

Seeking Alpha: Excess Risk Taking and Competition for Managerial Talent

Viral Acharya (New York University)

Marco Pagano (Università di Napoli Federico II)

Paolo Volpin (City University London)

April 2016

Abstract

We present a model where firms compete for scarce managerial talent (“alpha”) and managers are risk-averse. When managers cannot move across firms after being hired, employers learn about their talent, allocate them efficiently to projects and provide insurance to low-quality managers. When instead managers can move across firms, firm-level coinsurance is no longer feasible, but managers may self-insure by switching employer to delay the revelation of their true quality. However this results in inefficient project assignment, with low-quality managers handling projects that are too risky for them. (*JEL* D62, G32, G38, J33)

Send correspondence to Paolo Volpin, Cass Business School, City University London, 106 Bunhill Row, London, EC1Y 8TZ, UK; telephone: +44 (0) 2070400961. E-mail: Paolo.Volpin.1@city.ac.uk.

Ulf Axelson, Marco Battaglini, Philip Bond, Mike Burkart, Adolfo de Motta, David Dicks, Itay Goldstein (editor), Michael Fishman, Hendrik Hakesen, Hugo Hopenhayn, Jack Shim, Giancarlo Spagnolo, two anonymous referees, and participants in the 2010 AEA Meetings, 2012 AFA Meetings, the Third Conference on Financial Integration and Stability in Mannheim, 2010 EFA Meetings, 2010 Mitsui Conference at University of Michigan, the 2011 WFA Meetings, the 2011 ESSET in Gerzensee, the 2012 CSEF-IGIER Symposium on Economics and Institutions, the 2013 Econometric Society Meetings in San Diego, and in seminars at Arizona State University, Bocconi University, Boston College, INSEAD, New York University provided useful comments and discussions. We acknowledge financial support from the Fondation Banque de France, ERC (FINLAB project no. 295709), Europlace Institute of Finance, Inquire Europe, and BNP Paribas Hedge Fund Center (HEC Paris).

In the last few decades, the financial sector, and particularly investment banking, has featured increasing competition for managerial talent. As argued by Morrison and Wilhelm (2008), this development occurred in investment banking since the increased importance of economies of scale associated with new technologies made the partnership model obsolete, and induced investment banks to turn into corporate entities and go public. While partnerships encouraged close relationships between employees and posed a natural obstacle to their mobility, the greater transparency of corporations facilitated the poaching of star employees and decreased their corporate loyalty.¹ This development also occurred in commercial banking, which once entailed a great deal of local knowledge, so that over their careers bank managers developed employer- and location-specific skills; today banking is much less local, owing to the increasing role of large banks, greater distance between banks and customers, and reliance on hard rather than soft information in lending (Petersen and Rajan, 2002, and Berger, et al, 2005). In fact, an increase in managerial mobility occurred even beyond the boundaries of the financial industry, as witnessed by the historical trend towards outside CEO appointments (Huson, Parrino and Starks, 2001): in 1940-67, 70 percent of top U.S. executives worked for the same company throughout their careers, while in 1990-2003 their fraction was only 30 percent of the total (Frydman, 2007).

Most of the academic and media attention has focused on the spectacular growth of financial managers' pay associated with this increase in the competition for their talent, and on the resulting increase in income inequality (Philippon and Reshef, 2012, and Bell and van Reenen, 2013). In this paper, we argue that, beside its effect on income distribution, competition for managerial talent may also lead to misallocation of talent by hampering employers' ability to learn the true skills of

¹As argued by Smith (2009), "in time there was significant erosion of the simple principles of the partnership days. [...] Competition for talent made recruitment and retention more difficult and thus tilted negotiating power further in favor of stars. [...] You had to pay everyone well because you never knew what next year would bring, and because there was always someone trying to poach your best trained people, whom you didn't want to lose even if they were not superstars. Consequently, bonuses in general became more automatic and less tied to superior performance."

bankers, traders and managers.

We make this point in a setting where managers are risk-averse while risk-neutral firms compete for scarce managerial talent. We model managerial talent as “alpha”, the ability to generate high returns without incurring high risks: lacking such talent, managers can generate high returns only by exposing their firm to the risk of correspondingly high losses. But risk only materializes in the long run, so talent can be identified with certainty only if the managers entrusted with skill-sensitive projects stay with their employer long enough. If they leave earlier, it may be impossible to identify their contribution to the long-term performance of their projects.

In this setting, if managers were bound to their employer, then over time firms could determine which managers are talented, and so could also insure managers against the risk of being found to be untalented. There would therefore be two efficiency gains. First, efficient allocation of investment projects to managers: when managers’ skills are known, they can be assigned to the project they are best suited to manage. Second, efficient risk-sharing: managers who prove to be low-skill can be cross-subsidized at the expense of the more talented.

However, competition for managers can prevent both of these gains. If firms compete aggressively (“seeking alpha”), then managers can leave before the long-term risks associated with their projects materialize. Hence, the managers who are discovered to be high-alpha types will extract all rents from their firms by generating competitive offers that reward their talent, and so prevent firms from subsidizing low-alpha managers. Thus if the labor market is competitive, managers face skewed performance rewards once their types are revealed: high-alpha types extract all rents and low-alpha types get no subsidy. Now, if firms assign managers of unknown quality to skill-sensitive projects (which they will do if on average such projects outperform alternative projects by a large enough margin), then managers have the incentive to move to another firm before the risk materializes. There, they will replicate the same behavior. In the aggregate, many managers will churn from one

firm to the next, being assigned to skill-sensitive projects regardless of their true “alpha”, i.e., their ability to avoid the implied risks. Talented executives will be identified only in the long run: as managers proceed in their careers, their true quality gradually emerges anyway, so that their incentive to churn decreases. The end result is that competition for managers lowers efficiency: since types are not revealed quickly enough, the efficient allocation of managers to projects is delayed and too many projects fail; too many skill-sensitive projects are assigned to untalented managers compared to the case where the managerial labor market features no mobility.

The result is reminiscent of Rajan (2005), one of the first to warn of excessive risk-taking in financial institutions driven by “fake alpha”. In our model, fake alpha is identified slowly because when job churning is possible, competition for managerial talent induces a negative externality: every firm effectively offers an “escape route” to the others’ employees, thus slowing down learning of true alpha and the assignment of skill-sensitive projects to the few managers who can competently manage their risks, as well as preventing efficient insurance of low-alpha managers against their human capital risk.

When the model is extended to the infinite horizon case, it produces potentially testable predictions regarding the correlation between managerial reputation and mobility. Mobility is positively autocorrelated, and decreases over a manager’s career, as information about the manager’s ability becomes sharper over time. Specifically, if alpha is sufficiently rare, managers churn from firm to firm only if their reputation lies in an intermediate range: once it exceeds a certain ceiling or falls below a floor, they stop churning. If instead alpha is sufficiently widespread, managers prefer not to churn across employers.

The model generates several further results. First, since the benefit of churning is to delay the revelation of a manager’s true quality, a key parameter in the model is the sensitivity of project performance to the manager’s quality: the lower such sensitivity, the better the manager can cover his tracks, and thus the greater the

insurance benefit from churning. But by the same token, the greater the implied sacrifice of productive efficiency, which requires early learning of managers' quality. Second, the more risk-averse managers are, the stronger will be their incentive to churn across employers to benefit from the implied insurance, and thus the more likely that untalented managers are assigned to skill-sensitive projects: ironically, greater risk aversion by managers entails greater risk for society. Third, even though managers' mobility would be lower if firms made their compensation conditional on the actual project payoff or on the manager's decision to leave the firm, firms have no incentive to condition managerial compensation on these outcomes if the labor market is competitive. Fourth, frictions in the market for managers (e.g. search costs) and asymmetric information about the manager's quality can actually mitigate inefficiency by reducing managerial churning. Finally, we allow mobility to improve the efficiency of the match between firms and managers, and show that this benefit may outweigh the efficiency costs highlighted by the baseline model.

To summarize, competition in the market for managers generates an inefficiency due to the contractual externality among firms. The financial sector appears to fit our model particularly well since trading and sales skills are highly fungible, prompting firms to compete keenly for "alpha". And many financial sector products, from mortgage-backed securities to credit default swaps or longevity insurance, have the feature of earning a carry (interest or insurance premium) in the short run but with potential long-run risks (default or longevity). While there are other explanations for excess risk-taking, e.g., government guarantees for the financial sector without proper risk controls, our model may help explain why it occurred even in parts of the financial sector, such as investment banks and insurance, that were not apparently entitled to government guarantees, explicit or implicit.

The paper is organized as follows. Section 1 discusses the literature. Section 2 describes the overall setting. In Section 3 we analyze the two-period version of the model, solve for the equilibrium in the non-competitive and in the competitive labor market regime, and compare their efficiency properties. Section 4 analyzes the

infinite-horizon version of the model. In Section 5 we relax several of our assumptions. Section 6 concludes with a brief description of the model's policy implications. The proofs are in the Appendix.

1 Literature

Our model of the labor market is close to that by Harris and Holmstrom (1982). Workers are long-lived and their productivity is uncertain. Because workers are risk-averse and firms are risk-neutral, the first-best is for firms to fully insure workers and pay a constant wage; but, as noted by Harris and Holmstrom, full insurance is not feasible if there is labor market competition and worker mobility. The reason is that under full insurance, workers who turn out to be very productive will be paid less than their marginal product. So competing firms will want to hire them, leaving the original firm with only low-productivity workers.

With respect to this framework, our paper introduces two novel elements: a project choice by firms, and a decision to move by managers. The choice of projects allows the firm to control whether types can become observable: the managers' type becomes known only if they are assigned to a skill-sensitive project, unless the employee moves to another firm before the project's payoff becomes known. Alternatively, employers can assign managers to projects whose payoff is not skill-sensitive, hence solving the Harris-Holmstrom problem: insofar as productivity shocks are hidden, full insurance becomes possible. But this insurance comes at a cost, since knowing a worker's productivity is useful in selecting the most suitable project for him. Hence, our model features a trade-off between the two information effects discussed in Hirshleifer (1971): information revelation has a cost (destroying insurance possibilities) but also a benefit (enhancing production efficiency). However, in our model the firm considers only the efficiency benefit in assigning workers to projects: if a worker stays on for more than one period, the employer learns his type and thereafter assigns him to skill-sensitive, high-yield projects if he is good or to talent-insensitive,

low-yield projects otherwise. Thus if a worker wants to delay the revelation of his type, he will try to churn across firms. Such mobility provides insurance, but also produces inefficiency in worker-project matching.

Our results represent a countervailing force to the benefits arising from competitive labor markets through efficient matching. A vast literature in labor economics, starting with Jovanovic (1979), highlights the benefits of mobility to achieve efficient matches between employees and employers, on the assumption that job matches are experience goods. In our setting instead, mobility results in less efficient matching of managers to projects within each firm.

The fact that competition for scarce talent in our model introduces an externality in wage setting is reminiscent of the corporate governance externalities formalized by Acharya and Volpin (2009) and Dicks (2012). In these models, competition prompts firms to incentivize managers via higher salaries rather than better governance. In the same spirit, Thanassoulis (2012) shows that competition for bank executives generates a negative externality, driving up remuneration and hence increasing rival banks' default risk. In contrast to these studies on governance externalities, our paper posits a dynamic setting in which firms can learn about their employees and assign them to the right tasks, but such learning is hampered by managers' ability to generate offers from other firms before their type is revealed.

Labor market competition may also lead companies to rely too heavily on high-powered incentives, shifting effort away from the less easily contractible tasks, such as risk management, towards the contractible ones. This point is captured by Bénabou and Tirole (2016), in a multitasking model where workers differ in productivity in a rewardable task and in willingness to perform an unrewarded one (work ethic). When firms compete for workers, they use incentive pay also to attract or retain the most productive workers, and by doing so they reduce work ethic below the social optimum. Our model is complementary to that by Bénabou and Tirole: we focus on employees' firm-level insurance and on how labor-market competition, by eroding

such insurance, leads to churning as an alternative way of synthesizing insurance; in contrast, they focus on multi-tasking and on how competition reduces effort in non-contractible tasks.

