

Learning on the Job? Employee Mobility in the Asset Management Industry

Aaron K. Chatterji

Fuqua School of Business
Duke University
100 Fuqua Drive, Box 90120
Durham, NC 27708
ronnie@duke.edu

Rui J.P. de Figueiredo, Jr.

Haas School of Business
S-545 Student Services Bldg. #1900
UC Berkeley
Berkeley, CA 94720
rui@haas.berkeley.edu

Evan Rawley

Columbia Business School
Uris Hall #702
New York, NY 10027
erawley@columbia.edu

We present a new mechanism by which prior employment can influence transitions to other firms. We propose that some employees divert effort toward unproductive activities to learn about their own fitness for alternative employment. Based on the results of this costly learning experience, or “experiment,” some employees will transition into other firms or launch their own ventures, while others will remain at the incumbent firm. We develop a theoretical model to explicate these propositions, and test them using four datasets from the mutual fund and hedge fund industries. We find evidence that managers who engage in excessive risk-taking at mutual funds are subsequently more likely to join or start hedge funds, even though there is little evidence that this risk-taking is intended to signal quality to outside observers. Taken together, our findings suggest that learning about one’s own fitness for alternative employment, through experimentation on the job, is an important mechanism for enabling employee mobility.

Keywords: employee mobility, learning, financial services, employee entrepreneurship, spawning,

1. Introduction

The drivers and consequences of employee mobility have long been a lively topic of scholarly inquiry (Almeida & Kogut, 1999; Arrow, 1962; Marx, Strumsky, & Fleming, 2009; Rosenkopf & Almeida, 2003; Singh & Agrawal, 2011; Stephan, 1996). Employees might leave their employer to join a competitor, a firm in another industry or to start a new company (Campbell, Ganco, Franco, & Agarwal, 2012). Scholars have proposed many motivations for this kind of employee mobility, including better compensation (e.g., Zenger, 1992), lack of person-organization fit (e.g., Sheridan, 1992), disagreements

with top management (Klepper & Thompson, 2010) and the desire for greater autonomy and responsibility (e.g., Elfenbein, Hamilton, & Zenger, 2010). We build on this literature to propose a complementary mechanism by which prior employment can impact mobility—whereby select employees engage in costly learning activities, or “experiments,” to reveal their fitness for a new job. In particular, we argue that employment at existing firms can provide a laboratory for some employees to learn about their own capabilities and preferences for transitioning to another firm. We develop a formal model to demonstrate that based on the results of learning on the job through experimentation, some employees will transition while others will remain at their employer. Importantly, this argument extends the literature on employee mobility beyond only considering “why” individuals transition to understanding “how” they do so.

Using data from the asset management industry, we find evidence supporting the idea that learning on the job through experimentation facilitates employee mobility. We study the transitions of mutual fund managers to hedge funds from 1990-2011 with data from CRSP Survivorship-Bias-Free U.S. Mutual Fund (herein CRSP) Database and three hedge fund databases. As would be the case with most panel datasets with broad industry coverage, we cannot measure individual learning and experimentation directly. Instead, we operationalize learning by measuring managers’ risk-taking behavior, viewing excessive risk-taking—the kind that reduces risk-adjusted performance—as an “experiment” from which portfolio managers learn about themselves. We find that a subset of mutual fund managers experiment in their jobs by taking on greater risk, at levels more commensurate with hedge fund managers. These managers are more likely to transition to hedge funds: a one-standard-deviation increase in risk-taking corresponds to a 14% increase in the transition rate.¹ We are able to refute alternative explanations that could be consistent with similar empirical patterns, such as signaling, mean reversion or the argument that

¹ The base transition rate for hedge fund transitions in our sample is 3.8% (486 transitions/12,754 mutual fund managers). From Table 2, Column 2, doubling risk exposure, from a mean of 1.64%/month (Table 1a), increases the transition rate by 0.64%. Therefore, a one-standard-deviation increase in risk exposure, 1.35%/month (Table 1a), leads to a 14% increase in the transition rate ($0.64/3.8 \times 1.35/1.64$).

poorly performing mutual fund managers are simply “gambling” as opposed to learning from experimentation.

One important feature of our dataset is that it allows us to closely match, on observable characteristics, those individuals who transition into hedge funds to those who experimented with excessive risk-taking but stayed in the mutual fund industry. We find the matched set of individuals who remain employed in mutual funds reduce their risk levels by 24 basis points per month, a 15% reduction in risk-taking relative to the mean level, and improve their performance by 5 basis points per month per unit of risk, a pattern strikingly consistent with our proposed learning mechanism.

By proposing and finding support for a new within-firm mechanism by which prior employment influences employee mobility, this research contributes to the literatures on employee mobility and entrepreneurial spawning. Notably, our explanation can account for why some employees remain at a firm while others leave. Because we have pre-transition performance data on individuals, our analysis can deal directly with unobservable person-specific quality differences. We also use firm fixed effects to account for within-firm differences between employees who stay and those who transition into alternative employment.

In the following section, we recount the key implications from prior work on employee mobility and describe our theory. We then discuss our empirical setting, the asset management industry. Next, we develop a model tailored to that industry to derive testable implications in our empirical setting. We then outline the empirical methods and analyses, summarize the key results, and discuss implications from our work.

2. Understanding Employee Mobility

There is a sizable literature on employee mobility, both to existing firms (Almeida & Kogut, 1999; Marx et al., 2009; Rosenkopf & Almeida, 2003; Singh & Agrawal, 2011) and into entrepreneurship (Campbell, et al., 2012; Carnahan, Agarwal, & Campbell, 2012; Ganco, 2013). The motivations for employees to leave their employers include poor performance (Williams & Livingstone, 1994), issues regarding

compensation (Zenger, 1992), weak person-organization fit (Sheridan, 1992), disagreements with top management and the desire for greater autonomy (Elfenbein et al., 2010; Klepper & Thompson, 2010). Important contingencies that impact the decision to transition include the level of firm-specific human capital (Campbell, Coff, & Kruscynski, 2012) and the relative bargaining power between employee and employer (Campbell et al., 2012).²

We propose a novel theoretical mechanism to explain patterns of employee mobility: learning about oneself on the job through experimentation.³ Our theory does not exclude other explanations, but instead provides a complementary perspective to account for broader patterns in employee mobility and entrepreneurship. While prior work has carefully considered issues of economic and non-pecuniary incentives to stay at the employer or leave, we recast employment as a laboratory for experimentation and self-learning. Our view is that individuals can conduct experiments on the job to learn about themselves, and this knowledge can inform the decision to stay or leave the current employer. In this spirit, our explanation provides insight into how employees transition to alternative employment, an important complement to an extensive prior literature on why employees transition.

Employment as a Laboratory for Learning about Alternatives

Whereas the prior work on employee mobility and entrepreneurship allows for employees to learn about alternative employment in the industry and entrepreneurial opportunities (c.f. Gompers et al. 2005), we propose that taking costly actions on the job might provide an opportunity for an employee to learn specifically about their own fitness for a new job. As we discuss below, we expect that experimentation

² Cotton and Tuttle (1986) provide a useful categorization of the various motivations for employee turnover. Their three categories are external correlates (e.g., the unemployment rate), work-related correlates (e.g., task repetitiveness) and personal correlates (e.g., age). Within the expansive literature on turnover in management and economics, our work is most closely related to studies on the mobility of highly skilled workers so we limit our review of the literature to these papers for the sake of brevity.

³ As will become clear as we develop our theory, there are two key differences between our argument and the existing literature. First, learning in our model is about the innate characteristics of the individual, in contrast to what is acquired through work experience. These innate characteristics can encompass what might be deemed as “talents” or “preferences.” In our interviews within the asset management industry, managers emphasized both of these concepts, mentioning “talent or intuition for shorting” and “the comfort with running high risk portfolios.” A second unique feature of our work is that learning these characteristics may be relevant to the decision to transition to alternative employment but not to eventual success, conditional on transition. Thus, some of these personal characteristics may be relevant only to the individual herself.

would be most prevalent preceding transitions to different kinds of jobs. In practice, employment provides several opportunities for gaining insights into one's suitability for alternative jobs. For example, in asset management, a mutual fund employee may have the opportunity to employ particular trading strategies that allow her to learn about hedge fund investing—perhaps by short-selling, taking on more leverage, or investing in illiquid assets. Given an employee's existing knowledge base, the information gained about one's fitness for alternative employment through experimenting on the job could be a crucial factor in influencing the employee's decision to transition.

“Learning” is used in a wide variety of academic literatures. We define learning narrowly in our work to refer to activities that generate information about one's fitness for a new job, specifically information about (i) the employee's talents and (ii) the employee's preferences. Because learning is relatively more important for employees considering a transition to a new type of job, our theory suggests that employees will need to gather more information about their fitness for a job in a different industry segment than about their fitness for the same job at a different company in the same industry segment. We can test this ancillary prediction in our empirical analysis to verify the underlying mechanisms in our core argument.

