Calibrated Incentive Contracts, Supplementary Material: Extensions and Simulations

Sylvain Chassang*†

Princeton University

April 5, 2013

Abstract

This supplementary material to "Calibrated Incentive Contracts" provides several extensions as well as simulations illustrating key properties of calibrated contracts. In particular, extensions show how to allow for time discounting, describe a broader class of high-liability contracts that can be successfully calibrated under limited liability, and explore contract performance when the agent isn't fully rational.

OA 1 Extensions

For clarity, references internal to this supplementary material start with the prefix OA.

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

*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. I thank the co-editor and several anonymous referees for substantial help in revising the paper. I am grateful to 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 for many useful conversations. The paper greatly benefited from 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. SeHyoun Ahn and Juan Ortner provided excellent research assistance.

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

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

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

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

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

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

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

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

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

$$-\frac{w\overline{d}}{N_\delta} \leq \frac{1}{N_\delta} \sum_{t=1}^{N} \delta^{t-1}[\pi_t - \alpha(w_t - w_t^0)] \leq \frac{w\overline{d}}{\sqrt{N_\delta}}. \tag{2}$$

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 $||\rho_T||^2 \leq \delta^{2T} w\overline{d}$, which implies that

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

This implies the right-hand sides of (1) and (2). The left-hand side of (2) follows from a proof identical to that of the left-hand side of (13). □

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,*

$$r_{\lambda,\pi} \geq (1-\eta) r_{\alpha_0} - \frac{m}{\sqrt{N_\delta}} \tag{3}$$

$$r_{\lambda,\pi} \geq (1-\alpha) \sup_{\hat{c} \in [0,+\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} - \frac{3\overline{d}}{\sqrt{N_\delta}} \right). \tag{4}$$

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. □

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

$$\forall t \geq 1, \quad 0 \leq \pi_t \leq \overline{\pi}_t, \tag{4''}$$

such that for all t, $w_t \leq \overline{\pi}_t$, 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) $\pi_t^0 \leq \overline{\pi}_t$ and there exists $\overline{\pi} > 0$ independent of N such that, $\sup |\pi_t^0| \leq \overline{\pi}$.

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 $\kappa_G \kappa_\phi w_t \leq \overline{\pi}_t$. For instance, if for all t, $\overline{\pi}_t = w_t$, this includes contracts such that the agent gets paid a positive reward only when returns are above a threshold, i.e contracts such that

$$\Pi_N^0 = \begin{cases} \alpha \sum_{t=1}^N w_t - w_t^0 & \text{if } \sum_{t=1}^N w_t - w_t^0 < 0 \\ 0 & \text{if } \sum_{t=1}^N w_t - w_t^0 \in [0, \underline{W}] \\ \alpha\left(\left[\sum_{t=1}^N w_t - w_t^0\right] - \underline{W}\right) & \text{if } \sum_{t=1}^N w_t - w_t^0 > \underline{W}. \end{cases} \tag{5}$$

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, $|\pi_t^\eta| \leq \frac{1}{1+\eta}\overline{\pi} + \frac{\eta}{1+\eta}\overline{wd} \equiv \overline{\pi}^\eta$.

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:

$$\mathcal{R}_{1,T} = \sum_{t=1}^{T} \pi_t - \pi_t^\eta(\lambda_t) \qquad \text{(no excess rewards)} \tag{6}$$

$$\mathcal{R}_{2,T} = \max_{T' \leq T} \sum_{t=T'}^{T} \pi_t^\eta - \pi_t^\eta(\lambda_t) \qquad \text{(no foregone performance).} \tag{7}$$

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

$$\lambda_{T+1} = \frac{\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \mathcal{R}_{2,T}^+} \qquad \text{and} \qquad \pi_{T+1} = \begin{cases} \left[\pi_{T+1}^\eta\right]^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{8}$$

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*

$$\forall h_T, \qquad r_{\lambda,\pi}|h_T \geq \frac{1}{1+\eta} r_{\pi^0}|h_T - m\frac{1}{\sqrt{N}} \tag{9}$$

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

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

$$-\overline{\pi}^\eta \quad \leq \quad \sum_{t=1}^{T} \pi_t - \pi_t^\eta(\lambda_t) \quad \leq \quad \overline{\pi}^\eta \sqrt{T} \tag{10}$$

$$-\overline{\pi}^\eta \sqrt{T} \quad \leq \quad \max_{T' \leq T} \sum_{t=T'}^{T} \pi_t^\eta - \pi_t^\eta(\lambda_t) \quad \leq \quad \overline{\pi}^\eta \sqrt{T}. \tag{11}$$

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

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

Hence, the contract $(\lambda_t, \pi_t)_{t\geq 1}$ defined by (8) ensures that for all realizations of $\mathbf{r}_{T+1}$, $\left\langle \mathcal{R}_T^+, \rho_{T+1} \right\rangle = 0$.

