Occupational Choice, Human Capital and Learning: A Multi-Armed Bandit Approach*

Rafael Lopes de Melo† and Theodore Papageorgiou‡

University of Edinburgh and Boston College

August 2025

Abstract

This paper introduces a model of worker matching at the occupation level. In our setup, young workers, while employed in an occupation, accumulate human capital and also learn about their underlying productivity in that occupation. Human capital is partially transferable to other occupations and similarly, the information acquired in one occupation is useful for the worker's productivity elsewhere. Workers with low tenure levels, as well as low-paid workers, are the ones most likely to switch occupations, consistent with our empirical findings. Though the model is quite general, we show that Gittins indices can be used in this setup to preserve tractability. We discuss potential applications ranging from assessing the impact of AI and automation to the evaluation of policies such as unemployment benefits, sector-specific subsidies, or minimum wages.

Keywords: Human Capital, Occupations, Multi-armed Bandits, Worker Mobility, Learning, Information and Human Capital Spillovers, Wage Inequality, Gittins Index.

JEL Classification: J24, J31, J62

*We are grateful to Mike Dickstein, Fabian Lange, Zachary Mahone, Ariel Pakes, Elena Pastorino, Sven Rady, Anna Zaharieva and seminar participants for useful comments.

†University of Edinburgh, Department of Economics, 31 Buccleuch Pl, Edinburgh, EH8 9JS, United Kingdom, rafael.lopes@ed.ac.uk

‡Boston College, Department of Economics, Maloney Hall 332, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

1 Introduction

The substantial literature that has investigated the labor market outcomes of young workers has emphasized the importance of three factors: a) on-the-job human capital accumulation, b) selection and matching (e.g. occupational or job choice), and c) learning about unobserved productivity.[1] The most common approach is to investigate these mechanisms separately, often with simplifying assumptions, such as e.g. that all occupations are identical or that human capital is not transferable.[2] These restrictions have been necessary to construct tractable frameworks that are amenable to estimation. Nonetheless, they ignore salient features of the data, such as e.g. the path dependence of labor market outcomes or heterogeneity across occupations. These empirical patterns however may be key in understanding the impact of labor market shocks, such as the effects of automation and AI or the impact of trade shocks, or labor market policies.

This paper contributes to this literature in two ways. First, we document several facts relating to wages and occupational switching. These facts improve our understanding of the process through which workers allocate to different occupations. Second, motivated by these facts, we construct a model of occupational choice that has the following features: it allows for human capital accumulation (learning-by-doing), that is partially specific to the occupation and partially transferable. Workers have imperfect information about their ability in each occupation and learn about it from their performance in their current occupation. We allow for informational spillovers, so the information acquired in one occupation is useful about a worker's match in other occupations. The model features substantial heterogeneity, both on the worker, as well as on the occupational side. Despite the number of moving parts, we show that Gittins indices, most commonly employed in models with only information frictions, can in fact be used in this much more general setup to preserve tractability.

More specifically, we use data from the National Longitudinal Survey of Youth 1997 (NLSY97) where we focus on the nine major occupational groups. We first illustrate that there is substantial heterogeneity across occupations both in wage levels, as well as in how wage increases with tenure. In addition, we find that the variance of wages increases as a function of occupational tenure. We also confirm an older finding that occupational mobility declines with occupational tenure. In addition, we show that wages fall (or increase less) before switching occupations.

We then revisit how occupational mobility varies depending on the worker's wage rank within their

[1] See Rubinstein and Weiss (2006) for a survey, as well as the literature review in Section 2.

[2] Human capital is typically assumed to be either fully specific (to a job or occupation) or fully general.

occupation. The lowest-paid workers in an occupation are the ones more likely to switch to another one. The propensity to switch falls as we consider workers with higher wages and the highest-paid workers in an occupation are the ones who are least likely to switch. This pattern holds for every one of the nine major occupations. In addition, the same pattern emerges when using data from the 1996 panel of the Survey of Income and Program Participation (SIPP). This finding is related to the results of Groes et al. (2015) who use Danish data and document that occupational mobility is U-shaped: the lowest and the highest-paid workers in an occupation are most likely to switch, suggesting that as workers are revealed to be of higher ability they move up on a hierarchical occupational ladder. In contrast, we use U.S. data focusing on major occupational groups, and find a decreasing hazard with wages, which suggests that the occupation-specific match component is important: workers who are mismatched are paid less and are more likely to switch occupations, whereas well-matched workers are highly paid and are the least likely to switch.

These facts suggest that both occupation-specific matching and human capital accumulation are crucial ingredients for understanding the labor market outcomes of young workers. In particular, the high rates of occupational switching cannot be accounted for by human capital accumulation alone; learning, in the spirit of Jovanovic (1979), is perhaps the most natural explanation. Additionally, learning about occupational match quality is consistent with workers' wages falling before they switch occupations, and it can also explain why most occupational switching is driven by workers at the bottom of the within-occupation wage distribution.

Inspired by these facts, we introduce a model of occupational mobility that incorporates various forms of human capital accumulation, learning about unobserved productivities, as well as non-employment. Workers can work in only one occupation at a time. Each occupation belongs to a cluster, which contains groupings of occupations that use similar skill sets. For instance, journalists and public relations specialists might belong to the same cluster, since they both make heavy use of communication skills. While employed in an occupation, workers accumulate three forms of human capital: human capital that can be used in that occupation only, human capital that is transferable to other occupations within the same cluster, and general human capital that is fully transferable. In addition, workers learn about two components of their productivity, one that is specific to the occupation and one that is specific to the cluster, which captures learning spillovers across occupations. Continuing the previous example, if a worker realizes that they are well suited in journalism, then this is informative about how they may perform as a public relations specialist.

Our setup admits heterogeneity in several dimensions. In particular, occupations and clusters are allowed to differ in the importance of human capital, the rate of human capital accumulation, the uncertainty regarding the underlying match of each worker and also the speed at which workers learn about it.[3] Similarly, we allow for ex-ante heterogeneity in worker observable characteristics, which may interact with occupational heterogeneity. For example, we may have that in some occupations more educated workers may accumulate human capital faster compared to less educated workers, but not in others.

Every period each worker needs to decide between one of the occupations and non-employment. Their decision depends not only on their wage in each occupation but on dynamic considerations as well. In particular, when making their occupational choice, the worker takes into account a) this period's wage, b) human capital considerations, i.e. the different forms of human capital they will accumulate, c) learning considerations, i.e. the information they will receive about the different forms of unobserved productivity.

As the number of occupations/clusters grows past a handful, so does the state space of the worker, which includes their human capital levels and their beliefs. As a result, the curse of dimensionality quickly becomes binding. We are able to make progress as follows.

First, we show that Gittins indices are applicable in this context, despite it being much more general than the typical application of Gittins indices. In particular, almost all the uses of the Gittins index in economics have focused on models where the payoffs evolve exclusively due to learning (see discussion in Section 2). As noted however in Gittins et al. (2011), " ... the notion of learning is not really central to the story of the Gittins index and index theorem" (Chapter 2, page 29). We show that we can use this tool in this much more general setup where, in addition to information frictions, human capital accumulation is also a crucial determinant of worker compensation.

Second, we can relax two key restrictions of the Gittins index approach, namely the full independence across options and the absence of switching costs. For instance, we allow for the endogenous accumulation of specific human capital that would be lost following a switch, which the literature has argued may be a key component of switching costs. In addition, we allow for rich patterns of correlation across occupations, by grouping them in clusters, while preserving the tractability afforded by the Gittins index across clusters (see also Dickstein, 2021).

We then provide an illustrative calibration of the model. In the process, we provide a model-based approach to clustering occupations and we illustrate how the model's key mechanisms can be separately identified in the data. Moreover, we show how our theory can account for the empirical patterns we

[3]The idea of learning about match quality in the labor market dates back to Jovanovic (1979).

document in the data, including the negative relationship between the propensity to switch occupations and the worker's wage rank within their occupation.

Our setup is suitable both for assessing the effects of labor market shocks, as well as for policy evaluation. For instance, the labor market consequences of automation or a sectoral shock, such as a trade shock, depend on how transferable human capital is to new sectors and how informative previous experience is about potential new matches. Similarly, policies that may lead to lower employment can lead to less learning and worse matches, as well as lower human capital accumulation.

In the next section we discuss the related literature, while in Section 3, we document facts relating to wages, occupational switching, and tenure. Motivated by the empirical findings, in Section 4 we introduce our model of occupational choice and discuss how it remains tractable through the use of Gittins indexes. Section 5 discusses several extensions, while Section 6 concludes.

2 Related Literature

The debate about the importance of human capital accumulation vs. selection dates back to Abraham and Farber (1987), Altonji and Shakotko (1987) and Topel (1991) and focused mostly on the importance of firm-specific human capital.[4] Kambourov and Manovskii (2009b) using the methodology of Altonji and Shakotko (1987), found significant returns to occupational tenure, underlying the importance of occupation-specific human capital. The interested reader should also refer to the survey by Sanders and Taber (2012).

Following a more structural approach, Kahn and Lange (2014) estimate a model of learning with time-varying productivity. Pavan (2011) considers a model with general, career-specific and firm-specific human capital and sorting. Pastorino (2024) estimates a model of learning, job assignment and human capital accumulation within a firm and finds learning to account for about a quarter of wage growth, mostly accounted through improved job assignments. The present paper extends the above literature by considering several choices simultaneously, including non-employment, by allowing for heterogeneity across both occupations and individuals, partial transferability of human capital and information spillovers.

Gittins indexes were first used in a labor market setting in Miller (1984). In his setup, two jobs belong in the same occupation if their prior belief about unobserved match quality and the speed of learning is the same. A worker switches occupations once they have exhausted all jobs in an occupation. More recently Gittins indexes have been used by Papageorgiou (2022) in a model of local labor markets used to explain

[4] Altonji and Williams (2005) revisited and updated that literature.

agglomeration economies and by Silos and Smith (2015) who use it to examine the acquisition of pre-labor market skills. In contrast to these papers, we allow for informational spillovers across occupations.

Closer to our setup, Dickstein (2021) considers the choice of pharmaceutical care where physicians are uncertain about how a patient will respond to available treatments. As in our model, available choices are organized in clusters and within clusters choices are allowed to be correlated (see also Pandey et al., 2007). Pastorino (2015) introduces a setup where a worker's unobserved ability is firm specific, but each firm has a hierarchy of jobs and workers move up and down the firm hierarchy. Papageorgiou (2014) considers an economy with three occupations where workers learn about their comparative advantage.[5]

The present paper extends the above literature, by allowing not just for learning about match quality as in the typical multi-armed bandit problem, but also human capital accumulation which is partly transferable across occupations. The introduction of occupation-specific human capital is important both because of its direct relevance, and because it acts as a source of switching costs, which are typically not included in work that uses the bandit problem approach.[6] Traiberman (2019), in his study of the effects of import competition, shows that occupation specific human capital accumulation is as an important source of switching costs.

