

Calibrated Incentive Contracts

Sylvain Chassang*[†]

Princeton University

April 4, 2013

Abstract

This paper studies a dynamic agency problem which includes limited liability, moral hazard and adverse selection. The paper develops a robust approach to dynamic contracting based on calibrating the incentive properties of simple benchmark contracts that are attractive but infeasible, due to limited liability constraints. The resulting dynamic contracts are detail-free and satisfy robust performance bounds independently of the underlying process for returns, which need not be i.i.d. or even ergodic.

1 Introduction

This paper considers a dynamic agency problem in which a principal hires an agent to make investment decisions on her behalf.¹ The contracting environment includes limited liability, moral hazard, adverse selection, and makes few assumptions about the underlying process for returns and information. The paper develops a robust approach to dynamic contracting

*Contact: chassang@princeton.edu.

[†]I'm indebted to Abhijit Banerjee, Roland Benabou, Faruk Gul, Stephen Morris, Wolfgang Pesendorfer, David Sraer and Satoru Takahashi for time, encouragement, and advice. The paper greatly benefited from conversations with Manuel Amador, Bruno Biais, Markus Brunnermeier, Mike Golosov, Johannes Hörner, Augustin Landier, Thomas Mariotti, Ulrich Müller, Guillaume Plantin, Esteban Rossi-Hansberg, Adam Szeidl, Juuso Toikka, Wei Xiong, Rakesh Vohra, Muhamet Yildiz and John Zhu. I'm grateful to seminar audiences at the 2012 AEA meetings, Berkeley, the Institute for Advanced Study, Harvard, MIT, Microsoft New England, Northwestern, NYU, NYU Stern, Paris School of Economics, Penn State, Princeton, Toulouse School of Economics, Sciences Po Paris, UCL, Warwick, Wisconsin, and X-HEC for comments and suggestions. SeHyoun Ahn and Juan Ortner provided excellent research assistance.

¹Throughout the paper, the principal is referred to as she, while the agent is referred to as he.

whose main steps are as follows: 1) identify a simple class of *high-liability* benchmark contracts that satisfy attractive and robust efficiency properties; 2) construct *limited-liability* dynamic contracts whose terms are dynamically calibrated to ensure that key properties of benchmark contracts are approximately satisfied. The resulting dynamic contracts—referred to as calibrated contracts—perform approximately as well as benchmark contracts 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 principal is infinitely lived, while the agent has a large but finite horizon which need not be known to the principal. In every period a fixed amount of resources can be invested on behalf of the principal by the agent. The agent has private information about the process for returns and can exert costly effort to obtain additional information (e.g. collecting or purchasing data, running experiments. . .). The main constraint on contracts is limited liability: the agent cannot receive negative transfers, and rewards are bounded above by per-period resources, which rules out large deferred payments. The paper makes few assumptions on the underlying probability space and the agent may have arbitrary private information at the time of contracting. Furthermore, the process for information and returns need not be i.i.d. or even ergodic: it may be that with non-vanishing probability there is a large number of periods where returns happen to be negative, or where costly information turns out to be useless.

This is a difficult environment to contract in. The principal is facing both adverse selection (the agent may have persistent private information about returns, or about the cost-effectiveness of information acquisition) and moral hazard (the agent expends effort to acquire information and makes asset allocation decisions). 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 emphasizing performance bounds that hold over large classes of priors.

The first step of the approach relaxes limited liability constraints and identifies a suitable class of high-liability benchmark contracts. Benchmark contracts are simple linear contract which reward the agent a share of his externality on the principal. Such contracts exhibit high-liability since the agent is expected to provide compensation for the losses he causes. While this class of contracts is not necessarily optimal, it satisfies three important efficiency properties: *(i)* for any underlying environment, the principal is guaranteed positive expected surplus from the relationship; *(ii)* benchmark contracts are weakly renegotiation proof; *(iii)* benchmark contracts satisfy a robust efficiency bound, which is sufficiently tight to imply the max min optimality of linear contracts over appropriately chosen families of environments.

The second step of the approach is to develop a simple class of limited-liability dynamic contracts that robustly approximate the efficiency properties of linear high-liability contracts. The key insight is to calibrate both the rewards to the agent and the share of resources he is investing so that for all possible strategies of the agent and all realizations of uncertainty, the rewards obtained by the agent and his externality on the principal remain as tightly linked as they are under benchmark linear contracts. As the time horizon becomes large, these calibrated contracts induce performance approximately equal to that achieved by benchmark linear contracts. This result holds for continuation payoffs from the perspective of any history, which alleviates renegotiation concerns. Finally, for any positive level of liability available to the agent, calibrated contracts can be adjusted to ensure that fully uninformed agents do not participate.

The paper hopes to usefully complement the rich literature on optimal dynamic contracting (see for instance Rogerson (1985), Green (1987), Holmström and Milgrom (1987), Spear and Srivastava (1987), Laffont and Tirole (1988), and more recently Battaglini (2005), DeMarzo and Sannikov (2006), Biais et al. (2007, 2010), DeMarzo and Fishman (2007), Sannikov (2008), Edmans et al. (2012) or Zhu (forthcoming)). Because optimal contracts depend finely on the details of the underlying environment, this literature has delivered rich positive predictions on how contract form should vary with the circumstances. However, a limitation of the optimal contracting approach is that it provides little guidance on how

well those contracts perform if the environment is misspecified. The current paper gives up on optimality and develops a class of dynamic limited-liability contracts that satisfy attractive efficiency properties for a very broad class of stochastic environments. Notably, the performance bounds satisfied by these robust contracts hold in environments where solving for optimal contracts has proved particularly difficult. This includes non-stationary environments (as in Battaglini (2005), Tchisty (2006), He (2009), Pavan et al. (2010) or Garrett and Pavan (2010)), and settings with both moral hazard and adverse selection (such as Sannikov (2007) or Fong (2008)). Still, the contracts developed in the current paper are substantially connected to the optimal contracts derived by DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007) or Biais et al. (2007, 2010) in specific settings. The similarities as well as the differences are instructive and will be discussed in detail.

The paper also contributes to the literature on dynamic contracting and mechanism design with patient players. It is most closely related to the work of Rubinstein (1979), Rubinstein and Yaari (1983) and Radner (1981, 1985) which proves the existence of approximately first-best contracts in a dynamic moral hazard problem where the agent's production function is ergodic and common knowledge under first-best behavior. More recently Jackson and Sonnenschein (2007) and Escobar and Toikka (2009) propose simple quota mechanisms that approximately implement any Pareto efficient allocation rule in a class of dynamic multi-agent allocation problems where the agents have ergodic preferences. As in these previous approaches, the main idea of the current paper is to constrain payoffs to satisfy key properties that would hold under an ideal benchmark. The central difference is that previous work relies strongly on the assumption that the state of the world follows an ergodic process: the basic idea is to make sure that the empirical distribution of realized outcomes matches the expected distribution of outcomes under first-best behavior. This approach is not applicable in the current paper since the underlying environment need not be ergodic and no law of large numbers need apply.

The methods used in this 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)). However, the main question here is not whether good tests are available. Rather, this paper takes a principal-agent approach related to that of Echenique and Shmaya (2007), Olszewski and Peski (2011) or Gradwohl and Salant (2011). These papers show that in such environments there are satisfactory ways to identify experts that generate positive surplus. Olszewski and Peski (2011) rely on ex post high-liability contracts to incentivize truth telling. Gradwohl and Salant (2011) show it is possible to rely on upfront payments instead. Neither paper tackles incentive provision when information acquisition is costly.

The paper is structured as follows. Section 2 describes the framework. Section 3 introduces a class of benchmark linear contracts that satisfy attractive efficiency properties but require high liability. Section 4 is the core of the paper: it develops dynamic limited-liability contracts whose parameters are calibrated to ensure that key incentives properties of benchmark linear contracts are approximately satisfied. Section 5 shows how to adjust calibrated contracts to induce self-screening by uninformed agents. Section 6 concludes by discussing the paper’s assumptions and approach in further detail, and relating calibrated contracts to other contracts of interest. Proofs are given in Appendix A, unless mentioned otherwise. An Online Appendix provides extensions tackling issues such as time discounting, or how to calibrate a broader class of benchmark contracts, as well as simulations illustrating the main properties of calibrated contracts.

2 The Framework

Players, Actions and Payoffs. A principal hires an agent 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.²

In each period $t \in \{1, \dots, N\}$, the principal invests an amount of resources w at the

²Online Appendix OA 1 shows how to extend the analysis when future payoffs are discounted.

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, which rules out private savings. Both the principal and the agent are risk neutral. 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 $\sum_{k=1}^K a_{k,t} = 1$. Set A is convex and compact. It represents constraints on possible allocations. These constraints 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 dot product. Given an asset allocation a_t and returns \mathbf{r}_t , wealth at the end of period t is

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

By assumption, returns are bounded below by -1 so that $w_t \geq 0$ (there cannot be negative resources at the end of the period).

For any pair of allocations $(a, a') \in A^2$, the distance between a and a' is defined by

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

The following assumption puts joint restrictions on the set of permissible allocations A and returns R . It is maintained throughout the paper.

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

This assumption limits the magnitude of changes that can occur in a single period.

At the beginning of every period t , the agent can expend cost $c_t \in [0, +\infty)$ towards acquiring information. This cost can be the actual cost of purchasing or collecting data, an effort cost, or the opportunity cost of time when performing due diligence. The agent then makes an asset allocation suggestion $a_t \in A$ and receives a payment π_t depending on the realized public history at the end of period t . The agent's objective is to maximize his

expected average payoffs

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

Information. Information acquired at time $t \in \{1, \dots, N\}$ is represented as a random variable I_t from a measurable state space (Ω, σ) to a measurable signal space $(\mathcal{I}, \sigma_{\mathcal{I}})$. Publicly available information is denoted by I_t^0 . It includes (but need not be limited to) realized past returns $(\mathbf{r}_s)_{s < t}$, and corresponds to the information available to the principal. In each period, the agent can choose to acquire additional signals $I_t(c)$ at cost $c \in [0, +\infty)$ from a set of possible signals $\{I_t(c) | c \in [0, +\infty)\}$ indexed by their cost. By assumption I_t^0 is measurable with respect to $I_t(c)$ for any $c \geq 0$, so that the agent is more informed than the principal, regardless of the information he acquires. At time $t = 0$, before contracting occurs, the agent also observes an initial private and exogenous signal I_H , summarizing past private history. As a result the agent may know much more than the principal about the process for returns at the time of contracting. In addition, the agent's information and the agent's information acquisition strategy are private. As a result the framework exhibits both adverse selection and moral hazard. Given an information acquisition strategy $(c_t)_{t \geq 1}$, let $(\mathcal{F}_t)_{t \geq 1}$ be the agent's filtration (generated by $(I_H, I_t(c_t))_{t \geq 1}$), and let $(\mathcal{F}_t^0)_{t \geq 1}$ denote the public information filtration (generated by $(I_t^0)_{t \geq 1}$).

For simplicity, the paper assumes that at a sufficiently ex ante stage, the principal and the agent have a common prior P over state space (Ω, σ) . The corresponding probability space $\mathcal{P} = (\Omega, \sigma, P)$ will be referred to as the environment. This common prior assumption ensures that expected returns under the agent's beliefs would be evaluated in the same way by the principal if she had access to the same information. Said in other terms, the principal is willing to accept the agent's beliefs as a basis for the valuation of her expected returns. Note that this common prior assumption imposes little restrictions on beliefs at the time of contracting, since the agent can observe an arbitrary private initial signal I_H prior to contracting.

The only restriction placed on environment \mathcal{P} is Assumption 1. 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 non-vanishing probability that returns are below their period $t = 1$ expectation for an arbitrarily large number of periods; the value of the information that the agent can collect may be uncertain or vary in arbitrary ways; the agent may have to learn the informativeness of different signals through costly experimentation (i.e. by incurring the cost of observing signals and assessing their predictive power over returns); once valuable sources of information can become obsolete over time.

Strategies. Altogether, an agent's strategy consists of an information acquisition strategy $c = (c_t)_{t \geq 1}$, and an asset allocation strategy $a = (a_t)_{t \geq 1}$, where both c_t and a_t are adapted to the information available to the agent 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 :

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

Allocation a_t^0 is the allocation the principal could pick on her own, given public information \mathcal{F}_t^0 , while allocation a_t^* is the allocation that maximizes expected returns given the agent's private information \mathcal{F}_t . 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.

Contracts. Contracts $(\pi_t)_{t \geq 1}$ 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. The principal has commitment power but transfers are subject to the following constraints: in every period t ,

$$0 \leq \pi_t \leq w_t. \quad (4)$$

The constraint that $0 \leq \pi_t$ corresponds to a limited-liability constraint on the agent's side: the agent does not have access to side resources that can be pledged in the contract.

The constraint that $\pi_t \leq w_t$ corresponds to limited liability on the side of the principal: the principal does not have deep pockets and transfers in each period are limited by the resources available in each period. These constraints are at the origin of the contracting problem: 1) the agent does not share on the downside, and 2) large deferred payments are not feasible.

Clearly, these are demanding liability constraints. They can be thought of as a design objective which ensures that the contracts being constructed apply in broad classes of environments. While a narrow interpretation of these constraints is that they correspond to physical limitations (e.g. this is an autarkic economy with perishable resources), Section 6 suggests richer interpretations and identifies economic settings in which similar, although weaker, constraints arise endogenously.

3 A High-Liability Benchmark

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 information that the agent has or may acquire is private. At this level of generality, informative characterizations of optimal dynamic contracts are unlikely.

The paper embraces an alternative approach to dynamic contracting which aims to identify contracts satisfying robust efficiency properties over broad classes of environments. The first step of the analysis defines a class of benchmark contracts that have attractive efficiency properties, but violate limited liability condition (4). The second step of the analysis constructs a class of dynamic contracts that satisfy condition (4), and achieve performance approximately as good as that of the benchmark contracts, regardless of the underlying environment \mathcal{P} .

3.1 Benchmark contracts

The contracts used as benchmark are linear contracts in which the agent’s reward π_t in period t is a share α of the externality his decisions have on the principal:

$$\forall t \geq 1, \quad \pi_t = \alpha(w_t - w_t^0).^3 \tag{5}$$

These benchmark contracts are attractive for the following reasons:

- (i) they satisfy a demanding “no-loss” condition ensuring that the principal gets positive expected surplus whenever the agent gets positive expected rewards; they are the only class of contracts satisfying “no-loss” for all environments;
- (ii) they are weakly renegotiation proof in the sense of Bernheim and Ray (1989) and Farrell and Maskin (1989);
- (iii) they satisfy a robust performance bound which is sufficiently tight to imply that benchmark contracts are max min optimal over appropriately chosen families of environments.

Note that even though both parties are risk-neutral, the fact that the agent has significant private information means that fixed-price contracts in which all productive assets are sold to the agent need not be optimal.⁴

3.2 No-loss and renegotiation proofness

Benchmark contracts satisfy the following no-loss property, and they are the only class of contracts to do so.

³Recall that w_t and w_t^0 respectively denote realized wealth under the agent’s suggested asset allocation and under the default, public information, asset allocation. If $\alpha = 20\%$ 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.

⁴Consider the following stylized example: in each period t the agent can generate expected return $\nu_t = \mathbb{E}(w_t - w_t^0 | \mathcal{F}_t)$ at a cost $c(\nu_t) = \alpha_0 \nu_t$ for $\nu_t \in [0, \bar{\nu}_t]$, where $\bar{\nu}_t$ is an upper bound to feasible expected returns that follows a stochastic process privately observed by the agent. A linear contract with reward rate α_0 is an optimal contract since it induces efficient effort and extracts all the surplus. In contrast, any fixed price contract can cause inefficiencies since the agent will choose not to participate whenever process $(\bar{\nu}_t)_{t \geq 1}$ takes sufficiently low values in expectation.

Fact 1 (no loss). *Under the benchmark contract, for all environment \mathcal{P} and all strategy profiles (c, a) , $\mathbb{E}_{c,a} \left[\sum_{t=1}^N \pi_t \right] \geq 0 \iff \mathbb{E}_{c,a} \left[\sum_{t=1}^N w_t - w_t^0 - \pi_t \right] \geq 0$.*

The converse holds. If a contract $(\pi_t)_{t \geq 1}$ is such that for all \mathcal{P} and all strategy profiles (c, a) , $\mathbb{E}_{c,a} \left[\sum_{t=1}^N \pi_t \right] \geq 0 \iff \mathbb{E}_{c,a} \left[\sum_{t=1}^N w_t - w_t^0 - \pi_t \right] \geq 0$, then there exists $\alpha \in (0, 1)$ such that for all t , $\pi_t = \alpha(w_t - w_t^0)$.

