Rewarding Trading Skills Without Inducing Gambling

Igor Makarov* and Guillaume Plantin†

January 2009

VERY PRELIMINARY AND INCOMPLETE

Abstract

This paper develops a model in which traders may secretly take fair bets in complete markets in order to temporarily manipulate their reputation and attract more funds. It then solves for contracts that mitigate this friction. Such contracts must ensure that the present value of the trader's future payoffs is not too convex in her reputation. If the trader can commit to a long-term contract, she can achieve this by committing to leave a fraction of her future excess returns to investors if she turns out to be highly skilled. If the trader cannot commit to such behavior, then investors optimally commit to overcompensate the trader if she turns out to be unskilled. In both cases, the net excess returns earned by competitive, uninformed investors exhibit persistence.

*London Business School. E-mail: imakarov@london.edu.

†Toulouse School of Economics, London Business School, and CEPR. E-mail: guillaume.plantin@tse-fr.eu.

> "Hedge funds are investment pools that are relatively unconstrained in what they do. They are relatively unregulated, charge very high fees, will not necessarily give you your money back when you want it, and will generally not tell you what they do. They are supposed to make money all the time, and when they fail at this, their investors redeem and go to someone else who has recently been making money. Every three or four years they deliver a one-in-a-hundred year flood."
>
> -Cliff Asness, Journal of Portfolio Management, 2004.

Introduction

The last thirty years have witnessed two important evolutions in financial markets. First, a rapid pace of financial innovation has made it possible to slice and combine a large variety of risks by trading a rich set of financial instruments. Second, the management of increasing amounts of capital has been delegated to agents which have very few trading restrictions, and have to disclose little information about their trading strategies. As a result, investors seeking to allocate their funds across different investment vehicles have to rely increasingly on the history of realized returns to assess future performance.

Whether superior returns are due to skills or luck is difficult to assess in practice, however. In particular, it is difficult to statistically disentangle finite samples of excess returns that stem from true arbitrage opportunities from those that result from exposures to risk factors with rare negative realizations. This creates room for a particular type of agency problem. Traders running out of genuine arbitrage opportunities may find it tempting to take inefficient gambles that can temporarily improve their reputation. Sophisticated financial markets offer many opportunities to "pick up nickels in front of a steam roller", or generate small positive excess returns, and rarely experience very large losses. Examples include selling deep-out-of-the-money options, or selling protection against the default of very solvent entities.

The goal of this paper is to develop a new framework for the study of this risk-shifting problem that separates investors from sophisticated traders operating in financial markets. This is a timely topic. The spectacular demises of LTCM, Amaranth, and several other large hedge funds have left many investors with nothing, and suggested that this type of risk-shifting was at play. Moreover, following the recent severe financial crisis, many observers argue that perverse incentives lead to excessive and inefficient risk-taking by financial institutions. In particular, the impact of traders' compensation on their risk-taking incentives is fiercely debated. This question was at the top of the agenda of the September 2009 Pittsburgh G20 summit.

We believe that standard moral hazard models are not appropriate for the type of agency problems that separate trading skills from capital in modern financial markets. Standard moral hazard models typically assume that an agent can secretly "shirk" instead of working hard in the principal's interest. Formally, the agent can derive an exogenous private benefit from secret actions that reduce her output in the sense of first-order stochastic dominance. The moral hazard problem that we study differs from this standard one along two dimensions. First, the agent in our model can secretly add any mean-preserving noise to her realized returns. This captures the large set of trading opportunities available to a trader in sophisticated markets better than first-order stochastic dominance. Second, tampering the distribution of the output generates endogenous pecuniary benefits induced by a temporarily enhanced reputation. In other words, the agency problem that we study is the following. Traders can collect temporary high returns from exposures on skewed risk factors, and then pretend that these returns are "free lunches" that reflect a permanent ability to detect genuine arbitrage opportunities.

Our model builds upon the frictionless benchmark of Berk and Green (2004), who study career concerns in delegated fund management. As in their model, a trader and investors symmetrically discover the trader's skills by observing her realized returns. Expected excess rates of returns increase with the trader's skills, but decrease with the fund size. The fund size that optimally trades off scale and unit return increases with respect to skills. Competitive investors supply funds to the trader until they earn a zero net (after fees) expected return. The trader sets fees that enable her to reach the optimal fund size and extract all the surplus that she generates. Thus, competition among investors implies that both fund flows and trader compensation strongly depend on the trading record. We adopt Berk and Green's framework as our frictionless benchmark because it is a parsimonious model that matches well the empirical relationship between funds realized returns and inflows/outflows.

We add a gambling friction to this model: We assume that the trader can secretly add an orthogonal, arbitrarily distributed white noise to the profit she realizes at each period. We show that the trader finds it optimal to gamble when i) a highly skilled trader can generate significant gross (before fees) excess returns, ii) the volatility of returns is low, and iii) the trader's skills are scalable - that is, unit returns are not too sensitive to fund size. When these conditions are met, gambling pays off because a slightly positive excess return has a large impact on investors' beliefs about the trader's ability to generate excess returns, and thus on future fund sizes and profits. For a calibration similar to that by Berk and Green, we find that the equilibrium with short-term contracts that they study will always feature risk-shifting. In particular, "fallen stars" traders that follow high expectations with a stream of disappointing returns are particularly prone to gambling, consistent with anecdotal evidence on rogue traders.

We then study whether long-term arrangements can prevent gambling. Such long-term arrangements are risk-shifting proof if they reduce the sensitivity of a trader's continuation utility to her skills. There are two ways to ensure that this is the case: by overcompensating a low-skilled trader and/or undercompensating a highly skilled one. Which solution arises in equilibrium depends on the commitment power of each party. If the trader can commit to a long-term arrangement but investors cannot, for instance because they have strong liquidity needs, then the trader must commit to give up some excess returns if she turns out to be highly skilled. This ensures that she will not be tempted to gamble in the opposite case in which she disappoints. To implement this arrangement, the trader must contract not only on fees but also on her fund's size. She must turn down funds when successful. Conversely, if the trader cannot commit and is free to recontract with new investors at any date, then investors must commit to overpay a trader who does not meet expectations. Whether the trader or investors can commit, our model always predicts that this risk-shifting friction implies persistence in the net (after fees) excess rates of returns left to investors. Thus, the greater opaqueness and complexity of hedge funds strategies may explain why their performance seems more persistent than those of long-only mutual funds (see, for example, Jagannathan et al. (2006) and Kosowski et. al. (2007) for empirical evidence). Glode and Green (2009) and Hochberg et al. (2008) explain persistence in private equity returns with informational rents and/or bargaining power accruing to first-round investors in subsequent rounds. Our explanation differs in that investors in our environment are symmetrically informed and competitive throughout the entire timeline.

Our paper bridges two strands of literature: the literature on risk-shifting in complete markets and the literature on career concerns. A large finance literature studies the impact of (exogenous) nonconcave objective functions on the incentives of fund managers who have access to dynamically complete markets[1]. Like this literature, we seek to identify the risk-taking strategies that optimally respond to nonconcave objectives. Within a simpler framework of risk-neutral agents, we extend this line of research in two directions. First, nonconcavities in the trader's objective arise endogenously in our model from the trader's career concerns. Second, we address the question

[1]Contributions include Basak, Pavlova, and Shapiro (2007), Carpenter (2000), and Cuoco and Kaniel (2007).

of optimal contracting.

Our model of career concerns with gambling delivers different predictions from the well-known career concerns model studied by Holmstrom (1999). In Holmstrom's model, the agent has high incentives to "shirk" once her reputation - either good or bad - is made. This is not the case in our model: agents who acquire a bad reputation are the most prone to gambling. This happens because the incentive to gamble is not determined by the local properties of the objective function, but rather by its global behavior.

The paper is organized as follows. Section 1 studies gambling incentives in a static framework. We offer a simple and tractable characterization of the terminal payoff functions that lead a trader to gamble inefficiently. We then study the interplay of this risk-shifting friction with traders' reputational concerns. Section 2 introduces risk-shifting in the dynamic model of delegated asset management set forth by Berk and Green (2004). Section 3 analyzes contractual arrangements that mitigate the gambling problem.

