The Skill Paradox: Is Smarter Always Better? *

Professor Stephen Satchell
Trinity College, University of Cambridge
Trinity Street, Cambridge CB2 1TQ, United Kingdom.
Ph: +44 1223338409
Email: ses11@cam.ac.uk

Dr Oliver Williams
Scalpel Research Ltd
oliver.williams@scalpelresearch.com

Abstract
In the matter of investment forecasting skill it is often supposed that *more* is automatically preferable to *less*. In other words, the marginal utility of improving one's forecasting accuracy is monotonically positive. However it is also commonly believed that large scale speculative capital flows (presumably driven by skilled investors) are a major contributor to financial market volatility. The consequent impact of increased volatility on the real economy imposes disutility on non-investors and this exposes an interesting trade-off between the benefits of skill improvement which accrue to investors, and the costs suffered more broadly by society. This paper constructs a formal analytic framework in which to discuss these issues, questions whether the marginal utility of skill is in fact monotonic for the individual and considers implications for policy-makers.

*The authors thank Robert Kosowski and Hamish Low for many useful comments.

Abstract

In the matter of investment forecasting skill it is often supposed that *more* is automatically preferable to *less*. In other words, the marginal utility of improving one's forecasting accuracy is monotonically positive. However it is also commonly believed that large scale speculative capital flows (presumably driven by skilled investors) are a major contributor to financial market volatility. The consequent impact of increased volatility on the real economy imposes disutility on non-investors and this exposes an interesting trade-off between the benefits of skill improvement which accrue to investors, and the costs suffered more broadly by society. This paper constructs a formal analytic framework in which to discuss these issues, questions whether the marginal utility of skill is in fact monotonic for the individual and considers implications for policy-makers.

The Skill Paradox: Is Smarter Always Better?

S.E. Satchell
Trinity College
University of Cambridge

O.J. Williams*
Scalpel Research Ltd

1 Introduction

"Where ignorance is bliss, 'Tis folly to be wise" *Thomas Gray (1742)*

"All wisdom is foolish that does not adapt itself to the common folly" *Michel Eyquem de Montaigne (1533-1592)*[1]

Discussion of asset price predictability has raged in the economics and finance literature for centuries. Contributions range from the technical analysis methods devised in 18th century Japan[2] to the efficient market hypotheses debates throughout the 20th century (e.g. Bachelier (1900), Fama (1965), Samuelson (1965), Malkiel (1973)) and more recent work on behavioural finance (Kahneman and Tversky (1979), Simon (1987), Shleifer (1999)). Regardless of where one stands on this topic it is evident that investors' appetite for actively managed funds has continued unabated and their trading activities constitute a vast proportion of everyday financial market turnover.

Against this background an orthodoxy has developed in certain political-economic quarters as regards the preferred structure of financial markets. For instance arguments are frequently advanced in favour of wider household ownership of equities (often connected with privatization schemes, e.g Perotti (1995)), greater control by individuals of their pension portfolios (Emmerson and Wakefield (2003)) and investment by the general public in hedge funds (e.g. FSA (2005)). In many jurisdictions this has led to the implementation of specific policy measures with profound consequences for the asset management industry as well as the welfare of individuals themselves.

The topic of financial literacy is also becoming an increasingly prominent public agenda item globally. (Cox (2006), HM Treasury (2009) and EBF (2009) give insights into the treatment of this matter in the United States, United Kingdom and European Union respectively). Indeed, even when not stated explicitly, elements of investor education are often implicit in moves towards greater consumer protection in financial services, such as the Treating Customers Fairly principle devised by the UK Financial Services Authority (FSA (2005)).

*The authors thank Robert Kosowski and Hamish Low for many useful comments.

[1] quoted from 'The Complete Essays of Montaigne' translated by D. Frame (1957, book 3 chapter 3)

[2] the so-called Sakata Rules reputedly devised by rice trader Munehisa Homma

From the point of view of financial economics, a major substantive impact of these measures is to alter the distribution and levels of forecasting skill among investors. As the *dramatis personae* of markets change, so too does the character of their price dynamics. An extensive body of theoretical literature demonstrates how active trading can lead to increased volatility; this trading may itself be the result of superior infomation (Grossman and Stiglitz (1980)), opportunistic 'noisy' behaviour (Campbell and Kyle (1988), De Long *et al* (1990)) or fundamental analysis (Graham and Dodd (1934)), but all of these channels involve some component of skilled forecasting.[3] However when investors are relatively unskilled they may have no conscious forecast whatsoever or, if they do, update it erroneously and infrequently. Under relatively undemanding assumptions about rationality on the part of investors, low skill translates into negligible trading and, in turn, less volatile markets.

Given the high public profile of these important policy issues we have found a surprising lack of satisfactory theoretical frameworks with which to evaluate them. The supposition that such policies are a 'good thing' sits uneasily with the intuition that they may bring increased volatility in their wake. It is clear that a substantial proportion of the world's economic activity is concerned with the trading and management of financial products, executed by an increasingly-highly-trained technocracy. Recent months, however, have witnessed dramatic turbulence across these markets, with severe knock-on effects to the real sector, events which vocally beg the question of whether accumulating investment skill is 'worth all the trouble'.

In this paper we specifically set out to address two broad groups of policy questions which we categorise as follows;

Type I questions concern the *level* of skill in the market: *What is the optimal level of skill which an investor should attempt to attain? Is more skill always preferable to less skill? Is it socially desirable to increase the level of all investors' skill in general?* These are relevant to discussion of financial literacy and investor education.

Type II questions relate instead to the *distribution* of skill: *Is it in the interests of unskilled investors to delegate portfolio decisions to professional managers with superior skill? Is there a socially-optimal allocation of capital across differentially-skilled investor types?* These are relevant to financial product innovation (e.g. introduction of 130/30 funds) and broadening retail access to more sophisticated products such as hedge funds.

The importance of Type II questions is sharpened by the observation that many famous financial catastrophes (e.g. Barings, Long Term Capital Management and Madoff) have involved one group of investors accepting *uncritically* a real (or imaginary) complex, mathematical derivatives strategy purveyed by another group, *supposed* to have superior skill.[4]

Our questions are motivated by the very high degree of social anxiety caused by extreme financial market turbulence and the accompanying notion that highly-trained financial professionals are directly culpable. This view continues to enjoy a particularly high profile in the popular media, e.g. Trillin (2009) ('*...having smart guys there almost caused Wall Street to collapse.*') as well as in governments across the world, exemplified by the proposal of the 'Volcker

[3]Value investors used to claim that their skills involved no notion of forecasting outside pure arbitrage, but the notion of value involves some projection of future earnings in most cases.

[4]See Fay(1997), Lowenstein (2001) and Markopolos (2010) for explanatory accounts of Barings, LTCM and Madoff respectively.

Rule' by Obama (2010)). Similarly, within the academic and practitioner communities there is much debate about future performance prospects of quantitative funds given an increasing degree of herding by investors into a relatively small number of correlated strategies (see Fabozzi *et al* (2008)).

Our results are illuminating: we find that increasing the skill level of one group of investors can indeed increase their own utility but the induced incremental volatility can simultaneously reduce the utility of others with a net impact on social welfare which is negative. We also find that even if competing groups of independent investors increase their level of skill in tandem it is perfectly plausible for overall welfare to decline; once again the culprit is increased volatility.

Existing theory provides us with helpful foundations. While it is common practice to measure managers' *ex post* performance with metrics such as Sharpe ratios and Information Ratios these do not differentiate between returns achieved due to luck or forecasting skill. Fortunately, however, a range of alternatives is available which measure managers' forecasting ability directly; as well as being revealing statistics as regards skill we find that these parameters also have valuable economic content and indeed define skill as a quasi-commodity over which preferences can be established much like conventional goods.

Our approach is to first construct a general equilibrium model of agents with heterogeneous skill levels and solve for equilibrium prices and portfolios conditional on agents' respective proprietary forecasts. We then examine the resulting distributions of equilibrium price, wealth and utility over repeated forecasting instances. This gives us an *ex ante* picture of the relative merits of different market structures (in terms of agents' skill levels).

Existing literature treats the topic of performance measurement extensively; for instance Bain (1996) reviews methods from a practitioner viewpoint, Aragon and Ferson (2006) review theory and recent empirical evidence and Knight and Satchell (2002) covers a range of theoretical and practical themes. A separate branch of literature considers the appropriateness of performance measures and their relationship to the investor's optimal portfolio construction problem (e.g. Leland (1999)). Notwithstanding these contributions we believe that this paper is among the first to address the topic of skill from a marginal utility or welfare perspective.

Much of the literature on dynamic models with heterogeneous investors deals with the case where agents are all price-takers. Examples include Detemple and Murthy (1994) and Zapatero (1998). These models feature exogenous primitives such as stochastic production processes and derive equilibrium interest rates which typically have a weighted average form dependent on investors' asymmetric beliefs and evolution of wealth. However the effect of interaction on risky asset prices is limited or nonexistent. For example, agents decide what quantity of capital to commit to alternative available investments based on an imperfect returns forecast however the actual return does not respond to the quantity demanded, rather the demand effect shows up in the equilibrium rate of return of the risk-free asset which soaks up residual wealth.

For our purposes it is essential to capture the notion of competition between agents for returns. Introducing non-price-taking behaviour into conventional continuous-time dynamic models is a complex undertaking in its own right, and Basak (1997) remains one of the few treatments of this problem. In view of the rich heterogeneity which is central to our analysis we focus here on static models which nevertheless expose our key results. Heterogeneity in static

models has remained relatively unexplored for some time; Lintner (1969) covered much of the relevant ground in a non-price-taking model. We find that introducing forecasting skill into this context brings useful new insight and we look forward to extending this to a dynamic setting in future work.

The organisation of the paper is as follows: in section 2 we introduce our central proposition cast in a general risk-sharing framework, in section 3 we propose a model of two discrete states which demonstrates our central arguments using *hit-rate* as a skill measurement and in section 4 we apply the model to the policy questions outlined above. Section 5 concludes, considering the relevance of our findings for empirical work and areas for further study. Proofs and detailed derivations are relegated to the Appendix.

2 Marginal Expected Utility of Skill

2.1 Analytic Results

Suppose we have n agents: each has a private signal $S^{(i)}$ which is jointly distributed (according to their subjective belief) with the actual state of the world which we denote by z. This joint density is denoted $q(\mathbf{S}, z)$. The agents trade with each other and construct optimal portfolios of wealth in all states which we represent by the function $w^{(i)}(\mathbf{S}, z)$. We make no assumptions on the form which $w^{(i)}$ might take. Finally, each agent's joint distribution q depends on a skill parameter $\theta^{(i)}$. We do not explicitly solve here the optimization problem to determine $w^{(i)}(\mathbf{S}, z)$ (which is well-established elsewhere in the literature, dating back as far as Wilson (1968)).

