Recovery of Preferences from Observed Wealth in a Single Realization

Philip H. Dybvig and L. C. G. Rogers*

First draft: September, 1994.
This draft (sampreco9d.tex): December, 1995.

Abstract

Von Neumann-Morgenstern preferences over terminal consumption can be inferred from wealth on a single sample path when markets are complete and returns follow a known law in a neoclassical investment problem in either a discrete-time i.i.d. binomial model or a continuous-time diffusion model with a Gaussian state variable.

*Washington University in Saint Louis and University of Bath, respectively. We are grateful for helpful comments from Greg Willard and for partial support from EPSRC grant number GR/K00448.

Most of *demand theory* focuses on the derivation of optimal demand and its properties given preferences. The theory of *revealed preference*[1] focuses on the reverse problem: what does observation of optimal demand tell us about preferences? These results address both existence and uniqueness of preferences within a particular class that are consistent with a given demand function. Traditionally, revealed preference theory has related the entire demand function to preferences in a minimally restricted class. A newer literature on *recoverability* (or *identifiability*) of preferences, initiated by Green, Lau, and Polemarchakis (1978) has explored the question of whether we can learn preferences in a more narrowly defined class from more restricted information about optimal demand. In this present paper, sufficient conditions are given for the recoverability of von Neumann-Morgenstern preferences over terminal consumption from the wealth along a single realized sample path of stock prices. The main assumptions are complete markets and either returns are i.i.d. log-binomial in discrete time or there is a Gaussian state variable in continuous time.

Recoverability results are of interest in their own right because they help us to understand the assumptions we are making and their implications. Perhaps more importantly, recoverability results have implications for empirical demand analysis. In particular, recoverability results provide a guide to what can and cannot be learned from different types of data. For example, if the utility function cannot be recovered from perfect knowledge of certain demand observations, then we cannot hope to recover a nonparametric estimate of preferences from a finite number of those observations, and this will be known without any formal analysis of econometric identification. Conversely, if the utility function can be recovered from perfect knowledge of certain demand observations, then it is likely that identification will follow from finitely many observations given almost any parametric restriction or smoothing for interpolation. In fact, an early version of the argument in this paper helped to motivate a proposal by Hodges (1991) of an interesting procedure for evaluating portfolio performance.

In the binomial model, recoverability of preferences is closely tied to the path independence result of Cox and Leland (1982). They note that when returns

[1] Richter (1987) is a good brief review of revealed preference theory.

are i.i.d. in a binomial model, optimal wealth is path-independent, that is, it depends only on the reinvested value of the stock, not on the whole history of stock prices. Path-independence says that we need only solve for wealth on an "ingrown tree" that identifies nodes with identical stock prices at a point in time but different histories. We can then start from knowledge of the strategy on a single path to step through the entire tree using the one-period valuation equation (that follows from absence of arbitrage as derived by Cox, Ross, and Rubinstein (1979)). Knowing the entire wealth process is equivalent to knowing the entire portfolio strategy, and is also equivalent to knowing marginal rates of substitution of consumption at the wealth levels of any pair of terminal nodes, which determines the marginal utility at all terminal wealth levels that can be realized by the optimal strategy, up to rescaling by a constant that does not change preferences.

We have the intuition that recoverability is more common in continuous time than in discrete time. Our proof in continuous time admits a variety of non-i.i.d. returns including an embedded Gaussian term structure model. The critical assumption is that the conditional distribution at t of the terminal state-price density is log-Gaussian, although we conjecture recoverability is much more general.

For the continuous-time result, the proof relies on a general version of the path-independence result of Cox and Leland (1982), using a new state variable which is the expectation at t of the log of the terminal state-price density. The proof uses the properties of diffusion processes to show that all derivatives of wealth with respect to the state variable can be computed along the sample path, which gives the wealth as a function of the state variable at a point in time given analyticity. This analyticity follows from the smoothing effect of conditional expectations given the filtration generated by the Wiener process. By analyticity, this determines the wealth as a function of the state variable at a certain point in time, and Fourier analysis allows us to determine the terminal wealth. This paper is related to several other papers on recovery of preferences for uncertain consumption. Dybvig and Polemarchakis (1981) considers recoverability of twice differentiable von Neumann-Morgenstern preferences from portfolio choice for all wealth levels and all security prices in a one-period model with a riskless asset and a risky asset. Dybvig (1982) considers recoverability of preferences over gambles in

real wealth given preferences over gambles in nominal wealth that include price level uncertainty. The use of Fourier analysis in one part of the proof in this paper is reminiscent of the use of Fourier analysis in that paper. Wang (1993a, 1993b) are recoverability results in respectively continuous and discrete time that are closest to the spirit of this paper, although those results are significantly different since they are based on knowing the entire law for the portfolio process in a model with consumption withdrawal, while the present paper is based on knowing the portfolio process in a single realization in a model with or without consumption withdrawal. Another paper, He and Leland (1993), is less similar to this paper than might seem at first. The results in that paper concern recovery of preferences in an economy with a single agent. Unlike many results with additive preferences, those results do not apply to economies with multiple agents with additive preferences but different birth dates or horizons, nor do the results apply to individual agents. For these reasons, their results seem to be of limited empirical interest.

Section 1 contains the discrete-time results, and Section 2 contains the continuous-time results. Section 3 contains simulations indicating how much information can be drawn from noisy portfolio information in a single sample. Section 4 contains a discussion of how the paper's results can be extended to a model with consumption withdrawal. Section 5 briefly closes the paper.

1 Binomial result

First, we consider a binomial model of investment opportunities, which is essentially the one used by Cox, Ross, and Rubinstein (1979) (CRR) to price options. There is a riskless asset (the *bond*) paying r units at the end of the period for each unit invested in the riskless asset at the beginning of the period (r is one plus the interest rate), and there is a risky asset (the *stock*) paying a random amount x_t at time t per unit invested in the risky asset at time $t-1$. We assume that either $x_t = u$ (up), with probability $\pi_u \in (0,1)$, or $x_t = d$ (down), with probability $\pi_d = 1 - \pi_u$. The x_t's are i.i.d. over time. The riskless return r does not vary over time. We will assume that $u > r > d$ to avoid arbitrage and degeneracy, and that $\pi_u u + \pi_d d > r$, which

says that the expected return on the risky asset exceeds the riskless return. The underlying parameters u, d, r, π_u, and the maturity T are assumed known both to the agent making the choice and to the observer trying to infer the utility function.

The investor has a von Neumann-Morgenstern utility function $U : \Re_{++} \to \Re$ that is strictly increasing and strictly concave, with a continuous first derivative that takes on all positive values. These conditions imply the existence of an interior solution with positive consumption in all states. The investor chooses the adapted portfolio choice α_t (the proportion invested in the risky asset) that maximizes the expected utility of terminal wealth given initial wealth w_0, solving Problem 1.

Problem 1 *Choose the adapted process* α_t *to maximize* $E[U(w_0 \prod_{t=1}^{T}(r + \alpha_{t-1}(x_t - r)))]$.

