

Toward breaking the curse of dimensionality: an FPTAS for stochastic dynamic programs with multidimensional actions and scalar states

Nir Halman * Giacomo Nannicini †

August 18, 2018

Abstract

We propose a Fully Polynomial-Time Approximation Scheme (FPTAS) for stochastic dynamic programs with multidimensional action, scalar state, convex costs and linear state transition function. The action spaces are polyhedral and described by parametric linear programs. This type of problems finds applications in the area of optimal planning under uncertainty, and can be thought of as the problem of optimally managing a single non-discrete resource over a finite time horizon. We show that under a value oracle model for the cost functions this result for one-dimensional state space is “best possible”, because a similar dynamic programming model with two-dimensional state space does not admit a PTAS.

The FPTAS relies on the solution of polynomial-sized linear programs to recursively compute an approximation of the value function at each stage. Our paper enlarges the class of dynamic programs that admit an FPTAS by showing, under suitable conditions, how to deal with multidimensional action spaces and with vectors of continuous random variables with bounded support. These results bring us one step closer to overcoming the curse of dimensionality of dynamic programming.

1 Introduction

A dynamic program (DP) is a mathematical model for sequential decision making. DPs are widely used by the operations research community to model and solve a large variety of problems concerning optimal planning under uncertainty. Unfortunately, DPs are affected by the *curse of dimensionality* – an expression coined by Richard E. Bellman more than 50 years ago [Bel61] – that makes their solution very difficult in practice. There is a large body of work devoted to ways to circumvent the curse, possibly foregoing optimality or approximation guarantees: this is discussed in Section 1.1.

This paper deals with a class of discrete-time finite-horizon stochastic DPs characterized by a scalar state and a multidimensional action, where the optimal action at each stage and state can be computed as the solution of a linear program (LP). We now give a more detailed description of the underlying mathematical model. The evolution of the state of the DP is governed by the transition function f_t and the equation $I_{t+1} = f_t(I_t, \vec{x}_t, \vec{D}_t)$, $t = 1, \dots, T$, where: t is the discrete time index, $I_t \in \mathcal{S}_t$ is the state of the system at time t (\mathcal{S}_t is the *state space* at stage t), $\vec{x}_t \in \mathcal{A}_t(I_t)$ is the action or decision to be selected at time t after observing state I_t ($\mathcal{A}_t(I_t)$ is the *action space* at stage t and state I_t), \vec{D}_t is a vector of interstage independent random variables (r.v.s) over the sample space \mathcal{D}_t , and T is the number of time periods. The random vector \vec{D}_t represents an exogenous information flow, and can be continuous or discrete.

*Hebrew University of Jerusalem, Israel, E-mail: halman@huji.ac.il. Supported in part by the Israel Science Foundation, grant number 399/17

†IBM T. J. Watson, Yorktown Heights, NY, E-mail: nannicini@us.ibm.com

The cost function $g_t(I_t, \vec{x}_t, \vec{D}_t)$ gives the cost of performing action \vec{x}_t from state I_t at time t for each possible realization of \vec{D}_t . Costs are accumulated over all time periods: the total incurred cost is equal to $\sum_{t=1}^T g_t(I_t, \vec{x}_t, \vec{D}_t) + g_{T+1}(I_{T+1})$. We use \vec{x}_t, \vec{D}_t to emphasize vector quantities, whereas the other quantities are scalars. The problem is that of choosing policies $\pi_t : \mathcal{S}_t \rightarrow \bigcup_{I_t \in \mathcal{S}_t} \mathcal{A}(I_t)$ in order to determine a sequence of actions $\vec{x}_1 \in \mathcal{A}(I_1), \dots, \vec{x}_T \in \mathcal{A}(I_T)$ that minimizes the expectation of the total incurred cost. This problem is called a *stochastic dynamic program*. Formally, we want to determine:

$$z^*(I_1) = \min_{\pi_1, \dots, \pi_T} \mathbb{E} \left[\sum_{t=1}^T g_t(I_t, \pi_t(I_t), \vec{D}_t) + g_{T+1}(I_{T+1}) \right], \quad (\text{DP})$$

subject to: $I_{t+1} = f_t(I_t, \pi_t(I_t), \vec{D}_t), \quad \forall t = 1, \dots, T,$

where I_1 is the initial state of the system and the expectation is taken with respect to the joint probability distribution of the random variables \vec{D}_t . The seminal result of Bellman [Bel57] is that the solution to this problem is given by a recursion:

Theorem 1.1 *For every initial state I_1 , the optimal value $z^*(I_1)$ of the DP is given by $z_1(I_1)$, where z_1 is the function defined by $z_{T+1}(I_{T+1}) = g_{T+1}(I_{T+1})$ together with the recursion:*

$$z_t(I_t) = \min_{\vec{x}_t \in \mathcal{A}_t(I_t)} \mathbb{E} \left\{ g_t(I_t, \vec{x}_t, \vec{D}_t) + z_{t+1}(f_t(I_t, \vec{x}_t, \vec{D}_t)) \right\}, \quad t = 1, \dots, T. \quad (1)$$

When the state and action spaces are finite, and the expectations can be computed with a finite process, this recursion gives a finite algorithm to compute the optimal value. However, this algorithm may require exponential time in general. If the state and action spaces are not finite (e.g. when they are intervals or polyhedra as is discussed in the present paper), or if we cannot compute expectations in finite time (e.g. when there are continuous r.v.s that do not possess a closed-form formula for the expectation), computing $z^*(I_1)$ may be intractable.

This paper proposes an FPTAS for a specific class of DPs. A polynomial-time approximation scheme (PTAS) is an approximation scheme that returns a solution whose cost is at most $(1 + \epsilon)$ times the optimal cost, where $\epsilon > 0$ is a given parameter. A PTAS runs in time polynomial in the (binary) input size. If the algorithm also runs in time polynomial in $1/\epsilon$, we call it an FPTAS. To construct an FPTAS, we must impose additional structure on the DP, because it is known that not all DPs admit a polynomial-time approximation scheme (see e.g. [HKL⁺14, Thm. 9.2]). We assume that the state I_t is a scalar, whereas the action \vec{x}_t can be vector-valued. We show in Sect. 3 that if the cost functions are only accessible via value oracles, the restriction that I_t is a scalar cannot be relaxed. At each stage, the vector \vec{D}_t is composed of a given number of independent but not necessarily identically distributed r.v.s with bounded support, that can have continuous or discrete distribution. The vectors \vec{D}_t are also interstage independent. (If the vectors are not independent, it is known that the DP may be APX-hard and therefore cannot admit an FPTAS, see e.g. [HKL⁺14, Thm. 10.1, Cor. 10.2].) Furthermore, we assume that: the transition functions f_t are affine; the costs functions g_t decompose to the sum of a deterministic piecewise linear convex function and a convex function that may depend on the r.v.s via an affine transformation; and the action spaces $\mathcal{A}_t(I_t)$ are polyhedral sets, described as the feasible region of a parametric LP with right-hand side parameter I_t . A formal description of the assumptions is given in Sect. 2.2. DPs satisfying these assumptions can be seen as multistage stochastic LPs, see e.g. [BL11], with a single variable linking consecutive stages (the state variable) and stochastic r.h.s. vector. The constraint matrix and the matrix linking the state between consecutive stages (also called “technology matrix”) are deterministic. However, unlike the sample average approximation approach typically used in multistage stochastic LP, our algorithm is deterministic and its running time depends polynomially on the number of stages. Our framework allows some degree of convex nonlinearity in the objective function, see Condition 3(ii). One can think of the type of optimization problems addressed

here as stochastic resource management problems with a single resource, see Sect. 2.3. An example of a problem that exhibits the required characteristics in the context of energy resource allocation is described in [PGa⁺12]. More applications are described in [NP13]: management of water in a reservoir, of cement for a construction company, of the cash reserve for an investment bank.

Organization of the paper. This paper is organized as follows. In Sect. 2 we define our notation, state our assumptions, and present an example of an application. In Sect. 3 we discuss the necessity of our assumptions and show some related hardness results. Sect. 4 introduces an algorithm to deal with the issue of number sizes growing too rapidly. Sect. 5 contains an FPTAS for the sum of r.v.s, and shows how to efficiently compute the expectation in the DP in the presence of continuous r.v.s. Sect. 6 puts all the building blocks together to design the FPTAS for problem (DP) satisfying our assumptions. Sect. 7 concludes the paper. In the rest of the Introduction, we position our paper as compared to the existing literature, and provide a summary of our contributions.

1.1 Relevance to existing literature

Dynamic programming is an invaluable tool for sequential decision making under uncertainty, and it has received a large amount of attention. DPs are often very difficult to solve in practice; for this reason, several approximate solution methodologies have been developed. [Pow11] discusses explicitly the three *curses of dimensionality* of DPs, namely: the large dimensionality of the state space, of the action space, and the difficulty or impossibility of computing expectations.

To deal with these curses, [CDJ14] studies fixed-dimensional stochastic dynamic programs with discrete state and action spaces over a finite horizon. Assuming that the cost-to-go functions are discrete L^{β} -convex (classes of discrete convex functions are discussed later in this subsection), [CDJ14] proposes a pseudopolynomial-time approximation scheme that satisfies an arbitrary pre-specified additive error guarantee. The proposed approximation algorithm is a generalization of the explicit enumeration algorithm. The main differences between our paper and [CDJ14] are: (i) [CDJ14] considers discrete state and action spaces (as opposed to continuous); (ii) it considers fixed dimensional (as opposed to one-dimensional) state spaces; (iii) it gives additive (as opposed to relative) error approximation; (iv) the running time of the approximation algorithm in [CDJ14] is pseudopolynomial (as opposed to polynomial) in the binary size of the input. We note that both [CDJ14] and the current paper are based on generalization of the technique of K -approximation sets and functions.

Another relevant work in dealing with the curses of dimensionality in options pricing and optimal stopping is [CG18]. [CG18] assumes that the payout functions from executing the option at any time and any realization of the stochastic process have real value between zero and one. Under these assumptions, [CG18] provides rigorous theoretical guarantees on the additive approximation error of the solution they provide. Whenever the stochastic process consists of i.i.d. random variables, [CG18] also provides a relative error guarantee. The main differences between our paper and [CG18] are: (i) [CG18] considers a high-dimensional state space but scalar action space (as opposed to a high-dimensional action space and scalar state); (ii) it deals with a non-Markovian stochastic optimization problem in which the option payout function is path dependent; (iii) it provides algorithms that are completely data-driven and require no knowledge of any distribution beyond the ability to generate samples; (iv) it uses very different techniques from our paper (pure-dual methods, dual martingales and deeply nested conditional expectations).

The discipline known as Approximate Dynamic Programming (ADP) also strives to deal with the three curses of dimensionality, while hopefully providing theoretical guarantees on the solution quality. Comprehensive references in this area are [Ber95, Pow11]. In most cases, ADP approaches are content with proving asymptotic convergence to a “best possible” value function, e.g., the best value function approximation that

can be obtained from a given set of basis functions [dFVR03]. This is considerably different from the approach presented in this paper: we compute an ϵ -approximate solution in polynomial time, for any given $\epsilon > 0$.

Our paper gives a framework to construct an FPTAS for a continuous optimization problem. The literature contains only a few general frameworks that yield approximation schemes for non-discrete optimization problems. The framework in [SS12] deals with stochastic LPs, and is a randomized scheme, whereas the algorithm given in the present paper is deterministic. The framework of [HKL⁺14] deals with stochastic discrete DPs in which the r.v.s are described explicitly as lists of scenarios $(d, \Pr(D = d))$, and the single period cost functions possess either monotone or convex structure. [HNO15a] studies a subclass of the DP model of [HKL⁺14], in which the single period cost functions are assumed to possess convex structure, and provides a faster FPTAS from both theoretical worst-case upper bounds and practical standpoint. An extension of [HKL⁺14] to continuous state and action spaces is given in [HN16]: however, [HN16] still deals with scalar state and action spaces, and discrete (scalar) r.v.s. Our paper lifts some of those restrictions.

To construct an approximation scheme for (DP), we study the problem of approximating the cumulative distribution function (CDF) of $\sum_{i=1}^n X_i$ where X_1, \dots, X_n are given r.v.s. This problem is well known to be #P-hard, see e.g. [KRT97], [HKM⁺09, Thm. 4.1]. It plays a central role in stochastic optimization because many problems in this area inherit #P-hardness from it, e.g. [HKM⁺09, HKL⁺14]. The problem is related to counting knapsack solutions: given a binary knapsack $\sum_{i=1}^n a_i x_i \leq b$, a possible way to determine the number of feasible solutions is to define discrete r.v.s X_i equal to a_i or 0 with probability $\frac{1}{2}$ each, and then compute $\Pr(\sum_{i=1}^n X_i \leq b)$. Based on the technique for counting knapsack solutions developed in [GKM⁺11] (see also [Hal16]), [LS14] gives the first FPTAS to approximate the CDF of the sum of r.v.s. The approach presented in this paper to deal with vectors of r.v.s gives, as a by-product, an FPTAS for the same problem, under weaker assumptions: the difference between [LS14] and our paper is discussed toward the end of the next section.

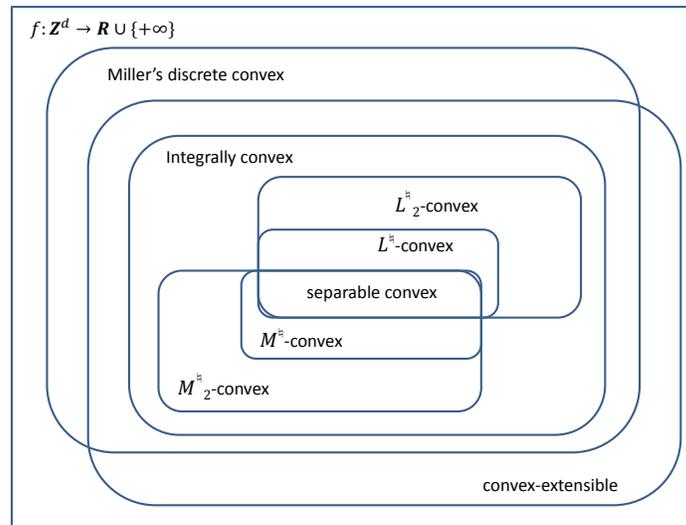


Figure 1: Classes of discrete convex functions (adapted from Figure 1.15 in [Mur03])

Another area of relevance to the present paper is that of discrete convexity, a discrete analog of convexity studied in the discrete optimization literature. The first investigation of a class of discrete functions for which local optimality implies global optimality is due to Miller [Mil71]. Favati and Tardella [FT90] considered a certain special way of extending functions defined over the integer lattice to piecewise-linear functions defined over the real space, and they introduced the concept of “integrally convex functions”.

[Mur03] introduced the concepts of “convex extensible”, “L-convexity” and “M-convexity,” in which “L” stands for “Lattice” and “M” stands for “Matroid”. The relationship between several classes of discrete convex functions is depicted in Fig. 1. [Hal15, Sec. 3] shows that separable convex functions in \mathbb{Z}^d admit polylogarithmic-space approximations, while any approximation of a two-dimensional nondecreasing convex function in the sense of Miller may require exponential space in the size of the domain of the function. Of the classes represented in Fig. 1, [Hal15] leaves open the approximability status of the class of convex-extensible functions; it is settled in the present paper, see Theorem 1.2 below.

1.2 Our contribution

Our paper makes progress toward overcoming some of the curses of dimensionality, while keeping the strong theoretical guarantee of finding an ϵ -approximate solution. First, the framework presented in this paper allows an action vector of arbitrary dimension, and the running time depends polynomially on such dimension. Second, we show how to handle vectors of continuous r.v.s with bounded support under suitable conditions using an oracle for their CDF, yielding an efficient approximate computation of expected values of the cost functions. We remark that continuous r.v.s are harder to handle than their discrete counterpart, because an exact computation of expected values cannot be carried out in finite time in general.

Thus, we summarize below our contribution regarding the three curses of dimensionality:

1. We can handle scalar state spaces in polynomial time, as was first shown in [HN16] under different assumptions. We show that if the cost functions are only accessible via value oracles, our continuous DP model with two-dimensional state space does not admit a PTAS.
2. We can handle vector action spaces of arbitrary dimension in polynomial time.
3. We can handle vectors of continuous r.v.s with bounded support in polynomial time. We can also handle vectors of discrete r.v.s, extending the work of [HN16] on scalar discrete r.v.s to the vector case. A more precise definition of the required conditions is given in Sect. 2.2.

Our approach to deal with vector action spaces is based on LP, and for this reason we impose structure on the problem that allows us to use an exact LP solver as a subroutine. Since our approximation scheme proceeds recursively, we must be careful to ensure that the size of the LPs does not grow too fast.

To prove that a DP with multidimensional state space does not admit a PTAS when cost functions are described by an oracle, we prove in Section 3.1 the following result (we define discrete convex-extensible in Section 3.1).

Theorem 1.2 *There exists a family of two-dimensional positive nondecreasing discrete convex-extensible functions for which any approximation (with approximation ratio below a certain threshold) requires exponential space in the size of the domain of the function, regardless of the scheme used to represent the approximation.*

This result is of independent interest and resolves the approximability status of discrete convex-extensible functions, see Fig. 1. Since discrete convex-extensible functions coincide with their convex envelope by definition, the above result implies that two-dimensional nondecreasing piecewise-linear convex functions do not admit an approximation that is polylogarithmic in the size of the domain (i.e., polynomial in the (binary) size of the domain’s description).

To handle vectors of r.v.s, we assume that such vectors appear in the cost and transition functions only via affine transformations. Our FPTAS to approximate the weighted sum of random variables has weaker assumptions than [LS14]. The work of [LS14] assumes that we are given integer discrete r.v.s and an exact oracle for the corresponding CDFs. It constructs an FPTAS for the sum of r.v.s by discretizing a

DP recurrence relation. We relax their assumptions: the approach described in this paper only requires an approximate oracle for the CDF of the r.v.s, and it can also handle non-integer discrete r.v.s and continuous r.v.s with bounded support. The running time is essentially the same. Our technique is not based on counting knapsack solutions, but rather on constructing compact representations for the value function of a DP: for this reason, we obtain a compact representation for the CDF of the sum of random variables over the entire domain, rather than its value at a single point. This contribution of our paper extends the applicability of approximation schemes for the sum of random variables.