In words, benchmark linear contracts are such that for any strategy (c, a) under which the agent obtains positive expected profit – even suboptimal ones – the principal must also obtain positive expected surplus.⁵

Admittedly, the requirement that “no loss” hold for all possible strategies of the agent (even suboptimal ones) is demanding, but this requirement can also be attractive in environments where the agent may not be fully optimizing (e.g. if he is satisficing): provided that a suboptimal strategy profile does not generate negative profit for the agent, it can only benefit the principal. The bulk of the paper assumes that agents are rational and fully optimizing, but the Online Appendix returns to the question of contract performance when agents can be temporarily irrational.

Another property of benchmark contracts is that since they are independent of history, the principal and the agent are never tempted to renegotiate to a continuation contract starting from a different history.

Fact 2. *Benchmark contracts are weakly renegotiation proof in the sense of Bernheim and Ray (1989) and Farrell and Maskin (1989).*

3.3 Performance bounds

Recent work by Rogerson (2003), Chu and Sappington (2007) and Bose et al. (2011) has emphasized that simple contracts can often guarantee large shares of the second-best surplus in contexts ranging from procurement to principal-agent problems. These approaches have

⁵As a comparison, note that this does not hold for log-scoring rules although they successfully elicit truthful information: log-scoring rules reward the agent for any information regardless of whether or not it is useful for investment purposes. The Online Appendix returns to this point.

focused on parametric models for which it is possible to compute the second-best explicitly. Theorem 1, stated below, contributes to this literature by providing a non-parametric bound for the surplus generated by linear contracts. This bound is sufficiently tight to imply that linear contracts are max min optimal over appropriately chosen families of environments.

Additional notation is needed. Given a contract $\pi = (\pi_t)_{t \geq 1}$, the agent solves optimization problem

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

The corresponding per-period excess returns r_π accruing to the principal (net of payments to the agent) are

$$r_\pi \equiv \inf \left\{ \mathbb{E}_{c,a} \left(\frac{1}{Nw} \sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \middle| (c,a) \text{ solves (P1)} \right\}.^6$$

Returns accruing to the principal when the contract is $\pi_t = \alpha(w_t - w_t^0)$ are denoted by r_α . 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 global incentive compatibility constraints, e.g. in problem (P1), the most tempting deviation of an agent exerting high effort may be to acquire no information at all.

For any $\hat{c} \in [0, +\infty)$, let $r_{\max}(\hat{c})$ denote the production function for returns, i.e. expected per-period returns generated when the agent: 1) incurs an expected per-period cost of information acquisition equal to \hat{c} ; 2) chooses optimal asset allocation a^* given information; and 3) requires no rewards.⁷ Formally we have

$$r_{\max}(\hat{c}) \equiv \sup_{\substack{c \text{ s.t.} \\ \mathbb{E}[\frac{1}{N} \sum_{t=1}^N c_t] \leq \hat{c}}} \mathbb{E}_{c,a^*} \left(\frac{1}{N} \sum_{t=1}^N \langle a_t^* - a_t^0, \mathbf{r}_t \rangle \right).$$

⁶By convention, since the agent may be indifferent between multiple strategy profiles, returns r_π focus on the inf of possible returns. This convention does not matter for the analysis.

⁷It is worth emphasizing that r_{\max} denotes gross returns, while r_α denotes returns net of payments to the agent.

Production frontier $r_{\max}(\cdot)$ should be viewed as a summary statistic of the underlying environment, and is a much simpler object than the full stochastic process for information and returns. The following performance bounds hold.

Theorem 1 (efficiency bounds). *(i) For any probability space \mathcal{P} ,*

$$wr_{\alpha} \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(wr_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha} \right). \quad (6)$$

(ii) For any \hat{c} and $\rho \in (0, 1)$ such that $\frac{\hat{c}}{wr_{\max}(\hat{c})} \leq \rho$, the linear contract of parameter $\alpha = \sqrt{\rho}$ satisfies

$$wr_{\alpha} \geq \left(1 - 2 \frac{\sqrt{\rho}}{1 + \sqrt{\rho}} \right) (wr_{\max}(\hat{c}) - \hat{c}). \quad (7)$$

Given a benchmark contract, point (i) provides a lower bound for the returns to the principal which holds in any environment \mathcal{P} . Given restrictions on aggregate production frontier $r_{\max}(\cdot)$, optimizing the right hand side of (6) provides a rationale for the choice of α . For instance, point (ii) shows that if one imposes the restriction that $\frac{\hat{c}}{wr_{\max}(\hat{c})} \leq \rho$, then the linear contract of parameter $\alpha = \sqrt{\rho}$ guarantees a fixed share of the maximum surplus obtainable at a per-period expected effort cost \hat{c} .

The next lemma builds on (6) to show these bounds are tight: linear contracts guarantee the highest possible share of first-best returns over appropriately chosen classes of environments. Note that first best surplus corresponds to solving $\max_{\hat{c} \in [0, +\infty)} wr_{\max}(\hat{c}) - \hat{c}$. Denote by r_{FB} and c_{FB} the first-best expected per-period returns and first-best expected per-period costs. By definition, it must be that $\frac{c_{FB}}{wr_{FB}} \leq 1$. For any $\rho \leq 1$, denote by

$$\mathbb{P}_{\rho} = \left\{ \mathcal{P} \mid \frac{c_{FB}}{wr_{FB}} \leq \rho \right\}$$

the set of environments such that the ratio of costs to returns at first-best is bounded above by ρ .

Corollary 1. *For any $\rho \in (0, 1)$, the benchmark contract of parameter $\alpha = \sqrt{\rho}$ satisfies*

$$\begin{aligned} \max_{\pi=(\pi_t)_{t \geq 1}} \min_{\mathcal{P} \in \mathbb{P}_\rho} \frac{wr_\pi}{wr_{FB} - c_{FB}} &= \min_{\mathcal{P} \in \mathbb{P}_\rho} \frac{wr_\alpha}{wr_{FB} - c_{FB}} \\ &= 1 - 2 \frac{\sqrt{\rho}}{1 + \sqrt{\rho}}. \end{aligned} \tag{8}$$

In words, the linear contract of parameter $\alpha = \sqrt{\rho}$ is the contract that guarantees the highest possible proportion of first-best surplus over environments $\mathcal{P} \in \mathbb{P}_\rho$. As in Rogerson (2003), Chu and Sappington (2007), Hartline and Roughgarden (2008) or Bose et al. (2011), Corollary 1 measures the performance of benchmark contracts as a ratio to a theoretical upper bound, which gives scale-free estimates of performance.

Corollary 1 can be interpreted as a motivation for the use of linear contracts: they are max min optimal over appropriately chosen classes of environments. Hurwicz and Shapiro (1978) and more recently Carroll (2012) also derive linear contracts as max min optimal in different contexts. However, because Corollary 1 depends on the class \mathbb{P}_ρ of environments over which the max min problem is defined, a more cautious interpretation of Corollary 1 is that it illustrates the tightness of performance bounds (6) and (7).⁸

Fact 1, Fact 2 and Theorem 1 motivate the use of linear contracts as a benchmark. Unfortunately benchmark linear contracts require high-liability from the agent. The next section constructs dynamic contracts that perform approximately as well as benchmark contracts, while also satisfying limited liability constraint (4).

4 Calibrated Contracts

The basic insight underlying calibrated contracts is that by dynamically adjusting contracting variables it is possible to approximate key aggregate incentive properties of the bench-

⁸Indeed the research agenda going forward is not to argue for a specific class of environments on which to perform max min optimization, but rather to understand how different non-parametric restrictions on the environment affect contract performance and contract form. Inequalities (7) and (8) establish tight performance bounds under restrictions on the ratio of costs to returns. Other restrictions may prove fruitful in future research.

mark linear contract. In turn this induces performance approximately as good as that of benchmark contracts: as the time horizon becomes large, calibrated contracts perform approximately as well as benchmark contracts for every underlying environment, and from the perspective of every history.

4.1 Approximating Incentive Properties

Calibrated contracts use the following contracting instruments: 1) in every period t , the agent only invests a share $\lambda_t \in [0, 1]$ of the principal's wealth, with the remaining share $1 - \lambda_t$ is invested in the default asset allocation a_t^0 ; 2) at the end of each period t , the agent receives a payment $\pi_t \in [0, w_t]$.

The contracting strategy used here is to choose process $(\lambda_t, \pi_t)_{t \geq 1}$ to approximate the following properties of linear contracts: for all $T \geq 1$,

$$\text{(correct rewards)} \quad \sum_{t=1}^T \pi_t - \alpha \lambda_t (w_t - w_t^0) = 0 \quad (9)$$

$$\text{(no foregone gains)} \quad \forall T' \leq T, \quad \sum_{t=T'}^T (1 - \lambda_t)(w_t - w_t^0) \leq 0. \quad (10)$$

Condition (9) states that in aggregate the agent is rewarded a share α of his aggregate externality on the principal (where coefficients $(\lambda_t)_{t \geq 1}$ scale the externality up or down). Condition (10) states that over any time interval $[T', T]$ resource allocation policy $(\lambda_t)_{t \geq 1}$ does not reduce potential gains from following the agent's allocation. It is trivially satisfied in benchmark contracts since $\lambda_t = 1$ for all t .⁹ Note that “no foregone gains” condition (10) is equivalent to

$$\forall T, \quad \max_{T' \leq T} \sum_{t=T'}^T (1 - \lambda_t)(w_t - w_t^0) \leq 0.$$

Any contract which satisfies these two properties induces performance at least as good as

⁹Note that under benchmark high-liability contracts, in equilibrium, it must be that for all $t \geq 1$, $\mathbb{E}(w_t - w_t^0) \geq 0$. Hence, parameter λ_t is optimally set to 1. The necessity of controlling λ_t under limited-liability will be discussed at length.

that of benchmark contracts. Manipulation shows that by construction, conditions (9) and (10) must hold starting from any interim period T' .¹⁰ As a result any contract satisfying (9) and (10) induces performance as good as that of linear contracts from the perspective of any history. The remainder of this section shows how to approximate properties (9) and (10) using only limited liability contracts.

Incentives as regrets. Define the following regrets,

$$\mathcal{R}_{1,T} \equiv \sum_{t=1}^T \pi_t - \alpha \lambda_t (w_t - w_t^0) \quad \text{and} \quad \mathcal{R}_{2,T} \equiv \max_{T' \leq T} \sum_{t=T'}^T (1 - \lambda_t) (w_t - w_t^0).$$

In addition, let us define aggregate regrets \mathcal{R}_T and marginal regrets ρ_T as follows,

$$\mathcal{R}_T \equiv \begin{bmatrix} \mathcal{R}_{1,T} \\ \alpha \mathcal{R}_{2,T}^+ \end{bmatrix} \quad \text{and} \quad \rho_T \equiv \begin{bmatrix} \mathcal{R}_{1,T} - \mathcal{R}_{1,T-1} \\ \alpha (\mathcal{R}_{2,T} - \mathcal{R}_{2,T-1}) \end{bmatrix},$$

where $x^+ = \max\{0, x\}$ for any $x \in \mathbb{R}$.¹¹ Regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ respectively correspond to the amount by which conditions (9) (“correct rewards”) and (10) (“no foregone gains”) are violated. As a result, conditions (9) and (10) are formally equivalent to $\mathcal{R}_T = 0$. The goal is now to pick a process $(\pi_t, \lambda_t)_{t \geq 1}$ satisfying limited liability condition (4) and such that aggregate regret \mathcal{R}_T is negligible compared to T as T grows large, i.e. of order $o(T)$. This regret minimization problem can be solved using the methods of Blackwell (1956) and Hannan (1957).¹² For any value \mathcal{R}_T , we want to pick $(\lambda_{T+1}, \pi_{T+1})$ such that

$$\forall w_{T+1}, w_{T+1}^0, \quad \langle \mathcal{R}_T, \rho_{T+1} \rangle \leq 0. \quad (11)$$

¹⁰Indeed, using condition (9) at $T' - 1$ and T yields that $\sum_{t=T'}^T \pi_t - \lambda_t (w_t - w_t^0) = 0$, i.e. condition (9) also holds starting from any period T' .

¹¹The vector of regrets \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 small improvement in performance bounds.

¹²See Foster and Vohra (1999) or Cesa-Bianchi and Lugosi (2006) for a more recent treatment.

Inequality (11), known as an approachability condition, 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.

Let us now find parameter values $(\lambda_{T+1}, \pi_{T+1})$ such that (11) holds. By construction, regret $\mathcal{R}_{2,T+1}$ —which measures maximum foregone gains—satisfies

$$\begin{aligned} \mathcal{R}_{2,T+1} &= \max_{T' \leq T+1} \sum_{t=T'}^{T+1} (1 - \lambda_t)(w_t - w_t^0) \\ &= (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \max_{T' \leq T+1} \sum_{t=T'}^T (1 - \lambda_t)(w_t - w_t^0) \\ &= (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T}^+. \end{aligned}$$

Using the identities $(\mathcal{R}_{2,T}^+ - \mathcal{R}_{2,T})\mathcal{R}_{2,T}^+ = 0$ and $\mathcal{R}_{1,T} = \mathcal{R}_{1,T}^+ + \mathcal{R}_{1,T}\mathbf{1}_{\mathcal{R}_{1,T} < 0}$, it follows that

$$\begin{aligned} \langle \mathcal{R}_T, \rho_{T+1} \rangle &= [\pi_{T+1} - \alpha \lambda_{T+1}(w_{T+1} - w_{T+1}^0)]\mathcal{R}_{1,T} + \alpha^2 [(1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0)] \mathcal{R}_{2,T}^+ \\ &= [\pi_{T+1} - \alpha \lambda_{T+1}(w_{T+1} - w_{T+1}^0)\mathbf{1}_{\mathcal{R}_{1,T} \leq 0}] \mathcal{R}_{1,T} \\ &\quad - \alpha [\lambda_{T+1}\mathcal{R}_{1,T}^+ - \alpha(1 - \lambda_{T+1})\mathcal{R}_{2,T}^+] (w_{T+1} - w_{T+1}^0). \end{aligned}$$

Hence approachability condition $\langle \mathcal{R}_T, \rho_{T+1} \rangle \leq 0$ can be satisfied for any realization of w_{T+1} and w_{T+1}^0 by setting

$$\lambda_{T+1} = \frac{\alpha \mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} \alpha \lambda_{T+1}(w_{T+1} - w_{T+1}^0)^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{if } \mathcal{R}_{1,T} > 0, \end{cases} \quad (12)$$

with the convention that $\frac{0}{0} = 1$.¹³

This defines a dynamic contract satisfying limited liability constraint (4); this is the calibrated contract of interest. The following lemma shows that incentive properties (9) and

¹³Note that because rewards $(\pi_t)_{t \geq 1}$ (and hence regret $\mathcal{R}_{1,T}$) are proportional to α , coefficient α can be simplified out of the expression for $(\lambda_t)_{t \geq 1}$.

Initial terms are set so that $\lambda_1 \in [0, 1]$ and $\pi_1 \in [0, \lambda_1 \alpha (w_1 - w_1^0)^+]$. Any such specification is consistent with the bounds given in Lemma 1 (below). A reasonable choice would be $\lambda_1 = 1$ and $\pi_1 = \alpha (w_1 - w_1^0)^+$.

(10) are indeed approximately satisfied.

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

$$-\alpha w \bar{d} \leq \sum_{t=1}^T \pi_t - \alpha \lambda_t (w_t - w_t^0) \leq \alpha w \bar{d} \sqrt{T} \quad (13)$$

$$\sum_{t=T'}^T (1 - \lambda_t) (w_t - w_t^0) \leq w \bar{d} \sqrt{T}. \quad (14)$$

Lemma 1 implies that incentive properties (9) and (10) hold up to an error term of order \sqrt{T} , which is small compared to the number of periods T . Note that this holds sample path by sample path, rather than in expectation or in equilibrium. The proof is instructive.

