Calibrated Incentive Contracts

Sylvain Chassang*

Princeton University

February 6, 2011

Abstract

This paper studies a financial agency problem which includes limited liability, moral hazard and adverse selection. The paper develops a robust approach to dynamic contracting based on calibrating the payoffs that would have been delivered by simple benchmark contracts that are attractive but infeasible, due to limited liability constraints. The resulting contracts are detail-free and perform well independently of the underlying process for returns. The paper discusses how these calibrated contracts relate and differ from contracts used in practice.

1 Introduction

This paper considers financial agency problem in which a principal hires an agent to make investment decisions on her behalf.¹ The contracting environment is delicate as it includes limited liability, moral hazard, adverse selection, and makes very few assumptions about the underlying process for returns and information. The paper's contribution is to develop a robust approach to dynamic contracting. The main steps are as follows: 1) identify a simple class of *high-liability* static linear contracts that satisfy attractive and robust efficiency properties; 2) construct *limited-liability* dynamic contracts that achieve the same performance

^{*}Contact: chassang@princeton.edu.

¹Throughout the paper, the principal is referred to as she, while the agent is referred to as he. Appendix A shows how to extend the analysis to a more general principal-agent framework not limited to asset allocation problems.

by calibrating the rewards to the agent so that they approximately satisfy key properties of the benchmark high-liability contracts. The resulting calibrated contracts perform well independently of the underlying process for returns. In particular, the results do not rely on any ergodicity or stationarity assumptions.

The model considers a risk-neutral principal and a risk-neutral agent. Both the principal and the agent are patient. The agent has a large but finite horizon which need not be know to the principal. For simplicity the principal has an infinite horizon. In every period a steady-state amount of wealth is to be invested on behalf of the principal by the agent. The agent can be of two types. If the agent is talented, he can invest in costly information acquisition. An untalented agent has no information beyond public knowledge. The agent privately observes his own type, the information he acquires, and his time horizon. The main constraint on contracts is limited liability: it is either impossible, or very difficult for the agent to receive negative transfers, and rewards must satisfy a pay-as-you-go constraint which rules out large deferred payments.

This is clearly a difficult environment to contract in. The principal is facing both adverse selection (is the agent talented, what is the cost effectiveness of information acquisition...) and moral hazard (is the agent acquiring information, is the agent making investment allocations that maximize expected returns). At this level of generality, characterizing optimal contracts is unlikely to be informative and may not actually be possible if the principal has poorly specified beliefs over the environment. Instead the paper develops a robust approach to dynamic contracting which emphasizes prior-free performance bounds.

The first step of the approach identifies a suitable – although infeasible – benchmark. The paper focuses on a simple class of linear contracts in which the agent is rewarded a share of his externality on the principal. These linear contract exhibit high-liability since the agent is expected to provide partial compensation for losses. Regardless of the underlying process for returns and information, they satisfy the following properties: (i) the agent can obtain positive expected rewards if and only the principal obtains positive expected surplus; as a consequence, untalented types weakly prefer screening themselves out; (ii) the agent chooses the optimal asset allocation given information; (iii) expected excess returns to the principal can be bounded below as a function of the maximum feasible expected excess returns; (iv)the previous properties for continuation play from the perspective of any realized history.

The second step of the approach, and the main theoretical contribution of the paper, is to develop a simple class of dynamic contracts that robustly approximate the performance of linear high-liability contracts while satisfying severe limited liability constraints. The key insight is to calibrate both the reward-rate of the agent and the share of total wealth he is investing, so that for all possible strategies and all realizations of uncertainty, the payoff obtained by the agent and the excess returns obtained by the principal remain as tightly linked as they are under benchmark linear contracts.

Taking the agent as given, these calibrated contracts induce performance approximately equal to that achieved by the benchmark linear contracts. A penalized version of the same contracts, achieves screening at a moderate performance loss. Under these penalized contracts, the agent is charged an initial performance fee and only obtains rewards in periods where his performance is above a hurdle which depends on the magnitude of his trading activity. Agents who trade often for moderate returns find it difficult to pass this hurdle, whereas agents that can deliver moderate returns while trading rarely are hardly impeded.

The calibrated contracts described in this paper share many features with existing financial contracts, notably those encountered in the hedge-fund industry. In particular, the agent only gets rewarded when the current surplus created for the principal is higher than its historical maximum, a feature shared by high-watermark contracts. However, the contracts described here exhibit important features that typical incentive contracts do not have and which matter essentially for performance:

1. The share of total wealth managed by the agent is calibrated jointly with his payoffs. This allows to keep tight the relation between excess returns and payoffs to the agent.

- 2. The returns generated by the agent are benchmarked using the counterfactual returns the principal would have obtained on her own. This allows for screening to hold under very general conditions.
- 3. The agent is only rewarded when his performance is above a hurdle which increases with the magnitude of positions he has been taking. This favors agents delivering good performance while taking relatively small positions, and allows the screening of untalented agents at a moderate incentive cost to talented agents.

Inversely, a surprising aspect of the contracts developed in this paper is that they do not exhibit some of the features typically deemed necessary for financial agents to deliver good performance: the agent is not required to hold the underlying asset allocation and there are no clawback provisions. Still, while the paper shows that these features are not necessary for good performance when the agent's horizon is long, it would be wrong to infer that they are not useful, for instance when the agent's horizon is short.

The approach of this paper is related to the work of Jackson and Sonnenschein (2007) who show how linking decisions across different states can be used in a general mechanism design context to relax incentive compatibility constraints (see also Radner (1981, 1985), Townsend (1982), McAfee (1992), or Casella (2005)). As in their work, the main idea of this paper is to constrain payoffs to satisfy key properties that would hold under an ideal benchmark. In both cases, this style of approach shows that agents will behave appropriately with high probability, but does not specify the rare circumstances in which agents might choose to deviate. An important difference with Jackson and Sonnenschein (2007) is that they assume the states of the world are i.i.d. In contrast, the contracts developed in this paper are designed to perform well even in environments where the process for the underlying state of the world is not ergodic. For instance, there can be non-vanishing probability that for a large number of periods, returns happen to be negative, or that even talented managers do not receive superior information. The methods used in the paper, as well as the emphasis on general stochastic processes, connect the paper to the literature on testing experts (see for instance Foster and Vohra (1998), Fudenberg and Levine (1999), Lehrer (2001) or more recently Al-Najjar and Weinstein (2008), Feinberg and Stewart (2008) and Olszewski and Sandroni (2008)). However, the main question is not whether good tests are available. Rather, this paper takes a principal agent approach related to that of Echenique and Shmaya (2007) and Olszewski and Peski (forthcoming). These papers both show that in such environments there are satisfactory ways to identify experts that generate positive surplus. Neither paper tackles incentive provisions when information acquisition is costly or the issue of self-screening by experts.

The paper is related to the work of Lo (2001), Goetzmann et al. (2007) and Foster and Young (2010) on appropriate performance measures and the difficulty of jointly rewarding and screening wealth managers. In particular Foster and Young (2010) describe environments in which rewarding and screening is in fact impossible. This occurs because in their environment, informed managers value income in early periods much more than in later periods and so that even talented managers are unwilling to pay the monetary cost needed to induce screening. In contrast, the current paper considers patient players with constant marginal utility for income.² In that case, self-screening can be obtained under severe limited liability constraints.

The paper hopes to usefully complement the rich literature on optimal dynamic contracting (see for instance Rogerson (1985), Spear and Srivastava (1987), and more recently DeMarzo and Sannikov (2006), Biais et al. (2007), DeMarzo and Fishman (2007), Sannikov (2008) or Edmans et al. (2009)). Indeed the optimal contracting approach delivers rich insights about how contracts should vary with the environment, but the performance of the resulting contracts is notoriously dependent on the underlying environment. In contrast, the contracts developed in this paper perform well independently of the underlying environment,

²Appendix A partially bridges the two sets of results, by extending performance bounds to the case where the agent does not have constant marginal utility of income.

and are well suited candidates whenever the principal has poorly specified beliefs. However, because these contracts are detail-free, this approach has little predictive power as to how contracts should change with the underlying environment.

Finally the paper is related to the literature on robust mechanism design that operationalizes the doctrine set by Wilson (1987), and attempts to characterize mechanisms that behave well under weak assumptions over payoff distributions and beliefs. A rich strand of that literature studies mechanisms that are robust with respect to the solution concept used to characterize the players' behavior.³ The paper is especially related to a recent strand in this literature, illustrated for instance by Hartline and Roughgarden (2008), which looks for mechanisms that satisfy robust performance bounds regardless of the underlying distribution of values.⁴ A tricky step, common to Hartline and Roughgarden (2008) and this paper, is to define appropriate benchmark performance measures that allow for informative worst-case analysis of mechanisms.

The paper is structured as follows. Section 2 describes the framework. Section 3 introduces a benchmark class of high liability linear contracts that satisfy a number of attractive efficiency properties but do require high levels of liability from the agent. Section 4 is the core of the paper: it develops the idea of calibrated contracts and analyzes their performance, taking the agent as given. Section 5 shows how to screen agents by means of an activity-based hurdle. Section 6 relates calibrated contracts to contracts used in practice and concludes. Appendix A extends the analysis to various environments. Proofs are given in Appendix B, unless mentioned otherwise.

³See for instance Dasgupta et al. (1979), Hagerty and Rogerson (1987), Eliaz (2002), Chung and Ely (2003, 2007) or Bergemann and Morris (2005).

⁴Local approaches are possible and informative. For instance Madarász and Prat (2010) consider screening mechanisms that satisfy strong efficiency bounds for all type distributions within a small neighborhood. Global incentive compatibility constraints play an important role in their analysis, and will also show up in this paper.

2 The Framework

Players, Actions and Payoffs. A principal (for instance, a representative investor, or a bank) hires an agent (say a wealth manager, a financial advisor, or a trader) to make investment allocations on her behalf. The agent is active for a large but finite number of periods N. The principal has an infinite horizon and need not know the agent's horizon N. Both the principal and the agent are patient and do not discount future payoffs.⁵

In each period $t \in \{1, \dots, N\}$, the principal invests an amount w at the beginning of the period. The amount of wealth w invested in each period is constant, and can be thought of as a steady state amount of wealth to be invested. The realized wealth w_t after investment is consumed at the end of the period. Both the principal and the agent are risk neutral over the range of flow payoffs.⁶ The agent's outside option is set to zero.

Wealth can be invested in one of K assets whose returns at time t are denoted by $\mathbf{r}_t = (r_{k,t})_{k \in \{1,\dots,K\}}$. Let R denote the set of possible returns \mathbf{r}_t . An asset allocation at time t is a vector $a_t \in A \subset \mathbb{R}^K$ such that $\sum_{k=1}^K a_t = 1$. Set A is convex and compact. It represents constraints on possible positions. These constraints on allocations can be thought of as a mandate set by the principal as in He and Xiong (2010). Let $\langle \cdot, \cdot \rangle$ denote the usual scalar product. Given asset allocation a_t and returns \mathbf{r}_t , the consumer's wealth at the end of period t is

$$w_t = w \times (1 + \langle a_t, \mathbf{r}_t \rangle).$$

It is assumed that, for any $a \in A$ and $\mathbf{r}_t \in R$, $w_t \ge 0.7$ For any pair of allocations $(a, a') \in A^2$, the distance between a and a' is defined by

$$d(a,a') \equiv \sup_{\mathbf{r}_t \in R} |\langle a - a', \mathbf{r}_t \rangle|.$$
(1)

⁵All results extend without difficulty to the case where future payoffs are discounted, but the performance bounds that one can derive only become attractive for discount factors sufficiently close to 1.

⁶Appendix A allows for varying amounts of wealth to be invested, and also considers extensions to cases where the principal is risk-averse or where the agent does not have constant marginal utility of income.

⁷This facilitates exposition but is inessential for the analysis.

The following assumption puts constraints on the set of permissible allocations A and is maintained throughout the paper.

Assumption 1. There exists $\overline{d} \in \mathbb{R}^+$ such that for all $(a, a') \in A^2$, $d(a, a') \leq \overline{d}$.

This assumption limits the magnitude of changes that can occur with the principal getting no feedback.

There are two types of managers: talented and untalented managers. Managers know their own type. At the beginning of every period t, talented managers can expend cost $c_t \in [0, \overline{c}]$ towards acquiring information. This cost can be the actual cost of obtaining data, an effort cost, or the opportunity cost of time. Untalented managers only have access to public information. Managers then make an asset allocation suggestion $a_t \in A$ (whether or not they have superior information) and receive a payment π_t depending on the realized public history at the end of period t. The manager's objective is to maximize his expected average payoffs

$$\mathbb{E}\left(\frac{1}{N}\sum_{t=1}^{N}\pi_t - c_t\right).$$
(2)

Information. Information acquired at time t is represented as random variables I_t from a measurable state space (Ω, σ) to a measurable signal space $(\mathcal{I}, \sigma_{\mathcal{I}})$. Untalented managers only have access to publicly available information I_t^0 (which includes realized past returns). In contrast, talented managers can acquire expert information I_t^c at cost $c \in [0, \overline{c}]$. For all $c \geq c' \geq 0$, $I_t^{c'}$ is measurable with respect to I_t^c , i.e. the more the agent invests, the more he knows. Given a sequence of information acquisition expenditures $(c_t)_{t\geq 1}$, let $(\mathcal{F}_t)_{t\geq 1}$ be the informed manager's filtration (generated by $(I_t^{c_t})_{t\geq 1}$), and let $(\mathcal{F}_t^0)_{t\geq 1}$ denote the uninformed manager's filtration (generated by $(I_t^0)_{t\geq 1}$).

For simplicity it is convenient to assume that the principal and the agent have a common prior P over the state space (Ω, σ) .⁸ Let $\mathcal{P} = (\Omega, \sigma, P)$ denote the resulting probability

⁸All results hold in a non-common prior setting, taking expectations under the agent's prior.

space. The paper *does not* assume that either information or returns follow an i.i.d. or ergodic process. This results in a very flexible model. For instance, there may be nonvanishing probability that returns are below their period t = 1 expectation for an arbitrarily large number of periods. Also, the value of information that talented managers can collect may vary in arbitrary ways. For instance, once valuable trading strategies can become obsolete over time.

Strategies. Altogether, an agent's strategy consists of an information acquisition strategy $c = (c_t)_{t \in \mathbb{N}}$, and an asset allocation strategy $a = (a_t)_{t \in \mathbb{N}}$, where both c_t and a_t are adapted to the information available to the manager at the time of decision. Let a_t^0 and a_t^* respectively denote efficient asset allocations under information \mathcal{F}_t^0 and \mathcal{F}_t :

$$a_t^0 \in \underset{a \in A}{\operatorname{arg\,max}} \mathbb{E}[\langle a, \mathbf{r}_t \rangle | \mathcal{F}_t^0] \quad \text{and} \quad a_t^* \in \underset{a \in A}{\operatorname{arg\,max}} \mathbb{E}[\langle a, \mathbf{r}_t \rangle | \mathcal{F}_t].$$
 (3)

Allocation a_t^0 is the allocation the principal could pick on her own, given public information \mathcal{F}_t^0 . Let $w_t^0 = w \times (1 + \langle a_t^0, \mathbf{r}_t \rangle)$ and $w_t = w \times (1 + \langle a_t, \mathbf{r}_t \rangle)$ denote realized wealth under allocation a_t^0 and under the allocation a_t actually chosen by the agent.

