Learning on the Job? Entrepreneurial Spawning 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

September 2013

Entrepreneurs often have prior experience at incumbent firms. We present a new mechanism by which prior employment can influence transitions into entrepreneurship. 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 spawn into related industry segments as entrepreneurs or employees. Others will remain at the incumbent firm or pursue entrepreneurship in the same industry segment. 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 individuals who engage in excessive risk-taking at mutual funds are more likely to transition into hedge funds. Taken together, our findings suggest that learning on the job through experimentation is an important mechanism for enabling entrepreneurial spawning.

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

* We thank Michael Ewens, Marina Halac, Bryan Hong, Aleksandra Kacperczyk, Dylan Minor, Damon Phillips, Tim Pollock, Andrea Prat, Christopher Rider, David Robinson, Scott Rockart, Sampsa Samila, Minyuan Zhao, seminar participants at Singapore National University, and conference participants at the 2013 Smith Entrepreneurship Research Conference, the 2013 Columbia Business School Strategy Cross Disciplinary Area Retreat, and the 2013 Strategy Research Forum for helpful comments and suggestions. Pascal Anders and Andrea Contigiani provided excellent research assistance. All mistakes are our own.

1. Introduction

Although the most celebrated examples of entrepreneurship often involve “2 guys in a garage or a dorm room,” (Audia & Rider, 2005) academic research indicates entrepreneurs typically have significant industry experience (e.g. Agarwal, Echambadi, Franco, & Sarkar, 2004; Bhide, 2000; Chatterji, 2009; Freeman, 1986; Gompers, Lerner, & Scharfstein, 2005; Hellmann, 2007; Klepper & Sleeper, 2005). An emerging line of research has sought to uncover how prior employment shapes entrepreneurial spawning, or transitions into entrepreneurship. Some influential papers have proposed that spawning can be explained by disagreements or breakdowns in bargaining between key employees and management at the incumbent firm (Anton & Yao, 1995; Klepper, 2007). Other work suggests particular corporate cultures may either provide opportunities for employees to learn about entrepreneurship and acquire valuable social networks, or alternatively pose bureaucratic hurdles that induce transitions into entrepreneurship (Gompers et al., 2005; Sørensen, 2007).

We build on this literature to propose another candidate mechanism by which prior employment can impact entrepreneurial transitions in the same or a related industry segment—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 entrepreneurship within the same or a related industry segment. We develop a formal model to demonstrate that based on the results of learning on the job through experimentation, some employees will transition to new positions in the same or a related industry while others will remain at the existing firm.

Using data from the asset management industry, we find evidence supporting the idea that learning on the job through experimentation facilitates entrepreneurship. 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. 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 13% increase in the transition rate.[1] We are able to refute alternative explanations that could be consistent with similar empirical patterns, such as mean reversion or the argument that poorly performing mutual fund managers are simply "gambling for resurrection" as opposed to learning from experimentation.

Importantly, one of the key advantages 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 17 basis points per month, a 10% reduction in risk-taking relative to the mean level, and improve their performance by 6 basis points per month per unit of risk, a pattern strikingly consistent with our proposed learning mechanism.

This research makes several contributions to the literature on entrepreneurship. First, we propose and find support for a new mechanism by which prior employment influences entrepreneurial spawning. Notably, our explanation can account for why some employees remain at the firm while others leave, a key gap in the prior literature. Relatedly, we also make an empirical contribution to the extant literature. One of the weaknesses of the literature on entrepreneurial spawning has been the difficulty of controlling for the heterogeneous unobserved quality of employee entrepreneurs, especially compared to the employees who remain at the firm. We have pre-entrepreneurship performance data on all individuals, which allows us to deal directly with unobservable quality. We also use firm fixed effects to account for "within-firm" differences between employees that stay and those who spawn into new ventures.

[1] The base transition rate in our sample is 3.8% (486 transitions/12,754 mutual fund managers). From Table 2, column 3, doubling risk exposure, from a mean of 1.64%/month (Table 1a), increases the transition rate by 0.62%. Therefore, a one-standard-deviation increase in risk exposure, 1.35%/month (Table 1a), leads to a 13% increase in the transition rate (0.62/3.8*1.35/1.64).

In addition, our findings relate to transitions into a related industry (from mutual funds to hedge funds), which distinguishes this paper from most others in the research tradition that has explored intra-industry entrepreneurship. We also use the distinction between intra-industry transitions and inter-industry entrepreneurship to empirically discern between (1) the effect of learning about one's fitness for an alternative position while on the job and (2) one's risk preferences for entrepreneurship. Since Frank Knight's (1921) seminal work, scholars have theorized about a connection between entrepreneurship and tolerance for risk and uncertainty. Though it is often implicitly assumed, most previous studies have not had the data to test the basic prediction that employees who are more risk-loving are more likely to transition into entrepreneurship. We find support for a connection between risk appetite and employee transitions into entrepreneurship. But perhaps more importantly for the theory we develop, we find that mutual fund managers who transition to hedge funds take on more risk ex ante compared to mutual fund managers who transition to mutual fund entrepreneurship, which suggests inter-industry entrepreneurship warrants special attention.

In the following section, we recount the key implications from prior work on employee entrepreneurship and describe our theory. Next, we present a model of learning on the job that provides general insights into entrepreneurial transitions. Third, we describe the salient features of the asset management industry and apply our general model directly to this empirical context, generating testable predictions. We then outline the empirical methods and analyses, summarize the key results, and discuss implications from our work.

2. Prior Work on Employee Entrepreneurship

There is a growing literature on employee entrepreneurs and their ventures, alternatively referred to as spinouts, spinoffs, and spawns. Most of this work focuses on how prior employment shapes entrepreneurial entry by employees into the same industry segment. For our purposes, the literature can be divided into two categories of mechanisms: (1) technical and market knowledge transfer and (2) bureaucratic constraints. In the spirit of Amit and Muller (1995), we view these two mechanisms as "pull" and "push" forces of entrepreneurial spawning.

The knowledge argument (Agarwal et al., 2004; Chatterji, 2009; R. J. de Figueiredo, Meyer-Doyle, & Rawley, 2013; Gompers et al., 2005), which can be classified as a "pull" mechanism, proposes that employment provides entrepreneurially inclined employees the opportunity to acquire technical and market knowledge they can bring with them to a new venture in the same industry. Although work in this tradition differs in terms of the kind of knowledge that is acquired, the assumption is generally that the knowledge is most appropriate for entrepreneurship in the same industry. Explaining why some employees have access to this knowledge while others do not, or why some employees are able to exploit this knowledge through entrepreneurship more effectively, has been difficult. One potential answer is provided by Kacperczyk (2013), who argues social networks may be critical in influencing employees to transition into entrepreneurship in the hedge fund industry. In particular, depending on the composition of university peer networks, access to entrepreneurial opportunities can be transferred across individuals and influence transition rates.

The bureaucratic constraints argument (Anton & Yao, 1995; Klepper, 2007; Klepper & Sleeper, 2005; Klepper & Thompson, 2010; Sørensen, 2007), which can be thought of as a "push" mechanism, proposes that numerous layers of decision making at the parent firm generally frustrate employees, or that employees have disagreements with top management over the value of a particular idea. These challenges induce the employee to leave the firm to found a new venture, presumably often to commercialize the idea at issue. This idea might fall outside the core business of the parent firm or compete with one of its existing products, but typically these employees also spinout into the same industry (though perhaps in sub-sectors adjacent to the parent firm).

These two general explanations have been successful in explaining the phenomenon of employee entrepreneurship, but do not provide much insight into why some employees leave and others stay at the incumbent firm. We propose a third theoretical mechanism to explain patterns of entrepreneurial spawning: learning about oneself on the job through experimentation. Our theory does not exclude the knowledge or bureaucratic constraints arguments, but instead provides a complementary perspective to account for broader patterns in employee mobility and entrepreneurship.

Employment as a Laboratory for Learning about Alternatives

Whereas the prior work on employee entrepreneurship allows for employees to learn about the industry and entrepreneurial opportunities generally (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 an alternative career. In practice, employment provides several opportunities for gaining insights into one's suitability for alternative jobs in the same or related industry. 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 spawn.

