

Optimal dynamic portfolio risk with first-order and second-order predictability

Christian Gollier¹
University of Toulouse

December 20, 2002

¹I thank Bruno Biais, Michael Brennan, Jerome Detemple, Bernard Dumas, Günter Franke, Pascal Maenhout, Richard Stapleton and participants of seminars at Toulouse, INSEAD, Hong Kong, Konstanz and at the 2002 Winter meeting of the Econometric Society (Atlanta) for helpful discussions. Comments are welcomed (gollier@cict.fr).

Abstract

We consider a two-period portfolio problem with predictable assets returns. First-order (second-order) predictability means that an increase in the first period return yields a first-order (second-order) stochastically dominated shift in the distribution of the second period state prices. Mean reversion in stock returns, Bayesian learning, stochastic volatilities and stochastic interest rates (bond portfolios) belong to one of these two types of predictability. We first show that a first-order stochastically dominated shift in the state price density reduces the marginal value of wealth if and only if relative risk aversion is uniformly larger than unity. This implies that first-order predictability increases the initial optimal portfolio risk if this condition is satisfied. A similar result is obtained with second-order predictability under the condition that absolute prudence be uniformly smaller than twice the absolute risk aversion. When relative risk aversion is constant, these two conditions are equivalent. We also examine the effect of exogenous predictability, i.e., when the information about the future opportunity set is conveyed by signals not contained in past asset prices. In this latter case, markets are incomplete.

1 Introduction

Since Merton (1969,1971), it has been recognized that the uncertainty surrounding the future opportunity set of investments should affect optimal portfolios. A recent trend of the finance literature tends to suggest that there is some predictability in the changes of the investment opportunity set. Investors are now advised to consider the possibility to hedge today against bad news about their future opportunity set. This is usually possible if random shifts in this set are statistically related to current and past returns. Several authors, considering various information structures about the stochastic opportunity set, characterized the optimal dynamic portfolio management. Pursuing operational goals, they did not often try to provide an intuition to their results. It is usually mentioned that a change in the opportunity set has a substitution effect and a wealth effect which makes the global effect ambiguous. Beside the fact that this argument is not really explicit, it fails to explain why the sign of the hedging demand depends upon whether relative risk aversion, assuming it to be constant, be smaller or larger than unity. Systematically, the intuitively appealing sign of the hedging demand is obtained when constant relative risk aversion is larger than unity. In other words, the sign of the hedging demand counter-intuition when "the investor is less risk-averse than the logarithmic agent". We believe that operational rules can be helpful only if one can explain their intuition to the decision makers. How could we explain to an investor with a constant relative risk aversion less than unity that mean-reversion in stock returns should make him to reduce, rather than to increase, his demand for stocks? Why is the unity of relative risk aversion critical for the sign of the hedging demand? Why can it be counter-intuitive? Why are we unable to determine the sign of the hedging demand for some specific forms of information structure? The main aim of this paper is to answer these questions.

We will provide a typology of the various concepts of predictability of asset prices. For each type of predictability, we will explain the underlying mechanisms affecting the optimal dynamic portfolio management. By doing so, we will not limit the analysis to preferences exhibiting constant relative (or absolute) risk aversion. It is useful not to rely on such condition in order to be able to disentangle the various effects of the predictability of returns. We consider the simplest dynamic structure to examine the effect of predictability by using a three-date two-period economy. Our objective is

to provide a structure and some intuition to this wide literature, rather than to improve the existing operational rules.

One of the earliest practical example of predictability in future returns is due to Detemple (1986) and Gennotte (1986) where some investors do not know the size of the equity premium, but they learn about it over time by observing stock returns. We can present the following simple parable to illustrate the problem: in a casino, there is an urn of indistinguishable coins, half of which are "good" and half "bad". The good and bad coins land heads with different probabilities that are known a priori. A single coin is picked at random from the urn that is used for n plays of the game. At each play of the game, you choose how much you want to bet. Because the same coin is tossed in the n plays, the first outcomes convey information about the distribution of future payoffs. In this specific example, a good (bad) draw in the first round is a good (bad) news for the future opportunity set. Investors can hedge against the risk on the future opportunity set by reducing their initial stake. The question is whether they should do so. More generally, would one adopt a different risk exposure in play 1 of the game if different coins would be used for the various plays?¹ Assuming constant relative risk aversion, Gennotte (1986) showed that players should hedge the risk by reducing their initial risk exposure if and only if relative risk aversion is larger than unity. Brennan (1998) and Barberis (2000) estimate the effect of learning on the optimal dynamic portfolio strategy.

Predictability can also come from the existence of serial correlation in stock returns. The predictability of stocks and bonds returns has recently been recognized. Barberis (2000) for example estimates significant predictability of US stocks returns. The implied standard deviation of ten-year returns is 23.7 percent, much smaller than the 45.2 percent value implied by the standard deviation of monthly returns.² More precisely, there is mean-reversion in stock returns: a high return of the risky portfolio today implies a lower portfolio return in expectation tomorrow: a good news today means a bad news for the future opportunity set. Because mean-reversion implies that stocks are safer in the long run, the intuition suggests that a long horizon agent should take more risk in the initial stage of the game. Kim and

¹McCardle and Winkler (1992) raised this question to over 200 students and obtained that most people prefer not to bet at first, in order to gather information about the coin.

²See also Poterba and Summers (1988), Campbell (1996), Campbell, Lo and MacKinlay (1997) and Cochrane (2001).

Omberg (1996) and Kogan and Uppal (2000) showed that this is indeed the case if constant relative risk aversion is larger than unity. Campbell and Viciera (1999) and Barberis (2000) numerically estimated this hedging demand. The effect of return predictability on the optimal structure of the initial portfolio is surprising large. For an agent with a relative risk aversion equaling 10 and a ten-year time horizon, the optimal investment in stocks is about 40% of current wealth without predictability. It goes up to 100% when mean-reversion is taken into account. Cochrane (2001) argues that this makes the equity premium puzzle even worse. This also shows that the time horizon of the investor is an important element to take into account to determine the optimal portfolio: when relative risk aversion is constant and larger than unity, a longer time horizon should induce investors to take more risk. The contrary is true if constant relative risk aversion is less than unity.

Predictability also arises from random time-varying volatility of stocks returns. There is ample evidence that large negative returns tend to be associated with increases in volatility over long periods of time (French, Schwert and Stambaugh (1987)). Chacko and Viciera (2000) shows that when changes in volatility are negatively correlated with excess stock returns, there is a negative (resp. positive) hedging demand when constant relative risk aversion is larger (resp. smaller) than unity. They also show that this generates a sizable hedging demand only when low frequency shocks to volatility exhibit enough persistence.

A fourth example of predictability is provided by the analysis of bond portfolios. Suppose that there is a short bond and a long bond, and that there is uncertainty about the future risk free rate. In such a situation, a good news in the second period (high risk free rate) is compatible only with bad news in the first period (low return of the long bond) because of the negative relationship between the interest rate and the price of the long-term bond. Again, because of the negative correlation between the current and the future return of long-term bonds, the intuition is that investors should hedge the future risk by raising their demand for the long-term bond. Brennan, Schwartz and Lagnado (1997) solve numerically a simple version of this model to motivate their paper. They show that the intuition is correct only if the constant relative risk aversion is larger than unity. Campbell and Viciera (2001) calibrated a more complex version of this model to quantify the optimal mix of short-term and long-term bonds.

Our contribution to this literature has two main components. First, we

propose a typology of the stochastic nature of the opportunity set. We show that its effect on optimal portfolios is in fact much different depending upon the type of predictability faced by investors. In spite of their apparent unity, the above-mentioned results relying on whether relative risk aversion is larger or smaller than unity hide a diversity of underlying phenomena. Following Lehmann (1966), we define two types of predictability, namely First-order Stochastic Dominance (FSD) predictability, and Mean-Preserving-Spread (MPS) predictability. The predictability of long-term bonds is in the class of FSD predictability. Predictability coming from either mean-reversion of stock returns, stochastic volatility of these returns, or learning about the size of the equity premium belongs to the class of MPS predictability. It is important to recognize this typology because, contrary to what is suggested in the above-mentioned literature, the necessary and sufficient conditions for an unambiguous effect of predictability have different natures depending upon the class to which they belong.

