

Predictable Forward Performance Processes: The Binomial Case*

Bahman Angoshtari[†] Thaleia Zariphopoulou[‡] Xun Yu Zhou[§]

March 22, 2017

Abstract

We introduce a new class of forward performance processes that are endogenous and predictable with regards to an underlying market information set, and are updated at discrete times. Such performance criteria accommodate short-term predictability of asset returns, sequential learning and other dynamically unfolding factors affecting optimal portfolio choice. We analyze in detail a binomial model whose parameters are random and updated dynamically as the market evolves. We show that the key step in the construction of the associated predictable forward performance process is to solve a single-period *inverse* investment problem, namely, to determine, period-by-period and conditionally on the current market information, the end-time utility function from a given initial-time value function. We reduce this inverse problem to solving a single variable functional equation, and establish conditions for the existence and uniqueness of its solutions in the class of inverse marginal functions.

Keywords: Portfolio selection, forward performance processes, binomial model, inverse investment problem, functional equation, predictability.

1 Introduction

The classical portfolio selection paradigm is based on three fundamental ingredients: a given investment horizon, $[0, T]$, a performance function (such as a utility or a risk-return trade-off), $U_T(\cdot)$, applied at the *end* of the horizon, and a market model which yields the random investment opportunities available over $[0, T]$. This triplet is exogenously and entirely specified at initial time, $t = 0$.

*This work was presented at the SIAM conferences in Financial Mathematics in Chicago and Austin, and seminars at Oxford University and University of Michigan, Ann Arbor. The authors would like to thank the participants for fruitful comments and suggestions, as well as the Oxford-Man Institute of Quantitative Finance for its support and hospitality. Zhou gratefully acknowledges financial support through a start-up grant at Columbia University and through Oxford–Nie Financial Big Data Lab.

[†]Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA. Email: bango@umich.edu.

[‡]Department of Mathematics and IROM, The University of Texas at Austin, Austin and the Oxford-Man Institute, University of Oxford, Oxford, UK. Email: zariphop@math.utexas.edu.

[§]Department of IEOR, Columbia University, New York, NY 10027. Email: xz2574@columbia.edu.

Once these ingredients are chosen, one then solves for the optimal strategy $\pi^*(\cdot)$, and derives the value function $U_0(\cdot)$ at $t = 0$ as the expectation of the terminal utility of optimal wealth. The value function thus stipulates the best possible performance value achievable from each and every amount of initial wealth, and hence it can be in turn considered as a performance criterion at $t = 0$. Here, the terminal performance function $U_T(\cdot)$ is exogenous, and the optimal strategy $\pi^*(\cdot)$ and the initial performance function $U_0(\cdot)$ are endogenous. The model therefore entails a *backward* approach in time, from $U_T(\cdot)$ to $U_0(\cdot)$. This is also in accordance with the celebrated Dynamic Programming Principle (DPP), or otherwise known as Bellman’s principle of optimality.¹

Despite its classical mathematical foundations and theoretical appeal, this approach nonetheless has several shortcomings, and hardly reflects what happens in investment practice.

Firstly, it relies heavily on the model selection for the entire investment horizon. As a result, once a model is chosen for $[0, T]$, no revisions are allowed, for it would lead to time-inconsistent decisions. Therefore, any additional information coming from realized returns, and other sources of learning, as the market evolves, cannot be incorporated in the investment decisions.

This issue has been partially addressed by allowing for optimization over a family of models, with a wealth of results on the so-called model robustness/ambiguity problem. But even with this extended modeling approach, one also has to pre-commit to a family of possible models, which itself may change as time evolves.² Furthermore, uncertainty might be, more generally, generated by both changes in asset prices and exogenous factors, and it is quite difficult to model them accurately, especially for a long time ahead.³

The second difficulty is the pre-commitment at the initial time to a terminal utility. Indeed, it might be difficult to assess and specify the performance function, especially if the investment horizon is sufficiently long. It is more plausible to know the utility or the resulting preferred allocations for now or the immediate future, and then to preserve it under certain optimality criteria. This was firstly pointed out by Fisher Black in 1968 (see Black (1988)) where he argued that, practically, investors choose their initial/current desired allocation and proposed a way to update this allocation in the future under time consistent expected utility criteria. He did so through an equation that the allocation must satisfy through time. This equation was much later extensively analyzed by Musiela and Zariphopoulou (2010b) and more recently by Geng and Zariphopoulou (2017) where its long-term (turnpike) behavior is studied.

Thirdly, it is very seldom the case that an optimal investment problem “terminates” at a single horizon T . In practice, investment decisions are made for a series of relatively small time horizons, say $[0, T_1]$, $[T_1, T_2]$, \dots , $[T_n, T_{n+1}]$, \dots , *forward in time*. Thus, a framework of “pasting” these individual investment problems in a time-consistent manner as the market evolves is needed. Establishing such a framework is by no means trivial as we will explain in the sequel, but it captures reality, for it offers substantial flexibility in terms of learning and model revision, dynamic risk preferences criteria and rolling horizons.

The above considerations have led to the development of the so-called *forward perfor-*

¹See Dreyfus (2002) for a historical account on dynamic programming.

²See Kallblad et al. (2017) for a critique on model ambiguity.

³ Nadochiy and Tehranchi (2015) argue that stochastic factors must be incorporated “forward in time” and propose a finite-dimensional model for it.

mance measurement, initially proposed by Musiela and Zariphopoulou (2006) and later extended by the same authors in a series of papers (see Musiela and Zariphopoulou (2009, 2011)) and by other authors (see El Karoui and Mrad (2013), and Nadtochiy and Tehranchi (2015)) in continuous-time market settings. The main idea of the forward approach is that instead of fixing, as in the classical setting, an investment horizon, a market model and a terminal utility, one starts with an initial performance measurement and updates it *forward in time* as the market and other underlying stochastic factors evolve. The evolution of the forward process is dictated by a forward-in-time version of the DPP and, thus, it ensures time-consistency across different times.

With the exception of a special case studied in Musiela and Zariphopoulou (2003) and Musiela et al. (2016), in the context of indifference prices under exponential forward criteria, the existing results have so far focused exclusively on continuous-time, Itô-diffusion settings, in which both trading and performance valuation are carried out continuously in time. It was shown in Musiela and Zariphopoulou (2010a) that the forward performance process is associated with an ill-posed infinite-dimensional stochastic partial differential equation (SPDE), the same way that the classical value function satisfies the finite-dimensional Hamilton-Jacobi-Bellman equation (HJB). This performance SPDE has been subsequently studied in El Karoui and Mrad (2013), Nadtochiy and Zariphopoulou (2014), Nadtochiy and Tehranchi (2015) and, more recently, in Shkolnikov et al. (2016) for asset price factors evolving at different time scales. Despite the technical challenges that this forward SPDE presents (ill-posedness, high or infinite dimensionality, degeneracies, and volatility specification), the continuous-time cases are tractable because stochastic calculus can be employed and infinitesimal arguments can be, in turn, developed.

However, the continuous-time setting has a major drawback in that it is hard to see how exactly the performance criterion evolves from one instant to another. This evolution is lost at the infinitesimal level and hidden behind the (generally intractable) stochastic PDE.

The aim of this paper is to initiate the development and a systematic study of forward investment performance processes that are *discrete* in time, while trading can be either discrete or continuous in time. We will introduce an iterative mechanism in which an investor updates/predicts her performance criterion at the next investment period, based on both her current performance and her assessment of the upcoming market dynamics in the next period. This predictability will be present in an explicit and transparent manner.

In addition to the conceptual motivation described above, there are also practical considerations in studying the discrete-time predictable forward performance. Indeed, in investment practice, trading occurs at discrete times and not continuously. More importantly, performance criteria are directly or indirectly determined by individuals, such as higher-level managers or clients, and not by the portfolio manager. These “performance evaluators” use information sets that are different, both in terms of contents and updating frequency, from the ones used by the portfolio manager. For example, a portfolio manager may have access to various data sets, proprietary forecasting models, and sophisticated trading strategies which are out of reach of (or simply deemed as “too detailed” to be considered by) the performance evaluator. Similarly, even if trading can happen at extremely higher frequencies (hence almost close to continuous trading), performance assessment/update takes place at a much slower pace, e.g., a senior manager will not keep track of the performance of a portfolio or update the performance criterion as frequently as the subordinate portfolio manager in charge of that portfolio.

In this paper, we will consider an (indefinite) series of time points, $0 = t_0, t_1, \dots, t_n$,

..., at which the performance measurement is evaluated and updated. The (short) period between any given two neighboring points is called an evaluation period. We then introduce forward performance processes that are *predictable* with regards to the information at the most recent evaluation time.

We are motivated to introduce this class of criteria for two reasons. Firstly, it is natural to infer at the beginning of the evaluation period the criterion we use for the end of it. Conceptually different from the expected utility framework, here the criterion to be used at the end of the period will be *endogenous*, not exogenous.

