EUROPEAN CENTRAL BANK

WORKING PAPER SERIES
NO. 520 / SEPTEMBER 2005

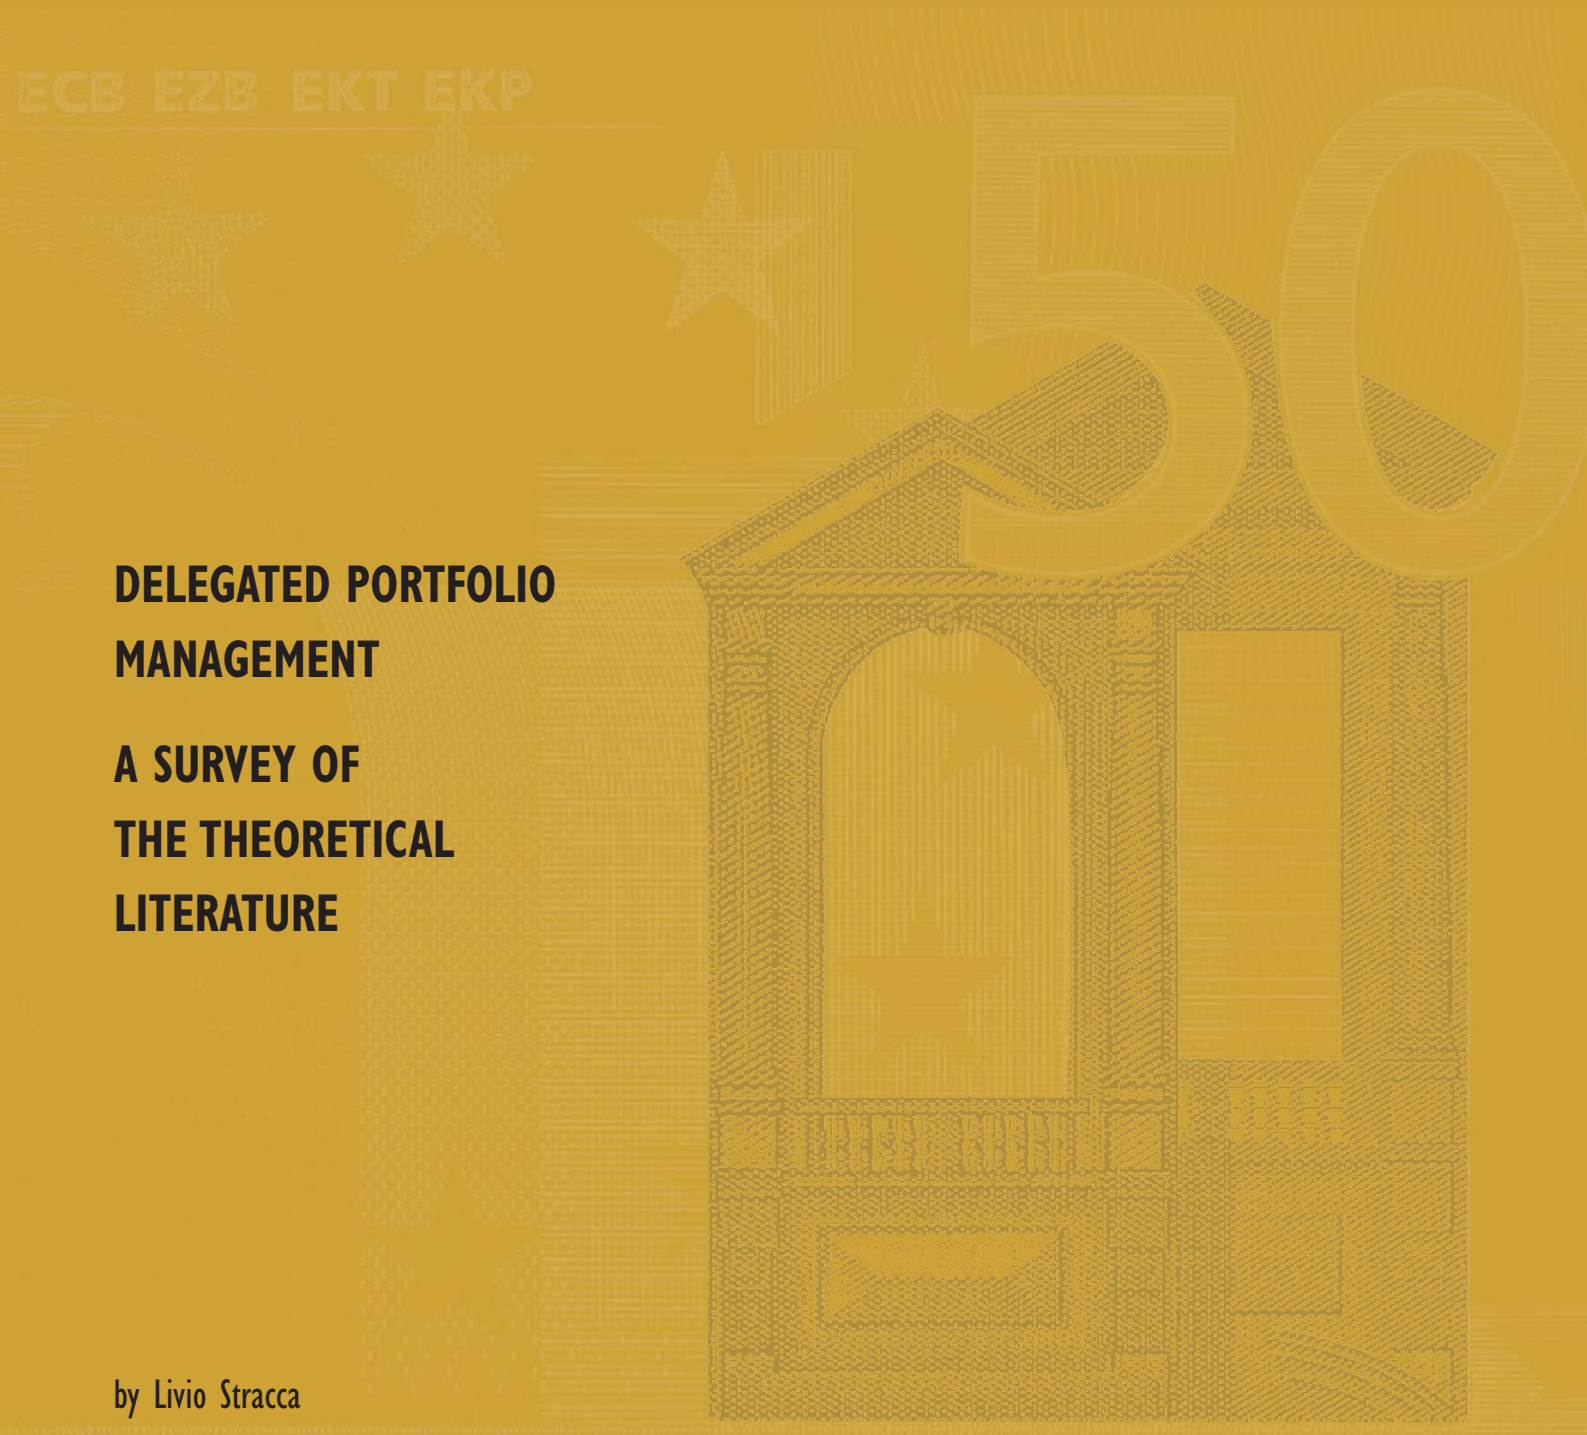

DELEGATED PORTFOLIO MANAGEMENT

A SURVEY OF THE THEORETICAL LITERATURE

by Livio Stracca

EUROPEAN CENTRAL BANK

WORKING PAPER SERIES
NO. 520 / SEPTEMBER 2005

DELEGATED PORTFOLIO MANAGEMENT

A SURVEY OF THE THEORETICAL LITERATURE [1]

by Livio Stracca [2]

In 2005 all ECB publications will feature a motif taken from the €50 banknote.

I The views expressed in this paper are only those of the author and are not necessarily shared by the European Central Bank. I thank Charles Goodhart for encouraging me to look at some of the issues dealt with in the paper. I also thank three referees for several useful suggestions.

2 European Central Bank, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany; e-mail: livio.stracca@ecb.int

Address
Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Internet
http://www.ecb.int

Fax
+49 69 1344 6000

Telex
411 144 ecb d

ISSN 1561-0810 (print)
ISSN 1725-2806 (online)

CONTENTS

Abstract

This paper provides a selective review of the theoretical literature on delegated portfolio management as a principal-agent relationship. The main focus of the paper is to review the analytical issues raised by the peculiar nature of the delegated portfolio management relationship within the broader class of principal-agent models. In particular, the paper discusses the performance of linear vs. nonlinear compensation contracts in a single-period setting, the possible effects of limited liability of portfolio managers, the role of reputational concerns in a multiperiod framework, and the incentives to noise trading. In addition, the paper deals with some general equilibrium dimensions and asset pricing implications of delegated portfolio management. The paper also suggests some directions for future research.

Keywords: Delegated portfolio management, agency, principal-agent models, adverse selection, moral hazard.

JEL codes: D82, G11

Non-technical summary

Recent decades have witnessed a sharp increase in the institutionally managed savings, both in absolute terms and relative to household financial wealth. As a result, institutional ownership is an increasingly dominant feature of developed financial markets. Delegated portfolio management is a complex phenomenon which encompasses different segments. The mutual fund industry is predominantly characterised by middle aged households investing individually in sometimes relatively standardised products. By contrast, pension funds are predominantly managed by corporate treasures, who often delegate the asset management to a third party, thus creating an additional layer of agency.

An important stylised fact of the delegated portfolio management industry is the poor performance of active management compared with a passive benchmark. Thus, active management appears to subtract, rather than to add value. Moreover, the evidence of performance consistency (i.e. managers who have performed well in the past are more likely to perform well in the future) is mixed at best, in particular in the mutual fund industry. The difficulty in ascertaining future performance has led part of the industry to focus on standard products, notably in the pension fund industry. Part of the industry, on the other hand, has remained specialized and displays a high instability of market shares and a distinct lack of concentration. It is mainly for this segment of the industry that issues of asymmetric information between the investor and the manager come to the fore. In this respect, it does not really matter if the financial intermediary involved in portfolio management is a bank, an insurance company, and so on. What is really important is whether the relationship involves asymmetric information (notably moral hazard and adverse selection) aspects which are interesting to study.