Limited liability contracts. Contracts $(\pi_t)_{t\in\mathbb{N}}$ are adapted to public histories observed by the principal, where public histories consist of past public information (including past returns) as well as past suggested asset allocations by the agent.⁹ The principal has commitment power but transfers are subject to the following constraints: in every period T,

$$\sum_{t=1}^{T} \pi_t \mathbf{1}_{\pi_t < 0} \geq -b, \quad \text{and} \tag{4}$$

$$\pi_T \leq \overline{\pi},\tag{5}$$

⁹To implement the contracts analyzed in the paper, it would be sufficient for the principal to observe $d(a_t, a_t^0)$ – the distance between a_t and a_t^0 – rather than a_t itself.

where $\overline{\pi}$ is a bound on rewards that is independent of N but weakly greater than wd, the maximum single-period gain in wealth that the agent can generate.

Condition (4) is a strong limited-liability constraint on the agent's side. The sum of punishments that the principal can inflict on the agent is bounded above by a fixed amount b. Punishments may correspond to monetary transfers, as well as non-monetary costs such as grueling work, sleep deprivation, or tedious and lengthy training.

Condition (5) puts an upper bound on the per-period transfers that the principal can make (or commit to make) to the agent. Importantly, this limits how long the payment of wages can be delayed and precludes the possibility of large deferred payments. The fact that the agent gets paid in real time complements the assumption that the agent is risk neutral over the range of payoffs.¹⁰

These constraints are at the origin of the contracting problem: the agent does not share on the downside, and rewards must be given in real time rather than delayed until the end.

3 The Benchmark: High-Liability Linear Contracts

The environment described in Section 2 involves both moral hazard and adverse selection: the agent must acquire information and makes asset allocation decisions that may or may not benefit the principal; in addition the agent's talent and the information he may acquire are private. At this level of generality, informative characterizations of optimal dynamic contracts are unlikely. Solving for optimal contracts may also be of limited use if the principal doesn't have well defined beliefs over the underlying environment.

The paper embraces an alternative robust approach to dynamic contracting. The first step is to define a class of benchmark contracts that have attractive properties, but violate the limited liability constraints (4) and (5). The second step of the analysis is to construct a class of dynamic contracts that do satisfy constraints (4) and (5), and also achieve performance

¹⁰Clearly, if it were possible, it would be optimal for the principal to delay payment until the end of the relationship. This would maximize degrees of freedom with respect to payment design. The use of large deferred payments would be a considerable stretch the assumption of risk-neutrality.

approximately as good as that of the benchmark contracts, regardless of the underlying environment \mathcal{P} .

The benchmark contracts of interest in the paper are simple linear contracts, loosely in the spirit of Vickrey-Clarke-Groves (VCG) mechanisms.¹¹ Specifically, in period t the agent's reward π_t is a share α of the externality his decisions have on the principal:

$$\forall t, \quad \pi_t = \alpha (w_t - w_t^0).^{12} \tag{6}$$

For instance, if $\alpha = .2$ and the default allocation a_t^0 is to invest all wealth in risk-free bonds, the benchmark contract pays the agent 20% of the excess-returns when he beats the risk-free rate, and charges him 20% of the foregone returns when he under-performs the risk-free rate. Throughout the paper, the working assumption is that the main problem with such a contract is not that it is too unsophisticated or provides insufficient incentives, but rather that it requires unrealistic levels of liability on the part of the agent.¹³

This working assumption is motivated in two ways. First, one can leverage the extensive literature pioneered by Holmstrom and Milgrom (1987), which makes the point that when the agent has a sufficiently rich action set, optimal contracts often must take a linear form.¹⁴ In fact, as the example below shows, it is easy to come up with reasonable environments under which such linear contracts are obviously optimal. Second, and more importantly for the purposes of this paper, Theorem 1 shows that even when linear contracts are not optimal, their VCG-like features guarantee a number of attractive properties, including a lower bound on their performance.

¹¹See Vickrey (1961), Clarke (1971) and Groves (1973).

¹²Recall that w_t and w_t^0 respectively denote final wealth under the agent's suggested asset allocation and under the default, public information, asset allocation. Appendix A extends the analysis to a broader class of benchmark contracts.

¹³In this paper the reward rate α is taken as given. Figuring out sensible ways to pick α is the object of ongoing research. In the spirit of Hartline and Roughgarden (2008), one can devise crude lower bounds by picking α at random, since this guarantees that at least sometimes, the reward rate will be appropriate.

¹⁴See also Sung (1995), Hellwig and Schmidt (2002), or Edmans and Gabaix (2009)

An example of optimality. Assume that in order to generate flow expected excess returns $\nu = \mathbb{E}(w_t - w_t^0 | \mathcal{F}_t)$ in period t, the manager must invest a cost

$$c_t = \alpha_0 \nu \mathbf{1}_{\nu \le \overline{\nu}_t} + \infty \mathbf{1}_{\nu > \overline{\nu}_t},\tag{7}$$

where $\overline{\nu}_t \in \mathbb{R}^+$ can be stochastic, is unknown to the principal, and observed by the agent. There is a constant marginal cost α_0 to generate expected excess returns up to an upper bound $\overline{\nu}_t$. Excess returns above bound $\overline{\nu}_t$ are infeasible. Note that because $\overline{\nu}_t$ can follow any stochastic process, this class of models includes environments in which talented managers can lose the ability to generate returns for extended periods of time. Clearly, any linear contract with reward rate $\alpha \geq \alpha_0$ can achieve the first best. The linear contract such that $\alpha = \alpha_0$ transfers all the surplus to the principal.

Robust Efficiency Properties. The main motivation for the use of linear contracts as a benchmark is that they satisfy a number of attractive properties regardless of the probability space \mathcal{P} .

Given a reward rate α , the agent solves optimization problem

$$\max_{c,a} \mathbb{E}\left(\frac{1}{N}\sum_{t=1}^{N}\alpha(w_t - w_t^0) - c_t\right).$$
(P1)

The corresponding per-period excess returns r_{α} accruing to the principal (net of payments to the agent) are

$$r_{\alpha} \equiv \inf \left\{ \mathbb{E}_{c,a} \left(\frac{1}{Nw} \sum_{t=1}^{N} w_t - w_t^0 - \pi_t \right) \left| (c,a) \text{ solves (P1)} \right\}.^{15} \right\}$$

The expression for returns r_{α} involves an inf since the agent may be indifferent between

 $^{^{15}}$ For concision, the paper focuses on net excess returns accruing to the principal. The analysis can be extended to total surplus without difficulties.

multiple policy profiles. In anticipation of technical subtleties to come, it is useful to note that because the underlying environment is very general, the paper cannot rule-out binding (or almost binding) global incentive compatibility constraints.

Let r_{\max} denote the per-period excess returns that can be generated when the agent invests \bar{c} towards information acquisition in every period and requires no rewards. Recalling that a^* is the surplus maximizing allocation strategy given information, we have $r_{\max} = \mathbb{E}_{\bar{c},a^*}\left(\frac{1}{N}\sum_{t=1}^N \langle a_t^* - a_t^0, \mathbf{r}_t \rangle\right).$

Theorem 1. Regardless of probability space \mathcal{P} , the benchmark contract defined by (6) satisfies

- (i) (no-loss): expected rewards to the agent are positive if and only if expected returns to the principal are positive; as a consequence, untalented managers screen themselves out;
- (ii) (truthfulness): given information \mathcal{F}_t , it is optimal for the manager to pick the efficient asset allocation a^* ;
- (iii) (lower bounds on returns): whenever the agent is rational,

$$r_{\alpha} \ge (1-\alpha) \left(r_{\max} - \frac{\overline{c}}{\alpha w} \right);$$
(8)

(*iv*) (*history independence*): the above properties hold for continuation behavior at any history.

Point (i) ensures that regardless of the environment \mathcal{P} , the principal can never lose surplus in expectation. Since untalented agents cannot generate positive surplus, this implies that they cannot obtain positive expected profit. Hence untalented agents are willing to screen themselves out.¹⁶

 $^{^{16}\}mathrm{Any}$ small participation fee would make it strictly optimal for uninformed agents to screen themselves out.

Point (ii) states that allocations are efficient conditional on information. Therefore, the only agency concern under benchmark contracts is information acquisition.

Point (*iii*) provides a lower bound for the returns that the principal obtains under the benchmark contract. For any α , as wealth under management w grows arbitrarily large, the contract becomes approximately efficient, and the principal obtains a share approximately $1 - \alpha$ of the maximum returns r_{max} . To illustrate this bound for a finite value of w, consider the specification $\bar{c} = \$5M$, $r_{\text{max}} = 5\%$, w = \$1Bn and $\alpha = 20\%$. In this case, bound (8) guarantees that the principal obtains at least 40% of the maximum possible excess returns. Note that this is more than a share 40% of the maximum surplus since the agent must incur information acquisition costs.

Finally, point (iv) states that the attractive properties of benchmark contracts hold from the perspective of any history. Although the paper assumes full commitment, this provides reassurance that renegotiation issues are limited under linear contracts. This directly echoes the point made by Holmstrom and Milgrom (1987) that because linear contracts apply constant incentive pressure across histories they limit the scope for gaming and manipulations.

The fact that these efficiency properties hold independently of probability space \mathcal{P} motivates the use of linear contracts as a robust benchmark. The contribution of the paper is to construct equally robust dynamic contracts that perform approximately as well as the benchmark contracts, while also satisfying limited liability constraints (4) and (5).¹⁷

4 Calibrated Contracts

This section introduces a novel class of dynamic "calibrated" contracts that robustly approximate the performance of linear contracts while satisfying limited liability constraints (4) and (5). This section focuses on providing a given agent with appropriate incentives.

¹⁷Appendix A shows that no static contract guarantees performance similar to linear contract while also satisfying limited liability constraints (4) and (5). For instance, option contracts of the form $\pi_t = \alpha (w_t - w_t^0)^+$ make the agent risk loving, which distorts investment and allocation decisions.

Section 5 deals with screening. Section 6 relates calibrated contracts to contracts used in practice.

4.1 The Contract

In every period t, the agent is allowed to invest a share $\lambda_t \in (0, 1)$ of the principal's wealth, while the remaining share $1 - \lambda_t$ is invested in the default asset allocation a_t^0 . At the end of the period, the agent receives a payment π_t .

Specifying investment shares and rewards $(\lambda_t, \pi_t)_{t \ge 1}$ requires additional notation. For all periods T and T' < T, define

$$\Pi_T = \sum_{t=1}^T \pi_t \; ; \quad \Sigma_T = \sum_{t=1}^T w_t - w_t^0 \; ; \quad S_T = \sum_{t=1}^T \lambda_t (w_t - w_t^0) \tag{9}$$

and

$$\Pi_{T \setminus T'} = \sum_{t=T'}^{T} \pi_t \; ; \quad \Sigma_{T \setminus T'} = \sum_{t=T'}^{T} w_t - w_t^0 \; ; \quad S_{T \setminus T'} = \sum_{t=T'}^{T} \lambda_t (w_t - w_t^0). \tag{10}$$

Value Π_T corresponds to the payoffs that the agent has obtained; Σ_T corresponds to the excess profits that would have been generated by fully investing according to the agent's suggested asset allocation; S_T corresponds to the actual excess profits that have been generated given that only a share λ_t of wealth w is invested according to the agent's suggestion. Values $\Pi_{T\setminus T'}$, $\Sigma_{T\setminus T'}$ and $S_{T\setminus T'}$ compute the same quantities over time range $\{T', \dots, T\}$.

The difference $\Sigma_{T\setminus T'} - S_{T\setminus T'}$ corresponds to the foregone gains from not investing entirely according to the agent's allocation between T' and T. The difference $\Pi_T - \alpha S_T$ corresponds to the agent's excess rewards, the target being to reward him a share α of his externality S_T on the principal. Using the notation $(x)^+ = \max\{0, x\}$, investment shares and rewards $(\lambda_t, \pi_t)_{t\geq 1}$ are defined recursively as follows: $\lambda_1 = 1, \pi_1 = 0$, and for all $T \geq 1$

$$\lambda_{T+1} \equiv \frac{\alpha \left[\max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'} \right]^+}{\left[\Pi_T - \alpha S_T \right]^+ + \alpha \left[\max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'} \right]^+}$$
(11)

 $\equiv \frac{\alpha \times \text{maximum foregone gain}}{\text{excess rewards} + \alpha \times \text{maximum foregone gain}}$

with the convention that $\frac{0}{0} = 1$, and

$$\pi_{T+1} \equiv \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if } \Pi_T \le \alpha S_T \\ 0 & \text{otherwise} \end{cases}$$

$$\equiv \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if rewards} \le \alpha \times \text{actual excess returns} \\ 0 & \text{otherwise} \end{cases}$$

$$(12)$$

Note that the contract specified above satisfies limited liability conditions (4) and (5): payments $(\pi_t)_{t\geq 1}$ are positive and bounded above by $\alpha w \overline{d}$. Theorem 2 (stated below) shows that this class of contracts approximates the performance of benchmark contracts. Some additional notation is needed. Given a contract specification $(\lambda, \pi) = (\lambda_t, \pi_t)_{t\geq 1}$, let $r_{\lambda,\pi}$ denote the net excess returns delivered by the agent under the corresponding contract:

$$r_{\lambda,\pi} = \inf \left\{ \mathbb{E}_{c,a} \left(\frac{1}{Nw} \sum_{t=1}^{N} \lambda_t (w_t - w_t^0) - \pi_t \right) \, \middle| \, (c,a) \text{ solves } \max_{c,a} \mathbb{E} \left(\frac{1}{N} \sum_{t=1}^{N} \pi_t - c_t \right) \right\}.$$

For any history h_T , net returns conditional on h_T are

$$r_{\lambda,\pi}|h_T = \inf\left\{ \mathbb{E}_{c,a}\left(\frac{1}{Nw}\sum_{t=T+1}^N \lambda_t (w_t - w_t^0) - \pi_t \middle| h_T\right) \middle| (c,a) \text{ solves } \max_{c,a} \mathbb{E}\left(\frac{1}{N}\sum_{t=1}^N \pi_t - c_t\right) \right\}$$

As in Section 3, when the contract in question is the benchmark contract with parameter α , net returns are denoted by r_{α} and $r_{\alpha}|h_{T}$. **Theorem 2** (approximate performance). Pick $\alpha_0 \in (0,1)$ and for any $\eta \in (0,1)$, let $\alpha = \alpha_0 + \eta(1-\alpha_0)$. Consider the calibrated contract (λ, π) defined by (11) and (12). There exists a constant m independent of time horizon N and probability space \mathcal{P} such that,

$$r_{\lambda,\pi} \geq (1-\eta)r_{\alpha_0} - m\frac{1}{\sqrt{N}}$$
(13)

$$\forall h_T, \quad r_{\lambda,\pi} | h_T \geq (1-\eta) r_{\alpha_0} | h_T - m \frac{1}{\sqrt{N}}.$$
(14)

It follows that for N large enough, the calibrated contract described by (11) and (12) generates a share approximately $1 - \eta$ of the returns the principal obtains under the benchmark contract of parameter α_0 . The mechanics underlying Theorem 2, and the reason why an additional wedge η is needed will be discussed in Section 4.2.