The main constructive result shown in this paper is that a stochastic DP satisfying some structural assumptions (Conditions 1-3) admits an FPTAS. The approximation scheme is remarkably short and simple: its pseudocode consists of less than ten lines that use three different subroutines. The majority of this paper is devoted to their analysis. While the problem that we solve is continuous in nature, the technique used to build the approximation is discrete in nature, and may be useful for the construction of approximation schemes for continuous optimization problems. We remark that our approximation scheme is fully deterministic and does not rely on any form of random sampling, unlike known sample average approximation schemes for multistage optimization problems.

While the work presented in this paper is still far from the generality of ADP approaches, there are several real-world applications that fit into our framework, see Sect. 2.3. Furthermore, our framework provides much stronger theoretical guarantees than ADP: despite the fact that an exact solution to the problem may be impossible to compute in finite time, we determine an ϵ -approximation of the optimum for arbitrary ϵ . To the best of our knowledge, the present paper is the first that uses FPTASs to tackle two curses of dimensionality of DP, i.e., multidimensional action spaces and r.v.s: other FPTASs for DPs mentioned in our literature review do not address these two aspects.

2 Preliminaries

In this section, we formally introduce our notation and give the necessary definitions, discussing the type of problems that can be handled by the algorithm we propose.

2.1 Definitions

Let $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}^+$ and \mathbb{R} be the sets of nonnegative integers, integers, rational numbers, nonnegative real numbers and real numbers respectively. We distinguish between three types of continuous random variables: continuous r.v.s that may have unbounded support and for which $\Pr(X = x) = 0$ for all x , mixed continuous r.v.s that may have $\Pr(X = x) > 0$ for some x , and truncated continuous r.v.s that have bounded support with strictly positive probability on the endpoints only. Vectors are denoted with an overhead arrow, such as in \vec{x} . Subscripts should generally be intended to index a stage, e.g. \vec{x}_t is the action vector at stage t rather than the t -th component of a vector \vec{x} . When it is necessary to index a component of a vector, it will be made explicit, e.g. $\vec{D}_{t,i}$ will be the i -th component of the vector \vec{D}_t . For $\ell, u \in \mathbb{Z}$, we denote a finite set of the form $\{\ell, \ell + 1, \dots, u\}$ by $[\ell, \dots, u]$, whereas $[\ell, u]$ denotes a real interval.

Given a real-valued function over a bounded linearly-ordered domain D , $\varphi : D \rightarrow \mathbb{R}$, such that φ is not identically zero, we denote $D^{\min} := \min_{x \in D} \{x\}$, $D^{\max} := \max_{x \in D} \{x\}$, $\varphi^{\min} := \min_{x \in D} \{|\varphi(x)| : \varphi(x) \neq 0\}$, and $\varphi^{\max} := \max_{x \in D} \{|\varphi(x)|\}$. We write $x^+ := \max\{0, x\}$. For any function $\varphi : D \rightarrow \mathbb{R}$, t_φ denotes an upper bound on the time needed to evaluate φ on a single point in its domain. A function $\varphi : D \subseteq \mathbb{R} \rightarrow \mathbb{R}$ is called Lipschitz continuous with a given Lipschitz constant κ (κ -Lipschitz continuous, in short), if $|\varphi(x) - \varphi(y)| \leq \kappa|x - y|$ for all $x, y \in D$. Given a Lipschitz continuous function $\varphi : D \rightarrow \mathbb{R}$, we denote by κ_φ its Lipschitz constant. Let X be a set, and let $Y(x)$ be a set for every $x \in X$. We denote by $X \otimes Y$ the set $\bigcup_{x \in X} Y(x)$. We write $\log z$ to denote $\log_2 z$. When indicating running times of algorithms,

we give bounds on the number of operations, that we assume to have unitary cost. For arithmetic operations, the number of bit operations depends on the size of the numbers: because we always ensure that the size of the numbers is polynomially bounded in the (binary) input size, the number of bit operations is at most a polylogarithmic factor larger than the number of arithmetic operations.

The approximation scheme discussed in this paper relies on the concept of K -approximation sets and functions, introduced in [HKM⁺09]. Intuitively, a K -approximation function is a function that has a multiplicative error of at most K with respect to a given function. A K -approximation set is a set of points that induces a K -approximation function via linear interpolation. These concepts are formalized below. The definitions are based on [HKL⁺14, HN16]: [HKL⁺14] deals with functions over discrete domains, [HN16] deals with functions over continuous domains but employs both multiplicative and additive error.

Definition 2.1 *Let $K \geq 1$ and let $\varphi : D \subset \mathbb{R} \rightarrow \mathbb{R}^+$. We say that $\tilde{\varphi} : D \rightarrow \mathbb{R}^+$ is a K -approximation function of φ (or more briefly, a K -approximation of φ) if for all $x \in D$ we have $\varphi(x) \leq \tilde{\varphi}(x) \leq K\varphi(x)$.*

Definition 2.2 *Let $\varphi : D \subset \mathbb{R} \rightarrow \mathbb{R}$ be a convex function. For every finite $E \subseteq D$, the convex extension of φ induced by E is the function $\hat{\varphi} : [E^{\min}, E^{\max}] \rightarrow \mathbb{R}$ defined as the lower envelope of the convex hull of $\{(x, \varphi(x)) : x \in E\}$.*

Definition 2.3 *Let $\varphi : D \subset \mathbb{R} \rightarrow \mathbb{R}$ be nondecreasing. For every finite $E \subseteq D$, the monotone extension of φ induced by E is the function $\hat{\varphi} : [E^{\min}, E^{\max}] \rightarrow \mathbb{R}$ defined as $\hat{\varphi}(x) := \min\{\varphi(y) : y \in E, y \geq x\}$.*

Definition 2.4 *Let $K \geq 1$ and let $\varphi : D \subset \mathbb{R} \rightarrow \mathbb{R}^+$ be a convex (resp. nondecreasing) function. We say that a finite set $W \subset D$ is a K -approximation set of φ if the convex extension (resp. the monotone extension) of φ induced by W is a K -approximation function of φ .*

For a monotone nondecreasing function, an algorithm to compute a K -approximation set in polynomial time is described in [HKL⁺14] if the function is defined over integers. If the function is defined over a real interval and is strictly positive, such an algorithm is given in [HN16], see Apx. A. We call this algorithm APXSETINC, consistently with these papers: whether the discrete or continuous version should be applied will be clear from the context. K -approximation functions can be combined as indicated below.

Proposition 2.5 (Calculus of K -approximation functions, Prop. 5.1 in [HKL⁺14]) *For $i = 1, 2$ let $K_i \geq 1$, let $\varphi_i : D \rightarrow \mathbb{R}^+$ be an arbitrary function over domain D , and let $\tilde{\varphi}_i : D \rightarrow \mathbb{R}$ be a K_i -approximation of φ_i . Let $\psi_1 : D \rightarrow D$, and let $\alpha, \beta \in \mathbb{R}^+$. Then:*

1. φ_1 is a 1-approximation of itself,
2. (linearity of appr.) $\alpha + \beta\tilde{\varphi}_1$ is a K_1 -approximation of $\alpha + \beta\varphi_1$,
3. (summation of appr.) $\tilde{\varphi}_1 + \tilde{\varphi}_2$ is a $\max\{K_1, K_2\}$ -approximation of $\varphi_1 + \varphi_2$,
4. (composition of appr.) $\tilde{\varphi}_1(\psi_1)$ is a K_1 -approximation of $\varphi_1(\psi_1)$,
5. (minimization of appr.) $\min\{\tilde{\varphi}_1, \tilde{\varphi}_2\}$ is a $\max\{K_1, K_2\}$ -approximation of $\min\{\varphi_1, \varphi_2\}$,
6. (maximization of appr.) $\max\{\tilde{\varphi}_1, \tilde{\varphi}_2\}$ is a $\max\{K_1, K_2\}$ -approximation of $\max\{\varphi_1, \varphi_2\}$,
7. (appr. of appr.) If $\varphi_2 = \tilde{\varphi}_1$ then $\tilde{\varphi}_2$ is a K_1K_2 -approximation of φ_1 .

We call *canonical representation* of an approximation set W of a function φ its representation as a linearly ordered set of points and the corresponding function values. Formally, we write the canonical representation as $\{(x, \varphi(x)) \mid x \in W\}$, represented as an array sorted by its x coordinate. Such a representation corresponds to the data that is normally stored in a computer implementation of the proposed algorithms. Let φ be a convex (resp. monotone) function, W a K -approximation set of φ , and let $\hat{\varphi}$ be the convex (resp. monotone) extension of φ induced by W . We say that $\hat{\varphi}$ is a K -approximation function of φ that is given in the form of a canonical representation. (Note that by Def. 2.4, $\hat{\varphi}$ is indeed a K -approximation function of φ .) Given the canonical representation and any query point x in the domain of φ , one can compute in $O(\log |W|)$ time the value of $\hat{\varphi}$ at x . Therefore, the canonical representation serves as a value oracle for $\hat{\varphi}$ with query time $O(\log |W|)$. Given a canonical representation, we can recover the approximation set W on which it was constructed in time linear in $|W|$.

We define a routine called `COMPRESSINC`($\varphi, [A, B], K$), that given a monotone nondecreasing function $\varphi : [A, B] \rightarrow \mathbb{R}^+$ and an approximation factor K , constructs a value oracle (in the form of a canonical representation) for a K -approximation function of φ . Details are given in Apx. A. Similar routines will be constructed in the remainder of this paper.

The main idea for the approximation algorithm proposed in this paper is to construct a value function approximation following the backward recursion (1). The value function approximation is stored as a K -approximation set. At each stage, we must be able to efficiently compute the expected value appearing in (1), and to perform the minimization over the action space. We will show how to achieve these goals under the assumptions detailed in the next section.

2.2 Assumptions

To construct an FPTAS, we require the following conditions to be satisfied:

Condition 1 [*Domains*] \mathcal{S}_{T+1} and \mathcal{S}_t for $t = 1, \dots, T$ are intervals on the real line. $\mathcal{A}_t(I_t) := \{\vec{x}_t : A_t \vec{x}_t \geq \vec{b}_t + \vec{\delta}_{b_t} I_t, \vec{x}_t \geq 0\} \subset \mathbb{R}^p$ is a p -dimensional polyhedral set that is expressed as the feasible set of a parametric LP with p variables, where: the right-hand side vector is an affine function of the parameter I_t ; $A_t \in \mathbb{Q}^{m \times p}$; and $\vec{b}_t, \vec{\delta}_{b_t} \in \mathbb{Q}^m$ for $t = 1, \dots, T$.

Condition 2 [*Description of random events*] For every $t = 1, \dots, T + 1$, we have $\vec{D}_t \in \mathbb{R}^\ell$. For every $t = 1, \dots, T + 1$ and $i = 1, \dots, \ell$, one of the following two conditions hold for the i -th random variable $\vec{D}_{t,i}$ of the vector \vec{D}_t :

- (i) The random variable $\vec{D}_{t,i}$ is truncated continuous with compact support $\text{support}(\vec{D}_{t,i}) = [\vec{D}_{t,i}^{\min}, \vec{D}_{t,i}^{\max}] \subset \mathbb{R}$ of length $n_{t,i} = \vec{D}_{t,i}^{\max} - \vec{D}_{t,i}^{\min}$, and its CDF is Lipschitz continuous.
- (ii) The random variable $\vec{D}_{t,i}$ is discrete, its support $\text{support}(\vec{D}_{t,i}) \subset [\vec{D}_{t,i}^{\min}, \vec{D}_{t,i}^{\max}] \subset \mathbb{R}$ consists of $n_{t,i} < \infty$ elements, and the k -th largest element x in the support can be identified in time polylogarithmic in $n_{t,i}$.

Furthermore, $\Pr(\vec{D}_{t,i} = \vec{D}_{t,i}^{\min}) > 0$, $\Pr(\vec{D}_t = \vec{D}_{t,i}^{\max}) > 0$, and the information about the random events is given via value oracles to the CDF. All $\vec{D}_{t,i}$ are independent.

Condition 3 [*Structure of the functions*] For every $t = 1, \dots, T$, the function $f_t(I_t, \vec{x}_t, \vec{D}_t) : \mathcal{S}_t \otimes \mathcal{A}_t \times \mathcal{D}_t \rightarrow \mathcal{S}_{t+1}$ is linear in its variables and is expressed as $f_t(I_t, \vec{x}_t, \vec{D}_t) := \theta^I I_t + \vec{\theta}^x \cdot \vec{x}_t + \vec{\theta}^D \cdot \vec{D}_t$. The function g_t can be expressed as $g_t(I_t, \vec{x}_t, \vec{D}_t) = g_t^I(I_t, \vec{x}_t) + g_t^D(f_t^g(I_t, \vec{x}_t, \vec{D}_t))$, where: $g_t^I : \mathcal{S}_t \otimes \mathcal{A}_t \rightarrow \mathbb{R}^+$ is a piecewise linear convex function described as the pointwise maximum of a set of $q_t < \infty$ given hyperplanes; and $f_t^g : \mathcal{S}_t \otimes \mathcal{A}_t \times \mathcal{D}_t \rightarrow \mathcal{G}_t$ is an affine function expressed as $f_t^g(I_t, \vec{x}_t, \vec{D}_t) := \sigma^I I_t + \vec{\sigma}^x \cdot \vec{x}_t + \vec{\sigma}^D \cdot \vec{D}_t$. One of the following conditions hold:

- (i) The function $g_{T+1} : \mathcal{S}_{T+1} \rightarrow \mathbb{R}^+$ is positive piecewise linear convex over \mathcal{S}_{T+1} with m_{T+1} given breakpoints and slopes, and minimum value $g_{T+1}^{\min} > 0$. Furthermore, for all $t = 1, \dots, T$ the function $g_t^D(\cdot) : \mathcal{G}_t \subset \mathbb{R} \rightarrow \mathbb{R}^+$ is a piecewise linear convex function over a compact interval \mathcal{G}_t with m_t given breakpoints and slopes.
- (ii) The function $g_{T+1} : \mathcal{S}_{T+1} \rightarrow \mathbb{R}^+$ is Lipschitz continuous and convex over \mathcal{S}_{T+1} and is described via a value oracle. There is a given positive bound $\tilde{g}_{T+1}^{\min} > 0$ on the minimum value of $g_{T+1}(\cdot)$. Furthermore, for all $t = 1, \dots, T$ the function $g_t^D(\cdot) : \mathcal{G}_t \subset \mathbb{R} \rightarrow \mathbb{R}^+$ is non-negative Lipschitz continuous and convex over a compact interval \mathcal{G}_t , and is described via a value oracle.

We give next a few remarks about our DP model. In Condition 2(i) and Condition 3(ii) the Lipschitz constant need not be given a-priori. However, the running time of our algorithms depends polylogarithmically on it. In Condition 2(ii), we use $n_{t,i}$ to indicate either the size of the support in terms of interval length for continuous random variables, or the cardinality of the support for discrete random variables: this allows us to write the running time in a more compact way. An example of a truncated continuous r.v. satisfying Condition 2 is that of a Gaussian r.v. Y with given mean and standard deviation, clipped to a bounded interval: for example, if Y represents a demand, we can truncate it from below at zero because demand cannot be negative ($\Pr(X = 0) = \Pr(Y < 0)$), and from above at some positive order capacity u because all demand above this value is treated identically ($\Pr(X = u) = \Pr(Y > u)$). Truncated mixed r.v.s can in principle be handled under assumptions similar to Condition 2, but their treatment is cumbersome and is not pursued here. We remark that the implicit model for the description of the r.v.s, as stated in Condition 2, is both more general and more compact than listing all possible values in the support of the r.v.s and the corresponding probabilities. The latter approach, while simpler and more commonly employed, would only be viable for discrete r.v.s with support of small size.

Condition 3(i) is the traditional model in which the cost functions are analytically known in explicit form. Condition 3(ii) is a more flexible model that allows some of the cost functions to be described by value oracles. The FPTAS proposed in this paper is essentially the same under both Condition 3(i) and 3(ii), but some of our hardness results assume the oracle model. Studying the oracle model is appealing because positive results that hold in this setting have weak assumptions and can therefore be broadly applied. Value oracles can also allow strong negative results, see Sect. 3.

The input data of the problem includes: the number of time periods T ; the initial state I_1 ; the constant $\gamma = \min\{\Pr(\vec{D}_{t,i}^{\max}), \Pr(\vec{D}_{t,i}^{\min}) \mid t = 1, \dots, T, i = 1, \dots, \ell\}$ (Condition 2 implies $\gamma > 0$); the maximum cost in a time period U_g ; the maximum Lipschitz constant κ of the cost functions g_t^I, g_t^D (and of the CDF of $\vec{D}_{t,i}$ in the case of Condition 2(i)) for $t = 1, \dots, T + 1$; the minimum value $\mu := g_{T+1}^{\min}$ of the terminal cost function; the maximum length U_S of the state spaces; the maximum size U_A of the coefficients in the constraint matrix or right-hand side vector of the LPs describing $A_t(I_t)$; the maximum size U_f of the coefficients in the vectors $\vec{\theta}^x, \vec{\sigma}^x, \vec{\theta}^D, \vec{\sigma}^D$ and of the scalars θ^I, σ^I in the description of f_t and f_t^g ; the maximum numbers of pieces $m^* = \max_t m_t$ and $q^* = \max_t q_t$ in the piecewise description of the cost functions; the maximum size of the support $n^* = \max_{t,i} n_{t,i}$ of the r.v.s. Our algorithms run in time polynomial in $\log I_1, \log 1/\gamma, \log 1/\mu, \log U_g, \log \kappa, \log U_S, \log U_A, \log U_f, m^*, q^*, p, \ell, T \log n^*$. Throughout this paper, the Lipschitz constants of the functions involved need not be known, but the running time of the algorithms depends on them: intuitively, a function with high Lipschitz constant may change quickly, and it may take longer for the algorithms to identify sudden ‘‘jumps’’ in function value with the required precision. The necessity of our assumptions is discussed in Sect. 3.