We will indicate the wealth at time t by

$$(1) \quad w_t \equiv w_0 \prod_{\tau=1}^{t}(r + \alpha_{\tau-1}(x_\tau - r)).$$

It is useful to review several well-known properties of the binomial option pricing model. We can represent the state of nature in the binomial model as $(x_1, x_2, ..., x_T)$, since the only uncertainty is the random payoff to the stock, and $(x_1, ..., x_t)$ represents what is known at t. Therefore, we can write wealth at t as $w_t(x_1, ..., x_t)$ and the portfolio choice at t as $\alpha_t(x_1, ..., x_t)$.[2] Define the *state-price density* process as starting at 1 and such that the stock price or bond price (and by extension the value of any portfolio) times the state-price density is a martingale. Then

$$(2) \quad \xi_t = \prod_{s=1}^{t} \psi_{x_s}$$

[2]We will use the following common conventions. A function with an empty argument list (e.g. $w_t(x_0, ..., x_t)$ when $t = 0$) is a constant. An empty sum is 0, and an empty product is 1.

where ψ_u and ψ_d are defined by

$$(3) \quad \psi_u \equiv \frac{r-d}{\pi_u r(u-d)}$$

and

$$(4) \quad \psi_d \equiv \frac{u-r}{\pi_d r(u-d)}$$

Then we have the following valuation equation. Whenever $s, t \in \{0, 1, ..., T\}$, $s > t$,

$$(5) \quad w_t = E_t[\frac{\xi_s}{\xi_t} w_s],$$

where E_t denotes expectation conditional on information at t. In terms of risk-neutral valuation (using equivalent martingale measures), ξ_t is the change-of-measure to risk-neutral probabilities times a discount factor. We use this representation of pricing because of its close connection to agents' first-order conditions.

For $s = t + 1$, (5) says that

$$(6) \quad w_t(x_0, ..., x_t) = \pi_u \psi_u w_{t+1}(x_0, ..., x_t, u) + \pi_d \psi_d w_{t+1}(x_0, ..., x_t, d)$$

which is easily verified from (1) and the definitions of ψ_u and ψ_d. For other s, the result follows by induction using the law of iterated expectations. For example,

$$(7) \quad \begin{aligned} w_t &= E_t[\frac{\xi_{t+1}}{\xi_t} w_{t+1}] \\ &= E_t[\frac{\xi_{t+1}}{\xi_t} E_{t+1}[\frac{\xi_{t+2}}{\xi_{t+1}} w_{t+2}]] \\ &= E_t[\frac{\xi_{t+2}}{\xi_t} w_{t+2}] \end{aligned}$$

proves the result for $s = t + 2$.

The valuation equation (5) is actually a complete characterization of feasible wealth processes. To replicate a payoff of x_u conditional on up and x_d conditional on down in one period, we invest wealth $w = (\pi_u \psi_u x_u + \pi_d \psi_d x_d)$ with a proportion $\alpha = (x_u - x_d)/w(u - d)$ invested in the stock. In general, a positive random terminal wealth w_T is feasible if and only if

$$(8) \quad w_0 = E_0[\xi_T w_T],$$

corresponding to a wealth process

$$(9) \quad w_t = E_t[\frac{\xi_T}{\xi_t} w_T].$$

This wealth process is then generated by a portfolio policy

$$(10) \quad \alpha_t(x_1, ..., x_t) = \frac{w_{t+1}(x_1, ..., x_t, u) - w_{t+1}(x_1, ..., x_t, d)}{w_t(x_1, ..., x_t)(u - d)}.$$

This completes our review of binomial option pricing.

Given the representation of terminal consumption, we can write the following reduced-form version of Problem 1.

Problem 2 *Choose random terminal wealth w_T to maximize $E[U(w_T)]$ subject to $E[\xi_T w_T] = w_0$.*

The choice of portfolio strategy is implicit in the problem and given by (9) and (10). The first-order condition for Problem 2 is the existence of λ such that

$$(11) \quad w_T = I(\lambda \xi_T),$$

where $I(\cdot)$ is the inverse function of marginal utility $U'(\cdot)$.

The following result due to Cox and Leland (1982) is central to our analysis. Cox and Leland refer to the property in the theorem as "path independence," since the various quantities depend only on the reinvested return on the stock, not the whole price path.

Theorem 1 *(Cox and Leland) The shared solution to Problems 1 and 2 has a wealth process that is* path-independent*, i.e., w_t depends on the state $(x_1, x_2, ..., x_t)$ only through the product $x_1 x_2 ... x_t$.*

PROOF It suffices to show that w_t depends only on ξ_t, since $\xi_t = {\psi_u}^{n_u} {\psi_d}^{t-n_u}$, where n_u is the number of ups before or at t, and $n_u = \log((x_1 x_2 ... x_t)/d^t)/\log(u/d)$. By (9) and (11), $w_t = E_t[I(\lambda \xi_T)\xi_T/\xi_t]$, which depends only on ξ_t and t, since ξ_t is Markovian, proving path-independence of w_t. ∎

Now we are prepared for our main result, which is the recoverability of the entire wealth process from its trajectory in a single realization.

Theorem 2 *Assume the return parameter values r, u, d, and π are known. Then, the wealth process for the solution to Problem 1 given a single realization of $x_1, ..., x_T$ determines the entire solution. Therefore we learn as much about preferences from a single realization as we do from seeing the entire strategy. This amount of information is the marginal utility at the different terminal wealths across states, up to the arbitrary normalization that does not affect preferences.*

PROOF Given path-independence (Theorem 1), we can write wealth at t as $w_t = \hat{w}_t(x_1 x_2 ... x_t)$, where the product $x_1 x_2 ... x_t$ can take on the $t+1$ values $u^{n_u} d^{t-n_u}$ for $n_u \in \{0, ..., t\}$, where n_u is the number of ups before or at t. The proof is by induction on t. The inductive hypothesis is that we can use the particular realization to construct the optimal strategy $\hat{w}_s(\cdot)$ for all $s \leq t$. The induction starts trivially, since w_0 is on the observed sample path.

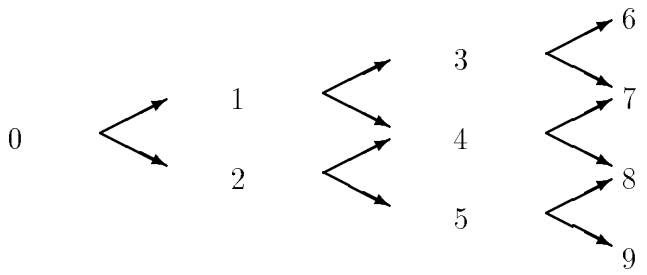

Figure 1 *Path-independence allows us to infer the entire portfolio strategy from the strategy in a single realization. Absence of arbitrage yields a pricing result that relates wealth at a node and the two following nodes. If we know the wealth at nodes* 0, 1, 4, *and* 7, *we can infer the wealth at node* 2 *from the wealth at nodes* 0 *and* 1, *at node* 3 *from nodes* 1 *and* 4, *and similarly all through the tree to determine the entire strategy.*

Now we need to show the induction step: given knowledge of the wealth process in one realization and the entire wealth function $\hat{w}_s$ for $s \leq t$, we must show how to infer the entire wealth function at $t+1$. From (6), we have

$$(12) \quad \hat{w}_t(u^{n_u}d^{t-n_u}) = \pi_u\psi_u\hat{w}_{t+1}(u^{n_u+1}d^{t-n_u}) + \pi_d\psi_d\hat{w}_{t+1}(u^{n_u}d^{t-n_u+1}).$$

Knowledge of wealth along the sample path gives us $\hat{w}_{t+1}(u^{n_u}d^{t+1-n_u})$ for one value of n_u and the inductive hypothesis tells us $\hat{w}_t(u^{n_u}d^{t-n_u})$ for all values of n_u. Furthermore, solving (12) for one of the $\hat{w}_{t+1}$ expressions on the right-hand side allows us to step up and down through all values of n_u and solve for $\hat{w}_{t+1}(u^{n_u}d^{t+1-n_u})$ for all n_u in {0,1,...,t+1}. This completes the induction step, and we have shown that the entire wealth process is determined by one sample path.