1 Risk Shifting: A Simple Characterization

This section studies the following problem. An agent seeks to maximize the expected value of an arbitrary continuous function of her terminal consumption. The agent can choose any distribution over (positive) final consumption subject to preserving its mean. Formally, let M denote the set of Borelian probability measures over $[0, +\infty)$. For each continuous function $U : [0, +\infty) \to \mathbb{R}$, we define $P(U)$ as

$$\begin{aligned} P(U) &\equiv \sup_{\mu \in M} \int_0^\infty U(R) d\mu(R) \\ s.t. \quad \int_0^\infty R d\mu(R) &= 1. \end{aligned} \tag{1.1}$$

We deem an objective function U to be conducive to risk-shifting if and only if $P(U) > U(1)$. Obviously, from Jensen's inequality, U is not conducive to risk-shifting if it is concave.

Aumann and Perles (1965) have studied a very similar class of problems, and have established conditions under which the problem associated with a given objective function U has the same solution as the one associated with the concavification of U.[2] This section shows that the dual approach generates a simple and practical determination of $P(U)$.

[2]The main difference is that they optimize over functions instead of measures.

Let $P^*(U)$ denote the solution of the dual problem. By definition,

$$\begin{aligned} P^*(U) &\equiv \inf_{(z_1,\, z_2)\in\mathbb{R}^2} z_1 + z_2 \\ s.t. \quad \forall y &\geq 0, z_1 + y z_2 \geq U(y). \end{aligned} \tag{1.2}$$

The dual problem seeks among all the straight lines that are above the graph of U the one that takes the minimal value in 1. The next proposition shows a condition under which these primal and dual problems have the same solution. In the language of optimization, there is no duality gap.

Proposition 1 *If*

$$\lim_{y\to+\infty} \frac{U(y)}{y} = 0, \tag{1.3}$$

then

$$P(U) = P^*(U).$$

Proof. See the Appendix.■

Proposition 1 offers an elementary and practical way to interpret the concavification result of Aumann and Perles (1965) and to apply it to our problem. It also yields a simple graphical characterization of the solution. Consider the elementary example in which

$$U(R) = u(\max(R - D, 0)),$$

where u is increasing, concave, $u(0) = 0$, and $D > 0$. This corresponds to the situation of an agent who is protected by limited liability, and faces a noncontingent liability equal to D at its consumption date. As is well-know, excessive leverage creates risk-shifting incentives. Figure 1 shows that if leverage is sufficiently low, then there is no risk-shifting. Conversely, as D increases, risk-shifting occurs. The optimal risk profile consists in a binary payoff either equal to 0 or strictly larger than 1 in this case. The solution to P^* is given by the value at 1 of the tangent to the objective that goes through the origin. It is graphically obvious that this tangent is the straight line above the graph of U that takes the smallest value in 1.

The following proposition shows that the generic solution to $P(U)$ also has this simple structure - either no gambling or a binary gamble.

Proposition 2 *Assume U satisfies (1.3), and $\lim_{+\infty} U = +\infty$. If there is risk-shifting, $P(U)$ can be attained with a binary payoff.*

Proof. See the Appendix.■

Thus, the solution to the general problem (1.1) boils down to finding an optimal binary gamble. Intuitively, as in the elementary example above, if the straight line corresponding to the solution of $P^*(U)$ is tangent to U in 1, then there is no risk-shifting. If this is not the case, then this straight line must have at least one intersection with the graph of U on the left of 1, and at least one on the right. In this case, a binary gamble that pays off the abscissae of two such intersections solves $P(U)$. In the particular case in which the intersection on the left of 1 is in 0, then the optimal gamble consists in losing everything with some probability or earn a superior return.

This general dual characterization is novel to the best of our knowledge. Propositions 1 and 2 have a number of interesting applications to static contracting problems. For instance, one could apply them to study the impact of secret gambles in environments in which the agent would optimally receive a convex payoff absent this generalized risk-shifting problem (e.g., in the costly state-verification problem of Townsend (1979)). For instance, Hellwig (2007) studies a static contracting problem which combines standard moral hazard and risk shifting, where risk shifting is restricted to binary bets. Our results could be useful to study the robustness of his findings to more general specifications. We leave these applications for future research. The remainder of the paper studies instead the impact of risk-shifting when nonconcavities in the agent's preferences stem from career concerns. As explained in the introduction, our focus on a model in which career concerns play a central role is motivated by the ample empirical evidence that investors chase past performances when delegating asset management.

2 Risk Shifting and Career Concerns

Subsection 2.1 develops a model of frictionless delegated asset management that has the same ingredients as the model of Berk and Green (2004). Subsection 2.2 introduces a risk-shifting friction in this environment.

2.1 The Berk and Green (2004) Model

Time is discrete and is indexed by $\{n\Delta t\}$, where $n \in \mathbb{N}$ and $\Delta t > 0$. There is a single consumption good which serves as the numéraire. Agents are of two types: a trader and several investors. Agents live forever, are risk-neutral, and discount future consumption at the instantaneous rate $r > 0$. The trader is protected by limited liability: She cannot have negative consumption. Investors receive a large endowment of the consumption good at each date $n\Delta t$, the trader does not. The trader has exclusive access to a storage technology. If the trader stores q consumption units at date t using her technology, she generates $\rho_{t+\Delta t}$ units at date $t+\Delta t$ such that

$$\rho_{t+\Delta t} = q_t e^{\left(r+a\theta-c(q_t)-\frac{\sigma^2}{2}\right)\Delta t+\sigma(B_{t+\Delta t}-B_t)}, \tag{2.1}$$

where $(B_t)_{t\geq 0}$ is a standard Wiener process, c is a continuous, nonnegative, and nondecreasing function with $c(0)=0$, $\lim_{q\to+\infty} c(q)=+\infty$, and a and σ are strictly positive numbers. The parameter $\theta\in\{0;1\}$ measures the trader's skills. This parameter is unobservable. All agents share the common date-0 *prior* that the trader is endowed with high skills with probability $\pi_0\in(0,1)$. The parameter a is the spread that a skilled trader can generate over the risk-free rate with her first dollar. As in Berk and Green (2004), the function c captures that many arbitrage opportunities or informational rents in financial markets are not scalable.

Except for the trader's skills, which nobody observes, each action and the trader's realized returns are publicly observable at each date $n\Delta t$. Let $(\pi_{n\Delta t})_{n\geq 0}$ denote the evolution of agents' beliefs about the probability that the trader is highly skilled - her perceived skills.

This baseline model is essentially identical to Berk and Green (2004). The main modelling difference is that the distribution of skills is binomial in our setup while it is Gaussian in theirs. As in Berk and Green, the social optimum is achieved when a trader with perceived skills π at date t receives a quantity of funds $q(\pi)$ that maximizes the net expected surplus that she creates over $[t,t+\Delta t]$:

$$q(\pi) = \arg\max_q q\left(\pi e^{(a-c(q))\Delta t}+(1-\pi)e^{-c(q)\Delta t}-1\right),$$

and thus

$$\lim_{\Delta t\to 0} q(\pi) = \arg\max_q q(\pi a - c(q)). \tag{2.2}$$

As in Berk and Green (2004), we assume that investors are competitive. Thus, the trader can extract the whole surplus generated by her skills by quoting linear fees in spot markets for asset management services. To see this, assume that the trader interacts with investors as follows.

Assumption 1 *(Market for asset management services). At each date, the trader quotes a linear fee and investors competitively supply funds.*

A fee quoted at date t determines the fraction of the date-$t+\Delta t$ assets under management (before new inflows/outflows of funds) that accrues to the trader. If a trader with perceived skills π quotes a fee $w\Delta t$, competitive investors will supply q as long as their net expected rate of return is equal to r. Thus, the fund supply $q(w)$

solves

$$(1 - w\Delta t)\left(\pi e^{(r+a-c(q))\Delta t} + (1-\pi)e^{(r-c(q))\Delta t}\right) = e^{r\Delta t},$$

therefore for a given w, $\lim_{\Delta t\to 0} q(w)$ is the solution (if any) to

$$\pi a = c(q) + w. \tag{2.3}$$

The trader chooses w such that:

$$w = \arg\max_w w\Delta t \times q(w) = \arg\max_w (\pi a - c(q(w))) \times q(w).$$

Further,

$$\max_w (\pi a - c(q(w))) \times q(w) = \max_q q(\pi a - c(q)),$$

meaning that the trader extracts the entire first-best surplus. The next proposition collects these properties.

