

NBER WORKING PAPER SERIES

SEEKING ALPHA: EXCESS RISK TAKING AND COMPETITION FOR MANAGERIAL
TALENT

Viral V. Acharya
Marco Pagano
Paolo Volpin

Working Paper 18891
<http://www.nber.org/papers/w18891>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2013

We are grateful to Ulf Axelson, Marco Battaglini, Philip Bond, Mike Burkart, Adolfo de Motta, David Dicks, Michael Fishman, Hendrik Hakesen, Hugo Hopenhayn, Giancarlo Spagnolo 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 CSEFIGIER 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 for helpful 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). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Viral V. Acharya, Marco Pagano, and Paolo Volpin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Seeking Alpha: Excess Risk Taking and Competition for Managerial Talent

Viral V. Acharya, Marco Pagano, and Paolo Volpin

NBER Working Paper No. 18891

March 2013

JEL No. D62,G01,G2,G32,G38,J38

ABSTRACT

We present a model in which managers are risk-averse and firms compete for scarce managerial talent (“alpha”). When managers are not mobile across firms, firms provide efficient compensation, which allows for learning about managerial talent and for insurance of low-quality managers. When instead managers can move across firms, firms cannot offer co-insurance among employees. In anticipation, risk-averse managers may churn across firms or undertake aggregate risks in order to delay the revelation of their true quality. The result is excessive risk-taking with pay for short-term performance and an accumulation of long-term risks. We conclude with a discussion of policies to address the inefficiency in compensation.

Viral V. Acharya
Stern School of Business
New York University
44 West 4th Street, Suite 9-84
New York, NY 10012
and CEPR
and also NBER
vacharya@stern.nyu.edu

Paolo Volpin
Finance Department
London Business School
Regent's Park, London NW1 4SA, UK
pvolpin@london.edu

Marco Pagano
Department of Economics
University of Naples Federico II
Via Cintia, Monte S. Angelo
80126 Napoli, ITALY
mrpagano@tin.it

“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. ... 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)

“Should any investor be prepared to bet on [Mexico’s] next 100 years - or that of any country?... Cynics suggest no one buys a century bond thinking further away than their next job move since it won’t be their problem when it does come due.” – Hughes (2010)

1 Introduction

Excessive risk-taking by financial institutions and overly generous executive pay are widely regarded as key factors in the 2007-09 crisis.¹ In particular, it has become commonplace to blame banks and securities companies for compensation packages that reward managers (and more generally, other risk-takers such as traders and salesmen) generously for making investments with high returns in the short run but 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 to cut executive pay packages and rein in incentives based on options and bonuses, making them more dependent on long-term performance and in extreme cases eliminating them outright.² It is natural to ask whether this is the right policy response to the problem. It is crucial to ask what is the root of the

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

²For 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 the pay of top executives at companies receiving significant federal assistance at \$500,000. The British, Swedish, and Swiss governments also set limits on financiers’ compensation in their efforts to rescue their banking systems.

problem – that is, precisely what market failure produced excessive rewards for short-term performance at the expense of a build-up of tail risks.

The thesis of this paper is that the root of the problem is the difficulty of rewarding managerial talent when projects can have long-term or tail risk and the market allows executives to move from firm to firm before that risk materializes. For instance, a trader in a financial firm may set up a “carry trade” and then leave the firm before it is known whether it can be determined whether the carry was an actual arbitrage opportunity or simply the reward for risk (so that the trade may eventually “blow up”). In this situation managers who take tail risks while moving rapidly between firms raise their short-term performance and pay, while reducing their accountability for failures. When such job churning is possible, competition for managerial talent induces a negative externality, every firm offering an “escape route” to the others’ employees. But if the market for managerial talent is not very competitive, managers are more likely to be stuck with their initial employer and so be held responsible for project failures. The contrast between these two executive labor market regimes recalls that between the current high-mobility scene and that prevailing around the middle of last century. As Frydman (2007) shows for a balanced panel of U.S. firms from 1936 to 2003, top executives who worked throughout their careers for the same company accounted for 30 percent of the total in 1990-2003, down from 70 percent in 1940-67.

More specifically, we consider a setting in which 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 taking correspondingly high risks. But risk only materializes in the long run, so managerial talent can be identified only if the managers

who have chosen risky projects stay with their employer for a long enough time. If they leave earlier, the long-term performance of their projects is never learnt, because it is more efficient for the firm to liquidate them.

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, there is better choice of investment projects: when managers' skills are known they can be assigned to the project they are best suited to manage. Second, there is better 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 that they have incurred materialize. This means that 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 where 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 firms assign managers of unknown quality to risky projects (which they will do if the risky projects outperform safe ones 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, choosing risky projects regardless of their ability to avoid the implied risks. Talented executives will be identified only in the long run: as managers approach the end of their careers, the residual risk of being exposed as low-alpha declines, and so also does

the demand for insurance via churning.

For young managers, the benefit of churning is to delay the revelation of their true quality. If projects carry aggregate risk that delays learning individual quality from realized outcomes, then designing such projects is an alternative way for managers to synthesize insurance. Regardless of the way in which managers synthesize insurance – by churning or by undertaking aggregate-risky projects – the end result is inefficiency relative to the case of no competition for managers: since types are not revealed quickly enough, efficient allocation of managers to projects does not take place in time and too many projects fail; along the way, managers’ pay is not commensurate with their actual performance.

The model generates several further results. First, if managers are sufficiently risk-averse, then an increase in the tail risk of projects can increase job churning, hence risk for society as a whole. Second, frictions in the market for managers (e.g. search costs) can actually mitigate inefficiency by reducing managerial churning. Conversely, easy interim liquidation of assets (e.g. securitization markets for loans) can aggravate inefficiency by prompting more churning. Third, limits to deferment of managerial compensation only make it harder for firms to keep employees, heightening the inefficiency that stems from competition for managers.

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 AAA-rated mortgage-backed securities to credit default swaps or longevity insurance, contain aggregate risks and have the flavor 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 in-

centives to engage in such 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 2 discusses the literature. Section 3 describes the model. In Section 4 we solve for the equilibrium and present its novel and testable empirical implications. In Section 5 we relax several of our assumptions to check robustness. Section 6 concludes with a brief description of the model's policy implications. The proofs are in the Appendix.

2 The literature

We study a model of the labor market à la Harris-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 (good types) 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 (bad types).

With respect to this framework, our paper introduces project choice, which allows the firm to control whether types become observable or not. If managers are assigned to the safe project their type stays hidden, while if they are assigned repeatedly the same risky project it becomes known. The option of the safe project eliminates the Harris-Holmstrom problem, since if productivity shocks are hidden, then 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 has an interest in re-assigning him to the same risky project in order to learn his type. 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 persistent inefficiency in worker-project matching. As an alternative to seeing its employees leave, a firm could assign them to projects whose outcome has low sensitivity to talent, i.e. those in which aggregate risk dominates. Here too, though, insurance sacrifices productive efficiency, which requires early learning of the workers' quality.