Our second contribution comes from extending the existing results that are limited to CRRA preferences to any increasing and concave utility function. This is not done for the beauty of the theoretical exercise. Rather, we show that allowing of more flexible functional forms is very helpful to extract intuition. In particular, we show how to disentangle the three possible effects of predictability: substitution effects, wealth effects and precautionary effects. These effects are not apparent with CRRA preferences because of their very specific properties.³

The choice of the initial portfolio risk is driven by the marginal value of wealth at the end of the initial period. This marginal value of wealth depends upon the future opportunity set. If predictability reduces the marginal value of wealth in states where it is large, and if it raises it in states where it is small, predictability has intuitively the same effect than a reduction in risk aversion. It must then raise the optimal portfolio risk. In consequence, we see that the central step of the analysis is to determine the effect of the change in the opportunity set on the marginal value of wealth. We have seen that this change in the opportunity set takes the form of a FSD or MPS shift in the distribution of the state contingent prices. We show that the mechanisms relating the shift in the distribution of prices and the marginal value of wealth

³In particular, the fact that relative prudence equals relative risk aversion plus one in the case of CRRA makes it difficult to separate the effects of risk aversion and prudence.

are much different depending upon whether we examine MPS shifts or FSD shifts. Their respective intuitions are much different.

In the case of a FSD-dominated shift in the distribution of state returns, there is a pure substitution effect and a wealth effect. Because wealth yields smaller returns in all states, its marginal value is smaller (substitution effect). But the shift also reduces the final consumption in each state. This raises the marginal utility of consumption and thereby it raises the marginal value of wealth (wealth effect). The strength of this wealth effect is proportional to the degree of concavity of the utility function. The net effect on the marginal value of wealth is negative only if relative risk aversion is sufficiently large. It happens that this is the case if it is larger than unity.

In the case of a MPS shift, there is a different wealth effect, and a precautionary saving effect. The MPS provides more risk taking opportunities. It therefore raises the *expected* final consumption, otherwise these additional risks would not be taken. This reduces the marginal value of wealth. The strength of this wealth effect is increasing with absolute risk aversion. But if the marginal utility of consumption is convex (prudence), the increased risk taking also raises the marginal value of wealth. This precautionary saving effect is increasing in the degree of absolute prudence. Thus a MPS shift in the distribution of state prices reduces the marginal value of wealth only if the ratio of absolute risk aversion to absolute prudence is sufficiently large. It happens that the critical level of this ratio is 2.

An interesting point is that in the case of constant relative risk aversion (CRRA), these two conditions are equivalent: under CRRA, the ratio of absolute risk aversion to absolute prudence is larger than 2 if and only if relative risk aversion is larger than unity! This is why the existing literature did not need to establish a typology of the predictability of returns in the various models. The conditions for an unambiguous effect of predictability are the same for MPS predictability and FSD predictability when relative risk aversion is constant. This is not the case for other preferences, as when absolute risk aversion is constant (CARA). For CARA preferences, MPS predictability yields an unambiguous hedging demand, but FSD predictability is ambiguous.

In section 2, we present the general model. In section 3, we show how to categorize the different predictability models that have been examined in the past. Section 4 is devoted to the analysis of FSD predictability. The effect of MPS predictability is characterized in section 5.

Up to now, we assumed that the future opportunity set is perfectly correlated with the first-period return, which implies that markets are complete. However, the assumption of perfect correlation is not empirically plausible. In section 6, we examine the effect of this market incompleteness on the hedging demand. More precisely, we examine the case where the information about the future opportunity set is conveyed by signals that are not contained in past returns. We say in this case that we have exogenous predictability. To illustrate, Fama and French (1988) showed that past dividend yields on common stocks is a good predictor of future asset returns. We can infer from the existing literature that investors with constant relative risk aversion should be myopic to exogenous predictability. We examine the effect of exogenous predictability for other preference functionals.

2 Description of the general problem

We consider the simplest model to deal with this question. There are two periods indexed $t = 0, 1$. Consumption takes place only at the end of the second period (investment problem). In each period, there is a complete set of markets contingent to the state of the world that will be revealed at the end of the period. $\tilde{\pi}_t$ is the price kernel in period t . At the end of the first period, not only $\tilde{\pi}_0$ is revealed, but also a signal \tilde{m} is received. Thus, the distribution of $\tilde{\pi}_1$ is conditional to the pair $(\tilde{\pi}_0, \tilde{m})$. This is the predictability hypothesis.

We solve the dynamic problem by backward induction. In the second period, given the realization (π_0, m) observed at the end of the first period, the investor selects his portfolio that maximizes the expected utility of his final consumption:

$$v(z; \pi_0, m) = \max_{c(\cdot)} E[u(c(\tilde{\pi}_1)) | \pi_0, m] \quad (1)$$

$$s.t. \quad E[\tilde{\pi}_1 c(\tilde{\pi}_1) | \pi_0, m] = z, \quad (2)$$

where z is the accumulated wealth at the beginning of the second period, and $v(z; \pi_0, m)$ is the maximum expected utility that one can obtain with wealth z , given information (π_0, m) . The optimal final consumption plan is given by $c(\pi_1, z; \pi_0, m)$. The maximum expected utility that one can extract

in turn from z prior to signal \tilde{m} is thus equal to $V(z; \pi_0)$, which is given by

$$V(z; \pi_0) = Ev(z; \pi_0, \tilde{m}). \quad (3)$$

At the beginning of the first period, the investor selects the portfolio $c_0(\cdot)$ which maximizes V :

$$\max_{c_0(\cdot)} EV(c_0(\tilde{\pi}_0); \tilde{\pi}_0) \quad s.t. \quad E\tilde{\pi}_0 c_0(\tilde{\pi}_0) = w_0, \quad (4)$$

where w_0 is the initial wealth of the investor. Notice that $c_0(\pi)$ is not a level of consumption, but rather the wealth accumulated at the end of the first period in all states of nature whose state price is π .

The aim of this paper is to examine the effect of the predictability hypothesis on the optimal risk exposure in the first period. The amount of risk borne by an agent is measured by the sensitivity of his state-contingent consumption to differences in the corresponding state prices. Of course, in the standard static Arrow-Debreu model, optimal consumption is larger in states with a lower price π . We say that a state-contingent consumption plan $c_0(\cdot)$ is riskier than another plan $c_0^*(\cdot)$ if curve c_0 is steeper than curve c_0^* . Or, more precisely, if curve c_0 single-crosses curve c_0^* from above.⁴ When the distribution of $\tilde{\pi}_0$ is binary, or when the risk on $\tilde{\pi}_0$ is small, this implies that the variance of c_0 is larger than the variance of c_0^* .

In the next three sections, we will assume that $\tilde{\pi}_1$ is independent of m . Section 6 is devoted to the case where $\tilde{\pi}_1$ is correlated to \tilde{m} . When the distribution of $\tilde{\pi}_1$ is correlated to $\tilde{\pi}_0$, we use the following definition.

Definition 1 *There is positive (negative) predictability when a worse news for the first period is a better (worse) news for the future opportunity set.*

A good news for the first period is equivalent to the realization of a cheaper state (π_0 low), since it implies a better first-period return. Under positive predictability, this implies that there opportunity set for the second period is improved. The intuition suggests that the hedging demand should be positive because of the time diversification that this negative correlation

⁴In the static Arrow-Debreu portfolio problem, this is the necessary and sufficient condition for plan c_0 to be selected by an agent that is less risk-averse than the one that selects plan c_0^* . See Franke, Stapleton and Subramanyam (2001) and Gollier (2001).

provides. At this stage, it is useful to observe that there are two ways to improve the opportunity set in period t . The first one is to deteriorate the distribution of $\tilde{\pi}_1$ in the sense of first-order stochastic dominance (FSD). This would imply that all contingent plans become less costly. The second one is to increase the riskiness of the distribution of $\tilde{\pi}_1$ in the sense of Rothschild and Stiglitz. This means introducing mean-preserving spreads (MPS), or zero-mean noises, to $\tilde{\pi}_1$. Because investors can always hedge these noises at an actuarially fair price, adding these noises enlarges the risk-taking opportunity set.