2.3 An application to resource management under uncertainty

We describe here an example of the type of models allowed under the assumptions given in the previous section. The example involves managing a resource of non-discrete nature, e.g., liquid natural gas or cement,

at a central storing facility. Let $G = (V, E)$ be a directed graph with $|V| = m + 1$ nodes with a special node v_0 that is the central storing facility; for $v \in V$, we denote by $\delta^+(v)$ its set of outgoing edges and by $\delta^-(v)$ its set of ingoing edges. The optimization is performed over a finite time horizon T . The state of the system I_t is the amount of resource available at the central storing facility at the beginning of time period t . At the end of every time period, there is an unknown arrival of the resource encoded by a r.v. $\vec{D}_{t,0}$, e.g., cargo ships (in which case, the r.v. is discrete). There are m locations, v_1, \dots, v_m that consume the resource but do not have (between time periods) storing capabilities. Each location has unknown demand at time period t , distributed according to a continuous r.v. $\vec{D}_{t,i}$ with $i = 1, \dots, m$. The demand cannot be transferred to other locations: unsatisfied demand or excess inventory at each of the locations is lost. The resource can be moved from the central storing facility to the m locations, as well as between locations, over a capacitated flow network described by G , potentially with losses $1 - a_{ij}$ over the arcs $(i, j) \in E$ (i.e., the flow variables may appear with coefficients < 1 in the flow balance constraints at some nodes). There is no backlogging at the central storing facility: inventory must stay nonnegative (but we could easily amend the formulation to allow backlogging). At each time period t , one observes the state of the system I_t , and based on the distribution of the demand in time periods $t, t + 1, \dots, T$ decides on the flow $f_{t,ij}$ over each arc $(i, j) \in E$. At the end of the period a cost is incurred that consists of the transportation cost given by the flow network with unit cost c_{ij} on each arc, and shortage (resp. disposal) cost s_i (resp. h_i) that is accounted for every unit by which the demand $\vec{D}_{t,i}$ at location i is not met (resp. is exceeded). At the final time period $T + 1$, there is a disposal cost h_0 for every unit left in the inventory. We denote by $y_{t,i}$, $i = 0, \dots, m$ the flow balance (exiting minus entering) at location i , therefore $y_{t,i}$ is positive if the node sends out more than it receives, negative otherwise. Furthermore, we denote by $w_{t,i}$ an auxiliary variable to encode the cost incurred at the end of time period $t - 1$ at location i , i.e., the maximum between the shortage and disposal costs.

This problem can be formulated as a multistage stochastic linear program as follows:

$$\begin{aligned} \min \mathbb{E}_{\vec{D}} \left[\sum_{t=1}^T \left(\sum_{(i,j) \in E} c_{ij} f_{t,ij}(\vec{D}_t) + \sum_{i=1}^m w_{t+1,i}(\vec{D}_t) \right) + h_0 I_{T+1}(\vec{D}_T) \right] \quad (2) \\ t = 1, \dots, T \quad -I_t(\vec{D}_{t-1}) + I_{t+1}(\vec{D}_t) + y_{t,0}(\vec{D}_t) = \vec{D}_{t,0} \quad (3) \\ t = 1, \dots, T, i = 0, \dots, m \quad \sum_{(i,j) \in \delta^+(i)} f_{t,ij}(\vec{D}_t) - \sum_{(j,i) \in \delta^-(i)} a_{ji} f_{t,ji}(\vec{D}_t) = y_{t,i}(\vec{D}_t) \quad (4) \\ t = 1, \dots, T \quad y_{t,0}(\vec{D}_t) \leq I_t(\vec{D}_{t-1}) \quad (5) \\ t = 2, \dots, T + 1, i = 1, \dots, m \quad -h_i y_{t-1,i}(\vec{D}_{t-1}) - h_i \vec{D}_{t-1,i} \leq w_{t,i}(\vec{D}_t) \quad (6) \\ t = 2, \dots, T + 1, i = 1, \dots, m \quad s_i y_{t-1,i}(\vec{D}_{t-1}) + s_i \vec{D}_{t-1,i} \leq w_{t,i}(\vec{D}_t) \quad (7) \\ t = 1, \dots, T, (i, j) \in E \quad f_{t,ij}(\vec{D}_t) \geq 0 \quad (8) \\ t = 1, \dots, T, i = 1, \dots, m \quad y_{t,i}(\vec{D}_t) \leq 0 \quad (9) \\ I_1(\vec{D}_0) = \bar{I}_1. \quad (10) \end{aligned}$$

In the above formulation, \bar{I}_1 is the initial inventory. The inventory level $I_{t+1}(\vec{D}_t)$ depends on the random variable at time period t , as indicated in constraint (3): I_{t+1} is automatically determined after all decisions at time period t are taken and \vec{D}_t is revealed. Similarly, the auxiliary variable $w_{t,i}(\vec{D}_t)$ encodes the cost at location i at the end of the previous time period, because such cost can only be computed after the demand at a given time period is revealed. \vec{D}_0 only serves the purpose of keeping this notation consistent for $I_1(\vec{D}_0)$, and is otherwise not used in the model. Constraint (5) implies $I_t(\vec{D}_{t-1}) \geq 0$.

If the $\vec{D}_{t,i}$ are all discrete r.v.s, this problem admits a deterministic equivalent formulation by introducing a copy of the decision variables for each sample path, but the number of sample paths is exponential in T . We now show that a DP formulation of this problem satisfies Conditions 1-3.

Let I_t be the state variable. The action vector is defined by $\vec{x}_t = ((y_{t,i})_{i=1,\dots,m}, (f_{t,ij})_{(i,j) \in E})$; for simplicity, we use $y_{t,i}, f_{t,ij}$ to refer to the variables in the action vector \vec{x}_t corresponding to the variables with the same name in (2)-(10). The transition function is $f_t(I_t, \vec{x}_t, \vec{D}_t) = I_t - y_{t,0} + \vec{D}_{t,0}$. The action space is $\mathcal{A}_t(I_t) = \{\vec{x}_t : \sum_{(i,j) \in \delta^+(i)} f_{t,ij} - \sum_{(j,i) \in \delta^-(i)} a_{ji} f_{t,ji} = y_{t,i} \forall i = 0, \dots, m; y_{t,0} \leq I_t; f_{t,ij} \geq 0 \forall (i,j) \in E; y_{t,i} \leq 0 \forall i = 1, \dots, m\}$. The cost functions are: $g_t^I(I_t, \vec{x}) = \sum_{(i,j) \in E} c_{ij} f_{t,ij}$, $g_t^D(I_t, \vec{x}, \vec{D}_t) = h_i(-y_{t,i} - \vec{D}_{t,i})^+ + s_i(y_{t,i} + \vec{D}_{t,i})^+$, $g_{T+1}(I_{T+1}) = h_0 I_{T+1}$. It is easy to verify that with the cost functions defined as given, we obtain an equivalent description to problem (2)-(10), and Conditions 1-3(i) are satisfied.

3 Hardness results

We now discuss the necessity of our assumptions. Regarding Condition 1, it is an open question under what additional assumptions the restriction that the state spaces \mathcal{S}_t are intervals on the real line can be lifted under our DP model. However, the next subsections show that constructing an approximation of the value function is difficult already whenever \mathcal{S}_t are two-dimensional boxes (Cor. 3.3). We show this by first proving the existence of a class of two-dimensional piecewise linear convex functions that are hard to approximate.

Regarding Condition 2, the conditions on truncated r.v.s seem difficult to relax under the oracle model for the CDFs: [HN16, Prop. 2.6] shows that a continuous function over a compact interval may not admit an efficient approximation unless its values at the endpoints of the domain are bounded away from zero. This result holds even if the function is monotone and Lipschitz continuous. Hence, approximating the CDF of a continuous r.v. that does not have positive probability mass at the endpoints seems difficult. ([HN16] also shows that for bounded Lipschitz continuous functions we can drop the strict nonnegativity requirement if we allow both relative and absolute error at the same time; however, in this paper we aim for an FPTAS in the usual sense.) Regarding Condition 3, the nonnegativity requirement on the cost functions g_t cannot be relaxed: [HNO15b] exhibits a DP that is #P-hard to approximate to any constant factor, and their problem can be cast into our framework if we allow the cost functions to be of either sign. In other words, the #P-hard DP in [HNO15b] satisfies Conditions 1-3 except for the nonnegativity requirement on g_t . It is an open question whether the restriction $g_{T+1}^{\min} > 0$ (or $\tilde{g}_{T+1}^{\min} > 0$ in the setting of Condition 3(ii)) can be relaxed to $g_{T+1}^{\min} \geq 0$ (to $\tilde{g}_{T+1}^{\min} \geq 0$ in the setting of Condition 3(ii)). The difficulty stemming from $g_{T+1}^{\min} = 0$ is that we would have to keep track of the exact location of the zeroes of the value function at each stage.

3.1 Hardness of approximation of multivariate convex functions

The DP model discussed in this paper naturally yields piecewise linear convex value functions. To extend the framework to a multidimensional state space setting we would have to efficiently build approximations of such functions. However, the next results show that some piecewise linear convex functions do not admit an efficient approximation. We first show that there is an exponential number of two-dimensional piecewise linear convex functions with distinct values at integer points, and use this to prove that approximations of such functions require space exponential in the size of the domain.

Theorem 3.1 *For any nonnegative integer A , there exist $\Omega(2^{\sqrt{U}})$ distinct nondecreasing piecewise linear convex functions $\varphi : [1, U]^2 \rightarrow [A, A+U]$ that attain the value A at distinct sets of points in $[1, \dots, U]^2$, and at least $A+1$ at the other sets of integer points. These functions are described as the pointwise maximum of $O(\sqrt{U})$ hyperplanes,*

Proof. Let

$$\Phi := \left\{ \varphi_{r_1, \dots, r_U} : [1, U]^2 \rightarrow [A, A+U] \mid r_1, \dots, r_U \in [0, U]; \sum_{i=1}^U r_i = U, r_1 \leq r_2 \leq \dots \leq r_U \right\}$$

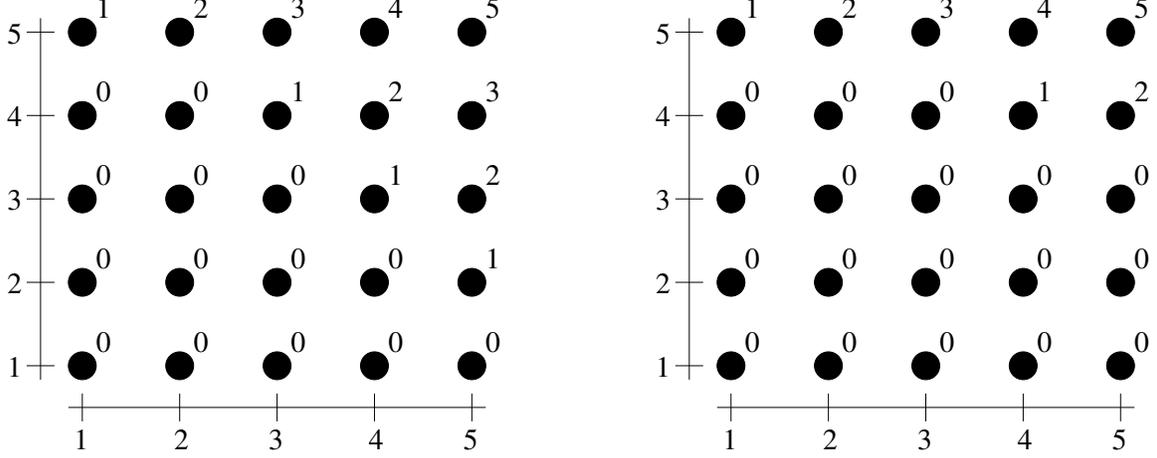


Figure 2: Values of the function $\varphi_{0,1,1,1,2}$ (left) and $\varphi_{0,0,0,2,3}$ (right) on the integer grid $[1, \dots, 5]^2$, with $A = 0$. At each integer point in the plane, we indicate the corresponding function value next to it.

be a family of bivariate discrete functions defined as follows: for $x, y \in [1, \dots, U]$,

$$\varphi_{r_1, \dots, r_U}(x, y) = \begin{cases} A & \text{if } x = 1, \dots, U - \sum_{i=1}^y r_i \\ A + x - U + \sum_{i=1}^y r_i & \text{otherwise.} \end{cases} \quad (11)$$

[Hal15, Prop. 3.4] shows (taking $A = 0$) that any pair of elements of Φ attains the value A at different points on the integer grid. Furthermore, it shows that the number of such functions is equal to the number of partitions of the integer U , which is bounded below by $\frac{H}{U} e^{2\sqrt{U}} \geq H2^{\sqrt{U}}$. In Fig. 2 we illustrate the shape of $\varphi_{r_1, \dots, r_U}$ by plotting its values for two choices of r_1, \dots, r_U .

Let $\text{diff}(r_1, \dots, r_U)$ be the number of distinct values among r_1, \dots, r_U . We show next that each function $\varphi_{r_1, \dots, r_U}$ can be described as the pointwise maximum of at most $\text{diff}(r_1, \dots, r_U) + 1$ hyperplanes, in a such a way that the value of the pointwise maximum on the integer grid matches that of $\varphi_{r_1, \dots, r_U}$. Note that the slopes of $\varphi_{r_1, \dots, r_U}$ satisfy the following relationship:

$$\varphi_{r_1, \dots, r_U}(x+1, y) - \varphi_{r_1, \dots, r_U}(x, y) = \begin{cases} 0 & \text{if } x < U - \sum_{i=1}^y r_i \\ 1 & \text{otherwise,} \end{cases} \quad (12)$$

$$\varphi_{r_1, \dots, r_U}(x, y+1) - \varphi_{r_1, \dots, r_U}(x, y) = \begin{cases} 0 & \text{if } x \leq U - \sum_{i=1}^{y+1} r_i \\ r_{y+1} + \min\{x - U + \sum_{i=1}^y r_i, 0\} & \text{otherwise.} \end{cases} \quad (13)$$

Furthermore, $\varphi_{r_1, \dots, r_U}(x, U) = A + x$ for all valid choices of r_1, \dots, r_U .

Claim. For $1 \leq k \leq U - 1$, there exists $S_k := \{(a, b, c) \in \mathbb{N}\}$, with $|S_k| \leq \text{diff}(r_1, \dots, r_U) + 1$, such that $\varphi_{r_1, \dots, r_U}(x, y) = \max_{(a, b, c) \in S_k} \{ax + by + c\}$ for all $1 \leq x \leq U, y \geq k$.

Proof of the claim. By backward induction. We start with $k = U - 1$. Consider the set $S_k := \{(0, 0, A), (1, r_U, -r_U U)\}$, which corresponds to two hyperplanes given in the form $ax + by + c$: one “flat” hyperplane at the value A , and one hyperplane with slope 1 along the x axis, r_U along the y axis. Consider the latter hyperplane, which we call H . By the choice of $c = -r_U U$, H interpolates $\varphi_{r_1, \dots, r_U}$ at the points (x, U) for all $1 \leq x \leq U$. Furthermore, by (13), it interpolates $\varphi_{r_1, \dots, r_U}$ at the points $(x, U - 1)$ for which $x \geq U - \sum_{i=1}^k r_i$, because for such points its slope on the y direction matches the slope of $\varphi_{r_1, \dots, r_U}$. It remains to analyze the points $(x, U - 1)$ with $x < U - \sum_{i=1}^k r_i$. Let P be the set of such points. For $(x, y) \in P$, $\varphi_{r_1, \dots, r_U}(x, y)$ has value A by definition, see (11). Furthermore,

$\varphi_{r_1, \dots, r_U}(x, y + 1) = \varphi_{r_1, \dots, r_U}(x, U) = x + A < U - \sum_{i=1}^k r_i + A = r_U + A$. Then, since H interpolates $\varphi_{r_1, \dots, r_U}$ at $(x, y + 1)$, and its slope is r_U in the y direction, we have that the value of H at (x, y) is equal to $x + A - r_U < A$. But then the maximum of the two hyperplanes in S_k at $(x, y) \in P$ is given by A (i.e., it is given by the “flat” hyperplane), which is exactly the value of $\varphi_{r_1, \dots, r_U}(x, y)$ by equation (11). This concludes the initial step of the induction.

We now need to prove the induction step k , assuming the claim is true for $k + 1$. Suppose first that $r_k = r_{k+1}$. In this case, a similar argument to the one above shows that no further hyperplane is needed in S_k , because the slope of $\varphi_{r_1, \dots, r_U}$ along the y axis for points with $x \geq U - \sum_{i=1}^k r_i$ and $y = k$ is the same as for $y = k + 1$, and for the remaining points, the “flat” hyperplane suffices. We need to analyze the case $r_k < r_{k+1}$. In this case, let $S_k = S_{k+1} \cup \{(1, r_{k+1}, \varphi_{r_1, \dots, r_U}(U, k + 1) - U - r_{k+1}(k + 1))\}$. Let H be the hyperplane defined by the coefficients $(1, r_{k+1}, \varphi_{r_1, \dots, r_U}(U, k + 1) - U - r_{k+1}(k + 1))$; it is easy to verify that H interpolates $\varphi_{r_1, \dots, r_U}$ at $(x, k + 1)$ for all $1 \leq x \leq U$ by construction. The argument now proceeds as before: the hyperplane interpolates $\varphi_{r_1, \dots, r_U}$ at points (x, k) for which $x \geq U - \sum_{i=1}^k r_i$, because for such points its slope on the y direction matches the slope r_{k+1} of $\varphi_{r_1, \dots, r_U}$. The remaining points with $y = k$ have value A , therefore the “flat” hyperplane interpolates them and H has value $< A$ at them. Furthermore, because the $r_{k+1} < r_{k+2} \leq \dots \leq r_U$, the value of H at all points (x, y) with $y > k$ is smaller than $\varphi_{r_1, \dots, r_U}(x, y)$. This concludes the inductive claim.