Our results represent 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 the scarcity of their talent and the fact that it is efficiently matched with larger firms. In our setting, instead, competition for talent 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 (2009). In these models, competition prompts firms to incentivate managers via higher salaries rather than better governance, a result supported empirically by Acharya, Gabarro and Volpin (2009). 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: optimal financial regulation should limit the proportion of resources going to bonuses. In contrast to these works on governance externalities, our paper posits a dynamic setting in which it takes time for firms to learn about their employees and assign them to the right tasks, but this is impeded 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 (2012), in a multitasking model where workers differ in productivity in a rewardable task and in willingness to perform an unrewarded one, i.e. in the strength of their work ethic. When these differences are unobservable, labor contracts are designed to screen workers. When firms compete for workers, however, they will use incentive pay also to attract or retain the most productive workers, and by doing so they reduce work ethics below the social optimum. Our model is complementary to theirs: we focus on employees' firm-level insurance and on how labor-market competition, by eroding such insurance, leads to churning and undertaking of aggregate risks as an alternative way of synthesizing insurance; in contrast, Bénabou and Tirole 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 (2009) show that smart workers may be "too hard to manage", because their high outside options make them insensitive to the threat of dismissal. De Marzo, Livdan and Tchisty (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. Similarly, 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 investment, with trading strategies that may 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.

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 financial firms' competitive "search for yield" (which we interpret as "seeking alpha"). Rajan (2005) was one of the first to warn of excessive risk-taking in financial institutions driven by short-termist pay packages, which he later called "fake alpha" (Rajan, 2008). In another thought-provoking piece, Smith (2009) refers to the role of managerial mobility in entrenching the culture of bonus without performance on Wall Street.³ Indeed, the argument could apply beyond the financial sector, considering that the mobility of U.S. top managers has increased in all industries since the 1970s (Frydman, 2007), while the idiosyncratic volatility of listed U.S. firms has risen considerably, be it gauged by real or financial variables (Campbell, Lettau, Malkiel and Xu, 2001, Comin and Mulani, 2006, Comin and Philippon, 2006, among others).

Admittedly, complete consensus on the role of pay packages in firms' risk-taking has not been reached. Fahlenbrach and Stulz (2009) present evidence that bank CEOs lost a

³Smith (2009) writes 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.*"

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 (2010) contest this thesis, showing that prior to the crisis bank CEOs had given themselves payoffs greatly in excess of the amounts that they lost eventually. So they argue that bank executives did benefit from short-term compensation that was not tied to long-run performance, as in our model with job churning and aggregate risks. 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 incentives for short-term risk. None of these papers, however, explicitly examines the role of employee turnover.

3 Model

There are K profit-maximizing firms (indexed by $k = 1, \dots, K$), which live forever and are owned by risk-neutral shareholders, and I risk-averse managers (indexed by $i = 1, \dots, I$), each living for T discrete periods. The analysis focuses 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 $U = E [u(\overline{W})]$, where $u(\cdot)$ is an increasing and concave function of final (period- T) wealth \overline{W} : managers are risk-averse as regards their lifetime compensation, and are therefore concerned about the disclosure of their managerial quality. The assumption that managers only care about final wealth avoids the need to deal with intertemporal optimization problems (which are not central to the analysis) and, more importantly, sets no limits to deferring compensation: payments can be deferred to the end of the employment period, at no cost to the employer. The case with partial deferral is discussed as an extension.

Each firm can make its compensation package conditional on the projects assigned to the manager and on the timing of his possible resignation. The firm does not have recourse to the manager's wealth outside the employment contract. That is, it cannot encroach on the compensation that he receives from other employers – a realistic enough assumption. Managers start their career with no initial wealth and enjoy limited liability. This implies that their total payoff from an employment contract can never 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, and each project lasts for two periods. Hence, a manager working with the same firm for his entire career is running two projects in every period but the first and the last. Managers are not equally talented: a fraction $p \in (0, 1)$ are good at managing risk and a fraction $1 - p$ are bad. Initially, the manager i does not know his own quality $q_i = \{G, B\}$.

Projects are of two types: safe projects, which only generate an immediate and certain payoff, and risky projects, which produce a large immediate payoff but, in the hands of low-quality managers, eventually entail a loss. Specifically:

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

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

Making the risky project's yield depend on the manager's type constitutes recognition of personal ability in managing it. Good managers add value by decreasing risk (for simplicity, to zero, the same as the safe project), without reducing expected revenue. In this sense, the

good manager generates “alpha”, that is, he improves the risk-return tradeoff of his firm. Conversely, bad managers generate the same short-run return x but only at the future cost c .⁴

The crucial assumption is that if a manager initiates a project of type R , his ability becomes known *only if* he remains in charge for *both* periods. Making first-period performance uninformative captures the idea that failure is infrequent (“tail risk”), so that it takes time to determine a person’s ability to manage a risky project. In fact, to reflect the possibility 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 have an uninformative outcome even after two periods, so that learning requires even 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. In this case the project is liquidated, and in the process the identity of the project’s initiator is lost.⁵ The incomplete projects are sold because their in-house completion is inefficient: using another manager from the firm to complete an unfinished project would prevent him from starting a new project of his own, while outside there are managers who can complete the project at zero cost. In other words, within the firm “creative managers” who can initiate new projects are scarce, while “non-creative managers” who can complete them outside are abundant.

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

⁴Project 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 and incurs no cost in the second period; the poor trader (who has no clue when to close) incurs a cost c in the second period.

⁵Avoiding such information loss would require an institution that is capable both of (i) pooling information about the identity of the departing manager (obtained from the first employer) and the eventual performance of the project (from the project’s buyer), and (ii) providing such information to the new employer. Establishing such an “information broker” would demand an unrealistic level of coordination.

every 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, then $\lambda = p$, so that the liquidation price of unfinished risky projects is $x - (1 - p)c$. We assume that

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

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 greater risk corresponds to 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, good ones only to risky projects. Assigning bad managers to risky projects would imply excessive risk-taking.

3.2 The market for managerial talent

We posit that 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 managers can start a new project.⁶

Let i denote a generic manager, k a generic firm and t a generic period. At the beginning of period t , the firm decides whether to make an offer to the manager. The offer is compensation W_{ikt} contingent on the projects $\{P_{ik\tau}\}_{\tau=t}^{\tau=T-1}$ to which manager i is assigned over his future career:

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

where $P_{ik\tau} \in \{R, S\}$ if in period τ manager i is assigned to a risky (safe) project at firm k ,

⁶Note that the managers hired by firms have the experience to initiate a project, even though their quality is unknown. They are not to be confused with inexperienced managers, who are indeed abundant but are able only to complete projects already started by someone else.