Finally, it follows from the first-order condition (11) that what we learn from this is exactly the marginal utility (up to the arbitrary rescaling) on the set of terminal wealth levels achieved by the optimal strategy. ∎

The intuition of the proof of Theorem 2 is shown in Figure 1. Suppose the stock price moves in order up, down, up. Path-independence implies we can consider an ingrown tree as in Figure 1. We are given the wealth at nodes 0,

1, 4, and 7. The valuation equation (or its underlying arbitrage) then imply the wealth at node 2 (from nodes 0 and 1), at node 3 (from nodes 1 and 4), at node 5 (from nodes 2 and 4), at node 8 (from nodes 4 and 7), at node 6 (from nodes 3 and 7), and at node 9 (from nodes 5 and 8). This is the entire wealth function. The first-order condition determines (up to the arbitrary constant) the marginal utility at the terminal wealth levels.

2 Continuous-time result

Our result for continuous time is based on the standard model first analyzed by Merton (1971). We assume that security returns are identically and independently distributed and that uncertainty is driven by an N-dimensional Wiener process Z_t. We assume there is a locally riskless asset bearing a continuously-compounded interest rate following the process r_t and that there are N risky assets whose rates of return over dt are given by $\mu_t dt + \sigma_t dZ_t$, where the adapted N-dimensional process μ_t is the local mean return and the adapted invertible $N \times N$-dimensional process σ_t gives the risk exposures of the assets. Note that invertibility of σ_t is the assumption that markets are locally complete. As before, the law defining returns and the maturity date T are known both to the agent making the choice and to the observer trying to infer preferences.

A portfolio strategy is an adapted N-dimensional process α_t. Initial wealth is exogenously given as w_0, and given α_t the wealth process solves

$$(13) \quad w_t = w_0 + \int_{s=0}^{t} w_s(r_s\, ds + \alpha_s'((\mu_s - r_s\mathbf{1})ds + \sigma_s dZ_s)),$$

where $\mathbf{1}$ indicates a vector of ones. A portfolio process is feasible if the wealth process satisfying (13) is positive for all time with probability 1. This restriction to positive wealth processes is consistent with our conditions on the utility function and also rules out doubling strategies and other arbitrages (see Dybvig and Huang (1988)).

As in the discrete-time case, we will assume the agent has a von Neumann-Morgenstern utility function U of terminal wealth w_T that is strictly increasing and strictly concave, with a continuous first derivative that takes on all positive values. In addition, to ensure existence of an optimum we assume that the inverse marginal utility function $I(\cdot)$ is bounded by some power function plus a constant, i.e., there exist constants k_0, k_1, and p such that, for all $m > 0$,

$$(14) \quad I(m) < k_0 + k_1 m^p.$$

The investor's decision problem follows.

Problem 3 *Choose adapted α_t to maximize $E[U(w_T)]$ s.t. $(\forall t \in [0,T])(w_t \geq 0)$, where the process w_t solves (13).*

The investor's task is to maximize expected utility of terminal wealth subject to wealth staying nonnegative. As in the discrete-time case, there is an equivalent asset-pricing version of the choice problem. Define the state-price density process

$$(15) \quad \xi_t \equiv \exp\left(-\int_0^t r_s ds - \int_0^t \left(\tilde{\mu}_s' dZ_s + \frac{1}{2}\tilde{\mu}_s'\tilde{\mu}_s ds\right)\right),$$

where $\tilde{\mu}_t \equiv {\sigma_t}^{-1}(\mu_t - r_t \mathbf{1})$ and the integrals defining the process are assumed to exist (are finite). Equivalently, $\xi_0 = 1$ and

$$(16) \quad d\xi_t = \xi_t(-r_t dt - \tilde{\mu}_t' dZ_t).$$

The state price density process ξ is designed so that for every feasible strategy the wealth times the state price density is a local martingale (and is a martingale if the strategy is not dominated, see Dybvig and Huang (1988)). The state-price density process is a discounted version of the change of probability to the risk-neutral probabilities.

Problem 4 *Choose random terminal wealth w_T to maximize $E[U(w_T)]$ subject to $E[\xi_T w_T] = w_0$.*

In the discrete problem, we assumed i.i.d. returns, which would be analogous to assuming that μ, r, and σ are constant in the continuous-time problem, which is much stronger than we need to assume. Our proof relies on conditional lognormality of ξ_T, but we conjecture that these conditions can be weakened much further and that recoverability is quite general.

Assumption 1 *Conditional on $\mathcal{F}_t$, $\log \xi_T$ has a Gaussian distribution with mean $y_t \equiv E_t(\log \xi_T)$ and variance $v_t \equiv \text{var}(\log \xi_T | \mathcal{F}_t)$, where v is a strictly decreasing continuous adapted process.*

We think of normality as the most important part of this assumption: given that v_t is the variance of ξ_T conditional on information at t, it is necessarily nondecreasing in time. In most applications, it will either automatically be strictly decreasing or will be strictly decreasing given some minimal nondegeneracy condition. Since the state-price density process is not a primitive in the problem, it is useful to provide a couple of examples of classes of processes for which the assumption is satisfied.

Example 1 Deterministic parameters: *Suppose that μ_t, r_t, and σ_t are deterministic, and for all t, $\mu_t - r_t \mathbf{1} \neq 0$ and σ_t is invertible, and that the integrals defining ξ_t in (15) exist. Then from (15), ξ_t is a lognormal diffusion and $\log(\xi_t)$ has independent and nondegenerate increments. If in addition μ_t, r_t, and σ_t are constant, returns are i.i.d. over time and so are the increments of $\log(\xi_t)$.*

The return processes in the second class can have embedded a model of bond pricing with a Gaussian interest rate process, as well as nontrivial models with both stocks and bonds.

Example 2 Gaussian interest rates: *Suppose that $r_t = r_t^e + \int_{s=0}^t h(s,t) dZ_s$ where r_t^e is deterministic, and $h(s,t)$ is continuous and bounded for all $t >$*

s. Suppose further that for all t, $\tilde{\mu}_t$ is deterministic, and that the integrals defining ξ_t in (15) exist. Then, excepting degenerate cases, $\log(\xi_t)$ *and* r_t *are Gaussian, but the returns and increments of* $\log(\xi_t)$ *are not independent over time.*[3] *Special cases of this model include some Gaussian term structure models of Vasicek (1977). For example, a one-factor mean-reverting Vasicek model with unpriced interest rate risk can be obtained by taking* $N = 1$, $\sigma_t > 0$ *constant,* $r_t^e = \overline{r} + (r_0 - \overline{r})\exp(-\kappa t)$, $h(s,t) = \sigma_0 \exp(-\kappa(t-s))$, *and* $\mu_s \equiv r_s$.