Implications

Our argument that prior employment can be a setting where employees learn about their suitability for alternative employment has three testable implications. First, a subset of employees at incumbent firms should engage in learning activities. Next, some proportion of these employees should transition to alternative employment. Finally, those employees that remain at incumbent firms should reduce their costly learning activities because the learning process reveals whether or not they are well suited for a transition. In the empirical section, we test these three propositions directly.

3. Empirical Context

To test our propositions about employee mobility, we explore the careers of individuals working in the asset management industry. In particular, we look at portfolio managers of mutual funds and study their career transitions to leadership positions in hedge funds. As of 2012, mutual funds held a total of \$24

trillion (\$13 trillion in the United States) in assets under management.⁴ Globally, hedge funds have over \$2 trillion under management, with the vast majority of the largest funds located in the United States.⁵ Within the broad umbrella of the asset management industry, mutual funds and hedge funds share many common features. These organizations assemble money from investors to invest in various financial instruments and typically charge a management fee that is a fixed percentage of assets under management.

However, mutual funds and hedge funds differ in important ways as well, primarily in terms of typical investing strategies, disclosure requirements, and business models. In terms of investing strategies, mutual funds invest in equities, bonds, and other financial instruments, typically taking long positions using investors' capital. Hedge funds employ a more diverse set of riskier strategies, including taking short positions, investing in illiquid assets, and amplifying returns through extensive leverage (i.e., debt financing). In addition, hedge fund managers are able, expected, and often willing to change their strategy (sometimes referred to as "strategy drift") whereas mutual fund managers tend to be much more circumscribed. Mutual fund managers need to better understand—both intellectually and psychologically—precisely these risky and more dynamic investment strategies in order to make an informed decision about whether to pursue a career in hedge fund management.

Because mutual funds are relatively transparent investment vehicles, they are permitted to market their services widely and can have large numbers of investors. However, in return for these privileges, U.S. mutual funds must register with the Securities and Exchange Commission (SEC), disclose their fees, and face strict limitations on short-selling, use of leverage, and incentive fees.⁶ The SEC regulates hedge funds but does not require them to register their financial statements or subject them to the investment and compensation requirements mutual funds face. As a result, hedge funds typically accommodate fewer investors (i.e., "qualified purchasers" or "accredited investors") with much higher minimum investment

⁴ <http://www.icifactbook.org/> (Last accessed May 15, 2013)

⁵ <http://www.economist.com/news/leaders/21568740-investors-have-paid-too-much-hedge-fund-expertise-better-focus-low-costs-star> (Last accessed May 15, 2013)

⁶ http://www.ici.org/files/faqs_hedge (Last accessed May 15, 2013)

thresholds and fees. Moreover, under most conditions, mutual funds cannot receive performance fees from investors. By contrast, approximately one half of hedge fund compensation comes from performance fees. Thus, mutual fund managers who become hedge fund managers are transitioning to a related, but different type of job with different investment strategies and compensation arrangements. For example, most hedge funds charge a fixed management fee on assets under management and an additional performance fee (e.g. 2% and 20% model), whereas most mutual funds charge only a fee based on assets under management. For a more comprehensive treatment of differences between hedge funds and mutual funds, see Fung and Hsieh (1999).

With an eye to our formalization that follows, it is worth noting that for asset managers a key choice variable is how much investment risk to take in managing a portfolio. Having determined the choice of risk, the manager will choose investments to maximize the expected returns given her beliefs concerning individual risk-return choices. Notably, at some point in the idiosyncratic risk dimension, the maximum excess return for a given level of risk is decreasing. While the existence of a downward sloping portion of the optimal risk-return frontier may seem to contradict the standard relationship between raw returns and systematic risk, the relationship between *excess* returns and *idiosyncratic* risk must have a downward sloping component given limits on leverage, and investment managers must choose lower excess return investments to generate risk levels commensurate with hedge fund investments (see Palomino and Prat 2003 for a formal derivation, and de Figueiredo, et. al. 2014 for a more detailed discussion in the context of the investment management industry). Moreover, it is likely that many hedge fund managers, who face strong incentives to maximize returns because of convex incentives structures, operate at the peak of the risk-return frontier. Indeed, de Figueiredo, et. al. (2014) provide evidence that hedge fund managers appear to be operating at risk levels that are at or very close to the peak return—when managers have incentives to take even small amounts of additional risk their returns tend to *go down*. These theoretical and empirical results imply that in this setting, the costly, but informative, experiment for mutual fund managers considering a pursuing a career in hedge funds involves increasing

portfolio risk exposure at the cost of harming performance. The result of the experiment reveals the mutual fund manager’s fitness for managing a hedge fund, which could be a better or worse career move for the manager depending on their type.

4. A Model of Learning for Employment Transitions in Asset Management

In order to make precise predictions about the behavior and performance of asset managers, we develop a simple reduced-form model of the learning and risk decisions of a mutual fund manager. The two primary features of our model are that (1) a manager may engage in activities on the job which provide information about their fitness for an employment transition,⁷ and (2) such activity is costly in the sense that it reduces her performance as an employee. When both (1) and (2) are present, we say the employee is “experimenting.” We are interested in a number of outcomes, including (i) who decides to learn on the job to discover their fitness for transition, (ii) who decides to make the transition, (iii) how those who make the transition behave once they have left their prior employer, and (iv) how the behavior of those who learn on the job but choose to remain with their employer changes over time.

Consider a representative mutual fund manager who must generate investment returns by choosing some level of risk v . In choosing v , the managers face an *efficient frontier* f —a mapping from risk into expected returns r ; in other words, we have: $f: v \rightarrow r$. We assume f is concave and single-peaked with maximum at v^{max} and corresponding risk r^{max} .⁸ The payoffs to the manager are a linear function of the returns they generate, so for simplicity, we say the payoff to the manager is r .⁹

⁷ One key implication of this assumption is that learning on the job will be the most efficient way for an employee to gain the knowledge they need to make a decision about a transition. For some types of “general” learning—for example, general risk aversion, comfort managing people in a start-up environment, and so on, existing employment may not be the best place to learn about one’s preferences and talents—thus, our model applies best when a manager needs to gain some “context specific” knowledge about their fitness for alternative employment in a setting that is related to, but not identical to their current job.

⁸ The assumption of single-peakedness follows Palomino and Prat (2003).

⁹ While we assume the agent is risk-neutral, this is primarily for analytic convenience. If the agent is risk-averse the results that follow will hold as long as the risk-aversion is “not too large” in the sense that the investment manager is not operating so far below the maximum of the efficient frontier that the increase in risk required for the experiment will result in higher expected returns.

The (portfolio) manager can either work for a mutual fund or become a portfolio manager of a hedge fund. There are two periods. In the first period, a mutual fund manager can decide to experiment on the job in order to learn something about her fitness for being a hedge fund manager. As a hedge fund manager, she would be one of two types $h \in \{H, L\}$, where a manager is of type H , a “high” type, with probability p , and a “low” type L otherwise. We assume that if f_M is the efficient frontier a portfolio manager faces when working for a mutual fund, then f_L and f_H are defined such that the maximum for f_h is located at the point $(\alpha_h v^{max}, \alpha_h r^{max})$ for $h \in \{H, L\}$ and where $\alpha_L < 1 < \alpha_H$. This setup means some managers are better off as mutual fund managers in the second period and some managers are better off as hedge fund managers.¹⁰

In the first period, if the mutual fund manager decides to experiment, she must take a level of risk at least as high as $v^{exp} > v^{max}$, with corresponding $r^{exp} < r^{max}$. If she experiments, she learns her type as a hedge fund manager and, in turn, can decide at the beginning of the second period whether to switch to being a hedge fund manager or remain a mutual fund manager.

To further analyze the decisions of the portfolio manager, we start by outlining some basic features of this model. First, without alternative incentives, portfolio managers will always choose a level of risk corresponding to the peak return of the efficient frontier they face. Second, if a mutual fund manager chooses to experiment, she will always choose exactly v^{exp} , because taking less risk will provide no information, and taking more risk will provide no additional information but will lead to lower returns.

Comparing the choice of experimenting to not, we have that a mutual fund manager will experiment if and only if:

$$2r^{max} < r^{exp} + p\alpha_H r^{max} + (1 - p)r^{max}. \quad (1)$$

¹⁰ In the model we assume an exact correspondence between returns and payoffs. In fact, the incentive schemes may be different across different employment types (e.g., mutual funds and hedge funds). While this could be incorporated with additional notation, for simplicity we assume that the returns correspond to the utility. If the payout rates based on performance differ, then the key assumption is that for some subset of employees the range of payoffs to being a hedge fund manager span the payoffs to being a mutual fund manager.