Finally, competition for talent may hinder firms’ ability to discipline managers, generating inefficient executive compensation in settings with moral hazard. Axelson and Bond (2015) show that smart workers may be “too hard to manage”, because their high outside options make them insensitive to the threat of dismissal. Makarov and Plantin (2015) develop a model of active portfolio management in which fund managers may secretly gamble in order to raise their reputation and attract investment, with trading strategies that expose investors to severe losses. Our analysis differs from these models insofar as excess risk-taking arises not from moral hazard but from inefficiently slow learning of employees’ skills.

2 Setting

There are K identical profit-maximizing firms, indexed by $k = 1, \dots, K$, owned by risk-neutral shareholders. Each firm employs I risk-averse managers. So managers are indexed by $i = 1, \dots, I \times K$. Both K and I are large: firms behave competitively, and each employs a large number of managers. Firms and managers have the same time horizon. Each manager i maximizes the discounted expected utility of future wages, conditional on current information:

$$V_{it} = \mathbb{E}_t \left[\sum_{s=0}^{T-1} \rho^s u(w_{it+s}) \right], \quad (1)$$

where $u(w_{it+s})$ is the utility of the wage w_{it+s} received in period $t+s$, ρ is the discount factor, $\mathbb{E}_t[\cdot]$ is the expectation conditional on the information available in period t , and T is the time horizon. In the simplest and most intuitive case, analyzed in Section 3, managers and firms have a two-period horizon ($T = 2$); Section 4 extends the analysis to the infinite-horizon case; and Section 5 contains extensions of the 2-period model. In all variants, $u(\cdot)$ is increasing and concave: managers are risk-

averse regarding their compensation. Moreover, they are born with no wealth and are impatient (their discount factor ρ being smaller than the market interest rate factor $1/(1+r)$), so that their consumption equals their wage at each date. Hence, managers do not insure themselves by saving against shocks to the value of their human capital due to changes in their reputation. This allows us to focus on the firm and on mobility across firms as the only sources of insurance against these shocks.²

Each firm can make its compensation conditional on the projects assigned to the manager and on past information about the manager. The results would not be affected if the firm could make pay conditional also (i) on the actual payoff of the project assigned to the manager or (ii) on the manager’s decision to resign and leave the firm. In both cases, in equilibrium firms will not make pay conditional on these additional outcomes, as shown in Section 5.2.

2.1 Projects and managers

Each manager can run a new project per period. The project produces its payoff at the end of the period. Managers are not all equally good: a fraction $p \in (0, 1)$ are high-quality managers, and a fraction $1 - p$ of them are low-quality. Moreover, high-quality managers are relatively scarce: $p \leq 1/2$. Initially, the manager i does not know his own quality $q_i = \{H, L\}$. Manager i starts working at any firm k and can move to another firm j before the project initiated in that period pays off.

Project payoffs are affected by two sources of risk: technological risk and managerial talent uncertainty. Some projects are exposed to both risks because the managers in charge of them have a degree of discretion, so that their skill affects the projects’ outcome: we label these α projects, to stress that their payoff is sensitive to the presence or absence of managerial “alpha”. Other projects feature purely technological risk: we refer to them as β projects.

²The impact of these shocks on consumption cannot be softened by borrowing either: the shocks that we analyze are not transitory ones, since they refer to the value of managers’ human capital.

Technological risk is firm-specific and diversifiable: it is captured by a random variable \tilde{y}_k affecting the payoff of all firm k 's projects, and is drawn from the same distribution for all firms, with mean $y > 0$.³ The payoff of β projects undertaken in firm k is $y_\beta = \tilde{y}_k$, that is, reflects only its technological risk, not the manager's skills: if assigned to such projects, managers have zero alpha, irrespective of their talent.

In contrast, the α projects of firm k produce positive alpha in the hands of good managers, and negative alpha in the hands of bad ones: in the first case, their payoff y_α is $\tilde{y}_k + (\bar{y} - y)$; in the second, it is $\tilde{y}_k + (\bar{y} - c - y)$, where $\bar{y} - y > 0 > \bar{y} - c - y$. Hence, when benchmarked against β projects in the same firm, α projects yield an extra gain $\bar{y} - y > 0$ when entrusted to good managers and a loss $\bar{y} - c - y < 0$ when entrusted to bad ones.⁴ This difference is illustrated in the upper panel of Figure 1.

Hence both projects are risky, but only the risk of α projects is affected by managerial talent. Conditioning on the manager's quality, the payoffs of α projects are identical to those of project β up to a positive or negative constant, while unconditionally they are a lottery that adds managerial risk to the payoffs of project β . So if the manager's type is uncertain, α projects feature both managerial quality risk and technological risk, while β projects feature only the latter. These assumptions imply that α projects are riskier than β projects. Yet this feature is inessential to our analysis: the results of the model would be unaffected if β projects were riskier than α projects due to greater technological risk.⁵ What matters is that learning about manager quality reduces the risk of mistakes in the assignment of α projects, thus raising their expected payoff, while it does not affect the payoff of β projects.

³Hence firms are assumed to be homogeneous in their average efficiency: an extension that allows for heterogeneous firms is presented in Subsection 5.5.

⁴Project α can be interpreted as a carry trade, which yields profit \bar{y} if closed in time. The skilled trader closes it in time; the unskilled trader, who does not know when to close, incurs a cost c .

⁵For instance, the model's results would be the same if firm k 's β projects were to yield $y_\beta = \tilde{y}_k + \tilde{\varepsilon}_k$, where $\tilde{\varepsilon}_k$ is an additional zero-mean technological shock such that $\text{var}(y_\beta) > \text{var}(y_\alpha)$, provided the realizations of $\tilde{\varepsilon}_k$ can be disentangled from those of \tilde{y}_k (being drawn from a different distribution or being common knowledge). If so, the firm could still learn the quality of the managers in charge of its α projects by benchmarking their payoff y_α against the payoff y_β of its β projects.

[Figure 1: Expected payoffs of project α]

A key assumption is that if a manager initiates a project of type α , his ability becomes perfectly known only if he remains in charge of it until the project pays off, that is, *until the end* of the corresponding period. If the manager leaves *before* the end of the period, the outcome of the project will reflect not only the manager's quality but also some noise, due to the fact that the project is no longer monitored by its initiator after his departure. This captures the idea that it takes time to determine a person's ability to manage such a project.

This assumption is illustrated in the lower panel of Figure 1: if the manager does not complete the project, with probability λ the project's expected payoff will reflect his type (\bar{y} if the manager is good, and $\bar{y} - c$ if he is bad), and with probability $1 - \lambda$ a noise factor that will make the project succeed (i.e., produce \bar{y}) with probability p , the same as if the initiator were randomly drawn from the managers' population. Hence, when noise intervenes the project's outcome is uninformative about the quality of its initiator. But the noise factor does not *per se* change the expected payoff of the project: as can be seen in Figure 1, even when the project is not completed by its initiator, project α succeeds with probability p and fails with probability $1 - p$.

To summarize, the payoff of project α depends both on the manager's type and on whether the manager stays or leaves. Defining the manager i 's type by the indicator $I_i = \mathbf{1}_{[q_i=H]}$ (equal to 1 if $q_i = H$ and 0 if $q_i = L$), α projects completed by manager i yield the following differential payoffs compared to β projects in the same firm:⁶

$$y_\alpha - \tilde{y} = \begin{cases} \bar{y} - y > 0 & \text{if } I_i = 1, \\ \bar{y} - c - y < 0 & \text{if } I_i = 0. \end{cases} \quad (2)$$

Instead, if left unfinished by manager i , α projects yield a differential payoff:

$$y_\alpha - \tilde{y} = \begin{cases} \bar{y} - y > 0 & \text{with probability } p_i, \\ \bar{y} - c - y < 0 & \text{with probability } 1 - p_i, \end{cases} \quad (3)$$

⁶Notice that, since these expressions are the same for all firms, for notational simplicity we have dropped the firm's subscript k from the realized payoff \tilde{y}_k : even though each firm has different realizations of the technological shock, these affect equally all of its projects.

where the probability of success p_i is

$$p_i = \lambda I_i + (1 - \lambda)p = \begin{cases} \lambda + (1 - \lambda)p & \text{if } I_i = 1, \\ (1 - \lambda)p & \text{if } I_i = 0. \end{cases} \quad (4)$$

Recall that if the manager leaves project α unfinished, the success probability p_i reflects his true quality (captured by the indicator function I_i) with probability λ and the noise factor with probability $1 - \lambda$. Hence, λ is the sensitivity of the project to its initiator's quality, or equivalently the informativeness of its outcome about the departed manager's quality. In the limiting case where $\lambda = 1$, the project always succeeds if initiated by a good manager and fails otherwise, so that its outcome is perfectly informative. In the polar opposite case where $\lambda = 0$, the project succeeds with the unconditional probability ($p_i = p$), irrespective of its initiator's quality.

The relative expected profitability of the two projects is assumed to satisfy the following condition:

$$\bar{y} - (1 - p)c > y > \bar{y} - c. \quad (5)$$

The left-hand side inequality indicates that, if the manager is of unknown quality, the expected payoff of project α exceeds that of project β : hence, on average managerial skills generate value – an assumption that will be relaxed in one of the extensions of the model. The right-hand side inequality states that, if the manager is bad, project β yields a greater expected return than project α . Assumption (5) implies that it is optimal to assign bad managers only to β projects, and good ones only to α projects: assigning bad managers to α projects would destroy value.

To characterize the difference between the two projects, it is convenient to define the variable $\eta \equiv (\bar{y} - y)/c$: $\bar{y} - y$ is the excess return that a good manager can generate if assigned to project α rather than β , while $c = \bar{y} - (\bar{y} - c)$ is the range of payoffs that project α produces in the hands of a good and a bad manager. Hence, we will refer to η as a measure of the risk-adjusted efficiency gain of project α compared to project β . Assumption (5) can thus be rewritten as:

$$1 - p < \eta < 1. \quad (6)$$

Another way of stating this assumption is that the probability p of finding a good manager must be large enough (i.e., exceed $1 - \eta$) as to induce firms to learn about managers' skills by assigning them to project α , and the project α should not be so efficient (i.e., $\eta < 1$) as to make it optimal for all types. Since we also assumed "alpha" to be an uncommon quality, the probability of a good manager must be $p \in (1 - \eta, 1/2)$. We extend the analysis to the case $p < (1 - \eta)$ in Section 5.1.

2.2 Market for managerial talent

We posit that in each period the pool of projects available to a firm includes at least one α and one β project per manager. Therefore, managers – not projects – are the scarce factor of production, since only managers can start a new project.

At the beginning of any period t , the firm decides whether to make an offer to the manager, who can accept or reject it. The offer consists of a sequence of wages $\{w_{ik\tau}\}_{\tau=t}^{\tau=T}$, where T is the maximum number of periods of employment. Being paid in advance, at the beginning of the relevant period, each wage $w_{ik\tau}$ reflects manager i 's expected productivity in period τ , and therefore is contingent on the project $P_{ik\tau}$ to which he will be assigned in period τ and on his perceived quality $\theta_{i\tau-1} \in [0, 1]$ conditional on the information available up to period $\tau - 1$:

$$w_{ik\tau} = w(P_{ik\tau}, \theta_{i\tau-1}), \quad (7)$$

where $P_{ik\tau} \in \{\alpha, \beta\}$ indicates whether manager i is assigned to project α or β in period τ . Since the belief $\theta_{i\tau-1}$ about the manager's quality evolves on the basis of his performance, the contract is effectively contingent on the payoffs of the past projects run by the manager at firm k and at previous employers. In the baseline version of the model, the period- τ wage cannot be contingent on the manager's decision to stay or leave the firm before the end of period τ : the maximum penalty for resignation is receiving no further wage payments from one's former employer. As already mentioned, this assumption is with no loss of generality (see Section 5).