Proof. Let $d_t = \sup_{\mathbf{r} \in R} | \langle a_t - a_t^0, \mathbf{r} \rangle |$ denote the magnitude of positions taken by the agent in period t . We first show that $\|\mathcal{R}_T\|^2 \leq \alpha^2 w^2 \sum_{t=1}^T d_t^2$. The proof is by induction. Let us first show the property holds at $T = 1$. Observe that $\mathcal{R}_{1,1} = \pi_1 - \alpha \lambda_1 (w_1 - w_1^0)$ and $\mathcal{R}_{2,1} = (1 - \lambda_1)(w_1 - w_1^0)$. Since $\lambda_1 \in [0, 1]$ and $\pi_1 \in [0, \alpha \lambda_1 (w_1 - w_1^0)^+]$, we obtain that $\|\mathcal{R}_1\|^2 \leq \alpha^2 \lambda_1^2 (w_1 - w_1^0)^2 + \alpha^2 (1 - \lambda_1)^2 (w_1 - w_1^0)^2 \leq \alpha^2 w^2 d_1^2$. Assume now that the induction hypothesis holds at $T \geq 1$. Consider the case where $\mathcal{R}_{2,T} \geq 0$ (i.e. there are some foregone returns). Since approachability condition (11) holds, we have that

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

Since $\pi_{T+1} \in [0, \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+]$, it must be that $\rho_{1,T+1}^2 \leq \alpha^2 \lambda_{T+1}^2 (w_{T+1} - w_{T+1}^0)^2$. Hence, it follows that $\|\rho_{T+1}\|^2 \leq \alpha^2 \lambda_{T+1}^2 (w_{T+1} - w_{T+1}^0)^2 + \alpha^2 (1 - \lambda_{T+1})^2 (w_{T+1} - w_{T+1}^0)^2 \leq \alpha^2 w^2 d_{T+1}^2$. Altogether this shows that the induction hypothesis holds when $\mathcal{R}_{2,T} \geq 0$. A similar proof holds when $\mathcal{R}_{2,T} < 0$, taking into account that in this case $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0)$. Altogether, this implies that the induction hypothesis holds, and for all $T \geq 1$, $\|\mathcal{R}_T^+\|^2 \leq \alpha^2 w^2 \sum_{t=1}^T d_t^2$. This proves (14) and the right-hand side of (13).

The left-hand side of (13) is also proven by induction. Recall that $\mathcal{R}_{1,T} = \sum_{t=1}^T \pi_t -$

$\alpha\lambda_t(w_t - w_t^0)$. If $\mathcal{R}_{1,T} \in [-\alpha w\bar{d}, 0]$, then $\lambda_{T+1} = 1$, and $\pi_{T+1} = \alpha(w_{T+1} - w_{T+1}^0)^+$. Hence by construction, $\mathcal{R}_{1,T+1} \geq -\alpha w\bar{d}$. If instead $\mathcal{R}_{1,T} > 0$, then by definition of \bar{d} , $\mathcal{R}_{1,T+1} \geq -\alpha w\bar{d}$. This implies the left-hand side of (13). \square

Necessity of varying investment scale $(\lambda_t)_{t \geq 1}$. The fact that property (9) holds approximately, i.e. that $\sum_{t=1}^T \pi_t - \alpha\lambda_t(w_t - w_t^0) = o(T)$, ensures that the agent's aggregate reward $\sum_{t=1}^T \pi_t$ is tightly linked to his externality $\sum_{t=1}^T \lambda_t(w_t - w_t^0)$ on the principal. This property cannot be achieved without varying investment scale $(\lambda_t)_{t \geq 1}$. A simple way to see this is by considering a situation in which starting from some period T_0 the agent's *potential* externality $w_t - w_t^0$ on the principal is on average negative for many successive periods.¹⁴ If λ_t is kept constant, say equal to 1, the relationship between rewards to the agent and his externality on the principal will break: because of limited liability, rewards to the agent $\sum_{t=1}^T \pi_t$ are weakly increasing with time T , whereas aggregate performance $\sum_{t=1}^T w_t - w_t^0$ is sharply decreasing. The time-varying investment policy $(\lambda_t)_{t \geq 1}$ described in (12) manages to shield the investor from such large downward deviations in the agent's performance without generating large foregone performance.¹⁵ In this respect, note that policies that fire the agent altogether—i.e. such that $\lambda_t = 0$ for all t after a sufficiently bad history—cannot satisfy “no foregone gains” condition (10). Indeed, since histories at which the agent is fired happen with positive probability on the equilibrium path, firing rules generate losses proportional to unknown time horizon N . Simulations provided in the Online Appendix illustrate the incentive alignment properties of calibrated contracts and the patterns of allocation rule $(\lambda_t)_{t \geq 1}$ used to guarantee them.

¹⁴This may be the result of poor luck, but may also happen for incentive reasons. Imagine that the agent learns that he can no longer acquire additional valuable information starting from period T_0 . Because of limited liability, it will be optimal for him to suggest asset allocations that are different from default allocation a_t^0 . Since allocation a_t^0 is optimal under public information, the agent will be suggesting allocations with negative expected returns.

¹⁵Because incentive properties (9) and (10) are only approximated up to a term of order \sqrt{T} , there are variants of process $(\lambda_t, \pi_t)_{t \geq 1}$ that also approximate (9) and (10) at similar but different speeds. For instance requiring that $\langle \mathcal{R}_T, \rho_{T+1} \rangle \leq m$ with m a fixed constant independent of T also yields bounds of order \sqrt{T} , but with worse coefficients. All such processes must vary resource allocation $(\lambda_t)_{t \geq 1}$.

4.2 Approximate Performance

Lemma 1 shows that as time horizon N grows large, the calibrated contract defined by (12) must approximate incentive properties (9) and (10). Theorem 2 (stated below) shows that approximately satisfying incentive properties (9) and (10) ensures performance approximately as good as that 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 contract (λ, π) :

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

For any history h_T observed by the agent, normalized net returns conditional on h_T are

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

When the contract in question is the benchmark linear contract of parameter α , net returns accruing to the principal continue to be denoted by r_α (similarly, let $r_\alpha | h_T$ denote conditional returns at history h_T). Note that returns $r_{\lambda, \pi}$ and r_{α_0} are computed under the assumption that the agent's behavior is an exact best-reply. No approximate best-reply assumption is made. An extension in the Online Appendix weakens rationality and studies the performance of calibrated contracts when the agent can behave suboptimally over an arbitrary interval of time.

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 (12). There exists a constant m independent of time horizon N and probability space \mathcal{P} such that for all histories h_T ,*

$$r_{\lambda, \pi} | h_T \geq (1 - \eta)r_{\alpha_0} | h_T - m \frac{1}{\sqrt{N}}. \quad (15)$$

In addition,

$$r_{\lambda,\pi} \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right) - \frac{3\bar{d}}{\sqrt{N}}. \quad (16)$$

It follows from (15) that for N large enough, the calibrated contract described by (12) generates a share approximately $1 - \eta$ of the returns the principal obtains under the benchmark contract of parameter α_0 . The result holds from the perspective of any history, which alleviates renegotiation concerns. Inequality (16) provides a performance bound analogous to (6) which holds independently of η .¹⁶ Note that for performance bounds (15) and (16) to hold, it is important that the principal actually allocate resources to the agent according to $(\lambda_t)_{t \geq 1}$. Reward scheme $(\pi_t)_{t \geq 1}$ does not induce perfectly good behavior from the agent. Rather, 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 $(\lambda_t)_{t \geq 1}$.

Theorem 2 does not follow immediately from Lemma 1 since approximation errors with respect to incentives may cause the agent to change his behavior significantly. Indeed, in this general environment global incentive constraints may be binding or almost binding under the benchmark contract. Hence getting incentives slightly wrong may result in large shifts in behavior and poor performance. For instance, this would be the case if under the benchmark contract, the agent were indifferent between working hard and not working at all. For this reason incentives must be reinforced: by sharing an additional fraction η of her returns, the principal ensures that potential changes in the agent's behavior do not compromise performance. Madarász and Prat (2010) make the same point in a screening context. Simulations presented in the Online Appendix highlight that when global incentive constraints are binding, η may need to be significant to guarantee meaningful efficiency bounds for realistic time horizons.

¹⁶Note that the current analysis does not attempt to optimize constant coefficients in error terms. For instance the constant 3 in (16) can be replaced by $2 + 1/\sqrt{N}$.

5 Inducing self-screening by uninformed agents

The framework of Section 2 allows for arbitrary adverse selection. In particular, the agent may be fully uninformed and know ex ante that he has no ability to generate valuable information: $\max_{c,a} \mathbb{E}_{c,a} \left(\sum_{t=1}^N w_t - w_t^0 \middle| I_H \right) = 0$. An issue with the calibrated contract defined by (12) is that by construction, rewards are positive and a sufficiently long-lived uninformed agent can obtain significant expected payoffs from luck and volatility alone.¹⁷ In order to induce entirely uninformed agents to self-screen, i.e. to not participate in the first place, some amount of liability is required. The question is how much? It turns out that for any liability level b available in the first period, the calibrated contract described in (12) can be adjusted to induce any uninformed agent to self-screen at a minimal efficiency cost, regardless of environment \mathcal{P} , or of the agent's time horizon.

Specifically, screening is induced by first imposing an initial participation cost $-b$ on the agent, and then only paying the agent when his performance is above a dynamic hurdle Θ_T which depends on a measure of how different from the default allocation the agent's allocation has been. Given a free parameter $M > 0$, define

$$\Theta_T \equiv 2w \left(1 + \sqrt{\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2} \right) \sqrt{M + \ln \left(\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2 \right)}, \quad (17)$$

where $d_t = \sup_{\mathbf{r}_t \in R} | \langle a_t - a_t^0, \mathbf{r}_t \rangle |$ 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, $d_t \leq \bar{d}$). 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. Denote by $S_T \equiv \sum_{t=1}^T \lambda_t (w_t - w_t^0)$ the surplus generated by the agent under resource allocation policy $(\lambda_t)_{t \geq 1}$. The quantity $\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2$ is a measure of time under which $(S_T)_{T \geq 1}$ will have at most the variation of a standard Brownian motion.

¹⁷Even if the agent has no information and all assets have the same expected returns, systematically picking assets different from the benchmark allocation will allow the agent to obtain rewards of order \sqrt{N} with non-vanishing probability.

Hurdled calibrated contracts are defined by an initial fee $\pi_0^\Theta = -b$ and a sequence $(\lambda_t, \pi_t, \pi_t^\Theta)_{t \geq 1}$. The sequence $(\lambda_t, \pi_t)_{t \geq 1}$ is still defined according to (12), and λ_t is still the share of wealth actually invested by the agent. However, for $T \geq 1$, reward π_T is no longer paid to the agent for sure. Rather, the agent is paid a hurdled reward π_T^Θ such that for $T \geq 1$,

$$\pi_T^\Theta = \begin{cases} \pi_T & \text{if } S_T \geq \Theta_T \\ 0 & \text{otherwise,} \end{cases}$$

i.e. potential reward π_T is paid to the agent if and only if the surplus S_T he has generated is greater than hurdle Θ_T .

An intuitive rationale for the form of hurdle Θ_T is as follows. Imagine the agent is uninformed, so that the process $(S_T)_{T \geq 1}$ is at best a martingale, and imagine that the agent is frequently active, i.e. $\sum_{t=1}^T \lambda_t^2 d_t^2$ is of order T . Then hurdle Θ_T is of order $\sqrt{T \ln T}$, whereas, by the law of the iterated logarithm, with probability 1 as T gets large, $\max_{T' \leq T} S_{T'}$ is of order $\sqrt{T \ln \ln T}$.¹⁸ Because $\frac{\sqrt{T \ln \ln T}}{\sqrt{T \ln T}}$ goes to 0 as T grows large, hurdle Θ_T insures that uninformed agents have very little hope to obtain unjustified returns. Indeed, the following result holds.

Lemma 2 (hurdle effectiveness). *If the agent is uninformed, then for any allocation strategy a , any environment \mathcal{P} and any horizon N ,*

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

where π is the constant 3.1415...

Because hurdles also reduce the payoffs accruing to informed agents, they carry an incentive cost. Still as the next theorem shows, this incentive cost is asymptotically moderate. Denote by r_{λ, π^Θ} the net expected per-period returns generated by the agent under the hurdled calibrated contract. The following result holds.

¹⁸See Billingsley (1995), Theorem 9.5.

Theorem 3 (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 space \mathcal{P} such that for all h_T ,*

$$r_{\lambda, \pi^\Theta} |h_T \geq (1 - \eta)r_{\alpha_0} |h_T - m\sqrt{\frac{\ln N}{N}} \quad (18)$$

Furthermore, whenever $-b + \alpha w \bar{d} \times \frac{\pi^2}{2} \exp(-2M) < 0$, it is strictly optimal for uninformed agents not to participate.¹⁹

The combination of initial fee $-b$ and hurdles $(\Theta_T)_{T \geq 1}$ induces self-screening by uninformed agents. Hurdles $(\Theta_T)_{T \geq 1}$ are large enough that uninformed agents have little hope to be rewarded by luck but small enough that they do 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 $\sqrt{\frac{\ln N}{N}}$. An extension in the Online Appendix shows that when expected returns are grainy, i.e. either zero or bounded away from 0, performance losses are still of order $\frac{1}{\sqrt{N}}$.

6 Discussion

This section revisits limited liability constraint (4), discusses in further detail how calibrated contracts relate to other contracts of interest, and delineates possible avenues for future research.

6.1 Alternative limited liability constraints

Limited liability constraint (4) imposes that for all $t \geq 1$, transfers π_t to the agent satisfy $0 \leq \pi_t \leq w_t$. One implication is that large deferred payments are not feasible. Clearly, greater

¹⁹This result contrasts with work by Lo (2001) and Foster and Young (2010) which emphasizes the difficulty of both rewarding and screening agents. In particular Foster and Young (2010) describe environments in which rewarding and screening is impossible. This occurs because their environment allows for a strong form of private savings such that informed agents value income in early periods much more than in later periods (consumption can be arbitrarily delayed and agents can save on their own at the same rate of returns they generate for the principal). As a result, talented agents are unwilling to pay the monetary cost needed to induce screening.

liability may be available to either the principal or the agent in many realistic settings but for the purpose of this paper, condition (4) remains a useful reduced-form design constraint: it ensures that the contracts being designed apply across a broad variety of environments. Indeed, limited liability constraints weaker than, but similar to (4), arise endogenously in more sophisticated settings.

As an example, consider the contracting problem under which there is no limited liability on the side of the principal, i.e. $\pi_t \in [0, +\infty)$, but the agent has the following stylized concave preferences over consumption $x \in \mathbb{R}$:

$$u(x) = \begin{cases} x & \text{if } x \in [0, \bar{x}] \\ \bar{x} & \text{if } x \geq \bar{x} \\ -\infty & \text{if } x < 0. \end{cases}$$

where $\bar{x} \geq w\bar{d}$ is a constant independent of time horizon N . In words, the agent cannot have a negative consumption, and has only bounded consumption needs within one period (e.g. consumption opportunities are limited by time constraints). In addition, the agent can save at zero interest rate, but cannot borrow against future labor income.

This contracting problem, with no limited-liability constraint on the principal and a risk-averse agent, is in fact equivalent to the contracting problem in which the agent is risk neutral but transfers are constrained to satisfy

$$0 \leq \pi_t \leq \bar{x}. \tag{4'}$$

Fact 3. *If a contract $\pi = (\pi_t)_{t \geq 1}$ satisfies (4') it induces the same performance from a risk neutral agent and from an agent with preferences u .*

Consider a contract $\pi = (\pi_t)_{t \geq 1}$ such that $\pi_t \in [0, +\infty)$ for all t . There exists a contract π' satisfying (4') such that for every underlying environment \mathcal{P} : π' and π induce the same behavior and the same payoffs for the agent—whether the agent is risk-neutral or has preferences u ; the principal obtains weakly greater payoffs under π' .

Note that calibrated contracts satisfy (4') since they satisfy $\pi_t \leq \alpha(w_t - w_t^0)^+ \leq w\bar{d}$. Here

limited liability on the side of the principal arises as a reduced form implication of limited consumption opportunities for the agent.

Similar limited liability constraints on the side of the principals may also arise for more direct reasons: for instance the principal may be able to renege on promised payments at a finite cost. Precluding deferred payments above renegeing costs ensures that the principal is never tempted to interrupt the relationship early, regardless of her beliefs over continuation values.

6.2 Relation to other contracts

High-water mark contracts. The calibrated contracts described in Section 4 are closely related to the high-water mark contracts frequently used in the financial industry. High-water mark contracts are structured as follows: at time T , the investment share λ_T is always 1, and the agent gets paid

$$\pi_T^{\text{hwmk}} = \alpha \left(\left[\sum_{t=1}^T w_t - w_t^0 \right] - \max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right] \right)^+ . \quad (19)$$

Quantity $\max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right]$ is referred to as the high-water mark and represents the maximum historical cumulated returns at time T . The agent only gets paid when he improves on his own historical performance.²⁰ Note that like calibrated contracts, high-water mark contracts are dynamic and satisfy limited liability constraint (4). In fact, high-water mark contracts approximately coincide with a calibrated contract in which for all T , the share λ_T of resources managed by the agent is kept constant and equal to 1.²¹

As was previously argued, because high-water mark contracts keep fixed the investment

²⁰For more on high-water mark contracts, see for instance Goetzmann et al. (2003) who develop an option-pricing approach to high-water mark contracts, or Panageas and Westerfield (2009) who show in a specific context that high-water mark contracts need not lead to excessive risk-taking.