The left-hand side of (1) is the payoff to the manager if she remains a mutual fund manager. The right-hand side is the payoff an experimenting mutual fund manager earns in the first period plus the expected outcome conditional on experimentation in the second period. Rearranging (1) and solving for p gives the following condition for experimentation:

$$p > \frac{1-r^{exp}/r^{max}}{\alpha_H-1} . \quad (2)$$

The expression reflects the tradeoff between the loss from experimenting in the first period—as given by the numerator—relative to the gain of being able to switch to running a hedge fund when the manager is a high type—as given by the denominator. (2) shows that if it is sufficiently likely the mutual fund manager will be a high type, she will pay the costs of experimentation and will not experiment otherwise.

Based on the logic of (2), we have Proposition 1:

Proposition 1. *In the model of learning and transitions for asset management, the following results will hold:*

- (a) *On average, mutual fund managers who become hedge fund managers will have higher risk and lower returns as mutual fund managers in the first period compared with those who do not. More generally, all else equal, all mutual fund managers who experiment will have lower returns than those who do not.*
- (b) *Mutual fund managers who experiment in the first period and do not switch will reduce risk and increase their performance in the second period.*
- (c) *Of mutual fund managers that become hedge fund managers, they will take higher risk and generate higher average performance as hedge fund managers than they did as mutual fund managers.*

Proposition 1 provides us with three testable predictions concerning the risk and return characteristics of variously situated mutual fund managers. The implication of the Proposition is that when there is sufficient reward (a high enough outcome in the good state), sufficient uncertainty (a spread between good and bad outcomes which make learning worthwhile) and some potential cost (a requirement to operate sub-optimally during the learning period), we can predict who will experiment and what they will do once the outcome of the experiment is known.

In addition, the model also provides direct implications for understanding employees who make transitions *without* experimenting. Intuitively, these will be employees where the expected value of transitioning without experimentation is higher than for experimenting or doing nothing. We summarize this in Corollary 1.

Corollary 1. *Transitions will not be preceded by experimentation if:*

- (a) *the cost of learning is high, and*
- (b) *the expected value of transitioning is positive.*

The additional prediction of Corollary 1 is that when (1) the value of learning is not very high (i.e., when the likelihood that one's fitness for a new job is high) and (2) the new job is attractive, conditional on good fit, then one will make the transition *without* experimenting. In the setting we study below, this suggests that transitions from mutual fund to mutual fund positions, which do not entail much fitness risk, will *not* be preceded by costly experimentation. Thus in these transitions we would not expect to see changes in risk *ex ante*.

5. Data and Empirical Methods

To test our theoretical arguments, we use the complete CRSP Survivorship-Bias-Free U.S. Mutual Fund (CRSP) Database, which contains monthly size, performance, and manager data for all 30,572 mutual funds opened in the United States 1964-2011. Because we rely on the individual characteristics that

influence transitions from mutual funds, we exclude funds that do not list individual managers and funds that do not report any specific portfolio manager(s) responsible for the fund for at least 12 months.¹¹ Using the 12,754 unique portfolio manager names in the CRSP dataset, we converted the data to 57,180 unique manager-fund dyads from 20,791 funds. These 57,180 observations form the core of our dataset, though we examine the relationship between risk-taking and transitions in the cross section and over time using four different units of analysis—the manager level, the manager-month level, the manager-fund dyad level, and the manager-fund-month level.

Our hedge fund dataset includes data from three major hedge fund data vendors: Hedge Fund Research, Barclays, and Morgan Stanley Capital International, which report fund size and performance over time, but only report concurrent information on their portfolio managers. We identified manager transitions by hand matching at the manager level. An extensive pre-cleaning of the name fields led to nearly perfect agreement between two independent coders. We researched ambiguous matches one by one online. Of the 558 matches identified, 486 mutual fund managers were in the core dataset.

Performance and risk are measured at the manager-fund dyad level in the usual way for the mutual fund industry using the three factors from Fama and French (1996) and Carhart's (1997) momentum factor. Raw returns are regressed on these four factors longitudinally, fund by fund, to estimate a fund's systematic risk exposure. Excess returns are computed as the difference between raw returns and the benchmark risk-adjusted return predicted by this four-factor asset-pricing model. Similarly, we follow the standard approach for measuring hedge excess returns using Fung and Hsieh's (2003) seven-factor model plus a liquidity risk-factor from Pastor and Stambaugh (2003).¹² Taking the difference every month, and computing a moving average over time, provides a time-varying performance measure, whereas average performance relative to the factor benchmark ("alpha") captures performance meaningfully in the cross section. Idiosyncratic risk is calculated as the standard deviation

¹¹ CRSP tracks data by share class. We consolidate share classes to funds. The results are robust to excluding very small share classes, funds and/or firms (e.g. those with \$10M in assets under management).

¹² Since hedge funds frequently hold illiquid assets and self-report returns, we also correct for serial correlation in hedge fund raw returns using an autoregressive lag one (AR1) model, as is standard in the hedge fund literature.

of excess returns, which is the residual risk a manager takes over and above the systematic market exposure a manager in an index fund would face. We use the 12-month moving-average standard deviation of excess returns to measure time-varying risk exposure, though the results are not sensitive to the number of months included in the rolling average. Performance is measured by the information ratio, which is equal to excess returns divided by idiosyncratic risk exposure.

Assets under management (AUM) are reported for most fund-months. When AUM is missing, we create a dummy variable, “missing AUM,” and impute the AUM level to be the mean of the rest of the core dataset. We also include year fixed effects in panel specifications, and control for the total number of months the manager is listed as the portfolio manager (“tenure”). At the firm level, we capture whether the mutual fund is headquartered in New York or elsewhere, and the total number of managers who transition from that firm to a hedge fund. Our strongest tests include firm, fund, manager-fund, or manager fixed effects, so we leave variables that do not vary within these categories aside in such tests.

Tables 1a and 1b show descriptive statistics for the key variables used in the mutual fund and hedge fund samples, respectively. The left panel of Table 1a displays data for the full sample, while the right panel shows data for the managers who transition (which we abbreviate as “MF->HF transitions” throughout the tables). Table 1b shows summary statistics for transitioning managers before (left panel) and after (right panel) their transition to hedge funds. Casual observation suggests that there may be some evidence for our hypotheses even in the raw data. Managers who transition from mutual funds to hedge funds take on more idiosyncratic risk *ex ante* than mutual fund managers who do not transition (2.04% per month vs. 1.64% per month, Table 1a), and these individuals increase their risk-taking, from 1.86% per month to 3.19% per month (Table 1b), and improve their performance, from negative 4 basis points per month (per unit of idiosyncratic risk) to 11 basis points per month (Table 1b).

We begin our analysis by showing the basic correlation between risk-taking and the probability of making a mutual fund to hedge fund transition at the manager level using each manager’s largest fund, by average (over time) AUM with and without firm fixed effects. Focusing only on the fund with the largest

AUM is a crude method of capturing manager behavior; however, it represents one plausible way of studying manager-level activity when some managers manage multiple funds. We verified that this approach did not bias the results, by replicating the results elsewhere using all manager-fund dyads, and we show the results are robust to using all the funds managed by all the managers in the full panel. We then extend the analysis to the manager-fund and manager-fund-month level with fund fixed effects. These tests are expected to show that transitioning managers take on more risk than non-transitioning managers, even when compared to other managers in the same firm, and even other managers of the same fund.

The ideal empirical test would be to randomly assign risk-taking to mutual fund managers so that we could directly observe the causal implication of risk-taking on learning and transitions. Although the correlation between risk-taking and transitions is informative, risk-taking may be correlated with unobservable person-specific characteristics that confound transition rates, but that we cannot control for directly even with fund fixed effects. Moreover, we do not know the precise period during which a mutual fund manager experiments, and do not want to assume our theory is true and that risk-taking is equivalent to experimentation.¹³ To account for person-specific heterogeneity directly and give us more insight into the learning process, we take two approaches. First, we analyze within-manager changes in risk-taking and performance for mutual fund managers who transition into hedge funds, an analysis that explicitly absorbs unobservable person-specific characteristics of the treatment (i.e., transitioning) group. Second, we use panel data to identify a set of managers who did not transition, but whose profiles were similar to managers who did transition, based on a matching algorithm, and analyze within-manager

¹³ We empirically observe when a manager leaves her mutual fund, and expect that in many cases the timing of the departure will be immediately subsequent to the key period of experimentation. However, it is probably also true that in many cases managers experiment multiple times, experiment for a prolonged period of time, or even experiment a single time earlier in their career. For example, a manager may experiment extensively early in her career, learn her type, spend a year or two searching for a partner with whom to start a hedge fund, and then experiment again just before leaving to make sure she feels comfortable with the move. In these cases we would not be able to sensibly test the correlation between within-manager changes in *ex ante* risk-taking related to a transition without putting more (ad hoc) structure on the analysis.

changes in that set.¹⁴ The results of the treatment on the treated analysis along with the analysis of the counterfactual group provide some insight into the experimentation and learning process behind the transitions we observe.