$P_{ik\tau} = 0$ if manager i does not work at firm k , and $W(\cdot)$ is a mapping $(0, R, S)^{T-t} \mapsto \mathbb{R}^+$.

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, i.e. it cannot exceed the expected revenue generated by manager i in his employment relationship with firm k . To save on notation, we set $W_{ikt} = 0$ when the firm chooses not to make an offer to the manager during the period.

The manager can accept or reject the offer W_{ikt} : let $F_{it} \in \{1, 2, \dots, K\}$ denote the firm 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 compensation W_{ikt} . As we will see, this precommitment prevents firms from exploiting any informational advantage that they might gain by gauging their own 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 of the expected profit that they will generate in their tenure with any employer. But firms do not precommit to any specific project assignment: once the wage contract is agreed upon, the firm assigns the manager to whatever projects $\{P_{ik\tau}\}_{\tau=t}^{\tau=T-1}$ maximize its expected profits.

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, impeding poaching by other firms. To bring out the implications of *ex-post* competition for managerial talent, we focus initially on the two polar cases of totally absent or prohibitively high switching costs – the “competitive” and the “non-competitive” regime, respectively. In both regimes, managerial performance is taken to be publicly observable: if a manager's ability becomes known to the current employer,

it is also known to other firms.⁷ In an extension, we shall consider the intermediate case of a managerial labor market with some frictions in the form of switching costs.

In the competitive regime, at the start of each period a manager chooses whether or not to leave his current employer. In the non-competitive regime, once he has accepted the initial offer, he cannot leave. Formally, $F_{it} = F_{it+1} = k$ if manager i employed by firm k in period t chooses to stay there in period $t + 1$, while $F_{it} \neq F_{it+1}$ if he leaves at the beginning of that period. When indifferent between staying and leaving, a manager is assumed to stay. This tie-breaking assumption can be thought as reflecting the presence of an arbitrarily small switching cost even in the competitive regime.

The difference between the two regimes may capture, for instance, the changing relationship between managers and their employers documented by Frydman (2007), their sharply rising mobility between companies over the last half-century as their skills grew less and less firm-specific. This has certainly been the case in 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, due to technological change and new financial products. And corporate loyalty has lost appeal in the world of finance, as Tett (2009) emphasizes in our epigraph.

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 he is employed in all periods. The sequence

⁷This assumption is not essential in our context, however, due to the multiperiod nature of the employment relationship. To see why, suppose 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 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 if the manager's performance were not publicly observed, outside offers would be effectively conditioned on his true type, if this has become known to the manager (and current employer).

of actions is as follows:

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

(ii) In period 2, the manager chooses whether to stay with employer k ($F_{i2} = k$) or to leave ($F_{i2} \neq k$). If he stays, he completes the project P_{ik1} started in period 1 and the employer assigns him to a new project $P_{ik2} \in \{R, S\}$.⁸ If the manager leaves, the project is sold for a price equal to its expected revenue (including any first-period revenue), while the manager is hired by a new firm j ($F_{i2} = j$) that pledges to pay final compensation W_{ij2} and assigns him to a new project $P_{ij2} \in \{R, S\}$.

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

(iv) In period T , the manager cannot leave (as he will not be starting a new project). He completes the project started in period $T - 1$ and consumes his final wealth, which is the sum of the compensations awarded by his series of employers, given by $\overline{W}_i = \sum_{k=1}^K \sum_{t=1}^{T-1} W_{ikt}$, where k is a generic firm and the terms inside the sum are zero for any firm k and period t in which no offer is made or the manager rejects the offer.

3.4 Learning managers' types

In any period t the employment history of manager i can be summed up in the belief θ_{it} that his type is good ($q_i = G$). Since in our setting information about the manager's quality is symmetrical, this belief is shared by all players. At the beginning of his career, the manager's quality is unknown: he is good with probability p and bad with probability

⁸Profit maximization requires the firm to assign the completion of the initial project P_{ik1} to manager i , since this allows it to learn his type in period 2 and improve its future project-manager matching.

$1 - p$. Hence, $\theta_{i0} = p$. In each period t , the belief θ_{it} is updated on the basis of the manager's performance in that period.

As the first-period payoff of a project is uninformative, 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 employer k ($F_{i2} \neq F_{i1}$) or if he managed the safe project in period 1 ($P_{ik1} = S$), i.e. $\theta_{i2} = p$. If instead the manager stayed with his employer ($F_{i2} = F_{i1}$) and managed a risky project in period 1 ($P_{ik1} = R$), the second-period revenue of the initial project reveals his quality: if the total revenue π_{ik1} is x , manager i is revealed to be good, so that $\theta_{i2} = 1$; but if $\pi_{ik1} = x - c$, he is revealed to be bad, so that $\theta_{i2} = 0$.

Under this logic, information on the manager's type is updated in all periods $t \geq 3$ as follows:

(i) $\theta_{it} = 0$ if 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 employer ($F_{it} = F_{it-1}$) and at $t - 1$ had managed a risky project ($P_{ikt-1} = R$) that produces low revenue ($\pi_{ikt-1} = x - c$).

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

(iii) $\theta_{it} = 1$ if 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 employer ($F_{it} = F_{it-1}$) and at $t - 1$ was assigned a risky project ($P_{ikt-1} = R$) that produces high revenue ($\pi_{ikt-1} = x$).

3.5 Strategies and payoffs

At the start of period t , firm k offers any manager i not currently employed by it a compensation based on its belief about his quality. This belief depends only on the 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 $W(\cdot | \theta_{it-1})$ to manager i .

If firm k employs manager i , in each period t it will assign him to a project, i.e. choose $P_{ikt} \in \{R, S\}$, as a function of the belief θ_{it-1} about the manager's quality, so as to maximize the expected revenue:

$$\pi(P_{ikt} | \theta_{it-1}) = \begin{cases} x - (1 - \theta_{it-1})c & \text{if } P_{ikt} = R, \\ y & \text{if } P_{ikt} = S. \end{cases} \quad (2)$$

The manager's strategy consists in his period-by-period choice of employer. Formally, manager i employed by firm k in period $t-1$ will choose which firm to work for in period t (F_{it}) as a function of the belief θ_{it-1} about his quality, so as to maximize the expected utility from his compensation $U(\overline{W}_i | \theta_{it-1})$.

4 Equilibrium

Here we solve for the equilibrium in each of the competitive and non-competitive labor market regimes described in Section 3.2. In the competitive regime, if he wishes a manager can choose to work in a different firm F_{it} in each period; in the non-competitive regime, he is constrained to remain with his initial employer F_{i1} , so that good managers cannot be poached by new employers, even if their talent has been revealed by their performance.