Our work complements the task-based approach to skills and occupations adopted in papers such as Poletaev and Robinson (2008), Yamaguchi (2012), Lise and Postel-Vinay (2020) and others. In fact, in our setup it would be possible to allow the initial distribution of worker productivities in every occupation to flexibly depend on multidimensional test scores, such as those available in the NLSY97.

Finally, the paper follows a recent growing literature that considers the importance of occupational choice in labor market outcomes. This includes work by Kambourov and Manovskii (2009a) and (2009b), Alvarez and Shimer (2009) and (2011), Antonovics and Golan (2012), Kramarz et al. (2014), Papageorgiou (2014) and (2018), Groes et al. (2015), Gervais et al. (2016), Gorry et al. (2019), Guvenen et al. (2020), Eeckhout and Weng (2022), Carrillo-Tudela et al. (2022), Carrillo-Tudela and Visschers (2023), Eckardt (2023), and others. We extend the literature by introducing a new model of occupational choice, consistent with the facts documented in the data.

[5] See also Felli and Harris (1996) and Klein and Rady (2011).

[6] The use of Gittins indices requires no fixed switching costs between options (Banks and Sundaram, 1994). We introduce specific human capital as a source of switching costs, which does not violate this requirement.

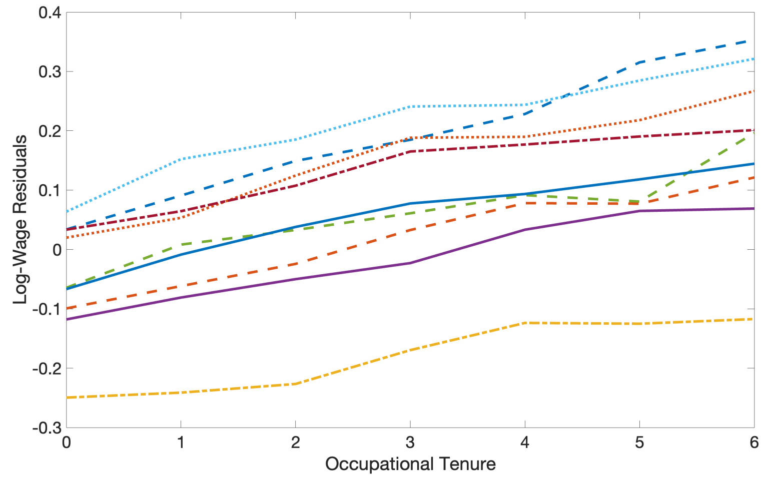

Figure 1: Log-Wage Occupational Tenure Profile by Occupation (Wage Residuals). Controls are a cubic polynomial in age, race, gender, education and year dummies. Each line corresponds to a different occupation. NLSY97.

3 Data/Facts

We use data from the National Longitudinal Survey of Youth 1997 (NLSY97). This survey is designed to be representative of the US population born between 1980 and 1984. Interviews are conducted every year starting in 1997. Our analysis uses data from the first 15 waves of the survey through 2011, so that its respondents are relatively young, which works well for our analysis, since most occupational switching takes place early in workers' careers.[7] The full sample includes 8,984 individuals and contains information on wages, occupations and the usual demographics, such as gender, age, race and education level. The NLSY97 serves our purposes well: unlike other commonly used datasets, such as the CPS or the SIPP, it is a long panel that allows us to track workers throughout their career and construct accurate measures of occupational tenure.

We only include observations on workers who have completed their education. We also do not include individuals from the oversample. Workers who report working fewer than 30 hours per week are considered non-employed. Wages are deflated to the real 2000 wage level using the Consumer Price index. We trim the top and bottom 0.5% of wage observations. The final sample contains 23,348 wage observations.

The NLSY97 uses the 2002 Census Occupational Classification. Our analysis focuses on the nine major occupational groups of the 2002 classification.[8] Since we are interested in examining different patterns for

[7] Kambourov and Manovskii (2008)

[8] The major occupational groups are 1) Management, business and financial operation occupations (10-950), 2) Professional and related occupations (1000-3540), 3) Service occupations (3600-4650), 4) Sales and related occupations (4700-4960), 5) Office and administrative support occupations (5000-5930), 6) Farming, fishing and forestry occupations (6000-6130), 7) Construction and extraction occupations (6200-6940), 8) Installation, maintenance and repair occupations (7000-7620), 9) Production occupations (7700-8960), 10) Transportation and material moving occupations (9000-9750).

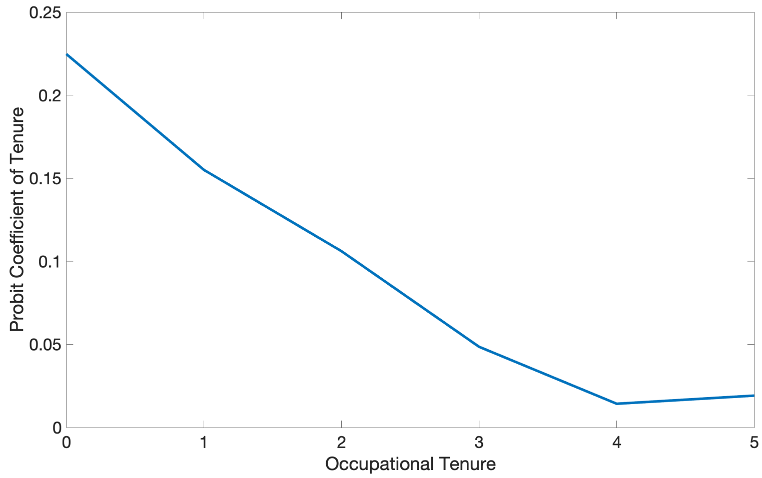

Figure 2: Impact of Tenure on Annual Probability of a 1-Digit Occupational Switch by Occupational Tenure. Probit regression of the probability of switching occupations on tenure levels zero through five, cubic polynomial in age, race, gender, education and year dummies. NLSY97.

each occupation, focusing on a finer partition is not possible due to sample size restrictions. It is worth noting however that 73% of all switches at the 3-digit level also involve a switch of a major occupational group as well. Table 5 in the Appendix contains descriptive statistics.

3.1 Wage Tenure Profiles

We begin by examining the wage profile as a function of occupational tenure.[9] Figure 1 presents the log-wage tenure profile broken down by major occupational groups. Wages here are residuals, after controlling for a cubic polynomial in age, race, education and year dummies. All occupations show an increase in residual log wages with tenure, but there is substantial heterogeneity, both in the wage level, but also in rate of wage increase: workers in some occupations enjoy significant wage gains, whereas others experience more moderate increases. Indeed, as shown in Table 6 of the Appendix which shows the slope coefficients for the different occupations, the rate of increase can vary significantly across occupations.

3.2 Probability of Occupational Switch with Tenure

We next investigate the relationship between occupational tenure and the probability of an occupational switch. In Figure 2 we present the results from a probit regression of the probability of switching occu-

We do not include observations for the sixth group, Farming, fishing and forestry occupations (6000-6130), because there are too few observations for a detailed analysis of that occupation (0.63% of all observations are in that group). The patterns we present below appear to hold in this occupation as well, but they are much noisier due to the small sample size.

[9]We compute occupational tenure, allowing for interruptions, so that if a worker leaves an occupation and returns to it at a later date, their tenure clock is not reset, but continues where it left off.

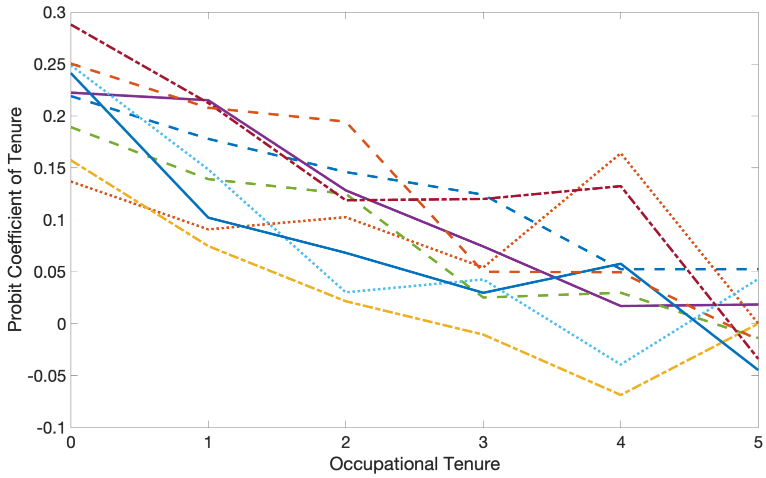

Figure 3: Impact of Tenure on Annual Probability of 1-Digit Occupational Switch by Occupational Tenure. Probit regressions of the probability of switching occupations on tenure levels zero through five, cubic polynomial in age, race, gender, education and year dummies. Each line corresponds to a different occupation. NLSY97.

pations on several tenure levels, demographic controls and year dummies. There is a very steep decline of the impact of tenure on the switching probability. In particular, workers with tenure less than a year have a 23% higher probability of switching occupations compared to workers with tenure six years or greater. The effect of tenure on switches seems to peak at year 4. Figure 3 plots the same graph for each of the major occupational groups separately. Again, the pattern holds qualitatively for all groups, with occupational switching falling rapidly with tenure, though there is substantial heterogeneity among occupations. This also evident in Table 7 of the Appendix, which regresses these coefficients on occupational tenure to measure the rate of decline.

These results are consistent with those in McCall (1990) who found that for workers who switched employers but remained in the same occupation, increased tenure in the previous employer lowers the probability of separation from current employer.

3.3 Wage Dispersion with Tenure

We now consider the within-occupation dispersion of wage residuals as a function of occupational tenure. As shown in Figure 4, the cross-sectional variance of log wage residuals increases with occupational tenure. In other words, within an occupation, inequality across workers with similar tenure levels, increases as tenure goes up. Inequality somewhat subsides for workers with 6 or more years of tenure. In Figure 5, we plot the same relationship within each occupational group separately. It is worth noting that there is substantial heterogeneity across occupations both in the level of inequality and also the rate of increase

Figure 4: Cross-Sectional Variance of Log-Wage Residuals as a function of Occupational Tenure. Controls are a cubic polynomial in age, race, gender, education and year dummies. Tenure indicator 6 includes workers with 6 or more years of tenure. NLSY97.