A firm’s strategy is a profit-maximizing choice of wage offers and project assignments. More precisely, the firm chooses its offer $\{w_{ik\tau}\}_{\tau=t}^{\tau=T}$ to each manager i and, upon hiring him, assigns him to project $P_{ikt} \in \{\alpha, \beta\}$, so as to maximize its expected revenue, conditional on the belief θ_{it-1} :

$$\pi(P_{ikt}|\theta_{it-1}) = \begin{cases} \bar{y} - (1 - \theta_{it-1})c & \text{if } P_{ikt} = \alpha, \\ y & \text{if } P_{ikt} = \beta. \end{cases} \quad (8)$$

Firms commit to pay the sequence of wages that they have offered, but not to a specific project assignment: once the contract is agreed upon, the firm assigns the manager to whatever project P_{ikt} maximizes its expected profits. However, as we shall see, in equilibrium firms pick the most appealing projects from the managers’ viewpoint (i.e., those yielding wages that maximize their expected utility), due to ex-ante competition and symmetric information. Therefore, even if the choice of projects were entrusted to managers, they would pick the same projects as firms. The assumption that firms, rather than managers, pick projects is irrelevant in our model, which features perfect congruence between their objectives.

The manager’s strategy consists of a period-by-period choice of employer: manager i employed by firm k in period t will choose whether to keep working at firm k or switch to a new employer in period $t + 1$ as a function of the belief θ_{it-1} about his quality, so as to maximize the expected utility (1) from his compensation.

We assume that in offering wage contracts, firms bid competitively for managers, anticipating their future performance: hence, managers extract all of the expected profit that they generate with an employer. But, while *ex ante* there is perfect competition for managerial talent, switching costs may prevent it *ex post*: over time, managers may make firm-specific investments or develop location-specific tastes, impeding poaching by other firms. To bring out the implications of *ex-post* competition for managerial talent, in the baseline model we focus on the two polar cases where switching costs are either prohibitively high – the “non-competitive” regime – or absent – the “competitive” regime. In an extension, we consider the intermediate case of a managerial labor market with some frictions in the form of switching costs.

In the non-competitive regime, once a manager accepts a firm's initial offer, he can no longer leave. In the competitive regime, at the start of each period a manager chooses whether or not to leave his current employer. When indifferent, he is assumed to stay – a tie-breaking rule that reflects the presence of an arbitrarily small switching cost even in the competitive regime.

In both regimes, managerial performance is publicly observable: if a manager's ability becomes known to the current employer, it is also known to other firms. This assumption is not essential in our context, however. To see why, suppose that a manager's performance is visible only to his current employer. Then, in the competitive regime a manager who turned out to be good could move to another firm and, if assigned to project α , would want to stay there for a whole period, to allow the new employer to verify his talent. So even if the manager's performance were not publicly observed, outside offers would be effectively conditioned on his true type, once this has become known to the manager.

2.3 Time line

Assuming without loss of generality that the representative manager i is employed in all periods, the sequence of his actions in a typical period t is as follows:

(i) At the start of period t , manager i accepts an offer from firm k (or renegotiates his previous contract with firm k), which assigns him to project $P_{ikt} \in \{\alpha, \beta\}$.

(ii) Before completion of the project, the manager chooses whether to stay with employer k also in period $t + 1$ or leave.

(iii) Project P_{ikt} is completed and produces its payoff y_{ikt} . If $P_{ikt} = \beta$, the observed payoff is \tilde{y}_{kt} . If $P_{ikt} = \alpha$ and manager i stays, the project's excess payoff $y_{ikt} - \tilde{y}_{kt}$ over the observed payoff of β projects in firm k reveals manager i 's quality, by (2); if instead he leaves, the project proceeds unsupervised, so that its excess payoff $y_{ikt} - \tilde{y}_{kt}$ is a noisy signal of the manager's quality, by (3).

(iv) At the end of period t , the belief θ_{it} that his quality is high ($q_i = H$) is updated on the basis of the available information.

(v) In any subsequent period, the sequence of moves is the same as in (i), (ii) and (iii), with appropriate changes in the firm and time indices.

3 Two-period model

Some of the key results of the model can be obtained in a simple two-period setting. In this case, manager i 's expected utility (1) reduces to $V_{i1} = E_0 [u(w_{i1}) + \rho u(w_{i2})]$. As mentioned in Section 2.2, we compare two regimes: a competitive labor market where managers can freely move between firms at the end of period 1, and a non-competitive one where they cannot, and thus effectively commit to work in the same firm in both periods. As the critical difference between the two regimes is how much firms learn about managers' quality, we start by characterizing this learning process.

3.1 Evolution of beliefs about managerial quality

At the beginning of his career, the manager's quality is unknown: he is good with probability p and bad with probability $1 - p$. Hence, the prior belief that manager i 's quality is high ($q_i = H$) is $\theta_{i0} = p$. At the end of period 1, this belief is updated to θ_{i1} on the basis of the manager's performance, depending on whether he was assigned to project α or β , and on whether he has chosen to stay with his employer until completion of the project or not.

Specifically, if in period 1 manager i is assigned to project β ($P_{ik1} = \beta$), there is no updating, as the project's payoff is independent of i 's quality: $\theta_{i1} = \theta_{i0} = p$. If instead the manager is assigned to project α ($P_{ik1} = \alpha$) and stays until completion of the project, his payoff y_{ik1} can be benchmarked against the current realization \tilde{y}_{k1} of project β 's payoff. The difference between the two payoffs reveals his quality, as shown by in (2), and all players update their beliefs accordingly: if $y_{ik1} - \tilde{y}_{k1} = \bar{y} - y$,

manager i is revealed to be good, so that $\theta_{i1} = 1$; if $y_{ik1} - \tilde{y}_{k1} = \bar{y} - c - y$, he is revealed to be bad, so that $\theta_{i1} = 0$. Finally, if the manager is assigned to project α but leaves before the project is completed, beliefs are updated using Bayes' rule:⁷

$$\theta_{i1} = \begin{cases} \theta_H = \lambda + (1 - \lambda)p > p = \theta_{i0} & \text{if } y_{ik1} - \tilde{y}_{k1} = \bar{y} - y, \\ \theta_L = (1 - \lambda)p < p = \theta_{i0} & \text{if } y_{ik1} - \tilde{y}_{k1} = \bar{y} - c - y. \end{cases} \quad (9)$$

Since $\theta_H > p$, after a high excess payoff ($\bar{y} - y$), the belief that the manager is good is revised upwards ($\theta_{i1} > \theta_{i0}$), the more so the greater is the sensitivity of the project's payoff to managerial quality (λ). Symmetrically, since $\theta_L < p$, after a low excess payoff ($\bar{y} - c - y$) the belief is revised downwards ($\theta_{i1} < \theta_{i0}$).

3.2 Non-competitive labor market

When there is no ex-post mobility of managers, any initial hire is expected to stay both in period 1 and 2. Therefore any firm k will offer to manager i the wages (w_{ik1}, w_{ik2}) that maximize the present discounted value of its two-period profits:

$$E_0 \left[\pi(P_{ik1}|\theta_{i0}) - w_{ik1} + \frac{1}{1+r} (\pi(P_{ik2}|\theta_{i1}) - w_{ik2}) \right]. \quad (10)$$

Since firms are risk neutral, compete initially for managers and employ a large number of them, they will bid wages up to the point where they earn zero expected profits:

$$w_{ik1} = E_0 [\pi(P_{ik1}|\theta_{i0})], \quad w_{ik2} = E_0 [\pi(P_{ik2}|\theta_{i1})]. \quad (11)$$

Hence, the equilibrium lifetime wage of manager i is the revenue he is expected to generate over his entire career at firm k . By symmetry, all firms pay an identical lifetime wage, implying that managers are indifferent between them. Moreover,

⁷Looking at Figure 1, one can easily compute the probabilities of the manager's type being good conditional on the two observed outcomes of the risky project:

$$\theta_H \equiv \Pr(q_i = G | y_{ik1} - \tilde{y}_{k1} = \bar{y} - y) = \frac{p\lambda + p^2(1-\lambda)}{p\lambda + p^2(1-\lambda) + p(1-p)(1-\lambda)} = \lambda + (1-\lambda)p,$$

$$\theta_L \equiv \Pr(q_i = G | y_{ik1} - \tilde{y}_{k1} = \bar{y} - c - y) = \frac{p(1-p)(1-\lambda)}{(1-p)\lambda + (1-p)^2(1-\lambda) + p(1-p)(1-\lambda)} = (1-\beta)p.$$

managers are perfectly insured against the risk arising from their unknown quality: equation (11) implies that good managers subsidize bad ones.

Even though firm k does not know its managers' quality when it sets wages, it anticipates that in choosing the period-2 project, P_{ik2} , it will be able to condition on the true manager's quality. This is because, under assumption (5), it is optimal to assign the manager to project α in period 1 ($P_{ik1} = \alpha$), and since the manager will stay until the completion of this project his quality will be known by the beginning of period 2: $\theta_{i1} = q_i$. Hence, in period 2 the firm will optimally assign project α to good managers and project β to bad ones. This yields expected revenues:

$$\mathbb{E}_0 [\pi(P_{ik1}|p)] = \bar{y} - (1 - p)c, \quad \mathbb{E}_0 [\pi(P_{ik2}|q_i)] = p\bar{y} + (1 - p)y, \quad (12)$$

where the first expression is the expected revenue of project α undertaken in period 1 by a manager of unknown type, and the second is the expected continuation revenue produced by the two (known) types in period 2, weighted by their frequencies.

Substituting (11) and (12) in (1) yields the manager's expected utility level:

$$V_{i0} = \mathbb{E}_0 [u(w_{i1}) + \rho u(w_{i2})] = u(\bar{y} - (1 - p)c) + \rho u(p\bar{y} + (1 - p)y). \quad (13)$$

This equilibrium outcome features both (i) optimal risk-sharing, i.e., complete insurance of managers by firms; and (ii) productive efficiency, i.e., optimal assignment of projects to managers. So in the non-competitive regime, the managers' expected utility is maximal:

Proposition 1 (Equilibrium under no competition) *Without ex-post competition for managers, the first-best outcome is attained in equilibrium.*

3.3 Competitive labor market

The regime where managers are free to move between firms at the end of period 1 is illustrated by the time line in Figure 2.

[Figure 2: Time line of the 2-period model]

In period 1, the manager's type is unknown: when he is assigned to the period-1 project the belief about his quality is the unconditional probability $\theta_{i0} = p$. His decision to stay with firm k or move to another firm h before the completion of the period-1 project does not affect the expected payoff of the project, but does affect how much is learnt about his type: if manager i assigned to project α stays with the initial employer k until the project pays off, his type q_i is perfectly learnt; if i leaves before the end of period 1, the updating is described by (9).

We solve the model by backward induction starting from the firm's choice of project in period 2. Since in that period the manager may be employed by firm k or h (depending on the manager's choice to stay or leave), for simplicity we drop the firm's subscript from the project assigned to manager i and from his wage.

3.3.1 Firm's project choice in period 2

The firm employing manager i in period 2 will assign him to the project that maximizes its profit $\pi(P_{i2}|\theta_{i1})$ in (8), which depends on the manager's reputation θ_{i1} :

$$P_{i2} = \begin{cases} \alpha & \text{if } \eta \geq 1 - \theta_{i1}, \\ \beta & \text{if } \eta < 1 - \theta_{i1}. \end{cases}$$

The manager will be assigned to project α only if his reputation is sufficiently good, so that the risk-adjusted efficiency gain η of project α exceeds the conditional probability of the manager being bad, $1 - \theta_{i1}$. Owing to competition, the wage paid to manager i in period 2 equals his expected productivity:

$$w_{i2} = \begin{cases} \bar{y} - (1 - \theta_{i1})c & \text{if } P_{i2} = \alpha \\ y & \text{if } P_{i2} = \beta \end{cases} \quad (14)$$

Notice that in period 1 the manager i , being of unknown quality, must have been assigned to project α ($P_{i1} = \alpha$), by assumption (5). Depending on the project's payoff and manager's decision to stay or leave in period 1, his reputation θ_{i1} will take one of four possible values:

- (i) $\theta = 1$ if manager i stayed and the project's excess payoff was $\bar{y} - y$;
- (ii) $\theta = \theta_H$ if manager i moved and the project's excess payoff was $\bar{y} - y$;
- (iii) $\theta = \theta_L$ if manager i moved and the project's excess payoff was $\bar{y} - c - y$; and,
- (iv) $\theta = 0$ if manager i stayed and the project's excess payoff was $\bar{y} - c - y$.