Recall that in both regimes firms compete for managers *ex ante*: they all bid, and managers take 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 following game:

(i) In any period t , firm k chooses to assign manager i to project P_{ikt} so as to maximize its expected profits conditional on the belief θ_{it-1} about manager i 's quality for all periods in which manager i works for firm k :

$$\max_{P_{ikt} \in \{S, R\}} \pi(P_{ikt} | \theta_{it-1}) \cdot I_{F_{it}=k} \equiv \pi_{ikt}^*, \quad (3)$$

where $\pi_{ikt}(P_{ikt} | \theta_{it-1})$ is defined in equation (2), and the matching indicator $I_{F_{it}=k} = 1$ if $F_{it} = k$ and $I_{F_{it}=k} = 0$ otherwise.

(ii) In any period t , firm k chooses W_{ikt} so as to maximize its expected profits from hiring manager i , conditional on the belief θ_{it-1} about manager i 's quality:

$$\max_{W_{ikt}} \left[E \left(\sum_{\tau=t}^{\tau=T-1} \pi_{ikt\tau}^* \mid \theta_{it-1} \right) - W_{ikt} \right] \cdot I_{F_{it}=k}, \quad (4)$$

taking as the manager's and the other firms' strategies as given.

(iii) In any given period t , manager i chooses his employer F_{it} so as to maximize his expected utility conditional on the belief θ_{it-1} :

$$\max_{F_{it}} U(\bar{W}_i \mid \theta_{it-1}) = E \left[u \left(\sum_{k=1}^K \sum_{s=1}^{T-1} W_{iks} \right) \mid \theta_{it-1} \right], \quad (5)$$

taking the firm's strategy as given.

(iv) 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 only in that the firm's problem (4) and the manager's problem (5) 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. Since in this case solving for the equilibrium is simpler, we shall begin with the analysis of the non-competitive regime.

First, however, let us note that, once manager i has chosen to work for firm k in period t , the solution to the optimal assignment problem (part (i) of the equilibrium as defined above) depends only on the belief concerning the manager's quality. Since θ_{it-1} can only take three values, namely 0, 1 and p , assumption (1) implies that:

$$(P_{ikt}^*, \pi_{ikt}^*) = \begin{cases} (R, x) & \text{if } \theta_{it-1} = 1, \\ (R, x - (1-p)c) & \text{if } \theta_{it-1} = p, \\ (S, y) & \text{if } \theta_{it-1} = 0. \end{cases} \quad (6)$$

That is, managers who are known to be good or of unknown quality are assigned to risky projects, those known to be bad to safe projects.

4.1 The non-competitive regime

When there is no ex-post mobility of managers, firm k 's problem (4) simplifies to:

$$\max_{W_{ik1}} \left[E \left(\sum_{t=1}^{t=T-1} \pi_{ikt}^* \mid p \right) - W_{ik1} \right] \cdot I_{F_{i1}=k}, \quad (7)$$

because the hiring decision is made 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

$$W_{ik1} = E \left(\sum_{t=1}^{t=T-1} \pi_{ikt}^* \mid p \right). \quad (8)$$

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, managers are perfectly

insured against the risk arising from their unknown quality: equation (8) implies that good managers subsidize bad ones.

The firm optimally chooses to assign the risky project to a manager of unknown quality in periods 1 and 2. From period 3 onwards, the firm makes the assignment conditional on the manager's true quality, assigning only risky projects to good managers, only safe projects to bad ones. Under this policy, over his career the manager will generate revenues

$$\Pi^* = 2[x - (1 - p)c] + (T - 3)[px + (1 - p)y]. \quad (9)$$

The first term in (9) is expected period-1 and period-2 profit from the risky project undertaken at $t = 1$ and $t = 2$ by a manager of unknown quality;⁹ the second term is the sum of the expected continuation revenues of the two (known) types of managers in periods 3 through T , weighted by their respective frequencies.

This equilibrium outcome is the first best: it features both (i) optimal risk-sharing, i.e. complete insurance of managers by firms (which are risk neutral) and (ii) productive efficiency, i.e. 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

$$U^* = u(\Pi^*), \quad (10)$$

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.*

⁹As 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.

Note that optimal risk-sharing requires the firm not to make salary conditional on employees' quality, even though this information is used in the matching of 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. And in fact in the competitive regime, to which we turn next, this cross-subsidization may break down.

4.2 A 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. 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 other employers can infer the manager's ability, they will bid the per-period compensation of good managers up to x , and offer y to bad ones. From expression (9), it follows immediately that the first-best compensation per period, $\Pi^*/(T-1)$, is lower than x and higher than y : thus if a firm offered this amount, its good managers would leave, and 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 penalizing good managers who leave. The most effective contract of this sort makes the entire date- T compensation Π^* contingent on the manager never leaving the firm: if the manager leaves at any time, the firm pays nothing.¹⁰

¹⁰Recall 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 :

$$W_{ik1} = \begin{cases} \sum_{t=1}^{t=T-1} E(\pi_{ikt}^* | p) = \Pi^* & \text{if } F_{it} = k \ \forall t, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

where π_{ikt}^* is the profit generated by the optimal project assignment at time t in firm k , from expression (6). 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 pays him nothing. But we must check whether the manager who signs such a contract actually has no incentive to leave.

First, we note that if a manager plans to leave the firm eventually, under contract (11) he will want to leave no later than the start of period 3, since staying longer only increases the penalty. Second, leaving at the start of period 2 is inefficient, because it entails no learning about the manager's quality, but a penalty equal to the first-period revenue. Third, a manager who is shown to be bad in period 2 has no incentive to leave. So we need only to consider a manager revealed to be good in period 2. If he stays with the firm, his final wealth is Π^* . If instead he leaves at the start of period 3, he earns 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:

$$\widehat{T} = 3 + 2 \frac{x - (1 - p)c}{(1 - p)(x - y)}. \quad (12)$$

If $T \leq \widehat{T}$ the first-best allocation can be sustained even in the competitive regime, but if $T > \widehat{T}$ it cannot. Intuitively, if the manager's career duration T is very short, then he must spend a large part of it with one employer merely in order to be recognized as good; he therefore loses an accordingly large fraction of his wealth if he chooses to leave. For instance, if his career spans three periods ($T = 3$), he loses $2/3$ of his lifetime stream of

revenue to the initial employer, and earns only $1/3$ 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. However, if career duration is longer, i.e., $T > \widehat{T}$, then contract (11) would not deter the manager from leaving. Intuitively, the penalty (the loss of the revenue produced in periods 1 and 2) is less than the gain in later periods. In this 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. 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 low because there are few of them. In Figure 2, instead, we see that an increase in the extra profitability of a well-managed risky project over a safe one, $x - y$, reduces the range of values of T for which the first-best allocation can be achieved: when these extra profits are large, outside employers can lure 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 (12).*

What happens when the first best cannot be attained, i.e. when $T > \widehat{T}$? Contract (11) 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 (11) by offering a higher compensation.