Proposition 3 *(First-Best).*

Under Assumption 1, let $w(\pi)\Delta t$, $q(\pi)$, and $v(\pi)\Delta t$ respectively denote the fee, fund size, and expected profit over $[t, t+\Delta t]$ when the trader's perceived skills at date t are π. We have

$$\begin{aligned}
\lim_{\Delta t\to 0} q(\pi) &= \arg\max_q q(\pi a - c(q)), \\
\lim_{\Delta t\to 0} w(\pi) &= \pi a - c(q(\pi)), \\
\lim_{\Delta t\to 0} v(\pi) &= \max_q q(\pi a - c(q)).
\end{aligned} \tag{2.4}$$

The trader extracts the whole social surplus created by investments in her storage technology.

Proof. See discussion above.■

The Berk and Green benchmark assumes seamless delegation of asset management. This assumption seems more natural in the case of mutual funds that take only long positions in publicly traded stocks - Berk and Green's focus - than in situations in which traders face more sophisticated instruments and fewer portfolio restrictions. Traders employed by hedge funds or banks' proprietary desks typically have a free hand at taking a very large variety of bets that outsiders do not continuously monitor. The remainder of the paper studies the impact of this risk-shifting friction in our baseline model. Before doing so, we outline two technical results. First, the following particular specification of c will play an important role in what follows.

Lemma 1 *If there exists $\alpha > 1$, $\beta > 0$ such that $c(q) = \beta q^{\frac{1}{\alpha-1}}$, then*

$$\lim_{\Delta t \to 0} v(\pi) = \beta^{2-\alpha} \frac{(\alpha-1)^{\alpha-1}}{\alpha^{\alpha}} (a\pi)^{\alpha}.$$

Proof. Straightforward computations.■

This power specification for c and thus v will generate simple intuitions on the impact of the scalability of trading skills on risk-shifting incentives. The larger α, the more scalable the trader's skills. In the hedge fund universe, global macro strategies would typically be quite scalable. On the other hand, strategies based on shareholder activism may be more difficult to spread over increasing amounts of capital.

Berk and Green specify a linear cost function c corresponding to $\alpha = 2$. In the limiting case in which $\alpha \downarrow 1$, the cost tends to 0 (respectively $+\infty$) for $q < 1$ (respectively $q > 1$), and $v(\pi)$ tends to $\beta\pi$. This linear surplus is the one used in Harris and Holmstrom (1982) model of wage dynamics.

Second, the next proposition is instrumental for our analysis. It provides a tractable representation of the expected life-time surplus of the trader as $\Delta t \to 0$.

Proposition 4 *Let v be a continuous function over $[0,1]$. Let*

$$V(\pi, \Delta t) = E_0 \left(\sum_{n=0}^{\infty} e^{-rn\Delta t} v(\pi_{n\Delta t}) \Delta t \right), \quad \pi_0 = \pi, \tag{2.5}$$

and

$$V(\pi) = \lim_{\Delta t \to 0} V(\pi, \Delta t).$$

We have

$$V(\pi) = \int_0^1 G(\pi, x) v(x) dx, \tag{2.6}$$

where

$$G(\pi, x) = \frac{2\sigma^2}{\psi a^2 x^2 (1-x)^2} \begin{cases} g(1-\pi) g(x) & \text{if} \quad 0 \leq x \leq \pi \\ g(\pi) g(1-x) & \text{if} \quad \pi \leq x \leq 1 \end{cases}, \tag{2.7}$$

and

$$\begin{aligned} \psi &= \sqrt{1 + \frac{8r\sigma^2}{a^2}}, \\ g(u) &= u^{\frac{1}{2} + \frac{1}{2}\psi} (1-u)^{\frac{1}{2} - \frac{1}{2}\psi}. \end{aligned} \tag{2.8}$$

Convergence of $V(\pi, \Delta t)$ to $V(\pi)$ when $\Delta t \to 0$ is uniform over $\pi \in (0,1)$.

Proof. See the Appendix.■

The function $G(\pi, x)$ has an intuitive interpretation: It is a weighted probability that the trader will have a skill level x if she starts with the skill level π.

Risk Shifting

Assumption 2 *(Risk shifting). In addition to the efficient storage technology described in (2.1), the trader has access to an alternative technology whose returns are perfectly scalable and independent from the returns on the efficient technology. This technology enables her to generate any arbitrary distribution over* $[0, \infty)$ *with mean* $e^{r'\Delta t}$ *for her one-period gross return, where* $r' \leq r - \frac{\sigma^2}{2}$. *How the trader chooses this distribution and allocates her funds between these two technologies is private information to her.*

Assumption 2 adds informational asymmetry to the baseline model. An interpretation of this setup is the following. The trader's efficient technology can be viewed as a risky arbitrage opportunity in a new compartment of the market which is still inefficient, and whose risk is not yet spanned by existing, fairly priced instruments. The trader can also secretly invest in a rich set of orthogonal liquid instruments, but camouflaging these trades comes at a cost $r - r'$. The investors observe returns realized at the reporting and contracting dates $n\Delta t$, at which the trader's position is marked-to-market.

Two remarks on this particular specification are in order.

Remark 1. We have been careful setting up a discrete-time model because a continuous-time environment with a continuous return process is not appropriate for this very general specification of risk-shifting. If (unrealistically) investors could continuously observe realized returns, then the discontinuities induced by risk-shifting would be detected, which would artificially restrict the set of feasible gambles.

Remark 2. That secretly gambling comes at a cost is a natural assumption. The lower bound on this cost $\frac{\sigma^2}{2}$ rules out situations in which the trader secretly invests at the risk-free rate. With $r' \geq r - \frac{\sigma^2}{2}$, the trader would be tempted to do so as her reputation deteriorates (π close to 0). We will show that a sufficiently low r' eliminates this " secret risk-reducing" problem, which seems of limited practical relevance, and is not in the scope of our paper. Importantly, we will see that the value of r' is completely immaterial for the determination of the trader's incentives to increase risk as $\Delta t \to 0$. The intuition is that risk-shifting amounts to adding an instantaneous discrete zero-mean jump to the return process in the limit of a small Δt. In other words, gambles become instantaneous fair lotteries as $\Delta t \to 0$.

We first outline a mathematical result that yields a tractable formulation of incentive-compatibility constraints in the presence of this generalized risk-shifting. Consider a situation in which investors believe that the trader does not gamble and always invests at the risk-free rate after receiving the return of the period. Assume that if investors believe that the trader is skilled with probability π at date t, then the trader collects a surplus for the $[t, t + \Delta t]$-trading round that has an expected present value equal to $v(\pi)\Delta t$ at date t, where v is a continuous function. Proposition 4 gives then the limit of the expected surplus of the trader over her life-time when $\Delta t \to 0$ if she starts out with skills π. Suppose now that the trader deviates, and gambles during her first trading round. Let $R \geq 0$ denote the return that she realizes between $\Delta t\,(1-\varepsilon)$ and Δt. Let $W(\pi, R, \Delta t)$ be the present value of her future earnings in the trader's eyes after the return R is realized at date Δt, and assuming she will no longer gamble from then on. We have

Proposition 5

$$\lim_{\Delta t \to 0} W(\pi, R, \Delta t) = W(\pi, R) = \int_0^1 G(\pi, x)\, v\left(\frac{xR^{\frac{a}{\sigma^2}}}{1 - x + xR^{\frac{a}{\sigma^2}}}\right) dx,$$

where G is defined in 2.7. Convergence is uniform over $(\pi, R) \in (0,1) \times [0, +\infty)$.