To identify the set of managers in the matched sample, we analyze the data at the manager-fund-month level, which gives equal weight to each of the dyadic manager-fund transitions that occurred more than one year before the end of our sample period. We then use Coarsened Exact Matching (Iacus, et. al., 2011; David, et. al., 2013) and propensity score matching (Rosenbaum and Rubin, 1983) to identify managers who are identical to the managers who transition based on all the observable characteristics of the individuals, their funds, and firms that might plausibly predict a transition.¹⁵ All of the dyad-months from non-transitioning managers are potential matches with the “treatment” group, excluding managers from funds in which other managers transition, and observations with less than 12 months of subsequent performance. Following de Figueiredo and Rawley (2011), we match treatment group observation to control-group observation one to one asynchronously to determine the control group’s “match date.” By isolating the point in time when each control group manager appears to be experimenting, we create a counterfactual for how risk-taking and performance changes for those who experiment but do not transition. We match each treatment observation to a single control-group observation—with a maximum of one match per manager—without replacement to obtain our matched sample.¹⁶

¹⁴ Some mis-measurement in the “true” experiment date is not a problem for the matched sample analysis, as long as there is enough information about transitions in the treatment group observations close to the transition date. Importantly, even if the last experiment occurs more than one year before the transition date, the bias in the match date would work against us because we would be generating spurious match dates that would create attenuation bias in the “second stage” of the matched sample analysis.

¹⁵ Coarsened Exact Matching (CEM) is similar to propensity score matching but requires fewer post-estimation assumptions about how to define a match by matching observations *exactly* across a number of dimensions simultaneously. CEM is also much faster computationally, which allowed us to easily verify that our results are not sensitive to a large number of variations in the matching approach.

¹⁶ The advantage of one-to-one matching is that we do not have to reweight our sample by the inverse probability of selection when specifying our model of *ex post* behavior. The disadvantage of one-to-one matching is that the matched sample generated by CEM is not identical across runs, since the program randomly selects which extra control group observations to drop. We verify that our results are not sensitive to randomly dropping different control group observations from the final match.

The base (CEM) match relies on coarsened measures of age, AUM, idiosyncratic risk, excess returns, firm spawns, time, a NYC dummy and a missing AUM dummy, resulting in 18,706 possible strata. 794 strata and 1,062 manager-fund dyads matched the treatment group. The “second stage” results are not sensitive to different levels of coarsening. However, match quality, as measured by differences in the means of the regressors, was lower with other partitions of the variables we match upon. We obtained similar levels of match quality and numbers of matches using propensity score matching.

To ensure that our subsequent analyses, meant to identify changes in behavior in the control group, are not mechanically picking up mean-reversion by matching on elevated levels of risk-taking, we also construct alternative “ex-risk” matched samples that are created using the base match variables *excluding* ex ante risk-taking.¹⁷ Similarly we also construct “ex-performance” matched samples by excluding the excess return variable from the base matching algorithm.

6. Results

Table 2 shows the basic manager, manager-fund and manager-fund-month-level tests of our baseline prediction that increased risk-taking leads to transitions. Column 1 shows the results from a probit model with one observation per manager in which the outcome variable is equal to 1 for any manager who makes a transition to a new or existing hedge fund and 0 otherwise. Standard errors are clustered at the firm level and marginal effects are reported. (For presentation purposes all coefficient estimates in this table are multiplied by 100). Doubling idiosyncratic risk exposure leads to a 0.59% increase in the rate of transitioning. Compared to the base transition rate of 3.8% , this effect represents a 15% increase in the probability of making a transition. Including 1,256 firm fixed effects in the OLS regression in Column 2

¹⁷ We address mean-reversion in the base match sample by employing a long-run measure of risk-taking in the second stage of the matched sample analyses, a measure which is less susceptible to mean reversion. Specifically, we compare the change in long-run, pre-match risk-taking with long-run, post-match risk-taking. The average control group observation is 46 months old at the time of the match date and is 103 months old at the end of the sample. Thus, we are comparing average risk-taking over 46 months with average risk-taking over 57 months, and not mechanically inducing mean reversion by matching on 12-month rolling-average risk-taking in the first stage. Moreover, closer inspection of the data suggests mean reversion is unlikely to explain our results. Only 2.5% of the matches occur at the control-group observation’s peak risk date. In fact, almost 60% of the control-group observations take their maximum risk, on a rolling-average basis, after the match date.

increases the point estimate on idiosyncratic risk to 0.64%, which means that a one-standard-deviation increase in risk exposure leads to a 14% increase in the probability of making a transition. After expanding the data to the manager-fund-level, Column 3 shows a within-fund analysis by including fund fixed-effects with very similar results. In Column 4 we show the equivalent test using the full panel at the manager-fund-month-level including year and fund fixed effects. The sign, and statistical significance of the estimate is very similar to the more highly aggregated analyses, and the point estimate is about 40% larger. Standard errors are clustered at the unit of analysis in all specifications.

In Table 3, we use the panel of mutual fund managers who become hedge fund managers before and after their transition date with tenure and size controls, as well as time and manager-fund fixed effects, to analyze whether risk-taking and performance increase for mutual fund managers who become hedge fund managers, as predicted by the model. While the analysis is only a within-manager estimate of the effect of the treatment on the treated it is, nevertheless, informative about how behavior changes in a mutual fund manager who makes a transition to a hedge fund. Column 1 shows within-manager changes in idiosyncratic risk increases 1.58%/month, an increase of 85% relative to the pre-transition mean level of risk for all treated managers (1.86%/month, Table 1b), while Column 2 shows within-manager performance (i.e., the information ratio) increases by 20 basis points/month per unit of risk, or 19% of one standard deviation of pre-transition performance (1.03 basis points/month, Table 1b). Standard errors are clustered at the manager-fund-level.

Before transitioning to a hedge fund, treated managers take on more risk and perform poorly, while after they transition they take on even more risk and perform better, which is suggestive of a process where managers experiment on the job to learn about their fitness for alternative employment. But the treatment group is self-selected, and the nature of their job changes as they transition, leaving open the question of whether there is really a causal effect of experimentation on learning and transitions. Table 4 shows the results of our efforts to construct a valid counterfactual by matching “treated” observations to similar “control group” observations at the dyad-month-level. Column 1 shows that

before the matching procedure, all the key variables from the non-transitioning population are statistically different from the transitioning group, and the joint significance of the differences is very large ($F=39.43$). In other words, mutual fund managers who become hedge fund managers are very different from the average mutual fund manager along observable dimensions. Columns 2 and 4 show the differences in the means once a control group is identified using coarsened exact matching and propensity score matching, respectively.

The results of the propensity score matching approach are displayed in Column 3. Besides reconfirming the results on risk-taking and transitioning in Table 2 we also see that that rolling average excess returns are negatively correlated with mutual fund to hedge fund transitions, suggesting that experimentation is costly in the sense of harming performance.¹⁸

After matching, all of the key observable differences between the treatment and control groups are eliminated, except on the variable “missing AUM” in the CEM match, which is significant at the 10% level, and the joint significance of all the differences becomes small for both matching regimes ($F<1.3$). Moreover, beyond the statistical tests on the means, the CEM procedure ensures that treatment and control groups are matched exactly on a highly discretized set of strata. In other words, we have a control group that is essentially identical to the set of mutual fund managers who become hedge fund managers based on a number of crucial observable characteristics.

Within-manager analyses of the control-group observations offer suggestive evidence of how mutual fund managers who experiment and decide not to become hedge fund managers behave after they learn more information about themselves (e.g., their “fitness” for a different job). Tables 5 and 6 show the results of changes in risk and performance before and after the match date controlling for fund size and age, year, manager fixed effects, and year fixed effects. Standard errors are clustered at the manager level. Table 5, Column 1 shows that those in the base CEM match, who appear to experiment but do not

¹⁸ In this specification, comparing the instantaneous (monthly) probability of making a transition (0.054%) to the marginal effect of changing idiosyncratic risk gives us the economic interpretation of the coefficient estimates. Doubling idiosyncratic risk leads to a 15% increase in the probability of making a transition.

become hedge fund managers, reduce their idiosyncratic risk-taking by -0.24% per month, or 17% of one standard deviation. The point estimate is halved with the propensity score matched control group (Column 2), but both results are precisely estimated. In Column 3 we show changes in risk-taking using an alternative matched sample where we match on all the observable characteristics from the base match *except* ex ante risk-taking, and the results are very similar. Thus, it appears that mutual fund managers who experiment with “excessive” risk-taking, but choose to remain mutual fund managers, reduce their risk-taking once they learn their type and choose not to pursue a career as a hedge fund manager.