For clarity in the algebra which follows we tend to suppress the (i) superscripts.

Proposition 2.1. *Marginal expected utility of θ^i is given by:*

$$\mathbb{E}\left[u_w(w(\mathbf{S}, z; \theta))w_\theta(\mathbf{S}, z; \theta)\right] + \mathbb{COV}\left[u(w(\mathbf{S}, z; \theta)), \frac{q_\theta(\mathbf{S}, z; \theta)}{q(\mathbf{S}, z; \theta)}\right] \tag{1}$$

Proof. See Appendix. □

Clearly a sufficient condition for positive marginal expected utility of skill will be having both terms positive and we consider these individually below.

The first term is the marginal expected utility of a skill improvement while keeping the subjective density unchanged, i.e. we focus only on the influence which the skill parameter has on position size (in terms of future state-dependent wealth). For instance: a skill improvement might take the form of a reduction in dispersion of an agent's future state distribution; depending on the precise form of the agent's utility function that increase in precision is likely to lead to larger position-taking *ceteris paribus*. Such an effect will be captured in the derivatives within the expectation operator, although the expectation itself is taken with respect to the unchanged density. For positiveness we need position sizes to be increased (in response to greater skill) at a rapid enough pace to outweigh the accompanying decline in marginal utility (due to risk aversion); evidently this depends on how agents use skill level to determine position size as well as the precise form of the utility function. We call this the *sizing* term.

The second term is the covariance between utility of state-dependent wealth and the ratio $\frac{q_\theta(\mathbf{S}, z; \theta)}{q(\mathbf{S}, z; \theta)}$ (this ratio could approximately be thought of as the proportionate 'distortion' applied

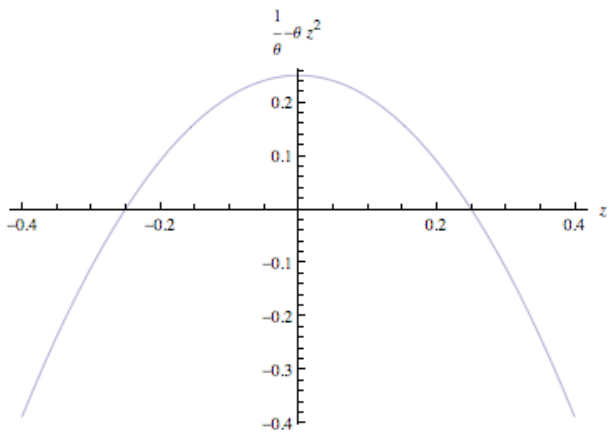

Figure 1: A simple example of a density 'distortion' function based on a standard normal where $\theta = 4$ plotted over the domain $(-0.4, +0.4)$

to the joint density function q at point (**S**,z) when the skill parameter is increased by one unit). For instance, suppose we consider the conditional $q(z|S)$ to be a standard univariate normal and treat the skill parameter as a precision adjustment to the (unity) variance, i.e. variance is defined as $\left(\frac{1}{\theta}\right)^2$, then we have:

$$q(z;\theta) = \frac{\theta}{\sqrt{2\pi}} \exp\left[-\frac{\theta^2 z^2}{2}\right]$$
$$q_\theta(z;\theta) = \frac{q(z;\theta)}{\theta} - \frac{2\theta z^2 q(z;\theta)}{2}$$
$$\frac{q_\theta(z;\theta)}{q(z;\theta)} = \frac{1}{\theta} - \theta z^2$$

which we illustrate in Figure 1 for the case $\theta = 4$. Positive values of the ratio occur where the density is 'pulled upwards' by the improvement in skill (indicating relatively more likely states) and negative values where the density is 'pushed downwards' (indicating the opposite). We are interested in the value of the covariance between this function and the agent's state dependent utility. Ideally, given an improvement in skill $d\theta$ and knowing their private signal S, the agent would adjust their portfolio in perfect sympathy with the 'distortion' function (e.g. overweighting positions where 'distortion' is positive and underweighting where negative), leading to a positive covariance. This term therefore represents how closely state-dependent utility is 'tuned' to the 'improved' subjective density belief which the agent obtains. We call this therefore the *tuning* term.

We next present an example of a special case which helps to provide some intuition.

Corollary 2.2. *Suppose that the agent's wealth is unconditionally normally-distributed, then the condition for positive marginal expected utility of θ^i is:*

$$\frac{\mathbb{COV}\left[w, \frac{q_\theta}{q}\right] + \mathbb{E}\left[w_\theta\right]}{\mathbb{COV}\left[w, w_\theta\right]} > R_A$$

where R_A is Rubinstein's (1973) coefficient of global risk-aversion.

Proof. See Appendix. □

Various remarks follow from this corollary:

(a) As regards $\mathbb{E}[w_\theta]$: positivity of this term by itself is neither necessary nor sufficient for (2.2) to hold. This is the very essence of our proposition: an increase in skill which increases future wealth on average is not guaranteed to deliver net positive marginal utility once other distributional characteristics are taken into account.

(b) The 'tuning' effect represented by $\mathbb{COV}\left[w, \frac{q_\theta}{q}\right]$ can take any sign depending on the precise relationship between the nature of the investor's skill (as it influences the subjective distribution q) and the flexibility which is available to construct the state-dependent portfolio w. If, for instance, the agent is in the unfortunate position of skillfully forecasting parts of the distribution which cannot be traded using available securities (perhaps, for instance, extreme events in the tails of the distribution) then this term could be negligible. In the opposite case of a complete market with no portfolio constraints we would expect this term to be positive.

In summarising these results it is clear that general arguments will be hard to find, therefore we consider further particular special cases in the section which follows. These will demonstrate the applicability of our proposition. Prior to that exposition however we briefly discuss parametric measurements of skill from a more practical viewpoint.

2.2 Skill Metrics in Practice

While a detailed discussion of skill measurement is beyond the scope of this paper it is nevertheless instructive to briefly review candidates for the parameter which we have thus far labeled as θ.[5] The fundamental requirement is that θ should be some measurement of dependence or association between actual future states of the world and a forecast which is made at the time when asset demands are determined.

2.2.1 Hit-Rate

Perhaps the simplest measure of association in this context is the *hit-rate* which we can define informally as the proportion of forecasts which are *directionally* correct. Here we ignore magnitudes of return entirely and focus simply on the signs of the forecasts and actual outcomes. In a basic incarnation we might consider a model of two discrete states which we represent by the outcomes 0 and 1 respectively and denote by the random variable Q. We also restrict forecasts to be correspondingly either 0 or 1 and denote these by the random variable S. For the sake of parsimony we assume that the hit-rate is not conditional on the direction of the forecast (i.e. 0 forecasts are equally-likely to be correct as 1 forecasts) and therefore we define the hit-rate as:

$$\mathbb{P}[Q = 1|S = 1] = P[Q = 0|S = 0] = h$$

2.2.2 Information Coefficient

Grinold and Kahn (1999) define the Information Coefficient (IC) as the correlation between an investor's *forecast* of a future asset return (in terms of a standardised score) and the *actual*

[5] More complete treatments of this topic are provided by Aragon and Ferson (2006), Granger and Machina (2006) and Williams (2009) among others.

return where return is defined as

$$r_{t+1} \equiv \frac{P_{t+1}}{P_t} - 1$$

Formally, if we were to represent the forecast as a standardised normal random variable, $S \sim N(0,1)$, unconditional expected future return as μ, and assume joint-normality between forecast and actual return r_{t+1}, then the conditional forecast return would be:

$$E[r_{t+1} \mid S] = \mu + IC\sigma S$$

where IC is the investor's Information Coefficient and σ is unconditional standard deviation of the asset price. Similarly we obtain the investor's conditional forecast variance:

$$Var[r_{t+1} \mid S] = \sigma^2(1 - IC^2)$$

The Information Coefficient features extensively in the portfolio management literature. By way of examples: Coggin and Hunter (1983) discuss empirical IC measurement issues, Dimson and Marsh (1984) refer to the concept in their empirical study of analysts' forecasting ability, Khan *et al* (1996) apply IC in global asset allocation, Khan (1998) discusses its applicability in bond investing, Lee (2000) uses it as a key parameter in a range of analytics pertaining to tactical asset allocation, Herold (2003) discusses its relevance in qualitative forecasting, Satchell and Williams (2007) consider the question of optimal forecasting horizon and Fishwick (2007) analyses the case of multiple models.

For purposes of reconciliation we provide the proposition below as a rule-of-thumb to aid in comparison of hit-rates and Information Coefficients.

Proposition 2.3. *For the case of forecasting a future state which takes only one of the two discrete values* 0 *or* 1*: if forecasts are also restricted to the discrete values* 0 *and* 1 *where the probability of a forecast of* 1 *is denoted by* g*, then the relationship between hit-rate* h *and Information Coefficient (IC) is given by:*

$$IC = \frac{h - g}{1 - g}$$

For example when $g = \frac{1}{2}$ *then* $IC = 2h - 1$.

Proof. See Appendix. □

2.2.3 Probabilistic Forecast Evaluation

Evidently hit-rate and IC constitute parametric measures of association taken directly from multivariate distribution functions and as such they have a natural interpretation as skill measurements. We also note however that there is an extensive literature which considers the measurement of accuracy of forecasters who articulate a *distribution* of possible outcomes. These measurements do not directly parameterise multivariate distributions in their own right but nevertheless we mention them here briefly since we believe they bring useful descriptive power

and with future work there may be potential to incorporate these more directly into our framework. We present two examples below which are discussed in detail in Granger and Pesaran (2000) and Granger and Machina (2006).

We suppose that there are n possible outcomes and f_i is the forecaster's subjective probability of outcome i taking place. We define j as the number of the outcome which actually does take place and then let $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

The *Brier Score* consists of applying a quadratic loss function:

$$BS = \frac{1}{n}\sum_{i=1}^{n}(f_i - \delta_{ij})^2$$

Evidently for the case of a perfect forecaster $BS = 0$.

Alternatively the *Kuipers Score* is well suited to measuring the performance of a forecaster called upon to predict whether a future event will be good or bad (e.g. in the meteorological literature this might be the event of rain or no-rain on a single well-defined time horizon). The measurement is defined as:

$$KS = H - F$$

i.e. it represents the difference between the proportion of correctly-forecast bad events H (the 'hit rate') and incorrectly-forecast good events F (the 'false alarm rate'). In this case an unskilled forecaster who makes random forecasts over time will achieve $KS = 0$.