Secondly, it is more feasible to estimate the market parameters for just one evaluation period ahead, than for longer periods. For example, the volatility can be reliably forecasted for a short time ahead using the so-called realized volatility introduced by the seminal work of Barndorff-Nielsen and Shephard (2002). On the other hand, short-term predictability of equity risk premium has been subject to extended studies in the last three decades (see, among others, Fama and French (1989)).

We stress that by looking at short-term predictability, we do not discard any long-term forecasting modeling input. On the contrary, the approach we propose accommodates both short- and long-term predictability modeling, by tracking the market and incorporating the unfolding information as time moves along. This cannot be done in the classical setting in which the market model, which may contain both short- and long-term factors, is pre-chosen at the initial time.

To highlight the key ideas of predictable forward performance processes, we start our analysis with a simple yet still rich enough setting. The market consists of two securities, a riskless asset and a stock whose price evolves according to a binomial model at times $0 = t_0, t_1, \dots, t_n, \dots$, at which the forward performance evaluation also occurs. The market model is more general than the standard binomial tree, in that the asset returns and their probabilities can be estimated/determined only one period ahead. Such a setting allows for genuine dynamic updating of the underlying parameters, as the market evolves from one period to the next.

In generating a predictable forward performance process, the investor starts at $t = 0$ and chooses (i.e. estimates) the market parameters for the upcoming trading/evaluation period $[0, t_1]$. She also chooses her initial performance function $U_0(\cdot)$, and derives a utility (performance) function $U_1(\cdot)$ that is deterministic, such that the pair (U_0, U_1) is consistent with the investment problem in $[0, t_1]$, with U_0, U_1 being respectively the value function and terminal utility function. Then, at t_1 , the agent repeats exactly the same procedure for the next period. Proceeding iteratively *forward in time*, and conditionally on the current information, a predictable performance process is constructed together with the optimal allocations and their wealth processes.⁴

Therefore, technically, we are left to solve a single-period investment problem where the value function is given and the terminal utility function is to be found. We term this as a *single-period inverse investment* problem, which needs to be solved sequentially “period-by-period,” conditionally on the information at the beginning of this trading period. It turns out that the key to solving this problem is a linear functional equation, which relates the

⁴Here we assume that the updating and trading take place at the same time. As discussed above, this does not have to be the case. However, we choose to study this parsimonious model in order to highlight the significance of updating the performance measurement in discrete times without getting into much technicality.

inverse marginal processes, at the beginning and the end of each evaluation period, with coefficients depending on the random market inputs. We analyze this equation in detail, and establish conditions for existence and uniqueness of the solutions in the class of inverse marginal functions.

The paper is structured as follows. In Section 2, we introduce the notion of predictable forward performance processes in a general market setting. We then formulate a binomial model with random, dynamically updated parameters in Section 3. In Section 4, we apply the definition of predictable forward performance processes to the binomial model, and show that their construction reduces to solving an inverse investment problem. In Section 5, this inverse problem is shown to be equivalent to a functional equation. We derive sufficient existence and uniqueness conditions as well as the explicit solution to the functional equation in Section 6. Finally, we present the general construction algorithm in Section 7, and conclude in Section 8. Proofs of the main results are relegated to an Appendix.

2 Predictable forward performance processes: A general definition

In this section, we define discrete-time predictable forward performance processes in a general market model. Starting from the next section, we will restrict the market setting to a binomial model with random parameters, and provide a detailed discussion on the existence and construction of such performance processes.

The investment paradigm is cast in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ augmented with a filtration (\mathcal{F}_t) , $t \geq 0$. We denote by $\mathcal{X}(t, x)$ the set of all the admissible wealth processes X_s , $s \geq t$, starting with $X_t = x$ and such that X_s is \mathcal{F}_s -measurable. The term “admissible” is for now generic and will be specified once a specific market model is introduced in the sequel.

We call a function $U : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ a *utility (or performance) function* if $U \in C^2(\mathbb{R}^+)$, $U' > 0$, $U'' < 0$, and satisfies the Inada conditions: $\lim_{x \rightarrow 0^+} U'(x) = \infty$ and $\lim_{x \rightarrow \infty} U'(x) = 0$.

For any σ -algebra $\mathcal{G} \subseteq \mathcal{F}$, the set of \mathcal{G} -measurable utility (or performance) functions is defined as

$$\mathcal{U}(\mathcal{G}) = \{U : \mathbb{R}^+ \times \Omega \rightarrow \mathbb{R} \mid U(x, \cdot) \text{ is } \mathcal{G} \text{-measurable for each } x \in \mathbb{R}^+, \text{ and } U(\cdot, \omega) \text{ is a utility function a.s.}\}.$$

In other words, the elements of $\mathcal{U}(\mathcal{G})$ are entirely known (predicted) based on \mathcal{G} , as they are predictable by the information contained in \mathcal{G} . Alternatively, we may think of $U \in \mathcal{U}(\mathcal{G})$ as a deterministic utility function, given the information in \mathcal{G} .

Next, we introduce the predictable forward performance processes. To ease the notation, we skip the ω -argument throughout.

Definition 1. *Let discrete time points $0 = t_0 < t_1 < \dots < t_n < \dots$ be given. A family of random functions $\{U_0, U_1, U_2, \dots\}$ is a predictable forward performance process with respect to (\mathcal{F}_t) if, for $X_n = X_{t_n}$ and $\mathcal{F}_n = \mathcal{F}_{t_n}$, $n = 1, 2, \dots$, the following conditions hold:*

- (i) U_0 is a deterministic utility function and $U_n \in \mathcal{U}(\mathcal{F}_{n-1})$.

(ii) For any initial wealth $x > 0$ and any admissible wealth process $X \in \mathcal{X}(0, x)$,

$$U_{n-1}(X_{n-1}) \geq E_{\mathbb{P}} [U_n(X_n) | \mathcal{F}_{n-1}].$$

(iii) For any initial wealth $x > 0$, there exists an admissible wealth process $X^* \in \mathcal{X}(0, x)$ such that

$$U_{n-1}(X_{n-1}^*) = E_{\mathbb{P}} [U_n(X_n^*) | \mathcal{F}_{n-1}].$$

We stress that there are no specific assumptions on the market model and how often trading occurs. The asset price processes can be discrete or continuous in time and, for the latter case, trading can be discrete or continuous. Furthermore, if trading takes place at discrete times, the rebalancing periods do not need to be aligned with the performance assessment times.

In practice, as mentioned in Introduction, it is typically the case that trading occurs more frequently than the performance evaluation, but the above definition accommodates cases when there is perfect alignment - as in the binomial model we will study herein - and the less realistic case when trading occurs less frequently than the performance measurement.

Compared with the continuous-time counterpart initially proposed by Musiela and Zariwopoulou (2009), the fundamentally distinctive element of Definition 1 is condition (i), which explicitly requires that the performance function at the *next* upcoming assessment time be *entirely determined* from the information up to the *present* time.

On the other hand, as in the continuous-time case, properties (ii)-(iii) draw from Bellman's principle of optimality, which stipulates that the processes $U_n(X_n)$ and $U_n(X_n^*)$, $n = 0, 1, \dots$, are, respectively, a supermartingale and a martingale with respect to the filtration (\mathcal{F}_n) . Since the Bellman principle underlines time-consistency, properties (ii)-(iii) directly ensure that the investment problem is time-consistent under the predictable forward performance criterion.

Hence, the above performance measurement is essentially *endogenized* by the time-consistency requirements (ii)-(iii).

We also note that the predictability of risk preferences is implicitly present in the classical expected utility in finite horizon settings, say $[0, T]$, in which a deterministic utility for T is pre-chosen at initial time $t_0 = 0$, and it is thus \mathcal{F}_0 -measurable. The fundamental difference, however, is that the terminal utility function in the classical theory is exogenous, instead of endogenous.

Definition 1 suggests a general scheme for constructing predictable forward performance functions in discrete times. Indeed, starting from an initial datum U_0 , given at time $t_0 = 0$, the entire family U_1, \dots, U_n, \dots , can be obtained by determining U_n from U_{n-1} iteratively, $n = 1, 2, \dots$, in the way described below.

Properties (ii)-(iii) dictate that, for each trading period $[t_{n-1}, t_n)$, we have

$$U_{n-1}(X_{n-1}^*) = \sup_{X_n \in \mathcal{X}(t_{n-1}, X_{n-1}^*)} E_{\mathbb{P}} [U_n(X_n) | \mathcal{F}_{n-1}]. \quad (1)$$

At instant t_{n-1} , since \mathcal{F}_{n-1} is realized, the random functions U_{n-1} and U_n are both deterministic and so is X_{n-1}^* . This, in turn, suggests that we should consider the following "*single-period*" investment problem

$$U_{n-1}(x) = \sup_{X_n \in \mathcal{X}_{n-1, n}(x)} E_{\mathbb{P}} [U_n(X_n) | \mathcal{F}_{n-1}], \quad (2)$$

for $x > 0$, where, with a slight abuse of notation, we use $\mathcal{X}_{n-1,n}(x)$ to denote the set of admissible wealths at t_n starting at t_{n-1} with wealth x .

