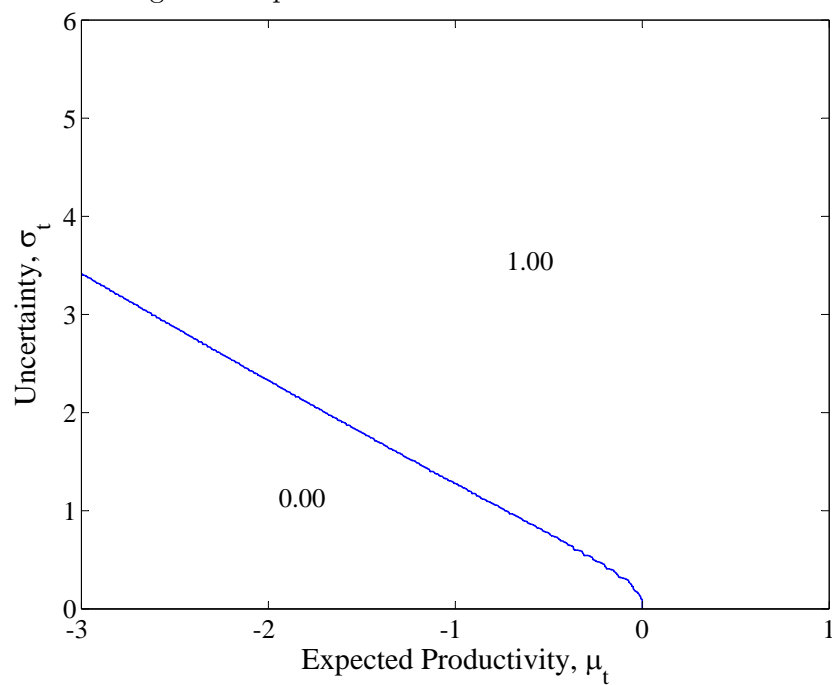


Figure 3: Optimal Job Choice: Three-Job Model




Figure 4: Optimal Job Choice: Five-Job Model




Figure 5: Job Assignments Over the Life-Cycle

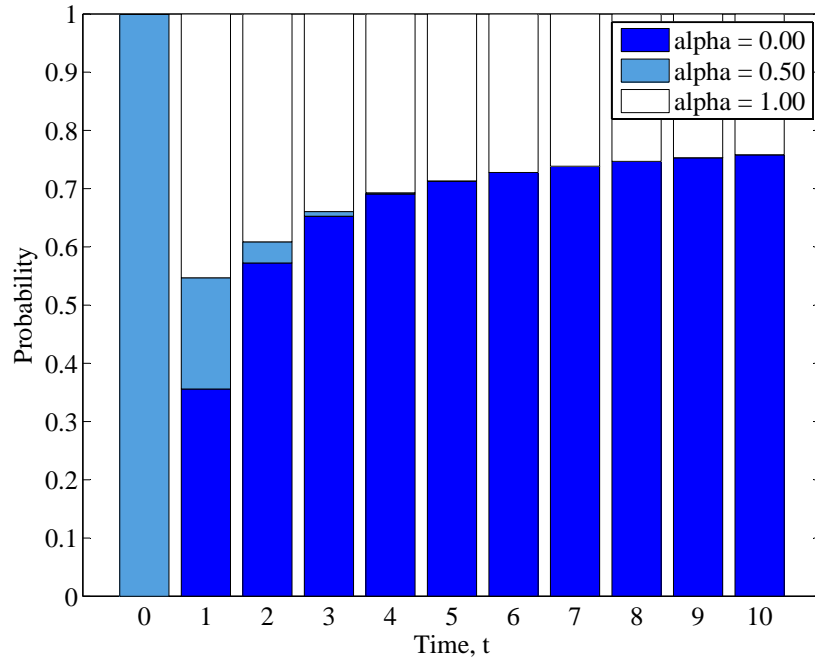


Figure 6: Percentiles of the Wage Distribution




Figure 7: Difference in Value Function Between a Three-Job and Two-Job Model

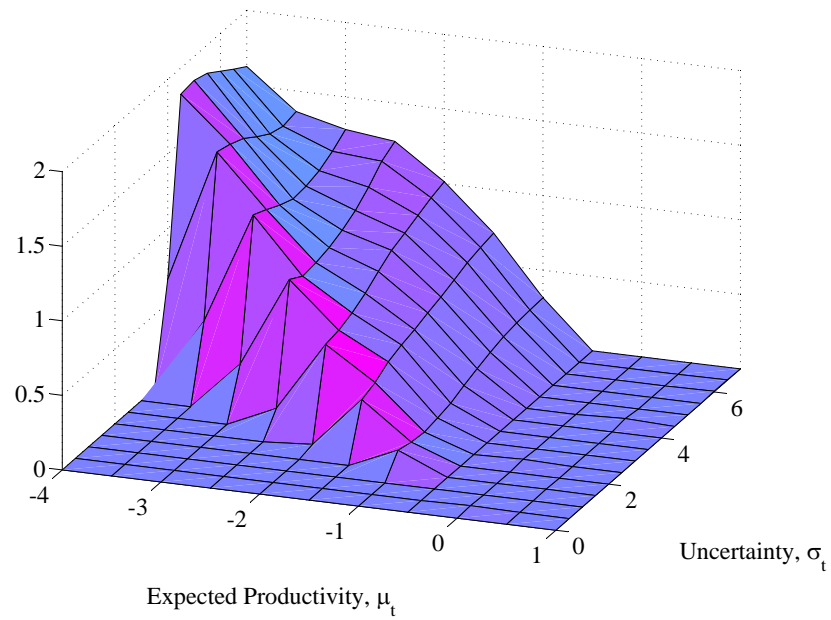


Figure 8: Comparison Between Job Choices in the Three-Job and Two-Job Model

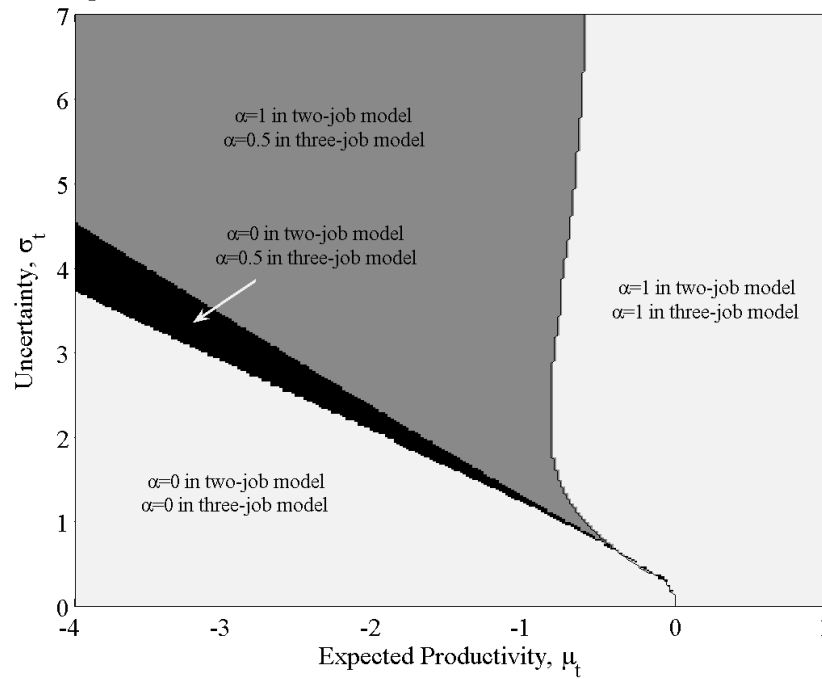


Table 1: Cumulative Probability of Permanent Mistakes

Time Period	Probability $\alpha=0$ and $\theta>k$	Marginal Change
$t=1$	0	
$t=2$	1.50%	1.50%
$t=3$	3.11%	1.61%
$t=4$	4.55%	1.44%
$t=5$	5.83%	1.28%
$t=6$	6.98%	1.15%
$t=7$	8.00%	1.03%
$t=8$	8.94%	0.94%
$t=9$	9.81%	0.87%
$t=10$	10.59%	0.78%

Table 2: Transition Probabilities for Inexperienced vs. Experienced Workers

	Inexperienced Workers ($\sigma=5$)			Experienced Workers ($\sigma=1$)		
	$\alpha=0$	$\alpha=0.5$	$\alpha=1$	$\alpha=0$	$\alpha=0.5$	$\alpha=1$
Probability of Assignment $t=0$	0%	100%	0%	0%	100%	0%
Probability of Assignment $t=1$	30%	17%	53%	51%	7%	42%