But is the ex-ante risk-taking excessive? In other words, is learning one’s type really an experiment—an act that is both costly and informative? If our characterization is correct, we should also see performance improve once a mutual fund manager learns her type and decides not to become a hedge fund manager. Table 6 shows within-manager test of changes in performance, as measured by the information ratio, which is computed as the average excess return from the 4-factor model divided by idiosyncratic risk-taking (before and after the match date).¹⁹ Performance increases by 4-5 basis points per month, per unit of systematic risk exposure in the second stage of each of the match specifications, (results are organized as in Table 5, with Column 3 showing the results after matching using the base match excluding performance), and all of the results are statistically significant, at the 5% level.

Addressing Alternative Hypotheses

At this point, we have found evidence consistent with Proposition 1. However, we would like to put the theory to more stringent tests wherever possible. Specifically, a number of alternative arguments may generate some of the results we find in our empirical analysis. In our theory, managers are taking more risk to learn about their fitness for a career in a new industry segment, which requires more learning than they would need if they were only planning on changing jobs within the mutual fund industry. Thus, managers are increasing risk before transitioning to a hedge fund to learn both about their fitness for a new job type. One logical implication of our theory, summarized by Corollary 1, therefore, is that

¹⁹ Because idiosyncratic risk is the denominator of the dependent variable we do not also include it as a control.

managers who make transitions to different kinds of jobs will take more risk before transitioning than managers who make transitions to similar kinds of jobs, because the latter group already has good information about their fitness for their current job. An alternative hypothesis about a correlation between risk-taking and risky types (i.e., those who like to transition) would posit that risk-taking behavior would be similar for transitions between firms within the same type of job as for transitions between firms and across different types of jobs. Because we observe both mutual fund to hedge fund and mutual fund to mutual fund transitions in the data, we can test this prediction directly.

A second alternative theory would be that we observe elevated risk-taking and declining performance prior to the transition as a result of involuntary exits rather than voluntary exits. Put another way, those mutual fund managers who perform poorly (i.e. generate poor risk-adjusted returns) are fired, and, faced with limited options, join or found hedge funds. To evaluate this possibility, we want to compare mutual-fund-to-hedge-fund transitions against all other mutual-fund exits to be sure that we are not picking up a more general phenomenon, whereby managers who leave their positions as portfolio managers take on a great deal of risk prior to leaving their posts (voluntarily or involuntarily).

Finally, the elevated pre-transition risk-taking we observe may also be consistent with a different “gambling” hypothesis. In this narrative, mutual fund managers who wish to become hedge fund managers, arguably attracted by the attractive economics of successful hedge funds, take on elevated levels of idiosyncratic risk in the hopes of fortuitously generating outsized returns. If successful, they become hedge fund managers. This alternative explanation is also observationally equivalent to a signaling story, where some mutual fund managers are essentially auditioning for hedge fund jobs, a role they can land by taking on risk and outperforming. Notably, if this type of selection is going on we would expect that the top performers, among those who experiment, would be the individuals who make the transition. If the learning is only “internal” to the manager, in the sense of one’s internal fitness for a different job (our argument), we would not expect those who transition to outperform those who experiment and choose to remain as mutual fund managers.

To test these alternative hypotheses, we conducted additional empirical tests. Table 7 shows the result of one test using a multinomial logit analysis in which the baseline case, coded as zero, is a manager who does not transition. A mutual fund to hedge fund transition is coded as a “1;” a mutual fund to mutual fund transition (a transition to a new mutual fund) is coded as a “2;” while exits from the mutual fund database by managers whose performance was in the bottom quartile of the distribution of the information ratio are coded as a “3.” Standard errors are clustered at the manager level and marginal effects are reported. Column 1 shows that ex-ante risk-taking is positive and significant for the mutual fund managers who become hedge fund managers. The magnitude of the effect, 0.5%, is very similar to the results reported in Table 2. Column 2 shows that risk-taking is indistinguishable from zero for mutual to mutual fund transitions, and, the point estimate, which is -0.1%, is statistically smaller than for those who transition from mutual funds to hedge funds (Column 4 reports the t-test on the difference between the two estimates). Column 3 tests whether those who leave are simply just poor performers who are “gambling for resurrection” and fail before eventually exiting the mutual fund industry and joining a hedge fund. The results reveal that those who experiment and transition to hedge funds behave quite differently from “typical” poor performers who leave the industry. Exits by bottom quartile performers are associated with an economically and statistically meaningful *reduction* in risk-taking. The point estimate in Column 3 is reliably statistically different from mutual fund managers who join hedge funds (Column 5).

The results in Table 7 provide two useful facts: (i) experimentation appears to be far more important for mutual fund to hedge fund transitions than for mutual fund to mutual fund transitions; and (ii) since poor-performing managers who subsequently leave the industry do not systematically take on high-levels of risk, it suggests that the transitioning managers we study are not merely representative of a broader “gambling for resurrection” effect.

While we already know that mutual fund managers who transition to hedge funds underperform in the year prior to leaving the mutual fund industry (Table 4, Column 3)—a result that is inconsistent

with “gambling and winning” or signaling—an even stronger test of these potential alternative explanations would be to compare changes in ex ante performance generated by mutual fund managers who subsequently transition to hedge funds against the changes in ex ante performance generated by mutual managers who experiment but do not become hedge fund managers. If managers who transition do not outperform managers who experiment and choose to stay at their mutual fund, then there can be little support for either the “gambling and winning” or signaling hypothesis—in order to plausibly “win” or send a meaningful signal, the treatment group would have to outperform the control group. Given the question at hand, the most natural matched sample for this analysis is precisely the sample generated by excluding ex ante performance (as in Table 6, Column 3).²⁰ We analyze the relationship between performance and treatment for a fixed window of time preceding the match date, and include manager and year fixed effects along with monthly event time dummies, which control for the average time path of returns for all managers during the experimental phase (i.e., trending downward), as in de Figueiredo and Rawley (2011). This approach delivers a meaningful estimate of the impact of “audition” performance on the likelihood of transitioning to alternative employment.

If mutual fund managers who are seeking to enter the hedge fund industry are simply “rolling the dice” with their mutual fund in the hopes of fortuitously signaling their quality to potential investors/partners, or creating a strong track record, we would expect that those who transition would experience a relative improvement in performance compared to those in the matched sample who do not transition. In contrast, the differences-in-difference results of (rolling average) performance on the treatment in Table 8 show that ex ante performance is actually no different for managers who experiment and transition to hedge funds, compared to those who experiment and remain mutual fund managers. Thus, the test shows that empirically it is impossible for an external observer to distinguish between managers who experiment and transition and those who experiment and do not transition, either on the actions they take (e.g., risk-taking) or the outcomes they achieve (e.g., performance). Since there must be

²⁰ We use the “ex-performance” matched sample for this test because the base match controls for ex ante returns directly, and thereby mechanically eliminates any variation in ex ante returns.

information in a signal for it to be a useful, and there appears to be no meaningful information transmitted in this context, the results are inconsistent with a signaling hypothesis.

We include one final robustness check in Table 8. If our theory is correct and managers experiment by taking on excessive risk before transitioning, we should observe a decrease in risk-taking within-fund after a transitioning manager's departure. In our main analysis we have shown that risk-taking is elevated in transitioning managers, relative to other mutual fund managers, relative to other managers at the same firm, and to other managers of the same fund (Table 2), and our risk-adjustment process controls for variation in investment style between funds. In Column (2) of Table 8 we also provide direct evidence that funds shift risk-taking downward following the departure of a manager for a hedge fund: a fund sheds 16 basis points of risk per month following the transition compared to the period during which the transitioning manager was at the helm.

In summary, there are several alternative explanations that could be consistent with some of our empirical findings. Our argument is that only our theoretical explanation about experimentation on the job can be reconciled with the full complement of the results.

7. Discussion

Contributions

We advance a novel mechanism by which prior employment facilitates employee mobility—learning on the job through experimentation. Our empirical findings from the asset management industry are consistent with the theoretical argument that some employees experiment on the job to learn about their fitness for an alternative career and, based on the results of these experiments, decide whether or not to leave their employer. The extant literature on employee mobility has offered several reasons why individuals may stay at or leave their employers. These drivers typically encompass financial incentives (Zenger, 1992) and non-pecuniary reasons related to culture (Sheridan, 1992) and disagreements (Klepper & Thompson, 2010). We take a different, but complementary, approach to explain mobility whereby the employee can choose to learn about herself on the job to inform the transition decision. We believe that this explanation can shed light on patterns of mobility across industries, within firms and also inform

future studies on employee recruitment and retention. Importantly, while much of the prior literature has focused on the motivation to transition to alternative employment (“why”), our work sheds light on the process individuals use to facilitate these transitions (“how”). This piece of the puzzle represents a promising new avenue for future research on employee mobility and other related phenomenon.