Definition 2 *There is FSD predictability if the conditional distributions of $\tilde{\pi}_1$ can be ordered according to the first-order stochastic dominance order. There is MPS predictability if the conditional distributions of $\tilde{\pi}_1$ can be ordered according to the Rothschild-Stiglitz increase in risk order.*

3 A typology of existing predictability models

All existing models of portfolio management with predictability take the distribution of *returns*, rather than the distribution of state prices, as their primitive random variables. We hereafter show that what really matters to determine the effect of predictability on the demand for assets is whether the conditional distributions of state prices can be ordered according one of the standard stochastic orders. If we want to link this work to the existing literature, we need to show how a shift in the conditional distribution of returns affects the conditional distribution of state prices. To do this, let us consider the simplest case with two states of nature in the second period. Let us denote them a and b , respectively with probability p_a and p_b , $p_a + p_b = 1$. This is an interesting case since the market is completed with only one risk free asset with gross return R_f , and a risky asset whose contingent gross returns are R_a and R_b . We assume without loss of generality that $R_a \leq R_f \leq R_b$ and $ER = p_a R_a + p_b R_b \geq R_f$.

Using the standard arbitrage method, we easily derive the state price π_a and π_b :

$$\pi_a = \frac{R_b - R_f}{p_a R_f (R_b - R_a)} \quad \text{and} \quad \pi_b = \frac{R_f - R_a}{p_b R_f (R_b - R_a)}.$$

Observe that π_a is larger than π_b and that $p_a\pi_a + p_b\pi_b = R_f^{-1}$. Suppose now that R_a , R_b and R_f can be affected by the first period return x of the risky asset. Let us consider in particular the following four examples.

1. $R_i(x) = r_i + g(x)$, $i = a, b, f$, $g' < 0$: this is a case where the excess return is stable, but the risk free rate is stochastic. One can check that

$$p_i \frac{\partial \pi_i}{\partial x} = -g'(x) \frac{p_i \pi_i}{R_f} > 0, \quad i = a, b.$$

Thus, when the excess return is constant, but the risk free rate is stochastic, the distribution of state prices can be ordered according to the FSD order. This is an example of positive FSD predictability that is typical of bond portfolios, since the instantaneous return of the long bond is negatively correlated to the future risk free rate.

2. $R_i(x) = r_i + g(x)$, $i = a, b$, $g' < 0$, and R_f is constant. This is the case of mean-reversion in stock returns. One can again show that

$$p_a \frac{\partial \pi_a}{\partial x} = \frac{g'(x)}{R_f(R_a - R_b)} < 0, \quad \text{and} \quad p_b \frac{\partial \pi_b}{\partial x} = \frac{-g'(x)}{R_f(R_a - R_b)} > 0.$$

Because $R_f = (p_a\pi_a + p_b\pi_b)^{-1}$ is left unchanged, mean-reversion in stock returns implies that an increase in the first period return (good news) yields a mean-preserving contraction in $\tilde{\pi}_1$ (bad news). Thus, mean-reversion in stock returns is an example of positive MPS predictability.

3. $R_a(x) = r_a - p_a^{-1}g(x)$, $R_b(x) = r_b + p_b^{-1}g(x)$, R_f constant, $g' < 0$: when the first period return is larger, volatility is reduced, whereas both the risk free rate and the excess return remain the same. This is the case of stochastic volatility. Observe that

$$-p_a \frac{\partial \pi_a}{\partial x} = p_b \frac{\partial \pi_b}{\partial x} = g'(x) \frac{E\tilde{R} - R_f}{p_a p_b R_f (R_a - R_b)^2} < 0.$$

The observed negative correlation between instantaneous returns and future volatility implies that an increase in the first-period return yields a mean-preserving spread in $\tilde{\pi}_1$, which is a good news. Stochastic volatility is a case of negative MPS predictability.




Figure 1: First-order predictability with negative correlation of returns.

4. In the learning case, R_a , R_b and R_f are constant, but the observed first period return x in period 0 affects state probabilities $p_a(x) = 1 - p_b(x)$ in period 1. Under Bayesian learning, a larger x reduces the likelihood of the bad state a : $p'_a < 0$. Because

$$\frac{\partial \pi_a}{\partial x} = -\frac{p'_a(x)\pi_a(x)}{p_a(x)} > 0, \quad \text{and} \quad \frac{\partial \pi_b}{\partial x} = \frac{p'_a(x)\pi_b(x)}{1 - p_a(x)} < 0,$$

Bayesian learning is an example of negative MPS predictability.

4 First-order predictability

We assume here that \tilde{m} and $\tilde{\pi}_1$ are independent, which implies that $V \equiv v$, and that an increase in π_0 generates a FSD-dominated shift in the distribution of $\tilde{\pi}_1$. Because this implies in particular that the second period risk-free rate $(E\tilde{\pi}_1)^{-1}$ be decreasing in the first period market return, this assumption is not compatible with a model where stock returns are mean-reverting but with a constant risk free rate. In this section, when the first period market return is larger, *all* second period Arrow-Debreu securities have a smaller return. Figure 1 illustrates positive FSD predictability.

As seen in the previous section, we get positive first-order predictability by considering the management of a portfolio of short and long bonds. A downward shift in the distribution of the future risk free rate is compatible with an upward jump in the current return of long bonds. The optimal investment in the short and long bonds must take account of the difference in current returns, but also of the recognition that long bonds have larger short-term returns when future short-term rates are expected to be lower. Thus, long bonds can hedge against changes in the future investment opportunity set. The hedging demand for long bonds is measured by the additional demand from a rational investor in comparison to what would demand a myopic investor treating the long bond as simply another risky asset. Brennan, Schwartz and Lagnado (1997) solve this problem by using a power utility function, together with assuming the pure expectations hypothesis.

4.1 FSD shifts and the marginal utility of wealth

The optimal portfolio at date 0 is governed by the comparison of the marginal value of wealth at date 1 in the various states. Because the first period state affects the distribution of the second period prices, it is useful to start with the analysis of the effect of a change in the distribution of $\tilde{\pi}_1$ on $\partial v/\partial z = \partial V/\partial z$.

At this stage, we simplify the above notation by considering the primitive problem

$$v(z) = \max E[u(c(\tilde{\pi}))] \quad s.t. \quad E[\tilde{\pi}c(\tilde{\pi})] = z. \quad (5)$$

Its first-order condition is written as

$$u'(c(\pi)) = \xi\pi. \quad (6)$$

Using the envelope theorem, we also obtain that the Lagrangian multiplier associated to the budget constraint is equal to the marginal value of wealth: $\xi = v'(z)$. We are interested in determining the effect of a shift in the distribution of $\tilde{\pi}$ on this marginal value of wealth. We have the following result which relies on relative risk aversion R defined as $R(c) = -cu''(c)/u'(c)$.

Proposition 1 *A first-order stochastically dominated shift in the distribution of state prices per unit of probability reduces (resp. increases) the marginal value of wealth if and only if relative risk aversion is uniformly larger (resp. smaller) than unity: $R(c) \geq (\leq) 1$ for all c in the domain of u .*

Proof: The value of $\xi = v'(z)$ can be obtained by combining the first-order condition (6) and the budget constraint. This yields

$$E\tilde{\pi}\phi\left(\frac{1}{\xi\tilde{\pi}}\right) = z, \quad (7)$$

where function ϕ is defined as $\phi(y) = u'^{-1}(1/y)$. Under risk aversion, this function is increasing in its argument. Define $\chi(\xi, \pi) = \pi\phi(1/\xi\pi)$. We have that χ is decreasing in ξ and

$$\frac{\partial\chi}{\partial\pi}(\xi, \pi) = \phi\left(\frac{1}{\xi\pi}\right) - \frac{1}{\xi\pi}\phi'\left(\frac{1}{\xi\pi}\right). \quad (8)$$

Now, observe that relative risk aversion larger than unity is equivalent to the right-hand side of the above equality to be nonnegative. Indeed, because $u'(\phi(y)) = y^{-1}$, $\phi'(y) = -u'^2(\phi(y))/u''(\phi(y))$, we have that

$$\frac{\partial\chi}{\partial\pi}(\xi, \pi) = \phi(y) - \frac{1}{u'(\phi(y))} \frac{-u'^2(\phi(y))}{u''(\phi(y))} = \phi(y) + \frac{u'(\phi(y))}{u''(\phi(y))} \quad (9)$$

$$= T(\phi(y)) [R(\phi(y)) - 1], \quad (10)$$

with $y = 1/\xi\pi$. $T(\phi) = -u'(\phi)/u''(\phi)$ is the degree of absolute tolerance evaluated at wealth ϕ . Thus χ is increasing in π if and only if relative risk aversion is larger than unity.