Note that for Theorem 2 to hold, it is not sufficient to just reward the agent according to the payment rule $(\pi_t)_{t\geq 1}$ defined by (11) and (12). It is important that the principal actually invest only shares $(\lambda_t)_{t\geq 1}$ of her wealth according to the agent's suggestion. Indeed the reward scheme $(\pi_t)_{t\geq 1}$ does not to induce perfectly good behavior from the agent.¹⁸ Rather, the payment scheme $(\pi_t)_{t\geq 1}$ reduces misbehavior to the point where it can be resolved by using the cautious investment rule specified by $(\lambda_t)_{t>1}$.

4.2 The Mechanics of Calibrated Contracts

This idea behind calibrated contracts is to identify key incentive properties that hold under the benchmark contract and calibrate payments $(\pi_t)_{t\geq 1}$ to the agent as well as investment shares $(\lambda_t)_{t\geq 1}$ so that the same incentive properties are approximately satisfied under the calibrated contract. The properties that calibrated contracts attempt to satisfy are as follows.

¹⁸For instance, an agent who has lost or never had any informational advantage may systematically pick allocations a_t that are inferior to a_t^0 , simply because they are different and, through volatility, induce a non-zero probability of reward.

For all histories h_T ,

$$\Pi_T = \alpha S_T \tag{15}$$

$$\forall T' \le T, \quad S_{T \setminus T'} \ge \Sigma_{T \setminus T'}. \tag{16}$$

In words, the agent receives a share α of his actual performance S_T , and over any time interval $\{T', \dots, T\}$, his actual performance $S_{T\setminus T'}$ (although potentially hindered by investment shares $\lambda_t \leq 1$) is at least as high as his potential performance $\Sigma_{T\setminus T'}$.¹⁹ Note that the family of inequalities (16) can be summarized by the single inequality

$$\max_{T' \le T} \Sigma_{T \setminus T'} - S_{T \setminus T'} \le 0.$$

Let us now show how (11) and (12) calibrate parameters $(\lambda_t, \pi_t)_{t\geq 1}$ so that these properties hold approximately, while satisfying limited liability constraints (4) and (5). Define regrets

$$\mathcal{R}_{1,T} \equiv \Pi_T - \alpha S_T$$
$$\mathcal{R}_{2,T} \equiv \max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}$$

Regret $\mathcal{R}_{1,T}$ measures how overpaid the agent has been, while regret $\mathcal{R}_{2,T}$ measures maximum foregone profits from not fully investing according to the agent's allocation. The goals are: (*i*) to keep $\mathcal{R}_{1,T}$ small so that the agent's reward Π_T is a share approximately α of his actual externality S_T on the principal; (*ii*) to keep $\mathcal{R}^+_{2,T}$ small, so that the foregone returns are not large.

These goals can be achieved by following the approach of Blackwell (1956) and Hannan

¹⁹To obtain only inequality (13) in Theorem 2, it would be sufficient to consider only inequality $\Sigma_T \leq S_T$ rather than the full family of inequalities described by (16). Considering the full family of inequalities (16) yields the history independent performance bounds (14).

(1957).²⁰ Define $\mathcal{R}_T \equiv (\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ and $\rho_T \equiv \mathcal{R}_T - \mathcal{R}_{T-1}$ the vector of flow regrets.²¹ In order to keep regrets $(\mathcal{R}_{1,t})_{t\geq 1}$ and $(\mathcal{R}_{2,t})_{t\geq 1}$ small, it is sufficient to keep vector \mathcal{R}_T small. This can be achieved by choosing sequences $(\pi_t)_{t\geq 1}$ and $(\lambda_t)_{t\geq 1}$ so that

$$\forall T \ge 1, \ \forall w_{T+1}, \ \forall w_{T+1}^0, \qquad \left\langle \mathcal{R}_T^+, \rho_{T+1} \right\rangle \le 0.$$
(17)

Inequality (17) is known as an approachability condition, and ensures that flow regrets ρ_{T+1} point in the direction opposite to that of aggregate regrets \mathcal{R}_T . This puts strong bounds on the speed at which aggregate regrets $(\mathcal{R}_T)_{T\geq 1}$ can grow.

By construction

$$\mathcal{R}_{2,T+1} = \begin{cases} (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T} & \text{if } \mathcal{R}_{2,T} \ge 0\\ (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) & \text{if } \mathcal{R}_{2,T} < 0 \end{cases}$$

Hence, it follows that $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T}^+$. Thus, condition (17) is equivalent to

$$\left[\pi_{T+1} - \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)\right] \mathcal{R}_{1,T}^+ + \alpha^2 \left[(1 - \lambda_{T+1}) (w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T}^+ - \mathcal{R}_{2,T} \right] \mathcal{R}_{2,T}^+ \le 0.$$

Noting the identity $(\mathcal{R}_{2,T}^+ - \mathcal{R}_{2,T})\mathcal{R}_{2,T}^+ = 0$, it follows that approachability condition (17) is equivalent to

$$[\pi_{T+1} - \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)] \mathcal{R}_{1,T}^+ + \alpha^2 [(1 - \lambda_{T+1}) (w_{T+1} - w_{T+1}^0)] \mathcal{R}_{2,T}^+ \leq 0$$

$$\iff \pi_{T+1} \mathcal{R}_{1,T}^+ + \alpha [\lambda_{T+1} \mathcal{R}_{1,T}^+ + (1 - \lambda_{T+1}) \alpha \mathcal{R}_{2,T}^+] (w_{T+1} - w_{T+1}^0) \leq 0$$

Hence approachability condition (17) can be satisfied for any realization of w_{T+1} and w_{T+1}^0

²⁰See also Foster and Vohra (1999) or Cesa-Bianchi and Lugosi (2006). Regret measure $\mathcal{R}_{2,T}$ is specifically related to "tracking" regrets, as discussed in Cesa-Bianchi and Lugosi (2006).

²¹Vector \mathcal{R}_T is defined as $(\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ rather than $(\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ only because it leads to a slight improvement in performance bounds.

by setting

$$\lambda_{T+1} = \frac{\alpha \mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if} \quad \mathcal{R}_{1,T} \le 0\\ 0 & \text{if} \quad \mathcal{R}_{1,T} > 0 \end{cases}$$

which corresponds to the calibrated contract defined by (11) and (12).

The following lemma shows that under the contract defined by (11) and (12), incentive properties (15) and (16) are approximately satisfied. Recall that $d_t = \sup_{\mathbf{r} \in R} |\langle a_t - a_t^0, \mathbf{r} \rangle|$ denotes the magnitude of positions taken by the agent in period t.

Lemma 1 (approximate incentives). For all T, all $T' \leq T$ and all possible histories,

$$\Sigma_{T\setminus T'} - S_{T\setminus T'} \leq w_{N} \sum_{t=1}^{T} d_{t}^{2}$$
(18)

$$-\alpha w \overline{d} \leq \Pi_T - \alpha S_T \leq \alpha w \sqrt{\sum_{t=1}^T d_t^2}.$$
 (19)

In words, Lemma 1 means that incentive properties (15) and (16) hold at any possible history h_T , up to an error term of order $O(\sqrt{T})$. Note that this holds sample path by sample path, rather than in expectation or in equilibrium.

Proof. Let us first show that $||\mathcal{R}_T^+||^2 \leq \alpha^2 w^2 \sum_{t=1}^T d_t^2$. The proof is by induction. The property clearly holds at T = 1. Assume it holds at T. Consider the case where $\mathcal{R}_{2,T+1} > 0$. Since approachability condition (17) holds, we have that

$$\begin{aligned} ||\mathcal{R}_{T+1}^+||^2 &\leq ||\mathcal{R}_T^+ + \rho_{T+1}||^2 = ||\mathcal{R}_T^+||^2 + 2\left\langle \mathcal{R}_T^+, \rho_{T+1} \right\rangle + ||\rho_{T+1}||^2 \\ &\leq ||\mathcal{R}_T^+||^2 + ||\rho_{T+1}||^2. \end{aligned}$$

In addition $||\rho_{T+1}||^2 = \alpha^2 \lambda_{T+1}^2 (w_{T+1} - w_{T+1}^0)^2 + \alpha^2 (1 - \lambda_{T+1})^2 (w_{T+1} - w_{T+1}^0)^2 \leq \alpha^2 w^2 d_{T+1}^2$. Altogether this shows that the induction hypothesis holds when $\mathcal{R}_{2,T+1} > 0$. A similar proof holds when $\mathcal{R}_{2,T+1} < 0$, taking into account that in this case, $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})(w_{T+1}^1 - w_{T+1}^0)$. Hence, by induction, this implies that for all $T \ge 1$, $||\mathcal{R}_T^+||^2 \le \alpha^2 w^2 \sum_{t=1}^T d_t^2$. This proves (18) and the right-hand side of (19).

The left-hand side of (18) is also proven by induction. If $\Pi_T \in [\alpha S_T - \alpha w \overline{d}, \alpha S_T]$, then $\mathcal{R}_{1,T} = 0, \ \alpha_{T+1} = \alpha \text{ and } \lambda_T = 1$. Hence by construction, $\Pi_{T+1} \ge \alpha S_{T+1} - \alpha w \overline{d}$. If instead $\Pi_T > \alpha S_T$, then by definition of $\overline{d}, \ \Pi_{T+1} \ge \alpha S_{T+1} - \alpha w \overline{d}$. This implies the left-hand side of (19).

As the next lemma shows, the approximate incentive conditions given by Lemma 1 imply performance bounds for the corresponding contracts.

Lemma 2. Pick $\alpha_0 \in (0,1)$ and for any $\eta \in (0,1)$ let $\alpha = \alpha_0 + \eta(1-\alpha_0)$. Consider a contract (λ, π) and numbers A, B and C such that for all final histories $h_N, \Sigma_N - S_N \leq A$ and $-B \leq \prod_N - \alpha S_N \leq C$. Then

$$r_{\lambda,\pi} \ge (1-\eta)r_{\alpha_0} - \frac{1}{Nw} \left[C + \frac{1-\eta}{\eta} (\alpha A + B + C) \right].$$

Theorem 2 is an immediate corollary of Lemmas 2 and 1. Intuitively, Lemma 1 shows that the calibrated contract (λ, π) defined by (11) and (12) gets incentives approximately right. Lemma 2 implies that when incentives are approximately right, then performance must be approximately right as well. While this last result seems natural, it isn't immediate. Whenever global incentive constraints are binding or almost binding under the benchmark linear contract of parameter α_0 getting incentives slightly wrong may result in dismal performance. This would be the case if under the benchmark contract, the agent is indifferent between working hard and not working at all. By sharing an additional fraction η of her returns, the principal ensures that (almost) binding global incentive compatibility constraints do not compromise performance. Madarász and Prat (2010) make the same point in a screening context.

A Simulation. Figure 1 illustrates the mechanics of calibrated contracts, and Lemma 1 in particular. Figure 1(a) plots a sample path for potential accumulated returns $(\Sigma_T)_{T\geq 1}$. There is significant variance and sharp drops in performance are possible. Figure 1(b) shows the induced patterns of investment shares $(\lambda_T)_{T\geq 1}$. When performance drops, investment shares diminish and when performance improves, investment shares grow. Note that shares λ_t are continuous rather than 0-1. This is essential for Lemma 1 to hold.²² As Figure 1(c) illustrates, this implies that actual excess returns $(S_T)_{T\geq 1}$ track the growth of potential returns $(\Sigma_T)_{T\geq 1}$ but do not fall as much as $(\Sigma_T)_{T\geq 1}$ when performance drops. This allows to keep tight the relationship between cumulated rewards $(\Pi_T)_{T\geq 1}$ —which are necessarily weakly increasing and scaled actual returns $(\alpha S_T)_{T\geq 1}$. A direct implication of this, illustrated in Figure 1(d), is that the effective reward rate of the agent $(\Pi_T/S_T)_{T\geq 1}$ stays close to the target reward rate α . More precisely, poor performance leads to slow divergence while good performance leads to quick convergence. Because the effective reward rate is only approximately equal to α , this perturbs incentives a little bit and it is necessary to use a target reward rate α strictly greater than α_0 to emulate the performance of the benchmark contract with parameter α_0 .

4.3 Robustness to Accidents

Before turning to screening, it is worth noting an additional property of calibrated contracts: they are robust to the possibility of "unexpected accidents" during which the agent performs particularly badly over an extended amount of time. Figure 1(c) illustrates this in a striking way: whenever potential performance $(\Sigma_T)_{T\geq 1}$ drops by a significant amount, calibrated contracts significantly limit the extent of the drop in actual performance $(S_T)_{T\geq 1}$. In turn, this makes recovering from large performance drops possible.

This section expands on this point. Imagine that an accident can occur over some unknown time interval $[T_1, T_2]$ of arbitrary length. For instance, there may be a mistake in

 $^{^{22}}$ See Foster and Vohra (1999).



(a) sample path of potential performance Σ_T



(b) induced investment shares λ_T



0.5 0.4 0.3 0.2 0.1 0.1 0.0 -0.1 -0.2 0 200 400 time T 600 800 1000

(c) scaled potential returns $\alpha \Sigma_T$, scaled actual returns αS_T , rewards Π_T



Figure 1: the behavior of calibrated contracts for a given sample path of potential returns $(\Sigma_T)_{T\geq 1}$, with target reward rate $\alpha = 20\%$.

the agent's trading strategy, a bias in his data, or the agent may be temporarily irrational. During time interval $[T_1, T_2]$ —in the accident state—the agent uses an exogenously specified allocation strategy a_t^{Δ} . This strategy may be arbitrarily bad (within the bounds imposed by Assumption 1) and need only be measurable with respect to \mathcal{F}_N . For instance, during the lapse of the accident, the agent could pick the worst ex post asset allocation in every period.