Switch Occupations $t+1$:	$\Delta \ln wage_t$	Residual $\Delta \ln wage_t$
Yes	4.29%	-0.82%
No	6.49%	1.95 %

Table 1: Annual Wage Change Before Occupational Switching. Wage residual controls are a cubic polynomial in age, race, gender, education and year dummies. NLSY97.

with tenure. Indeed, as shown in Table 8 of the Appendix which shows the coefficients from a regression of cross-sectional dispersion on tenure for the different occupations, while, with only one exception, the coefficients are positive, their magnitude varies significantly across occupations and in several cases it is not statistically significant.

3.4 Wage Change Before Switching Occupations

We also investigate the path of wages prior to an occupational switch. As shown in the first column of the top panel of Table 1, workers who switch occupations in between periods t and $t+1$, experience a wage increase of 4.3% between periods $t-1$ and t. On the contrary, workers who remain the same occupation, enjoy a 6.5% increase in their wages. The same pattern holds when considering residual wages: workers who switch occupations in the following year, experience declining wages before the switch; workers who remain in the same occupation have a wage increase of almost 2%. This is consistent with the notion that workers switch out of an occupation following a series of relatively bad realizations.

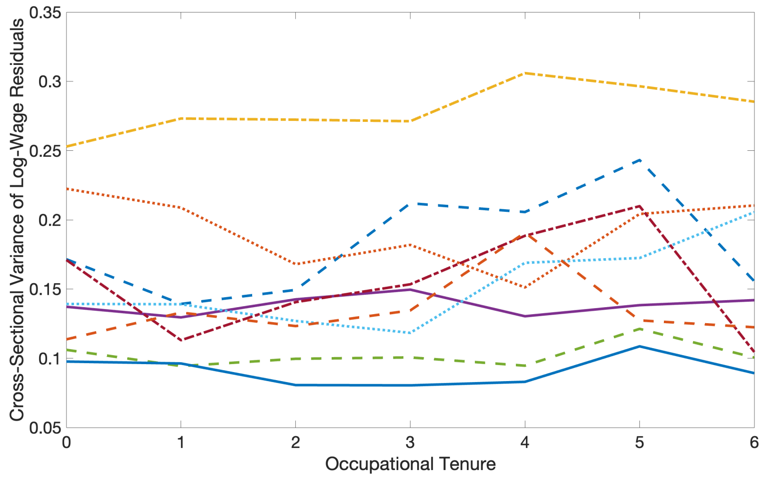

Figure 5: Cross-Sectional Variance of Within-Occupation Log-Wage Residuals as a function of Occupational Tenure. Controls are a cubic polynomial in age, race, gender, education and year dummies. Tenure indicator 6 includes workers with 6 or more years of tenure. Each line corresponds to a different occupation. NLSY97.

3.5 Probability of Occupational Switch by Wage Percentile

We next turn to investigate how a worker's relative position within an occupation affects the probability of an occupational switch. In Figure 6 we plot the probability of an occupational switch as a function of the worker's within occupation wage percentile, using wage residuals as the base of the rank.[10] The results suggest that the lowest paying workers have the highest probability of an occupational switch. Workers further up the occupation-specific wage distribution are less likely to switch, and the highest paid workers are the ones least likely to switch occupations. In Figure 7 we present the same relationship for each major occupational group separately. It is clear that the result is not driven by a particular group; instead, in every major group, the probability of an occupational switch is declining in the worker's within-occupation wage percentile. In Table 9 of the Appendix we present the coefficients from the regression of the annual occupational switching probability on the within-occupation residual wage percentile. For all occupations the slope is negative and statistically significant. In order to make sure that our result is not driven by some particular feature of the NLSY97 dataset, we compute the same set of moments using the 1996 panel of the Survey of Income and Program Participation (SIPP). Unlike the NLSY97 that follows a single cohort of workers, who are still relatively young, the SIPP is a nationally representative sample of households in the civilian non-institutionalized U.S. population. In addition, the 1996 SIPP's occupational classification system is different from the NLSY97's. Interviews were conducted every four

[10]The results remain qualitatively unchanged if we consider raw wages, use residual wages using the worker's position within an occupation in a particular year, include only occupational switches within the same employer or consider only switches involving an employer switch.

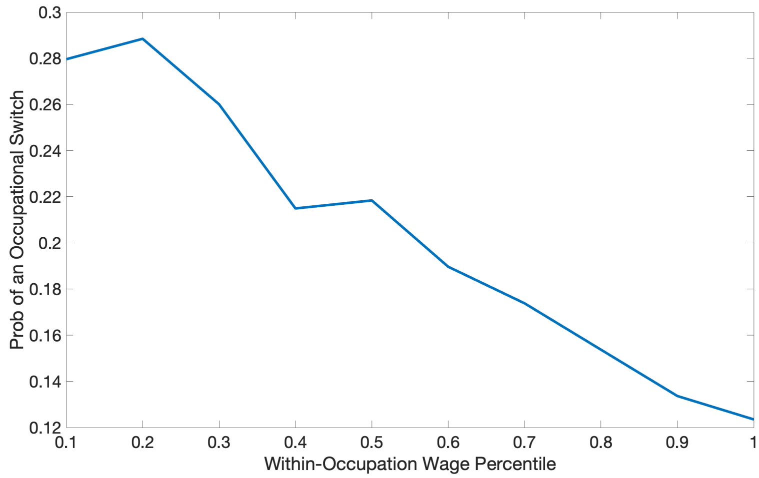

Figure 6: Annual Probability of a 1-Digit Occupational Switch by Within-Occupation Wage Percentile (residuals). Controls are a cubic polynomial in age, race, gender, education and year dummies. NLSY97.

months instead of annually, for four years and included approximately 36,000 households.[11]

As shown in Figure 8, the same relationship holds in the 1996 SIPP as well. Again the lowest-paid workers in a occupational group are the ones most likely to switch, whereas the highest-paid ones are the ones most likely to stay. When examining each major occupation separately (Figure 9), again the same pattern emerges for each one of them. Thus, we conclude that the pattern documented here is not due to some particular feature of the dataset used.

This particular moment matters, since, as Groes, Kircher, and Manovskii (2015) argue, it informs on the modeling of occupational choice. Groes et al. (2015) use Danish data and document that occupational mobility is U-shaped: The lowest and the highest paid workers in an occupation are most likely to switch. Workers paid near the occupation mean wage have the lowest probability of switching. Groes et al. (2015) interpret their finding as evidence for vertical occupational mobility: There are better and worse occupations; as workers learn about their ability, they climb up or down the occupational hierarchy. In contrast, our results that focus on major occupational groups in the U.S., suggest that occupation-specific match quality is important: workers who are mismatched, i.e. those with low wages and low occupational tenure, switch occupations, whereas well-matched workers, i.e. those with high wage and tenure as a result, are the least likely to switch. This conclusion has been subsequently confirmed by Carrillo-Tudela et al. (2022), who analyze patterns of occupation mobility using SIPP data from 1990 until 2008.

[11]Hourly wages are deflated to real 1996 dollars using the Consumer Price Index.

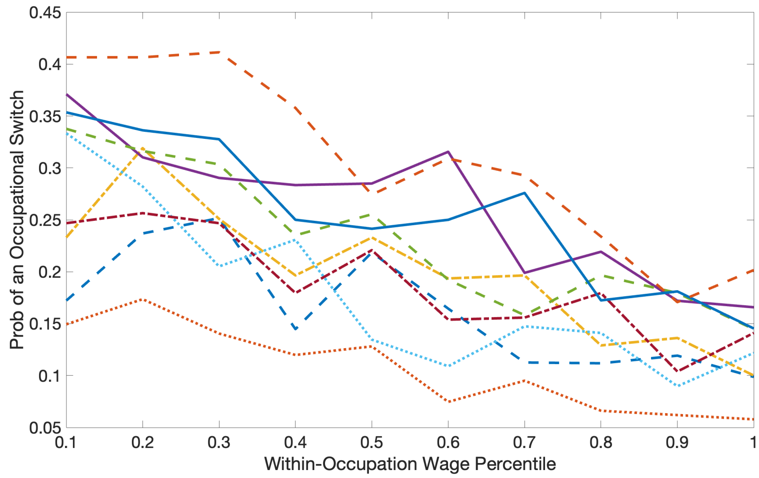

Figure 7: Annual Probability of a 1-Digit Occupational Switch by Within-Occupation Wage Percentile by Major Occupation (residuals). Controls are a cubic polynomial in age, race, gender, education and year dummies. Each line corresponds to a different occupation. NLSY97.

3.6 Summary and Discussion

To summarize, some of the key facts characterizing workers' occupational switching probability, wages, and occupational tenure are: (i) wages are increasing in occupational tenure; (ii) occupational switching probability declines with tenure; (iii) wage dispersion increases with tenure; (iv) wages fall (or increase less) before switching occupations, and (v) within an occupation, the lowest paid workers are most likely to switch occupations.

Next, we argue that both human capital accumulation, as well as learning and selection, are potentially important for understanding these facts jointly. For instance, the increase in wages with occupational tenure could be explained by workers accumulating human capital in their given occupation, as well as by worker selection: workers who are well-suited in a particular occupation remain there, whereas the ones less suited are more likely to explore a different occupation.[12] Human capital, as well as learning, are also consistent with the decline in occupational switching as workers' tenure lengthens, the former because human capital acts as an endogenously accumulated switching cost, and the latter because of the interaction between learning and selection.

However, the high rates of occupational switching cannot be accounted for by human capital accumulation alone: if workers simply became better over time in their chosen profession they would have no

[12] Job-ladder models, where workers move up the ladder through job-to-job transitions also predict wage increases over time (e.g. Burdett and Mortensen, 1998). However one might expect wage inequality to become more compressed with tenure rather than increase as in Figure 4: wages in those models can be quite dispersed early on, as many workers are at the bottom of the wage distribution, but some, "lucky" ones, might be closer to the top. Over time, if the separation rate is sufficiently low, continuing workers concentrate towards the top of the ladder, thus reducing wage dispersion. We are grateful to an anonymous referee for making this point.