Suppose that $\tilde{\pi}_b$ is dominated by $\tilde{\pi}_a$ in the sense of first-order stochastic dominance. We have to show that $\xi_b \leq \xi_a$ if $R \geq 1$. Suppose by contradiction that ξ_b be larger than ξ_a . If relative risk aversion is larger than unity, we would have a contradiction, since condition (7) yields

$$\begin{aligned} z &= E\chi(\xi_b, \tilde{\pi}_b) \\ &< E\chi(\xi_a, \tilde{\pi}_b) \\ &\leq E\chi(\xi_a, \tilde{\pi}_a) = z. \end{aligned} \quad (11)$$

The first inequality comes from the assumption that $\xi_b > \xi_a$, whereas the second inequality comes from the definition of first-order stochastic dominance together with the property that χ is increasing in π if $R \geq 1$. This proves that $R \geq 1$ is sufficient for the result. The proof of necessity follows the standard method to prove necessity in the theory of stochastic dominance and is therefore omitted. ■

The intuition of the result is best understood by considering the two-period saving-portfolio problem. The agent faces uncertainty about the state of nature that will prevail in the second period. We assume that there exists a complete set of contingent markets. In period 0, the agent is endowed with wealth w_0 from which he consumes $w_0 - s$ and saves s . Savings in period 1 is used to purchase a portfolio of Arrow-Debreu securities as described above. The problem is to maximize $u_0(w_0 - s) + \beta v(s)$ where v is defined by (5). Obviously, under the concavity of the u_0 function, anything that reduces the marginal value of savings, $v'(s)$, also reduces the optimal savings in period 0. Thus Proposition 1 implies that a FSD-dominated shift in the distribution of state-prices reduces the optimal saving if and only if relative risk aversion is larger than unity.

Here is our intuition about Proposition 1: a FSD-dominated shift in the state-price density has a substitution effect and a wealth effect. The substitution effect comes from the increased efficiency of savings due to the reduction in state prices, i.e., to the increase in the rate of return on savings. It raises the marginal value of wealth. But for a given s , the FSD-dominated shift in the state-price density increases consumption in all states in period 1. Because the marginal utility on consumption is decreasing ($u'' < 0$), this tends to reduce the marginal value of wealth. The wealth effect is increasing with the speed at which u' is decreasing, i.e., it is increasing in the degree of relative risk aversion. Thus, if we want the wealth effect to dominate the substitution effect, we need to put a lower bound to relative risk aversion in order to guarantee that a FSD-dominated shift in the state price density reduces the marginal value of wealth.

4.2 The investment problem with first-order predictability

We can now go back to the three-date two-period investment problem presented in section 2. We assume that

$$\tilde{\pi}_1 \mid \pi_0 =_d g(\pi_0, \tilde{\varepsilon}), \quad (12)$$

where $\tilde{\varepsilon}$ is independent of $\tilde{\pi}_0$. Function g is decreasing in its first argument, which implies that $\tilde{\pi}_1 \mid \pi_0$ is FSD-dominated by $\tilde{\pi}_1 \mid \pi'_0$ whenever π_0 is larger

than π'_0 . The initial portfolio problem in this predictable world is written as

$$\max_{c_0(\cdot)} Ev(c_0(\tilde{\pi}_0); \tilde{\pi}_0) \quad \text{s.t.} \quad E\tilde{\pi}_0 c_0(\tilde{\pi}_0) = w_0, \quad (13)$$

where function v is defined by program (1). The first-order condition for program (13) yields

$$v'(c_0^p(\pi_0); \pi_0) = \lambda^p \pi_0 \quad \forall \pi_0. \quad (14)$$

The superscript p is to refer to the solution for the predictable world. The prime is used for the derivative of the function with respect to its first argument.

We compare the optimal portfolio $c_0^p(\cdot)$ in the predictable world to the optimal portfolio $c_0^u(\cdot)$ in an unpredictable world. In the unpredictable world, the second period distribution of price is independent of the first period state: the investment set is fixed. Without loss of generality, let us suppose that there exists $\hat{\pi}$ such that $\tilde{\pi}_1$ in the unpredictable economy is distributed as $\tilde{\pi}_1 \mid \hat{\pi}$ in the predictable economy. Equivalently, this means that $\tilde{\pi}_1$ is distributed as $g(\hat{\pi}, \varepsilon)$. It implies that the date-0 portfolio problem in the unpredictable world can be written as

$$\max Ev(c_0(\tilde{\pi}_0); \hat{\pi}) \quad \text{s.t.} \quad E\tilde{\pi}_0 c_0(\tilde{\pi}_0) = w_0. \quad (15)$$

It yields the first-order condition

$$v'(c_0^u(\pi_0); \hat{\pi}) = \lambda^u \pi_0 \quad \forall \pi_0, \quad (16)$$

where $c_0^u(\cdot)$ is the optimal portfolio in the unpredictable world and λ^u is the associated Lagrangian multiplier of the budget constraint. Combining the two first-order conditions (14) and (16) yields

$$v'(c_0^p(\pi_0); \pi_0) = \frac{\lambda^p}{\lambda^u} v'(c_0^u(\pi_0); \hat{\pi}) \quad (17)$$

The intuition suggests that FSD predictability should induce risk-averse investors with $R > 1$ to take more risk in period 0. This would be true if $c_0^p(\cdot)$ single-crosses $c_0^u(\cdot)$ from above. Since Merton (1971), we know that this property certainly holds in the case of preferences with constant relative risk

aversion: $u'(c) = c^{-\gamma}$. Indeed, it is well-known that the value function inherits this property from the original utility function: $v'(z; \pi_0) = h(\pi_0)u'(z)$. Condition (17) is then rewritten as

$$u'(c_0^p(\pi_0))h(\pi_0) = \frac{\lambda^p}{\lambda^u}u'(c_0^u(\pi_0))h(\hat{\pi}). \quad (18)$$

Then, define π^* such that $h(\pi^*) = \lambda^p h(\hat{\pi})/\lambda^u$. The above equality is thus equivalent to

$$u'(c_0^p(\pi_0))h(\pi_0) = u'(c_0^u(\pi_0))h(\pi^*). \quad (19)$$

Obviously, if h is decreasing (increasing), c^p single-crosses c^u from above (below) at π^* . But using Proposition 1, we know that h is increasing when relative risk aversion γ is less than 1, whereas h is decreasing when $\gamma > 1$. Combining these remarks yields the following well-known result.

Proposition 2 *Assume FSD predictability with negative serial correlation of returns and constant relative risk aversion. If relative risk aversion is larger than unity, the initial optimal risk exposure is larger than in the unpredictable economy: for any pair (a, b) such that $a \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when relative risk aversion is uniformly less than unity.*

In the limit case, the logarithmic investor has no hedging demand at all, as first proved by Mossin (1968).

When relative risk aversion is not constant, we cannot guarantee the single-crossing property anymore. But we get a similar property, which is presented in the next Proposition. This property states that if c_0^u is already above c_0^p at $\hat{\pi}$, it must remain above c_0^p to the right of $\hat{\pi}$. Reciprocally, if c_0^u is still below c_0^p at $\hat{\pi}$, it must remain below c_0^p to the left of $\hat{\pi}$.

Proposition 3 *Assume FSD predictability with negative correlation of returns. Suppose also that relative risk aversion is uniformly larger than unity. Then, for any pair (a, b) such that $a \leq \hat{\pi} \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when relative risk aversion is uniformly less than unity.*

Proof: Suppose that relative risk aversion is uniformly larger than unity. Using Proposition 1, we know that $v'(c; b) \leq v'(c; \hat{\pi}) \leq v'(c; a)$ for all c .

Suppose by contradiction that $c_0^u(a) \geq c_0^p(a)$, but $c_0^p(b) > c_0^u(b)$. Because v' is decreasing in its first argument, we have that

$$\frac{v'(c_0^p(a); a)}{v'(c_0^u(a); \hat{\pi})} \geq \frac{v'(c_0^u(a); a)}{v'(c_0^u(a); \hat{\pi})} \geq 1. \quad (20)$$

Similarly, we obtain that

$$\frac{v'(c_0^p(b); b)}{v'(c_0^u(b); \hat{\pi})} < \frac{v'(c_0^u(b); b)}{v'(c_0^u(b); \hat{\pi})} \leq 1. \quad (21)$$

This is a contradiction, since the first-order condition (17) implies that

$$\frac{v'(c_0^p(a); a)}{v'(c_0^u(a); \hat{\pi})} = \frac{v'(c_0^p(b); b)}{v'(c_0^u(b); \hat{\pi})}. \quad (22)$$

This concludes the proof. ■

This result is weaker than what we obtained in the CRRA case, as it may be possible in the general case with $R(\cdot) \geq 1$ that c_0^p crosses c_0^u more than once either to the right or to the left of $\hat{\pi}$. We hereafter propose two solutions to this problem. A first solution is provided in the following Corollary, where we make the additional assumption that the initial propensity to consume is unaffected by predictability. This means that $\lambda^p = \lambda^u$. In that case, we get exactly the same single crossing property than in the CRRA case.

