DISCUSSION PAPER SERIES

No. 8905

SEEKING ALPHA: EXCESS RISK TAKING AND COMPETITION FOR MANAGERIAL TALENT

Viral V Acharya, Marco Pagano
and Paolo Volpin

FINANCIAL ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP8905.asp**

ISSN 0265-8003

SEEKING ALPHA: EXCESS RISK TAKING AND COMPETITION FOR MANAGERIAL TALENT

Viral V Acharya, New York University and CEPR
Marco Pagano, Università di Napoli Federico II
Paolo Volpin, London Business School

Discussion Paper No. 8905
March 2012

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 8905

March 2012

ABSTRACT

Seeking Alpha:
Excess Risk Taking and Competition for Managerial Talent*

We present a model of labor market equilibrium in which managers are risk-averse, managerial talent (‘alpha’) is scarce, and firms seek alpha, that is, compete for this talent. When managers are not mobile across firms, firms provide efficient long-term compensation, which allows for learning about managerial talent and insures low-quality managers. In contrast, when managers can move across firms, high-quality managers can fully extract the rents arising from their skill, which prevents firms from providing co-insurance among their employees. In anticipation, risk-averse managers may churn across firms before their performance is fully learnt and thereby prevent their efficient choice of projects. The result is excessive risk-taking with pay for short-term performance and build up of long-term risks. We conclude with analysis of policies to address the resulting inefficiency in firms' compensation.

JEL Classification: D62, G32, G38 and J33
Keywords: executive compensation, managerial talent, managerial turnover, short-termism

Viral V Acharya
Stern School of Business
New York University
44 West 4th St., Suite 9-84
New York, NY - 10012
USA

Email: vacharya@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=147901

Marco Pagano
Department of Economics
Università di Napoli Federico II
Via Cinthia 45
80126 Napoli
ITALY

Email: mrpagano@tin.it

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=104123

Paolo Volpin
Institute of Finance and Accounting
London Business School
Sussex Place
Regent's Park
London
NW1 4SA

Email: pvolpin@london.edu

* We are grateful to Ulf Axelson, Philip Bond, Mike Burkart, David Dicks, Michael Fishman, Hendrik Hakesen, Giancarlo Spagnolo and participants in the 2010 AEA 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, and in seminars at Arizona State University, Bocconi University, Boston College, INSEAD, New York University for helpful comments and discussions. We acknowledge financial support from the Fondation Banque de France, Europlace Institute of Finance, Inquire Europe and BNP Paribas Hedge Fund Center at HEC Paris.

Submitted 1 January 2012

"*The dirty secret of bank bonuses is that these practices have arisen not merely due to a culture of arrogance; the more pernicious problem is a sense of insecurity. 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. ... [A]s one banker says: "These bonuses are crazy - we all know that. But we don't know how to stop paying them without losing our best staff." Against that background, what the members of some compensation committees are quietly starting to conclude is that the only real solution is to start clamping down on the whole "transfer" game. "If Fifa can stop clubs poaching other players and ripping up contracts, then why can't the banks do the same?" asks one... It is time, in other words, for bankers and regulators to take a leaf out of football's book and start debating not just the issue of pay, but also the poaching culture that is at the root of those huge bonus figures.*" – Tett (2009)

1 Introduction

Excess risk-taking by financial institutions and overly generous managerial pay are regarded by many as key factors contributing to the 2007-09 crisis.[1] In particular, it has become commonplace to blame banks and securities companies for offering compensation packages that reward managers (and more generally, other risk-takers such as traders and salesmen) generously for undertaking investments with high returns in the short run but with large "tail risks" that emerge only in the long run. As governments have been forced to rescue failing financial institutions, politicians and the media have stressed the need that managerial pay packages be cut and incentive systems based on options and bonuses be reined in, made more sensitive to long-term performance, and in some extreme cases be outright eliminated.[2] It is natural to ask whether these limitations to managerial pay are the right policy response to the problem. Indeed, it is crucial to ask what is the root of the problem, that is, which market

[1] See, for example, Rajan (2005, 2008), Richardson and Walter (2009), although there is less than perfect agreement on the effect of managerial compensation on risk-taking (see Section 2).

[2] For instance, the 2008 German bailout plan required banks accepting state aid to cap annual salaries of their executives at €500,000, and to forgo bonuses and dividend payments. Similarly, in early 2009 the U.S. government capped at $500,000 the pay of top executives at companies that received significant federal assistance. Also in the U.K., Sweden and Switzerland, governments set limits on financiers' compensation in their efforts to rescue their banking systems.

failure in compensation practices has led to rewards for short-term performance at the expense of a build-up of tail risks.

The argument that we explore in this paper is that the root of the problem is the difficulty of rewarding managerial talent when managers can pick projects with long-term or tail risk and the market allows them to move across firms before that risk materializes. For instance, a trader in a financial firm can set up a "carry trade" and leave the firm before it is known whether the carry represented an arbitrage opportunity or simply reward for risk (so that the trade eventually "blows up"). In this situation managers have an incentive to take tail risks in order to raise their short-term performance and pay, while moving rapidly across firms, reducing their effective tenure at any firm and thereby the extent to which they can be held responsible for project failures. When such job churning is possible, competition for managerial talent induces a negative externality, insofar as each firm provides an "escape route" to the managers of others. This is to be contrasted with the case where the market for managerial talent is not very competitive, so that managers are more likely to be stuck with their initial employer and held responsible for project failures.

More specifically, we consider a setting in which managers are risk-averse and firms compete for scarce managerial talent. We model managerial talent as "alpha", that is, the ability to generate high returns without incurring high risks: lacking such talent, managers can generate high returns only by taking correspondingly high risks. However, risk emerges only in the long run. So managerial talent can be identified only if managers who have chosen potentially risky projects remain for a sufficiently long period of time with their initial employers: if they leave earlier, the long-term performance of the projects that they have initiated is never learnt, because it is more efficient for the firm to liquidate them.

In this setting, if managers were tied to their initial employer, then over time firms could tell apart the talented from those which are not, and could also insure managers against the risk of finding out that they are not talented. So there would be two efficiency gains. First, a gain in the choice of investment projects: once managers learn about their skills, they will pick the project that they are best suited to manage.

Second, there is a risk-sharing gain: managers who are revealed to be low-skills are cross-subsidized at the expense of the talented ones.

However, competition for managers can prevent both of these welfare gains from being fully realized. If firms compete aggressively in the labor market ("seeking alpha"), then managers can leave before the long-term risks that they have incurred materialize. In particular, the managers who are discovered to be the high-alpha type will extract all rents from their firms by generating competitive offers rewarding their alpha. This would prevent firms from subsidizing the other managers. Therefore, if the labor market is competitive, managers face skewed performance rewards before their types are revealed: high-alpha types extract all rents, and low-alpha types get no subsidy. Now, if risky projects have a greater expected return (even when chosen by a manager of unknown quality) than safe ones, then risk-averse managers are driven to choose risky rather than safe projects, get a higher pay than they would from the safe project, but then move to another firm before the risk of their project has materialized. They are then going to replicate the same behavior in other firms. In the aggregate, managers will churn continuously from one firm to the next, choosing risky projects irrespective of their ability to avoid the implied risks. Talented managers will be identified only in the long run: as managers approach the end of their careers, residual risk from being revealed as the low-alpha type declines, so that the demand for insurance through churning wanes.

To summarize, competition in the market for managers generates an inefficiency due to the contractual externality among firms. We show that frictions in the labor market for managers (e.g., search costs) can reduce the inefficiency by lowering managerial churning. Conversely, ease in interim liquidation of assets (e.g., opening up of securitization markets for loans) can exacerbate the inefficiency by inducing greater churning. We also show that if managers are sufficiently risk-averse, then an increase in the tail risk of projects can lead to greater churning. Finally, we show that if projects carry aggregate risk that delays learning about managerial type, then undertaking such projects becomes a way for managers to synthesize insurance (as an alternative to churning).

The financial sector appears to fit our model particularly well since trading and sales skills are highly fungible across firms, inducing them to compete keenly for "alpha". Further, much risk-taking in the financial sector, ranging from making mortgages or holding AAA-rated mortgage-backed securities to selling credit default swaps or longevity insurance, has the flavor of earning a carry (interest or insurance premium) in the short run but with potential long-run risks (default risk or longevity risk). Hence, we bring our analysis to bear on current policy proposals in the financial sector. We show that constraints on compensation deferrals and clawbacks aggravate managerial churning, while an appropriately chosen salary cap restores the employers' ability to cross-subsidize less talented managers at the expense of more talented ones, even in the presence of a competitive managerial market. The same outcome could be obtained by "taxing mobility", namely, charging a sufficiently large tax on the income of managers who switch employers. This would effectively eliminate ex-post competition for managerial talent (so that the tax would not be paid in equilibrium), but may be hard to implement in practice.

Finally, another interesting application of our model is in understanding why business schools adopt the practice of non-disclosure of grades of Masters in Business Administration (MBA) students to their potential employers. Disclosure of grades would lead to far superior job-market outcomes for student with high grades compared to others. If student quality is uncertain at the admission stage, then such job-market risk would induce students to undertake relatively easy course-work so as to ace grades, resulting in delayed revelation of their type. Interestingly, the grade non-disclosure policy is only prevalent at highly-ranked MBA programs and not in most other professional degree programs (Gottlieb and Smetters, 2011). This is consistent with the model, since the financial upside to job-market success is greatest in highly-ranked programs where employers such as top investment banks and consulting companies queue up for "seeking alpha".

The paper is organized as follows. Section 2 discusses the related literature. Section 3 lays out the structure of the model. In Section 4 we solve for the equilibrium. In Section 5 we relax several of assumptions to explore the robustness of our results. In Section 6 we examine the effects of various policy interventions. Section 7 concludes. Proofs are in the Appendix.

2 Related literature

Our results presents a countervailing force to the benefits arising from competitive labor markets through efficient matching. Gabaix and Landier (2008) present matching models à la Rosen (1981) in which the rise in CEO pay is attributed to their scarce talent and its efficient matching to larger firms. In contrast, in our setting competition for talent among firms results in less efficient matching of managers to projects within each firm.