This paper offers a selective survey and an assessment of the existing theoretical literature on delegated portfolio management. The more general literature on asymmetric information and agency relationships, as reviewed for example in a textbook such as Salanie (1996), is taken as the starting point. The features of the typical delegated portfolio management relationship, which makes it essentially different from a standard agency contract traditionally studied in the earlier literature, are then analysed more carefully.

Generally speaking, there are two main differences between a delegated portfolio management problem and a standard agency problem. First, the delegated portfolio management problem is one of information acquisition rather than of direct performance. This implies that the typical timing of a portfolio management problem implies some effort on the side of the agent in order to receive an information signal, and subsequently an (unobservable) action based on the realisation of the signal. Second, a portfolio manager can control the scale of the response to the signal,

effectively influencing both the level and the variance of the portfolio returns. In a typical agency problem, the agent controls either the return or the variance, but not both.

A main conclusion of the literature survey is that, while this peculiar form of agency relationship shares many features with a traditional principal-agent model, it also presents its own challenges. The fact that in a delegated portfolio management setting the agent controls effort and can influence risk makes it more difficult for the principal to write incentive compatible contracts which are optimal from her standpoint. In particular, it is shown how the fact that the portfolio manager can control the scale of his response to the information signals in a linear (and potentially also nonlinear) way makes the quest for an optimal linear (and perhaps also nonlinear) contract for the principal very difficult. Indeed, this is a literature where negative results tend to prevail over constructive ones.

Reputation concerns in a multi-period setting may also affect the incentives faced by portfolio managers and in some cases make the job of explicit incentives. At the same time, reputation concerns may also have distortionary effects insofar as they may lead managers to take on more risk or to discard private information and herd with the market, none of which necessarily goes to the benefit of investors. Finally, the literature has emphasised that (implicit or explicit) benchmarking might have significant implications for asset prices and volatilities from a general equilibrium standpoint.

There are several directions in which this theoretical literature could be fruitfully extended. Delegated portfolio management often implies more than one layer of agency, especially in the pension fund industry: how is this likely to affect incentives and outcomes? Another interesting extension appears to be considering less standard utility functions for principals, say shortfall risk, which may be again particularly relevant in the pension fund industry. More generally, gaining a better understanding of the general equilibrium implications of the agency aspects of delegated portfolio management should be a paramount objective in future research. This is, in particular, a topic which should be interesting and relevant for policy-makers, given the importance of delegated portfolio management relationships in all developed financial markets.

1 Introduction

In most industrialised countries, a substantial part of financial wealth is not managed directly by savers, but through a financial intermediary, which implies the existence of an agency contract between the investor (the principal) and a portfolio manager (the agent). Therefore, delegated portfolio management is arguably one of the most important agency relationships intervening in the economy, with a possible impact on financial market and economic developments at a macro level. Although there are no harmonised data across countries, the general view is that the trend towards delegated portfolio management has not been interrupted by the increased direct accessibility to financial markets witnessed in recent years, for example through the internet. Davis and Steil (2001) report that the share of household wealth managed by financial institutions has increased sharply in recent decades, in particular in the Anglo-Saxon countris but also in Europe and Japan. The growth of institutional assets has been particularly visible in relation to pension funds. These developments suggest that gaining a deeper understanding of the nature and consequences of delegated portfolio management contracts is interesting and relevant, for academics and policy-makers alike (as an indication of the interest by the latter, see for example the recent survey by the Bank for International Settlements, 2003).

Against this background, the original contribution of this paper is to offer a selective survey and an assessment of the existing theoretical literature on delegated portfolio management. In particular, we concentrate on the peculiar nature of delegated portfolio management as an agency relationship. Therefore, we do not discuss – and take for granted – the more general literature on asymmetric information and agency relationships as reviewed for example by Sappington (1991) or in a textbook such as Salanie (1996). Rather, we concentrate on the features of the typical delegated portfolio management relationship which makes it essentially different from a standard agency contract traditionally studied in the earlier literature, such as, for example, the classic sharecropping model of Stiglitz (1974). Moreover, our focus is on the theoretical contributions to the literature and we will touch upon the empirical literature on portfolio delegation only occasionally.

The following key elements of a traditional agency setting provide the essential traits against which we discuss the peculiar features of delegated portfolio management contracts. There is a principal, say a landlord, who wishes to delegate the management of her property to an agent, say a tenant. The agent is better informed than the principal before the contract is signed about its own quality (adverse selection), or can acquire superior information (and therefore deliver a better result

than the principal acting alone) through costly effort after the contract is signed, which is unobservable for the principal (moral hazard). Nonetheless, the outcome does not depend only on the agent's effort, but also on environmental factors outside his control. The task of the landlord is to design a contract which encourages the best tenants to participate and, at the same time, gives the tenant the right incentives to work hard to achieve the result after the contract is signed.

If the agent is risk neutral, the best contract makes the agent the residual claimant of the contract and the principal is paid a flat fee by the agent. These performance-related contracts normally ensure that 'good' agents enter the contract and charlatans are left out, and at the same time provides the best incentives to the agent for expending costly effort. Moreover, if the principal is risk averse, the contract is also optimal for him from the standpoint of risk minimisation.

If, more realistically, the agent is risk averse and has no access to credit markets on the same terms as the principal, there is an obvious trade-off between inducing effort (which requires the agent to be exposed to the risky outcome) and providing insurance (which goes in the opposite direction). For this situation, the literature has reached two main results. First, it may be convenient for the principal to isolate the agent from the factors on which he has no control, for example by defining performance in terms of difference from a benchmark catering for a common environmental shock which affects all agents indistinctly. In this way, the principal may save on the remuneration she has to provide to the agent, at the risk however of exposing himself entirely to the effect of the common shock. Second, under reasonably general assumptions, a linear sharing rule is an optimal contract as it induces an optimal trade-off between risk-sharing and effort inducement (Holmstrom and Milgrom, 1987). Hence, the optimal contract has a linear form in which the agent is paid a flat fee plus a share of the outcome, possibly defined in terms of a spread against a certain benchmark value. The share is optimally set taking into account the relative risk aversion of the principal and the agent, as well as the need to motivate the agent to undertake costly effort.

After this admittedly oversimplified description of a standard agency setting, we set out to discuss how the basic features outlined above can be extended to the particular form of agency relationship which is the delegated portfolio management contract. The paper is structured as follows. In Section 2 we present the standard framework of the problem, and we then deal with the optimal compensation contract in a single-period set-up in Section 3. In Section 4 we extend the discussion to deal with a multi-period setting in which reputation concerns play an

important role. Some general equilibrium considerations, including the possible impact of agency on financial market prices and volatility, are in Section 5. Finally, Section 6 concludes.

2 A benchmark single-period setting

2.1 Some stylised facts about delegated portfolio management

Recent decades have witnessed a sharp increase in the institutionally managed savings, both in absolute terms and relative to household financial wealth (Davis and Steil, 2001; BIS, 2003). As a result, institutional ownership is an increasingly dominant feature of developed financial markets.

Delegated portfolio management is a complex phenomenon which encompasses different segments. The mutual fund industry is predominantly characterised by middle aged households investing individually in sometimes relatively standardised products. By contrast, pension funds are predominantly managed by corporate treasures, who often delegate the asset management to a third party, thus creating an additional layer of agency. As emphasised by Lakonishok, Shleifer and Vishny (1992b), corporate treasures often make recourse to investment counsellors for reasons which go beyond the optimization of asset allocation. Non-economic factors such as hand-holding and generally direct interaction are likely to play an important role in the pension fund industry, while funds allocation is more based on past performance in the mutual fund industry.

Another important stylised fact of the delegated portfolio management industry is the poor performance of active management compared with a passive benchmark (Malkiel, 1995; Gruber, 1996). Thus, active management appears to subtract, rather than to add value. Moreover, the evidence of performance consistency (i.e. managers who have performed well in the past are more likely to perform well in the future) is mixed at best, in particular in the mutual fund industry (Berk and Green, 2004). The difficulty in ascertaining future performance has led part of the industry to focus on standard products, notably in the pension fund industry (Lakonishok, Shleifer and Vishny, 1992b). Part of the industry, on the other hand, has remained specialized and displays a high instability of market shares and a distinct lack of concentration. It is mainly for this segment of the industry that the theoretical literature reviewed in this paper has something interesting to say, given that non-standardised portfolio management raises obvious issues of asymmetric information between the investor and the manager. In this respect, it

does not really matter if the the financial intermediary involved in portfolio management is a bank, an insurance company, and so on. What is really important is whether the relationship involves asymmetric information (notably moral hazard and adverse selection) aspects which are interesting to study.

2.2 The standard setting

Generally speaking, there are two main differences between a delegated portfolio management problem and a standard agency problem. First, the delegated portfolio management problem is one of information acquisition rather than of direct performance. This implies that the typical timing of a portfolio management problem implies some effort on the side of the agent in order to receive an information signal, and subsequently an (unobservable) action based on the realisation of the signal. Second, a portfolio manager can control the scale of the response to the signal, effectively influencing both the level and the variance of the portfolio returns. In a typical agency problem, the agent controls either the return or the variance, but not both.

The seminal contribution to the literature on delegated portfolio management relationships is due to Bhattacharya and Pfleiderer (1985), who proposed a model where a better informed agent must be solicited to reveal superior information on the rate of return on a risky financial asset to the principal. Bhattacharya and Pfleiderer find that a quadratic contract penalising deviations between the ex post realisation of the return and the signal released by the agent to the principal ensures a truthful revelation.

Bhattacharya and Pfleiderer's model is one of hidden information (security analyst framework), but the literature on delegated portfolio management which has evolved since then often proposes model in which there is also hidden action, i.e. an information asymmetry also after the contract has been signed. In fact, most portfolio managers directly manage the financial portfolios of their clients, rather than provide them with information which they can then use in trading directly in the financial market. Transaction costs and other frictions often make it inconvenient or even impossible for clients to trade directly.