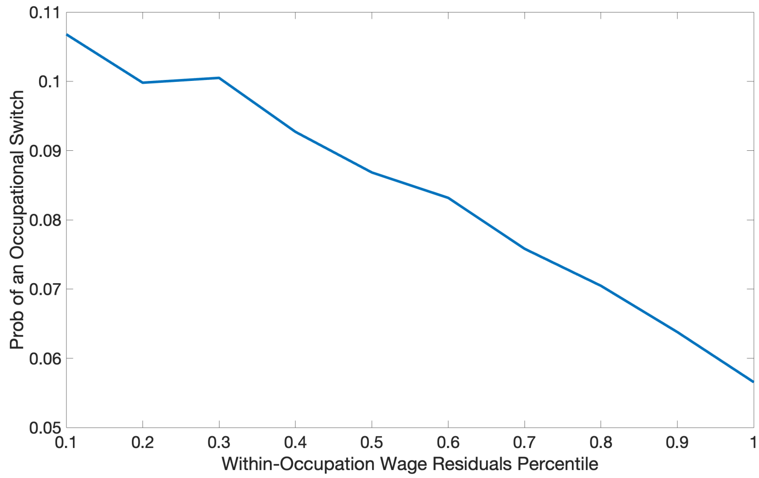

Figure 8: Four-Month Probability of Switching Major Occupations by Within-Occupation Wage Percentile - SIPP 1996 data (residuals). Controls are a cubic polynomial in age, race, gender, education and year dummies.

incentives to switch. Therefore a "shock" is necessary to induce workers to switch.[13] The most natural explanation is perhaps that workers are learning about their fit in different occupations, as originally proposed by Jovanovic (1979) in the context of matching with firms. Learning is also consistent with workers' wages falling before they switch occupations since their expected productivity falls as they learn that they are not as well suited in that occupation. It would also explain why most of the occupational switching is driven by workers at the bottom of the within-occupation wage distribution since those are the ones with low (expected) productivity, where a negative signal could push them to try out other occupations.

Thus human capital accumulation, as well as learning appear to be crucial ingredients of any framework built to account for these facts and explore shocks or policies affecting worker occupational choices, such as automation or trade shocks. While human capital or learning could be specific to an occupation, it is also possible that there are "spillovers" of both to other occupations. For instance, the human capital that a worker accumulates in an occupation might be (partially) transferable to other occupations. Similarly, what a worker learns about their match in an occupation might be informative about their fit in other occupations as well.[14]

[13]Stochastic human capital accumulation could also be an alternative explanation for the flows between occupations. This might be the case for instance, when e.g. automation makes some workers less productive in some occupations, but more productive in others, therefore changing the value of their human capital. However given that the vast majority of occupational flows are offsetting, rather than net flows (Kambourov and Manovskii, 2008), these types of "aggregate" shocks can only potentially account for a small fraction of the high rates of occupational switching that we observe in the data. In the conclusion of the paper we revisit these shocks and how our proposed framework might be able capture them.

[14]These observations have motivated the task-based approach to skills and occupations adopted in papers such as Poletaev and Robinson (2008), Yamaguchi (2012), Lise and Postel-Vinay (2020) and others, where different occupations make use of similar skills, albeit at different intensities.

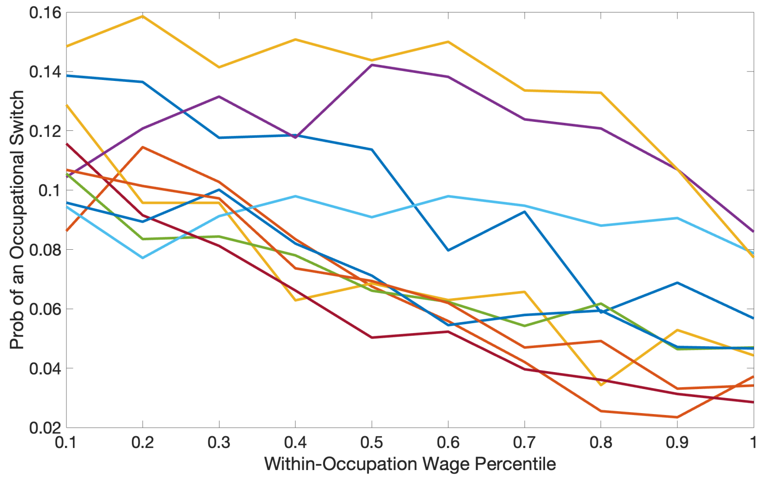

Figure 9: Four-Month Probability of Switching Major Occupations by Within-Occupation Wage Percentile by Major Occupation - SIPP 1996 data (residuals). Controls are a cubic polynomial in age, race, gender, education and year dummies. Each line corresponds to a different occupation.

Motivated by these observations, we present a tractable framework that includes both learning, as well as human capital accumulation, both of which can be partially specific and partially transferable to other or even all occupations. Adding all these ingredients, while conceptually straightforward, can also increase the state space quite a bit rendering the model computationally intractable, especially if we consider more than a few occupations. However as we show below, our setup can maintain tractability, despite allowing for human capital and beliefs in all occupations, as well as spillovers across them.

4 Model

4.1 Environment

Time is discrete and individuals have log utility with discount rate β. At each point in time individuals can be employed or non-employed.[15] Each individual dies with probability ξ every period and a new individual takes their place ensuring that the total population remains constant. The model includes both ex ante and ex post individual heterogeneity. There are two dimensions of ex ante heterogeneity: first individuals are heterogeneous in their fixed observable characteristics, such as education, gender, etc., denoted by the vector X. Second, there is also fixed heterogeneity captured by the vector θ. This vector is perfectly observable by all market participants.[16] In what follows, individual subscripts are suppressed

[15] In what follows we do not distinguish between unemployment and out of the labor force. In Section 6.2 we discuss an extension where that distinction would be meaningful.

[16] This is meant capture worker characteristics that are not usually collected in the data, so they are unobserved to the econometrician, but are observable by market participants, such as coding skills, knowledge of foreign languages, social skills

throughout to reduce notation congestion. The ex post heterogeneity comes from the accumulation of human capital, and the beliefs regarding individual productivity in each occupation which we discuss in detail below.

Individuals work in one out of O available occupations or they can be non-employed. Occupations are organized into one of C clusters and each occupation belongs to exactly one cluster. The purpose of this classification is to allow for a pattern of dependence for occupations within a cluster, so that a worker's productivities across occupations within the same cluster are not independent, while keeping the assumption of independence between clusters, which is required for our use of Gittins indices as explained in detail below.

Log output of a worker with characteristics X, in occupation o of cluster c at time t is given by

$$q_{oct}(\cdot) = hc_o(\tau_{ot}, X) + hc_c(\tau_{ct}, X) + y_{ot} + y_{ct}$$

where τ_{ot} is tenure in occupation o and $\tau_{ct} = \sum_{o \in c} \tau_{ot}$ is tenure in cluster c. $hc_o(\tau_{ot}, X)$ captures occupation-specific human capital which depends on occupation-specific tenure. This function can differ across occupations. Similarly, $hc_c(\tau_{ct}, X)$ captures cluster-specific human capital which evolves as a function of cluster-specific tenure and again can differ across clusters.[17] The last two terms of the output equation capture the worker's current period productivity, net of human capital. In particular,

$$y_{ot} = \eta_o + \varepsilon_{ot}$$

and

$$y_{ct} = \eta_c + \varepsilon_{ct}$$

where $\varepsilon_{ot} \sim N\left(0, \sigma_{yo}^2\right)$ and $\varepsilon_{ct} \sim N\left(0, \sigma_{yc}^2\right)$ are i.i.d. Both y_{ot} and y_{ct} are separately observed, but η_o and η_c are unobserved and are drawn independently before a worker enters the labor force from normal distributions,

$$\eta_o \sim N\left(\mu_o(X, \theta), (\sigma_o^\eta)^2\right), \tag{1}$$

and

$$\eta_c \sim N\left(\mu_c(X, \theta), (\sigma_c^\eta)^2\right), \tag{2}$$

etc.

[17] It would be straightforward to include stochastic human capital accumulation, but we prefer to have the uncertainty in individual's histories to be captured by the learning processes.

where $\mu_o(X,\theta)$ and $\mu_c(X,\theta)$ are observed by the agents.[18] Information is symmetric, so workers and firms share the same beliefs about one's abilities.[19] Their dependence on both X and θ, implies that two workers with the same observable characteristics, X, can have different priors, if they differ in θ. To ease notation congestion, in what follows we drop the X argument where obvious.

Every period, agents observe y_{ot} and y_{ct} and update their beliefs regarding η_o and η_c according to Bayes' rule, using the usual updating formulas. A worker's posterior beliefs regarding η_o and η_c are $N\left(\mu_{ot}^p,(\sigma_{ot}^p)^2\right)$ and $N\left(\mu_{ct}^p,(\sigma_{ct}^p)^2\right)$ respectively. μ_{ot}^p (μ_{ct}^p) is a function of all past output realizations in that occupation (cluster) and σ_{ot}^p (σ_{ct}^p) is updated deterministically based on tenure in the occupation (cluster). It is worth noting that a single state variable (tenure) is the sole determinant of both the variance of the posterior distribution of beliefs and the amount of specific human capital accumulation, which helps with the tractability of the problem.

There are competitive markets within occupations and following Kahn and Lange (2014) log wages are given by:

$$w_{oct} = E_t(q_{oct}) = hc_o(\tau_{ot}) + hc_c(\tau_{ct}) + \mu_{ot}^p + \mu_{ct}^p.$$

Finally, log home production of a non-employed worker with characteristics X is $b(z_t,X)$ where z_t is a latent variable that is observed by the worker and follows a Markov process that evolves only when the worker is non-employed. New workers draw their initial value of z_0 from $b_0(\cdot)$. There is no cost to switching occupations/clusters or from movement between employment and non-employment.[20]

The sequence of actions is the following: every period a workers picks an occupation or non-employment and receives their wage if employed and their value of home production otherwise. If employed, they then produce output, observe y_{ot} and y_{ct} and update their beliefs regarding η_o and η_c. At the same time their cluster and occupation specific human capital increase. Finally, with probability ξ they die and in the next period a new worker takes their place, with the same characteristics X and θ who draws their z_0 and has zero tenure in all occupations.

It is also straightforward to incorporate general human capital into the model by adding an additive

[18]Note that the relative magnitude of the cluster variables, $hc_c(\tau_{ct},X)$ and y_{ct}, essentially captures how correlated a worker's output is across the occupations of the cluster: if these variables are large relative to the occupation-specific variables ($hc_o(\tau_{ot},X)$ and y_{ot}), then the output correlation will be high; conversely, if the cluster variables are relatively small, then a worker's output across the cluster's occupations will be close to orthogonal.

[19]We implicitly assume that firms also have symmetric information about the worker's productivity. Scheonberg (2007) finds that learning is mostly symmetric across employers, whereas Kahn (2013) does find evidence for asymmetric information. See also Eeckhout (2006) for a model with asymmetric information across employers.

[20]The specification of non-employment is similar in spirit to the rest unemployment state in Alvarez and Shimer (2011) and also Alvarez et al. (2023). Workers costlessly move out of non-employment when that state is no longer attractive relatively to employment.

term to log output that increases both when the worker is employed as well as when non-employed. As long as general human capital depends on age and the rate of general human capital accumulation is everywhere the same, this does not affect worker behavior and the solution we present below holds. This extension requires including tenure in non-employment as a state variable as well.