The fact that managerial turnover introduces an externality across firms in setting their compensation can be considered as a corporate governance externality. Such externalities have been formalized in Acharya and Volpin (2009) and Dicks (2009) in models where a firm's corporate governance is a strategic substitute for governance in other firms (as governance lowers a manager's reservation wages), a result supported empirically by Acharya, Gabarro and Volpin (2009).

In contrast to these papers on governance externalities, our focus is on a dynamic setting in which firms need time to learn about their employees and allocate them to proper tasks, but this is hindered by managers' ability to generate offers from other firms before their type is fully learnt. In this sense, it is reminiscent of Harris and Holmstrom (1982), where an employer designs long-term contracts for risk-averse workers with unknown ability, though our focus is on project choice and the endogenous revelation of manager's talent.

Another recent strand of papers studies inefficiencies in managerial compensation in dynamic settings with moral hazard. Axelson and Bond (2009) show that smart workers may be "too hard to manage", because their high outside options make them insensitive to firing incentives. De Marzo, Livdan and Tchistyi (2011) show that in a dynamic moral hazard model limited liability may make it too costly for the firm to restrain managers from taking tail risks. Relatedly, Makarov and Plantin (2010) develop a model of active portfolio management in which fund managers may secretly gamble in order to manipulate their reputation and attract more funds, resulting in trading strategies that expose investors to severe losses. Our analysis differs from these models because in our context excess risk-taking does not arise from dynamic moral hazard, but from inefficiently slow learning of employees' skills.

Our paper is motivated by the anecdotal evidence of trader churning in the financial sector (see Tett, 2009, cited in the introductory quote) and the competitive "search for yield" (which we interpret as "seeking alpha") by financial firms. Rajan (2005) was one of the first to warn about excessive risk taking in financial institutions driven by short-termist pay packages, labeled as "fake alpha" in Rajan (2008). In another thought-provoking piece, Smith (2009) refers to the role of turnover in entrenching the culture of bonus without performance on Wall Street.[3]

However, there is still no full agreement on the role of pay packages in financial firms' risk taking. Fahlenbrach and Stulz (2009) present evidence that bank CEOs lost a significant portion of their stock-based pay and conclude that pay excesses were not the likely cause of risk-taking at financial firms. Bebchuk, Cohen and Spamann (2009) contend this view, by documenting that bank CEOs, including those of Bear Stearns and Lehman Brothers, had paid out to themselves huge payoffs prior to the crisis, greatly in excess of the amounts they lost eventually. So they argue that bank management did benefit from short-term compensation that was not tied to long-run performance, as is the case in our model with managerial churning. Chen, Hong and Scheinkman (2009) also present evidence linking compensation and risk-taking at financial firms in 1992-2008 that is consistent with payouts to top management being tied to short-term risk-taking incentives. None of these papers examines explicitly the role of employee turnover in generating risk-taking incentives, as we do.

[3] An extended quote borrowed from Smith (2009) runs as follows: "*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. Henry Paulson, when he was CEO of Goldman Sachs, once remarked that Wall Street was like other businesses, where 80% of the profits were provided by 20% of the people, but the 20% changed a lot from year to year and market to market. 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.*"

3 Model

There are K profit-maximizing firms (indexed by $k = 1, ..., K$), which live forever and are owned by risk-neutral shareholders, and I risk-averse managers (indexed by $i = 1, ..., I$), each living for T discrete periods. The analysis will focus on a generation of managers who start their career in period $t = 1$ and retire in period $t = T$.

Firms are competitive and maximize their expected profits. Managers maximize their expected utility $U = E\left[u(\overline{W})\right]$, where $u(\cdot)$ is an increasing and concave function of final (period-T) wealth $\overline{W}$. The assumption that managers only care about final wealth not only avoids dealing with intertemporal optimization problems (which are not central to the analysis), but more importantly puts no limits on deferring compensation: payments can be deferred to the end of the employment period, at no cost for the employer. The case with partial deferral or intermediate consumption is discussed as an extension.

Each firm can condition its own compensation package on the manager's resignation date and on the type of projects that he manages during the employment relationship. But the firm does not have recourse to manager's wealth outside the employment contract. In particular, the firm cannot encroach on the compensation that the manager has received or will receive from other employers – a realistic assumption about the legal reach of each employment contract. Managers start their career with no initial wealth and have limited liability. This implies that their total payoff from each employment contract cannot be negative. For simplicity, there is no discounting: the interest rate is normalized to zero.

3.1 Projects and managers

Managers can run one new project per period. Each project is "long term", that is, lasts for two periods. Hence, a manager that works with the same firm for his entire career runs two projects in each period, except in the first and the last period of employment. Not all managers are equally talented: a fraction $p \in (0, 1)$ of managers are good (G) and a fraction $1 - p$ are bad (B). Each manager i initially does not know his own quality $q_i = \{G, B\}$.

Firms are endowed with a continuum of projects of two types:

(i) safe projects S yielding y at the end of the first period and 0 at the end of the second period, irrespective of the ability of the manager in charge of it;

(ii) risky projects R yielding x in the first period and either 0 or $-c$ in the second period, depending on whether they are matched with a good or bad manager.

The dependence of the risky project's revenue on the manager's type can be interpreted as a reflection of his ability in managing the risky project. Good managers add value to a risky project by reducing its risk (for simplicity, to zero, which is the same level of risk as the safe project), without reducing its expected revenue. In this sense, good managers generate "alpha", namely they improve the risk-return tradeoff of the firm that employs them. Conversely, bad managers can generate the same short-run return x but only at the future cost c.[4]

A key assumption is that if a manager initiates a project of type R, his ability becomes known *only if* he remains in charge of it for *both* periods. The assumption that the project's first-period performance is uninformative captures the idea that failure is an infrequent event ("tail risk"), so that it takes time to screen a person's ability to manage a risky project. Indeed, to capture the fact that the wait to ascertain the quality of a match can be considerable, in an extension we generalize the model to the case in which the project may also have uninformative outcomes, so that learning typically requires more than two periods.

By the same token, if a manager leaves after one period, the quality of the project can no longer be gauged. We assume that in this case the project is liquidated, and that in the process the information about the identity of the project's initiator is lost.[5] The reason why incomplete projects are sold is that their in-house completion

[4]Project R can be interpreted as a carry trade. To generate a profit x the trade needs to be closed in time. So the skilled trader chooses the right time to close the trade and incurs no cost in the second period; while the unskilled traded (who has no clue when to close out the trade) incurs a cost c in the second period.

[5]Avoiding such information loss would require an institution that is capable both of (i) pooling

is inefficient: using another manager from the firm to complete an unfinished project would prevent him from starting a new project of his own. In contrast, outside the firm there are managers who can complete the project at zero cost. In other words, within the firm there is a scarce supply of "creative managers" who can initiate new projects, while outside the firm there is abundant supply of "non-creative managers" who can complete them.

If the project is liquidated, it is sold for its expected value $x - (1 - \lambda)c$, where λ denotes the probability that the risky project was initiated by a good manager. We assume that each firm has a large number of managers, so that one can apply the law of large numbers to compute λ: for instance, if the pool of departing managers is representative of the population, then $\lambda = p$.

We assume that

$$x - (1 - p)c > y > x - c. \qquad (1)$$

The left-hand side inequality indicates that the expected revenue of project R exceeds that of project S if the manager is of unknown quality: this captures the idea that accepting greater risk entails higher expected return. The right-hand side inequality indicates that the expected revenue of a safe project exceeds that of a risky one if the manager is known to be bad. The implication of assumption (1) is that it is optimal to assign bad managers only to safe projects, and good ones only to risky projects. Assigning bad managers to risky projects would imply excessive risk-taking.

3.2 Market for managerial talent

In each period, the pool of projects available to a firm includes at least one safe and one risky project per manager. Therefore, managers – not projects – are the scarce factor of production, since only they can start a new project.

Let i denote a generic manager, k a generic firm and t a generic period. At the

information about the identity of the departing manager (obtained from his initial employer) and the eventual performance of the project (from the project's buyer), and (ii) reselling such information to the new employer of the manager. Establishing such an "information broker" would require a unrealistic level of coordination.

beginning of period t, the firm decides whether to make an offer to the manager. The offer consists of a compensation W_{ikt} contingent on manager i's choice of projects $\{P_{ikt}\}_{t=1}^{t=T-1}$ over his employment life, where $P_{ikt} \in \{R, S\}$ if manager i is employed by firm k in period t and $P_{ikt} = 0$ otherwise:

$$W_{ikt} = W\left(\{P_{ikt}\}_{t=1}^{t=T-1}\right),$$

where $W(\cdot)$ is a mapping $(0, R, S)^{T-1} \mapsto \mathbb{R}^+$, 0 indicates that manager i is not working for firm k in period t, and R (S) indicates that manager i undertakes a risky (safe) project for firm k in period t. The only constraints on the firm's choice of compensation are that it must be non-negative ($W(\cdot) \geq 0$ because of managers' limited liability) and feasible ($W(\cdot)$ cannot exceed the revenues generated by manager i in his employment relationship with firm k). To save on notation, we set $W_{ikt} = 0$ when firm k chooses not to make an offer to manager i in period t.

The manager can accept or reject the offer W_{ikt}: let $F_{it} \in \{1, 2, ..., K\}$ denote the employer he works for in period t. Hence, $F_{it} = k$ means that manager i works for firm k in period t.

It is important to notice that firms can precommit to the wage contracts W_{ikt}. As we will see, this precommitment prevents firms from exploiting any informational advantage that they may gain over their competitors by gauging their employees' ability. We also assume that, in offering such long-term wage contracts, firms bid competitively for managers, anticipating their future performance: hence, managers extract all the expected profits that they will generate in their tenure with any employer.