In the analysis which follows we have chosen to focus on the use of *hit-rate* as a skill measurement within the context of a two-state model. Equivalent analysis for the case of the Information Coefficient (with a continuum of states) is more involved but raises several further subtle insights. We consider this latter case in a separate paper[6] where, despite the incremental complexity, we find several results which are qualitatively similar to those which follow in this paper.

3 A Model with Two Discrete States

3.1 The Model and Equilibrium Solution

Suppose we have K agents, divided into two types (A and B), two states of the world and a single risky asset which pays off one unit in state 1 and zero units in state 0. We denote the actual probability of state 1 by g and the proportion of agents who are type A by f. Each agent receives a type-dependent private signal S^A or S^B and has a hit-rate h^A or h^B which represents the probability of their signal being correct. Therefore for each agent we need to solve the expected utility maximisation problem to obtain the quantity of risky asset which they buy *given* their signal. We denote this quantity x_0 (when $S = 0$) or x_1 (when $S = 1$). We assume the agents have identical exponential utility functions (with common risk aversion λ) and to ease the algebra we make the additional assumption that each agent can borrow and

[6] Satchell and Williams (2010)

Table 1: Equilibrium in equifrequent two-agent two-state model, conditional on each agent having formed their forecasts but before state of the world is revealed

S^A, S^B	P	p
$0,0$	$\frac{1}{M}$	$\frac{1}{1+M}$
$0,1$	$\frac{1}{Q}$	$\frac{1}{1+Q}$
$1,0$	Q	$\frac{1}{1+\frac{1}{Q}}$
$1,1$	M	$\frac{1}{1+\frac{1}{M}}$

lend unlimited amounts at a risk-free rate of zero.[7]

Furthermore all settlement and payoffs take place in the same time period so the assets have the character of futures-type contracts.

$$\mathbb{E}\left[U|S=0\right] = -\frac{1}{\lambda}(1-h)\exp\left[-\lambda x_0(1-p)\right] - \frac{1}{\lambda}h\exp\left[\lambda x_0 p\right]$$
$$\mathbb{E}\left[U|S=1\right] = -\frac{1}{\lambda}h\exp\left[-\lambda x_1(1-p)\right] - \frac{1}{\lambda}(1-h)\exp\left[\lambda x_1 p\right]$$

Optimal position amounts x are easily found by differentiation:

$$\frac{dU_0}{dx_0} = (1-p)(1-h)\exp\left[-\lambda x_0(1-p)\right] - ph\exp\left[\lambda x_0 p\right] = 0$$
$$x_{S=0} = \frac{1}{\lambda}\log\left[\frac{1}{HP}\right]$$
$$\frac{dU_1}{dx_1} = (1-p)h\exp\left[-\lambda x_1(1-p)\right] - p(1-h)\exp\left[\lambda x_1 p\right] = 0$$
$$x_{S=1} = \frac{1}{\lambda}\log\left[\frac{H}{P}\right]$$

where for convenience we refer to the probabilities in terms of *odds*:

$$H = \frac{h}{1-h}$$
$$P = \frac{p}{1-p}$$

and for notational clarity these quantities will often feature in our expressions instead of the underlying probabilities. If the signal $S = 1$, for instance, then the odds ratio $\frac{H}{P}$ can be interpreted as indicating the agent's level of optimism or pessimism relative to the market.

The equilibrium price p is now found by imposing market clearing. We have four cases to consider depending on values of S^A and S^B and the corresponding solutions appear in Table 1 where in the interests of tidy notation we define $M = (H^A)^f(H^B)^{1-f}$ and $Q = (H^A)^f(H^B)^{f-1}$ (and superscripts here index H according to agent type).

[7] While clearly unrealistic this allows us to highlight key details in the model which concern forecast formulation, price determination and utility realisation. Including more complex heterogeneity in wealth and preferences is straightforward and does not significantly alter the basic results we present here.

3.2 Equilibrium Price and Wealth Dynamics

Although our model is inherently static it can be thought of as a transmission mechanism where agents' signals (either 0 or 1) are converted into equilibrium price. In this sense it is apparent that equilibrium price dynamics are driven by the changing nature of agents' beliefs over time. Given knowledge of the agents' skill levels (as embodied in h) along with further details of the probability distribution of forecasts and outcomes we can gain insight into plausible characteristics which the unconditional price distribution may have over time. (Here we interpret time to mean a series of repeated forecasting instances.[8])

In order to fully specify the joint distribution of (S^A, S^B, z) it remains for us to specify the marginal univariate probabilities of S^A and S^B and also the nature of dependency between them. We consider simple extreme examples below but note that more complexity could potentially be incorporated here, for instance via copula-type representations (e.g. Tajar *et al* (2001)).

It is essential to note that the dependence structure assumed between the agents' signals places strict constraints on values which h^A and h^B can take, e.g. it is nonsense for both h^A and h^B to equal unity and still be independent. This topic is considered for general multivariate Bernoulli distributions by Chaganty and Joe (2006); for the trivariate case they find necessary and sufficient conditions on the bivariate marginals such that the overall joint distribution is valid. Here we make the plausible assumption that unconditional signal probabilities exactly equal the actual probability of the state of the world which, in turn, we fix at $\frac{1}{2}$, i.e.

$$\mathbb{P}\left[S^A = 1\right] = \mathbb{P}\left[S^B = 1\right] = g = \frac{1}{2} \tag{2}$$

We further note that a hit-rate of $h = \frac{1}{2}$ $(H = 1)$ represents the *least* level of possible skill (greatest uncertainty) since for $h < \frac{1}{2}$ an agent with signal S will invest identically as he would if receiving the signal $(1 - S)$ and having hit-rate $(1 - h) > \frac{1}{2}$. Therefore we restrict our exploration to cases where $h > \frac{1}{2}$.

(1) In the case of perfect correlation (which we might see in, for instance, a herding scenario) we first note that $H^A = H^B$ and there is in fact zero trade. Hence we trivially find that:

$$\mathbb{E}\left[p\right] = (1-g)\frac{1}{1+M} + g\frac{M}{1+M}$$
$$\mathbb{E}\left[p^2\right] = (1-g)\frac{1}{1+2M+M^2} + g\frac{1}{1+2\frac{1}{M}+\frac{1}{M^2}}$$

and when hit-rates are both equal to $\frac{1}{2}$ $(H = M = 1)$ we have expectation of $\frac{1}{2}$ and variance of 0. As the hit-rate increases towards unity $(H \to \infty)$ the mean price tends to g (the frequency of forecasts equal to 1) and variance $g(1-g)$ as we would expect from a Bernoulli distribution.

From this we note that (quite intuitively) the equilibrium price is a convex combination of a 'bullish' price and a 'bearish' price and the variance depends on both the degree of switching between bullish/bearish states (represented by g) and the profile of agents' skill (upon which

[8] This approach is similar in spirit to techniques employed in the market microstructure literature where the process by which trading orders arrive asynchronously into the market is modeled and the pace of this flow is a more appropriate scale on which to measure time than wall-clock time.

Figure 2: Variance of price with two equifrequent independent forecasters

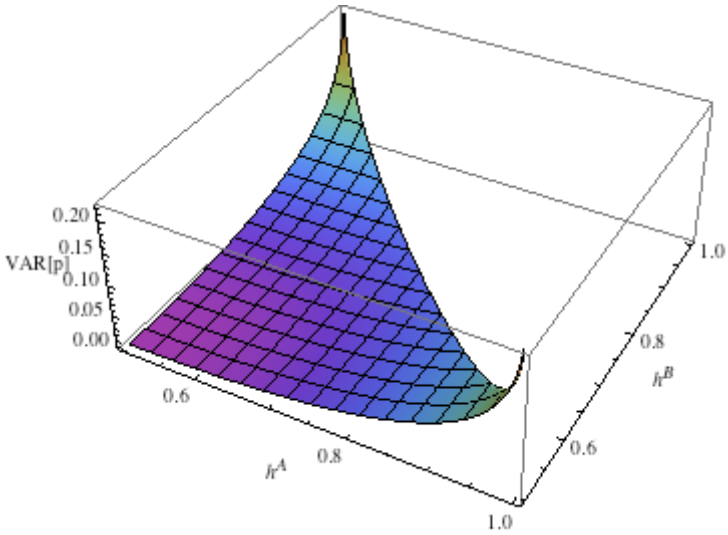

the size of position-taking depends).

(2) We now consider independence. Again for comparison we consider the case of equal skill. Hence we have:

$$\mathbb{E}\left[p\right] = (1-g)^2\frac{1}{1+M} + g(1-g)\frac{1}{1+Q} + g(1-g)\frac{1}{1+\frac{1}{Q}} + g^2\frac{1}{1+\frac{1}{M}}$$

$$\mathbb{E}\left[p^2\right] = (1-g)^2\left[\frac{1}{1+M}\right]^2 + g(1-g)\left[\frac{1}{1+Q}\right]^2 + g(1-g)\left[\frac{1}{1+\frac{1}{Q}}\right]^2 + g^2\left[\frac{1}{1+\frac{1}{M}}\right]^2$$

In this case we must apply the conditions of Chaganty and Joe to determine the admissible range of hit-rates. Given assumptions in (2) these reduce to

$$h^A + h^B < \frac{3}{2}; h^A < 1, h^B < 1$$

demonstrating that if $h^A = h^B$ then the maximum acceptable common hit rate will be $h = \frac{3}{4}$ where $H = M = 3$.

As illustration here we consider the case of equifrequent types (i.e. $f = 0.5$). This simplification conveniently means that when both agents are equally uninformed we have $H^A = H^B = 1 = M = Q$. Once again in this case the mean price is $\frac{1}{2}$ and variance is 0 and as both $h \to 1$ we find $\mathbb{E}[p] \to g$ (mean price simply represents the average of the agents' forecasts) and variance increases. The latter is illustrated in Figure 2 where we carefully restrict the plot to the region of (h^A, h^B) points consistent with independence. It is apparent, however, that variance in price is at its highest with a single skilled forecaster (e.g. $h^A \to 1$ and $h^B = \frac{1}{2}$).