"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 skills and (ii) the employee's preferences. Our theory seeks to explain intra-industry entrepreneurship as well as transitions to other industry segments through learning on the job. We assume that because the employee is currently employed in the focal industry, she will have relatively more information about her fitness there than in a related industry segment, and we can test this assumption in our empirical analysis. A key distinction between our argument and the knowledge-transfer work discussed above is that we are proposing that employees can also learn about their own fitness for another job through working at an incumbent firm, as opposed to acquiring market and technical knowledge about the industry in which they are working.

Our argument is also related to the literature that casts jobs as an "experience good" (Jovanovic, 1979:973; Nelson, 1970). For example, Johnson (1978) explores "job shopping" whereby a given worker's suitability for a job can only be ascertained after some work experience in the job. This body of work contrasts with models of job search, whereby individuals can learn the characteristics of jobs by searching and do not require direct experience. In Johnson's model, the worker has both general ability and job-specific ability, about which employment provides the worker with some insight. Although the literature cited above does not explicitly consider learning as a costly experiment, or entrepreneurial spawning, it does promote the idea of learning one's type through experience. And indeed, consistent with the literature on jobs as experience goods, we find that those individuals who experiment and transition to hedge funds appear to receive a stronger "performance signal"[2] compared to those who experiment and do not transition. However, this "performance signal" is costly for both types, and in the cross-section those individuals who transition do not outperform those individuals who experiment and remain as mutual fund managers.

Implications

Our argument that prior employment can be a setting in which employees can 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 spawn, either in the same or a related industry segment. One testable assumption in our framework is that all else being equal, more learning should be required as an antecedent to transitions into a related industry, because it is less familiar. Finally, those employees that remain at incumbent firms should reduce their learning activities, which are costly, 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.

[2] We discuss this result in detail when reviewing Table 8. The "performance signal" in our context is a positive change in performance relative to a meaningful control group.

3. A Formal Model of On-the-Job Learning for Entrepreneurship

Next, we develop a general—if reduced-form—model that captures both of the main features of our argument and generates predictions. Later, we will apply this general framework to our empirical context to develop industry specific predictions that we can test with data from the asset management industry. The two primary features of our general model are that (1) an employee may engage in activities on the job, which provides information about their fitness for launching an entrepreneurial venture, and (2) such activity is costly in the sense that it reduces her productivity 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 entrepreneurship, (ii) who decides to become an entrepreneur, (iii) how those who become entrepreneurs 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.

To implement these basic features, we employ a workhorse framework for studying incentives in the employment relation: the principal-agent model (see e.g. Laffont & Martimort, 2002). Our model deviates from the standard model in that ours has two periods, and the employee (the agent) can decide to spend some of her effort in the first period to gain information about her fitness for entrepreneurship. In the second period, if an employee has learned about her fitness, she can decide to launch a new firm—or simply remain an employee. This simple adaptation allows us to study learning on the job and the implications for selection into entrepreneurship.

Model Setup. Consider the standard principal-agent model in which an owner (or principal) initially employs an individual (or agent). The employee's output is a noisy function of her effort, given by:

$$y = e + \varepsilon$$

where y is output, e is the effort the employee chooses, and ε is a noise term. We assume e is weakly positive, that is, $e \geq 0$, and the noise is mean zero, that is, $E(\varepsilon) = 0$. When the employee works within a firm, she faces a standard linear incentive contract, which we assume is exogenously determined in the market.[3] Specifically, as an employee, she earns $w(y) = \gamma_M y$, where γ_M is an incentive payment based on a share of the output.[4] The employee is risk-neutral, and her effort comes at a cost of $\frac{e^2}{2}$, so her utility of effort is:

$$\pi_M = \gamma_M y - \frac{e^2}{2}, \tag{1}$$

which is simply the expected wage minus the cost of effort. The employee chooses her effort to maximize (1), which implies her optimal effort and corresponding expected utility as an employee are given by

$$e_M^* = \gamma_M \tag{2}$$

$$E(\pi_M^*) = \frac{\gamma_M^2}{2}. \tag{3}$$

Our model departs from the standard model by introducing a second period in which an employee may decide to leave the firm. To find a new job in the second period, an employee must spend part of her effort during the first period to learn more about her fitness for the alternative position either within the industry or in a related industry.[5] Specifically, to learn her type, the employee must spend c units of effort, so her expected output in the first period would be $y(e|l = 1) = e - c + \varepsilon$, where l is an indicator

[3] The exogeneity of the wage contract is a reduced form assumption for market-based contracts that may be optimal across a large set of potential employees. Implicitly, we assume that the these contracts are not designed to screen employees based on their propensity to experiment—a dimension which we introduce later. In this sense, we assume the principal establishes an optimal "generic" contract for the market, so our model reduces to one which focuses solely on the agent. While it may be possible to design contracts to induce experiment-prone agents to act optimally as managers—for example, incentives such as retention bonuses—these contracts would be inefficient for those who never intended to experiment. Thus, we assume the standard linear contract is optimal in the context of the broader market. In sum, in many contexts, where the baseline probability of being an entrepreneur is low, including the one we study empirically, this assumption would be a reasonable approximation of reality.

[4] For the purposes of this analysis, we normalize the "base pay" in the standard model to 0 to simplify notation. The implication in what follows is that we assume all possible activities of the individual generate the same base amount of compensation, and the differences come from differences in the incentive-based components of the payout.

[5] Our model may be considered a case of the standard multi-tasking model (see for example Holmstrom and Milgrom 1991, 1994), given that agents can decide between two tasks in the first period. That said, we adapt the standard model to allow for future employment-relevant information, which allows us to analyze the behavior of employees who are in the risk set for entrepreneurship.

variable equal to 1 if the employee expends effort to learn, and 0 otherwise.[6] When l = 1, the employee maximizes:

$$\pi_M(e|l=1) = \gamma_M(e-c+\varepsilon) - \frac{e^2}{2} \qquad (4)$$

which means the employee exerts the same effort when she is learning on the job as in (2), and the expected utility when learning is given by:

$$E(\pi_M^*|l=1) = \frac{\gamma_M^2}{2} - \gamma_M c. \qquad (5)$$

If the employee has exerted effort to learn on the job in the first period, they learn how much utility they would derive from an alternative position. With probability *p*, they are a high type, and with probability *1-p*, they are a low type, where type determines the marginal productivity of effort in the alternative position, as we explain below.

If an employee has exerted effort to learn in the first period, at the beginning of the second period, she can either remain an employee or she can spawn into entrepreneurship.[7] If the employee decides to become an entrepreneur, she keeps all of the output she produces. Although the cost of effort is the same as when the individual is an employee, we assume an individual's marginal utility as an entrepreneur will depend on her type. Specifically, her output will be given by $\gamma_i(e+\varepsilon)$, where $i\epsilon\{H,L\}$. So the expected utility for an individual becoming an entrepreneur is given by:

$$E(\pi_i) = \gamma_i(e+\varepsilon) - \frac{e}{2} \quad i\epsilon\{H,L\}. \qquad (6)$$

Given (6) and following the same logic as in (2) and (3), we have that the optimal effort and corresponding expected utility an entrepreneur will receive is simply

$$e_i^* = \gamma_i \qquad i\epsilon\{H,L\}, \qquad (7)$$

[6] For ease of notation, we suppress time subscripts when the payoff calculations are clear.
[7] In our model, all individuals begin as employees of the firm. This condition implies that employees can *only* transition *if* they have experimented. In reality, some individuals begin their careers as entrepreneurs. That said, because we are focusing on entrepreneurial spawns, we are particularly interested in the choice by individuals who begin as employees of firms but who may eventually decide to become entrepreneurs. Although beyond the scope of this paper, a more general model would be of interest theoretically. A version of such a model, available upon request, provides more nuance on potential outcomes but also confirms the general results we present here.

$$E(\pi_i^*) = \frac{\gamma_i^2}{2} \qquad i\epsilon\{H,L\}. \tag{8}$$

To capture the idea that career transitions may be rewarding, we assume a low-type entrepreneur has a marginal utility less than the marginal utility of the individual as an employee, and a high-type entrepreneur has a marginal utility that is greater. In other words, we assume:[8]

$$\gamma_L < \gamma_M < \gamma_H. \tag{9}$$

Figure 1 summarizes the timing of the game. In the first period, an individual is an employee and must make two choices: whether to learn, and how much effort to exert. If the employee chooses to learn her type, she spends *c* units of effort. In the second period, if an employee has not engaged in learning, she remains an employee and exerts effort under the employment contract. If the employee engages in learning, she can choose to either spawn or remain an employee. The payoffs to the individual are simply the sum of the payoffs from each period.