In this environment, the benchmark linear contract is no longer sufficient to guarantee

good performance. Accidents can undo all the profit generated by the well incentivized agent in his normal state. Strikingly, in spite of accidents, calibrated contracts are such that the excess returns generated by the agent will be approximately as high as the returns he could generate when accidents are lucky, i.e. when the exogenous allocation during accident states is

$$\forall T \in [T_1, T_2], \quad a_T^{\Delta \Delta} = \begin{cases} a_T^0 & \text{if } \sum_{t=T_1}^{T_2} w_t^{\Delta} - w_t^0 < 0\\ a_T^{\Delta} & \text{if } \sum_{t=T_1}^{T_2} w_t^{\Delta} - w_t^0 > 0 \end{cases}$$

where w_t^{Δ} is the realized wealth under the a_t^{Δ} at time t. Denote by $r_{\lambda,\pi}^{\Delta}$ the net expected returns to the principal when accidental behavior is a_t^{Δ} and the calibrated contract is used. Denote by $r_{\alpha}^{\Delta\Delta}$ the net expected returns to the principal when accidental behavior is $a_t^{\Delta\Delta}$ and the benchmark contract of parameter α is used. The following holds.

Theorem 3 (accident proofness). Pick α_0 and for any $\eta > 0$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists a constant m, independent of N and \mathcal{P} such that,

$$r_{\lambda,\pi}^{\Delta} \ge (1-\eta)r_{\alpha_0}^{\Delta\Delta} - \frac{m}{\sqrt{N}}$$

5 Screening

As it is, the calibrated contract defined by (11) and (12) does not induce untalented managers to screen themselves. Rewards are positive, and a sufficiently long-lived uninformed agent can obtain large expected payoffs from luck and volatility alone. Indeed, imagine that the agent has no information and all assets have the same expected returns. By systematically picking assets different from the benchmark allocation, the agent ensures that $(\Sigma_T)_{T\geq 1}$ is a martingale with volatility bounded away from 0. Hence, under appropriate time normalization, $(\Sigma_T)_{T\geq 1}$ behaves like a Brownian motion.²³ Lemma 1 implies that the agent's payoff in period Tsatisfies $\Pi_T \geq \alpha \Sigma_T - \alpha w \overline{d} (1 + \sqrt{T})$. In addition, since Π_T is weakly increasing in T, it

²³See, for instance, Billingsley (1995), Theorem 35.12.

follows that $\Pi_N \geq \max_{T \leq N} \Sigma_T - \alpha w \overline{d} (1 + \sqrt{N})$. Hence, given that $\max_{T \leq N} \Sigma_T$ behaves approximately like the maximum of a Brownian motion, the agent can obtain rewards of order \sqrt{N} with non-vanishing probability.

A simple modification of the contract described by (11) and (12) achieves screening by imposing an initial participation cost -b and only paying the agent when his performance is above a dynamic hurdle Θ_T which depends on the size of positions he has been taking. Given a free parameter M > 0, define

$$\Theta_T \equiv 2w \left(1 + \sqrt{\overline{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2} \right) \sqrt{M + \ln\left(\overline{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2\right)},$$
(20)

where $d_t = \sup_{\mathbf{r}_t \in \mathbb{R}} |\langle a_t - a_t^0, \mathbf{r}_t \rangle|$ and $\lambda_t d_t$ measures the size of the agent's effective bet $\lambda_t (a_t - a_t^0)$ away from the default allocation a_t^0 (note that by Assumption 1, $d_t \leq \overline{d}$). Hurdle Θ_T is an aggregate measure of how active the agent has been. If the agent makes significant bets away from a_t^0 in every period then Θ_T will be of order $\sqrt{T \ln T}$. If the agent makes few bets, hurdle Θ_T will remain small. The quantity $\overline{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2$ is a measure of time under which $(\Sigma_T)_{T\geq 1}$ will have at most the variation of a standard Brownian motion.

Hurdled calibrated contracts are defined by a sequence $(\lambda_t, \hat{\pi}_t, \pi_t^{\Theta})_{t\geq 1}$. Value λ_t is still the share of wealth actually invested by the agent. For t > 1, reward π_t is paid to the agent if and only if $S_t \ge \Theta_t$. For t > 1, the actual (hurdled) reward π_t^{Θ} is therefore $\pi_t^{\Theta} = \mathbf{1}_{S_t \ge \Theta_t} \pi_t$.²⁴ For t = 1, $\pi_1^{\Theta} = -b$.

This hurdled contract coincides with the baseline calibrated contract defined in Section 4, except that: (i) the agent must pay a participation fee -b in the first period, and (ii) the agent obtains rewards only when actual performance S_T is above a hurdle Θ_T which grows at a rate at most $\sqrt{T \ln T}$. Theorem 4 will shows that this contract induces uninformed agents to screen themselves in the first period, and imposes only a moderate incentive cost

²⁴Note that recursion equations (11) and (12) still use $\Pi_T = \sum_{t=1}^T \pi_t$ rather than hurdled aggregate payment $\Pi_T^{\Theta} = \sum_{t=1}^T \pi_t^{\Theta}$.

on informed agents.

An intuitive rationale for the form of hurdle Θ_T is as follows. Imagine for simplicity that the agent is frequently active, i.e. $\sum_{t=1}^{T} d_t^2$ is of order T. Then hurdle Θ_T is of order $\sqrt{T \ln T}$. As has been discussed, an uninformed agent can guarantee that Σ_T is comparable to the maximum of a Brownian motion. The law of the iterated logarithm implies that with probability 1, as T gets large, $\max_{T' \leq T} \Sigma_{T'}$ is of order $\sqrt{T \ln \ln T}$.²⁵ Because $\frac{\sqrt{T \ln \ln T}}{\sqrt{T \ln T}}$ goes to 0 as T grows large, hurdle Θ_T insures that uniformed agents have very little hope to obtain unjustified returns. Indeed, the following result holds.

Lemma 3 (hurdle effectiveness). If the agent is uninformed, then for any allocation strategy a,

$$\mathbb{E}_a\left(\sum_{t=1}^N \mathbf{1}_{S_t \ge \Theta_t}\right) \le \frac{\pi^2}{2} \exp(-2M),$$

where π is the constant 3.1415...

As the next lemma shows, the use of hurdles comes only at a moderate incentive cost.

Lemma 4 (approximate incentives). For all T, T' < T, and all paths of play, we have that

$$\Sigma_{T\setminus T'} - S_{T\setminus T'} \leq w \sqrt{\sum_{t=1}^{T} d_t^2}$$
(21)

$$-\alpha\Theta_T - \alpha w\overline{d} - b \leq \Pi_T - \alpha S_T \leq \alpha w \sqrt{\sum_{t=1}^T d_t^2}.$$
 (22)

Combining Lemmas 2, 3 and 4 yields the main result of this section. Denote by $r_{\lambda,\pi\Theta}$ the net expected per-period returns generated by the agent under the hurdled calibrated contract.

Theorem 4 (performance with screening). Pick $\alpha_0 \in (0,1)$ and for any $\eta \in (0,1)$, let $\alpha > \alpha_0 + \eta(1-\alpha_0)$. There exists a constant m independent of time horizon N and probability

²⁵See Billingsley (1995), Theorem 9.5.

space \mathcal{P} such that for all h_T ,

$$r_{\lambda,\pi^{\Theta}}|h_T \ge (1-\eta)r_{\alpha_0}|h_T - m\sqrt{\frac{\ln N}{N}}$$
(23)

Furthermore, whenever $-b + \alpha w \overline{d} \times \frac{\pi^2}{2} \exp(-2M) < 0$, it is strictly optimal for uninformed agents not to participate.

The combination of initial fee -b and hurdle Θ_t induces early screening by uninformed agents. Hurdle Θ_t is large enough that uninformed agents have little hope to be rewarded by luck but small enough that it does not significantly affect the incentives of informed agents. The penalty which was of order $\frac{1}{\sqrt{N}}$ in Theorem 2 is now of order $\sqrt{\frac{\ln N}{N}}$. The next lemma provides conditions under which the performance loss from screening is in fact of order $\frac{1}{\sqrt{N}}$.

Assumption 2 (grainy returns). Let (c, a^*) denote the agent's policy under the benchmark contract with rate α_0 . There exists $\xi > 0$ such that whenever $\mathbb{E}_{c,a^*}[w_t - w_t^0 | \mathcal{F}_t] > 0$, then $\mathbb{E}_{c,a^*}[w_t - w_t^0 | \mathcal{F}_t] > \xi$.

Theorem 5. Pick α_0 and for any $\eta > 0$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. If Assumption 2 holds, there exists a constant m such that for all N and all probability spaces \mathcal{P} ,

$$r_{\lambda,\pi^{\Theta}} \ge (1-\eta)r_{\alpha_0} - m\frac{1}{\sqrt{N}}$$

Indeed, whenever Assumption 2 holds it can be shown that hurdles grow at a slow rate compared to expected excess returns. In fact the expected number of payments that are omitted because of hurdles is bounded above independently of N.

6 Discussion

6.1 Relation to High-Watermark Contracts

High-watermark contracts. The calibrated contracts described in Section 4 and 5 are closely related to the high-watermark contracts frequently used in the financial industry (see for instance Goetzmann et al. (2003) or Panageas and Westerfield (2009)). High-watermark contracts are structured as follows: at time T, the investment share λ_T is always 1, and the agent gets paid

$$\pi_T^{\text{wmk}} = \alpha \left(\sum_{t=1}^T w_t - w_t^0 - \max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right] \right)^+.^{26}$$
(24)

Quantity $\max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right]$ represents the maximum historical returns at time T i.e the high-watermark. The agent only gets paid when he improves on his own historical performance. Note that high-watermark contracts are dynamic and satisfy limited liability constraints (4) and (5). In particular, for all $T, \pi_T \in [0, \alpha w \overline{d}]$.

High-watermark contract, as well as calibrated contracts, attempt to reward the agent a share α of his externality on the principal. In other words, both types of contracts attempt to keep aggregate rewards Π_T close to αS_T . Lemma 1 shows that calibrated contracts achieve this goal for any realization of uncertainty and any allocation strategy. High-watermark contracts also do well in this respect, provided that the sequence of returns does not have prolonged downturns. For instance, if an agent has performance $(1, 1, 1, 1, \dots)$, then $S_N = N$ and $\Pi_N^{\text{wmk}} = \alpha N = \alpha S_N$. If the agent has performance $(1, -1, 1, -1, \cdot)$, then $S_N \in \{0, 1\}$ and $\Pi_N^{\text{wmk}} = \alpha$, so that $\Pi_N^{\text{wmk}} = \alpha S_N + o(N)$. In this respect, high watermark contracts do much better than static option contracts of the form $\pi_t = \alpha (w_t - w_t^0)^+$, which would reward the agent $\alpha N/2$.

²⁶This is in fact slightly different from standard high-watermark contracts. Payments under highwatermark contracts are typically based on gross performance $w_t - w$ rather than on performance net of default returns $w_t - w_t^0$.

The value of jointly calibrating rewards and investment shares. An important issue with high-watermark contracts—which calibrated contracts resolve—is that they do not behave well if the agent's performance goes through an extended downturn. This is illustrated by Figure 2(a). Whenever there is an extended drop in performance, the relationship between rewards Π_T and performance αS_T breaks down. Indeed Π_T is by construction weakly increasing while under the high-watermark contract, αS_T can decrease in arbitrary ways. This has two implications. First, because the agent does not suffer from extended downturns, an agent who has lost the ability to generate positive return (e.g. his information has become unreliable) may prefer to make negative expected value investments that generate variance, rather than admit he has become uninformed. Second, if a talented agent has been unlucky and experienced a drop in returns, the difficulty of catching up with a high watermark may discourage investment altogether.

As Figure 2(b) shows, calibrated contracts keep tight the relationship between Π_T and αS_T by calibrating investment shares so that actual performance S_T has limited downturns without losing too much on the growth of potential returns Σ_T . As a result, extended downturns have a much more limited impact on incentives. As was noted previously, investment shares λ_t must move smoothly with performance. Rather than a stop-loss provision, it is more accurate to think of the calibrated investment shares $(\lambda_t)_{t\geq 0}$ as continuously implementing a robust option on the agent's potential performance Σ_T .²⁷

Screening. Unlike high-watermark contracts, the contracts described in this paper achieve screening by using hurdles that link payoffs to the size of positions that the agent has been taking. Given an initial participation cost -b, these activity based hurdles insure that it is not viable strategy for uninformed agents to participate. Note that these hurdles depend only on activity measure $d_t = \sup_{\mathbf{r}_t \in R} |\langle a_t - a_t^0, \mathbf{r}_t \rangle|$. In particular the actual asset allocation

²⁷See DeMarzo et al. (2009) for a discussion of the relation between approachability methods and robust option pricing.



wards Π_T

Figure 2: high watermark and calibrated contracts for a sample path of potential returns $(\Sigma_T)_{T\geq 1}$, with target reward rate $\alpha = 20\%$.

chosen by the agent need not be made public, which matters if the agents are worried about privacy.

Note also that while such hurdles reduce incentives, they are specifically chosen to have minimal impact on talented agents. In particular, whenever returns are grainy, talented agents can deliver significant returns without taking large aggregate positions and the use of hurdles hardly reduces their payoffs.

In addition note that the use of an accurate counter-factual performance measure w_t^0 is essential for the success of the screening strategy. Indeed, it ensures that whenever the agent is uninformed, then $\Sigma_T = \sum_t^T w_t - w_t^0$ is a martingale with weakly negative drift (i.e. a surmartingale). Not all contracts have this property. Imagine for instance that the an ex ante counter factual performance measure $\overline{w}_t^0 = \mathbb{E}(w_t^0 | \mathcal{F}_1)$ is used. Because, $\overline{\Sigma}_T = \sum_{t=1}^T w_t - \overline{w}_t^0$ is *not* a martingale with weakly negative drift, there can be non-vanishing probability that a fixed asset allocation beat period 1 expectations \overline{w}_t^0 arbitrarily many times. As a result screening cannot be achieved using ex ante counterfactual performance measures.

Absent Features. Finally, a surprising property of the calibrated contracts developed in Sections 4 and 5 is that they do not require the use of contractual provisions often deemed necessary for good performance: agents are not required to hold a significant amount of the asset allocation they are suggesting and there is no use of clawbacks or deferred payments. Of course this doesn't mean that such provision aren't useful, especially if the agent's horizon is small.

6.2 Future Work

This paper develops a robust approach to dynamic contracting in two steps: the first step identifies high-liability linear contracts that satisfy attractive efficiency properties regardless of the underlying environment; the second step shows how to approximate the performance of benchmark contracts using limited-liability dynamic contracts. The contracting strategy is to calibrate rewards to the agent as well as the share of wealth he manages, so that key properties of the benchmark contract are approximately replicated. The resulting calibrated contracts are simple, and perform approximately as well as an attractive benchmark under very general conditions. The simplicity of the results is encouraging and suggests that the approach might be fruitfully applied in other settings.