The above claim shows that we can construct a set S_1 containing hyperplanes such that for every $(x, y) \in [1, \dots, U]^2$, the pointwise maximum of the hyperplanes at (x, y) is equal to $\varphi_{r_1, \dots, r_U}(x, y)$. Note that all hyperplanes except the “flat” hyperplane have positive integer slopes. Thus, the pointwise maximum of the hyperplanes at integer points is integer-valued. This implies that the value of the pointwise maximum of such hyperplanes on the integer grid is either A or $\geq A + 1$. To conclude the proof, we need a bound on $|S_1| = \text{diff}(r_1, \dots, r_U) + 1$. Recall that $\sum_{i=1}^U r_i = U, r_1 \leq r_2 \leq \dots \leq r_U$. The largest possible number of distinct values of r_i is attained when all r_i that increase, increase by exactly one. This implies that it is the largest number d such that $\sum_{j=1}^d j = \frac{d}{2}(d + 1) \leq U$. Then, $\text{diff}(r_1, \dots, r_U) = O(\sqrt{U})$. \square

Corollary 3.2 *Let A and U be nonnegative integers, and let Φ be a family of nondecreasing convex functions $\varphi : [1, U]^2 \rightarrow [A, A + U]$ that are described as a pointwise maximum of hyperplanes. Then, any approximation of $\varphi \in \Phi$ that attains relative error less than $\frac{A+1}{A}$, or absolute error less than 1, requires $\Omega(\sqrt{U})$ space regardless of the scheme used to represent the function.*

Proof. By Thm. 3.1, there are $\Omega(2^{\sqrt{U}})$ different nondecreasing convex functions that attain the value A at distinct points in $[1, \dots, U]^2$, and value at least $A + 1$ at the remaining integer points. Any approximation of φ that attains relative error less than $\frac{A+1}{A}$, or absolute error less than 1, characterizes exactly the points at which φ attains the value A . Thus, the approximation must distinguish between all functions in the family Φ described in the proof of Thm. 3.1. This implies that every function in Φ must have a different description, regardless of the scheme used to describe a function. Since $|\Phi| \in \Omega(2^{\sqrt{U}})$, it follows that any approximation of φ must take $\Omega(\sqrt{U})$ space. \square

As mentioned in Sect. 1.1, there exist several special classes of multivariate discrete convex functions; the inclusion relationships among these classes are depicted in Figure 1. Corollary 3.2 has implications on the approximability of convex-extensible functions, which we define next. Let $f : \mathbb{Z}^d \rightarrow \mathbb{R}^+$ be a function defined over the d -dimensional lattice of integer points. The convex closure of f is defined to be a function $\bar{f} : \mathbb{R}^n \rightarrow \mathbb{R}^+$ given by

$$\bar{f}(x) = \sup_{p \in \mathbb{R}^n, \alpha \in \mathbb{R}} \{ \langle p, x \rangle + \alpha \mid \langle p, y \rangle + \alpha \leq f(y), \forall y \in \mathbb{Z}^n \} \quad \forall x \in \mathbb{R}^n.$$

If this function \bar{f} coincides with f on integer points, i.e., if $\bar{f}(x) = f(x) \quad \forall x \in \mathbb{Z}^n$, we say that f is *convex extensible* and call \bar{f} the *convex extension of f* [Mur03, Chap. 3.4].

The starting point for the proof of Theorem 3.1 is the family Φ of bivariate discrete convex functions (in the sense of Miller [Mil71]) over domain $[1, \dots, U]^2$ as defined in the proof of [Hal15, Prop. 3.4]. For each function $f \in \Phi$ we construct a set S_1 of $O(\sqrt{U})$ hyperplanes and define a (continuous) convex function $\bar{f} : [1, U] \rightarrow \mathbb{R}^+$ as the pointwise maximum of the hyperplanes in S_1 . We then show that f coincides with \bar{f} on the integer square $[1, \dots, U]^2$ and therefore the family Φ consists of discrete functions that are convex-extensible (and are also convex in the sense of Miller at the same time). Using the fact that $|\Phi| = \Omega(\sqrt{U})$ [Hal15, Prop. 3.4], and following similar arguments as in the proof of Corollary 3.2, Theorem 1.2 follows. We remark that the functions in [Hal15, Prop. 3.4] are discrete convex as opposed to convex in the ordinary sense, and the two concepts do not necessarily coincide; this paper settles the status of discrete convex-extensible and continuous (ordinary) piecewise linear convex functions, showing their inapproximability in space polynomial in the (binary) size of the domain.

3.2 Hardness of multidimensional DP

In this section we make some remarks about the approximability of (DP) when functions g_{T+1}, g_t^D are described by a value oracle and are not necessarily piecewise linear, as in Condition 3(ii). We focus on these two components of the cost function because the remaining terms g_t^I appear in the minimization over an ℓ -dimensional polyhedron; achieving an FPTAS for the case with g_t^I described by a value oracle would therefore require many additional mathematical tools, and is outside the scope of this paper.

As a consequence of Cor. 3.2, a two-dimensional generalization of problem (DP) is hard to approximate under the oracle model. We formalize this below. Consider the problem:

$$z^*(\vec{I}_1) = \min_{\pi_1, \dots, \pi_T} \mathbb{E} \left[\sum_{t=1}^T g_t(\vec{I}_t, \pi_t(\vec{I}_t), \vec{D}_t) + g_{T+1}(\vec{I}_{T+1}) \right], \quad (\text{MD-DP})$$

subject to: $\vec{I}_{t+1} = \vec{f}_t(\vec{I}_t, \pi_t(\vec{I}_t), \vec{D}_t), \quad \forall t = 1, \dots, T,$

where $\vec{I}_t \in \mathbb{R}^d$, \vec{f}_t is a vector-valued function, each component of \vec{f}_t is affine, g_{T+1} is strictly positive convex, and g_t can be expressed as $g_t(\vec{I}_t, \vec{x}_t, \vec{D}_t) = g_t^I(\vec{I}_t, \vec{x}_t) + g_t^D(\vec{f}_t^g(\vec{I}_t, \vec{x}_t, \vec{D}_t))$, where g_t^I and $g_t^D(\cdot)$ are nonnegative convex, and \vec{f}_t^g has d components each of which is affine. Then, if g_{T+1} or g_t^D are described by value oracles, Cor. 3.2 immediately implies the following.

Corollary 3.3 *There does not exist a PTAS for continuous dynamic programs of the form (MD-DP) with cost functions described by a value oracle, even if $d = 2, T = 0$, g_1 is nonincreasing, piecewise linear convex and bounded from below by a given positive number.*

By contrast, when $d = 1$ this paper provides an FPTAS under the same assumptions.

3.3 Hardness of approximating the value function of two-stage stochastic linear programs

In addition to DP, a widely used paradigm for optimization under uncertainty is stochastic programming. Due its importance and applicability, we find it worthy to report that Thm. 3.1 has implications on the hardness of approximation of stochastic programs. The concept of value functions is widely used in stochastic programming just as in DP. It is well known that the value function of a stochastic linear program may have a number of hyperplanes that is proportional to the number of values in the support of the r.v.s. In this section we show that if there are two variables linking consecutive stages, even an approximation of the value function may require a large number of hyperplanes. In particular, if the r.v.s are described in a

compact form, i.e., as a value oracle to their CDF, then such number is exponential in the input size: this is formalized in the next theorem. By contrast, if there is only one linking variable, we can construct such an approximation, with the approach discussed in the previous section.

Theorem 3.4 *There does not exist a PTAS to compute the optimal value function for the second stage of a two-stage stochastic linear program with a single discrete r.v. described via an oracle to its CDF, even if there are only two variables linking the first and second stage, and the value function is bounded from below by a given positive number.*

Proof. We show how to construct a two-stage stochastic linear program such that, for any partitioning r_1, \dots, r_U of a given integer U , the optimal second-stage value function $Q(x, y)$ is exactly the function $\varphi_{r_1, \dots, r_U}$ described in the proof of Thm. 3.1. At the same time, the input size is $O(\log U)$. By Cor. 3.2, this implies that there cannot be a PTAS for the problem.

Let the first stage of the problem be the following problem:

$$\begin{aligned} \min \quad & c_1 x + c_2 y + Q(x, y) \\ & B \begin{pmatrix} x \\ y \end{pmatrix} = \vec{b} \\ & x, y \in [1, U], \end{aligned}$$

where $Q(x, y)$ is the second-stage value function, B is a suitable constraint matrix and \vec{b} is a corresponding r.h.s. vector. Define the second stage problem as:

$$\begin{aligned} Q(x, y) := \min \quad & z(D) + A \\ \text{s.t.} \quad & w(D) \geq y - D \\ & z(D) \geq U\mathbb{E}_D[w(D)] + x - U \\ & w(D) \geq 0 \\ & z(D) \geq 0. \end{aligned} \tag{14}$$

This is a stochastic linear program with an expected-value constraint. In (14), $A > 0$ is the minimum of the value function.

For a given partition r_1, \dots, r_U of the integer U , the slopes of $\varphi_{r_1, \dots, r_U}$ along its two dimensions are given in (12) and (13). Notice that the x -axis slope is always 0 or 1. Let $b_1 = 1, \dots, b_m \in [1, \dots, U - 1]$ be the integer values at which the slope of $\varphi_{r_1, \dots, r_U}(U, y)$ changes along the y axis. We define $b_{m+1} = U$ for notational convenience. For $i = 1, \dots, m$, let Δ_i be the y -slope of $\varphi_{r_1, \dots, r_U}(U, y)$ in $[b_i, b_{i+1}]$, which is constant by definition of b_i . We define $\Delta_0 = 0$ for notational convenience. Let D be a discrete r.v. with support $\{b_1, \dots, b_{m+1}\}$ and $\Pr(D = b_i) = \frac{\Delta_i - \Delta_{i-1}}{U}$ for $i = 1, \dots, m$, $\Pr(D = b_{m+1}) = \frac{U - \Delta_m}{U}$. Notice that problem (14) can be described in $O(\log U)$ bits.

With this construction, the optimal value function $Q(x, y)$ is the following:

$$\begin{aligned} Q(x, y) &= \min \{z(D) + A : w(D) \geq y - D, z(D) \geq U\mathbb{E}_D[w(D)] + x - U, w(D) \geq 0, z(D) \geq 0\} \\ &= \min \{(U\mathbb{E}_D[w(D)] + x - U)^+ + A : w(D) \geq y - D, w(D) \geq 0\} \\ &= \left(U \sum_{i=1}^{m+1} \Pr(D = b_i)(y - b_i)^+ + x - U \right)^+ + A \\ &= \left(\sum_{i=1}^m (\Delta_i - \Delta_{i-1})(y - b_i)^+ + x - U \right)^+ + A. \end{aligned}$$

In the above expression, the last equality is due to the fact that for $i = m + 1$ we have $(y - b_i)^+ = 0$, so the last term of the summation vanishes. We claim that $Q(x, y) = \varphi_{r_1, \dots, r_U}(x, y)$. Clearly $Q(x, y)$ is piecewise

```

1: Let  $A' \leftarrow \frac{100\kappa_\varphi}{\epsilon^2\varphi^{\min}}A$ ,  $B' \leftarrow \frac{100\kappa_\varphi}{\epsilon^2\varphi^{\min}}B$ .
2: Define  $\rho(x) := \frac{10\varphi(\frac{\epsilon^2\varphi^{\min}}{100\kappa_\varphi}x)}{\epsilon\varphi^{\min}}$ 
3:  $W' \leftarrow \{(A', \lceil \rho(A') \rceil)\}$ 
4:  $x \leftarrow A'$ 
5: while  $(1 + \epsilon)\rho(x) < \rho(B')$  do
6:   Find a point  $y' \in [A', B']$  such that  $(1 + \frac{3}{10}\epsilon)\rho(x) \leq \rho(y') \leq (1 + \frac{4}{10}\epsilon)\rho(x)$ 
7:   Find a point  $y'' \in [A', B']$  such that  $(1 + \frac{6}{10}\epsilon)\rho(x) \leq \rho(y'') \leq (1 + \frac{7}{10}\epsilon)\rho(x)$ 
8:   Find an integer point  $y \in [y', y'']$ 
9:    $W' \leftarrow W' \cup \{(y, \lceil \rho(y) \rceil)\}$ 
10:   $x \leftarrow y$ 
11:  $W' \leftarrow W' \cup \{(B', \lceil \rho(B') \rceil)\}$ 
12:  $W \leftarrow \left\{ \left( \frac{\epsilon^2\varphi^{\min}}{100\kappa_\varphi}y, \frac{\epsilon\varphi^{\min}}{10}v^* \right) : (y, v^*) \in W' \right\}$ 
13: return  $W$  as an array of tuples sorted by their first coordinate

```

Algorithm 1: Function SCALEDCOMPRESSCONVINC($\varphi, [A, B], K = 1 + \epsilon$).

linear convex, nondecreasing, and its range is $[A, A + U]$ over the domain $[1, U]^2$. Since its slope along the x axis is 1 whenever the function value is $> A$, we just need to show that $Q(U, y) = \varphi_{r_1, \dots, r_U}(U, y)$: equality over the rest of the domain then follows from the analysis carried out in the proof of Thm. 3.1. We show by induction on $k = 1, \dots, m + 1$ that $Q(U, y) = \varphi_{r_1, \dots, r_U}(U, y)$ for $y \in [1, b_k]$. For $k = 1$ this is trivial because $Q(U, 1) = A = \varphi_{r_1, \dots, r_U}(U, 1)$. For the induction step, we need to ensure that the y -slope of $Q(U, y)$ over the interval $[b_{k-1}, b_k]$ matches that of $\varphi_{r_1, \dots, r_U}(U, y)$. The terms in the summation $\sum_{i=1}^{m+1} (\Delta_i - \Delta_{i-1})(y - b_i)^+ + x - U$ for $i \geq k$ are clearly 0, whereas for $i < k$ the slope is $\Delta_i - \Delta_{i-1}$. So, the overall slope is $\sum_{i=1}^{k-1} (\Delta_i - \Delta_{i-1}) = \Delta_{k-1}$, as desired. This concludes the proof. \square

We remark that the stochastic LP (14) does not satisfy the conditions of our DP framework because of the term $(U\mathbb{E}_D[(y - D)^+] + x - U)^+$ in the objective function, which is unusual for DP. More specifically, Condition 3 allows a cost function g_t of the form $(U(y - D)^+ + x - U)^+$, and our DP model (as is customary in DP) would take the expectation of the whole expression: $\mathbb{E}_D \left[(U(y - D)^+ + x - U)^+ \right]$. In other words, it is unclear how to formulate as a DP the cost $(U\mathbb{E}_D[(y - D)^+] + x - U)^+$, where a piecewise linear convex function is applied to the expectation. In any case, Thm. 3.4 is indicative of the difficulties faced when dealing with two-dimensional value functions.

4 Constructing approximation sets with bounded-size numbers

The FPTAS that we are about to build (Sect. 6) relies on solving several LPs to compute an approximation of the value function at each stage $t = T, \dots, 1$. These LPs have the property that the input to the LPs at stage t depends on data computed by the LPs at stage $t + 1$ — the data are the domain and codomain values of a K -approximation set for the value function. It is well known that the size of the output of an LP (i.e., the optimal solution and its value) is, in the worst case, superlinear in the size of the input data, although still polynomial. In other words, every time an LP is solved, the size of the output numbers may increase polynomially compared to the size of the input numbers. We therefore need a device to limit the growth of the number sizes, to ensure that they do not take an exponential number of bits. As will become clearer in Sect. 6, to limit the growth it is sufficient to have the ability to compute K -approximation sets that involve numbers that are “not too large”. The next result shows that this can be done efficiently (recall from Sect. 2.1 that t_φ denotes an upper bound on the time needed to evaluate φ on a single point in its domain).

Proposition 4.1 *Let $\varphi : [A, B] \rightarrow \mathbb{R}^+$ with $A, B \in \mathbb{Z}$ be a positive convex nondecreasing function with Lipschitz constant κ_φ . Then for any $K := 1 + \epsilon$, Alg. 1 returns a canonical representation of a K -approximation set of φ of size $O(\log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ consisting of points that have domain and codomain values representable in $O(\log \frac{1}{\epsilon} + \log \frac{\varphi^{\max}}{\varphi^{\min}} + \log \kappa_\varphi + \log |A| + \log |B|)$ space. The algorithm runs in $O(t_\varphi \log \frac{\kappa_\varphi(B-A)}{\epsilon \varphi^{\min}} \log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ time.*

Proof. To simplify the exposition, we assume $\epsilon \leq 1/4$. The idea for the proof is to construct an appropriately scaled version of φ that admits a fully integer K -approximation set, i.e. with integer domain and codomain values. After constructing such an integer K -approximation set, we apply the inverse scaling to obtain a K -approximation set for the original function, and show that the size of the points is bounded. The construction that we follow is summarized in Alg. 1, called SCALEDCOMPRESSCONVINC. The function ρ defined therein (step 2) is the scaled version of φ mentioned above.

We claim that a canonical representation W' of a K -approximation set for ρ that is fully integer (except at its endpoints A', B') is constructed in time $O(t_\varphi \log \frac{\kappa_\varphi(B-A)}{\epsilon \varphi^{\min}} \log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ by routine SCALEDCOMPRESSCONVINC. We start by showing that in every iteration of the “while” loop at lines 5-10, we find a point $y \in [x, B']$ and a value v^* such that: (i) $\rho(y) \geq (1 + \frac{3}{10}\epsilon)\rho(x)$, (ii) $\rho(x) \leq v^* \leq (1 + \epsilon)\rho(x)$, and (iii) $\rho(x) \leq \rho(w) \leq (1 + \epsilon)\rho(x)$ for all $w \in [x, y]$.