Model Results. To solve this game, we use backward induction. Starting with those who engage in learning in the first period, given (3), (8), and (9), only high-type entrepreneurs will choose to become entrepreneurs; low types will remain employees.

Based on this framework and the solutions to the individual's optimal-effort problems in various scenarios, we can now analyze the decision to learn in the first period. More specifically, we have that the employee will choose to learn if and only if

$$p\left(\frac{\gamma_M^2+\gamma_H^2}{2}\right) + (1-p)\gamma_M^2 - \gamma_M c > \gamma_M^2.$$

The left-hand side of the expression is the payoff to learning: the employee gets the "on-the-job" payoff less the cost of learning in the first period, and in the second period, with probability *p*, she will become a high-type entrepreneur, and with probability *1-p*, she will stay on the job. The right-hand side is the

[8] The key assumption we make is that there is uncertainty about whether the manager will do better or worse as an entrepreneur. In our model, the key drivers of manager behavior will be their own relative levels of productivity. Thus, even if there was heterogeneity in the values of γ_i across individuals it would not change the basic results we derive below. Heterogeneity in quality will be an important empirical question which we address later in the paper.

payoff to staying on the job in both periods with the full (costless) on-the-job payoff. This setup implies the employee will decide to learn if and only if:

$$p\left(\frac{\gamma_H^2-\gamma_M^2}{2\gamma_M}\right) > c^*. \qquad (10)$$

The inequality in (10) indicates a critical value $c*$ exists such that the employee will decide to experiment only if the actual value of c is below that cutoff. The intuition is straightforward: employees will only decide to pay the price to learn when the costs are sufficiently low. The left-hand side of the expression is more instructive. The critical value represents the probability-weighted relative gain from being a successful entrepreneur against being an employee in the second period, weighed against the cost of learning as an employee in the first period. Thus the higher the critical value is (i.e., learning rates increase), the higher the relative payoff to entrepreneurship (i.e., the larger the difference between γ_H and γ_M) and the higher the probability of being a high type. Finally, as the payoff to being on the job increases—for example, as the share of output going to the employee increases—the likelihood of learning goes down.

This analysis allows us now to characterize the primary predictions of our model, which Proposition 1 summarizes (Proofs in the Appendix).

Proposition 1. *In the model of learning and entrepreneurship, the following ceteris paribus statements hold:*

(a) On average, those who become entrepreneurs will have worse performance in the first period than those who do not; more generally, those who experiment will have worse performance than those who do not experiment in the first period.

(b) Those who experiment and do not become entrepreneurs will subsequently improve their performance.

The proposition follows almost directly from the analysis above. First, as described in part (a), given that learning is both a precursor to entrepreneurship and is costly in the sense that it reduces output y on average, entrepreneurial transitions will be associated with “underperformance” in a job relative to the employees' potential (or, alternatively, given the effort they expend). Indeed, even those who learn they should be lifetime employees will suffer a temporary performance decline. Second, the fact that learning is costly further implies that when the period of learning ends, the employee will focus all of her on-the-job effort on her current job. Therefore, employees who have engaged in learning but decided against becoming entrepreneurs will see their performance (in terms of expected output y) improve.[9]

4. Empirical context

To test our propositions about entrepreneurial spawning, 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 new mutual funds and hedge funds.[10] As of 2012, mutual funds held a total of $24 trillion ($13 trillion in the United States) in assets under management.[11] Globally, hedge funds have over $2 trillion under management, with the vast majority of the largest funds located in the United States.[12] 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

---

[9] There are a number of other interesting implications of our model. For example, the model suggests that only a subset of high type entrepreneurs become entrepreneurs because some high types will never learn their type given the incentives to learn. In this sense, the costs of learning reduce expected output by stifling entrepreneurship. Additionally, the model also suggests that because marginal productivity is higher, those who do switch to entrepreneurship will expend more effort as entrepreneurs than as employees. This suggests output gains for two reasons—higher marginal productivity and more effort. In Proposition 1 we focus on what we are able to subsequently test given our empirical setting and leave these additional results to future exploration.

[10] To connect our theory to our empirical context, a mutual fund manager corresponds to an "employee" in the general model, and an "entrepreneur" includes all mutual fund managers who become hedge fund managers.

[11] http://www.icifactbook.org/ (Last accessed May 15, 2013)

[12] 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)

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.[13] 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") with much higher minimum investment 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 most likely transitioning to a more entrepreneurial job—one with more variable compensation. 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).

We explore several kinds of career transitions within these industries, though our data do not allow us to observe all kinds of transitions. For example, a mutual fund portfolio manager could stay in the same role throughout her career, transition to another managerial role within the same organization, or launch a new mutual fund or hedge fund within the same organization. The manager could also leave her

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

organization to join a new mutual fund, either taking over an existing fund or launching a new fund. Finally, the manager could launch a new organization with multiple mutual funds.

For within-industry transitions, we treat the subset of mutual-fund-to-mutual-fund transitions as entrepreneurial only where the focal mutual fund manager becomes the first portfolio manager at a new mutual fund firm. Because we want to compare intra-industry transitions with inter-industry transitions in a setting where inter-industry transitions are inherently riskier, we use this somewhat narrow definition of an intra-industry entrepreneurial transition to be sure we are capturing the biggest risk-takers in the population of intra-industry transitions. An alternative definition of an intra-industry entrepreneurial transition could include those cases in which a mutual fund manager becomes the portfolio manager at a new fund in an existing firm. Our results are qualitatively unchanged if we use this broader definition of intra-industry entrepreneurial transitions.

The other set of transitions in which we are interested involve the hedge fund industry. A former mutual fund manager could join an existing hedge fund, open a new fund within an established organization, or launch a new hedge fund. For the purposes of this study, we treat all mutual fund managers who transition to hedge funds as inter-industry entrepreneurial spawns. We use a somewhat broader definition of entrepreneurial spawning in this particular context because mutual-fund-to-hedge-fund transitions are uniformly to senior management positions in a hedge fund—all are called "principals" in the hedge fund data set—which suggests the transitions result in managers being residual claimants, and are thereby entrepreneurial in the spirit of our argument above[14]. However, the results are not sensitive to a narrower definition of spawning. We discuss exactly what we observe about mutual fund managers and how we classify their various career transitions below in our discussion of the data.

While the model presented above provides a general framework to understand experimentation and employee entrepreneurship, it may be helpful to consider the theoretical arguments with the specific domain of asset managers in mind. Next, we use the general framework above to develop specific

[14] More generally, hedge fund managers typically have greater personal stakes in firm profits and higher carried interest than mutual fund managers which is consistent with the notion of entrepreneurship in our model.

predictions for the asset management industry that we can test with our data. The central additional maintained hypothesis is that asset managers' action space concerns risk (rather than simply effort): they decide how to allocate risk to various investments in their opportunity set. In this setting, the experiment involves increasing risk exposure, and the results of the experiment reveal the mutual fund manager's ("employee" in the general model) fitness for managing a hedge fund. Of course, as mentioned above, we do not directly observe the experiment itself or the learning that comes from it. Instead, our analyses allow us to draw strong and reliable inferences into this process.

Our approach assumes the key uncertainty mutual fund managers interested in joining a hedge fund face is whether they will be successful taking on new and large risks. Although the economic benefits to running a hedge fund may be large, many (even successful) mutual fund managers may not be successful hedge fund managers. Mutual fund managers contemplating this transition, therefore, may benefit from "testing the waters" or experimenting by running different and more material forms of risk to inform their decision. These risks can come in the form of using higher than average leverage, increasing portfolio concentration, investing in less liquid securities (e.g. more thinly traded equities), greater variation in exposure to their benchmark and taking other actions more consistent with hedge fund investment management.