From a theoretical perspective, a first valuable extension would be to allow for riskaversion on both sides. Appendix A develops extensions to environments where the principal is risk averse or the agent may have varying marginal utility of income, but more works remains to be done on that issue. For instance, the calibration strategy described in Section 4 results in unnecessary variation in the rewards to the agent: if he has been performing well, but underperforms one period, he does not receive rewards in the following period. When the agent is risk-averse, such a calibration strategy leads to inefficiencies and smoother calibration techniques become desirable.

Another avenue for research is to develop robust approaches to pick an appropriate reward

rate α for the benchmark linear contract. In this respect, it may be fruitful to consider multiagent mechanisms that may help extract information from agents. Considering multi-player environments may also be interesting beyond the principal-agent setting that this paper focuses on. For instance many attractive allocation mechanisms, such as VCG, require agents to make significant payments and are therefore ill-suited in environments where agents are severely cash constrained. A dynamic calibration approach such as the one developed in this paper may help relax such limited liability constraints.

Finally, with respect to applications, it seems important to determine whether calibrated contracts really do induce approximately good behavior from agents. Indeed, Theorems 2 and 4 may place unrealistic confidence in the agent's ability to understand the incentive properties of calibrated contracts. This is ultimately an empirical question. An advantage of the calibrated contract approach is that it lends itself naturally to realistic experiments using actual returns data, since the contracts should perform well regardless of the agent's beliefs over the process for returns.

Appendix

A Additional Results and Extensions

A.1 Static Contracts under Limited Liability

This appendix shows that there is no limited-liability static contract that can approximate the performance of the high-liability benchmark contract described in Section 3 (even though the benchmark contract is itself static). Option-like static contracts that reward the agent according to $\alpha(w_t - w_t^0)^+$ have well-known issues: uninformed agents can obtain large payoffs, and talented agents may be induced to choose asset allocations with large variance and negative expected value. All static limited-liability contracts suffer from similar issues. Consider contracts such that for all $t \ge 1$ reward π_t depends only on returns at time t (i.e. $\pi_t = \pi_t(w_t, w_t^0)$) and $\pi_t \in [-b, w_t]$. Note that this is actually a weaker limited liability constraint than that imposed by conditions (4) and (5). The following lemma uses an argument similar to that of Foster and Young (2010) to show that such contracts cannot simultaneously screen agents and reward them a share of the surplus they create. The proof requires that the ex post optimal asset allocation be uncertain from the perspective of public information, i.e. any selection $a_t \in \arg \max_{a \in A} \langle a, \mathbf{r}_t \rangle$ is not \mathcal{F}_t^0 -measurable.

Lemma A.1. Consider a one-shot reward function π_t satisfying limited liability, and such that for all distributions w_t of realized wealth,

$$\mathbb{E}\pi_t \ge \alpha \mathbb{E}[w_t - w_t^0].$$

There exists $\underline{\alpha} > 0$ such that for w large enough, an untalented manager can obtain expected profit at least $\underline{\alpha}w$.

Proof. Let \overline{a}_t denote an expost optimal asset allocation, i.e. $\overline{a}_t \in \arg \max_{a \in A} \langle a, \mathbf{r}^t \rangle$. By assumption, we have that $\overline{r} \equiv \mathbb{E} \langle \overline{a}_t - a_t^0, \mathbf{r}^t \rangle > 0$. Since A is a compact set, for M large enough, there exists M points $\{a_1, \dots, a_M\} \subset A$ such that for all $a \in A$, there exists $k \in \{1, \dots, M\}$ satisfying for all \mathbf{r}_t , $| < a_k - a, \mathbf{r}_t > | \leq \overline{r}/2$. Hence there is a selection \overline{a}_t^M of assignments in $\{a_1, \dots, a_M\}$ such that $\mathbb{E} \langle \overline{a}_M - a_t^0, \mathbf{r}_t \rangle > \overline{r}/2$.

Consider the allocation strategy \tilde{a} , consisting of independently and uniformly picking an allocation $a \in \{a_1, \dots, a_M\}$. The agent's expected payment satisfies

$$\mathbb{E}_{\tilde{a}}\pi_t \ge \frac{1}{M}\mathbb{E}_{\overline{a}_t^M}\pi_t \ge \frac{1}{2M}\alpha\overline{r}w.$$

This concludes the proof.

A corollary of this is that if the agent receives significant rewards using static limited liability contracts, there can be no screening. Such misaligned incentives can also affect

talented managers. Imagine that at each time t there is positive probability that the information I_t^c acquired by the manager is such that the optimal portfolio is the same under \mathcal{F}_t and \mathcal{F}_t^0 , i.e. expert information is not helpful. When this happens, the manager cannot deliver excess returns. However, by deviating from truthtelling, the manager can obtain an expected payoff of at least $\underline{\alpha}w$.

A.2 A More General Principal-Agent Framework

This appendix extends the analysis to a principal-agent framework more general than the financial contracting problem studied in the paper. The principal and the agent are still risk-neutral, but in every period the agent suggests and implements an action $a_t \in A$, where A is a potentially non-convex set of actions. Every period, a state of the world \mathbf{r}_t is drawn, which given action a yields observable payoffs $w(a, \mathbf{r}_t)$ to the principal. Cost c_t may now represent the cost of information, as well as the cost of making a specific action available. Action a_t^0 is the action that the principal would (could) implement on her own. The main difference is that because set A need not be convex, the principal must use randomized strategies to calibrate her contract with the agent.

The calibrated contract of Section 4 can be adapted as follows. Parameter λ_t now denotes the probability that the principal follow the action suggested by the agent.²⁸ Let a_t^{λ} denote the action actually taken at time t. Denote by $\psi_t \equiv w(a_t, \mathbf{r}_t) - w(a_t^0, \mathbf{r}_t)$ the potential excess returns and by $\psi_t^{\lambda} \equiv w(a_t^{\lambda}, \mathbf{r}_t) - w(a_t^0, \mathbf{r}_t)$ the realized excess returns. As in Section 4, define

$$\Sigma_T = \sum_{t=1}^T \psi_t, \ \Pi_T = \sum_{t=1}^T \pi_t, \ S_T = \sum_{t=1}^T \psi_t^{\lambda},$$

as well as $\Sigma_{T\setminus T'} = \Sigma_T - \Sigma_{T'-1}$, $\Pi_{T\setminus T'} = \Pi_T - \Pi_{T'-1}$ and $S_{T\setminus T'} = S_T - S_{T'-1}$. As in Section

²⁸For calibration results to hold, it is important that the agent not be able to condition his suggested action on the outcome of the principal's randomization. Note that the agent may also take the action on behalf of the principal. In that case λ_t should be interpreted as the probability that the principal approve the agent's proposed course of action.

4, regrets are defined by

$$\mathcal{R}_{1,T} \equiv \Pi_T - \alpha S_T$$
 and $\mathcal{R}_{2,T} \equiv \max_{T' < T} \Sigma_{T \setminus T'} - S_{T \setminus T'}.$

Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$. Contract $(\lambda_t, \pi_t)_{t \in \mathbb{N}}$ is defined by

$$\lambda_{T+1} = \frac{\alpha \mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} 0 & \text{if} \quad \mathcal{R}_{1,T} > 0\\ \alpha \psi_{T+1}^+ & \text{if} \quad \mathcal{R}_{1,T} \le 0 \end{cases}$$
(25)

with the convention that $\frac{0}{0} = 1$. Lemma 1 extends as follows.

Lemma A.2 (approximate incentives). For all T, and any strategy (c, a) of the agent, we have that

$$\mathbb{E}_{c,a}\Sigma_T - \mathbb{E}_{c,a}S_T \leq w\overline{d}\sqrt{T}$$
(26)

$$-\alpha w \overline{d} \leq \Pi_T - \alpha \mathbb{E}_{c,a} S_T \leq \alpha w \overline{d} \sqrt{T}.$$
(27)

Proof. The left-hand side of (27) follows from a proof identical to that of the left-hand side of (19).

Let us turn to the other inequalities. Let $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$ denote flow regrets, and observe that $\mathbb{E}_{c,a}\left(\left\langle \mathcal{R}_{T-1}^+, \rho_T \right\rangle\right) \leq 0$. Hence, a proof identical to that of Lemma 1 yields that

$$\forall (c,a), \quad \mathbb{E}_{c,a} ||\mathcal{R}_T^+||^2 \le (\alpha \overline{d}w)^2 T.$$

Finally note that by Jensen's inequality, for all $i \in \{1, 2\}$,

$$\mathbb{E}_{c,a}(\mathcal{R}_{i,t}^{+}) \leq \mathbb{E}_{c,a}\left(\sqrt{\left[\mathcal{R}_{i,t}^{+}\right]^{2}}\right) \leq \sqrt{\mathbb{E}_{c,a}\left(\left[\mathcal{R}_{i,t}^{+}\right]^{2}\right)} \leq w\overline{d}\sqrt{T}.$$

This implies (26) and the right-hand side of (27).

Given Lemma A.2, a proof identical to that of Lemma 2 yields the following performance bound: pick α_0 , $\eta > 0$ and let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists *m* independent of *N* and \mathcal{P} such that

$$r_{\lambda,\pi} \ge (1-\eta)r_{\alpha_0} - \frac{m}{\sqrt{N}}$$

where returns r generated by some contract are given by $r = \mathbb{E}\left[\frac{1}{N}\sum_{t=1}^{N}w(a_t^{\lambda},\mathbf{r}_t) - w(a_t^{0},\mathbf{r}_t)\right]$, and the expectation depends on the policy induced by the relevant contract.

A.3 Risk Aversion

This paper considers the case where both the principal and the agent have quasilinear preferences. Extending the analysis to the case where either the principal or the agent are risk averse presents a number of challenges, most of which are left for future research. A significant difficulty relates to the provision of insurance by the agent. Indeed, if a significant portion of first best surplus is derived by having the agent provide insurance to the principal, it seems unlikely that such surplus can be generated using only positive transfers from the principal to the agent. This section is able to provide a partial extension of the calibrated contracts analyzed in the paper to the case where the agent is risk neutral while the principal is risk averse.

Consider an increasing concave utility function u. This section considers the case where the agent is risk-neutral while the principal has utility function u over flow wealth. This section shows how to construct calibrated contracts such that payoffs to the agent and residual utility to the principal satisfy

$$\sum_{t=1}^{N} \pi_t \simeq \nu \sum_{t=1}^{N} [u(w_t - \pi_t) - u(w_t^0)], \qquad (28)$$

where $\nu > 0$ is a design parameter used to shift surplus between the principal and the agent. If this condition holds, it ensures that whenever the agent gets positive surplus, the principal must obtain a commensurate expected payoff.²⁹ Note that if u(w) = w, then condition (28) simply boils down to the condition that $\sum_{t=1}^{N} \pi_t \simeq \alpha \sum_{t=1}^{N} w_t - w_t^0$ with $\alpha = \nu/(1+\nu)$, i.e. the framework considered in the bulk of the paper.

Preliminaries. Let us begin by providing a generalization of Assumption 1.

Assumption 3. There exists \overline{d} such that for all For all $(a, a') \in A^2$ and all $\mathbf{r}_t \in R$,

$$|u(w\langle a, \mathbf{r}_t \rangle) - u(w\langle a', \mathbf{r}_t \rangle)| \le \overline{d}.$$

In addition, it is assumed that there exists $\kappa > 0$ such that for any possible realized wealth $w_t, u'(w_t) \leq \kappa$. Let $\phi(\cdot, \cdot)$ denote the implicit function uniquely defined by

$$\forall w_1, w_0, \quad \phi(w_1, w_0) = \nu \left[u(w_1 - \phi(w_1, w_0)) - u(w_0) \right], \tag{29}$$

Note that by construction, $|\phi(w_t, w_t^0)| \leq \nu \overline{d}$. The following properties will be useful in the analysis.

Lemma A.3. (i) $\phi(w, w) = 0$ for all w;

(ii) $\phi(\cdot, \cdot)$ is increasing and concave in its first argument;

Proof. Point (i) and the fact that ϕ is increasing in its first argument follow immediately from (29). The fact that $\phi(w_1, w_0)$ is concave in w_1 follows from concavity of u. For any values, w_0 , w_1 , w_2 and $\rho \in (0, 1)$, let us define

$$\phi_1 = \nu[u(w_1 - \phi_1) - u(w_0)] \quad ; \quad \phi_2 = \nu[u(w_2 - \phi_2) - u(w_0)]$$

$$\phi_\rho = \rho\phi_1 + (1 - \rho)\phi_2 \quad \text{and} \quad w_\rho = \rho w_1 + (1 - \rho)w_2.$$

²⁹An important motivating example for this extension is the case in which $u = \log$ and $w_t^0 = w \langle a_t^0, \mathbf{r}_t \rangle$, with $a_t^0 \in \arg \max_{a \in A} \mathbb{E} \left[\log(w \langle a, \mathbf{r}_t \rangle) | \mathcal{F}_t^0 \right]$. The corresponding calibrated contract would be appropriate if wealth is accumulated (with compounded returns) and the principal has log utility over the final outcome.

By concavity of u we have that

$$\nu[u(w_{\rho} - \phi_{\rho}) - u(w_{0})] \geq \nu[\rho u(w_{1} - \phi_{1}) + (1 - \rho)u(w_{2} - \phi_{2}) - u(w_{0})]$$

$$\geq \rho \phi_{1} + (1 - \rho)\phi_{2}.$$

Hence, it must be that $\phi(w_{\rho}, w_0) \ge \rho \phi(w_1, w_0) + (1 - \rho) \phi(w_2, w_0)$, i.e. $\phi(\cdot, \cdot)$ is concave in its first argument.

Calibrated contracts. For any sequence of adapted investment shares $\lambda = (\lambda_t)_{t \ge 0}$ and actual payments $(\pi_t)_{t \ge 0}$, let $w_t^{\lambda} = \lambda_t w_t + (1 - \lambda_t) w_t^0$, and

$$\Sigma_T = \sum_{t=1}^T \nu[u(w_t - \pi_t) - u(w_t^0)]; \quad \Pi_T = \sum_{t=1}^T \pi_t; \quad S_T = \sum_{t=1}^T \nu[u(w_t^\lambda - \pi_t) - u(w_t^0)]. \quad (30)$$

For any T' < T, let $\Sigma_{T \setminus T'} = \Sigma_T - \Sigma_{T'-1}$, $\Pi_{T \setminus T'} = \Pi_T - \Pi_{T'-1}$ and $S_{T \setminus T'} = S_T - S_{T'-1}$. In addition, let

$$\mathcal{R}_{1,T} = \Pi_T - S_T$$
 and $\mathcal{R}_{2,T} = \max_{T' < T} [\Sigma_{T \setminus T'} - S_{T \setminus T'}]$

The objective is to calibrate payments $(\pi_t)_{t\geq 0}$ and investment shares $(\lambda_t)_{t\geq 0}$ so that $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}^+$ be remain small compared to T.

As in Section 4, this can be achieved by setting

$$\lambda_{T+1} = \frac{\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} 0 & \text{if} \quad \mathcal{R}_{1,T} > 0\\ \phi(w_t^\lambda, w_t^0)^+ & \text{if} \quad \mathcal{R}_{1,T} \le 0 \end{cases}$$
(31)

Indeed, under this calibrated contract the following extension of Lemma 1 holds.