Therefore, if we are able to determine, for each $n = 1, 2, \dots$, a performance function $U_n \in \mathcal{U}(\mathcal{F}_{n-1})$, such that the pair (U_n, U_{n+1}) satisfies (2), then we will have an iterative scheme to construct the *entire* predictable forward performance process, starting from U_0 .

One readily recognizes that (2) would be the classical expected utility problem if the objective were to derive U_{n-1} from U_n , with U_n being a deterministic utility function. Therefore, what we consider now is an *inverse* investment problem in that we are given its initial value function and we seek a terminal utility that is consistent with it, with both these functions being deterministic (conditionally on \mathcal{F}_{n-1}).

To our best knowledge, such inverse discrete-time problems have not been considered in the literature. The aim herein is to initiate a concise study of such performance criteria for general market settings. We start with the binomial case in which, however, the parameters - including the transition probabilities and price levels - are not known a priori but are updated as the market moves. Recall that, while this is very much in accordance with real investment practice, such a model is not implementable in the classical expected utility settings because model commitment occurs once, at the initial time. As we will see, while the binomial case is one of the simplest discrete-time market models, its analysis is sufficiently rich and its results reveal the key economic insights regarding the predictable performance criteria.

3 A binomial market model with random, dynamically updated parameters

We consider a market with two traded assets, a riskless bond and a stock. The bond is taken to be the numeraire and assumed to offer, without loss of generality, zero interest rate. The stock price at times t_0, t_1, \dots , evolves according to a binomial model that we now specify.

Let R_n be the total return of the stock over period $[t_{n-1}, t_n)$. Here, R_n is a random variable with two values $R_n^u > R_n^d$. We assume that R_n, R_n^u , and R_n^d , $n = 1, 2, \dots$, are all random variables in a measurable space (Ω, \mathcal{F}) augmented with a filtration (\mathcal{F}_n) , $n = 1, 2, \dots$, with \mathcal{F}_n representing the information available at t_n . Moreover, we assume that R_n is \mathcal{F}_n -measurable, and that its values, R_n^u and R_n^d , are taken to be \mathcal{F}_{n-1} -measurable. In other words, the high and low return levels for each investment period are known at the beginning of this period, while the realized return is known at its end.

Finally, the historical measure \mathbb{P} is a probability measure on (Ω, \mathcal{F}) and the following standard no-arbitrage assumption is satisfied.

Assumption 2. For all $n = 1, 2, \dots$:

- (i) $0 < R_n^d < 1 < R_n^u$, \mathbb{P} -almost surely; and,
- (ii) $0 < E_{\mathbb{P}}(\mathbf{1}_{\{R_n=R_n^u\}}|\mathcal{F}_{n-1}) = 1 - E_{\mathbb{P}}(\mathbf{1}_{\{R_n=R_n^d\}}|\mathcal{F}_{n-1}) < 1$, \mathbb{P} -almost surely.

For $n = 1, 2, \dots$, the \mathcal{F}_{n-1} -measurable random variable

$$p_n := E_{\mathbb{P}}[\mathbf{1}_{\{R_n=R_n^u\}}|\mathcal{F}_{n-1}] = 1 - E_{\mathbb{P}}[\mathbf{1}_{\{R_n=R_n^d\}}|\mathcal{F}_{n-1}],$$

represents the *best estimate* of the probability of an upward jump over $[t_{n-1}, t_n]$, *given the information* available at t_{n-1} . In practice, p_n corresponds to the outcome of a sequential learning procedure that is conducted at t_{n-1} .

We assume no further information about the physical measure \mathbb{P} . In particular, we do not assume that \mathbb{P} is known, other than it satisfies Assumption 2.

The investor trades between the stock and the bond using self-financing strategies. She starts at $t_0 = 0$ with total wealth $x > 0$ and rebalances her portfolio at times t_n , $n = 1, 2, \dots$. At the beginning of each period, say $[t_n, t_{n+1})$, she chooses the amount π_{n+1} to be invested in the stock (and the rest in the bond) for this period. In turn, her wealth process, denoted by X_n^π , $n = 1, 2, \dots$, evolves according to the wealth equation

$$X_{n+1}^\pi = X_n^\pi + \pi_{n+1}(R_{n+1} - 1),$$

with $X_0 = x$.

The investor is allowed to short the stock but her wealth can never become negative; thus π_{n+1} must satisfy

$$-\frac{X_n^\pi}{R_{n+1}^u - 1} \leq \pi_{n+1} \leq \frac{X_n^\pi}{1 - R_{n+1}^d}; \quad n = 1, 2, \dots \quad (3)$$

We call an investment strategy $\pi = \{\pi_n\}_{n=1}^\infty$ *admissible* if it is self-financing, π_n is \mathcal{F}_{n-1} -measurable, and (3) is satisfied \mathbb{P} -almost surely. A wealth process $X = \{X_n^\pi\}_{n=0}^\infty$ is then admissible if the strategy π that generates it is admissible.

We recall that $\mathcal{X}(n, x)$ is the set of admissible wealth processes $\{X_m\}_{m=n}^\infty$, starting with $X_n = x$.

We also introduce the auxiliary "single-step" set of admissible portfolios π_{n+1} , chosen at t_n for the trading period $[t_n, t_{n+1})$ assuming wealth x at t_n , by

$$\mathcal{A}_{n,n+1}(x) = \left\{ \pi_{n+1} : \pi_{n+1} \text{ is } \mathcal{F}_n\text{-measurable, } -\frac{x}{R_{n+1}^u - 1} \leq \pi_{n+1} \leq \frac{x}{1 - R_{n+1}^d}, x > 0 \right\},$$

as well as the corresponding set of admissible wealth processes

$$\mathcal{X}_{n,n+1}(x) = \{x + \pi_{n+1}R_{n+1} : \pi_{n+1} \in \mathcal{A}_{n,n+1}(x), x > 0\}.$$

4 Problem statement and reduction to the single-period inverse investment problem

In this section, we consider predictable forward performance processes in the binomial model, and show that their construction reduces to solving a series of single-period inverse investment problems.

The investor starts with an initial utility U_0 and updates her performance criteria at times t_1, t_2, \dots , with the associated performance functions U_1, U_2, \dots satisfying Definition 1.

We now present the procedure that yields the construction of a predictable forward performance process starting from U_0 , and determining U_n from U_{n-1} , iteratively for $n = 1, 2, \dots$.

At $t_0 = 0$, equation (1) becomes

$$U_0(x) = \sup_{X_1 \in \mathcal{X}(0,x)} E_{\mathbb{P}} \left[U_1(X_1) \middle| \mathcal{F}_0 \right] = \sup_{\pi_1 \in \mathcal{A}_{0,1}(x)} E_{\mathbb{P}} \left[U_1 \left(x + \pi_1(R_1 - 1) \right) \right]; \quad x > 0. \quad (4)$$

Since the market parameters (R_1^u, R_1^d, p_1) and the initial datum U_0 are known at t_0 , finding a deterministic (\mathcal{F}_0 -measurable) U_1 reduces to the single-period inverse investment problem discussed in Section 2. Let us for the moment assume that we are able to solve this inverse problem to obtain U_1 .

At $t = t_1$, the investor observes the realization of the stock return R_1 and estimates the parameters (R_2^u, R_2^d, p_2) for the second trading period $[t_1, t_2]$. Setting $n = 2$ in (1) then yields

$$U_1(X_1^*(x)) = \sup_{X_2 \in \mathcal{X}(1, X_1^*(x))} E_{\mathbb{P}} [U_2(X_2) | \mathcal{F}_1], \quad (5)$$

where $X_1^*(x)$ is the optimal wealth generated at t_1 , starting at x at $t_0 = 0$, from the previous period.

It follows from the classical expected utility theory (see also Theorem 4 below) that $X_1^*(x) = I_1(\rho_1 U_0'(x))$, $x > 0$, where $I_1 = (U_1')^{-1}$ and ρ_1 is the pricing kernel over the period $[0, t_1]$, given by

$$\rho_1 = \frac{1 - R_1^d}{p_1(R_1^u - R_1^d)} \mathbf{1}_{\{R_1=R_1^u\}} + \frac{R_1^u - 1}{(1 - p_1)(R_1^u - R_1^d)} \mathbf{1}_{\{R_1=R_1^d\}}.$$

The mapping $x \rightarrow X_1^*(x)$ is strictly increasing for each $x > 0$ and of full range, since I_1 and U_0' are both strictly decreasing functions, $\rho_1 > 0$, and the Inada conditions yield $X_1^*(0) = 0$ and $X_1^*(\infty) = \infty$.