Model Setup. 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}.[15] 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.[16]

Following the structure of the general model in Section 3 above, we now assume the (portfolio) manager can either work for a mutual fund or become a residual claimant of a hedge fund. The former is

[15] The assumption of single-peakedness follows Palomino and Prat (2003).
[16] For a discussion of this type of risk-return set up see Palomino and Prat (2003) and de Figueiredo, Rawley and Shelef (2013).

analogous to an employee in our general model and the latter to an entrepreneur. Although these different forms of asset managers have different incentive schemes and payouts, we focus on the payoff functions to the manager as a function of the risk they take. Thus, as long as these incentive schemes are ordered as in the earlier case, the specific schemes are not important.

As before, two periods exist. In the first period, a 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 with probability p. We assume, as Figure 2 illustrates, 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 $h \epsilon \{H, L\}$ 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.

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.

Following the same logic as above, we have that a mutual fund manager will experiment if and only if:

$$2r^{max} < r^{exp} + p\alpha_H r^{max} + (1-p)r^{max}. \qquad (11)$$

The left-hand side of (11) 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 (11) and solving for p gives the following condition for experimentation:

$$p > \frac{1 - r^{exp}/r^{max}}{\alpha_H - 1} \quad . \qquad (12)$$

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. Analogous to the result in (10), (12) 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 (12), we have the following proposition, which follows closely on the results in Proposition 1:

Proposition 2. *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 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, those that take higher risks as mutual fund managers will take higher risks as hedge fund managers.

5. Data and empirical methods

To test our theoretical arguments, we use the complete CRSP Survivorship-Bias-Free U.S. Mutual Fund (herein 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 single portfolio manager responsible for the fund for at least 12 months.[17]

Using the 12,754 unique portfolio manager names in the CRSP dataset, we converted the data to 57,202 unique manager-fund dyads from 20,807 funds. These 57,202 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. We identified manager transitions between industries by hand matching at the manager level. An extensive pre-cleaning of the name fields led to nearly 100% 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, accounting for 1,794 manager-mutual-fund dyads. Three hundred and twenty-five of the mutual fund managers left their position at the mutual fund before at least one of the hedge funds they subsequently managed started operations, whereas 161 joined hedge funds that had already started.

Performance and risk are measured at the 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. Performance is based on the difference between raw returns and the benchmark risk-adjusted return predicted by this four-factor asset-pricing model. Taking the difference every month, and

[17] 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).

computing a moving average over time, provides a time-varying performance measure, whereas average performance relative to the four-factor benchmark captures performance meaningfully in the cross section. Idiosyncratic risk is calculated as the standard deviation 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.

CRSP records assets under management (AUM) 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 controls for the first date the manager is listed as the portfolio manager in the CRSP dataset ("start date"), and 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 cross-sectional tests include firm-fixed effects, so we leave other firm-level variables aside in the baseline tests. We also run a number of specifications with manager fixed effects.

For ease of exposition, we show our baseline mutual-fund-to-hedge-fund transition results at the manager level using each manager's largest fund, by average (over time) AUM. Focusing only on the fund with the largest AUM is a crude method of capturing manager performance; however, it represents one plausible way of studying manager-level performance when some managers manage multiple funds. We verified that this approach did not bias the results, by replicating the results elsewhere using all 57,202 dyads, and we show the results are robust to using all the funds managed by all the managers in the full panel in Table 4.

In the ideal experiment, we would randomly assign risk-taking to mutual fund managers so that we could directly observe the causal implication of risk-taking on spawning. Although the correlation between risk-taking and spawning is informative, risk-taking may be correlated with unobservable person-specific characteristics that we cannot control for directly in a cross-sectional regression. To create a valid counterfactual that can be used to address person-specific heterogeneity directly, 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 propensity-score-matching algorithm (Rosenbaum & Rubin, 1983), and use that set to make inferences about learning and spawning. In particular, by isolating the point in time when each "control group" manager appears to be experimenting with risk in a meaningful way, we can study how their risk-taking behavior changes following their learning experience.

To identify the set of managers in the control group, we analyze the data at the dyadic level, which gives equal weight to each of the 1,355 dyadic transitions that occurred more than one year before the end of our sample period. We then estimate a probit for dyad i at time t: $\Pr(Transition_{it} = 1|\boldsymbol{X_{it}})$, where, in our base specification, $\boldsymbol{X_{it}}$ includes all the observable characteristics of individuals, funds, and firms that might plausibly predict a transition, and the residual from the probit is the propensity score of the odds of transitioning. All of the other dyad-months 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 asynchronously to determine the control group's "match date."[18] After trimming extreme values and observations off the common support of the propensity-score distribution, 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, resulting in a set of 1,341 "treated" observations and 1,341 control-group observations.

6. Results

Table 2 shows manager-level tests of our baseline prediction that increased risk-taking leads to transitions. Column 1 shows the raw correlation from a probit model in which the outcome variable is equal to 1 for any manager who transitions and 0 otherwise. Standard errors are clustered at the firm level and marginal effects are reported. Doubling idiosyncratic risk exposure leads to a 0.72% increase in the

[18] See also de Figueiredo, Meyer-Doyle and Rawley (2013).

probability of transitioning. Including controls in column 2 reduces the point estimate to 0.59%. The inclusion of 1,256 firm fixed effects (in the OLS regression) generates a coefficient estimate of 0.62%, which means that compared to a 3.8% base transition rate, a one-standard-deviation increase in risk exposure leads to a 13% increase in the probability of transitioning (0.62*1.351/1.644/3.8). We find similar results when excluding transitions to existing hedge funds.

In Table 3, we analyze whether the level of risk-taking before a transition predicts the level of risk taking after the transition, by focusing on the 499 hedge funds launched by 442 mutual fund managers, where we have at least one year of hedge fund performance data. Column 1 shows a positive raw correlation of 0.17 between risk-taking ex ante and risk-taking ex post, a result consistent with Proposition 2(c). The result is similar when we control for the managers' historical mutual fund experience and their contemporaneous hedge fund performance, size, and tenure in columns 2 and 3. These results suggest that a one-standard-deviation in risk-taking before the transition is associated with an 11% increase in risk-taking after the transition (1.644*.16/2.04).

In an attempt to construct a valid counterfactual, Table 4 shows the results of matching "treated" observations to similar "control group" observations using dyad-month data. Column 1 shows that before the matching procedure, all the key variables from the non-transitioning population are statistically different from the transitioning group, which underscores the importance of matching in this context. In other words, mutual fund managers who become hedge fund managers are very different from the average mutual fund manager along all observable dimensions. Column 2 shows the marginal effects from a probit predicting a mutual-fund-to-hedge-fund transition. As in the manager-level tests, idiosyncratic risk is positively correlated with transitioning and precisely estimated, though in this table, risk is measured as a 12-month rolling average to better capture the relationship between transitioning and temporal variation in risk-taking. Doubling risk-taking leads to a 0.008% increase in the probability of transitioning in any given month, a 9% increase in the transition rate compared to a 0.089% base rate. After matching dyads that did not result in a transition to those that did using the propensity scores from the probit in column 2 and limiting the matched sample to a single observation per manager, all of the key differences between the "treatment" and "control" groups are eliminated (see column 3). Figures 3a and 3b show the effectiveness of the base match visually. The sign, significance and point estimate on risk-taking were similar under a wide range of alternative matching regimes

Using our 1,341 control-group observations, we can test how managers who experiment and decide not to transition behave after they have learned their type. Tables 5 and 6 show the change in risk and performance for the matched sample, respectively. Table 5, Column 1 shows the change in idiosyncratic risk-taking before and after the match date, for the 1,297 matched managers for whom we have at least 12 monthly observations before and after the match date, based on a simple test that uses a single observation before and after the match date. The coefficient on the "ex-post indicator" is -0.16, or a 10% reduction in risk-taking compared to the average level (0.16/1.644), and is precisely estimated. Column 3 examines the same ("base") matched sample at the manager-month level including manager fixed effects, year fixed effects, a control for performance (excess returns), and additional controls for fund size and age. Standard errors are clustered at the manager level. The estimated change in idiosyncratic risk-taking is -0.17 (column 3) and is precisely estimated. In columns 2 and 4 we show changes in risk-taking using alternative matched samples. When excluding performance from the first stage matching regime the second stage point estimate is lower (in absolute value), but still significant at the 5% level (column 2, "x-performance" matching). The second stage results are slightly larger (in absolute value) to the base match when including a full set of idiosyncratic risk times year interactions in the first stage (column 4, "saturated" matching). 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 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 manager learns her type and decides not to transition. Table 6 shows the same set of tests on 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 6-7 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).