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 $\mathcal{R}_{2,T} > 0$.

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

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

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

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 $|\pi_t^\eta| \leq \overline{\pi}^\eta$, this implies inequality (11) and the right-hand side of (10). The left-hand side of (10) 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 (9) starting from initial history h_0. Let $(\widehat{c}, \widehat{a})$ 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 $(\widehat{c}, \widehat{a})$ under (λ, π), we obtain that

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

By (10) this implies that

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

By (11) we obtain

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

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

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

Thus, using (10) 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

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

Inequality (9) at h_0 follows from normalizing by $\frac{1}{wN}$.

Inequality (9) continues to hold conditional on any history because the incentive bounds provided by Lemma OA 3 hold starting from any interim period T. □

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 $\overline{l} > \underline{l} > 0$ such that for every history,*

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

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 1. *Truthtelling, 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 $0 \leq \pi_t \leq \overline{\pi} = \alpha \log \left(\overline{l}/\underline{l}\right)$, Theorem OA 2 applies, and the contract (λ, π) derived from $(\pi_t^{ls})_{t\geq 1}$ according to (8) satisfies performance bound (9), 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 $\{\frac{1}{100}, \frac{2}{100}, \cdots, 1\}$ in every period t, and a bad asset 1 with i.i.d. returns $r_{1,t}$ uniformly distributed over$\{-\frac{99}{100}, -\frac{98}{100}, \cdots, 0\}$. 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 it considerably reduces uncertainty. However this information is of no value to the principal since the good asset always dominates the bad asset.

[1] Fact 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.

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 $a^{\triangle} = (a_t^{\triangle})_{t \geq 1}$.[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 $a^{\triangle}$ during the accident. Strategy $a^{\triangle}$ 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 $\mathbb{E}_{a^{\triangle}}(w_t - w_t^0)$ can be negative in an accident period, accidents can undo all the profit generated by the well incentivized

[2]The 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.

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 $a^{\triangle}$, the exogenous allocation during accident states is

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

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

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

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

$$-\alpha w\overline{d} \leq \sum_{t=1}^{N} \pi_t^{\triangle} - \alpha\lambda_t(w_t^{\triangle} - w_t^0) \leq \alpha w\overline{d}\sqrt{N}. \tag{12}$$

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

$$\left(\sum_{t=1}^{N} w_t^{\triangle\triangle} - w_t^0\right) - 4w\overline{d}\sqrt{N} \leq \sum_{t=1}^{N} \lambda_t(w_t^{\triangle} - w_t^0). \tag{13}$$

We have that

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

Applying inequality (14), we obtain

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

By (12), it follows that

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

Subtracting these two inequalities yields that,

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

Since flow rewards $\pi_t^{\triangle}$ are weakly positive, this implies that for any realization of returns,

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

This implies (13). Given (12) and (13), Theorem OA 3 follows by applying Lemma A.1. □

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 $j \in \{1, \cdots, J\}$, 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 $(I_t^0, I_c^j(c_{j,t}))_{j \in \{1,\cdots,J\}}$ are arbitrary random variables from an underlying measurable state space (Ω, σ) to a measurable 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

$$\forall j \in \{1, \cdots, J\}, \quad \pi_{j,t} = \alpha(w_{j,t} - w_{j-1,t}). \tag{14}$$

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 $(c, a) = (c_j, a_j)_{j\in\{1,\cdots,J\}}$ and $\mathcal{C} \times \mathcal{A} = \prod_{j\in\{1,\cdots,J\}} \mathcal{C}_j \times \mathcal{A}_j$ 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

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

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

[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 $j > j'$.