Since $X_1^*(x)$ is \mathcal{F}_1 -measurable and the parameters (R_2^u, R_2^d, p_2) together with U_1 are all known at $t = t_1$, we deduce that (5) reduces, with a slight abuse of notation, to finding $U_2(\cdot) \in \mathcal{U}(\mathcal{F}_1)$ such that

$$U_1(x) = \sup_{\pi_2 \in \mathcal{A}_{1,2}(x)} E_{\mathbb{P}} [U_2(x + \pi_2(R_2 - 1)) | \mathcal{F}_1]; \quad x > 0,$$

with U_1 given. In other words, one needs to solve yet another inverse investment problem that is mathematically identical to (4).

At $t = t_n$, in exactly the same manner as above, we have to solve

$$U_n(x) = \sup_{\pi_{n+1} \in \mathcal{A}_{n,n+1}(x)} E_{\mathbb{P}} [U_{n+1}(x + \pi_{n+1}(R_{n+1} - 1)) | \mathcal{F}_n]; \quad x > 0,$$

thereby deriving U_{n+1} from U_n , with $U_{n+1} \in \mathcal{U}(\mathcal{F}_{n+1})$ and with the parameters (R_n^u, R_n^d, p_n) known.

Thus, all the terms of a predictable forward performance process can be obtained, starting from any arbitrary initial wealth $x > 0$ and proceeding iteratively solving a “period-by-period” inverse optimization problem. Moreover, as we will show in the next section, we will also derive the optimal portfolio and wealth processes at the same time.

To summarize, the crucial step in the entire predictable forward construction is to solve this *single-period inverse investment problem*. We do this in the next section.

5 The single-period inverse investment problem

We focus on the analysis of the inverse investment problem (4). To ease the presentation, we introduce a simplified notation. We set $t_0 = 0, t_1 = 1$ and $R_1 = R$ taking values u and d , $u > 1$ and $0 < d < 1$, with probability $0 < p < 1$ and $1 - p$, respectively. We recall the risk neutral probabilities

$$q = \frac{1-d}{u-d} \quad \text{and} \quad 1-q = \frac{u-1}{u-d},$$

and the pricing kernel

$$\rho_1 = \rho^u \mathbf{1}_{\{R=u\}} + \rho^d \mathbf{1}_{\{R=d\}} := \frac{q}{p} \mathbf{1}_{\{R=u\}} + \frac{1-q}{1-p} \mathbf{1}_{\{R=d\}}. \quad (6)$$

The investor starts with wealth $X_0 = x > 0$, and invests the amount π in the stock. Her wealth at $t = 1$ is then given by the random variable $X = x + \pi(R - 1)$. The no-bankruptcy constraint (3) becomes $\underline{\pi}(x) \leq \pi \leq \bar{\pi}(x)$, with

$$\underline{\pi}(x) = -\frac{x}{u-1} < 0 \quad \text{and} \quad \bar{\pi}(x) = \frac{x}{1-d} > 0.$$

We denote the set of admissible portfolios as

$$\mathcal{A}(x) = \{\pi \in \mathbb{R}, \text{ and } \underline{\pi}(x) \leq \pi \leq \bar{\pi}(x), x > 0\}.$$

Given an initial utility function U_0 , we then seek another performance function U_1 , such that

$$U_0(x) = \sup_{\pi \in \mathcal{A}(x)} E_{\mathbb{P}} [U_1(x + \pi(R - 1))]; \quad x > 0. \quad (7)$$

Let \mathcal{U} be the set of deterministic utility functions. We introduce the set of *inverse marginal functions* \mathcal{I} ,

$$\mathcal{I} := \left\{ I \in C^1(\mathbb{R}^+) : I' < 0, \lim_{y \rightarrow \infty} I(y) = 0, \lim_{y \rightarrow 0^+} I(y) = \infty \right\}. \quad (8)$$

Note that if functions U and I satisfy $I = (U')^{-1}$, then U is a utility function if and only if I is an inverse marginal function.

Assuming for now that a utility function U_1 satisfying (7) exists, we consider the inverse marginal functions

$$I_0 = (U_0')^{-1} \quad \text{and} \quad I_1 = (U_1')^{-1}.$$

Our main goal in this section is to show that the inverse investment problem (7) reduces to a functional equation in terms of I_0 and I_1 ; see (9) below.

The following theorem is one of the main results herein, establishing a direct relationship between the inverse marginals at the beginning and at the end of the trading period $[0, 1]$, when the corresponding utilities are related by (7).

Theorem 3. *Let $U_0, U_1 \in \mathcal{U}$ satisfy the optimization problem (7). Then, their inverse marginals I_0 and I_1 must satisfy the linear functional equation*

$$I_1(ay) + bI_1(y) = (1+b)I_0(cy); \quad y > 0, \quad (9)$$

where

$$a = \frac{1-p}{p} \frac{q}{1-q}, \quad b = \frac{1-q}{q} \quad \text{and} \quad c = \frac{1-p}{1-q}. \quad (10)$$

Proof. From standard arguments, for all $x > 0$, there exists an optimizer $\pi^*(x)$ for (7) satisfying the first-order condition

$$p(u-1)U_1'(x + \pi^*(x)(u-1)) + (1-p)U_1'(x + \pi^*(x)(d-1)) = 0. \quad (11)$$

On the other hand, it follows from (7) that

$$pU_1(x + \pi^*(x)(u-1)) + (1-p)U_1(x + \pi^*(x)(d-1)) = U_0(x).$$

Differentiating the above equation and using (11) yield

$$pU_1'(x + \pi^*(x)(u-1)) + (1-p)U_1'(x + \pi^*(x)(d-1)) = U_0'(x). \quad (12)$$

Solving the linear system (11)-(12) gives

$$U_1'(x + \pi^*(x)(u-1)) = \frac{(1-d)}{p(u-d)}U_0'(x)$$

and

$$U_1'(x + \pi^*(x)(d-1)) = \frac{(u-1)}{(1-p)(u-d)}U_0'(x).$$

Therefore, the optimal allocation function $\pi^*(x)$ satisfies

$$\begin{cases} x + \pi^*(x)(u-1) &= I_1 \left(\frac{1-d}{p(u-d)}U_0'(x) \right), \\ x + \pi^*(x)(d-1) &= I_1 \left(\frac{u-1}{(1-p)(u-d)}U_0'(x) \right), \end{cases} \quad (13)$$

from which we obtain the solution

$$\pi^*(x) = \frac{1}{u-d} \left(I_1 \left(\frac{1-d}{p(u-d)}U_0'(x) \right) - I_1 \left(\frac{u-1}{(1-p)(u-d)}U_0'(x) \right) \right); \quad x > 0.$$

Substituting the above in either of the equations in (13) yields

$$\frac{1-d}{u-d} I_1 \left(\frac{1-d}{p(u-d)}U_0'(x) \right) + \frac{u-1}{u-d} I_1 \left(\frac{u-1}{(1-p)(u-d)}U_0'(x) \right) = x.$$

Changing variables $x = I_0 \left(\frac{(1-p)(u-d)}{u-1}y \right)$, $y > 0$, the above becomes

$$I_1 \left(\frac{(1-p)(1-d)}{p(u-1)}y \right) + \frac{u-1}{1-d} I_1(y) = \frac{u-d}{1-d} I_0 \left(\frac{(1-p)(u-d)}{u-1}y \right); \quad y > 0.$$

Noting (10) we conclude. □

Next, we show by an explicit construction how to recover U_1 from I_1 . At the same time we derive the optimal portfolio $\pi^*(x)$ and its wealth $X^*(x)$.

Theorem 4. *Let U_0 be a utility function and I_0 be its inverse marginal, and I_1 be an inverse marginal solving the functional equation (9). Let also ρ_1 be the pricing kernel given by (6). Then, the following statements hold.*

(i) The function U_1 defined by

$$U_1(x) := U_0(1) + E_{\mathbb{P}} \left[\int_{I_1(\rho_1 U'_0(1))}^x I_1^{-1}(\xi) d\xi \right]; \quad x > 0 \quad (14)$$

is a well-defined utility function.

(ii) We have

$$U_0(x) = \sup_{\pi \in \mathcal{A}(x)} E_{\mathbb{P}} [U_1(x + \pi(R - 1))]; \quad x > 0.$$

(iii) The optimal wealth $X_1^*(x)$ and the associated optimal investment allocation $\pi^*(x)$ are given, respectively, by

$$X_1^*(x) = I_1(\rho_1 U'_0(x)) = X^{*,u}(x) \mathbf{1}_{\{R=u\}} + X^{*,d}(x) \mathbf{1}_{\{R=d\}}$$

and

$$\pi^*(x) = \frac{X^{*,u}(x) - X^{*,d}(x)}{u - d},$$

with

$$X^{*,u} = I_1 \left(\frac{q}{p} U'_0(x) \right) \quad \text{and} \quad X^{*,d} = I_1 \left(\frac{1-q}{1-p} U'_0(x) \right).$$

Proof. See Appendix A. □

Remark 5. As shown in the proof of Theorem 4, we can replace (14) with

$$U_1(x) := U_0(c) + E_{\mathbb{P}} \left[\int_{I_1(\rho_1 U'_0(c))}^x I_1^{-1}(\xi) d\xi \right]; \quad x > 0,$$

for any arbitrary constant $c > 0$. The choice of c does not change the value of $U_1(x)$, neither the optimal policies.