Extensions and robustness checks

At this point, we have found evidence consistent with our general Proposition 1 and the context-specific predictions of Proposition 2. However, we would like to put the theory to more stringent tests wherever possible. One logical implication of our theory is that managers who make inter-industry transitions take more risk before transitioning than managers who make intra-industry transitions, because the latter group is not engaged in learning about a new industry segment. Because we observe both inter- and intra-industry transitions in the data, we can test this prediction by comparing mutual-fund-to-hedge-fund transitions with mutual-fund-to mutual-fund transitions. Additionally, 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 (voluntarily or involuntarily) their posts.

Table 7 shows the result of this 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 entrepreneurial transition (a transition to a new mutual fund) is coded as a "2;" all other exits from the mutual fund database by managers whose performance was below the median of the full-sample 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 486 managers who make mutual-fund-to-hedge-fund transitions. The magnitude of the effect, 0.8%, is slightly larger though broadly consistent with the results reported in Table 2. Column 2 shows that risk-taking is also elevated for the 582 mutual fund managers who transition to become the

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

founding portfolio managers of a new mutual fund in a new firm. However, the size of the effect, which is 0.2%, is noisy and is statistically smaller than the mutual-fund-to-hedge-fund effect (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" before eventually exiting the industry. 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 below-the-median performers are associated with a *reduced* level of risk-taking ex ante. The point estimate in column 3 is precisely estimated and reliably statistically different from managers who spawn from the mutual fund industry to the hedge fund industry (column 5), and is robust to alternative definitions of "other exits" (e.g. using the bottom quartile of the performance distribution or even all other exits).

The results in Table 7 provide two useful facts: (i) although experimentation on the job can facilitate learning about entrepreneurship in general, experimentation is far more important for inter-industry transitions than for intra-industry 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 mangers we study are not merely representative of a broader "gambling for resurrection" effect.

We also study and reject a different variation of the "gambling" hypothesis, whereby mutual 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, and, if successful, become hedge fund managers. We test this potential alternative explanation by comparing ex ante excess returns generated by transitioning managers against those managers in a matched sample identified by excluding ex ante performance (as in Table 5 and Table

6 column 2).[20] 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 creating a strong track record, we would expect that those who transition would have the highest performance from this sample. In contrast, the results show that managers who transition actually *underperform* managers who take on the same level of idiosyncratic risk ex ante, but do not transition, by 3.6 basis points per month in Table 8 column 1.

Another empirical concern is that the matched-sample results may simply reflect reversion to the mean. We attempt to address this concern by employing a long-run measure of risk-taking in the analyses in Table 5, which is less susceptible to mean reversion. Specifically, we measure 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. Although we identify the matched sample using 12-month rolling-average risk-taking, we do not use this measure in the "second stage" regression analysis.

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 (40/1,589). In fact, 939 of the 1,589 control-group observations take their maximum risk, on a rolling-average basis, after the match date. Thus, the fact that the control group decreases its risk-taking on average is not being driven by reversion to the mean.

Finally, we would like to understand more precisely what mutual fund managers learn when they experiment. In our theory we model managers experimenting as a way to learn their type, but we know that in practice, type is a multidimensional concept that includes both ability and preference. Although we suspect that many mutual fund managers who would be high-ability hedge fund managers simply prefer the lower-risk mutual fund environment, on the margin, ability in the alternative career should

[20] We use the "x-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.

matter too. Thus we construct a test that allows us to observe whether managers who transition get a more positive signal about their type during the experimental phase compared to managers who experiment but do not transition. To measure the signal about ability, we take all the treatment and control-group observations from the baseline match in Table 4 and regress the transition dummy on 12-month rolling-average excess returns in a differences-in-differences specification with year and manager-fund fixed effects on all the observations before the match/treatment date. This specification shows whether managers, who receive a better signal about their ability while experimenting, are more likely to transition. And, indeed, the results in Table 8 column 2 confirm that managers who transition improve their performance by 0.11%/month compared to managers who experiment but do not transition.[21] Thus at least part of the signal mutual fund managers pick up from experimenting with risk-taking on the job appears to be associated with their performance ability in a higher-risk investing environment.[22]

7. Discussion

Contributions

We advance a new mechanism by which prior employment facilitates entrepreneurship into the same or a related industry segment—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 enter entrepreneurship.

This work contributes to the growing literature on employee entrepreneurship, which has found that employees acquire valuable knowledge and/or run into bureaucratic constraints that operate as "pull/push" factors into entrepreneurship. Our key point is that employment also provides an opportunity

[21] This analysis is a comparison within manager for the matched sample. The average cross-sectional rolling-average excess return is -0.017%/month. This analysis is a test of *changes* in rolling-average excess returns. The result is robust to the definition of entrepreneurial spawning.
[22] One common empirical test in the spawning literature compares the performance of employee entrepreneurs, in this case former mutual fund managers, to other entrants in the same industry. Since we do not have background information on most of the entrants in the hedge fund industry, we cannot examine this question systematically with our data.

to experiment, such that employees can learn about their skills and preferences for a new position, with implications for both inter- and intra-industry 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.

While the prior literature has rightfully emphasized entrepreneurs as organizational products (Audia & Rider, 2006; Freeman, 1986), it has been difficult to account for individual heterogeneity in entrepreneurial spawning. It is not at all obvious whether it is the most talented employees, the worst matches (Åstebro, Chen, & Thompson, 2011), or some other category of employees who is most likely to leave to start a new venture. We introduce the notion that experimentation is a mechanism select employees can use to discover their own fitness for entrepreneurship, and depending on the results, decide whether to leave or stay at their employer.

Our argument has some similarities with anecdotal accounts about employees and entrepreneurship that are important to clarify. It is common for incipient entrepreneurs to work on a business idea during "nights and weekends" while maintaining their current salaried position, what might be referred to as "moonlighting". In 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).

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. One of the key contributions of this work is that our research design allows us to address the potential unobserved heterogeneity of entrepreneurial spawns that has been a concern in previous work. Most datasets only allow researchers to observe entrepreneurs and their prior employment affiliations, and rarely facilitate a matched set of non-entrepreneurs who continue at the same employer or transition into different kinds of careers, as we use in this study. While we do not claim to have the ideal experiment, our methodological contribution is important because estimating the treatment effect of prior employment on entrepreneurship, when sampling only on the dependent variable, introduces bias. Furthermore, 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.

Limitations

Our empirical context also introduces some limitations to the analysis however. While knowledge transfer, the inheritance of routines, and different levels of bureaucracy are all present in the asset management context, defining the focal industry and the entrepreneurial event itself are more challenging than in previous studies. For example, mutual funds and hedge funds are sometimes referred to as different asset classes, as opposed to different industries. Because of differences in investing strategies, disclosure requirements, business models, and many other dimensions, we treat them as distinct industries in the present study. However, even if we consider hedge funds and mutual funds to be within the same industry, our data still allow us to observe entrepreneurship in two different settings. Thus, interpreting the results as comparing employee entrepreneurship in a focal sub-sector (mutual funds) and adjacent sub-sector (hedge funds) would still constitute a contribution to the prior literature, though more in line with work on which markets employee spinoffs choose to enter (Klepper and Sleeper, 2005). Future scholarship should seek to explore entrepreneurial transitions into related industries with clearer boundaries, such as biotechnology and medical devices.

Another limitation is that institutional characteristics of the asset management industry make entrepreneurship a more complex outcome to measure than in previous studies. Starting a new firm is clearly an act of entrepreneurship, but joining a young firm as a principal with rights to the residual claims is more ambiguous. Due to the nature of this empirical context and the spirit of our model, we treat all mutual-fund-to-hedge-fund transitions as entrepreneurial transitions. Although our key results are qualitatively unchanged when we use alternative definitions of inter-industry entrepreneurship, the event of interest (entrepreneurship) may be easier to observe distinctly in other industry contexts.