The choice of projects in period 2 is as follows:

Lemma 1 *There are two cases to consider:*

1. if $\eta \geq 1 - \theta_L$, then $P_{i2} = \begin{cases} \alpha & \text{if } \theta \in \{1, \theta_H, \theta_L\}, \\ \beta & \text{otherwise;} \end{cases}$
2. if $\eta < 1 - \theta_L$, then $P_{i2} = \begin{cases} \alpha & \text{if } \theta \in \{1, \theta_H\}, \\ \beta & \text{otherwise.} \end{cases}$

3.3.2 Manager's decision to move or stay

We proceed backwards to the manager's period-1 decision whether to stay with the current employer (firm k) or to move to firm j . If the manager stays, his period-2 wage w_{i2} will equal \bar{y} if he is found to be a good type ($q_{i1} = G$), which happens with probability p ; or y if he is found to be a bad type ($q_{i1} = B$), which occurs with probability $1 - p$. Hence, his expected continuation utility is

$$pu(\bar{y}) + (1 - p)u(y). \quad (15)$$

If he moves, his reputation will be θ_H if project P_{ik1} succeeds, and θ_L if it fails. From Lemma 1, a manager with reputation θ_H is always assigned to project α in period 2; one with reputation θ_L is assigned to project α only if $\eta \geq 1 - \theta_L$. Hence:

1. if $\eta \geq 1 - \theta_L$, then the expected utility from moving is:

$$pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\bar{y} - (1 - \theta_L)c) \quad (16)$$

2. if $\eta < 1 - \theta_L$, then the expected utility from moving is:

$$pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(y) \quad (17)$$

Comparing the continuation payoffs from moving and staying, one obtains:

Proposition 2 (Decision to move in period 1) *Manager i switches firm at the end of period 1 if and only if*

$$(1 - p) [u(\bar{y} - (1 - \theta_L)c) - u(y)] \geq p [u(\bar{y}) - u(\bar{y} - (1 - \theta_H)c)], \quad (18)$$

where $\theta_H \equiv \lambda + (1 - \lambda)p$ and $\theta_L \equiv (1 - \lambda)p$.

Switching firms before the project terminates provides insurance to the manager, in the form of a less variable continuation wage: instead of the payoffs \bar{y} and y , the manager receives the less extreme payoffs $\bar{y} - (1 - \theta_H)c$ and $\bar{y} - (1 - \theta_L)c$, as $\bar{y} > \bar{y} - (1 - \theta_H)c > \bar{y} - (1 - \theta_L)c \geq y$. By moving, the manager trades a wage reduction $(1 - \theta_H)c$ in the state in which his type is good with a wage increase in the state in which it is bad ($\bar{y} - (1 - \theta_L)c - y$). The manager decides to move only when the expected benefit if he is good exceeds the expected cost if it is bad. But this insurance comes at the cost of a lower expected wage, because a manager who does not move – being of known quality – is always assigned efficiently (to project α if good and to project β if bad), while a manager who moves may be assigned inefficiently (to project α even if he is actually bad). This expected efficiency loss is an increasing function of the frequency of bad managers $(1 - p)$, since these managers are inappropriately assigned to project α when they move.

Hence, the choice between moving and staying involves a trade-off between the insurance benefit of mobility and its efficiency cost. The manager's risk aversion is therefore the key parameter in the decision to move: managers prefer mobility if they are sufficiently risk averse. Indeed, if they were risk neutral, they would choose not to move: by moving, they would suffer a reduction in the expected wage, but they would not value the implied insurance.

The trade-off is also affected by other parameters, besides risk aversion. Mobility is more attractive if η is high, i.e., if project α is much more efficient than project

β , even considering the losses from assigning it to bad managers. Conversely, an increase in the sensitivity of project α to its initiator’s quality, λ , makes mobility less attractive: intuitively, when project α ’s payoff is very informative about its initiator’s talent even when he does not complete it, moving does not allow him to cover his track, and therefore provides little insurance.

We can characterize the decision to move in period 1 as follows:

Proposition 3 (Characterizing the decision to move) *(i) If a manager moves, his period-2 wage has lower mean and lower variance than if he does not. (ii) The expected gain from moving is increasing in the efficiency gain (η) from project α , and is decreasing in the informativeness of project α ’s payoff (λ). (iii) The expected gain from moving is increasing in the manager’s risk aversion.*

Interestingly, in the proof of this proposition the assumption $p \leq 1/2$ guarantees that greater risk aversion makes mobility more attractive: intuitively, when “alpha” is not widespread, each manager will worry about not being one of the talented few, and therefore an increase in his risk aversion will lead him to value mobility more. As risk aversion increases, the trade-off gradually tilts in favor of managerial mobility, since the reduction in the variance of future compensation gets an increasing weight compared to the reduction of its expected value.

[Figure 3: Moving decision and risk aversion in the 2-period model]

To illustrate this point, in Figure 3 we assume constant relative risk aversion (CRRA) utility $u(w_2) = (w_2^{1-\gamma} - 1)/(1 - \gamma)$, and vary risk aversion γ while holding the other parameters fixed at $\bar{y} = 3$, $y = 1$, $c = 2.5$, $p = 0.4$ and $\lambda = 0.2$. As can be seen from the figure, moving dominates staying only if relative risk aversion γ exceeds 1.4. Ironically, as managers become more risk averse, society takes a greater amount of risk, since when they move across firms they are all assigned to the project α : mobility gives managers insurance, at the cost of greater risk taking for the economy.

To compute the expected utility of managers in the competitive labor market, notice that the first-period wage of the manager is the same as in the non-competitive case, i.e., the expected payoff from project α undertaken by a manager of unknown quality: $w_{i1} = \bar{y} - (1 - p)c$. Substituting the implied expression for period-1 utility and the continuation utilities (15), (16) and (17) in (1) yields the manager's maximum expected utility level:

$$V_{i0} = u(\bar{y} - (1 - p)c) + \rho \max \{ pu(\bar{y}) + (1 - p)u(y), [pu(\bar{y} - (1 - \theta_H)c) + (1 - p) \max(u(y), u(\bar{y} - (1 - \theta_L)c))] \}. \quad (19)$$

3.4 Comparing labor market regimes

It is easy to see that the expected utility (19) achieved under competition is lower than the first-best level (13) achieved if the labor market is not competitive: the period-1 utility, $u(\bar{y} - (1 - p)c)$, is the same, while the period-2 expected utility is lower, because without competition the manager obtains for sure the wage $p\bar{y} + (1 - p)y$ corresponding to the expected profits with complete learning. This implies optimal risk sharing, as the wage is not conditional on employees' quality, even though in period 2 this information is used to match managerial talent to projects. In other words, good managers subsidize bad ones: this cross-subsidy is feasible only because in the non-competitive regime good managers cannot leave for higher pay at other firms. Under the assumption of ex-post competition maintained in Section 3.3, instead, this cross-subsidization cannot be achieved, as any firm offering the wages (12) would lose all its good managers in period 2 and hence make losses: once the true quality of managers is known, other firms would offer the competitive wage $w_{ik2} = \bar{y}$ to good managers, outbidding the period-2 wage $p\bar{y} + (1 - p)y$ in (12). Hence, a firm offering the wages in (12) would be left only with overpaid low-quality managers in period 2. Therefore, ex-post competition destroys risk sharing, as in Harris and Holmstrom (1982).

It is worth noticing that under competition the first-best outcome is unattainable not only when in the competitive equilibrium managers move across companies ac-

according to Proposition 2, but also when they do not, condition (18) being violated: also in that case, in equilibrium good managers are paid the period-2 wage $w_{ik2} = \bar{y}$ in line with their quality, because ex-post competition bids it to that level, even if they do not move to another firm. So, even when a competitive labor market features no mobility, optimal risk sharing cannot be achieved. But at least in that case managers are efficiently allocated, since without mobility firms can learn their true quality in period 1 and allocate them efficiently to projects in period 2. When instead a competitive labor market features mobility, i.e., condition (18) holds, there is both inefficient assignment of managers and incomplete risk sharing, even though mobility provides some insurance compared to the case of no mobility. To summarize:

Proposition 4 (Inefficiency of the competitive labor market) *The competitive equilibrium features inefficient project assignment and partial risk-sharing if managers move across firms, and efficient project assignment but no risk sharing if they do not.*

In principle, firms might constrain themselves to play the efficient, non competitive equilibrium rather than the inefficient, competitive one, by signing no-compete clauses with each other. The situation is akin to a “prisoner’s dilemma,” as no individual firm has the incentive to abstain from poaching other firm’s managers, but social welfare would be higher if they all together credibly commit not to hire other firms’ managers. This suggests that policies that “throw sand in the wheels” of ex-post competition in the managerial market may increase welfare, effectively forcing firms to behave as if they had signed a binding no-compete agreement. We will return to the policy implications of the model in the conclusions.

4 Infinite-horizon model

As shown above, the analysis becomes quickly more complex if the manager’s horizon increases while staying finite. This is because the decision problem faced by

the manager is not stationary: as the number of periods increases, the number of contingencies to be considered in previous decisions escalates. In contrast, when the manager’s horizon becomes infinite, the problem is stationary, so that one can define stationary cutoffs for the manager’s reputation that determine his decision to move or stay. The key additional insight from this analysis is that mobility occurs only if his reputation lies in an intermediate range: for extreme values of his reputation, the insurance gain stemming from mobility is too low, because the information publicly available about the manager’s ability is already quite precise – another instance of the Hirshleifer effect.

If the horizon is infinite, in each period t the manager maximizes the expected utility from his future wages conditional on his past reputation, that is, on the common belief about his quality as of period $t - 1$, θ_{t-1} .⁸

$$V(\theta_{t-1}) = \mathbb{E} \left[\sum_{s=0}^{\infty} \rho^{t+s} u(w_{t+s}) | \theta_{t-1} \right] = u(w_t) + \rho \mathbb{E} [V(\theta_t) | \theta_{t-1}], \quad (20)$$

where in the second step the manager’s expected utility is shown in recursive form.

We analyze the model by considering a generic period t as described in Figure 4:

[Figure 4: Time line of the model with infinite horizon]

At the beginning of the period, the manager’s reputation coincides with the common belief about his quality θ_{t-1} . His current employer, firm k , assigns the manager to project α or β . Before completing the project, the manager can move to firm j . At the end of period t , the project’s excess payoff $y_t - \tilde{y}_t = \{\bar{y} - c - y, 0, \bar{y} - y\}$ is realized and the manager’s reputation is updated.

We proceed in three steps. First, we show how the manager’s reputation evolves over time. Second, we consider which project the current employer assigns to the manager at the beginning of period t , based on his reputation θ_{t-1} (dropping the

⁸Recall that in each period the manager is assumed to consume all of his wage income, so that his consumption equals his wage.

manager's index i to simplify notation). Third, we analyze his decision to stay or move to a new firm, based on how this choice is expected to impact his future reputation and continuation utility.

4.1 Manager's reputation

At the end of any period t , the common belief θ_{it} that manager i 's quality is high ($q_i = H$) is a sufficient statistic of manager i 's past employment history. In each period t the belief θ_{it} is updated on the basis of the manager's previous performance, depending on whether he is assigned to project α or β , and on whether he has ever chosen to stay with his employer for an entire period or not. As illustrated by Section 3.1 with reference to the updating of beliefs in period 1, if manager i is assigned to project β , there is no updating; if he is assigned to project α and stays until completion of the project, the manager's quality is revealed by his performance, so that the belief is updated either to $\theta_{i1} = 1$ or to $\theta_{i1} = 0$; if instead the manager is assigned to project α but leaves before the completion of the period-1 project, beliefs are updated according to expression (9).

