Записки научных
семинаров ПОМИ
Том 312, 2004 г.

M. A. H. Dempster

SEQUENTIAL IMPORTANCE SAMPLING ALGORITHMS FOR DYNAMIC STOCHASTIC PROGRAMMING

ABSTRACT. This paper gives a comprehensive treatment of EVPI-based sequential importance sampling algorithms for dynamic (multistage) stochastic programming problems. Both theory and computational algorithms are discussed. Under general assumptions it is shown that both an expected value of perfect information (EVPI) process and the corresponding marginal EVPI process (the supremum norm of the conditional expectation of its generalized derivative) are nonanticipative nonnegative supermartingales. These processes are used as importance criteria in the class of sampling algorithms treated in the paper. When their values are negligible at a node of the current sample problem scenario tree, scenarios descending from the node are replaced by a single scenario at the next iteration. High values on the other hand lead to increasing the number of scenarios descending from the node. Both the small sample and asymptotic properties of the sample problem estimates arising from the algorithms are established and the former are evaluated numerically in the context of a financial planning problem. Finally, current and future research is described.

1. INTRODUCTION

This paper contains a comprehensive treatment of sequential importance sampling schemes for dynamic stochastic programmes based on *expected value of perfect information* (EVPI) processes as importance criteria. The term *dynamic* stochastic programme is used to refer to a multi-stage recourse formulation of the problem to be solved in which the data – and corresponding model coefficient – vector stochastic process is known to be evolving in real time, but is sampled discretely at the first instant of each period and *a fortiori* of each stage. This is the true situation in most practical applications of multistage stochastic programming, e.g. in finance, about which more will be said in §8. Since the pioneering work of Rockafellar, Wets and King [41, 26, 29] and Dantzig and Glynn [11], interest in – and the sophistication of – sampling schemes for stochastic optimization problems has grown rapidly. This development has been intertwined with a parallel and also rapidly growing literature on the sampling properties of dynamic simulation (see e.g. Rachev [38], Rubinstein and Shapiro [43] and Pflug [36]). In stochastic programming, however, sampling schemes have mainly been applied to 2-stage recourse problems [23–25] for which the concept of sampling "scenarios" is somewhat of a misnomer. From a statistician's viewpoint this paper concerns more-or-less classical sequential (importance) sampling; not of a random variable, vector, or even process, but rather of a full dynamic stochastic optimization problem in discrete time. The contribution of the paper in this context is an attempt to address the vexed question in applications, "How many scenarios are enough to make the sample problem sufficiently representative of the real-world situation?" This is a classical question of adequate sample size and is addressed here by sequential importance sampling of the vector data process underlying a specific problem. The sophisticated importance criteria proposed in the sequel utilize dynamic EVPI process information calculated from the optimal solutions of successive scenario-based sample problems. These sample problems are approximations to the full problem and they are generated by sampling the continuous time, continuous state vector data process – in applications usually via a corporate dynamic Monte Carlo simulator operated in conditional mode. Thus it is not surprising that all the basic considerations of classical statistical estimation theory apply to this extended situation in which dynamic stochastic optimization problems are being sampled.

The next two sections (§§2 and 3) of the paper set out respectively theory and computational techniques for both full and marginal EVPI processes – the latter being essentially directional derivatives of the former. Both are seen to be supermartingales – an essential property for their use in sequential importance sampling schemes for dynamic stochastic programming. The (full) EVPI process can be reasonably efficiently calculated by post-processing the optimal solution of the deterministic equivalent of the current sample problem in standard form after its computation by nested Benders decomposition, while marginal EVPI process values may be immediately extracted as dual variables from an appropriate split variable representation of the sample problem solved by a primal-dual interior point algorithm. Section 4 of the paper describes a class of sequential sampling algorithms utilizing these importance criteria. The basic idea of these procedures is that when the expected value of information – full or marginal – at a node of the scenario tree of the current sample problem is *high*, the number of scenarios descending from this node should be increased in the next sample problem. On the other hand, when EVPI is *negligible* the descendant scenarios should be decreased to one – a single scenario – which is resampled to increase robustness. The EVPI process value at a node measures the full – or marginal – impact of knowledge of future scenarios on the optimal value of the stochastic optimization problem remaining at the node through the decisions of this problem. Consequently, when the EVPI process value at a node is negligible it essentially means that the remaining decisions cannot effectively utilize scenario information to influence forward objective function values, so that a deterministic problem (single scenario) forward from the node will suffice. It follows that full or marginal EVPI processes are the natural importance sampling criteria in this context and we shall see in §8 that the variance reduction achieved through their use can exceed an order of magnitude.

As noted above, following classical statistics we may consider a sample problem to be a realization of a problem estimator. The three sections of the paper following §4 treat the usual properties expected from a statistical estimator – albeit involving much more difficult mathematics than in classical cases. In §5 it is shown that under reasonable assumptions, and in a technical sense to be made precise, these problem estimators are consistent, i.e. as the number of data process sample scenarios tends to infinity in a suitable way, the solution to the original stochastic optimization problem is recaptured. Section 6 shows that in *very* large samples the distribution of the optimal value of the problem estimator is approximately Gaussian. However, this central limit theory result is of *very* limited practical value in the context of a stochastic optimization problem, where the bias in the optimal value estimator is one-sided due to optimization and the small sample distribution of problem optimal values is highly skewed. Evaluating this bias is the topic of §7. All the concepts introduced in the paper are illustrated computationally in terms of a strategic financial investment planning problem in §8. Section 9 contains conclusions and directions for further research. In the remainder of this section the dynamic stochastic programming problem is formally introduced.

Consider now a vector stochastic *data process* $\boldsymbol{\omega} := \{\boldsymbol{\omega}_t : t = 1, \ldots, T\}$ whose realizations are *data paths* (vector sequences) in a (canonical) probability space $(\Omega, \mathcal{F}, \mu)$. We shall consider a vector *decision process* $\mathbf{x} := \{\mathbf{x}_t : t = 1, \ldots, T\}$ to be a measurable function $x : \omega \mapsto x(\omega)$ mapping a data trajectory to a (vector) decision trajectory. Thus $\mathbf{x}$ is a contingency plan in the sense that whatever data path ω is realized $x(\omega)$ gives the corresponding sequence of decisions. The temporal evolution of data and decisions up to the (infinite) *horizon* T in all the problems treated in this paper is $\boldsymbol{\omega}_1, \mathbf{x}_1, \boldsymbol{\omega}_2, \mathbf{x}_2, \boldsymbol{\omega}_3, \mathbf{x}_3, \ldots, \boldsymbol{\omega}_T, \mathbf{x}_T$ (where ω_1 and x_1 may be deterministic). (Here we consider measurable functions as both random vectors, functions or processes and elements of appropriate (Lebesgue) function spaces, distinguishing random entities from their analytic counterparts notationally by the use of bold symbols.)

We shall be concerned for theoretical purposes with the following very general *stochastic optimization* problem

$$\begin{aligned} \text{(SP)} \quad & \pi_0 := \sup \mathbb{E} f(\boldsymbol{\omega}, \mathbf{x}) \\ & s.t. \;\; \mathbf{x} \in P(\boldsymbol{\omega}) \;\; a.s., \quad g(\boldsymbol{\omega}, \mathbf{x}) \in Q \;\; a.s. \end{aligned}$$

Here f is a real valued measurable *objective functional* of the data and decision trajectories, P is a measurable multifunction (point-to-set map), g is a measurable vector trajectory valued *constraint function* of the data and decision trajectories, $\mathbb{E}$ denotes $(\mu-)$ expectation and the constraints are required to hold *almost surely* (*a.s.*), i.e. with $(\mu-)$ probability one. An analytic formulation of (*SP*) will be given in §2. Note first that (following the notational convention for random entities stated above) we may write (SP) as

$$\sup \mathbb{E}\, \mathbf{f}(\mathbf{x}) \qquad s.t. \;\; \mathbf{x} \in \mathbf{P} \quad a.s., \qquad \mathbf{g}(\mathbf{x}) \in Q \quad a.s.,$$

where now the multifunction P is interpreted as a *random set* $\mathbf{P}$. Although several alternative constraint formulations of equivalent generality are available, that of (SP) fits most naturally with the models in the sequel.

2. EVPI PROCESSES: THEORY

We first give a precise analytic formulation of (SP) and then turn to the definitions and properties of the full and marginal EVPI processes.

To this end, let us take the space of decision processes of (SP) to be the Banach space $L_p^n := X_{t=1}^T L_p(\Omega, \mathcal{F}, \mu; \mathbb{R}^n)$ equipped with the topology defined by the norm $||x|| := (\sum_{t=1}^T ||x_t||_p^p)^{1/p}$ for some p, $1 \leqslant p \leqslant \infty$. Since T is finite the given topology is equivalent to the product topology and we may consider L_p^n to be $L_p(\Omega, \mathcal{F}, \mu; \mathbb{R}^{nT})$.

Let $\mathbb{F} := \{\mathcal{F}_t\}_{t=1}^T$ be the filtration defined by the data process $\boldsymbol{\omega}$, where $\mathcal{F}_t := \sigma(\boldsymbol{\omega}^t)$ is the $\mu-$completed $\sigma-$field defined by the *history* $\boldsymbol{\omega}^t$ of the process, $t = 1, \ldots, T$, and (without loss of generality) we assume that the sequence $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \cdots \subset \mathcal{F}_t \subset \cdots \subset \mathcal{F}_T := \mathcal{F}$ defines the *filtration* $\mathbb{F} := \{\mathcal{F}_t\}$. We specify $\mathbf{x}_1$ (by means of $\boldsymbol{\omega}_1$) to be deterministic by setting $\mathcal{F}_1 := \{\emptyset, \Omega\}$. We say that the decision process $\mathbf{x}$ is *nonanticipative* (alternatively, that $\mathbf{x}$ is *adapted* to the data process $\boldsymbol{\omega}$ or filtration $\mathbb{F}$) if, and only if, the current decision $\mathbf{x}_t$ depends only on the data history $\boldsymbol{\omega}^t$ to date, viz. $\mathbf{x}_t := x_t(\boldsymbol{\omega}^t)$ *a.s.*, where x_t is measurable. Equivalently, $\mathbf{x}_t := x_t(\boldsymbol{\omega})$ is $\mathcal{F}_t-$ measurable or

$$\mathbf{x}_t = \mathbb{E}\{\mathbf{x}_t | \mathcal{F}_t\} \qquad a.s., \quad t = 1, \dots, T, \tag{1}$$

where $\mathbb{E}\{\cdot|\mathcal{F}_t\}$ denotes conditional expectation with respect to the $\sigma-$field $\mathcal{F}_t \in \mathbb{F}$, i.e. generated by the data available at time t. More analytically, in terms of a current decision $z \in L_p(\Omega, \mathcal{F}, \mu; \mathbb{R}^n)$ $(1 \leqslant p \leqslant \infty)$, Π_t : $z \mapsto \Pi_t z := \mathbb{E}\{z|\mathcal{F}_t\}$, $t = 1, \dots, T$, defines a closed projection. Hence (1) becomes a sequence of linear operator constraints

$$(I - \Pi_t)x_t = 0, \qquad t = 1, \dots, T, \tag{2}$$

involving complementary projections on L_p^n. We shall denote by $\mathcal{N}$ the closed linear subspace of nonanticipative decision processes in L_p^n and in the sequel require that all feasible decision processes for instances of (SP) lie in this subspace.

Before considering in detail the implications of this requirement, it will be useful to state the analytic version of (SP) as the *abstract optimization* problem

$$\begin{aligned} &\sup_{x \in \mathcal{N}} \int_\Omega f(\omega, x(\omega))\mu(d\omega) \\ &s.t.\ x(\omega) \in P(\omega), \qquad g(\omega, x(\omega)) \in Q \qquad a.s.[\mu], \end{aligned} \tag{3}$$

where, for each $x \in \mathcal{N} \subset L_p^n$, $f(\cdot, x(\cdot)) : \Omega \to \mathbb{R}$ is integrable and $g(\cdot, x(\cdot))$ measurable, $P : \Omega \rightrightarrows X_{t=1}^T \mathbb{R}^n$ is a measurable multifunction and $Q \subset X_{t=1}^T \mathbb{R}^n$. If all the problem functions in (3) are concave, the multifunction P is convex set valued and Q is a convex set, we say that (SP) is *convex*. If the supremum in (3) is achieved by some not necessarily unique feasible decision process $\mathbf{x}_0$ we term (SP) *solvable*.