4.2 Behavior

Every period a worker needs to decide between O occupations and non-employment. The worker's state space is

$$s_t = \left\{ [\mu^p_{ot}, \tau_{ot}]_{o\in\{1,\dots,O\}}, [\mu^p_{ct}]_{c\in\{1,\dots,C\}}, z_t \right\} \tag{3}$$

in other words, their belief, μ^p_{ot}, and tenure level, τ_{ot}, for every occupation o, their belief, μ^p_{ct}, for every cluster c and the latent non-employment variable z_t. Note that by keeping track of τ_{ot} and using $\tau_{ct} = \sum_{o\in c} \tau_{ot}$, the worker effectively also keeps track of their tenure in each cluster. Tenure affects both the level of specific human capital at the occupational or cluster level, as well as beliefs regarding η_o or η_c.

If there are more than a handful of occupations available to the worker, then the worker's state space is quite large and solving this problem numerically becomes extremely difficult, as one quickly runs into the "curse of dimensionality". For illustration, let $O = 10$ and $C = 5$ and assume we discretize the state space to allow for 10 tenure levels and 10 levels of beliefs for each occupation and each cluster, as well as the latent non-employment variable. Then state space of this problem contains $10^{10\times2+5+1} = 10^{26}$ elements. Below we discuss how we use Gittins indexes to solve the problem, without being constrained by these issues.

We first reformulate the worker's problem as a choice between one of C clusters and non-employment, and then picking an occupation within a cluster. In particular, the worker's problem is given by

$$W(s_t, X, \theta) = \max \{W_1(s_t), W_2(s_t), \dots, W_C(s_t), W_U(s_t)\} \tag{4}$$

where $W_c(s_t, X, \theta)$ is the value of a worker with state variables, s, defined in (3) and characteristics (X, θ) employed in cluster c at time t and optimally choosing an occupation within that cluster. Similarly $W_U(s_t, X, \theta)$ is the value of a non-employed worker.

In addition, the workers picks an occupation within each cluster c, so that

$$W_c(s_t, X, \theta) = \max_{o\in c} \{W_{co}(s_t, X, \theta)\}$$

where $W_{co}(s_t, X, \theta)$ is the value of a worker employed in occupation o of cluster c.

Note that when deciding between the clusters and also non-employment, a worker takes into account not only the cluster level variables, but also the state variables of the occupations within each cluster. For instance, some clusters may contain occupations where a worker has high levels of occupation-specific human capital, $hc_o(\tau_{ot}, X)$ and/or optimistic beliefs, μ^p_{ot}, regarding their ability in that occupation and of course this will influence the worker's cluster choice. In addition, it is important to point out that the worker's problem is dynamic and a worker does not necessarily choose to work in the choice with the highest contemporaneous wage. In particular, they also take into account the impact of their choice on a) the accumulation of different forms of human capital and b) the option value of learning about their occupational and cluster match.

This problem is a multi-armed bandit problem at the cluster level. We use Gittins indices to solve it, since it satisfies the necessary conditions for the Gittins theorem to hold.[21] In particular, a) there are no costs to switching clusters (arms), b) the state variables of each cluster evolve only when the worker is employed in that cluster and are otherwise unchanged, c) there is constant exponential discounting, d) the clusters are independent, so that the payoff from each cluster depends only the cluster-specific variables and e) the process governing the evolution of the state variables is Markovian.[22,23]

Lemma 1. *The solution to the worker's problem (4) takes the form of an index policy, whereby every period the worker chooses the cluster with the highest index.*

Proof. Gittins (1979), DAI Theorem. □

In order to compute the Gittins index for each cluster and for non-employment, we follow the approach employed by Whittle (1980) and (1982), Karatzas (1984) and Pandey et al. (2007). We first consider the clusters and then discuss non-employment.

In particular, for each of the C clusters, consider a "transformed" problem whereby the worker from now on, has only two choices: either work in one of the occupations of that cluster (and continue accumulating human capital and learning) or retire and receive value M_c. Note that this problem ignores the other clusters and non-employment. The Gittins index of that cluster is the retirement value at which the

[21]The formulation of the bandit problem dates back to the 1940s. Gittins and Jones (1974) showed that the problem can be solved through the use an index that is now known as the Gittins index (see also (Bergemann and Valimaki, 2008)).

[22]See Mahajan and Teneketzis (2007), Bergemann and Valimaki (2008) and Gittins et al. (2011) (Chapters 2 and 3).

[23]For the general human capital extension mentioned above, the value of an individual would be the sum of two components: the first, coming from general human capital that depends only on age; the second would be just as described above. The two components have no interdependence so the value net of general human capital can be computed with Gittins indices as described above.

worker is exactly indifferent between retiring and continuing to work in that cluster. More specifically, the value function of the worker in that case is given by

$$V_c\left(M_c, s_t^c, X, \theta\right) = \max\left\{M_c, \max_{o\in c} w_{oct}\left(s_t^c, X\right) + \beta\left(1-\xi\right) E_t V_c\left(M_c, s_{t+1}^c, X, \theta\right)\right\}, \tag{5}$$

where

$$s_t^c = \left\{\left[\mu_{ot}^p, \tau_{ot}\right]_{o\in c}, \mu_{ct}^p\right\}$$

is the worker's state in the above problem which includes their tenure level for each occupation, which determines their level of human capital and precision of their beliefs, both at every occupation, but also at the cluster level, since $\tau_{ct} = \sum_{o\in c} \tau_{ot}$. In addition, it includes the worker's mean beliefs regarding η_c and η_o for every o within the cluster.

As shown in Whittle (1982), the Gittins index of cluster c is given by

$$\gamma_c\left(s_t^c, X, \theta\right) = \inf\left\{M_c | V_c\left(M_c, s_t^c, X, \theta\right) = M_c\right\}, \tag{6}$$

i.e. the retirement value for which the worker, in the transformed problem, is indifferent between working in one of the cluster's occupation or retiring and receiving value $\gamma_c\left(s_t^c, X, \theta\right)$.

For the case of non-employment the "transformed" problem assumes that the worker can either remain non-employed and receive the value of non-production while their latent variable, z_t evolves, or retire and receive value M_n. In this case their value function becomes

$$V_n\left(M_n, z_t, X\right) = \max\left\{M_n, b_t\left(z_t, X\right) + \beta\left(1-\xi\right) E_t V_n\left(M_n, z_{t+1}, X\right)\right\}. \tag{7}$$

The Gittins index of non-employment is similarly given by

$$\gamma_n\left(z_t, X\right) = \inf\left\{M_n | V_n\left(M_n, z_t, X\right) = M_n\right\}, \tag{8}$$

i.e. it is the retirement value such as that the worker is indifferent between remaining the in the non-employment state or retiring and receiving value $\gamma_n\left(z_t, X\right)$.

Given the above, we have

Proposition 1. *The optimal strategy of a worker is to choose the option (cluster c or non-employment) that has the highest index, where the indices are given by (6) and (8) respectively.*

Proof. Follows from Lemma 1, equations (5) and (7) and Whittle (1982). □

In other words, a worker solves the above problem for each cluster and for non-employment, computes the associated indices and picks the option with the highest index. Within a cluster, the choice of occupation is given by the occupational choice o in (5), evaluated at $M_c = \gamma_c\left(s_t^c, X, \theta\right)$ (Pandey et al., 2007).

4.3 Discussion

If we let O_c denote the number of occupations within cluster c, then in our setup, Gittins indices transforms the problem as follows

$$\begin{gathered}1 \text{ problem with } (2\times O)+C+1 \text{ state variables}\\ \Rightarrow C \text{ problems, each with } 2\times O_c+1 \text{ state variables and 1 problem with 1 state variable}\end{gathered}$$

In other words, using Gittins indices transforms the original problem into $C+1$ separate problems, one for each cluster and one for non-employment. The goal of each one of these $C+1$ "transformed" problems is to compute an index for that cluster/non-employment.

The advantage of Gittins indexes is that they substantially reduce the dimensionality of the original problem. The state variables of each "transformed" problem consist of the state variables of that cluster alone. For instance, following the example mentioned above, if $O = 10$, $C = 5$ and $O_c = 2$ with the same number of grid points as before, the Gittins formulation consists of 5 problems with 10^5 elements each and 1 problem with 10 elements. The reduction in the state space is by 21 orders of magnitude. The reduction becomes even larger, once one considers even more clusters.

It is worth emphasizing the above setup allows for flexible heterogeneity, both at the individual level, as well as the cluster/occupation level. In particular, clusters and occupations are allowed to differ in the importance of human capital, the rate of human capital accumulation and the learning processes. On the worker side, individual observable characteristics, captured by X, such as education or observed ability may similarly affect how quickly individuals accumulate human capital or discover their unobserved productivities. In addition, heterogeneity at the occupational level may flexibly interact with heterogeneity at the individual level, so that for instance certain occupations may allow rapid human capital accumulation for more educated workers, but not less educated ones.

In our model, we allow for correlation across occupations both in levels, as well as innovations. In terms of levels, workers' productivity may be arbitrarily correlated across occupations through the dependence of η_o and η_c on X and θ. For instance, workers with certain characteristics, such as particular parental background or education level may be more likely to be productive in certain occupations, but less likely to be productive in others. There is no restriction on the structure of the correlation in levels.

In terms of innovations, we allow workers' human capital accumulation and learning to be correlated across occupations in the same cluster, but not with occupations in other clusters. The common component of human capital, $hc_c(\cdot)$, implies that part of the human capital that a worker accumulates in one occupation may be used in other occupations in the same cluster. Similarly, the information a worker acquires about their unobserved cluster mean productivity, η_c, is informative regarding their productivity in other occupations in the same cluster. Finally, as discussed in Section 4.1 above, it is straightforward to allow for general human capital, which is equally valued across all occupations and also implies that a part of accumulated human capital can be equally used across all occupations, regardless of cluster.

In the baseline setup introduced above, we have not allowed for an explicit cost of switching occupations. These are easy to incorporate at the within-cluster level, but not when workers are switching across clusters, since then the Gittins indices solution will no longer be optimal (Banks and Sundaram, 1994). While this assumption appears restrictive, arguably the biggest cost when switching across clusters is the foregone specific human capital in the old cluster that can no longer be used in the new one, and also the need to accumulate specific human capital anew. This cost however is included in the model.[24]

We next present a simple illustrative calibration of the model where we shut down heterogeneity across occupations and clusters, as well as heterogeneity across individuals in ex ante characteristics, X and θ. We also shut down the non-employment margin. This parsimonious calibration is useful in a) providing a model-based approach to clustering occupations and b) illustrating how the key mechanisms can be separately identified in the data. This calibration is also useful to illustrate how our theory is suited to explain the empirical patterns highlighted in Section 3.