The standard delegated portfolio management problem involves a game between an investor, the principal (henceforth P and "she") and a portfolio manager, the agent (henceforth A or "he"). Both can be either risk averse or risk neutral (the principal is more often assumed to be risk neutral). The basic problem can be illustrated in a simple financial market where there is one risky asset and a riskless asset with rate of return normalised to zero, with both P and A price-taker, i.e. having

no impact on asset returns. To illustrate the (Rotschild-Stiglitz) adverse selection aspect of the relationship, we assume that A can be either a 'good' agent – who can have private information about the probability distribution of the rate of return of the risky asset – or a 'bad' agent – who has no access to such information. It is also normally assumed that A has no access to insurance markets, nor can he invest in financial assets on his own account (for example in order to hedge its compensation). This assumption is required to study the incentives provided to A by the compensation contract in itself in a tractable manner, abstracting from the complexities of optimal hedging against a stochastic performance evaluation. This, however, does not necessarily make this assumption very realistic.

The good type of A can learn about a signal S which is correlated with the return on the risky asset upon expending an effort e which costs him a disutility $h(e)$. For simplicity, we are assuming that the variable e can take values 1 (effort is spent) or 0 (no effort). If the agent expends effort, he has therefore access to the *conditional* probability distribution of the return on the risky asset.

Let the conditional return on the risky asset, y/S, be determined in the following linear way:

$$y/S = \mu_u + S + v, \tag{1}$$

where v is an environmental shock which is uncorrelated with S and e, and which has variance σ_v^2, and μ_u is the unconditional mean of y (i.e. without observing S). This specification captures the notion that the return on the risky asset depends on both the signal observed by A, which is under his control, and an environmental shock, which is not.

Based on this information, A selects the weight of the risky asset in the portfolio, α_c, and consequently the rate of return on the conditional portfolio is:

$$x_c = \alpha_c(S)y = \alpha_c(S)(\mu_u + S + v) \tag{2}$$

For simplicity of illustration, we assume here that both P and A have a mean-variance utility function. The optimal conditional portfolio, α_c, is then the portfolio which maximises the Sharpe ratio.

For P and for the bad type of A, S is not observable and the probability distribution of y is the unconditional one:

$$y = \mu_u + q, \tag{3}$$

where q is a (for P and the bad A) unpredictable component.

The return on the unconditional portfolio is:

$$x_u = \alpha_u y, \tag{4}$$

where α_u is the optimal unconditional mean-variance portfolio.

In the standard specification of the game, P generally proposes the compensation contract[1], then A expends effort, learns the signal S (if he is of the good type) and either reveals it to P or (more commonly) makes the portfolio choice directly. Subsequently, the shock v is realized, and A is paid a remuneration $f(x)$, while the rest of the portfolio return goes back to P.

Most models assume an asymmetric information setting in which P cannot monitor the activities of A. While this is realistic with a view to monitoring costs, there is nevertheless a variety of ways in which P can gain at least a partial and imperfect monitoring of A's activities and this might be alleviate the costs of the agency relationship (Almazan et al, 2004). This possibility, and the interesting trade-off which arises between the monitoring costs and the costs related to the information asymmetries in the agency relationship is, however, normally not addressed in the theoretical literature.

Let EU_P be the expected utility of P. In this framework, by construction, the expected utility of the conditional portfolio is always higher than that of the unconditional one:

$$EU_P(1 + x_c) > EU_P(1 + x_u) \tag{5}$$

However, P has also to pay a stochastic fee $f(x)$ to A in order to accede to the conditional portfolio, and we turn to discuss the role of this variable in the next section.

2.3 Functions of the compensation contract

Broadly speaking, a compensation function $f(x)$ serves three main purposes. First, it may have a *risk-sharing function* if P and A are risk averse. Second, it will affects the *portfolio selection* of A, including the incentives it provides to expend effort and learn about the signal S, and therefore the distribution of returns. Third, it might send a *signal about the type* of agent participating in the contract. In particular, the compensation contract is *separating* if it reveals the type, and *pooling* otherwise.

From a conceptual standpoint, the optimal compensation contract is the solution to the following decision problem:

$$Max_f EU_P(1 + x - f(x)) \tag{6}$$

$$x = \alpha(S)y(e) \tag{7}$$

[1] Although there are exceptions in the literature, for example the signalling model in Das and Sundaram (2002).

$$e, \alpha(S) = \arg\max[EU_A f(x) - h(e)] \tag{8}$$

$$EU_A f(x) - h(e) \geq U_A^* \tag{9}$$

$$EU_P(1 + x - f(x)) \geq EU_P(1 + x_u) \tag{10}$$

Equation (6) states that the compensation contract has to be chosen so as to maximise the expected utility of P. Equation (7) describes the probability distribution of returns, which hinges on whether the agent decides to put some effort in obtaining the information signal S. Equation (8) is the incentive compatibility constraint, which states that effort e and the portfolio allocation $\alpha(S)$ are chosen by the agent so as to maximise his own utility. Equation (9) is the participation constraint of the agent, which states that the expected utility of the agent, EU_A, has to be at least equal to his reservation utility U_A^*.[2] Finally, equation (10) is the participation constraint of P. It should be noted that here, for simplicity of illustration, we are assuming that P can invest in the financial market on the same terms as A, for example as regards transaction costs. In practice, P is likely to face higher transaction costs, and this will be reflected in a level of P's reservation utility which is lower than $EU_P(1 + x_u)$.

A remuneration contract $f(x)$ is said to be *feasible* if equations (7), (8), (9) and (10) are satisfied. Several papers also include the additional constraint that $f(x) \geq 0$, because it is often realistic to assume that A will not accept a negative remuneration at the end of the period (limited liability constraint). As we will see later, the assumption of limited liability of A has very significant implications for the nature of the contract and the incentives it provides.

A further important aspect to consider is the competitive situation in the market for portfolio managers, which is of crucial importance to determine who extracts the surplus from the agency contract. It is somewhat more common in the theoretical literature on delegated portfolio management to consider a situation in which there is perfect competition among portfolio managers and all the surplus is accrued to P as in the problem described in (6)-(10). This situation implies that equation (9) is valid with perfect equality (A is pushed to his reservation utility), while equation (10) is valid with inequality (P extracts the whole information rent). Of course, also the opposite situation (all the surplus accrues to A) is possible if there is enough competition among principals, or an intermediate case where the surplus is split between P and A. Both variants have been duly taken into consideration in the literature.

[2] Here we are assuming that the good and the bad agent have the same reservation utility, which may be unrealistic. Relaxing this assumption may have important implications on the nature of the problem, as will become apparent later on.

3 The optimal compensation contract

3.1 The first best case when effort is observable

If the effort of A to acquire information is observable and there is no information asymmetry before signing the contract, then principal and agent act as an aggregate entity. Therefore, the optimal amount of effort is determined at the level where the aggregate marginal benefit for both P and A is equal to the marginal cost for A. This is usually referred to as the *first best case*.

In the first best case, the delegated portfolio management contract is reduced to one of optimal risk sharing between P and A, if they are risk averse (Stoughton, 1993). Under the assumption of preference similarity and constant absolute risk aversion, the optimal payment $f(x)$ is a linear sharing rule (Wilson, 1968), plus a fixed payment dependent on effort. In particular, the optimal sharing rule requires that each individual receives a fraction of the return on the risky asset equal to his risk tolerance divided by the aggregate social risk tolerance.

3.2 Second best contracts when effort and risk are unobservable

In a realistic setting, however, there is information asymmetry both before and after signing the contract, which makes the first best contract unattainable.

As regards the moral hazard dimension of the problem, the agency literature has emphasised that a linear sharing rule is no more first best since it leads to an underinvestment in effort, due to the inability of A to internalize the benefits accruing to P from the expenditure of effort. Nonetheless, a linear contract might remain optimal in a second best sense since it may lead to an optimal compromise between effort inducement and risk sharing (Holmstrom and Milgrom, 1987; Sappington, 1991).[3]

The typical linear compensation contract is specified as follows:

$$f(x) = C + B(x - b), \tag{11}$$

where $C, B > 0$ and b is a benchmark value, known in advance and which can be fixed or stochastic. In a standard agency setting – say a landlord-tenant relationship – the crucial parameter is B, while C is

[3] In particular, the linear form of the contract can be obtained if the principal is risk neutral or risk averse with a CARA utility function, the agent has a CARA utility function, and the shock is additive and Normally distributed (Holmstrom and Milgrom, 1987).

chosen so as to satisfy A's participation constraint.[4] Crucially, a higher value of B pushes A to expend more effort, but also to take on more risk which will increase the required flat payment C.

However, the situation is substantially different in a delegated portfolio management contract. As we will see shortly, the linear contract cannot be optimal in a general setting of the problem, nor can B used to motivate the agent to expend effort as in a standard agency contract. More fundamentally, a compensation contract which is optimal (from a second-best perspective) in a general class of delegated portfolio management problems is not known even under the assumptions of Holmstrom and Milgrom (1987). Generally speaking, the literature has reached more negative rather than constructive results, and the search for an optimal contract has proved to be inconclusive even in the most simple settings.[5]

3.2.1 The irrelevance result

As noted, it is a standard conclusion of agency models that increasing the share of output faced by A makes him work harder to the benefit of P. A main feature of the delegated portfolio management relationship is that this basic force does not work, a conclusion known as the *irrelevance result* and which is due to Stoughton (1993) and Admati and Pfleiderer (1997). So, A always behaves as he was acting only on his own, and cannot be pushed (by the compensation contract) to internalise any positive externality from his behaviour onto the welfare of P.

Intuitively, the reason behind the irrelevance result is that in a delegated portfolio management contract A has both the incentive and the possibility to undo the incentive effect of the linear contract. In fact, as explained in detail by Admati and Pfleiderer (1997), in a delegated portfolio management setting such as the one described in Section 2 A has complete control over the *scale* of his response to the signal. Assume a simple agency setting where output, x, is determined by a variable controlled by A, S, and an error term v:

[4] Note that the choice of b will matter in the determination of C. For example, if b is fixed in advance, the flat component of the fee is, in fact, $C - Bb$.