Corollary 1 *Suppose that there is FSD predictability with negative correlation and that the initial marginal propensity to consume is the same in the predictable economy and in the unpredictable economy. If relative risk aversion is larger than unity, the initial optimal risk exposure is larger than in the unpredictable economy: for any pair (a, b) such that $a \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when relative risk aversion is uniformly less than unity.*

Proof: Observe that when $\lambda^p = \lambda^u$, c_0^u and c_0^p cross at $\hat{\pi}$. Applying the property of Proposition 3 respectively for $a = \hat{\pi}$ and $b = \hat{\pi}$ yields that the two curves cannot cross anywhere else. And c_0^p must cross c_0^u from above at $\hat{\pi}$. ■

Forcing the preservation of the initial marginal propensity to consume in the two economies can be seen as a ceteris paribus assumption. It requires to

properly select the distribution of $\tilde{\pi}_1$ in the unpredictable economy. A more natural ceteris paribus assumption would be to consider that the distribution of $\tilde{\pi}_1$ in the unpredictable world be equal to the unconditional distribution of $\tilde{\pi}_1$ of the predictable world. However, it is a complex problem to determine the effect of predictability on the initial marginal propensity to consume in that case. An alternative solution is to consider the case where $\tilde{\pi}_0$ has a two-point distribution.

Corollary 2 *Suppose that there is FSD predictability with negative correlation and that there are only two possible states ($\pi_0 = a$ or $\pi_0 = b > a$) in the first period. If relative risk aversion is larger than unity, the initial optimal risk exposure is larger than in the unpredictable economy: $c_0^p(a) \geq c_0^u(a) \geq c_0^u(b) \geq c_0^p(b)$. The opposite result holds when relative risk aversion is uniformly less than unity.*

Proof: By the concavity of v with respect to its first argument, we know that $c_0^u(a)$ is larger than $c_0^u(b)$. Suppose by contradiction that $c_0^p(a) > c_0^u(a)$, which implies that $c_0^p(b) > c_0^u(b)$ in order to satisfy the budget constraints in the first period. This case is prohibited by Proposition 1. ■

After these technicalities, it is noteworthy to recall our main results and their basic intuition. When relative risk aversion is larger than unity, the presence of mean-reversion in the return of capital induces agents with a long time horizon to take more risk. This is intuitive, since taking risk in period 1 is a way to reduce the risk taken in period 0. For example, in the bond portfolio problem presented earlier, investors will raise their demand for the long-term bond due to recognition of the negative correlation of its immediate return with future returns. However, when relative risk aversion is constant and less than one, the hedging demand is negative, thereby inducing investors to reduce his risk exposure at date 0. This is a case where agents want to be wealthier when markets perform badly in the first period, because of the expectation of a larger return on that wealth in the second period. Indeed, remember that the substitution effect (which raises the marginal value of wealth) is larger than the wealth effect (which reduces v') in this case.

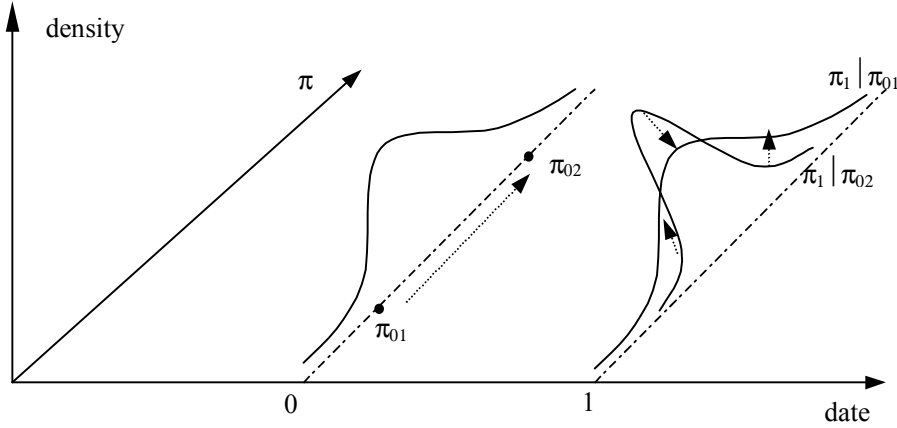


Figure 2: Positive MPS predictability.

5 Second-order predictability

In this section, we examine the effect of second-order predictability, or MPS predictability. We assume that a change in the first period state transforms the distribution of $\tilde{\pi}_1$ according to a mean-preserving spread (MPS). This is possible only if the risk-free rate is unaffected by the first period return.

The classical example of positive second-order predictability is obtained with the one-riskfree-one-risky-asset portfolio problem with a constant risk free rate r and a risky asset whose returns exhibit mean-reversion, as in Kim and Omberg (1996) and Campbell and Viciera (1999). The intuition suggests that positive MPS predictability should induce investors to take more portfolio risk initially, because of the positive time diversification effect. Merton (1973) shows that this is correct when constant relative risk aversion is larger than unity. Positive MPS predictability is depicted in Figure 2.

A second example of MPS predictability comes from a learning process in dynamic portfolio decisions, as studied by Gennotte (1986), Brennan (1998), Brennan and Xia (1999), Barberis (2000) and Xia (2000). In this case we say that we have a negative MPS predictability. The intuition suggests that negative MPS predictability should induce agents to be more prudent in the initial stage in order to hedge against bad news. Gennotte (1986) shows that

this is true when constant relative risk aversion is larger than unity. Negative predictability is also present when there is a negative correlation between the current asset return and the future volatility, in which case more prudence is often suggested (Chacko and Viciara (2000)).

Kandel and Stambaugh (1996), Barberis (2000) and Xia (2000) combine different forms of MPS predictability by looking at models where there is mean-reversion in stock returns, but with some estimation risk on the parameters of the mean-reversion. Barberis (2000) shows that the negative MPS predictability due to the learning process may dominate the positive MPS predictability due to the mean reversion. In that case, the overall MPS predictability reduces the initial demand for stocks when constant relative risk aversion is larger than unity.

5.1 MPS shifts and the marginal value of wealth

As seen before, the analysis of the effect of predictability on the initial portfolio requires first to examine the effect of the shift in the distribution of state prices on the marginal value of wealth in a static framework. Consider again the static portfolio problem (5). We examine the effect of a mean-preserving spread of the state price density on the marginal value of wealth, $\xi = v'(z)$. Let us assume that $\tilde{\pi}_b$ is a mean-preserving spread of $\tilde{\pi}_a$, i.e., that $\tilde{\pi}_b$ is obtained from $\tilde{\pi}_a$ by adding a white noise to it, as defined by Rothschild and Stiglitz (1970). The following result determines the necessary and sufficient condition for $\xi_b \leq \xi_a$. It relies on absolute risk aversion A defined as $A(c) = -u''(c)/u'(c)$ and on absolute prudence P defined as $P(c) = -u'''(c)/u''(c)$, a concept introduced by Kimball (1990).