Proof. Assume the trader gambles and realizes a return R over $[0, \Delta t]$, and from then on invests in her own storage technology. Let $\left(\pi^R_{n\Delta t}\right)_{n\in\mathbb{N}}$ denote the process - *under the trader's filtration* - of her skills as perceived by investors who believe instead that she has invested in her storage technology at date 0. These investors believe that

$$R = e^{\left(r+\theta a - c(q_0) - \frac{\sigma^2}{2}\right)\Delta t + \sigma W_{\Delta t}},$$

and thus

$$\begin{aligned}
\pi^R_{\Delta t} &= \frac{\pi_0 R^{\frac{a}{\sigma^2}} e^{\frac{a}{\sigma^2}\left(\frac{\sigma^2}{2} + c(q_0) - r - \frac{a}{2}\right)\Delta t}}{1 - \pi_0 + \pi_0 R^{\frac{a}{\sigma^2}} e^{\frac{a}{\sigma^2}\left(\frac{\sigma^2}{2} + c(q_0) - r - \frac{a}{2}\right)\Delta t}}, \\
\forall n &\geq 0, \quad \pi^R_{(n+1)\Delta t} = \frac{\pi^R_{\Delta t}\frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}}}{1 - \pi^R_{\Delta t} + \pi^R_{\Delta t}\frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}}},
\end{aligned} \tag{2.9}$$

As $\Delta t \to 0$

$$\lim_{\Delta t \to 0} \pi^R_{\Delta t} = \frac{\pi_0 R^{\frac{a}{\sigma^2}}}{1 - \pi_0 + \pi_0 R^{\frac{a}{\sigma^2}}}.$$

From (A2),

$$\frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}} = \frac{\pi_{(n+1)\Delta t}\left(1-\pi_{\Delta t}\right)}{\pi_{\Delta t}\left(1-\pi_{(n+1)\Delta t}\right)}.$$

Injecting in (2.9) and taking the limit yields

$$\lim_{\Delta t \to 0,\, n\Delta t \to t} \pi^R_{n\Delta t} = \frac{\pi_t R^{\frac{a}{\sigma^2}}}{1-\pi_t+\pi_t R^{\frac{a}{\sigma^2}}}. \tag{2.10}$$

The result now follows from Proposition 4.■

In the remainder of the paper, all results will be established for Δt sufficiently small. All our results hinge on properties satisfied by the continuation value functions $V(\pi,\Delta t)$ and $W(\pi,R,\Delta t)$ for Δt sufficiently small. To ease the exposition and focus on economic intuitions, we will only establish that these properties hold for their respective limits as $\Delta t \to 0$, $V(\pi)$ and $W(\pi,R)$. That the functions $V(\pi,\Delta t)$ and $W(\pi,R,\Delta t)$ also satisfy these properties if Δt is sufficiently small is a simple consequence of the uniform convergence properties established in Propositions 4 and 5. We first exhibit conditions under which this risk-shifting friction does not upset the Berk and Green efficient equilibrium.

Proposition 6 *Let $v(x) = kx^\alpha$, where $v(\cdot)$ is the expected surplus that the trader generates per period, and $\alpha, k > 0$. If $\sigma^2 > \alpha a$ then for sufficiently small Δt the trader does not engage in risk-shifting.*

Proof.

Step 1. We first show that one-period risk-increasing gambles are not desirable in this case. We omit the proof that $W(\pi,R,\Delta t)$ is twice differentiable in R, and that the second partial derivative converges uniformly to $\frac{\partial^2 W(\pi,R)}{\partial R^2}$ over $(\pi,R)\in(0,1)\times[R_0,+\infty)$ when $\Delta t\to 0$ for all $R_0>0$. It is essentially similar to the proof of Proposition 4. We have,

$$\frac{\partial^2 W(\pi,R)}{\partial R^2} = -\int_0^1 G(\pi,x)\,\frac{\alpha a(1-x)R^{-2+\frac{\alpha a}{\sigma^2}}}{\sigma^2\left(1-x+xR^{\frac{a}{\sigma^2}}\right)^{2+\alpha}}\left((1+\frac{a}{\sigma^2})xR^{\frac{a}{\sigma^2}}+(1-x)(1-\frac{\alpha a}{\sigma^2})\right)dx,$$

Thus if $\sigma^2 > \alpha a$, $W(\pi,\cdot)$ is concave for all π and gambles are undesirable.

Step 2. It is easy to see that if one-period gambles are not desirable, then no multi-period risk-shifting strategy can be. Backward induction implies that if one-period gambles are undesirable, then so is any series of T consecutive gambles. Since this holds for all T, it must also be true for perpetual gambling strategies. This is because "bubbly" strategies such as doubling are not feasible given that maximal fund size is bounded from above.

Step 3. Finally we show that risk-reducing gambles are not desirable. If the trader invests in her storage technology, then her continuation utility process is a martingale. On the other hand, if she invests in the risk-free asset, the investors will update π in the following way:

$$d\pi = \frac{a}{\sigma}\pi(1-\pi)\left(r' - r + \frac{\sigma^2}{2} - (\pi a - c(q_t))\right)dt,$$

and the continuation utility is a surmartingale.■

From Lemma 1, the per period surplus is proportional to π^α as soon as the cost $c(q)$ is proportional to $q^{\frac{1}{\alpha-1}}$. The efficient spot market for asset management services with linear fees described in Berk and Green (2004) still holds in the presence of a risk-shifting friction if $\sigma^2 > \alpha a$. The following Proposition establishes conditions under which this is no longer true.

Proposition 7 *Suppose that there exists* $\phi_0 \geq 0$ *and* $\alpha > 0$ *such that*

$$v(x)x^{-\alpha} \underset{x\to 0}{\to} \phi_0, \tag{2.11}$$

where $v(\cdot)$ *is the expected surplus that the trader generates per period. If* $\sigma^2 < \alpha a$ *and* $r\sigma^2 > \frac{a^2}{2}\alpha(\alpha-1)$, *then for* π_0 *sufficiently small, the trader will engage in risk-shifting if she contracts with investors according to Assumption* 1 *and* Δt *is sufficiently small.*

Proof. It suffices to show that there exists a one-period gamble which makes the trader better off than never gambling for π_0 and Δt sufficiently small. Let $R > 1$, consider the following gamble:

$$\begin{cases} R & \text{Prob.} \quad 1/R \\ 0 & \text{Prob.} \quad 1 - 1/R. \end{cases}$$

Let $R = (1-\rho)^{-\sigma^2/a}$ and $\phi(x) = v(x)x^{-\alpha}$. For small Δt, the expected net gain from the above one-shot gamble over perpetual investment in the efficient technology converges uniformly over π to

$$\int_0^1 G(\pi, x)\, x^\alpha u(x,\rho)dx, \tag{2.12}$$

where

$$u(x,\rho) = \phi\left(\frac{x}{1-\rho(1-x)}\right)\frac{(1-\rho)^{\sigma^2/a}}{(1-\rho(1-x))^\alpha} - \phi(x). \tag{2.13}$$

Consider first the case $\phi_0 > 0$. Fix $\widehat{\rho} \in (0,1)$, and let

$$\varepsilon = \frac{(1-\widehat{\rho})^{\sigma^2/a}}{(1-\widehat{\rho}(1-x))^{\alpha}} - 1.$$

Since $\sigma^2 < \alpha a$, $\varepsilon > 0$. From (2.11), there exists $\bar{x}$ such that for all $x \in [0, \bar{x}]$, $u(x, \widehat{\rho}) > \frac{\phi_0 \varepsilon}{2} > 0$. Thus for π small enough

$$\int_0^1 G(\pi, x) x^{\alpha} u(x, \widehat{\rho}) dx > \int_{\pi}^1 G(\pi, x) x^{\alpha} u(x, \widehat{\rho}) dx.$$

Using (2.7) we have

$$\int_{\pi}^1 G(\pi, x) x^{\alpha} u(x, \widehat{\rho}) dx = \frac{2\sigma^2}{\psi a^2} g(\pi) \int_{\pi}^1 x^{\alpha - \frac{3}{2} - \frac{1}{2}\psi} (1-x)^{-\frac{3}{2} + \frac{1}{2}\psi} u(x, \widehat{\rho}) dx > 0 \qquad (2.14)$$

That $r\sigma^2 > \frac{a^2}{2}\alpha(\alpha - 1)$ implies that $\psi > 3$. Thus the integral on the right-hand side of (2.14) diverges as $\pi \to 0$. In this case its sign is determined by the sign of $u(\cdot, \widehat{\rho})$ in the neighborhood of 0, which we just found to be positive.