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

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

Summing over j, this implies that

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

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

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

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

□

Calibrated contracts. The high-liability multi-agent contract described in (14) can be calibrated using the methods of Section 4. The main difference is that there is now a vector $\lambda_t = (\lambda_{j,t})_{j\in\{1,\cdots,J\}} \in [0,1]^J$ 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

$$\mathcal{R}_{j,T}^1 = \sum_{t=1}^{T} \pi_{j,t} - \alpha(w_{j,t}^\lambda - w_{j-1,t}^\lambda) \qquad \text{(correct rewards)} \qquad (15)$$

$$\mathcal{R}_{j,T}^2 = \max_{T'\leq T}\sum_{t=T'}^{T} w_{j,t} - w_{j,t}^\lambda \qquad \text{(no foregone returns).} \qquad (16)$$

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 $(\mathcal{R}^1_{j,T}, \mathcal{R}^2_{j,T})_{j\in\{1,\cdots,J\}}$ can be kept small by choosing contract parameters $(\lambda_j, \pi_j)_{j\in\{1,\cdots,J\}}$ according to,

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

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 $(\lambda, \pi) = (\pi_j, \lambda_j)_{j\in\{1,\cdots,J\}}$ of parameter α satisfies*

$$\forall h_T, \quad r_{\lambda,\pi}|h_T \;\geq\; (1-\eta) r_{\alpha_0}|h_T - \frac{m}{\sqrt{N}}. \tag{17}$$

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

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 $\mathbb{E}_{c,a^*}[w_t - w^0_t|\mathcal{F}_t] > 0$, then $\mathbb{E}_{c,a^*}[w_t - w^0_t|\mathcal{F}_t] > \xi$.*

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

$$-B \leq \mathbb{E}_{c,a^*}\left[\Pi_N^\Theta - \alpha S_N\right], \tag{18}$$

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

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

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

Let $\Delta_t \equiv w_t - w_t^0 - \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t]$ and $\chi_T = \overline{d}^2 + \sum_{t=1}^T d_t^2$. 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 $\sum_{t=1}^T \mathbb{E}_{c,a^*}(w_t - w_t^0 | \mathcal{F}_t) \geq \xi(\chi_T / \overline{d}^2 - 1)$. By (30), for any T,

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

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

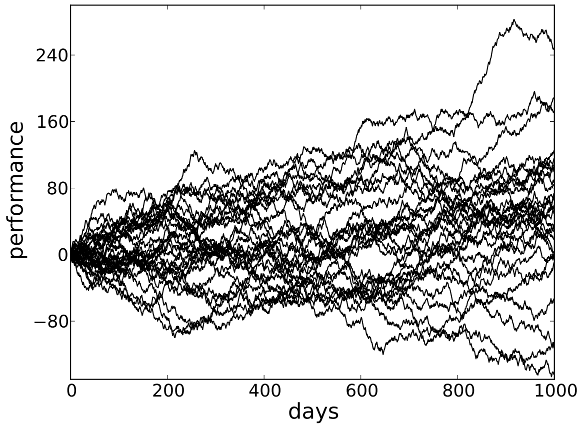

Figure 1: Example of sample paths for returns process.

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.

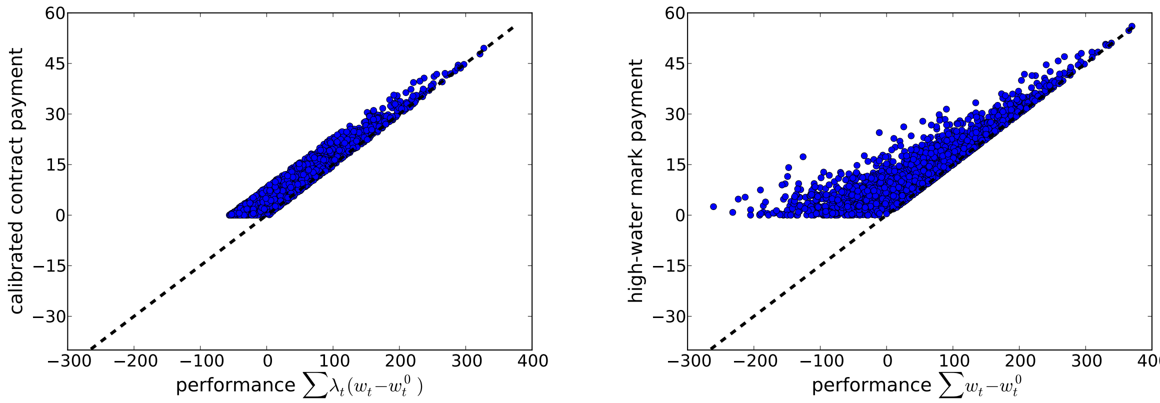

(a) (b)

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.