Now, suppose that also after period 1 manager i keeps moving across firms: the in each subsequent period t the belief about his quality will keep being updated according to Bayes' rule. , As long as he moves across firms, the odds ratio $\theta_t/(1-\theta_t)$ of his type can be shown to evolve according to the following law of motion (dropping the manager's and firm's subscripts to simplify notation):

$$\frac{\theta_t}{1-\theta_t} = \frac{\theta_{t-1}}{1-\theta_{t-1}} \times \begin{cases} 1 + \frac{\lambda}{(1-\lambda)^p} \equiv 1 + \delta^+ > 1 & \text{if } y_t - \tilde{y}_t = \bar{y} - y, \\ 1 & \text{if } y_t - \tilde{y}_t = 0, \\ 1 - \frac{\lambda}{1-(1-\lambda)^p} \equiv 1 - \delta^- < 1 & \text{if } y_t - \tilde{y}_t = \bar{y} - c - y. \end{cases} \quad (21)$$

where δ^+ indicates the size of upward revisions of the ratio upon "good news" and δ^- the size of downward revisions of the ratio upon "bad news". For instance, the reputation of a manager whose first project did well improves if his second project does well too, and deteriorates otherwise; symmetrically, the reputation of a manager whose first project did badly improves if his second project does well, and deteriorates

otherwise. The size of upward revisions δ^+ is increasing in λ and decreasing in p : when the manager leaves the firm, good news have a large positive impact on his reputation if the project's outcome is very sensitive to the manager's quality (large λ), and if the chance of a lucky outcome is low (small p). The size of downward revisions δ^- is also increasing in λ but is increasing in p : bad news have a large negative impact on the manager's reputation if the project's outcome is very sensitive to his quality, and if the chance of a lucky outcome is high.

By iterating expression (21), the odds ratio at any future date $t + T$ is seen to be increasing in the odds ratio in period t : denoting the number of upward and downward revisions by U and D (where $U + D = T$), respectively, we can write it as

$$\frac{\theta_{t+T}}{1 - \theta_{t+T}} = \frac{\theta_{t-1}}{1 - \theta_{t-1}} \times (1 + \delta^+)^U \times (1 - \delta^-)^D, \quad (22)$$

so that manager's future reputation θ_{T+t} is increasing in his current reputation θ_{t-1} . Expression (21) can also be used to compute the law of motion of the manager's reputation itself:

$$\theta_t = \begin{cases} \theta_t^U \equiv \theta_{t-1} \times \frac{1+\delta^+}{1+\theta_{t-1}\delta^+} > \theta_{t-1} & \text{if } y_t - \tilde{y}_t = \bar{y} - y, \\ \theta_{t-1} & \text{if } y_t - \tilde{y}_t = 0, \\ \theta_t^D \equiv \theta_{t-1} \times \frac{1-\delta^-}{1-\theta_{t-1}\delta^-} < \theta_{t-1} & \text{if } y_t - \tilde{y}_t = \bar{y} - c - y. \end{cases} \quad (23)$$

Hence, the manager's reputation conditional on good news at t , θ_{it}^U , is increasing and concave in his past reputation θ_{it-1} : good news are less informative for already reputable managers. Conditional upon receiving bad news at t , the manager's reputation, θ_{it}^D , is increasing and convex in his past reputation θ_{it-1} : bad news are more informative if they concern reputable managers.

4.2 Project choice

The project $P_{kt} = \{\alpha, \beta\}$ to which the manager is assigned by firm k depends on the manager's reputation as of the previous period, θ_{t-1} :

$$P_{kt} = \begin{cases} \alpha & \text{if } \eta \geq 1 - \theta_{t-1}, \\ \beta & \text{if } \eta < 1 - \theta_{t-1}. \end{cases}$$

Because of perfect competition for managers, the manager is paid his expected productivity:

$$w_t = \begin{cases} \bar{y} - (1 - \theta_{t-1})c & \text{if } P_{kt} = \alpha \text{ and manager expected to stay at } t, \\ \bar{y} - [1 - \lambda\theta_{t-1} - (1 - \lambda)p]c & \text{if } P_{kt} = \alpha \text{ and manager expected to move at } t, \\ y & \text{if } P_{kt} = \beta. \end{cases} \quad (24)$$

4.3 Manager's decision to move or stay

When he takes his decision to move or stay in period t , the manager conditions on his past reputation θ_{t-1} , but takes into account that his decision will affect his future reputation θ_t .

If $\theta_{t-1} < 1 - \eta$, he does not benefit from moving, since his current employer assigns him to project β . Hence, by (23) his reputation remains unchanged: $\theta_t = \theta_{t-1}$.

If instead $\theta_{t-1} \geq 1 - \eta$, the current employer decides to assign the manager to project α , so that some updating of the manager's reputation occurs by the end of period t . Hence, in every period the manager can decide to delay learning about his true quality: it is revealed if he stays, while it may not be if he moves.

Specifically, if the manager stays, his continuation utility is

$$\theta_{t-1}V_H + (1 - \theta_{t-1})V_L, \quad (25)$$

where $V_H \equiv \frac{u(\bar{y})}{1-\rho}$ and $V_L \equiv \frac{u(y)}{1-\rho}$ are the present discounted utilities from being identified as type H and a type L , respectively. As $V_H > V_L$, his continuation utility from staying is strictly increasing in his past reputation θ_{t-1} .

If the manager moves instead, his continuation utility can be written as

$$[\lambda\theta_{t-1} + (1 - \lambda)p]V(\theta_t^U) + [1 - \lambda\theta_{t-1} - (1 - \lambda)p]V(\theta_t^D). \quad (26)$$

Hence, the manager's utility in (20) must be rewritten taking into account that his continuation utility takes two different forms depending on whether he stays or

moves. It is the sum of the utility from consuming his current wage, $u(w_t)$, and the discounted value of the maximum of the continuation utilities if he stays or moves:

$$V(\theta_{t-1}) = u(w_t) + \rho \max\{\theta_{t-1}V_H + (1 - \theta_{t-1})V_L, \quad (27)$$

$$[\lambda\theta_{t-1} + (1 - \lambda)p]V\left(\theta_{t-1}\frac{1 + \delta^+}{1 + \theta_{t-1}\delta^+}\right) + [1 - \lambda\theta_{t-1} - (1 - \lambda)p]V\left(\theta_{t-1}\frac{1 - \delta^-}{1 - \theta_{t-1}\delta^-}\right)\}.$$

To derive the manager's optimal decision regarding moving or staying, it is useful to characterize the function $V(\theta)$:

Lemma 2 *The manager's utility $V(\theta)$ is increasing in his reputation θ , and is bounded between V_L and V_H .*

Using these results, we can now establish the manager's optimal stopping rule:

Proposition 5 (Manager's reputation and mobility) *Define the upper and lower bounds for the manager's reputation:*

$$\underline{\theta} = \frac{[\lambda\underline{\theta} + (1 - \lambda)p][V\left(\frac{\underline{\theta}(1 + \delta^+)}{1 + \underline{\theta}\delta^+}\right) - V_L]}{V_H - V_L}$$

and

$$\bar{\theta} = \frac{[1 - \lambda\bar{\theta} - (1 - \lambda)p][V\left(\frac{\bar{\theta}(1 - \delta^-)}{1 - \bar{\theta}\delta^-}\right) - V_L](1 + \bar{\theta}\delta^+)}{(V_H - V_L)\{(1 + \bar{\theta}\delta^+) - [\lambda\bar{\theta} + (1 - \lambda)p](1 + \delta^+)\}},$$

where $0 < \underline{\theta} < p$ and $\bar{\theta} \in (0, 1]$. Then, if $p \leq \bar{\theta}$, the manager moves in period t if and only if $\theta_{t-1} \in [\underline{\theta}, \bar{\theta}]$; if $p > \bar{\theta}$, the manager never moves.

Hence, if the probability of being a good manager is sufficiently low ($p \leq \bar{\theta}$), i.e., if “alpha” is sufficiently rare, the manager chooses to buy insurance by moving across firms only when his reputation has “intermediate” values, namely falls in the interval $(\underline{\theta}, \bar{\theta})$. Intuitively, when his reputation drops to the lower bound $\underline{\theta}$, he stops moving because the wage that he would get by moving to a new firm is close to the wage that he would get if he stays with his current employer and is revealed as a bad type: hence, the insurance gain from moving is too modest compared with the

implied inefficiency in project assignment, a result already found in the two-period model. When instead the manager’s reputation rises to the upper bound $\bar{\theta}$, he stops moving because he is sufficiently likely to be revealed as a good type, so that the wage that he can expect if his true quality is revealed is likely to be the high wage \bar{y} : also in this case, the insurance gain from moving is too modest compared with the implied inefficiency in project assignment.

If instead the probability of being a good manager is sufficiently high ($p > \bar{\theta}$), i.e., if “alpha” is sufficiently widespread, the manager prefers not to buy insurance by moving across firms, because the risk of being revealed to be a bad type is low enough to be borne by him.

It is also interesting to note that the upper bound $\bar{\theta}$ defined by Proposition 5 may equal or even exceed 1, so that mobility will occur in the interval $(\underline{\theta}, 1)$. In this case, while the manager will eventually stop moving if his reputation becomes sufficiently bad, a manager with good enough reputation will never stop moving across firms.

5 Extensions

In this section we consider several extensions of the baseline two-period model. First, we relax the assumption that on average managers add value, i.e., consider the case in which $\bar{y} - (1 - p)c < y$: we shall see that in this parameter region ex-post competition may induce firms to assign managers to the skill-insensitive project β and thus forgo learning about managers’ quality. Second, we allow for pay to be conditional on the actual payoff of the project and on the decision to leave or not, and show that under competition firms will never exploit such conditionality in their offers. Third, we allow for the presence of switching costs in the managerial labor market, and show that they reduce mobility. Fourth, we consider the case in which managers have superior information about their skills, and show that also the presence of asymmetric information reduces managerial mobility. Finally, we allow mobility to improve the efficiency of the match between firms and managers, and show that this

“bright side of mobility” may outweigh its efficiency costs emphasized in the baseline model. In all of these extensions, to simplify notation we assume the payoff of project β to be riskless: $\tilde{y} = y$. This simplification entails no loss of generality.

5.1 Insuring human capital risk without mobility

Our model is related to that by Harris and Holmstrom (1982), who show that labor market competition and worker mobility prevent full insurance by firms. In our model, however, employers can overcome the Harris-Holmstrom problem by assigning managers to projects whose payoff is not skill-sensitive: this enables them to provide full insurance, because it hides workers’ productivity to competitors, although it also creates an efficiency cost for the firm itself, as it prevents the firm from assigning the employee to the most suitable project.

In the baseline version of the model, we effectively assumed this cost to be so high as to make this option a dominated one: by (5), $\bar{y} - (1 - p)c > y$, or equivalently $p > 1 - \eta$, implying that even managers of unknown quality produce on average a larger payoff when assigned to project α than to project β , so that firms always prefer to assign managers to project α rather than β . Formally, in this case the manager’s expected utility from being assigned to project β when his quality is unknown, $(1 + \rho)u(y)$, is strictly smaller than the expected utility under competition (19), and thus *a fortiori* also smaller than the expected utility under no competition (13).

In this subsection, we consider the opposite case where $p < 1 - \eta$, so that on average project β dominates project α , although we maintain the assumption that, when entrusted to a good manager, project α is still superior to project β (i.e., $\bar{y} > y$). Specifically, suppose that the firm assigns project β to a manager of unknown quality, so that it does not learn anything about his talent. Then, with $p < (1 - \eta)$, the project β may become preferable to project α even in the absence of competition, which implies that assigning project β to all managers becomes the first-best project allocation. This happens when the probability p of a good manager is smaller than

the threshold \underline{p} such that

$$(1 + \rho)u(y) = u(\bar{y} - (1 - \underline{p})c) + \rho u(\underline{p}\bar{y} + (1 - \underline{p})y), \quad (28)$$

where the right-hand side is the utility that the manager obtains from project α under no competition, and is increasing in p . Intuitively, for $p < \underline{p}$ learning requires such a high likelihood of failure that it is more efficient to forgo it.

However, the more interesting case is that in which learning would be efficient in the absence of competition but the firm prefers to forgo it in the presence of competition. This happens if the probability p of a good manager is in the range $[\underline{p}, \bar{p}]$, where \bar{p} is such that the expected utility under competition (19) equates $u(y)(1 + \rho)$: notice that \bar{p} exceeds \underline{p} because for any p expression (19) is smaller than (13), and they are both increasing in p . Intuitively, for $p \in (\underline{p}, \bar{p})$ the choice of project β enables the firm to provide full insurance to its employees, which under competition would be impossible to achieve if employees were assigned to project α . This comes at the cost of forgoing learning about employees' productivity. But if the probability of good managers is below \bar{p} , the cost of forgone learning is worth bearing when compared to the benefit of providing insurance to employees.