Given our assumptions, there exists a unique optimum.[4] The first-order condition for the solution is the budget constraint $E[\xi_T w_T] = w_0$ and

$$(17) \quad U'(w_T) = \lambda \xi_T$$

for some Lagrange multiplier $\lambda > 0$. This first-order condition and our regularity conditions on U imply together that the solution is a positive random variable that has full support on $\Re_{++}$ (as does ξ_t). It is a standard result that for the solution, adjusted wealth $\xi_t w_t$ is given by the martingale

$$(18) \quad \xi_t w_t = E_t[\xi_T w_T].$$

The portfolio choice can be inferred from the adjusted wealth process. The portfolio choice α_t must make the local change in $\xi_t w_t$ implicit in (13) and (15) the same as what comes from (18). Specifically, adjusted wealth has a martingale representation

$$(19) \quad d(\xi_t w_t) = \beta_t' dZ_t.$$

[3]By the definition of ξ_t and the assumptions made here, $\log(\xi_t)$ is a stochastic integral with fixed coefficients, and is therefore Gaussian. The conditional variance can never be decreasing (by the law of iterated expectations); the slight non-degeneracy is needed to imply conditional variance is actually increasing.

[4]With a small amount of work we could apply a result in Cox and Huang (1990) given the bound on I and the fact that ξ_T possesses all moments. These facts imply that $E_0[\xi_T I(\lambda \xi_T)]$ is continuous and bounded and ranges through all positive reals as λ varies. When this equals w_0, then $w_T = I(\lambda \xi_T)$ is a solution satisfying the first-order conditions.

But we also know from (13) and (16) that

$$(20)\quad d(\xi_t w_t) = \xi_t w_t(\alpha_t'\sigma_t - \tilde{\mu}_t')dZ_t$$

and therefore

$$(21)\quad \alpha_t = (\sigma_t')^{-1}\left(\tilde{\mu}_t + \frac{\beta_t}{\xi_t w_t}\right)$$

can be inferred from the adjusted wealth process and the return processes. This shows how to construct the supporting portfolio strategy given a wealth process satisfying (18).

If we had assumed a single security and independent returns or something slightly weaker, we would have another path-independence result from Cox and Leland (1982), as in the binomial case, saying that wealth depends only on the stock price and time and not on the whole history. Given our Assumption 1, we have another type of path-independence result saying that adjusted wealth depends only on $y_t \equiv E_t[\log \xi_T]$ and $v_t \equiv \text{var}(\log \xi_T|\mathcal{F}_t)$.

Theorem 3 *The shared solution to Problems 3 and 4 has an adjusted wealth process that is path-independent in the sense that adjusted wealth $\xi_t w_t$ depends on the state $\{Z_s\}$, $s \in [0,t]$, only through y_t and v_t,*[5] *i.e. $\xi_t w_t = W(y_t, v_t)$ for some function W possessing all own and cross partial derivatives.*

PROOF From Assumption 1, (17) and (18),

$$(22)\quad \begin{aligned} \xi_t w_t &= E_t[\xi_T I(\lambda\xi_T)] \\ &= \int e^x I(\lambda e^x) p(v_t, x - y_t)dx \\ &\equiv W(y_t, v_t), \end{aligned}$$

[5]In many cases of interest, v_t is deterministic and y_t is the only state variable.

say, where $p(t,z) = (2\pi t)^{-1/2}\exp(-z^2/2t)$ is the Gaussian density. To prove that all derivatives of W exist, we notice that for any non-negative integers n and k, and for any positive C and ε,

$$\begin{aligned}\sup_{|z|\le C}\sup_{\varepsilon\le t\le\varepsilon^{-1}}\left|\frac{\partial^{n+k}p}{\partial y^n\partial t^k}(t,x-y)\right| \;&\le\; P(x)(\exp(-\varepsilon(x-C)^2/2)\\ &\quad+\exp(-\varepsilon(x+C)^2/2))\end{aligned}$$

where P is some polynomial (whose coefficients depend on n, k, C and ε). Now in conjunction with the bound (14) we can deduce that on any compact set all derivatives of W exist. ∎

Now we are prepared to state our main result for the continuous model. In the case that there is a constant spot rate and a single stock with i.i.d. returns, knowing y_t is the same as knowing the accumulated stock return, and the result is fully analogous to Theorem 2. However, much less structure is required for this result.

Theorem 4 *Under Assumption 1, the wealth process for the solution to Problem 3 is determined (a.s.) by a single realization of adjusted wealth $\xi_t w_t$ and the state variable y_t in an open interval $N = (0, t^*)$ of times $t < T$ for some $t^* > 0$. Then, this determines the utility function $U(\cdot)$ up to an affine transform which does not affect preferences.*

PROOF Recall that $\langle y\rangle$, the quadratic variation process of the martingale y, is the unique continuous increasing adapted process A vanishing at 0 for which $y_t^2 - A_t$ is a martingale. Consider the process $v_t + \langle y\rangle_t$, which we will show to be a finite-variation martingale and therefore constant (e.g. by Theorem IV.30.4 in Rogers and Williams (1987)). The process $v_t + \langle y\rangle_t$ has finite variation since v_t is decreasing (by Assumption 1) and $\langle y\rangle_t$ is nondecreasing (by definition), and $v_t + \langle y\rangle_t$ is a martingale since we can write it as the sum of two martingales: $v_t + \langle y\rangle_t = (v_t + y_t^2) + (\langle y\rangle_t - y_t^2) = E_t[(\log\xi_T)^2] + (\langle y\rangle_t - y_t^2)$. Since $v_T = 0$ and $\langle y\rangle_0 = 0$, we have $\langle y\rangle_T = v_0$. It is well known that there exists some Brownian motion $\mathcal{Z}$ such that the continuous martingale y is expressed as $y_t = \mathcal{Z}(\langle y\rangle_t) = \mathcal{Z}(v_0 - v_t)$; see, for example, Theorem IV.34.11

of Rogers and Williams (1987). We shall write γ for the (continuous) inverse, state-by-state, to the process $t \mapsto \langle y \rangle_t$.

Since we see the entire process y on $N = (0, t^*)$, we may deduce the quadratic variation process $\langle y \rangle$ there almost surely, by taking discrete approximations to the quadratic variation; the proof of this well-known fact is given in various places in the literature, for example, Theorem IV.30.1 of Rogers and Williams (1987).

Note that since N is open and γ is continuous, $\gamma^{-1}(N)$ is open. Fix any $t \in N$. By a standard application of the law of the iterated logarithm (e.g. Arnold (1974, section 3.1)), there exists a sequence of times $t_n \in N$, $t_n \downarrow t$, such that

$$
\begin{aligned}
(23) \quad 1 &= \limsup_{s \downarrow t} \frac{\mathcal{Z}(\langle y \rangle_s) - \mathcal{Z}(\langle y \rangle_t)}{h(\langle y \rangle_s - \langle y \rangle_t)} \\
&= \lim_{n \uparrow \infty} \frac{\mathcal{Z}(\langle y \rangle_{t_n}) - \mathcal{Z}(\langle y \rangle_t)}{h(\langle y \rangle_{t_n} - \langle y \rangle_t)} \\
&= \lim_{n \uparrow \infty} \frac{y(t_n) - y(t)}{h(v_{t_n} - v_t)}
\end{aligned}
$$

where we have used the abbreviation $h(\Delta) \equiv (2\Delta \log\log(1/\Delta))^{1/2}$. Since $\lim_{\Delta \downarrow 0} \Delta / h(\Delta) = 0$, it follows that

$$
(24) \quad \lim_{n \uparrow \infty} \frac{v_{t_n} - v_t}{y_{t_n} - y_t} = 0.
$$

Pick any such sequence, by "inspection" of y in N. Along this sequence, we have by a.s.-continuity of y_t that $y_{t_n} \to y_t$ in $\Re$ and $(y_{t_n}, t_n) \to (y_t, t)$ in $\Re^2$.