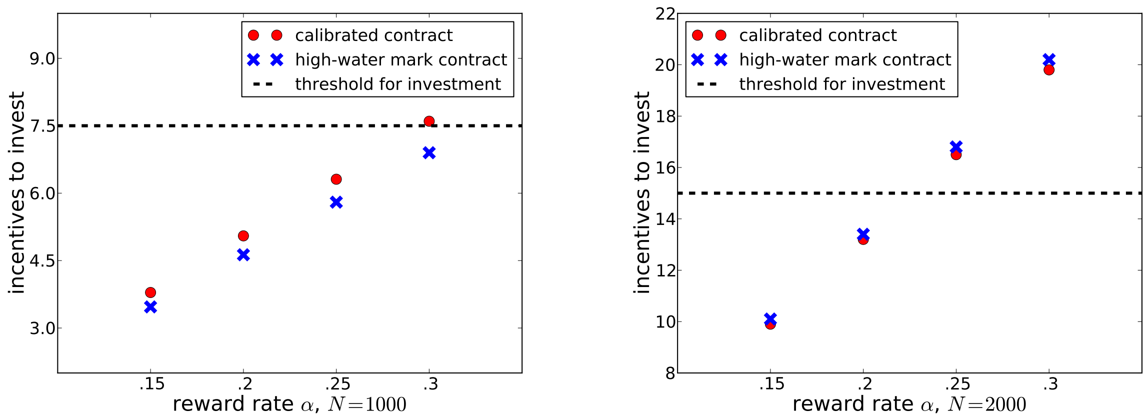

(a) (b)

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

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 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.

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

[4]Under 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.

scaled performance $\sum_{t=1}^{T} \lambda_t(w_t - w_t^0)$ decreases by only 30.

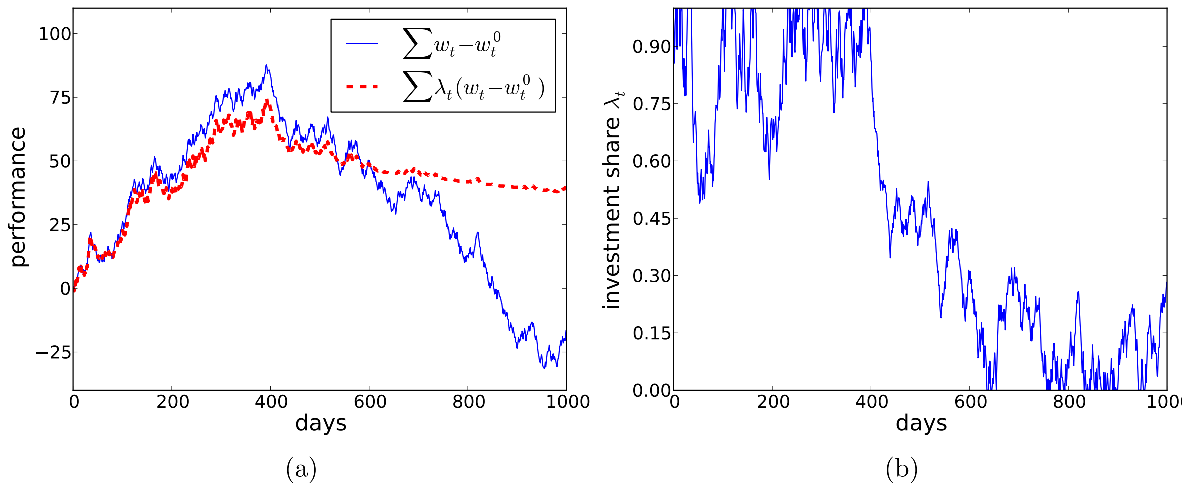

Figure 4: Damage control upon large downward deviation (a) and resource allocation $(\lambda_t)_{t\geq 1}$ (b).

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 history: 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.

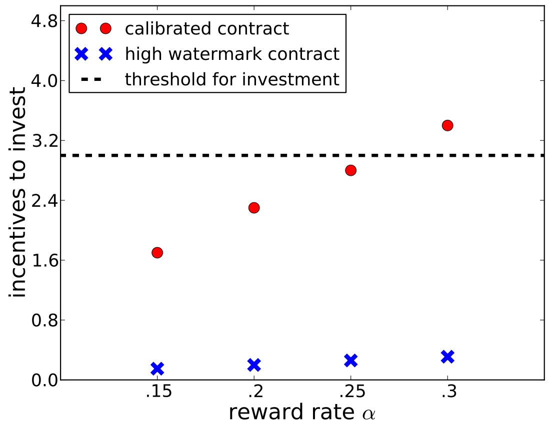

Figure 5: Incentives to invest conditional on large downward deviation.