Suppose now that $\phi_0 = 0$. Then unless $\phi(x) \equiv 0$ there exists $\bar{x}$ such that $\phi(\bar{x}) = \widehat{\phi} > 0$ and for all $x \in [0, \bar{x}]$ $\phi(x) < \widehat{\phi}$. Then by choosing $\widehat{\rho}$ large enough so that $\frac{\pi}{1-\widehat{\rho}(1-\pi)} = \bar{x}$ we can ensure that (2.14) still holds.■

Note that v satisfies (2.11) in particular if the cost c is a power function as in Lemma 1. The broad intuition for Proposition 7 is quite simple in the particular case in which $v(\pi) = cst \times \pi^{\alpha}$. In this case, the continuation utility of the trader V behaves as π^{α} for π small. Assume that a trader with such a small π "picks up nickels in front of a steamroller", or generates an instantaneous return of $1+\varepsilon$ with probability $\frac{1}{1+\varepsilon}$ or a large loss. Then in case of success her new reputation is approximately

$$\pi' \simeq \pi \left(1 + \frac{a}{\sigma^2}\varepsilon\right).$$

Her net expected gain from such a gamble is thus close to (up to a multiplicative constant)

$$\frac{1}{1+\varepsilon}\pi'^{\alpha} - \pi^{\alpha} \simeq \pi^{\alpha}\left(\frac{\alpha a}{\sigma^2} - 1\right)\varepsilon > 0.$$

Propositions 6 and 7 can be interpreted as follows. There are two ratios with simple economic interpretations that determine whether gambling occurs or not in the Berk and Green equilibrium with linear fees:

$$R_1 = \frac{a}{\sigma^2}, R_2 = \frac{r}{a}.$$

Risk-shifting does not occur when $R_1 \leq \frac{1}{\alpha}$, and occurs if $R_1 > \frac{1}{\alpha}$ and $2R_2 > R_1\alpha(\alpha - 1)$.The condition on R_1 means that risk-shifting occurs if R_1 is sufficiently large, and the trader's skills are sufficiently scalable (α large). The ratio R_1 is large when skilled traders generate high excess returns with relatively low risks. The condition on R_2 implies that given that R_1 satisfies this condition, it must also be the case that the trader is sufficiently impatient. The intuitions for these results are simple. If skilled traders generate high unit excess returns with low risk, and their strategies are scalable, then small good news about their skills have an important impact on their continuation utilities. This is because it entails that they will receive large inflows of funds in the short term. If the trader is sufficiently patient, however, she cares only for the long run in which she will end up with the reputation that she deserves regardless of earlier gambling attempts. Note that R_1 can also be interpreted as the agents' *prior* about the Sharpe ratio of a skilled trader, and that immediate gambling occurs only if π_0 is sufficiently small. This suggests that fallen traders, that are initially viewed as high-potential individuals, and then disappoint with a series of low returns, would be particularly prone to gambling. This is consistent with anecdotal evidence from rogue trading.

Calibration

Berk and Green assume $c(q) = q/2\gamma$, which yields $v(\pi) = \gamma\pi^2a^2/2$. Proposition 7 shows that in this case risk-shifting occurs whenever $\sigma^2 < 2a$ and $r\sigma^2 > a^2$. A simple calibration exercise shows that situations in which risk-shifting takes place are quite plausible. For example, broadly consistent with Berk and Green, take $a = 5\%$, $\sigma = 25\%$, and $r = 5\%$ then $a/\sigma^2 = 0.8$ and $a/r = 1$. Berk and Green show that such orders of magnitude lead to realistic sensitivities of fund flows to realized performance. Proposition 7 predicts that there should be gambling in this case.

3 Contracting Risk Shifting Away

The previous section shows that under plausible parameter values, the spot markets with linear fees described in Berk and Green induce gambling. Situations in which gambling occurs cannot be equilibrium outcomes. Investors are strictly better off investing directly in the risk-free benchmark when they anticipate risk shifting by the trader because $r' < r$.

This section studies feasible contracts in the presence of risk-shifting. Note first that the contracting problem is trivial under full commitment. If both the trader and investors can fully commit to long-term contracts, then coping with risk-shifting is easy. Both parties can agree for example on a contract of infinite duration that features a unique upfront payment to the trader, and specifies future fund sizes as a function of future perceived skills. This does not induce risk-shifting while generating the same *ex ante* utilities as in the first-best described in Proposition 3. This "slavery" arrangement is of limited practical interest. Here, we study more practically relevant situations in which commitment is one-sided. Section 3.1 tackles the case in which the trader can commit but not investors. Section 3.2 studies the opposite situation in which the commitment problem is on the trader's side. Each situation seems empirically relevant. Limited commitment on the investors' side represents the situation of a hedge fund that has only access to investors with short horizons, e.g. because of important liquidity needs. The situation in which the trader cannot commit while investors can may be interpreted as one in which a proprietary trader in a large bank with commitment power is free to seize better opportunities in the job market.

3.1 Investors Cannot Commit

We modify the model in Section 2 as follows. We relax Assumption 1 (linear fees), maintain Assumption 2 (risk-shifting), and introduce the following.

Assumption 3 *At each date, a new cohort of investors enters the economy. Investors are competitive and live for one period. The trader can fully commit to a plan of future actions.*

Our first result shows that with the parameter values identified in Proposition 7, the only contract that does not induce gambling must be such that the trader earns zero surplus.

Proposition 8 *Suppose that* $c(q) = cq^{\frac{1}{\alpha-1}}$*, where* $c > 0$*,* $\alpha > 1$*. If* $\sigma^2 < \alpha a$ *and* $r\sigma^2 > \frac{a^2}{2}\alpha(\alpha - 1)$*, then for* Δt *sufficiently small, the trader cannot extract any surplus.*
.

Proof. We first establish the following lemma.

Lemma 2 *There exists* $k > 0$ *s.t.* $r\sigma^2 > \frac{a^2}{2}\alpha(\alpha - 1) \rightarrow V(\pi) \sim_0 k\pi^\alpha$

Proof. See the Appendix.■

Now, let $\chi(\pi, \Delta t)$ denote the expected surplus of the trader if she starts out with perceived skills $\pi \in (0, 1)$. We have

Lemma 3 *For all* $\pi \in (0,1)$,

$$\chi(x,\Delta t) - \chi(\pi,\Delta t) \underset{\substack{x\to\pi \\ \Delta t\to 0}}{\longrightarrow} 0$$

Proof. See the Appendix.■

Consider a gamble that generates R with prob. $1/R$ and 0 otherwise. Define ρ as $R = (1-\rho)^{-\sigma^2/a}$. Let $U(\pi,\rho,\Delta t)$ denote the net expected gain from such a gamble for a trader with $\pi_0 = \pi$. Let

$$\omega(\pi,\Delta t) = \chi(x,\Delta t)/V(\pi,\Delta t).$$

Note that $0 \leq \omega \leq 1$. Lemma 3 implies that

$$U(\pi,\rho,\Delta t) \sim_{\Delta t\to 0} \left[(1-\rho)^{\frac{\sigma^2}{a}}\omega\left(\frac{\pi}{1-\rho(1-\pi)},\Delta t\right)V\left(\frac{\pi}{1-\rho(1-\pi)},\Delta t\right) - \omega(\pi,\Delta t)V(\pi,\Delta t)\right]. \tag{3.1}$$