Since all derivatives exist, the total derivative of W is tied to the partial derivatives, and we have

$$
(25) \quad \lim_{n \uparrow \infty} \frac{W(y_{t_n}, v_{t_n}) - W(y_t, v_t) - (y_{t_n} - y_t) W_y(y_t, v_t) - (v_{t_n} - v_t) W_t(y_t, t)}{\left((y_{t_n} - y_t)^2 + (v_{t_n} - v_t)^2\right)^{1/2}} = 0.
$$

Now, the choice of t_n's to satisfy (23) implies by (24) that

$$(26)\quad \lim_{n\uparrow\infty} \frac{v_{t_n} - v_t}{((y_{t_n} - y_t)^2 + (v_{t_n} - v_t)^2)^{1/2}} = 0$$

and

$$(27)\quad \lim_{n\uparrow\infty} \frac{y_{t_n} - y_t}{((y_{t_n} - y_t)^2 + (v_{t_n} - v_t)^2)^{1/2}} = 1.$$

Therefore, (25) becomes

$$(28)\quad \lim_{n\uparrow\infty} \frac{W(y_{t_n}, v_{t_n}) - W(y_t, v_t)}{y_{t_n} - y_t} = W_y(y_t, v_t).$$

Since $W(y_s, v_s) = \xi_s w_s$, the left-hand expression is known along the sample path, and therefore $W_y(y_t, v_t)$ is known almost surely there. This procedure computes $W_y(y_t, v_t)$ almost surely for a single $t \in N$, and sigma-additivity says that it will compute $W_y(y_t, v_t)$ almost surely for a countable dense set of $t \in N$, and for all $t \in N$ by continuity. Therefore, $W_y(y_t, v_t)$ is determined almost surely from knowing $\xi_t w_t$ and y_t on N in one realization. The same construction can be repeated by induction to obtain all the derivatives $\partial^k W/\partial y^k$. Since there are countably many derivatives, the construction will work almost surely for all of them at once.

Now pick any $t \in N$. To show that knowledge of all partials with respect to y at one value of v and one value of y implies knowledge of $W(y, v)$ for all y, we need to show that $W(y, v)$ is entire (analytic on the entire complex plane). From (22), we can write $W(y, v)$ as the product of

$$(29)\quad e^{-\frac{y^2}{2v}},$$

which is entire in y, and

$$(30) \quad \int_{x=-\infty}^{\infty} e^{-\frac{xy}{v}} H(x)dx$$

where

$$(31) \quad H(x) \equiv I(\lambda e^x) \frac{e^{-\frac{x^2}{2v}+x}}{\sqrt{2\pi v}}.$$

For each real y, the bound (14) on $I(\cdot)$ implies that the expression (30) is finite, so Lemma 1 (in the Appendix) implies that the expression (30) is entire in y. As a product of two entire functions, $W(y, v_t)$ is also entire in y, which shows that we can infer the function $W(y, v_t)$ of y on the entire complex plane from all its partials with respect to y, and therefore (for any $t \in N$) from wealth and the stock process in N for a single sample path.

To infer the function $H(x)$ of x from the function $W(y, v)$ of y for some $v > 0$, we will make use of the theory of Fourier transforms (as in Rudin (1974, chapter 9)). Recall that the Fourier transform of a function $f : \Re \to \Re$ is given by[6]

$$(32) \quad \hat{f}(q) \equiv \int_{x=-\infty}^{\infty} e^{iqx} f(x) \frac{dx}{\sqrt{2\pi}}.$$

Since $W(y, v)$ is entire, it is well-defined to use a complex argument y. Since $W(y, v)$ is the product of expressions (29) and (30), it follows that

$$(33) \quad \frac{W(iy, v)}{e^{-y^2/2v}} = \hat{H}\left(-\frac{y}{v}\right).$$

We established earlier that knowledge of $\xi_t w_t$ and y_t on N implies knowledge of $W(y, v)$ for all complex y and some $v > 0$. Hence, by (33), this also

[6]Some authors, including Rudin, substitute $-i$ for i in the definition. (Perhaps they are labelling the two roots of -1 differently than we do!) The exact definition does not matter so long as we are consistent.

implies knowledge of $\hat{H}$. But the bound (14) on $I(\cdot)$ implies that $H \in L^2(\Re)$. Since the Fourier transform is an L^2-isometry, knowledge of $\hat{H}$ is equivalent to knowledge of H (via the inversion formula). But knowledge of H is knowledge of the marginal utility $U'(\cdot)$ up to the constant λ, variation of which is a rescaling that does not change preferences. Obviously, knowledge of the preferences implies knowledge of the whole solution to the choice problem, and we are done. ∎

3 Implications for Empirical Research

Our recoverability results (Theorems 2 and 4) in previous sections are suggestive that it may be possible to obtain a reasonable estimate of preferences from one or a few sample paths. In the binomial model, the result is only suggestive because actual stock returns are not in the support of the model and because the result only implies knowledge of marginal utilities of realized wealths. In the continuous model, the result is only suggestive because observations are discrete while the result assumes observation of the continuous sample path. In both cases, transaction costs or other trading considerations may introduce small errors that are mathematically important even if they are not very important economically. However, none of these problems are necessarily important barriers to estimation in a statistical model with a restriction to a parametric class of utility functions:[7] the parametric restriction and the statistical analysis can smooth over any gaps or error terms in the observations. This section reports simulations that illustrate useful estimation of preferences from 5 years of monthly observations of a single sample path in continuous time. More technical details are contained in the Appendix. Our results are consistent with the findings of Hodges (1991), who discusses not only estimation of the preferences of an agent following an optimal strategy, but also estimation of the degree of inefficiency of an agent following a possibly suboptimal strategy. In general, our simulations are intended to persuade the reader that this is a fruitful area; ultimate application will be

[7] Or, estimation should indeed be possible in a nonparametric model given sufficient smoothing in the estimation procedure.

influenced by the institutions and the nature of the available data.

One case that would be very easy to analyze is the case of constant relative risk aversion with lognormal stock returns. In that case, the portfolio choice is a fixed proportion of wealth, and estimation of preferences would be trivial. Another case, translated power utility, is also easy to analyze, since the risky portfolio holding is linear in wealth and the discount factor (as in the examples in Dybvig, Rogers, and Back (1995)). The simulations here indicate that estimation is possible for more general preferences. Specifically, the simulations assume a utility function that exhibits possibly different levels of relative risk aversion γ_0^{-1} below and γ_1^{-1} above the initial wealth level. This particular specification could be used to test whether risk preferences are the same for increases and decreases compared to initial wealth. This is a simple example to exposit; it is not much more complicated to have multiple breakpoints.[8] Depending on the location and number of breakpoints, useful estimation may or may not still be possible; we expect that useful estimation should be possible if there are not too many breakpoints and they are neither too close together nor too far from typical values of terminal wealth.