Just as prior papers have considered the distinction between moving to a competitor firm as opposed to launching a startup (Campbell et al., 2012; Carnahan et al., 2012; Ganco, 2013), our work can also contribute to the growing literature on employee entrepreneurship. The work on entrepreneurial spawning has found that employees acquire valuable knowledge and/or run into bureaucratic constraints that operate as “pull/push” factors into entrepreneurship (e.g. Agarwal, Echambadi, Franco, & Sarkar, 2004; Bhide, 2000; Chatterji, 2009; Freeman, 1986; Gompers, Lerner, & Scharfstein, 2005; Hellmann, 2007; Klepper & Sleeper, 2005). Our key point is that employment also provides an opportunity to experiment, such that employees can learn about their skills and preferences for a new position, with implications for entrepreneurship. Our proposed mechanism can explain why some employees remain at the incumbent firm while others spawn into the same or a different industry; though it does not preclude other explanations for how prior employment influences patterns of entrepreneurship. Our characterization of experimentation and job matching as being critical to understanding entrepreneurship are themes that can be traced back to seminal work on this topic, such as Jovanovic (1979). As described above, the “how” aspect of employee transitions into entrepreneurship specifically is an important priority for future work. Our view is that this line of inquiry can inform the choices and activities of incipient entrepreneurs as they “prepare to found” new businesses.

Further, our study supports the generalizability of entrepreneurial spawning beyond high technology, in the spirit of Kacperczyk’s (2012, 2013) studies of the asset management industry, which focus on social networks and intrapreneurship inside existing firms. Finally, this paper deepens our understanding of how individual differences explain variation in entrepreneurial activity among employees at the same organization, a key objective of the broader research agenda on employee

entrepreneurship, especially as it relates to considering “trait-based” versus “contextual” drivers of entrepreneurial entry (Ganco, 2013).

Our argument has some similarities with anecdotal accounts about employees considering alternative employment that are important to clarify. It is common for incipient entrepreneurs to “moonlight”, working on a business idea nights and weekends while maintaining their current salaried position. In a similar spirit, some firms, most notably Google, provide employees a fixed amount of time during the workweek to work on their own projects. Further, casual observation also suggests that employees may not be equally productive at all times, and might “slack off” during certain periods. Our argument is distinct from each of these examples, but does characterize features of Google’s policy in the sense that experimentation in our context is directly associated with reduced productivity at work but also builds future option value for the employee. It appears that Google and other firms have recognized a tendency toward on the job learning through experimentation among skilled employees and has developed a policy to direct and capitalize upon such behavior (e.g. by marketing the policy as a perk as part of its larger human resources strategy).

Limitations

A conceptual limitation of our theory is that we do not offer a general equilibrium model of learning through experimentation. Indeed, our model does not feature contract design by principals to discourage experimentation, though we show that such experimentation is costly to the firm and its investors. However, given the relatively modest number of mutual fund managers who experiment in practice, the expected cost to principals is likely quite low. Indeed, the low average cost of experimentation under the standard contract (which is, presumably, optimal for the vast majority of managers) across all managers supports the notion that this assumption is not a serious limitation of our empirical study. However, scholars who wish to study learning on the job as a precursor to mobility events in other contexts should be mindful of this assumption.

Finally, using risk-taking as a proxy for experimentation rather than observing the experimentation process directly represents another limitation of this research. For example, our inability to observe the exact experimentation dates makes it more difficult to parse the results from employee experimentation more finely to better understand the various components of learning. While we exploit a relatively comprehensive industry dataset and a matched sample counterfactual to generate evidence that is consistent with our theoretical model, future work may consider more micro-level evidence to carefully document exactly how employees experiment prior to transitioning. Researchers may have to tradeoff high-quality panel data that allows for counterfactual comparisons against having much more detailed information on the intentions and activities of employees.

Based on discussions with industry participants, we believe experimentation is primarily related to short-selling strategies, adding leverage, and trading a more concentrated portfolio. Along these dimensions and many others, operating a mutual fund differs considerably from operating a hedge fund, and it seems plausible that experimenting with these approaches might inform a transition decision by providing insights into individual skills and preferences. Differences in investment strategies between the two industries also suggest that success in mutual funds and hedge funds would not necessarily be positively correlated, and underscores the importance of learning through experimentation in this context. Of course, for our argument to apply, experimentation need not even be intentional unless one chose to make the strong assumption that learning cannot occur from “accidental” experiments.

8. Conclusion

Our work casts employment as a laboratory where individuals can experiment and learn about alternative jobs. Through experimentation, employees learn how their skills and preferences match to new opportunities. Adopting this lens, we make precise predictions about the drivers of employee mobility and find evidence in support of our arguments. In doing so, we hope to open up new avenues of research in explaining the process behind employee transitions more broadly.

Appendix: Proofs of propositions

Proof of Proposition 1. Part (a) follows from two facts. First, in the absence of experimentation all mutual fund managers will choose v^{max} . Second, only those who experiment can switch to being a hedge fund manager. Thus, all switchers are experimenters. Since all experimenters choose v^{exp} , and $r^{exp} < r^{max}$, then all those who switch have higher risk and lower returns on average in the first period. Moreover the same will hold for the identical reason for all experimenters. Part (b) follows from the fact that the non-experimenting mutual fund manager's problem is trivially solved by choosing a risk of v^{max} . Part (c) follows from the fact that $\alpha_L < 1 < \alpha_H$. Because only high types switched as shown in Proposition 2(b), then the result is obtained by noting the solution to the hedge fund manager's problem is $\alpha_H v$.

Proof of Corollary 1. That transitions are sequentially rational requires part (b). That experimentation is dominated follows from equation (2).

References

- Agarwal, R., Echambadi, R., Franco, A. M., & Sarkar, M. (2004). Knowledge Transfer Through Inheritance: Spin-out Generation, Development, and Survival. *Academy of Management Journal*, 47(4), 501-522.
- Almeida, P., & Kogut, B. (1999). Localization of knowledge and the mobility of engineers in regional networks. *Management Science*, 45(7), 905-917.
- Arrow, K. (1962). Economic welfare and the allocation of resources for invention *The rate and direction of inventive activity: Economic and social factors* (pp. 609-626): Nber.
- Bhide, A. V. (2000). *The origin and evolution of new businesses*: Oxford University Press, USA.
- Campbell, B. A., Coff, R., & Kryscynski, D. (2012). Rethinking sustained competitive advantage from human capital. *Academy of Management Review*, 37(3), 376-395.
- Campbell, B. A., Ganco, M., Franco, A. M., & Agarwal, R. (2012). Who leaves, where to, and why worry? Employee mobility, entrepreneurship and effects on source firm performance. *Strategic Management Journal*, 33(1), 65-87.
- Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, 52(1), 57-82.
- Carnahan, S., Agarwal, R., & Campbell, B. A. (2012). Heterogeneity in turnover: The effect of relative compensation dispersion of firms on the mobility and entrepreneurship of extreme performers. *Strategic Management Journal*, 33(12), 1411-1430.
- Chatterji, A. K. (2009). Spawned with a silver spoon? Entrepreneurial performance and innovation in the medical device industry. *Strategic Management Journal*, 30(2), 185-206.
- Cotton, J. L., & Tuttle, J. M. (1986). Employee turnover: A meta-analysis and review with implications for research. *Academy of Management Review*, 11(1), 55-70.
- David, G., Rawley, E., & Polsky, D. (2013). Integration and Task Allocation: Evidence from Patient Care. *Journal of Economics & Management Strategy*, 22(3), 617-639.