Lemma A.4. For all T and T' < T, we have that

$$\Sigma_{T \setminus T'} - S_{T \setminus T'} \leq \nu \overline{d} \sqrt{T} \tag{32}$$

$$-\nu \overline{d} \leq \Pi_T - S_T \leq \nu \overline{d} \sqrt{T}.$$
 (33)

Proof. The first part of the proof shows that as T grows large, $||\mathcal{R}_T^+||$ remains small compared to T. Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ and $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$. We have that

$$\mathcal{R}_{2,T+1} = \nu \left[u(w_t - \pi_t) - u(w_t^0) - \left[u(w_t^\lambda - \pi_t) - u(w_t^0) \right] \right] + \mathcal{R}_{2,T}^+$$

In addition $\mathcal{R}^+_{2,T}(\mathcal{R}^+_{2,T} - \mathcal{R}_{2,T}) = 0$. Hence, it follows that

$$\left\langle \mathcal{R}_{T}^{+}, \rho_{T+1} \right\rangle = \mathcal{R}_{1,T}^{+} \pi_{T+1} + \nu \left(\mathcal{R}_{2,T}^{+} [u(w_{t} - \pi_{t}) - u(w_{t}^{0})] - [\mathcal{R}_{1,T}^{+} + \mathcal{R}_{2,T}^{+}] [u(w_{t}^{\lambda} - \pi_{t}) - u(w_{t}^{0})] \right).$$

We want to show that $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle \leq 0$. Consider first the case where $\pi_{T+1} = 0$. Given concavity of u, simple algebra yields that

$$\langle \mathcal{R}_{T}^{+}, \rho_{T+1} \rangle \leq \mathcal{R}_{1,T}^{+} \pi_{T+1} + \nu \left(\mathcal{R}_{2,T}^{+} - \lambda_{T+1} [\mathcal{R}_{1,T}^{+} + \mathcal{R}_{2,T}^{+}] \right) \left[u(w_{t} - \pi_{t}) - u(w_{t}^{0}) \right]$$

$$\leq 0,$$

where the last inequality follows from the fact that $(\lambda_t, \pi_t)_{t\geq 0}$ satisfies (31). Now consider the case where $\pi_{T+1} > 0$. By construction, we must have that $w_{T+1} > w_{T+1}^0$ and $\pi_{T+1} = \phi(w_{T+1}^{\lambda}, w_{T+1}^0)$. By definition of ϕ and using the concavity of ϕ in its first argument (Lemma A.3), as well as the fact that $(\lambda_t, \pi_t)_{t\geq 0}$ satisfies (31), we obtain that

$$\langle \mathcal{R}_{T}^{+}, \rho_{T+1} \rangle = \mathcal{R}_{1,T}^{+} \pi_{T+1} + \mathcal{R}_{2,T}^{+} \phi(w_{T+1}, w_{T+1}^{0}) - [\mathcal{R}_{1,T}^{+} + \mathcal{R}_{2,T}^{+}] \phi(w_{T+1}^{\lambda}, w_{T+1}^{0})$$

$$\leq \mathcal{R}_{1,T}^{+} \pi_{T+1} + \left(\mathcal{R}_{2,T}^{+} - \lambda_{T+1} [\mathcal{R}_{1,T}^{+} + \mathcal{R}_{2,T}^{+}] \right) \phi(w_{T+1}, w_{T+1}^{0}) \leq 0.$$

We now prove by induction that $||\mathcal{R}_T^+||^2 \leq (\nu \overline{d})^2 T$. The property clearly holds for T = 1. Assume that it holds for $T \geq 1$ and let us show it must hold for T + 1. Consider first the case where $\mathcal{R}_{2,T} > 0$. We have that

$$||\mathcal{R}_{T+1}^{+}||^{2} \leq ||\mathcal{R}_{T}^{+}||^{2} + 2\left\langle \mathcal{R}_{T}^{+}, \rho_{T+1} \right\rangle + ||\rho_{T+1}||^{2} \leq ||\mathcal{R}_{T}^{+}||^{2} + (\nu \overline{d})^{2}$$
(34)

where we used the fact that $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle \leq 0$, and

$$\begin{aligned} ||\rho_{T+1}||^2 &\leq \nu^2 \left([u(w_{T+1}^{\lambda} - \pi_{T+1}) - u(w_{T+1}^0)]^2 \\ &+ [u(w_{T+1}^{\lambda} - \pi_{T+1}) - u(w_{T+1}^0) - (u(w_{T+1} - \pi_{T+1}) - u(w_{T+1}^0))]^2 \right) \\ &\leq \nu^2 \left([u(w_{T+1} - \pi_{T+1}) - u(w_{T+1}^0)]^2 + 2[u(w_{T+1}^{\lambda} - \pi_{T+1}) - u(w_{T+1}^0)]^2 \\ &- 2[u(w_{T+1} - \pi_{T+1}) - u(w_{T+1}^0)][u(w_{T+1}^{\lambda} - \pi_{T+1}) - u(w_{T+1}^0)] \right) \\ &\leq \nu^2 [u(w_{T+1} - \pi_{T+1}) - u(w_{T+1}^0)]^2. \end{aligned}$$

The last inequality uses the fact that $\phi(w_t^{\lambda}, w_t^0)^2 \leq \phi(w_t, w_t^0)\phi(w_t^{\lambda}, w_t^0)$. This is implied by the fact that: ϕ is increasing in its first argument; $\phi_{\nu}(w, w) = 0$; and for all $t, w_t - \pi_t$ and $w_t^{\lambda} - \pi_t$ are on the same side of w_t^0 . Altogether, this implies that the induction hypothesis holds for T + 1 when $\mathcal{R}_{2,T} > 0$. A similar proof holds in the case where $\mathcal{R}_{2,T} \leq 0$. Hence, for all $T \geq 1$, $||\mathcal{R}_T||^2 \leq (\nu \overline{d})^2 T$.

Inequality $||\mathcal{R}_T|| \leq \nu \overline{d} \sqrt{T}$ implies (32) and the right-hand side of (33). The left-hand side of (33) follows from a proof identical to that of the left-hand side of (19).

Let us denote by $r_{\lambda,\pi} = \frac{1}{N} \mathbb{E}_{\lambda,\pi} \left(\sum_{t=1}^{N} u(w_t^{\lambda} - \pi_t) - u(w_t^0) \right)$ the average expected utility gain when the agent is given the calibrated contract of parameter ν . Similarly, denote by $r_{\max} = \frac{1}{N} \mathbb{E}_{\bar{c},a^*} \left(\sum_{t=1}^{N} u(w_t) - u(w_t^0) \right)$ the maximum feasible utility gain, i.e. the utility gain when the agent invest \bar{c} in every period, chooses the allocation a^* that maximizes the principal's expected utility, and does not get paid.

Theorem A.1. There exists m > 0 such that for all N and all \mathcal{P} ,

$$r_{\lambda,\pi} \ge \frac{1}{\nu} \mathbb{E}_{\lambda,\pi} \left(\frac{\Pi_N}{N} \right) - \frac{m}{\sqrt{N}} \quad and \quad r_{\lambda,\pi} \ge \frac{1}{1 + \nu\kappa} r_{\max} - \frac{c}{\nu} - \frac{m}{\sqrt{N}}.$$
 (35)

Proof. The first inequality follows from (33). Indeed, we have that

$$\begin{split} S_N + \nu \overline{d}\sqrt{N} \geq \Pi_N &\iff \sum_{t=1}^N \nu [u(w_t^\lambda - \pi_t) - u(w_t^0)] \geq \Pi_N - \nu \overline{d}\sqrt{N} \\ &\iff \frac{1}{N} \sum_{t=1}^N u(w_t^\lambda - \pi_t) - u(w_t^0) \geq \frac{1}{\nu N} \Pi_N - \frac{\overline{d}}{\sqrt{N}} \end{split}$$

The first part of (35) follows directly by taking expectations.

Let us turn to the second inequality. Combining (32) and (33) it follows that

$$\Pi_{N} \geq \Sigma_{N} - \nu \overline{d}(1 + \sqrt{N})$$

$$\geq \sum_{t=1}^{N} \nu [u(w_{t} - \pi_{t}) - u(w_{t}^{0})] - \nu \overline{d}(1 + \sqrt{N})$$

$$\geq \sum_{t=1}^{N} \nu [u(w_{t}) - u(w_{t}^{0}) - \kappa \pi_{t}] - \nu \overline{d}(1 + \sqrt{N})$$

$$\Rightarrow \Pi_{N} \geq \frac{\nu}{1 + \nu \kappa} \sum_{t=1}^{N} [u(w_{t}) - u(w_{t}^{0})] - \frac{\nu \overline{d}}{1 + \nu \kappa} (1 + \sqrt{N}).$$

The agent's optimal policy (\tilde{c}, \tilde{a}) under calibrated contract (λ, π) must provide the agent with greater utility than (\bar{c}, a^*) . Hence

$$\mathbb{E}_{\tilde{c},\tilde{a}}\left(\Pi_N - \sum_{t=1}^N \tilde{c}_t\right) \geq \mathbb{E}_{\bar{c},a^*}\left(\Pi_N - N\bar{c}\right) \\
\Rightarrow \frac{1}{N} \mathbb{E}_{\tilde{c},\tilde{a}} \Pi_N \geq \frac{\nu}{1 + \nu\kappa} r_{\max} - \bar{c} - \frac{\nu \overline{d}}{1 + \nu\kappa} \frac{1 + \sqrt{N}}{N}.$$

This last inequality and the first part of (35) implies the second part of (35).

A.4 Varying Wealth

The calibrated contracts described in Section 4 performs equally well if the invested wealth in each period varies within some set $[0, \overline{w}]$. Let w_t^i denote the initial invested wealth in period *i*. Given a contract (λ, π) , quantities Σ_T, S_T and Π_T are defined as

$$\Pi_{T} = \sum_{t=1}^{T} \pi_{t} ; \quad \Sigma_{T} = \sum_{t=1}^{T} w_{t}^{i} \left\langle a_{t} - a_{t}^{0}, \mathbf{r}_{t} \right\rangle ; \quad S_{T} = \sum_{t=1}^{T} \lambda_{t} w_{t}^{i} \left\langle a_{t} - a_{t}^{0}, \mathbf{r}_{t} \right\rangle.$$

Similarly, let $\Sigma_{T\setminus T'} = \Sigma_T - \Sigma_{T'-1}$, $\Pi_{T\setminus T'} = \Pi_T - \Pi_{T'-1}$, $S_{T\setminus T'} = S_T - S_{T'-1}$. As in Section 4, regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ are defined by

$$\mathcal{R}_{1,T} = \Pi_T - \alpha S_T$$
 and $\mathcal{R}_{2,T} = \max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}.$

Contract (λ, π) is unchanged:

$$\lambda_{T+1} = \frac{\alpha \mathcal{R}_{2,T}}{\alpha \mathcal{R}_{2,T} + \mathcal{R}_{1,T}} \quad \text{and} \quad \pi_{T+1} = \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if} \quad \mathcal{R}_{1,T} \le 0\\ 0 & \text{otherwise} \end{cases}$$

Under this adjusted contract, Theorem 2 extends as is, with an identical proof.

A.5 Perturbed Preferences

A more challenging extension is to allow for the agent's marginal utility for money to vary over time. In particular, assume that the agent has preferences $\sum_{t=1}^{N} \mu_t \pi_t$, where $\mu = (\mu_t)_{t\geq 1}$ is bounded below by $\underline{\mu} > 0$. Contract (λ, π) is the same as in Section 4. Hence, it does not correct for varying marginal utility, which is consistent with the idea that $(\mu_t)_{t\geq 1}$ is an unobserved nuisance parameter. Let $v \equiv \sum_{t=1}^{N} |\mu_t - \mu_{t+1}|$ denote the total variation of sequence $(\mu_t)_{t\geq 1}$ and \hat{v} the total variation of sequence $(\mu_t^{-1})_{t\geq 1}$. This section provides an adequate extension of Lemma 1 involving the agent's perturbed utility. For all $T \geq 1$, let $\Pi_N^{\mu} = \sum_{t=1}^{N} \mu_t \pi_t$, $\Sigma_T^{\mu} = \sum_{t=1}^{T} \mu_t (w_t - w_t^0)$ and $S_T^{\mu} = \sum_{t=1}^{T} \mu_t \lambda_t (w_t - w_t^0)$. The following result holds. Lemma A.5 (perturbed incentives).

$$\Sigma_N^{\mu} - S_N^{\mu} \leq w \overline{d} (\mu_N + \upsilon) \sqrt{N}$$
(36)

$$-\alpha w \overline{d} (1 + \upsilon \sqrt{N}) \le \Pi_N^{\mu} - \alpha S_N^{\mu} \le \alpha w \overline{d} (\mu_N + \upsilon) \sqrt{N}.$$
(37)

Proof. Let us begin with inequality (36). Using summation by part and Lemma 1, we have that

$$S_{N}^{\mu} = \sum_{t=1}^{T} \mu_{t} \lambda_{t} (w_{t} - w_{t}^{0}) = \mu_{N} S_{N} + \sum_{t=1}^{N-1} (\mu_{t} - \mu_{t+1}) S_{t}$$

$$\geq \mu_{N} \Sigma_{N} + \sum_{t=1}^{N} (\mu_{t} - \mu_{t+1}) \Sigma_{t} - w \overline{d} \mu_{N} \sqrt{N} \left(\mu_{N} + \sum_{t=1}^{N-1} |\mu_{t} - \mu_{t+1}| \right).$$

$$\geq \Sigma_{N}^{\mu} - w \overline{d} (\mu_{N} + v) \sqrt{N}.$$

Let us turn to (37). Again, Lemma 1 and summation by part implies that

$$\sum_{t=1}^{N} \mu_t \pi_t = \mu_N \Pi_N + \sum_{t=1}^{N-1} (\mu_t - \mu_{t+1}) \Pi_t$$

$$\leq \alpha \mu_N S_N + \alpha \sum_{t=1}^{N-1} (\mu_t - \mu_{t+1}) S_t + \alpha w \overline{d} \mu_N \sqrt{N} \left(\mu_N + \sum_{t=1}^{N-1} |\mu_t - \mu_{t+1}| \right).$$

$$\leq \alpha S_N^{\mu} + \alpha w \overline{d} (\mu_N + v) \sqrt{N}.$$

Finally, a similar argument implies that

$$\sum_{t=1}^{N} \mu_t \lambda_t \pi_t \ge \alpha S_N^{\mu} - \alpha w \overline{d} (1 + \upsilon \sqrt{N}).$$

This concludes the proof.

Lemma A.5 implies that under the original calibrated contracts of Section 4, even though the agent's payoff are perturbed, they still approximate the aggregate payoffs the agent

would have received under the benchmark linear contract. Lemma A.5 can be used to derive efficiency bounds.