Theorem 4 reduces the inverse investment problem (7) to the functional equation (9). We study this functional equation in the next section.

6 A functional equation for inverse marginals

In this section, we analyze the linear functional equation (9), with I_0 given and I_1 to be found, for positive constants a, b, c , given by (10). We provide conditions for the existence and uniqueness of its solutions and, in particular, solutions in the class of inverse marginal functions.

When $a = 1$, the unique solution is trivially $I_1(y) = I_0(y)$. This is economically intuitive. If $p = q$, then essentially there is no risk premium to exploit. As a result, when $r = 0$ as assumed herein, the pricing kernel becomes a constant, $\rho = 1$, and the optimal wealth reduces to $X^*(x) = x$. In turn, the value function (at $t = 0$) coincides with the terminal utility. So the forward performance remains constant, $U_0(x) = U_1(x)$, or their inverse marginals I_0

and I_1 coincide.⁵ Indeed, there is no reason to modify the performance function in a market with no investment opportunities.

Henceforth we assume that $a \neq 1$. We start with an example showing that a general solution of (9) may not be unique, even if we restrict the solutions to inverse marginals.

Example 6. Let $\log_a b < 0$ and $I_0(y) = y^{\log_a b}$, $y > 0$. It is easy to check that the function $I_1(y) = \delta y^{\log_a b}$, $y > 0$, with $\delta = \frac{(1+b)}{2b c^{-\log_a b}} > 0$, is a solution to (9).

However, this particular solution is not the only solution. Indeed, consider any differentiable anti-periodic function, say $\Theta(z) = -\Theta(z + \ln a)$, for which there exists a constant $M > 0$ such that

$$\sup_{z \in \mathbb{R}} (|\Theta(z)|, |\Theta'(z)|) < M < \delta \frac{-\log_a b}{1 - \log_a b} = -\frac{(1+b) \log_a b}{2b c^{-\log_a b} (1 - \log_a b)}.$$

For instance, $\Theta(x) = M \sin(\frac{x}{\ln a} \pi)$ is such a function. One can then directly check that the function

$$\tilde{I}_1(y) = y^{\log_a b} (\delta + \Theta(\ln y)); \quad y > 0$$

is a solution.

As a matter of fact, both solutions I_1 and \tilde{I}_1 are inverse marginals. This is obvious for I_1 . As for \tilde{I}_1 , we have $\lim_{y \rightarrow \infty} \tilde{I}_1(y) = 0$ since $\log_a b < 0$. Moreover, it follows from the inequality $\tilde{I}_1(y) \geq y^{\log_a b} (\delta - M)$, $y > 0$, that $\lim_{y \rightarrow 0^+} \tilde{I}_1(y) = \infty$. Furthermore,

$$\begin{aligned} \tilde{I}'_1(y) &= y^{\log_a b - 1} \log_a b \left(\delta + \Theta(\ln y) + \frac{\Theta'(\ln y)}{\log_a b} \right) \\ &\leq y^{\log_a b - 1} \log_a b \left(\delta - \frac{M \log_a b - M}{\log_a b} \right) < 0; \quad y > 0. \end{aligned}$$

Thus, in general, there is no uniqueness even among inverse marginals.

The above example suggests that we need additional conditions to ensure uniqueness. To identify these conditions, we first note that (9) is a functional equation of the more general form

$$F(f(y)) = g(y)F(y) + h(y), \tag{15}$$

with f , g , and h given functions, $y \in \mathcal{Y} \subseteq \mathbb{R}$ and F to be found. The equations of this type have been studied in the literature; see Kuczma et al. (1990) and the references therein for a general exposition.

In general, such equations have many solutions. A trivial example is $F(y+1) = F(y)$, $y \in \mathbb{R}$, for which any periodic function with period 1 is a solution. Such non-uniqueness often renders the underlying equation inapplicable for concrete problems, where a single well-defined solution is usually needed. For the general equation (15), conditions for the uniqueness of solutions usually limit the set of solutions by imposing additional assumption on $F(y_0)$, where y_0 is a fixed point for f : $f(y_0) = y_0$. In the example of the equation

⁵This is also in accordance with the so-called time-monotone forward processes in the continuous-time setting. For example, in Musiela and Zariphopoulou (2010b), it is shown that this forward performance is given by $U(x, t) = u\left(x, \int_0^t |\lambda_s|^2 ds\right)$, with $u(x, t)$ a deterministic function and the process λ being the market price of risk. If $\lambda \equiv 0$, then $U(x, t) = u(x, 0) = U(x, 0)$, for all $t > 0$.

$F(y+1) = F(y)$, $y \in \mathbb{R}$, if we require a solution to be such that $\lim_{y \rightarrow \infty} F(y)$ exists, then $F \equiv 0$ becomes the only possible solution. Note here that ∞ is a fixed point of the function $f(y) = y + 1$.

For equation (9), $f(y) = ay$, $g(y) = -b$ and $h(y) = (1+b)G(cy)$. Therefore, uniqueness conditions should impose additional assumptions on F at $y_1 = 0$ and $y_2 = \infty$, which are the fixed points of $f(y) = ay$.

We start with the following auxiliary result in which we provide *general* uniqueness conditions for equation (9). Afterwards, we will strengthen the results for the family of inverse marginals.

Lemma 7. *Let I_0 be given. Then there exists at most one solution to (9), say I , satisfying $\lim_{y \rightarrow 0^+} y^{-\log_a b} I(y) = 0$. Similarly, there exists at most one solution satisfying $\lim_{y \rightarrow \infty} y^{-\log_a b} I(y) = 0$.*

Proof. Let F_1 and F_2 be two solutions of (9) that both satisfy either conditions given in the lemma. We show that their difference $w := F_1 - F_2 \equiv 0$.

The function w satisfies the homogenous equation $w(ay) = -bw(y)$, $y > 0$. Therefore, for $k = 1, 2, \dots$,

$$w(y) = \frac{w(ay)}{-b} = \frac{w(a^2y)}{(-b)^2} = \dots = \frac{w(a^k y)}{(-b)^k},$$

and

$$w(y) = -bw\left(\frac{y}{a}\right) = (-b)^2 w\left(\frac{y}{a^2}\right) = \dots = (-b)^k w\left(\frac{y}{a^k}\right).$$

It then follows that for $k = \pm 1, \pm 2, \dots$ and $y > 0$,

$$\begin{aligned} |w(y)| &= b^k \left| w\left(\frac{y}{a^k}\right) \right| = y^{\log_a b} \left(\frac{y}{a^k}\right)^{-\log_a b} \left| w\left(\frac{y}{a^k}\right) \right| \\ &\leq y^{\log_a b} \left(\frac{y}{a^k}\right)^{-\log_a b} \left(\left| F_1\left(\frac{y}{a^k}\right) \right| + \left| F_2\left(\frac{y}{a^k}\right) \right| \right). \end{aligned}$$

The right side vanishes as either $k \rightarrow \infty$ or $k \rightarrow -\infty$, and we conclude. \square

We note that the function \tilde{I}_1 in Example 6 satisfies neither conditions in Lemma 7, and thus uniqueness fails.

Next, we state the main result for this section, which provides sufficient conditions for existence and uniqueness of solutions to (9) that are inverse marginal functions.

Theorem 8. *Let I_0 in (9) be an inverse marginal utility, i.e. $I_0 \in \mathcal{I}$ with \mathcal{I} defined in (8). Define the functions*

$$\Phi_0(y) = I_0(acy) - bI_0(cy) \quad \text{and} \quad \Psi_0(y) = y^{-\log_a b} I_0(cy); \quad y > 0. \quad (16)$$

The following assertions hold:

- (i) *If Φ_0 is strictly increasing and, either $a > 1$ and $\lim_{y \rightarrow \infty} \Psi_0(y) = 0$ or $a < 1$ and $\lim_{y \rightarrow 0^+} \Psi_0(y) = 0$, then a solution of (9) is given by*

$$I_1(y) = \frac{1+b}{b} \sum_{m=0}^{\infty} (-1)^m b^{-m} I_0(a^m cy); \quad y > 0. \quad (17)$$

(ii) If Φ_0 is strictly decreasing and, either $a > 1$ and $\lim_{y \rightarrow 0^+} \Psi_0(y) = 0$ or $a < 1$ and $\lim_{y \rightarrow \infty} \Psi_0(y) = 0$, then a solution of (9) is given by

$$I_1(y) = (1 + b) \sum_{m=0}^{\infty} (-1)^m b^m I_0(a^{-(m+1)} c y); \quad y > 0. \quad (18)$$

(iii) In parts (i) and (ii), the corresponding I_1 satisfies the uniqueness condition(s) of Lemma 7 and, moreover, $I_1 \in \mathcal{I}$, i.e., I_1 preserves the inverse marginal properties.