To summarize:

Proposition 6 (Insurance vs. learning) *If $p < 1 - \eta$, three cases can occur: (i) if $p < \underline{p}$, then firms will assign employees to project β and therefore not learn their ability but fully insure them, irrespective of labor market competition; (ii) if $p \in (\underline{p}, \bar{p})$, then firms will assign employees to project β under competition and to project α under no competition, so that in both cases employees are fully insured but there is learning only under no competition; and, (iii) if $p > \bar{p}$, firms will assign employees to project α and the results are as in the baseline model.*

5.2 Conditional pay

In the baseline version of the model, we assumed that the wage in period 1 cannot be contingent (i) on the actual payoff of the project assigned to the manager, or (ii) on the manager's decision to resign and leave the firm. In this section we remove this assumption and show that in equilibrium firms will not make pay conditional on these additional outcomes.

With conditional pay, the employer can defer compensation after the realization of the cash-flows and can choose a different pay when the manager stays or leaves. First, it is easy to show that even if managerial pay could be conditioned on the actual payoff of the project assigned to the manager, competition will induce firms to set pay equal to the manager's expected payoff from the project, given the manager's perceived quality: ex-ante competition for risk-averse managers will lead risk-neutral firms to offer contracts that are not performance-based.

Next, the employer (firm k) may want to choose a different pay when the manager stays or leaves. By doing so, firm k can increase the chances of retaining manager i by paying him a salary $w_{ik1} = 0$ if he leaves, and a fixed wage equal to the expected output $w_{ik1} = \bar{y} - (1 - p)c$ if he does not leave.

Given this contract, if the manager stays, his expected utility is

$$u(\bar{y} - (1 - p)c) + \rho[pu(\bar{y}) + (1 - p)u(y)], \quad (29)$$

since he is paid $\bar{y} - (1 - p)c$ at the end of the first period and his continuation utility in the second period is $u(\bar{y})$ with probability p (when his type is found to be G) and $u(y)$ with probability $1 - p$ (when his type is found to be B).

If he moves, his expected utility depends on whether θ_L is large enough that the manager is assigned to project α even when the period-1 payoff is low, which happens only if $\eta \geq 1 - \theta_L$. Hence, if $\eta \geq 1 - \theta_L$, the expected utility from moving is:

$$\rho[pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\bar{y} - (1 - \theta_L)c)]; \quad (30)$$

if $\eta < 1 - \theta_L$, then the expected utility from moving is:

$$\rho[p u(\bar{y} - (1 - \theta_H)c) + (1 - p)u(y)]. \quad (31)$$

Comparing the equations above we can show:

Proposition 7 (Decision to move with conditional pay) *Manager i moves if and only if:*

$$(1 - p)[u(\bar{y} - (1 - \theta_L)c) - u(y)] \geq p[u(\bar{y}) - u(\bar{y} - (1 - \theta_H)c)] + \frac{u(\bar{y} - (1 - p)c)}{\rho},$$

where $\theta_H \equiv \lambda + (1 - \lambda)p$ and $\theta_L \equiv (1 - \lambda)p$. In all other cases, manager i does not move.

Comparing Proposition 2 and 6, it is immediate that moving is less likely with conditional pay than without it. Moreover, only a sufficiently patient manager (one with sufficiently high ρ) chooses to move.

Will firms use conditional pay? In the model, there is ex-ante competition for managers, who are a scarce resource. Hence, managers extract all the surplus. Whether firms use conditional pay depends on whether this contract clause increases managers' expected utility.

Quite clearly, conditional pay will not be used when moving is optimal in Proposition 2. As a matter of fact, the expected utility with conditional pay in equation (30) is strictly lower than the expected utility without conditional pay in equation (16). Hence, competition will drive firms to offer pay that is not conditional on their moving decision.

The manager's expected utility is also strictly greater when condition (18) is met. To see this, notice that the expected utility from moving if there is no conditional pay is:

$$u(\bar{y} - (1 - p)c) + \rho[p u(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\bar{y} - (1 - \theta_L)c)], \quad (32)$$

which is strictly larger than (29) whenever condition (18) is met.

Finally, the manager's expected utility is identical with and without conditional pay when (18) is violated. To summarize:

Proposition 8 (Equilibrium compensation) *In equilibrium, no firm will condition pay on the manager's decision to move to another firm.*

5.3 Switching costs

In Section 3.4 we have compared two extreme labor market regimes: one in which there is perfect ex-post competition and another where there is no competition at all. In this section we consider the intermediate case in which managers suffer a switching cost s if they switch employers. The cases analyzed so far correspond to the case in which $s = 0$ (perfect competition) and $s > u(\bar{y})$ (no competition).

With switching costs $s \in (0, u(\bar{y}))$, if the manager stays, his continuation utility is as in equation (15). If he moves, his expected utility depends on both θ_L and s . If $\eta \geq 1 - \theta_L$, then the expected utility from moving is:

$$\max\{pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\bar{y} - (1 - \theta_L)c) - s, 0\}; \quad (33)$$

if $\eta < 1 - \theta_L$, then the expected utility from moving is:

$$\max\{pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(y) - s, 0\} \quad (34)$$

Hence:

Proposition 9 (Decision to move with switching costs) *Manager i moves if and only if*

$$(1 - p)[u(\bar{y} - (1 - \theta_L)c) - u(y)] - s \geq p[u(\bar{y}) - u(\bar{y} - (1 - \theta_H)c)],$$

where $\theta_H \equiv \lambda + (1 - \lambda)p$ and $\theta_L \equiv (1 - \lambda)p$. In all other cases, manager i does not move.

Comparing the condition in Proposition 8 with that in Proposition 2, it is immediate that the higher the switching costs s , the smaller the parameter region in which managerial mobility is worthwhile.

5.4 Asymmetric information

The assumption of symmetric information between firms and managers is critical to our results. If all managers knew their type, then in equilibrium no insurance could be obtained by moving: good managers would stay with their firms to reveal themselves as good and get higher pay. Bad managers would then also be revealed and be assigned to projects of type β from period 2 onwards.

A less extreme assumption is one where only a fraction ϕ of managers know their type from the start. In this case, in equilibrium mobility decreases for two reasons: (i) mechanically, the fraction $p\phi$ of managers who know they are good will stick with their employer to demonstrate their type; and (ii) managers of unknown type will get pooled with those who know they are bad, and so will be less willing to move than in the baseline model.

This happens because the probabilities of the manager's type being good, conditional on the two observed outcomes of project α , change as follows:

$$\theta_H = \frac{(1 - \phi)[\lambda + p(1 - \lambda)]}{(1 - \phi)[\lambda + p(1 - \lambda)] + (1 - p)(1 - \lambda)} \leq \lambda + (1 - \lambda)p, \quad (35)$$

and

$$\theta_L = \frac{(1 - \phi)p(1 - \lambda)}{(1 - \phi)p(1 - \lambda) + \lambda + (1 - p)(1 - \lambda)} \leq (1 - \lambda)p, \quad (36)$$

where both θ_H and θ_L are decreasing in ϕ . Hence, condition (18) in Proposition 1 is less likely to be met. To summarize:

Proposition 10 (Decision to move with asymmetric information)

Managers are less likely to move as the degree of asymmetric information ϕ increases.

5.5 Allowing for the matching gains from competition

We now consider the case in which there is heterogeneity of *both* firms and workers, assuming that there is a chance $m_i = m < 1$ that worker i is well matched with his initial employer. Before the end of the first period (and before moving), the manager learns whether there was a good match or not, as in Jovanovic (1979) and the vast subsequent literature on learning in labor economics. For simplicity, we assume that this information increases the chance of a good match with the next employer from m to 1, as the employee perfectly learns which type of firm is appropriate for him. The payoffs for project α are as in the baseline case in case of a good match, and are the same as with project β in case of mismatch.⁹ Specifically, if manager i stays with firm k the payoffs are

$$y_{\alpha i} = \begin{cases} \bar{y} & \text{with probability } m_i I_i, \\ \bar{y} - c & \text{with probability } m_i(1 - I_i), \\ y & \text{with probability } 1 - m_i, \end{cases} \quad (37)$$

where, as before, the indicator $I_i = \mathbf{1}_{[q_i=H]}$ denotes manager i 's type and $m_i = \{0, m, 1\}$, depending on whether manager i is a bad, unknown or good match for firm k . If instead the manager moves the payoffs are

$$y_{\alpha i} = \begin{cases} \bar{y} & \text{with probability } m_i[(1 - \beta)p + \beta I_i], \\ \bar{y} - c & \text{with probability } m_i[(1 - \beta)(1 - p) + \beta(1 - I_i)], \\ y & \text{with probability } 1 - m_i. \end{cases} \quad (38)$$

Notice that, as in the baseline case, staying or moving does not affect the expected payoff of the project already initiated by the manager, but moving reduces the probability of learning about his managerial talent.

Consider first the benchmark case with no ex-post competition. In this case, as in the baseline model, there is full insurance but now there is a cost arising from the possible mismatch between manager i and firm k . The expected utility is:

$$V_0 = u(mp\bar{y} + m(1 - p)(\bar{y} - c) + (1 - m)y) + \rho u(mp\bar{y} + (1 - mp)y). \quad (39)$$

⁹This assumption is without loss of generality. The only requirement is that the payoff in case of mismatch is independent of (and thus uninformative about) managerial quality.

Specifically, the manager is assigned to the project α in the first period and is paid the expected payoff that he produces. In the second period, the manager's assignment depends on the first period's payoffs: the manager is assigned again to project α only if in the first period $y_{\alpha i} = \bar{y}$; in all other cases, the manager is assigned to project β . As in the first period, the employer insures the manager by paying him a fixed wage equal to the expected payoff.

In the case of ex-post competition, the manager now faces a greater benefit from moving than in the baseline case: in case of mismatch, by moving the manager can find a better match. As shown in the following proposition, the benefit from moving may be so large that the expected utility under ex-post competition may exceed that obtained in the absence of ex-post competition:

Proposition 11 (Bright side of competition) *If firms and workers are sufficiently heterogeneous (i.e., if the probability m of a good match is low enough), the equilibrium with ex-post competition dominates that with no ex-post competition.*

Intuitively, when a random allocation of employees across firms would feature a high degree of mismatch, the welfare gain from reallocating employees across firms dominates that from learning about the talent of the employees of each firm and reallocating them across its projects.

6 Conclusions

The efficient allocation of talent is also considered to be the prime function of a competitive market for managers (see Gabaix and Landier, 2008, among others). Here, however, we show that when projects have risks that materialize only in the long term, there may be a dark side to competition for managers: by destroying the boundary of the firm that encapsulates its employees, short-run labor market opportunities interfere with the long-run information-gathering function of the firm.

Competition hampers each firm’s ability to provide insurance to risk-averse employees, and at the same time allows managers to churn across employers so as to delay the resolution of uncertainty about their talent; but by doing so they also hinder their employers’ ability to allocate them efficiently across projects.

Our model has important policy implications for the financial sector, where projects with long-run risk are often available. In our inefficient churning equilibrium, no individual financial institution has the incentive to deviate and unilaterally stop competing for the others’ managers, so that only intervention by a public authority can stop banks from poaching one another’s managers.¹⁰ No employer can insulate itself from such competition unless all its employees signed a no-compete clause that is enforceable – a possibility that is precluded in our regime with ex-post competition. Our model implies that discouraging managerial mobility – say, taxing relatively young managers who switch jobs – can improve efficiency: if such a surtax were high enough, it would effectively move the economy to the first best (although in equilibrium it would not be paid, since managers would not switch jobs). In short, a policy prescription deriving from the model is to “throw sand in the wheels” of the managerial labor market.¹¹

Another policy implication of the model is capping managerial compensation in banks. How would this change the equilibrium with managerial competition? Capping managers’ pay at the first-best level would prevent employers from poaching good managers in the competitive regime and make the perfect risk-sharing and no-churning outcome sustainable in equilibrium. Hence, capping the pay of top

¹⁰This idea is captured well by Tett (2009): “Banks operate in a world where their star talent is apt to jump between different groups, whenever a bigger pay-packet appears, with scant regard for corporate loyalty or employment contracts. The result is that the compensation committees of many banks feel utterly trapped. [...] It is time, in other words, for bankers and regulators to [...] start debating not just the issue of pay, but also the poaching culture that is at the root of those huge bonus figures.”