Now we consider an agent's unconditional expected *wealth* over repeated forecasting instances. We compute by the weighted sum of wealth levels which we list in Table 2 from the perspective of type A (the equivalent values for type B agents are straightforward to determine

Figure 3: Distribution of future wealth over repeated forecasting instances (two equifrequent independent agents)

(a) Expected value of wealth of agent A

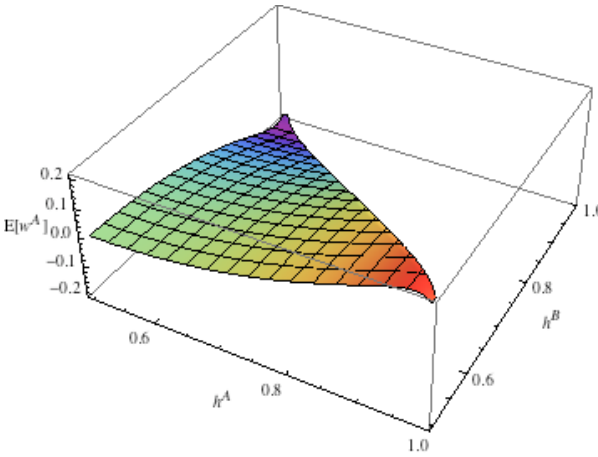

(b) Variance of wealth of agent A

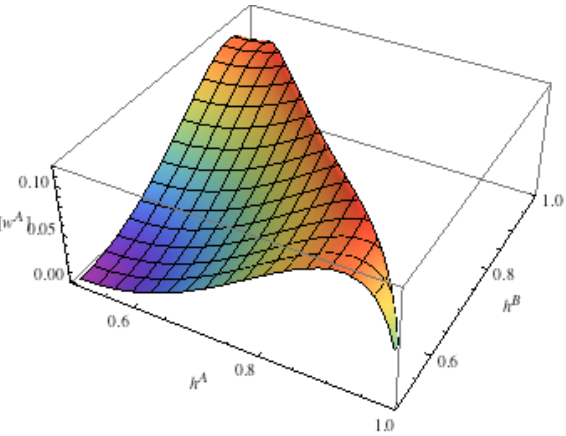

but we omit the details in the interests of brevity).

$$\mathbb{E}\left[w^A\right] = \sum_{(S^A,S^B,z)} w^A(S^A, S^B, z)q(S, z; h^A) \tag{3}$$

We plot the overall expected level of wealth in Figure 3. This plot at once highlights the challenge we face in determining the marginal benefit of skill improvement. Although expected wealth gradually increases for the agent if we hold the other agent's skill level constant, if we increase both together it is apparent that there is an offsetting effect. Although the agent gets better at forecasting future market states, the competition which he faces from the other agent diminishes the returns which can be captured.

We also note that increased skill levels in general lead to increased variance in equilibrium price over repeated forecasting instances. This raises the question of whether the increased variance in returns may also outweigh the improved mean. We investigate this next by computing expected utility.

3.3 Unconditional Expected Utility

In Table 2 we have presented the possible levels of utility which an agent can achieve together with the probability of each state and from this it is straightforward to compute expected utility over repeated realisations of forecasts and states. We therefore plot these levels in Figure 4 against agents' skill levels

It is apparent from the plot that the marginal expected utility of skill is positive for most skill levels. In principle the optimal level of skill can be computed analytically however we do not present that analysis here as the expression is rather unwieldy and we do not find it particularly insightful.

Table 2: *Ex-post* characteristics of equilibrium in equifrequent two-agent two-state model

S^A, S^B, z	$w^A(S^A, S^B, z)$	$U(w^A)$	$q(S, z; h^A)$	$\frac{q_{h^A}(S,z;h^A)}{q(S,z;h^A)}$
$0, 0, 0$	$-\frac{1}{\lambda}\frac{1}{1+M}\log\frac{1}{Q}$	$-\frac{1}{\lambda}Q^{\frac{-1}{1+M}}$	$(1-g)h^A h^B$	$\frac{1}{h^A}$
$0, 0, 1$	$\frac{1}{\lambda}\frac{M}{1+M}\log\frac{1}{Q}$	$-\frac{1}{\lambda}Q^{\frac{M}{1+M}}$	$g(1-h^A)(1-h^B)$	$-\frac{1}{1-h^A}$
$0, 1, 0$	$-\frac{1}{\lambda}\frac{1}{1+Q}\log\frac{1}{M}$	$-\frac{1}{\lambda}M^{\frac{-1}{1+Q}}$	$(1-g)h^A(1-h^B)$	$\frac{1}{h^A}$
$0, 1, 1$	$\frac{1}{\lambda}\frac{Q}{1+Q}\log\frac{1}{M}$	$-\frac{1}{\lambda}M^{\frac{Q}{1+Q}}$	$g(1-h^A)h^B$	$-\frac{1}{1-h^A}$
$1, 0, 0$	$-\frac{1}{\lambda}\frac{1}{1+\frac{1}{Q}}\log M$	$-\frac{1}{\lambda}M^{\frac{Q}{1+Q}}$	$(1-g)(1-h^A)h^B$	$-\frac{1}{1-h^A}$
$1, 0, 1$	$\frac{1}{\lambda}\frac{\frac{1}{Q}}{1+\frac{1}{Q}}\log M$	$-\frac{1}{\lambda}M^{\frac{-1}{1+Q}}$	$gh^A(1-h^B)$	$\frac{1}{h^A}$
$1, 1, 0$	$-\frac{1}{\lambda}\frac{1}{1+\frac{1}{M}}\log Q$	$-\frac{1}{\lambda}Q^{\frac{M}{1+M}}$	$(1-g)(1-h^A)(1-h^B)$	$-\frac{1}{1-h^A}$
$1, 1, 1$	$\frac{1}{\lambda}\frac{\frac{1}{M}}{1+\frac{1}{M}}\log Q$	$-\frac{1}{\lambda}Q^{\frac{-1}{1+M}}$	$gh^A h^B$	$\frac{1}{h^A}$

Figure 4: Expected utility for two equifrequent independent agent types

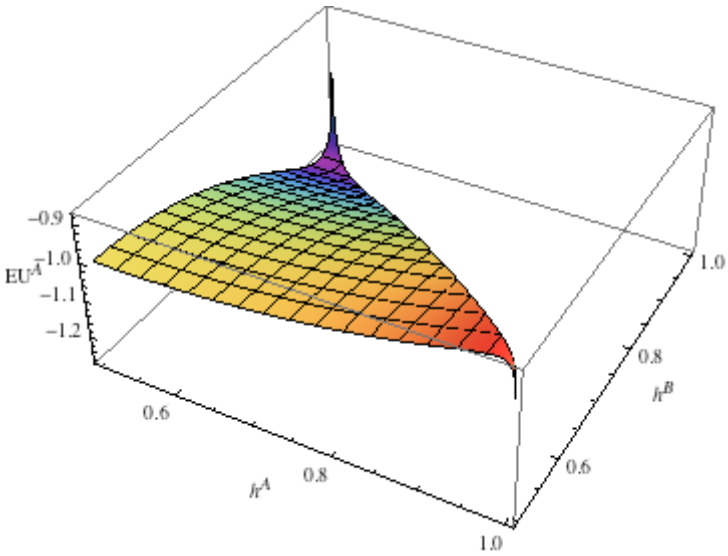

4 Application to Policy Issues

Equipped with knowledge of the distribution of future wealth we can now turn attention to the questions we presented in the introduction.

4.1 Type I: Effects of the Level of Skill

First of all we consider matters regarding the overall level of skill in the market, e.g. effects of changing the level of skill and whether we can say anything about an 'optimal' level of skill. While there are clearly many approaches to tackling this question, the approach which we follow here demonstrates the usefulness of our simple model along with the intuition provided by Proposition 2.1. In order to keep our analysis separate from skill distributional factors (which we consider Type II issues) we begin by assuming an equal number of investors of each type in the market, i.e. the equifrequent case to which we have already made some reference in the previous section. We now address three related questions.

(1) Suppose both agent types always have equal skill and we increase the shared skill level in tandem. What are the consequences of skill adjustment for expected utility?

In this case $h^A = h^B = h$, thus $H^A = H^B = H = M$ and $Q = 1$. The equilibrium properties of this structure are shown in Table 3. We find that expected utility in this case has a convenient analytical form which we reach with the aid of hyperbolic functions as follows:

$$\begin{aligned}
\mathbb{E}[U] &= -\frac{1}{\lambda}\left[h^2 + (1-h)^2 + h(1-h)M^{-\frac{1}{2}} + (1-h)hM^{\frac{1}{2}}\right] \\
\mathbb{E}[U] &= -\frac{1}{\lambda}\left[1 - 2h(1-h) + h(1-h)\left[\exp(-0.5\log M) + \exp(0.5\log M)\right]\right] \\
\mathbb{E}[U] &= -\frac{1}{\lambda}\left[1 + 2h(1-h)\left[\cosh 0.5\log M - 1\right]\right]
\end{aligned}$$

Using the fact that

$$\frac{dM}{dh} = \frac{1}{(1-h)^2}$$

it is straightforward to now derive

$$\frac{d\mathbb{E}[U]}{dh} = -\frac{1}{\lambda}\left[\sinh 0.5\log H - (2-4h) + (2-4h)\cosh 0.5\log H\right]$$

We can compare this expression with the results of applying (1) from section 2, which enables us to break the marginal utility down into intuitive components. First we have the 'sizing' component which we find to be *strictly negative* for $h > 0.5$:

$$\mathbb{E}\left[u_w(w(S,z;h))w_h(S,z;h)\right] = \frac{1}{2\lambda}H^{-\frac{1}{2}} - \frac{1}{2\lambda}H^{\frac{1}{2}} = -\frac{1}{\lambda}\sinh 0.5\log H$$

and separately we have the 'tuning' component which is *strictly positive* for $h > 0.5$:

$$\begin{aligned}
\mathbb{E}\left[u(w(S,z;h))\frac{q_h(S,z;h)}{q(S,z;h)}\right] &= -\frac{1}{\lambda}\left[2h - 2(1-h) + (1-2h)H^{-\frac{1}{2}} + (1-2h)H^{\frac{1}{2}}\right] \\
&= -\frac{1}{\lambda}\left[-(2-4h) + (2-4h)\cosh 0.5\log H\right]
\end{aligned}$$

Table 3: Characteristics of equilibrium in two-agent two-state model where $h^A = h^B = h$