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 allow the principal to design contracts that discourage experimentation. Given the relatively modest number of mutual fund managers who experiment in practice (16%), the expected cost to principals is quite low according to our calculations. Thus, the low cost of experimentation under the standard contract (which is optimal for the vast majority of managers) supports the notion that this assumption is not a limitation of our empirical study. However, scholars who wish to study learning on the job as a precursor to entrepreneurship in other contexts should be mindful of the need to test this assumption carefully.

Finally, using risk-taking as a proxy for experimentation rather than observing the experimentation process directly represents another limitation of this research. While we are able to utilize a relatively comprehensive industry dataset 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 engaging in entrepreneurship. Researchers will likely have to tradeoff high quality panel data that allows for counterfactual comparisons against having much more information on the intentions and activities of incipient entrepreneurs.

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 incipient entrepreneurs can experiment and learn about alternative careers. Through experimentation, employees learn how their skills and preferences match to new opportunities, a process that is particularly important for entrepreneurial transitions. Adopting this lens, we make precise predictions about the drivers of employee entrepreneurship in the same or a related industry segment and find evidence in support of our arguments. In doing so, we hope to open up new avenues of research in explaining patterns of employee entrepreneurship.

Appendix. Proofs of Propositions.

Proposition 1. Part (a) follows from two facts. First, only those who experiment can switch to entrepreneurship. Thus, all switchers are experimenters, meaning all switchers will have less output per unit of effort. Second, from (4), we have that the effort for experimenters is the same as non-experimenters in the first period. Part (b) follows from (2), (4) and the fact that $c>0$.

Proposition 2. 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 $\alpha_H v$.

$\alpha_L < 1 < \alpha_H v^{max} v^{exp} r^{exp} < r^{max} v^{max} \alpha_L < 1 < \alpha_H \alpha_H v f_M v^{exp} > v^{max} f_M \frac{\partial r^{exp}}{\partial v^{exp}} < 0$

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.

Amit, R., & Muller, E. (1995). " Push" and" Pull" Entrepreneurship. *Journal of small business and entrepreneurship, 12*, 64-80.

Anton, J. J., & Yao, D. A. (1995). Start-ups, spin-offs, and internal projects. *JL Econ. & Org., 11*, 362.

Åstebro, T., Chen, J., & Thompson, P. (2011). Stars and Misfits: Self-Employment and Labor Market Frictions. *Management Science, 57*(11), 1999-2007.

Audia, P. G., & Rider, C. I. (2005). A garage and an idea: what more does an entrepreneur need? *California Management Review, 48*(1), 6.

Audia, P. G., & Rider, C. I. (2006). Entrepreneurs as organizational products: revisited. *The psychology of entrepreneurship*, 1-33.

Bhide, A. V. (2000). *The origin and evolution of new businesses*: Oxford University Press, USA.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance, 52*(1), 57-82.

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.

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. (2013). Bad Bets: Excessive Risk-taking, Convex Incentives and Performance. *Unpublished manuscript.*

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

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.

Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org., 7*, 24.

Holmstrom, B., & Milgrom, P. (1994). The firm as an incentive system. *The American Economic Review*, 972-991.

Johnson, W. R. (1978). A Theory of Job Shopping. *The Quarterly Journal of Economics, 92*(2), 261-278. doi: 10.2307/1884162

Jovanovic, B. (1979). Job matching and the theory of turnover. *The Journal of Political Economy*, 972-990.

Kacperczyk, A. (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.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston: Houghton Mifflin.

Laffont, J., & Martimort, D. (2002). The Theory of Incentives: The Principal Agent Model: Princeton University Press.

Nelson, P. (1970). Information and Consumer Behavior. *Journal of Political Economy, 78*(2), 311-329. doi: 10.2307/1830691

Palomino, F., & Prat, A. (2003). Risk taking and optimal contracts for money managers. *RAND Journal of Economics*, 113-137.

Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika, 70*(1), 41-55.

Sørensen, J. B. (2007). Bureaucracy and entrepreneurship: Workplace effects on entrepreneurial entry. *Administrative Science Quarterly, 52*(3), 387-412.

Figure 1. Model of "On the Job" Learning about Entrepreneurship

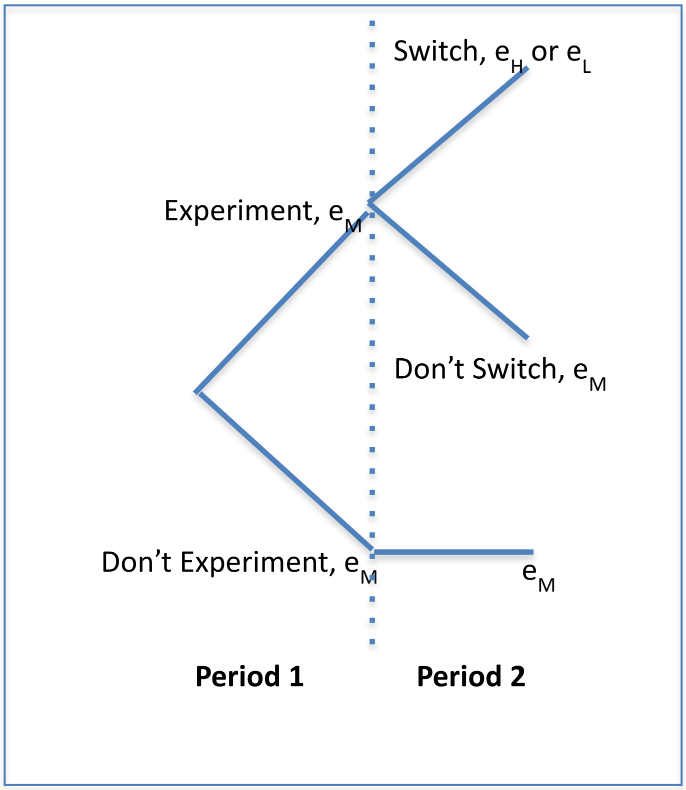

Figure 2. Efficient Frontiers by Type and Industry

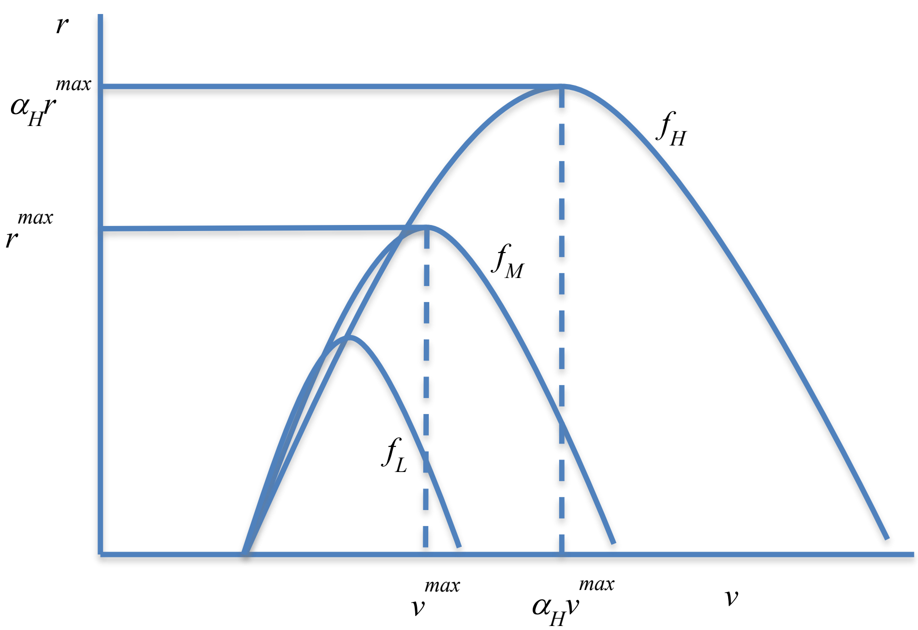

Figure 3a: Kernel density plot of the propensity score of a mutual-fund-to-hedge-fund (MF->HF) transition before matching (1st-95th percentile)