[5] A notable exception is a recent paper by Ou-Yang (2003). Ou-Yang proposes a continuous time model where A, the manager, decides how much to invest in bonds and stocks and continuously adjusts his portfolio positions. The optimal contract is derived in closed form and is found to be symmetric. The manager should be paid a fixed fee, a fraction of the total assets under management, plus a bonus or a penalty depending upon the excess return on the portfolio compared with an appropriate benchmark. However, in Ou-Yang's model A has no superior information compared with P, and trades only because of lower transaction costs. This rules out most of the interesting information asymmetry between P and A which is discussed in this paper.

$$x = \alpha(S + v), \tag{12}$$

where α is a positive scalar. If the agent cannot control α, this is a standard agency problem (say, a landlord-tenant relationship) where, in a linear contract, the parameter B can be used to motivate the agent to work hard. However, if the agent has complete and costless control of α, the linear compensation contract becomes:

$$f(x) = C + KS + Kv, \tag{13}$$

where $K = B\alpha$ is a real scalar completely controlled by A and which can be varied at his choice after the realisation of the variable S. Therefore, the variable part of the compensation depends on K (controlled by A) and not on B (controlled by P), which consequently cannot be used to motivate effort. A key conclusion of Stoughton (1993) and Admati and Pfleiderer (1997) is that a delegated portfolio management relationship resembles this second type of agency problem. This implies that linear contracts tend to lead to underinvestment of effort by A, which is detrimental to P.

Admati and Pfleiderer (1997) note that this result holds irrespective of whether the contract contains a benchmark b different from zero, which is evident from expression (13). So, the idea that including benchmarks in the compensation contract may motivate agents to undertake costly effort is misplaced, at least when using a linear compensation contract and without making very restrictive assumptions. Moreover, Admati and Pfleiderer also emphasise that benchmarks do not play any useful role in risk-sharing in delegated portfolio management, since optimal risk-sharing should always be on the *total* return of the risky portfolio, not on the return *relative* to a benchmark. As regards the possible advantage often mentioned in the agency literature that benchmarks can serve as screening devices in order to keep bad agents away from signing the contract (adverse selection), a benchmark b is certainly not an optimal way to assess performance, unless the benchmark is the optimal conditional portfolio.[6] Overall, Admati and Pfleiderer conclude that the benchmark (unless it is the optimal conditional portfolio, which is generally not accessible to P) is a distorting factor and should at best be ignored in a delegated portfolio management contract.

A conclusion of this analysis is that moral hazard appears to be an important reason speaking for a recourse to non-linear contracts in

[6] Moreover, in a linear contract the flat fee parameter C can be used to keep bad agents away from the contract as well as, if not better than, measuring performance in terms of the spread of x with a benchmark b.

delegated portfolio management relationships, and that a linear contract cannot be used to motivate the agent due to his ability to undo its incentive structure by varying the scale of the response to the signal in a linear way.[7]

For example, a quadratic contract can take the following form:

$$f(x) = C + B(x - b)^2 \tag{14}$$

Based on the idea that a quadratic contract might succeed where the linear contract fails, Stoughton (1993) has shown that, in a security analyst framework (where P trades for herself in a direct revelation mechanism conditional on the report of a signal by A), the quadratic contract of Bhattacharya and Pfleiderer (1985) (where $B < 0$) is able to elicit truthful elicitation of the signal observed by A. Moreover, Stoughton finds that a higher B (in absolute value) does provide the agent with incentives to undertake effort. In addition, the quadratic contract is able to reach the first best if P is risk neutral (in fact, the quadratic contract is not optimal from a risk sharing perspective, but this is not important for a risk neutral P).

It should be emphasised that, apart from the quite restrictive assumption that P trades on herself (which is not very realistic in today's financial markets), Stoughton's quadratic contract does not solve the irrelevance problem if A is allowed to respond in a *nonlinear* way to the signal. More generally, it is doubtful that nonlinear contracts may solve the irrelevance problem highlighted by Stoughton (1993) and Admati and Pfleiderer (1997) unless there are restrictions imposed on A's trading set. Consider, in fact, a nonlinear compensation scheme $f(x)$ and the simple agency problem described above, where A has a complete and costless control of a nonlinear reaction function to the signal S, $g(S+v)$. We obtain:

$$f(x) = f(g(S + v)) \tag{15}$$

It is immediate to conclude that if $g = f^{-1}$ is within the trading possibilities of A, the irrelevance result still holds. So, A's incentives to

[7] It is important to note, however, that the irrelevance result is weakened once we allow for short-selling bounds on α. Notably, Gomez and Sharma (2003) find that linear performance-adjusted contracts do provide managers with good incentives for gathering better information. This is quite intuitive since bounds on α do not allow A to fully undo the incentive structure of the linear contract by setting $\alpha = \frac{K}{B}$ (see eq. (13)) for any signal S received. Moreover, Gomez and Sharma find that with short-selling bounds A's share in the portfolio return is higher than the first best and decreases as the leverage constraint is relaxed. In addition, Ozerturk (2004) has shown that if A is a large trader who has an impact on asset prices, linear contracts do give A the incentives to expend costly effort. The irrelevance result only holds when A is completely price-taker in the financial market.

expend effort in order to increase S will be entirely self-serving with no consideration of the positive externality on P's utility.[8]

3.2.2 Linear symmetric vs. nonlinear asymmetric contracts

In the absence of a generally accepted optimal compensation contract valid in all settings, some contributions in the literature have investigated the property of specific contracts, normally without making any claim of general optimality. An important issue dealt with in the literature is whether the compensation contract should be symmetric (penalising negative outcomes in the same way as it rewards good outcomes) or asymmetric and convex (i.e., where the *marginal* reward is higher, the better the performance compared with a benchmark). In this respect, it should be recalled that in the United States the SEC prohibits asymmetric and nonlinear compensation contracts in the mutual fund industry, i.e. contracts which are not linear and symmetric around a prespecified benchmark, although recent changes in the regulation have weakened the prohibition to some extent.

Starks (1987) is an important paper in this respect as it analyses the relative desirability of a symmetric, fulcrum performance fee compared with an asymmetric bonus contract. The main finding of the analysis is that the symmetric performance fee is to be preferred to the bonus contract because it can at least align A's attitude towards risk to that of P (although, as we have seen, it does not solve all the agency problems in a delegated portfolio management relationship).

As noted by Garcia (2001), symmetric linear contracts or, more frequently, contracts with no explicit performance-related compensation are typical of the (regulated) mutual fund industry, while in the (unregulated) hedge fund industry call-option type contracts are widespread (Goetzmann, Ingersoll and Ross, 1998). Call-option compensation contracts are also popular in corporate finance mainly because they are thought to make managers (e.g., CEOs) less conservative, which should ultimately benefit shareholders. On the other hand, it is generally not

[8] This possibility is indeed hinted at in the conclusions of Stoughton's (1993) paper, also suggesting that limiting the recourse of portfolio managers to complex and nonlinear financial products such as derivatives might be desirable from an agency standpoint. Dybig, Farnsworth and Carpenter (2004) and Almazan et al (2004) also justify the recourse to contraints on A's trades as part of the portfolio delegation contract on similar grounds. Interestingly, a study of the Bank for International Settlements (2003) reports an increased use of constraints on investment, such as limits on investing in specific securities, in the mutual fund industry (though Almazan et al report the opposite conclusion). The real question, however, is whether such constraints are enforceable in practice.

known (even in a general principal-agent setting) whether there are situations in which these compensation contracts may be optimal.[9]

A call-option compensation contract has the following expression:

$$f(x) = \begin{cases} C + B(x-b) \text{ if } x \geq b \\ 0, \text{ if } x < b \end{cases} \tag{16}$$

In fact, the conviction that option-like contracts make agents more risk-loving is only partially founded, as discussed by Chen and Pennacchi (2002) and Ross (2004). Indeed, without further conditions on utility functions of A, it is not correct to conclude that giving options to managers makes them more risk-loving (and more willing to work hard). The overall effect of the compensation scheme has to be evaluated taking into account both the utility function U_A as well the functional form of the compensation contract, so the *derived* utility of the agent $U_A(f(x))$. So, a convex compensation scheme does not necessarily convexify A's utility function, nor a concave compensation schedule necessarily concavify it and makes A more risk averse.

Ross (2004) demonstrates that there are indeed two additional effects to consider beyond the fact that the compensation schedule is convex (*convexity effect*). First, the *translation effect* shifts the evaluation of any risky bet to a different portion of the domain of the utility function. In practice, if A has an increasing absolute risk aversion, he might behave in a more risk-averse manner the more asymmetric the compensation contract is compared with his initial level of wealth. This happens because with a call-option compensation A evaluates the lottery in a portion of the domain of the utility function where he is more risk averse. Second, the *magnification effect*, which depends on the size of the parameter B, raises the risk of the gamble for A and therefore lowers his willingness to take it on.[10] A consequence of this analysis is that also in a delegated portfolio management context it is not a foregone conclusion that granting the portfolio manager a call option should increase his willingness to bear risks.

Carpenter (2000) proposes a model providing a rigorous description of the optimal investment choice of a risk averse manager who is compensated with a call-option contract such as that in (16). On the one

[9] Garcia (2001) attempts at providing a theoretical justification of the optimality of call-option contracts, showing that they may be an optimal way (at least among linear contracts) to let fund managers with low prospects get a fixed wage, and managers with high prospects get a contingent compensation. Garcia's model crucially depend on the assumption that A's participation constraint has to be satisfied *ex post* and not only *ex ante*.

[10] Ross (2004) shows that even for a manager with a CARA utility function (for which there is obviously no translation effect) the magnification effect alone is able to offset the convexity effect.

hand, the convex nature of the compensation encourages A to look for portfolio returns that are very away from the money (the manager either significantly outperforms the benchmark or incurs in severe losses), which increases return variance. This tendency is aggravated for options which are deep out of the money. On the other hand, A moderates the return variance if the value of the asset portfolio is large enough in order to lock in his gains, owing to his risk aversion. An interesting result – also in the light of the subsequent analysis of Ross (2004) – is that an option-like compensation contract might under certain conditions lead A to choose a level of return volatility which is *below* the one he would choose when trading on his own. In the terminology of Ross (2004), this result appears to be related to the magnification effect implicit in the leverage of the option contract. Overall, Carpenter concludes that the effects of an option-like compensation on the manager attitude towards risk are more complex than simple intuition might suggest.