5 Calibration

We first discuss how we cluster occupations in the data using our theory's predictions; we then present our calibration procedure and discuss which features of the data identify the key model parameters; finally we present the results. In what follows we use the NLSY97 data that we used earlier in the empirical

[24]It is worth noting that we cannot allow occupation or cluster-specific human capital to depreciate once a worker has left the cluster, as we would no longer be able to use Gittins indices to solve the worker's problem.

Cluster 1	Cluster 2
Management, Business and Financial (10-950)	Professional and Related Occupations (1000-3540)
Sales and Related Occupations (4700-4960)	Service Occupations (3600-4650)
Cluster 3	**Cluster 4**
Office and Administrative Support (5000-5930)	Installation, Maintenance and Repair (7000-7620)
Construction and Extraction (6200-6940)	Production Occupations (7700-8960)
Cluster 5	
Transportation & Material Moving (9000-9750)	

Table 2: Clusters of Occupations. See text for details on the clustering.

analysis and focus on a sub-sample of white men with a high school, but no college education.

5.1 Clustering

We cluster occupations based on the following prediction of our setup: ignoring general human capital accumulation, a worker who switches to an occupation in a different cluster and then returns to their original occupation, should receive the same wage as before switching.[25] For example, if a worker is employed in occupation A that belongs to cluster 1, and switches to occupation B that belongs to cluster 2, but then returns to occupation A at a later date, they should receive the wage they received prior to switching. This also holds if the worker is employed in several occupations before returning to occupation A, as long as these occupations do not belong to cluster 1. This speaks to one of the key role of clusters: how past experience in certain occupations, but not others, matters for worker's wage in their current occupation.

Given the above prediction, we cluster occupations using the following algorithm. First assign occupations into clusters. We then find in the data instances where a worker switches out of an occupation to work in occupations in different clusters and then returns to the original occupation. Afterwards, using only the above switches, compute the variance of the wage differences between the wage prior to switching out and the wage upon returning.[26] Finally repeat this process for all possible cluster configurations and find the configuration that minimizes the variance of these log wage differences. Table 2 contains the resulting occupational clustering, where we restrict each cluster to contain at most two occupations.

[25]In discrete time, this does not hold exactly because after being employed in an occupation, immediately before switching, a worker increases their human capital and receives information that cause them to update his beliefs which may not show up in their wage. However if the time period is not too large, then the wage difference are small.

[26]We control for age, by taking out a quartic polynomial in age and using wage residuals.

5.2 Calibration Procedure

We take our model to the data by pre-setting the value of some of the parameters and then calibrating the remaining ones to match moments in the NLSY97 data that we used earlier in the empirical analysis. We use simulated method of moments, whereby we pick a vector of parameters, simulate a panel of data from the model and compute the predicted moments, compare them to the empirical ones and then update our parameters until the theoretical and data moments converge. To keep the calibration tractable we assume that occupations and clusters are symmetric, and focus on within-occupation/cluster patterns. We calibrate our model to four clusters with two occupations each, to mimic the configuration described on Table 2 (minus the cluster with a single occupation). It is worth noting that even in this symmetric occupations/cluster configuration, if we were to solve the value function of an individual worker, the problem would have twenty state variables, whereas using Gittins indices we only need to keep track of five state variables per cluster.

We set the discount rate, β to 0.96 and also set the death/retirement probability to 2.5% per year, which implies an expected working life of 40 years. We assume that occupation and cluster human capital both increase linearly at rate λ_o and λ_c annually up to 10 years of accumulated occupation/cluster tenure, and then remains constant afterward.

We then proceed to calibrate the following seven parameters: the human capital accumulation parameters, λ_o and λ_c, the standard deviation of the occupation and cluster signal parameters, σ_{yo} and σ_{yc}, the mean prior belief of the log of the worker's occupational ability, μ_o, and the dispersion of a worker's ability in an occupation and cluster, σ_o^η and σ_c^η respectively.[27] We use the following moments in order to identify them: the coefficients on occupational and cluster log tenure in a regression with log wages as the dependent variable (controlling for a cubic in age); the annual probability of occupational (cluster) switch for workers in their first year in an occupation (cluster), as well as those with four years of tenure; the average log wage of workers in their first year in the labor market.

As is typical of models with many moving parts, we cannot provide a formal identification argument, but instead we discuss informally how the different mechanisms are identified in the data. The standard deviation of a worker's ability in an occupation, σ_o^η, as well as the standard deviation of the signal, σ_{yo}, are identified by the occupational switching probability for workers in their first and fifth year in an occupation, respectively. The first parameter determines the precision of workers' beliefs before they enter

[27]Since we cannot separately identify the mean prior belief of the log of the worker's cluster ability, μ_c, from the log mean of the occupational ability, μ_o, we set the former to zero.

	Data	Model	No Clusters	No HC
Occupation Tenure Coefficient on Log Wage Regression	0.034	0.034	0.069	0.054
Cluster Tenure Coefficient on Log Wage Regression	0.018	0.018	-0.016	-0.014
Annual Prob of Occup Switch for Workers in First Year in Occup	0.333	0.32	0.705	0.321
Annual Prob of Occup Switch for Workers with Four Years in Occup	0.115	0.115	0.005	0.058
Annual Prob of Cluster Switch for Workers in First Year in Cluster	0.314	0.3	0.676	0.299
Annual Prob of Cluster Switch for Workers with Four Years in Cluster	0.128	0.115	0.041	0.058
Average Log Wage of Workers in First Year in Labor Market	2.229	2.229	2.229	2.229

Table 3: Empirical and Simulated Moments. Empirical moments computed using the NLSY97 for white, males with a high school degree, but no college education. Simulated moments computed using the procedure described in the main text. No Clusters sets $\lambda_c = \sigma_c^\eta = \sigma_{yc} = 0$, while No HC sets $\lambda_o = \lambda_c = 0$, while keeping the remaining parameters at their baseline calibration values.

λ_o	0.0001
λ_c	0.013
σ_{yo}	0.325
σ_{yc}	0.973
σ_o^η	3.4947
σ_c^η	0.464
μ_o	2.229

Table 4: Estimated Parameters

the market and therefore the level of occupational switching early on: if workers are confident in their beliefs, they are more likely to start off in the "correct" occupation. Conversely, if they are fairly unsure, they are quite likely to switch occupations, so we should observe high levels of occupational switching for workers with low tenure. The speed at which they learn about their productivities-and, consequently, the reduction in occupational switching as well as their switching rate after five years in a given occupation-is determined by σ_{yo}, the standard deviation of the noise variable. A similar argument applies regarding the identification of σ_c^η and σ_{yc} using the respective cluster moments.

Moreover, the human capital accumulation parameters, λ_o and λ_c, are pinned down using the coefficients of occupational and cluster tenure in the log wage regression: conditional on the parameters determining the speed of learning and therefore selection, any remaining increase in wages is driven by human capital accumulation. Finally, the mean prior belief regarding the log of the worker's occupational ability, μ_o is pinned down by the average log wage of workers in their first year in the labor market.

5.3 Results

Table 3 presents the data and model predicted moments, while Table 4 shows the resulting parameter estimates. Overall, the model matches the targeted moments well, though the model tends to underpredict occupational switches at higher levels of tenure (see Figure 10 discussed below). It is worth noting that measurement error can affect occupational switching measures (Kambourov and Manovskii, 2008), so the relatively high number of switches observed at higher levels of tenure may partly reflect this.

The parameters suggest that there is significant dispersion in the occupation-specific component of productivity, η_o, as captured by the high value of σ_o^η. In addition, there is sizable human capital accumulation of cluster-specific human capital (λ_c), but virtually no human capital accumulation for the occupation-specific component (λ_o). Given that the model predicts significant returns to both occupational and cluster tenure, as shown in the first two rows of Table 3, this suggests that a) the returns to occupational tenure are entirely driven by learning and selection of workers in occupations where their match quality is the highest and b) more than half of the returns to cluster tenure are driven by the accumulation of cluster-specific human capital, whereas the remainder are driven by learning and selection to clusters. This is consistent with the high estimated dispersion of the occupation-specific component of the match, σ_o^η, which implies that there are significant returns to finding a good occupational match.

If we shut down clusters (i.e. set $\lambda_c = \sigma_c^\eta = \sigma_{yc} = 0$) and simulate the model, then occupational mobility in the first year approximately doubles (0.705) and then goes to zero (0.005) five years later. This seems to happen for two reasons: first, since there is no longer sorting for clusters, workers can sort into occupations without worrying about losing a good cluster match. Second, since there is no cluster human capital, this allows workers to select into the occupations they believe they are best matched with, without worrying about foregone human capital. Both of these forces, can lead to more occupational sorting early on, but by year five, almost all workers have learned about their occupational matches and there is no more occupational switching. This suggests that allowing for a cluster component is key in understanding mobility patterns.[28]

Having said that, if we shut down human capital accumulation, rather than clusters (i.e. set $\lambda_o = \lambda_c = 0$), then the total returns to occupational and cluster tenure fall by approximately $\lambda_o + \lambda_c$, suggesting that the interaction between human capital and selection discussed in the second reason is not as strong.[29]

Finally, in Figures 10 through 11, we illustrate how the calibrated model can match some of the key facts we documented in Section 3, namely the relationship between log wages and occupational

[28]If instead we shut down clusters and also recalibrate all model parameters, then the recalibrated model again cannot account for occupational switching at four years of occupational tenure, predicting that it goes to zero. This happens, despite this being one of the targeted moments.

[29]Somewhat surprisingly, we obtain positive returns to occupational tenure, but negative returns to cluster tenure. This is driven by the high returns to finding a good occupational match. Conditional on occupational tenure, if a worker has high cluster tenure, suggesting that they're likely to have selected based on their cluster match rather than their occupational match, they will have lower wages.