Proposition 4 *A mean-preserving spread in the of state-price density reduces (resp. increases) the marginal value of wealth if and only if the ratio of absolute prudence to absolute risk aversion is uniformly smaller (resp. larger) than 2: $2A(c) \geq (\leq)P(c)$ for all c in the domain of u .*

Proof: Consider again function χ with $\chi(\xi, \pi) = \pi\phi(1/\xi\pi)$. Differentiating equation (8) yields

$$\frac{\partial^2 \chi}{\partial \pi^2}(\xi, \pi) = \frac{1}{\xi^2 \pi^3} \phi''\left(\frac{1}{\xi\pi}\right). \quad (23)$$

Now, observe that from $\phi'(y) = -u^2(\phi(y))/u''(\phi(y))$, we have that

$$\begin{aligned}\phi''(y) &= \phi'(y) \left[\frac{-2u'(\phi)(u''(\phi))^2 + (u'(\phi))^2 u'''(\phi)}{(u''(\phi))^2} \right] \\ &= [\phi']^2 [-2A(\phi) + P(\phi)].\end{aligned}\tag{24}$$

Thus, ϕ and χ are concave in π when $2A \geq P$. Suppose now that $\tilde{\pi}_b$ is a mean-preserving spread of $\tilde{\pi}_a$, which means that $Ef(\tilde{\pi}_b) \leq Ef(\tilde{\pi}_a)$ for all concave real-valued function f . We have to show that $\xi_b \leq \xi_a$. Suppose by contradiction that ξ_b be larger than ξ_a . If $2A \geq P$, we would have a contradiction, since

$$\begin{aligned}z &= E\chi(\xi_b, \tilde{\pi}_b) \\ &< E\chi(\xi_a, \tilde{\pi}_b) \\ &\leq E\chi(\xi_a, \tilde{\pi}_a) = z.\end{aligned}\tag{25}$$

The first inequality comes from the assumption that $\xi_b > \xi_a$, whereas the second inequality comes from the definition of a mean-preserving spread together with the property that χ is concave in π if $2A \geq P$. This proves sufficiency. Again, the proof of necessity is omitted. ■

The intuition of this result is best understood in the framework of the saving-portfolio problem. If the distribution of state prices per unit of probability undergoes a mean-preserving spread, the consumer will consume more in the first period if $2A > P$.⁵ The intuition of this is easier to extract for the special case where the initial distribution is degenerated. Suppose that we start with $\tilde{\pi}_a = 1$ almost surely. The agent will consume z with certainty in that case. The agent does not face any risk. Consider alternatively state prices $\tilde{\pi}_b$, with $E\tilde{\pi}_b = 1$. It is optimal for the agent to select a consumption plan $c(\cdot)$ that is negatively correlated with π in that new risky environment. This has two effects on the marginal propensity to consume. First, because the agent consumes more in cheaper states, the mean consumption in the second period will be larger than in the risk free environment. Because the agent wants to smooth consumption over time, this induces him to increase his consumption in the first period. The size of this wealth effect is approximately proportional to the degree of aversion to consumption fluctuations

⁵Gollier and Kimball (1996) obtained a similar result by considering the one-riskfree-one-risky-asset portfolio problem. They showed that the introduction of the risky asset raises initial consumption if and only if $2A > P$.

over time, which is given by A . But there is a second effect, which is a *precautionary saving effect*. Because the agent takes risk in the second period, he wants to accumulate precautionary saving in the first period. This negative effect on early consumption is approximately proportional to P , as shown by Kimball (1990). The net effect of the introduction of risk-taking opportunities on early consumption is positive only if the wealth effect is larger than the precautionary effect, i.e., if $2A$ is larger than P . We see that the mechanisms acting in the case of a MPS shift of $\tilde{\pi}$ are quite different than in the case of a FSD shift.

What can we say about whether $2A$ be larger or smaller than P ? This is a difficult question, because of the absence of any estimation of absolute prudence. However, the following result is helpful for solving this problem.

Proposition 5 *Suppose that absolute risk aversion tends to infinity when consumption tends to zero. Then, $2A > P$ uniformly implies that relative risk aversion is uniformly larger than unity, whereas $2A < P$ uniformly implies that relative risk aversion is uniformly less than unity.*

Proof: $A(0) = +\infty$ is equivalent to $T(0) = 0$, where T is the degree of absolute risk tolerance. Observe that $T'(c) = -1 + (P(c)/A(c))$. Thus, $2A(c) > P(c)$ is equivalent to $T'(c) < 1$. Combining this with $T(0) = 0$ implies that $T(c) < c$, or that relative risk aversion is larger than unity. ■

Of course, the opposite results are not true. In particular, if relative risk aversion is uniformly larger than unity, that does not imply that absolute prudence is uniformly less than twice the degree of absolute risk aversion.

In the special case of constant relative risk aversion $2A > P$ is equivalent to $R > 1$. It is widely believed that relative risk aversion is larger than unity. Thus, in this special case, FSD-shifts and mean-preserving spreads of state prices have the same qualitative effect to raise early consumption, in spite of the important differences of the underlying mechanisms leading to these results.

5.2 The investment problem with MPS predictability

In this section, we reexamine the two-period investment problem that we presented in section 4.2. In that section, we assumed that an increase in π_0 yields a FSD-dominated shift in $\tilde{\pi}_1$. We now alternatively assume that

it yields a mean-preserving spread in the distribution of state prices. This is what we called a positive MPS predictability, as in the learning process. Remember that the case of mean reversion in stock returns corresponds to a negative MPS predictability. The next results can be proved by using the same method as in section 4.2. The first Proposition is specific to the HARA case. The second Proposition provides the best characterization of the solution when we only make the assumption that $2A$ is uniformly larger (or smaller) than P . The two Corollaries show that positive MPS predictability raises the optimal portfolio risk when $2A \geq P$ in two special cases: (i) when the initial marginal propensity to consume is preserved, or (ii) when the first period distribution of state prices is binary.

Proposition 6 *Assume positive MPS predictability and HARA preferences: $u'(c) = \left[\eta + \frac{c}{\gamma}\right]^{-\gamma}$. If γ is larger than unity, the initial optimal risk exposure is larger than in the unpredictable economy: for any pair (a, b) such that $a \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when γ is less than unity, or when there is negative MPS predictability.*

Proposition 7 *Consider the case of a positive MPS predictability. Suppose also that $2A \geq P$ uniformly. Then, for any pair (a, b) such that $a \leq \hat{\pi} \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when $2A \leq P$ uniformly, or under negative MPS predictability.*

Corollary 3 *Suppose that there is positive MPS predictability and that the initial marginal propensity to consume is the same in the predictable economy and in the unpredictable economy. If $2A \geq P$ uniformly, the initial optimal risk exposure is larger than in the unpredictable economy: for any pair (a, b) such that $a \leq b$, $c_0^u(a) \geq c_0^p(a)$ implies $c_0^u(b) \geq c_0^p(b)$. The opposite result holds when $2A \leq P$, or when there is negative MPS predictability.*

Corollary 4 *Suppose that there is positive MPS predictability and that there are only two possible states ($\pi_0 = a$ or $\pi_0 = b > a$) in the first period. If $2A \geq P$ uniformly, the initial optimal risk exposure is larger than in the unpredictable economy: $c_0^p(a) \geq c_0^u(a) \geq c_0^u(b) \geq c_0^p(b)$. The opposite result holds when $2A \leq P$ uniformly, or under negative MPS predictability.*

It is noteworthy that Proposition 6 is more general than Proposition 2 since the former covers all HARA functions whereas the latter covers only

the subset of CRRA functions ($\eta = 0$). In particular, Proposition 6 covers the case with constant absolute risk aversion (CARA) that has been examined by Merton (1971). When absolute risk aversion is constant, the ratio of absolute risk aversion to absolute prudence is equal to unity, which implies that the mean reversion of stock returns induces long term CARA investors to increase their demand for stock. For the sake of comparison, notice that the effect of FSD predictability is ambiguous in the CARA case, since relative risk aversion goes from zero to infinity when wealth increases. It cannot be uniformly larger than unity.

We also conclude that there is some fallacy in claiming that positive MPS predictability raises the demand for stock if investors are more risk-averse than log utility investors, as stated for example by Brennan (1998) and Barberis (2000). This statement is true only when relative risk aversion is assumed to be constant. When relative risk aversion is not constant,⁶ the relevant condition is that the ratio of absolute risk aversion to absolute prudence be larger than 2. The intuition of this condition is easy to understand. Under this condition, a low return in the first period reduces the large marginal value of wealth due to the mean-preserving spread in $\tilde{\pi}_1$, because the precautionary effect of the increased risk-taking is not large enough to compensate for the wealth effect. The value function is therefore in a sense less concave in the first period, thereby inducing investors to increase their demand for stocks. It is not possible to get this intuition when, limiting the analysis to power utility functions, we end up with the condition that "investors must be more risk-averse than log utility investors".