(iv) The function I_1 in parts (i) and (ii), respectively, is the only positive solution of (9). It is also the only inverse marginal that solves (9).

Proof. See Appendix B. □

Next, we apply the above result to the case when the initial utility is a power function. The following example provides results complementary to the ones in Example 6 where uniqueness lacks as the result of not satisfying the conditions of Lemma 7.

Corollary 9. Let $U_0(x) = \left(1 - \frac{1}{\theta}\right)^{-1} x^{1-\frac{1}{\theta}}$, $x > 0$, and assume that $1 \neq \theta > 0$, $\theta \neq -\log_a b$, with $a, b, c > 0$ given by (10). Then, the following assertions hold:

(i) The unique marginal utility function that satisfies the functional equation (9) with the initial $I_0(y) = y^{-\theta}$ is given by

$$I_1(y) = \delta y^{-\theta}; \quad y > 0, \quad (19)$$

where $\delta = \frac{1+b}{c^\theta(a^{-\theta}+b)}$.

(ii) The unique utility function U_1 that satisfies the inverse investment problem (7) is given by

$$U_1(x) = \delta^{\frac{1}{\theta}} \left(1 - \frac{1}{\theta}\right)^{-1} x^{1-\frac{1}{\theta}} = \delta^{\frac{1}{\theta}} U_0(x); \quad x > 0.$$

(iii) The corresponding optimal allocation is given by

$$\pi^*(x) = \frac{\delta(p/q)^\theta - 1}{u - 1} x; \quad x > 0.$$

So, if we start with an initial power utility U_0 , then the forward utility at $t = 1$ is a multiple of the initial datum, with the constant given by $\delta^{\frac{1}{\theta}}$. Proceeding iteratively, the utilities for all the future periods remain to be power functions. In other words, in the binomial setting, the (predictable) power utility preferences are preserved throughout.

7 Construction of the predictable forward performance process

We are now ready to present the general algorithm for the construction of forward performance processes as well as the associated optimal investment strategies and their wealth processes. We stress that one of the main strengths of our approach is that for every trading period, say $(t_n, t_{n+1}]$, we do not update the model parameters $(p_{n+1}, R_{n+1}^u, R_{n+1}^d)$ for this period until t_n . Thus we take full advantage of the incoming information up to time t_n .

The algorithm is based on repeatedly applying the following result on the single-period inverse investment problem (7).

Theorem 10. *For the inverse investment problem (7), assume that the initial inverse marginal $I_0 = (U'_0)^{-1}$ satisfies condition (i) (resp. condition (ii)) in Theorem 8, and define I_1 by (17) (resp. (18)). Then, the unique solution to (7) is given by*

$$U_1(x) = U_0(1) + E_{\mathbb{P}} \left[\int_{I_1(\rho_1 U'_0(1))}^x I_1^{-1}(\xi) d\xi \right]; \quad x > 0,$$

where ρ_1 as in (6). Moreover, the optimal wealth $X_1^*(x)$ and the associated optimal investment allocation $\pi^*(x)$ are given, respectively, by

$$X_1^*(x) = I_1(\rho_1 U'_0(x)) = X_1^{*,u}(x) \mathbf{1}_{\{R_1=u\}} + X_1^{*,d}(x) \mathbf{1}_{\{R_1=d\}},$$

$$\pi^*(x) = \frac{X_1^{*,u}(x) - X_1^{*,d}(x)}{u - d},$$

where

$$X_1^{*,u}(x) = I_1 \left(\frac{q}{p} U'_0(x) \right) \quad \text{and} \quad X_1^{*,d}(x) = I_1 \left(\frac{1-q}{1-p} U'_0(x) \right).$$

Proof. The results follow directly from Theorem 8 and Theorem 4. □

Given an initial performance function U_0 and initial wealth X_0 , the following is the algorithm for constructing the predictable forward performance process $\{U_1, U_2, \dots\}$ along with the associated optimal portfolio process $\{\pi_1^*, \pi_2^*, \dots\}$ and the wealth process $\{X_1^*, X_2^*, \dots\}$ in the binomial market model.

- At $t = 0$: Assess the market parameters (R_1^u, R_1^d, p_1) for the first investment period, $[0, t_1)$. Compute

$$q_1 = \frac{1 - R_1^d}{R_1^u - R_1^d}, \quad a_1 = \frac{q_1(1 - p_1)}{p_1(1 - q_1)}, \quad b_1 = \frac{1 - q_1}{q_1}, \quad \text{and} \quad c_1 = \frac{1 - p_1}{1 - q_1},$$

and

$$\rho_1^u = \frac{q_1}{p_1} \quad \text{and} \quad \rho_1^d = \frac{1 - q_1}{1 - p_1}.$$

Using (a_1, b_1, c_1) check the conditions in part (i) (resp. (ii)) of Theorem 8, and obtain I_1 from (17) (resp. (18)). Then, apply Theorem 10 to compute

$$\begin{aligned} U_1(x) &= U_0(1) + p_1 \int_{I_1(\rho_1^u U_0'(1))}^x I_1^{-1}(\xi) d\xi \\ &\quad + (1 - p_1) \int_{I_1(\rho_1^d U_0'(1))}^x I_1^{-1}(\xi) d\xi; \quad x > 0, \\ \pi_1^* &= \frac{X_1^{*,u}(X_0) - X_1^{*,d}(X_0)}{u - d}, \end{aligned}$$

and

$$X_1^* = X_0 + \pi_1^*(R_1 - 1),$$

where

$$X_1^{*,u}(x) = I_1 \left(\frac{q_1}{p_1} U_0'(x) \right) \quad \text{and} \quad X_1^{*,d}(x) = I_1 \left(\frac{1 - q_1}{1 - p_1} U_0'(x) \right); \quad x > 0.$$

- At $t = t_n$ ($n = 1, 2, \dots$): We have already obtained $\{U_1, \dots, U_n; I_1, \dots, I_n\}$, $\{\pi_1^*, \dots, \pi_n^*\}$ and $\{X_1^*, \dots, X_n^*\}$. Estimate the market parameters $(R_{n+1}^u, R_{n+1}^d, p_{n+1})$ for the upcoming investment period $[t_n, t_{n+1})$. Let

$$q_{n+1} = \frac{1 - R_{n+1}^d}{R_{n+1}^u - R_{n+1}^d}, \quad a_{n+1} = \frac{q_{n+1}(1 - p_{n+1})}{p_{n+1}(1 - q_{n+1})}, \quad b_{n+1} = \frac{1 - q_{n+1}}{q_{n+1}}, \quad \text{and} \quad c_{n+1} = \frac{1 - p_{n+1}}{1 - q_{n+1}},$$

and

$$\rho_{n+1}^u = \frac{q_{n+1}}{p_{n+1}} \quad \text{and} \quad \rho_{n+1}^d = \frac{1 - q_{n+1}}{1 - p_{n+1}}.$$

Check the conditions in part (i) (resp. (ii)) in Theorem 8, using $(a_{n+1}, b_{n+1}, c_{n+1})$ (instead of (a, b, c)) and I_n instead of I_0 , and obtain I_{n+1} from (17) (resp. (18)).⁶

Compute

$$\begin{aligned} U_{n+1}(x) &= U_n(1) + p_{n+1} \int_{I_{n+1}(\rho_{n+1}^u U_n'(1))}^x I_{n+1}^{-1}(\xi) d\xi \\ &\quad + (1 - p_{n+1}) \int_{I_{n+1}(\rho_{n+1}^d U_n'(1))}^x I_{n+1}^{-1}(\xi) d\xi; \quad x > 0, \\ \pi_{n+1}^* &= \frac{X_{n+1}^{*,u}(X_n^*) - X_{n+1}^{*,d}(X_n^*)}{R_{n+1}^u - R_{n+1}^d}, \end{aligned} \tag{20}$$

and

$$X_{n+1}^* = X_n^* + \pi_{n+1}^*(R_{n+1} - 1) = X_0 + \sum_{i=1}^{n+1} \pi_i^*(R_i - 1),$$

⁶If both conditions in part (i) and (ii) do not hold, then the functional equation (9) may not have a solution, or the solution may not be unique. For the case of initial power utility $U_0(x) = \frac{x^{1-1/\theta}}{1-1/\theta}$, $\theta > 0$, Example 6 and Corollary 9 show that both condition fail at t_n if and only if $\theta = -\log_a b > 0$, in which case the solution exists but is not unique. This case is pathological, but to solve it remains a technically interesting question.

where,

$$X_{n+1}^{*,u}(x) = I_{n+1} \left(\frac{q_{n+1}}{p_{n+1}} U'_n(x) \right) \quad \text{and} \quad X_{n+1}^{*,d}(x) = I_{n+1} \left(\frac{1 - q_{n+1}}{1 - p_{n+1}} U'_n(x) \right); \quad x > 0.$$

In summary, starting with an initial datum U_0 , we have constructed for (the end of) each trading period, say $(t_n, t_{n+1}]$, $n = 1, 2, \dots$, a performance criterion U_{n+1} at t_{n+1} that is indeed \mathcal{F}_n -measurable. This measurability is inherited by the same measurability of the inverse marginal I_{n+1} that enters in the lower part of the integration in (20). Moreover, as expected, the optimal wealth X_{n+1}^* is \mathcal{F}_{n+1} -measurable, given that the pricing kernel ρ_{n+1} is \mathcal{F}_{n+1} -measurable. The optimal portfolio π_{n+1}^* is \mathcal{F}_n -measurable, chosen at the beginning of the period $(t_n, t_{n+1}]$.