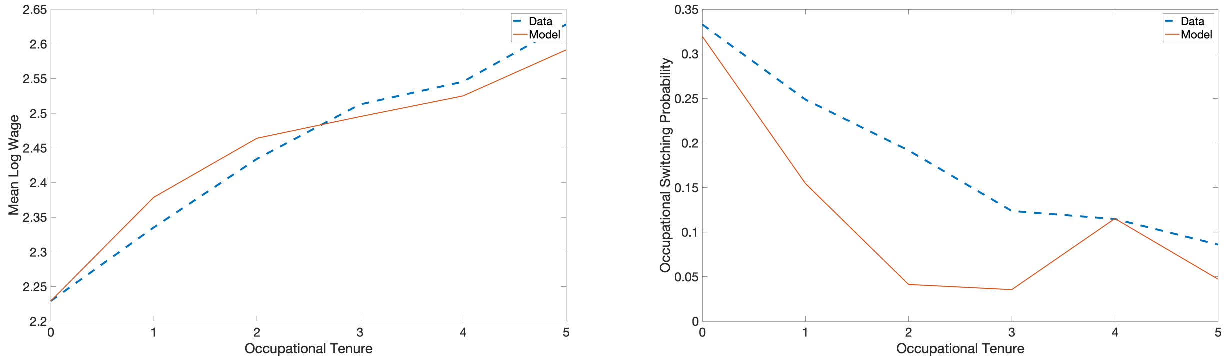

Figure 10: Other Moments

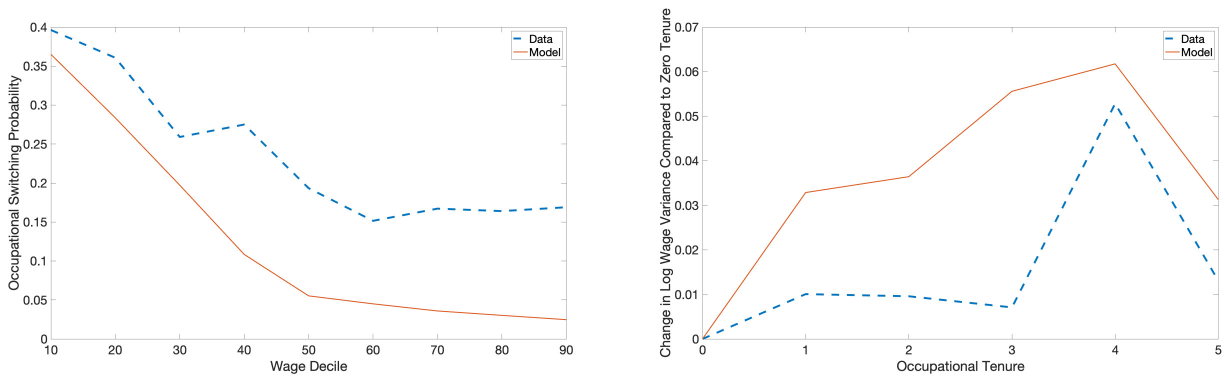

Figure 11: Other Moments 2

tenure, occupational switching and occupational tenure, occupational switching and within-occupation wage percentile, and finally the increase in log wage dispersion and occupational tenure.[30] The fit is generally good, given that many of these moments are not targeted in the calibration. In particular, it is worth noting that our model is able to produce a declining probability of occupational switching in ranked wage deciles, consistent with our earlier empirical results (Figure 6), even though this was not targeted in the calibration.

6 Extensions

Before concluding, we discuss some potential extensions of the model introduced here.

[30]We look at changes in the variance of log wages compared to the variance of log wages at zero tenure, because in our calibration, since we do not allow for heterogeneity in initial beliefs, there is no wage dispersion at zero tenure. When computing the variance of log wages in the data, we first compute the within-occupation variance for each occupation and then compute a weighted average, since our calibration abstracts from across-occupation differences.

6.1 Clusters Available over Time (Arm-Acquiring Bandit)

In the baseline model, all clusters are always available to the worker. However we can extend it to allow for some clusters to become available only later in a worker's career. New workers have access to only a subset of the available clusters. Over time, perhaps through the use of social networks, a worker stochastically learns the location of additional clusters. As shown in Whittle (1981), the use of Gittins indices is optimal in this modified setup as well, so long as the arrival of the new options does not depend on the past choices.[31]

6.2 Search Frictions

Our setup can be extended to allow for search frictions. In this formulation, individuals are either employed, unemployed (searching) in a particular occupation or out of the labor force. The first time a worker visits an occupation they are unemployed. An unemployed worker produces $b_0(\cdot)$ which can depend on various cluster and occupation specific variables, as well as individual characteristics X. They meet firms with some probability λ_o, while the employed lose their jobs exogenously with some probability δ_o. When employed a worker receives a wage $w_0(.)$.

There are a couple of challenges in the setup. First, in terms of the wage determination, wages cannot depend on state variables of other clusters or occupations belonging to other clusters, as then Gittins indices would be inapplicable. For instance, Nash bargaining cannot be used, since the worker's outside option would include the option of working in other clusters. An exception here would be the case where the worker has all the bargaining power. One possible solution is to model the wage determination as in Postel-Vinay and Robin (2002), by allowing for within-occupation on-the-job search. Similarly it is possible to allow for a distribution of posted wages and allow workers to obtain wage increases through on-the-job search as in Burdett and Mortensen (1998).

The second challenge is that when a worker leaves an occupation we need to allow the possibility of recall, so that if and when a worker returns to that occupation, they are immediately employed in their old firm. Otherwise, if workers started unemployed again, then this would effectively introduce a cost to switching clusters, and therefore precluding the use of Gittins indices.

[31] This modified problem is known in the literature as an "arm-acquiring bandit" and it was first considered in Nash (1973). See also discussion in Mahajan and Teneketzis (2007).

6.3 Time-Varying Unobserved Heterogeneity

In the baseline model, a worker's unobserved abilities, η_o and η_c are constant throughout. It is possible to allow them to evolve over time, as in Kahn and Lange (2014) for instance. Whereas in the baseline model learning is relevant mostly early in workers' careers, now learning becomes relevant over their entire lifecycle. In this case, workers' beliefs now evolve according to a Kalman filter.

6.4 Choice of Hours Worked

Another extension that is possible is to allow workers a choice of how many hours to work. In that case, the worker's static choice of hours, h is given by

$$\max_h w_{oc}h - \psi\left(h\right),$$

where $\psi'\left(h\right) > 0$ and $\psi''\left(h\right) > 0$ and $\psi\left(h\right)$ captures a worker's disutility from labor. Part-time employment can be thought of as jobs with few hours. The above setup is able to generate a positive relationship between hours worked and tenure which is present in our data.

6.5 General Equilibrium

The baseline setup models the supply side of the labor market, while the demand side is exogenous. It is possible to extend the model to a general equilibrium setup, by having each occupation produce a differentiated product. A representative consumer derives utility from the consumption of the final good that is given by CES aggregate of the intermediate goods produced by the different occupation using labor. In this case, each product operates in a monopolistically competitive environment and faces a downward slopping demand curve. In such a setup, the price for each occupation's output is now endogenous and therefore so is labor demand and thus the wage. In order to use Gittins indices, we would need to consider a stationary environment in which, from the worker's perspective, prices and wages are exogenous and constant. Consequently, any counterfactuals must be both unexpected and permanent.

6.6 More general within cluster specification

Our formulation of the within cluster setup is not restrictive, in that it does not affect the use of Gittins indices across clusters. The key assumption is that innovations, in this case in the accumulation of human capital and in learning, are independent across clusters (with the exception of general human capital).

As long as this assumption holds, the within-cluster structure can be modified to allow for different forms of human capital formation and learning. In addition, it is possible to allow for human capital accumulation to depend on current occupation or cluster beliefs: for instance, workers with higher beliefs could accumulate more human capital, as would be the case if human capital were endogenous.

7 Conclusion

We introduce a flexible model of the labor market that allows for transferable human capital, information spillovers, and various dimensions of heterogeneity. Despite its generality, our setup remains tractable through the use of Gittins indices. We also document several facts related to wages, as well as occupational mobility.

Our framework could be suited to study several counterfactuals.[32] For instance, it can be used to understand how trade liberalization affects the allocation of workers across different sectors of the economy (Cosar, 2013, Ritter, 2014): consider a one-time unexpected change in the demand for labor services of particular sectors of the economy. In the context of a small open economy that takes world prices as given, this corresponds to a trade liberalization episode. As some sectors suddenly become more attractive than others, there is net worker reallocation. Previously accumulated human capital may be partially transferable to the new sectors, but a part of it is also lost. Similarly, previous work experience is informative about a worker's potential match in the new sectors. Both of these mechanisms affect reallocation, as well as the composition of workers who choose to reallocate. Along similar lines, one could consider the impact of automation and AI on different types of workers, where similar mechanisms are potentially at play (see e.g. Acemoglu and Restrepo, 2019).

In the same spirit, it is possible to consider the cost of entering the labor market in a recession (Kahn, 2010, Wee, 2013). The relative value of non-employment may be particularly high for a subset of workers temporarily, so they spend a considerable amount of time unemployed/out of the labor force. Then, unexpectedly, its value drops, and these workers become employed. The loss of labor market experience in the earlier years may have lasting effects, as it implies both foregone human capital, as well as more dispersed beliefs regarding their productivity in different occupations. This, in turn, leads to lower productivity and more mismatch later in the working life and potentially also more reallocation.

Finally, the above setup can be used to conduct policy experiments, such as the impact of changes in unemployment benefits, which here correspond to an increase in the attractiveness of non-employment,

[32]In what follows, the suggested shocks need to be unanticipated for the use of Gittins indices to still be valid.

the consequences of minimum wages, which might lead to lower employment, but also reduce human capital accumulation and learning or the impact of sector-specific subsidies.

Appendix

A Additional Tables and Figures

	Average
Mean Age	24.52
Standard Deviation Age	2.96
Fraction White	0.71
Fraction Female	0.41
Highest Degree High School	0.51
Highest Degree College	0.20
Mean Real Wage	$12.19
Standard Deviation Real Wages	$6.81
Average Prob of Switch 1-Digit Occupation (annual)	0.20
Average 1- Digit Occupational Tenure	1.79
Standard Deviation 1-Digit Occupational Tenure	1.94

Table 5: Descriptive Statistics, NLSY97

Aggregate	0.035
	(0.001)
Management, Business and Financial (10-950)	0.051
	(0.005)
Professional and Related Occupations (1000-3540)	0.041
	(0.004)
Service Occupations (3600-4650)	0.022
	(0.004)
Sales and Related Occupations (4700-4960)	0.029
	(0.004)
Office and Administrative Support (5000-5930)	0.034
	(0.003)
Construction and Extraction (6200-6940)	0.034
	(0.003)
Installation, Maintenance and Repair (7000-7620)	0.026
	(0.005)
Production Occupations (7700-8960)	0.031
	(0.003)
Transportation & Material Moving (9000-9750)	0.034
	(0.004)

Table 6: Residual Log Wages on Occupational Tenure, Regression Coefficients and associated Standard Errors. NLSY97.

Aggregate	-0.043
	(0.006)
Management, Business and Financial (10-950)	-0.035
	(0.004)
Professional and Related Occupations (1000-3540)	-0.015
	(0.014)
Service Occupations (3600-4650)	-0.036
	(0.011)
Sales and Related Occupations (4700-4960)	-0.048
	(0.006)
Office and Administrative Support (5000-5930)	-0.041
	(0.005)
Construction and Extraction (6200-6940)	-0.045
	(0.016)
Installation, Maintenance and Repair (7000-7620)	-0.053
	(0.012)
Production Occupations (7700-8960)	-0.046
	(0.011)
Transportation & Material Moving (9000-9750)	-0.056
	(0.008)

Table 7: Annual Occupational Switching Probability on Occupational Tenure. First we run a probit regression of the probability of switching occupations on occupational tenure levels zero through five, cubic polynomial in age, race, gender, education and year dummies. Then we regress the coefficients of the occupational tenure levels (zero through five) on occupational tenure. The table presents their coefficients and associated standard errors. NLSY97.