While *ex ante* there is perfect competition for managerial talent, *ex post* switching costs may prevent it: over time, managers may make location- or firm-specific investments or develop location- or firm-specific tastes, so that other firms cannot poach them. To bring out the implications of *ex-post* competition for managerial talent, we will focus initially on the two polar cases where switching costs are either totally absent – the "competitive regime" – or prohibitively high – the "non-competitive regime". In both regimes, managerial performance is assumed to be publicly observable: if a manager's ability becomes known to the current employer, it becomes equally known to outside employers.[6] In an extension, we shall consider an intermediate case where the managerial labor market features some frictions.

In the competitive regime, at the end of each period managers can choose whether to leave their current employer or not. In the non-competitive regime, they cannot leave the initial employer once they have accepted the initial offer. Formally, $F_{it} = F_{it+1} = k$ if manager i employed by firm k in period t chooses to remain there also in period $t+1$, while $F_{it} \neq F_{it+1}$ if the manager leaves firm k at the end of period t. When indifferent between staying with the current firm and leaving, a manager is assumed to stay with his current employer. This tie-breaking assumption can be motivated with the presence of a tiny switching cost even in the competitive regime.

The difference between the non-competitive and the competitive regime may capture, for instance, the changing relationship between bank managers and their employers: in the past, banking used to entail much local knowledge, so that over their careers bank managers developed employer- and location-specific skills; currently, banking is less local, due to technological change and new financial products. In turn, company loyalty has lost appeal in the world of finance, as noticed by Tett in the introductory quote of this paper.

3.3 Time line

A representative manager i lives for $T \geq 2$ periods. Because managers are scarce, in what follows we assume without loss of generality that manager i is employed in all periods. The sequence of actions is as follows:

(i) In period 1, manager i is hired by firm k ($F_{i1} = k$) that pledges to pay him a final compensation W_{ik1}. The manager then chooses a project: $P_{ik1} \in \{R, S\}$. Then

[6] However, note that this assumption is inessential in our context, due to the multiperiod nature of the employment relationship. To see why, suppose that a manager's performance were visible only to his current employer. Then, in the competitive regime a manager who turned out to be good could be hired by an outside employer, who could condition his pay on his subsequent performance. The manager would have the incentive to choose a risky project and remain with the same employer for at least two periods, to allow him to verify that he is good. So even in an opaque labor market, outside offers would be effectively conditioned on the manager's true type, if this has become known to the manager (and current employer).

the manager chooses whether to stay with employer k ($F_{i2} = k$) or to leave ($F_{i2} \neq k$). If he leaves, the project is sold for a price equal to its expected revenues (including any first-period revenue).

(ii) In period 2, two cases can occur: (a) If the manager stays ($F_{i2} = k$) with employer k, he chooses a new project $P_{ik2} \in \{R, S\}$. (b) If the manager leaves ($F_{i2} \neq k$), he is hired by a new firm j ($F_{i2} = j$) that pledges to pay him a final compensation W_{ij2} and chooses a project $P_{ij2} \in \{R, S\}$. In both cases, at the end of period 2 the manager decides whether to stay with the current employer or to leave.

(iii) In any subsequent period from $t = 3$ to $t = T - 2$, the sequence of moves is the same as under (ii) with appropriate change of time indices.

(iv) The sequence of events is also the same in the penultimate period $T - 1$, with the only exception that the manager cannot leave (as he will not be starting a new project in period T).

(v) In period T, the manager will complete the project started in period $T - 1$ and will consume his final wealth, which is the sum of the compensations awarded by the various employers that have hired him: $\overline{W}_i = \sum_{k=1}^{K} \sum_{t=1}^{T-1} W_{ikt}$, where k is a generic firm and the terms inside the sum are zero for any firm k and period t in which either no offer is made or the manager rejects the offer.

3.4 Learning about managers' types

In any period t the employment history of manager i can be summarized by the belief θ_{it} that his type is good ($q_i = G$). Since in our setting information about the manager's quality is symmetric, the belief θ_{it} is shared by all players. At the beginning of his career, the manager's quality is unknown: he is good with probability p or bad with probability $1 - p$. Hence, $\theta_{i0} = p$. In each period t, the belief θ_{it} is updated on the basis of manager i's performance in period t.

As projects last for two periods, there is no updating of beliefs in period 1: $\theta_{i1} = p$. In period 2, there is no change in belief if manager i left his initial employer k ($F_{i2} \neq F_{i1}$) or if he chose the safe project in period 1 ($P_{ik1} = S$), that is, $\theta_{i2} = p$. If instead the manager did not leave his past employer ($F_{i2} = F_{i1}$) and he chose the risky project in period 1 ($P_{ik1} = R$), then the second-period revenue of the initial project reveals his quality: if the total revenue π_{ik1} from the project chosen in period 1 equals x, manager i is revealed to be good and therefore $\theta_{i2} = 1$; if instead $\pi_{ik1} = x - c$, manager i is revealed to be bad, so that $\theta_{i2} = 0$.

Following the same logic, information about the manager's type is updated in all periods $t \geq 3$ as follows:

(i) $\theta_{it} = 0$ if either the manager is already known to be bad ($\theta_{it-1} = 0$) or if his quality was unknown in period $t-1$ ($\theta_{it-1} = p$) but is revealed to be bad in period t, which happens if he remains with his previous employer ($F_{it} = F_{it-1}$) and at $t-1$ had chosen a risky project ($P_{ikt-1} = R$) that produces a low revenue over its lifetime ($\pi_{ikt-1} = x - c$).

(ii) $\theta_{it} = p$ if previously the manager's type was uncertain ($\theta_{it-1} = p$) and in period $t-1$ he chose the safe project ($P_{ikt-1} = S$), or chose the risky one ($P_{ikt-1} = R$) and left his previous employer ($F_{it} \neq F_{it-1}$).

(iii) $\theta_{it} = 1$ if either the manager is already known to be good ($\theta_{it-1} = 1$) or if his quality was unknown in period $t-1$ ($\theta_{it-1} = p$) but is revealed to be good in period t, which happens if he remains with his previous employer ($F_{it} = F_{it-1}$) and at $t-1$ had chosen a risky project ($P_{ikt-1} = R$) that produces a high revenue over its lifetime ($\pi_{ikt-1} = x$).

3.5 Strategies and payoffs

At the start of each period t, firm k offers to any manager i not currently employed in the firm a compensation based on its belief about the manager's quality. This belief is conditional only on information available as of period $t-1$, since the offer is made before period-t revenues are realized. Formally, the firm's strategy is an offer of the compensation schedule $W(\cdot \mid \theta_{it-1})$ to manager i.

The strategy of a generic manager i in period t is a choice of employer and a project. Formally, manager i employed by firm k in period $t-1$ will choose (i) which firm to work for in period t (F_{it}), and (ii) which project $P_{ikt} \in \{R, S\}$ to carry out in that firm, as a function of the belief θ_{it-1} about his quality, so as to maximize the expected utility from his compensation $U(\overline{W}_i \mid \theta_{it-1})$.

The expected revenue that firm k obtains from the project started in period t by manager i is:

$$\pi_{ikt} = \begin{cases} x & \text{if } P_{ikt} = R, F_{it} = F_{it+1} \text{ and } q_i = G, \\ x - c & \text{if } P_{ikt} = R, F_{it} = F_{it+1} \text{ and } q_i = B, \\ x - (1 - \theta_{it})c & \text{if } P_{ikt} = R \text{ and } F_{it} \neq F_{it+1}, \\ y & \text{if } P_{ikt} = S, \\ 0 & \text{if } F_{it} \neq k \end{cases}$$

where the first (second) line corresponds to the case of a good (bad) manager who chooses the risky project and stays with his current employer in periods t and $t + 1$; the third line corresponds to the case of a manager who chooses the risky project and at the end of period t leaves firm k (which sells the project at a price reflecting the belief θ_{it} about its initiator's quality); the fourth line refers to the case of a manager who chooses the safe project; and the fifth line refers to the case where manager i is not employed by firm k in period t.

Hence, the payoff to firm k from hiring manager i in period τ equals the sum of the revenues generated by the manager over his remaining employment career net of the promised wage:

$$\sum_{t=\tau}^{t=T-1} \pi_{ikt} - W_{ik\tau}.$$

Feasibility of the manager's compensation requires this expression to be non-negative.

4 Equilibrium

In this section we solve for the equilibrium in each of the two alternative labor market regimes described in Section 3.2: the competitive and non-competitive regime, respectively. If there is *ex-post* competition, a manager can choose to work in a different firm F_{it} in each period, if he wishes to do so; in contrast, in the non-competitive regime a manager is constrained to remain with his initial employer F_{i1}, so that good managers cannot be poached by outside employers even if their talent has been revealed by their performance with the current employer. Thus, in the competitive regime managers choose both their preferred employer and project in each period; in contrast, in the non-competitive regime they choose their preferred project in every period, and their preferred employer only in the first period.

Recall however that in both regimes firms are assumed to compete for managers *ex ante*: they all bid for managers, and managers choose the highest bid. Even though in equilibrium this drives their expected profits to zero, we make the usual tie-breaking assumption that they prefer to attract as many managers as possible.

Formally, we solve for the perfect Bayesian equilibrium of the game:

(i) in any period τ, firm k chooses $W_{ik\tau}$ to maximize its expected profits from hiring manager i, where the expectation is conditional on the belief $\theta_{i\tau-1}$ about manager i's quality:

$$\max_{W_{ik\tau}} \left[\sum_{t=\tau}^{t=T-1} E\left(\pi_{ikt} \mid \theta_{i\tau-1}\right) - W_{ik\tau} \right] \cdot I_{F_{i\tau}=k}, \tag{2}$$

where $I_{F_{i\tau}=k} = 1$ if $F_{i\tau} = k$ and $I_{F_{i\tau}=k} = 0$ otherwise, taking as given the strategy of the manager and of other firms;

(ii) in any given period t, manager i chooses his employer F_{it} and the project P_{ikt} so as to maximize his expected utility conditional on the belief θ_{it-1}:

$$\max_{F_{it}, P_{ikt}} U\left(\overline{W}_i \mid \theta_{it-1}\right) = E\left[u\left(\sum_{k=1}^{K} \sum_{s=1}^{T-1} W_{iks} \right) \mid \theta_{it-1} \right]. \tag{3}$$

taking as given the firm's strategy;

(iii) beliefs are updated as described in Section 3.4.