Because φ is nondecreasing, lines 6 and 7 of SCALEDCOMPRESSCONVINC can be executed by binary search. We are looking for a point within an interval of size at least 1 starting with an interval of size $\frac{100\kappa_\varphi}{\epsilon^2 \varphi^{\min}}(B - A)$, hence the running time for each of the two steps is $O(t_\varphi \log \frac{\kappa_\varphi(B-A)}{\epsilon \varphi^{\min}})$. Line 8 can be performed in constant time by setting $y = \lceil y' \rceil$, as we will show next. By construction, $\rho(y'') - \rho(y') \geq \frac{2}{10}\epsilon\rho(x)$. Furthermore, $\kappa_\rho \leq \frac{\epsilon}{10}$ by construction and definition of κ_φ , and $\rho(x) \geq \frac{10}{\epsilon} \geq 1$. Therefore, $y'' - y' \geq \frac{2}{10}\epsilon \frac{\rho(x)}{\kappa_\rho} \geq 2$, and there must exist at least two integer points between y' and y'' . The chain $y' \leq \lceil y' \rceil < y''$ implies $(1 + \frac{3}{10}\epsilon)\rho(x) \leq \rho(y') \leq \rho(\lceil y' \rceil) < \rho(y'') \leq (1 + \frac{7}{10}\epsilon)\rho(x)$, showing that $y = \lceil y' \rceil$ satisfies property (i). We claim that $\lceil \rho(y) \rceil$, used in line 9 of SCALEDCOMPRESSCONVINC, is such that $\rho(y) \leq \lceil \rho(y) \rceil \leq \rho(y)(1 + \frac{1}{10}\epsilon)$. Indeed,

$$\frac{\lceil \rho(y) \rceil}{\rho(y)} \leq \frac{1 + \rho(y)}{\rho(y)} \leq 1 + \frac{\epsilon \varphi^{\min}}{10 \varphi(\frac{\epsilon^2 \varphi^{\min}}{100 \kappa_\varphi} y)} \leq 1 + \frac{1}{10}\epsilon,$$

where the last inequality follows by definition of φ^{\min} . Finally, we have:

$$\frac{\lceil \rho(y) \rceil}{\rho(x)} \leq (1 + \frac{1}{10}\epsilon) \frac{\rho(y)}{\rho(x)} \leq (1 + \frac{1}{10}\epsilon)(1 + \frac{7}{10}\epsilon) \leq 1 + \epsilon,$$

showing property (ii). For the last inequality we used the assumption that $\epsilon \leq 1/4$; larger ϵ can be handled adjusting lines 6-7 of the algorithm. Property (iii) now follows immediately because ρ is convex nondecreasing.

By our discussion above, all (domain) points in W' except the two endpoints A' and B' have integer values. By properties (i), (ii) and (iii), the set W' is a canonical representation of a K -approximation set for ρ (the ceiling function at the endpoints A', B' adds only $\frac{1}{10}\epsilon$ relative error, by the same argument used in showing property (ii)). Clearly, the “while” loop at lines 5-10 is executed at most $O(\log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ times because the value of $\rho(x)$ increases by at least a bounded below fraction of the multiplication factor K . Thus, W' contains $O(\log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ points and can be constructed in $O(t_\varphi \log \frac{\kappa_\varphi(B-A)}{\epsilon \varphi^{\min}} \log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ time.

The final step in SCALEDCOMPRESSCONVINC consists of applying to W' the inverse scaling that transforms φ into ρ , thus obtaining W . It follows that the points in W belong to the domain and codomain of φ . The error bound, i.e. the K -approximation property, is trivially satisfied because it is satisfied for ρ

1: $x^* \leftarrow \arg \min_{x \in [A, B]} \varphi(x)$ /* We assume that a minimum of φ can be found efficiently */
2: $W_A \leftarrow \text{SCALEDCOMPRESSCONVDEC}(\varphi, [A, x^*], K)$
3: $W_B \leftarrow \text{SCALEDCOMPRESSCONVINC}(\varphi, [x^*, B], K)$
4: $W \leftarrow (W_A \cup W_B) \setminus \{(x^*, \lceil \varphi(x^*) \rceil)\}$
5: **return** W as an array of tuples sorted by their first coordinate

Algorithm 2: Function $\text{SCALEDCOMPRESSCONV}(\varphi, [A, B], K = 1 + \epsilon)$.

and the relative error is invariant to multiplicative scaling. To conclude the proof, note that the size of the numbers in W' , both domain and codomain, is $O(\log \frac{1}{\epsilon} + \log \frac{\varphi^{\max}}{\varphi^{\min}} + \log \kappa_\varphi + \log |A| + \log |B|)$ because the domain is $[A', B']$ and the codomain is $[\frac{10}{\epsilon}, \frac{10\varphi^{\max}}{\epsilon\varphi^{\min}}]$; the multiplication factor when rescaling W' to W is of the same size, which gives the desired bound. \square

We define Function $\text{SCALEDCOMPRESSCONVDEC}$ for positive convex nonincreasing Lipschitz continuous functions in a similar way. We define Function $\text{SCALEDCOMPRESSCONV}$ for (not necessarily monotone) convex Lipschitz continuous functions as described in Alg. 2. The algorithm assumes that a minimum of φ can be found exactly; in the FPTAS, this will be the case, as the minima will be computable by solving an LP.

Theorem 4.2 *Let $\varphi : [A, B] \rightarrow \mathbb{R}^+$ with $A, B \in \mathbb{Z}$ be a positive convex function with Lipschitz constant κ_φ , and assume that $x^* \in \arg \min_{x \in [A, B]} \varphi(x)$ can be efficiently computed. Then for any $K := 1 + \epsilon$, Alg. 2 returns a canonical representation of a K -approximation set of φ of size $O(\log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ consisting of points that have domain and codomain values representable in $O(\log \frac{1}{\epsilon} + \log \frac{\varphi^{\max}}{\varphi^{\min}} + \log \kappa_\varphi + \log |A| + \log |B|)$ space. The algorithm runs in $O(t_\varphi \log \frac{\kappa_\varphi(B-A)}{\epsilon\varphi^{\min}} \log_K \frac{\varphi^{\max}}{\varphi^{\min}})$ time.*

Proof. Let $\hat{\varphi}$ be the convex extension of φ induced by W . Let $\ell := \max\{x \in W : x \leq x^*\}, r := \min\{x \in W : x \geq x^*\}$. Clearly $\varphi(x) \leq \hat{\varphi}(x) \leq K\varphi(x)$ for $x \in [A, \ell] \cup [r, B]$ because W_A and W_B are K -approximation sets over the respective domains. By construction (see the proof of Prop. 4.1), we have $\varphi(\ell) \leq K\varphi(x^*)$ and $\varphi(r) \leq K\varphi(x^*)$. But then the convex extension $\hat{\varphi}$ induced by W is such that $\varphi(x) \leq \hat{\varphi}(x) \leq K\varphi(x)$ also for $x \in [\ell, r]$. The size of the numbers follows directly from Prop. 4.1. \square

We remark that the main reason for discarding $(x^*, \lceil \varphi(x^*) \rceil)$ from W in Alg. 2 is to ensure that the size of the numbers remains bounded as indicated by Prop. 4.1; in a practical implementation of the routine, it is recommended to keep in the canonical representation some bounded-size rounding of $(x^*, \lceil \varphi(x^*) \rceil)$.

5 Approximating expectations with vectors of continuous random variables

To construct an approximation of the value function (1), we must be able to compute expected values efficiently using the given representation of the r.v.s. Under Condition 2, we can compute a K -approximation set for the CDF of each of the components of the vectors of r.v.s. More precisely, in case (i) of Condition 2 we can use the algorithm given in [HN16, Sec. 8.2], and in case (ii) we can use the algorithm given in [HKL⁺14, Sec. 4.1] (see Apx. A). Within a DP setting, the expression of the value function contains the expectation of a composition of functions, and we need to find a way to efficiently compute this expectation. Under Conditions 2 and 3, at each stage the vector of random variables \vec{D}_t is mapped to a scalar r.v. via a linear transformation, i.e., $\vec{\theta}^D$ or $\vec{\sigma}^D$. However, we do not have explicit access to the CDF of the corresponding scalar r.v., but only to the CDF of each component of the sum. We now show how to efficiently construct an approximate oracle for the CDF of a sum of r.v.s, using only oracles to their CDFs. We focus on the case of an unweighted sum of continuous r.v.s, remarking that a similar construction for a weighted

sum of continuous or discrete r.v.s is straightforward. At the same time, our approach can be applied to the case where instead of having exact access to the CDFs of the r.v.s, one has an FPTAS to it. Below, by “generalized PDF” we mean a probability distribution function (PDF) that may include the Dirac delta function δ in its definition. Recall that the Dirac delta function can be informally thought of as a function such that $\delta(0) = \infty$, $\delta(x) = 0 \forall x \neq 0$ and $\int_{-\infty}^{+\infty} \delta(x) dx = 1$; a formal definition is out of the scope of this paper.

Proposition 5.1 *For $i = 1, \dots, \ell$, let X_i be a truncated continuous r.v. with support $[X_i^{\min}, X_i^{\max}]$. Suppose X_i admits a (generalized) PDF f_i and κ -Lipschitz continuous CDF F_i . Assume $\gamma = \min_{i=1, \dots, \ell} \{\Pr(X_i = X_i^{\min}), \Pr(X_i = X_i^{\max})\} > 0$. Denote $X^{\min} = \min_{i=1, \dots, \ell} X_i^{\min}$, $X^{\max} = \max_{i=1, \dots, \ell} X_i^{\max}$, and let $t_F = \max_{i=1, \dots, \ell} t_{F_i}$. If X_1, \dots, X_ℓ are independent, then for any $K = 1 + \epsilon$, we can compute a K -approximation set with $O(\frac{\ell}{\epsilon} \log \frac{1}{\gamma})$ points for the CDF of $\sum_{i=1}^{\ell} X_i$, in $O(\frac{t_F \ell^2}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa \ell (X^{\max} - X^{\min})))$ time.*

The proof technique is based on building an approximation set of the sum of r.v.s adding one variable at a time: assuming we have an approximation set W for the CDF of $\sum_{i=1}^{k-1} X_i$, we show that an approximation for the CDF of $\sum_{i=1}^k X_i$ can be constructed iterating over the points in W .

Proof. Let $\bar{K} := \sqrt[\ell]{1 + \epsilon}$. We proceed by induction on $k = 1, \dots, \ell$, computing a \bar{K}^k -approximation set for the CDF of $\sum_{i=1}^k X_i$ that contains $O(\frac{\ell}{\epsilon} \log \frac{1}{\gamma})$ points and requires $O(\frac{t_F \ell^2}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa k (X^{\max} - X^{\min})))$ time at each induction step. In these expressions, the factor $\frac{\ell}{\epsilon}$ comes from the fact that $\log_{\bar{K}}(\cdot)$ is $O(\frac{\ell}{\epsilon} \log(\cdot))$ because $K = \sqrt[\ell]{1 + \epsilon}$. For $k = 1$, we simply apply COMPRESSINC, defined in [HKL⁺14] if D_t is a discrete random variable, and in [HN16] if D_t is a continuous random variable, see Apx. A. This yields a \bar{K} -approximation set W_1 of X_1 with $O(\frac{\ell}{\epsilon} \log \frac{1}{\gamma})$ points in $O(\frac{t_F \ell^2}{\epsilon^2} \log \frac{1}{\gamma} \log(\kappa (X^{\max} - X^{\min})))$ time.

Denote $Y = \sum_{i=1}^{k-1} X_i$. Now assume that we have a \bar{K}^{k-1} -approximation set $W_{k-1} := \{b_1, \dots, b_m\}$ for the CDF F_Y of Y , and let \tilde{F}_Y be the corresponding monotone extension. Recall that \tilde{F}_Y can be seen as a summation of step functions. Let $h(x)$ be the step function $h(x) = 1$ if $x \geq 0$, $h(x) = 0$ otherwise. Using this notation, we can express $\tilde{F}_Y(x) = \sum_{i=1}^m a_i h(x - b_i)$ for some $a_i \in \mathbb{R}^+$.

We want to compute $\varphi(z) := \Pr(X_k + Y \leq z)$ for any given z . Because X_k and Y are independent, we can write:

$$\begin{aligned} \int_{\mathbb{R}} \int_{-\infty}^{z-x} f_{X_k, Y}(x, y) dy dx &= \int_{\mathbb{R}} F_Y(z-x) f_k(x) dx \leq \\ \int_{\mathbb{R}} \tilde{F}_Y(z-x) f_k(x) dx &\leq \bar{K}^{k-1} \int_{\mathbb{R}} F_Y(z-x) f_k(x) dx, \end{aligned}$$

where the inequalities follow from the facts that \tilde{F}_Y is a \bar{K}^{k-1} -approximation of F_Y , and f_k is nonnegative. This shows that $\tilde{\varphi}(z) := \int_{\mathbb{R}} \tilde{F}_Y(z-x) f_k(x) dx$ is a \bar{K}^{k-1} -approximation of $\varphi(z)$, i.e. the CDF of $X_k + Y$. Breaking down \tilde{F}_Y into a summation of step functions, we can expand this integral as:

$$\int_{\mathbb{R}} \sum_{i=1}^m a_i h(z-x-b_i) f_k(x) dx = \sum_{i=1}^m a_i \int_{-\infty}^{z-b_i} f_k(x) dx = \sum_{i=1}^m a_i \Pr(X_k \leq z-b_i).$$

It follows that $\tilde{\varphi}(z)$ can be computed in time $O(mt_{F_k})$. By the induction hypothesis, this is $O(\frac{t_F \ell}{\epsilon} \log \frac{1}{\gamma})$. We now apply COMPRESSINC to $\tilde{\varphi}$ with approximation factor \bar{K} . By approximation of approximation (Prop. 2.5(7)), this yields an approximation set W_k that induces a \bar{K}^k -approximation function of φ . W_k has $O(\frac{\ell}{\epsilon} \log \frac{1}{\gamma})$ points, and the computation takes $O(\frac{t_F \ell^2}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa k (X^{\max} - X^{\min})))$ time. This concludes the inductive claim.

```

1:  $\bar{K} \leftarrow \sqrt[\ell]{K}$ 
2:  $W_1 \leftarrow \text{COMPRESSINC}(F_1(\frac{\cdot}{c_1}), [X_1^{\min}, X_1^{\max}], \bar{K})$  /*  $F_1(\cdot)$  is the CDF of  $X_1$  */
3: for  $k = 2, \dots, \ell$  do
4:   Denote  $\{b_1, \dots, b_m\}$  the points in the approximation set  $W_{k-1}$ 
5:   Let  $a_1, \dots, a_m$  be the coefficients of the monotone extension  $\tilde{F}_{k-1}(x) = \sum_{i=1}^m a_i h(x - b_i)$  of  $W_{k-1}$ 
   /*  $h(x)$  is the unit step function:  $h(x) = 1$  for  $x \geq 0$ , 0 otherwise */
   /*  $\tilde{F}_{k-1}(\cdot)$  is a  $\bar{K}^{k-1}$ -approximation of the CDF of  $\sum_{j=1}^{k-1} c_j X_j$  */
6:   Define  $\tilde{\varphi}(z) := \sum_{i=1}^m a_i F_k(\frac{z-b_i}{c_k})$  /*  $F_k(\cdot)$  is the CDF of  $X_k$  */
7:    $W_k \leftarrow \text{COMPRESSINC}(\tilde{\varphi}, [\sum_{j=1}^k X_j^{\min}, \sum_{j=1}^k X_j^{\max}], \bar{K})$ 
8: return  $W_\ell$  as an array of tuples sorted by their first coordinate

```

Algorithm 3: Function COMPRESSCONVOLUTION($(X_1, \dots, X_\ell), (c_1, \dots, c_\ell), K$).

To conclude the proof, we note that the loop discussed above is repeated for $k = 1, \dots, \ell$, therefore the total runtime is $O(\frac{t_F \ell^3}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa \ell (X^{\max} - X^{\min})))$. \square

Prop. 5.1 is given for truncated continuous r.v.s, but it is straightforward to note that the same approach works for discrete r.v.s. In this case, instead of the PDF of the r.v.s we simply employ the corresponding PMF, obtaining the same result. Furthermore, we note that the proof can be adapted to the case in which only approximate access to the CDF of the r.v.s is available. Indeed, suppose we have an FPTAS for the CDF of each the r.v.s. In this case, we simply need to decrease the approximation factor \bar{K} used in the proof to, say, $\bar{K} = \sqrt[2\ell]{1 + \epsilon}$, and replace each call to the oracles for the CDF to a $\sqrt[2\ell]{1 + \epsilon}$ -approximate call to the FPTAS for the CDF. In this respect we improve upon [LS14]. The running time given by [LS14] to compute $\Pr(\sum_{i=1}^{\ell} X_i \leq \rho)$ for discrete integer r.v.s is $O(\frac{t_F \ell^3}{\epsilon^2} \log^2 \frac{1}{\gamma} \log \rho)$, essentially the same as ours if we restrict our algorithm to the discrete case: we have $\log \ell (X^{\max} - X^{\min})$ rather than $\log \rho$ because we compute the CDF for all values of ρ simultaneously.

We define in Alg. 3 a routine COMPRESSCONVOLUTION(\vec{X}, \vec{c}, K) that implements the constructive proof of Prop. 5.1. The routine takes as input a vector of random variables \vec{X} and returns a canonical representation of a K -approximation oracle function for $\varphi(z) := \Pr(\vec{c} \cdot \vec{X} \leq z)$. The query time of the oracle is $O(\log(\frac{\ell}{\epsilon} \log \frac{1}{\gamma}))$.

We can now assume that we have approximate access to the CDF of the random variables defined by $\vec{\theta}^D \cdot \vec{D}_t$ and $\vec{\sigma}^D \cdot \vec{D}_t$ that appear in the cost functions, as detailed in Condition 3. We then need to compute the expectation of a composition of functions using the approximate CDF. If $\vec{\theta}^D \cdot \vec{D}_t$ and $\vec{\sigma}^D \cdot \vec{D}_t$ have a discrete distribution, i.e. case (ii) of Condition 2, we can rely on [Hal17, Prop. 6.8]. The rest of this section details how to efficiently compute expectations involving continuous r.v.s, thereby taking care of case (i) of Condition 2.