8 Conclusion

We have introduced a new approach to optimal portfolio management that allows for dynamic model specification and adaptation, flexible investment horizons, and stochastic risk preferences. These risk preferences are modeled as a discrete-time predictable process, which is a rather natural and intuitive property of performance measurement criteria in practical applications. The frequency of performance evaluation is allowed to be different from or the same as the one at which the portfolio is rebalanced.

Specifically, at the beginning of each evaluation period, the investor assesses the market parameters only for this period (during which trading may take place many times, in both discrete or continuous fashion). Then, she solves an inverse single-period investment model which yields the utility at the end of the period, given the one at the beginning. The martingality and supermartingality requirements of the forward performance process ensure that this construction, "period-by-period forward in time" and adapted to the new market information, yields time-consistent policies.

We have implemented this new approach in a binomial model with random parameters, including both the probabilities and the levels of the stock returns. Such a setting is considerably flexible, as it accommodates short-term predictability of the asset returns, sequential learning and other dynamically evolving factors affecting optimal investments. We have then discussed in detail how the construction of predictable forward performance processes essentially reduces to a single-period inverse investment problem. We have, in turn, shown that the latter is equivalent to solving a functional equation involving the inverse marginal functions at the beginning and the end of trading period, and have established conditions for the existence and uniqueness of solutions in the class of inverse marginal functions.

We have finally provided an explicit algorithm that yields the forward performance process, the optimal portfolio and the associated optimal wealth processes.

There are a number of possible future research directions. Firstly, one may depart from the binomial model to study general discrete-time models allowing for trading to be discrete or continuous. Such models are inherently incomplete and additional difficulties are expected to arise with regards to the derivation of the functional equation for the inverse marginals as well as the existence and uniqueness of its solutions among suitable classes of functions.

A second direction is to enrich the predictable framework by incorporating model ambiguity. This will allow for the specification of all possible market models only one evaluation period ahead, thus offering flexibility to narrow down the most realistic models period-by-period as the market evolves.

From the theoretical point of view, an interesting question is to investigate whether predictable forward performance processes converge to their continuous-time counterparts. While this is naturally and intuitively expected, conditions on the appropriate convergence scaling need to be imposed, which might be quite challenging due to the ill-posedness of the problem. Such results may also shed light to deeper questions on the construction of continuous-time forward performance criteria related to the appropriate choice of their volatility, finite-dimensional approximations, Markovian or path-dependent cases, among others.

A Proof of Theorem 4

We start with the following auxiliary result, showing that (7) is equivalent to

$$U_0(I_0(y)) = E_{\mathbb{P}}(U_1(I_1(\rho_1 y))); \quad y > 0. \quad (21)$$

Lemma 11. *Suppose that $U_0, U_1 \in \mathcal{U}$ and let I_0 and I_1 be respectively their inverse marginals. Then, (7) holds if and only if (21) holds.*

Proof. We first show that (7) implies (21). Indeed, standard results yield that (7) implies

$$U_0(x) = E_{\mathbb{P}}\left[U_1\left(I_1(\rho_1 U'_0(x))\right)\right]; \quad x > 0,$$

and (21) is obtained by the change of variable $y = U'_0(x)$.

Next, we show that (21) yields (7). Define the value function \tilde{U} by

$$\tilde{U}(x) = \sup_{\mathcal{A}(x)} E_{\mathbb{P}}[U_1(X)]; \quad x > 0.$$

We claim that $\tilde{U} \equiv U_0$. Let \tilde{I} be the inverse marginal of \tilde{U} . By (i), one must then have

$$\tilde{U}(\tilde{I}(y)) = E_{\mathbb{P}}\left[U_1(I_1(\rho_1 y))\right]; \quad y > 0,$$

and it follows that $\tilde{U}(\tilde{I}(y)) = U_0(I_0(y))$, for $y > 0$.

Differentiating with respect to y yields $\tilde{I}' \equiv I'_0$. Therefore $\tilde{I}(y) = I_0(y) + C$, $y > 0$, for some constant C . Taking the limit as $y \rightarrow \infty$ and using the Inada condition $\tilde{I}(\infty) = I_0(\infty) = 0$ yields $C = 0$. Therefore, we obtain $\tilde{I} \equiv I_0$, which implies $\tilde{U}'(x) = U'_0(x)$, for all $x > 0$. Finally, we obtain

$$\tilde{U}(x) = E_{\mathbb{P}}\left[U_1(I_1(\rho \tilde{U}'(x)))\right] = E_{\mathbb{P}}\left[U_1(I_1(\rho U'_0(x)))\right] = U_0(x); \quad x > 0. \quad \square$$

Proof of Theorem 4. (i): From (14) it follows that

$$U_1(x) := U_0(1) + p \int_{x_u(1)}^x I_1^{-1}(\xi) d\xi + (1-p) \int_{x_d(1)}^x I_1^{-1}(\xi) d\xi; \quad x > 0,$$

where $x_u(\cdot)$ and $x_d(\cdot)$ are given by

$$x_i(c) = I_1(\rho^i U'_0(c)); \quad c > 0, \quad i = u, d. \quad (22)$$

Thus,

$$U'_1(x) = p I_1^{-1}(x) + (1-p) I_1^{-1}(x) = I_1^{-1}(x); \quad x > 0.$$

It then follows that I_1 is the inverse marginal of U_1 and that U_1 is a utility function.

(ii): Define the function F by

$$F(x, c) := U_0(c) + p \int_{x_u(c)}^x I_1^{-1}(\xi) d\xi + (1-p) \int_{x_d(c)}^x I_1^{-1}(\xi) d\xi; \quad (x, c) \in \mathbb{R}^+ \times \mathbb{R}^+, \quad (23)$$

with $x_u(c)$ and $x_d(c)$ as in (22). We claim that

$$\frac{\partial F}{\partial c}(x, c) = 0; \quad x, c > 0.$$

Differentiating (23) with respect to c and then using that $I_1^{-1}(x_i(c)) = \rho^i U'_0(c)$, for $c > 0$, we have

$$\begin{aligned} \frac{\partial F}{\partial c}(x, c) &= U'_0(c) - p x'_u(c) G(x_u(c)) - (1-p) x'_d(c) G(x_d(c)) \\ &= U'_0(c) - p x'_u(c) \rho^u U'_0(c) - (1-p) x'_d(c) \rho^d U'_0(c) \\ &= U'_0(c) \left(1 - p \rho^u x'_u(c) - (1-p) \rho^d x'_d(c) \right) = 0. \end{aligned}$$

To obtain the last equation, note that (9) is equivalent to

$$I_0(y) = p \rho^u I_1(y \rho_u) + (1-p) \rho^d I_1(y \rho_d); \quad y > 0.$$

Therefore, substituting $y = U_0(c)$ and differentiating with respect to c yield

$$\begin{aligned} 1 &= \frac{d}{dc} \left(I'_0(U'_0(c)) \right) = \frac{d}{dc} \left(p \rho^u I_1(\rho_u U'_0(c)) + (1-p) \rho^d I_1(\rho_d U'_0(c)) \right) \\ &= p (\rho^u)^2 I'_1(\rho_u U'_0(c)) U''_0(c) + p (\rho^d)^2 I'_1(\rho_d U'_0(c)) U''_0(c) \\ &= p \rho^u x'_u(c) + (1-p) \rho^d x'_d(c). \end{aligned}$$

Note that, by definition, $U_1(x) = F(x, 1)$. Since we have showed that $\frac{\partial F}{\partial c} \equiv 0$, we must have $U_1(x) = F(x, c)$, for all $x > 0$ and $c > 0$. In other words, for all $x, c \in \mathbb{R}^+$, U_1 satisfies

$$U_1(x) = U_0(c) + p \int_{x_u(c)}^x I_1^{-1}(\xi) d\xi + (1-p) \int_{x_d(c)}^x I_1^{-1}(\xi) d\xi.$$

On the other hand, as it was shown in (i), $U'_1 \equiv I_1^{-1}$. Therefore, for all $x > 0$ and $c > 0$,

$$U_1(x) = U_0(c) + p \left(U_1(x) - U_1(x_u(c)) \right) + (1-p) \left(U_1(x) - U_1(x_d(c)) \right),$$

which, in turn, yields that

$$U_0(c) = p U_1(x_u(c)) + (1-p) U_1(x_d(c)) = E_{\mathbb{P}} \left[U_1(I_1(\rho_1 U'_0(c))) \right]; \quad c > 0.$$

This is equivalent to (21). Hence, (ii) follows from Lemma 11.