This defines the equilibrium for the competitive regime. The equilibrium for the non-competitive regime differs from this only because the firm's problem (2) and the manager's problem (3) are solved under the additional constraint $F_{it} = F_{i1}$ for all t. In other words, either the firm succeeds in hiring manager i in period 1 ($I_{F_{i1}=k} = 1$) or it never does. Hence, the equilibrium allocation of managers across firms is irrevocably set in period 1, and only the choice of projects can change over time. Since solving for the equilibrium in this case is simpler, in the next section we start from the analysis of the non-competitive regime.

4.1 Non-competitive regime

When there is no ex-post mobility of managers, firm k's problem (2) simplifies to:

$$\max_{W_{ik1}} \left[\sum_{t=1}^{t=T-1} E\left(\pi_{ikt} \mid p\right) - W_{ik1} \right] \cdot I_{F_{i1}=k}, \tag{4}$$

because the hiring decision is done only in period 1, where the belief $\theta_{i0} = p$ is based on the unconditional distribution of managers' quality. Due to *ex-ante* competition, the solution to this problem is simply

$$W_{ik1} = \sum_{t=1}^{t=T-1} E\left(\pi_{ikt} \mid p\right). \tag{5}$$

Hence, the equilibrium lifetime wage of manager i is the revenue that he is expected to generate over his entire career at firm k. By their symmetry, all firms pay an identical lifetime wage, implying that managers are indifferent between them. Moreover, managers are perfectly insured against the risk arising from their unknown quality: equation (5) implies that good managers subsidize bad ones.

Given this result and considering that manager i will be employed by the same firm k throughout his career, the problem (3) simplifies to

$$\max_{P_{ikt}} E\left[u\left(\sum_{t=1}^{t=T-1} \pi_{ikt} \mid \theta_{it-1} \right) \right]. \tag{6}$$

In other words, in each period t manager i chooses projects so as to maximize their expected revenue, conditional on the belief θ_{it-1} about his quality as of period $t-1$.

To solve problem (6), the manager must learn his own quality as early as possible, by choosing the risky project in period 1. In period 2 he will not have learnt his quality yet, but by assumption 1 he will still want to choose the risky project. From period 3 onwards, he will able to condition project choice on his true quality: he must choose only risky projects if he discovers to be good, and only safe projects otherwise.

Under this policy, over his career the manager will generate revenues

$$\Pi^* = 2\left[x - (1-p)c\right] + (T-3)\left[px + (1-p)y\right]. \tag{7}$$

The first term in (7) is the expected period-1 and period-2 profits from the risky project undertaken at $t = 0$ and $t = 1$ by a manager of unknown quality (because it takes two periods to learn his type, the manager's quality is still unknown at $t = 1$, so that assigning him to the risky project yields the highest profit by assumption 1); while the second term is the sum of the expected continuation revenues of the two types of managers in periods 3 through T, weighted by their respective frequencies.

This equilibrium outcome coincides with the first best: it features both (i) optimal risk-sharing, that is, complete insurance of managers by firms (as the latter are risk neutral) and (ii) productive efficiency, that is, optimal choice of projects conditional on managers' quality. So in the non-competitive regime, the managers' equilibrium final wealth is $\overline{W} = \Pi^*$ and their utility is

$$U^* = u(\Pi^*), \tag{8}$$

while firms earn zero expected profits.

This argument establishes the following result:

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

Note that optimal risk-sharing requires the firm not to condition the salary on the quality of the employees, even though this information is used in the matching of managerial talent to projects. This implies that good managers subsidize bad ones: this cross-subsidy is feasible only because in the non-competitive regime good managers cannot leave the company to get higher pay at other firms. Indeed this cross-subsidization breaks down in the competitive regime, to which we turn next.

4.2 Competitive market for managers

When there is *ex-post* competition for managerial talent, the first-best allocation characterized above may no longer be an equilibrium. The key observation is that competition changes the outside options for managers who choose the risky project and remain at least two periods with an employer: since in this case outside employers can infer the manager's ability, they will bid up to x the per-period compensation of good managers, and offer y to bad ones. From expression (7), it is immediate that the first-best compensation per period, $\Pi^*/(T-1)$, is smaller than x and greater than y: hence, if a firm were to offer this compensation, its good managers would leave, while the bad ones would stay. Hence, paying Π^*would entail losses, and the cross-subsidization required to provide optimal risk-sharing would become infeasible.

However, the initial employer may offer a contract that still provides optimal risk sharing and deters managerial mobility by imposing a penalty on good managers if they leave the firm. The most effective such contract is one that makes the entire date-T compensation Π^* contingent on the manager never leaving the firm: the firm will pay nothing if the manager leaves at any time in his career.[7] Formally, at any time $\tau \in (1, ..., T-1)$ firm k offers the following contract to manager i:

$$W_{ik\tau} = \begin{cases} \sum_{t=\tau}^{t=T-1} E\left(\pi_{ikt} \mid \theta_{i\tau-1}\right) & \text{if } F_{it} = k \ \forall t \in (\tau, ..., T-1), \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

Given this contract, a manager who at the beginning of period τ knows to be good ($\theta_{i\tau-1} = 1$) will choose the risky project in all subsequent periods, and will be paid x per period. Similarly, a manager who at beginning of period τ knows to be bad ($\theta_{i\tau-1} = 0$) will subsequently choose the safe project , and will be paid y per period.

Now consider the optimal choice for a manager of unknown type ($\theta_{i\tau-1} = p$): if he never leaves, contract (9) gives him insurance against the risk about his own quality. Moreover, the contract gives him the incentive to choose the risky project at time τ so as to learn his quality, and thereafter choose the risky project if he learns to be good, and the safe one otherwise. Notice that, if the manager plans not to leave the firm, under this contract it is best for him to learn about his quality as early as possible, so as to maximize his compensation. Hence, we can focus on a manager who takes the decision about staying or leaving the employer that hired him in period 1.

[7] Recall that, having zero initial wealth and limited liability, the manager cannot be penalized more than this.

Formally, at time 1 firm k offers the following contract to manager i:

$$W_{ik1} = \begin{cases} \sum_{t=1}^{t=T-1} E\left(\pi_{ikt} \mid p\right) & \text{if } F_{it} = k \ \forall t, \\ 0 & \text{otherwise.} \end{cases} \tag{10}$$

The firm that offers this contract earns zero expected profits only if the manager does not leave the company: if he does, it makes positive profits because it earns the revenues produced by the manager but does not pay him anything. But we must check if the manager who accepts this contract has no incentive to leave.

First of all, notice that if a manager plans to eventually leave the firm, under contract (10) he will want to leave no later than period 2, since staying longer would only increase the penalty for resigning. Second, leaving in period 1 is inefficient, because it entails no learning about the manager's quality, yet it implies a penalty equal to the first period's revenue. Third, a manager who was revealed to be bad in period 2 has no incentive to leave the firm. Hence, we need only to consider a manager revealed to be good in period 2. If he were to stay with the initial firm, his final wealth would be Π^*. If instead he were to leave at the end of period 2, he would earn a final wealth $(T-3)x$ from the new employer, as shown above.

The comparison between $(T-3)x$ and Π^* yields a cutoff value $\widehat{T}$, which defines the maximum career duration that allows the firm to retain its managers through the contract just described:

$$\widehat{T} = 3 + 2\frac{x-(1-p)c}{(1-p)(x-y)}. \tag{11}$$

If $T \leq \widehat{T}$ the first-best allocation can be sustained even in the competitive regime, while if $T > \widehat{T}$ it cannot. Intuitively, if the manager's career duration T is very short, then he must spend a large fraction of his career with an employer just to be recognized as being of good quality and therefore loses a large fraction of his wealth if he chooses to leave. For instance, if his career were to span three periods ($T = 3$), he would lose 2/3 of his lifetime stream of revenue to the initial employer, and only earn 1/3 with the new one. So leaving would not be optimal, as witnessed by the fact that $\widehat{T} > 3$. In this case, the first-best would be feasible.

If instead the manager's career duration is longer, i.e., $T > \widehat{T}$, then contract (10) would not deter the manager from leaving. Intuitively, the penalty for leaving (which is the loss of the revenue produced in periods 1 and 2) is small compared to the gain in later periods. In such case, the first-best would not be feasible.

It is instructive to see how the cutoff value $\widehat{T}$ responds to changes in the other two main parameters of the problem. In Figure 1, we show that an increase in the fraction of good managers, p, expands the range of values of T for which the first-best allocation can be achieved (for instance, for p very close to 1 it can be achieved even for very large T): intuitively, the cost of subsidizing bad managers is quite low because there are few of them. In Figure 2, instead, we see that an increase in the excess profitability of a well-managed risky project over that of a safe one, $x - y$, reduces the range of values of T for which the first-best allocation can be achieved: when these excess profits are large, outside employers can lure away a good manager even if his remaining job tenure is relatively short.

The following proposition summarizes the discussion up to this point:

Proposition 2 (First-best region under competition) *In a competitive managerial market, the first-best outcome can be attained in equilibrium if and only if the manager's career duration is sufficiently short, i.e.* $T \leq \widehat{T}$, *where* $\widehat{T}$ *is defined by (11).*

What happens when the first best cannot be attained, that is, when $T > \widehat{T}$? In this case, contract (9) cannot be offered in equilibrium because managers would leave and firms would make profits. This is inconsistent with equilibrium, because it would lead firms to deviate from contract (9) by offering a higher compensation.

To find the equilibrium, we start by noticing that, due to competition for managers, equilibrium contracts must lead to zero expected profits, conditional on the *current* belief about the manager's quality θ_{it-1}. Formally, at any time $\tau \in (1, ..., T-1)$ firm k offers the following contract to manager i:

$$W_{ik\tau} = \sum_{t=\tau}^{t=T-1} E\left(\pi_{ikt} \mid \theta_{it-1}\right). \qquad (12)$$

Because managers are paid the entire stream of revenues, managers choose the project that maximizes revenue conditional on their beliefs about their quality. Each manager who learn to be good (bad) will choose the risky project (safe), and each manager of uncertain quality will choose the risky project, from assumption (1). Hence, a manager of good quality is paid x per period; while a manager of bad quality is paid y per period. Instead, a manager of unknown quality is paid $x-(1-p)c$.