Let us rewrite the right-hand side of (3.1) as $A+B$, with

$$\begin{aligned} A &= \left[V\left(\frac{\pi}{1-\rho(1-\pi)},\Delta t\right) - \frac{V(\pi,\Delta t)}{(1-\rho(1-\pi))^{\alpha}}\right](1-\rho)^{\frac{\sigma^2}{a}}\omega\left(\frac{\pi}{1-\rho(1-\pi)},\Delta t\right), \\ B &= V(\pi,\Delta t)\left[\omega\left(\frac{\pi}{1-\rho(1-\pi)},\Delta t\right)\frac{(1-\rho)^{\frac{\sigma^2}{a}}}{(1-\rho(1-\pi))^{\alpha}} - \omega(\pi,\Delta t)\right]. \end{aligned}$$

Assume first that $\omega(\pi,\Delta t) \underset{\pi,\Delta t\to 0}{\nrightarrow} 0$. Then, ω being bounded one can assume w.l.o.g. that there exists two sequences $\pi_n \to 0$, $\delta_n \to 0$, and $l \in (0,1]$ s.t.

$$\omega(\pi_n,\delta_n) \underset{n\to\infty}{\longrightarrow} l.$$

Pick ρ s.t.

$$(1-\rho)^{\frac{\sigma^2}{a}-\alpha} > 1/l.$$

Define the sequence (π_n') as

$$\forall n, \pi_n = \frac{\pi_n'}{1-\rho(1-\pi_n')}.$$

Then $U(\pi_n',\rho,\delta_n) > 0$ for n sufficiently large because $V(\pi_n',\delta_n)$ dominates A for n large from Lemma 2, and the term between square brackets in B is strictly positive for such n because $\omega(\pi,\Delta t) \leq 1$.

Thus we have established that $\omega(\pi, \Delta t) \underset{\pi,\Delta t \to 0}{\to} 0$. Assume by contradiction that for all Δt there exists $p(\Delta t)$ s.t. $\omega(p(\Delta t), \Delta t) \neq 0$. For all π sufficiently small define then $\rho(\pi, \Delta t)$ as

$$\frac{\pi}{1-\rho(\pi,\Delta t)(1-\pi)} = p(\Delta t)$$

Then, $U(\pi, \rho(\pi, \Delta t), \Delta t) > 0$ for π and Δt sufficiently small. This is again because B dominates A and is strictly positive for $\pi, \Delta t$ sufficiently small. Thus $\omega = 0$ for Δt sufficiently small.■

Thus, for a range of parameter values that are empirically plausible according to the Berk and Green calibrations, the trader must commit to work for free if she wants to credibly tie her hands and not gamble. Otherwise stated, any arrangement in which the trader has nonzero utility must feature gambling.

This result is driven the behavior of the continuation utility $V(\pi)$ for "bad" traders with a π close to zero. Under the assumptions of Proposition 8, the continuation utility of the trader is too convex for π small and induces gambling as soon as the trader has a hope for some surplus if she reaches some higher π. To illustrate the key role that the curvature of V and thus v in 0 plays in this result, consider the situation - arguably realistic - in which the excess returns generated by the trader decrease with her fund's size only when the assets under management exceed a threshold $\overline{q}$.

Proposition 9 *Assume $a < \sigma^2 < 2a$, $r\sigma^2 > a^2$ and $c(q) = (q - \overline{q})^+$ where $\overline{q} > 0$. For Δt sufficiently small, the trader can commit to not gamble and still extract an expected surplus at least equal to $\chi(\pi)\Delta t = \pi a \overline{q} \Delta t$ per period, where π are her perceived skills at the beginning of the period. If $\overline{q}$ and Δt are sufficiently small, the trader cannot extract the whole surplus that she generates.*

Proof. With this cost function, the first-best surplus has the following limit as $\Delta t \to 0$:

$$\begin{array}{ccc} \pi a \overline{q} & \text{if} & \pi \leq \frac{\overline{q}}{a}, \\ \left(\frac{\pi a + \overline{q}}{2}\right)^2 & \text{if} & \pi \geq \frac{\overline{q}}{a}. \end{array}$$

Propositions 6 implies that a linear surplus does not induce risk-shifting for these values of a, r, and σ. The largest linear surplus that can be implemented is $\pi a \overline{q}$. Proposition 7 implies that a quadratic surplus induces gambling for π and Δt sufficiently small. It implies that for $\overline{q}$ sufficiently small any contract that yields the total surplus will then induce risk-shifting as soon as the trader's skills become sufficiently small. ■

Proposition 9 shows two results. First, it shows an example of parameter values for which the trader can credibly commit to not gamble and, unlike in Proposition 8, still earn some of the surplus. This is so even though parameter values in Proposition 9 satisfy the conditions in Proposition 8, except for the fact that the cost c is equal to 0 on an arbitrarily small right-neighborhood of 0. This illustrates that with short-termist investors, how much the trader can extract crucially depends on the behavior of the trading surplus in the right-neighborhood of 0.

Second, Proposition 9 shows that the maximal surplus the trader can earn must be strictly lower than the total surplus generated by the trader. This implies that the trader is forced in this case to leave some surplus to investors. The empirical implication is that in the presence of risk-shifting and short-term contracts, competitive and uninformed investors earn persistent excess returns from investing in the best funds.

A particular implementation of the contract with linear surplus is as follows. The trader commits to a fee per period that is smaller than the first-best one in Proposition 3 so as to leave some of the total surplus to investors. With such smaller fees, investors would supply more funds than the optimum fund size in Proposition 3. Thus the trader must also impose a restriction on the fund's size, and turn down inflows when he has a good reputation. In other words, by committing to fees and fund size, the trader can implement the first-best in the Berk and Green model with risk-shifting. Some rents must be left to investors, however. In particular, this contract with linear surplus is not necessarily the optimal one from the trader's standpoint.

3.2 The Trader Cannot Commit (INCOMPLETE)

We now consider the symmetric situations in which investors can commit to a contract, while the trader cannot.

Assumption 4 *We relax Assumptions 1 and 3 and maintain Assumption 2 (risk-shifting). We assume that investors are long-lived, competitive, and can fully commit to a contract, while the trader cannot. At the end of each trading round, she can terminate a contract, and recontract with new investors with impunity.*

Such a "no-slavery" assumption is commonplace in the career-concern literature (see, e.g., Harris and Holmstrom (1982)). Our problem in this case can be described as a variant of Harris and Homstrom (1982) where the agent is risk-neutral but can secretly take fair gambles. We also make the following additional assumption on the shape of the expected first-best surplus $V(\pi)$ when the trader starts out with perceived skills π.

Assumption 5 *The function*

$$\begin{array}{rcl} (0,+\infty) & \to & (0,+\infty) \\ x & \to & \log(V(e^{-x})) \end{array}$$

is concave.

When the cost c has the form imposed in Lemma 1, V clearly satisfies this assumption in the limiting case in which $\alpha = 1$. More generally, this condition clearly holds when V is a power function.[3] Under these conditions, gambling-proof contracts have the following features.

Proposition 10

There exists $\pi^* \in [0, 1]$ such that a trader with initial skills π receives funds if and only if $\pi \geq \pi^*$. If the trader receives funds, she is promised the first-best expected surplus $V(\pi)$.

Proof.

Note that without loss of generality, we can consider only contracts in which payments to the trader are indefinitely postponed. Such contracts are weakly dominant since all agents are risk-neutral and equally impatient.

Assume first that a trader is employed. If her initial expected surplus is strictly less than the first-best, a competitive investor could offer the same contract plus an upfront payment equal to the difference. This would entail no risk-shifting and the trader would prefer it. Thus if a trader is employed, she initially receives the first-best surplus.

TO BE COMPLETED■

References

[1] Aumann, R.J. and Perles, M., 1965. "A variational problem arising in economics," *Journal of Mathematical Analysis and Applications* 11, 488-503.

[2] Basak, S., Pavlova, A., and Shapiro, A., 2007. "Optimal Asset Allocation and Risk Shifting in Money Management", *Review of Financial Studies* 20(5), 1583-1621.

[3] Berk, J.B. and Green, R.C., 2004. "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269-1295.