Conformal with (1) the constraints of (SP) may be written in a form which exhibits their temporal structure as

$$x_t(\boldsymbol{\omega}) \in P_t(\boldsymbol{\omega}^t), \qquad g_t(\boldsymbol{\omega}^t, x^t(\boldsymbol{\omega})) \in Q_t \quad a.s., \quad t = 1, \dots, T. \tag{4}$$

Assuming the objective f to be *separable*, i.e.

$$f(\boldsymbol{\omega}, x(\boldsymbol{\omega})) := \sum_{t=1}^{T} f_t(\boldsymbol{\omega}^t, x^t(\boldsymbol{\omega})),$$

we may write (SP) in the more compact *dynamic programming* representation in terms of a set of path dependent problems

$$\begin{aligned}
\pi_t(\boldsymbol{\omega}^t) &:= \pi_t(\boldsymbol{\omega}^t, x^{t-1}(\boldsymbol{\omega})) \\
&:= \sup_{x_t(\boldsymbol{\omega})} \left[f_t(\boldsymbol{\omega}^t, x^{t-1}(\boldsymbol{\omega}), x_t(\boldsymbol{\omega})) + \mathbb{E}\{\pi_{t+1}(\boldsymbol{\omega}^{t+1}, x^t(\boldsymbol{\omega})) | \mathcal{F}_t\} \right] \\
&s.t. \quad x_t(\boldsymbol{\omega}) \in P_t(\boldsymbol{\omega}^t) \qquad a.s., \\
&\qquad\;\; g_t(\boldsymbol{\omega}^t, x^t(\boldsymbol{\omega})) \in Q_t \qquad a.s., \\
&\qquad\;\; x_t(\boldsymbol{\omega}) = \mathbb{E}\{x_t(\boldsymbol{\omega}) | \mathcal{F}_t\} \quad a.s., \quad t = 1, \ldots, T,
\end{aligned} \tag{5}$$

where π_{t+1} expresses the optimal expected value for the remaining optimization problem of the same form for the stages from $t+1$ to T. Clearly $\pi_1(\omega^1) = \pi_1(\omega_1) = \pi_0$ and at the horizon $\pi_{T+1}(\boldsymbol{\omega}^T, \mathbf{x}^T) = \pi_{T+1}(\boldsymbol{\omega}, \mathbf{x}) :\equiv 0$. In (5) the dependence of the decision vectors $\mathbf{x}_t$ on the filtrations $\mathcal{F}_t$ is expressed *implicitly* through the third set of constraints. Thus at the t^{th} stage nonanticipative decisions $x_t(\boldsymbol{\omega}^t)$ are chosen from the *feasible set* $X_t(\boldsymbol{\omega}^t) := X_t(\boldsymbol{\omega}^t, \mathbf{x}^{t-1})$ defined by the constraints of (5).

Notice that without loss of generality any instance of (SP) may be considered to have a separable objective with $f_t :\equiv 0$ for all $t = 1, \ldots, T-1$ and

$$f_T(\boldsymbol{\omega}^T, x^T(\boldsymbol{\omega})) := f(\boldsymbol{\omega}, x(\boldsymbol{\omega})).$$

Following control theory, such a version of the problem is said to be in *Mayer form.*

EVPI process

An *expected value of perfect information (EVPI) process* is defined by

$$\eta_t(\boldsymbol{\omega}^t) := \phi_t(\boldsymbol{\omega}^t) - \pi_t(\boldsymbol{\omega}^t), \quad t = 1, \ldots, T, \tag{6}$$

where $\phi_t(\boldsymbol{\omega}^t)$ denotes the value function of the *distribution* problem associated with the relaxation of the nonanticipativity conditions (1) of (SP) to the case of *perfect foresight* to the horizon T relative to a specific decision process $\mathbf{x}$ optimal for (SP). Thus

$$\begin{aligned}
\phi_t(\boldsymbol{\omega}^t) &:= \phi_t(\boldsymbol{\omega}^t, \mathbf{x}^{t-1}) \\
:= \mathbb{E}\{ \sup_{\mathbf{z}_t \in \mathbf{Z}_t} &\left[f_t(\boldsymbol{\omega}^t, \mathbf{x}^{t-1}, \mathbf{z}_t) + \phi_{t+1}(\boldsymbol{\omega}^{t+1}, \mathbf{x}^{t-1}, \mathbf{z}_t) \right] | \mathcal{F}_t \},
\end{aligned} \tag{7}$$

where $\mathbf{Z}_t$ represents the feasible set of the first two constraints of (5) given $\mathbf{x}$, i.e.

$$\begin{aligned} Z_t(\boldsymbol{\omega}) &:= Z_t(\boldsymbol{\omega}, \mathbf{x}^{t-1}) \\ &= \{z_t(\boldsymbol{\omega}) : z_t(\boldsymbol{\omega}) \in P_t(\boldsymbol{\omega}^t),\ g_t(\boldsymbol{\omega}^t, z^{t-1}(\boldsymbol{\omega}), z_t(\boldsymbol{\omega})) \in Q \quad a.s.\}. \end{aligned} \tag{8}$$

If the decision process $\mathbf{x}$ optimal for (SP) is unique, $\boldsymbol{\eta}$ is uniquely defined. Otherwise, the EVPI process $\boldsymbol{\eta}$ may depend upon the specific optimal decision process $\mathbf{x}$ with respect to which it is defined, since in period t the decision opportunities afforded by perfect foresight may in general depend upon the history $\mathbf{x}^t$ of the decision process to date as well as upon that $\boldsymbol{\omega}^t$ of the data process $\boldsymbol{\omega}$.

Notice that by construction an EVPI process $\boldsymbol{\eta}$ is nonanticipative. Based on the behaviour of such a process we can both assess the level of *stochasticity* of the problem (SP) and define a sampling procedure for the selection of a sample set of *relevant representative* data paths in a sequential procedure. From the definition of $\boldsymbol{\eta}$ we have by construction that at the horizon $\boldsymbol{\eta}_T :\equiv 0$.

Theorem 1. *For solvable* (SP) *and a specific optimal decision process* $\mathbf{x}$ *the corresponding unique EVPI process* $\boldsymbol{\eta}$ *is a nonanticipative nonnegative supermartingale , i.e.*

$$\boldsymbol{\eta}_t \geqslant \mathbb{E}\{\boldsymbol{\eta}_s | \mathcal{F}_t\} \geqslant 0, \qquad 1 \leqslant t \leqslant s, \quad s, t \in \mathcal{T} := \{1, \ldots, T\}. \tag{9}$$

Proof. Let $\mathbf{x}$ be an optimal decision process for (5) and let $\mathbf{z}(\mathbf{x})$ be a corresponding perfect foresight process optimal almost surely along data paths for (7) defining uniquely the EVPI process $\boldsymbol{\eta}(\mathbf{x})$ through (6).

The proof is by induction on the finite horizon T.

Consider first the case $T{:=}2$, where we may simplify the problem notationally to

$$\sup_{x \in X} \left[f(x) + \mathbb{E}\pi(\boldsymbol{\omega}, x) \right].$$

Define

$$x \in \arg\max_{x \in X} \left[f(x) + \mathbb{E}\pi(\boldsymbol{\omega}, x) \right]$$

and for each $\omega \in \Omega$

$$z(\omega) \in \arg\max_{x \in X} \left[f(x) + \pi(\omega, x) \right].$$

Thus for all $\omega \in \Omega$

$$f(z(\omega)) + \pi(\omega, z(\omega)) \geqslant f(x) + \pi(\omega, x)$$

and taking expectations

$$\varphi := \mathbb{E} f(z(\boldsymbol{\omega})) + \mathbb{E}\pi(\boldsymbol{\omega}, z(\boldsymbol{\omega})) \geqslant f(x) + \mathbb{E}\pi(\boldsymbol{\omega}, x) := \pi,$$

i.e. $\eta := \varphi - \pi \geqslant 0$. Hence $\eta := \eta_1 \geqslant \mathbb{E}\{\boldsymbol{\eta}_2 | \mathcal{F}_1\} := \mathbb{E}\boldsymbol{\eta}_2 := 0$.

Now consider an arbitrary finite horizon T and assume the result holds for horizon $T-1$.

Since we may reformulate the solvable problem (SP), assumed separable without loss of generality, as

$$\begin{aligned} & \max\ \mathbb{E}\boldsymbol{\xi}_T \\ s.t.\qquad & \sum_{t=1}^{T} f_t(\boldsymbol{\omega}^t, x^t(\boldsymbol{\omega})) = \boldsymbol{\xi}_T \end{aligned} \tag{10}$$

and the constraints of (5), we may without loss of generality assume that $\mathbf{f}_t \equiv 0$ for all $t = 1, \ldots, T-1$, if this is not already the case.

Then, since again the first stage deterministic decision constraint $\mathbf{x}_1 = x_1$ a.s. in (5) has been removed to give (7), and making use of the induction hypothesis,

$$\begin{aligned} \eta_1 &:= \varphi_1 - \pi_1 \\ &= \mathbb{E}\{\varphi_2(\boldsymbol{\omega}_2, \mathbf{z}_1, \mathbf{z}_2) - \pi_2(\boldsymbol{\omega}_2, x_1, \mathbf{x}_2)\} && (f_1 :\equiv 0) \\ &\geqslant \mathbb{E}\{\varphi_2(\boldsymbol{\omega}_2, x_1, \mathbf{z}_2) - \pi_2(\boldsymbol{\omega}_2, x_1, \mathbf{x}_2)\} && (\text{as for } T := 2) \\ &:= \mathbb{E}\boldsymbol{\eta}_2 := \mathbb{E}\{\boldsymbol{\eta}_2 | \mathcal{F}_1\} && (\text{by definition}) \\ &\geqslant \mathbb{E}\{\mathbb{E}\{\boldsymbol{\eta}_t | \mathcal{F}_2\} | \mathcal{F}_1\} = \mathbb{E}\{\boldsymbol{\eta}_t | \mathcal{F}_1\} \geqslant 0 && (\text{by induction}) \end{aligned}$$

for all $t = 3, \ldots, T$ as required.

Thus we have established the supermartingale property (9) by backwards induction and, for example, a similar argument applies conditional on each realization $\boldsymbol{\omega}_2 = \omega_2$. $\square$

Marginal EVPI process

We next introduce the multiplier process $\boldsymbol{\rho}' := (\boldsymbol{\rho}_1', \boldsymbol{\rho}_2', \ldots, \boldsymbol{\rho}_T')$ on the nonanticipativity constraints of (SP) and show that it is also a supermartingale. (Here prime denotes vector transpose.) First recall from above that for $1 \leqslant p \leqslant \infty,\ 1 \leqslant t < \infty$

$$\mathcal{N} := \{x \in L_p^n : (I - \Pi_t)x_t = 0,\ t = 1, \ldots, T\} \tag{11}$$

is a closed linear subspace. Moreover, we shall see that the process $\boldsymbol{\rho}'$ in L_p^n $(1 \leqslant q := p/(p-1) \leqslant \infty)$ can always be chosen to be nonanticipative, i.e.

$$\boldsymbol{\rho}_t' = \{\boldsymbol{\rho}_t' | \mathcal{F}_t\} \quad a.s. \qquad t = 1, 2, \ldots, T. \tag{12}$$

Since $\boldsymbol{\rho}_t'$ represents the Lagrange multiplier (row) vector on the nonanticipative constraint $(I - \Pi_t)x_t = 0$, $t = 1, \ldots, T$ – i.e. the marginal value at the optimum of a perturbation of this constraint with respect to information not available at time t (more precisely, with respect to some σ-measurable function z in L_p^n) – and is nonanticipative (i.e. known at time t), we may interpret it as the *marginal expected value of perfect information (EVPI)* at time t. We shall refer to $\boldsymbol{\rho}'$ as the *marginal EVPI process.*

More formally, we consider the *perturbed* abstract problem

$$\begin{array}{ll} (\mathrm{SP}_z) & \pi(z) := \sup\ \mathbb{E} f(\mathbf{x}) \\ & s.t.\ \ \mathbf{x} \in \mathbf{P}, \qquad g(\mathbf{x}) \in Q, \qquad \mathbf{x} \in \mathcal{N} + \mathbf{z} \qquad a.s., \end{array}$$

where the *anticipative perturbation* $\mathbf{z}$ in L_p^{nT} is such that for $t = 1, \ldots, T$, z_t is $\mathcal{F}$-measurable, i.e. dependent on values $\boldsymbol{\omega}_s$ of the data process at some *future* time(s) $s > t \in \mathcal{T}$.