Our next result shows that we can efficiently compute an expected value of a composition of functions involving a continuous r.v. Essentially, the proposition is the analog of [Hal17, Prop. 6.8] for continuous r.v.s, but the proof technique has some significant differences, because for continuous r.v.s we do not know how to write an easy-to-evaluate closed-form expression for the expected value.

Proposition 5.2 *Let $\xi : [A, B] \rightarrow \mathbb{R}^+$ be a (not necessarily monotone) convex function. Let $K_1, K_2 \geq 1$. Let $\psi : [A, B] \rightarrow \mathbb{R}^+$ be an increasing piecewise linear convex function with breakpoints $A = a_1 < \dots < a_n < B$ and slopes $0 = \Delta_0 \leq \Delta_1 < \dots < \Delta_n$ that K_1 -approximates ξ . Let D be a (not necessarily nonnegative) truncated continuous r.v. with support $[D^{\min}, D^{\max}]$ and such that $\min\{\Pr(D = D^{\min}), \Pr(D = D^{\max})\} > 0$. Suppose D admits a (generalized) PDF F' , and let F be the corresponding CDF. Let $\tilde{F}(\cdot)$ be a monotone K_2 -approximation of F with breakpoints at $D^{\min} = d_1 < \dots < d_m = D^{\max}$, and let $\tilde{F}(d_0) = 0$. Let $f(x, d) = bx + e - d$ for some given $b, e \in \mathbb{R}$, and $m_i(x) = \max\{j \mid d_j \leq$*

$bx + e - a_i\}$. Then

$$\tilde{\xi}(x) = \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \left((bx + e - d_{m_i(x)} - a_i) \tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k) \tilde{F}(d_{k+1}) \right) \quad (15)$$

is a convex piecewise linear $K_1 K_2$ -approximation function of $\mathbb{E}_D(\xi(f(x, D)))$ with $O(mn)$ pieces. Furthermore, one can construct in $O(mn \log n)$ time an oracle for $\tilde{\xi}(\cdot)$ in the form of either a canonical representation, or a representation that consists of breakpoints and slopes, each of size $O(mn)$.

The proof technique involves decomposing $\mathbb{E}_{D_t}(\xi(f(x, D)))$ as a finite sum involving \tilde{F} for each piece of the piecewise linear function φ . The error bound is obtained by constructing “easy-to-handle” discrete random variables \hat{D}, \check{D} such that $\hat{D} \preceq D_t \preceq \check{D}$ in the usual stochastic order.

Proof. For $i = 0, \dots, n$, define:

$$\psi_i(y) := \begin{cases} (\Delta_i - \Delta_{i-1})(y - a_i) & \text{if } y \geq a_i, \\ 0 & \text{otherwise.} \end{cases}$$

Because ψ is piecewise linear, for all $y \in [A, B]$ we have $\psi(y) = \psi(A) + \sum_{i=1}^n \psi_i(y)$. We have:

$$\begin{aligned} \mathbb{E}_D(\xi(f(x, D))) &\leq \mathbb{E}_D(\psi(f(x, D))) \\ &= \int_{d_1}^{d_m} \psi(f(x, d)) F'(d) dd \\ &= \psi(A) + \int_{d_1}^{d_m} \sum_{i=1}^n \psi_i(f(x, d)) F'(d) dd \\ &= \psi(A) + \sum_{i=1}^n \left(\int_{d_1}^{d_m} \psi_i(f(x, d)) F'(d) dd \right) \\ &= \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \int_{d_1}^{d_m} \max\{0, (bx + e - d - a_i)\} F'(d) dd \\ &= \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i) F'(d) dd, \end{aligned} \quad (16)$$

where the first inequality is due to ψ being a K_1 -approximation function of ξ , and the rest are algebraic manipulations.

We construct a discrete r.v. \hat{D} that takes values d_1, \dots, d_{m-1} with:

$$\Pr(\hat{D} = d_j) = \begin{cases} F(d_2) & \text{if } j = 1 \\ F(d_{j+1}) - F(d_j) & \text{if } j = 2, \dots, m-1. \end{cases}$$

It follows that the CDF \hat{F} of \hat{D} is $\hat{F}(d) = \max\{F(d_{j+1}) : d_j \leq d, j = 1, \dots, m-1\}$. In order to compute expected values of continuous functions of \hat{D} using the classical integration approach, we define the generalized PDF of \hat{D} as follows:

$$\hat{F}'(d) := \delta(d - d_1)F(d_2) + \sum_{j=2}^{m-1} \delta(d - d_j)(F(d_{j+1}) - F(d_j)).$$

Notice that $\hat{D} \preceq D$ in the usual stochastic order, because $\Pr(\hat{D} > d) \leq \Pr(D > d)$ for all d . Since $\hat{D} \preceq D$ and $(bx + e - d - a_i)$ is a decreasing function in d , it follows that:

$$\int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i) F'(d) dd \leq \int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i) \hat{F}'(d) dd =$$

$$\begin{aligned}
& (bx + e - d_1 - a_i)F(d_2) + \sum_{j=2}^{m_i(x)} (bx + e - d_j - a_i)(F(d_{j+1}) - F(d_j)) = \\
& (bx + e - d_1 - a_i)\tilde{F}(d_2) + \sum_{j=2}^{m_i(x)} (bx + e - d_j - a_i)(\tilde{F}(d_{j+1}) - \tilde{F}(d_j)) = \\
& (bx + e - d_{m_i(x)} - a_i)\tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k)\tilde{F}(d_{k+1}).
\end{aligned}$$

Putting everything together in (16), we obtain:

$$\mathbb{E}_D(\xi(f(x, D))) \leq \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \left((bx + e - d_{m_i(x)} - a_i)\tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k)\tilde{F}(d_{k+1}) \right).$$

Because $\psi(x) \leq K_1\xi(x)$ for all x and F' is nonnegative, we can write:

$$\begin{aligned}
\mathbb{E}_D(\psi(f(x, D))) &= \int_{d_1}^{d_m} \psi(f(x, d))F'(d) dd \\
&\leq K_1 \int_{d_1}^{d_m} \xi(f(x, d))F'(d) dd = K_1\mathbb{E}_D(\xi(f(x, D))).
\end{aligned} \tag{17}$$

We now construct a discrete r.v. \check{D} that takes values d_1, \dots, d_m with $\Pr(\check{D} = d_j) = F(d_j) - F(d_{j-1}) = \tilde{F}(d_j) - \tilde{F}(d_{j-1})$. It follows that the CDF \check{F} of \check{D} is $\check{F}(d) = \max\{F(d_j) : d_j \leq d, j = 1, \dots, m\}$. As before, we define the generalized PDF of \check{D} as follows:

$$\check{F}'(d) := \sum_{j=1}^m \delta(d - d_j)(F(d_j) - F(d_{j-1})).$$

Notice that $D \preceq \check{D}$ in the usual stochastic order, because $\Pr(D > d) \leq \Pr(\check{D} > d)$ for all d . Since $D \preceq \check{D}$ and $(bx + e - d - a_i)$ is a decreasing function in d , it follows that:

$$\begin{aligned}
& \int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i)F'(d) dd \geq \int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i)\check{F}'(d) dd = \\
& \sum_{j=1}^{m_i(x)} (bx + e - d_j - a_i)(F(d_j) - F(d_{j-1})) = \sum_{j=1}^{m_i(x)} (bx + e - d_j - a_i)(\tilde{F}(d_j) - \tilde{F}(d_{j-1})) = \\
& (bx + e - d_{m_i(x)} - a_i)\tilde{F}(d_{m_i(x)}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k)\tilde{F}(d_k) \geq \\
& \frac{1}{K_2}(bx + e - d_{m_i(x)} - a_i)\tilde{F}(d_{m_i(x)+1}) + \frac{1}{K_2} \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k)\tilde{F}(d_{k+1}),
\end{aligned}$$

where the last inequality follows from the fact that $F(d_{j+1}) \leq K_2F(d_j)$ for $j = 1, \dots, m-1$ by definition of K -approximation set for monotone function. Then we can write:

$$\mathbb{E}_D(\psi(f(x, D))) = \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \int_{d_1}^{\max\{d: bx+e-d \geq a_i\}} (bx + e - d - a_i)F'(d) dd \geq$$

$$\frac{1}{K_2} \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \left((bx + e - d_{m_i(x)} - a_i) \tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k) \tilde{F}(d_{k+1}) \right) \geq \frac{\tilde{\xi}(x)}{K_2}. \quad (18)$$

By combining inequalities (17) and (18) we get the desired approximation ratio.

It is easy to verify that $\tilde{\xi}$ has increasing slopes and is therefore a convex piecewise linear increasing function. To conclude, we discuss how to compute a representation of $\tilde{\xi}$ in terms of breakpoints and slopes. By examining the expression:

$$\tilde{\xi}(x) = \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \left((bx + e - d_{m_i(x)} - a_i) \tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k) \tilde{F}(d_{k+1}) \right),$$

we see that the slope of each term of the summation changes whenever $m_i(x)$ changes. There are at most m such changes for each term, and their location can be computed in $O(m)$ time because for term i the breakpoints are of the form $d_j - e + a_i$ for $j = 1, \dots, m$. We then obtain, in $O(mn)$ time, n sorted lists with m elements each. These lists can be merged in $O(mn \log n)$ time, yielding a superset of the breakpoints of $\xi(x)$. To compute the slopes we only need an additional $O(m)$ time to preprocess the $m - 1$ partial sums $\sum_{k=1}^{m-1} (d_{k+1} - d_k) \tilde{F}(d_{k+1})$, since the value of \tilde{F} at all queried points is known and available in the approximation set that induces \tilde{F} . The overall time requirement is therefore $O(mn \log n)$. \square

As in [Hal17], it is easy to see that Prop. 5.2 can be applied to separable linear transition functions $f(x, d) = bx + e + cd$ for any coefficient c , using a transformed r.v. D' with $D' := D/(-c)$, and a similar result holds for convex functions approximated by a decreasing piecewise linear convex function. As a consequence, we obtain a modified version of [Hal17, Cor. 6.9], stated below.

Proposition 5.3 *Let $\xi : [A, B] \rightarrow \mathbb{R}^+$ be a convex function. Let $K_1, K_2 \geq 1$. Let $\psi : [A, B] \rightarrow \mathbb{R}^+$ be a piecewise linear convex function with n breakpoints that K_1 -approximates ξ over $[A, B]$. Let D be a (not necessarily nonnegative) continuous r.v. satisfying the conditions stated in Prop. 5.2. Let $\tilde{F}(\cdot)$ be a monotone nondecreasing K_2 -approximation of the CDF of D with m breakpoints. Let $f(x, D) = bx + e + cd$, with $b, e, c \in \mathbb{R}$. Then we can construct in $O(mn \log n)$ time a convex piecewise linear $K_1 K_2$ -approximation function of $\mathbb{E}_D(\xi(f(x, D)))$ in the form of either a canonical representation, or a representation that consists of breakpoints and slopes, each of size $O(mn)$.*

The proof of Prop. 5.3 follows from an application of Prop. 5.2 twice, to the nondecreasing and nonincreasing part of ψ . In Alg. 4 we define a routine COMPRESSEXPVAL($\psi, (b, e, c), \tilde{F}$) that takes as input ψ, b, e, c, \tilde{F} as defined in Prop. 5.3, and returns an oracle for a $K_1 K_2$ -approximation function of $\mathbb{E}_D(\xi(f(\cdot, D)))$ in the form of a canonical representation. The query time for the oracle is $O(\log mn)$. COMPRESSEXPVAL uses one subroutine for the increasing part of the convex function ψ , called COMPRESSEXPVALINC and defined in Alg. 5, and one for the decreasing part, called COMPRESSEXPVALDEC. The pseudocode for COMPRESSEXPVALDEC is the same as COMPRESSEXPVALINC after “mirroring” the function ψ given as the first argument to transform it from decreasing to increasing.

6 The approximation scheme

We now show how to approximate a DP that satisfies Conditions 1-3. We give a full proof of the correctness of the algorithm for the simpler Condition 3(i), according to which the cost functions are known explicitly.

- 1: Let $[A, B]$ be the domain of ψ and x^* one of its minimizers
- 2: Define ψ_1 as the restriction of ψ to $[A, x^*]$, ψ_2 as the restriction of ψ to $[x^*, B]$
- 3: $W_1 \leftarrow \text{COMPRESSEXPVALDEC}(\psi_1, (b, e, c), \tilde{F})$
- 4: $W_2 \leftarrow \text{COMPRESSEXPVALINC}(\psi_2, (b, e, c), \tilde{F})$
- 5: **return** $W_1 \cup W_2$ as an array of tuples sorted by their first coordinate

Algorithm 4: Function $\text{COMPRESSEXPVAL}(\psi, (b, e, c), \tilde{F})$.

- 1: Using the approximation set for r.v. D inducing \tilde{F} , compute an approximation set $W_{D'}$ for the transformed random variable $D' = -D/c$
- 2: Let \tilde{F}' be the monotone extension of $W_{D'}$, d_1, \dots, d_m be the breakpoints of $W_{D'}$
- 3: Let a_1, \dots, a_n be the breakpoints of ψ , $\Delta_0, \dots, \Delta_n$ its slopes
- 4: $S \leftarrow \text{sort}(\bigcup_{i=1}^n \{d_j - e + a_i : j = 1, \dots, m\})$
- 5: Compute the $m - 1$ partial sums $\sum_{k=1}^j (d_{k+1} - d_k) \tilde{F}'(d_{k+1})$ for $j = 1, \dots, m - 1$
- 6: $W \leftarrow \{(x, \tilde{\xi}(x)) : x \in S\}$, where $\tilde{\xi}$ is defined in (15)
- 7: Loop over the elements of W to eliminate redundant breakpoints where the slope does not change
- 8: **return** W as an array of tuples sorted by their first coordinate

Algorithm 5: Function $\text{COMPRESSEXPVALINC}(\psi, (b, e, c), \tilde{F})$.

The FPTAS under Condition 3(ii) is similar, and its proof only highlights the differences from Condition 3(i).

6.1 Condition 3(i): explicit functions

Recall that by Condition 3(i) we have $g_t(I_t, \vec{x}_t, \vec{D}_t) = g_t^I(I_t, \vec{x}_t) + g_t^D(f_t^g(I_t, \vec{x}_t, \vec{D}_t))$; to avoid additional notation, we will not specify the expression of g_t^I and g_t^D in terms of breakpoints and slopes, but we will exploit their structure to construct an FPTAS. (Hereafter, given any function f we use the notation $f(\cdot)$ where “ \cdot ” stands for the varying argument of f .)

Proposition 6.1 *Suppose the DP formulation (DP) satisfies Conditions 1-3. Let $K_1, K_2, K_3 > 1$. Let \hat{z}_{t+1} be a piecewise linear convex K_1 -approximation of the value function z_{t+1} . For a given $I_t \in S_t$, let $\tilde{Z}_{t+1}(\cdot)$ be a piecewise linear convex K_2 -approximation of $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \tilde{\theta}^D \cdot \vec{D}_t)]$, and $\tilde{G}_t^D(\cdot)$ be a piecewise linear convex K_3 -approximation of $\mathbb{E}[g_t^D(\cdot + \sigma^I I_t + \tilde{\sigma}^D \cdot \vec{D}_t)]$. Then the value of the following mathematical program:*

$$\left. \begin{aligned} \tilde{z}_t(I_t) := \min_{\vec{x}_t} \quad & g_t^I(I_t, \vec{x}_t) + \tilde{G}_t^D(w) + \tilde{Z}_{t+1}(y) \\ & A_t \vec{x}_t \geq \vec{b}_t + \vec{\delta}_{b_t} I_t \\ & \vec{\sigma}^x \cdot \vec{x}_t - w = 0 \\ & \vec{\theta}^x \cdot \vec{x}_t - y = 0 \\ & \vec{x}_t \geq 0, \end{aligned} \right\} \quad (19)$$

is a $\max\{K_1 K_2, K_3\}$ -approximation of the optimal value function $z_t(I_t)$. (19) can be cast as an LP using the standard min max reformulation of piecewise linear convex functions.

Proof. The expression for the optimal value function is:

$$z_t(I_t) = \min_{\vec{x}_t \in \mathcal{A}_t(I_t)} \mathbb{E}\{g_t(I_t, \vec{x}_t, \vec{D}_t) + z_{t+1}(f_t(I_t, \vec{x}_t, \vec{D}_t))\}.$$

By linearity of the expectation and by definition of g_t^I , \tilde{G}_t^D , \tilde{Z}_{t+1} , f_t , f_t^g and $A_t(I_t)$, it is straightforward to check that (19) corresponds exactly to the expression of the value function, with two of the terms replaced

```

1:  $K \leftarrow \sqrt[2T]{1 + \epsilon}$ ,  $\hat{z}_{T+1} \leftarrow g_{T+1}$ 
2: for  $t := T$  downto 1 do
3:    $\tilde{F}_g \leftarrow \text{COMPRESSCONVOLUTION}(\vec{D}_t, \vec{\sigma}^D, K)$ 
4:    $\tilde{F}_z \leftarrow \text{COMPRESSCONVOLUTION}(\vec{D}_t, \vec{\theta}^D, K)$ 
5:   For fixed  $I_t$ , define  $\tilde{G}_t^D \leftarrow \text{COMPRESSEXVAL}(g_t^D, (1, \sigma^I I_t, 1), \tilde{F}_g)$ 
   /*  $\tilde{G}_t^D(\cdot)$  is an oracle for a  $K$ -approximation of  $\mathbb{E}[g_t^D(\cdot + \sigma^I I_t + \vec{\sigma}^D \cdot \vec{D}_t)]$  */
6:   For fixed  $I_t$ , define  $\tilde{Z}_{t+1} \leftarrow \text{COMPRESSEXVAL}(\hat{z}_{t+1}, (1, \theta^I I_t, 1), \tilde{F}_z)$ 
   /*  $\tilde{Z}_{t+1}(\cdot)$  is an oracle for a  $K$ -approximation of  $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \vec{\theta}^D \cdot \vec{D}_t)]$  */
7:    $\hat{z}_t \leftarrow \text{SCALEDCOMPRESSCONV}(\bar{z}_t, [\mathcal{S}_t^{\min}, \mathcal{S}_t^{\max}], K)$ , where  $\bar{z}_t$  is defined as in (19)
8: return  $\hat{z}_1$ 

```

Algorithm 6: Procedure APXScheme1(ϵ) for Condition 3(i)

by approximations. The fact that the resulting value is a $\max\{K_1 K_2, K_3\}$ approximation of $z_t(I_t)$ follows from Prop. 2.5 (summation of approximation and approximation of approximation). \square

We are now ready to state our main result. The proof consists of an application of the different results already stated in this paper, being careful to keep track of the runtime of the algorithm.