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

$$W_{ikt} = E \left(\sum_{\tau=t}^{T-1} \pi_{ik\tau}^* \mid \theta_{it-1} \right). \quad (13)$$

What remains to be determined in order to characterize the equilibrium is the managers' choice to stay or leave. We analyze the following candidate equilibrium: the manager changes employers in each of the first K periods, earning the expected revenue $x - (1-p)c$ per period, with $K \in [0, T-3]$, and from period $K+1$ onwards he remains with the same employer. Since the employer optimally chooses project R in periods $K+1$ and $K+2$, by period $K+3$ the manager's quality is known, so he will then be assigned project R if he is good, project S if not. Hence, the manager's problem in (5), substituting for the compensation (13) and for the optimal choice of project described above, can be rewritten simply as:

$$\max_K pu(W_G) + (1-p)u(W_B), \quad (14)$$

where

$$W_G \equiv (K+2)[x - (1-p)c] + (T-3-K)x \quad (15)$$

is the final wealth of a good manager, and

$$W_B \equiv (K+2)[x - (1-p)c] + (T-3-K)y \quad (16)$$

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 produces greater inefficiency, as bad managers should optimally be assigned only to safe projects. So

the trade-off is between insurance, obtained by delaying quality revelation (larger K) and productive efficiency, which comes with earlier revelation (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 does achieve productive efficiency; in the second, he achieves perfect insurance by churning jobs forever, at the cost of low efficiency. The optimal K maximizes expression (14), and is defined implicitly by the first order condition:

$$\frac{u'(W_B)}{u'(W_G)} = \frac{pc}{x - y - (1 - p)c}, \quad (17)$$

where W_G and W_B are given by (15) and (16) and the fraction is positive by assumption (1). Intuitively, increasing 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 $T > \hat{T}$ in equilibrium the manager switches firms in each of the first K^* periods, and subsequently remains with the same firm, where K^* satisfies condition (17).*

Figure 3 describes the equilibrium in the space (W_G, W_B) . Point A on the 45° line represents the final wealth obtained by churning for $T - 3$ periods: in this case the manager obtains the same wealth independently of his type. Point B represents the case in which the manager elects not to churn. In this case, if his type is good his final wealth (W_G) is much greater than 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, that is, it depends on the marginal rate of substitution between the two states of the world (good type and bad type) and hence on the manager's risk aversion. Intuitively, the more risk-averse will choose a higher K to smooth consumption between the two states more fully. As the graph shows, the solution is the point of tangency between the manager's indifference curve and \overline{AB} .

It is worth noticing that churning for $T - \hat{T}$ periods, where \hat{T} is given in equation (12), and then staying with a firm for the remaining \hat{T} periods cannot be an equilibrium if there is competition for managers. Intuitively, under competition a firm can only punish departing managers by appropriating the revenue that they have produced within the firm; it cannot get any part of what managers have received from previous employers. This destroys the bonding mechanism at the basis of the derivation of \hat{T} . With competition, insurance can be achieved only by churning every period for up to K^* periods, and this is a dominant strategy when $T > \hat{T}$.

Finally, one wonders whether firms might offer managers the possibility of delaying revelation by allowing "internal churning" across projects without switching to other employers. This could be achieved by assigning managers of unknown quality to a new risky project in every period and liquidating all the resulting incomplete projects. But under our assumptions the firm cannot commit to such a rule: once the wage contract is signed, the firm assigns managers to projects so as to maximize its expected profits, which calls for learning managers' quality as fast as possible. Hence, a manager of unknown quality spending more than a single period with an employer would be assigned to the same risky project twice, and his type would be revealed.

4.3 Comparative statics

The time that the typical manager spends churning, K^* , can be taken as a measure of the inefficiency generated by *ex-post* competition for managers. Hence, it is interesting to investigate how K^* responds to changes in parameters. In the simple case where managers have negative exponential (CARA) utility, predictions are unambiguous:

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

$$K^* = \max \left\{ T - 3 - \frac{\log(g)}{x - y}, 0 \right\}, \quad (18)$$

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 and in the probability of being a good manager p , and decreasing in the magnitude of tail risk c .¹¹

These results are intuitive. A longer employment horizon T makes the manager more averse to revealing his type, because the risk refers to a larger future cash flow; he accordingly churns over a longer interval. By the same token, a more risk-averse manager will seek more insurance, hence churn longer: ironically, greater risk aversion by managers actually implies more risk-taking by society! 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 $1 - p$.

The result that a longer employment horizon T implies a longer churning interval K^* extends beyond the CARA case: any risk averse manager will churn longer if his career

¹¹Expression (18) follows from replacing $u(w) = -e^{-\gamma w}$ in the first-order condition (17) and solving 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 .

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

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

Intuitively, an increase in c increases the cost of getting 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. The substitution effect reduces the demand for insurance (and thus induces a reduction in K^*), but the wealth effect (because higher c implies lower average payoff for the manager) may increase the demand for insurance if risk aversion is decreasing in wealth. The proposition identifies cases in which the substitution effect dominates.

However, there are circumstances in which tail risk affects the churning period K^* in the opposite direction from Proposition 4. This occurs if managers are highly risk-averse and if the parameter c is large, so that the associated wealth effect is sizeable. 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 high c it leads to a greater K^* . (More precisely, on the horizontal axis the tail risk parameter c is standardized by $(x - y)/(1 - p)$, which is the maximum value of c consistent with assumption (1)).¹² On the left-hand portion of Figure 4 the substitution effect dominates, so managers churn less as tail risk

¹²If 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

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

increases; on its right-hand portion the wealth effect dominates, so they churn more if tail risk increases.

Accordingly, if managers are highly risk-averse and projects feature large tail risk, as in the right-hand part of Figure 4, they may respond to an increase in tail risk by taking more insurance (churning) rather than less. This would exacerbate the inefficiency due to delayed allocation of good managers to the risky projects. Hence, paradoxically, where managers are strongly risk-averse, an increase in the tail risk of projects would stimulate tail-risk seeking.

A limitation of the model considered so far is that it assumes that a risky project is perfectly informative about the quality of a manager when he is assigned to it for two consecutive periods. In the next section, we consider a more general version, expanding the choice of projects to include a risky project that is not always informative about the manager’s quality, even when he is assigned to it twice or more.

4.4 Aggregate risk as a source of insurance

We now consider a setting where managers can be lucky for a time (that is, their type is not recognized even if they stay with their employer for two or more periods). This is because the firm has wider choice of projects: in addition to the risky and safe projects described above, the firm can choose a new type of risky project, whose payoff depends not only on the manager’s quality but also on the realization of an aggregate shock. We denote this as project A (for “aggregate risk”). Formally, at time t firm k can assign manager i to project

$$P_{ikt} \in \{A, R, S\}.$$

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).