²¹More precisely, consider the contract $(\hat{\lambda}, \hat{\pi})$ defined by $\forall T \geq 1, \hat{\lambda}_T = 1$ and $\hat{\pi}_T = \alpha(w_T - w_T^0)^+ \mathbf{1}_{\hat{\mathcal{R}}_{1,T-1} \leq 0}$ with $\hat{\mathcal{R}}_{1,T-1} \equiv \sum_{t=1}^{T-1} \hat{\pi}_t - \alpha(w_T - w_T^0)$. For any time T , aggregate rewards $\sum_{t=1}^T \pi_t^{\text{hwmk}}$ and $\sum_{t=1}^T \hat{\pi}_t$ to the agent differ by at most $\alpha w d$.

scale $(\lambda_t)_{t \geq 1}$ the relationship between rewards to the agent and his externality on the principal will break down following large drops in performance.²² This has two implications. First, an agent who has lost the ability to generate positive return (e.g. his information has become unreliable) will cause large losses by choosing suboptimal allocations (choosing the default allocation, which is optimal under public information, guarantees him zero rewards). Second, if a talented agent has been unlucky and experienced a drop in returns, the difficulty of catching up with a high-water mark may discourage investment altogether. As a result high-water mark contracts exhibit large gains to renegotiation. If an agent performs well for an extended amount of time, following which he experiences sharp losses, the principal and the agent may both benefit strongly from forgiving the losses and pretending that the current high-water mark is lower than it really is.

By choosing appropriate investment shares $(\lambda_t)_{t \geq 1}$, calibrated contracts are able to keep tight the relationship between rewards and externality starting from every history. As a result, extended drops in performance have a limited impact on payoffs to the principal, and a limited impact on continuation incentives for the agent (see Online Appendix OA 2 for illustrative simulations). The fact that calibrated contracts do not generate large foregone performance (Lemma 1) implies that along parts of the path of play where the agent is generating positive returns, investment shares $(\lambda_t)_{t \geq 1}$ will be close to one. Inversely, investment shares may be significantly below one along portions of the path where the agent is not generating positive returns.²³

²²For instance, imagine the agent delivers performance $(w_t - w_t^0)_{t \geq 1}$ equal to $(1, 1, -1, 1, 1, -1, \dots, 1, 1, -1)$ so that total surplus is $N/3$. Cumulated value-added $\sum_{t=1}^T w_t - w_t^0$ is on average increasing and under the high-water mark contract the agent obtains a reward $\alpha N/3 + o(N)$. If instead the agent delivers returns $(1, 1, \dots, 1)$ for the first $2N/3$ periods, followed by $(-1, -1, \dots, -1)$ for the last $N/3$ periods, then the surplus generated by the agent is still $N/3$, but the high-water mark contract now gives him a payoff $2\alpha N/3 + o(N)$.

²³Note that for Lemma 1 to hold sample path by sample path, investment shares λ_t must move smoothly with performance instead of taking only values 0 or 1. Rather than a stop-loss provision, it is more accurate to think of the calibrated investment shares $(\lambda_t)_{t \geq 1}$ as continuously implementing a robust option on the agent's potential performance $\sum_{t=1}^T w_t - w_t^0$. See DeMarzo et al. (2009) for work on the relation between approachability and robust option pricing.

Connection with optimal contracting. It is instructive to note that DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007) as well as Biais et al. (2007, 2010) derive high-water mark contracts as optimal contracts in their environments. The link is not entirely obvious because their optimal contracts are described in the standard (forward looking) language of continuation values. Because calibrated contracts and high-water mark contracts are detail-free, they can only be described in reference to (backward looking) realized observables. This difference however is superficial and the connection between the two approaches is significant.²⁴ To a first order, DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007) and Biais et al. (2007, 2010) find that in their environment, under optimal contracts, the agent's continuation value follows a random walk, proportional to the agent's performance, and reflected at some upper bound \bar{W} . Whenever the agent's continuation value hits this upper bound, he is paid a fixed proportion of the surplus he generates. This in fact encodes for a high-water mark contract. Imagine that at time t , the agent is promised value \bar{W} , and that he starts losing money. Then, his continuation value moves in a way proportional to his performance, and he is only paid again when his performance covers his losses so that his continuation value climbs back to \bar{W} . This coincides with the reward profile of a high-water mark contract: the agent only gets paid once he has recouped his losses. This connection should not be entirely surprising: DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007) as well as Biais et al. (2007, 2010) consider environments with a linear production technology in which the benchmark high-liability contracts of Section 3 are close to optimal for patient players; calibrated contracts are specifically designed to approximate the performance of such contracts.

The connection is particularly strong with Biais et al. (2007) and especially Biais et al. (2010) who emphasize the role of downsizing the project managed by the agent as a function of his performance. This is related to varying investment shares $(\lambda_t)_{t \geq 1}$ in the current paper. The use of downsizing in Biais et al. (2007, 2010) however is slightly different. In their work,

²⁴Indeed, the optimal contracts derived by DeMarzo and Sannikov (2006), DeMarzo and Fishman (2007) and Biais et al. (2007, 2010) can be given a backward looking description since there is a one-to-one mapping between realized payoffs and continuation values.

downsizing occurs when continuation values are so low that at the current size of the project, optimal behavior can no longer be enforced. Downsizing allows to deliver the promised low values while maintaining appropriate incentive compatibility conditions in the continuation game. As a result, downsizing occurs only after sufficiently long strings of poor performance. In the current paper, scaling rule $(\lambda_t)_{t \geq 1}$ can be seen as a *preventive* downsizing scheme, which rules out continuation values so low that incentive provision becomes problematic.

6.3 Future Work

The relative simplicity of the analysis presented in the paper gives reasonable hope that parts of the approach may be used in other settings. Three directions for further research seem promising.

A first direction for research is to reintroduce prior restrictions on the environment. The current analysis puts few such restrictions and is meant to stand in sharp contrast to the usual optimal contracting approach which fully specifies an underlying environment \mathcal{P} . One challenge going forward is to bridge the two approaches and explore how additional non-parametric restrictions on the environment map into richer contractual designs.²⁵

A second challenge is to allow for risk-aversion. Some suggestions are offered in Chassang (2011), but more work remains to be done. One difficulty is to characterize the amount of co-insurance between principal and agent that can be sustained using prior-free approaches, and specify what constitutes an appropriate performance benchmark.

Finally, a third avenue for research is to extend the incentive calibration approach of this paper to study dynamic mechanism design under limited liability constraints. For instance, Vickrey-Clarke-Groves mechanisms require agents to make significant payments and are therefore ill-suited to environments where agents are severely cash constrained. The incentive calibration approach developed in this paper can help relax such limited liability constraints.

²⁵For instance, further restrictions could be imposed on aggregate production function $r_{\max}(\cdot)$ without specifying the detailed stochastic process giving rise to this aggregate production function.

Appendix

A Proofs

A.1 Proofs for Section 3

Proof of Fact 1: The fact that benchmark contracts satisfy no-loss is immediate: rewards to the agent and payoffs to the principal are proportional under the benchmark contract. Let us turn to the converse.

A contract $(\pi_t)_{t \geq 1}$ induces indirect vNM preferences for the agent and the principal over lotteries with outcomes $(w_t, w_t^0)_{t \geq 1}$. Given such a lottery L , the principal and the agent respectively have expected utilities

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \quad \text{and} \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right).$$

Because no-loss must hold for every underlying environment \mathcal{P} and every strategy of the agent, it implies that for every probability distribution L over outcomes $(w_t, w_t^0)_{t \geq 1}$,

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \geq 0 \quad \iff \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) \geq 0.$$

If $\mathbb{E}_L(\sum_{t=1}^N w_t - w_t^0) = 0$, then $\mathbb{E}_L(\sum_{t=1}^N \pi_t)$ and $-\mathbb{E}_L(\sum_{t=1}^N \pi_t)$ must have the same sign, which implies that

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 \right) = 0 \quad \Rightarrow \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) = 0.$$

Consider the deterministic sequence such that: for all $t > 1$, $w_t = w_t^0 = 0$; $w_1 = 0$; $w_1^0 = 1$. Let us define $\alpha \equiv -\sum_{t=1}^N \pi_t$ for this deterministic sequence of outcomes. Let L_{-1} denote the lottery putting unit mass on this sequence. For any lottery L such that $\mathbb{E}_L(\sum_{t=1}^N w_t - w_t^0) \geq 0$, consider the compound lottery $\widehat{L} = pL_{-1} + (1-p)L$, with $p/(1-p) = \mathbb{E}_L(\sum_{t=1}^N w_t - w_t^0)$.

By construction, $\mathbb{E}_{\hat{L}}(\sum_{t=1}^N w_t - w_t^0) = 0$ so that necessarily,

$$\begin{aligned} \mathbb{E}_{\hat{L}}\left(\sum_{t=1}^N \pi_t\right) = 0 &\iff -p\alpha + (1-p)\mathbb{E}_L\left(\sum_{t=1}^N \pi_t\right) = 0 \\ &\iff \mathbb{E}_L\left(\sum_{t=1}^N \pi_t\right) = \alpha\mathbb{E}_L\left(\sum_{t=1}^N w_t - w_t^0\right). \end{aligned}$$

Since this must hold for all lotteries L , it must be that for all t , $\pi_t = \alpha(w_t - w_t^0)$. Finally it is immediate that in order to satisfy no-loss, it must be that $\alpha \in (0, 1)$. \square

Proof of Theorem 1: Let us begin with point (i). Let (c_α, a^*) denote the agent's policy under the benchmark contract with reward rate α . Recall that a^* denotes the strategy that maximizes expected returns conditional on information. Recall also that $r_{\max}(\hat{c})$ denotes potential gross returns, i.e. the production function for returns, whereas r_α denotes returns net of payments to the agent under the benchmark contract. Pick $\hat{c} \in [0, +\infty)$ and denote by (c, a^*) a policy that maximizes surplus $\mathbb{E}_{c, a^*}\left(\sum_{t=1}^N w_t - w_t^0\right)$ conditional on average expected cost constraint $\frac{1}{N}\mathbb{E}_c\left(\sum_{t=1}^N c_t\right) \leq \hat{c}$. Since policy (c, a^*) guarantees the agent an expected per-period payoff of $\alpha w r_{\max}(\hat{c}) - \hat{c}$, it must be that $\frac{\alpha}{1-\alpha} w r_\alpha - \mathbb{E}c_\alpha \geq \alpha w r_{\max}(\hat{c}) - \hat{c}$. This implies that $\frac{\alpha}{1-\alpha} w r_\alpha \geq \alpha w r_{\max}(\hat{c}) - \hat{c}$, which yields point (i).

Let us now turn to point (ii) and assume that $\frac{\hat{c}}{w r_{\max}(\hat{c})} \leq \rho$. Applying point (i), we get that

$$w r_\alpha \geq (1-\alpha)\left(w r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha}\right) \geq (1-\alpha)\left(1 - \frac{\rho}{\alpha}\right) w r_{\max}(\hat{c}). \quad (20)$$

Setting $\alpha = \sqrt{\rho}$ maximizes the right-hand side of inequality (20) and yields that

$$\begin{aligned} w r_\alpha &\geq w r_{\max}(\hat{c})(1 - \sqrt{\rho})^2 = w r_{\max}(\hat{c})(1 - \rho)\frac{(1 - \sqrt{\rho})^2}{1 - \rho} \\ &\geq (w r_{\max}(\hat{c}) - \hat{c})\frac{1 - \sqrt{\rho}}{1 + \sqrt{\rho}} = \left(1 - 2\frac{\sqrt{\rho}}{1 + \sqrt{\rho}}\right)(w r_{\max}(\hat{c}) - \hat{c}). \end{aligned}$$

This proves point (ii). \square

Proof of Corollary 1: Using the bound given in point (ii) of Theorem 1 for $\hat{c} = c_{FB}$ and

$\alpha = \sqrt{c_{FB}/wr_{FB}}$ yields that

$$\frac{wr_\alpha}{wr_{FB} - c_{FB}} \geq 1 - 2\frac{\sqrt{\rho}}{1 + \sqrt{\rho}}.$$

We now show that this bound is tight: no contract can improve on this bound over the class of environments \mathbb{P}_ρ . For this it is sufficient to show that no contract can improve on this bound for some subclass of environments included in \mathbb{P}_ρ . We consider the following family of settings.

There are two assets, 1 and 2. Asset 1 is riskless with returns $r_{1,t} = 0$ every period. Asset 2 is risky and i.i.d. with negative expected value. Specifically, $r_{2,t} = 1$ with probability $1/3$ and $r_{2,t} = -1$ with probability $2/3$. The agent can only acquire information in period $t = 1$, but that information is valuable over the entire course of the relationship. Expending cost Nc in the first period implies that with probability $p(c)$ the agent learns the entire profile of realizations $(r_{2,t})_{t \geq 1}$. With probability $1 - p(c)$ the agent does not observe any information and there are no more information acquisition opportunities.

Environments \mathcal{P} in this subclass of interest differ by the probability $p(c)$ with which the agent can learn the profile of returns $(r_{2,t})_{t \geq 1}$. This is equivalent to per-period expected returns $r_{\max}(c) = p(c)/3$. Attention is restricted to expected returns that are step functions of the form $r_{\max}(c) = r_{\max}(0) \geq 0$ for $c \in [0, c_{FB})$ and $r_{\max}(c) = r_{FB} > 0$ for $c \in [c_{FB}, +\infty)$. Furthermore we impose the restriction that $\frac{c_{FB}}{wr_{FB}} = \rho$. In this environment, since we are not imposing limited liability condition (4), one can restrict attention to contracts in which aggregate payments $\Pi_N \equiv \sum_{t=1}^N \pi_t$ are decided and transferred in the last period. Define $\pi = \Pi_N/N$ the corresponding per-period reward. In this environment, reward π need only be conditioned on the following events:

- the agent only invests in asset 1 (event 0)
- the agent invests in asset 2 and only obtains returns equal to 1 when he does (event 1)
- the agent invests in asset 2 and obtains returns equal to -1 at one history (event -1).

Indeed, the agent can be discouraged to choose asset 2 when uninformed by setting $\pi(-1)$ arbitrarily low. In equilibrium event -1 will not occur and these off-path incentives have no efficiency cost. Given these incentives, if event 1 occurs, it must be that the agent has invested in information acquisition, and incentives to do so depend only on expected payoffs conditional on events 0 and 1. Hence it is sufficient to condition π only on events 0, 1 and -1.

Define $\Delta = \pi(1) - \pi(0)$ the difference in per-period rewards between events 1 and 0. The agent's per-period expected payoff from putting effort c is

$$\begin{aligned} p(c) \left[(1 - (2/3)^N) \pi(1) + (2/3)^N \pi(0) \right] + (1 - p(c)) \pi(0) - c \\ = \pi(0) + 3r_{\max}(c) \left[1 - (2/3)^N \right] \Delta - c, \end{aligned}$$

while the principal's per-period payoff is

$$-\pi(0) + r_{\max}(c) \left(w - 3 \left[1 - (2/3)^N \right] \Delta \right).$$

Let us first show that any contract such that $\pi(0) \neq 0$ cannot guarantee the principal a positive share of first-best surplus. Indeed, if $\pi(0) < 0$, then for values of r_{FB} low enough, the agent's payoff is strictly negative for all values of $c \in [0, +\infty)$, which implies that the agent doesn't participate in the first place and the principal gets profits equal to 0. If instead $\pi(0) > 0$, then for values of r_{FB} low enough, the principal will get negative profits.

Now consider the case where $\pi(0) = 0$. If $\Delta < \frac{c_{FB}}{3r_{FB}[1-(2/3)^N]}$, then the agent chooses cost level $c = 0$, which leads to zero profits in environments where $r_{\max}(0) = 0$. Assume now that $\Delta \geq \frac{c_{FB}}{3r_{FB}[1-(2/3)^N]}$. For any $\epsilon > 0$, in environments such that $3r(0)\Delta [1 - (2/3)^N] = 3r_{FB}\Delta [1 - (2/3)^N] - c_{FB} + \epsilon$, the agent chooses to expend cost $c = 0$, and the principal obtains payoff

$$\left(r_{FB} - \frac{c_{FB} - \epsilon}{3[1 - (2/3)^N] \Delta} \right) \left(w - 3[1 - (2/3)^N] \Delta \right).$$

Maximizing over Δ and letting ϵ go to 0 (which yields $\Delta = \frac{w}{3[1-(2/3)^N] \sqrt{\rho}}$) implies an upper bound to the principal's guaranteed payoff: the principal can guarantee himself a payoff of at most

$$\begin{aligned} \left(wr_{FB} - \frac{c_{FB}}{\sqrt{\rho}} \right) (1 - \sqrt{\rho}) &= wr_{FB}(1 - \sqrt{\rho})^2 = (wr_{FB} - c_{FB}) \frac{(1 - \sqrt{\rho})^2}{1 - \rho} \\ &= (wr_{FB} - c_{FB}) \left(1 - \frac{2\sqrt{\rho}}{1 + \sqrt{\rho}} \right). \end{aligned}$$

We know from point (i) of Theorem 1 that the linear contract of parameter $\alpha = \sqrt{\rho}$ guarantees this performance, which is therefore the maxmin efficiency ratio over the class of environments \mathbb{P}_ρ . This concludes the proof. \square

A.2 Proofs for Section 4

The proof of Lemma 1, which shows that incentive properties (9) and (10) hold approximately under calibrated contracts, was given in the text.