For security prices, we assume the standard lognormal model with a single stock for which μ, r, and σ are constant across time and states of nature and satisfy $\mu > r$ and $\sigma > 0$. To make the estimation problem more realistic, we add noise to the portfolio choice at each data point (consistent with practice given avoidance of transaction costs and model misspecification) and, less importantly, we add noise to the theoretical portfolio return (consistent with practice given tracking error). Specifically, the portfolio choice at each time period is multiplied by a unit-mean lognormal variate whose log has a standard deviation of 2% (200 basis points) or 5% (500 basis points) and wealth after each time increment is subject to a unit-mean lognormal shock whose log has a standard deviation of 0.5% (50 basis points). These standard deviations are intended to be conservative (relatively large numbers), and the 5% error in the portfolio choice implies that the manager is doing a very

[8] As usual, it is more problematic to have a breakpoint in an unknown location since, for example, the location of the breakpoint is not identified when the coefficients of relative risk aversion are equal. This invalidates the usual asymptotic analysis of the null hypothesis of equal risk aversion above and below the breakpoint, since not all parameters are identified under the null hypothesis.

imprecise job of following the strategy. The portfolio value to which noise is added is the correct one given the actual wealth and time to maturity, not what would have been predicted without tracking error for that level of the stock price at that time–this is consistent with the forward-looking perspective of managers, who base their portfolio on how much wealth they actually have and not some hypothetical amount they should have. For simplicity, we take all the noise terms and the normalized stock price innovations to be mutually independent. This is like assuming that the portfolio deviations are very short-lived, since otherwise the tracking error in a period would be related to the stock performance times the persistent part in the portfolio error. For formal empirical work with actual data, it might be nice to estimate the covariance structure of the error terms; for the current purpose of showing that there is useful information in even a single observation of a portfolio problem, our assumptions should suffice.

Estimation in the simulations chooses the risk aversion parameters γ_0^{-1} and γ_1^{-1} that minimize the nonlinear sum across time of the squared errors in predicting at each of the 60 months to maturity the portfolio holding in the risky asset as a proportion of wealth, where the prediction uses the theoretical optimum for that level of wealth and time to maturity. More details are given in an Appendix. The simulation results are summarized in Tables 1 (less noise) and 2 (more noise). In both cases, useful parameter estimates are obtained from a single realization!

4 Continuous Consumption

While a problem with no consumption up until the horizon is a good representation for problems like investing for retirement, consumption withdrawal is an important feature of many portfolio problems. Fortunately, almost the same recoverability result holds as in Theorem 4, provided preferences are smooth enough, even without such a strong assumption about returns. The proof would be similar to the proof of Theorem 4; here is a sketch of how it would work. Assume that the agent's von Neumann-Morgenstern utility function is

$$(34)\quad \int_{t=0}^{T} e^{-\delta t} u(c_t) dt,$$

where $u(\cdot)$ is the felicity function and δ is the pure rate of time preference.[9] The budget constraint is $w_T \geq 0$ where

$$(35)\quad w_t = w_0 + \int_{s=0}^{t} w_s(r_s\, ds + \alpha_s'((\mu_s - r_s\mathbf{1})ds + \sigma_s dZ_s)) - \int_{s=0}^{t} c_s ds$$

and $(\forall t \in [0,T])w_t \geq 0$. This is equivalent to the following reduced problem.

Problem 5 *Choose an adapted consumption process c_t, $t \in [0,T]$, to maximize $E[\int_{t=0}^{T} e^{-\delta t} u(c_t)dt]$ subject to $E[\int_{t=0}^{T} \xi_t c_t] = w_0$.*

This problem will have existence of a solution under conditions parallel to those of Theorem 3.

In the consumption withdrawal problem, path independence works in a more subtle way, although fortunately we do not need path independence to recover preferences. The first-order condition for Problem 5 is

$$(36)\quad e^{-\delta t} u'(c_t) = \lambda \xi_t.$$

Up to a constant of proportionality λ that does not affect preferences, this expresses u' as a function of ξ_t and t, and we can infer all the derivatives of this function with respect to ξ_t along the sample path using the same approach as in the proof of Theorem 4. If we take u' very smooth (for example if $\log(u'(c))$ is analytic in $\log(c)$, then we can recover the entire utility function from the one sample path. Note that this argument requires less structure than we assumed before and that ξ need not be Gaussian. So long as a solution exists it suffices that ξ is locally random ($\tilde{\mu} \neq 0$ almost everywhere).

[9] The text assumes a finite horizon; the only material difference given an infinite horizon is that the existence of a solution is more subtle.

Of course, it is a strong assumption to take $\log(u'(c))$ analytic in c. However, the parametric restriction in actual estimation should suffice.

Another extension that is possible is to incorporate random income over time. Provided the income process has as finite number of state variables and is known (resp. estimable), preferences should be recoverable (resp. estimable). Of course, some care would have to write down precise sufficient conditions.

5 Conclusion

We have shown that an investor's entire portfolio strategy is revealed by the dynamic strategy on a single sample path. The main assumptions are that the agent has von Neumann-Morgenstern preferences over terminal wealth and faces complete markets with i.i.d. investment returns. This result helps us to understand better the nature of the assumptions we make, and our simulations suggest the result may have useful implications for empirical analysis.

Appendix

Technical Lemma

There follows a technical lemma that was used in the proof of Theorem 4. It provides a sufficient condition for the moment generating function (or equivalently the Fourier transform) to be entire (analytic on the whole complex plane).

Lemma 1 *Let m be a non-negative measure on $\Re$. Then the moment-generating function $M(s) \equiv \int_{x=-\infty}^{\infty} e^{sx} m(dx)$ is analytic on the entire complex plane if and only if $M(s)$ is finite for all $s \in \Re$.*

PROOF Obviously, an entire function is finite everywhere in $\Re$. Conversely, suppose M is finite on $\Re$. Then we want to show that M is entire. Since M is finite on $\Re$, it is clear that all moments of m are finite, so we may consider the approximating (polynomial) functions

$$(37) \quad M_N(s) \equiv \int_{x=-\infty}^{\infty} (\sum_{n=1}^{N} \frac{(sx)^j}{j!}) m(dx),$$

for $N = 0, 1, \ldots$ We want to show that the M_N's converge uniformly to M on compact subsets of the complex plane. Fix $C > 0$, and take any complex s, $|s| \leq C$, and positive integers $K < N$. Then

$$\begin{aligned} (38) \quad |M_N(s) - M_K(s)| &\leq \int_{x=-\infty}^{\infty} \sum_{j=K+1}^{N} \left(\frac{|sx|^j}{j!} \right) m(dx) \\ &\leq \int_{x=-\infty}^{\infty} \sum_{j=K+1}^{N} \left(\frac{C^j |x|^j}{j!} \right) m(dx) \\ &\leq \int_{x=-\infty}^{\infty} \sum_{j=K+1}^{\infty} \left(\frac{C^j |x|^j}{j!} \right) m(dx) \end{aligned}$$

Now as K increases, this decreases to 0 by monotone convergence, provided that for some K the integral is finite. But the integral is bounded by

$$\begin{aligned} \int_{x=-\infty}^{\infty} \sum_{j=0}^{\infty} \left(\frac{C^j |x|^j}{j!} \right) m(dx) &= \int_{x=-\infty}^{\infty} e^{C|x|} m(dx) \\ &\leq \int_{x=-\infty}^{\infty} (e^{Cx} + e^{-Cx}) m(dx) \\ &= M(C) + M(-C) \\ &< \infty \end{aligned}$$

Thus the Cauchy sequence $(M_N(s))_{N \geq 0}$ converges uniformly in $\{|s| \leq C\}$; the limit, $M(s)$, is therefore analytic, being the uniform limit of analytic functions. ∎

More Simulation Details

We shall firstly (a) describe the set up in a general context, then (b) explain the steps involved in the simulation, and finally (c) give more detail for the specific example considered here.