Specifically, we make several changes to the model. First, the risk of project R does not depend only on the possible mismatch with managers; there is also intrinsic risk: the first-period payoff is a random variable $\tilde{x} = x + \tilde{u}$, where \tilde{u} is a zero-mean project-specific shock, so that its expected value is x as in the baseline model; its second-period payoff is exactly as in the baseline model: 0 if the manager is good and $-c$ if he is bad.

Second, project A has the same expected payoff as project R but different risk characteristics.¹³ With probability β , it reflects aggregate risk, not the manager's quality: the first-period payoff is a random variable $\tilde{a} = x + \tilde{v}$, where \tilde{v} is a zero-mean economy-wide shock; the second-period payoff does not depend on the manager's type but has the same mean value as project R , i.e. $-(1-p)c$. With complementary probability $1 - \beta$, project A has the same payoffs as project R and is therefore sensitive to the manager's quality. Hence, β captures project A 's sensitivity to aggregate risk, and $1 - \beta$ its sensitivity to managerial quality.

Employers can correctly identify instances in which project A reflects aggregate risk, since they can observe the performance of other such projects in their own and other firms, and the shock \tilde{v} is common to all. In these instances, firms will consider the payoff of project A as uninformative about the manager's quality: if it is unknown, it remains unknown, even if he does not move to a new employer. Managers will then be indifferent between staying and moving, since in either case their type will not be revealed. But assuming a tiny moving cost to break the tie, they will prefer to remain with the initial employer. Instead, when project A turns out to be sensitive to managerial quality, it is completely

¹³If project A had a lower expected payoff than project R , from the firms' viewpoint it would be dominated, and therefore would never be chosen. If instead it had a higher expected payoff, firms would prefer it to project R in the interval K^* when managers do not want their identity disclosed. If instead project A had a greater expected payoff than project R , the results would be similar to those reported in this section, except that the value of K^* chosen by managers would generally differ from that of the baseline model, since their equilibrium wealth level would be greater.

equivalent to project R , so that managers who wish their type not to be revealed will have to switch to a new employer.

Consider now the firm's optimal project assignment. When the manager does not want his type to be revealed, project A and project R have the same expected payoff, so that from the firm's viewpoint they are equivalent. Again, with a tiny moving cost, it is natural to break the tie by assuming that firms will pick project A , which with probability β allows managers to save the moving cost. But when the manager is willing to have his type disclosed, so that he does not keep churning if assigned to project R , it is more efficient to assign him to project R , since this entails faster learning about his type than project A . Once the manager's type is revealed, it becomes optimal to assign good managers to project R , by the left-hand side inequality in condition (1); it will also be optimal to assign bad managers to project S if $y > x - c + \beta pc$, which is a more stringent condition than the right-hand side inequality in (1). Otherwise, project A will dominate project S , which will therefore never be chosen. To preserve the similarity with the equilibrium in the baseline model, we assume that this stricter condition holds: it effectively requires an upper bound on project A 's exposure to aggregate risk: $\beta < (y + c - x)/pc$.

To summarize, from the manager's standpoint the presence of a project with aggregate risk is a source of insurance that substitutes for switching to a new employer whenever the project's payoff is determined by aggregate risk, which happens with probability β . Under the same tie-breaking conditions as above, this insurance is superior to that obtained by churning, so the mobility of managers is lower than in the baseline model and firms assign managers to project A rather than to project R until their types become known.

How long will managers want to delay this revelation? It turns out that they do so for the same number of periods, K^* , for which they choose to churn in the baseline model.

To see this, consider that the manager's payoff in each of these K^* periods is $x - (1 - p)c$ as in the baseline model, since it is efficient for firms to insure him against the aggregate shock \tilde{v} . The payoff once the manager's type is revealed is also the same as in the baseline model, that is, x for good managers and y for bad. Hence, their lifetime payoffs W_G and W_B are again defined by expressions (15) and (16) and their choice problem by expression (14). The only difference is that the solution K^* to this problem is no longer the churning interval, but the interval over which managers prevent their type from becoming known: the difference from the baseline model is that in these K^* periods managers stay with the same firm with frequency β and churn with frequency $1 - \beta$. Now the firm partly synthesizes an insurance for the manager using the project with aggregate risk, so the manager needs to churn less. As this occurs the more frequently, the greater the project's aggregate risk sensitivity β , the expected number of periods in which managers churns is decreasing in β .

The analysis of the first-best case is identical to that in Section 4.1: it is optimal to discover the type as soon as possible and then assign project R to the good and project S to the bad manager. Similarly, the cutoff \hat{T} – which defines the longest career duration that allows the firm to retain its managers in the competitive regime – does not change. Hence the inefficiency due to delayed disclosure of managers' types is the same as in the baseline model – only the way this outcome is achieved is different.

The foregoing can be summarized as follows:

Proposition 5 (Project-level risk) *In a competitive managerial market with $T > \hat{T}$, if firms can choose projects involving aggregate risk and whose payoff is uninformative about managers' quality with frequency β and is on average equal to that of the risky and informative project, then on average managers churn for $(1 - \beta)K^*$ periods and stay with the same firm for βK^* periods, where K^* satisfies condition (17).*

This proposition implies that if the aggregate risk exposure β of the financial sector increases, managers are less inclined to churn. The inefficiency in managerial assignment is the same as in the baseline model: rather than only via managerial mobility, the delay in recognition of managers' true skill now occurs partly via aggregate risk-taking. If managers could not move, such aggregate risk would not be undertaken, as firms would learn managers' types as quickly as possible and would never assign managers to project A . Hence, the key implication that competition in the managerial labor market leads to inefficiency is present also in this more general version of the model – indeed, under the above assumptions it is the same as in the baseline model.

4.5 Empirical predictions

The model delivers several important predictions that could be tested using micro data on the compensation and mobility of low- and middle-level managers or security traders.

First, given low mobility, pay should be high on average and insensitive to individual ability. This follows from Proposition 1: with low labor mobility, the average pay level should be high because employees are allocated efficiently across projects, while the insensitivity of pay to ability reflects employees' being insured against uncertainty over their quality.

Second, to the extent that there is labor mobility across firms, employees should switch jobs early in their career, when the risk of being revealed to be low-quality would affect employees' income for a longer part of their entire career, as implied by Propositions 2 and 3. Further, since junior managers are more likely than senior ones to switch jobs, they are more likely to be associated with excessive risk-taking.

Third, once employees stop switching across firms – in case of senior employees – wages

should become more highly differentiated and more sensitive to individual ability. In other words, given labor mobility there should be greater cross-sectional dispersion of pay within older than younger cohorts.

Fourth, since by Proposition 5 undertaking projects with high aggregate risk is an alternative to churning, the model predicts that junior employees of financial firms will be assigned to the projects that are more exposed to aggregate risk and less sensitive to individual managerial ability, and the more so, the greater labor mobility or competition for talent there is.