- de Figueiredo, R. J., Meyer-Doyle, P., & Rawley, E. (2013). Inherited Agglomeration Effects in Hedge Fund Spawns. *Strategic Management Journal*, 34(7), 843-862.
- de Figueiredo, R. J., & Rawley, E. (2011). Skill, Luck, and the Multiproduct Firm: Evidence from Hedge Funds. *Management Science*, 57(11), 1963-1978.
- de Figueiredo, R.J., Rawley, E., & Shelef, O.(2014) Bad Bets: Excessive Risk Taking, Convex Incentives, and Performance. No. 13-002. Stanford Institute for Economic Policy Research.
- Elfenbein, D. W., Hamilton, B. H., & Zenger, T. R. (2010). The small firm effect and the entrepreneurial spawning of scientists and engineers. *Management Science*, 56(4), 659-681.
- Fama, E. F., & French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1), 55-84.
- Freeman, J. (1986). Entrepreneurs as organizational products: Semiconductor firms and venture capital firms. *Advances in the study of entrepreneurship, innovation, and economic growth*, 1(33-52).
- Fung, W., & Hsieh, D. A. (1999). A primer on hedge funds. *Journal of Empirical Finance*, 6(3), 309-331. doi: [http://dx.doi.org/10.1016/S0927-5398\(99\)00006-7](http://dx.doi.org/10.1016/S0927-5398(99)00006-7)
- Fung, W., & Hsieh, D.A. (2003). The risk in hedge fund strategies, theory and evidence from trend followers. *Review of Financial Studies* 14(2), 313-341
- Ganco, M. (2013). Cutting the Gordian knot: The effect of knowledge complexity on employee mobility and entrepreneurship. *Strategic Management Journal*, 34(6), 666-686.
- Gompers, P., Lerner, J., & Scharfstein, D. (2005). Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. *The Journal of Finance*, 60(2), 577-614.
- Hellmann, T. (2007). When do employees become entrepreneurs? *Management Science*, 53(6), 919-933.
- Iacus, S. M., King, G., & Porro, G. (2011) "Causal inference without balance checking: Coarsened exact matching." *Political analysis* : mpr013.
- Kacperczyk, A. J. (2013). Social Influence and Entrepreneurship: The Effect of University Peers on Entrepreneurial Entry. *Organization Science*, Forthcoming.
- Kacperczyk, A. J. (2012). Opportunity Structures in Established Firms Entrepreneurship versus Intrapreneurship in Mutual Funds. *Administrative Science Quarterly*, 57(3), 484-521.
- Klepper, S. (2007). Disagreements, spinoffs, and the evolution of Detroit as the capital of the US automobile industry. *Management Science*, 53(4), 616-631.
- Klepper, S., & Sleeper, S. (2005). Entry by spinoffs. *Management Science*, 51(8), 1291-1306.
- Klepper, S., & Thompson, P. (2010). Disagreements and intra-industry spinoffs. *International Journal of Industrial Organization*, 28(5), 526-538.
- Laffont, J., & Martimort, D. (2002). *The Theory of Incentives: The Principal Agent Model*: Princeton University Press.
- Marx, M., Strumsky, D., & Fleming, L. (2009). Mobility, skills, and the Michigan non-compete experiment. *Management Science*, 55(6), 875-889.
- Palomino, F., & Prat, A. (2003). Risk taking and optimal contracts for money managers. *RAND Journal of Economics*, 113-137.
- Pastor, L., R.F. Stambaugh. (2003). Liquidity risk and expected stock returns. *Journal of Political Economy* 111(3), 642-685.
- Rosenbaum, Paul R., and Donald B. Rubin.(1983) "The central role of the propensity score in observational studies for causal effects." *Biometrika* 70.1: 41-55.
- Rosenkopf, L., & Almeida, P. (2003). Overcoming Local Search Through Alliances and Mobility. *Management Science*, 49(6), 751-766.
- Sheridan, J. E. (1992). Organizational culture and employee retention. *Academy of Management Journal*, 35(5), 1036-1056.

- Singh, J., & Agrawal, A. (2011). Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science*, 57(1), 129-150.
- Sørensen, J. B. (2007). Bureaucracy and entrepreneurship: Workplace effects on entrepreneurial entry. *Administrative Science Quarterly*, 52(3), 387-412.
- Stephan, P. E. (1996). The economics of science. *Journal of Economic Literature*, 1199-1235.
- Williams, C. R., & Livingstone, L. P. (1994). Another look at the relationship between performance and voluntary turnover. *Academy of Management Journal*, 37(2), 269-298.
- Zenger, T. R. (1992). Why do employers only reward extreme performance? Examining the relationships among performance, pay, and turnover. *Administrative Science Quarterly*, 198-219.

Table 1a: Key descriptive statistics for mutual fund managers (one observation per manager)

	Full sample 12,754 mutual fund managers from 1,272 firms				MF->HF transitions only 486 mutual fund managers from 252 firms			
	<u>Mean</u>	<u>Stdv</u>	<u>Min</u>	<u>Max</u>	<u>Mean</u>	<u>Stdv</u>	<u>Min</u>	<u>Max</u>
MF->HF transition	0.04	0.19	0	1	1	0	0	1
Idiosyncratic risk (% monthly)	1.64	1.35	0.02	6.80	2.04	1.50	0.02	6.80
Alpha (4 factor)	-0.06	0.50	-1.54	1.25	-0.02	0.53	-1.54	1.25
Information ratio	-0.12	0.68	-9.11	5.19	-0.06	0.37	-2.62	1.17
Start date	2/01	64 mos.	12/78	1/11	9/98	45 mos.	12/88	12/10
Tenure (months)	45	37	12	229	54	45	12	227
Assets under management (\$M)	967	3,234	1	74,403	724	1,670	1	18,302
Firm transitions (total count)	9.08	15.96	0	98	11.85	19.29	0	98
NYC HQ dummy	0.22	0.42	0	1	0.22	0.41	0	1

Table 1b: Key descriptive statistics for MF->HF transitions (unit of analysis is the manager-fund-month)

	Prior to transition to HF				After transition to HF			
	<u>Mean</u>	<u>Stdv</u>	<u>Min</u>	<u>Max</u>	<u>Mean</u>	<u>Stdv</u>	<u>Min</u>	<u>Max</u>
Idiosyncratic risk (% monthly)	1.86	1.52	0.00	21.93	3.19	2.82	0.06	42.60
Alpha (% monthly)	-0.00	2.51	-30.01	52.85	0.27	4.49	-49.29	95.79
Information ratio	-0.04	1.03	-7.86	7.92	0.11	1.00	-8.44	8.51
Start date	2/98	38 mos.	1/92	1/09	12/97	41 mos.	1/92	2/10
Tenure (months)	75	49	11	228	120	54	22	204
Assets under management (\$M)	562	1,651	1	24,139	378	5,961	0.04	225,141
Missing AUM	0.03	0.18	0	1	0.11	0.31	0	1

There are 33,942 mutual fund manager-fund-months prior to the MF->HF transition and 138,715 hedge fund manager-fund-months after the MF->HF transition reported in this table. The 486 MF->HF transitioners managed 1,790 mutual funds and 689 hedge funds.

Table 2: Ex ante risk-taking in mutual fund to hedge fund managers*Dependent variable: MF->HF transition*

<i>Unit of analysis:</i>	Manager	Manager	Manager-fund	Manager-fund-month
<i>Model:</i>	<u>Probit</u>	<u>OLS firm f.e.</u>	<u>OLS fund f.e.</u>	<u>OLS fund f.e.</u>
	(1)	(2)	(3)	(4)
<i>Idiosyncratic risk</i>	<i>0.59*</i> <i>(0.10)</i>	<i>0.64*</i> <i>(0.17)</i>	<i>0.64*</i> <i>(0.24)</i>	<i>1.07*</i> <i>(0.23)</i>
Excess returns	0.53 ⁺ (0.31)	0.63 (0.41)	-0.02 (0.24)	0.01 (0.01)
Log AUM	-0.06 (0.09)	0.23* (0.11)	-0.75* (0.16)	-0.15* (0.06)
Missing AUM	-2.21* (0.81)	-1.91 (1.54)	0.41 (1.42)	0.72* (0.33)
Manager log tenure	0.30 (0.23)	0.16 (0.31)	0.97* (0.19)	0.17* (0.08)
Firm transitions	0.04* (0.01)			
NYC dummy	-0.11 (0.39)			
Firm fixed effects	N	Y	N	N
Fund fixed effects	N	N	Y	Y
Year fixed effects	N	N	N	Y
Constant	Y	Y	Y	Y
N	12,754	12,754	57,180	2,448,405
Adjusted R ²	0.04	0.04	0.13	0.55

* Significant at the 5% level, ⁺ Significant at the 10% level.

This table shows how the behavior of mutual fund managers who become hedge fund managers differs from that of other mutual fund managers at the manager level, manager-fund, and the manager-fund-month level. MF->HF transition is a dichotomous variable that is equal to one when a mutual fund manager is/becomes a hedge fund manager and zero otherwise. Risk and performance are measured net of the four standard “common factors.” There are 1,256 firms and 20,791 funds in the analysis. All coefficient estimates and standard errors are multiplied by 100 for presentation purposes. Standard errors are clustered at the manager, manager-fund or manager-fund level, based on the unit of analysis.

Table 3: Changes in risk-taking and performance for mutual fund managers who become hedge fund managers

<i>Dependent variable:</i>	<u>Risk</u>	<u>Inf. ratio</u>
	(1)	(2)
<i>MF->HF transition</i>	<i>1.58*</i> <i>(0.12)</i>	<i>0.20*</i> <i>(0.01)</i>
Log manager tenure	0.28* (0.09)	0.07* (0.02)
Assets under management	-0.10* (0.02)	-0.02* (0.00)
Missing AUM	0.63* (0.17)	0.06* (0.02)
Constant	Y	Y
Manager-fund fixed effects	Y	Y
Year fixed effects	Y	Y
N	172,657	172,657
Adjusted R ²	0.63	0.03

* Significant at the 5% level, + Significant at the 10% level.