(a) Under our assumption that the spot rate, return and volatility of the stock are constant, the stock price process is given by

$$S_t = S_0 \exp[\sigma Z_t + (\mu - \sigma^2/2)t]$$

and the state-price density is given by

$$\xi_t = \exp[-rt - \tilde{\mu} Z_t - \tilde{\mu}^2 t/2]$$

with $\tilde{\mu} \equiv (\mu - r)/\sigma$. As is well known, an agent maximising $Eu(w_T)$ subject to initial wealth w_0 invests so as to achieve terminal wealth

$$w_T = I(\lambda \xi_T)$$

where the constant λ is chosen to satisfy the budget constraint:

$$w_0 = E[\xi_T I(\lambda \xi_T)].$$

If we abbreviate $\xi_T/\xi_t \equiv \xi_{tT}$ for $0 \leq t \leq T$, we have by the independent increments property of Z and () that

$$w_t = E_t[\xi_{tT} I(\lambda \xi_T)] = V(t, \lambda \xi_t),$$

where

$$V(t,x) \equiv E \xi_{tT} I(x \xi_{tT}).$$

It is a simple exercise in Itô calculus to verify that the optimally-investing agent holds wealth

$$-\frac{\tilde{\mu}}{\sigma} \lambda \xi_t V_\xi(t, \lambda \xi_t)$$

in the risky asset, and therefore holds a proportion

$$\alpha_t \equiv - \frac{\lambda \xi_t V_\xi(t, \lambda \xi_t)}{V(t, \lambda \xi_t)} \frac{\tilde{\mu}}{\sigma}$$

of his wealth in the risky asset.

(b) In our particular application, the utility (and hence the function V) depends on several parameters, θ; we write $V(t,x;\theta)$ to make this explicit when necessary. We now make some choice of θ and w_0, and simulate a sequence $(w_t,\alpha_t)_{t=0,\Delta,\ldots,N\Delta}$ of (wealth, portfolio) pairs in the following way. In what follows, the simulated random variables X_j, Y_j, and Z_j are all drawn independently.

(i) From $w_{j\Delta}$ we compute the multiplier λ_j by solving

$$w_{j\Delta} = V(j\Delta, \lambda_j \xi_{j\Delta}).$$

(ii) We then simulate $\alpha_{j\Delta}$ by

$$\alpha_{j\Delta} = -\lambda_j \xi_{j\Delta} \frac{V_\xi}{V}(j\Delta, \lambda_j \xi_{j\Delta}) \cdot \frac{\tilde{\mu}}{\sigma} \cdot X_j$$

where the X_j are IID log normal random variables of mean 1.

(iii) We simulate an increment $Z_{j+1} \sim N(0,\Delta)$ of the Wiener process, and a log normal random variable Y_{j+1} of mean 1, and use these to define $w_{j\Delta+\Delta}$ via

$$w_{j\Delta+\Delta} = V(j\Delta+\Delta, \lambda_j \xi_{j\Delta+\Delta}) Y_{j+1}$$

where $\xi_{j\Delta+\Delta} = \xi_{j\Delta} \exp(-(r+\frac{1}{2}\tilde{\mu}^2)\Delta - \tilde{\mu} Z_{j+1})$.

(iv) Repeat.

Once we have simulated the sequence $(w_t,\alpha_t)_{t=0,\Delta,\ldots,N\Delta}$, we estimate the parameters θ by minimising

$$\sum_{j=0}^{N-1} \left(\alpha_{j\Delta} + \lambda_j \xi_{j\Delta} \frac{V_\xi}{V}(j\Delta, \lambda_j \xi_{j\Delta}; \theta) \frac{\tilde{\mu}}{\sigma} \right)^2$$

where the λ_j are determined from the $w_{j\Delta}$ via

$$w_{j\Delta} = V(j\Delta, \lambda_j \xi_{j\Delta}; \theta),$$

that is, the model portfolio choice for simulated observation j is the the optimal continuation given wealth at that point in time (which is different from what the model would have predicted given w_0 and $\xi_{0,j\Delta}$).

(c) The preferences are characterised by

$$I(z) = \begin{cases} w_0 z^{-\gamma_1} & \text{for } z \leq 1 \\ w_0 z^{-\gamma_0} & \text{for } z \geq 1 \end{cases},$$

where $1/\gamma_0$ is the relative risk aversion coefficient for wealth less than w_0, and $1/\gamma_1$ is the relative risk aversion coefficient for wealth greater than w_0.

Step (i) of the simulation requires the inversion of the function $V(t,\cdot)$. We can here compute $V(t,\cdot)$ and its derivatives explicitly, although the inversion is done numerically. Since $\log \xi_{tT} \sim N(-\eta\tau, \tilde{\mu}^2\tau)$, where $\eta \equiv r + \tilde{\mu}^2/2$, $\tau \equiv T - t$, we have

$$\begin{aligned} V(t,x) &= \int e^{-(y+\eta\tau)^2/2\tilde{\mu}^2\tau} e^y I(xe^y) \frac{dy}{(2\pi\tilde{\mu}^2\tau)^{\frac{1}{2}}} \\ &= w_0 \int_{-\log x}^{\infty} \exp\left(-\frac{(y+\eta\tau)^2}{2\tilde{\mu}^2\tau} + y - \gamma_0 y\right) \frac{x^{-\gamma_0} dy}{(2\pi\tilde{\mu}^2\tau)^{\frac{1}{2}}} \\ &\quad + w_0 \int_{-\infty}^{-\log x} \exp\left(-\frac{(y+\eta\tau)^2}{2\tilde{\mu}^2\tau} + y - \gamma_1 y\right) \frac{x^{-\gamma_1} dy}{(2\pi\tilde{\mu}^2\tau)^{\frac{1}{2}}} \\ &= w_0 x^{-\gamma_0} \exp\left(-\frac{\tau(1-\gamma_0)(2\eta - \tilde{\mu}^2(1-\gamma_0))}{2}\right) \bar{\Phi}(z_0) \\ &\quad + w_0 x^{-\gamma_1} \exp\left(-\frac{\tau(1-\gamma_1)(2\eta - \tilde{\mu}^2(1-\gamma_1))}{2}\right) \Phi(z_1) \end{aligned}$$

where $\Phi(x) \equiv 1 - \bar{\Phi}(x) = P(Z_1 \leq x)$ is the standard normal distribution function, and where

$$z_i \equiv \frac{-\log x + \eta\tau - (1-\gamma_i)\tilde{\mu}^2\tau}{\tilde{\mu}\sqrt{\tau}}.$$

As with most numerical optimizations, some care is required to avoid infeasible and numerically stable regions. Our estimators (based on the netlib

program acm500 and transformed γ's) start from points far from the optimum and always converge under the criterion given the algorithm. Included in the algorthim is an ability to restart the algorithm (with different starting values) when it wanders into part of the parameter space where the numerical calculation of derivatives has an overflow, and this was done manually in fewer than 1% of the draws in the most extreme case.[10] Any remaining shortcomings of our optimization procedure (for example because of numerical rounding error or too forgiving a convergence criterion) probably give us less precision than the full optimum, and therefore our results potentially understate the potential precision of the estimates.

References

Arnold, Ludwig, 1974, *Stochastic Differential Equations: Theory and Applications*, Wiley: New York.

Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, "Option Pricing: A Simplified Approach" *Journal of Financial Economics*, 7, 229–263.

Cox, John C., and Hayne Leland, 1982, "On Dynamic Investment Strategies" *Proceedings, Seminar on the Analysis of Security Prices*, 26, Center for Research in Security Prices, University of Chicago.