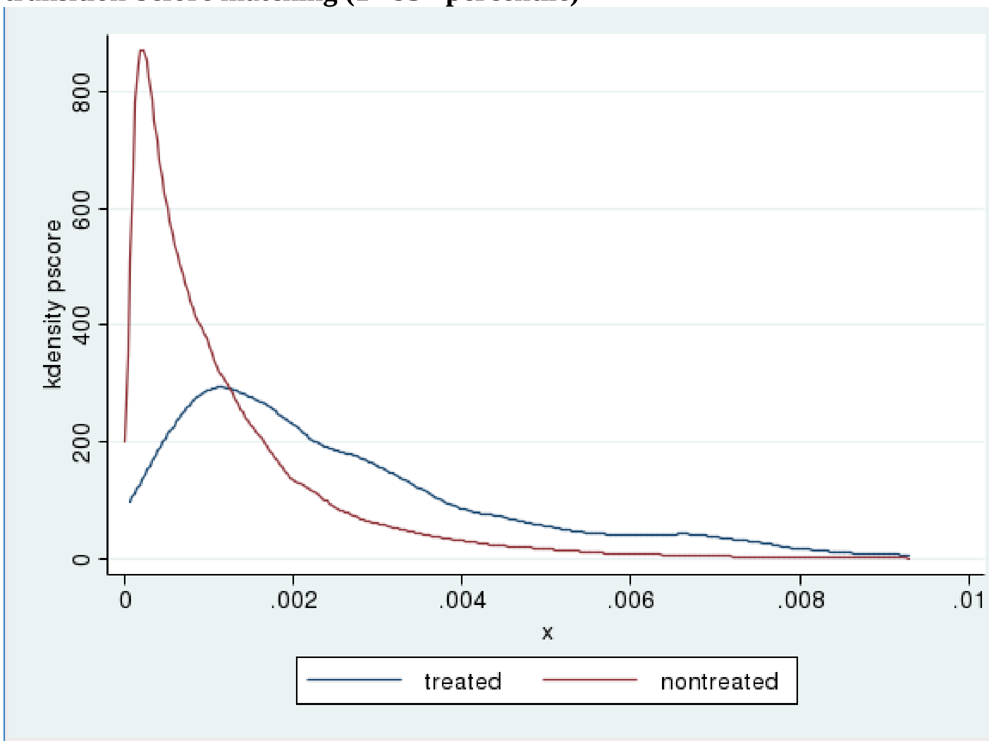

Figure 3b: Kernel density plot of the propensity score of a MF->HF transition after matching

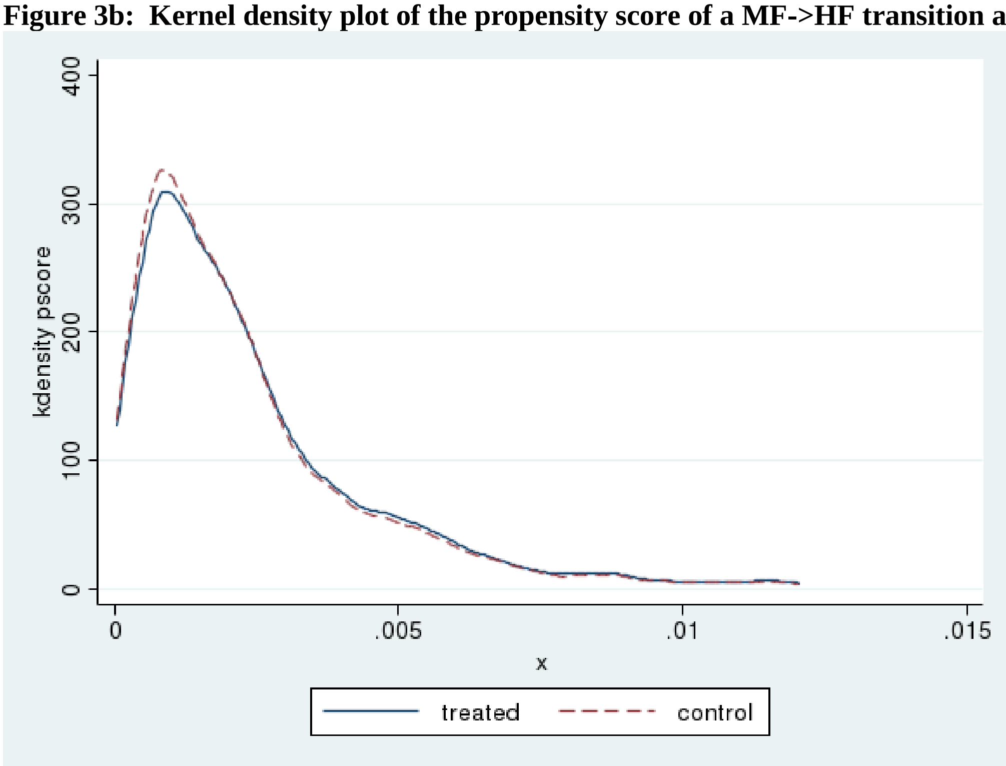

Table 1a: Key descriptive statistics for the mutual fund sample

N=12,754 mutual fund managers from 1,256 firms

	Mean	Std dev	Min	Max
Mutual fund to hedge fund transition	0.038	0.191	0	1
Idiosyncratic risk (% monthly)	1.644	1.351	0.021	6.796
Alpha (4 factor)	-0.059	0.495	-1.544	1.250
Start date	2/2001	64 mos.	12/1978	1/2011
Tenure as a portfolio manager (months)	45	37	12	229
Assets under management ($M)	967	3,234	1	74,403
Firm transitions (total count)	9.029	16.093	0	99
NYC HQ dummy	0.227	0.419	0	1

486 mutual-fund-to-hedge-fund (MF->HF) transitions.
Idiosyncratic risk is measured by the standard deviation of (monthly) excess returns.

Table 1b: Key descriptive statistics for hedge fund sample of MF->HF transitions

N=499 hedge funds (from 418 hedge fund firms, managed by 442 managers)

	Mean	Std dev	Min	Max
Mutual fund idiosyncratic risk (% monthly)	2.04	1.43	0.02	6.80
Hedge fund idiosyncratic risk (% monthly)	3.34	2.34	0.38	15.47
Mutual fund alpha (% monthly)	-0.08	0.54	-1.54	1.25
Hedge fund alpha (% monthly)	0.33	0.97	-2.45	4.80
Mutual fund manager tenure (months)	56.3	46.4	12	229
Hedge fund manager tenure (months)	80.5	59.5	11	321
Hedge fund start date	3/02	5 yrs.	1/80	1/10
Hedge fund assets under management ($M)	157	1,400	10	92,900
Missing hedge fund AUM	0.40	0.49	0	1
Mutual fund AUM ($M)	100	264	0	2,940
Mutual fund firm transitions (total count)	13.2	18.7	0	99

44 managers who transition have fewer than 12 monthly observations at the hedge fund they manage and so are dropped from “hedge fund” analyses.

Table 2: Mutual-fund-to-hedge-fund transitions – managers in cross section
Dependent variable: mutual-fund-to-hedge-fund transition (0,1)

	(1) Probit	(2) Probit	(3) OLS fixed effects
Idiosyncratic risk	***0.0072****	***0.0059****	***0.0062****
	(0.0010)	***(0.0096)***	***(0.0017)***
4-factor alpha		0.0052^{+}	0.0066
		(0.0029)	(0.0041)
Log AUM		-0.0007	0.0022^{+}
		(0.0009)	(0.0011)
Missing AUM		-0.0221*	-0.0190
		(0.0082)	(0.0153)
Manager start date		-0.0003*	-0.0003*
		(0.0000)	(0.000)
Manager log tenure		0.0031	0.0018
		(0.0023)	(0.0031)
Parent firm spawns		0.0004*	
		(0.0001)	
NYC dummy		-0.0010	
		(0.0037)	
Firm fixed effects	N	N	Y
Constant	Y	Y	Y
N	12,754	12,754	12,754
Pseudo/Adjusted R^2	0.01	0.04	0.04

* Significant at the 5% level, $^{+}$ Significant at the 10% level.
The unit of analysis is the manager.
Standard errors are clustered at the firm level.
All explanatory variables pertain to mutual funds.
In probit specifications marginal effects are reported.
There are 486 MF->HF transitions.