S^A, S^B, z	$w^A(S^A, S^B, z)$	$U(w^A)$	$q(S, z; h)$	$\frac{q_h(S,z;h)}{q(S,z;h)}$	w_h
$0, 0, 0$	0	$-\frac{1}{\lambda}$	$(1-g)h^2$	$\frac{2}{h}$	0
$0, 0, 1$	0	$-\frac{1}{\lambda}$	$g(1-h)^2$	$-\frac{2}{1-h}$	0
$0, 1, 0$	$\frac{1}{2\lambda}\log H$	$-\frac{M^{-\frac{1}{2}}}{\lambda}$	$(1-g)h(1-h)$	$\frac{1-2h}{h(1-h)}$	$\frac{1}{2\lambda h(1-h)}$
$0, 1, 1$	$-\frac{1}{2\lambda}\log H$	$-\frac{M^{\frac{1}{2}}}{\lambda}$	$g(1-h)h$	$\frac{1-2h}{h(1-h)}$	$-\frac{1}{2\lambda h(1-h)}$
$1, 0, 0$	$-\frac{1}{2\lambda}\log H$	$-\frac{M^{\frac{1}{2}}}{\lambda}$	$(1-g)(1-h)h$	$\frac{1-2h}{h(1-h)}$	$-\frac{1}{2\lambda h(1-h)}$
$1, 0, 1$	$\frac{1}{2\lambda}\log H$	$-\frac{M^{-\frac{1}{2}}}{\lambda}$	$gh(1-h)$	$\frac{1-2h}{h(1-h)}$	$\frac{1}{2\lambda h(1-h)}$
$1, 1, 0$	0	$-\frac{1}{\lambda}$	$(1-g)(1-h)^2$	$-\frac{2}{1-h}$	0
$1, 1, 1$	0	$-\frac{1}{\lambda}$	gh^2	$\frac{2}{h}$	0

We plot marginal expected utility for this case in Figure 5, showing the breakdown between the 'sizing' and 'tuning' components. Evidently increasing skill has a negative effect on each agent's utility. This immediately addresses one of our Type I questions: increasing the skill of competing investor types in tandem may not lead to utility benefits overall.

(2) Suppose only one type is skilled. What is the effect on their ex ante utility as skill is increased, assuming the other type remains unskilled?

In this case the analytic expression is more complex and so we resort to inspection of the expected utility plot in Figure 4 and the cross-section of this surface in Figure 6. In contrast to the previous example we find here that marginal expected utility is positive at all but the very highest levels of skill from the perspective of the skilled investor (shown in panel (A)). Note however that the expected utility of the unskilled investor declines (panel (B)). One policy issue here is that targeting particular investor groups for skill improvements can easily have negative spillover effects on others.

We do not go to the extent of decomposing marginal utility in this case but use this as a demonstration of how sensitive the impact of skill improvement is to the exact structure of skill in the market. This raises the important question of how altering individual investors' skill levels may affect overall social welfare. This constitutes our third question.

(3) Is it socially desirable to increase skill levels?

To address this question we adopt the device of a Social Welfare Function (SWF). We refer to Figure 7 which plots both Utilitarian (equally-weighted) and Rawlsian (maximin) social welfare functions for our model market (see Mas-Colell *et al* 1995). An immediate consequence of the assumptions of our model is that the surface is downward-sloping. For all plausible skill levels (i.e. with the exception of extreme levels bordering on perfect foresight) we find that improving skill of one type results in decreasing social welfare overall (if we keep the other type's skill constant). Indeed the optimal social allocation of skill in this economy occurs when both types are perfectly unskilled, hence our original quotation by Thomas Gray.

Figure 5: Marginal expected utility for two identically-skilled independent types

(a) positive 'tuning' component (thick line) and negative 'sizing' component (dotted line)

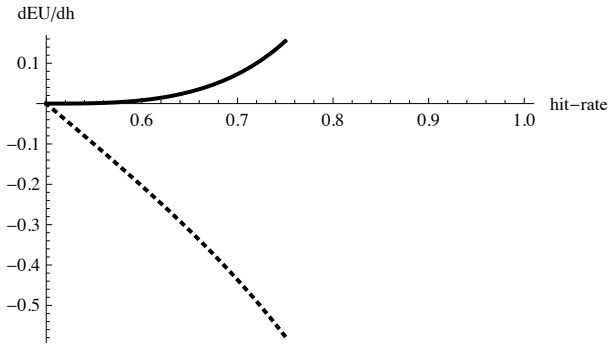

(b) net overall marginal utility

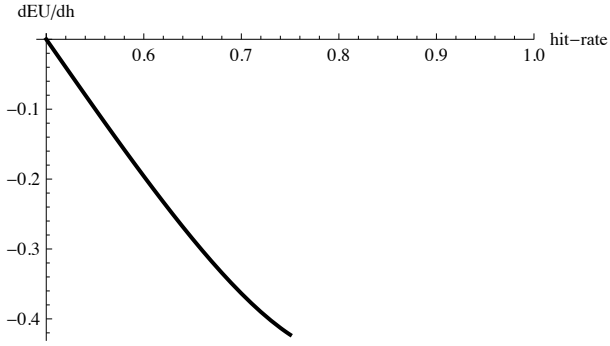

4.2 Type II: Effects of Distribution of Skill

In the preceding section we have demonstrated the ambiguous effects of raising investor skill while keeping the distribution of investor types constant (with an equal weighting between the two types). We now consider the effects of adjusting the proportion of investors of each type while keeping skill levels constant. As a reminder, we are concerned here with questions of the form:

Is it in the interests of unskilled investors to delegate portfolio decisions to professional managers with superior skill? Is there a socially-optimal allocation of capital across differentially-skilled investor types?

A central aspect of our analysis so far was typically whether or not incremental skill-induced price volatility was justified by increased investor utility. However from a policy perspective a drawback of this approach is its assumption that all members of the population are investors, with heterogenenity characterised only by their differential levels of forecasting skill. This might be justified (at least on a *ceteris paribus* basis) if we could demonstrate that the non-investors were an entirely independent group whose welfare was unaffected by events taking place in financial markets. Naturally this is an impossible case to make; in practice we see ubiquitous reminders of the (largely negative) welfare effects which asset price volatility have on non-investors via the real economy.[9]

[9]This theme is prevalent in many strands of the macroeonomics literature: by way of example: Agenor and Aizenman (1999) and Caballero (2000) consider volatility in the context of emerging financial markets, Aghion

Figure 6: Expected utility for skilled agents of each type

(a) $\mathbb{E}(U^A)$ with h^B fixed at 0.5 (dotted), 0.6 (solid) or 0.75 (dashed)

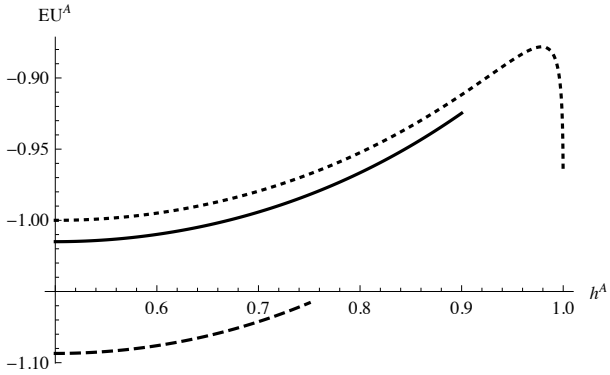

(b) $\mathbb{E}(U^B)$ with h^B fixed at 0.5 (dotted), 0.6 (solid) or 0.75 (dashed)

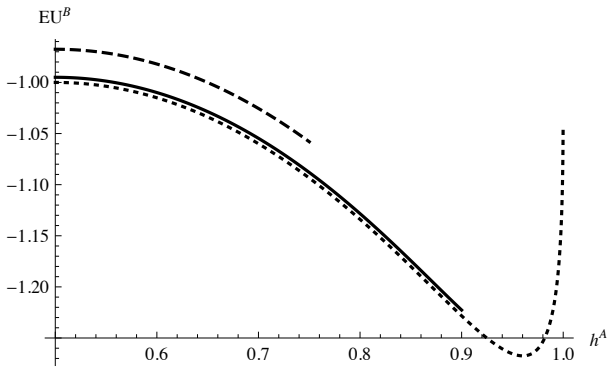

Figure 7: Social Welfare Functions for two skill agents

(a) Utilitarian

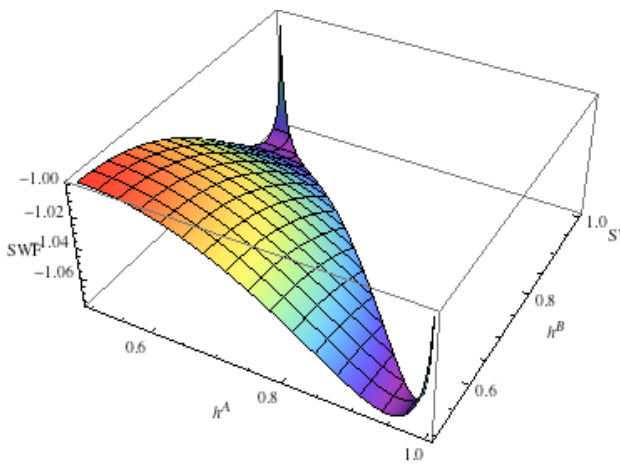

(b) Rawlsian (maximin)

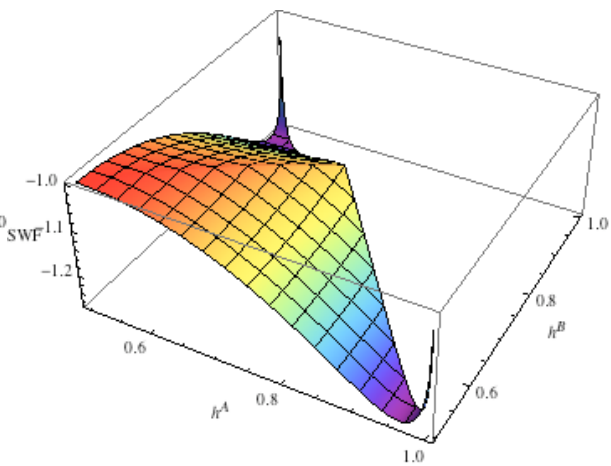

We now recognise that increased volatility is a 'bad' from the point of view of *both* investors *and* non-investors, however the latter group enjoy the offsetting benefit of (potentially) increased wealth due to the trading opportunities which volatility presents. From a quantitative point of view we can apply the analysis of the preceding sections to summarise the average welfare of the *investor* group while the disutility of the *non-investor* group can be represented by the standard deviation of the equilibrium price over repeated forecasting instances.