Theorem A.2. There exists m > 0 such that for all N and all \mathcal{P} ,

$$r_{\lambda,\pi} \ge \left(\frac{1-\alpha}{\alpha\overline{\mu}}\right) \mathbb{E}_{\lambda,\pi}\left(\frac{\Pi_N^{\mu}}{N}\right) - \frac{m\upsilon}{\sqrt{N}} \quad and \quad r_{\lambda,\pi} \ge \frac{(1-\alpha)\mu}{\overline{\mu}} r_{\max} - \frac{\overline{c}}{\alpha\overline{\mu}} - \frac{m\upsilon}{\sqrt{N}}.$$
 (38)

Proof. Let us begin with the first part of (38). Let (\tilde{c}, \tilde{a}) denote the agent's optimal policy under the calibrated contract (λ, π) . Let a^* denote the allocation policy that maximizes expected final wealth $(w_t)_{t\geq 1}$ given information. By (37), and by definition of a^* we have that

$$\mathbb{E}_{\tilde{c},\tilde{a}}\Pi_N^{\mu} - \alpha w \overline{d}(\mu_N + \upsilon) \sqrt{N} \leq \alpha \mathbb{E}_{\tilde{c},\tilde{a}} S_N^{\mu} \leq \alpha \mathbb{E}_{\tilde{c},a^*} S_N^{\mu}.$$

By definition of \tilde{a} and (37), we also have that

$$\mathbb{E}_{\tilde{c},\tilde{a}}\Pi_{N}^{\mu} \geq \mathbb{E}_{\tilde{c},a^{*}}\Pi_{N}^{\mu} \geq \alpha \mathbb{E}_{\tilde{c},a^{*}}S_{N}^{\mu} - \alpha w \overline{d}(\mu_{N} + v)(1 + \sqrt{N})$$

hence
$$\mathbb{E}_{\tilde{c},a^{*}}S_{N}^{\mu} - \mathbb{E}_{\tilde{c},\tilde{a}}S_{N}^{\mu} \leq \alpha w \overline{d}(\mu_{N} + v)(1 + 2\sqrt{N}).$$

By definition of a^* , for all t, $\mathbb{E}_{\tilde{c},a^*}\lambda_t(w_t - w_t^0) \ge \mathbb{E}_{\tilde{c},\tilde{a}}\lambda_t(w_t - w_t^0)$. Therefore,

$$\begin{aligned} \mathbb{E}_{\tilde{c},a^*}S_N - \mathbb{E}_{\tilde{c},\tilde{a}}S_N &= \sum_{t=1}^N \mathbb{E}_{\tilde{c},a^*}\lambda_t(w_t - w_t^0) - \mathbb{E}_{\tilde{c},\tilde{a}}\lambda_t(w_t - w_t^0) \\ &\leq \frac{1}{\underline{\mu}}\sum_{t=1}^N \mu_t \left[\mathbb{E}_{\tilde{c},a^*}\lambda_t(w_t - w_t^0) - \mathbb{E}_{\tilde{c},\tilde{a}}\lambda_t(w_t - w_t^0) \right] \leq \frac{1}{\underline{\mu}} \left[\mathbb{E}_{\tilde{c},a^*}S_N^{\mu} - \mathbb{E}_{\tilde{c},\tilde{a}}S_N^{\mu} \right] \\ &\leq \frac{1}{\underline{\mu}}\alpha w \overline{d}(\mu_N + v)(1 + 2\sqrt{N}). \end{aligned}$$

This implies that

$$\mathbb{E}_{\tilde{c},\tilde{a}}S_N + \frac{1}{\underline{\mu}}\alpha w \overline{d}(\mu_N + \upsilon)(1 + 2\sqrt{N}) \geq \mathbb{E}_{\tilde{c},a^*}S_N \geq \frac{1}{\underline{\mu} + \upsilon} \mathbb{E}_{\tilde{c},a^*}S_N^{\mu} \\
\geq \frac{1}{\alpha(\underline{\mu} + \upsilon)} \left[\mathbb{E}_{\tilde{c},\tilde{a}}\Pi_N^{\mu} - \alpha w \overline{d}(\mu_N + \upsilon)\sqrt{N} \right].$$

Since (19) still holds, this yields the first part of (38). Let us turn to the second part of (38). By definition of (\tilde{c}, \tilde{a}) , it must be that

$$\mathbb{E}_{\tilde{c},\tilde{a}}\left(\Pi_{N}^{\mu}-\sum_{t=1}^{N}\tilde{c}_{t}\right) \geq \mathbb{E}_{\bar{c},a^{*}}\left(\Pi_{N}^{\mu}-N\bar{c}\right) \geq \alpha\mathbb{E}_{\bar{c},a^{*}}\Sigma_{N}^{\mu}-N\bar{c}-\alpha w\bar{d}(\mu_{N}+\upsilon)(1+2\sqrt{N}) \\ \geq \alpha\underline{\mu}\mathbb{E}_{\bar{c},a^{*}}\Sigma_{N}-N\bar{c}-\alpha w\bar{d}(\mu_{N}+\upsilon)(1+2\sqrt{N}).$$

This and the first part of (38) yields the second part of (38).

Note that this provides a connection with the results of Foster and Young (2010) whose counter-example is based on an agent whose marginal utility for income becomes arbitrarily large in early periods as the agent's time horizon increases.

B Proofs

B.1 Proofs for Section 3

Proof of Theorem 1: Points (i) and (ii) follow immediately from the fact that

$$\pi_t = \alpha (w_t - w_t^0) = \frac{\alpha}{1 - \alpha} (w_t - w_t^0 - \pi_t)$$

Let us turn to point (*iii*). Let $\hat{r}_{\alpha} \equiv \frac{1}{1-\alpha}r_{\alpha}$ denote gross returns under the benchmark contract of parameter α . Let (c, a) denote the agent's policy under the benchmark contract. Since policy (\bar{c}, a^*) guarantees the agent a per-period payoff of $\alpha w r_{\text{max}} - \bar{c}$, it must be that

 $\alpha w \hat{r}_{\alpha} - \mathbb{E}c \geq \alpha w r_{\max} - \bar{c}$. Since the agent must expend weakly positive effort, this implies that $\alpha w \hat{r}_{\alpha} \geq \alpha w r_{\max} - \bar{c}$, which yields point *(iii)*.

Point (iv) follows immediately from the static nature of the benchmark contract.

B.2 Proofs for Section 4

Proof of Lemma 2: Under any benchmark linear contract, the agent is truthful – i.e. uses allocation policy a^* . Let (c, a^*) denote the agent's policy under the benchmark contract of parameter α , (\tilde{c}, \tilde{a}) his policy under contract (λ, π) , and (c_0, a^*) the agent's policy in the benchmark contract of parameter α_0 .

By optimality of (\tilde{c}, \tilde{a}) under contract (λ, π) , we have that $\mathbb{E}_{\tilde{c}, \tilde{a}} \left[\Pi_N - \sum_{t=1}^N \tilde{c}_t \right] \geq \mathbb{E}_{c, a^*} \left[\Pi_N - \sum_{t=1}^N c_t \right]$. We obtain that

$$\mathbb{E}_{\tilde{c},\tilde{a}}\left[\alpha S_N - \sum_{t=1}^N \tilde{c}_t\right] + C \ge \mathbb{E}_{c,a^*}\left[\alpha S_N - \sum_{t=1}^N c_t\right] - B \ge \mathbb{E}_{c,a^*}\left[\alpha \Sigma_N - \sum_{t=1}^N c_t\right] - B - \alpha A.$$
(39)

By optimality of (c, a^*) under the benchmark contract of parameter α , we have that

$$\mathbb{E}_{c,a^*}\left[\alpha\Sigma_N - \sum_{t=1}^N c_t\right] \ge \mathbb{E}_{c_0,a^*}\left[\alpha\Sigma_N - \sum_{t=1}^N c_{0,t}\right].$$
(40)

By optimality of (c_0, a^*) under the benchmark contract of parameter α_0 we obtain

$$\mathbb{E}_{c_0,a^*}\left[\alpha_0\Sigma_N - \sum_{t=1}^N c_{0,t}\right] \ge \mathbb{E}_{\tilde{c},a^*}\left[\alpha_0\Sigma_N - \sum_{t=1}^N \tilde{c}_t\right].$$

Note that by definition of a^* and S_N , $\mathbb{E}_{\tilde{c},a^*}\Sigma_N \geq \mathbb{E}_{\tilde{c},\tilde{a}}S_N$. Indeed, under a^* , Σ_T delivers positive expected returns every period, while S_T (under any allocation policy) provides at best a fraction of these returns. This implies that

$$\mathbb{E}_{c_0,a^*}\left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t}\right] \ge \mathbb{E}_{\tilde{c},\tilde{a}}\left[\alpha_0 S_N - \sum_{t=1}^N \tilde{c}_t\right].$$
(41)

Combining (39), (40) and (41) yields

$$\mathbb{E}_{\tilde{c},\tilde{a}}\left[\alpha S_N - \sum_{t=1}^N \tilde{c}_t\right] + \alpha A + B + C \geq \mathbb{E}_{c_0,a^*}\left[\alpha \Sigma_N - \sum_{t=1}^N c_{0,t}\right]$$
$$\geq (\alpha - \alpha_0)\mathbb{E}_{c_0,a^*}\Sigma_N + \mathbb{E}_{c_0,a^*}\left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t}\right]$$
$$\geq (\alpha - \alpha_0)\mathbb{E}_{c_0,a^*}\Sigma_N + \mathbb{E}_{\tilde{c},\tilde{a}}\left[\alpha_0 S_N - \sum_{t=1}^N \tilde{c}_t\right].$$

Altogether, this implies that $(\alpha - \alpha_0) [\mathbb{E}_{c_0,a^*} \Sigma_N - \mathbb{E}_{\tilde{c},\tilde{a}} S_N] \leq \alpha w \overline{d} (2\sqrt{N} + 1)$. Hence we obtain that

$$\mathbb{E}_{\tilde{c},\tilde{a}}[S_N - \Pi_N] \ge (1 - \alpha) \mathbb{E}_{c_0, a^*} \Sigma_N - (1 - \alpha) \frac{\alpha A + B + C}{\alpha - \alpha_0} - C.$$

Dividing by Nw, this yields that

$$r_{\lambda,\pi} \ge (1-\eta)r_{\alpha_0} + \frac{1}{Nw} \left[C + \frac{1-\eta}{\eta} (\alpha A + B + C) \right].$$

Proof of Theorem 3: Let $w_t^{\Delta\Delta}$ and $\Sigma_N^{\Delta\Delta} = \sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0$ denote potential realized wealth and aggregate excess returns when accidents are lucky. The notation of Section 4 extends, adding superscripts Δ and $\Delta\Delta$ to denote relevant values under the original accidental allocation a^{Δ} , and under the lucky accidental allocation $a^{\Delta\Delta}$. The key step is to provide an adequate extension of Lemma 1. Inequality (19) still applies, and we necessarily have that

$$-\alpha w \overline{d} \le \Pi_N^{\Delta} - \alpha S_N^{\Delta} \le \alpha w \overline{d} \sqrt{N}.$$
(42)

In turn let us show that for any investment strategy of the agent,

$$\Sigma_N^{\Delta\Delta} - 4w\overline{d}\sqrt{N} \le S_N^{\Delta} \tag{43}$$

i.e. up to an order \sqrt{N} , given any investment strategy, the actual excess returns generated under the responsive calibrated contract are at least as high as the returns generated when accidents are lucky. We have that $\Sigma_N^{\Delta\Delta} = \Sigma_{N\setminus T_2+1}^{\Delta} + \Sigma_{T_2\setminus T_1}^{\Delta,+} + \Sigma_{T_1-1}^{\Delta}$. Because inequality (18) still holds, this implies that

$$\Sigma_N^{\Delta\Delta} \le \begin{cases} S_N^{\Delta} + w \overline{d} \sqrt{N} & \text{if } \Sigma_{T_2 \setminus T_1}^{\Delta} > 0 \\ S_{N \setminus T_2 + 1}^{\Delta} + S_{T_1 - 1}^{\Delta} + 3w \overline{d} \sqrt{N} & \text{otherwise} \end{cases}$$

By (42), it follows that

$$\Pi_{T_2}^{\Delta} - \alpha w \overline{d} \sqrt{T_2} \le \alpha S_{T_2}^{\Delta} \le \Pi_{T_2}^{\Delta} + \alpha w \overline{d}$$
$$\Pi_{T_1-1}^{\Delta} - \alpha w \overline{d} \sqrt{T_1 - 1} \le \alpha S_{T_1-1}^{\Delta} \le \Pi_{T_1-1}^{\Delta} + \alpha w \overline{d}.$$

Subtracting these two inequalities yields that,

$$\Pi_{T_2 \setminus T_1}^{\Delta} - \alpha w \overline{d} (1 + \sqrt{T_2}) \le \alpha (S_{T_2}^{\Delta} - S_{T_1-1}^{\Delta}) = \alpha (S_{T_2 \setminus T_1}^{\Delta}).$$

By construction, $\Pi_{T_2 \setminus T_1}^{\wedge} \geq 0$, which implies that for any realization of returns,

$$\begin{split} \Sigma_N^{\Delta\Delta} &\leq S_{N\setminus T_2+1}^{\Delta} + S_{T_2\setminus T_1}^{\Delta} + S_{T_1-1}^{\Delta} + 4w\overline{d}\sqrt{N} \\ &\leq S_N^{\Delta} + 4w\overline{d}\sqrt{N}. \end{split}$$

This implies (43). Given (42) and (43) Theorem 3 follows from a reasoning identical to that of Lemma 2. $\hfill \Box$

B.3 Proofs for Section 5

The proof of Lemma 3 (as well as that of Theorem 5) requires the following extension of the Azuma-Hoeffding inequality.

Lemma B.1 (an extension of Azuma-Hoeffding). Consider a martingale with increments Δ_t such that $|\Delta_t| \leq \overline{\gamma}$. Let $\gamma_t \equiv \sup |\Delta_t| |\mathcal{F}_t$ and $T_m \equiv \inf \left\{ T | \overline{\gamma}^2 + \sum_{t=1}^T \gamma_t^2 \geq m \right\}$. The following hold.

(i)
$$\forall \kappa > 0, \ Prob\left(\sum_{t=1}^{T_m} \Delta_t \ge \kappa\right) \le \exp\left(-2\frac{\kappa^2}{m}\right)$$

(ii) $\forall \kappa > 0, \ Prob\left(\max_{T \le T_m} \sum_{t=1}^T \Delta_t \ge \kappa\right) \le 2\exp\left(-2\frac{\kappa^2}{m}\right)$.