The condition of the following proposition holds for all special cases of (SP) under suitable regularity conditions (see Dempster [17, p. 30]).

Theorem 2 (Dempster [16, 17]). *Let* $\mathbf{x}$ *be an optimal decision process for* (SP) *over* L_p^n $(1 \leqslant p \leqslant \infty,\ 1 \leqslant t \leqslant \infty)$ *and suppose that* (SP_z) *is solvable and proper (i.e. has a finite value) for all anticipative perturbations* $\mathbf{z}$ *in a neighbourhood of 0 in* L_p^n. *Then under the above regularity conditions the marginal EVPI process* $\boldsymbol{\rho}' := (\boldsymbol{\rho}_1', \boldsymbol{\rho}_2', \ldots, \boldsymbol{\rho}_t')$ *in* L_q^n $(1 \leqslant q \leqslant \infty)$ *corresponding to* $\mathbf{x}$ *is a nonanticipative nonnegative supermartingale, i.e.*

$$\boldsymbol{\rho}_t' \geqslant \{\boldsymbol{\rho}_s' | \mathcal{F}_t\} \geqslant 0' \quad a.s., \qquad 1 \leqslant t \leqslant s,\ \ s, t \in \mathcal{T}.$$

Proof. See Dempster [17, pp. 30–31], where it is shown that ρ' is the adaptation of $\nabla\pi(0)$, the (unique) Fréchet derivative of π evaluated at $z = 0$, to the filtration $\mathbb{F}$. $\square$

We may interpret the supermartingale property of the EVPI processes $\boldsymbol{\eta}$ and $\boldsymbol{\rho}'$ as representing the fact that anticipative decision processes for (SP) have the obvious property that anticipative information is potentially the more valuable in terms of the objective the sooner it is known.

It can be shown (Dempster [17]) that $\boldsymbol{\rho}'$ and the dual process $\mathbf{p}'$ (Rockafellar and Wets [40]) – which is the multiplier process on the explicit subspace constraint $\mathbf{x} \in \mathcal{N}$ – lie in the complementary subspaces $R(\Pi_t')$ and $R(\Pi_t)^0 = R(I' - \Pi_t')$ respectively of $L_q(\mathcal{T} \times \Omega, \mathcal{P}(\mathcal{T}) \times \mathcal{F}, \# \times \mu; \mathbb{R}^n)$. The process $\mathbf{p}'$ is a martingale difference involving a martingale $\mathbf{m}'$ with values in $\mathbb{R}^n$, not necessarily centered at the origin, and hence $\mathbb{E}\mathbf{p}' = 0'$.

The properties of the marginal EVPI process $\boldsymbol{\rho}'$ stated in Theorem 2 can be expected to hold for optimal control of *continuous time* processes such as diffusions. In this situation, partial results have been obtained by Back and Pliska [1] regarding the dual process $\mathbf{p}'$ (which they term the *shadow price of information* process A) and further results have been given by Davis, Dempster and Elliott [14]. In the continuous time case one can capture the *full* EVPI process by suitable time integrals of the *marginal* EVPI process, as conjectured by Dempster [16].

Before turning to a discussion of numerical methods, it should be pointed out that techniques for calculating (or bounding) the norm $||\boldsymbol{\rho}'||$ of the marginal EVPI process $\boldsymbol{\rho}'$ can be expected to prove of great value in computation. Indeed, small values of $||\boldsymbol{\rho}'||$ would justify ignoring random variables rather than computing – at great effort – full stochastic solutions to (remaining) dynamic stochastic problems with a low EVPI. In such cases the corresponding deterministic *expected value* problem (in which random variables are replaced by their expectations) would provide – at greatly reduced complexity – a sufficiently accurate approximation to the optimal expected value and decisions. These are ideas fundamental to the sampling schemes of §4.

3. EVPI PROCESSES: COMPUTATION

To make numerical solution of (SP) tractable it is commonly assumed that the number of possible data process paths ω – termed *scenarios* – is finite and that they can be conceptually arranged in a tree structure – called an *event* or *scenario tree* – whose nodes correspond to the optimal decision problems in specific time periods (*cf.* Raiffa [39]). For the purposes of this paper such scenario based instances of (SP) will constitute *sample* problems which must be solved numerically in a suitable form and the corresponding EVPI process values computed from their solutions.

Although with the new generation of fully convex interior point computer codes such as LOQO it will soon be possible to solve numerically large scale fully convex scenario based instances of (SP), in the remainder of this paper we specialize to the linearly constrained convex scenario based problem given by

$$\max_{x_1} f_1(x^1) + \mathbb{E}_{\boldsymbol{\omega}^2}\left\{\max_{\mathbf{x}_2} f_2(\boldsymbol{\omega}^2,\mathbf{x}^2) + \mathbb{E}_{\boldsymbol{\omega}^3|\boldsymbol{\omega}^2}\left\{\max_{\mathbf{x}_3} f_3(\boldsymbol{\omega}^3,\mathbf{x}^3) + \right.\right.$$
$$\left.\left. \dots + \mathbb{E}_{\boldsymbol{\omega}^T|\boldsymbol{\omega}^{T-1}}\{\max_{\mathbf{x}_T} f_T(\boldsymbol{\omega}^T,\mathbf{x}^T)\}\dots\right\}\right\}$$

$$\begin{array}{llllll}
\text{s.t. } A_{11}x_1 & & & & = b_1 & \\
A_{21}(\boldsymbol{\omega}^2)x_1 & + A_{22}(\boldsymbol{\omega}^2)\mathbf{x}_2 & & & = b_2(\boldsymbol{\omega}^2) & a.s. \\
A_{31}(\boldsymbol{\omega}^3)x_1 & + A_{32}(\boldsymbol{\omega}^3)\mathbf{x}_2 & + A_{33}(\boldsymbol{\omega}^3)\mathbf{x}_3 & & = b_3(\boldsymbol{\omega}^3) & a.s. \\
\vdots & \vdots & \ddots & & \vdots & \\
A_{T1}(\boldsymbol{\omega}^T)x_1 & + A_{T2}(\boldsymbol{\omega}^T)\mathbf{x}_2 + \cdots & & + A_{TT}(\boldsymbol{\omega}^T)\mathbf{x}_T & = b_T(\boldsymbol{\omega}^T) & a.s.
\end{array}$$
$$\begin{array}{ccc}
l_1 \le & x_1 & \le u_1 \\
l_t(\boldsymbol{\omega}^t) \le & \mathbf{x}_t & \le u_t(\boldsymbol{\omega}^t), \quad t = 2,\dots,T.
\end{array} \tag{13}$$

Here the notation, e.g. $\mathbb{E}_{\boldsymbol{\omega}^3|\boldsymbol{\omega}^2}(\cdot) := \mathbb{E}\{\cdot|\boldsymbol{\omega}^2\}$, is used for simplicity, $\mathbf{P}_t := [\mathbf{l}_t, \mathbf{u}_t]$ and $Q := \{0\}$.

Problem (13) is often referred to as the convex *multistage recourse model* with linear constraints.

For scenario based versions of (SP) we may replace consideration of the σ-fields $\mathcal{F}_t$ of the filtration $\mathbb{F}$ generated by the data process $\boldsymbol{\omega}$ by *partitions* of the finite path space

$$\Omega = \{\omega(1),\dots,\omega(k),\dots,\omega(K)\}$$

of scenarios. (See Billingsley [4] for a discussion of generating σ-fields from vector interval partitions in the continuous state case.) While $\mathcal{F}_1 := \{\Omega, \emptyset\}$, or the corresponding partition

$$\mathcal{A}_1 = \{\{\omega(1),\dots,\omega(k),\dots,\omega(K)\}\},$$

represents no information beyond the identity of the path space, the field $\mathcal{F}_T$, or the corresponding partition

$$\mathcal{A}_T = \{\{\omega(1)\},\dots,\{\omega(k)\},\dots,\{\omega(K)\}\},$$

represents full information.

A partition $\mathcal{A}_t$ of the sample path space Ω is *finer* than the partition $\mathcal{A}_{t-1}$ ($\mathcal{A}_{t-1}$ is *coarser* than $\mathcal{A}_t$) iff for all $A \in \mathcal{A}_t$ there is an $A' \in \mathcal{A}_{t-1}$ such that $A \subseteq A'$. A finer (coarser) partition represents more (less) information available to the decision maker.

An *information structure* is a sequence of partitions $\{\mathcal{A}_t\}$ such that $\mathcal{A}_t$ is finer than $\mathcal{A}_{t-1}$ for $t = 2,\dots,T$.

At time t, it is possible to distinguish sets $A \in \mathcal{A}_t$ each of which has as elements those scenarios continuing the same data path history ω^t. However, it is not possible at time t to distinguish amongst the scenarios in a given A. By considering a graph with all sets in $\mathcal{A}_t$, $t = 2, \ldots, T$, as nodes and each set in $\mathcal{A}_{t-1}$ joined to its corresponding subsets in $\mathcal{A}_t$ with arcs, the information structure can be represented as a *tree* structure τ. In terms of scenarios, each node of τ represents a *nodal optimization subproblem* and each *branch* a random event $\boldsymbol{\omega}^t$, conditioned on the realized past ω^{t-1}.

Thus there is a bijective relation between each data path history ω^t and each set in $\mathcal{A}_t$. Let $A_t(\omega^t)$ denote the set of scenarios with the same data history path ω^t, *viz.*

$$A_t(\omega^t) \;:=\; \left\{\omega \in \Omega \mid \omega \succ \omega^t\right\}. \tag{14}$$

Here $\omega \succ \omega^t$ denotes the fact that ω continues ω^t.

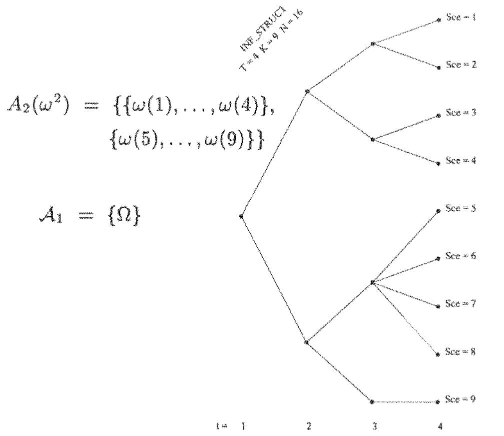

Fig. 1. Graphical representation of an information structure in terms of scenarios.

Figure 1 illustrates the tree structure of a 4 period, 9 scenario, 16 node problem, with associated partition at time $t = 3$ given by

$$\mathcal{A}_3 := \{\{\omega(1), \omega(2)\},\ \{\omega(3), \omega(4)\},\ \{\omega(5), \omega(6), \omega(7), \omega(8)\},\ \{\omega(9)\}\}.$$

Before setting the decisions x_3, the decision maker is at the most uncertain amongst four scenarios (depending on the data path history ω^3).

If the realisation of the data process is $\omega(9)$, this will be already known by the decision maker by $t = 3$.

Thus the *tree* τ may also be considered to be a structured collection of non-disjoint history paths ω^t termed a *scenario tree*. The hierarchical structure of nodes (paths) uses the same descriptors as for family trees (ancestor, successor, parent, child, sibling, cousin) and in §§4 through 7 it will be useful to use the notation $\omega \in \tau$ for scenarios $\omega := \omega^T$ in τ (considered as a sequence of nodes).