We illustrate this approach in Figures 8 to 11 where we measure utility and social welfare in terms of certainty equivalent wealth. In Figure 8 we consider the perspective of a skilled type (A) while we keep type B's skill level fixed with $h^B = 0.5$. The proportion of skilled types in the population is denoted by f. The skilled types enjoy maximum welfare when they are as small as possible a group in the overall population. In this case they can derive positive benefit from their superior forecasting ability. However the magnitude of this benefit declines monotonically as they become more frequent (f increases) and we find that when they constitute the entire population ($f = 1.0$) they obtain zero incremental utility. Reminiscent of the case in Section 4.1 we find that the increased welfare of the skilled types comes at the cost of declining welfare for the unskilled types and monotonically increasing volatility.

For comparison in Figure 9 we illustrate the case where type B actually does have some forecasting skill (h^B is fixed at 0.6) but this is nevertheless lower than type A. Whilst at first glance the overall picture is similar we find that in this case the switching of investors from type B to type A (i.e. f increasing from zero) actually reduces volatility at first. This effect occurs because the two types have independent forecasts (by the deliberate construction of our model) and equilibrium price is, loosely speaking, a weighted average of the two types' forecasts.

In Figures 10 and 11 we show society's trade-off between investors' welfare and standard deviation of price (which we use as a measure of the volatility which affects investors and non-investors alike). The plots consist of two separate overlaid components:

(1) In the background of each figure we have sketched convex social indifference curves which are a stylized representation of society's preferences over investor-welfare/volatility pairs; these curves are for indicative purposes only and have not been chosen with any particular preferences in mind, although we feel that convexity is a reasonable assumption on the usual basis that average combinations are (plausibly) preferable to extremes. These curves radiate outwards from the origin, with the optimal point being the origin itself. Hence curves further from the origin represent less desirable outcomes.

(2) Overlaid on the indifference curves are - in each case - three separate curves, each of which is the locus of the feasible investor-welfare/volatility pairs which can be achieved *for given fixed levels of* h^A *and* h^B as the proportion of types (f) is varied. For the sake of argument we call these curves *isoskill* curves since each one is drawn by keeping available skill levels fixed and varying only f. All the *isoskills* have a point in common which is where all investors are type B (the relatively unskilled type whose skill level is always kept the same across *isoskills* in each figure). Note that in Figure 10 (where $h^B = 0.5$ representing no skill) both investors'

and Banerjee (2005) model relationships between volatility and production and Ul Haq *et al* (1996) collect various perspectives on the celebrated Tobin Tax, specifically intended to combat the damaging real effects of spurious volatility.

Figure 8: Investor certainty-equivalent wealth and volatility as the mixture of types is varied from fewer skilled to more skilled types; h_B fixed at 0.5 (i.e. unskilled) and h_A set at 0.7 (dashed line), 0.6 (solid line) and 0.55 (dotted line); f represents the proportion of Type A's in the investor population

(a) Certainty Equivalent Wealth (Type A)

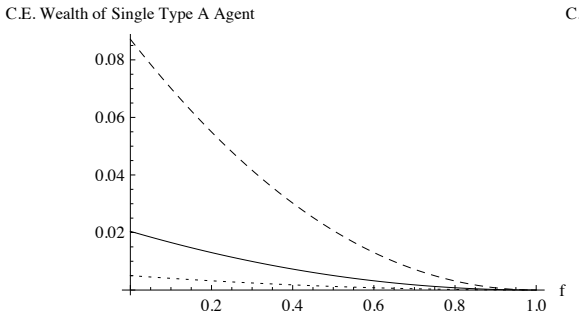

(b) Certainty Equivalent Wealth (Type B)

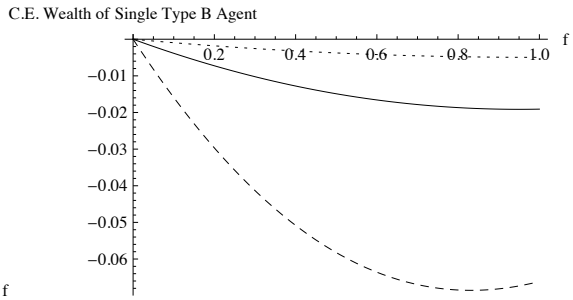

(c) Standard Deviation of Price

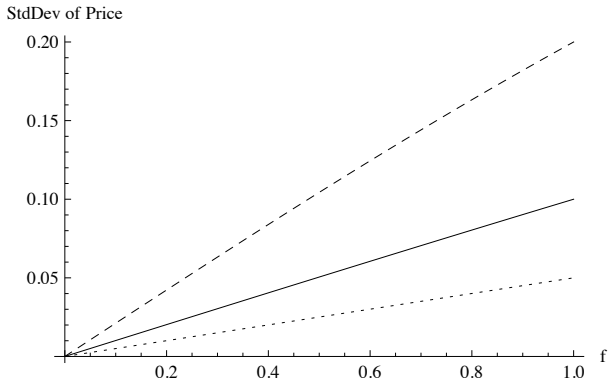

welfare and volatility are zero at this common point, as would be expected.[10]

In view of our previous results it is unsurprising to find that investors' *average* welfare is negative (in certainty equivalent wealth terms) irrespective of the particular levels of skill h^A and h^B, i.e. all points on all *isoskills* lie below the x-axis. As the proportion of (relatively higher-skilled) type A's is increased (f increasing) the market moves along the isoskill away from the common point *initially* in the direction of declining average investor welfare. From the perspective of investors, this trajectory is akin to the 'tragedy of the commons': although they may continue to enjoy positive welfare individually, this declines in magnitude *per capita*.

While the average investor welfare falls, the effect on volatility is interesting. In Figure 10 (where type B's are unskilled with $h^B = 0.5$) we find that volatility increases throughout this process, however in Figure 11 (where type B's are somewhat skilled but less so than type A's) we can see conditions where volatility initially declines along with investor welfare (specifically on the *isoskill* where $h^A = 0.65$ and $h^B = 0.6$). Indeed one might imagine circumstances where society views this relationship as an acceptable trade-off, i.e. tangential to a social indifference curve.

In due course, as f continues to rise, a turning point is reached beyond which average investor welfare increases back towards zero; this comes about since the proportion of less skilled

[10]We might informally describe this as the 'ignorance is bliss' point.

Figure 9: Investor certainty-equivalent wealth and volatility as the mixture of types is varied from fewer skilled to more skilled; h_B fixed at 0.6 and h_A set at 0.8 (dashed line), 0.7 (solid line) and 0.65 (dotted line)

(a) Certainty Equivalent Wealth (Type A)

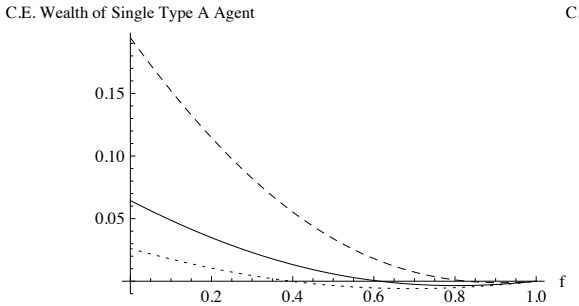

(b) Certainty Equivalent Wealth (Type B)

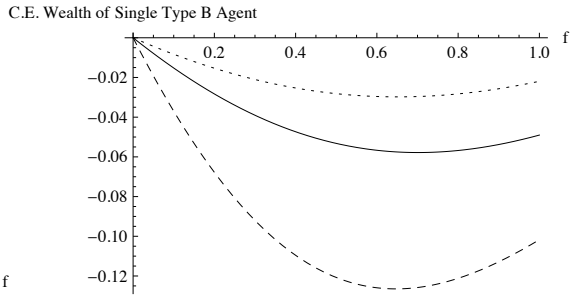

(c) Standard Deviation of Price

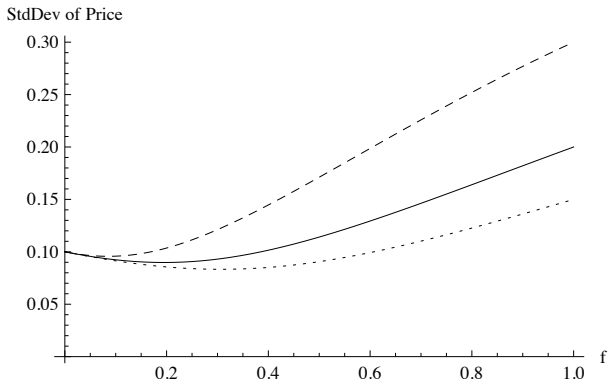

Figure 10: Social trade-off between investors' certainty-equivalent wealth and asset price volatility as proportion of skilled investors increases; h_B is fixed at 0.5 (i.e. unskilled) for all the curves and h_A is fixed at either 0.55 (thinnest curve), 0.6 (middle curve) or 0.7 (thickest curve); arrows indicate direction of increasing proportion of skilled types