¹¹One such proposal is currently being considered by the Bank of England, which wants to tighten up rules around so-called “bonus buyouts”, whereby banks compensate newly hired employees for any remuneration cancelled by their previous employer when they changed jobs (Bank of England Prudential Regulation Authority, 2016).

financial managers may respond not only to ethical or political concerns but also to an efficiency rationale: reducing the excessive risk-taking associated with churning. Indeed, according to the model, an appropriate pay cap would raise the expected utility of managers themselves.¹²

Admittedly, in more elaborate models some of these policy interventions would entail efficiency costs. Either a salary cap or an equivalent surtax on managerial mobility would redistribute income from good to bad managers, which could decrease efficiency in a model in which managers themselves invest in their own quality *ex ante* – by investing in better education, say. In this case, capping their salary would reduce the “average alpha” of managers in equilibrium. Moreover, preventing the reallocation of managerial talent could have other efficiency costs: as we have shown in the last of the extensions of the model, if both managers *and* firms are sufficiently heterogeneous, so that allowing bad matches to be dissolved and new ones formed can dominate those from the appropriate allocation of talent to projects within each firm. Finally, limiting managerial mobility may give market power to firms and create hold-up problems. In our setting, this is inconsequential because of *ex-ante* competition, but in reality this assumption too might not hold. While all this suggests the need for caution in drawing policy conclusions, our analysis highlights that the competition for managerial talent may generate inefficiencies that have been so far neglected and are potentially policy relevant.

¹²Interestingly, also in the setting of Bénabou and Tirole (2015) a cap on managerial pay, hence a reduction in its sensitivity to performance, can restore the first-best outcome.

Appendix: Proofs

Proof of Lemma 1. If the manager stays, his type is revealed. In this case, if $\theta = 1$, the manager is assigned to project α ; if instead $\theta = 0$, he is assigned to project β . If the manager moves, the project allocation depends on the realized period-1 payoff. If the period-1 project was successful ($y_{ik1} = \tilde{y} + \bar{y} - y$), so that his reputation is $\theta_H > p$, the manager is assigned to project α , because the assumption $\eta > 1 - p$ in (6) implies $\eta > 1 - \theta_H$. If the period-1 project failed ($y_{ik1} = \tilde{y} + \bar{y} - y - c$), so that the manager's reputation is θ_L , he is assigned to project α only if $\eta \geq 1 - \theta_L$. ■

Proof of Proposition 2. Comparing (15) with (17), it is immediate that the manager does not move if $\eta < 1 - \theta_L$. This happens because the payoff in the good state (which happens with probability p) is strictly lower if the manager moves, while the payoff in the bad state (which happens with probability $1 - p$) is the same.

The manager moves only if $\eta \geq 1 - \theta_L$ and the expected utility in (16) exceeds the expected utility in (15). This second condition, as stated in the Proposition, is more stringent than the first: to see this, consider that the right-hand side of (18) is positive, $u(\cdot)$ being an increasing function; but then the left-hand expression is also positive, which requires that $\bar{y} - (1 - \theta_L)c > y$, or equivalently to $\eta > 1 - \theta_L$. ■

Proof of Proposition 3. We assume that $\eta \geq 1 - \theta_L$, since otherwise the manager never moves, by Proposition 2. If the manager moves his period-2 wage is

$$w_M = \begin{cases} w_{MH} = \bar{y} - (1 - \lambda - (1 - \lambda)p)c & \text{if } y_{i1} = \tilde{y} + \bar{y} - y, \\ w_{ML} = \bar{y} - (1 - (1 - \lambda)p)c & \text{if } y_{i1} = \tilde{y} + \bar{y} - y - c, \end{cases} \quad (40)$$

where we substituted θ_H and θ_L from (9), the subscript M stands for “moving”, and the subscripts H and L refer to the high and low payoffs of the period-1 project α , respectively (while we have dropped the time and the manager's subscripts to simplify notation). If instead the manager stays, his period-2 wage is

$$w_S = \begin{cases} w_{SH} = \bar{y} & \text{if } y_{i1} = \tilde{y} + \bar{y} - y, \\ w_{SL} = y & \text{if } y_{i1} = \tilde{y} + \bar{y} - y - c, \end{cases} \quad (41)$$

where the subscript S stands for “staying”, and the subscripts H and L are defined as in the previous expression.

(i) To show that by moving a manager receives a payoff with lower expected value and lower variance, it is convenient to introduce the following notation for the expected value and the variance of period-2 wage when moving and staying:

$$\bar{w}_M \equiv pw_{MH} + (1-p)w_{ML}, \quad \sigma_M^2 \equiv p(w_{MH} - \bar{w}_M)^2 + (1-p)(w_{ML} - \bar{w}_M)^2,$$

$$\bar{w}_S \equiv pw_{SH} + (1-p)w_{SL}, \quad \sigma_S^2 \equiv p(w_{SH} - \bar{w}_S)^2 + (1-p)(w_{SL} - \bar{w}_S)^2.$$

If the manager moves, the expected payoff of the project is $\bar{w}_M = \bar{y} - (1-p)c$, as can be seen by using (9) in (16), while it is $\bar{w}_S = p\bar{y} + (1-p)y$ if he stays with his former employer: the difference between these two expressions is $\bar{w}_M - \bar{w}_S = (1-p)(\bar{y} - c - y) < 0$ by assumption (5). The absolute value of this difference is increasing in the frequency of bad managers, $1-p$. To establish that $\sigma_M^2 < \sigma_S^2$, it is sufficient to show that $w_{MH} - \bar{w}_M < w_{SH} - \bar{w}_S$ and that $\bar{w}_M - w_{ML} < \bar{w}_S - w_{SL}$. The first inequality can be rewritten as $(1-p)(\lambda - \eta)c < 0$, and the second as $p(\lambda - \eta)c < 0$, and both of these inequalities hold since we are assuming $\eta \geq 1 - \theta_L = (1-p) + \lambda p > \lambda$.

(ii) To perform comparative statics, let us denote by Δ the gain from moving, i.e., the change in the expected continuation utility when moving (16) relative to staying (15):

$$\begin{aligned} \Delta &\equiv \text{E}[u(w^M)] - \text{E}[u(w^S)] \\ &= (1-p)[u(\bar{y} - (1-\theta_L)c) - u(y)] - p[u(\bar{y}) - u(\bar{y} - (1-\theta_H)c)]. \end{aligned} \quad (42)$$

To show that Δ is increasing in $\eta \equiv (\bar{y} - y)/c$, notice that it is increasing in \bar{y} and decreasing in y and c :

$$\frac{\partial \Delta}{\partial \bar{y}} = p[u'(w_{MH}) - u'(w_{SH})] + (1-p)u'(w_{ML}) > 0,$$

since $w_{MH} < w_{SH}$;

$$\frac{\partial \Delta}{\partial y} = -(1-p)u'(w_{SL}) < 0;$$

and

$$\frac{\partial \Delta}{\partial c} = - [pu'(w_{MH})(1-\lambda) + (1-(1-\lambda)p)u'(w_{ML})](1-p) < 0.$$

Moreover, Δ is decreasing in λ because

$$\frac{\partial \Delta}{\partial \lambda} = [u'(w_{MH}) - u'(w_{ML})]p(1-p)c < 0,$$

since $w_{MH} > w_{ML}$.

(iii) Finally, we wish to identify whether the gain from moving Δ in equation (42) is increasing in the manager's risk aversion. Using the mean value theorem (as Δ is a continuous function),

$$\begin{aligned} \Delta &= (1-p)u'(w_1)(\eta-1+\theta_L)c - pu'(w_2)(1-\theta_H)c \\ &= c[(1-p)u'(w_1)(\eta-1+(1-\lambda)p) - pu'(w_2)(1-p)(1-\lambda)] \\ &= c(u'(w_1) - u'(w_2))p(1-p)(1-\lambda) - c(1-p)(1-\eta)u'(w_1) \end{aligned}$$

where $w_1 \in [y, \bar{y} - (1-\theta_L)c]$, $w_2 \in [\bar{y} - (1-\theta_H)c, \bar{y}]$, and therefore $u'(w_1) > u'(w_2)$ in case of risk-averse managers.

Therefore,

$$\text{Sign}(\Delta) = \text{Sign} \left(\frac{u'(w_1) - u'(w_2)}{u'(w_1)} p(1-\lambda) - (1-\eta) \right).$$

This expression is increasing in $\frac{u'(w_1) - u'(w_2)}{u'(w_1)}$, which is itself increasing the risk aversion, as it is the slope of the marginal utility. ■

Proof of Lemma 2. To show that $V(\theta)$ is increasing in θ , note that it increases both the manager's future wages and the conditional probabilities of good outcomes relative to bad ones. First, each of the instantaneous utility functions $u(w_t)$ in expression (20) is increasing in w_t , and the wages w_t are increasing in the manager's reputation θ_t by (24). Second, an increase in θ_{t-1} raises reputation at each future date $\{\theta_t, \theta_{t+1}, \theta_{t+2}, \dots\}$, by expression (22). Hence, it increases the conditional probability of the good outcome \bar{y} and decreases that of the bad outcome $\bar{y} - c$, which raises expected utility.

To show that $V(\theta)$ is bounded above by V_H , notice that, by expressions (20) and (27), $V(\theta_{t-1}) - V_H$ can be written as follows:

$$\begin{aligned} V_H - V(\theta_{t-1}) &= [u(\bar{y}) - u(w_t)] + \rho \max\{(1 - \theta_{t-1})(V_H - V_L), \\ &\quad [\lambda\theta_{t-1} + (1 - \lambda)p]E \left[\sum_{s=0}^{\infty} \rho^s [u(\bar{y}) - u(w_{t+1+s})] |\theta_t^U \right] \right. \\ &\quad \left. + [1 - \lambda\theta_{t-1} - (1 - \lambda)p]E \left[\sum_{s=0}^{\infty} \rho^s [u(\bar{y}) - u(w_{t+1+s})] |\theta_t^D \right] \right\}. \end{aligned}$$

This expression is strictly positive, as each of the differences that enter it is strictly positive, recalling expression (24) for the wages. Symmetrically, $V(\theta_{t-1}) - V_L$ can be written as follows:

$$\begin{aligned} V(\theta_{t-1}) - V_L &= [u(w_t) - u(y)] + \rho \max\{\theta_{t-1}(V_H - V_L), \\ &\quad [\lambda\theta_{t-1} + (1 - \lambda)p]E \left[\sum_{s=0}^{\infty} \rho^s [u(w_{t+1+s}) - u(y)] |\theta_t^U \right] \right. \\ &\quad \left. + [1 - \lambda\theta_{t-1} - (1 - \lambda)p]E \left[\sum_{s=0}^{\infty} \rho^s [u(w_{t+1+s}) - u(y)] |\theta_t^D \right] \right\}. \end{aligned}$$

Also this expression is strictly positive, as each of the differences entering it is strictly positive, again using expression (24) for the wages. ■

Proof of Proposition 5. Consider first a value $\theta_{t-1} < 1 - \eta$ so low that the manager will be assigned to the β project if he moves. In this case, his continuation utility from moving is V_L , while his utility from staying is $V_L + \theta_{t-1}(V_H - V_L) > V_L$, so that the manager chooses to stay. This suggests that there is a value $\theta = \underline{\theta}$ that makes the manager indifferent between staying and moving. This value is such that $\frac{\theta(1-\delta^-)}{1-\underline{\theta}\delta^-} < 1 - \eta < \frac{\theta(1-\delta^+)}{1-\underline{\theta}\delta^+}$, so that, if he moves, the manager will be assigned to the β project if his output in period t is low and the α project if his output is high. The associated continuation utility is V_L in the former case and $V\left(\frac{\theta(1-\delta^+)}{1-\underline{\theta}\delta^+}\right) > V_L$ in the latter case. Hence, the value $\underline{\theta}$ at which the manager is indifferent between moving and staying is such that $\underline{\theta}(V_H - V_L) = [\lambda\underline{\theta} + (1 - \lambda)p][V\left(\frac{\theta(1-\delta^+)}{1-\underline{\theta}\delta^+}\right) - V_L]$. Notice that $\underline{\theta} > 0$, as $V\left(\frac{\theta(1-\delta^+)}{1-\underline{\theta}\delta^+}\right) - V_L > 0$; and $\underline{\theta} < p$, since $\underline{\theta}(V_H - V_L) < [\lambda\underline{\theta} + (1 - \lambda)p](V_H - V_L)$

and $V\left(\frac{\theta(1-\delta^+)}{1-\theta\delta^+}\right) < V_H$. Because $V(\theta)$ is increasing in θ , for all $\theta_{t-1} < \underline{\theta}$, staying dominates moving.