The unit of analysis is the fund-month.

This table shows how risk and performance of mutual fund managers who become hedge fund managers changes as they transition to running a hedge fund. MF->HF transition is a dichotomous variable that is equal to one when a mutual fund manager becomes a hedge fund manager and zero otherwise. Risk and performance are measured net of the four standard “common factors.” Standard errors are clustered at the manager-fund level.

Table 4: Matching managers who make mutual fund to hedge fund transitions to non-transitioning mutual fund managers

	t-tests before <u>matching</u>	t-tests after CEM <u>matching</u>	Probit for p-score <u>matching</u>	t-tests after p- <u>score matching</u>
	(1)	(2)	(3)	(4)
<i>Idiosyncratic risk (rolling average)</i>	5.46*	1.26⁺	0.008* (0.000)	-0.50
Excess return (rolling average)	-4.28*	-0.83	-0.016* (0.003)	-0.03
Log age	-1.92 ⁺	0.07	-0.011* (0.001)	0.73
Log AUM	-5.49*	0.04	0.056* (0.020)	0.94
Missing AUM	1.75 ⁺	1.70 ⁺	0.017* (0.003)	-0.20
Firm transitions	9.67*	1.49	0.001* (0.000)	0.15
NYC dummy	-1.31	-0.84	-0.011 ⁺ (0.006)	0.59
Date	-13.40*	-1.07	Year f.e.	0.88
N	906,302	2,124	906,302	2,282
Pseudo-R ²			0.08	
F-test on joint signif.	39.43*	1.29		0.31

* Significant at the 5% level, ⁺ Significant at the 10% level.

The unit of analysis is the manager-fund-month. t-statistics are reported on the difference in means between the “treated” (i.e., MF->HF transitions) population and the “control” group population before and after matching. 1,141 treated and control dyads matched using propensity score matching, and 1,062 treated and control dyads matched using Coarsened Exact Matching (CEM). Manager-fund dyads cannot be matched in their first year of operation (because risk and return rolling averages are not defined), and control group observations cannot be matched in their last year of operation (because there would not be enough “second stage” observations). Matching treated managers to control group managers who managed the same fund is also disallowed. Marginal effects are reported in Column (3). Match quality was improved by including a risk x age interaction and a risk² term in the propensity score matching algorithm, but the “second stage” results were unchanged. In the CEM specification risk, excess return, date and firm spawns were coarsened at the decile level; age and AUM were coarsened into 25 buckets, and the dichotomous variables (missing AUM and NYC dummy) entered as zeroes and ones, resulting in 794 matched strata out of 18,706 possible strata.

Table 5: Matched sample behavior – changes in risk-taking

Dependent variable: idiosyncratic risk

<i>Match type:</i>	<u>CEM</u>	<u>p-score</u>	<u>ex-risk</u>
	(1)	(2)	(3)
<i>Ex post indicator</i>	-0.24* (0.03)	-0.11* (0.03)	-0.21* (0.03)
Excess return	0.01* (0.00)	0.01* (0.00)	0.00 (0.00)
Log AUM	0.03+ (0.02)	0.02 (0.01)	0.03+ (0.01)
Missing AUM	-0.17+ (0.10)		-0.05 (0.11)
Log age	-0.04+ (0.02)	-0.02 (0.02)	-0.01 (0.02)
Year fixed effects	Y	Y	Y
Manager fixed effects	Y	Y	Y
Constant	Y	Y	Y
N	84,435	104,411	84,737
Adjusted R ²	0.87	0.89	0.90

* Significant at the 5% level, + Significant at the 10% level.

This table shows how risk-taking changes in the matched sample of non-transitioning mutual fund managers who are observationally equivalent to mutual fund managers who become hedge fund managers. *Ex-post indicator* is equal to 1 for all dates after the “match date,” determined by matching transitioning managers to non-transitioning managers, based on their observable characteristics, using Coarsened Exact Matching (CEM), or propensity score matching (pscore), as in Table 4, and 0 otherwise. In the “CEM” and “p-score” matching regime *ex ante* risk-taking (12-month rolling average standard deviation of returns net of common risk factors) enters as a variable in the matching algorithm (see Table 4). In the “ex-risk” matching regime *ex ante* risk does not enter the (CEM) matching algorithm. Standard errors are clustered at the manager-fund level.

Table 6: Matched sample behavior – changes in performance*Dependent variable: information ratio*

<i>Match type</i>	<u>CEM</u>	<u>p-score</u>	<u>ex-performance</u>
	(1)	(2)	(3)
<i>Ex post indicator</i>	0.05* (0.02)	0.04* (0.02)	0.05* (0.03)
Log AUM	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)
Missing AUM	-0.01 (0.03)		0.01 (0.04)
Log age	0.01 (0.01)	0.00 (0.01)	0.01 (0.01)
Year fixed effects	Y	Y	Y
Manager fixed effects	Y	Y	Y
Constant	Y	Y	Y
N	84,435	104,411	85,832
Adjusted R ²	0.76	0.79	0.77

* Significant at the 5% level, † Significant at the 10% level.

This table shows how risk-taking changes in the matched sample of non-transitioning mutual fund managers who are observationally equivalent to mutual fund managers who become hedge fund managers. *Ex-post indicator* is equal to 1 for all dates after the “match date,” determined by matching transitioning managers to non-transitioning managers, based on their observable characteristics, using Coarsened Exact Matching (CEM), or propensity score matching (pscore), as in Table 4, and 0 otherwise. In the “CEM” and “p-score” matching regime a 12-month rolling average of excess returns enters as a variable in the matching algorithm (see Table 4). In the “ex-performance” matching regime performance does not enter the (CEM) matching algorithm. Standard errors are clustered at the manager-fund level.

Table 7: Risk-taking and transitions—comparing risk-taking amongst transitioning mutual fund managers, and poor performing mutual fund managers who subsequently exit

Multinomial logit analysis

<i>Dependent variable:</i>	<u>MF->HF transitions</u>	<u>MF->MF transitions</u>	<u>Other exit</u>	t-stat on difference <u>(1) – (2)</u>	t-stat on difference <u>(1) – (3)</u>
	(1)	(2)	(3)	(4)	(5)
<i>Idiosyncratic risk</i>	<i>0.005*</i> <i>(0.001)</i>	<i>-0.001</i> <i>(0.001)</i>	<i>-0.006*</i> <i>(0.001)</i>	<i>3.68*</i>	<i>7.87*</i>
Controls	Yes	Yes	Yes		
N		12,754			
Pseudo R ²		0.24			

* Significant at the 5% level, † Significant at the 10% level.

This table compares the *ex ante* risk taking of different types of mutual fund managers who leave the industry: mutual fund managers who become hedge fund managers, mutual fund managers who become mutual fund managers of new mutual funds, and mutual fund managers who leave the industry after performing poorly (i.e., were in the bottom quartile of the distribution of the information ratio over all mutual fund managers). Categories are mutually exclusive. The baseline observation is “no transition.” The unit of analysis is the manager. Standard errors are clustered at the manager level. Marginal effects are reported. Controls include: 4-factor alpha, log AUM, Missing AUM, start date, log tenure, firm transitions, and NYC dummy (as in Table 2). The specification passed all the standard IIA tests.

Table 8: Robustness checks on mobility

	Change in performance of MF->HF compared to the matched sample	Change in fund-level risk taking after MF->HF period
	(1)	(2)
<i>Dependent variable</i>	<i>Information ratio</i>	<i>Idiosyncratic risk</i>
<i>MF->HF transition</i>	<i>0.01</i> <i>(0.02)</i>	<i>-0.16*</i> <i>(0.02)</i>
Manager fixed effects	Y	N
Fund fixed effects	N	Y
Year fixed effects	Y	Y
Event time fixed effects	Y	N
Other controls	Y	Y
N	39,245	243,858
Adjusted R ²	0.81	0.86

* Significant at the 5% level, + Significant at the 10% level.

The unit of analysis is the manager-fund-month. Column (1) is a differences-in-differences (manager fixed effects) analysis of changes performance for the treatment and control group during a thirty-six month window leading up to the focal experiment date after being matched using the “ex-performance” (excluding performance measures) matching specification (see Table 6). Column (2) is a within-fund fixed effects analysis of changes in risk-taking during the period that the transitioning MF->HF manager managed the mutual fund. Other controls include assets log under management, log tenure and excess returns. Standard errors are clustered at the manager-fund-level in Column (1) and at the fund-level in Column (2).