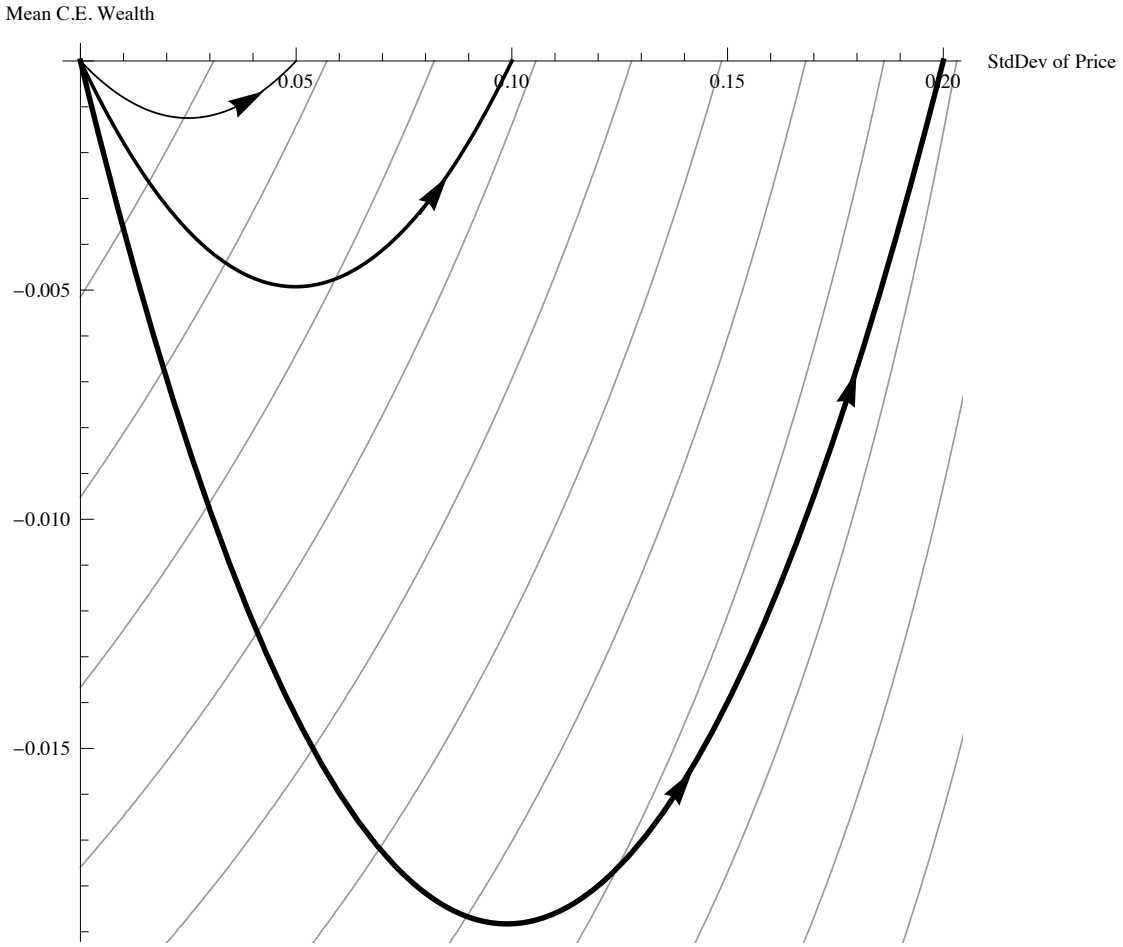

Figure 11: Social trade-off between investors' wealth and asset price volatility as proportion of skilled investors increases; h_B is fixed at 0.6 for all the curves and h_A is fixed at either 0.65 (thinnest curve), 0.7 (middle curve) or 0.8 (thickest curve); arrows indicate direction of increasing proportion of skilled types; point a represents entire population at the lower hit-rate (which is always 0.6), while b, c and d represent entire population at the higher hit-rates (0.65, 0.7 and 0.8 respectively)

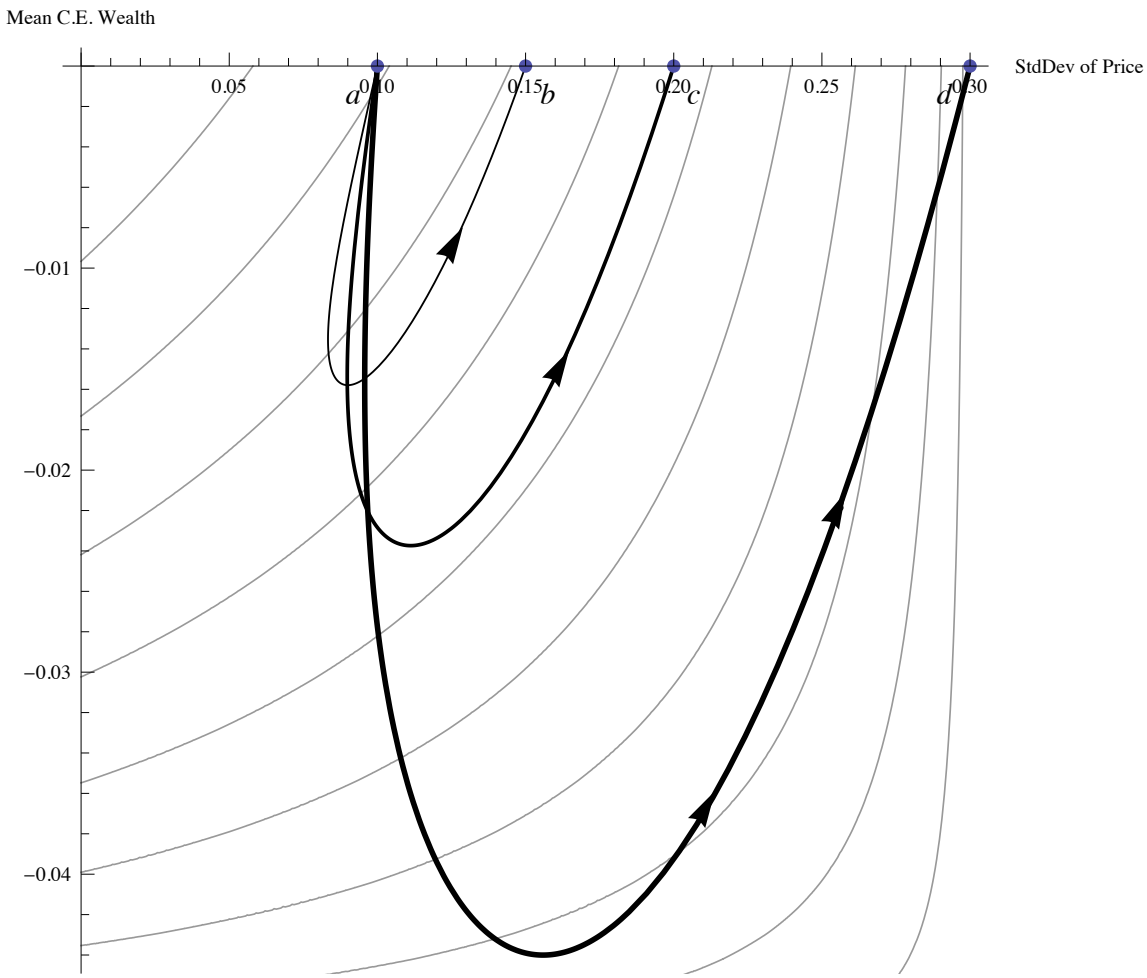

Type B's is a gradually reducing part of the average welfare calculation. Nevertheless, even as average investor welfare improves beyond the turning point, we find that volatility continues to climb. Given the convex shape of the *isoskill* there is therefore a segment representing investor-welfare/volatility outcomes where equal investor-welfare could be achieved but with a lower volatility. This is loosely reminiscent of the concept of efficient versus inefficient portfolios as described by Markowitz (1952) and leads to a notion of a skill-efficient set which will correspond to the solution of a Mean-Variance Social Welfare Function.

We now apply these tools to our Type II policy questions but find that direct answers are hard to come by without more specific knowledge of skill levels and market structure.

Although these examples indicate that a 'first best' allocation of skill across investors would be our 'ignorance is bliss' point, unless society explicitly outlaws the accumulation of investor skill it seems likely that wealth incentives for skill accumulation *on an individual basis* will persist (as depicted, for instance, in Figure 8). This means that to a large extent society may have to take the location of the prevailing *isoskill* curve as exogenously given, depending on the limitations of investors' own human capital rather than any structural arrangements in the economy itself.

We cautiously hypothesize therefore that policy interventions intended to shift the location of the *isoskill* are likely to be slow and costly. In contrast, governments may be able to use faster-acting policy tools to influence the exact location *on* the *isoskill*, e.g. supply side measures aimed at adjusting the structure of the asset management industry, targeted tax incentives to reward specific types of investment, etc..

In a 'second best' world, this highlights the need for policy-makers to attempt to establish where on the *isoskill* their economy is located: if on an 'inefficient' segment then there are clear welfare incentives to reducing the proportion of skilled types (f), thereby reducing volatility while maintaining (or increasing) average investor welfare. However unless the economy can be immediately relocated to the new 'efficient' point then the challenge here is that a gradual reduction in f may lead to reduced investor welfare during the transition period: potentially an unpopular consequence.

If the economy is on the 'efficient' segment there may be a perfectly socially-acceptable trade-off between investor-welfare and volatility; this depends on the precise levels of skill in the market as well as society's preferences, for instance we see a possible examples in Figure 11 but not in Figure 10. It is also possible in this case that an increase in the skilled population (f) may move society onto a higher indifference curve.

5 Concluding Remarks

To some degree governments have the policy tools to influence the distribution of investor skill, albeit rather crudely. Preliminary tentative observations from our analysis here are that circumstances favouring skill improvement may be less commonplace than conventionally believed. We have demonstrated several scenarios where skill improvement has questionable value. Although our models are highly simplified versions of reality it is intriguing that finding robust examples to the contrary is actually rather challenging.

As regards Type I policies: we would argue that the benefits of investor education must be weighed against the effects which higher levels of skilled trading can have on volatility. Clearly investor education which does not increase skill would be a suboptimal social investment, but even increased skill combined with greater volatility may have negative welfare implications for society as a whole as we demonstrated in Section 4.

Considering Type II policies: we have shown that *in appropriate circumstances* it can indeed be socially desirable to switch investors between skill levels. This might be achieved by - for instance - encouraging broader participation in actively-managed funds (including 130/30 and hedge funds). However we have been careful to demonstrate that this depends on the levels and distribution of skill in the market. It is perfectly possible for such policies to have adverse effects on volatility.

Although these results apparently introduce new complexity into policy making, there is great scope for helpful empirical modeling of markets along the lines we have described. In particular a wealth of analysts' forecast data is available from which forecasting skill can be measured, along with reports of institutional Assets Under Management which give some sense of distribution of investors across levels of skill. The frameworks we have introduced in this paper enable the debate to benefit from proper quantitative treatment and raise it from the level of folk myths and value judgements.

This leads us to empirical applications. Assuming access to a suitably broad database of forecasts, trading positions and outcomes, our results suggest a variety of testable hypotheses:

(a) Do investors typically operate close to their optimal skill level? (In section 2 we specialised our proposition to the case of normally-distributed wealth to find simple conditions for the turning point in marginal utility).

(b) Do more complete markets lead to higher skill levels in general? (Our example in section 4 hints at this possibility due to more positive 'tuning' effects).

(c) Do markets segment into separating equilibria of skilled and unskilled types? (This is suggested by the negative marginal utility achieved when all agents increase skill together).

Finally, appropriate analysis of skill parameters should enable society to detect the extreme implausibility (even impossibility) of fraudulent investment offerings which are indeed 'to good to be true'.

The notion of the 'transfer paradox' and 'technology transfer paradox' are already well-established in the general equilibrium literature. These concern the counter-intuitive effects on equilibrium and welfare brought about when endowments or technology are transferred between agents (see for instance Leontief (1936) and Polemarchakis (1983)). We hope that this paper lays some preliminary foundations for further analysis of the 'skill paradox', both theoretical and empirical.