Let $A_t(\omega^t) \in \mathcal{A}_t$ as in (14). The *branching degree* at node ω^t is equal to the number of sets of the partition $\mathcal{A}_{t+1}$ that are subsets of $A_t(\omega^t)$. In Figure 1, let $A_2(\omega^2)$ be $\{\omega(1), \omega(2), \omega(3), \omega(4)\}$, and $\mathcal{A}_3$ as above. Notice that $\{\omega(1), \omega(2)\}$ and $\{\omega(3), \omega(4)\}$ are the two subsets of $A_2(\omega^2)$ corresponding to the branching degree at node ω^2.

A *sub-tree* τ_1 of τ is a tree in which all its paths are also paths of τ. In a *scenario-sub-tree*, all scenarios of τ_1 are also scenarios of τ.

A *uniform tree* of depth T is a tree where all nodes at periods $t = 1, \dots, T-1$ have the same branching degree. A *balanced tree* is a tree where all nodes at the same period have the same branching degree.

At the root node in Figure 1 the decision maker faces an instance of (SP) and if the path realised in the following stages is denoted by ω^3 and previous decisions were x^2, at the third period the remaining problem to solve will be given by

$$\pi_3(\omega^3, x^2) \;=\; \min_{x_3 \in X_3(\omega^3)} \; \left\{ f_3(\omega^3, x^3) \;+\; \mathbb{E}\left\{ \pi_4(\boldsymbol{\omega}^4, x^3) \mid \boldsymbol{\omega}^3 = \omega^3 \right\} \right\}. \tag{15}$$

To aid in processing the generation from a simulator of vector data paths consistent with the form of a specific scenario tree, the bijection between data path histories ω^t – i.e. nodes of the tree – and sets of the information partition may be defined numerically in terms of a *scenario partition matrix* L (Lane and Hutchinson [30]). This $|\Omega| \times T$ matrix has general *scenario number* entry

$$l(\omega, t) = l(\omega', t) \qquad \forall \omega' \in A_t \in \mathcal{A}_t \tag{16}$$

for $\omega = 1, \dots, |\Omega| := K$, $t = 1, \dots, T$. More useful for the conditional generation of scenarios from a simulator however is the $|\Omega| \times T$ *nodal partition* matrix M whose general entry $m(\omega, t)$ is the *node number* assigned to ω^t in the tree.

Figure 2 provides an example of the matrix specifications associated with an arbitrary tree structure.

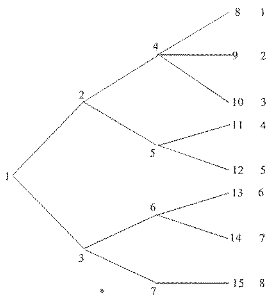

$$L := \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 2 \\ 1 & 1 & 1 & 3 \\ 1 & 1 & 4 & 4 \\ 1 & 6 & 5 & 5 \\ 1 & 6 & 6 & 6 \\ 1 & 6 & 6 & 7 \\ 1 & 6 & 8 & 8 \end{bmatrix} \qquad M := \begin{bmatrix} 1 & 2 & 4 & 8 \\ 1 & 2 & 4 & 9 \\ 1 & 2 & 4 & 10 \\ 1 & 2 & 5 & 11 \\ 1 & 2 & 5 & 12 \\ 1 & 3 & 6 & 13 \\ 1 & 3 & 6 & 14 \\ 1 & 3 & 7 & 15 \end{bmatrix}$$

Fig. 2. Example of a scenario tree schema with corresponding scenario and nodal partition matrices.

Either matrix identifies uniquely the tree structure for the associated stochastic program. The nodal partition matrix may be used by a data generator in order to derive the states of the data process in conditional mode, and by the model generator *STOCHGEN* [8, 9] for the definition of the corresponding *SMPS* input files [5] necessary for the numerical solution of the problem. See Chen *et al.* [7] for more details.

A scenario tree can also be seen as a measure of the size of the corresponding scenario based stochastic programming problem, and therefore of the complexity and difficulty of solving the problem. To see this note that a problem equivalent to (13) can be constructed by introducing for each possible realisation $\omega^t \in \Omega^t$, $t = 1, \dots, T$, vectors of decision variables $x_t(\omega^t)$, i.e. a unique set of decision vectors for each node of the tree (Dantzig and Madansky [13]). This leads to the *standard* form of the *deterministic equivalent* of the linearly constrained stochastic convex problem (13) as

$$\max \Bigg[f(x_1) + \sum_{\omega^2 \in \Omega^2} p(\omega^2) f_2(\omega^2, x_2(\omega^2)) + \sum_{\omega^3 \in \Omega^3} p(\omega^3) f_3(\omega^3, x_3(\omega^3)) + \cdots + \sum_{\omega^T \in \Omega^T} p(\omega^T) f_T(\omega^T, x_T(\omega^T)) \Bigg]$$

s.t.

$$\begin{array}{llllll}
A_{11}x_1 & & & & = b_1 & \\
A_{21}(\omega^2)x_1 & +A_{22}(\omega^2)x_2(\omega^2) & & & = b_2(\omega^2) & \forall \omega^2 \in \Omega^2 \\
A_{31}(\omega^3)x_1 & +A_{32}(\omega^3)x_2(\omega^2) & +A_{33}(\omega^3)x_3(\omega^3) & & = b_3(\omega^3) & \forall \omega^3 \in \Omega^3 \\
\vdots & \vdots & \ddots & & \vdots & \\
A_{T1}(\omega^T)x_1 & +A_{T2}(\omega^T)x_2(\omega^2) & + \cdots & + A_{TT}(\omega^T)x_T(\omega^T) & = b_T(\omega^T) & \forall \omega^T \in \Omega^T
\end{array}$$

$$\begin{aligned} l_1 \leq x_1 \quad &\leq u_1 \\ l_t(\omega^t) \leq x_t(\omega^t) &\leq u_t(\omega^t), \quad \forall\, \omega^t \in \Omega^t,\ t = 2, \ldots, T. \end{aligned} \tag{17}$$

This deterministic equivalent formulation allows the use of standard optimization algorithms and software to solve problem (13).

Furthermore, in most applications problem (17) involves a very sparse lower triangular matrix. Figure 3 shows the bitmap of the constraint matrix of the deterministic equivalent of an instance of the FRC_20 portfolio problem treated in §8 of 20 periods and 32 scenarios. The bitmap shows the non-zero entries of the matrix.

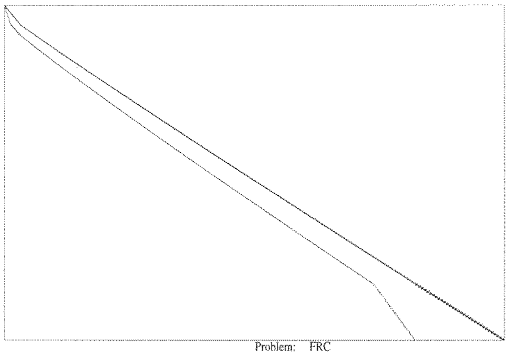

Fig. 3. Bitmap of the FRC_20 deterministic equivalent constraint matrix based on 32 scenarios.

EVPI process

To calculate the EVPI process variables at the nodes of the tree for a given problem (refer Fig. 2), it is necessary to solve, for each node ω^t, $t = 1, \ldots, T$, the perfect foresight (deterministic) dynamic optimization problems corresponding to each descendant scenario $\omega \succ \omega^t$ to obtain the optimal value $\phi_t(\omega)$ and then to compute

$$\phi_t(\omega^t) := \sum_{\omega \succ \omega^t} p(\omega)\phi_t(\omega), \tag{18}$$

the discrete version of (7). When (in the linear objective case) the optimal value π_0 of the problem has been computed by a nested Benders decomposition code such as MSLiP-OSL (Thompson [48]), this is easily accomplished by backwards recursion on partial scenarios with warm starting from the optimal bases of the nodal subproblems. Subtracting from these values $\phi_t(\omega^t)$ the corresponding optimal values $\pi_t(\omega^t)$ of the remaining problems – computed in solving (17) – yields the EVPI process values $\eta(\omega^t)$ at the nodes ω^t of the tree. (This calculation is performed by all variants of the MSLiP code family, see Gassmann [21].)

Marginal EVPI process

To calculate the values of the marginal EVPI process $\rho(\omega^t)$ on the other hand, we consider the *deterministic equivalent* of (13) in *conditional expectation split variable* form. This problem is formed by adding to the so-called *split variable* problem (Dempster [16]),

$$\max \sum_{\omega\in\Omega} p(\omega)\left[f(x_1(\omega))+f_2(\omega^2,x_2(\omega))+f_3(\omega^3,x_3(\omega))+\cdots+f_T(\omega^T,x_T(\omega))\right]$$

$$\begin{array}{llll}
\text{s.t.} & & & \\
A_{11}x_1(\omega) & & & =b_1 \\
A_{21}(\omega^2)x_1(\omega) + A_{22}(\omega^2)x_2(\omega) & & & =b_2(\omega^2)\ \forall\omega^2\in\Omega^2 \\
A_{31}(\omega^3)x_1(\omega) + A_{32}(\omega^3)x_2(\omega) + A_{33}(\omega^3)x_3(\omega) & & & =b_3(\omega^3)\ \forall\omega^3\in\Omega^3 \\
\vdots \qquad\qquad \vdots & & \ddots & \vdots \\
A_{T1}(\omega^T)x_1(\omega) + A_{T2}(\omega^T)x_2(\omega) + \cdots & & + A_{TT}(\omega^T)x_T(\omega) & =b_T(\omega^T)\ \forall\omega^T\in\Omega^T \\
l_1 \le x_1(\omega) \le u_1 & & & \\
l_t(\omega^t) \le x_t(\omega) \le u_t(\omega^t), & \forall\,\omega^t\in\Omega^t, & t=2,\dots,T, & \forall\,\omega\in\Omega:=\Omega^T,
\end{array} \tag{19}$$

the nonanticipativity constraints in *conditional expectation* form given by

$$x_t(\omega) = x_t(\omega^t), \quad \omega \succ \omega^t \in \Omega^t, \quad t = 1,\dots,T. \tag{20}$$

Note that (20) implies

$$x_t(\omega^t) = \sum_{\omega\succ\omega^t} p(\omega)x_t(\omega) / \sum_{\omega\succ\omega^t} p(\omega), \quad \omega^t\in\Omega^t, \quad t=1,\dots,T. \tag{21}$$

Here $p(\omega)$ denotes the path probability of scenario ω.

The number of variables in the problem and the number of constraints in (20) can be reduced by noting that *nodal* probabilities are given by

$$p(\omega^t) = \sum_{\omega\succ\omega^t} p(\omega) = \sum_{\omega\in A(\omega^t)} p(\omega), \qquad \omega^t\in\Omega^t, \quad t=1,\dots,T,$$

so that (21) can be written in terms of conditional probabilities of the branches (nodes) ω^{t+1} descending from ω^t (in an obvious notation) as

$$\begin{aligned}
x_t(\omega^t) &= \sum_{\omega^{t+1}\succ\omega^t} p(\omega^{t+1})x_t(\omega^{t+1}) / \sum_{\omega^{t+1}\succ\omega^t} p(\omega^{t+1}) \\
&:= \sum_{\omega^{t+1}\succ\omega^t} p(\omega^{t+1}|\omega^t)x_t(\omega^{t+1})
\end{aligned} \tag{22}$$

and (20) is replaced by

$$x_t(\omega^{t+1}) = x_t(\omega^t), \qquad \omega^{t+1} \succ \omega^t \in \Omega^t, \quad t = 1, \dots, T. \qquad (23)$$

The constraints (23), after substituting from (22), are exactly the discrete analogue of a realization of the continuous nonanticipativity constraint in conditional expectation form (1) – hence the terminology.

At an optimum of the problem (19), (20), the nonnegative Lagrange multiplier row vectors $\rho_t'(\omega) \in \mathbb{R}_n'$ are *anticipative marginal EVPI* values whose *nonanticipative* counterpart values $\rho_t'(\omega^t)$ are given by $\sum_{\omega \succ \omega^t} p(\omega|\omega^t)\rho_t'(\omega)$. Similarly, relative to (22) and (23),

$$\rho_t'(\omega^t) = \sum_{\omega^{t+1} \succ \omega^t} p(\omega^{t+1}|\omega^t)\rho_t'(\omega^{t+1}), \qquad \omega^t \in \Omega^t, \quad t = 1, \dots, T. \quad (24)$$

Note that in either case the appropriate multipliers are nonnegative since, as in the proof of Theorem 1, relaxing (20) or (23) can only increase the problem value.