Table 3: Mutual fund managers as hedge fund managers

Dependent variable: hedge fund idiosyncratic risk

	(1) OLS	(2) OLS	(3) OLS
Mutual fund idiosyncratic risk	***0.16****	***0.17****	***0.16****
	(0.07)	***(0.07)***	***(0.07)***
Mutual fund alpha			0.03
			(0.19)
Hedge fund alpha			0.52*
			(0.11)
Log mutual fund AUM		0.00	0.01
		(0.07)	(0.07)
Log mutual fund firm AUM		-0.01	-0.02
		(0.05)	(0.05)
Log hedge fund AUM		-0.14+	-0.12
		(0.08)	(0.08)
Missing hedge fund AUM		0.43+	0.30
		(0.22)	(0.22)
Hedge fund start date		0.01*	0.01*
		(0.00)	(0.00)
Log manager tenure (mutual fund)		0.00	0.00
		(0.00)	(0.00)
Log hedge fund age		0.02*	0.01*
		(0.00)	(0.00)
Firm spawns		-0.00	-0.00
		(0.01)	(0.01)
Constant	Y	Y	Y
N	499	499	499
Adj. R^2	0.01	0.04	0.08

* Significant at the 5% level, + Significant at the 10% level.
The unit of analysis is the hedge fund.

Table 4: Mutual-fund-to-hedge-fund transitions – panel of manager-funds (“base match”)

Dependent variable: mutual-fund-to-hedge-fund transition (0,1)

	(1) t-tests before	(2) Probit	(3) t-tests after (3)
Idiosyncratic risk (rolling average)	***5.29***	***0.00008**** ***(0.00001)***	***-1.18***
Excess return (rolling average)	-4.67	-0.00016* (0.00004)	-0.44
Log age	-3.15	0.00031* (0.00005)	0.19
Log AUM	-10.00	-0.00016* (0.00002)	-0.17
Missing AUM	2.15	0.00055* (0.00022)	0.19
Firm spawns	14.45	0.00002* (0.00000)	-0.34
NYC dummy	-1.64	-0.00012+ (0.00007)	0.35
Year fixed effects	Y*	Y*	Y
N		906,425	
Pseudo R^2		0.06	
Prob(treated)		0.00095	

* Significant at the 5% level, + Significant at the 10% level.
The unit of analysis is the manager-fund-month.
Standard errors are clustered at the firm level.
Marginal effects reported in column 2.

There are 1,355 manager-fund transitions, and 1,341 “treated” observations on the common support of the joint distribution of the propensity score, after trimming at the 1st and 99th percentile of the control and treatment groups, respectively. After matching there are 1,341 matched “control” group observations (ties are broken randomly).

t-statistics are on the difference in means between the transitioning and non-transitioning populations reported in columns 1 and 4. In column (1) sixteen of the differences in the means of the year fixed effects are statistically significant; in column (3) only one is statistically significant.

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

Dependent variable: mutual fund idiosyncratic risk

Unit of analysis	Manager	Manager-month	Manager-month	Manager-month
Match type:	base	x-performance	base	saturated
	(1) OLS	(2) OLS fixed effects	(3) OLS fixed effects	(4) OLS fixed effects
Ex post indicator	***-0.16****	***-0.07****	***-0.17****	***-0.21****
	(0.03)	***(0.03)***	***(0.03)***	***(0.08)***
Excess return		0.01*	0.00*	-0.01
		(0.00)	(0.00)	(0.01)
Log AUM		0.04	0.01	0.02
		(0.02)	(0.02)	(0.01)
Missing AUM		-0.11	0.01	-0.10
		(0.12)	(0.10)	(0.08)
Log age		0.01	-0.03*	-0.01
		(0.01)	(0.01)	(0.01)
Year fixed effects	N	Y	Y	Y
Manager fixed effects	N	Y	Y	Y
Constant	Y	Y	Y	Y
N	2,594	124,494	125,931	125,098
Adjusted R^2	0.00	0.90	0.88	0.79

* Significant at the 5% level, $^{+}$ Significant at the 10% level.
Standard errors are clustered at the manager level.

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 (see Table 4 for specific details on the base match), and 0 otherwise.

In the base match (column 3) 44 of the 1,341 matched managers do not have 12 observations before and after the match date and so are dropped from the analysis.

Table 6: Matched sample behavior – changes in performance

Dependent variable: mutual fund information ratio

Unit of analysis	Manager	Manager-month	Manager-month	Manager-month
Match type:	base	x-performance	base	saturated
	(1) OLS	(2) OLS fixed effects	(3) OLS fixed effects	(4) OLS fixed effects
Ex post indicator	***0.07**** ***(0.02)***	***0.07**** ***(0.02)***	***0.06**** ***(0.02)***	***0.07**** ***(0.02)***
Log AUM		-0.01 (0.01)	-0.01* (0.01)	-0.01+ (0.01)
Missing AUM		0.05 (0.04)	0.01 (0.05)	0.04 (0.03)
Log age		-0.01 (0.01)	0.00* (0.01)	-0.00 (0.01)
Year fixed effects	N	Y	Y	Y
Manager fixed effects	N	Y	Y	Y
Constant	Y	Y	Y	Y
N	2,594	124,494	125,931	125,098
Adjusted R^2	0.00	0.81	0.81	0.82

* Significant at the 5% level, + Significant at the 10% level.
Standard errors are clustered at the manager level.

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 (see Table 4 for specific details), and 0 otherwise.

In the base match (column 3) 44 of the 1,341 matched managers do not have 12 observations before and after the match date and so are dropped from the analysis.

Table 7: Risk-taking and transitions – comparing MF->HF, MF->MF entrepreneurs, and other MF exits

Multinomial logit analysis

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	MF->HF	MF->MF entrepreneurs	Other exit	t-stat on difference (1) – (2)	t-stat on difference (1) – (3)
Idiosyncratic risk	***0.008**** ***(0.001)***	***0.002*** ***(0.001)***	***-0.013**** ***(0.004)***	***3.53****	***5.48****
Controls	Yes	Yes	Yes		
N		12,754			
Pseudo R^2		0.15			

* Significant at the 5% level, + Significant at the 10% level.
The unit of analysis is the manager.
Standard errors are clustered at the manager level.
Marginal effects are reported.
The baseline observation is "no transition."
Controls include: 4-factor alpha, Log AUM, Missing AUM, Start date, Log tenure, Parent firm spawns, and NYC dummy

There are 486 mutual fund managers who transition to being a portfolio manager at a new hedge fund (MF->HF), 582 mutual fund managers who become mutual fund entrepreneurs (MF->MF entrepreneurs), and 3,388 "other exits"—mutual fund managers who left the industry after performing poorly (i.e., were in the bottom half of the distribution of the information ratio over all mutual fund managers) and did not join a hedge fund in our sample. Categories are mutually exclusive: MF->HF; MF->MF entrepreneurs; "other exit"; "no transition"

Table 8: *Ex-ante* performance – MF->HF transitions versus the matched sample

Dependent variable: mutual-fund-to-hedge-fund transition (0,1)

Matching basis	(1) *x-performance* OLS cross-section	(2) *base* OLS fixed effects
Excess return	***-0.036****	***0.0011****
(rolling average)	***(0.013)***	***(0.0004)***
Idiosyncratic risk	0.012	-0.0009*
(rolling average)	(0.007)	(0.0003)
Log age	0.006	0.0333*
	(0.017)	(0.0014)
Log AUM	0.001	-0.0026*
	(0.017)	(0.0003)
Missing AUM	n/a	-0.0011
		(0.0034)
Year fixed effects	Y	Y
Manager-fund fixed effects	N	Y
Constant	Y	Y
N	2,680	144,877
Adjusted R^2	0.01	0.10

* Significant at the 5% level, + Significant at the 10% level.
The unit of analysis is the manager in column 1 and the manager-fund-month in column 2.
The sample includes matched sample observations *ex ante* (i.e., before the transition/match date.)
Standard errors are clustered at the manager-fund level.

After matching excluding ex ante performance, in column (1), there are 1,340 "treatment group" observations and 1,340 matched sample control-group observations.

After matching as in Table 4, in column (2), there are 2,682 manager-fund fixed effects—1,341 for the "treatment group" observations and 1,341 for the matched-sample control-group observations.