(iii): This part follows easily from existing results in the classical expected utility problems, if we view (7) as a terminal expected utility problem with U_1 now given and U_0 being its value function. \square

B Proof of Theorem 8

We only show part (i) and the corresponding statements in parts (iii) and (iv), since (ii) follows from similar arguments.

(i) Direct substitution shows that if the infinite series in (17) converges, then I_1 satisfies (9). Thus, to show (i), it only remains to show that the series converges. Note that (17) can be written, for $y > 0$, as

$$I_1(y) = \frac{b}{1+b} y^{\log_a b} \sum_{m=0}^{\infty} (-1)^m \Psi_0(a^m y), \quad (24)$$

which, by the Leibniz test for alternating series, converges if $\lim_{m \rightarrow \infty} \Psi_0(a^m y) = 0$ monotonically. The fact that $\lim_{m \rightarrow \infty} \Psi_0(a^m y) = 0$ follows directly from either of the conditions in (i) on a and Ψ_0 . To show that the convergence is monotonic, note that by (16)

$$\Psi_0(a^{m+1} y) - \Psi_0(a^m y) = b^{-m-1} y^{-\log_a b} \Phi_0(a^m y); \quad y > 0, \quad m = 0, 1, \dots \quad (25)$$

On the other hand, since Φ_0 is increasing and $\lim_{y \rightarrow \infty} \Phi_0(y) = \lim_{y \rightarrow \infty} I_0(a c y) - b I_0(c y) = 0$ by the Inada condition, we must have $\Phi_0(y) < 0$, for $y > 0$. Thus, by (25), we have that $\Psi_0(a^m y) > \Psi_0(a^{m+1} y)$ and $\lim_{m \rightarrow \infty} \Psi_0(a^m y) = 0$ monotonically.

(iii) First, we prove that I_1 is strictly decreasing. Indeed, (24) and (25) yield

$$I_1(y) = \frac{b}{1+b} y^{\log_a b} \sum_{m=0}^{\infty} \left(\Psi_0(a^{2m} y) - \Psi_0(a^{2m+1} y) \right) = -\frac{1}{1+b} \sum_{m=0}^{\infty} b^{-2m} \Phi_0(a^{2m} y).$$

It then follows that, for $y < y'$,

$$I_1(y') - I_1(y) = \frac{1}{1+b} \sum_{m=0}^{\infty} b^{-2m} \left(\Phi_0(a^{2m} y) - \Phi_0(a^{2m} y') \right) < 0,$$

where the inequality holds because Φ_0 is strictly increasing.

Using equation (9), that $a, b, c > 0$ and $\lim_{y \rightarrow \infty} I_0(y) = 0$, and the monotonicity of I_1 , we deduce that $\lim_{y \rightarrow \infty} I_1(y) = 0$, and, hence, that $I_1(y) > 0$, $y > 0$. Similarly, the fact that $\lim_{y \rightarrow 0^+} I_0(y) = \infty$ yields $\lim_{y \rightarrow 0^+} I_1(y) = \infty$. Thus, we have shown that $I_1 \in \mathcal{I}$.

Finally, conditions in Lemma 7 follow from $\Psi_0(y) \rightarrow 0$, as either $y \rightarrow 0^+$ or $y \rightarrow \infty$, and from

$$0 < y^{\log_a b} I_1(y) = \frac{I_1(y)}{I_0(c y)} \Psi_0(y) < \frac{b+1}{b} \Psi_0(y); \quad y > 0,$$

where we used (9) and that $I_1(y) > 0$ to obtain

$$\frac{I_1(y)}{I_0(c y)} = \frac{(1+b)I_1(y)}{I_1(a y) + b I_1(y)} < \frac{1+b}{b}.$$

(iv) Repeating the last part of the argument in part (iii) for any solution $\tilde{I} > 0$ yields that \tilde{I} satisfies the same uniqueness condition for (9) as I_1 . The result then follows directly from Lemma 7.

C Proof of Corollary 9

Assertion (ii) follows from (i) and Theorem 4. Also, one can easily check that I_1 given by (19) is an inverse marginal satisfying (9).

It only remains to show the uniqueness of solutions that are inverse marginals. To this end, it suffices to check that the condition of Theorem 8 holds for all possible values of the parameters. Setting $G(y) = y^{-\theta}$, $y > 0$, in (16) yields

$$\Phi_0(y) = (a^{-\theta} - b)c^{-\theta}y^{-\theta} \quad \text{and} \quad \Psi_0(y) = y^{-(\theta + \log_a b)}.$$

Since $\theta \neq -\log_a b$ and $a \neq 1$, we have the following dichotomy:

- a) Either $\theta < -\log_a b$ and $a < 1$ or $\theta > -\log_a b$ and $a > 1$. Then, one can show that conditions (i) of Theorem 8 hold.
- b) Either $\theta < -\log_a b$ and $a > 1$ or $\theta > -\log_a b$ and $a < 1$. Then, one can show that conditions (ii) of Theorem 8 hold.

References

- O.E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(2):253–280, 2002.
- F. Black. Individual investment and consumption under uncertainty. In D. L. Luskin, editor, *Portfolio Insurance: A Guide to Dynamic Hedging*, pages 207–225. John Wiley and Sons, New York, 1988. First version: November 1, 1968, Financial Note No. 6B, Investment and consumption through time.
- S. Dreyfus. Richard Bellman on the birth of dynamic programming. *Operations Research*, 50(1):48–51, 2002.
- N. El Karoui and M. Mrad. An exact connection between two solvable SDEs and a nonlinear utility stochastic PDE. *SIAM Journal on Financial Mathematics*, 4(1):697–736, 2013.
- E. F. Fama and K. R. French. Business conditions and expected returns on stocks and bonds. *Journal of financial economics*, 25(1):23–49, 1989.
- T. Geng and T. Zariphopoulou. Temporal and spatial turnpike-type results under forward time-monotone performance criteria. submitted for publication, 2017.
- S. Kallblad, J. Oblój, and T. Zariphopoulou. Time-consistent investment under model uncertainty: the robust forward criteria. *to appear in Finance and Stochastics*, 2017.
- M. Kuczma, B. Choczewski, and R. Ger. *Iterative Functional Equations*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1990.
- M. Musiela and T. Zariphopoulou. Backward and forward utilities and the associated pricing systems: The case study of the binomial model. In R. Carmona, editor, *Indifference Pricing*. Princeton University Press, 2003.

- M. Musiela and T. Zariphopoulou. Investments and forward utilities. Technical report, 2006. URL <http://www.ma.utexas.edu/users/zariphop/pdfs/TZ-TechnicalReport-4.pdf>.
- M. Musiela and T. Zariphopoulou. Portfolio choice under dynamic investment performance criteria. *Quantitative Finance*, 9(2):161–170, 2009.
- M. Musiela and T. Zariphopoulou. Stochastic partial differential equations in portfolio choice. In C. Chiarella and A. Novikov, editors, *Contemporary Quantitative Finance*, pages 195–216. Springer-Verlag Berlin Heidelberg, 2010a.
- M. Musiela and T. Zariphopoulou. Portfolio choice under space-time monotone performance criteria. *SIAM Journal on Financial Mathematics*, 1(1):326–365, 2010b.
- M. Musiela and T. Zariphopoulou. Initial investment choice and optimal future allocations under time-monotone performance criteria. *International Journal of Theoretical and Applied Finance*, 14(01):61–81, 2011.
- M. Musiela, E. Sokolova, and T. Zariphopoulou. Forward exponential indifference valuation in an incomplete binomial model. In J. Kallsen and A. Papapantoleon, editors, *Advanced Modelling in Mathematical Finance, In Honour of Ernst Eberlein*, pages 277–302. Springer, 2016.
- S. Nadtochiy and M. Tehranchi. Optimal investment for all time horizons and Martin boundary of space-time diffusions. *Mathematical Finance*, 2015.
- S. Nadtochiy and T. Zariphopoulou. A class of homothetic forward investment performance processes with non-zero volatility. In Y. Kabanov, M. Rutkowski, and T. Zariphopoulou, editors, *Inspired by Finance: The Musiela Festschrift*, pages 475–504. Springer, 2014.
- M. Shkolnikov, R. Sircar, and T. Zariphopoulou. Asymptotic analysis of forward performance processes in incomplete markets and their ill-posed HJB equations. *SIAM Journal on Financial Mathematics*, 7(1):588–618, 2016.