Das and Sundaram (2002) explicitly compare the symmetric contracts dictated by the SEC ('fulcrum' contracts) with an asymmetric, convex one ('incentive' contract). In a signalling model including a good (informed) and a bad (uninformed) A, Das and Sundaram find that, quite intuitively, incentive fees lead to the adoption of more risky portfolios than fulcrum fees. So, fulcrum fees provide a better risk-sharing than incentive fees, and reduce A's incentive to choose overly risky portfolios. However, they also find that equilibrium investor welfare may be strictly higher under incentive fees than fulcrum fees if the 'bad' (uninformed) agent has a reservation utility (U_A^*) which is low enough. In fact, an asymmetric contract may lure uninformed agents in the business of delegated portfolio management who would otherwise stay out of the market were the fees of the symmetric type. Because fulcrum fees make it easier for the good type of A to separate himself from the bad one, they make mimicking by the uninformed agent more expensive. To the extent that a pooling contract is desirable for P and that the reservation utility of the good agent is high enough, it may be preferable to offer a contract which makes separation less easy for good agents and the incentive contract is preferable to the fulcrum contract in this respect.

Clearly, however, in a competitive market for portfolio managers where the reservation utility of the good A is pushed to zero and all the surplus is accrued to P, there is no need to make separation less easy for good agents, and there is consequently no reason to consider incentive fees, which are strictly dominated by fulcrum fees.

3.2.3 Possible effects of limited liability

An asymmetric and convex nature of the compensation contract may also arise not as an explicit and built-in feature of the contractual relationship, but rather as a *de facto* consequence of limited liability of A. This is very realistic in a delegated portfolio management setting, where for example the manager can sabotage the portfolio ex post (Palomino and Prat, 2003). Under limited liability of A, a linear symmetric contract can de facto become an option contract like the one in (16).

This possibility has attracted some attention in the literature due to its empirical plausibility. Grinblatt and Titman (1989) have showed in a general agency context that if there is limited liability, the agent has an incentive to take on a riskier portfolio than otherwise (see also Gollier, Koehl and Rochet, 1997). Hence, the relevant question is whether, due to limited liability, A has an incentive to take on too much risk, and whether P can do anything (in the structure of the compensation contract that he proposes) to offset this incentive, especially if P is risk averse. It should be emphasised, however, that what it is claimed here is that A is more risk-loving under limited liability *than otherwise*, not that he is risk-loving in general or that he necessarily takes on more risk than it is desirable for P (see the discussion in the previous section on convex compensation schemes).

Rajan and Srivastava (2000) demonstrate that if A's utility function is unbounded (for example, A is risk neutral), the compensation function must be bounded, otherwise A would prefer to take an infinite amount of risk, and the delegated portfolio management problem would have no solution. Therefore, a linear contract is fundamentally incompatible with limited liability. The alternative considered by Rajan and Srivastava is the *bonus contract*, whereby:

$$f(x) = C, \text{ if } x \geq b, \tag{17}$$

and zero if $x \leq b$, where b is a benchmark which plays the role of a threshold level. Assume that portfolio allocation takes place in a stochastic context where a number of potential states of the world lead to different portfolio returns. Rajan and Srivastava prove that the optimal portfolio choice by A consists of a short position in *one* state of nature and a long position in all other states (which they label the 'bang-bang' portfolio). P is able to influence the level of risk through the choice of the threshold b, subject to A's participation constraint. In this context, a main result of Rajan and Srivastava is that financial innovation – which allows A to take even more aggressive gambles by enlarging the menu of available assets – may be detrimental to P as it induces excessive risk-taking by A for any given threshold b.

Also Palomino and Prat (2003) find that, with limited liability, there is no optimal linear contract. Therefore, they analyse the bonus contract and model the interaction between a risk-neutral P and a risk-neutral A. They find that, under certain assumptions, the set of optimal contracts contains a bonus contract.[11] The threshold b above which the agent receives the bonus is determined as a compromise between the need not to induce a too high risk level (which calls for a lower b) and the need to reduce the rents for A by making his job too easy (which calls for a higher b). Palomino and Prat conclude that if it is cheap for A to play with risk (e.g. high-powered financial instruments are available) we should expect a low threshold b and an inefficiently conservative behaviour by A.

Another element which can be induced by A's limited liability in a delegated portfolio management setting and which has attracted some attention in the literature is *churning* or *noise trading*, i.e. trading in the financial market without any particular reason (such as superior information or hedging motives) for doing so. The seminal paper on this topic is Allen and Gorton (1993), who assume a competitive industry for managers where good managers are in short supply and bad managers are around. Because of limited liability, agents are (de facto) paid with a call option which attracts charlatans and induces risk-loving behaviour by them. So, these agents just churn in the market even though they lose on average to the good managers. As a result of this behaviour and under certain assumptions, the agency contract may create 'churning bubbles' and mispricing.

In contrast with Allen and Gorton (1993) where the *bad* agents churn, Dow and Gorton (1997) present a model where churning is consequence of the actions by the *good* agents. In a moral hazard context, because of asymmetric information, P cannot distinguish 'actively doing nothing' by A (managers do not trade because no mis-priced security is available) from 'simply doing nothing' (managers do not look for mis-priced securities and just sit idle). In a context in which this makes it impossible to reward inactivity and if limited liability prevents punishing ex post incorrect decisions, the contract may induce the good portfolio manager to churn.[12] In particular, good managers churn because this may be a

[11] It should be noted that a bonus contract would not be optimal (even under the assumptions set out by Palomino and Prat) if at least one of the two parties were not risk neutral. In fact, a bonus contract does not provide an optimal risk sharing, and with an insurance motive the optimal contract will tend to be smoother than the bonus contract.

[12] Theoretically, a contract involving a large bonus for taking correct trading decisions (such as that proposed by Palomino and Prat, 2003) and a small lump-sum payment for inactivity (which, if sufficiently large, would encourage good agents to work and bad agents to stay away from the contract) would solve the problem. Dow

valid way to signal their talent and superior knowledge to the principals, given that it is difficult to distinguish talent and effort just in the portfolio return, especially at relatively short evaluation horizons such those which prevail in the mutual fund industry. Dow and Gorton also stress that for P noise trading lowers the expected rate of return on the portfolio. On the other hand, churning may benefit hedgers and therefore reduce the cost of hedging, which may have a beneficial impact on the functioning of the market as a whole by increasing its liquidity.

4 Multiperiod games and reputation

4.1 The role of reputation as an implicit incentive

There is an adage in the mutual fund industry that "the real business of money management is not managing money, it is getting money to manage". The underlying idea is that new money is expected to flow into the fund if the manager has performed well relative to a certain point of reference (such as the average performance of his competitors). Since compensation contracts usually include a fee which is a share on the assets managed, this creates an implicit incentive for managers to perform well which adds to the explicit incentives set out in the compensation contract.[13] In the real world, delegated portfolio management is a fundamentally multi-period game even if the vast majority of contracts have a relatively short horizon (Goldman and Sleazak, 2003). In this situation, career concerns of portfolio managers related to the willingness to attract new funds in the future may motivate them to undertake costly effort, which creates information rents for the principals in the same way as explicit performance-based compensation. More in general, reputation may have a significant impact on the incentives faced by agents in a delegated portfolio management setting.

In this respect, it is striking to note that there appears to be little explicit relation between fees and performance in the mutual fund industry (Lakonishok, Shleifer and Vishny, 1992a; Bank for International Settlements, 2003). In fact, most contracts do not even have a performance-based fees component, and the compensation is simply a fixed fraction of the assets managed (Heinkel and Stoughton, 1994).[14] This phenomenon

and Gorton assume that this compensation contract is not available or not feasible. See Bhattacharya (1999) for a criticism of the assumptions of Dow and Gorton and of the conclusion that churning is an unavoidable consequence of limited liability.

[13] The idea of implicit incentives in the principal-agent literature is well known at least since Fama (1980).

[14] However, the unit fee is normally a concave function of assets managed, which might be related to the existence of fixed costs for managers. Moreover, Christoffersen (2001) document that fund managers often voluntarily waive their fixed fees in order

might, at least in part, be related to the difficulty of writing an optimal contract in a single-period delegated portfolio management setting, as discussed extensively in the previous section. So, implicit rather than explicit incentives appear to matter in practice. This has led to some attention in the literature on the possible alternative incentives provided by reputation concerns.

In addition, several papers have pointed to the fact that there seems to be a fundamental asymmetry in the payoffs related to reputation concerns. Sirri and Tufano (1998), among others, have reported that net investment flows are much less sensitive to past returns when these are bad, and more sensitive when these are good, leading to an implicit asymmetric compensation pattern. Lynch and Musto (2003) provide a rationale for that pattern, by pointing to the possibility that funds respond to bad – but not good – performance by replacing the personnel or techniques that produced it. Therefore, a good outcome for returns may be more informative for the future performance of a certain portfolio manager than a bad outcome.

To rationalise the recourse to implicit, rather than explicit incentives by investors, Heinkel and Stoughton (1994) analyse the incentives for managers in a two-period model in which managers are compared to a sample of other managers, and retained only if their performance is good enough. The retention decision is assumed to be taken by the investor at the end of the first period. In their model, investors (P) cannot observe neither A's quality (adverse selection), nor his effort level (moral hazard). Heinkel and Stoughton's main finding is that the contract taken in the early stages of a client-manager relationship is not fully separating and has a weak performance-based component. This is in striking difference from the performance-based, single-period contracts which has been discussed in the literature (but which does not seem to be very popular in the industry).

Importantly, in Heinkel and Stoughton's paper P's interests are best served by providing incentives through a threat of dismissal following a performance evaluation, rather than with a large performance-based fee component in the first-period contract. This contract creates an implicit incentive for managers to perform well compared with their peers. By contrast, the second-period contract (conditional on the retention decision) has a strong performance-based fee as in any single-period contract.[15] Other papers (e.g., Farnsworth, 2003) emphasise the stronger

to improve their performance and thereby their reputation (which is important in order to get assets to manage in the future). This practice may make fixed fees more performance-related, albeit in a way which is not controlled by the investors.