By standard Lagrangean duality theory, the nonnegative row vector $\rho_t'(\omega^t)$ is an element of the (convex) subgradient of $\pi(\omega^t)$. Choosing a norm, e.g. $||\rho'||_\infty := \max_j \rho_j$, for $\rho_t'(\omega^t)$ gives a single *marginal EVPI* number – comparable to the full EVPI value $\eta(\omega^t)$ – at node ω^t. In light of Theorem 2, the resulting values clearly define a discrete supermartingale $||\rho'||_\infty$.

Although the split variable representation (19) and (20) – or its reduction corresponding to (23) – is considerably larger than the standard deterministic equivalent representation (17) of the scenario based (SP) problem (13), a good preprocessor will essentially reduce the larger problem to the smaller by eliminating variables using the nonanticipativity constraints. If a primal-dual interior point algorithm is used to solve the preprocessed problem, upon solution and postprocessing the anticipative marginal EVPI values may be immediately extracted from the solution and the nonanticipative values $\rho_t'(\omega^t)$ and the scalar $||\rho_t'(\omega^t)||_\infty$, $\omega^t \in \Omega^t$, $t = 1, \dots, T$, calculated.

As a consequence of the supermartingale property of $\boldsymbol{\eta}$ and $||\boldsymbol{\rho}'||_\infty$ we have the following important result whose proof is immediate.

Theorem 3. *For the scenario based (SP) problem,* $\eta_t(\omega^t) = 0$ *implies* $\eta_{t+s}(\omega) = 0$ *for all* $\omega \succ \omega^t$ *and* $s = 1, \dots, T-t$. *A similar result holds when* $||\rho'(\omega^t)||_\infty = 0$. $\square$

As a consequence, all optimal decisions $x_{t+s}(\omega)$ for such scenarios can be taken to agree and a single scenario $\omega \succ \omega^t$ and a deterministic (perfect foresight) decision process $x(\omega) \succ x(\omega^t)$ serve to represent the *remaining* decision problem at ω^t, $\omega^t \in \Omega^t$, $t = 1, \ldots, T$.

4. Sequential importance sampling with EVPI criteria

In this section a parametric class of sequential importance sampling estimation procedures for the dynamic stochastic programming problem (SP) will be briefly described. As discussed in the previous sections, at each sampling iteration $\nu = 1, 2, \ldots$ a sample data path process $\hat{\omega}_\nu$ is generated by a conditional simulator in the form of a scenario tree with branching structure specified by the current nodal partition matrix M_ν. In turn $\hat{\omega}_\nu$ specifies – through coefficient generation and a modelling language – the sample problem $SP(\hat{\omega}_\nu)$ in an appropriate format for optimization (in the linear case as SMPS or MPS solver input files). Accordingly, either a full EVPI sample process $\hat{\mu}_\nu$ or its marginal equivalent $||\hat{\rho}'_\nu||$ is extracted from the solution of $SP(\hat{\omega}_\nu)$. We now specify rules by which the current sample scenario tree $\hat{\omega}_\nu$ may be updated to the new tree $\hat{\omega}_{\nu+1}$ (initially through updating M_ν to $M_{\nu+1}$). It was mentioned in the introduction that at nodes where the appropriate importance sampling criterion value is high we wish to expand the number of descendent scenarios, while – in light of Theorem 3 – at those at which it is negligible we wish to collapse the current scenarios to a single scenario which should be resampled for robustness.

Before specifying the detailed importance sampling algorithm iteration rules, we note that the following overall considerations will be critical to the variance reduction performance of the resulting algorithm.

1. "Negligible" must be specified for the importance sampling criterion used.

The "zero" of Theorem 3 must be defined as a suitable ϵ for computational purposes. For $\boldsymbol{\eta}$ we define

$$\epsilon := \alpha \hat{\eta}_{1\nu}, \tag{25}$$

i.e. the product of the root node EVPI of the current sample problem and a parameter $\alpha := 0.1, 0.05, 0.01, 0.005, \ldots$ which may be tuned to specific problems. A value of $\eta_\nu(\omega^t_\nu) < \epsilon$ is deemed negligible, i.e. " zero", for the purposes of applying Theorem 3. A similar procedure is used for $||\boldsymbol{\rho}'||$. Note that for some purposes it is useful to specify equivalently the ϵ tolerance for $\boldsymbol{\eta}$ in terms of the *relative EVPI* or *stochasticity* of the current sample problem – i.e. the ratio of its root node EVPI to its optimal value

– given by $\epsilon := \alpha\,[\hat{\eta}_{1\nu}/\pi(\hat{\omega}_\nu)]$ and to divide all nodal EVPI values by the problem value $\pi(\hat{\omega}_\nu)$ before comparison.

2. The initial sample scenario tree $\hat{\omega}_1$ branching structure must be specified by the nodal partition matrix M_1 as input.

In general, a reasonable number of scenarios should descend from the root node in the initial balanced tree (*cf.* §3) and at least binary branching should occur for the next several periods. A useful shorthand notation for specifying balanced trees is to give the $T-1$ uniform branching numbers between successive stages. For example, in the experiments with 12 stage financial problems reported in §8, we have specified the initial tree of 32 scenarios in the form $8 \cdot 2 \cdot 2 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot 1$ or simply $8 \cdot 2^2 \cdot 1^8$.

3. The sample scenario tree enrichment by adding further descendent scenarios or collapsing existing scenarios at a node should be considered for all current nodes at a time stage.

Successive stages are considered at successive iterations at least until the period before the horizon, at which point processing returns to stage 2, where the initial sample tree nodal stage processing also begins. This has the effect of creating richer branching at earlier stages of successive sample trees unless the EVPI based importance criterion utilized dictates otherwise.

4. A termination rule for sequential sampling must be specified.

For both importance sampling criteria to date we have used a stopping criterion of the form e.g. $\hat{\eta}_{1\nu} = \hat{\eta}_{1\nu+1}$ to within 3 decimal places, together with a maximum iteration count ν_{max}.

Sample tree modification from iteration to iteration may be effected by variations on the following themes.

5. Differential scenario branching expansion of high importance criterion value nodes.

The difference between the highest criterion value in the sample tree and ϵ is divided into equal intervals and nodes with criterion values falling in each interval are branched the same number of times. For example, we might divide this interval into 3 and branch 2, 3 and 4 times for nodal values falling in successively higher thirds.

6. "Zero" importance criterion nodes with binary branching have their two descendent scenarios resampled.

This ensures enhanced robustness of the sample problem at early iterations.

7. Multibranched "zero" nodes have all descendent scenarios collapsed into a single resampled scenario.

This removes portions of the current sample tree which are irrelevant from the point-of-view of usefully exploiting sample data process information with the remaining decisions. Thereby sample problem dimensions are reduced and only "relevant" scenarios remain in the subsequent sample problem.

8. When the current stage expansion produces only "zero" nodes the stage considered is stepped backwards in time to reconsider a stage considered previously.

This procedure of stepping back in the tree – particularly for the marginal EVPI importance sampling criterion – keeps the procedure running and results in sample problems with highly unbalanced trees (see Fig. 5). Currently *step back* is set at 2 stages.

Figure 4 shows an initial sample scenario tree of 32 scenarios with the marginal EVPI computed for the FRC problem of §8. Notice the supermartingale property of the computed $||\rho(\omega^t)||$ values. Figure 5 illustrates the highly unbalanced nature of the sample scenario tree after several iterations. Note that the path probabilities of all scenarios generated by sampling from a simulator will be *equal* – as for an empirical probability distribution.

5. Consistency of EVPI-based problem estimators

In the next three sections of this paper we study the sampling properties of the (SP) problem estimator $SP(\hat{\boldsymbol{\omega}}_N)$ produced at termination by the EVPI-based sequential importance sampling algorithms introduced in the previous section. Following classical statistics, what is needed is a version of *central limit theory* – extending functional central limit theory – appropriate to finite scenario based sampling of a continuous state dynamic stochastic optimization problem. Following early relevant work on the *distribution problem* (see e.g. Bereanu [2, 3] and Prékopa [37] the literature on this topic has mainly been confined to two stage problems. An exception is the approach we shall take to demonstrating the *consistency* – i.e. asymptotic accuracy – of $SP(\hat{\boldsymbol{\omega}}_N)$ as $N \to \infty$ in the sense that both optimal decisions – and *a fortiori* (optimal) values – of the scenario based finite dimensional sample problem estimator converge almost surely to those of the continuous infinite dimensional "true" problem (SP) in the strictly convex linearly constrained case. This is of course a *strong law of large numbers* for dynamic stochastic optimization problems which

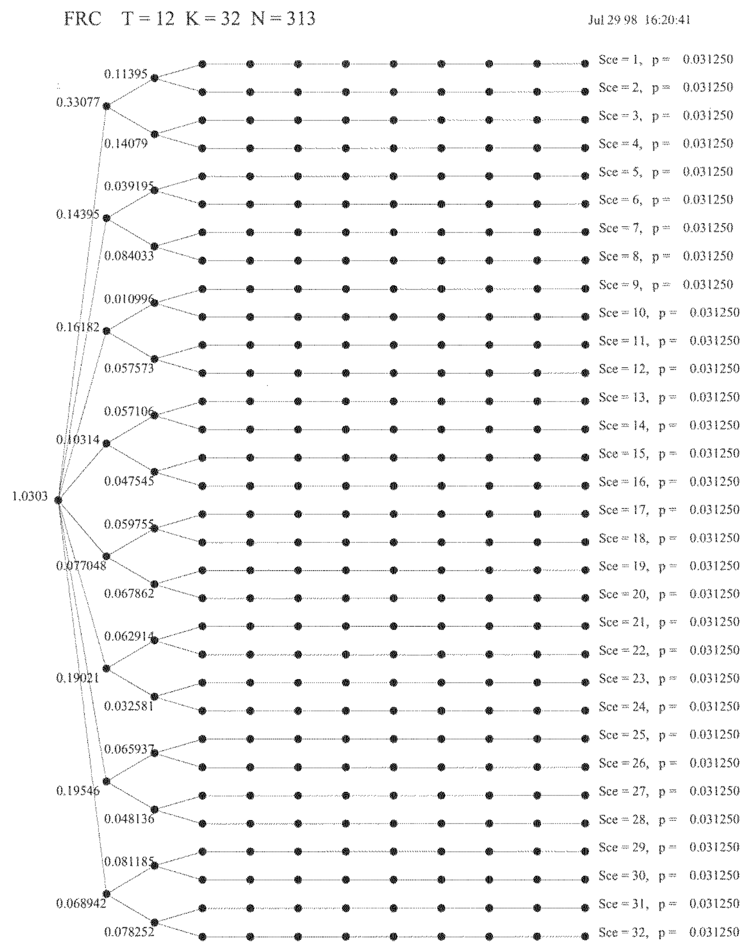

Fig. 4. Initial 32 scenario tree for the FRC problem with computed marginal EVPI values marked on its nodes.

is based here on earlier work of Daniel [10] and Olsen [35]. In the next section we treat the question of the asymptotic distribution of sample