Let us turn to the proof of Theorem 2. The following lemma—which will be used again—provides the missing step. It shows that approximately satisfying (9) and (10) implies approximate performance bounds.

Lemma A.1. *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 ,*

$$\sum_{t=1}^N (1 - \lambda_t)(w_t - w_t^0) \leq A \quad \text{and} \quad -B \leq \sum_{t=1}^N \pi_t - \alpha \lambda_t (w_t - w_t^0) \leq C. \quad (21)$$

Then

$$r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - \frac{1}{Nw} \left[C + \frac{1 - \eta}{\eta} (\alpha A + B + C) \right] \quad (22)$$

$$r_{\lambda, \pi} \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right) - \frac{1}{Nw} \left[C - \frac{1 - \alpha}{\alpha} (\alpha A + B + C) \right]. \quad (23)$$

Proof of Lemma A.1: Let us first prove (22). Recall that a^* denotes the allocation strategy that maximizes expected returns conditional on acquired information. Under any benchmark linear contract, the agent uses conditionally optimal 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 .

It is convenient to introduce the following notation:

$$\Pi_N = \sum_{t=1}^N \pi_t, \quad \Sigma_N = \sum_{t=1}^N w_t - w_t^0 \quad \text{and} \quad S_N = \sum_{t=1}^N \lambda_t (w_t - w_t^0).$$

Π_N denotes aggregate rewards, Σ_N potential surplus created by the agent, and S_N actual surplus created given cautious investment rule $(\lambda_t)_{t \geq 1}$.

By optimality of (\tilde{c}, \tilde{a}) under contract (λ, π) , we have that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left[\Pi_N - \sum_{t=1}^N \tilde{c}_t \right] \geq \mathbb{E}_{c, a^*} \left[\Pi_N - \sum_{t=1}^N c_t \right].$$

Using (21), we obtain that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha S_N - \sum_{t=1}^N \tilde{c}_t \right] + C \geq \mathbb{E}_{c, a^*} \left[\alpha S_N - \sum_{t=1}^N c_t \right] - B \geq \mathbb{E}_{c, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_t \right] - B - \alpha A. \quad (24)$$

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

$$\mathbb{E}_{c, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_t \right] \geq \mathbb{E}_{c_0, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_{0,t} \right]. \quad (25)$$

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

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

Note that by definition of a^* and S_N , $\mathbb{E}_{\tilde{c}, \tilde{a}} \Sigma_N \geq \mathbb{E}_{\tilde{c}, \tilde{a}} S_N$. Indeed, under a^* , potential returns $w_t - w_t^0$ have positive expectation every period, and $\lambda_t(w_t - w_t^0)$ (under any allocation policy) provides at best a fraction of these returns. This implies that

$$\mathbb{E}_{c_0, a^*} \left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t} \right] \geq \mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha_0 S_N - \sum_{t=1}^N \tilde{c}_t \right]. \quad (26)$$

Combining (24), (25) and (26) yields

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

Regrouping terms appropriately, this implies that $(\alpha - \alpha_0) [\mathbb{E}_{c_0, a^*} \Sigma_N - \mathbb{E}_{\tilde{c}, \tilde{a}} S_N] \leq \alpha A + B + C$. This inequality implies a lower bound for $\mathbb{E}_{\tilde{c}, \tilde{a}} S_N$. Using this lower bound and the fact that by (21), $\Pi_N \leq \alpha S_N + C$, we obtain that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} [S_N - \Pi_N] \geq (1 - \alpha) \mathbb{E}_{c_0, a^*} (\Sigma_N) - (1 - \alpha) \frac{\alpha A + B + C}{\alpha - \alpha_0} - C. \quad (27)$$

Recall that $\alpha = \alpha_0 + \eta(1 - \alpha_0)$ so that $1 - \alpha = (1 - \alpha_0)(1 - \eta)$ and $\alpha - \alpha_0 = \eta(1 - \alpha_0)$. By dividing (27) with Nw we obtain that $r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - \frac{1}{Nw} \left[C + \frac{1 - \eta}{\eta}(\alpha A + B + C) \right]$. This proves inequality (22).

We now turn to performance bound (23). Recall that $r_{\max}(\hat{c})$ denotes potential gross returns, i.e. the production function for returns, whereas $r_{\lambda, \pi}$ denotes returns net of payments to the agent under contract (λ, π) . Continue to denote by (\tilde{c}, \tilde{a}) the policy chosen by the agent under calibrated contract (λ, π) . For any $\hat{c} \in [0, +\infty)$, let (c, a^*) denote a policy that maximizes expected returns $\mathbb{E}_{c, a^*} \left(\sum_{t=1}^N w_t - w_t^0 \right)$ conditional on $\frac{1}{N} \mathbb{E}_c \left(\sum_{t=1}^N c_t \right) \leq \hat{c}$.

By optimality of (\tilde{c}, \tilde{a}) under contract (λ, π) it follows that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left(\sum_{t=1}^N \pi_t - \tilde{c}_t \right) \geq \mathbb{E}_{c, a^*} \left(\sum_{t=1}^N \pi_t - c_t \right).$$

Using (21), we obtain that

$$\begin{aligned} \mathbb{E}_{\tilde{c}, \tilde{a}} \left(\sum_{t=1}^N \alpha \lambda_t (w_t - w_t^0) - \tilde{c}_t \right) + C &\geq \mathbb{E}_{c, a^*} \left(\sum_{t=1}^N \alpha \lambda_t (w_t - w_t^0) - c_t \right) - B \\ &\geq \mathbb{E}_{c, a^*} \left(\sum_{t=1}^N \alpha (w_t - w_t^0) - c_t \right) - B - \alpha A, \end{aligned}$$

where we first used that by (21), $-B + \sum_{t=1}^N \alpha \lambda_t (w_t - w_t^0) \leq \sum_{t=1}^N \pi_t \leq C + \sum_{t=1}^N \alpha \lambda_t (w_t - w_t^0)$, and then used that by (21), $\sum_{t=1}^N (1 - \lambda_t)(w_t - w_t^0) \leq A$. Dividing by α and using the definition of $r_{\max}(\hat{c})$, this implies that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0) \right) \geq N \left(w r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha} \right) - \frac{\alpha A + B + C}{\alpha}.$$

Hence it follows that

$$\begin{aligned} r_{\lambda, \pi} &= \frac{1}{Nw} \mathbb{E}_{\tilde{c}, \tilde{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0) - \pi_t \right) \\ &\geq \frac{1}{Nw} \left[(1 - \alpha) \mathbb{E}_{\tilde{c}, \tilde{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0) \right) - C \right] \\ &\geq (1 - \alpha) \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right) - \frac{1}{Nw} \left[C + \frac{1 - \alpha}{\alpha} (\alpha A + B + C) \right]. \end{aligned}$$

□

Proof of Theorem 2: Theorem 2 follows directly from applying Lemma A.1 to the family of incentive bounds of Lemma 1, noting that incentive bounds continue to hold starting from any history. □

A.3 Proofs for Section 5

The proof of Lemma 2 requires the following extension of the Azuma-Hoeffding inequality.

Lemma A.2 (an extension of Azuma-Hoeffding). *Consider a martingale with increments Δ_t such that $|\Delta_t| \leq \bar{\gamma}$. Filtration $(\mathcal{F}_t)_{t \geq 1}$ corresponds to the information available at the beginning of period t . Let $\gamma_t \equiv \sup |\Delta_t| | \mathcal{F}_t$ and $T_m \equiv \inf \left\{ T \mid \bar{\gamma}^2 + \sum_{t=1}^T \gamma_t^2 \geq m \right\}$. The following hold:*

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

Proof of Lemma A.2: 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

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

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

$$\begin{aligned} \text{Prob} \left(\sum_{t=1}^{T_m} \Delta_t \geq \kappa \right) &\leq \exp(-\lambda \kappa) \mathbb{E} \left(\prod_{t=1}^{T_m} \exp(\lambda \Delta_t) \right) \\ &\leq \exp(-\lambda \kappa) \mathbb{E}(\exp(\lambda \Delta_1)) \mathbb{E}(\exp(\lambda \Delta_2)) \cdots \mathbb{E}(\exp(\lambda \Delta_{T_m}) | \mathcal{F}_{T_m}) | \cdots | \mathcal{F}_2) \\ &\leq \exp(-\lambda \kappa) \mathbb{E} \left(\exp \left(\frac{\lambda^2}{8} \sum_{t=1}^{T_m} \gamma_t^2 \right) \right) \leq \exp(-\lambda \kappa) \exp \left(\frac{\lambda^2}{8} m \right). \end{aligned}$$

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 $\text{Prob}(\max_{T \leq T_m} B_T \geq \kappa)$. Consider the process $\tilde{B}_T = \sum_{t=1}^T \epsilon_t \Delta_t$, where $\epsilon_t = \mathbf{1}_{[\max_{s < t} B_s] < \kappa} - \mathbf{1}_{[\max_{s < t} B_s] \geq \kappa}$. Process

\tilde{B}_T is a martingale, corresponding to reflecting B_T the first time it crosses level κ . Note also that $|\epsilon_t \Delta_t| = |\Delta_t|$. We have that

$$\begin{aligned} \text{Prob} \left(\max_{T \leq T_m} B_T \geq \kappa \right) &= \text{Prob}(B_{T_m} \geq \kappa) + \text{Prob}(B_{T_m} < \kappa \text{ and } \max_{T \leq T_m} B_T \geq \kappa) \\ &\leq \text{Prob}(B_{T_m} \geq \kappa) + \text{Prob}(\tilde{B}_{T_m} \geq \kappa). \end{aligned} \quad (28)$$

Note that (28) 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, $\text{Prob} \left(\max_{T \leq T_m} \sum_{t=1}^T \Delta_t \geq \kappa \right) \leq 2 \exp \left(-2 \frac{\kappa^2}{m} \right)$. This concludes the proof. \square

Proof of Lemma 2: Recall that S_T is defined by $S_T \equiv \sum_{t=1}^T \lambda_t (w_t - w_t^0)$. 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 is uninformed, by definition of w_t^0 , we have that for all allocation strategies a , $\mathbb{E}_a [w_t - w_t^0 | \mathcal{F}_t^0] \leq 0$. Define $\Delta_t \equiv \lambda_t (w_t - w_t^0 - \mathbb{E}_a [w_t - w_t^0 | \mathcal{F}_t^0]) / w$. Δ_t is a martingale increment such that $|\Delta_t| \leq 2\lambda_t d_t$.

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

$$\begin{aligned} \text{Prob} \left(S_{T_m} \geq 2w \sqrt{\chi_{T_m}} \sqrt{M + \ln \chi_{T_m}} \right) &\leq \text{Prob} \left(\sum_{t=1}^{T_m} \Delta_t \geq 2\sqrt{\chi_{T_m}} \sqrt{M + \ln \chi_{T_m}} \right) \\ &\leq \exp(-2(M + \ln m)) \leq \exp(-2M) \frac{1}{m^2}. \end{aligned}$$

In addition, conditional on $S_{T_m} \leq 2w \sqrt{\chi_{T_m}} \sqrt{M + \ln \chi_{T_m}}$, Lemma A.2 implies that the probability that there exists $T \in [T_m, T_{m+1} - 1]$ such that $S_T \geq \Theta_T$ is less than

$$\text{Prob} \left(\sup_{T \in \{T_m, \dots, T_{m+1} - 1\}} \sum_{t=T_m}^T \Delta_t \geq 2\sqrt{M + \ln m} \right) \leq 2 \exp(-2M) \frac{1}{m^2}.$$

Hence it follows that

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

This concludes the proof. \square

Let us now turn to the proof of Theorem 3. Let $\Pi_T^\Theta = \sum_{t=1}^T \pi_t^\Theta$ denote actual rewards, up to time T . The following lemma extends Lemma 1.

Lemma A.3 (approximate incentives). *For all $T, T' < T$, and all paths of play, we have that*

$$-\alpha\Theta_T - \alpha w \bar{d} - b \leq \sum_{t=1}^T \pi_t^\Theta - \alpha \lambda_t (w_t - w_t^0) \leq \alpha w \sqrt{\sum_{t=1}^T d_t^2} \quad (29)$$

$$\sum_{t=T'}^T (1 - \lambda_t) (w_t - w_t^0) \leq w \sqrt{\sum_{t=1}^T d_t^2}. \quad (30)$$

Proof. A proof identical to that of Lemma 1 yields the right-hand sides of (30) and (29).

Let us turn to the left-hand side of (29). We know from (13) that for all $T \geq 1$, $-\alpha w \bar{d} \leq \sum_{t=1}^T \pi_t - \alpha \lambda_t (w_t - w_t^0)$. We now show by induction that for all $T \geq 1$, $\sum_{t=1}^T \pi_t^\Theta \geq \left(\sum_{t=1}^T \pi_t \right) - \Theta_T$. This holds for $T = 1$. Assume it holds at time $T \geq 1$. If $\pi_{T+1} = 0$, then $\pi_{T+1}^\Theta = 0$ and the induction hypothesis holds. If $\sum_{t=1}^{T+1} \pi_t \leq \Theta_{T+1}$ the induction hypothesis also holds since $\pi_t^\Theta \geq 0$. Consider now the case where $\pi_{T+1} > 0$ and $\sum_{t=1}^{T+1} \pi_t > \Theta_{T+1}$. Since $\pi_{T+1} > 0$, it must be that $\alpha S_{T+1} \geq \sum_{t=1}^{T+1} \pi_t$, which implies that $S_{T+1} > \Theta_{T+1}$ and hence $\pi_{T+1}^\Theta = \pi_{T+1}$. This implies that the induction hypothesis holds at time $T + 1$, which implies the left-hand side of (29). \square

Proof of Theorem 3: Combining Lemma A.3 and Lemmas A.1 and 2 yields Theorem 3. \square

Proof of Fact 3: The first result is immediate: whenever contract π_t satisfies condition (4) then $u(\pi_t) = \pi_t$. Hence, contract π_t induces the same behavior and the same surplus whether the agent is risk-neutral or has preferences u .

Inversely, consider a contract $\pi = (\pi_t)_{t \geq 1}$ such that for all t , $\pi_t \geq 0$. Because of borrowing constraints, the agent's consumption profile $(x_t)_{t \geq 1}$ must be such that for all $T \geq 1$,

$\sum_{t=1}^T x_t \leq \sum_{t=1}^T \pi_t$. We now show that given contract $(\pi_t)_{t \geq 1}$ one can solve the agent's optimal saving problem in a prior-free way. This allows to build a contract equivalent to $(\pi_t)_{t \geq 1}$ and satisfying (4') by letting the principal do all savings. Consider the consumption profile $(x_t^*)_{t \geq 1}$ defined recursively by setting

$$\forall T \geq 1, \quad x_T^* = \min \left\{ \bar{x}, \pi_T + \sum_{t=1}^{T-1} \pi_t - x_t^* \right\}.$$

It is easy to verify that for all T , $(x_t^*)_{t \geq 1}$ solves the optimal consumption problem

$$\max_{(x_t)_{t \geq 1}} \sum_{t=1}^T x_t \quad \left| \quad \forall T' \leq T, \quad x_{T'} \in [0, \bar{x}] \quad \text{and} \quad \sum_{t=1}^{T'} x_t \leq \sum_{t=1}^{T'} \pi_t.$$

By construction, offering the agent contract $(\tilde{\pi}_t)_{t \geq 1} \equiv (x_t^*)_{t \geq 1}$ instead of $(\pi_t)_{t \geq 1}$ induces the same behavior from the agent, and yields greater payoffs to the principal since by constraint $\sum_{t=1}^N x_t^* \leq \sum_{t=1}^N \pi_t$. \square

References

- BATTAGLINI, M. (2005): "Long-term contracting with markovian consumers," *The American economic review*, 95, 637–658.
- BERNHEIM, B. D. AND D. RAY (1989): "Collective dynamic consistency in repeated games," *Games and Economic Behavior*, 1, 295–326.
- 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.
- BIAIS, B., T. MARIOTTI, J. ROCHET, AND S. VILLENEUVE (2010): "Large risks, limited liability, and dynamic moral hazard," *Econometrica*, 78, 73–118.
- 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.
- BOSE, A., D. PAL, AND D. SAPPINGTON (2011): "On the performance of linear contracts," *Journal of Economics & Management Strategy*, 20, 159–193.