What still remains to be pinned down to characterize the equilibrium is the managers' choice whether to stay with their initial employer or to leave. We focus on the following candidate equilibrium: the manager changes employer (and chooses project R) in each of the first K periods, earning the expected revenue $x-(1-p)c$ per period, with $K \in [0, T-3[$. From period $K+1$ onwards, he remains with the same employer. Since the manager will optimally choose project R in periods $K+1$ and $K+2$, by period $K+3$ his quality will be known, so that subsequently he will choose project R if good, and project S otherwise. Hence, the manager's problem in (3), upon substituting for the compensation (12) and for the optimal choice of project described above, can be rewritten simply as:

$$\max_{K} pu\left(W_G\right)+(1-p)\,u\left(W_B\right), \tag{13}$$

where

$$W_G \equiv (K+2)\left[x-(1-p)c\right]+(T-3-K)x \tag{14}$$

is the final wealth of a good manager, and

$$W_B \equiv (K+2)\left[x-(1-p)c\right]+(T-3-K)y \tag{15}$$

is the final wealth of a bad manager. Hence, the manager's problem reduces to the choice of K, namely, the number of periods in which he "churns" jobs: churning is a way for the manager to delay the revelation of his type and thus obtain insurance, but this comes at the cost of greater inefficiency, as bad managers should be assigned to the safe project rather than the risky one. Therefore, the trade-off is between insurance, which is obtained by delaying the revelation of the manager's quality (a larger K) and productive efficiency, which comes with earlier revelation (a smaller K). The two polar cases are $K=0$ and $K=T-3$: in the first case, the manager never leaves his initial employer, and thus obtains no insurance (except in periods 1 and 2), but achieves productive efficiency; in the second case, the manager achieves perfect insurance by churning jobs all the time, at the cost of low productive efficiency. The optimal K maximizes expression (13), and is defined implicitly by the first order condition:

$$\frac{u'(W_B)}{u'(W_G)} = \frac{pc}{x - y - (1-p)c}, \tag{16}$$

where W_G and W_B are given by (14) and (15) and the fraction is positive by assumption (1). Intuitively, an increase in K transfers wealth from the state in which the manager is revealed to be good (W_G being decreasing in K) to that in which he is revealed to be bad (W_B being increasing in K). Hence:

Proposition 3 (Churning equilibrium) *In a competitive managerial market, if $T > \widehat{T}$ in equilibrium the manager switches firm in every period for the first K^* periods, and subsequently remains with the same firm, where K^* satisfies condition (16).*

Figure 3 describes the equilibrium in the space (W_G, W_B). Point A on the 45^o line represents the final wealth obtained by churning for $T-3$ periods: in this case the manager obtains the same wealth independently of its type. Point B in the figure represents instead the case in which the manager chooses not to churn. In this case, if his type is good his final wealth (W_G) is much larger than his wealth if his type is bad (W_B). By setting the number of churning periods K between 0 and $T-3$, the manager can choose any point on the segment $\overline{AB}$: this line, whose slope is $-p/(1-p)$, illustrates the extent to which the manager can self-insure by churning. The optimal choice on that line depends on the probability p of being a good type and on the utility function $u(\cdot)$: in particular, it depends on the marginal rate of substitution between the two states of the world (the state in which the type is good and the state in which the type is bad) and thus on the degree of risk aversion of manager. Intuitively, a more risk-averse manager will choose a higher K to smooth consumption more between the two states. As shown in the graph, the solution is the tangency point between the manager's indifference curve and the segment $\overline{AB}$.

4.3 Comparative statics

Proposition 3 yields a testable cross-sectional prediction: all else equal, that is, with same residual uncertainty about type, junior managers are more likely to churn than senior ones, and therefore more likely to be associated with excess risk-taking by firms (indeed, if type uncertainty were greater for juniors, it would only strengthen their incentives to churn). Since in the equilibrium with competition K^* can be taken as a measure of the pervasiveness of churning, it is interesting to investigate how it responds to changes in the parameters of the problem.

To illustrate comparative statics in a simple example where predictions are unambiguous, it is worth focusing on the case where managers have negative exponential (CARA) utility:[8]

Example (Comparative statics in the churning equilibrium: CARA utility) *If managers have CARA utility $u(w) = -e^{-\gamma w}$ (with $\gamma \geq 0$), then the optimal number of churning periods is*

$$K^* = \max\left\{T - 3 - \frac{\log(g)}{x-y}, 0\right\}, \tag{17}$$

where $g \equiv \{pc/[x-y-(1-p)c]\}^{1/\gamma} > 1$. K^* is increasing in the managers' employment horizon T, in the degree of risk-aversion γ, the probability of being a good manager p, and is decreasing in the magnitude of tail risk c.

These results are intuitive. A longer employment horizon T makes the manager more averse to revealing his type, because the implied risk refers to a larger future cash flow, and therefore induces him to churn across firms for a longer interval. By the same token, a more risk-averse manager will seek more insurance, and therefore churn longer. Finally, the demand for insurance decreases in its cost, which is increasing in the tail risk c and in the probability of being a bad manager $1-p$.

---

[8] Expression (17) follows from replacing $u(w) = -e^{-\gamma w}$ in the first-order condition (16) and solving for for K^*. The expression immediately implies that K^* is increasing in T. To establish the other comparative statics results, notice that K^* is decreasing in g, and that in turn g is decreasing in γ and p, and is increasing in c.

In general, comparative statics depend on how risk aversion behaves as a function of wealth. In particular, the response of job churning to a change in tail risk can be characterized as follows:

Proposition 4 (Effect of tail risk on job churning) *The length of equilibrium churning period K^* is decreasing in the tail risk parameter c if the manager's utility function features constant or increasing absolute risk aversion, or constant relative risk aversion equal or less than 1.*

Intuitively, an increase in c raises the cost of obtaining insurance by churning, and this greater cost has both a substitution effect and a wealth effect on the manager's desired level of self-insurance via churning. The substitution effect will lead to a reduction in the demand for insurance (and thus induce a reduction in K^*), but the wealth effect (due to the fact that a larger value of c implies a lower average payoff for the manager) may increase the demand for insurance if risk aversion is increasing in wealth. The proposition identifies cases in which the substitution effect dominates.

However, there are circumstances in which the effect of tail risk on the churning period K^* goes in the opposite direction relative to what is predicted by Proposition 4. This occurs if managers are very risk-averse and if the parameter c is large, so that the associated wealth effect is sizeable. Indeed in Figure 4, where γ is assumed to be equal to 7.5, an increase in tail risk c is initially associated with a shorter churning period K^*, but for sufficiently large c it leads to a longer churning period K^*. (More precisely, on the horizontal axis of Figure 4 the tail risk parameter c is standardized by $(x-y)/(1-p)$, which is the maximum value of c consistent with assumption (1)).[9] In the left-side portion of Figure 4, the substitution effect dominates, so that

[9] If managers have CRRA utility $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$ (with $\gamma \geq 0$), then the optimal number of churning periods can be shown to be

$$K^* = \max\left\{ \frac{(T-3)(x-gy) - 2(g-1)[x-(1-p)c]}{g[x-y-(1-p)c] + (1-p)c}, 0 \right\},$$

where $g \equiv \left[\frac{pc}{x-y-(1-p)c}\right]^{\frac{1}{\gamma}} > 1$. Figure 4 plots this expression for K^*, assuming $x = 10$, $y = 1$, $\gamma = 7.5$, $p = 0.99$ and c ranging between $x-y$ and $(x-y)/(1-p)$, i.e. the bounds defined by assumption (1).

managers churn less as tail risk increases, while in its right-side portion the wealth effect dominates, so that managers actually churn more if tail risk increases.

Hence, if managers are highly risk-averse and projects feature large tail risk, as in the right portion of Figure 4, they may respond to an increase in tail risk by taking more insurance in the form of churning, rather than less. This would exacerbate the inefficiency arising from delayed allocation of good managers to the risky projects. Hence, paradoxically, in a situation where managers are very risk-averse, an increase in the tail risk of projects would lead to more tail-risk seeking.

5 Extensions

In this section we extend the model analyzed so far, in order to investigate how its insights change when some of its key assumptions are modified. In Section 5.1 we consider a setting where risky projects are not always informative, even when managers stay in the same firm for two or more periods: we do so by letting the payoff of risky projects not reflect solely the manager's quality. In Section 5.2 we explore how the model's results change when firms are not allowed to defer all managerial compensation until the end of the employment relationship. Finally, in Section 5.3 we consider how frictions in the labor market or in financial markets affect the extent to which managers wish to churn across firms. Specifically, we consider first an informational friction in the market for managers, arising from the presence of adverse selection, and then a friction arising from search costs in the managerial labor market or from liquidation costs in the market for incomplete projects.

5.1 Project risk as a source of insurance

A stark assumption of the baseline model is that, as soon as a manager's performance is observed for two periods, his type becomes known for sure. We now consider a setting where managers can be lucky for some time, in the sense that their type is not recognized even if they stick with the same employer for two or more periods. This occurs because, besides the risky and safe projects described in the previous section, the manager can now choose a new type of risky project, whose payoff does depend only on the manager's quality, but also on the realization of an aggregate shock. We refer to this as project A, as a mnemonic for "aggregate risk". Formally, the choice of manager i employed by firm k at time t is now $P_{ikt} \in \{A, R, S\}$.