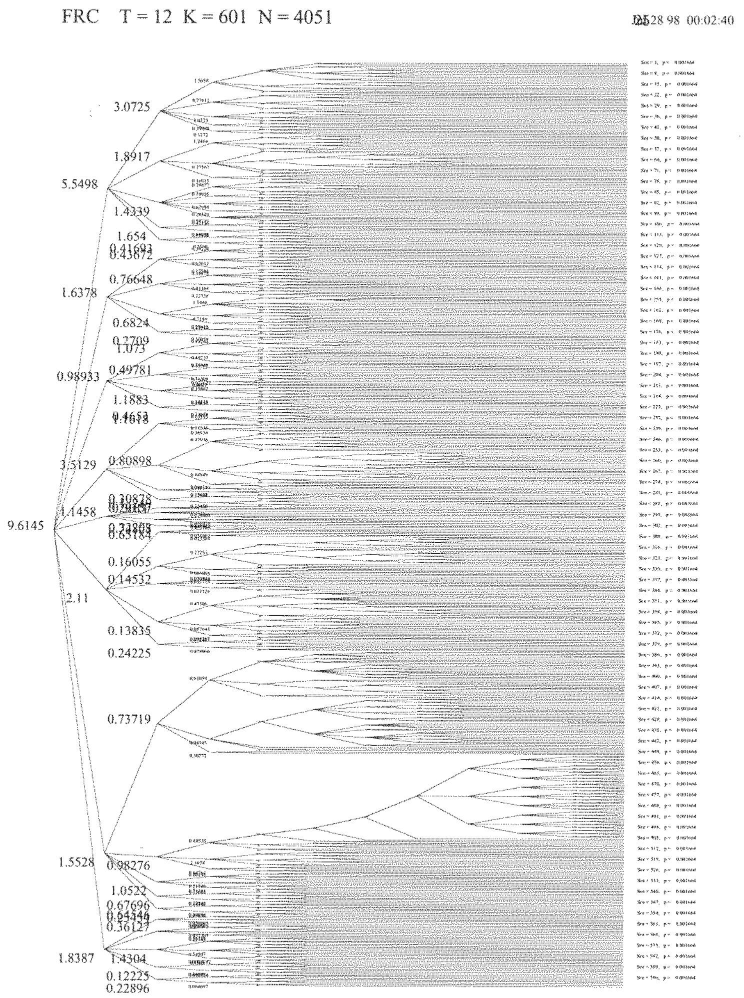

Fig. 5. FRC problem sample scenario tree after 6 iterations with the marginal EVPI importance sampling criterion.

problem estimators, which requires a corresponding *central limit theorem*. These hard-to-obtain results – while reassuring – are unfortunately of little or no practical relevance. In §7 we therefore set out what is known theoretically about the small sample behaviour of the problem estimators and in §8 give empirical results on a specific 12 stage financial problem.

For the purposes of studying the sample properties of scenario based problem estimators $SP(\hat{\boldsymbol{\omega}}_N)$ we shall express the linearly constrained strictly convex version of (SP) given by (13) variously as

$$\pi_0 := \sup_{x \in \mathbf{X}} \mathbb{E}\ f(\mathbf{x}, \boldsymbol{\omega}) \tag{26}$$

$$:= \sup_{x_1 \in X_1} \left[f(x_1) + \mathbb{E}\ \pi_2(\boldsymbol{\omega}_2, x_1) \right] \tag{27}$$

as appropriate, where $X_1 := \{x \in \mathbb{R}_n : A_{11}x = b_1\}$. Correspondingly, the finite dimensional problem estimators at termination of the sampling algorithm after N iterations will be denoted by

$$\hat{\boldsymbol{\pi}}_N := \sup_{\hat{\mathbf{x}}_N \in \mathbf{X}_N} \mathbb{E}\ f(\hat{\mathbf{x}}_N, \hat{\boldsymbol{\omega}}_N) \tag{28}$$

$$:= \sup_{\hat{x}_{1N} \in X_1} \left[f(\hat{x}_{1N}) + \mathbb{E}\ \pi_2(\hat{\boldsymbol{\omega}}_{2N}, \hat{x}_{1N}) \right]. \tag{29}$$

We wish to establish in this section that $\hat{\mathbf{x}}_N \to \mathbf{x}_0$ – and in particular $\hat{x}_{1N} \to x_{10}$ – almost surely so that *a fortiori* $\hat{\boldsymbol{\pi}}_N \to \pi_0$ as $N \to \infty$ in an appropriate sense.

To this end we recall the ingenious construction due to Olsen [35] of (essentially) a generalized scenario tree appropriate to the study of the asymptotic behaviour of a *multistage* scenario based stochastic programming sample problem as the branching number of the sample problem scenario tree tends to infinity *at every stage*. Clearly this latter condition is necessary if the *empirical measures* on $(\Omega, \mathcal{F})$

$$\mu_N(\cdot) := \sum_{\omega \in \hat{\omega}_N} \delta_\omega(\cdot)/|\hat{\omega}_N| \tag{30}$$

corresponding to the realized sample scenario trees $\hat{\omega}_N$ are to converge weakly to the underlying probability measure μ as $N \to \infty$. Here $|\hat{\omega}_N|$ denotes the number of distinct scenarios in $\hat{\omega}_N$.

A *sampling tree* τ_N for the data process $\boldsymbol{\omega}$ of (SP) is defined as $\tau_N := [\hat{\omega}_N; B_N]$, where $\hat{\omega}_N$ is an ordinary (finite) scenario tree with σ-field $\mathcal{F}_N := \mathcal{P}\{\omega : \omega \in \hat{\omega}_N\}$ given by all subsets of scenarios in $\hat{\omega}_N$ and $B_N : \hat{\omega}_N \rightrightarrows \Omega$ is a measurable multifunction such that:

a) $B_N(\omega)$ is an element of the Borel σ-field $\mathcal{B}(\Omega)$ of data paths in $\mathbb{R}^{dT}$ for all $\omega \in \hat{\omega}_N$ and suitable d,

b) $B_N\,[\hat{\omega}_N] = \Omega$,
c) $\omega, \omega' \in \hat{\omega}_N$ and $\omega \neq \omega'$ imply $B_N(\omega) \cap B_N(\omega') = \emptyset$,
d) $\xi \in B_N(\omega), \xi' \in B_N(\omega')$ and $\xi = \xi'$ imply $\omega = \omega'$.

Condition a) states that measurability of the multifunction B_N is with respect to $(\mathcal{F}_N, \mathcal{B}(\Omega))$. Conditions b) and c) state that the images of B_N are Borel *partitions*, while d) is a *generalized branching* condition which states that the element of the partition corresponding to a given scenario $\omega \in \hat{\omega}_N$ is unique. If Ω is finite, B_N is an identity map and τ_N is equivalent to an ordinary scenario tree. The general situation for scalar datapaths $(r := 1)$, $T := 5$ and a 4 scenario tree $\hat{\omega}$ is illustrated schematically in Figure 6. Here $\hat{\omega} := \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $B := \{B_1, B_2, B_3, B_4\}$ and the sets B_i are depicted as four factor Cartesian products of half open intervals.

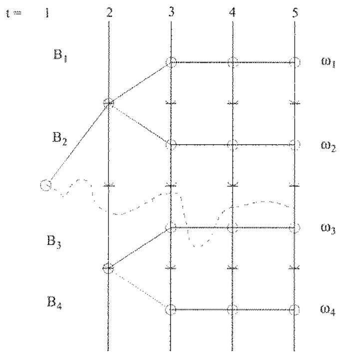

Fig. 6. Illustration of an Olsen sampling tree.

An *empirical probability measure* μ_N is defined on the finite scenario tree $\hat{\omega}_N$ of the sampling tree τ_N by

$$\mu_N(\omega) := \mu(B_N(\omega)) \qquad \forall \omega \in \hat{\omega}_N. \tag{31}$$

For (conditional) Monte Carlo path generation from a simulator of course $\mu_N(\omega) = 1/|\hat{\omega}|$ as in (30).

Theorem 4 (Consistency of problem estimators)**.** *Consider the linearly constrained strictly convex problem (SP) given by* (17) *and suppose* $\boldsymbol{\omega}$ *has* $p > 1$ *moments finite and* $\Omega \subset \mathbb{R}^{rT}$ *is bounded. Let* $\mathbf{x}_N$ *solve the* *sample problem $SP(\hat{\boldsymbol{\omega}}_N)$. If the branch numbers at every node of the sample scenario trees $\hat{\boldsymbol{\omega}}_N$ tend to infinity almost surely as $N \to \infty$, then $\hat{\boldsymbol{\pi}}_N \to \pi_0$ and $\mathbf{x}_N \to \mathbf{x}_0$ almost surely for $\mathbf{x}_0$ the unique solution of (SP).*

Proof. We only sketch the proof whose details will appear elsewhere. (The linear case was treated in Corvera Poiré [8].)

The proof consists of the following steps.

First we observe that the stated convergence condition on the structure of the sample scenario tree $\hat{\boldsymbol{\omega}}_N$ implies that

$$\max\{\mathrm{diam} B_N(\omega) : \omega \in \hat{\omega}_N\} \to 0 \quad \text{as} \quad N \to \infty$$

for the corresponding Olsen sampling trees τ_N.

Next, following Olsen [35], we use the sampling tree τ_N to construct a finite dimensional discrete approximation of (SP) in $L_p(\Omega, \mathcal{F}, \mu; \mathbb{R}^{nT})$ for $p > 1$.

Results of Daniel [10] then imply that $\hat{\mathbf{x}}_N \to \mathbf{x}_0$ in p^{th} norm as $N \to \infty$ (see [35]).

Finally, using Skorohod's theorem (Billingsley [4, p. 337–340]) we may conclude that $\hat{\mathbf{x}}_N \to \mathbf{x}_0$ almost surely.

Regularity of the appropriate abstract version (3) of (SP) and uniqueness of its solution in the strictly convex linearly constrained case are standard matters (see e.g. Dempster [16]). □

Two other approaches to establishing a similar result are possible. The most straightforward is to notice that the "deterministic" equivalent of the sample problem estimator $SP(\hat{\boldsymbol{\omega}}_N)$ has variables and constraints tending to infinity with $|\hat{\boldsymbol{\omega}}_N|$ as $N \to \infty$ and to generalize the results of Prékopa [37] appropriately. The other, perhaps more elegant, approach is to study the sample problem estimator as a random *perturbation* of (SP) in terms of the empirical sample measure $\boldsymbol{\mu}_N$ (see Fiedler and Römisch [20], Dentcheva and Römisch [19]).

6. Asymptotic distribution of EVPI-based problem estimators

Turning to the question of the asymptotic distributions associated with the scenario based sample problem estimator $SP(\hat{\boldsymbol{\omega}}_N)$ we know from the important work of King and Rockafellar [26, 27, 29] – and as originally observed by Bereanu [2, 3] – that it is too much to hope in general that in nontrivial – i.e. constrained – cases that the distribution of $\hat{\mathbf{x}}_N$ is asymptotically Gaussian. However, this does not preclude the asymptotic sampling distribution of $\hat{\boldsymbol{\pi}}_N$ from being Gaussian in the *complete recourse* case appropriate to financial modelling. We shall make use of results of Shapiro [45–47] to prove the following.

Theorem 5 (Asymptotic normality of the sample problem value for complete recourse). *Under the conditions of Theorem 4 for the complete recourse case,*

$$\sqrt{N}(\hat{\boldsymbol{\pi}}_N - \pi_0) \rightarrow \mathcal{N}(0, \sigma^2(x_{10})) \tag{32}$$

weakly, i.e. in distribution.

Proof. The proof is by induction on the finite horizon T.

First note that since $\Omega \subset \mathbb{R}^{rT}$ is bounded it may be demonstrated by a backwards recursive argument that both $\pi_2(\cdot, x_1)$ and $\pi_{2N}(\cdot, \hat{x}_{1N})$ of (27) and (29) respectively are Lipschitz continuous – with bounded Lipschitz constant – uniformly in the original problem and sample problem decisions over $X_1 \subset \mathbb{R}_n$ providing that the terms in the separable objective function have a similar property (Römisch [42]).

Next observe that as a result of Monte Carlo independent identically distributed scenario sampling between the first two stages

$$\hat{\boldsymbol{\pi}}_N := \sup_{\hat{x}_{1N} \in X_1} \{f(\hat{x}_{1N}) + \frac{1}{K(N)} \sum_{k=1}^{K(N)} \pi_{2N}(\hat{\boldsymbol{\omega}}_2^k, \hat{x}_{1N})\}, \tag{33}$$

where $K(N) = O(|\hat{\boldsymbol{\Omega}}_N|^{1/T})$ almost surely.

Now consider $T := 2$. The conditions of Shapiro [47, p. 438], are met and we may conclude the required result.