These predictions differ from those of a search model, where workers stop searching upon finding a better match, and therefore a higher wage, than in their previous job. In our model, when workers stop churning, their quality is revealed and their salary becomes permanently higher or lower, depending on whether they are of the good or bad type.

5 Extensions

We can extend the model to investigate how its insights change when two of its key assumptions are modified. In Section 5.1 we explore how the results change when firms are not allowed to defer all managerial compensation until the end of the employment relationship. In Section 5.2 we consider how the labor or financial market frictions affect managers' desire to churn jobs: we consider first informational friction in the market for managers, due to adverse selection and then friction due to search costs in the labor market or liquidation costs in the market for incomplete projects.

5.1 Limits to deferring compensation

A key assumption for all the results derived so far was that there are no constraints on withholding compensation to a manager who resigns. In practice, however, this assumption may not be realistic: employment contracts where the manager is denied compensation for past work because he switches to a new employer may be illegal. In practice, at least part of the total compensation takes the form of salary.

Limited liability would prevent the initial employer from reclaiming such interim salary payments: hence, limits to deferred compensation shrink the parameter region where the first best is attainable compared to the region described in Proposition 2. Intuitively, the more constrained the firm is in deferring compensation, the smaller the penalty it can threaten on resigning managers, and hence the smaller the parameter region where it can attain the same employee 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 consists in a non-recoverable per-period salary $\underline{w} > 0$, the maximum career duration for which the first-best outcome can be attained is:

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

which is strictly decreasing in \underline{w} .

5.2 Imperfections in the labor or asset market

In this section we consider the effects of imperfections in the labor and the financial markets that directly or indirectly increase the cost of churning, to see how they affect the extent of churning to delay type revelation.

5.2.1 Asymmetric information

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

A less extreme case is one where only a fraction ϕ of managers know their type from the start. In this case, churning should decrease in equilibrium for two reasons: (i) mechanically, the fraction $p\phi$ of managers who know they are good will stick with their employer; (ii) managers of unknown type will get pooled with those who know they are bad, and so will want to churn for less time 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 becomes $x - (1 - \hat{p})c$, where \hat{p} is the updated probability that the project was started by a bad manager:

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

Since \hat{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 $x - (1 - \hat{p})c$. Formally, the manager's problem is identical to the case described in Section 4.3 to consider the effect of a change in c : the conflict between substitution and wealth effect keeps us from determining the sign of the effect on K . If we assume a CARA utility function, Proposition 4 implies that the optimal churning period K^* shortens with the severity of adverse selection ϕ , as the cost of insurance increases.

5.2.2 Search costs

Consider next the impact of search costs: leaving his job, the manager is unemployed for some time before finding a new job. If we denote this search cost by s , the payoff in case of churning falls 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 to liquidate them. This affects the manager's payoff in case of churning: it decreases from $x - (1 - p)c$ to $x - (1 - p)c - s$ because of the liquidation cost. These two imperfections have the same effects on the churning equilibrium. The only change in the manager's problem (14) is that now his final wealth is defined as follows:

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

for a good manager, and

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

for a bad manager.

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

$$\frac{u'(\widehat{W}_B)}{u'(\widehat{W}_G)} = \frac{p[(1 - p)c + s]}{(1 - p)[x - y - (1 - p)c - s]}. \quad (19)$$

This expression is greater than the right-hand side of equation (17) for any $s > 0$ and is strictly increasing in s . Also, if $s > x - y - (1 - p)c$ then the first-order condition (19) cannot hold: the optimal choice is to set $\widehat{K} = 0$. Intuitively, if search costs are very high, there is no churning and no excessive risk-taking.

As before, the conflict between the substitution and wealth effects prevents us from determining whether in general \widehat{K} is smaller or greater than K^* . Following the steps in

Proposition 4, this ambiguity disappears in the case of CARA utility, where:

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

and

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

Hence, $\widehat{K} < K^*$ and $\partial \widehat{K} / \partial s < 0$: an increase in search costs leads managers to churn less.

The same holds 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 poorly developed securitization markets.

However, as we know from Proposition 4, with different utility functions the effect may be 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 the tail risk of projects 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 Conclusions

Firms are strongly motivated to gather information about their employees' talents and use it to allocate them efficiently to projects. 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. And managers can

exploit this dark side by taking on projects with tail risk and using the labor market to move from firm to firm to delay the resolution of uncertainty about their talent.

In addition to producing a number of testable predictions, our model presented also carries policy implications for the financial sector, where projects with tail risk are often available. In our inefficient churning equilibrium, no individual financial institution has the incentive to deviate and unilaterally stop competing for other the others' managers: as in the epigraph from Tett (2009), banks "feel utterly trapped", and only the intervention of a public authority (like the FIFA for soccer or the US major league baseball organization) can stop banks from poaching one another's managers. 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 ex-post competition regime. And in fact, in the real world we do not find such no-compete clauses in finance, presumably because of a scarcity of talented managers and their low loyalty to employers. The outcome is that in our setting policies that discourage managerial mobility – say, taxing managers who switch jobs at a higher rate than loyal ones – can improve efficiency: if such a surtax were high enough, it would effectively move the economy to the first best, although in equilibrium it not be paid, since managers would not switch jobs. In short, one policy prescription deriving from the model is to "throw sand in the wheels" of the managerial labor market.

Another policy proposal is capping managerial compensation in banks. How would this change the equilibrium in our model with managerial competition? Would it make churning – and the associated excessive risk-taking – less attractive to managers? In the model, 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 the top financial managers may

respond not only to ethical or political concerns but also to an efficiency rationale, not just a basis in (though this is yet to be spelled out by the caps' proponents). Indeed, according to the model, an appropriate pay cap would raise the expected utility of managers themselves. Interestingly, also in the setting of Bénabou and Tirole (2012) a cap on managerial pay, hence a reduction in its sensitivity to performance, can restore the first-best outcome.

Another common proposal in the debate on executive compensation is partial deferral (“claw back”) and indexation of the deferred portion to long-term performance. Our model provides a rationale for this approach: as is shown in Section 5.1, anything that constrains firms' ability to defer compensation is inefficient.

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 taking an MBA, say. In this case, capping their salary (or concealing MBA grades from employers) would reduce the “average alpha” of managers in equilibrium. Moreover, preventing the reallocation of managerial talent could have efficiency costs that are not captured by our 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 formed. 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. Such considerations warrant further modeling in our framework, which was limited exclusively on one dark side to managerial mobility.