More specifically, the payoff structure of project A is as follows: with probability $1-\beta$ its payoff is just as that of project R in the baseline model (namely, it equals x if the manager is good and $x-c$ if he is bad); with complementary probability β its payoffs are independent of the quality of the manager, and only reflect an economy-wide risk factor that leads its payoff to be high (h) or low (l), with probabilities q and $1-q$ respectively. Note that we assume that employers are able to correctly identify the source of the shock based on its idiosyncratic or aggregate nature, and therefore consider aggregate shocks as uninformative about managers' quality (even if the payoffs determined by the aggregate shock happened to coincide with those determined by the manager's quality, that is, $h = x$ and $l = x - c$). Hence, β captures project A's aggregate risk component. We assume the average payoff of project A in the states of the world where it is driven by the aggregate shock to be the same as in those where it is driven by the manager's quality and is initiated by a manager of average quality:

$$qh + (1-q)l = x - (1-p)c.$$

This ensures that the A project's average payoff is $x-(1-p)c$, the same as that of the R project, and makes the analysis closely comparable to that of the previous section.

From the manager's standpoint, the presence of a project with aggregate risk is a source of insurance, since it delays the recognition of his type without requiring the switch to a new employer. Specifically, consider a manager of unknown quality who has picked the A project: if the payoff happens to be driven by aggregate risk (which happens with probability β), he will not need to switch to a new employer to avoid his type being revealed. In either case, his compensation is $x-(1-p)c$. We break the tie in this choice by assuming that the manager prefers to carry out the project with his current employer, rather than switching to a new firm: this involves an arbitrarily small switching cost, that the manager avoids by continuing the project already undertaken when its payoff is uninformative.

Thus, if a manager of unknown quality has undertaken the A project and wants his type not to be revealed yet, he will stay in the firm if the project's payoff happens to be determined by aggregate risk (with probability β) and will switch to a new firm otherwise (with probability $1-\beta$). Moreover, he will strictly prefer project A to project R, since it affords him some insurance about his type risk without requiring him to move with probability 1 in the subsequent period. He will switch to project R only when he wants his type to be revealed and has decided to stop churning.

Once a manager's type is revealed, good managers will no longer undertake the A project, since in their hands R projects are more profitable, their payoff being always x irrespective of the realization of aggregate risk. Managers who learn to be bad will instead prefer the S project to the A project, provided the safe project's payoff y fulfills a condition more restrictive than (1):

$$x-(1-p)c>y>x-c-\beta pc,$$

which can be restated as an upper bound on the A project's systematic risk component: $\beta<(y+c-x)/pc$. Hence, under these assumptions, managers will undertake A projects for the K periods by which they wish to delay the revelation of their type, sticking with their employer whenever their payoff are uninformative about their quality, and churning otherwise. After K periods, they will switch to project R if they are revealed to be good and to project S if they are revealed to be bad.

This optimal strategy implies the same final wealth W_G and W_B for good and bad managers respectively as in the baseline model, and therefore the same amount of insurance. This equivalence implies that also in this new setting, if the managerial market is competitive and $T>\widehat{T}$, managers will want to delay the revelation of their true quality for K^* periods, where K^* is defined by expression (16). The difference from the baseline model is that now the manager partly synthesizes this insurance using the project with aggregate risk, and partly by churning. Since the frequency with which the payoff of A projects is uninformative is β, the expected number of periods in which managers undertake uninformative projects is βK^*, and the expected number of periods in which managers will churn is $(1-\beta)K^*$. This is summarized in the following proposition:

Proposition 5 (Project-level risk) *In a competitive managerial market with* $T > \widehat{T}$*, if managers (i) can choose projects with aggregate risk whose payoff is uninformative about managers' quality with frequency* β*, and (ii) face a small job switching cost, then on average they churn for* $(1-\beta)K^*$ *periods and undertake uninformative projects for* βK^* *periods, where* K^* *is defined by expression (16).*

This proposition predicts that, if the aggregate risk β exposure of the financial sector increases, managers are less inclined to churn across firms. But the inefficiency in managerial assignment stays unchanged: rather than via managerial mobility, delayed recognition of managers' true skill occurs by aggregate risk-taking.

5.2 Limits to deferring compensation

Recall that an important assumption made in deriving all the results so far is that there are no constraints on withholding compensation to a manager who resigns. In practice, however, this assumption may neglect legal restrictions: it may be illegal to write an employment contract where the manager is denied compensation for past employment because he chooses to switch to a new employer. In practice, at least a portion of the total compensation is paid in the form of salary, to fund intermediate consumption (possibly because otherwise managers would be unable to achieve the desired consumption smoothing due to borrowing constraints).

Limited liability would prevent the initial employer from reclaiming such interim salary payments: hence, limits to deferred compensation reduce the parameter region where the first-best can be attained, compared to the region described in Proposition 2. Intuitively, the more the firm is constrained in deferring compensation, the lower is the penalty that it can threaten to inflict on resigning managers, and therefore the smaller is the parameter region where it can attain the same employees' loyalty as in the non-competitive regime – and offer risk-sharing to them. Specifically, it is easy to show that, if part of the total compensation is paid as non-recoverable per-period salary $\underline{w} > 0$, the maximum career duration for which the first-best outcome can be attained is:

$$\widehat{T}(\underline{w}) = 3 + 2\frac{x - (1-p)c - \underline{w}}{(1-p)(x-y)},$$

which is strictly decreasing in $\underline{w}$.

5.3 Imperfections in the labor or asset markets

In this section we will consider the effects of imperfections in either the labor or the financial markets that (directly or indirectly) increase the cost of churning.

5.3.1 Asymmetric information

Our assumption of symmetric information between firms and managers is critical. If managers knew their type, then in equilibrium no insurance can be obtained through churning: good managers would stay in their initial firm so that they are revealed as good and can enjoy higher pay. Bad managers would then also be revealed and assigned to safe projects from period 2 onwards.

A less extreme case is that where only a fraction ϕ of managers who know their type from the start. In this case, we expect churning to decrease in equilibrium for two reasons: (i) mechanically, the fraction $p\phi$ of managers who know to be good will stick with their initial employer; (ii) managers of unknown type will get pooled with those who know to be bad, and therefore will wish to churn for a shorter period than in the baseline model.

Since by churning a manager of unknown type is pooled with the bad type, the price for an unfinished project is: $x - (1-\widehat{p})c$ where $\widehat{p}$ is the updated probability that the project was started by a bad manager:

$$\widehat{p} = \frac{p(1-\phi)}{(1-\phi)+\phi(1-p)} = p\frac{1-\phi}{1-p\phi} < p.$$

Since $\widehat{p}$ is decreasing in the severity of the asymmetric information ϕ, with $\phi > 0$ the payoff in case of churning decreases from $x - (1-p)c$ to $x - (1-\widehat{p})c$. Mathematically, the manager's problem is identical to the case described in Section 4.3

when we considered the effect of a change in c. As in that case, the conflict between substitution and wealth effect prevents us to sign the effect on K of this case. If we assume CARA utility function, it follows from Proposition 4 that the optimal length of the churning period K^* declines with the severity of the adverse selection problem ϕ, as the cost of insurance increases.

5.3.2 Search costs

Consider next the impact of search costs: when the manager leaves his employer, he must hire a headhunter or be unemployed for some time before finding a new job. If we denote this search cost by s, the payoff in case of churning drops from $x-(1-p)c$ to $x-(1-p)c-s$. Similarly, the market for incomplete projects could be illiquid, in which case firms would have to accept a discount s when selling these projects. This will affect the payoff of the manager in case of churning: his payoff would decrease from $x-(1-p)c$ to $x-(1-p)c-s$ because of the search cost. These two imperfections have a similar effect on the churning equilibrium. The only change in the manager's problem (13) is that now his final wealth is defined as follows:

$$\widehat{W}_G = K\left[x-(1-p)c-s\right] + 2\left[x-(1-p)c\right] + (T-3-K)x$$

for a good manager, and

$$\widehat{W}_B = K\left[x-(1-p)c-s\right] + 2\left[x-(1-p)c\right] + (T-3-K)y$$

for a bad manager.

The optimal churning interval $\widehat{K}$ solves the first order condition:

$$\frac{u'(\widehat{W}_B)}{u'(\widehat{W}_G)} = \frac{p\left[(1-p)c+s\right]}{(1-p)\left[x-y-(1-p)c-s\right]}. \tag{18}$$

Notice that this expression exceeds the right-hand side of equation (16) for any $s > 0$ and is strictly increasing in s. Also note that if $s > x-y-(1-p)c$ then the first order condition (18) cannot hold. Then, the optimal choice is to set $\widehat{K} = 0$. Intuitively, when the search cost is very high, there is no more churning and no more excessive risk taking. In such case managers are better off obtaining insurance by choosing the safe project (instead than churning).

As before, the conflict between substitution and wealth effects prevents us from assessing whether in general $\widehat{K}$ is smaller or greater than K^*. Following the steps in Proposition 4, this ambiguity disappears in the CARA utility case, where:

$$\widehat{K} = \max\left\{T - 3 - \frac{\log(\widehat{g})}{x-y}, 0\right\},$$

and

$$\widehat{g} \equiv \left\{\frac{p\left[(1-p)c+s\right]}{(1-p)\left[x-y-(1-p)c-s\right]}\right\}^{1/\gamma}.$$

Hence, $\widehat{K} < K^*$ and $\partial\widehat{K}/\partial s < 0$: an increase in search costs in their job market leads managers to churn less. The same holds true if the parameter s is interpreted as capturing frictions in the secondary market for projects, such as illiquidity in the market for loan sales or lack of well-developed securitization markets.

However, as we know from Proposition 4, with different utility functions the effect may go in the opposite direction: with constant relative risk aversion, churning may actually increase in response to greater search frictions if managers are highly risk-averse and these frictions are already severe or project tail-risk is already high (a large s being equivalent to a large c). This indicates that frictions are not necessarily stabilizing in the presence of high tail risk.

6 Policy interventions

The model presented in the previous sections highlights that competition for managerial talent induces inefficiencies in two ways: first, it limits risk-sharing opportunities that firms can offer to managers; second, it induces excess risk taking and therefore a loss of productive efficiency. In this section we consider which policy interventions can limit or eliminate these inefficiencies. Such public interventions are warranted by the fact that in our churning equilibrium, no individual bank has the incentive to deviate and unilaterally stop competing for other banks' managers: in the words of the initial quote by Tett (2009), banks "feel utterly trapped", and only the intervention of a public authority (such as FIFA for soccer) can stop banks from poaching employees from each other.