To demonstrate the induction step, suppose the results holds for horizon $T-1$. From (33) we observe that $\hat{\boldsymbol{\pi}}_N$ is an equiprobable mixture of asymptotically Gaussian variables and hence is asymptotically Gaussian with finite variance $\sigma^2(x_{10})$ as required. □

The difficulty in the present context with this asymptotic result is that N – or in terms of scenarios $|\hat{\omega}_N|$ – is likely to have to be *very* large before sampling error in $SP(\hat{\omega}_N)$ is suppressed.

For the sequential importance sampling algorithms of §4 this is of course an empirical matter – to which we turn in §§7 and 8. First however note that we may use Theorem 5 to construct in the usual way an *approximate* $100\alpha\%$ confidence interval for π_0 as

$$\left[\hat{\pi}_N - z_{\alpha/2}\hat{\sigma}_N/\sqrt{N}, \hat{\pi}_N + z_{\alpha/2}\hat{\sigma}_N/\sqrt{N}\right], \tag{34}$$

where $\hat{\pi}_N$ is the realized sample problem optimum at sequential algorithm termination after N iterations, $\hat{\sigma}_N$ is the sample standard deviation of successive sample problems prior to termination and $z_{\alpha/2}$ is the standard normal z-score at level $\alpha/2$.

7. Evaluating small sample bias

It has long been a folk theorem that due to optimization the sample problem value estimator $\hat{\boldsymbol{\pi}}_N$ has unidirectional bias. Recently, Norkin *et al.* [34] and Mak *et al.* [31] have come up with a simple demonstration of this fact in the two stage case which is easily extended to multistage (SP).

Theorem 6. *The finite sample problem estimator $\hat{\boldsymbol{\pi}}_N$ of the (SP) problem value π_0 has one-sided bias, i.e. $\mathbb{E}\hat{\boldsymbol{\pi}}_N \geqslant \pi_0$.*

Proof. We again make use of induction on the horizon T. Using (33) the result is easily seen to be valid for $T := 2$, due to the independent identically distributed (i.i.d.) sampling of "scenarios" in this case.

Similarly, for the general induction step for horizon T,

$$\begin{aligned}
\mathbb{E}\hat{\boldsymbol{\pi}}_N &= \mathbb{E} \sup_{\hat{x}_{1N} \in X_1} \left[f_1(\hat{x}_{1N}) + \frac{1}{K(N)} \sum_{k=1}^{K(N)} \pi_{2N}(\hat{\boldsymbol{\omega}}_2^k, \hat{x}_{1N}) \right] \\
&\geqslant \sup_{\hat{x}_{1N} \in X_1} \left[f_1(\hat{x}_{1N}) + \frac{K(N)}{K(N)} \mathbb{E}\pi_{2N}(\hat{\boldsymbol{\omega}}_2^k, \hat{x}_{1N}) \right] \\
&\geqslant \sup_{\hat{x}_{1N} \in X_1} \left[f_1(\hat{x}_{1N}) + \mathbb{E}\pi_{2N}(\hat{\boldsymbol{\omega}}_2, \hat{x}_{1N}) \right] \\
&= \pi_0.
\end{aligned}$$

Here the second inequality makes use of the induction hypothesis and i.i.d. sampling of second stage data path values. □

Theorem 6 prompts the search for an easily computed lower bound for the value π_0 of (SP) in which a large enough scenario sample can be generated to effectively suppress sampling error. With such a lower bound, maximum small sample bias $\hat{\boldsymbol{\pi}}_N - \pi_0$ can be conservatively estimated. It is easily seen that any variant of (SP) in which the decision process $\mathbf{x}$ is further restricted provides such a lower bound. In the next section, which treats a strategic financial asset allocation problem, we report bias estimates based on restricting *all* elements of the decision process to be constant – i.e. restricting decisions to the Markowitz [32] buy-and-hold portfolio strategy.

8. A financial example

The strategic asset allocation *(CALM-)FRC model* applies scenario based dynamic stochastic linear – and quadratic – programming to long term financial planning. It incorporates uncertainty in unknown rates of return, and produces investment decision strategies contingent on these rates. The stochastic data process paths were generated by a simulation code described in Hicks Pedrón [22] and based on a vector diffusion process model of Brennan, Schwartz and Lagnado [6] calibrated to recent US market history (see also Chen *et al.* [7]).

The investor is assumed to invest in three assets – an instantaneously riskless security, *cash*, a long term (consol) *bond* and an equity *portfolio*. The model assumes that the investor has a long term horizon (20 years) and no liabilities. Borrowing is not allowed.

The specification of the stochastic programming strategic asset allocation-FRC-model is presented in detail in Hicks Pedrón [22]. The model assumes 20 annual time periods in which random variables are realized at the beginning of each period. There are no stochastic parameters in the first period and period 21 is just a balance closing period in which no decisions take place.

The stages of the stochastic programming formulation – corresponding to annual portfolio rebalancing – are defined by associating a recourse problem with each realization of the rates of return on the different assets, computed by compounding biweekly returns to the end of every year from the twenty-year simulation. Figure 7 shows the bitmap of the base scenario for the resulting FRC_20 model.

For the EVPI-based sequential importance sampling experiments described here the last 12 annual rebalancing decisions were grouped into two year stages to make a 12 stage approximation FRC_12 of the full FRC_20 model.

Table 1 shows the results of 10 runs of the full EVPI importance sampling algorithm on the FRC_12 problem and Fig. 8 depicts the progress of the sequential algorithm in terms of the value of the expected terminal wealth objective of successive sample problems for the 10 runs. Table 2 and Fig. 9 show comparable results for the marginal EVPI sampling algorithm, while Fig. 10 shows the progress of the corresponding root node marginal EVPI values. Note that a high value of marginal EVPI is quickly reduced by further branching so that any arbitrages generated are handled by the importance sampling algorithm directly.

The sampling standard deviation across the runs is in each case re-

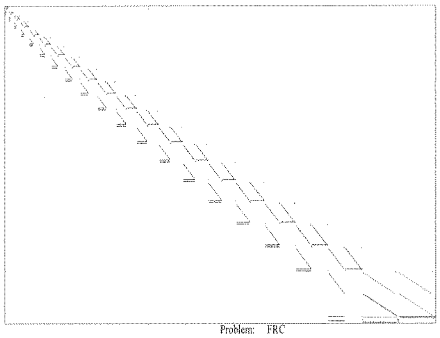

Fig. 7. Nonzero entries in FRC_20 base scenario coefficient matrix.

Table 1. Full EVPI sampling – FRC_12 problem – initial tree $32 = 8^1 2^2 1^8$ scenarios – sequential sampling with stopping criterion – 10 runs – tolerance 0.005 – max number of iterations in each sample is 15 – stopping tolerance is 3 decimal places – IBM RS6000/590/AIX 4.2.

Sample:	SCE	GEN	ITER	OPT	μ_{opt}	EVPI	%EVPI	CPU time	x_1^T	x_2^T	x_3^T
$FRC\_12\_32.1$	146	852	9	7.391	7.689	0.280	3.79%	$62'10''$	0.0	1.0	0.0
$FRC\_12\_32.2$	359	2124	9	7.850	7.789	0.483	6.15%	$105'15''$	0.0	1.0	0.0
$FRC\_12\_32.3$	237	1479	10	7.532	7.417	0.316	4.19%	$133'55''$	0.0	1.0	0.0
$FRC\_12\_32.4$	168	1454	12	8.004	7.935	0.524	6.55%	$116'52''$	0.0	1.0	0.0
$FRC\_12\_32.5$	216	2077	14	8.090	7.689	0.458	5.66%	$181'12''$	0.0	1.0	0.0
$FRC\_12\_32.6$	219	1976	14	7.497	7.587	0.411	5.48%	$171'45''$	0.0	1.0	0.0
$FRC\_12\_32.7$	283	2114	13	8.102	7.957	0.296	3.65%	$127'40''$	1.0	0.0	0.0
$FRC\_12\_32.8$	247	1829	13	6.949	7.395	0.307	4.42%	$105'23''$	0.0	1.0	0.0
$FRC\_12\_32.9$	182	1425	12	7.413	7.372	0.422	5.69%	$109'37''$	1.0	0.0	0.0
$FRC\_12\_32.10$	249	1484	10	8.081	7.760	0.711	8.79%	$136'03''$	0.0	1.0	0.0
μ	230	1681	12	7.691	7.658	0.421	5.44%	$135'12''$	0.2	0.8	0.0
σ	61	412	2	0.392	0.215	0.133	1.55%	$46'56''$	0.42	0.42	0
$\sigma/\mu\sqrt{10}$	8.38%	7.74%	5.28%	1.61%	0.88%	10.01%	9.01%	9.59%	66.7%	16.8%	0

spectable – 1.61% and 1.06% respectively for the full and marginal criteria. While this corresponds to an approximate 95% confidence interval of about 5% of objective value, much further experimentation with algo-

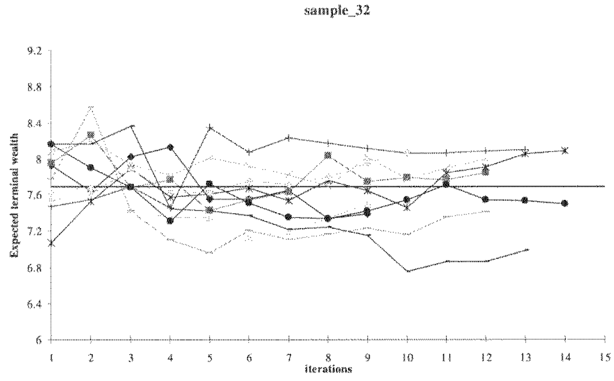

Fig. 8. Progress of the full EVPI importance sampling algorithm for 10 runs on FRC_12 with 32 scenario initial tree.

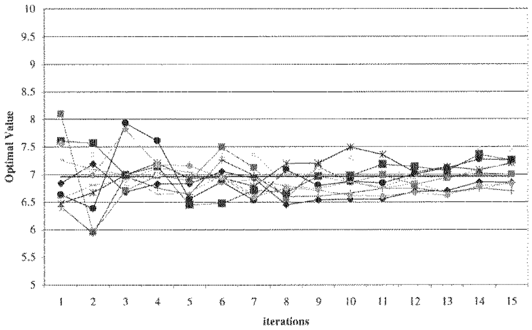

Fig. 9. Progress of the marginal EVPI importance sampling algorithm for 10 runs on FRC_12 with 32 scenario initial tree.

rithm strategies and parameters is required. The instability of the terminal first stage decisions $\hat{x}_{1N}$ shown for full EVPI in Table 1 is perhaps worrying. Nevertheless, in both cases variance reduction – in the sense of number of scenarios retained in the final sample tree versus (partial) scenarios generated – is about an order of magnitude (see Tables 1 and

Table 2. Marginal EVPI sampling – FRC_12 problem – initial tree 64 = $8^1 2^3 1^7$ scenarios – sequential sampling with stopping criterion – 10 runs – tolerance 0.05 – max number of iterations in each sample is 15 – Intel PII 400/Debian Linux 2.0.34.

Sample:	SCE	GEN	ITER	OPT	μ_{opt}	μ EVPI	CPU time
FRC_12_64.1	798	5199	15	6.81	6.78	0.85	18'27"
FRC_12_64.2	980	4806	15	6.94	6.97	0.87	27'31'
FRC_12_64.3	854	4704	15	7.37	7.15	0.92	19'36"
FRC_12_64.4	792	5099	15	7.35	7.06	0.92	17'37"
FRC_12_64.5	988	4463	15	7.17	7.06	0.89	18'48"
FRC_12_64.6	902	4830	15	7.27	7.01	0.91	18'30"
FRC_12_64.7	808	5193	15	6.66	6.70	0.83	17'55"
FRC_12_64.8	806	5100	15	6.97	6.94	0.87	17'52"
FRC_12_64.9	816	5275	15	6.84	6.83	0.86	18'22"
FRC_12_64.10	956	4871	15	6.64	6.91	0.83	19'25"
μ	860.44	4963.22	15	7.04	6.95	0.88	19'21"
σ	77.95	274.46	0	0.26	0.15	0.03	3'07"
$\sigma/\mu\sqrt{10}$	2.62%	1.60%	0	1.06%	0.62%	1.06%	4.65%

2). This empirical fact holds out the promise of the implicit generation of tens of thousands of scenarios in strategic financial planning problems over long term horizons.