Proof of Lemma B.1: Let us begin with point (i). By Hoeffding's Lemma, (see Hoeffding (1963) or Cesa-Bianchi and Lugosi (2006), Lemma 2.2) we have that

$$\mathbb{E}(\exp(-\lambda\Delta_t)|\mathcal{F}_t) \le \exp\left(\frac{\lambda^2\gamma_t^2}{8}\right).$$

By construction $\sum_{t=1}^{T_m} \gamma_t^2 \leq m$. Hence, using Chernoff's method, we have that for any $\lambda > 0$

$$\operatorname{Prob}\left(\sum_{t=1}^{\tau_m} \Delta_t \ge \kappa\right) \le \exp(-\lambda\kappa) \mathbb{E}\left(\prod_{t=1}^{\tau_m} \exp(\lambda\Delta_t)\right)$$
$$\le \exp(-\lambda\kappa) \mathbb{E}\left(\exp(\lambda\Delta_1) \mathbb{E}\left(\exp(\lambda\Delta_2) \cdots \mathbb{E}\left(\exp(\lambda\Delta_{\tau_m} |\mathcal{F}_{\tau_m}) | \cdots |\mathcal{F}_2\right)\right)$$
$$\le \exp(-\lambda\kappa) \mathbb{E}\left(\exp\left(\frac{\lambda^2}{8} \sum_{t=1}^{\tau_m} \gamma_t^2\right)\right) \le \exp(-\lambda\kappa) \exp\left(\frac{\lambda^2}{8} m\right).$$

Minimizing over λ (i.e. setting $\lambda = 4\kappa/m$) yields point (i).

Point (*ii*) follows from point (*i*) by adapting the standard reflection techniques used for Brownian motions. Let $B_T = \sum_{t=1}^T \Delta_t$. Pick $\kappa > 0$. We want to evaluate $\operatorname{Prob}(\max_{T \leq T_m} B_T \geq$ κ). Consider the process $\widetilde{B}_T = \sum_{t=1}^T \epsilon_t \Delta_t$, where $\epsilon_t = \mathbf{1}_{\max_{s < t} B_s < \kappa} - \mathbf{1}_{\max_{s < t} B_s \geq \kappa}$. Process \widetilde{B}_T is a martingale, corresponding to reflecting B_T once it crosses level κ . Note also that $|\epsilon_t \Delta_t| = |\Delta_t|$. We have that

$$\operatorname{Prob}\left(\max_{T \leq T_m} B_T \geq \kappa\right) = \operatorname{Prob}(B_{T_m} \geq \kappa) + \operatorname{Prob}(B_{T_m} < \kappa \text{ and } \max_{T \leq T_m} B_T \geq \kappa)$$
$$\leq \operatorname{Prob}(B_{T_m} \geq \kappa) + \operatorname{Prob}(\widetilde{B}_{T_m} \geq \kappa). \tag{44}$$

Note that (44) is an inequality, rather than an equality as in the case of a Brownian motion, because of the discreteness of martingale increments. Still this suffices for our purpose. Indeed, by applying point (i) to both B_{T_m} and \tilde{B}_{T_m} , we obtain that indeed, $\operatorname{Prob}\left(\max_{T\leq T_m}\sum_{t=1}^T \Delta_t \geq \kappa\right) \leq 2\exp\left(-2\frac{\kappa^2}{m}\right)$. This concludes the proof. \Box

Proof of Lemma 3: We have that

$$S_T = \sum_{t=1}^T \lambda_t \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0] + \sum_{t=1}^T \lambda_t (w_t - w_t^0 - \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0])$$

Since the agent only has access to public information, by definition of w_t^0 , we have that for all allocation strategies a, $\mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0] \leq 0$. In addition, $\Delta_t \equiv \lambda_t (w_t - w_t^0 - \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0])$ is a martingale increment such that $|\Delta_t| \leq 2\lambda_t d_t$.

Let us define $\chi_T = \overline{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2$. For all $m \in \mathbb{N}$, let T_m denote the stopping time inf $\{T | \chi_T \ge m\}$. Using Lemma B.1, we obtain that for all m

$$\operatorname{Prob}\left(S_{T_m} \ge 2\sqrt{\chi_{T_m}}\sqrt{M+\ln\chi_{T_m}}\right) \le \operatorname{Prob}\left(\sum_{t=1}^{T_m} \Delta_t \ge 2\sqrt{\chi_{T_m}}\sqrt{M+\ln\chi_{T_m}}\right)$$
$$\le \exp\left(-2(\ln m+M)\right) \le \exp(-2M)\frac{1}{m^2}.$$

Given that $S_{T_m} \leq 2\sqrt{\chi_{T_m}}\sqrt{M + \ln \chi_{T_m}}$, the probability that there exists $T \in [T_m, T_{m+1} - 1]$

such that $S_T \ge \Theta_T$ is less than

$$\operatorname{Prob}\left(\sup_{T\in\{T_m,\cdots,T_{m+1}-1\}}\sum_{t=T_m}^T \Delta_t \ge 2\sqrt{M+\ln m}\right) \le 2\exp(-2M)\frac{1}{m^2}.$$

Hence it follows that

$$\mathbb{E}_a\left(\sum_{t=1}^N \mathbf{1}_{S_t > \Theta_t}\right) \leq 3\exp(-2M)\sum_{m \in \mathbb{N}} \frac{1}{m^2} \leq \frac{\pi^2}{2}\exp(-2M).$$

This concludes the proof.

Proof of Lemma 4: A proof identical to that of Lemma 1 yields the left-hand side of (21) and the right-hand side of (22). The left-hand side of (22) is proven by induction. If $\alpha S_T - \alpha \Theta_T \leq 0$, then the inequality holds trivially. Consider now the case where $\alpha S_T - \alpha \Theta_T > 0$. If $\Pi_T \geq \alpha S_T - \alpha \Theta_T$ then we necessarily have $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha \Theta_{T+1} - \alpha w \overline{d}$ since Θ_T is increasing in T. If instead, $\Pi_T \in [\alpha S_T - \Theta_T - \alpha w \overline{d}, \alpha S_T - \Theta_T]$, then necessarily, $\Pi_T < \alpha S_T$, so that $\lambda_T = 1$ and $\pi_T = \alpha (w_T - w_T^0)^+$. It follows that $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha \Theta_{T+1} - \alpha w \overline{d}$. \Box

Proof of Theorem 5: Consider the case where the agent is informed. As in the case of Theorem 2, 3, and 4, the proof strategy is to adapt the the bounds of Lemma 1 and the reasoning of Lemma 2. Let (c, a^*) denote the agent's optimal strategy under the benchmark contract of parameter α . To exploit the reasoning of Lemma 2 it is sufficient to prove a bound of the form

$$-B \le \mathbb{E}_{c,a^*} \left[\Pi_N - \alpha S_N \right],\tag{45}$$

where B is a number independent of N and \mathcal{P} . By construction, we have that

$$\mathbb{E}_{c,a^*} \Pi_N \geq \mathbb{E}_{c,a^*} \alpha S_N - \alpha w \overline{d} - \alpha w \overline{d} \mathbb{E}_{c,a^*} \left[\sum_{T=1}^N \mathbf{1}_{S_T < \Theta_T} \right].$$

Hence, it is sufficient to show that under (c, a^*) , the expected number of periods where the hurdle is not met is bounded above by a constant independent of N.

Let $\Delta_t = w_t - w_t^0 - \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t]$ and $\chi_T = \sum_{t=1}^T d_t^2$. Note that under allocation strategy a^* , By (21), for any T,

$$\begin{aligned} \operatorname{Prob}_{c,a^*}(S_T < \Theta_T) &\leq \operatorname{Prob}_{c,a^*}\left(\Sigma_T < \Theta_T + w\sqrt{\chi_T}\right) \\ &\leq \operatorname{Prob}_{c,a^*}\left(\sum_{t=1}^T \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t] + \sum_{t=1}^T \Delta_t < \Theta_T + w\sqrt{\chi_T}\right) \\ &\leq \operatorname{Prob}_{c,a^*}\left(\frac{\xi}{\overline{d}^2}\chi_T + \sum_{t=1}^T \Delta_t < \Theta_T + w\sqrt{\chi_T}\right) \\ &\leq \operatorname{Prob}_{c,a^*}\left(\sum_{t=1}^T \Delta_t < -\frac{\xi}{\overline{d}^2}\chi_T + \Theta_T + w\sqrt{\chi_T}\right).\end{aligned}$$

An argument similar to that of Lemma 3 yields that $\sum_{T=1}^{+\infty} \operatorname{Prob}\left(\sum_{t=1}^{T} \Delta_t < -\frac{\xi}{\overline{d}^2}\chi_T + \Theta_T + w\sqrt{\chi_T}\right)$ is bounded above by some constant. This concludes the proof.

References

- AL-NAJJAR, N. AND J. WEINSTEIN (2008): "Comparative Testing of Experts," *Econometrica*, 76, 541–559.
- BERGEMANN, D. AND S. MORRIS (2005): "Robust mechanism design," *Econometrica*, 73, 1771–1813.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J. ROCHET (2007): "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies*, 74, 345–390.

BILLINGSLEY, P. (1995): Probability and Measure, John Wiley & Sons, New York.

- BLACKWELL, D. (1956): "An Analog of the Minimax Theorem for Vector Payoffs," Pacific Journal of Mathematics, 6, 1–8.
- CASELLA, A. (2005): "Storable Votes," Games and Economic Behavior, 51, 391–419.
- CESA-BIANCHI, N. AND G. LUGOSI (2006): *Prediction, Learning, and Games*, Cambridge University Press.
- CHUNG, K. AND J. ELY (2003): "Implementation with Near-Complete Information," *Econometrica*, 71, 857–871.
- CHUNG, K. S. AND J. ELY (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74, 447–476.
- CLARKE, E. (1971): "Multipart Pricing of Public Goods," Public Choice, 11, 17–33.
- DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *The Review of Economic Studies*, 46, 185–216.
- DEMARZO, P. AND M. FISHMAN (2007): "Optimal Long-Term Financial Contracting," *Review of Financial Studies*, 20, 2079.
- DEMARZO, P., I. KREMER, AND Y. MANSOUR (2009): "Hannan and Blackwell meet Black and Scholes: Approachability and Robust Option Pricing," *mimeo*.
- DEMARZO, P. AND Y. SANNIKOV (2006): "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 61, 2681–2724.
- ECHENIQUE, F. AND E. SHMAYA (2007): "You Won't Harm Me if You Fool Me," Working Papers.
- EDMANS, A. AND X. GABAIX (2009): "Tractability in Incentive Contracting," *NBER* Working Paper.

- EDMANS, A., X. GABAIX, T. SADZIK, AND Y. SANNIKOV (2009): "Dynamic Incentive Accounts," *NBER Working Paper*.
- ELIAZ, K. (2002): "Fault Tolerant Implementation," *Review of Economic Studies*, 69, 589–610.
- FEINBERG, Y. AND C. STEWART (2008): "Testing Multiple Forecasters," *Econometrica*, 76, 561–582.
- FOSTER, D. AND R. VOHRA (1998): "Asymptotic Calibration," Biometrika, 85, 379.
- (1999): "Regret in the Online Decision Problem," *Games and Economic Behavior*, 29, 7–35.
- FOSTER, D. AND P. YOUNG (2010): "Gaming Performance Fees By Portfolio Managers," The Quarterly Journal of Economics.
- FUDENBERG, D. AND D. LEVINE (1999): "An Easier Way to Calibrate," Games and economic behavior, 29, 131–137.
- GOETZMANN, W., J. INGERSOLL, AND S. ROSS (2003): "High-Water Marks and Hedge Fund Management Contracts," *The Journal of Finance*, 58, 1685–1718.
- GOETZMANN, W., J. INGERSOLL, M. SPIEGEL, AND I. WELCH (2007): "Portfolio Performance Manipulation and Manipulation-Proof Performance Measures," *Review of Financial Studies*.
- GROVES, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617–631.
- HAGERTY, K. AND W. ROGERSON (1987): "Robust Trading Mechanisms," Journal of Economic Theory, 42, 94–107.
- HANNAN, J. (1957): "Approximation to Bayes Risk in Repeated Play," Contributions to the Theory of Games, 3, 97–139.

- HARTLINE, J. D. AND T. ROUGHGARDEN (2008): "Optimal Mechanism Design and Money Burning," in *STOC*, 75–84.
- HE, Z. AND W. XIONG (2010): "Delegated Asset Management and Investment Mandates," mimeo.
- HELLWIG, M. AND K. SCHMIDT (2002): "Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision," *Econometrica*, 70, 2225–2264.
- HOEFFDING, W. (1963): "Probability Inequalities for Sums of Bounded Random Variables," Journal of the American Statistical Association, 58, 13–30.
- HOLMSTROM, B. AND P. MILGROM (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303–328.
- JACKSON, M. AND H. SONNENSCHEIN (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75, 241–257.
- LEHRER, E. (2001): "Any Inspection is Manipulable," *Econometrica*, 69, 1333–1347.
- Lo, A. (2001): "Risk Management for Hedge Funds: Introduction and Overview," Financial Analysts Journal, 57, 16–33.
- MADARÁSZ, K. AND A. PRAT (2010): "Screening with an Approximate Type Space," working paper.
- MCAFEE, P. (1992): "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," *Journal of Economic Theory*, 56, 266–293.
- OLSZEWSKI, W. AND M. PESKI (forthcoming): "The Principal-Agent Approach to Testing Experts," *American Economic Journals: Microeconomics*.

- OLSZEWSKI, W. AND A. SANDRONI (2008): "Manipulability of Future-Independent Tests," *Econometrica*, 76, 1437–1466.
- PANAGEAS, S. AND M. WESTERFIELD (2009): "High-Water Marks: High Risk Appetites? Convex Compensation, Long Horizons, and Portfolio Choice," *The Journal of Finance*, 64, 1–36.
- RADNER, R. (1981): "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica*, 49, 1127–1148.
- (1985): "Repeated Principal-Agent Games with Discounting," *Econometrica*, 53, 1173–1198.
- ROGERSON, W. (1985): "Repeated Moral Hazard," Econometrica, 53, 69–76.
- SANNIKOV, Y. (2008): "A Continuous-Time Version of the Principal–Agent Problem," *Review of Economic Studies*, 75, 957–984.
- SPEAR, S. AND S. SRIVASTAVA (1987): "On Repeated Moral Hazard with Discounting," The Review of Economic Studies, 54, 599–617.
- SUNG, J. (1995): "Linearity with Project Selection and Controllable Diffusion Rate in Continuous-Time Principal-Agent Problems," RAND Journal of Economics, 26, 720–743.
- TOWNSEND, R. (1982): "Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information," *The Journal of Political Economy*, 90, 1166–1186.
- VICKREY, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," Journal of finance, 16, 8–37.
- WILSON, R. B. (1987): "Game-Theoretic Analyses of Trading Processes," in Advances in Economic Theory: Fifth World Congress, ed. by T. Bewley, Cambridge UK: Cambridge University Press, chap. 2, 33–70.