Aggregate	0.0054
	(0.0014)
Management, Business and Financial (10-950)	0.0077
	(0.0072)
Professional and Related Occupations (1000-3540)	-0.0022
	(0.0053)
Service Occupations (3600-4650)	0.0063
	(0.0023)
Sales and Related Occupations (4700-4960)	0.0007
	(0.0014)
Office and Administrative Support (5000-5930)	0.0011
	(0.0018)
Construction and Extraction (6200-6940)	0.0110
	(0.0040)
Installation, Maintenance and Repair (7000-7620)	0.0015
	(0.0079)
Production Occupations (7700-8960)	0.0001
	(0.0022)
Transportation & Material Moving (9000-9750)	0.0029
	(0.0051)

Table 8: Cross-Sectional Variance of Within-Occupation Log-Wage Residuals on Occupational Tenure, Regression Coefficients and associated Standard Errors. NLSY97.

Aggregate	-0.192
	(0.012)
Management, Business and Financial (10-950)	-0.142
	(0.042)
Professional and Related Occupations (1000-3540)	-0.128
	(0.017)
Service Occupations (3600-4650)	-0.189
	(0.036)
Sales and Related Occupations (4700-4960)	-0.206
	(0.033)
Office and Administrative Support (5000-5930)	-0.213
	(0.028)
Construction and Extraction (6200-6940)	-0.233
	(0.045)
Installation, Maintenance and Repair (7000-7620)	-0.151
	(0.029)
Production Occupations (7700-8960)	-0.221
	(0.031)
Transportation & Material Moving (9000-9750)	-0.275
	(0.033)

Table 9: Annual Occupational Switching Probability on Within-Occupation Residual Wage Percentile, Regression Coefficients and associated Standard Errors. NLSY97.

References

Abraham, K. and H. Farber (1987): "Job Duration, Seniority, and Earnings," *American Economic Review*, 77, 278–297.

Acemoglu, D. and P. Restrepo (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33, 3–30.

Altonji, J. and N. Williams (2005): "Do Wages Rise with Job Seniority? A Reassessment," *ILR Review*, 58, 370–397.

Altonji, J. G. and R. A. Shakotko (1987): "Do Wages Rise with Job Seniority?" *The Review of Economic Studies*, 54, 437–459.

Alvarez, F., K. Borovickova, and R. Shimer (2023): "Decomposing Duration Dependence in a Stopping Time Model," *The Review of Economic Studies*, 91, 3151–3189.

Alvarez, F. and R. Shimer (2009): "Unemployment and Human Capital„" *mimeo, University of Chicago.*

——— (2011): "Search and Rest Unemployment," *Econometrica*, 79, 75–122.

Antonovics, K. and L. Golan (2012): "Experimentation and Job Choice," *Journal of Labor Economics*, 30, 333–366.

Banks, J. S. and R. K. Sundaram (1994): "Switching Costs and the Gittins Index," *Econometrica*, 62, 687–694.

Bergemann, D. and J. Valimaki (2008): "Bandit Problems," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Macmillan.

Burdett, K. and D. T. Mortensen (1998): "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39, 257–273.

Carrillo-Tudela, C. and L. Visschers (2023): "Unemployment and Endogenous Reallocation over the Business Cycle," *Econometrica*, 91, 1119–1153.

Carrillo-Tudela, C., L. Visschers, and D. Wiczer (2022): "Cyclical Earnings, Career and Employment Transitions," .

Cosar, A. K. (2013): "Adjusting to trade liberalization: Reallocation and labor market policies," *mimeo, University of Chicago, Booth School of Business.*

Dickstein, M. J. (2021): "Efficient Provision of Experience Goods: Evidence from Antidepressant Choice," *mimeo, NYU Stern.*

ECKARDT, D. (2023): "Training Specificity and Occupational Mobility: Evidence from German Apprenticeships," *mimeo, University of Warwick*.

EECKHOUT, J. (2006): "Employer Learning and General Human Capital," *mimeo, University of Pennsylvania, Department of Economics*.

EECKHOUT, J. AND X. WENG (2022): "Assortative Learning," *Economica*, 89, 647–688.

FELLI, L. AND C. HARRIS (1996): "Learning, Wage Dynamics, and Firm-Specific Human Capital," *Journal of Political Economy*, 104, 838–868.

GERVAIS, M., N. JAIMOVICH, H. E. SIU, AND Y. YEDID-LEVI (2016): "What Should I Be When I Grow Up? Occupations and Unemployment over the Life Cycle," *Journal of Monetary Economics*, 83, 54–70.

GITTINS, J. C. (1979): "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society. Series B (Methodological)*, 41, 148–177.

GITTINS, J. C., K. GLAZEBROOK, AND R. WEBER (2011): *Multi-armed bandit allocation indices*, John Wiley & Sons.

GITTINS, J. C. AND D. M. JONES (1974): "A Dynamic Allocation Index for the Sequential Design of Experiments," in *Progress in Statistics,*, ed. by J. M. Gani, K. Sarkadi, and I. Vincze, North-Holland, Amsterdam, 241–266.

GORRY, A., D. GORRY, AND N. TRACHTER (2019): "Learning and Life Cycle Patterns of Occupational Transitions," *International Economic Review*, 60, 905–937.

GROES, F., P. KIRCHER, AND I. MANOVSKII (2015): "The U-Shapes of Occupational Mobility," *Review of Economic Studies*, 82, 659–692.

GUVENEN, F., B. KURUSCU, S. TANAKA, AND D. WICZER (2020): "Multidimensional Skill Mismatch," *American Economic Journal: Macroeconomics*, 12, 210–44.

JOVANOVIC, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, Part 1, 972–990.

KAHN, L. B. (2010): "The long-term labor market consequences of graduating from college in a bad economy," *Labour Economics*, 17, 303–316.

——— (2013): "Asymmetric Information between Employers," *American Economic Journal: Applied Economics*, 5, 165–205.

KAHN, L. B. AND F. LANGE (2014): "Employer Learning, Productivity and the Earnings Distribution: Evidence from Performance Measures," *Review of Economic Studies*, 81, 1575–1613.

Kambourov, G. and I. Manovskii (2008): "Rising Occupational and Industry Mobility in the United States: 1968-1997," *International Economic Review*, 49, 41–79.

——— (2009a): "Occupational Mobility and Wage Inequality," *Review of Economic Studies*, 76, 731–759.

——— (2009b): "Occupational Specificity of Human Capital," *International Economic Review*, 50, 63–115.

Karatzas, I. (1984): "Gittins Indices in the Dynamic Allocation Problem for Diffusion Processes," *The Annals of Probability*, 12, 173–192.

Klein, N. and S. Rady (2011): "Negatively Correlated Bandits," *The Review of Economic Studies*, 78, 693–732.

Kramarz, F., F. Postel-Vinay, and J.-M. Robin (2014): "Occupational Mobility and Wage Dynamics Within and Between Firms," *mimeo, Science Po.*

Lise, J. and F. Postel-Vinay (2020): "Multidimensional skills, sorting, and human capital accumulation," *American Economic Review*, 110, 2328–2376.

Mahajan, A. and D. Teneketzis (2007): "Multi-Armed Bandit Problems," in *Foundations and Applications of Sensor Management*, ed. by A. O. H. III, D. A. Castanon, D. Cochran, and K. Kastella, Springer-Verlag, 121–151.

McCall, B. P. (1990): "Occupational Matching: A Test of Sorts," *Journal of Political Economy*, 98, 45–69.

Miller, R. A. (1984): "Job Matching and Occupational Choice," *Journal of Political Economy*, 92, 1086–1120.

Nash, P. (1973): "Optimal allocation of Resources Between Research Projects." Ph.D. thesis, University of Cambridge.

Pandey, S., D. Chakrabarti, and D. Agarwal (2007): "Multi-armed bandit problems with dependent arms," in *Proceedings of the 24th international conference on Machine learning*, 721–728.

Papageorgiou, T. (2014): "Learning Your Comparative Advantages," *Review of Economic Studies*, 81, 1263–1295.

——— (2018): "Large Firms and Within Firm Occupational Reallocation," *Journal of Economic Theory*, 174, 184–223.

——— (2022): "Occupational Matching and Cities," *American Economic Journal: Macroeconomics*, 14, 82–132.

Pastorino, E. (2015): "Job Matching Within and Across Firms," *International Economic Review*, 56, 647–671.

——— (2024): "Careers in Firms: The Role of Learning about Ability and Human Capital Acquisition," *Journal of Political Economy*, 132, 1994–2073.

Pavan, R. (2011): "Career Choice and Wage Growth," *Journal of Labor Economics*, 29, 549–587.

Poletaev, M. and C. Robinson (2008): "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000," *Journal of Labor Economics*, 26, 387–420.

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

Ritter, M. (2014): "Offshoring and occupational specificity of human capital," *Review of Economic Dynamics*, 17, 780–798.

Rubinstein, Y. and Y. Weiss (2006): "Post schooling wage growth: Investment, search and learning," *Handbook of the Economics of Education*, 1, 1–67.

Sanders, C. and C. Taber (2012): "Life-cycle wage growth and heterogeneous human capital," *Annual Review of Economics*, 4, 399–425.

Scheonberg, U. (2007): "Testing for Asymmetric Employer Learning," *Jounal of Labor Economics*, 25, 651–691.

Silos, P. and E. Smith (2015): "Human Capital Portfolios," *Review of Economic Dynamics*, 18, 635–652.

Topel, R. (1991): "Specific capital, mobility, and wages: Wages rise with job seniority," *Journal of political Economy*, 99, 145–176.

Traiberman, S. (2019): "Occupations and Import Competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.

Wee, S. L. (2013): "Born under a bad sign: The cost of entering the job market during a recession," *University of Maryland mimeo.*

Whittle, P. (1980): "Multi-Armed Bandits and the Gittins Index," *Journal of the Royal Statistical Society. Series B (Methodological)*, 42, 143–149.

——— (1981): "Arm-Acquiring Bandits," *The Annals of Probability*, 9, 284–292.

——— (1982): *Optimization over Time: Programming and Stochastic Control*, New York: Wiley.

Yamaguchi, S. (2012): "Tasks and Heterogeneous Human Capital," *Journal of Labor Economics*, 30, 1–53.