Theorem 6.2 *Given a stochastic DP satisfying Conditions 1-3(i), APXScheme1(ϵ) (Algorithm 6) computes a $(1 + \epsilon)$ -approximation of the optimal value function z_1 , and runs in polynomial time in the binary input size and $1/\epsilon$.*

Proof. The proof proceeds by induction for $t = T + 1, \dots, 1$, showing that at stage t we obtain a piecewise linear convex $K^{2(T+1-t)}$ -approximation \hat{z}_t of the value function z_t via an approximation set of cardinality $O(\log_K \frac{(T+2-t)U_g}{g_{T+1}^{\min}}) = O(\frac{T}{\epsilon} \log \frac{(T+2-t)U_g}{g_{T+1}^{\min}})$ since $K = \sqrt[2T]{1 + \epsilon}$. The argument of the log follows from the fact that the largest value of the value function at stage t is bounded above by $(T + 2 - t)U_g$, and its smallest value is bounded below by g_{T+1}^{\min} . The step $t = T + 1$ is trivial, as we have $\hat{z}_{T+1} = g_{T+1}$, which is exactly the value function at the terminal stage, and we are given its piecewise linear convex description in the input. We now show the induction step.

At step 3, using the routine COMPRESSCONVOLUTION, we compute a K -approximation \tilde{F}_g of the CDF of $\vec{\sigma}^D \cdot \vec{D}_t$ in the form of a canonical representation. By Prop. 5.1, the cardinality of the canonical representation is $O(\frac{T\ell}{\epsilon} \log \frac{1}{\gamma})$, and this step requires $O(\frac{tF T \ell^3}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa \ell n^* U_f))$ time. Step 4 is similar and computes a K -approximation \tilde{F}_z of the CDF of $\vec{\theta}^D \cdot \vec{D}_t$.

At step 5, for a fixed value of I_t we define a K -approximation $\tilde{G}_t^D(\cdot)$ of $\mathbb{E}_{\vec{D}_t}[g_t^D(\cdot + \sigma^I I_t + \vec{\sigma}^D \cdot \vec{D}_t)]$, using Prop. 5.3 with parameters set to $\xi = \psi = g_t^D$, $n = O(m_t)$, $m = O(\frac{T\ell}{\epsilon} \log \frac{1}{\gamma})$, $K_1 = 1$, $K_2 = K$. This is possible because by Condition 3(i) we have a description of g_t^D in terms of breakpoints and slopes, which is equivalent to a 1-approximation set of g_t^D with $O(m_t)$ points. We remark that no actual computation is involved in this step because we did not fix a value of I_t yet (it will be determined later in step 7), but including step 5 in the algorithm helps us for the analysis. By Prop. 5.3, given a value for I_t , computing \tilde{G}_t^D in the form of a canonical representation of size $O(\frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma})$, requires $O(\frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} \log m_t)$ time.

At step 6, for a fixed value of I_t we define a K -approximation $\tilde{Z}_{t+1}(\cdot)$ of $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \vec{\theta}^D \cdot \vec{D}_t)]$, using Prop. 5.3 with parameters set to $\xi = \psi = \hat{z}_{t+1}$, $n = O(\frac{T}{\epsilon} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}})$, $m = O(\frac{T\ell}{\epsilon} \log \frac{1}{\gamma})$, $K_1 = 1$ and $K_2 = K$. Given a value for I_t , computing \tilde{Z}_{t+1} in the form of a canonical representation of size $O(\frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma})$, requires $O(\frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma} \log \frac{T}{\epsilon} \log \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}})$ time.

The main workhorse of the FPTAS is step 7, that computes an approximation \hat{z}_t for the value function z_t , using the function \bar{z}_t as defined in (19). Note that by Prop. 6.1 with parameters set to $K_1 = K^{2(T-t)}$

and $K_2 = K_3 = K$, \bar{z}_t is $K^{2(T-t)+1}$ -approximation of z_t . The routine SCALEDCOMPRESSCONV can be used because a minimum of \bar{z}_t can be found exactly with the solution of a single LP. Applying Thm. 4.2 to \bar{z}_t , coupled with Def. 2.4 and approximation of approximation (Prop. 2.5(7)) we obtain a $K^{2(T+1-t)}$ -approximation \hat{z}_t of z_t in the form of a canonical representation of size $O(\frac{T}{\epsilon} \log \frac{(T+2-t)U_g}{g_{T+1}^{\min}})$. This concludes the induction claim. It remains to show that the runtime of this step is polynomial. By Thm. 4.2, the runtime of step 7 is $O(t_{\bar{z}} \log \frac{\kappa_{\bar{z}_t} U_S}{\epsilon g_{T+1}^{\min}} \log_K \frac{(T+2-t)U_g}{g_{T+1}^{\min}})$. Note that by (19) the Lipschitz constant $\kappa_{\bar{z}_t}$ is bounded above by $(3U_f U_A^2)^{(T+1-t)} \kappa$. We now proceed to bound $t_{\bar{z}}$, i.e., the time to construct and compute (19) at a given value of I_t .

We first analyze the size of the LP (19). The variables y, w used in (19) can be substituted out. Thus, (19) has $p + 3$ variables: p variables in the description of $\mathcal{A}_t(I_t)$, plus one variable for the minmax formulation of each of the convex functions $g_t^I(I_t, \cdot)$, $\tilde{G}_t^D(\cdot)$ and $\tilde{Z}_{t+1}(\cdot)$. The LP has $O(m + q_t + \frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} + \frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma})$ constraints plus p nonnegativity constraints: m constraints come from $\mathcal{A}_t(I_t)$, q_t from the minmax formulation of the piecewise linear convex cost g_t^I , the remaining constraints come from the minmax formulation of \tilde{G}_t^D and \tilde{Z}_{t+1} , whose number of pieces is discussed above. The size of the numbers in the LP is bounded by $O(\max\{U_A, U_f, \log \frac{1}{\epsilon} + \log \frac{(T+2-t)U_g}{g_{T+1}^{\min}}\})$. An LP with these characteristics is solved in $O((m + q_t + \frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} + \frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma} + p)^{1.5} p^2 \max\{U_A, U_f, \log \frac{1}{\epsilon} + \log \frac{(T+2-t)U_g}{g_{T+1}^{\min}}\})$ arithmetic operations, following [Ren88]. In addition, the piecewise linear representation for the functions \tilde{G}_t^D and \tilde{Z}_{t+1} is computed for each fixed I_t , which requires additional $O(\frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} \log m_t + \frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma} \log \frac{T}{\epsilon} \log \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}})$ time.

Hence, the total runtime to compute a single value of \bar{z}_t is

$$t_{\bar{z}_t} = O\left(\left(\left(m + q_t + \frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} + \frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma} + p\right)^{1.5} p^2 \max\left\{U_A, U_f, \log \frac{1}{\epsilon} + \log \frac{(T+2-t)U_g}{g_{T+1}^{\min}}\right\}\right) + \frac{m_t T \ell}{\epsilon} \log \frac{1}{\gamma} \log m_t + \frac{T^2 \ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}} \log \frac{1}{\gamma} \log \frac{T}{\epsilon} \log \log \frac{(T+1-t)U_g}{g_{T+1}^{\min}}\right). \quad (20)$$

Step 7 has to be repeated T times. Adding the time to compute \tilde{F}_z and \tilde{F}_g , we obtain the total runtime of the approximation scheme (Algorithm 6):

$$O\left(t_{\bar{z}_1} \frac{T^2}{\epsilon} \log \frac{\kappa_{\bar{z}_1} U_S}{\epsilon g_{T+1}^{\min}} \log \frac{(T+1)U_g}{g_{T+1}^{\min}} + \frac{t_F T^2 \ell^3}{\epsilon^2} \log^2 \frac{1}{\gamma} \log(\kappa \ell n^* U_f)\right),$$

where $t_{\bar{z}_0}$ is as in (20). If we perform the substitution, the dependency on ϵ is at most $\frac{1}{\epsilon^5} \log^2 \frac{1}{\epsilon}$, the dependency on T is at most $T^6 \log^3 T$, the dependency on ℓ is at most ℓ^3 , and the dependency on p is at most $p^{3.5}$. If κ, U_A and U_f are small enough so that the Lipschitz constant $\kappa_{\bar{z}_1}$ is bounded by a constant, independent of the problem instance, then the dependency on T decreases to at most $T^5 \log^3 T$. \square

6.2 Condition 3(ii): oracle model

We provide an FPTAS via a reduction to the FPTAS for Condition 3(i). To do so, for $t = 1, \dots, T$, we want to construct piecewise linear convex approximations of g_{T+1} and g_t^D with bounded error. In general, relative

approximations for Lipschitz continuous univariate convex functions that are described via value oracles may not be computable in finite time, cfr. [HN16, Cor. 2.5]. However, approximations that attain small multiplicative plus additive error (i.e., with both additive and relative error) can be found in polynomial time [HN16, Prop. 3.7]. The optimum of (DP) has value of at least $g_{T+1}^{\min} > 0$ by assumption, therefore a small additive error in each of the T backward recursion steps can be absorbed without forsaking the approximation guarantee. In particular, we slightly reduce the relative approximation factor at each stage, while allowing an additive error of the order of $\epsilon g_{T+1}^{\min}/2(2T+1)$. This is the main idea to adapt the FPTAS to the case in which g_{T+1}, g_t^D are Lipschitz continuous univariate convex functions accessed by value oracles.

We now give more detail. We start by formally defining (Σ, Π) -approximation functions, as introduced in [HN16, Def. 2.10]. Σ denotes an additive approximation error, $\Pi = 1 + \epsilon$ a multiplicative approximation error; Π has the same role as K in this paper, and after reporting a few basic results from [HN16] we revert to our usual notation using K for relative error.

Definition 6.3 *Let $\epsilon > 0, \Sigma > 0, \Pi = 1 + \epsilon$ and let $\varphi : [A, B] \rightarrow \mathbb{R}^+$. We say that $\tilde{\varphi} : [A, B] \rightarrow \mathbb{R}^+$ is a (Σ, Π) -approximation function of φ (or more briefly, a (Σ, Π) -approximation of φ) if for all $x \in [A, B]$ we have $\varphi(x) \leq \tilde{\varphi}(x) \leq \Pi\varphi(x) + \Sigma$.*

As stated in the following proposition, it is possible to efficiently construct piecewise linear convex (Σ, Π) -approximation functions for Lipschitz continuous univariate convex functions that are described via value oracles.

Proposition 6.4 (Adapted from Prop. 3.7 in [HN16]) *Let $\varphi : [A, B] \rightarrow \mathbb{R}^+$ be a κ -Lipschitz continuous convex function. Then, for every constants $\Sigma > 0$ and $\Pi = 1 + \epsilon > 1$, one can construct a piecewise-linear convex (Σ, Π) -approximation function $\check{\varphi}$ for φ with $p := O\left(\frac{1}{\epsilon} \log \frac{\epsilon \varphi^{\max}}{\Sigma}\right)$ pieces in $O\left((1 + t_\varphi)\left(\frac{1}{\epsilon} \log \frac{\epsilon \varphi^{\max}}{\Sigma}\right) \log \frac{\kappa(B-A)}{\Sigma}\right)$ time, with explicitly computed breakpoints and slopes. Moreover, the value of $\check{\varphi}(\cdot)$ can be determined in $\log p$ time at any point in $[A, B]$.*

Proposition 6.5 (Calculus of (Σ, Π) -approximation Functions) *For $i = 1, 2$ let $\Sigma_i \geq 0, \Pi_i \geq 1$, let $\varphi_i : D \rightarrow \mathbb{R}^+$ be an arbitrary function over continuous domain D , and let $\tilde{\varphi}_i : D \rightarrow \mathbb{R}^+$ be a (Σ_i, Π_i) -approximation of φ_i . Let $\psi_1 : D \rightarrow D$, and let $\alpha_i \in \mathbb{R}^+$. The following rules hold:*

1. φ_1 is a $(0, 1)$ -approximation of itself,
2. (linearity of appr.) $\alpha_1 \tilde{\varphi}_1 + \alpha_2$ is a $(\alpha_1 \Sigma_1, \Pi_1)$ -approximation of $\alpha_1 \varphi_1 + \alpha_2$,
3. (summation of appr.) $\tilde{\varphi}_1 + \tilde{\varphi}_2$ is a $(\Sigma_1 + \Sigma_2, \max\{\Pi_1, \Pi_2\})$ -approximation of $\varphi_1 + \varphi_2$,
4. (composition of appr.) $\tilde{\varphi}_1(\psi_1)$ is a (Σ_1, Π_1) -approximation of $\varphi_1(\psi_1)$,
5. (minimization of appr.) $\min\{\tilde{\varphi}_1, \tilde{\varphi}_2\}$ is a $(\max\{\Sigma_1, \Sigma_2\}, \max\{\Pi_1, \Pi_2\})$ -approximation of $\min\{\varphi_1, \varphi_2\}$,
6. (maximization of appr.) $\max\{\tilde{\varphi}_1, \tilde{\varphi}_2\}$ is a $(\max\{\Sigma_1, \Sigma_2\}, \max\{\Pi_1, \Pi_2\})$ -approximation of $\max\{\varphi_1, \varphi_2\}$,
7. (approximation of appr.) If $\varphi_2 = \tilde{\varphi}_1$ then $\tilde{\varphi}_2$ is a $(\Sigma_2 + \Pi_2 \Sigma_1, \Pi_1 \Pi_2)$ -approximation of φ_1 .

Proposition 6.6 *Let $\xi : [A, B] \rightarrow \mathbb{R}^+$ be a (not necessarily monotone) convex function. Let $K_1, K_2 \geq 1$ and $\Sigma > 0$. Let $\psi : [A, B] \rightarrow \mathbb{R}^+$ be an increasing piecewise linear convex function with breakpoints $A = a_1 < \dots < a_n < B$ and slopes $0 = \Delta_0 \leq \Delta_1 < \dots < \Delta_n$ that (Σ, K_1) -approximates ξ . Let D be a (not necessarily nonnegative) truncated continuous r.v. with support $[D^{\min}, D^{\max}]$ and such that $\min\{\Pr(D = D^{\min}), \Pr(D = D^{\max})\} > 0$. Suppose D admits a (generalized) PDF F' , and let F be the corresponding CDF. Let $\tilde{F}(\cdot)$ be a monotone K_2 -approximation of F with breakpoints at $D^{\min} = d_1 < \dots < d_m = D^{\max}$, and let $\tilde{F}(d_0) = 0$. Let $f(x, d) = bx + e - d$ for some given $b, e \in \mathbb{R}$, and*

$m_i(x) = \max\{j \mid d_j \leq bx + e - a_i\}$. Then

$$\tilde{\xi}(x) = \psi(A) + \sum_{i=1}^n (\Delta_i - \Delta_{i-1}) \left((bx + e - d_{m_i(x)} - a_i) \tilde{F}(d_{m_i(x)+1}) + \sum_{k=1}^{m_i(x)-1} (d_{k+1} - d_k) \tilde{F}(d_{k+1}) \right) \quad (21)$$

is a convex piecewise linear $(\Sigma K_2, K_1 K_2)$ -approximation function of $\mathbb{E}_D(\xi(f(x, D)))$ with $O(mn)$ pieces. Furthermore, one can construct in $O(mn \log n)$ time an oracle for $\tilde{\xi}(\cdot)$ in the form of either a canonical representation, or a representation that consists of breakpoints and slopes, each of size $O(mn)$.

The proof of this proposition is identical to the one of Proposition 5.2, except that, instead of using inequality (17), we use

$$\begin{aligned} \mathbb{E}_D(\psi(f(x, D))) &= \int_{d_1}^{d_m} \psi(f(x, d)) F'(d) dd \\ &\leq \Sigma + K_1 \int_{d_1}^{d_m} \xi(f(x, d)) F'(d) dd = \Sigma + K_1 \mathbb{E}_D(\xi(f(x, D))). \end{aligned} \quad (22)$$

As a consequence of the above proposition, we obtain the following.

Proposition 6.7 *Let $\xi : [A, B] \rightarrow \mathbb{R}^+$ be a convex function. Let $K_1, K_2 \geq 1$ and $\Sigma > 0$. Let $\psi : [A, B] \rightarrow \mathbb{R}^+$ be a piecewise linear convex function with n breakpoints that (Σ, K_1) -approximates ξ over $[A, B]$. Let D be a (not necessarily nonnegative) continuous r.v. satisfying the conditions stated in Prop. 5.2. Let $\tilde{F}(\cdot)$ be a monotone nondecreasing K_2 -approximation of the CDF of D with m breakpoints. Let $f(x, D) = bx + e + cd$, with $b, e, c \in \mathbb{R}$. Then we can construct in $O(mn \log n)$ time a convex piecewise linear $(K_2 \Sigma, K_1 K_2)$ -approximation function of $\mathbb{E}_D(\xi(f(x, D)))$ in the form of either a canonical representation, or a representation that consists of breakpoints and slopes, each of size $O(mn)$.*

Similarly, we can modify Prop. 6.1 to allow for a small additive error. The proof for the result below is the same as for Prop. 6.1, using Prop. 6.5 for the approximation bound.