[4] Carpenter, J.N., 2000. "Does Option Compensation Increase Managerial Risk Appetite?," *Journal of Finance*, 55(5), 2311-2331.

[5] Cuoco, D. and Kaniel, R., 2007. "Equilibrium Prices in the Presence of Delegated Portfolio Management", Working paper.

[3] Equation (A5) gives the per period surplus v corresponding to a power function V.

[6] Harris, M. and Holstrom, B., 1982. "A Theory of Wage Dynamics," *Review of Economic Studies*, 49(3), 315-33.

[7] Holstrom, B., 1999. "Managerial Incentive Problems - A Dynamic Perspective", *Review of Economic Studies*, 66 (1), 169-82.

[8] Jagannathan, R., Malakhov, A., and Novikov, D., 2006. "Do Hot Hands Exist Among Hedge Fund Managers? An Empirical Evaluation," NBER Working Papers 12015.

[9] Driver, K.B., 2003. Analysis Tools With Applications, Lecture Notes.

[10] Kosowski, R., Naik, N.Y., and Teo, M., 2007. "Do hedge funds deliver alpha? A Bayesian and bootstrap analysis," *Journal of Financial Economics* 84, 229—264.

[11] Liptser, R. Sh. and Shiryaev, A. N., 1978. Statistics of random processes, Springer.

[12] Townsend, R., 1979. "Optimal Contracts and Competitive Markets with Costly State Verification", *Journal of Economic Theory*, 21(2), 265-293.

Appendix

Proof of Proposition 1

Lemma 4 *There exists* $(z_1, z_2) \in \mathbb{R} \times [0, +\infty)$ *such that* (z_1, z_2) *satisfies (1.2) and* $P^*(U) = z_1 + z_2$.

Proof. Condition (1.3) implies that the set $Z = \{(z_1, z_2) : \forall y \geq 0, z_1 + yz_2 \geq U(y)\}$ is nonempty. It is closed, and there exists K such that

$$(z_1, z_2) \in Z \rightarrow z_1 \geq K, z_2 \geq K.$$

The function $(z_1, z_2) \rightarrow z_1 + z_2$ is continuous. Thus, there exists $(z_1,\, z_2) \in Z$ such that $P^*(U) = z_1 + z_2$. Condition (1.3) readily implies that $z_2 \geq 0$.■

Let μ a probability measure that satisfies (1.1), and $(z_1,\, z_2) \in \mathbb{R}^2$ defined as in Lemma 4. We have

$$z_1 + z_2 = \int_0^\infty (z_1 + Rz_2)\, d\mu(R) \geq \int_0^\infty U(R) d\mu(R).$$

This implies that

$$P^*(U) \geq P(U).$$

Let us show that the reverse inequality also holds. Establishing the reverse inequality for U with compact support is without loss of generality: for all U satisfying (1.3), there clearly exists $V \in C_c([0, \infty))$ such that $V \leq U$ and $P^*(V) = P^*(U)$. We omit the straightforward proof of the following lemma.

Lemma 5

a. $P^*(U_1) \leq P^*(U_2)$ *for* $U_1,\, U_2 \in C_c([0, +\infty))$ *such that* $U_1 \leq U_2$,

b. $P^*(\alpha U) = \alpha P^*(U)$ *for* $U \in C_c([0, +\infty))$ *and* $\alpha \in [0, +\infty)$,

c. $P^*(U_1 + U_2) \leq P^*(U_1) + P^*(U_2)$ *for* $U_1,\, U_2 \in C_c([0, +\infty))$.

From Lemma 5, $P^*(\cdot)$ is a positively homogeneous and subadditive functional on $C_c([0, +\infty))$. The Hahn-Banach Theorem therefore implies that for any $U \in C_c([0, \infty))$, there exists a positive linear functional L_U on $C_c([0, \infty))$ such that $L_U \leq P^*$ and $L_U(U) = P^*(U)$. By the Riesz representation Theorem, there exists a Borelian measure μ_U on $[0, \infty)$ such that for all $V \in C_c([0, \infty))$

$$L_U(V) = \int_0^\infty V(R) d\mu_U(R).$$

For $M > 1$, let $u_M, v_M \in C_c([0,\infty)) \times C_c([0,\infty))$ such that

$$\begin{aligned} u_M(x) &= 1 \text{ on } [0,M], \, x \geq M \to u_M \leq 1, \\ v_M(x) &= x \text{ on } [0,M], \, x \geq M \to v_M \leq M. \end{aligned}$$

Clearly,

$$P^*(u_M) = P^*(v_M) = 1$$

Then

$$\begin{aligned} L_U(u_M) &= \int_0^\infty u_M(R)\, d\mu_U(R) \leq P^*(u_M) = 1, \\ L_U(v_M) &= \int_0^\infty v_M(R)\, d\mu_U(R) \leq P^*(v_M) = 1. \end{aligned}$$

Letting $M \to +\infty$ implies

$$\int_0^\infty d\mu_U(R) \leq 1, \int_0^\infty R d\mu_U(R) \leq 1,$$

and thus

$$P^*(U) = L_U(U) = \int_0^\infty U(R) d\mu_U(R) \leq P(U).$$

■

Proof of Proposition 2

Let U satisfying condition (1.3) and such that $\lim_{+\infty} U = +\infty$. Let $(z_1,\, z_2) \in \mathbb{R}^2$ defined as in Lemma 4. That $\lim_{+\infty} U = +\infty$ implies that for such a $(z_1,\, z_2)$, $z_2 > 0$. (Note that the assumption $\lim_{+\infty} U = +\infty$ is only a sufficient condition for $z_2 > 0$, which is all that is required to establish the proposition). Let

$$S = \{y \geq 0 : z_1 + z_2 y = U(y)\}.$$

That U is continuous together with condition 1.3 implies that S is a nonempty compact set. Let

$$y_1 = \min S, \;\; y_2 = \max S.$$

We now proceed in two steps.

Step 1. We show that $y_2 \geq 1 \geq y_1$.

Proof. We show that $y_2 \geq 1$. The proof that $y_1 \leq 1$ is symmetric. Assume that

$y_2 < 1$. For some $\varepsilon \in (0, 1 - y_2)$, let

$$\eta(\varepsilon) = \min_{y \geq y_2 + \varepsilon} \left\{ \frac{z_1 - U(y)}{y} + z_2 \right\}.$$

By condition 1.3,

$$\lim_{y \to +\infty} \frac{z_1 - U(y)}{y} + z_2 = z_2 > 0,$$

and by definition of y_2

$$y \geq y_2 + \varepsilon \rightarrow \frac{z_1 - U(y)}{y} + z_2 > 0,$$

continuity of U therefore implies that $\eta(\varepsilon) > 0$. Let (z_1', z_2') defined as

$$z_1' = z_1 + (y_2 + \varepsilon)\, \eta(\varepsilon), \;\; z_2' = z_2 - \eta(\varepsilon).$$

First, the pair (z_1', z_2') satisfies (1.2). To see this, notice that

$$z_1' + y z_2' = z_1 + y z_2 + \eta(\varepsilon)(y_2 + \varepsilon - y).$$

Thus $z_1' + y z_2' > z_1 + y z_2 \geq U(y)$ for $y < y_2 + \varepsilon$.

Further $z_1' + y z_2' \geq z_1 + y z_2 - \eta(\varepsilon)\, y \geq U(y)$ for $y \geq y_2 + \varepsilon$ by definition of $\eta(\varepsilon)$.