[15]It should be noted, however, that Heinkel and Stoughton's two period contract

role of reputation at the beginning of a multiperiod setting. In most reputation models, reputation effects are important in the first period but die out when reputation becomes very good. In this later stage, explicit incentives must be provided substituting for the incentives previously provided by reputation concerns. It should be noted that these models, however, abstract from the transition costs of switching managers after the single period performance tests, which are likely to matter considerably in the real world.

With the above as background, two sets of issues have been dealt with extensively in the literature. The first is whether the asymmetry of implicit incentives might be the trigger of possible *excessive risk-taking* by portfolio managers. Suppose, in fact, that remuneration is given simply by a fixed fraction of the total assets under management. In this setting, there is clearly a situation of limited liability which generally encourages agents to take on more risk, as we have seen in the previous section. The second issue is whether the fact of competing with other managers might encourage portfolio managers to *discard private information and herd with the market*, depending on the precise incentive structures related to reputation and managers' degree of risk aversion. In other words, reputation concerns may lead A to lean towards either *anti-herding* or *pro-herding* behaviour, depending on his utility function and the payoff structure, and these are tendencies which do not necessarily benefit P or the overall efficiency of the financial market. We consider the first possibility in Section 4.2, and the second one in Section 4.3.

4.2 Reputation and excessive risk-taking

With a focus on the effects of reputation concerns on risk-taking, Huberman and Kandel (1993) present an adverse selection model where a manager's objective is the sum of the expected utility of the investment immediate outcome and a reward entailed by the market's inference regarding his ability. The amount of wealth A manages and the fees he receives for it depend on his reputation. Huberman and Kandel find that the portfolio weights may be used to signal the managers' ability,

is optimal under the assumption that, once A has been revealed to be of the good type, he gains all the bargaining power in the relationship (and all information rents). This is the fundamental reason why implicit incentives succeed (in the first period) where explicit incentives fail from the standpoint of P. The only way for P to gain a surplus is to make the more able type of A face a significant risk of not being retained at the time of the performance evaluation. For this reason, the pooling contract in the first period works to the mutual benefit of both the informed A and P. On the other hand, a multi-period model with increased competition among managers will become similar to a single-period model.

and this may distort the asset allocation strategy of both the low and high quality managers. In particular, under a certain parameter configuration, if only few bad managers are present in the population, a pooling equilibrium may arise, which entails exaggerated investment of both types of managers. Under an alternative parameter configuration, only the better type exaggerates his position in the risky asset compared with the optimal level, in order to deter the bad manager from imitating it (separating equilibrium). In either case, reputation concerns lead to an over-investment in the risky asset.

In the same vein, Huddart (1999) also proposes a adverse selection, two-period model. Huddart emphasises that when asset fees are a fraction of assets under management, managers will have incentives to distort their asset allocation in order to appear to be informed at the end of the first period, which leads to a suboptimal situation for P. In particular when the barriers to entry are low, uninformed managers enter in the relationship and are willing to take on high risks in order to produce a track record which suggests possession of superior information (which in reality they do not have). A main conclusion of the analysis is that the variance of fund returns are higher in the presence of a reputation effect than in a single-period setting with no reputation, and this hold both in a pooling and in a separating equilibrium.

However, Huddart also finds that explicit performance fees (e.g., symmetric and linear) mitigate such excess risk-taking by A, irrespective of whether he is good or bad. In fact, a performance fee may in part undo the effects of reputation, even when Stoughton (1993)'s irrelevance result holds as regards information acquisition effort. Since A is risk averse in Huddart's model, the possible reward in terms of getting additional funds to manage in the future if his future reputation is good has to be weighted against the cost of taking on a risky bet today. A performance fee provides exactly that risk to A, which then mitigates his risk-taking behaviour.

Several contributions in the literature have emphasised the risk-taking behaviour of mutual funds in a *tournament* setting, i.e. depending on their interim performance in the middle of the year compared with their peers (given that relative performance is normally calculated over the calendar year). In a situation where the amount of remuneration that a fund receives for winning a tournament depends upon its performance relative to other participants, theory (Basak, Pavlova and Shapiro, 2003) and empirical evidence (Brown, Harlow and Starks, 1996; Chevalier and Ellison, 1997) tend to suggest that funds which trail the benchmark at mid-year will subsequently take on more risk (and gamble for resurrection), while funds which are ahead of the competition will become

more conservative in order to secure their position. This asymmetric behaviour of portfolio managers is rational if it is due to an asymmetric structure of incentives whereby a good performance is rewarded more than a bad performance is penalised. This appears to be particularly true for young portfolio managers, for whom reputation concerns matter most (Chevalier and Ellison, 1999).

As shown by Chen and Pennacchi (2002), however, portfolio managers do not really have any incentive to increase the *overall* volatility of their portfolios in a tournament setting, but only the variance of the *tracking error*, i.e. the departure of their portfolios from the benchmark portfolio. In fact, Chen and Pennacchi show that a worsening performance (which, under certain assumptions, leads A to gamble for resurrection, as we have seen) might actually move the portfolio closer to the one which minimises absolute return variance.

In addition, ranking concerns may exacerbate the possible distortion of reputation concerns in inducing excessive risk-taking, as emphasised by Goriaev, Palomino and Prat (2001). In a two-period model with two risk-neutral competing fund managers with ranking objectives, these authors show (and find empirical confirmation that) a good (top decile) interim performance generates strong incentives to take on more, rather than less, risks in order to end the year ranked first. So, if performance is defined in terms of ranking and if this creates a "winner takes all" situation, the mutual funds with the highest chance of being the interim winner, and not the interim losers, will take on the most risk.

4.3 Reputation concerns and herding behaviour

Another possible dimension of the tournament aspect of the delegated portfolio management problem related to reputation concerns is the possibility that mutual funds tend to discard private information and herd, as suggested by Scharfstein and Stein (1990). This tendency, in particular, might be related to two aspects of the typical delegated portfolio management relationship. First, most contracts have a relatively short evaluation period, such as one year, and can be terminated at a very short notice (Goldman and Sleazak, 2003). Therefore, mutual fund managers cannot afford to wait until their private information (even if eventually correct) is revealed and incorporated in asset prices (Froot, Scharfstein and Stein, 1992). Second, performance is evaluated against an implicit benchmark often given by the average peer performance. In this situation and abstracting from the possible asymmetry of the compensation structure and risk-taking which we have discussed in the previous section, risk-averse portfolio managers will tend to stick to the portfolio allocation decisions of the herd in order to minimise the performance risk.[16]

[16] This point is, of course, well known also in a broader agency theory setting; see, in particular, the discussion on corporate conservatism in Zwiebel (1995).

The former point concerning the short investment horizon is especially emphasised by Goldman and Sleazak (2003), who consider a sequence of managers, where each subsequent manager inherits the portfolio of their predecessor. A key assumption of the model is that the tenure of managers is shorter than the investment horizons of the principals. Managers are, quite realistically, assumed to be remunerated according to the change in the value of the portfolio (compared with the portfolio they inherited). A main conclusion of the model by Goldman and Sleazak is that the inherited portfolio distorts the subsequent manager's incentive to trade on long-term information, and this allows erroneous information and mispricing to persist. So, the inherited position may create a 'lock-in' effect. For example, if a manager inherits a long position from his predecessor, in an agency context and with a short time horizon his incentives are tilted against selling the asset even when he has negative private information on it.

On the latter point related to implicit benchmarks, Maug and Naik (1996) note that contracts for delegated portfolio management always contain relative performance elements. They show that this relative element biases managers to deviate from return-maximising portfolio allocations and follow those of their benchmark (herding). In their model, the smart manager mimics the dumb and neglects private information, possibly also implying underinvestment in the effort necessary to acquire that information. Thus, although a relative compensation structure is not necessarily suboptimal from the point of view of P because it allows her to insulate A from the market risk, therefore being able to hire him cheaper, it may ultimately lead to a situation where each fund managers may decide to acquire the same information and the same assets as their peers, even if these decisions do not maximise risk adjusted expected returns. In this vein, Arora and Ou-Yang (2001b) present a two-period continuous-time model where a risk-averse manager wishes to track the benchmark against which he is being measured because he is risk averse. Arora and Ou-Yang develop a dynamic principal-agent model which focuses on the interplay between implicit incentives (given by reputation concerns) and explicit incentives (given by the compensation contract). In this model, the optimal portfolio policy suggests some herding behaviour. Arora and Ou-Yang (2001b) also conclude that the manager tends to follow the benchmark portfolio more closely in the earlier stage of the career, when reputation is more important, than later when he grows older. Moreover, the incentives to herd increase with the volatility of the underlying assets.

5 General equilibrium dimensions of delegated portfolio management

The literature reviewed thus far has focused only, or at least mainly, on a partial equilibrium setting represented by a game between an investor and a portfolio manager, taking asset prices and returns as given. But there is also an interesting and important aspect of delegated portfolio management which has to do with the impact that this type of agency relationship may have on asset prices and volatilities from a general equilibrium standpoint. In fact, the type of compensation contract chosen by each investor and manager, only with a view to maximising his/her own welfare, can clearly have repercussions on the equilibrium at the aggregate level. Although this aspect is quite underdeveloped in the literature, there have been a few papers which have explicitly looked at the general equilibrium dimension of delegated portfolio management. Most of these papers have concentrated on a particular feature of delegated portfolio management contracts, i.e. the presence of an explicit or implicit benchmark in the incentive structure of portfolio managers which stems from the agency relationship.

An important paper in this direction is Roll (1992), who showed that when performance is defined in terms of an exogenous benchmark and if the benchmark portfolio is not mean-variance efficient (which is highly plausible in practice), the portfolio choice by A consistent with a benchmark-adjusted compensation structure will not be mean-variance efficient. Drawing on this idea, Brennan (1993) analysed the market equilibrium implications of portfolios managed on an agency basis. The main finding of the simple analysis by Brennan is that if the benchmark portfolio is chosen exogenously (e.g., the Standard and Poor 500), equilibrium expected returns are a linear function of covariances (or betas) with respect to the residual portfolio of the assets *not* contained in the aggregate benchmark portfolio.[17] As a result, *only* the covariance with the residual portfolio is priced in the model. Unless principals know the structure of expected returns and choose the benchmark portfolio optimally, equilibrium expected returns will not be efficient.