5.3 A simple numerical illustration with learning

To illustrate the effect of second-order predictability on the optimal risk-taking strategy, let us examine the learning problem of the gambler who does not know the probability of heads (H) and tail (T). Consider an urn of coins, half of which are "good" and half "bad". The good coins land heads with probability $0.5 + k$, whereas the bad coins land heads with probability $0.5 - k$, $k \in [0, 0.5]$. A single coin is picked at random from the urn at the beginning of the game, and is not changed afterwards. At each play of the

⁶Detemple, Garcia and Rindisbacher (2001) are among the few who examines the structure of optimal portfolios when investors have wealth-dependent relative risk aversion.

game, the gambler choose how much he wants to bet (positive or negative). If the coin lands heads, the gambler gets 300% ($a = 2$) of the initial stake back. But if it lands tail, he loses his stake ($b = -1$). We assume that the agent has a constant relative risk aversion equaling 2. This problem is formally equivalent to the one-riskfree-one-risky-asset model with negative MPS predictability.

Because of constant relative risk aversion, we know that myopia is optimal in the unpredictable world, i.e., in a world where a different coin is taken for each repetition of the game. Because the agent has the choice between two assets in this game with two states of the world, this is an Arrow-Debreu portfolio problem (5). The distribution of $\tilde{\pi}_0$ is $(2/3, 1/2; 4/3, 1/2)$. Indeed, the probability of H and T are 1/2. The price of the contingent claim associated to H is 1/3, because of the 300% payoff per dollar invested in that state. The state price per unit of probability is thus 2/3. A similar computation yields a state-price per unit of probability 4/3 for state T. For $\gamma = 2$, it is straightforward to check that the optimal stake when one may gamble only on the first draw is equal to 12.1% of total wealth. The same strategy would be followed by the long horizon gambler when different coins are used for each play, because we know that myopia is optimal in that case.

Suppose alternatively that the agent can gamble on the first two draws of the same coin. Should he revise how much to invest on the first draw? Using Bayes' rule, one can check that the updated probabilities of H are $0.5 + 2k^2$ and $0.5 - 2k^2$ respectively if H or T is observed at the first draw. For example, for $k = 0.25$, the underlying conditional distributions of $\tilde{\pi}_1$ are

$$\begin{aligned} H & : & \tilde{\pi}_1 \mid \pi_0 = 2/3 & \sim & (0.533, 5/8; 1.777, 3/8) \\ T & : & \tilde{\pi}_1 \mid \pi_0 = 4/3 & \sim & (0.888, 3/8; 1.066, 3/8). \end{aligned}$$

Obviously, the increase in π_0 from 2/3 to 4/3 generates a mean-preserving contraction in the conditional distribution of $\tilde{\pi}_1$. From Proposition 6, this predictability should induce a reduction in the optimal risk exposure, i.e., a reduction in the bet. Using Mathematica[©], we obtain that the optimal initial bet is 9.49% of initial wealth, which is less than the original 12.1% of the unpredictable economy.

The intensity of predictability is here measured by parameter k : when $k = 0$, there is no predictability at all, as all coins are unbiased. On the contrary, when $k = 0.5$, there is full predictability, since the observation of




Figure 3: Optimal stake as a percentage of wealth at the initial draw of the coin.

the first draw eliminates the uncertainty for the second draw. In Figure 3, we represented the optimal percentage of wealth gambled on the first draw as a function of the degree of predictability. We see that a low degree of predictability has almost no effect on the optimal initial risk exposure. But introducing a sizeable degree of predictability ($k > 0.4$) should induce the gambler not to take any risk at the initial stage of the game. It is better to wait the next round to get information about the type of the coin.

6 Exogenous predictability and incomplete markets

Up to now, we assumed that the future opportunity set is perfectly correlated to the first-period return. In this section, we take the extreme view that the source of predictability is exogenous to financial markets, i.e., that there is no serial correlation in assets returns. This is a model where the investor does not know the true distribution of returns. But he obtains exogenous information about it over time. This information is not conveyed by the observation of past returns. For example, the exogenous information about future returns could be the outcome of an election, the occurrence of

a technical innovation, or some signal contained in recent dividends.⁷ Signal \tilde{m} summarizes this information. How does the existence of such uncertainty about the distribution of future returns affect the optimal portfolio prior to the observation of the information? Because \tilde{m} and $\tilde{\pi}_0$ are not statistically related, there is no way to hedge the risk of change of the opportunity set. Markets are thus incomplete. It is thus a priori unclear how exogenous predictability should affect the first period portfolio.

We solve this problem by backward induction. At date $t = 1$, the investor observes signal m . He then solves his second period portfolio problem:

$$v(z; m) = \max E [u(c(\tilde{\pi}_1)) | m] \quad s.t. \quad E [\tilde{\pi}_1 c(\tilde{\pi}_1) | m] = z. \quad (26)$$

The maximum expected utility that one can extract in turn from z prior to signal \tilde{m} is thus equal to $V(z)$, which is given by

$$V(z) = Ev(z; \tilde{m}). \quad (27)$$

At the beginning of the first period, the investor selects the portfolio $c_0(\cdot)$ which maximizes $EV(c_0(\tilde{\pi}_0))$ subject to the budget constraint $E[\tilde{\pi}_0 c_0(\tilde{\pi}_0)] = w_0$. We compare the optimal first period portfolio risk to the one that would be optimal in an unpredictable economy where the distribution of $\tilde{\pi}_1$ is signal independent and is equal to the unconditional distribution of $\tilde{\pi}_1$ of the predictable world.

To solve this problem, we just have to examine how the concavity of the single-variable function V is affected by this uncertainty. Fully differentiating the first order condition $u'(c(\pi_1, z; m)) = \xi(z; m)\pi_1$ with respect to z yields

$$\frac{\partial c(\pi_1)}{\partial z} = T(c(\pi_1)) \frac{-\partial \xi / \partial z}{\xi},$$

where $T(c) = -u'(c)/u''(c)$ is the degree of absolute tolerance. We know from the budget constraint that $E\tilde{\pi}_1 \frac{\partial c}{\partial z} = 1$, which implies that the degree of absolute risk tolerance of the value function equals

$$T_v(z; m) = -\frac{v'(z; m)}{v''(z; m)} = -\frac{\xi(z; m)}{\frac{\partial \xi}{\partial z}(z; m)} = E[\tilde{\pi}_1 T(c(\tilde{\pi}_1, z; m)) | m]. \quad (28a)$$

⁷Fama and French (1988) showed that the dividend yield on common stocks is a good predictor of stock returns. Several recent papers examine the effect of this kind of predictability on the optimal dynamic portfolio management. See for example Campbell (1996), Brennan, Schwartz and Lagnado (1997) and Barberis (2000).

Suppose first that u is CRRA, i.e., that $T(c) = c/\gamma$. It yields

$$T_v(z; m) = \gamma^{-1} E [\tilde{\pi}_1 c(\tilde{\pi}_1, z; m) | m] = z/\gamma.$$

Notice that it is independent of m . Since V is a weighted sum of the same function v , we conclude that the absolute risk tolerance of the value function V is independent of the structure of future information. We conclude that the information structure has no effect on the optimal portfolio ex ante in the CRRA case. The CRRA agent is myopic to exogenous predictability. The treatment of the more general HARA case is more tricky. When $T(c) = \eta + c/\gamma$, condition (28a) implies that

$$T_v(z; m) = E [\tilde{\pi}_1 | m] T\left(\frac{E [\tilde{\pi}_1 c(\tilde{\pi}_1, z; m) | m]}{E [\tilde{\pi}_1 | m]}\right) = E [\tilde{\pi}_1 | m] T\left(\frac{z}{E [\tilde{\pi}_1 | m]}\right).$$

We see that we need to assume that signal \tilde{m} does not affect the risk free rate $[E\tilde{\pi}_1]^{-1}$ to obtain optimal myopia to exogenous predictability in the HARA case.

Proposition 8 *CRRA investors are myopic to exogenous predictability. HARA investors are myopic to exogenous predictability which does not affect the risk free rate.*

It is noteworthy that several authors as Campbell (1996) and Barberis (2000) derived optimal portfolio strategies for CRRA investors when future returns are predictable by using the dividend yield as a predictor. The above Proposition claims that information contained in future dividends have no effect on the portfolio strategy of these investors. However, because the dividend yield is correlated with realized returns, using the dividend yield as a predictor combines exogenous predictability to a positive MPS predictability. Only this second form of predictability matters for CRRA investors to determine the hedging demand.⁸ The existence of the potential innovation contained in the dividend yield that is not in realized returns has no impact on the hedging demand.

⁸When the level of the exogenous signal follows an autoregressive process $\tilde{m}_{t+1} = a + b\tilde{m}_t + \tilde{\varepsilon}_t$, the level of the signal at $t = 0$ may affect the optimal initial portfolio. But it is not different from what a myopic agent would do.