Cox, J. C., and C. Huang, 1991, "A variational problem arising in financial economics" *Journal of Mathematical Economics*, 20, 465–87.

Dybvig, Philip H., 1982, "Recovering Preferences from Preferences over Nominal Gambles" *Journal of Economic Theory*, 28, 354–360.

Dybvig, Philip H. and Chi-fu Huang, 1988, "Nonnegative Wealth, Absence of Arbitrage, and Feasible Consumption Streams" *Review of Financial Studies*, 1, 377–401.

[10]The most extreme case has $\gamma_1 = 1.0$, implying, for large wealth, borrowing of 150% of wealth to buy a position of 250% of wealth in stock.

Dybvig, Philip H., and Heraklis Polemarchakis, 1981, "Recovering Cardinal Utility" *Review of Economic Studies*, 48, 159–166.

Dybvig, Philip H., Rogers, L. C. G., and Kerry Back, 1995, "Portfolio Turnpikes: Synthesis and Critique," working paper, Washington University in Saint Louis.

Green, Jerry, Lawrence Lau, and Heraklis Polemarchakis, 1978, "Identifiability of the von Neumann-Morgenstern Utility Function from Asset Demands" in Green, Jerry Green and José Scheinkman (eds.) *General Equilibrium, Growth, and Trade* Academic Press: New York.

Hodges, Stewart, 1991, "Ex-Post Performance Measurement of Dynamic Portfolio Strategies (or How to Tell Whether a Million Dollars Has Been Thrown Away)," Financial Options Research Centre Pre-print 91/23, University of Warwick.

Merton, Robert C., 1971, "Optimal Consumption and Portfolio Rules in a Continuous-time Model," *Journal of Economic Theory*, 3, 373–413.

Richter, Marcel K., 1987, "Revealed Preference Theory," an entry in Eatwell, John, Murray Milgate, and Peter Newman, eds., *The New Palgrave: a Dictionary of Economics*, MacMillan: London.

Rogers, L. C. G., and David WIlliams, 1987, *Diffusions, Markov Processes, and Martingales, Volume 2: Itô Calculus*, Wiley: Chichester.

Rudin, Walter, 1974, *Real and Complex Analysis*, McGraw-Hill: New York.

Vasicek, Oldrich, 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177–188.

Wang, Susheng, 1993a, "The Integrability Problem of Asset Prices," *Journal of Economic Theory*, 59, 199–213.

Wang, Susheng, 1993b, "The Local Recoverability of Risk Aversion and Intertemporal Substitution," *Journal of Economic Theory*, 59, 333–363.

	true	mean	stddev	min	5%	25%	50%	75%	95%	max
γ_0	0.2	0.22	0.028	0.06	0.18	0.20	0.22	0.24	0.26	0.31
γ_1	0.2	0.20	0.058	0.17	0.18	0.19	0.20	0.20	0.20	1.94
γ_0	0.2	0.20	0.029	0.15	0.18	0.19	0.20	0.20	0.21	2.05
γ_1	0.3	0.30	0.009	0.08	0.29	0.30	0.30	0.30	0.31	0.33
γ_0	0.2	0.20	0.017	0.17	0.19	0.20	0.20	0.21	0.22	0.67
γ_1	0.5	0.50	0.019	0.25	0.48	0.49	0.50	0.50	0.51	0.56
γ_0	0.3	0.29	0.038	0.06	0.21	0.28	0.30	0.31	0.33	0.39
γ_1	0.2	0.21	0.063	0.19	0.20	0.20	0.20	0.21	0.21	1.95
γ_0	0.3	0.30	0.012	0.07	0.28	0.29	0.30	0.30	0.31	0.39
γ_1	0.3	0.30	0.011	0.25	0.29	0.30	0.30	0.30	0.31	1.26
γ_0	0.3	0.30	0.010	0.23	0.29	0.30	0.30	0.30	0.31	0.48
γ_1	0.5	0.50	0.009	0.33	0.49	0.50	0.50	0.50	0.51	0.59
γ_0	0.5	0.49	0.062	0.06	0.46	0.49	0.50	0.51	0.53	0.58
γ_1	0.2	0.21	0.089	0.18	0.19	0.20	0.20	0.20	0.21	1.96
γ_0	0.5	0.50	0.017	0.08	0.48	0.49	0.50	0.51	0.52	0.59
γ_1	0.3	0.30	0.022	0.24	0.29	0.30	0.30	0.30	0.31	1.30
γ_0	0.5	0.50	0.011	0.44	0.48	0.49	0.50	0.51	0.52	0.62
γ_1	0.5	0.50	0.006	0.40	0.49	0.50	0.50	0.50	0.51	0.54

Table 1: These simulations illustrate estimation of two preference parameters from observations of a single portfolio in each of 60 months before maturity. The first parameter, γ_0, is the inverse of the relative risk aversion parameter below the initial wealth, and the second parameter, γ_1, is the inverse of the relative risk aversion parameter above the initial wealth. The interest rate is 6.5%, and the market portfolio has a mean return of 16.5% and a standard deviation of 20%, all on an annual basis. Each simulation is based on 10,000 draws, and assumes a noise with standard deviation 2% in the portfolio strategy.

	true	mean	stddev	min	5%	25%	50%	75%	95%	max
γ_0	0.2	0.22	0.031	0.07	0.17	0.20	0.22	0.24	0.27	0.34
γ_1	0.2	0.20	0.040	0.17	0.18	0.19	0.20	0.20	0.21	1.23
γ_0	0.2	0.20	0.019	0.12	0.17	0.19	0.20	0.21	0.23	0.33
γ_1	0.3	0.30	0.014	0.21	0.28	0.29	0.30	0.31	0.32	0.39
γ_0	0.2	0.21	0.029	0.14	0.18	0.20	0.20	0.21	0.23	2.04
γ_1	0.5	0.49	0.027	0.13	0.47	0.49	0.50	0.51	0.52	0.61
γ_0	0.3	0.29	0.042	0.06	0.21	0.27	0.29	0.31	0.34	0.50
γ_1	0.2	0.21	0.056	0.17	0.19	0.20	0.20	0.21	0.21	1.33
γ_0	0.3	0.30	0.029	0.06	0.26	0.29	0.30	0.31	0.34	2.00
γ_1	0.3	0.30	0.012	0.06	0.28	0.29	0.30	0.31	0.32	0.39
γ_0	0.3	0.30	0.022	0.20	0.27	0.29	0.30	0.31	0.33	0.74
γ_1	0.5	0.50	0.018	0.28	0.47	0.49	0.50	0.51	0.52	0.66
γ_0	0.5	0.49	0.076	0.06	0.39	0.47	0.50	0.52	0.56	1.98
γ_1	0.2	0.21	0.099	0.17	0.19	0.20	0.20	0.20	0.22	1.99
γ_0	0.5	0.50	0.033	0.05	0.45	0.48	0.50	0.52	0.55	0.65
γ_1	0.3	0.30	0.031	0.24	0.28	0.29	0.30	0.31	0.32	2.04
γ_0	0.5	0.50	0.028	0.35	0.46	0.48	0.50	0.52	0.55	0.75
γ_1	0.5	0.50	0.016	0.38	0.47	0.49	0.50	0.51	0.52	0.64

Table 2: These simulations are similar to the simulations in Table 1 except that there is more noise (standard deviation 5% instead of 2%) in the portfolio strategy (either in the observer's measurement of the market risk exposure or in the manager's compliance to the strategy). Useful estimation can still be obtained from a single draw.