- CARROLL, G. (2012): “Robustness and Linear Contracts,” *Stanford University Working Paper*.
- CESA-BIANCHI, N. AND G. LUGOSI (2006): *Prediction, Learning, and Games*, Cambridge University Press.
- CHASSANG, S. (2011): “Calibrated Incentive Contracts,” *Princeton University Economic Theory Center Working Paper No. 013-2011*.
- CHU, L. AND D. SAPPINGTON (2007): “Simple cost-sharing contracts,” *The American Economic Review*, 97, 419–428.
- 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,” *Working paper, Stanford University*.
- 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 Paper, Caltech*.
- EDMANS, A., X. GABAIX, T. SADZIK, AND Y. SANNIKOV (2012): “Dynamic Incentive Accounts,” *Journal of Finance*, 67, 1603–1647.
- ELIAZ, K. (2002): “Fault Tolerant Implementation,” *Review of Economic Studies*, 69, 589–610.
- ESCOBAR, J. AND J. TOIKKA (2009): “A Folk Theorem with Markovian Private Information,” *Working Paper, Massachusetts Institute of Technology*.
- FARRELL, J. AND E. MASKIN (1989): “Renegotiation in repeated games,” *Games and economic behavior*, 1, 327–360.
- FONG, K. (2008): “Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons,” *Working Paper, Stanford University*.
- 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.
- GARRETT, D. AND A. PAVAN (2010): “Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes,” .
- GOETZMANN, W., J. INGERSOLL, AND S. ROSS (2003): “High-Water Marks and Hedge Fund Management Contracts,” *The Journal of Finance*, 58, 1685–1718.
- GRADWOHL, R. AND Y. SALANT (2011): “How to Buy Advice,” *Working Paper, Northwestern University*.
- GREEN, E. (1987): “Lending and the smoothing of uninsurable income,” in *Contractual arrangements for intertemporal trade*, ed. by E. Prescott and N. Wallace, University of Minnesota Press, 3–25.
- HANNAN, J. (1957): “Approximation to Bayes Risk in Repeated Play,” in *Contributions to the Theory of Games*, ed. by M. Dresher, A. Tucker, and P. Wolfe, vol. 3, 97–139.
- HARTLINE, J. D. AND T. ROUGHGARDEN (2008): “Optimal Mechanism Design and Money Burning,” in *Symposium on Theory Of Computing (STOC)*, 75–84.
- HE, Z. (2009): “Optimal executive compensation when firm size follows geometric brownian motion,” *Review of Financial Studies*, 22, 859.
- HE, Z. AND W. XIONG (2010): “Delegated Asset Management and Investment Mandates,” *Working Paper, Princeton University*.
- HOEFFDING, W. (1963): “Probability Inequalities for Sums of Bounded Random Variables,” *Journal of the American Statistical Association*, 58, 13–30.
- HOLMSTRÖM, B. AND P. MILGROM (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55, 303–328.
- HURWICZ, L. AND L. SHAPIRO (1978): “Incentive structures maximizing residual gain under incomplete information,” *The Bell Journal of Economics*, 9, 180–191.
- JACKSON, M. AND H. SONNENSCHNEIN (2007): “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75, 241–257.
- LAFFONT, J. AND J. TIROLE (1988): “The dynamics of incentive contracts,” *Econometrica*, 1153–1175.
- 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, London School of Economics*.
- OLSZEWSKI, W. AND M. PESKI (2011): “The principal-agent approach to testing experts,” *American Economic Journal: Microeconomics*, 3, 89–113.
- 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.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2010): “Infinite-Horizon Mechanism Design,” *Working Paper, Northwestern University*.
- 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.
- (2003): “Simple menus of contracts in cost-based procurement and regulation,” *The American economic review*, 93, 919–926.
- RUBINSTEIN, A. (1979): “An optimal conviction policy for offenses that may have been committed by accident,” in *Applied game theory*, ed. by J. Brams, A. Schotter, and G. Schwodiauer, Physica-Verlag, Wurzburg, 406–413.
- RUBINSTEIN, A. AND M. YAARI (1983): “Repeated insurance contracts and moral hazard,” *Journal of Economic Theory*, 30, 74–97.
- SANNIKOV, Y. (2007): “Agency problems, screening and increasing credit lines,” *Working paper, Princeton University*.
- (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.
- TCHISTYI, A. (2006): “Security design with correlated hidden cash flows: The optimality of performance pricing,” *Working Paper, University of California at Berkeley*.
- ZHU, J. (forthcoming): “Optimal contracts with shirking,” *Review of Economic Studies*.

Online Appendix for “Calibrated Incentive Contracts”

Section OA 1 of this Online Appendix extends the analysis of Sections 3 and 4. Section OA 2 provides simulations illustrating key properties of calibrated contracts.

OA 1 Extensions

Appendix OA 1.1 extends the analysis to the case where principal and agent discount future payoffs. Appendix OA 1.2 shows how to calibrate a broader class of high-liability contracts, including log-scoring rules. Appendix OA 1.3 shows that the calibrated contracts of Section 4 perform well even if the agent isn’t rational and behaves suboptimally over any arbitrary interval of time. Appendix OA 1.4 considers the case where the principal can use more than one agent. Finally, Appendix OA 1.5 proves improved performance bounds for the screening contracts introduced in Section 5 when expected returns are grainy. Chassang (2011) contains additional extensions dealing with varying wealth, varying preferences, risk-aversion and non-convex action spaces.

OA 1.1 Discounting

The analysis of Section 4 can be extended to environments where principal and agent discount the future by a factor δ so that the agent’s payoffs are $\mathbb{E} \left(\sum_{t=1}^N \delta^{t-1} (\pi_t - c_t) \right)$ and the principal’s surplus is $\mathbb{E} \left(\sum_{t=1}^N \delta^{t-1} (w_t - w_t^0) \right)$. Let $N_\delta = \sum_{t=1}^N \delta^t$. This appendix shows that under discounting, the performance bound of Theorem 2 extends with a loss of order $\sqrt{1/N_\delta}$ instead of $\sqrt{1/N}$.

Benchmark contract. The benchmark contract still gives the agent reward $\pi_t = \alpha(w_t - w_t^0)$ in every period t . This linear contract guarantees the principal a payoff bound similar to that of Theorem 1(i). For any contract (λ, π) , where sequence $\lambda = (\lambda_t)_{t \geq 1}$ may be constant and equal to 1, define

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

the average discounted per-period returns accruing to the principal under contract (λ, π) . Let r_α denote returns accruing to the principal under the benchmark contract. In addition

define

$$r_{\max}(\hat{c}) \equiv \sup_{\substack{c \text{ s.t.} \\ \mathbb{E}\left[\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} c_t\right] \leq \hat{c}}} \mathbb{E}_{c, a^*} \left(\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} \langle a_t^* - a_t^0, r_t \rangle \right)$$

the maximum discounted per-period returns that can be generated at an expected discounted per-period cost of \hat{c} .

Lemma OA 1. *For all environments \mathcal{P} ,*

$$r_\alpha \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right).$$

Proof. The proof is identical to that of Theorem 1(i). □

Calibration. The calibrated contract is built using the following regrets

$$\mathcal{R}_{1,T} = \sum_{t=1}^T \delta^{t-1} (\pi_t - \alpha(w_t - w_t^0)) \quad \text{and} \quad \mathcal{R}_{2,T} = \max_{T \leq T'} \sum_{t=T'}^T \delta^{t-1} (1 - \lambda_t) (w_t - w_t^0)^+.$$

Contract parameters $(\lambda_t, \pi_t)_{t \geq 1}$ are computed recursively according to

$$\lambda_t = \frac{\alpha \mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_t = \begin{cases} \alpha(w_t - w_t^0)^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise} \end{cases}.$$

The following result extends Lemma 1, showing that incentives are approximately correct.

Lemma OA 2 (approximate incentives). *For all T , and all possible histories,*

$$\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} (1 - \lambda_t) (w_t - w_t^0) \leq \frac{w\bar{d}}{\sqrt{N_\delta}} \tag{31}$$

$$-\frac{w\bar{d}}{N_\delta} \leq \frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} [\pi_t - \alpha(w_t - w_t^0)] \leq \frac{w\bar{d}}{\sqrt{N_\delta}}. \tag{32}$$

Proof. Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ denote the vector of regrets, and $\rho_{T+1} = \mathcal{R}_{T+1} - \mathcal{R}_T$. Contract (λ, π) is calibrated so that in every period $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$. It follows that

$$\|\mathcal{R}_N^+\|^2 \leq \sum_{t=1}^N \|\rho_T\|^2.$$

Furthermore, we have that $\|\rho_T\|^2 \leq \delta^{2T} w\bar{d}$, which implies that

$$\|\mathcal{R}_T^+\|^2 \leq w\bar{d} \sum_{t=1}^N \delta^{2(t-1)} \leq w\bar{d} \sum_{t=1}^N \delta^{t-1}.$$

This implies the right-hand sides of (31) and (32). The left-hand side of (32) follows from a proof identical to that of the left-hand side of (13). \square

This implies the following bounds for returns $r_{\lambda,\pi}$.

Theorem OA 1. *Pick $\alpha_0 \in (0, 1)$ and for $\eta > 0$, let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists $m \geq 0$ such that for all environments \mathcal{P} , all δ and all N ,*

$$r_{\lambda,\pi} \geq (1 - \eta)r_{\alpha_0} - \frac{m}{\sqrt{N_\delta}} \quad (33)$$

$$r_{\lambda,\pi} \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} - \frac{3\bar{d}}{\sqrt{N_\delta}} \right). \quad (34)$$

Proof. The proof follows the same steps as that of Theorem 2, with the bounds provided in Lemma OA 2 replacing those provided in Lemma 1. \square

OA 1.2 Calibrating a broader class of contracts

This section provides sufficient conditions ensuring that a benchmark high-liability contract can be calibrated using limited-liability contracts. Fix a family of limited liability constraints

$$\forall t \geq 1, \quad 0 \leq \pi_t \leq \bar{\pi}_t, \quad (4'')$$

such that for all t , $w_t \leq \bar{\pi}_t$, and take as given a contract with aggregate final rewards denoted by Π_N^0 (where Π_N^0 is adapted to the principal's information at time N). It turns out that contract Π_N^0 can be calibrated by a dynamic contract satisfying limited liability constraint (4'') whenever the following assumption holds.

Assumption OA 1. *Benchmark contract Π_N^0 can be written as $\Pi_N^0 = \sum_{t=1}^N \pi_t^0$, with $(\pi_t^0)_{t \geq 1}$ such that*

- (i) π_t^0 is adapted to the information available to the principal at time t ;
- (ii) $w_t = w_t^0$ implies $\pi_t^0 \geq 0$;

(iii) $\pi_t^0 \leq \bar{\pi}_t$ and there exists $\bar{\pi} > 0$ independent of N such that, $\sup |\pi_t^0| \leq \bar{\pi}$.

Note that π_t^0 may be negative and require liability from the agent. It is immediate that Assumption OA 1 holds for all contracts of the form $\Pi_N^0 = \sum_{t=1}^N \alpha_t^0 (w_t - w_t^0)$ where $\alpha_t^0 \in (0, 1)$ is adapted to public information $(\mathcal{F}_t^0)_{t \geq 1}$. Assumption OA 1 also holds for contracts of the form

$$\Pi_N^0 = G \left(\sum_{t=1}^N \phi(w_t - w_t^0) \right)$$

where $\phi(0) = G(0) = 0$ and G and ϕ are Lipschitz, with constants κ_G and κ_ϕ such that $\kappa_G \kappa_\phi w_t \leq \bar{\pi}_t$. For instance, if for all t , $\bar{\pi}_t = w_t$, this includes contracts such that the agent gets paid a positive reward only when returns are above a threshold, i.e contracts such that

$$\Pi_N^0 = \begin{cases} \alpha \sum_{t=1}^N w_t - w_t^0 & \text{if } \sum_{t=1}^N w_t - w_t^0 < 0 \\ 0 & \text{if } \sum_{t=1}^N w_t - w_t^0 \in [0, \underline{W}] \\ \alpha \left(\left[\sum_{t=1}^N w_t - w_t^0 \right] - \underline{W} \right) & \text{if } \sum_{t=1}^N w_t - w_t^0 > \underline{W}. \end{cases} \quad (35)$$

Another example of alternative benchmark contract is to reward the agent for probability assessments according to a log-scoring rule. This example will be discussed in further detail after stating the main calibration result.

Calibration. Theorem OA 2, stated below, shows that the performance of any contract satisfying Assumption OA 1 can be approximated in a prior-free way using dynamic limited liability contracts.

As in Section 4 an additional incentive wedge is necessary to take care of potentially binding global incentive constraints. For any $\eta > 0$ define the auxiliary contract

$$\pi_t^\eta \equiv \pi_t^0 + \eta(w_t - w_t^0 - \pi_t^\eta) = \frac{1}{1+\eta} \pi_t^0 + \frac{\eta}{1+\eta} (w_t - w_t^0).$$

If contract $(\pi_t^0)_{t \geq 1}$ satisfies Assumption OA 1, then so does contract $(\pi_t^\eta)_{t \geq 1}$. In particular, $|\pi_t^\eta| \leq \frac{1}{1+\eta} \bar{\pi} + \frac{\eta}{1+\eta} w \bar{d} \equiv \bar{\pi}^\eta$.

The approach consists in calibrating the incentives provided by contract $(\pi_t^\eta)_{t \geq 1}$. Once again, the two instruments used are rewards $(\pi_t)_{t \geq 1}$ and the proportion of resources $(\lambda_t)_{t \geq 1}$ managed by the agent. Define $\pi_t^\eta(\lambda_t) = \lambda_t \pi_t^\eta$. The regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ to be minimized

are:

$$\mathcal{R}_{1,T} = \sum_{t=1}^T \pi_t - \pi_t^\eta(\lambda_t) \quad (\text{no excess rewards}) \quad (36)$$

$$\mathcal{R}_{2,T} = \max_{T' \leq T} \sum_{t=T'}^T \pi_t^\eta - \pi_t^\eta(\lambda_t) \quad (\text{no foregone performance}). \quad (37)$$

The usual approachability condition yields contract parameters $(\lambda_t, \pi_t)_{t \geq 1}$ of the form,

$$\lambda_{T+1} = \frac{\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} [\pi_{T+1}^\eta]^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (38)$$

As in Section 4 this ensures that the vector of regrets $(\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ remains of order \sqrt{T} , so that incentives are approximately correct. The following performance bounds obtain.

Theorem OA 2. *There exists a constant m independent of environment \mathcal{P} and time horizon N , such that under contract $(\lambda_t, \pi_t)_{t \geq 1}$, returns accruing to the principal satisfy*

$$\forall h_T, \quad r_{\lambda, \pi} | h_T \geq \frac{1}{1 + \eta} r_{\pi^0} | h_T - m \frac{1}{\sqrt{N}} \quad (39)$$

Proof. The proof uses the following extension of Lemma 1.

Lemma OA 3 (incentive approximation). *For any realization of uncertainty,*

$$-\bar{\pi}^\eta \leq \sum_{t=1}^T \pi_t - \pi_t^\eta(\lambda_t) \leq \bar{\pi}^\eta \sqrt{T} \quad (40)$$

$$-\bar{\pi}^\eta \sqrt{T} \leq \max_{T' \leq T} \sum_{t=T'}^T \pi_t^\eta - \pi_t^\eta(\lambda_t) \leq \bar{\pi}^\eta \sqrt{T}. \quad (41)$$

Proof. Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ denote the vector of regrets and $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$ the vector of flow regrets. Using the fact that $\mathcal{R}_{2,T+1} = \mathcal{R}_{2,T}^+ + (1 - \lambda_{T+1})\pi_{T+1}^\eta$, and exploiting the equality $\mathcal{R}_{2,T}^+(\mathcal{R}_{2,T} - \mathcal{R}_{2,T}^+) = 0$, we have

$$\begin{aligned} \langle \mathcal{R}_T^+, \rho_{T+1} \rangle &= \mathcal{R}_{1,T}^+ [\pi_{T+1} - \lambda_{T+1} \pi_{T+1}^\eta] + \mathcal{R}_{2,T}^+ (1 - \lambda_{T+1}) \pi_{T+1}^\eta \\ &= \mathcal{R}_{1,T}^+ \pi_{T+1} + [(1 - \lambda_{T+1}) \mathcal{R}_{2,T}^+ - \lambda_{T+1} \mathcal{R}_{1,T}^+] \pi_{T+1}^\eta. \end{aligned}$$

Hence, the contract $(\lambda_t, \pi_t)_{t \geq 1}$ defined by (38) ensures that for all realizations of \mathbf{r}_{T+1} ,

$$\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0.$$

We now prove by induction that $\|\mathcal{R}_T^+\|^2 \leq \sum_{t=1}^T (\pi_t^\eta)^2$. The property clearly holds for $T = 1$. We now assume that it holds at T and show it must hold at $T + 1$. Consider first the case where $\mathcal{R}_{2,T} > 0$.