Proofs

Proof of Proposition 3. Since in this setting the only reason for switching employers is to preserve uncertainty about one's type, in a given period $t \in [2, T - 1]$ a manager leaves the current employer only if he has also done so in previous periods $t' \in [1, t)$. If not, his type is already known and there is no reason to churn. Conversely, if a manager chooses to stay with his employer in a given period $t \in [2, T - 1]$, he has no reason to leave in subsequent periods $t'' \in (t, T - 1]$, again because his quality is already known. Thus the equilibrium simplifies to the choice of the length of the churning period K that maximizes the manager's expected utility in (14). This is defined by the first-order condition (17). 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 (17) 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, it is the sign of the expression:

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

Dividing this by $u'(W_B)$, dividing and multiplying the second term by $u'(W_G)$, and substituting from (17), we get:

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

where $A(W)$ is the absolute risk aversion (ARA) coefficient for wealth W . The first term is negative; 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 feature constant or increasing ARA (i.e., that it be 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, as long as 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 less 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, “Wall Street Occupations: An Equilibrium Theory of Overpaid Jobs,” Working Paper, University of Pennsylvania.
- Bebchuk, Lucian, Alma Cohen and Holger Spamann, 2010, “The Wages of Failure: Executive pay in Lehman and Bear Stearns 2000-2008,” *Yale Journal on Regulation*, 27, 257-282.
- Bénabou, Roland, and Jean Tirole (2012), “Bonus Culture: Competitive Pay, Screening, and Multitasking,” Working Paper, Princeton University and Toulouse School of Economics.
- Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu (2001), “Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk,” *Journal of Finance*, 56(1), 1-43.
- Cheng, Ing-Haw, Harrison Hong and Jose Scheinkman, 2009, “Yesterday’s Heroes: Compensation and Creative Risk-Taking,” Working Paper, Princeton University.
- Comin, Diego, and Sunil Mulani (2006), “Diverging Trends in Aggregate and Firm Volatility.” *Review of Economics and Statistics*, 88(2), 374-83.

- Comin, Diego, and Thomas Philippon (2006), “The Rise in Firm-Level Volatility: Causes and Consequences,” in *NBER Macroeconomics Annual 2005*, Vol. 20, ed. Mark Gertler and Kenneth Rogoff, 167–201. Cambridge, MA: MIT Press.
- DeMarzo, Peter, Dmitry Livdan and Alexei Tchisty, 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,” *Review of Financial Studies*, forthcoming.
- Fahlenbrach, Rudiger, and Rene Stulz, 2009, “Bank CEO Incentives and the Credit Crisis,” *Journal of Financial Economics*, forthcoming.
- Frydman, Carola (2007), “Rising Through the Ranks: The Evolution of the Market for Corporate Executives, 1936-2003,” Working Paper, Boston 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.
- Hirshleifer, Jack, 1971, “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review* 61, 561-574.
- Hughes, Jennifer, 2010, “Mexico’s sale of the century,” *Financial Times*, October 6.
- 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.
- Thanassoulis, John, 2012, "The Case of Intervening in Bankers' Pay," *Journal of Finance* 67, 849-895.




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


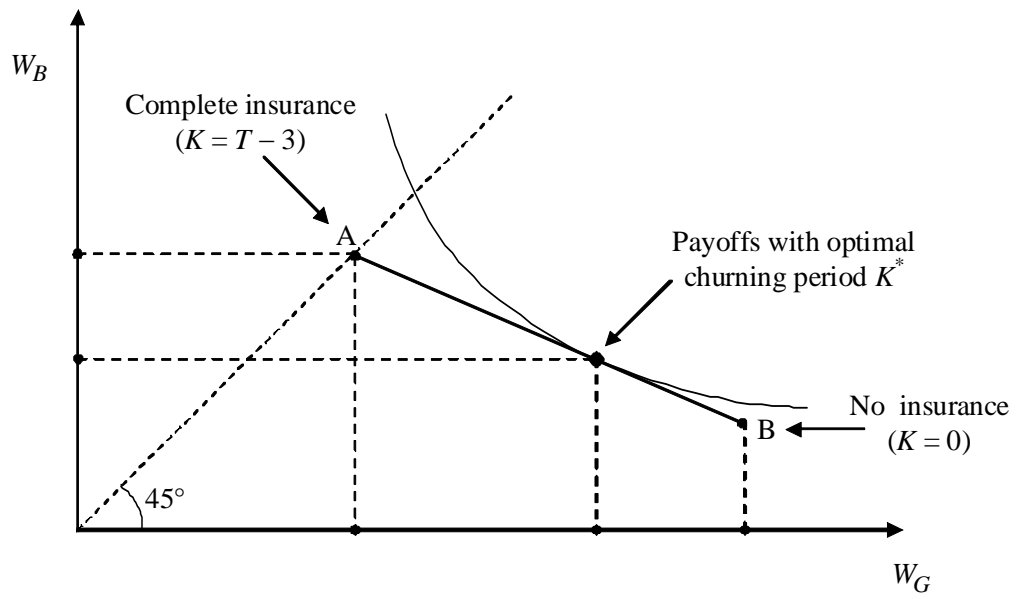


Figure 2. First-best equilibrium: career duration T and high payoff x of risky project



State-space representation of the equilibrium with churning

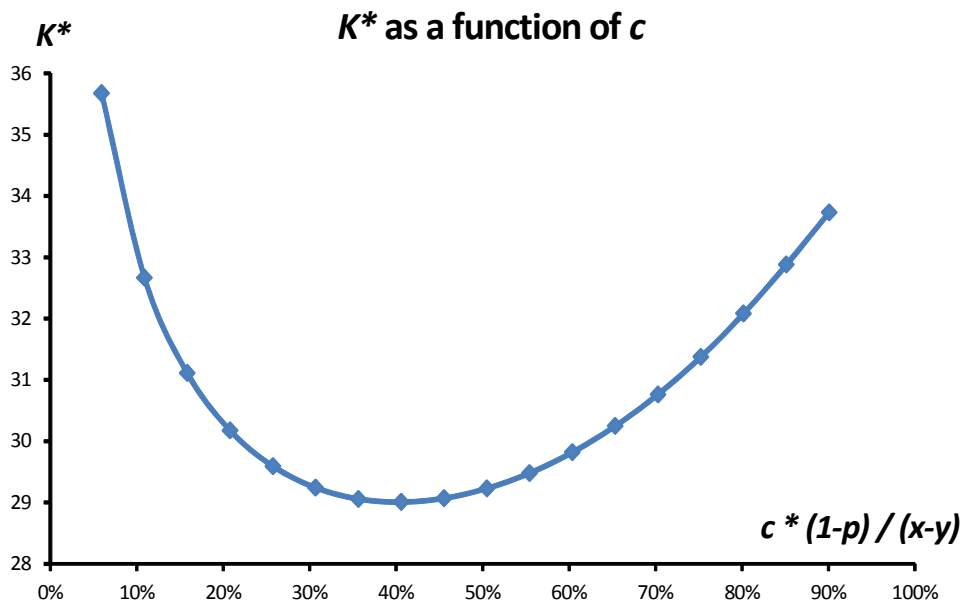


Figure 4. Churning period K^* as a function of tail risk c