A Derivations and Proofs

A.1 Proofs

A.1.1 Proposition 2.1

Proof. Agent i has unconditional expected utility given by:

$$\int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} u(w(\mathbf{S}, z; \theta)) q(\mathbf{S}, z; \theta) dS^{(1)} dS^{(2)} \cdots dS^{(n)} dz$$

Differentiating with respect to the agent's skill parameter θ^i gives:

$$\int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} u_w(w(\mathbf{S}, z; \theta)) w_\theta(\mathbf{S}, z; \theta) q(\mathbf{S}, z; \theta) dS^{(1)} \cdots dz + \int_{-\infty}^{+\infty} \cdots \int_{-\infty}^{+\infty} u(w(\mathbf{S}, z; \theta)) q_\theta(\mathbf{S}, z; \theta) dS^{(1)} \cdots dz$$
$$= \mathbb{E}\left[u_w(w(\mathbf{S}, z; \theta)) w_\theta(\mathbf{S}, z; \theta)\right] + \mathbb{E}\left[u(w(\mathbf{S}, z; \theta)) \tfrac{q_\theta(\mathbf{S}, z; \theta)}{q(\mathbf{S}, z; \theta)}\right]$$

Since q is a probability density we can think of $\frac{q_\theta(\mathbf{S}, z; \theta)}{q(\mathbf{S}, z; \theta)}$ as if it were a score function; from well-known properties of the score function it therefore follows that

$$\mathbb{E}\left[\tfrac{q_\theta(\mathbf{S}, z; \theta)}{q(\mathbf{S}, z; \theta)}\right] = 0 \tag{4}$$

□

A.1.2 Corollary 2.2

Proof. We deploy Stein's Lemma to rewrite (1) in more familiar terms. This results in a condition for positive marginal utility of skill as follows:

$$\begin{aligned}
\mathbb{COV}\left[u_w(w), w_\theta\right] + \mathbb{COV}\left[u(w), \tfrac{q_\theta}{q}\right] + \mathbb{E}\left[u_w(w)\right]\mathbb{E}\left[w_\theta\right] &> 0 \\
\mathbb{E}\left[u_{ww}(w)\right]\mathbb{COV}\left[w, w_\theta\right] + \mathbb{E}\left[u_w(w)\right]\mathbb{COV}\left[w, \tfrac{q_\theta}{q}\right] + \mathbb{E}\left[u_w(w)\right]\mathbb{E}\left[w_\theta\right] &> 0 \\
\frac{\mathbb{COV}\left[w, \frac{q_\theta}{q}\right] + \mathbb{E}\left[w_\theta\right]}{\mathbb{COV}\left[w, w_\theta\right]} &> R_A
\end{aligned}$$

where

$$R_A \equiv -\frac{\mathbb{E}\left[u_{ww}(w)\right]}{\mathbb{E}\left[u_w(w)\right]}$$

is the coefficient of risk aversion defined by Rubinstein (1973). □

Table 4: Joint distribution of forecast S and actual outcome Q where h denotes hit-rate and g denotes marginal probability of a forecast equal to 1

S	Q	SQ	$\mathbb{P}[S,Q]$
0	0	0	$(1-g)h$
0	1	0	$(1-g)(1-h)$
1	0	0	$g(1-h)$
1	1	1	gh

A.1.3 Proposition 2.3

Proof. Suppose we make a discrete forecast S of either 0 or 1 in respect of a discrete outcome Q which is also either 0 or 1. We define the following probabilities:

$$\mathbb{P}[S=1]=g$$
$$\mathbb{P}[Q=1|S=1]=P[Q=0|S=0]=h$$

We list the possible outcomes and joint probabilities of (S,Q) pairs in Table 4 and hence compute the following:

$$\mathbb{E}[SQ]=gh$$
$$\mathbb{E}[S]=\mathbb{E}[Q]=g$$
$$\mathbb{V}[S]=\mathbb{V}[Q]=g(1-g)$$
$$\mathbb{COV}[S,Q]=gh-g^2=g(h-g)$$
$$IC=\mathbb{COR}[S,Q]=\frac{\mathbb{COV}[S,Q]}{\sqrt{\mathbb{V}[S]\,\mathbb{V}[Q]}}=\frac{h-g}{1-g}$$

For example when $g=\frac{1}{2}$ we have $IC=2h-1$. □

References

Agenor, P.-R. and J. Aizenman (1999): "Volatility and the welfare costs of financial market integration," in *The Asian financial crisis: causes, contagion and consequences*, ed. by P.-R. Agenor, M. Miller, D. Vines, and A. Weber, Cambridge University Press.

Aghion, P. and A. Banerjee (2005): *Volatility and Growth*, Oxford University Press.

Aragon, G. and W. Ferson (2006): "Portfolio Performance Evaluation," *Foundations and Trends in Finance*, 2, 83–190.

Bachelier, L. (1900): "Theory of Speculation," in *The Random Character of Stock Market Prices*, ed. by P. Cootner, MIT Press, 17–78.

Bain, W. (1996): *Investment Performance Measurement*, Woodhead Publishing.

Basak, S. (1997): "Consumption choice and asset pricing with a non-price-taking agent," *Economic Theory*, 10, 437–462.

Caballero, R. (2000): "Aggregate Volatility in Modern Latin America: Causes and Cures," MIT Department of Economics Working Paper No 00-11.

Campbell, J. Y. and A. S. Kyle (1988): "Smart Money, Noise Trading and Stock Price Behavior," NBER Technical Working Papers 0071, National Bureau of Economic Research, Inc.

Chaganty, N. and H. Joe (2006): "Range of correlation matrices for dependent Bernoulli random variables," *Biometrika*, 93, 197–206.

Coggin, T. and J. Hunter (1983): "Problems in Measuring the Quality of Investment Information: The Perils of the Information Coefficient," *Financial Analysts Journal*.

Cox, C. (2006): "Improving Financial Literacy in the United States," Testimony before the U.S. Senate Committee on Banking, Housing and Urban Affairs (May 23rd).

DeLong, J. B., A. Shleifer, L. Summers, and R. Waldmann (1990): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*, 98, 703–738.

Detemple, J. and S. Murthy (1994): "Intertemporal Asset Pricing with Heterogeneous Beliefs," *Journal of Economic Theory*, 62, 294–320.

Dimson, E. and P. Marsh (1984): "An Analysis of Brokers' and Analysts' Unpublished Forecasts of UK Stock Returns," *Journal of Finance*, 39, 1257–1292.

EBF (2009): "EBF Report on Financial Literacy," European Banking Federation (19th May 2009).

Emmerson, C. and M. Wakefield (2003): "Achieving Simplicity, Security and Choice in Retirement? An Assessment of the Government's Proposed Pension Reforms," Briefing Note 36, IFS London.

Fama, E. (1965): "The Behavior of Stock-Market Prices," *Journal of Business*, 38, 34–105.

Fay, S. (1997): *The Collapse of Barings*, W.W. Norton.

Fishwick, E. (2007): "The information horizon - optimal holding period, strategy aggression and model combination in a multi-horizon framework," in *Forecasting Expected Returns in the Financial Markets*, ed. by S.Satchell, Academic Press, chap. 10, 215–226.

Frame, D. (1957): *The Complete Essays of Montaigne*, Stanford University Press.

FSA (2005): "Wider-range Retail Investment Products," FSA Discussion Paper 05/3.

Graham, B. and D. Dodd (1934): *Security Analysis*, McGraw-Hill.

Granger, C. and M. Machina (2006): "Forecasting and Decision Theory," in *Handbook of Economic Forecasting*, ed. by G. Elliott, C. Granger, and A. Timmermann, North-Holland.

Granger, C. and M. Pesaran (2000): "Economic and Statistical Measures of Forecast Accuracy," *Journal of Forecasting*, 19, 537–560.

Grinold, R. and R. Kahn (1999): *Active Portfolio Management*, McGraw Hill, second ed.

Grossman, S. and J. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393–408.

Haq, M. U., I. Kaul, and I. Grunberg (1996): *The Tobin tax: coping with financial volatility*, Oxford University Press.

Herold, U. (2003): "Portfolio Construction with Qualitative Forecasts," *Journal of Portfolio Management*, 30, 61–72.

Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–291.

Knight, J. and S. Satchell (2002): *Performance Measurement in Finance*, Butterworth-Heinemann.

Lee, W. (2000): *Theory and Methodology of Tactical Asset Allocation*, Wiley.

Leland, H. (1999): "Beyond Mean-Variance: Performance Measurement in a Nonsymmetrical World," *Financial Analysts Journal*, 27–36.

Lintner, J. (1969): "The Aggregation of Investors' Diverse Judgements and Preferences in Purely Competitive Security Markets," *Journal of Financial and Quantitative Analysis*, 4, 347–400.

Lowenstein, R. (2001): *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House.

Malkiel, B. G. (1973): *A Random Walk Down Wall Street*, W.W. Norton and Co.

Markopolos, H. (2010): *No One Would Listen: A True Financial Thriller*, Wiley.

Mas-Colell, A., M. Whinston, and J. Green (1995): *Microeconomic Theory*, Oxford University Press.

Obama, B. (2010): "Remarks by the President on Financial Reform," Http://www.whitehouse.gov/the-press-office/remarks-president-financial-reform.

Perotti, E. (1995): "Credible Privatization," *American Economic Review*, 85, 847–859.

Rubinstein, M. (1973): "A Comparative Statics Analysis of Risk Premiums," *Journal of Business*, 46, 605–615.

Satchell, S. and O. Williams (2007): "Optimal Forecasting Horizon for Skilled Investors," in *Forecasting Expected Returns in the Financial Markets*, ed. by S. Satchell, Academic Press.

——— (2010): "On the difficulty of measuring forecasting skill in financial markets," Unpublished.

Shleifer, A. (1999): *Inefficient markets: An introduction to behavioral finance*, Oxford University Press.

Simon, H. (1987): "Rationality in Psychology and Economics," in *Rational Choice*, ed. by R.Hogarth and M.Reder, University of Chicago Press, 25–40.

Tajar, A., M. Denuit, and P. Lambert (2001): "Copula-type representation for random couples with Bernoulli margins," University Catholique de Louvain Institut De Statistique Discussion Paper 118.

Treasury, H. M. (2009): "Reforming financial markets," CM 7667.

Trillin, C. (2009): "Wall Street Smarts," The New York Times (14th October 2009).

Williams, O. (2009): "Asset Pricing With Heterogeneous Skilled Agents," PhD Dissertation, University of Cambridge Faculty of Economics.

Zapatero, F. (1998): "Effects of Financial Innovations on Market Volatility when Beliefs are Heterogeneous," *Journal of Economic Dynamics and Control*, 22, 597–626.