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

where we used the fact that by construction, $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$. Furthermore, we have that

$$\begin{aligned} \|\rho_{T+1}\|^2 &\leq (\pi_{T+1} - \pi_{T+1}^\eta(\lambda_{T+1}))^2 + (\mathcal{R}_{2,T}^+ + (1 - \lambda_{T+1})\pi_{T+1}^\eta - \mathcal{R}_{2,T})^2 \\ &\leq \lambda_{T+1}^2 (\pi_{T+1}^\eta)^2 + (1 - \lambda_{T+1})^2 (\pi_{T+1}^\eta)^2 \\ &\leq (\pi_{T+1}^\eta)^2. \end{aligned}$$

Using the induction hypothesis, this implies that $\|\mathcal{R}_{T+1}^+\|^2 \leq \sum_{t=1}^{T+1} (\pi_t^\eta)^2$. A similar proof holds when $\mathcal{R}_{2,T} < 0$, taking into account that in this case, $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})\pi_{T+1}^\eta$. Hence, by induction, this implies that for all $T \geq 1$, $\|\mathcal{R}_T^+\|^2 \leq \sum_{t=1}^T (\pi_t^\eta)^2$. Since $|\pi_t^\eta| \leq \bar{\pi}^\eta$, this implies inequality (41) and the right-hand side of (40). The left-hand side of (40) follows from an induction identical to that used to prove the left-hand side of (13). \square

We can now conclude the proof of Theorem OA 2.

Let us begin by proving (39) starting from initial history h_0 . Let (\hat{c}, \hat{a}) denote an optimal strategy for the agent under calibrated contract (λ, π) , and let (c, a) denote an optimal strategy for the agent under benchmark contract $\pi^0 = (\pi_t^0)_{t \geq 1}$. By optimality of (\hat{c}, \hat{a}) under (λ, π) , we obtain that

$$\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \pi_t - \hat{c}_t \right) \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t - c_t \right).$$

By (40) this implies that

$$\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \pi_t^\eta(\lambda_t) - \hat{c}_t \right) + \bar{\pi}^\eta \sqrt{N} \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^\eta(\lambda_t) - c_t \right) - \bar{\pi}^\eta.$$

By (41) we obtain

$$\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^\eta - \hat{c}_t \right) + \bar{\pi}^\eta \sqrt{N} \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^\eta - c_t \right) - \bar{\pi}^\eta (1 + \sqrt{N}).$$

Using the fact that (c, a) is optimal under contract $(\pi_t^0)_{t \geq 1}$, and that necessarily, $\mathbb{E}_{c, a}(\pi_t^0) \geq 0$, this implies that

$$\begin{aligned} \mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^0 + \lambda_t \eta (w_t - w_t^0 - \pi_t^\eta) - \hat{c}_t \right) &\geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^0 + \eta (w_t - w_t^0 - \pi_t^\eta) - c_t \right) - \bar{\pi}^\eta (2\sqrt{N} + 1) \\ &\geq \mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^0 - \hat{c}_t \right) + \mathbb{E}_{c, a} \left(\sum_{t=1}^N \eta (w_t - w_t^0 - \pi_t^\eta) \right) - \bar{\pi}^\eta (2\sqrt{N} + 1). \end{aligned}$$

Thus, using (40) and the fact that $w_t - w_t^0 - \pi_t^\eta = \frac{1}{1+\eta}(w_t - w_t^0 - \pi_t^0)$, we obtain that

$$\begin{aligned} &\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0 - \pi_t^\eta) \right) - \mathbb{E}_{c, a} \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t^\eta \right) \geq -\frac{\bar{\pi}^\eta}{\eta} (2\sqrt{N} + 1) \\ \Rightarrow &\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0) - \pi_t \right) \geq \frac{1}{1+\eta} \mathbb{E}_{c, a} \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t^0 \right) - \frac{\bar{\pi}^\eta}{\eta} (1 + (2+\eta)\sqrt{N}). \end{aligned}$$

Inequality (39) at h_0 follows from normalizing by $\frac{1}{wN}$.

Inequality (39) continues to hold conditional on any history because the incentive bounds provided by Lemma OA 3 hold starting from any interim period T . \square

The following example applies this analysis to the calibration of log-scoring rules.

Calibrating log-scoring rules. The benchmark linear contract $\pi_t = \alpha(w_t - w_t^0)$ ensures that the agent has incentives to make allocation decisions that maximize expected returns conditional on information. A potential alternative is to elicit truthful beliefs over returns from the agent using a log-scoring rule, and implement the allocation that maximizes surplus under these beliefs.

Log-scoring rules take the following form. Assume for simplicity that the set R of possible returns \mathbf{r}_t is finite. In each period t , the agent gets rewarded according to

$$\pi_t^{ls} = \gamma \log \left(\frac{f_t(\mathbf{r}_t)}{f_t^0(\mathbf{r}_t)} \right) \quad \text{with } \gamma > 0,$$

where f_t is a distribution over realized returns \mathbf{r}_t stated by the agent in period t , $f_t^0 = P(\cdot|\mathcal{F}_t^0)$ is the principal's belief conditional on public information \mathcal{F}_t^0 , and \mathbf{r}_t are the realized returns. Given f_t , the allocation a_t is chosen to maximize expected returns $\mathbb{E}_{f_t}(w_t - w_t^0)$ under belief f_t . To insure that rewards π_t^{ls} are bounded, the following restriction is imposed.

Assumption OA 2 (bounded likelihood ratio). *There exists $\bar{l} > \underline{l} > 0$ such that for every history,*

$$\forall \mathbf{r}_t \in R, \quad \frac{P(\mathbf{r}_t|\mathcal{F}_t)}{P(\mathbf{r}_t|\mathcal{F}_t^0)} \in [\underline{l}, \bar{l}].$$

It is well known that log-scoring contracts $(\pi_t^{ls})_{t \geq 1}$ induce truthful revelation of beliefs. In addition, the agent can expect positive expected rewards if and only if his belief is different from that of the principal.

Fact 4. *Truth-telling, i.e. sending message $f_t = P(\cdot|\mathcal{F}_t)$, maximizes the agent's payoff conditional on information. An agent whose belief $P(\cdot|\mathcal{F}_t)$ coincides with that of the principal conditional on public information $P(\cdot|\mathcal{F}_t^0)$ gets an expected payoff of zero.*

The proof of this fact is standard and omitted. Noting that $0 \leq \pi_t \leq \bar{\pi} = \alpha \log(\bar{l}/\underline{l})$, Theorem OA 2 applies, and the contract (λ, π) derived from $(\pi_t^{ls})_{t \geq 1}$ according to (38) satisfies performance bound (39), i.e. it successfully approximates the performance of the benchmark log-scoring rule while requiring no liability from the agent and only limited liability from the principal.

Note that this result should be viewed as an illustration of the broader applicability of the contract calibration approach developed in the paper, rather than an endorsement of log-scoring rules as an appropriate benchmark contract. Indeed, contrary to benchmark linear contracts of the form $\pi_t = \alpha(w_t - w_t^0)$, log-scoring rules do not guarantee that the principal must be getting positive surplus out of the relationship, i.e. it does not satisfy the “no loss” property emphasized in Fact 1.¹ The following example illustrates the problem in a stark manner.

There are two assets: a good asset 0, with i.i.d. returns $r_{0,t}$ uniformly distributed over $\{\frac{1}{100}, \frac{2}{100}, \dots, 1\}$ in every period t , and a bad asset 1 with i.i.d. returns $r_{1,t}$ uniformly distributed over $\{-\frac{99}{100}, -\frac{98}{100}, \dots, 0\}$. The principal has no further information about returns, whereas the agent observes returns $(r_{0,t}, r_{1,t})$ without noise. Clearly, the agent has a lot of information, and under the log-scoring rule, he will be rewarded for this information since

¹Fact 1 also shows that the only contracts satisfying “no loss” for all possible strategies giving the agent positive surplus are in fact benchmark linear contracts.

it considerably reduces uncertainty. However this information is of no value to the principal since the good asset always dominates the bad asset.

Contrary to benchmark linear contracts, log-scoring rules reward the agent for any information, regardless of whether it is valuable or not. Note also that while potential losses could be controlled by letting coefficient γ go to zero, this is not generally helpful since this also implies that the agent has no incentives to exert effort and acquire information.

OA 1.3 Robustness to Accidents

The analysis presented in the main text of the paper assumes that the agent is rational. It turns out that calibrated contracts are robust to the possibility of “accidents” during which the agent behaves in arbitrary, possibly suboptimal, ways over an extended amount of time.

An accident may correspond to a temporary mistake in the agent’s trading strategy or an error in his data; alternatively, the agent may be temporarily irrational or have unmodeled incentives to misbehave (e.g. he is bribed to unload bad risks on the principal). Formally, this is modeled by assuming that during a random time interval $[T_1, T_2]$ —in the accident state—the agent is constrained to use an exogenously specified allocation strategy $a^\Delta = (a_t^\Delta)_{t \geq 1}$.²

The agent takes into account the possibility of such accidents when choosing his strategy and has an ex ante belief over the interval $[T_1, T_2]$ and over his prescribed behavior a^Δ during the accident. Strategy a^Δ may be arbitrarily bad (within the bounds imposed by Assumption 1) and need only be measurable with respect to final information \mathcal{F}_N . For instance, during the lapse of the accident, the agent could pick the worst ex post asset allocation in every period. Robustness to accidents of this kind is related to Eliaz (2002) which studies how well mechanisms perform if some players are faulty, i.e. if they use non-optimal strategies. Here, robustness to accidents can be thought of as fault tolerance with respect to the agent’s selves over $[T_1, T_2]$.

It should be noted that in this environment, the benchmark linear contract is no longer sufficient to guarantee good performance. Since expected returns $\mathbb{E}_{a^\Delta}(w_t - w_t^0)$ can be negative in an accident period, 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 instead of a^Δ , the exogenous allocation

²The analysis given here allows accidents to occur over a single interval of time. The analysis extends without change to environments with a fixed number of accident intervals independent of horizon N .

during accident states is

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

where w_t^Δ is the realized wealth under allocation a_t^Δ at time t . Denote by $r_{\lambda, \pi}^\Delta$ the net expected per-period returns to the principal when accidental behavior is $(a_t^\Delta)_{t \geq 1}$ and the calibrated contract (λ, π) defined in (12) is used. Denote by $r_\alpha^{\Delta\Delta}$ the net expected per-period returns to the principal when accidental behavior is $(a_t^{\Delta\Delta})_{t \geq 1}$ and the benchmark contract of parameter α is used. The following holds.

Theorem OA 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,*

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

Proof. The notation of Section 4 is extended by adding superscripts $^\Delta$ and $^{\Delta\Delta}$ to denote relevant objects under the original accidental allocation a^Δ , and under the lucky accidental allocation $a^{\Delta\Delta}$. For instance, let $(w_t^{\Delta\Delta})_{t \geq 1}$ and $\Sigma_N^{\Delta\Delta} = \sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0$ denote potential realized wealth and aggregate performance when accidents are lucky. For concision this extension is done for all time periods $t \in \{1, \dots, N\}$, with the understanding that the allocation is exogenous over interval $[T_1, T_2]$, but endogenous in other time periods; i.e. an allocation policy $(a_t^\Delta)_{t \geq 1}$ corresponds to endogenous allocations a_t for $t \notin [T_1, T_2]$ and coincides with a_t^Δ for $t \in [T_1, T_2]$.

Given these adjustments, the proof of Theorem OA 3 is analogous to that of Theorem 2, the key step being to provide an adequate extension of Lemma 1. Because inequality (13) still applies, we have that

$$-\alpha w \bar{d} \leq \sum_{t=1}^N \pi_t^\Delta - \alpha \lambda_t (w_t^\Delta - w_t^0) \leq \alpha w \bar{d} \sqrt{N}. \quad (42)$$

This corresponds to “correct rewards” condition (9). In addition, let us show that for any investment strategy of the agent, the following variant of “no foregone gains” condition (10) must hold

$$\left(\sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0 \right) - 4w \bar{d} \sqrt{N} \leq \sum_{t=1}^N \lambda_t (w_t^\Delta - w_t^0). \quad (43)$$

We have that

$$\sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0 = \left[\sum_{t=1}^{T_1-1} w_t^\Delta - w_t^0 \right] + \left[\sum_{t=T_1}^{T_2} w_t^\Delta - w_t^0 \right]^+ + \left[\sum_{t=T_2+1}^N w_t^\Delta - w_t^0 \right].$$

Applying inequality (14), we obtain

$$\sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0 \leq \begin{cases} \left[\sum_{t=1}^N \lambda_t(w_t^\Delta - w_t^0) \right] + w\bar{d}\sqrt{N} & \text{if } \sum_{t=T_1}^{T_2} w_t^\Delta - w_t^0 > 0 \\ \left[\sum_{t=1}^{T_1-1} \lambda_t(w_t^\Delta - w_t^0) \right] + \left[\sum_{t=T_2+1}^N \lambda_t(w_t^\Delta - w_t^0) \right] + 3w\bar{d}\sqrt{N} & \text{otherwise.} \end{cases}$$

By (42), it follows that

$$\begin{aligned} -\alpha w\bar{d}\sqrt{T_2} &\leq \sum_{t=1}^{T_2} \alpha \lambda_t(w_t^\Delta - w_t^0) - \pi_t^\Delta \leq \alpha w\bar{d} \\ -\alpha w\bar{d}\sqrt{T_1-1} &\leq \sum_{t=1}^{T_1-1} \alpha \lambda_t(w_t^\Delta - w_t^0) - \pi_t^\Delta \leq \alpha w\bar{d}. \end{aligned}$$

Subtracting these two inequalities yields that,

$$-\alpha w\bar{d}(1 + \sqrt{T_2}) \leq \sum_{t=T_1}^{T_2} \alpha \lambda_t(w_t^\Delta - w_t^0) - \pi_t^\Delta.$$

Since flow rewards π_t^Δ are weakly positive, this implies that for any realization of returns,

$$\sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0 \leq \left(\sum_{t=1}^N \lambda_t(w_t^\Delta - w_t^0) \right) + 4w\bar{d}\sqrt{N}.$$

This implies (43). Given (42) and (43), Theorem OA 3 follows by applying Lemma A.1. \square

OA 1.4 Multi-agent contracts

The analysis presented in the paper focused on contracting with a single agent. This appendix shows how to extend the logic of Sections 3 and 4 to environments with multiple agents. The framework is identical to that of Section 2 except that there are now J agents denoted by $j \in \{1, \dots, J\}$, each of whom makes private information acquisition decisions $c_{j,t} \in [0, +\infty)$, inducing a filtration \mathcal{F}_t^j . In each period t , agent j suggests an asset allocation $a_{j,t}$ inducing potential wealth $w_{j,t} = w(1 + \langle a_{j,t}, \mathbf{r}_t \rangle)$.

As in Section 2 the environment is general. Public and private signals $(I_t^0, I_c^j(c_{j,t}))_{j \in \{1, \dots, J\}}$ are arbitrary random variables from an underlying measurable state space (Ω, σ) to a mea-

surable signal space $(\mathcal{I}, \sigma_{\mathcal{I}})$. The environment $\mathcal{P} = (\Omega, \sigma, P)$ is specified by defining a probability measure P on (Ω, σ) . This probability measure is unrestricted: the agents may have access to different information, their respective ability to generate information may differ, vary over time, and be correlated in arbitrary ways. Filtration $(\mathcal{F}_t^0)_{t \geq 1}$ still denotes the public information filtration available to the principal.

The first step of the analysis extends the high-liability benchmark contract of Section 3. The second step of the analysis shows how to calibrate this high-liability contract.

Multi-agent benchmark contracts. The multi-agent contract described here is a direct extension of the linear contract described in Section 3. Each agent $j \geq 1$ is paid according to a linear contract in which the allocation of agent $j - 1$ serves as the default allocation previously corresponding to a_t^0 , i.e. each agent is paid a share α of his externality on the principal, taking into account the information provided by previous agents. Resources are invested according to the allocation $a_{J,t}$ suggested by the last agent.