6.1 Clawbacks and long-term indexing

Several recent proposals to reform managerial compensation in financial institutions are based on the idea that it would be desirable to defer ("claw back") a part of the managerial compensation and index this deferred compensation to long-term managerial performance. The idea behind such proposals is to address excess risk-taking. Note that excess risk-taking also arises in our "churning equilibrium". Hence, it is desirable to discourage managers from taking projects that are likely to be highly profitable in the short run but feature "tail risk".

However, in our benchmark setting deferring compensation would be inconsequential. The model places no constraints on deferral of managerial compensation: indeed, in the above analysis compensation is already assumed to be paid at the end of the manager's career. Even in the churning equilibrium, it is inessential whether in each period the employer pays the manager's compensation for that period or defers it to some future date: the essential point is that the compensation cannot be made contingent on the manager's type. In such an equilibrium, long-term indexing would be ineffective, because the past performance of the manager is uninformative about his type (his "true alpha").

It is true instead that anything that constrains the firms' ability to defer compensation is inefficient. As shown by Corollary 3, if for some exogenous reason firms cannot defer compensation entirely and make payments contingent on the employees' loyalty, then the parameter region where the first-best outcome obtains shrinks.

6.2 Salary caps

Another very frequently mentioned policy proposal is to impose a cap on managerial compensation. How would such a policy change the equilibrium in our model with managerial competition? Specifically, would it make churning – and the associated excess risk taking – less attractive to managers?

Suppose that policy-makers were to introduce a salary cap on the per-period compensation of managers, at the first-best level w^*. Such a cap would indeed prevent employers from poaching high-quality managers from each other in the competitive regime, and make the perfect risk-sharing and no-churning outcome sustainable in equilibrium. To see this, consider the candidate equilibrium where each employer offers the wage w^* to all his managers, and assigns them optimally once their type becomes known. Then, due to the salary cap, a competing employer could not poach the managers who have proved to be good from their current employer. Moreover, churning for K periods would not be an equilibrium: in that case, on a per-period basis he would earn utility (13) which is smaller than the first-best utility $u(\Pi^*)$, so that he would not deviate from an employer who offered him w^*.

So a binding price cap would guarantee efficient risk-sharing between employees by shutting down competition for good managers. It would also simultaneously ensure the avoidance of excess risk-taking by firms, since it would discourage managers from churning across firms to avoid revealing their true ability. This highlights that current policy proposals about caps on the pay of top managers of financial institutions may have an efficiency rationale, not just a basis in ethical and political concerns (though this efficiency rationale is yet to be spelled out by those proposing caps). Indeed, according to the model, an appropriately set pay cap would raise the expected utility of managers themselves.

6.3 Taxing mobility or FIFA-style no-compete clauses

An effect similar to that of a salary cap could be achieved by a tax on managerial mobility: suppose that the compensation of a manager who switches to a new employer were taxed at a higher rate than that of a loyal manager. If the tax is set at a sufficiently high rate, it would effectively move the economy to the first-best even if the managerial labor market is competitive, as it would effectively block *ex-post* competition for managerial talent. Such a tax would not be paid in equilibrium, since managers would not switch to other employers. Therefore, the policy prescription from the model is to "throw sand in the wheels" of the managerial labor market.

To see this, consider the equilibrium where each employer pays the first-best compensation Π^* to his managers, and assigns them optimally once their type is revealed. After the first two periods, managers learn their type. Hence, the good manager could leave and obtain utility $u((T-3)\,x)$. As shown in Section 4, this deviation is profitable if $T > \widehat{T}$, where $\widehat{T}$ is given in (11). Then, a tax on mobility

$$\tau \geq u((T-3)\,x) - u(\Pi^*) \tag{19}$$

would prevent this deviation. Notice this condition would also ensure that there is no deviation after the third period because the benefits of deviating in period $H > 2$ (i.e. $u((T-1-H)\,x))$ decreases in H while the cost of deviating (i.e. the loss of $u(\Pi^*)$) does not change. With such a tax on mobility (19), a competing employer could not poach the managers who have proved to be good from their current employer.

Similarly, as discussed in Section 5.3.2, when search costs are sufficiently high (that is, when the search costs $s > x - y - (1-p)c$), then there is no more churning and no more excessive risk taking as managers are better off obtaining insurance by choosing the safe project (instead than churning).

6.4 Investing in "alpha"

As discussed above, both a salary cap and the equivalent tax on managerial mobility would redistribute income from good to bad managers. In the current setting this redistribution prevents managerial churning and facilitates productive allocation of talent. We note, however, that the redistribution could have a negative effect on efficiency in a richer setting in which managers invest in their quality *ex ante* at a private cost – for instance, by taking an MBA (or taking tougher courses in the MBA), they can raise their probability p of being a good manager. In this case, capping their salary (or not revealing MBA grades to employers) would reduce the "average alpha" of managers in equilibrium.

Moreover, in the real world preventing reallocation of managerial talent may have efficiency costs that are not captured by the present model: if both managers *and* firms are heterogeneous, they may both learn gradually about the quality of their match, so that it may be efficient for bad matches to be dissolved and new ones be formed. Also, limiting or preventing managerial mobility may confer market power to firms, and thereby create holdup problems. In our setting, this would be inconsequential because of *ex-ante* competition, but in reality this assumption may not hold either. Such considerations are worthy of further modeling in the context of our setup which focused exclusively on one dark side to managerial mobility.

7 Conclusions

An important economic purpose of the firm is to gather information about its employees' talents and use it to allocate them efficiently to projects. Such efficient allocation of talent is also considered to be the key role of a competitive market for managers (see Gabaix and Landier, 2008, among others). In this paper we show, however, 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. Indeed, this dark side gets exploited by managers as they prefer to take on projects with tail risks and use the labor market to move across firms, delaying the resolution of uncertainty about their talent.

This theoretical contribution is especially suited to understand risk-taking in the financial sector. It also yields several testable implications. An immediate prediction of our model is that there should be a positive correlation between the mobility of managers and traders across financial institutions and their risk-taking. Moreover, according to the model, cross-sectional differences between managers can make some of them more prone to switch jobs than others, for instance, because – in keeping with the model – they are at the start of their careers. It appears promising to pursue empirically such testable implications of our model.

Proofs

Proof of Proposition 3. Since in this setting the only reason for switching employer is to preserve uncertainty about one's type, in a given period $t \in [2, T-1]$ a manager will leave the current employer only if he has done so also in previous periods $t' \in [1, t)$. Otherwise, his type is already known and there is no reason to churn. Conversely, if a manager chooses to stay with the same employer in a given period $t \in [2, T-1]$, he has no reason to leave in subsequent periods $t'' \in (t, T-1]$. This is because his quality is already known and again there is no reason to churn. Therefore, the equilibrium simplifies to the choice of the length of the churning period K that maximizes the manager's expected utility in (13). This is defined by the first-order condition (16). The second order condition is satisfied, since

$$pu''(W_G)(1-p)c^2 + u''(W_B)[x-y-(1-p)c]^2 < 0,$$

recalling that $u''(\cdot) < 0$. ■

Proof of Proposition 4. Total differentiation of the first-order condition (16) with respect to K and c yields:

$$\frac{dK^*}{dc} = \frac{pu'(W_G) + (1-p)u'(W_B) + (1-p)(K+2)\{u''(W_B)[x-y-(1-p)c] - u''(W_G)pc\}}{(1-p)pu''(W_G)c^2 + u''(W_B)[x-y-(1-p)c]^2}$$

Since the denominator is negative, the sign of dK^*/dc is the opposite of that of the numerator, that is, is the sign of the expression:

$$-pu'(W_G) - (1-p)u'(W_B) + (1-p)(K+2)\{u''(W_B)[x-y-(1-p)c] - u''(W_G)pc\}.$$

Upon dividing this expression by $u'(W_B)$, dividing and multiplying the second term by $u'(W_G)$, and substituting from (16), one obtains:

$$\text{sign}\left(\frac{\partial K^*}{\partial c}\right) = \text{sign}\left\{-\frac{x-y}{c} - (1-p)[x-y-(1-p)c](K+2)[A(W_G) - A(W_B)]\right\}, \tag{20}$$

where $A(W)$ is the absolute risk aversion (ARA) coefficient for wealth W. The first term is negative, while the second is negative, zero or positive depending on whether the manager's ARA is increasing, constant or decreasing in wealth. So a sufficient condition for K^* to be a decreasing function of c is that the manager's utility function features constant or increasing ARA (i.e., is CARA or IARA). But this is a sufficient, not a necessary condition: it may be satisfied even if ARA decreases with wealth. In particular, it is satisfied for constant relative risk aversion (CRRA) utility, provided the relative risk aversion coefficient γ is equal to 1 (log utility) or less than 1, as can be seen by rewriting expression (20) as follows:

$$\text{sign}\left(\frac{\partial K^*}{\partial c}\right) = \text{sign}(1-p)\left\{\frac{x-y}{(1-p)c} - \frac{W_B-(T-1)y}{W_B}\frac{W_G-W_B}{W_G}\gamma\right\}.$$

The first term in curly brackets exceeds 1 (by assumption), while the two fractions in the second term are smaller than 1: hence, if $\gamma \leq 1$, K^* is decreasing in c. ■

References

Acharya, Viral, Marc Gabarro and Paolo Volpin, 2009, "Competition for Managers, Corporate Governance and Incentive Compensation," Working Paper, New York University.

Acharya, Viral, and Paolo Volpin, 2010, "Corporate Governance Externalities," *Review of Finance*, 14(1), 1-33.

Axelson, Ulf and Philip Bond, 2009, "Investment Banking Careers," Working Paper, University of Pennsylvania.

Bebchuk, Lucian, Alma Cohen and Holger Spamann, 2009, "The Wages of Failure: Executive pay in Lehman and Bear Stearns 2000-2008," Working Paper, Harvard University.

Cheng, Ing-Haw, Harrison Hong and Jose Scheinkman, 2009, "Yesterday's Heroes: Compensation and Creative Risk-Taking," Working Paper, Princeton University.