Proposition 6.8 *Suppose the DP formulation (DP) satisfies Conditions 1-3. Let $K_1, K_2, K_3 > 1$, $\Sigma > 0$. Let \hat{z}_{t+1} be a piecewise linear convex K_1 -approximation of the value function z_{t+1} . For a given $I_t \in S_t$, let $\tilde{Z}_{t+1}(\cdot)$ be a piecewise linear convex K_2 -approximation of $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \bar{\theta}^D \cdot \bar{D}_t)]$, and $\tilde{G}_t^D(\cdot)$ be a piecewise linear convex (Σ, K_3) -approximation of $\mathbb{E}[g_t^D(\cdot + \sigma^I I_t + \bar{\sigma}^D \cdot \bar{D}_t)]$. Then the value of the following mathematical program:*

$$\left. \begin{aligned} \bar{z}_t(I_t) := \min_{\vec{x}_t} \quad & g_t^I(I_t, \vec{x}_t) + \tilde{G}_t^D(w) + \tilde{Z}_{t+1}(y) \\ & A_t \vec{x}_t \geq \vec{b}_t + \bar{\delta}_{b_t} I_t \\ & \bar{\sigma}^x \cdot \vec{x}_t - w = 0 \\ & \bar{\theta}^x \cdot \vec{x}_t - y = 0 \\ & \vec{x}_t \geq 0, \end{aligned} \right\} \quad (23)$$

is a $(\Sigma, \max\{K_1 K_2, K_3\})$ -approximation of the optimal value function $z_t(I_t)$. (23) can be cast as an LP using the standard min max reformulation of piecewise linear convex functions.

Equipped with these concepts, we can analyze the FPTAS for the case of Condition 3(ii).

Theorem 6.9 *Given a stochastic DP satisfying Conditions 1-3(ii), APXSCHEME2(ϵ) (Algorithm 7) computes a $(1 + \epsilon)$ -approximation of the optimal value function z_1 , and runs in polynomial time in the binary input size and $1/\epsilon$.*

1: $\bar{\epsilon} \leftarrow \frac{\epsilon}{2T+1}, K \leftarrow 1 + \bar{\epsilon}, \bar{K} \leftarrow 1 + \frac{\bar{\epsilon}}{2}$ 2: Using Prop. 6.4, let \hat{z}_{T+1} be a $(\frac{g_{T+1}^{\min} \bar{\epsilon}}{2}, \bar{K})$ -approximation of g_{T+1} 3: for $t := T$ downto 1 do 4: Compute a $(\frac{g_t^{\min} \bar{\epsilon}}{2}, K)$ -approximation of g_t^D , called \tilde{g}_t^D , using Prop. 6.4 5: $\tilde{F}_g \leftarrow \text{COMPRESSCONVOLUTION}(\vec{D}_t, \vec{\sigma}^D, K)$ 6: $\tilde{F}_z \leftarrow \text{COMPRESSCONVOLUTION}(\vec{D}_t, \vec{\theta}^D, \bar{K})$ 7: For fixed I_t , define $\tilde{G}_t^D \leftarrow \text{COMPRESSEXVAL}(\tilde{g}_t^D, (1, \sigma^I I_t, 1), \tilde{F}_g)$ /* $\tilde{G}_t^D(\cdot)$ is an oracle for a $(\frac{g_t^{\min} \bar{\epsilon} K}{2}, K^2)$ -approximation of $\mathbb{E}[g_t^D(\cdot + \sigma^I I_t + \vec{\sigma}^D \cdot \vec{D}_t)]$ */ 8: For fixed I_t , define $\tilde{Z}_{t+1} \leftarrow \text{COMPRESSEXVAL}(\hat{z}_{t+1}, (1, \theta^I I_t, 1), \tilde{F}_z)$ /* $\tilde{Z}_{t+1}(\cdot)$ is an oracle for a \bar{K} -approximation of $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \vec{\theta}^D \cdot \vec{D}_t)]$ */ 9: $\hat{z}_t \leftarrow \text{SCALEDCOMPRESSCONV}(\tilde{z}_t, [\mathcal{S}_t^{\min}, \mathcal{S}_t^{\max}], K)$, where \tilde{z}_t is defined as in (23) 10: return \hat{z}_1
--

Algorithm 7: Procedure APXScheme2(ϵ) for Condition 3(ii)

Proof. We outline here only the differences from the proof of Thm. 6.2, focusing on the computation of the approximation factor. The runtime analysis is very similar to that of Thm. 6.2, with one main modification: g_t^D is not described by m_t breakpoints and slopes, but by an approximation set of size $O\left(\frac{T}{\epsilon} \log \frac{U_g}{g_{T+1}^{\min} \epsilon}\right)$ computed in the course of the algorithm. This has repercussions through the analysis; however, it is easy to verify that it yields only a polynomial difference with respect to the runtime of Thm. 6.2.

The proof proceeds by induction for $t = T + 1, \dots, 1$, showing that at stage t we obtain a piecewise linear convex $K^{2(T+1-t)+1}$ -approximation \hat{z}_t of the value function z_t via an approximation set of cardinality $O(\log_K \frac{(T+2-t)U_g}{g_{T+1}^{\min} \epsilon}) = O(\frac{T}{\epsilon} \log \frac{(T+2-t)U_g}{g_{T+1}^{\min} \epsilon})$. The last equality is due to the fact that $K = 1 + \frac{\epsilon}{2(2T+1)}$. The error bound $K^{2T+1} \leq 1 + \epsilon$ is a consequence of the inequality $(1 + \frac{x}{n})^n \leq 1 + 2x$, which holds for every $0 \leq x \leq 1$ and $n \in \mathbb{N}$.

The step $t = T + 1$ is trivial, because \hat{z}_{T+1} is a $(\frac{g_{T+1}^{\min} \bar{\epsilon}}{2}, \bar{K})$ -approximation of g_{T+1} . By definition of (Σ, Π) -approximation functions $\hat{z}_{T+1}(I_{T+1}) \geq z_{T+1}(I_{T+1})$ for all I_{T+1} . Furthermore,

$$\begin{aligned} \hat{z}_{T+1}(I_{T+1}) &\leq \frac{g_{T+1}^{\min} \bar{\epsilon}}{2} + \left(1 + \frac{\bar{\epsilon}}{2}\right) z_{T+1}(I_{T+1}) = \frac{g_{T+1}^{\min} \epsilon}{2(2T+1)} + \left(1 + \frac{\epsilon}{2(2T+1)}\right) z_{T+1}(I_{T+1}) \\ &\leq \left(1 + \frac{\epsilon}{2T+1}\right) z_{T+1}(I_{T+1}) \leq K z_{T+1}(I_{T+1}), \end{aligned}$$

where the second inequality follows from the fact that \tilde{g}_{T+1}^{\min} is a lower bound for g_{T+1} and consequently for z_{T+1} . By Prop. 6.4 we obtain a piecewise linear convex explicit description of \hat{g}_{T+1} via an approximation set of the stated cardinality. We now show the induction step.

At step 7, for a fixed value of I_t we define a $(\frac{g_t^{\min} \bar{\epsilon} K}{2}, K^2)$ -approximation $\tilde{G}_t^D(\cdot)$ of $\mathbb{E}_{\vec{D}_t}[g_t^D(\cdot + \sigma^I I_t + \vec{\sigma}^D \cdot \vec{D}_t)]$, using Prop. 6.7 with parameters set to $\xi = g_t^D$, $\psi = \tilde{g}_t^D$, $n = O\left(\frac{T}{\epsilon} \log \frac{U_g}{g_{T+1}^{\min} \epsilon}\right)$, $m = O(\frac{T\ell}{\epsilon} \log \frac{1}{\gamma})$, $K_1 = K$, $\Sigma = \frac{g_{T+1}^{\min} \bar{\epsilon}}{2}$, $K_2 = K$.

At step 8, for a fixed value of I_t we define a \bar{K} -approximation $\tilde{Z}_{t+1}(\cdot)$ of $\mathbb{E}[\hat{z}_{t+1}(\cdot + \theta^I I_t + \vec{\theta}^D \cdot \vec{D}_t)]$, using Prop. 5.3 with parameters set to $\xi = \psi = \hat{z}_{t+1}$, $n = O(\frac{T}{\epsilon} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min} \epsilon})$, $m = O(\frac{T\ell}{\epsilon} \log \frac{1}{\gamma})$, $K_1 = 1$ and $K_2 = \bar{K}$. Given a value for I_t , we compute \tilde{Z}_{t+1} in the form of a canonical representation of size $O(\frac{T^2\ell}{\epsilon^2} \log \frac{(T+1-t)U_g}{g_{T+1}^{\min} \epsilon} \log \frac{1}{\gamma})$.

At step 9 we compute an approximation \hat{z}_t for the value function z_t , using the function \tilde{z}_t as defined in

(23). By Prop. 6.8 with parameters set to $K_1 = K^{2(T-t)+1}$, $K_2 = \bar{K}$, $K_3 = K^2$, and $\Sigma = \frac{g_{T+1}^{\min} \bar{\epsilon} K}{2}$, we have that \bar{z}_t is a $(\frac{g_{T+1}^{\min} \bar{\epsilon} K}{2}, K^{2(T-t)+1} \bar{K})$ -approximation of z_t . This is also a $K^{2(T+1-t)}$ -approximation of it, because:

$$\begin{aligned} \hat{z}_t(I_t) &\leq \frac{g_{T+1}^{\min} \bar{\epsilon} K}{2} + K^{2(T-t)+1} \left(1 + \frac{\bar{\epsilon}}{2}\right) z_t(I_t) \leq K^{2(T-t)+1} \left(\frac{\bar{\epsilon}}{2K^{2(T-t)}} + 1 + \frac{\bar{\epsilon}}{2}\right) z_t(I_t) \\ &\leq K^{2(T-t)+1} (1 + \bar{\epsilon}) z_t(I_t) = K^{2(T-t)+2} z_t(I_t), \end{aligned}$$

where we used the facts that $g_{T+1}^{\min} \leq z_t(I_t)$ for all I_t , and $K^{2(T-t)} \geq 1$ since $t \leq T$. The rest of the proof is identical to Thm. 6.2, updating the runtime as described at the beginning of this proof. \square

7 Concluding remarks

This paper presents an FPTAS for stochastic DPs with continuous scalar state spaces and polyhedral action spaces. To construct an approximation algorithm, we introduce several tools within the framework of K -approximation sets and functions. More specifically, we show how to approximately compute the convolution of a finite number of continuous random variables, the expectation of a convex function applied to a linear transformation of a continuous random variable, and how to bound the size of the numbers of a K -approximation set. Combining these tools, we obtain an FPTAS for a general class of DP models. These models can be seen as multistage stochastic LPs with one variable linking the stages. The most important open question is whether our results can be extended to the case of two (or more) variables linking the stages, or, in other words, two-dimensional state spaces for the DP. Under an oracle model for the cost functions, this paper shows that the problem is intractable even for the two-dimensional case. If the cost functions are known explicitly, the tools developed here do not suffice to settle the approximability status of the problem. This is left for future research.

References

- [Bel57] Richard E. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957.
- [Bel61] Richard E. Bellman. *Adaptive control processes: a guided tour*. Princeton university press, 1961.
- [Ber95] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 1995.
- [BL11] John R. Birge and Francois Louveaux. *Introduction to stochastic programming*. Springer Science & Business Media, 2011.
- [CDJ14] Wei Chen, Milind Dawande, and Ganesh Janakiraman. Fixed-dimensional stochastic dynamic programs: An approximation scheme and an inventory application. *Operations Research*, 62(1):81–103, 2014.
- [CG18] Yuan Chen and David Goldberg. Beating the curse of dimensionality in options pricing and optimal stopping. *ArXiv e-prints*, July 2018.

- [dFVR03] Daniela Pucci de Farias and Benjamin Van Roy. The Linear Programming approach to approximate Dynamic Programming. *Operations Research*, 51(6):850–865, 2003.
- [FT90] Paola Favati and Fabio Tardella. Convexity in nonlinear integer programming. *Ricerca Operativa*, 53:3–44, 1990.
- [GKM⁺11] Parikshit Gopalan, Adam Klivans, Raghu Meka, Daniel Štefankovic, Santosh Vempala, and Eric Vigoda. An FPTAS for # knapsack and related counting problems. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 817–826. IEEE, 2011.
- [Hal15] Nir Halman. Approximating convex functions via non-convex oracles under the relative noise model. *Discrete Optimization*, 16:1–16, 2015.
- [Hal16] Nir Halman. A deterministic fully polynomial time approximation scheme for counting integer knapsack solutions made easy. *Theoretical Computer Science*, 645:41–47, 2016.
- [Hal17] Nir Halman. Provably near-optimal approximation schemes for implicit stochastic and for sample-based dynamic programs. Technical Report 4952, Optimization Online, 2017.
- [HKL⁺14] Nir Halman, Diego Klabjan, Chung-Lun Li, James Orlin, and David Simchi-Levi. Fully polynomial time approximation schemes for stochastic dynamic programs. *SIAM Journal on Discrete Mathematics*, 28(4):1725–1796, 2014.
- [HKM⁺09] Nir Halman, Diego Klabjan, Mohamad Mostagir, James Orlin, and David Simchi-Levi. A fully polynomial time approximation scheme for single-item inventory control with discrete demand. *Mathematics of Operations Research*, 34(3):674–685, 2009.
- [HN16] Nir Halman and Giacomo Nannicini. Fully polynomial-time (Σ, Π) -approximation schemes for continuous nonlinear newsvendor and continuous stochastic dynamic programs. Technical Report 5726, Optimization Online, 2016.
- [HNO15a] Nir Halman, Giacomo Nannicini, and James Orlin. A computationally efficient FPTAS for convex stochastic dynamic programs. *SIAM Journal on Optimization*, 25(1):317–350, 2015.
- [HNO15b] Nir Halman, Giacomo Nannicini, and James Orlin. On the complexity of energy storage problems. Technical report, Singapore University of Technology and Design, 2015.
- [KRT97] Jon Kleinberg, Yuval Rabani, and Eva Tardos. Allocating bandwidth for bursty connections. In *Proceedings of the 29th ACM Symposium on the Theory of computing (STOC)*, pages 664–673, 1997.
- [LS14] Jian Li and Tianlin Shi. A fully polynomial-time approximation scheme for approximating a sum of random variables. *Operations Research Letters*, 42(3):197 – 202, 2014.
- [Mil71] Bruce L. Miller. On minimizing nonseparable functions defined on the integers with an inventory application. *SIAM Journal on Applied Mathematics*, 21:166–185, 1971.
- [Mur03] Kazuo Murota. *Discrete Convex Analysis*. SIAM, Philadelphia, 2003.
- [NP13] Juliana M. Nascimento and Warren B. Powell. An optimal approximate dynamic programming algorithm for concave, scalar storage problems with vector-valued controls. *IEEE Transactions on Automatic Control*, 58(12):2995–3010, 2013.

- [PGa⁺12] Warren B. Powell, Abraham George, Hugo Sim ao, Warren Scott, Alan Lamont, and Jeffrey Stewart. SMART: A stochastic multiscale model for the analysis of energy resources, technology, and policy. *INFORMS Journal on Computing*, 24(4):665–682, 2012.
- [Pow11] Warren B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition*. Wiley, 2011.
- [Ren88] James Renegar. A polynomial-time algorithm, based on Newton’s method, for linear programming. *Mathematical Programming*, 40(1):59–93, 1988.
- [SS12] Chaitanya Swamy and David B. Shmoys. Sampling-based approximation algorithms for multistage stochastic optimization. *SIAM Journal on Computing*, 41(4):975–1004, 2012.

A Additional routines

For the sake of self-containment, we report here the pseudocode and running time of the routines taken from [HKL⁺14, HN16] referenced in the main text of the paper. Given a monotone nondecreasing function φ , we define a routine `FUNCSEARCHINC`(φ, D, ℓ, u) that looks for a point $x \in D$ such that $\ell \leq \varphi(x) \leq u$. Implementing this function is straightforward for both discrete domains and real intervals, using binary search.

In Alg. 8 we formally define the routine `COMPRESSINC` that constructs an oracle for a K -approximation function of a monotone nondecreasing φ in the form of a canonical representation.

The main routine used in Algorithm 8 is `APXSETINC`. We give its pseudocode in Algorithm 9 for continuous functions that are bounded away from zero; its counterpart for functions over discrete domains is straightforward. It is shown in [HN16] that `APXSETINC` determines a K -approximation set with $O\left(\frac{1}{\epsilon} \log \frac{\epsilon \varphi^{\max}}{\varphi^{\min}}\right)$ points in $O\left((1 + t_\varphi) \left(\frac{1}{\epsilon} \log \frac{\epsilon \varphi^{\max}}{\varphi^{\min}}\right) \log((B - A)\kappa)\right)$ time. For a function defined over a discrete domain D , the running time of `APXSETINC` becomes $O\left((1 + t_\varphi) \left(\frac{1}{\epsilon} \log \frac{\epsilon \varphi^{\max}}{\varphi^{\min}}\right) \log |D|\right)$, see [HKL⁺14].

1: $W \leftarrow \text{APXSETINC}(\varphi, [A, B], K)$ 2: return $\{(x, \varphi(x)) \mid x \in W\}$ as an array of tuples sorted by their first coordinate
--

Algorithm 8: Function `COMPRESSINC`($\varphi, [A, B], K$).

1: $x \leftarrow A, W \leftarrow \{A, B\}$ 2: while $x < B$ do 3: $x \leftarrow \text{FUNCSEARCHINC}(\varphi, [x, B], \frac{K+1}{2}\varphi(x), K\varphi(x))$ 4: $W \leftarrow W \cup \{x\}$ 5: return W
--

Algorithm 9: FUNCTION `APXSETINC`($\varphi, [A, B], K$)