Second, we have

$$z_1' + z_2' = z_1 + z_2 + (y_2 + \varepsilon - 1)\, \eta(\varepsilon) < z_1 + z_2,$$

which contradicts the definition of (z_1, z_2). Thus it must be that $y_2 \geq 1$.□

Step 2. Now, if $y_1 = y_2$, then step 1 implies that $S = \{1\}$, and there is no gambling. If $y_1 < y_2$,

from

$$\begin{aligned} z_1 + y_1 z_2 &= U(y_1), \\ z_1 + y_2 z_2 &= U(y_2), \end{aligned}$$

we have

$$z_1 + z_2 = \frac{1 - y_1}{y_2 - y_1} U(y_2) + \frac{y_2 - 1}{y_2 - y_1} U(y_1). \tag{A1}$$

From (A1), $P(U) = P^*(U)$ is attained with a payoff equal to y_1 with probability $\frac{y_2 - 1}{y_2 - y_1}$ and y_2 with probability $\frac{1 - y_1}{y_2 - y_1}$. These probabilities are well-defined from step 1. Notice that there is gambling if and only if $y_1 < 1$ and $y_2 > 1$.■

Proof of Proposition 4

We first show pointwise convergence. That is, we establish (2.6) for fixed $\pi_0 = \pi$. By Bayes' theorem, $\pi_{n\Delta t}$, the perceived skills at date $n\Delta t$, satisfy

$$\pi_{n\Delta t} = \frac{\pi_0 \varphi_{n\Delta t}}{1 - \pi_0 + \pi_0 \varphi_{n\Delta t}}, \tag{A2}$$

where

$$\varphi_{n\Delta t} = \exp\left\{ \frac{a}{\sigma^2} \left(a \left(\theta - \frac{1}{2} \right) n\Delta t + \sigma B_{n\Delta t} \right) \right\} \tag{A3}$$

is the likelihood ratio process. Let us introduce the continuous-time process $(\pi_t)_{t\geq 0}$ that obeys

$$d\pi_t = \frac{a}{\sigma}\pi_t(1-\pi_t)d\overline{B}_t, \quad \pi_0 = \pi,$$

where $\overline{B}_t = \frac{1}{\sigma}\left(\theta a t + \sigma W_t - a\int_0^t \pi_s ds\right)$. Then $(\overline{B}_t)_{t\geq 0}$ is a standard Wiener process under the agents' filtration (see Liptser and Shiryaev (1978)). Further, as $\Delta t \to 0$ and $n\Delta t \to t$, $\pi_{n\Delta t} \to \pi_t$ a.s. (see Liptser and Shiryaev (1978)). Hence, $V(\pi)$ can be re-written as

$$\begin{aligned} V(\pi) &= E_0 \int_0^\infty e^{-rt} v(\pi_t) dt, \\ s.to \quad d\pi_t &= \frac{a}{\sigma}\pi_t(1-\pi_t)d\overline{B}_t, \quad \pi_0 = \pi, \end{aligned} \tag{A4}$$

By the Feynman-Kac formula, the function V solves the following linear second-order differential equation:

$$\frac{a^2}{2\sigma^2}\pi^2(1-\pi)^2 V''(\pi) - rV(\pi) + v(\pi) = 0. \tag{A5}$$

From (A4) it follows that

$$V(0) = v(0)/r, \quad V(1) = v(1)/r. \tag{A6}$$

The corresponding homogeneous equation

$$\frac{a^2}{2\sigma^2}\pi^2(1-\pi)^2 V''(\pi) - rV(\pi) = 0 \tag{A7}$$

has two regular singular points at 0 and 1. The solutions of the homogeneous equation are linear combinations of the two independent solutions

$$\begin{aligned} g(\pi) &= (1-\pi)^{\frac{1}{2}+\frac{1}{2}\psi}\pi^{\frac{1}{2}-\frac{1}{2}\psi}, \\ h(\pi) &= g(1-\pi). \end{aligned}$$

From here, formulas (2.6) and (2.7) are standard results in the theory of inhomogeneous differential equations. The function G is the Dirichlet-Green function for the differential operator associated with the homogeneous differential equation (see, e.g., Driver (2003)).

We now show that $V(\pi_0, \Delta t)$ converges to $V(\pi_0)$ uniformly in π_0 as $\Delta t \to 0$. We have

$$V(\pi_0) - V(\pi_0, \Delta t) = E_0 \sum_{n=0}^{\infty} \int_{n\Delta t}^{(n+1)\Delta t} \left(e^{-rt}v(\pi_t) - e^{-rn\Delta t}v(\pi_{n\Delta t})\right) dt.$$

Thus it suffices to show that $\forall \varepsilon > 0$, $\exists \overline{\Delta} t$ such that $\forall \Delta t < \overline{\Delta} t$ and $\forall \pi \in [0,1]$

$$\sup_{s \leq \Delta t} \sup_{\pi_0 \in [0,1]} \left| E\left(v(\pi_s) - v(\pi_0)\right) \right| < \varepsilon. \tag{A8}$$

By change of variables (A8) can be re-written as

$$\sup_{s \leq \Delta t} \sup_{\pi_0 \in [0,1]} \left| E\left(\widehat{v}(\pi_0, B_s, s) - \widehat{v}(\pi_0, 0, 0)\right) \right| < \varepsilon,$$

where

$$\widehat{v}(\pi_0, x, t) = v\left(\frac{\pi_0 \exp\left\{\frac{a}{\sigma^2}\left(a\left(\theta - \frac{1}{2}\right)t + \sigma x\right)\right\}}{1 - \pi_0 + \pi_0 \exp\left\{\frac{a}{\sigma^2}\left(a\left(\theta - \frac{1}{2}\right)t + \sigma x\right)\right\}}\right). \tag{A9}$$

Since v is continuous over [0,1] and thus uniformly continuous, it is enough to show that $\forall \pi_0 \in [0,1]$ and $\forall \varepsilon > 0$, $\exists \overline{\Delta} t$ such that $\forall \Delta t < \overline{\Delta} t$

$$\sup_{s \leq \Delta t} \left| E\left(\widehat{v}(\pi_0, B_s, s) - \widehat{v}(\pi_0, 0, 0)\right) \right| < \varepsilon, \tag{A10}$$

which follows from the weak converge of the measures induced by B_s to the measure concentrated at 0 as $s \to 0$.■

3.3 Proof of Lemma 2

We have

$$\int_0^1 G(\pi, x)\, x^\alpha dx = \frac{2\sigma^2}{a^2\psi}\left[\begin{array}{l} (1-\pi)^{\frac{1}{2}(1+\psi)}\pi^{\frac{1}{2}(1-\psi)}\int_0^\pi (1-x)^{-\frac{1}{2}(\psi+3)}x^{\alpha+\frac{1}{2}(\psi-3)}dx \\ +(1-\pi)^{\frac{1}{2}(1-\psi)}\pi^{\frac{1}{2}(1+\psi)}\int_\pi^1 (1-x)^{\frac{1}{2}(\psi-3)}x^{\alpha-\frac{1}{2}(3+\psi)}dx \end{array}\right],$$

where ψ is defined in (2.8). Further

$$\sigma^2 > \frac{a^2}{2}\alpha(\alpha-1) \to \alpha < \frac{1}{2}(1+\psi),$$

and thus

$$(1-\pi)^{\frac{1}{2}(1-\psi)}\pi^{\frac{1}{2}(1+\psi)}\int_\pi^1 (1-x)^{\frac{1}{2}(\psi-3)}x^{\alpha-\frac{1}{2}(3+\psi)}dx \sim_0 cst \times \pi^\alpha$$

because the integral diverges, and

$$(1-\pi)^{\frac{1}{2}(1+\psi)}\pi^{\frac{1}{2}(1-\psi)}\int_0^\pi (1-x)^{-\frac{1}{2}(\psi+3)}x^{\alpha+\frac{1}{2}(\psi-3)}dx \sim_0 cst \times \pi^\alpha$$

because the integral converges.■

3.4 Proof of Lemma 3

Assume conversely that there exists a sequence $(x_n, \delta_n)_{n\in\mathbb{N}}$ s.t.

$$x_n \to \pi, \delta_n \to 0, \text{ and } \chi(x_n, \delta_n) - \chi(\pi, \delta_n) \nrightarrow 0.$$

Since χ is bounded, we can assume w.l.o.g. that $\chi(x_n, \delta_n) - \chi(\pi, \delta_n) \to l \neq 0$. But then there would be risk shifting. For n arbitrarily large, the trader could take an arbitrarily safe bet to switch from x_n to π (or reciprocally depending on the sign of l). This generates a strict utility gain of order l and comes at a gambling cost of order δ_n only.■