Consider next managers with high reputation. There is a reputation level $\theta_{t-1} = \bar{\theta} \leq 1$ such that the manager is indifferent between staying and leaving. For $\theta > \bar{\theta}$, because $V(\theta)$ is increasing in θ , staying dominates moving and thus $V\left(\frac{\bar{\theta}(1-\delta^+)}{1-\bar{\theta}\delta^+}\right) = V_L + \frac{\bar{\theta}(1+\delta^+)}{1+\bar{\theta}\delta^+}(V_H - V_L)$. Hence, $\bar{\theta}$ is such that $\bar{\theta}(V_H - V_L) = [\lambda\bar{\theta} + (1-\lambda)p]\frac{\bar{\theta}(1+\delta^+)}{1+\bar{\theta}\delta^+}(V_H - V_L) + [1 - \lambda\bar{\theta} - (1-\lambda)p][V\left(\frac{\bar{\theta}(1-\delta^-)}{1-\bar{\theta}\delta^-}\right) - V_L]$. As the right-hand side of this equation is strictly positive and the equation is trivially met when $\bar{\theta} = 1$, it follows that $\bar{\theta} \in (0, 1]$.

Combining the two above results, churning occurs at time t if $\theta_{t-1} \in [\underline{\theta}, \bar{\theta}]$, provided churning occurred in every previous period. For this to happen, churning must occur in the first period. Since, $\theta_0 = p$, this happens if and only if $\bar{\theta} \geq p$, which also implies that $\bar{\theta} > \underline{\theta}$. If instead $\bar{\theta} < p$, churning never occurs. ■

Proof of Proposition 7. Comparing (29) with (31), it is immediate that the manager does not move if $\eta < 1 - \theta_L$. This happens because the payoff in the good state (which occurs with probability p) is strictly lower if the manager moves, while the payoff in the bad state (which occurs with probability $1 - p$) is the same.

The manager moves only if $\eta \geq 1 - \theta_L$ and the expected utility in (30) exceeds the expected utility in (29). The second condition simplifies to the one stated in the Proposition, which implies $\eta > 1 - \theta_L$ by the same argument used in the proof of Proposition 2. ■

Proof of Proposition 9. As before, the manager does not move if $\eta < 1 - \theta_L$. If instead $\eta \geq 1 - \theta_L$ and $s < pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(y)$, the manager moves only if (33) exceeds the expected utility in (15). The latter condition simplifies to the one stated in the Proposition, which implies $\eta > 1 - \theta_L$ by the same argument used in the proof of Proposition 2. ■

Proof of Proposition 11. In case of a good match, the expected utility from

staying is $pu(\bar{y}) + (1 - p)u(y)$, while the expected utility from moving is $pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\max(y, \bar{y} - (1 - \theta_L)c))$, as in the baseline model. The choice whether to stay or to move is then as characterized in Proposition 2. In case of a mismatch, the expected utility from staying is $u(y)$, while the expected utility from moving is $u(p\bar{y} + (1 - p)(\bar{y} - c))$ because the manager is offered a new start and full insurance. Since by assumption (5), $p\bar{y} + (1 - p)(\bar{y} - c) > y$, the manager always moves in case of a mismatch.

To compare the cases with and without ex-post competition, notice that the first-period expected utility is the same in the two cases. Following from the analysis in the previous paragraph, the second-period expected utility in the competitive equilibrium is

$$m \max\{pu(\bar{y} - (1 - \theta_H)c) + (1 - p)u(\max(y, \bar{y} - (1 - \theta_L)c)), \\ pu(\bar{y}) + (1 - p)u(y)\} + (1 - m)u(p\bar{y} + (1 - p)(\bar{y} - c)).$$

This expression exceeds the second-period utility in the equilibrium without ex-post competition, which is $u(mp\bar{y} + (1 - mp)y)$ as shown in equation (39), for m sufficiently close to 0. The opposite happens in a neighborhood of $m = 1$, which represents the baseline case analyzed in the previous sections. ■

References

- Acharya, V., and P. Volpin. 2010. Corporate Governance Externalities. *Review of Finance*, 14:1–33.
- Axelson, U., and P. Bond. 2015. Wall Street Occupations. *Journal of Finance*, 70:1949–1996.
- Bank of England Prudential Regulation Authority. 2016. Buy-outs of Variable Remuneration. Consultation Paper CP2/16, London.
- Bell, B. D., and J. Van Reenen. 2013. Extreme wage inequality: pay at the very top. *American Economic Review* 103:153–157.
- Bénabou, R., and J. Tirole. 2016. Bonus Culture: Competitive Pay, Screening, and Multitasking. *Journal of Political Economy*, 124:305–370.
- Berger, A. N., N. H. Miller, M. A. Petersen, R. G. Rajan, and J. C. Stein. 2005. Does function follow organizational form? Evidence from the lending practices of large and small banks. *Journal of Financial Economics* 76:237–269.
- Dicks, D. 2012. Executive Compensation and the Role for Corporate Governance Regulation. *Review of Financial Studies*, 25:1971–2004.
- Frydman, C. 2007. Rising Through the Ranks: The Evolution of the Market for Corporate Executives, 1936-2003. Working Paper, Boston University.
- Gabaix, X., and A. Landier. 2008. Why Has CEO Pay Increased So Much? *Quarterly Journal of Economics* 123:49–100.
- Harris, M., and B. Holmstrom. 1982. A Theory of Wage Dynamics. *Review of Economic Studies* 49:315–333.
- Hirshleifer, J. 1971. The Private and Social Value of Information and the Reward to Inventive Activity. *American Economic Review* 61:561–574.

- Huson, M. R., R. Parrino, and L. T. Starks. 2001. Internal Monitoring Mechanisms and CEO Turnover: A Long-Term Perspective. *Journal of Finance* 56:2265–2297.
- Jovanovic, B. 1979. Job Matching and the Theory of Turnover. *Journal of Political Economy* 87:972–990.
- Makarov, I., and G. Plantin. 2015. Rewarding Trading Skills Without Inducing Gambling. *Journal of Finance* 70:925–962.
- Morrison, A. D., and W. J. Wilhelm Jr. 2008 The Demise of Investment Banking Partnerships: Theory and Evidence. *Journal of Finance* 63:311–350.
- Petersen, M. A., and R. G. Rajan. 2002. Does Distance Still Matter? The Information revolution in Small Business Lending. *Journal of Finance* 57:2533–2570.
- Philippon, T., and A. Reshef. 2012. Wages and Human Capital in the U.S. Finance Industry: 1909–2006. *Quarterly Journal of Economics* 127:1551–1609.
- Rajan, R. 2005. Has Financial Development Made the World Riskier? *Proceedings, Federal Reserve Bank of Kansas City*, August:313–369.
- Richardson, M., and I. Walter. 2009. Rethinking Compensation Practices in Financial Firms. Chapter 8 in Acharya V. and M. Richardson, eds., *Restoring Financial Stability: How to Repair a Failed System*, John Wiley & Sons.
- Smith, R. 2009. Greed is Good. *Wall Street Journal*, February 7.
- Tett, G. 2009. What Bankers Can Learn From Chelsea Football Club. *Financial Times*, September 11.
- Thanassoulis, J. 2012. The Case of Intervening in Bankers’ Pay. *Journal of Finance* 67:849–895.

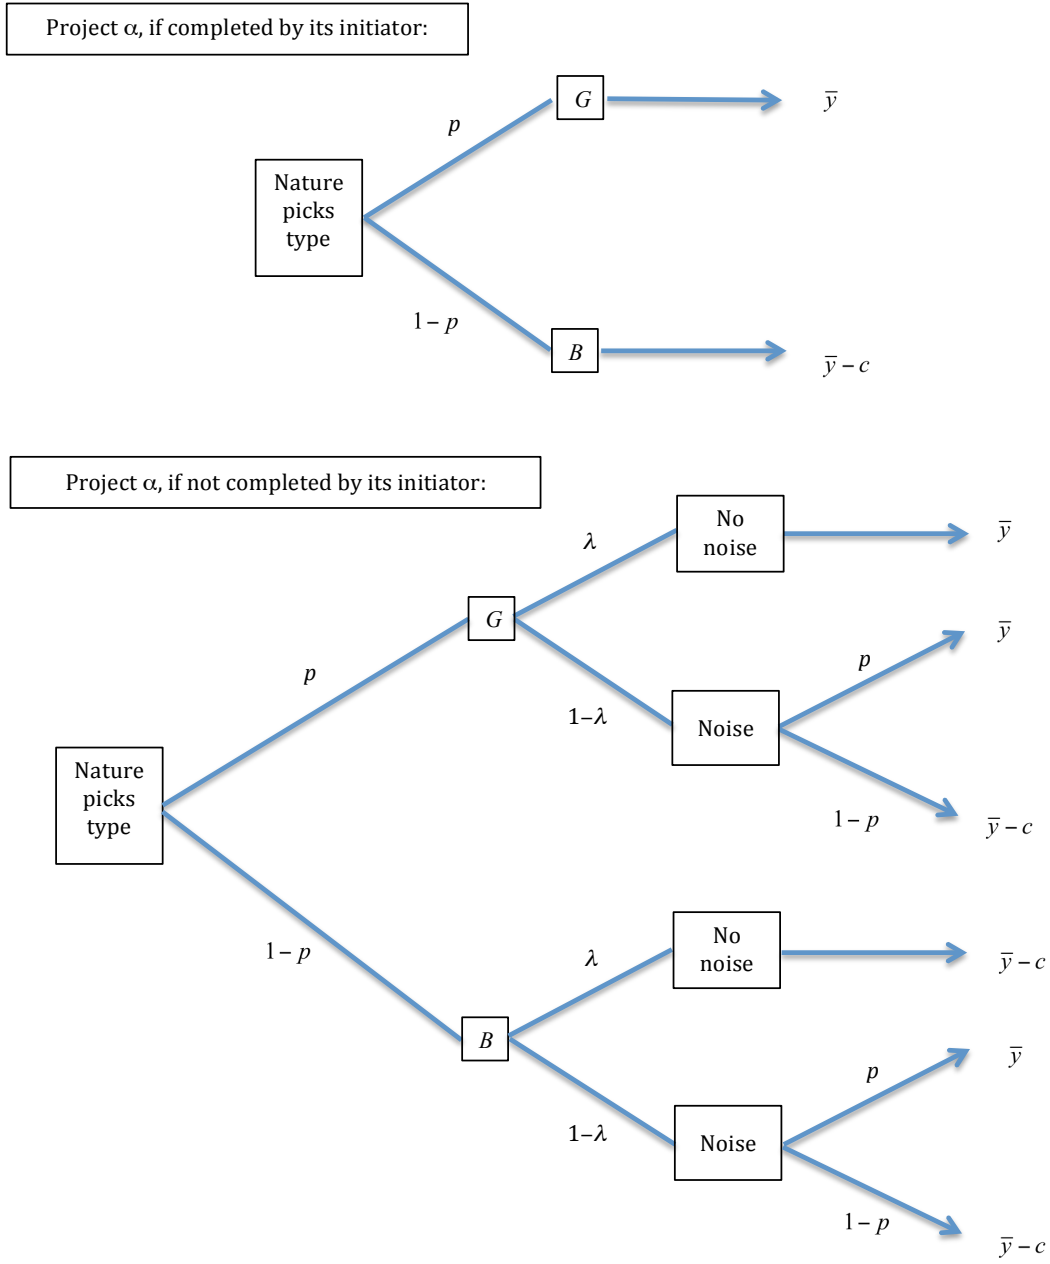


Figure 1: Expected payoffs of project α




Figure 2: Time line of the 2-period model




Figure 3: Moving decision and risk aversion in the 2-period model




Figure 4: Time line of the infinite-horizon model