More precisely, in each period t , allocations $a_{j,t}$ are submitted by agents in increasing order of rank j . This ordering is a constraint imposed by the mechanism. The mechanism informs each agent j of the allocations $(a_{j',t})_{j' < j}$ chosen by agents $j' < j$. Agent j receives no information about the allocations chosen by agents $j'' > j$. Under the benchmark contract, payments $\pi_{j,t}$ to agent j are defined by

$$\forall j \in \{1, \dots, J\}, \quad \pi_{j,t} = \alpha(w_{j,t} - w_{j-1,t}). \quad (44)$$

The strategy profile (c_j, a_j) of agent j must be adapted to the information available to the agent (by construction this includes allocations by previous agents). The set of such adapted strategies is denoted by $\mathcal{C}_j \times \mathcal{A}_j$.³ Furthermore define $(c, a) = (c_j, a_j)_{j \in \{1, \dots, J\}}$ and $\mathcal{C} \times \mathcal{A} = \prod_{j \in \{1, \dots, J\}} \mathcal{C}_j \times \mathcal{A}_j$ the set of adapted strategy profiles. For any $\hat{c} \in [0, +\infty)$, the maximum returns that can be obtained at an expected per-period cost of \hat{c} are denoted by

$$r_{\max}(\hat{c}) = \max_{\substack{(c,a) \in \mathcal{C} \times \mathcal{A} \\ \frac{1}{N} \mathbb{E}(\sum_{j,t} c_{j,t}) \leq \hat{c}}} \frac{1}{wN} \mathbb{E}_{c,a} \left(\sum_{t=1}^N w_t^J - w_t^0 \right).$$

Denote by r_α the average returns accruing to the principal under this benchmark contract.

³ Because of the hierarchical structure of the mechanism, agent $j' < j$ is indifferent about whether or not to send information to agent j . For simplicity it is assumed that agent j' shares his information with agents $j > j'$.

The following bound extends point (i) of Theorem 1.

Lemma OA 4. *For any environment \mathcal{P} ,*

$$wr_\alpha \geq (1 - \alpha) \max_{\hat{c} \in [0, +\infty)} \left(wr_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha} \right).$$

As in Theorem 1, given restrictions on $r_{\max}(\cdot)$, a rationale for choosing α is to maximize this lower bound. Note that similarly to the benchmark contract of Section 3, this contract also satisfies no-loss.

Proof. Optimal strategies for the agents $(c^*, a^*) = (c_j^*, a_j^*)_{j \in \{1, \dots, J\}}$ are such that for any other profile of strategies $(c, a) = (c_j, a_j)_{j \in \{1, \dots, J\}}$, and for all $j \in \{1, \dots, J\}$,

$$\mathbb{E}_{c_j^*, a_j^*} \left[\sum_{t=1}^N \alpha(w_{j,t} - w_{j-1,t}) - c_{j,t}^* \right] \geq \mathbb{E}_{c_j, a_j} \left[\sum_{t=1}^N \alpha(w_{j,t} - w_{j-1,t}) - c_{j,t} \right].$$

Summing over j , this implies that

$$\mathbb{E}_{c^*, a^*} \left[\sum_{t=1}^N \alpha(w_{J,t} - w_{0,t}) - \sum_{j \in J} c_{j,t}^* \right] \geq \mathbb{E}_{c, a} \left[\sum_{t=1}^N \alpha(w_{J,t} - w_{0,t}) - \sum_{j \in J} c_{j,t} \right].$$

$$\text{Hence, } \mathbb{E}_{c^*, a^*} \left[\sum_{t=1}^N (1 - \alpha)(w_{J,t} - w_{0,t}) \right] \geq \frac{1 - \alpha}{\alpha} \mathbb{E}_{c, a} \left[\sum_{t=1}^N \alpha(w_{J,t} - w_{0,t}) - \sum_{j \in J} c_{j,t} \right].$$

Since this holds for any strategy profile (c, a) , we obtain that indeed

$$wr_\alpha \geq (1 - \alpha) \max_{\hat{c} \in [0, +\infty)} \left(wr_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha} \right).$$

□

Calibrated contracts. The high-liability multi-agent contract described in (44) can be calibrated using the methods of Section 4. The main difference is that there is now a vector $\lambda_t = (\lambda_{j,t})_{j \in \{1, \dots, J\}} \in [0, 1]^J$ of J scaling factors used to define adjusted allocations $a_{j,t}^\lambda$ in the following recursive manner:

$$a_{1,t}^\lambda = \lambda_{1,t} a_{1,t} + (1 - \lambda_{1,t}) a_t^0 \quad \text{and} \quad \forall j > 1, \quad a_{j,t}^\lambda = \lambda_{j,t} a_{j,t} + (1 - \lambda_{j,t}) a_{j-1,t}^\lambda.$$

Let $w_{j,t}^\lambda$ denote the corresponding wealth realizations. For all $j \geq 1$, define regrets

$$\mathcal{R}_{j,T}^1 = \sum_{t=1}^T \pi_{j,t} - \alpha(w_{j,t}^\lambda - w_{j-1,t}^\lambda) \quad (\text{correct rewards}) \quad (45)$$

$$\mathcal{R}_{j,T}^2 = \max_{T' \leq T} \sum_{t=T'}^T w_{j,t} - w_{j,t}^\lambda \quad (\text{no foregone returns}). \quad (46)$$

Keeping these regrets small corresponds to implementing appropriate generalizations of incentive properties (13) and (14) for all agents. The usual approachability condition implies that regrets $(\mathcal{R}_{j,T}^1, \mathcal{R}_{j,T}^2)_{j \in \{1, \dots, J\}}$ can be kept small by choosing contract parameters $(\lambda_j, \pi_j)_{j \in \{1, \dots, J\}}$ according to,

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

Under this calibrated multi-agent contract the following extension of Theorem 2 obtains.

Theorem OA 4. *Pick $\alpha_0 > 0$ and for $\eta \in (0, 1)$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists a constant m independent of environment \mathcal{P} , time horizon N such that the multi-agent calibrated contract $(\lambda, \pi) = (\pi_j, \lambda_j)_{j \in \{1, \dots, J\}}$ of parameter α satisfies*

$$\forall h_T, \quad r_{\lambda, \pi} | h_T \geq (1 - \eta) r_{\alpha_0} | h_T - \frac{m}{\sqrt{N}}. \quad (47)$$

Proof. The result follows from applying Theorem 2 iteratively over agents $j \in \{1, \dots, J\}$. \square

OA 1.5 Screening when returns are grainy

This appendix shows that the efficiency bound given in Theorem 3 can be improved when expected returns are either zero or bounded away from 0: performance losses are of order $\sqrt{1/N}$ rather than $\sqrt{\ln N/N}$.

Assumption OA 3 (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 OA 5. *Pick α_0 and for any $\eta > 0$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. If Assumption OA 3 holds, there exists a constant m such that for all N and all probability spaces \mathcal{P} ,*

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

Proof. The proof strategy is identical to that of Theorem 3. The missing step is to improve the left-hand side of bound (29). Let (c, a^*) denote the agent's optimal strategy under the benchmark contract of parameter α . Recall that $\Pi_T^\Theta = \sum_{t=1}^T \pi_t^\Theta$ and $S_T = \sum_{t=1}^T \lambda_t(w_t - w_t^0)$. It is sufficient to prove a bound of the form

$$-B \leq \mathbb{E}_{c, a^*} [\Pi_N^\Theta - \alpha S_N], \quad (48)$$

where B is a number independent of N and \mathcal{P} . We show that this is indeed the case. By construction, we have that

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

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 $\Delta_t \equiv w_t - w_t^0 - \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t]$ and $\chi_T = \bar{d}^2 + \sum_{t=1}^T d_t^2$. Note that under strategy (c, a^*) , Assumption OA 3 implies that if $d_t > 0$, then $\mathbb{E}_{c, a^*}(w_t - w_t^0 | \mathcal{F}_t) > \xi$. Hence $\sum_{t=1}^T \mathbb{E}_{c, a^*}(w_t - w_t^0 | \mathcal{F}_t) \geq \xi(\chi_T / \bar{d}^2 - 1)$. By (30), for any T ,

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

An argument similar to that used in the proof of Lemma 2 yields that $\sum_{T=1}^{+\infty} \text{Prob} \left(\sum_{t=1}^T \Delta_t < -\frac{\xi}{\bar{d}^2} \chi_T + \xi + \Theta_T + w\sqrt{\chi_T} \right)$ is bounded above by a constant. \square

OA 2 Simulations

This appendix provides simulations illustrating key properties of calibrated contracts, and contrasts them with properties of high-water mark contracts that do not adjust the share of resources $(\lambda_t)_{t \geq 1}$ invested by the agent as a function of past history. Throughout, time periods are referred to as days, and the returns processes' ratio of standard-deviation to drift (which matters for the speed at which incentives are approximated) is kept large (comparable to that of stock market returns). This makes the calibration exercise realistically difficult.

Incentive alignment. This first simulation illustrates Lemma 1: calibrated contracts approximately align performance and rewards to the agent. In this simulation 1000 paths for returns process $(w_t - w_t^0)_{t \geq 1}$ are sampled from a random walk with per-period standard deviation $\sigma = 3$, and a stochastic drift $(\nu_t)_{t \geq 1}$ following Markov chain:

$$\nu_{t+1} = \begin{cases} \nu_t & \text{with prob. 98\%} \\ \sim \mathcal{N}(\mu_\nu = 0.05, \sigma_\nu = .2) & \text{with prob. 2\%.} \end{cases}$$

Example of sample paths are illustrated in Figure 1. Note that the process generating these paths need not be the process for returns at equilibrium. Rather, it is meant to generate enough variety in sample paths to illustrate the incentive alignment properties of calibrated contracts on a sample path by sample path basis.




Figure 1: Example of sample paths for returns process.

Figure 2 illustrates the incentive alignment properties of calibrated and high-water mark

contracts. It plots realized payments to the agent against total surplus generated for 1000 sample path realizations. In each case the dashed line corresponds to the benchmark linear contract with reward parameter parameter $\alpha = 15\%$. Both calibrated and high-water mark contracts achieve significant alignment between rewards and performance when performance is high. This is because realizations for which final performance is high are also realizations for which aggregate performance is on average increasing. In contrast, this continues to hold for calibrated contracts even if the path of returns has significant downward deviations, but not for high-water mark contracts.




Figure 2: Incentive alignment for (a) calibrated and (b) high-water mark contracts.

Correct performance and correct ex-ante investment. As Theorem 2 shows, the fact that calibrated contracts approximately align performance and rewards implies that asymptotically, they also induce performance close to that of benchmark linear contracts. In particular, the agent should be making similar returns-generating investments. A caveat to this result is that for finite time horizons incentive alignment is only approximate, and approximation errors can distort investment behavior. As a result, to guarantee performance close to that of a benchmark contract with reward rate α_0 , calibrated contracts must use a reward rate $\alpha > \alpha_0$, that can approach α_0 as the time horizon becomes large.

In this simulation, the agent can make a lumpy initial investment in information at a fixed cost. If he makes the investment, the surplus maximizing investment strategy under that information generates a process for returns $(w_t - w_t^0)_{t \geq 1}$ that is a random walk with




Figure 3: Incentives to invest under calibrated and high-water mark contracts, for investment horizons of (a) 1000 and (b) 2000 days.

drift .05 and standard deviation 1. If the agent doesn't make the investment, limited-liability implies that his optimal strategy is to pick allocations that are different from the optimal allocation under public information: choosing the default allocation would ensure rewards equal to 0. This results in a process for returns that is a random walk with drift $-.01$ and standard deviation 2.⁴ The difference in expected rewards under the two returns processes corresponds to the agent's incentives to invest. Figure 3 contrasts the need for more highly powered incentives as the time horizon goes from 1000 periods (and a fixed cost of 7.5) to 2000 periods (and a proportional fixed cost of 15): in both cases a linear benchmark contract with reward rate $\alpha_0 = 15\%$ is sufficient to induce investment; for 1000 periods, a calibrated contract with reward rate $\alpha = 30\%$ is needed to induce investment; for 2000 periods a calibrated contract with reward rate $\alpha = 25\%$ will induce investment. Note that in this example, the agent never loses access to valuable private information and high-watermark contracts also provide adequate ex ante incentives to invest.

⁴Under limited-liability, it is always in the interest of an uninformed agent to choose suboptimal asset allocations in order to get some rewards. The precise equilibrium returns processes used here can be micro-founded by the following environment: there are three assets numbered 1, 2 and 3. Asset 1 offers a risk free return equal to 0. Assets 2 and 3 have the following correlation structure: each period one of the two assets is "good" with probability .5 while the other asset is "bad", and vice versa. Returns for the good asset have mean .05 and standard deviation 1. Returns for the bad asset have mean $-.07$ and standard deviation $\sqrt{7}$. Investing in information allows the agent to perfectly predict which asset is good and which asset is bad. Under public information the optimal allocation is to pick asset 1, but an uninformed agent will pick either asset 2 or asset 3, since picking asset 1 ensures 0 rewards. An informed agent would pick the good asset in every period.

Damage control upon large downward deviations. One key difference between calibrated and high-water mark contracts is that under high-water mark contracts, the agent’s reward and the agent’s performance cease to be tightly linked if there is a large downward deviation in performance. The reason for this is that performance can decrease arbitrarily while aggregate rewards must be weakly increasing. This can have a large effect on equilibrium performance since agents that become uninformed after some period will choose suboptimal strategies in order to get rewarded through luck.

In contrast calibrated contracts limit large downward deviations by controlling the share of resources $(\lambda_t)_{t \geq 1}$ that the agent manages in each period. This is a form of damage control that allows the agent’s aggregate reward to remain linked to his aggregate performance. Figure 4(a) illustrates an instance of such damage control: although potential performance $\sum_{t=1}^T w_t - w_t^0$ falls by approximately 100 between periods 400 and 1000, the dynamically scaled performance $\sum_{t=1}^T \lambda_t(w_t - w_t^0)$ decreases by only 30.




Figure 4: Damage control upon large downward deviation (a) and resource allocation $(\lambda_t)_{t \geq 1}$ (b).

This damage control is achieved by sharply reducing the fraction of resources $(\lambda_t)_{t \geq 1}$ managed by the agent (Figure 4(b)).

Continuation behavior and performance after a large downward deviation. An important property of calibrated contracts emphasized in Theorem 2 is that, unlike high-water mark contracts, their continuation performance does not depend significantly on his

tory: from the perspective of any history, they induce performance approximately as good as the performance of history-independent, weakly renegotiation proof, benchmark linear contracts. In contrast, under a high-water mark contract, agents that have just experienced a large downward deviation may not find it worthwhile to continue investing in information acquisition since they have to compensate for previous large downward deviations before they get rewarded again.

The simulation takes as given the history of raw returns $(w_t - w_t^0)_{t \geq 1}$ from period 1 to period 1000—it is the one corresponding to Figure 4(a)—and considers incentives to invest in further information that is valuable over the next 1000 periods. The contingent investment problem in period 1000 is similar to that presented in Figure 3: the agent must expend a fixed cost of 3 to acquire further information; if the agent acquires information, the surplus maximizing allocation yields a returns process following a random walk with i.i.d. increments of mean .05 and standard deviation 1; if the agent does not acquire information, the agent no longer has valuable information, and his optimal strategy is to choose suboptimal allocations that yield a returns process following a random walk with i.i.d. increments of mean $-.01$ and standard deviation 2.

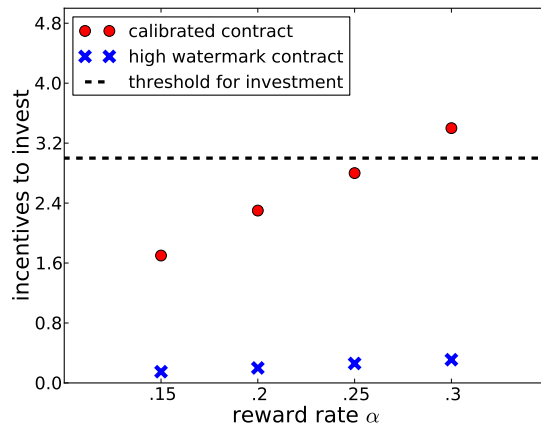


Figure 5: Incentives to invest conditional on large downward deviation.

As Figure 5 highlights, from the perspective of period 1000, calibrated contracts still provide incentives for contingent investment whereas high-water mark contracts do not. Indeed, as Figure 6(c) illustrates, under a high-water mark contract, it is very unlikely—even with additional investment in information—that the agent can compensate for past losses and get significant continuation rewards. In contrast, as Figures 6(a) and 6(d) show,




Figure 6: Performance and conditional payments under calibrated and high-watermark contracts.

calibrated contract manage to control resource allocation $(\lambda_t)_{t \geq 1}$ in a way that limits large downsides, but still capture large upsides. This requires process $(\lambda_t)_{t \geq 1}$ to reduce exposure to the agent's performance upon large downward deviations, and restore exposure to the agent's performance when the agent starts generating positive returns again (Figure 6(b)). As a result, the agent can get significant continuation rewards even conditional on poor past performance.