A more developed analysis of the general equilibrium implications of delegated portfolio management as an agency relationship is Cuoco and Kaniel (2001), who study an economy where a fraction of investors delegate their investment decisions to fund managers. The fund managers' compensation is assumed to include (with no claim that this is an opti-

[17] See also Stutzer (2003).

mal contract) a flat fee, a proportional fee depending on the total value of the managed assets, and a performance fee which depends on the spread between the return on the managed portfolio and a benchmark portfolio. Cuoco and Kaniel study the general equilibrium implications of two types of performance-based components, namely a fulcrum fee and an asymmetric fee where good performance is more important than bad performance. Importantly, Cuoco and Kaniel only consider explicit incentives and do not deal with the kind of (possibly asymmetric) implicit incentives provided by reputation concerns, as discussed in Section 4.

As regards the symmetric performance contract, Cuoco and Kaniel find that individual risk-averse managers ceteris paribus prefer to over-invest in stocks included in the benchmark portfolio and under-invest in stocks excluded from it. At a macro level, this leads to higher equilibrium prices and Sharpe ratios of the stocks in the benchmark portfolio, which is also consistent with the available empirical evidence (for example, stocks included in the Standard and Poor 500, the most widely used benchmark portfolio, tend to experience an increase in price). With asymmetric performance fees, the situation is more complex since, on the one hand, risk averse agents want to hold assets having a high correlation with the benchmark in order to minimise the variance of their compensation, but, on the other hand, they may also want a low correlation with the index to maximize the expected return of the performance fee which is, by assumption, a convex function of the excess return. In fact, either incentive can dominate and, in general, with asymmetric performance fees differentials in prices and Sharpe ratios tend to be smaller than with fulcrum fees.

Gomez and Zapatero (2003) consider a financial market with two types of investors, *absolute* (caring for absolute returns) and *benchmark* (caring for returns relative to a benchmark). Portfolio managers are evaluated with respect to a given, exogenous benchmark portfolio. Based on this set-up, Gomez and Zapatero derive a two-factor CAPM model where the equilibrium relationship between expected excess returns and covariance with both the market portfolio (absolute) and the difference between the market and benchmark portfolios (benchmark) matters. In line with Brennan (1993), their empirical evidence also suggests that, especially in the more recent period, the benchmark risk appears to be priced in the market.

In contrast with the previous papers (which overall tend to stress that agency and dependency on a benchmark leads to a bias in equilibrium asset prices), Kapur and Timmerman (2003) have a more benign interpretation of the role of the delegated portfolio management agency

relationship from a general equilibrium standpoint. Kapur and Timmerman propose a model where, by assumption, managerial remuneration is linked in a linear way (which, as seen, is not necessarily optimal) to the fund's absolute performance as well as to its performance relative to rival funds. A main conclusion of the analysis is that delegated portfolio management is likely to lead to a larger demand for risky assets than if investors invested directly, for essentially two reasons. First, the compensation contract leads to risk sharing between P and A, so each of them is individually willing to take on more risk. Second, fund managers may be better informed than direct investors, and so their information advantage indirectly lowers the risk of investing in risk-bearing assets also for investors. The overall result of a more widespread recourse to delegated portfolio management might be a lower required equity premium, which is broadly consistent with the recent empirical evidence.

Finally, an important paper stressing the general equilibrium dimensions of delegated portfolio management is Berk and Green (2004). The starting point of this article is the empirical evidence reported by a large number of studies (as already touched upon in Section 2) suggesting that, first, mutual fund managers do not appear to outperform passive benchmarks (at least beyond very short horizons), and, second, that the relative performance of fund managers at a certain point in time is found not to be predicted by past relative performance. *Prima facie*, this evidence would be inconsistent with a lot of what the theoretical literature has been based on, such as the idea that A has superior information and has to be compensated in order to reveal or use it. Berk and Green argue, however, that this empirical evidence is easily explained once one takes into account the general equilibrium dimension of delegated portfolio management. Berk and Green's model is one in which managers do have superior and differential ability to generate information about future asset returns, but there are decreasing returns to scale in doing that. Managed assets flow towards managers which have had superior performance, and principals learns rationally about managers' abilities. In equilibrium, and assuming perfect competition among principals for these abilities, all principals earn zero extra profits, which exactly explains the empirical evidence just mentioned.[18]

6 Conclusions

In this paper we have selectively reviewed the theoretical literature dealing with the analytical issues arising from delegated portfolio manage-

[18] On the other hand, this conclusion would still invalidate a good part of the delegated portfolio management literature which is based on the idea that principals can extract a surplus from the agency relationship.

ment as a principal-agent relationship between an investor (the principal) and a portfolio manager (the agent). We have argued that, while this peculiar form of agency relationship shares many features with a traditional principal-agent model, it also presents its own challenges. The fact that in a delegated portfolio management setting the agent controls effort and can influence risk makes it more difficult for the principal to write incentive compatible contracts which are optimal from her standpoint. In particular, we have shown how the fact that the portfolio manager can control the scale of his response to the information signals in a linear (and potentially also nonlinear) way makes the quest for an optimal linear (and perhaps also nonlinear) contract for the principal very difficult. Indeed, this is a literature where negative results tend to prevail over constructive ones.

We have also seen how reputation concerns in a multi-period setting may affect the incentives faced by portfolio managers and in some cases make the job of explicit incentives. At the same time, reputation concerns may also have distortionary effects insofar as they may lead managers to take on more risk or to discard private information and herd with the market, none of which necessarily goes to the benefit of investors. Finally, the literature has emphasised that (implicit or explicit) benchmarking might have significant implications for asset prices and volatilities at a macro level.

Needless to say, there are several directions in which this literature could be fruitfully extended. Delegated portfolio management often implies more than one layer of agency, especially in the pension fund industry: how is this likely to affect incentives and outcomes? Another interesting extension appears to be considering less standard utility functions for principals, say shortfall risk, which may be again particularly relevant in the pension fund industry.[19]

More generally, gaining a better understanding of the general equilibrium implications of the agency aspects of delegated portfolio management should be a paramount objective in future research. This is, in particular, a topic which should be interesting and relevant for policymakers, given the importance of delegated portfolio management relationships in all developed financial markets. Notably, the possible impact of delegated portfolio management on the emergence of asset price bubbles and on excessive trading in capital markets is an issue on which the theoretical literature reviewed in this paper has definitely shed some light and which would deserve further research.

Ideally, general equilibrium models of delegated portfolio management should be able to determine the optimal compensation structure in

[19] I am grateful to a referee for this suggestion.

a principal-agent setting and its general equilibrium implications *jointly*. Although there is no reason to think that developing such models will be an easy task, since the literature has not been able to determine the optimal structure of the compensation structure in a delegated portfolio management context even in a partial equilibrium framework, our conclusion is that this research agenda should have a high priority in financial economics.

References

[1] Admati, A. R. and P. Pfleiderer (1997): "Does it all add up? Benachmarks and the compensation of active portfolio managers", Journal of Business, 70, 3, pp. 323-350.

[2] Allen, F. and G. Gorton (1993): "Churning bubbles", Review of Economic Studies, 60, pp. 813-836.

[3] Almazan, A., Brown, K. C., Carlson, M. and D. A. Chapman (2004): "Why constrain your mutual fund manager?", Journal of Financial Economics, 73, 2, pp. 289-321.

[4] Arora, N. and H. Ou-Yang (2001a): "A model of asset pricing and portfolio delegation", paper presented at the AFA 2002 annual conference.

[5] Arora, N. and H. Ou-Yang (2001b): "Explicit and implicit incentives in a delegated portfolio management problem: theory and evidence", working paper.

[6] Bailey, J. V. and D. E. Tierney (1993): "Gaming manager benchmarks", Journal of Portfolio Management, 19, pp. 37-40.

[7] Baker, M. (1998): "Fund managers' attitudes to risk and time horizons: the effect of performance benchmarking", European Journal of Finance, 4, pp. 257-278.

[8] Bank for International Settlements (2003): "Incentive structures in institutional asset management and their implications for financial markets", report submitted by a Working Group established by the Committee on the Global Financial System.

[9] Basak, S., Pavlova, A. and A. Shapiro (2003): "Offsetting the incentives: risk shifting, and benefits of benchmarking in money management", paper prepared for the EFA 2003 annual conference.

[10] Berk, J. B. and R. C. Green (2004): "Mutual fund flows and performance in rational markets", Journal of Political Economy, 12, 6, pp. 1269-1295.

[11] Bhattacharya, S. and P. Pfleiderer (1985): "Delegated portfolio management", Journal of Economic Theory, 36, pp. 1-25.

[12] Bhattacharya, S. (1999): "Delegated portfolio management, no churning, and relative performance-based incentive/sorting schemes", working paper, London School of Economics.

[13] Blake, C., Elton, E. and M. Gruber (2003): "Incentive fees and mutual funds", Journal of Finance, 58, 2, pp. 779-804.

[14] Brennan, M. J. (1993): "Agency and asset prices", working paper 6/93, Anderson Graduate School of Management.

[15] Brown, K. C., Harlow, W. V. and L. T. Starks (1996): "Of tournaments and temptations: an analysis of managerial incentives in the mutual fund industry", Journal of Finance, 51, pp. 85-110.