As to bounding sampling error, Table 3 shows that the lower bound buy-and-hold strategy with reinvestment of dividends based on 10,000 sample paths gives a maximum sampling error bias of 114 basis points (1.14%) per annum for a single FRC_12 sample problem based on 1024 scenarios. By comparison, collapsing the time stages from an arbitrary instance of FRC_20 to this problem gives a difference of 32 basis points per annum.

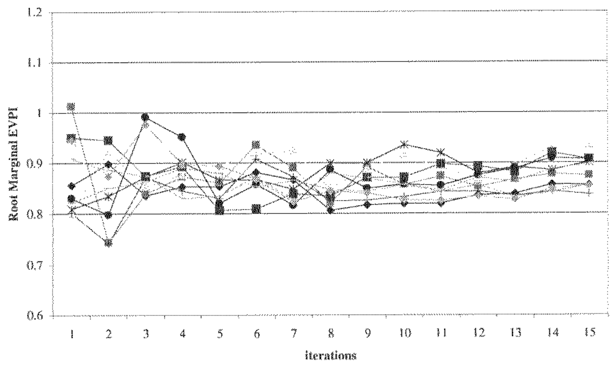

Fig. 10. Root node marginal EVPI process sample paths corresponding to Figure 9.

Table 3. Characteristics of the terminal wealth distributions.

Strategy	μ	σ	$\mathbf{S}R_T$ $\left[\frac{1+\mu-(1+0.05)^{20}}{\sigma}\right]^{\frac{1}{20}}$	Annual Average Compound Rate of Return
Stochastic Programming FRC_20.1024 with reinvestment	9.823	0.961	1.113	12.65% p.a.
Stochastic Programming FRC_12-20.1024 with reinvestment	9.210	0.916	1.111	12.32% p.a.
100/0 Stocks/Bonds Buy-and-hold with reinvestment	5.951	0.635	1.055	10.18% p.a.

9. Conclusion

This paper has attempted to give a comprehensive treatment of EVPI-based sequential importance sampling algorithms for dynamic stochastic programming problems. While the theory – at least for the linearly constrained convex case – is essentially complete, much experimentation with sampling strategies remains to be carried out.

The principal impediment to implementing this programme fully is the current running time of the complex software needed to effect sequential sampling (Chen *et al.* [7]). However, considerable progress – a factor of about 30 for a comparable number of scenarios in the final tree – has already been made in this direction (*cf.* Tables 1 and 2) and a final reduction in runtime of the algorithms to termination of several orders of magnitude is expected. To further reduce computation time we have already experimented with parallel full EVPI-based sequential importance sampling (Dempster and Thompson [18]) with some success and further progress is certainly possible. An additional promising line of attack on speedup and increased problem size is to use sparse memory interior point codes to enable larger sample problems with many more scenarios to be rapidly computed.

Finally, following the elegant recent paper of King [28] we hope ultimately to use these methods to produce accurate valuations of derivative financial instruments and portfolios.

Acknowledgments.

The ideas reported in this paper have developed over a long period and benefitted greatly from collaboration and discussion with Horand Gassmann, Zhao Yang, Mark Davis, Xavier Corvera-Poiré, Giorgio Consigli, Robin Thompson, John Mulvey, Nieves Hicks Pedrón and James Scott. The research has over the years enjoyed the generous support of NSERC Canada, EPSRC UK, 3rd Directorate EU and a number of Universities on both sides of the Atlantic – for all of which the author is grateful. The development of the FRC strategic asset allocation model as an instance of the more general CALM asset-liability management model was partially supported by the Frank Russell Company.

References

1. K. Back and S. R. Pliska, *The shadow price of information in continuous-time decision problems.* — Stochastics, **22** (1987), 151–186.
2. B. Bereanu, *On stochastic linear programming II: Distribution problems. Non-stochastic technology matrix.* — Rev. Roumaine Math. Pures Appl., **11** (1966), 713–725.
3. B. Bereanu, *On stochastic linear programming: Distribution problems. Stochastic technology matrix.* — Z. Wahrscheinlichkeitstheorie und Verw. Gebiete, **8** (1967), 148–152.
4. P. Billingsley, *Probability and Measure*, 2nd edition, Wiley, New York (1980).
5. J. R. Birge, M. A. H. Dempster, H. I. Gassmann, E. A. Gunn, A. J. King, and S. Wallace, *A standard input format for multiperiod stochastic linear programs.*

Mathematical Programming Society. — Committee on Algorithms Newsletter, **17** (1987), 1–20.

6. M. J. Brennan, E. S. Schwartz, and R. Lagnado, *Strategic asset allocation.* — Journal of Economic Dynamics and Control, **21** (1997), 1377–1403.
7. Z. Chen, G. Consigli, M. A. H. Dempster, and N. Hicks Pedrón, *Towards sequential sampling algorithms for dynamic portfolio management.* — In [49], pp. 197–211.
8. X. Corvera Poiré, *Model Generation and Sampling Algorithms for Dynamic Stochastic Programming.* PhD Thesis, Departmentt of Mathematics, University of Essex, UK (1995).
9. X. Corvera Poiré, STOCHGEN User's Manual. Department of Mathematics, University of Essex, UK (1995).
10. J. W. Daniel, *The Approximate Minimization of Functionals.* Prentice-Hall, Englewood Cliffs, NJ (1971).
11. G. B. Dantzig and P. W. Glynn, *Parallel processors for planning under uncertainty.* — Annals of Operations Research, **22** (1990), 1–21.
12. G. B. Dantzig and G. Infanger, *Large-scale stochastic linear programs: Importance sampling and Benders decomposition.* — Technical Report SOL 91-4, Department of Operations Research, Stanford University, Stanford, CA (1991).
13. G. B. Dantzig and A. Madansky, *On the solution of two-stage linear programs under uncertainty.* — In: Proceedings of the Fourth Symposium on Mathematical Statistics and Probability, Vol. I. Univ. of California, Berkeley (1961), pp. 165–176.
14. M. H. A. Davis, M. A. H. Dempster, and R. J. Elliott, *On the value of information in controlled diffusion processes.* — In [33], pp. 125–138.
15. M. A. H. Dempster (ed.), *Stochastic Programming*, Academic Press, London (1980).
16. M. A. H. Dempster, *The expected value of perfect information in the optimal evolution of stochastic systems.* — In: Stochastic Differential Systems, M. Arato, D. Vermes, and A. V. Balikrishnan (eds.), Lecture Notes in Information and Control **36**, Springer, Berlin (1981), pp. 25–40.
17. M. A. H. Dempster, *On stochastic programming: II. Dynamic problems under risk.* — Stochastics, **25** (1988), 15–42.
18. M. A. H. Dempster and R. T. Thompson, *EVPI-based importance sampling solution procedures for multistage stochastic linear programmes on parallel MIMD architectures.* — Annals of Operations Research **90** (1999), 161–184.
19. D. Dentcheva and W. Römisch, *Differential stability of two-stage stochastic programs.* Preprint No. 96-36, Institute of Mathematics, Humboldt University, Berlin (1996).
20. O. Fiedler and W. Römisch, *Stability in multistage stochastic programming.* — Annals of Operations Research, **56** (1995), 79–93.
21. H. I. Gassmann. *MSLiP: A computer code for the multistage stochastic linear programming problem.* — Mathematical Programming, **47** (1990), 407–423.
22. N. Hicks Pedrón, *Model-Based Asset Management: A Comparative Study.* PhD Thesis, Judge Institute of Management Studies, University of Cambridge, UK (1998).
23. J. L. Higle and S. Sen, *Stochastic decomposition: An algorithm for two stage linear programs with recourse.* — Mathematics of Operations Research, **16** (1991), 650–669.

24. J. L. Higle and S. Sen, *Stochastic Decomposition: A Statistical Method for Large Scale Stochastic Linear Programming.* Kluwer Academic Publishers, Dordrecht (1996).
25. G. Infanger, *Planning under Uncertainty: Solving Large-Scale Stochastic Linear Programs.* Boyd & Fraser, Danvers, MA (1994).
26. A. J. King, *Asymptotic Behaviour of Solutions in Stochastic Optimization: Nonsmooth Analysis and the Derivation of Non-normal Limit Distributions.* PhD Thesis, Department of Mathematics, University of Washington (1986).
27. A. J. King, *Generalized delta theorems for multivalued mappings and measurable selections.* — Mathematics of Operations Research, **14** (1989), 720–736.
28. A. J. King, *Martingales and duality in contingent claims analysis: The discrete case.* Research Report RC21153, IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY (1998).
29. A. J. King and R. T. Rockafellar, *Asymptotic theory for solutions in statistical estimation and stochastic programming.* — Mathematics of Operations Research, **18** (1993), 148–162.
30. M. Lane and P. Hutchinson, *A model for managing a certificate of deposit portfolio under uncertainty.* — In [15], pp. 473–493.
31. W. K. Mak, D. P. Morton, and R. K. Wood, *Monte Carlo bounding techniques for determining solution quality in stochastic programs.* Research Report, Department of Computer Science, University of Texas at Austin (1997).
32. H. M. Markowitz, *Portfolio Selection, Efficient Diversification of Investments*, 2nd edition, Blackwell, Oxford (1987).
33. E. Mayer-Wolf, E. Merzbach, and A. Schwartz, *Stochastic Analysis.* Academic Press, New York (1991).
34. V. I. Norkin, G. C. Pflug, and A. Ruszcynski, *A branch and bound method for stochastic global optimization.* Working Paper WP-96-065, International Institute for Applied Systems Analysis, Laxenburg, Austria (1996).
35. P. Olsen, *Discretizations of multistage stochastic programming problems.* — Mathematical Programming Studies, **6** (1976), 111–124.
36. G. Ch. Pflug, *Optimization of Stochastic Models: The Interface Between Simulation and Optimization.* Kluwer, Boston (1996).
37. A. Prékopa, *Laws of large numbers for random linear programs.* — Math. Systems Theory, **6** (1972), 277–288.
38. S. T. Rachev, *Probability Metrics and the Stability of Stochastic Models.* Wiley, New York (1991).
39. H. Raiffa, *Decision Analysis.* Addison-Wesley, Reading, MA (1968).
40. R. T. Rockafellar and R. J.-B. Wets, *Nonanticipativity and L^1-martingales in stochastic optimization problems.* — Mathematical Programming Studies, **6** (1976), 170–186.
41. R. T. Rockafellar and R. J.-B.Wets, *A Lagrangian finite generation technique for solving linear-quadratic problems in stochastic programming.* — Mathematical Programming Studies, **28** (1986), 63–93.
42. W. Römisch, *Quantitative stability of stochastic programs.* Presented at 16th International Symposium on Mathematical Programming, Lausanne (1997).
43. R. Y. Rubinstein and A. Shapiro, *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method.* Wiley, New York (1993).

44. G. Salinetti and R. J.-B.Wets, *On the distribution of measurable multifunctions (random sets), normal integrands, stochastic processes and stochastic infima.* — Mathematics of Operations Research, **11** (1986), 385–419.
45. A. Shapiro, *Asymptotic analysis of stochastic programs.* — Annals of Operations Research, **30** (1991), 169–186.
46. A. Shapiro, *Asymptotic behaviour of optimal solutions in stochastic programming.* — Mathematics of Operations Research, **18** (1993), 829–845.
47. A. Shapiro, *Simulation-based optimization: Convergence analysis and statistical inference.* — Commun. Statist. – Stochastic Models, **12** (1996), 425–454.
48. R. T. Thompson, *MSLiP-OSL 8.3 User's Guide.* Judge Institute of Management Studies, University of Cambridge, UK (1997).
49. C. Zopounidis (ed.), *Operational Tools in the Management of Financial Risks.* Kluwer, Dordrecht (1997).

Centre for Financial Research
Judge Institute of Management Studies
University of Cambridge,
Cambridge, England
and
Cambridge Systems Associates Limited
E-mail: mahd2@cam.ac.uk

Поступило 23 мая 2004 г.