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, 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

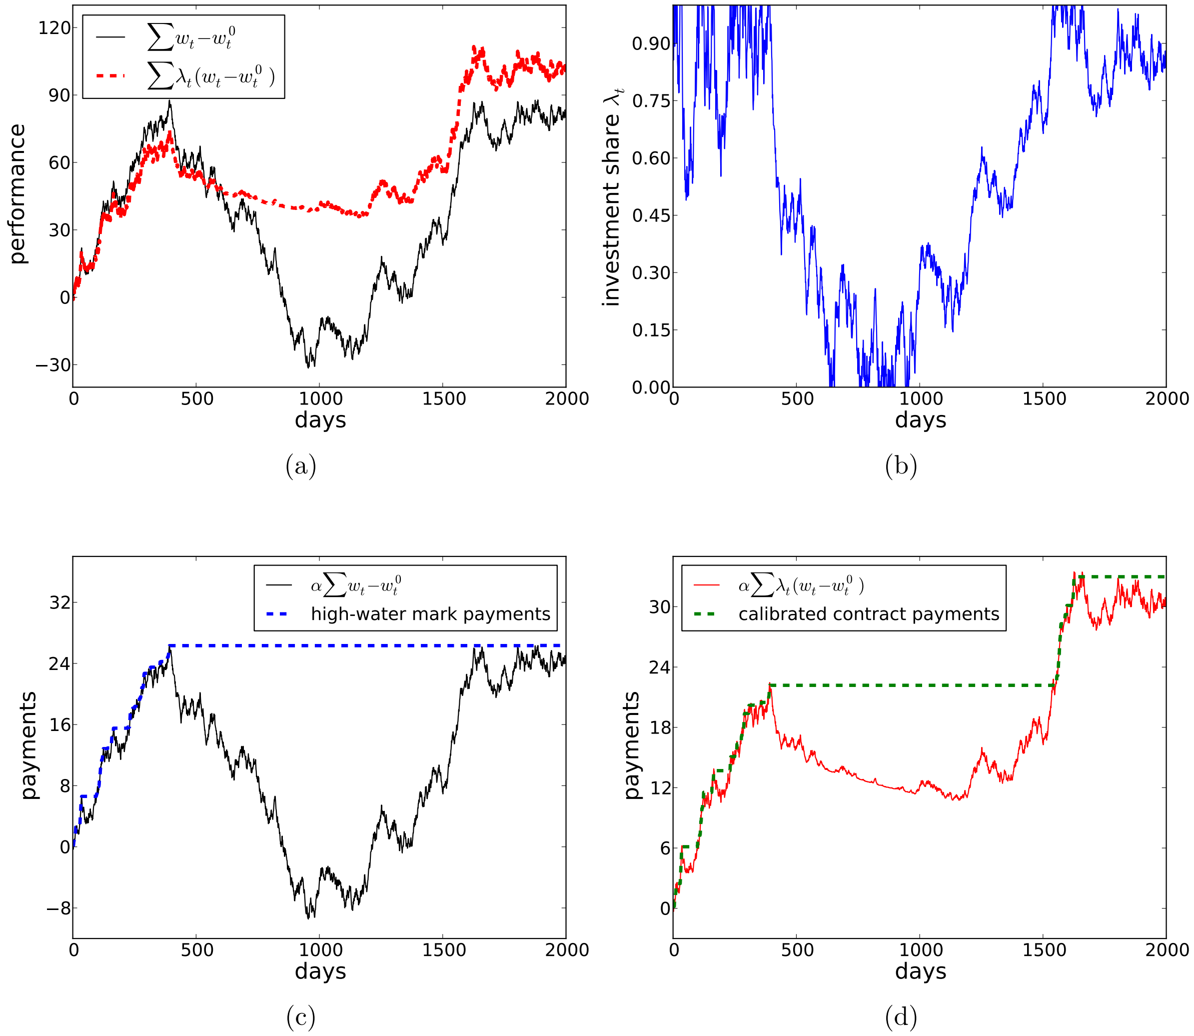

Figure 6: Performance and conditional payments under calibrated and high-watermark contracts.

performance.

References

Chassang, S. (2011): "Calibrated Incentive Contracts," *Princeton University Economic Theory Center Working Paper No. 013-2011*.

Eliaz, K. (2002): "Fault Tolerant Implementation," *Review of Economic Studies*, 69, 589–610.