[16] Carpenter, J. N. (2000): "Does option compensation increase managerial risk appetite?", Journal of Finance, pp. 2311-2331.
[17] Chen, H. L. and G. Pennacchi (2002): "Does prior performance affect a mutual fund's choice of risk? Theory and further empirical evidence", working paper.
[18] Chevalier, J. and G. Ellison (1997): "Risk taking by mutual funds as a response to incentives", Journal of Political Economy, 105, 6, pp. 1167-1200.
[19] Chevalier, J. and G. Ellison (1999): "Career concerns of mutual fund managers", Quarterly Journal of Economics, 114, pp. 389-432.
[20] Christoffersen, S. E. K. (2001): "Why do money fund managers voluntarily waive their fees?", Journal of Finance, 56, 3, pp. 1117-1140.
[21] Cuoco, D. and R. Kaniel (2001): "Equilibrium prices in the presence of delegated portfolio management", working paper.
[22] Das, S. R. and R. K. Sundaram (2002): "Fee speech: signalling, risk sharing, and the impact of fee structures on investor welfare", Review of Economic Studies, 15, 5, pp. 1465-97.
[23] Davis, E. P. and B. Steil (2001): *Institutional Investors*, MIT Press.
[24] Dow, J. and G. Gorton (1997): "Noise trading, delegated portfolio management, and economic welfare", Journal of Political Economy, 105, 5, pp. 1024-1050.
[25] Dybig, P. H., Farnsworth, H. K. and J. N. Carpenter (2004): "Portfolio performance and agency", NYU Law and Economics working paper n. 04-010.
[26] Fama, E. (1980): "Agency problems and the theory of the firm", Journal of Political Economy, 88, 2, pp. 288-307.
[27] Farnsworth, H. (2003): "Reputation effects in portfolio management", working paper.
[28] Froot, K. A., Scharfstein, D. S. and J. C. Stein (1992): "Herd on the street: informational inefficiencies in a market with short-term speculation", Journal of Finance, 47, 4, pp. 1461-84.
[29] Garcia, D. (2001): "Optimal menus of linear contracts in portfolio selection problems", working paper.
[30] Goetzmann, W. N., Ingersoll, J. and S. A. Ross (1998): "High water marks", NBER working paper n. 6413.
[31] Goldman, E. and S. L. Slezak (2003): "Delegated portfolio management and rational prolonged mispricing", Journal of Finance, 58, 1, pp. 283-311.
[32] Gollier, C., Koehl, P. F. and J. C. Rochet (1997): "Risk-taking behavior with limited liability and risk aversion", Journal of Risk and Insurance, 64, 2, pp. 347-70.

[33] Gomez, J.-P. and J. Sharma (2003): "Portfolio delegation under short-selling constraints ", UPF Economics and Business working paper n. 695.
[34] Gomez, J.-P. and F. Zapatero (2003): "Asset pricing implications of benchmarking: a two-factor model", European Journal of Finance, 9, pp. 343-357.
[35] Goriaev, A., Palomino, F. and A. Prat (2001): "Mutual fund tournament: risk-taking incentives induced by ranking objectives", CEPR discussion paper n. 2794.
[36] Grinblatt, M. and S. Titman (1989): "Adverse risk incentives and the design of performance based contracts", Management Science, 35, pp. 807-822.
[37] Gruber, M. (1996): "Another puzzle: the growth of actively managed mutual funds", Journal of Finance, pp. 783-810.
[38] Heinkel, R. and N. M. Stoughton (1994): "The dynamics of portfolio management contracts", Review of Financial Studies, 7, 2, pp. 351-387.
[39] Holmstrom, B. and P. Milgrom (1987): "Aggregation and linearity in the provision of intertemporal incentives", Econometrica, 55, 2, pp. 303-328.
[40] Huberman, G. and S. Kandel (1993): "On the incentives for money managers: a signalling approach", European Economic Review, 37, pp. 1065-1081.
[41] Huddart, S. (1999): "Reputation and performance fee effects on portfolio choice by investment advisers", Journal of Financial Markets, 2, 3, pp. 227-271.
[42] Kapur, S. and A. G. Timmermann (2003): "Relative performance evaluation contracts and asset market equilibrium", CEPR discussion paper n. 4038, and forthcoming in the Economic Journal.
[43] Lakonishok, J., Shleifer, A. and R. W. Vishny (1992a): "The impact of institutional trading on stock prices", Journal of Financial Economics, 32, pp. 23-43.
[44] Lakonishok, J., Shleifer, A. and R. W. Vishny (1992b): "The structure and performance of the money management industry", Brookings Papers on Economic Activity, Microeconomics, pp. 339-391.
[45] Malkiel, B. G. (1995): "Returns from investing in equity mutual funds 1971 to 1991", Journal of Finance, 50, pp. 549-572.
[46] Maug, E. G. and N. Y. Naik (1996): "Herding and delegated portfolio management: the impact of relative performance evaluation on asset allocation", IFA working paper n. 223/1996.
[47] Ou-Yang, O. (2003): "Optimal contracts in a continuous-time delegated portfolio management problem", The Review of Financial Studies, 16, 1, pp. 173-208.

[48] Ozerturk, S. (2004): "Equilibrium incentives to acquire precise information in delegated portfolio management", Journal of Financial Services Research, 25:1, pp. 25-36.
[49] Palomino, F. and A. Prat (2003): "Risk taking and optimal contracts for money managers", RAND Journal of Economics, 34, 1, pp. 113-137.
[50] Palomino, F. and A. Sadrieh (2004): "Overconfidence and delegated portfolio management", CEPR Discussion Paper No. 4131.
[51] Rajan, U. and S. Srivastava (2000): "Portfolio delegation with limited liability", GSIA Working Paper No. 2000-E16.
[52] Roll, R. (1992): "A mean/variance analysis of the tracking error", Journal of Portfolio Management, 8, 4, pp. 13-22.
[53] Ross, S. (2004): "Compensation, incentives, and the duality of risk aversion and riskiness", Journal of Finance, 59, 1, pp. 207-225.
[54] Salanie, B. (1996): *The economics of contracts: a primer*, MIT Press.
[55] Sappington, D. E. M. (1991): "Incentives in principal-agent relationships", Journal of Economic Perspectives, 5, 2, pp. 45-66.
[56] Scharfstein, D. S. and J. C. Stein (1990): "Herd behaviour and investment", American Economic Review, 80, pp. 465-479.
[57] Sirri, E. R. and P. Tufano (1998): "Costly search and mutual fund flows", Journal of Finance, 54, 1, pp. 359-375.
[58] Starks, L. T. (1987): "Performance incentive fees: an agency theoretic approach", Journal of Financial and Quantitative Analysis, 22, pp. 17-32.
[59] Stiglitz, J. (1974): "Incentives and risk sharing in sharecropping", Review of Economic Studies, 41, pp. 219-255.
[60] Stoughton, N. M. (1993): "Moral hazard and the portfolio management problem", Journal of Finance, 48, pp. 2009-2028.
[61] Stutzer, M. (2003): "Fund managers may cause their benchmarks to be priced "risks"", working paper, University of Colorado.
[62] Wilson, R. B. (1968): "The theory of syndicates", Econometrica, 36, pp. 119-132.
[63] Zwiebel, J. (1995): "Corporate conservatism and relative compensation", Journal of Political Economy, 103, 1, pp. 1-25.

European Central Bank working paper series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (http://www.ecb.int)

490 "Unions, wage setting and monetary policy uncertainty" by H. P. Grüner, B. Hayo and C. Hefeker, June 2005.

491 "On the fit and forecasting performance of New-Keynesian models" by M. Del Negro, F. Schorfheide, F. Smets and R. Wouters, June 2005.

492 "Experimental evidence on the persistence of output and inflation" by K. Adam, June 2005.

493 "Optimal research in financial markets with heterogeneous private information: a rational expectations model" by K. Tinn, June 2005.

494 "Cross-country efficiency of secondary education provision: a semi-parametric analysis with non-discretionary inputs" by A. Afonso and M. St. Aubyn, June 2005.

495 "Measuring inflation persistence: a structural time series approach" by M. Dossche and G. Everaert, June 2005.

496 "Estimates of the open economy New Keynesian Phillips curve for euro area countries" by F. Rumler, June 2005.

497 "Early-warning tools to forecast general government deficit in the euro area: the role of intra-annual fiscal indicators" by J. J. Pérez, June 2005.

498 "Financial integration and entrepreneurial activity: evidence from foreign bank entry in emerging markets" by M. Giannetti and S. Ongena, June 2005.

499 "A trend-cycle(-season) filter" by M. Mohr, July 2005.

500 "Fleshing out the monetary transmission mechanism: output composition and the role of financial frictions" by A. Meier and G. J. Müller, July 2005.

501 "Measuring comovements by regression quantiles" by L. Cappiello, B. Gérard, and S. Manganelli, July 2005.

502 "Fiscal and monetary rules for a currency union" by A. Ferrero, July 2005

503 "World trade and global integration in production processes: a re-assessment of import demand equations" by R. Barrell and S. Dées, July 2005.

504 "Monetary policy predictability in the euro area: an international comparison" by B.-R. Wilhelmsen and A. Zaghini, July 2005.

505 "Public good issues in TARGET: natural monopoly, scale economies, network effects and cost allocation" by W. Bolt and D. Humphrey, July 2005.

506 "Settlement finality as a public good in large-value payment systems" by H. Pagès and D. Humphrey, July 2005.

507 "Incorporating a "public good factor" into the pricing of large-value payment systems" by C. Holthausen and J.-C. Rochet, July 2005.

508 "Systemic risk in alternative payment system designs" by P. Galos and K. Soramäki, July 2005.

509 "Productivity shocks, budget deficits and the current account" by M. Bussière, M. Fratzscher and G. J. Müller, August 2005.

510 "Factor analysis in a New-Keynesian model" by A. Beyer, R. E. A. Farmer, J. Henry and M. Marcellino, August 2005.

511 "Time or state dependent price setting rules? Evidence from Portuguese micro data" by D. A. Dias, C. R. Marques and J. M. C. Santos Silva, August 2005.

512 "Counterfeiting and inflation" by C. Monnet, August 2005.

513 "Does government spending crowd in private consumption? Theory and empirical evidence for the euro area" by G. Coenen and R. Straub, August 2005.

514 "Gains from international monetary policy coordination: does it pay to be different?" by Z. Liu and E. Pappa, August 2005.

515 "An international analysis of earnings, stock prices and bond yields" by A. Durré and P. Giot, August 2005.

516 "The European Monetary Union as a commitment device for new EU member states" by F. Ravenna, August 2005.

517 "Credit ratings and the standardised approach to credit risk in Basel II" by P. Van Roy, August 2005.

518 "Term structure and the sluggishness of retail bank interest rates in euro area countries" by G. de Bondt, B. Mojon and N. Valla, September 2005.

519 "Non-Keynesian effects of fiscal contraction in new member states" by A. Rzońca and P. Ciżkowicz, September 2005.

520 "Delegated portfolio management: a survey of the theoretical literature" by L. Stracca, September 2005.