In the remaining of this section, we explore exogenous predictability when preferences are not HARA. We claim that convex (concave) absolute risk tolerance is necessary for exogenous predictability to raise (reduce) the optimal exposure to risk ex ante. To show this, we consider the special case where the unconditional distribution is degenerated to a single value which is normalized to unity. This is equivalent to say that the unconditional expected excess return of stocks is zero. It implies that in the unpredictable world, investors do not take any risk in the second period, and $V^u(z) = u(z)$. In the predictable world, the distribution of $\tilde{\pi}_1$ will generally not be degenerated ex post. Assuming the convexity of absolute risk tolerance on consumption, Jensen inequality yields

$$T_v(z; m) = E[\tilde{\pi}_1 T(c(\tilde{\pi}_1, z; m)) | m] \geq E[\tilde{\pi}_1 | m] T\left(\frac{E[\tilde{\pi}_1 c(\tilde{\pi}_1, z; m) | m]}{E[\tilde{\pi}_1 | m]}\right).$$

Consider the case of an exogenous predictability which preserves the risk free rate: $E[\tilde{\pi}_1 | m] = 1$ for all m . It implies that

$$T_v(z; m) \geq T(z).$$

This is true for all m . Because V is a weighted average of functions $v(\cdot; m)$ whose degree of absolute risk tolerance is larger than T , the absolute risk tolerance must be larger than T , the degree of absolute risk tolerance of the value function in the unpredictable world. It implies that this predictability raises the initial portfolio risk.

Proposition 9 *Consider the case of exogenous predictability which preserves the risk free rate. Suppose also that the unconditional expected excess return of stocks is zero. This exogenous predictability raises (resp. reduces) the initial optimal portfolio risk if the absolute risk tolerance on final consumption is convex (resp. concave).*

Treich (1997) obtained a similar result in the special case of small standard portfolio risks. Gollier (2000) and Guiso and Paiella (2000) provide arguments for and against the convexity of absolute risk tolerance with respect to the level of consumption. The intuition for the above result is as follows: because in the unpredictable world, the investor does not take any portfolio risk in the second period, introducing predictability in this model is equivalent to enlarging his time horizon of risky investment. Gollier and Zeckhauser (1999) showed that doing so makes agent more risk prone if and only if their absolute risk tolerance is convex.

7 Conclusion

The analysis of dynamic risk taking is a fascinating area of research. How does the opportunity to take risk in the future affect the willingness to take risk today? This question is particularly difficult when there is time variation in the distribution of assets returns. All papers addressing this question since the seminal work by Merton (1973) focussed on the special case of power, logarithm and exponential utility functions. By relaxing this constraint, we have been able to provide additional insights to the solution of this problem. First, we have shown that the kind of optimal dynamic risk management depends upon whether the predictability of future returns is of the first stochastic order or of the second stochastic order. In the case of FSD predictability, we only need to know whether relative risk aversion is uniformly smaller or uniformly larger than unity to determine the sign of the hedging demand for risk. But if we have MPS predictability, we need to compare the intensity of absolute prudence to absolute risk aversion to get an equivalent comparative statics property. It is only by chance that the two conditions are equivalent in the standard case of constant relative risk aversion. They differ in all other cases, as preferences exhibit constant absolute risk aversion. More importantly, our analysis allows us to provide simple intuitions for the comparative statics of the various forms of predictability.

References

- Arrow, K.J., (1963), Liquidity preference, Lecture VI in " *Lecture Notes for Economics 285, The Economics of Uncertainty*", pp 33-53, undated, Stanford University.
- Barberis, N., (2000), Investing for the long run when returns are predictable, *Journal of Finance*, 55, 225-64.
- Brennan, M.J., (1998), The role of learning in dynamic portfolio decisions, *European Finance Review*, 1, 295-306.
- Brennan, M.J., E.S. Schwartz and R. Lagnado, (1997), Strategic asset allocation, *Journal of Economic Dynamics and Control*, 21, 1377-1403.
- Brennan, M.J., and Y. Xia, (1999), Assessing asset pricing anomalies, Discussion paper #19-99, The Anderson School at UCLA.
- Campbell, J.Y., (1996), Understanding risk and return, *Journal of Political Economy*, 104, 298-345.
- Campbell, J.Y, A.W. Lo and A. C. MacKinlay, (1997), *The Econometrics of Financial Markets*, Princeton University Press, Princeton.
- Campbell, J., and L. Viciera, (1999), Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics*, 114, 433-95.
- Campbell, J.Y., and L. Viciera, (2001), Who should buy long-term bonds?, *American Economic Review*.
- Chacko, G., and L.M. Viciera, (2000), Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets, mimeo, Harvard Business School.
- Cochrane, J.H., (2001), *Asset pricing*, Princeton University Press, Princeton.
- Detemple, J.B., (1986), Asset pricing in an economy with incomplete information, *Journal of Finance*, 61, 383-392.
- Detemple, J.B., R. Garcia and M. Rindisbacher, (2001), A Monte Carlo method for optimal portfolios, mimeo, *Journal of Finance*, forthcoming.

- Fama, E.F. and K.R. French, (1988), Dividend yields and expected stock returns, *Journal of Financial Economics*, 22, 3-26.
- Franke, G., R.C. Stapleton and M.G. Subrahmanyam, (2001), Standard risk aversion and the demand for risky assets in the presence of background risk, unpublished paper, University of Konstanz.
- French, K.R., G.W. Schwert, and R.F. Stambaugh, (1987), Expected stock returns and volatility, *Journal of Financial Economics*, 19, 3-29.
- Genotte, G., (1986), Optimal portfolio choice under incomplete information, *Journal of Finance*, 41, 733-749.
- Gollier, C., (2000), Wealth inequality and asset prices, *Review of Economic Studies*, forthcoming.
- Gollier, C., (2001), *The economics of risk and time*, MIT Press, forthcoming.
- Gollier, C., and M.S. Kimball, (1996), Toward a systematic approach to the economic effects of uncertainty: characterizing utility functions, Discussion paper, University of Michigan.
- Gollier, C. and R.J. Zeckhauser, (1999), Horizon length and portfolio risk, Discussion Paper, University of Toulouse.
- Guiso, L. and M. Paiella, (2000), Risk aversion, wealth and financial market imperfections, Discussion Paper, University of Sassari, Italy.
- Kandel, S., and R. Stambaugh, (1996), On the predictability of stock returns: An asset allocation perspective, *Journal of Finance*, 51, 385-424.
- Kim, T.S., and E. Omberg, (1996), Dynamic nonmyopic portfolio behavior, *Review of Financial Studies*, 9, 141-61.
- Kimball, M.S., (1990), Precautionary savings in the small and in the large, *Econometrica*, 58, 53-73.
- Kogan, L., and R. Uppal, (2000), Risk aversion and optimal portfolio policies in partial and general equilibrium economies,

<http://finance.commerce.ubc.ca/~uppal/papers.html>.

- Lehmann, E.L., (1966), Some concepts of dependence, *Annals of Mathematical Statistics*, 37, 1137-1153.
- McCardle, K.F., and R.L. Winkler, (1992), Repeated gambles, learning, and risk aversion, *Management Science*, 38, 807-818.
- Merton, R.C., (1969), Lifetime portfolio selection under uncertainty: The continuous time case, *Review of Economics and Statistics*, 51, 247-257.
- Merton, R.C., (1971), Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory*, 3, 373-413.
- Merton, R.C., (1973), An intertemporal capital asset pricing model, *Econometrica*, 41, 867-887.
- Mossin, J., (1968), Optimal multiperiod portfolio policies, *Journal of Business*, 215-229.
- Poterba, J., and L.H. Summers, (1988), Mean reversion in stock returns: Evidence and implications, *Journal of Financial Economics*, 22, 27-60.
- Pratt, J., (1964), Risk aversion in the small and in the large, *Econometrica*, 32, 122-136.
- Rothschild, M. and J. Stiglitz, (1970), Increasing risk: I. A definition, *Journal of Economic Theory*, 2, 225-243.
- Treich, N., (1997), Risk tolerance and value of information in the standard portfolio model, *Economics Letters*, 55, 361-363.
- Xia, Y., (2000), Learning about predictability: The effects of parameter uncertainty on dynamic asset allocation, Discussion paper #11-00, The Anderson School at UCLA.