DeMarzo, Peter, Dmitry Livdan and Alexei Tchistyi, 2011, "Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives," Working Paper, Stanford University.

Dicks, David, 2009, "Executive Compensation, Incentives, and the Role for Corporate Governance Regulation," Working Paper, University of North Carolina.

Gottlieb, Daniel and Kent Smetters, 2011, "Grade Non-Disclosure," NBER Working Paper, #17465.

Fahlenbrach, Rudiger and Rene Stulz, 2009, "Bank CEO Incentives and the Credit Crisis," Working Paper, Ohio State University.

Gabaix, Xavier, and Augustin Landier, 2008, "Why Has CEO Pay Increased So Much?," *Quarterly Journal of Economics* 123, 49-100.

Harris, Milton, and Bengt Holmstrom, 1982, "A Theory of Wage Dynamics," *Review of Economic Studies* 49, 315-333.

Makarov, Igor and Guillaume Plantin, 2010, "Rewarding Trading Skills Without Inducing Gambling, " Working Paper, London Business School.

Rajan, Raghuram, 2005, "Has financial development made the world riskier?," *Proceedings, Federal Reserve Bank of Kansas City*, August, 313-369.

Rajan, Raghuram, 2008, "Bankers' Pay Is Deeply Flawed," *Financial Times*, January 9.

Richardson, Matthew and Ingo Walter, 2009, "Rethinking Compensation Practices in Financial Firms", Chapter 8 in Acharya and Richardson, eds., *Restoring Financial Stability: How to Repair a Failed System*, John Wiley & Sons.

Rosen, Sherwin, 1981, "The Economics of Superstars," *American Economic Review* 71, 845-858.

Smith, Roy, 2009, "Greed is Good," *Wall Street Journal*, February 7.

Tett, Gillian, 2009, "What Bankers Can Learn From Chelsea Football Club," *Financial Times*, September 11.

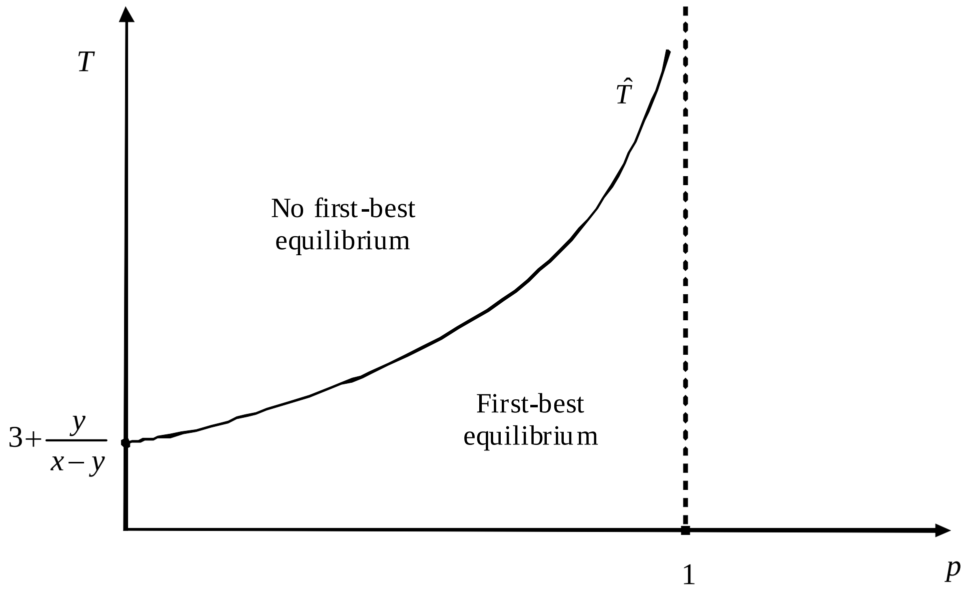

Figure 1. First-best equilibrium: career duration T and fraction of good managers p

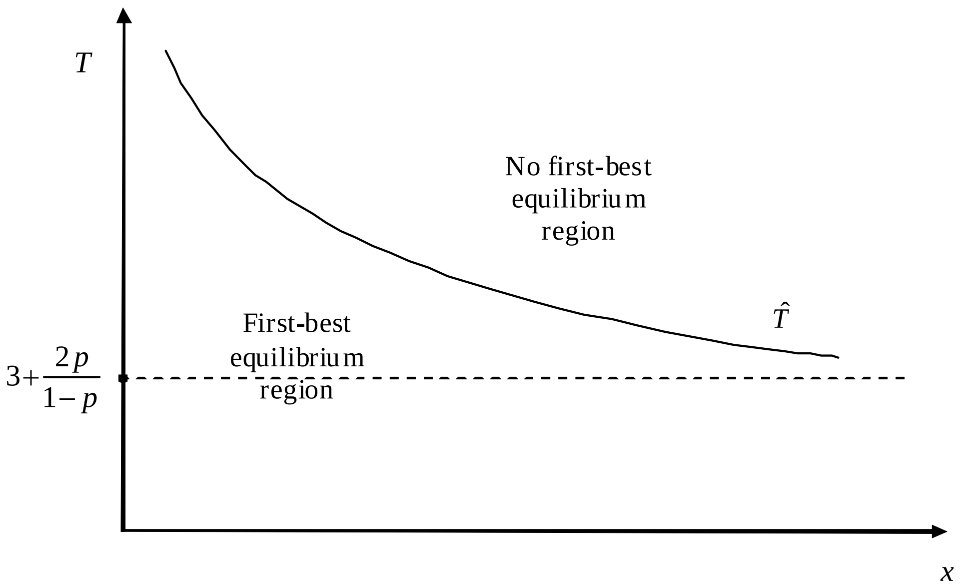

Figure 2. First-best equilibrium: career duration T and high payoff x of risky project

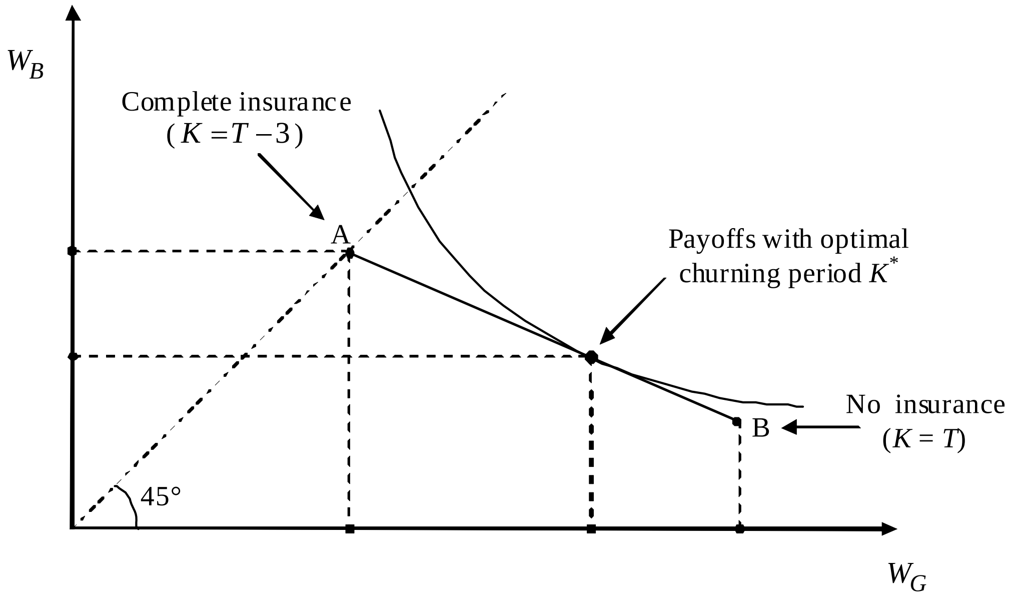

Figure 3. State-space representation of the equilibrium with churning

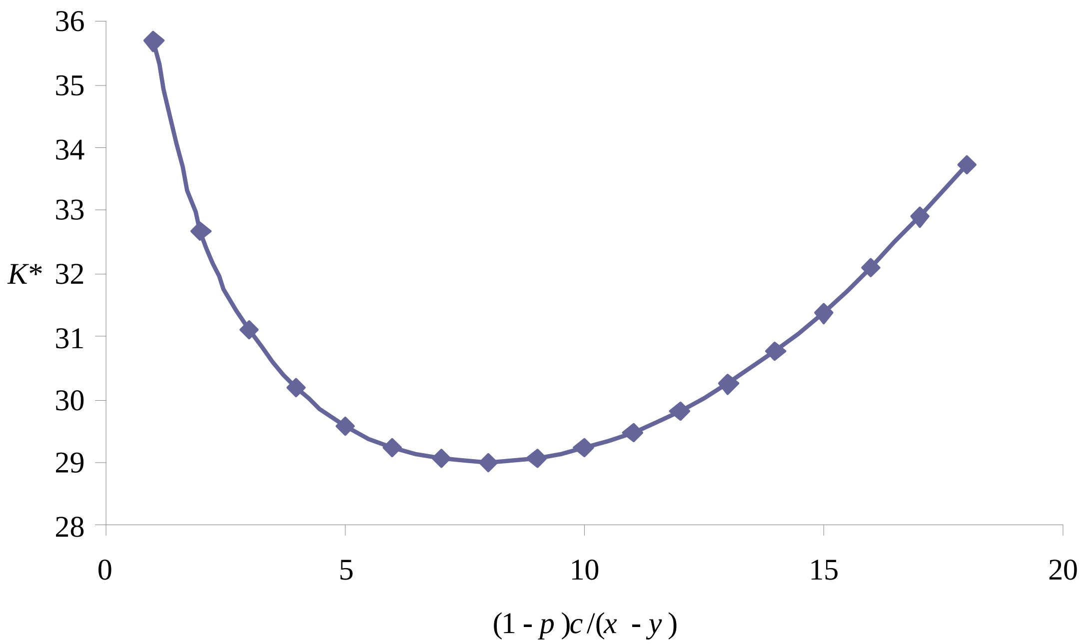

Figure 4. Churning period K^* as a function of tail risk c